

What Works and What Backfires: Mode Shift, Rerouting and Departure-Time Policy on a Saturated Kathmandu Corridor

Subash Singh Dhama

Department of Software Engineering
Nepal College of Information Technology
Lalitpur, Nepal
subashdhamee@gmail.com

Abstract—Travel-demand management offers several instruments, and which of them a city should use is an empirical question about its demand and its network. This paper tests four on the Tripureshwor–Koteshwor corridor in Kathmandu, using a SUMO testbed, run mesoscopically, whose demand matches 2019 classified counts at 42 locations ($GEH < 5$ at 95.2%), at 24 treatment-assignment seeds per grid point. Two work and two backfire. Removing 15% of motorcycle trips reduces network delay by $55.1 \pm 2.8\%$ and diverting 20% of corridor trips onto alternative routes reduces it by $42.6 \pm 2.5\%$; both raise cordon throughput in the mean. The first is a ceiling rather than an option: the buses needed to carry those riders reclaim as much carriageway as the motorcycles vacate. Retiming departures fifteen minutes earlier raises delay at every share tested, and a school-hours shift raises it at both shares tested. The failures have a measured cause: hourly counts from three Department of Roads stations put the morning peak at 6.8% of daily vehicle traffic against the 10–12% national design practice assumes, so the profile is a plateau and retiming moves departures into hours already as busy. A third result concerns method. The calibrated network’s critical accumulation falls at about half the counted demand, so it has no stable congested regime, and effects read near that boundary depend on the loading they are read at. We report the diagnostic sequence, including a network-build defect in which unmarked six-metre carriageways are encoded as single lanes.

Keywords—travel demand management, traffic simulation Kathmandu, motorcycle-dominant traffic, network calibration

I. INTRODUCTION

Kathmandu Valley’s road network carries a fleet in which motorcycles are the majority mode, on an arterial system that has not expanded in proportion to demand. Proposed responses concentrate on redistributing travel: staggering school and office hours, shifting travellers to public transport, and directing traffic onto alternative routes. Successive valley transport plans have proposed instruments of each kind [1], [2]. Each assumes a property of the demand or the network that has not been measured for this corridor.

This paper tests four such interventions on the Tripureshwor–Koteshwor corridor using a SUMO testbed [3] built from official survey data [4], [5]: departure retiming, mode shift from motorcycle to bus, a school-hours shift, and spatial redistribution onto alternative routes.

The contribution is threefold. First, we rank them on one calibrated network under identical settings. Second, we measure the corridor’s morning demand profile from hourly count records rather than inferring it from person-trip generation, and show that the difference explains why the departure-time instruments fail. Third, we locate the network’s critical accumulation at about half the counted demand, above which it has no stable congested regime and the rerouting effect is no longer measurable. That instability is the congested branch of the network fundamental diagram [24], [25] not anything peculiar to this corridor; what is specific is where the boundary falls.

II. BACKGROUND

Peak-spreading schemes reward travel outside the peak, as in the Dutch Spitsmijden trials [6] and the Bangalore INSTANT pilot [7]. Their theoretical basis is the bottleneck model [8], in which travellers trade schedule delay against queueing delay. That schedules can be moved, however, is not evidence that moving them relieves congestion, and the recent record is mixed. A comparative study of Paris and San Francisco associates flexible working hours with a *higher* probability of peak-period arrival in Paris and with no significant effect in San Francisco [17]; an ex-ante evaluation of staggered work hours for the Paris region bounds the relief at a fifth of total crowding cost [18]; and a randomised controlled trial of information and incentives found no significant difference between treatment arms and control [19]. School travel is a recurring target because it is concentrated and administratively movable.

How much any of this can achieve scales with how concentrated the peak is: a sharp peak leaves room to move demand into adjacent hours, a flat profile does not. For Kathmandu no published corridor-level hourly vehicle profile was located; studies of the corridor’s intersections report saturation and delay [9]–[11] without the underlying temporal distribution of arrivals.

Directing travellers onto alternative paths raises the classical question of whether user-optimal routing improves system performance [12], and the possibility that reassigned capacity

degrades it [13]. The result extends to information itself, which can raise equilibrium travel times when only part of the population receives it [20]; simulation of navigation-app uptake correspondingly finds network performance improving over a band of partial adoption rather than at full compliance [21]. Whether alternatives help depends on their capacity relative to the corridor, which is an empirical property of the network.

