Transportation Research Part C 95 (2018) 248–260

Contents lists available at ScienceDirect

Transportation Research Part C

journal homepage: www.elsevier.com/locate/trc

Scheduling lane conversions for bus use on city-wide scales and in time-varying congested traffic

Nathalie Saade, Jean Doig, Michael J. Cassidy*

Department of Civil and Environmental Engineering and the Institute of Transportation Studies, University of California, Berkeley, CA 94720, United States

ARTICLE INFO

Keywords:
Bus Lanes
Scheduled lane conversion
Road allocation
Public transport
Network-wide modeling

ABSTRACT

The paper explores what can occur when select street lanes throughout a city are periodically reserved for buses. Simulations of an idealized city were performed to that end. The city's time-varying travel demand was studied parametrically. In all cases, queues formed throughout the city during a rush, and dissipated during the off-peak period that followed. Bus lanes were activated all at once across the city, and were eventually deactivated in like fashion. Activation and deactivation schedules varied parametrically as well. Schedules that roughly balanced the trip-time savings to bus riders against the added delays to car travelers were thus identified.

Findings reveal why activating conversions near the start of a rush can degrade travel, both by car and by bus. Balance was struck by instead activating lane conversions nearer the end of the rush, when vehicle accumulation in the city was at or near its maximum. Most of the time savings to bus riders accrued after the conversions had been left in place for only 30 min. Leaving them for longer durations often brought modest additional savings to bus travelers. Yet, the added delays to cars often grew large as a result.

These findings held even when buses garnered high ridership shares. This was the case when lane conversions gradually induced new bus trips among residents who formerly did not travel. It was also true when high ridership was a pre-existing feature of the city. Activating conversions a bit earlier in a rush was found to make sense only if commuters shifted from cars to buses in very large numbers. Findings also unveiled how to fine-tune activation and deactivation schedules to suit a city's congestion level. Guidelines for scheduling conversions in real settings are furnished. So is discussion on how these schedules might be adapted to daily variations in city-wide traffic states. Roles for technology are discussed as well.

1. Introduction

Where unused space is in short supply, cities often convert regular street lanes to exclusive bus lanes during select times of day. By traveling in converted lanes, buses can bypass car queues. A converted lane will reduce the street's car-carrying capacity, however. The challenge thus lies in converting lanes in ways that strike a balance between the benefits extended to buses, and the damage done to cars.

Cities presently strike that balance by limiting the spatial extents over which conversions are deployed; i.e. lanes are converted on a few select streets, rather than throughout a city in wholesale fashion (e.g. Zheng et al., 2017). Yet as the world's population marches

* Corresponding author.
E-mail address: cassidy@ce.berkeley.edu (M.J. Cassidy).

https://doi.org/10.1016/j.trc.2018.07.025
Received 5 January 2018; Received in revised form 24 July 2018; Accepted 24 July 2018

toward greater urbanization, crowded cities of the future might benefit from more physically-expansive networks of converted bus lanes. Importantly, some of the technologies needed for these wholesale conversions already exist, while others now loom on the horizon. Our interest lies in scheduling these city-wide conversions in time, and in harnessing technologies to vary these schedules across days in response to daily variations in city-wide traffic conditions.

Current practice is to put conversions in place for pre-defined periods that either span a morning or evening rush, or persist for longer durations that include both rush periods and the hours between (Agrawal et al., 2013). It turns out that neither of the two above-cited periods is appropriate for deploying conversions on large, city-wide scales. Schedules for activating and deactivating lane conversions in this wholesale fashion call instead for fuller consideration of traffic variations within and across days.

Present findings came through simulation, and rely upon the Passenger Hours Traveled (PHT) by bus and by car as the performance metrics. The outcomes show the value of fine-tunning conversion schedules to accommodate a city's level of rush-period congestion. Yet for every congestion level studied, activating lane conversions nearer the end of a rush balanced benefits and costs better than did scheduling activations near the rush's start. Findings were insensitive to the ridership shares that buses garnered when those riders were not pulled from cars. Activating conversions at slightly earlier times during a rush balanced benefits and costs only when commuters had shifted from cars to buses in very large numbers.

