Spatiotemporal Effects of Segregating Different Vehicle Classes on Separate Lanes

Michael J. Cassidy, Carlos F. Daganzo, and Kitae Jang
WORKING PAPER
UCB-ITS-VWP-2008-8

September 2008

Spatiotemporal Effects of Segregating Different Vehicle Classes on Separate Lanes

Michael J. Cassidy, Carlos F. Daganzo and Kitae Jang, Department of Civil and Environmental Engineering and Institute of Transportation Studies, University of California, Berkeley, California, USA; K. Chung, California Department of Transportation, USA

Abstract Spatiotemporal analysis of real freeway traffic reveals that carpool lanes are not as damaging as previously reported. To the contrary, the analysis unveils a surprising benefit of carpool lanes that should be even greater when special lanes are used to segregate very different vehicle classes, such as buses and cars. The paper pursues this finding and shows how reserving lanes on freeways and city streets for bus-use only can favorably affect not just buses, but also cars.

1. Introduction

Empirical assessments of road traffic are numerous and date back at least as far as the 1930's. Many, if not most of these studies examine time series of vehicle flows, speeds, etc. from a single milepost on a road, while ignoring (or not collecting) measurements at neighboring mileposts.

The present paper illustrates the problems inherent in this kind of purely temporal analysis. It does so by revisiting in spatiotemporal fashion a study of freeway carpool lanes published in earlier proceedings of this symposium series (Chen, et al, 2005).[1] Purely temporal in its approach, that earlier study was titled an Empirical Assessment of Traffic Operations, and will be referred to here as EATO. After examining some freeways in California's San Francisco Bay Area, EATO concluded that carpool lanes were creating congestion; and recommended that carpool restrictions be rescinded everywhere on the study sites, so that these lanes would always be available for general use. EATO recognized, however, that its analyses were fragmentary, and therefore also recommended that its findings be

---

[1] The carpool lanes are set aside during rush hours to serve only vehicles (mostly cars) that carry more than a predetermined number of occupants.

solidified through further assessments. Thus, further assessments are done in sections 2 and 3 of the present paper; they lead to a conclusion that is different from EATO, and unveil new phenomena indicating that special-use lanes (such as those reserved for carpools, buses, etc.) can benefit all modes.

Section 2 provides background on the EATO study. We will show that by ignoring the spatial component of the system, EATO reached conclusions that cannot be supported.

Section 3 reassesses the EATO sites in spatiotemporal detail, and over multiple days. We find no evidence that the carpool lanes triggered congestion on any of the sites (in all cases presented in EATO, queues formed first at bottlenecks for reasons unrelated to the carpool lanes); and no evidence that the carpool lanes reduced bottleneck flows or prolonged the queues, even though these lanes were underutilized. To the contrary, our detailed analysis shows both: that a severely underutilized carpool lane passing through a bottleneck was actually increasing the bottleneck's total discharge flow; and the reason for this counterintuitive effect.

Section 4 extends this finding to other lane reservation strategies: theory is used to quantify the benefits of setting aside special lanes to segregate very different vehicle classes, such as buses and cars. Section 5 discusses the implications of our findings.

2. Background

The EATO study used time series of vehicle speeds and flows (by lane) as metrics to assess the carpool lanes' impacts on freeway traffic. These data were measured at single detector stations, and were not analyzed together with data from neighboring detectors. As a result, EATO identified periods when queues persisted atop a detector, and correlated these with carpool-lane operating times, but could not verify the locations where these queues initially formed. To see how the missing information colored the analysis, refer to Figs. 1(a) – (f). These charts reproduce the speed time series data in Fig. 8 of EATO and characterize all of the carpool facilities it studied. Our charts include additional annotations to aid in their interpretation.

On all the freeways of Figs. 1(a) – (f), the median lane is reserved for carpools on weekdays during the afternoon rush, from 15:00 to 19:00. This period is demarcated by vertical lines in each chart. In all cases, speeds are lower during that period than outside it, both in the carpool lanes (dark curves) and in the adjacent General Purpose (GP) lanes. Our concern is with the impacts that special lanes have on those vehicle classes excluded from using them, and we therefore focus on the (larger) speed drops in the GP lanes.[2]

---

[2] Note as an aside that the speed reductions in the carpool lanes were found to have insignificant impacts on carpool-vehicle delay (see Cassidy, et al, 2006).

The EATO study claims (without examining demand effects) that queues arose in these GP lanes because they were eventually in short supply; i.e., demand among Low Occupancy Vehicles (LOVs) presumably grew while the median lane was unavailable for their use, and this supposedly pushed the GP lanes into the congested regime. The evidence of this mechanism is said in EATO to come after 19:00 hrs because by this time, when each lane-use restriction had been lifted, the speeds reportedly increased. EATO's conclusion is that speeds rose because the median lane was no longer squandered on carpools.