Kathmandu's traffic operates without lane discipline: motor-cycles filter through gaps, right-of-way at minor approaches is negotiated rather than assigned, and several vehicles occupy the width of a nominal lane [22]. Mid-block speeds are governed as much by roadside friction, parking and pedestrian encroachment as by the traffic stream [14], and capacity in such streams is conventionally referred to carriageway width rather than to marked lanes [26]. Calibrating a lane-based simulator to this traffic is established but not routine: SUMO has been calibrated to heterogeneous flow in Khulna [23], where the simulator's default parameters had to be replaced with locally observed ones. SUMO offers a sublane model for lateral freedom within a lane [15].

At network scale, production rises with accumulation to a critical point and falls beyond it [25], and the congested branch is unstable where congestion distributes unevenly, so such networks gridlock instead of settling [24]. Section IV-D reports where that boundary falls for this network, and it falls low.

III. METHOD

A. Network

The study corridor runs from Tripureshwor to Koteshwor, including the Maitighar, Thapathali, New Baneshwor, Tinkune and Jadibuti junctions. The network was built from an OpenStreetMap extract [16] with netconvert, restricted to the largest connected component reachable by passenger vehicles.

Lane counts were derived from carriageway width rather than from the `lanes` tag. The corridor's arterials carry no lane markings, so OpenStreetMap records them as single-lane and places the carriageway in a separate `width` tag: Ram Shah Path is recorded as a one-way primary road with one lane and a width of six metres. Both tags are accurate. Taking the lane tag at face value builds one 3.2 m lane out of six metres of road. Width is present on 1,578 road ways in the extract against 635 for lanes, so lane count is taken as carriageway width divided by 3.0 m per traffic stream, halved between directions on two-way links.

The 3.0 m divisor is the lower of the national standard's own two values. Nepal Urban Road Standard 2076 Table 17 [27] gives 3.5 m per lane for arterial roads and 3 m for local streets; the lower figure is taken because what is being divided is an unmarked carriageway carrying non-lane-disciplined traffic, not a marked arterial lane. Deriving capacity from carriageway width rather than from markings is established practice for mixed traffic [26]. Width is divided and rounded to the nearest whole stream, so for the corridor's binding link the choice is

immaterial: a 6 m one-way carriageway yields two lanes at either value.

B. Data sources

Three datasets underlie the model, and their sampling differs.

Origin–destination structure. Zone-to-zone vehicle matrices from the 2012 valley-wide survey [4], digitised from the published tables. These fix the spatial pattern of demand but not its level, which is dated; they are scaled forward as described below.

Calibration targets. Classified counts at nine intersections along the corridor from the 2019 survey [5], giving 77 leg-direction records. Each leg was observed for 14.5 to 15.0 hours and reported as a 24-hour passenger car unit total. One of the nine, Tinkune, is surveyed as three sub-sites, so the published tables carry eleven site labels. Calibration uses the 38 inbound legs, since outbound legs would count the same vehicles twice at the same junction; those 38 legs map onto 42 distinct network edges because four of them are divided carriageways that split across parallel edges. Those 42 edges are the calibration targets. Counts are published in passenger car units. The 2019 report gives the fleet mix at these junctions but not the factors it used, so conversion to vehicle counts applies the 2012 survey's factors [4] at the reported mix.

Temporal profile. Hourly counts from three Department of Roads permanent stations: 64 Manohara Bridge, on the corridor's own road at the Jadibuti end; 65 Ring Road (Sina-mangal); and 58 Satdobato South. Each station was surveyed on three consecutive days in FY 2024/25 over all 24 hours, giving 216 station-day-hours [28].

Two limitations bound what that profile can support. All three stations are highway or Ring-Road cross-sections, not corridor-interior, and the counts used are both-direction totals, which flatten a tidal profile; the per-direction columns exist on the same pages and are the registered sensitivity. A corridor-interior morning peak could therefore be sharper than the pooled 6.8%, by an amount this study does not measure. The profile is used only to distribute demand across hours, not to set its level.

C. Demand construction and calibration

The 2012 matrices were scaled to the study year using Department of Roads station growth records, cut at the count cordon, and distributed over the morning using the measured hourly profile. Routes were generated under randomised edge weights so that travellers between the same origin and destination do not all receive the same path, and the route set was then sampled to match the counted volumes.

