The 11th International Conference on Ambient Systems, Networks and Technologies (ANT)
April 6 - 9, 2020, Warsaw, Poland

Door-to-door transit accessibility using Pareto optimal range queries

Thomas Koch[a,b,*], Luk Knapen[a,b,c], Elenna Dugundji[a,b]

[a]Centrum Wiskunde en Informatica, Science Park 123, Amsterdam 1098 XG, The Netherlands
[b]Vrije Universiteit Amsterdam, De Boelelaan 1111, Amsterdam 1081 HV, The Netherlands
[c]Universiteit Hasselt, Agoralaan building D, 3590 Diepenbeek, Belgium

Abstract

Public transit is a backbone for well-functioning cities, forming a complicated system of interconnecting lines each with their own frequency. Defining accessibility for public transit is just as complicated, as travel times can change every minute depending on location and departure time. With Pareto optimal journeys it is possible to look beyond the earliest arrival times and also optimize for the shortest travel time, as travellers base their departure time on the start time given by their smartphone app, especially when service frequencies are low. By querying for all Pareto optimal journeys in a time range it becomes possible to get a grasp of what passengers see as their choice set when it comes to transit route choice. Based on the averages of the Pareto optimal journeys it should become possible to calculate more realistic skim matrices for traffic analysis zones, including reliability factors such as frequencies and the number of transfers. In this study we calculate Pareto optimal journeys in the area in and around Amsterdam, looking at how travel times are distributed and what factors impact them.

Public transportation is an important travel mode that keeps cities liveable. Determining travel times for public transport alternatives is a more difficult task than for pedestrians, bicyclists and even cars. This has multiple reasons, starting with the fact that public transportation always involves another modality such as walking or cycling, meaning that public transportation accessibility is heavily dependant on the distance from the origin to the nearest location to board a transit vehicle and the distance between the destination and the nearest location to disembark transit. A main reason that analysis of public transportation is difficult is that we are dealing with a time-dependent network, as transit is an intricate system of buses, trains, trams, metros that drive in frequencies that change depending on the time-of-day and are affected by external factors such as traffic and weather. To compensate for these external factors, timetables often include extra time in the travel time and transfer times, so passengers will still be able to make transfers in case of minor delays and the vehicle will just wait in locations where possible and convenient. Finally there are situations where there is trade-off between the access and egress distance and the total travel time. For example, a bus stop next

---

* Corresponding author. Tel.: +31 20 592 4132
  E-mail address: koch@cwi.nl

to the actual point of origin of a journey might be serviced by a bus every hour while a metro station a kilometer away might bring you to your destination every 5 minutes, in half the time; in extreme cases the nearest transit facilities may not even take your destination at all. All-in-all this means that public transportation users take into account a lot more variables that affect mode choice for public transportation including departure time, number of transfers [16], aspects such as reliability [4], transfers to high or low frequency services [3]. In the reported research we consider the following factors:

- total travel time from door to door
- access time to first transit stop, which can be bicycling, walk or a car ride to a Park+Ride location
- egress time to the last transit stop, which can be bicycling, walk or car, and is not necessarily the same modality used to access the public transit network
- number of transfers required
- average adaption time: the average time a passenger needs to wait to catch the next shortest journey to the destination
- perceived service frequency.

Specific to the Dutch context we also include:

- total distance travelled on the transit network, as the fare is based on distance
- fraction of distance travelled on the train network, due to a lack of fare integration.

The results of this research are aimed to be fed into an *activity based travel demand model*. Such models predict the number of trips by generating a daily plan for a set of synthetic individuals. They include underlying choice models for activity type and duration, activity location, travel departure time, transportation mode and sometimes other choices, such as coordination with other persons for activities or travel. In the prediction stage, travel times need to be estimated for several options for the daily schedule being constructed: this is often based traffic analysis zones (TAZ) in order to limit the computational effort. In many cases a skim matrix is computed. Such matrix provides *point estimates* for the inter-zonal travel time. We propose a method that is still TAZ based but takes travel time variability into account. This provides more information for each skim matrix cell.