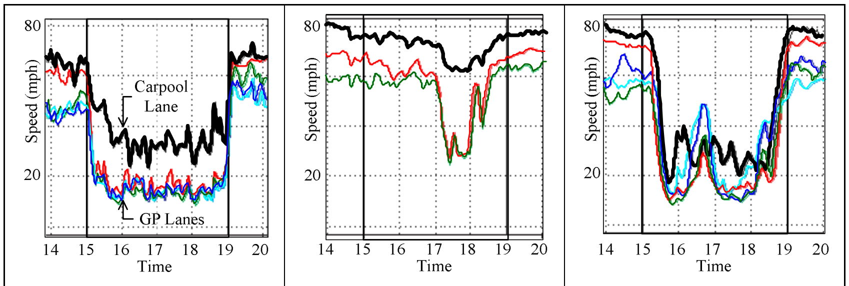

(a) I-80E, Aug 8, 2004 (b) SR-237E, Aug 20, 2004 (c) I-880N, Aug 23, 2004

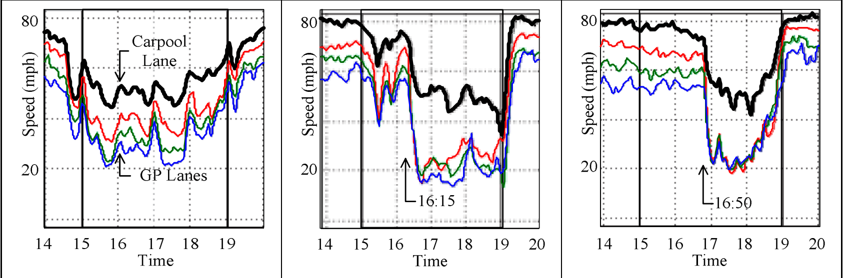

(d) SR-101S, Aug 18, 2004 (e) I-880S, Aug 18, 2004 (f) I-880S, Aug 18, 2004

Fig. 1. Time-Series Diagrams of Speed Furnished in EATO

The evidence, however, is not as stated in EATO. In four of the cases (Figs. 1(b), (c), (d) and (f)) speeds began to recover *before* the carpool restriction expired at 19:00, and in three of these (Figs 1(b), (d) and (f)) recoveries began around 18:00 -- an hour before the restriction's expiration time. We cannot conclude that lifting the carpool restriction increased the speeds (removed the queues) in these cases: an effect cannot precede its cause. The more plausible explanation is that congestion eased because demand declined at the end of the rush. And of course, lack of evidence around 19:00 does not mean that the carpool lanes can not be causing other problems at other times. These matters can only be sorted-out by examining the events of Fig. 1 in spatiotemporal detail.

3. Spatiotemporal Assessments

The cases in Fig. 1 are considered here: section 3.1 examines cases (e) and (f); section 3.2 case (c); and section 3.3 cases (a), (b) and (d). We will pinpoint the real causes of congestion by unveiling the effects of both the carpool restrictions and the traffic demands.

3.1. Cases (e) and (f)

The last two charts of Fig. 1 come from neighboring detectors on the same freeway (southbound Interstate 880, labeled I-880S in the charts) during the same day and time, though EATO did not analyze their spatiotemporal relationship. (It mistakenly states that the data were from different freeways.) Note how speed dropped much earlier at one detector location than at the other, suggesting that the queue formed at a specific location.

Fig. 2 reveals that location: it presents a spatiotemporal plot of occupancy from all the detectors along the relevant stretch of I-880, including the two detectors in EATO. Note in interpreting the figure that: (i) Fig. 2 spans the same observation period as Figs. 1(e) and (f); (ii) occupancies of about 20 percent or more denote the presence of queues; and (iii) occupancies below 20 percent imply no queues; i.e. that flow *is* demand. The downward-slanting, diagonal pattern separating light and dark shadings (labeled queue growth in Fig. 2) shows that the queue started locally at Post Mile (PM) 18.7, around 15:30 hrs. The queue then grew, eventually causing the speed reductions visible in Fig. 1(e) around 16:15, and then the reductions of Fig. 1(f) around 16:50.

The accident log maintained by the state police indicates that this queue was triggered by a vehicle collision, and not by the carpool lane. An archival record of that collision can be found at http://pems.eecs.berkeley.edu.

Fig. 2 also shows that after the collision was removed, a second bottleneck became active at PM 26.7 from 17:30 onward: the shading reveals a queue upstream of this location, with freely flowing conditions downstream. Note that PM 26.7 is the location of a merge. Later in the rush, but still prior to 19:00, the back of the queue gradually receded forward toward this second bottleneck for lack of demand, and eventually dissipated. Thus, the gradual speed recoveries seen in Figs. 1(e) and (f) were due to a reduction in traffic demand, and not to the expiration of the carpool restriction. The speed recovery at PM 26.7 (Fig. 1(e)) coincided with the expiration of the carpool restriction only by chance, and this could have been a confounding factor in EATO.[3]

---

[3] Although the expiration of the lane-use restriction could have slightly accelerated the queue's dissipation, Fig. 2 shows that the queue was already well on its way to dissipating, and likely would have done so at around 19:00 hrs – even if the lane-use restriction had not expired.

We examined this freeway for nine additional weekdays (in July and August, 2004). On four of these days, a queue did not arise at all. On each of the five other days, a small queue did form, and always did so locally at the merge bottleneck near PM 26.7.