Calibration is assessed with the GEH statistic at the hourly scale it is defined for, against the 42 counted edges. The demand used throughout this paper matches the counted volumes at $GEH < 5$ on 95.2% of those locations. Two qualifications belong with that figure. It is measured on the sampled demand as written, before simulation, so it states that the trips loaded onto the network reproduce the counts, not that the flows the

network delivers do; Section IV-D reports the latter, and it is much worse. And the hourly targets are the 24-hour counts distributed by the profile of Section IV-A, so the same profile appears on both sides of the comparison: the calibration does not test the profile, it assumes it.

D. Scenarios and metrics

Four demand transforms were applied to the calibrated route set. Each is a pure function of the route file, so network, physics and loading are identical across arms and any difference comes from the demand change under test.

Departure retiming shifts a seeded random share p_t of trips departing in the peak hour to fifteen minutes earlier. Only the departure time changes; the route and vehicle type are untouched.

Mode shift converts a share m of motorcycle trips on an origin–destination pair into bus trips, adding a bus once the converted trips on that pair reach a busload within the analysis window.

School-hours shift moves a seeded random share of peak-hour trips one hour earlier, standing in for the 2020 proposal to advance school start times from 10:00 to 09:00. It is a rigid translation of a cohort, not a stagger across a range, and although the share comes from school travel’s share of person-trips the transform does not identify school trips or restrict itself to them.

Spatial redistribution diverts a seeded share p_r of corridor trips onto alternative paths, found by shortest-path search with the primary path excluded and searched under the vehicle’s own class, since a route legal for a motorcycle may be illegal for a bus.

Two outcomes are reported over the analysis window: D_{net} , total network time loss against desired speed in vehicle-hours, standing for network delay; and H , cordon throughput, the passenger-car-unit flow across the counted edges. Each grid point was run at 24 treatment-assignment seeds, which vary which trips are treated rather than which demand realisation is drawn.

IV. RESULTS

Fig. 1 summarises the four interventions on one axis. The subsections below take each in turn.

A. The morning peak is a plateau

Hourly counts from three Department of Roads stations across three survey days each [28], 216 station-day-hours in total, give shares of daily vehicle traffic of 3.8%, 4.7%, 5.5%, 6.8%, 6.7% and 6.4% for the hours 06 to 11. The busiest morning hour carries 6.8%.

That is well below either available proxy. Where no hourly vehicle counts exist, a peak share is commonly inferred from survey person-trip generation, and the 2012 Kathmandu survey reports a morning generation peak near 20% of daily person-trips [4]. Trip generation counts when people begin journeys, walking and school trips included, so using it as a vehicle-departure share overstates this corridor’s peak by about a factor

TABLE I
CHANGE AGAINST THE UNTREATED BASELINE AT 0.55 LOADING, 24 TREATMENT-ASSIGNMENT SEEDS PER ROW, MEAN $\pm$ STANDARD ERROR. D_{net} IS NETWORK TIME LOSS AND H IS CORDON THROUGHPUT; A LEVER WORKS WHEN IT REDUCES THE FIRST AND RAISES THE SECOND. ALL RETIMING GRID POINTS ARE SHOWN.

Lever	Share	ΔD_{net} (%)	ΔH (%)
Rerouting	5%	-14.4 ± 3.1	+2.4
	10%	-25.7 ± 3.1	+3.7
	20%	-42.6 ± 2.5	+5.8
Mode shift	5%	-13.4 ± 3.8	+1.8
	10%	-38.9 ± 3.4	+5.7
	15%	-55.1 ± 2.8	+7.9
Retiming	5%	$+7.9 \pm 3.1$	−1.0
	10%	$+15.7 \pm 4.7$	−2.5
	15%	$+9.9 \pm 4.0$	−1.4
	20%	$+26.3 \pm 4.6$	−3.8
	25%	$+36.1 \pm 5.1$	−6.0
School	25%	$+26 \pm 16$	−1.2
	46%	$+319.5 \pm 4.7$	−38.1

of three. The second proxy, national design practice, offers 10–12% of daily flow [27], still nearly twice the measured value. Neither substitutes for counted hourly flow on the corridor itself.