1. Background

Computing a skim matrix for many traffic zones pairs is computationally very heavy, as for each sampled point in each zone a complete "shortest path" tree has to be constructed from that point to all other points. Due to this computational complexity, initially this problem has been addressed by creating static models of the transportation network, simplifying the network from a complex time-dependant network to a very simple network with static costs between each transit stop. Studies such as Blanchard and Waddell [2], which published their work as an open source Python library, transform the timetable at a specific time-of-day and day-of-week into a simplified graph. Each vertex is a transit stop and each edge is the in-vehicle time derived from the timetable. Modelling the transit network does not model dwell times, boarding and alighting times and traffic congestion nor transfers. The pedestrian network consists out of three components: pedestrian access to and egress from the transit network and pedestrian-pedestrian for uni-modal pedestrian trips. To model the access to transit and egress from transit, each transit stop is connected to the nearest pedestrian node. As not every transit stop is serviced every minute, waiting time is modelled with the assumption that passengers arrive at transit stops randomly, this way the average waiting time can be formulated as the average interval between each departure. Other studies using similar representations are O'Sullivan et al. [14], Curtis and Scheurer [6], Delmelle and Casas [8], Tribby and Zandbergen [20]. This static representation is limited as it does not model the concept of transfers. The number of transfers is important in two ways: it affects the travel time and the reliability of the journey and is an important characteristic that influences mode choice. Hadas and Ranjitkar [11] proposes an extension where the number of required transfers is modelled along with the nature of each transfer: e.g. transfers that do not involve walking, street crossing transfers, same sidewalk transfers and direct trips. In their variants the timetable is still statically defined but in a more detailed manner than just between pairs of stops.

With the increasing availability of data, computational power and improved algorithms, the focus has shifted to even more detailed representations of transit accessibility. Work by Benenson et al. [1], Salonen and Toivonen [17], [19] use transit timetables to give better approximations of total journey duration at a given moment of time and day based on earliest arrival. However these methods are still limited and do not capture two other aspects important

to transit users: the frequency of connections and the wait times involved. In their study [19] calculate all optimal journeys from all transit stops to all other transit stops in the Helsinki study area. Kujala et al. [13] use recent work on transit routing algorithms to perform a range query of Pareto optimal journeys. This list of Pareto optimal journeys describe all journeys departing from a given time period. The simplest version just optimizes on travel time with the tuple $(T_{departure}, T_{arrival})$, where for each advice applies that for each advice 1 and 2 in the list applies that $(T1_{arrival} < T2_{arrival}$ or that $T2_{departure} > T1_{departure}$ allowing for a later departure. A range query will look at all possible departure possibilities within the given time range, excluding departure possibilities where passengers will just have to unnecessarily wait at transfer locations. Typically the number of transfers is also included as a property, allowing for longer travel time as long as it reduces the number of transfers, as an additional transfer might only reduce the travel time by a couple of minutes. For most public transit passengers in the Netherlands this list forms their choice set for transit route choice, as many will use the route suggestions from their favourite public transport app. Kujala et al. [13] use the Connection Scan Algorithm (CSA) by Dibbelt et al. [9], a very simple algorithm that puts all possible connections between transit stops in an ascending departure time order in a single array, to optimize cache locality. This might be counter intuitive as it increases the number of operations, however due to faster access to data stored in the CPU cache the algorithm performs very well or better than algorithms that access the memory locations randomly. Another algorithm that performs well for computing a range of Pareto optimal journeys is the Round-Based Public Transit Optimized Router (RAPTOR) algorithm by Delling et al. [7]. RAPTOR is a dynamic programming algorithm that is organized into rounds, where in each round n the algorithm finds the earliest arrival time for stops reachable in n transfers. By using the transit lines reachable from a transfer in round $n − 1$ or those accessible from the origin, the algorithm runs until no better arrival time is found at any stop. At that point a shortest path tree is complete, with the additional bonus of journeys with a longer travel time but a lower number of transfers.

With this full list of Pareto optimal journeys it is possible to define accessibility beyond just travel times. It is also possible to compute variables such as the service frequency by the number of Pareto optimal journeys within the time range and the average adaption time metric, using the rooftop method defined in Guis and Nijënstein [10]. We take all the Pareto optimal connections between Origin and Destination found using the range query, calculate for each minute in the time range the required adaption time that a passenger has to wait for the next best connection and average the adaption time per minute. This forms a better metric for the travel time that a passenger has to experience including waiting time before and after the journey. The higher the frequency, the higher the probability a journey will fit well to their schedule and thus reduce the time it will take between origin and destination.

2. Case study

To calculate Pareto optimal journeys, we used an existing extension to MATSim[12] developed by the Swiss railways[18]. This library uses the RAPTOR[7] algorithm and the MATSim framework for transit schedules to calculate all Pareto optimal journeys departing in a given time range from and to zones. The library samples origin and destination locations from each zone, to get a realistic view of each zone, as the travel time within each zone has a variance depending on the proximity of the origin/destination to the nearest public transport stop. In this case study we opted for the zoning system defined by Statistics Netherlands (CBS) at the smallest level of *buurt* or *neighbourhood* level [5]. This is the highest resolution level of detail with a significant number of (openly) available statistics such as demographics and income. Descriptive statistics on the size of these zones are listed in table 1, both for the zones exclusively in the municipality of Amsterdam and all selected zones for this study. We limited the study area to the area of the '*Vervoerregio Amsterdam*', the transportation authority in Amsterdam and surrounding municipalities in the greater metropolitan region. For the transit schedules, we used the open data General Transit Feed Specification feed published for the Netherlands.[15]: for performance reasons we removed all services not relevant for our study area; specifically we removed all services not traversing the metropolitan region of Amsterdam. We selected the day with the most services as our study day, a Friday with night bus and train services and we selected a 2 hour time range in the morning rush hour between 07:30 and 09:30 for our range queries. We selected this time period specifically to capture more than one hour to capture hourly frequencies and the specific time period to capture additional bus services in the morning rush hour.