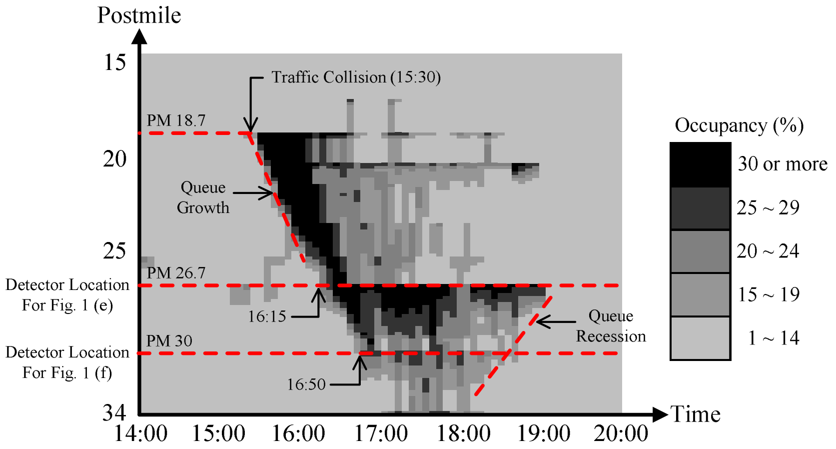

Fig. 2. Time-Space-Occupancy Plot, I-880S (Aug 18, 2004)

It may be tempting to blame this bottleneck (near PM 26.7) on the carpool lane because the lane was underutilized: flow in the lane remained at, or slightly below, 1500 vph while the carpool restriction was in force. However, assigning culpability to the carpool lane would have been premature because the bottleneck discharge flows in the other lanes actually increased while the carpool restriction was in force. We investigate this phenomenon next by using video data from another merge bottleneck where the carpool lane was even more underutilized.

3.2 Case (c)

Fig. 3 displays a time-space-occupancy plot for a long stretch of the freeway I-880 (northbound) that includes the detector station used in Fig. 1(c), and spans the same period. Once again, the occupancies show that a queue started at definite points in space and time: namely, at PM 26, the location of a merge, and 15:00 hrs.[4] Since the queue formed when the carpool restriction took effect, and since

---

[4] The lightly shaded rectangle pinned at the bottleneck was the result of a collision that occurred at 15:49 and PM 25; see again http://pems.eecs.berkeley.edu.

there was a good vantage point near the merge, we used videos to determine both: the mechanism of queue formation; and the subsequent effect of the carpool lane.

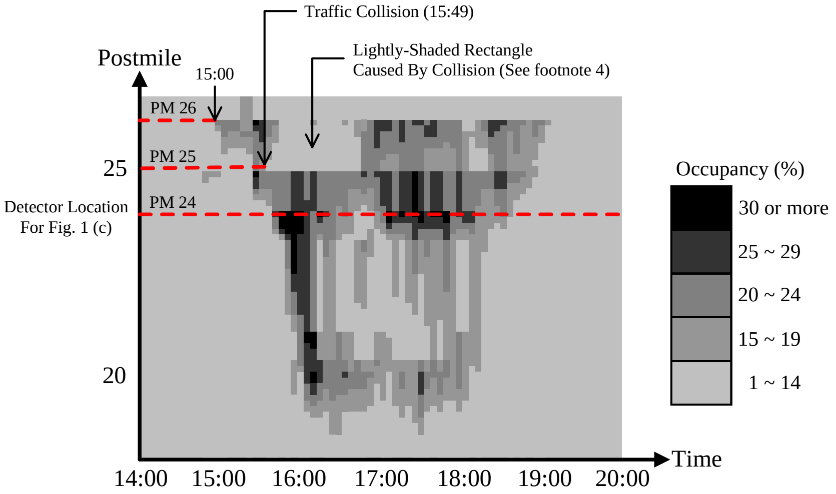

Fig. 3. Time-Space-Occupancy Plot, I-880N (Aug 23, 2004)

<u>Queue Formation:</u> Fig. 4 displays the freeway geometry in the bottleneck's vicinity. Video cameras were erected on the pedestrian over-crossing, and these recorded traffic during part of an afternoon rush (on July 19, 2006). Vehicle arrival times at locations X_1, X_2 and X_3 were manually extracted from the videos and, as is customary, cumulative curves of vehicle count were plotted on an oblique coordinate system (O-curves), as shown in Fig. 5. Note that the slopes of the O-curves are the excess flows over a background flow, which is 6800 vph in the present case; and that the curves in Fig. 5 were constructed in such ways that superimposed curves indicate free-flow traffic (flow = demand) and separated curves indicate delays: the wider the separation the longer the delays (see Cassidy and Windover, 1995; and Munoz and Daganzo, 2002).

In Fig. 5, curves 2 and 3 are superimposed, and below curve 1. Thus, traffic was freely flowing between X_2 and X_3, but delays existed between X_1 and X_2. These two curves diverged for good at about 14:43 hrs when a disruption reduced the flow at X_2. Less than 3 minutes later (at approximately 14:45:30) flow dropped further to about 7020 vph. The video data establish that this flow reduction was triggered by a queue that first formed in the shoulder lane due to pulses of merging vehicles, and then spread to all lanes; the carpool lane had no role in this.