The profile is close to flat across 09:00–11:00 (Fig. 2). Departure-time redistribution moves demand into adjacent hours that are almost as busy.

B. Rerouting and mode shift both work

Alternative paths were found for essentially every corridor origin and destination pair when the primary path was excluded. Diverting a share of corridor trips onto them reduces network delay and raises cordon throughput, monotonically in the share diverted (Table I). Shifting motorcycle trips to bus reduces delay more than any other intervention tested.

Mode shift is the only lever that removes vehicles rather than rearranging them, and it is the strongest. Its caveat is capacity, and it is severe enough to change what the number means. The transform adds a bus only where a single origin–destination pair accumulates a busload inside the analysis window; across the corridor’s 51,612 pairs none does, so at every share tested it adds no buses at all. The run measures the removal of 1.8, 3.6 and 5.3% of vehicles, with the person-trips they carried leaving the network without boarding anything. Section V-A puts the buses back and finds they reclaim the carriageway the motorcycles vacate, so the -55.1% is a ceiling: what removing that share of motorcycles achieves, not what a bus service able to replace them would.

Retiming raises delay at every share tested at a fifteen-minute shift, and a 46% school-hours shift raises it by more than a factor of three. The two mechanisms differ, and both follow from the measured profile.

The school transform moves a share of peak-hour trips one hour earlier, into 08:00, which already carries 5.5% of daily traffic against the peak hour’s 6.8%. Shifting 25% of the peak hour produces a new maximum of 7.20% at 08:00, 5.9% above the peak it replaced; shifting 46% produces 8.63%, which is




Fig. 1. Change in network delay against the untreated baseline, by share of trips treated. Bars are the standard error of the mean over 24 treatment-assignment seeds. The school-hours shift is on its own axis.




Fig. 2. Measured share of daily vehicle traffic by clock hour, against the share implied by treating person-trip generation as vehicle departures and the 10–12% national design practice assumes.

26.9% above it. The intervention does not flatten the peak, it relocates and enlarges it.

Retiming moves departures fifteen minutes earlier, and the demand file's finest structure is a fifteen-minute slice, so most treated trips stay within the peak hour and advance by one slice. Its adverse result is a within-hour loading effect, not the relocation the school shift produces. A second arm separates the two: moving departures thirty minutes rather than fifteen reduces delay by $8.2 \pm 3.8\%$ at a 15% share ($p = 0.04$), and turns adverse again at larger shares, $45.3 \pm 7.8\%$ at 25%. Moving departures far enough to clear the hour can help; moving them within it does not.

That distinction is what separates the two rows. At 25% the relocated peak is close enough to the original that the effect is not separable from zero. At 46% the network is asked to carry a morning maximum roughly a quarter higher than the one it was calibrated against, which is enough to move it past the loading threshold described in Section IV-D, where delivered flow falls instead of levelling off. The +319% figure

TABLE II
DELAY RESPONSE PER ONE PERCENT OF TRIPS TREATED, ORDINARY LEAST SQUARES OVER EVERY TREATED GRID POINT AND SEED. NEGATIVE IS A REDUCTION. R^2 IS COMPARABLE ACROSS ROWS ONLY TO THE EXTENT THE DOSE RANGES ARE, WHICH THEY ARE NOT; THE SPREAD WITHIN A DOSE IS GIVEN IN THE TEXT INSTEAD.

Lever	Slope (pp)	95% CI	R^2
Mode shift	-4.18	[-5.12, -3.23]	0.53
Rerouting	-1.85	[-2.39, -1.32]	0.41
Retiming	+1.34	[+0.79, +1.89]	0.16
School shift	+13.98	[+12.41, +15.54]	0.88

is therefore a consequence of the regime boundary, not a linear extrapolation of the 25% result.

C. Dose–response and significance

Each grid point in Table I is a mean over 24 independent treatment assignments, so the effects can be tested rather than read off. Against no change, every point is significant at $p < 0.05$ except the 25% school shift ($t = 1.66$, $p = 0.11$). Holm correction across the family of thirteen grid-point tests leaves that one non-significant and moves the 5% and 15% retiming points to $p = 0.054$; the rest are unaffected. The 25% school shift is not a small adverse effect but an undetermined one: its median across assignments is -8.8% and 13 of the 24 assignments reduce delay, so the sign is not established at that share.