Table 1: Descriptive statistics of area size in used zones: Statistics Netherlands Buurt level. In this study we used the zones of the Amsterdam Transportation authority.

	N	Min	Area size Average	Median	Max
City of Amsterdam	478	$0.015\ km^2$	$0.411\ km^2$	$0.215\ km^2$	$10.121\ km^2$
Amsterdam Transportation authority	802	$0.015\ km^2$	$1.004\ km^2$	$0.340\ km^2$	$26.078\ km^2$

In this case study we made a number of modifications to the MATSim library. Firstly to get a realistic distribution of origin and destinations, we extended the library to load a list of addresses from the *Dutch building registry* in order to sample buildings weighted based on their function and usable surface. This way we get a good sample of both residential locations and prominent locations such as hospitals, schools and retail locations. Secondly we modified the library to write all results between each sampled origin and destination point to disk, so we can build up a high number of observations for each zone pair.

To get a sufficiently high number of samples we ran the computation on the Dutch research compute cluster LISA, where running the query with 20 sample points in 802 zones took approximately 3,5 hours on a 16 core CPU node, including the time to write 1 gigabyte of gzip compressed CSV with the intermediate results to disk. Running the query with 5 sampled points in each zone took 25 minutes for each run and with 10 sampled points took 67 minutes. With these computations times in mind, we conducted 100 runs with 5 points in each zone and 100 runs with 10 points in each zone, to get a good number of runs. To look at the effects of a very high number of points per zone, we conducted 100 runs of 500 points per zone. Due to the quadratic growth of computational complexity and data output, we limited the 500 points per zone run to a group of 16 zones. Additionally to look at the difference between public transport accessibility and car accessibility we used the same library to calculate skim matrices to provide us with comparable data on car accessibility. Finally we calculated a single run to produce a skim matrix based on 75 points sampled in each zone. This took approximately 44 hours to produce on a machine with two Intel Xeon(R) CPU E5-2620 (32 threads in total).

3. Results

In figure 1 we have visualized the travel times from a single origin zone, calculated in a single run with 75 nodes sampled in each zone. As we are more interested in the travel times to popular places such as hospitals and schools we applied sampling with replacement weighted based on surface and building function, dividing the surface of residential addresses by a factor of 8. This way when selecting 75 points, only 58 unique coordinates function as origin location in figure 1. With the stored intermediate results for each single origin-destination pair between zones, we analyzed the data that forms the average travel times in the skim matrices. This is a large collection of travel times as we have run the same query 100 times each and for each run we calculate $N * N$ travel times for the N points per zone. For an initial data exploration we looked at the histograms of travel times of a series with random origin-destination pairs, of which some are included in figure 2. In this paper we only included the histograms for the 100 runs with 10 sampled points in each zone, as the 100 runs with 5 sample points in each zone had identical shapes.

What we discovered in the histograms was that the travel times for a large number of origin-destination pairs are bi-modally distributed and that only a very small number of origin-destination pairs are normally distributed. For our smaller focus group of 16 selected zones, we did not spot any differences with regard to distribution when comparing the 100 runs of 5, 10 and 500 points per zone, beyond missing a few outliers in the runs with the smaller number of points. We included one example in figure 3.

