A Scalable Framework for Public Transit Assignment

Prateek Agarwal[1], Tarun Rambha[1,2], Dhanus M Lal[3], and Ayush Kumar[4]

[1]Department of Civil Engineering, Indian Institute of Science
[2]Center for infrastructure, Sustainable Transportation and Urban Planning (C*i*STUP), Indian Institute of Science
[3]Department of Mathematics, Indian Institute of Science
[4]Department of Computer Science and Engineering, Manipal Institute of Technology

1 Introduction

Public transit systems form a vital component of urban mobility. They offer a sustainable, equitable, and economical solution to transporting people. For transit to be an attractive option for travel and commuting, operators constantly aim to improve service quality. However, the benefits of such supply-side improvements heavily depend on inherently complex user behavior. Unlike personal mobility users, transit travelers typically use several criteria to choose their routes/journeys. Example criteria include in-vehicle travel time, waiting duration, walking time, cost, number of transfers, type of service (metro vs. buses), and crowding/comfort. Operators, on the other hand, have to decide routes, schedules, and fares (or modifications to existing ones) to improve key performance indicators while accounting for the changes in passenger choices that their actions may trigger. These choices are typically made in an ad-hoc manner, particularly in data-sparse environments such as cities in developing countries where smart cards and real-time GPS data are less prevalent. These cities are also plagued by crowding and reliability issues. A broad question of interest in this context is — "How can operators make better decisions to improve public transit?". This research aims to answer a critical facet of this question — the transit assignment problem.

The transit assignment problem predicts travelers' route choices for a given supply configuration in a transit network. While these choices can be observed using technological solutions such as onboard cameras and mobile apps, answers to "what-if" questions remain elusive. The problem requires addressing several issues summarized below, which are elaborated on in the next section.

- Given a set of transit schedules and network characteristics, can we determine what routes will be chosen by a particular individual?
- How do collective selfish choices of agents translate to flows and crowding in transit systems?
- What actions can an operator take to improve service levels, ridership, and revenue?

The assignment problem has been widely studied as network games in road routing, and solution concepts and algorithms that settle at a Nash equilibrium have been proposed. A related social optimum problem attempts to distribute flows to minimize system-level objectives. While similar ideas are in transit literature in two varieties, frequency-based (Schmöcker et al., 2011) and schedule-based assignment (Hamdouch and Lawphongpanich, 2008), they have been relatively less successful in terms of accuracy, scale, and computational efficiency. (Literature also points to Braess-like paradoxes where an increase in transit capacity can worsen system performance (Jiang and Szeto, 2016).) Frequency-based models are helpful for smaller subsets of routes where wait times for transfers are not critical. Scheduled-based approaches, on the other hand, have gained traction recently with increased data availability. A key difference in road routing and transit is that the supply exhibits a staccato-like time dependency. Capacity between stops exists only when buses pass through them. Further, transit-capacity constraints are notoriously hard to model. A passenger may choose a journey with a transfer, but the chances of boarding the transfer bus or finding a seat in it may be affected by passenger trips that originate at other locations in the future. Bus capacities are also small in magnitude compared to roadway capacities, making it challenging to formulate continuous mean-field approximations that can be cast as convex optimization problems. While there are a few variational inequality approximation models, they remain highly intractable for medium to large networks (Hamdouch et al., 2011). The proposed work addresses the gaps in this classic problem by taking advantage of recent advances in transit routing.

The full gamut of public transportation planning necessitates understanding the role of demand elasticity that supply changes may induce, departure time choices, the effect of new-age mobility solutions such as ride-sharing and hailing services, impacts of real-time information, bundled multi-modal mobility options, etc. While important, these problems are excluded from the current scope because the algorithms for transit assignment are not computationally quick and mature enough to embed them within other larger mathematical models. This research will serve as a conduit for addressing such problems in the future.

2 Methodology

The first step of the proposed transit assignment model involves predicting the route choices of a single transit traveler. To this end, we use data from public sources and the Bengaluru Metropolitan Transport Corporation (BMTC), one of India's largest public transportation systems with over 6000 buses, 5000 routes, and 27,000 trips. This information is fused with land-use data to build a multinominal logit model that determines passenger preferences and weights for different journey attributes such as time, cost, and transfers. However, the application of these models requires a choice set. Existing literature uses simple heuristics that do not discover journeys that passengers might find on routing apps such as Google Maps. They are also computationally inefficient since they build time-expanded graphs and run a variant of Dijkstra's algorithm. Instead, we use state-of-the-art transit routing algorithms such as the multi-criteria RAPTOR (and McRAPTOR) Delling et al. (2015) to generate a set of Pareto-optimal journeys that form the choice set.