Fitting the delay change against the share treated gives a slope per lever (Table II), which quantifies the response rather than tabulating it. Ordering is unchanged, but the magnitudes separate: mode shift returns about 2.3 times the delay reduction per unit of demand treated that rerouting does. The school row is fitted through two doses and is given for completeness; Section IV-D is why its second point should not be read as a linear extrapolation of its first.

How much *which* trips are treated matters is better read from the spread within a single dose than from $1 - R^2$, which also absorbs the width of the dose range a lever was swept over.

TABLE III

DELIVERED FLOW AS A SHARE OF COUNTED VOLUME, BY LOADING. D_{net} IS NETWORK TIME LOSS AND <1 M/S THE MEAN SHARE OF TRAFFICKED EDGES BELOW WALKING PACE, OVER THE 08:00–10:00 WINDOW. ALL ROWS ARE ONE BATCH.

Loading	08:00	09:00	10:00	D_{net}	<1 m/s
50%	43.3%	47.9%	42.2%	3,054	0.6%
55%	47.2%	50.5%	32.5%	7,577	1.8%
60%	51.4%	46.6%	16.5%	15,542	4.8%
65%	55.6%	48.1%	13.8%	18,479	5.5%
70%	59.8%	43.3%	8.5%	24,582	7.7%
100%	56.0%	10.5%	8.3%	65,592	20.2%

Pooled across doses, that spread is 14.3 percentage points for rerouting, 16.3 for mode shift, 21.3 for retiming and 56.6 for the school shift. Assignment matters least for rerouting and most for the levers with the largest effects. A scheme that could select which travellers divert is still a different instrument from one broadcast to all of them, and this study measures the latter.

These slopes describe this corridor's response under one network, one demand and one loading. The seeds randomise treatment assignment, not the corridor, so the intervals bound the effect of assignment and not the transferability of the result to other networks.

They are also conditioned on a single draw of the simulator's own random number generator, and that draw matters. Replicating the untreated baseline under ten simulator seeds gives network time loss from 6,355 to 8,737 veh h, a coefficient of variation of 11.2%; the single realisation every delta in Table I is measured against sits 5% above that mean. Repeating the 20% rerouting arm at five simulator seeds, each paired against the baseline at its own seed, reduces delay in every one: -15.1 , -28.5 , -32.5 , -44.5 and -46.4% . The direction of the headline result holds under simulator randomness; its magnitude is far less precisely determined than the ± 2.5 of Table I, which bounds assignment variance alone. Cordon throughput over the same five draws runs from -0.4 to $+7.0\%$, so the throughput half of the claim is the weaker one.

D. The model has no stable congested regime

Scenario results depend on the loading at which the network is run (Table III). Total delivered flow across the morning peaks at 50% loading (152,907 cordon crossings) and drops by roughly a sixth once the tipping point is passed: 125,684 at 60%, 130,060 at 65% and 126,185 at 70%. Past the tipping point the level fluctuates within run-to-run variation, not monotonically. Below the tipping point the network runs close to free flow, with 0.6% of trafficked edges below 1 m/s; at 60% and above it gridlocks within the analysis window, reaching 20.2% at full counted demand.

Run at a loading it can carry, the network moves about 44,000 veh/h through the count locations against a counted peak-hour 92,073. Part of that gap is not a property of the model. Ram Shah Path, the 6 m one-way arterial of Section III-A, carries a peak demand of 4,563 veh/h, about 4,300 PCU/h, against the 1,200 to 2,400 PCU/h Nepal Urban Road Standard 2076 Table 5 [27] allows a two-lane one-way urban

carriageway. A PCU-based capacity derived for a car-dominant fleet is a poor instrument for a stream that is 70% motorcycles, so this bounds the gap without explaining it.

The absence of a congested-but-flowing band is the stronger limitation. Real networks spend the peak in that regime, and it is the regime demand-management measures act on. At full counted demand the rerouting arm gave delay changes of $+0.15$, $+0.16$ and $+0.31\%$, indistinguishable from zero against a treatment-assignment spread of 12 to 15 points, and read in an earlier draft as evidence that rerouting does not work on this corridor.

Scenario results are reported at 0.55 loading: the most congested setting at which the 08:00 and 09:00 hours still gain delivered flow. The 10:00 hour is already losing it there, so the operating point sits on the shoulder of the regime, not safely inside it.