The idealized city-street network and our simulation-based methods are described in Section 2. Outcomes from parametric tests, including tests on the bus system's modal share, are furnished in Section 3. Physical explanations are offered as the outcomes are unveiled. Matters pertaining to real-world implementation are discussed in Section 4. This will include discussion on key roles to be played by technology to: (i) monitor city-wide traffic conditions to determine when to activate and deactivate converted bus lanes; (ii) inform motorists and bus drivers of those actions; and (iii) enforce motorist compliance with lane-use restrictions via warnings, citations or the like. A summary of findings and final thoughts on implementation are offered in Section 5. Before proceeding further, we review earlier studies on the subject of large-scale lane conversions.

1.1. Background

The literature includes several pertinent references. Research in Mesbah et al. (2011), for example, explored how to optimally select lanes to be converted across large, city-scale networks. The work assumed that network traffic conditions were uncongested, and therefore always in steady-state.

Research in Gonzales and Daganzo (2012) pursued city-wide designs assuming that lanes could be converted on an intermittent, as-needed basis, meaning that: conversions occurred one link at a time and persisted only for short periods needed to individually serve a buses' passage across that link. The work assumed that city traffic was congested, but with conditions that changed little over time. The rates by which trips were completed on the network therefore changed slowly, and depended only on the network's vehicle accumulations.

To our knowledge, the only previous attempt to design city-wide conversions in non-steady-state traffic appears in Zheng and Geroliminis (2013). It formulated analytical models to determine not only which lanes should be converted during a 1-h rush, but also when to schedule activations and deactivations. The models used a Network Exit Function (NEF) to estimate trip-completion rates on a street network, as a function of the network's vehicle accumulations (only). Yet an NEF assumes that network traffic conditions are in steady-state. Thus, the fractions of trips on the network that have only recently begun, that are near their mid-points, and that will soon be completed are all assumed to remain fixed over time.

To appreciate why use of an NEF is problematic in non-steady-state traffic, consider a formerly-empty street network that is suddenly loaded with traffic. All vehicles in the early-loading stage will have only recently started their trips. An NEF does not account for this disproportionately-high fraction of new trips, and will therefore over-estimate trip-completion rates in the early going. When a network is instead emptying itself (and no new vehicles arrive), an NEF would under-estimate trip completions. Similar prediction problems occur, albeit to lesser degree, in more ordinary non-steady-state conditions, when traffic rapidly transitions from lower to higher (or from higher to lower) densities.

Further discussion on this concern can be found in on pp. 51–53 of Daganzo (2007). The proposed approach to address the concern is described below.

2. Methods

Simulation was used in the present work as means to reliably estimate trip completions in non-steady-state, even fast-changing traffic. The model was developed in-house to track all trips individually. A trip was completed when the traveler arrived at her destination. A street network described in Section 2.1 was modeled as an aspatial queueing system. Macroscopic Fundamental Diagrams (MFDs) shown in Fig. 1 were used to that end. The simulation model was therefore a hybrid one in which traffic was represented on both the micro- and macro-levels; i.e. the model kept tabs on each traveler's individual trip, without worrying about specific routes traveled, vehicular interactions and various other details that might have otherwise characterized each of those trips. The macro-level approximations enhance computational efficiency. They also enhance the generality of our findings, which should qualitatively hold for most any street network where traffic can be described by MFDs.[1]

Importantly, the degrees of detail used in modeling bus and car travel were distinct. For cars, the use of MFDs alone sufficed, as

[1] This same hybrid approach was adopted in Daganzo and Lehe (2015) to develop congestion-pricing schemes in non-steady-state traffic.