The basic version of RAPTOR finds Pareto-optimal journeys based on two criteria: arrival time and the number of transfers. It does not involve any preprocessing stage and works directly on the timetable network by storing the transit timetable in the form of trips. RAPTOR works in rounds, where each round is divided into two phases: a *route scan* and a *transfer scan*. The k^{th} round computes arrival time at reachable stops with exactly $k-1$ transfers. Labels at every stop (initialized to infinity) are of the form: $(\tau_0(s), \tau_1(s), \ldots, \tau_k(s))$, where $\tau_k(s)$ denotes the earliest arrival time at stop s with k trips. A round k involves the following two phases:

1. *route scan* phase: This phase starts by collecting all routes serving the stops whose labels were updated in the previous round (known as marked stops). Next, labels of the stops along the route are updated (if possible) using the earliest trip that can be boarded on it. A stop whose label is updated is added to the marked stop list.
2. *transfer scan* phase: This step involves scanning the footpaths (or walking links) originating from the stops present in the marked stop list. Stops whose arrival time is improved using these footpath connections are added to the marked stop list.

Likewise, Round $k+1$ starts its route scan phase using the marked stop list obtained from the transfer scan phase of Round k. The algorithm proceeds similarly and terminates when the maximum transfer limits are reached, or when no stop labels are updated. Stop labels can then generate Pareto-optimal journeys to all the stops using a quick post-processing step.

To find the efficiency frontier of an individual, we use McRAPTOR (an extension of RAPTOR, which can handle other objectives) with the following criteria: arrival time, number of transfers, in-vehicle travel time, walking time, and cost.

To address the second question, we build a digital twin of a transit system by simulating transit travelers, i.e., their origins, destinations, and departure times. To assign them to routes, we create a network loading process in which the logit model is employed to load travelers one by one to journeys on their respective efficiency frontiers. Figure describes the framework for the proposed methodology. McRAPTOR is used for *Choice set generation*. The passenger exits the system if no journey is possible for the given origin, destination, and departure time. Else, an MNL model is employed to select the preferred journey. If the passenger cannot board the preferred journey (due to overcrowding), the departure time is updated, and the choice set is generated again. Otherwise, passengers are assumed to board the bus until the next transfer stop. Operators can then use the post-processing function to answer crowding, congestion, and ridership-related questions.

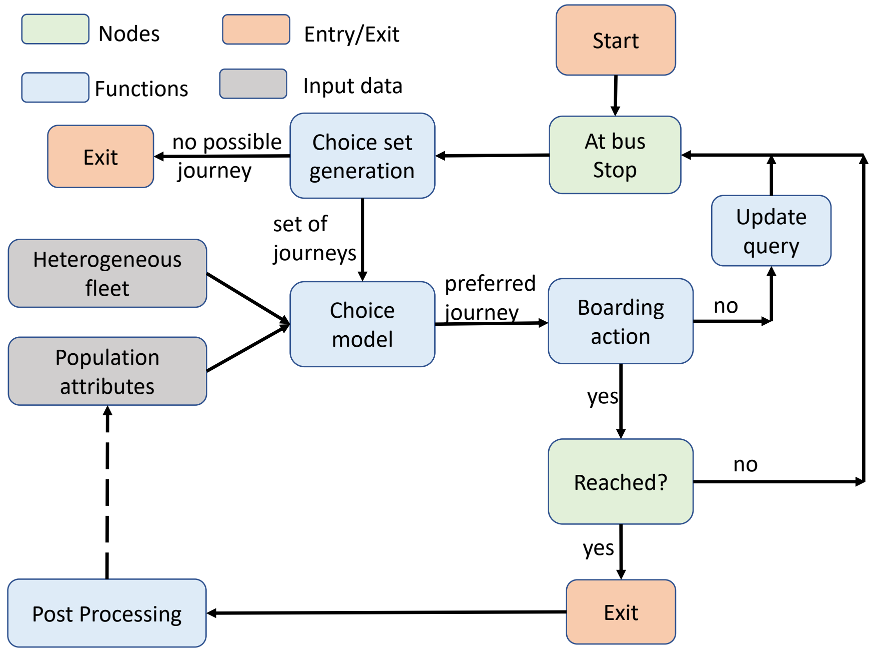

Figure 1: Framework used for scalable public transit assignment

3 Results and Future Work

All algorithms were implemented in C++ (with OpenMP parallelization) and compiled with optimization flags on a 128-core Intel Xeon Gold CPU clocked at 3.0 GHz with 512 GB RAM. The choice set is generated based on five criteria mentioned earlier. The maximum transfer limit was set to three. The input data can be classified into two categories: supply and demand side. Supply data consists of fleet and service characteristics of the transit agency – BMTC. Three types crush loads/capacity were assumed (60, 80, and 100). The demand data (i.e., origin, destination, and departure time) was generated synthetically using data from a four-step model.