E. Diagnostic sequence

Five explanations for the network's failure to carry its calibrated demand were tested and rejected before the lane-count defect was identified. *Traffic signals*: disabling all 39 reproduced the baseline exactly (92,319 vehicles inserted, 54,881 running, 89,959 waiting), because mesoscopic junction control is off by default and the actuated programs were never applied. *Junction control*: enabling it made throughput worse, 23.8% delivered at 07:00 against 56.8%. *Network storage*: the network holds 281,248 vehicles at jam density and peaked at 19.5% occupancy, so it was never globally full. *Route concentration*: all 51,612 origin–destination pairs had exactly one route, and randomising edge weights spread traffic over 1,951 additional edges while moving peak-hour delivery only from 6.6% to 6.7%, though the concentration was real, the busiest edge carrying 6,067 veh/h against a counted 6,313. *Simulation mode*: microscopic simulation with the sublane model was worse than mesoscopic at every hour, 14.4% delivered at 07:00 against 58.9%.

That last result is explained by the lane-count defect. Sublane gives vehicles lateral freedom inside a lane, and the lane was 3.2 m wide, so there was nothing to filter into. With carriageway-derived lanes, peak-hour delivery rose from 6.7% to 10.5% and the 08:00 hour from 29.4% to 56.0% (Fig. 3).

V. DISCUSSION

The measured peak shape bears on proposals to stagger working or school hours. Such schemes redistribute departures out of the busiest hour, and that hour carries 6.8% of daily traffic against 6.7% in the next. The instrument and the problem are mismatched, which is a property of the demand, not of any scheme's design.

Studies building simulation networks for unmarked urban carriageways should check that lane counts reflect carriageway width rather than markings, and that the model has a stable congested regime at the loading used: results past a gridlock threshold measure distance from the threshold, not the intervention.




Fig. 3. Delivered flow as a share of counted volume, with lane counts taken from lane markings and from carriageway width. The gap is the network-build defect of Section III-A.

A. What this implies for Kathmandu

Two of the four measures reduce delay here, but they are not equally available. Mode shift is the stronger lever and the harder one to use. The transform assumes buses exist to absorb the shifted trips, and the corridor's public transport is already constrained at exactly the hour the shift would load it: a passenger survey reports that 30.5% of users wait more than ten minutes for a ride during the morning peak, and 61.7% of female respondents report discomfort from overcrowding [30]. Valley transport plans have proposed capacity expansion for more than a decade [1], [2], [4], [29].

The size of that constraint can be stated. At the 15% share the shift strands about 8,800 morning person-trips, some 585 bus trips through the corridor in the morning peak alone. At the survey's equivalency factors those buses occupy about as much carriageway as the motorcycles they replace ($-2,395$ PCU against $+2,342$), so the road-space gain the run reports depends on a service that does not appear in it. Whether a real service nets a gain turns on bus occupancy and on the equivalency factors assumed, neither of which this study varies.

Rerouting has the opposite profile, and it is well supported here: at a 20% diverted share corridor travel time falls by $9.6 \pm 1.6\%$ eastbound and $3.6 \pm 0.9\%$ westbound alongside the network-level reduction, so the diverted driver and the network gain together, though at 5% neither direction separates from zero. It is also the lowest-capital of the four, since it builds nothing and the alternatives already exist, though route guidance still requires an information channel this study does not cost. Its constraint is behavioural: the model diverts a chosen share, whereas a real scheme can only inform. Outcomes at a fixed share vary across assignments by 14.3 percentage points,

so which travellers divert matters alongside how many, and neither compliance nor the selection mechanism is established here.

Three recommendations follow. First, advancing school start times uniformly by an hour, as was proposed for the valley in 2020, is not supported on congestion grounds for this corridor: the measured profile shows it relocates the peak rather than flattening it. This is a result about a rigid one-hour advance, not about staggering in general, which distributes starts across a range and is untested here. Second, any peak-spreading proposal for a Kathmandu corridor should be preceded by hourly counts on that corridor, since both proxies in common use overstate its peak. Third, the first thing worth establishing about route guidance is what share of drivers would follow a route they did not already know.

B. Limitations

Absolute delay levels are not offered as field predictions. The network delivers about half the counted peak throughput at its ceiling, and that gap was not tested against an alternative simulator.