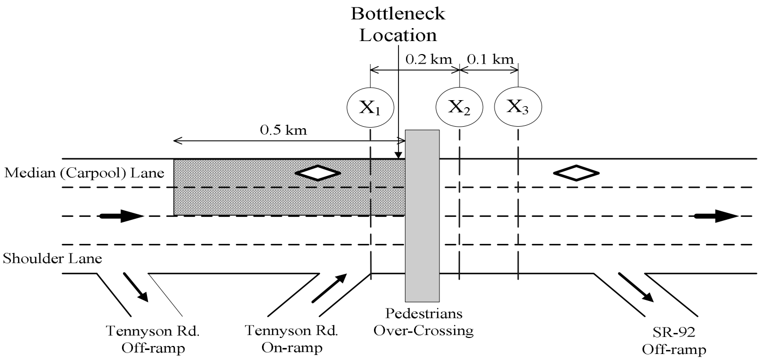

Fig. 4. I-880N Freeway Geometry

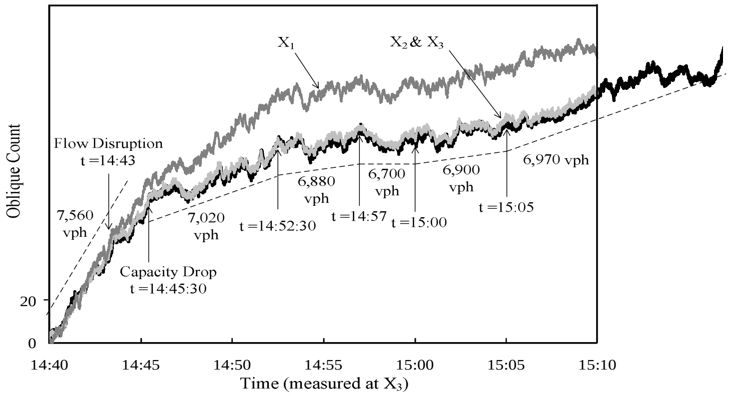

Fig. 5. O-Curves at X_1 through X_3 (July 19, 2006)

To illustrate, Fig. 6 presents two O-curves of the bottleneck discharge flows at location X_2: one for the shoulder lane (in boldface) and another for the three remaining lanes combined. Note how the flow in the shoulder lane suddenly diminished from 2060 vph to 1,790 vph at 14:43 hrs (the time when queuing began in Fig. 5) without any effect on the remaining lanes. The videos clearly reveal that disruptions began in the shoulder lane at 14:43 because vehicles decelerated to make room for merging traffic from the on-ramp; this is the cause of the first reduction in flow.

Fig. 6 also shows that the capacity drop that occurred around 14:45:30 coincided with a flow reduction in the adjacent lanes, signifying that the queue had by then spread across the entire width of the freeway. This traffic pattern is typical of merge bottlenecks *without* carpool lanes (see Cassidy and Rudjanakanoknad, 2005.) In short, the carpool lane did not trigger the bottleneck. As we show below, moreover, the lane did not impede bottleneck flow and prolong congestion, despite being underutilized.

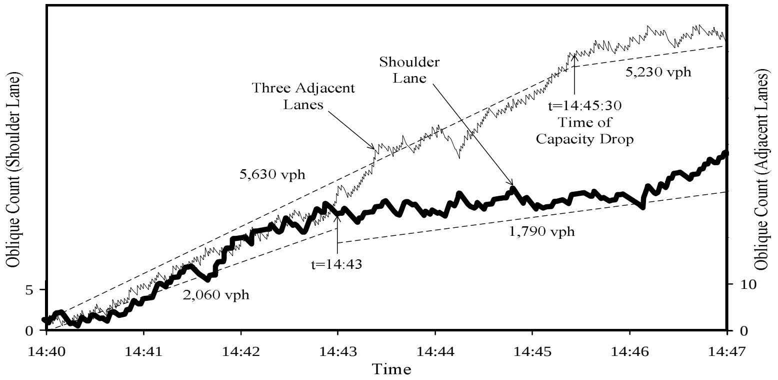

Fig. 6. O-Curves of Shoulder Lane and Adjacent Lanes at X_2 (July 19, 2006)

Subsequent effect of the carpool lane: The video data reveal that drivers (of LOVs) began avoiding the median lane shortly before the carpool restriction went into effect; and that these driver responses began only after the capacity drop had already occurred. This is evident in Fig. 7, which displays the cumulative number of vehicular lane changes, both in and out of the median (carpool) lane, as counted from the videos over the 0.5-km-long shaded segment in Fig. 4. Note from the boldface curve how maneuvers out of the median lane were steady while the capacity was dropping during the period from 14:43 to 14:45:30 hrs; and then how these maneuvers increased from 14:52:30 to 15:00 hrs. This later period likely marks when LOVs began migrating from the lane due to the impending carpool restriction – particularly since the curve also shows that the rate of this lane changing returned to the earlier low value (210/hr/km) once the carpool restriction went into effect at 15:00. Note too that maneuvers into the carpool lane (shown by the thin curve in Fig. 7) were also steady while the capacity was dropping; and that these maneuvers diminished over time: the rate eventually declined from 380/hr/km, to 60/hr/km soon after the carpool restriction took effect.

Fig. 8 shows how these lane changes affected carpool-lane utilization. It displays the lane's O-curve measured at X_3. Note that flow in the lane was 1930 vph immediately following the capacity drop, and that flows steadily diminished thereafter. Note too how the times marking the onsets of reduced flows (from 14:52:30 to 15:05 hrs) coincide with the lane-changing patterns of Fig. 7. So, what effect did the lane changing and flow reductions in the carpool lane have on the bottleneck?