Table 1 summarizes the results. The demand column shows the number of passengers simulated in the system. The next column shows the total runtime in seconds. The values in parenthesis show the total time required for choice set generation, choice model, and boarding action. The third column shows the number of iterations. An iteration comprises three steps: choice set generation + choice model application + boarding action. The maximum number of iterations was set to 200. Columns failure and reached show the total number of people who failed to reach and those who reached their destination, respectively. A passenger is said to have failed if the choice set generation results in no journeys (e.g., when no trip is available from the stop after the passenger's departure time). Queries that fail thus indicate the lack of public transit and such passengers are assumed to switch to personalized modes of travel. The last column shows the number of people who have not reached their destination within the limit of maximum iterations due to capacity restrictions . Figure 2 shows a heat map of stops belonging to overcrowded trips in the time intervals 6–7 PM, 7–8 PM, and 8–9 PM, respectively. A trip is classified as overcrowded if the volume-to-capacity ratio of 70% of its link is more than 0.8.

4 Future Work

The proposed framework assumed fixed capacity constraints, and updates routes in an online manner. It however fails to account for day-to-day learning resulting from crowded experiences. To overcome this issue, we plan to initialize subsequent rounds with an additional criterion — a penalty function that reflects the crowding levels. This criterion would prevent overcrowded trip segments from being part of Pareto-optimal journeys and thus would have reduced ridership in subsequent rounds. However, the solutions from this procedure can oscillate. To converge to a

Demand	Runtime	Iterations	Failure	Reached	Denied
100,000	393 (99+159+135)	200	23,598	76,400	2
500,000	1767 (727+586+453)	200	195,678	304,177	145
1,000,000	3520 (1,879+935+713)	200	512,051	487,668	281

Table 1: Assignment Statistics (*Demand*: Number of passengers, *Runtime*: Total runtime (in seconds), *Iterations*: Total number of iterations required, *Failure*: Passengers who failed to board, *Reached*: Passengers who reached their destination, *Denied*: Passengers who are denied boarding and are still in the system)

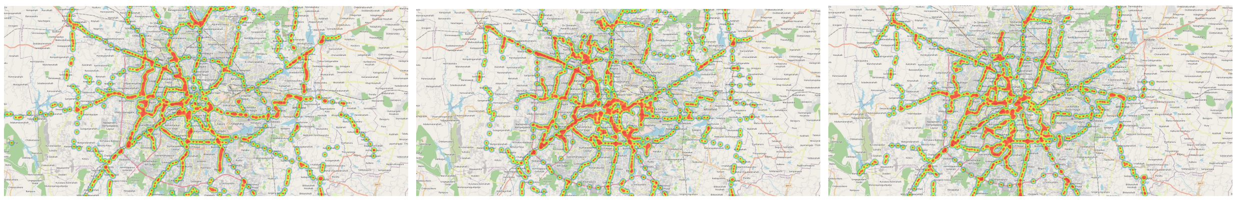

(a) 6–7 PM (b) 7–8 PM (c) 8–9 PM

Figure 2: Stops belonging to overcrowded trips

steady state, we plan to experiment with two methods — (1) use the crowding criteria for only those travelers who were denied boarding or had to spend a significant duration of their journey in crowded buses, and (2) design penalty functions that get "damped" in subsequent iterations akin to step-sizes in stochastic approximation algorithms. The proposed methods will be validated using ticketing data which is available on a subset of route segments. This process will likely produce a feasible solution, but a notion of a stochastic user equilibrium is needed to characterize its properties. (We call it stochastic since a logit probability distribution governs the route choice.) Finally, the model will be tested to predict ridership from changes to routes and schedules (also called route rationalization) and discounted fare strategies using variable neighborhood searches. Since closed-form expressions for the operator's objectives cannot be conceived without making the model overly stylized, we plan to build a simulation-assisted heuristic that perturbs the current supply levels using different exchange/randomization techniques.

Acknowledgments: The authors would like to thank interns – Aditya Shankar Garg and Kalash Abhay Kala for initial assistance with setting up parts of the code.

References

Delling, D., T. Pajor, and R. F. Werneck (2015). Round-Based Public Transit Routing. *Transportation Science 49*(3), 591–604.

Hamdouch, Y., H. Ho, A. Sumalee, and G. Wang (2011). Schedule-based transit assignment model with vehicle capacity and seat availability. *Transportation Research Part B: Methodological 45*(10), 1805–1830.

Hamdouch, Y. and S. Lawphongpanich (2008). Schedule-based transit assignment model with travel strategies and capacity constraints. *Transportation Research Part B: Methodological 42*(7-8), 663–684.

Jiang, Y. and W. Szeto (2016). Reliability-based stochastic transit assignment: Formulations and capacity paradox. *Transportation Research Part B: Methodological 93*, 181–206.

Schmöcker, J.-D., A. Fonzone, H. Shimamoto, F. Kurauchi, and M. G. Bell (2011). Frequency-based transit assignment considering seat capacities. *Transportation Research Part B: Methodological 45*(2), 392–408.