Every reported run is mesoscopic, where SUMO does not apply junction control: the corridor's 39 signal programs are present in the network and inactive in the results. That matters here because the survey evidence identifies intersection saturation as this corridor's binding constraint, and Section IV-E shows enabling control changes throughput sharply. The microscopic comparison there predates the lane-count correction and was not repeated on the rebuilt network, so it is not evidence about microscopic simulation on this corridor.

The model has no departure-time choice, no trip abandonment and fixed routes outside the treated share, so it measures a one-shot loading response, not an equilibrium. Effects are read at a single loading on a steep part of the network's response, where a 10% demand increase raises delay by 148% between 0.50 and 0.55, and are not separated from what an equivalent change in loading alone would produce. Sweeping loading with no intervention and comparing each lever against that curve is the test this design most needs, and is the registered next step.

The 3.0m stream width bears directly on capacity and is an assumption rather than a measurement: no Kathmandu measurement of effective stream width under non-lane-based operation was located. A sensitivity sweep at 2.75 and 3.5 m is the registered next step.

Passenger car unit factors follow the survey source [4]. Nepal Urban Road Standard 2076 gives different values, and on a motorcycle-dominant fleet the newer factors would raise the PCU equivalent of the same stream. The two sets were not reconciled here.

Scenario effects vary with which travellers are treated. That spread is a property of the interventions: the effect of retiming a given share of trips depends on which trips comply. At the 25% school shift it is large enough that the sign is undetermined, not merely imprecise: 13 of 24 assignments reduce delay.

The model fuses three vintages: the 2011–12 origin–destination survey for spatial structure, the 2019 counts for level, and an FY 2024/25 hourly profile for temporal distribution. Only the morning is modelled, and the same records put the evening peak higher, 7.4% of daily traffic against the morning’s 6.8%, with a 6.0–7.4% plateau between. The policy readings are therefore morning-only, though an all-day plateau strengthens the argument against peak-spreading instruments.

VI. CONCLUSION

Two of four interventions tested reduce network delay on this corridor and two increase it. Removing motorcycle trips is the strongest and diverting trips onto alternative routes is second, and the delay reduction is the better-determined half of each. A fifteen-minute departure retiming and a school-hours shift raise delay at every share tested, because the corridor’s morning profile is a plateau rather than a peak: instruments that move departures between adjacent hours have nothing to move demand into. The third result is methodological. Critical accumulation falls at about half the counted demand, so the network has no stable congested regime, and effects read near that boundary are sensitive to the loading they are read at.

ACKNOWLEDGMENT

The author used Claude (Anthropic) to draft and edit prose, to write and debug the simulation and analysis code in the accompanying repository, and to assist with the diagnostic procedure reported in Section IV-E. The author reviewed and edited the content and takes full responsibility for it. Data, network build and testbed are available at github.com/mesubash/mod.