A visual comparison of Figs. 8 and 5 from 14:45:30 hrs onward reveals that the reductions in carpool-lane use, though substantial, had almost no effect on bottleneck discharge rate. The carpool-lane flow (Fig. 8) eventually dropped to 1370 vph, and yet the total flow across all lanes (including the carpool lane) shown in Fig. 5 remained quite steady (at rates approaching 7000 vph). The bottleneck's total discharge flow returned to nearly its highest rate (6970 vph) after 15:05, when carpool-lane flow was lowest (1370 vph). This indicates that the diminished carpool-lane flow was compensated by increased queue discharge flow (capacity) in the adjacent GP lanes. Note too from Fig. 7 that during this period (from 15:05 onward), lane changing to and from the carpool lane was lowest (270/hr/km), indicating that lane changing played a role in producing the higher GP-lane flows.

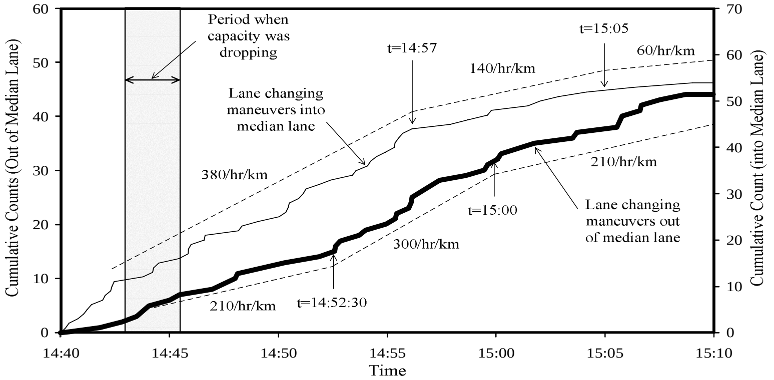

Fig. 7. Cumulative Curves of Lane Changing into and out of Median (Carpool) Lane (July 19, 2006)

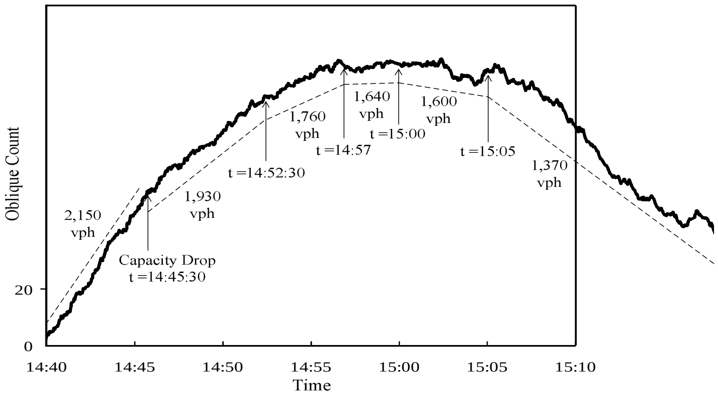

Fig. 8. O-Curve of Median (Carpool) Lane at X_3 (July 19, 2006)

This effect was predicted in Menendez and Daganzo (2007), where it was shown through simulation experiments that a carpool lane's presence can diminish disruptive vehicle lane changes, and that this in turn can smooth (and increase) bottleneck flows. This prediction, moreover, was consistent with an earlier analysis, which had shown that disruptive lane changes were a main cause for the capacity drop at bottlenecks without carpool lanes (Laval and Daganzo, 2006). The present findings confirm that the so-called smoothing effect arises in real traffic; that it is linked to lane changing; and moreover, that it can persist for extended durations. To underscore the latter point, some of the curves in Figs. 5 and 8 are shown for an extended period beyond 15:10. The effect was observed on other days as well.

3.3 Cases (a), (b) and (d)

As in the previous cases, spatiotemporal analyses of the three final sites (in Fig. 1(a), (b) and (d)) uncovered no evidence that the carpool lanes are adding to freeway congestion. To the contrary, the loop detector data from the site in Fig. 1(d) suggest that the carpool lane is reducing congestion slightly by inducing the smoothing effect. Analysis also shows that the site of Fig. 1(b) became congested because of construction activity downstream of the chart's measurement location; and that the site was not congested on days when this roadwork did not occur. For further details, see Cassidy, et al (2006).

In summary, we find that underutilized carpool lanes can sometimes be installed without creating new bottlenecks, thanks to the smoothing effect. And since the smoothing effect is due, at least in part, to reductions in disruptive lane changing, the effect should be even more pronounced for special-use lanes that segregate buses from cars. (Note, for example, that Laval and Daganzo, 2006 shows that lane-changing involving vehicles with low acceleration capabilities, such as buses, create large capacity drops.) We pursue this matter next.

4. Theoretical Generalization: The Effect of Bus Lanes

It is known (Daganzo and Cassidy, 2008) that reserving a freeway lane for carpool use reduces People-Hours of Travel (PHT) if the average passenger occupancy of carpools exceeds that of LOVs by a critical amount that depends on the lane's utilization; and that this happens even in the absence of any smoothing effect. Since the on-board occupancies of buses are on average much greater than those of cars, lane reservation strategies for buses (or trams) should yield benefits even if the special-use lanes are severely underutilized, and even if there is no smoothing effect.

With the smoothing effect, however, comes an added benefit: reserving lanes for buses can improve the flow of cars; and as we will now show, can not only reduce overall PHT, but also the PHT of car users. Because we do not yet know how to predict the magnitude of the smoothing effect for roadways with cars and buses (this is a topic of ongoing research), we will predict its impact on system PHT parametrically. We will do this for rotationally symmetric closed-loop beltways with access and egress via on- and off-ramps because as explained in Daganzo and Cassidy (2008), this is the least favorable environment for a special-use lane.

4.1 Beltway capacities with and without bus lanes

Consider an L-lane, uncongested and rotationally symmetric beltway with on-ramp/off-ramp pairs and with a fixed fraction β of the flow (downstream of each on-ramp) exiting via each off-ramp. We assume that the transit agency supplies enough buses and drivers to sustain the same fixed service frequency, q_B, whether or not any of the beltway's L lanes are reserved for buses. In this way, bus passengers experience the same out-of-vehicle delay in both cases (mixed vs segregated vehicle classes), and we can focus on bus passenger in-vehicle travel time: the PHT. We also assume that buses tend to remain on the beltway during our analysis period, and therefore do not create significant cross-modal conflicts by entering or exiting.

Take q_{max} to be the capacity of a (single) lane carrying only cars (cars/hr). The capacity of a lane carrying only buses (buses/hr) will in general be smaller, since buses make stops and are larger and less maneuverable than cars. Hence, we introduce a constant, p that converts buses into "passenger-car-equivalents" (pce's) still with capacity $q_{max.}$ (We might estimate $p \cong 2$ or 3.) [5] This means that the number of lanes allocated to buses, $l < L$, must satisfy $lq_{max} \geq pq_B$ to prevent buses from being delayed.

We now compare the beltway's maximum possible steady-state car outflows that can exit via all the off-ramps (the beltway capacity) under the two scenarios – segregated vehicle classes with bus lanes and mixed classes without. Since these total outflows are a fixed fraction of the beltway's circulating car flows (measured downstream of each on-ramp merge), we focus for the moment on the latter. Consider first the segregated scenario, and look for the best l. If $lq_{max} \geq pq_B$, the maximum car flow downstream of each beltway merge is $q_{max} \cdot (L-l)$. To maximize it, choose the smallest integer l that satisfies: $lq_{max} \geq pq_B$. The result, l^*, should leave a gap in the inequality smaller than the capacity of one lane, such that $l^* q_{max} - pq_B = uq_{max}$, where $u \in [0, 1)$ is the underutilization level of one of the bus lanes. Since all other lanes, including those devoted to cars, can operate at capacity, the beltway's maximum pce flow would be $q_{max} \cdot (L-u)$ in the segregated scenario.

If vehicle classes are mixed on the other hand, we expect the capacity of a lane (again counting each bus as p cars) to be significantly smaller than q_{max} (in pce's) due to the smoothing effect. This capacity reduction shall be denoted rq_{max}, where $r < 1$ is a positive parameter.[6] So, the beltway's maximum pce flow in this scenario is $(1-r) \cdot q_{max} \cdot L$.

Thus by setting aside bus lanes to segregate vehicle classes, the extra pce flow circulating on the beltway can be as large as $q_{max} \cdot (Lr - u)$. This extra flow is

---

[5] Unfortunately, the literature does not furnish information to help us choose suitable values for p. Handbooks like the Highway Capacity Manual (2000), for example, furnish pce's for buses operating in mixed traffic (only), which is not what we seek here.

[6] This parameter should depend on the mix of buses vs. cars; and be largest when the traffic stream includes significant numbers of both.

composed of cars only, since bus flow, q_B, is fixed; and can produce up to $\beta q_{max}\cdot(Lr - u)$ extra units of car outflow per off-ramp. Deploying l^* bus lanes therefore can improve the beltway's ability to serve cars if $Lr > u$. We expect values of r comparable with 0.2 to arise when the traffic stream contains a significant fraction of buses that make many stops. (This is common in cities that rely heavily on buses to meet their transportation needs.) With r this large, we see that separating modes has the potential for increasing the rate at which a beltway serves cars, which would reduce the PHT of car users.

Recall that our segregation strategy does not affect the PHT of bus users because it keeps invariant both: the bus service frequency, q_B, and the bus speed on the (uncongested) beltway. Therefore, if a beltway can be metered to operate at capacity, segregating street space when modes are very different can improve mobility for everyone, even in the worst-case situation of a symmetric beltway.

4.2 Bus lanes on beltways with queues

Suppose that a beltway cannot be metered very restrictively, such that queues form on it. We show here that even in this case, segregation can increase the beltway's flow (and therefore its input and output flows). This is a good thing because delays and queues would then diminish outside the beltway (e.g. on its on-ramps and connecting streets), without increasing on the beltway itself; and bus users would also benefit by enjoying higher travel speeds.

We now evaluate the circulating car flow in both scenarios, holding the bus flow and the car density constant across scenarios. To this end, we assume the existence of fundamental diagrams that define beltway flows in pce's as a function of density, also in pce's. We assume that the fundamental diagram, Q, for a single segregated lane is identical for cars and buses; and that it is approximately $(1–r)\cdot Q$ for a single lane in the mixed scenario.