REFERENCES

- [1] Japan International Cooperation Agency, “Project on Urban Transport Improvement for Kathmandu Valley: Final Report,” JICA, 2017.
- [2] Asian Development Bank, “Kathmandu Sustainable Urban Transport Project: Completion Report,” ADB, Loan 2656, 2020.
- [3] P. Alvarez Lopez et al., “Microscopic traffic simulation using SUMO,” in *Proc. IEEE Intell. Transp. Syst. Conf. (ITSC)*, 2018.
- [4] Japan International Cooperation Agency, “Data Collection Survey on Traffic Improvement in Kathmandu Valley,” JICA, 2012.
- [5] Japan International Cooperation Agency, “Data Collection Survey on Urban Transport in Kathmandu Valley: Final Report,” JICA, 2019.
- [6] S. Donovan, “Spitsmijden peak-avoidance experiments, Netherlands,” Victoria Transport Policy Institute, 2011.
- [7] D. Merugu, B. Prabhakar, and N. S. Rama, “An incentive mechanism for decongesting the roads: a pilot program in Bangalore,” in *Proc. ACM NetEcon*, 2009.
- [8] Z.-C. Li, H.-J. Huang, and H. Yang, “Fifty years of the bottleneck model: a bibliometric review and future research directions,” *Transp. Res. Part B*, vol. 139, pp. 311–342, 2020.
- [9] K. Timala, A. Marsani, and H. Tiwari, “Impact of traffic bottleneck on urban road: a case study of Maitighar–Tinkune,” in *Proc. IOE Graduate Conf.*, 2017.
- [10] S. Shrestha and A. Marsani, “Saturation flow and delay models for signalised intersections in Kathmandu,” in *Proc. IOE Graduate Conf.*, 2014.
- [11] L. Maharjan and A. Marsani, “Performance evaluation and improvement of an intersection: a case study of Thapathali intersection,” in *Proc. 13th IOE Graduate Conf.*, vol. 13, 2023.
- [12] J. G. Wardrop, “Some theoretical aspects of road traffic research,” *Proc. Inst. Civil Engineers*, vol. 1, no. 3, pp. 325–362, 1952.
- [13] D. Braess, “Über ein Paradoxon aus der Verkehrsplanung,” *Unternehmensforschung*, vol. 12, pp. 258–268, 1968.
- [14] R. Dhimas and A. Marsani, “Analysis of traffic speed on urban mid-block section under the influence of roadside frictions,” in *Proc. 13th IOE Graduate Conf.*, vol. 13, 2023.
- [15] Eclipse SUMO, “SUMO documentation,” v1.27.1, 2026. [Online]. Available: <https://sumo.dlr.de/docs/>
- [16] OpenStreetMap contributors, “OpenStreetMap,” 2026. [Online]. Available: <https://www.openstreetmap.org/>
- [17] E. Munch and L. Proulhac, “Is work hours’ flexibility really a solution to morning peak period congestion? Comparative analysis between Paris and San Francisco,” *J. Transp. Geogr.*, vol. 113, art. 103712, 2023.
- [18] N. Coulombel, E. Munch, and C. Pivano, “Travel demand management: the solution to public transit congestion? An ex-ante evaluation of staggered work hours schemes for the Paris region,” *Transp. Policy*, vol. 137, pp. 48–66, 2023.
- [19] A. Rosenfield, J. P. Attanucci, and J. Zhao, “A randomized controlled trial in travel demand management,” *Transportation*, vol. 47, no. 4, pp. 1907–1932, 2020.
- [20] D. Acemoglu, A. Makhdomi, A. Malekian, and A. Ozdaglar, “Informational Braess’ paradox: the effect of information on traffic congestion,” *Oper. Res.*, vol. 66, no. 4, pp. 893–917, 2018.
- [21] A. R. Bagabaldo, Q. Gan, A. M. Bayen, and M. C. González, “Impact of navigation apps on congestion and spread dynamics on a transportation network,” *Data Sci. Transp.*, vol. 6, no. 2, art. 12, 2024.
- [22] G. Asaithambi, V. Kanagaraj, and T. Toledo, “Driving behaviors: models and challenges for non-lane based mixed traffic,” *Transp. Developing Economies*, vol. 2, no. 2, art. 19, 2016.
- [23] Md. M. H. Chowdhury and T. Chakraborty, “Calibration of SUMO microscopic simulation for heterogeneous traffic condition: the case of the city of Khulna, Bangladesh,” *Transp. Eng.*, vol. 18, art. 100281, 2024.
- [24] C. F. Daganzo, V. V. Gayah, and E. J. Gonzales, “Macroscopic relations of urban traffic variables: bifurcations, multivaluedness and instability,” *Transp. Res. B*, vol. 45, no. 1, pp. 278–288, 2011.
- [25] N. Geroliminis and C. F. Daganzo, “Existence of urban-scale macroscopic fundamental diagrams: some experimental findings,” *Transp. Res. B*, vol. 42, no. 9, pp. 759–770, 2008.
- [26] S. Chandra and U. Kumar, “Effect of lane width on capacity under mixed traffic conditions in India,” *J. Transp. Eng.*, vol. 129, no. 2, pp. 155–160, 2003.
- [27] Government of Nepal, Ministry of Urban Development, “Nepal Urban Road Standard, 2076,” 2020.
- [28] Department of Roads, Nepal, “Traffic count station records,” 2025. [Online]. Available: <https://ssrn.dor.gov.np/>
- [29] UNESCAP, “Comprehensive Public Transport and Mass Transit Plan for Kathmandu Valley,” United Nations ESCAP, 2022.
- [30] Clean Energy Nepal, “Public transportation in Kathmandu Valley,” MaYA Fact Sheet No. 4 (CEN/CANN passenger survey), 2011.