Consider first the mixed scenario. Take q_B and q^M as given, where q^M is the beltway's total flow of cars and buses (expressed in pce's). Since bus flow, q_B, is again invariant to the scenario, the queued car flow *per lane* in the mixed scenario is $q_C^M = (q^M - pq_B)/L$. We further define for this scenario k_C^M and k^M as the car density per lane and the combined density of buses and cars (in pce's) per lane, respectively. We know that $k^M = Q^{-1}(q^M/L) / (1- r)$; and since queued traffic in the mixed scenario is first-in, first-out, $k_C^M = L\cdot k^M\cdot q_C^M / q^M$.

If the beltway is converted to the segregated scenario, we would again deploy a sufficient number of bus-only lanes to prevent bus queues from forming: $l^* = [q_B \cdot p / q_{max}]^+$, where $[\ ]^+$ is the ceiling operator. If we define k_C^S as the car density in each of the beltway's $L - l^*$ lanes allocated to cars, we have $k_C^S = k_C^M \cdot L /(L–l^*)$, since car density must be invariant to the scenario. Thus, the total flow of cars in the segregated scenario is: $Q(k_C^S)\cdot(L - l^*)$. To see how and when bus-only lanes favorably affect cars, we now compare the above flow with its counterpart in the mixed scenario, $q_C^M \cdot L$, under varying conditions.

Figs. 9(a) and (b) display $\Delta = (Q(k_C^S)\cdot(L - l^*) - q_C^M\cdot L) / q_C^M\cdot L$, the percent increase in beltway car flow when operation is converted to a segregated scenario, vs $\rho = q^M / (L\cdot q_{max}\cdot(1- r))$, the percent of capacity utilized by queued cars and buses in the mixed scenario. The curves are given for various bus flows, expressed as percentages of beltway capacity in mixed traffic, $s = q_B/(L\cdot q_{max}\cdot(1-r))$, for $p = 2.5$, $L = 3$, $r = 0.1$, 0.2, and based on a triangular Q typical of freeway lanes. These figures show that if bus flow is sufficiently high, a bus-only lane increases car flows; particularly of course when the smoothing effect is large, as in Fig. 9(b). Curiously, for $s \geq 1\%$, bus lanes have greater effects as ρ moves further from 100%, indicating that reserving lanes for buses can be especially beneficial *to cars* when traffic is very congested. The qualitative reason for these gains is that to maintain bus flow when buses have been released from the grip of congestion, one needs fewer buses, and therefore fewer lanes to accommodate them. This leaves proportionally more room for cars.

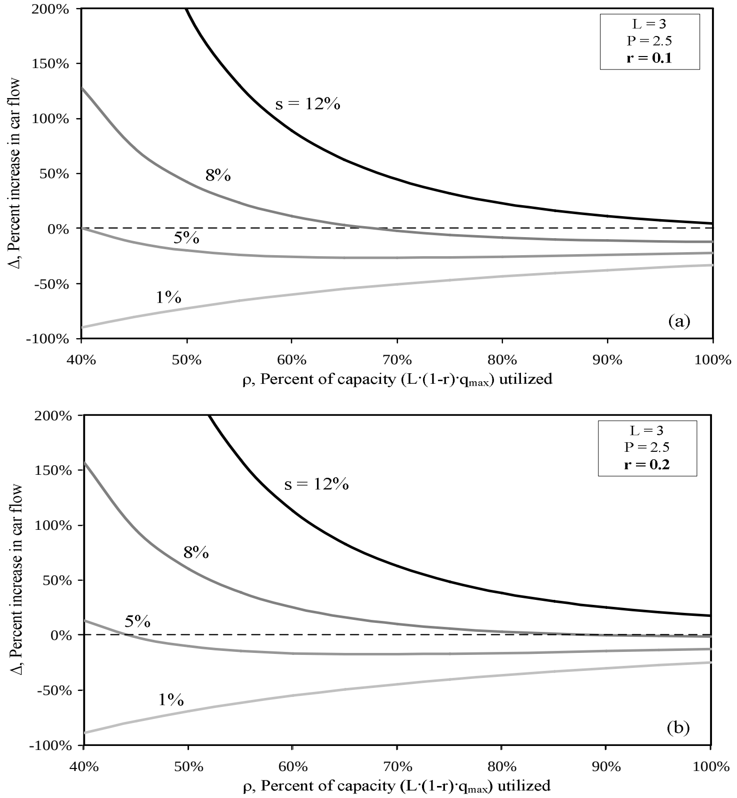

Fig. 9. Curves of ρ vs Δ for a congested freeway. The selected Q has a free flow vehicle speed = 95 km/hr, q_{max} = 2000 cars/hr, and a congested wave speed = 25 km/hr: (a) $r = 0.1$, (b) $r = 0.2$

Segregation can offer even greater benefits on a network of (many) city streets, which we roughly analyze as a beltway with L approaching infinity. Figs. 10(a) and (b) show curves of Δ vs ρ for large L, using a triangular Q suitable for city streets. (In this city-street context, car density could be held invariant to scenario via perimeter control strategies, such as signal metering or pricing, as proposed in Daganzo, 2007, for example.) The figures show that bus-only lanes produce significantly higher car flows even for relatively small q_B. And once again the gains increase as ρ diminishes: the more congested the street network, the greater the attractiveness of bus-only lanes.

Of course, the bus-side of the system benefits even more from segregation. Not only does the bus agency benefit by maintaining the stipulated q_B with fewer vehicles and drivers, but by bypassing the car queue, the bus passengers enjoy a reduction in PHT. What we have shown is that these bus benefits can sometimes be achieved while benefiting car-users as well.

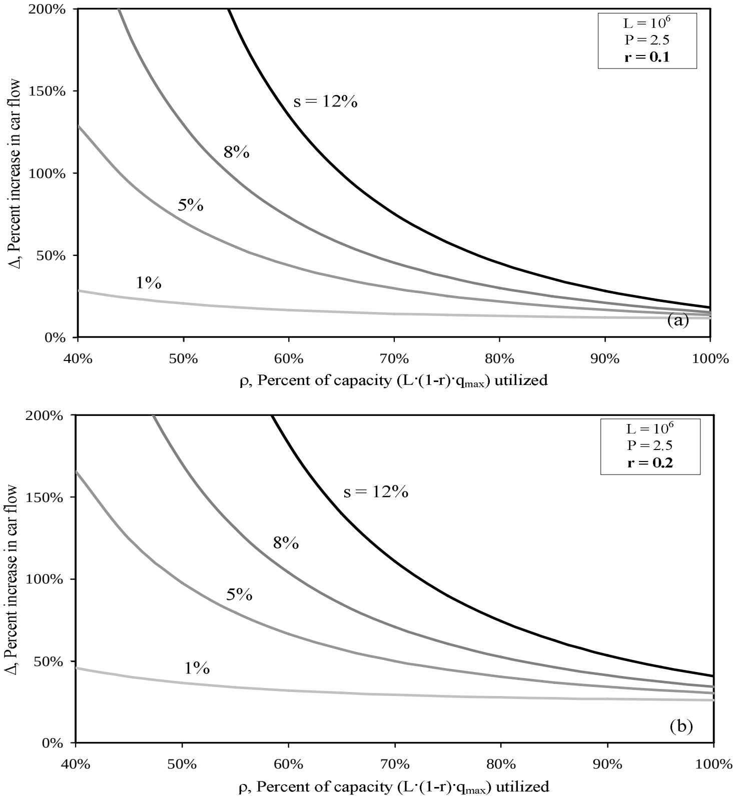

Fig. 10. Curves of ρ vs Δ for a congested city street network. The selected Q has a free flow vehicle speed = 55 km/hr, q_{max} = 1800 cars/hr, and a congested wave speed = 25 km/hr: (a) $r = 0.1$, (b) $r = 0.2$

5. Summary and Conclusions

There are many possible causes of roadway traffic congestion, including accidents, roadwork activity, high merge demands and special-use lanes; and one needs to rule-out all other possibilities before attributing congestion to any one cause. This paper has shown that analysis of time series data alone, without also considering a system's spatial component, will not provide a complete picture of traffic and can produce misleading results. Contrary to an earlier study, we found that carpool lanes are not creating congestion on five freeway sites in the San Francisco Bay Area. These lanes may instead be reducing congestion: spatiotemporal analysis of real data showed that underutilized carpool lanes that run thorough bottlenecks can *increase* the bottleneck discharge flows by smoothing them, as predicted in Menendez and Daganzo (2007). A carpool lane with this desirable property would not only reduce total PHT, but also the PHT among LOVs, and could therefore become a win-win proposition for society.

Since the smoothing effect is at least partly due to vehicular lane changing, it should be stronger when special-use lanes are deployed to segregate vehicle classes with markedly different performance characteristics. Findings from our parametric analyses in sec. 4 are cause for optimism: they reveal that bus-only lanes can not only benefit bus operation, but can also improve car travel, particularly in cases of severe congestion. Field experiments to confirm this are now underway.

6. References

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 Rudjanakanoknad, J. (2005) Increasing the capacity of a merge by metering its on-ramp. *Transpn Res Part B 39*, 896-913.

Cassidy, M.J., Daganzo, C.F., Jang, K. and Chung, K. (2006) Empirical reassessment of traffic operations: Freeway bottlenecks and the case for HOV lanes. *UCB-ITS-RR2006-6,* Institute of Transportation Studies, University of California, Berkeley, CA.

Chen, C., Varaiya, P. and Kwon, J. (2005) An empirical assessment of traffic operations. In: Mahmassani, H.S. editor, Transpn and Traffic Theory, *Procs of 16th International Symp on Transpn and Traffic Theory*.

Daganzo, C.F. (2007) Urban gridlock: Macroscopic modeling and mitigation approaches. *Transpn Res Part B 41*, 49-62.

Daganzo, C.F. and Cassidy, M.J. (2008) Effects of high occupancy vehicle lanes on freeway congestion, *Transpn Res Part B (in press)*.

Highway Capacity Manual (HCM) (2000) Transportation Research Board, National Research Council, Washington, DC, USA

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

Menendez, M. and Daganzo, C.F. (2007) Effects of HOV lanes on freeway bottleneck. *Transpn Res Part B 41*, 809-822.

Munoz, J.C. and Daganzo, C.F. (2002) Fingerprinting traffic from static freeway sensors. *Cooperative Transpn Dynamics 1*, 1-11.