4. Discussion

We see three main factors that reduce the number of normally distributed travel times, which play a role in both car and public transit but in different ways and intensity: uniformity of travel speed; weighting of points per zone; and number of points per zone. We discuss each of these in turn below.
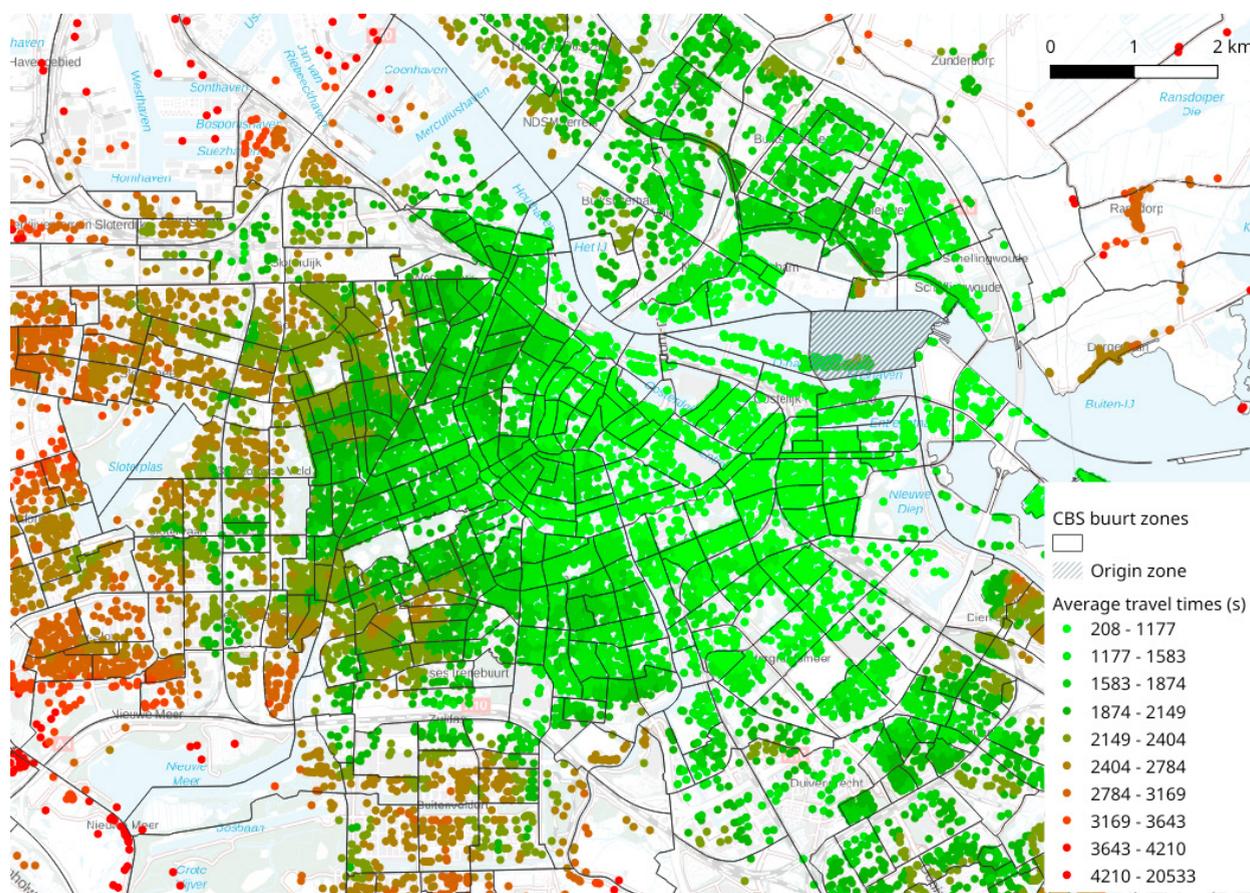
Fig. 1: Travel times originating from KNSM-eiland neighbourhood (indicated on map) based on sample size of 75 points per zone. Travel times for each destination are based on the average travel time from 58 origin points on KNSM-eiland.

4.1. Uniformity of travel speed

Travel speed with car across the neighbourhood will be uniform as long as all roads are accessible by car and one can drive directly from door-to-door and park there. If both origin and destination zone are structured in this way, then the distance would theoretically be ranged between $Length_{shortestPath}$ and $Length_{shortestPath} + Diameter_{origin} + Diameter_{destination}$ where $Length_{shortestPath}$ is the length of the path with the shortest distance between any two points in the respective zones. Thus the total distance cannot be lower than the shortest path between any point in each zone and cannot be higher than the sum of the shortest path with the diameter of the origin zone and destination zone. Assuming a single modality with a low variance in speed this means that travel times are bounded by the same range. On the other hand public transit is by definition multi-modal with walk, bicycle or car as access and egress modes. And even if the origin and destination are the transit stop itself, there will be a high variability in travel speed, depending on mode, dwell times and the number of transfers; furthermore, depending on how well connecting services are timed there will be waiting involved as well. Transit times in the worst case would involve traversing the diameter of the zone with walk speed and in the best case there could be a high-speed train directly between origin and destination. This effect has consequences however for studies focusing on infrastructure changes such as new metro stations, comparing the accessibility before and after; looking only at the average travel duration (i.e. a single number estimate) may give the wrong picture. An example of this effect is given in figure 4, where a 700 meter distance between two origin points can increase the travel time to the destination zone by a factor of 2 in some parts of that zone. This is because one point has a direct tram service with limited walking whereas the other point is 700 meter from that tram stop. However the difference near the metro stations is smaller as the metro is connected by a bus that passes both points. We see this as the cause of the bi-modality in figure 2f, which describes this zone pair.
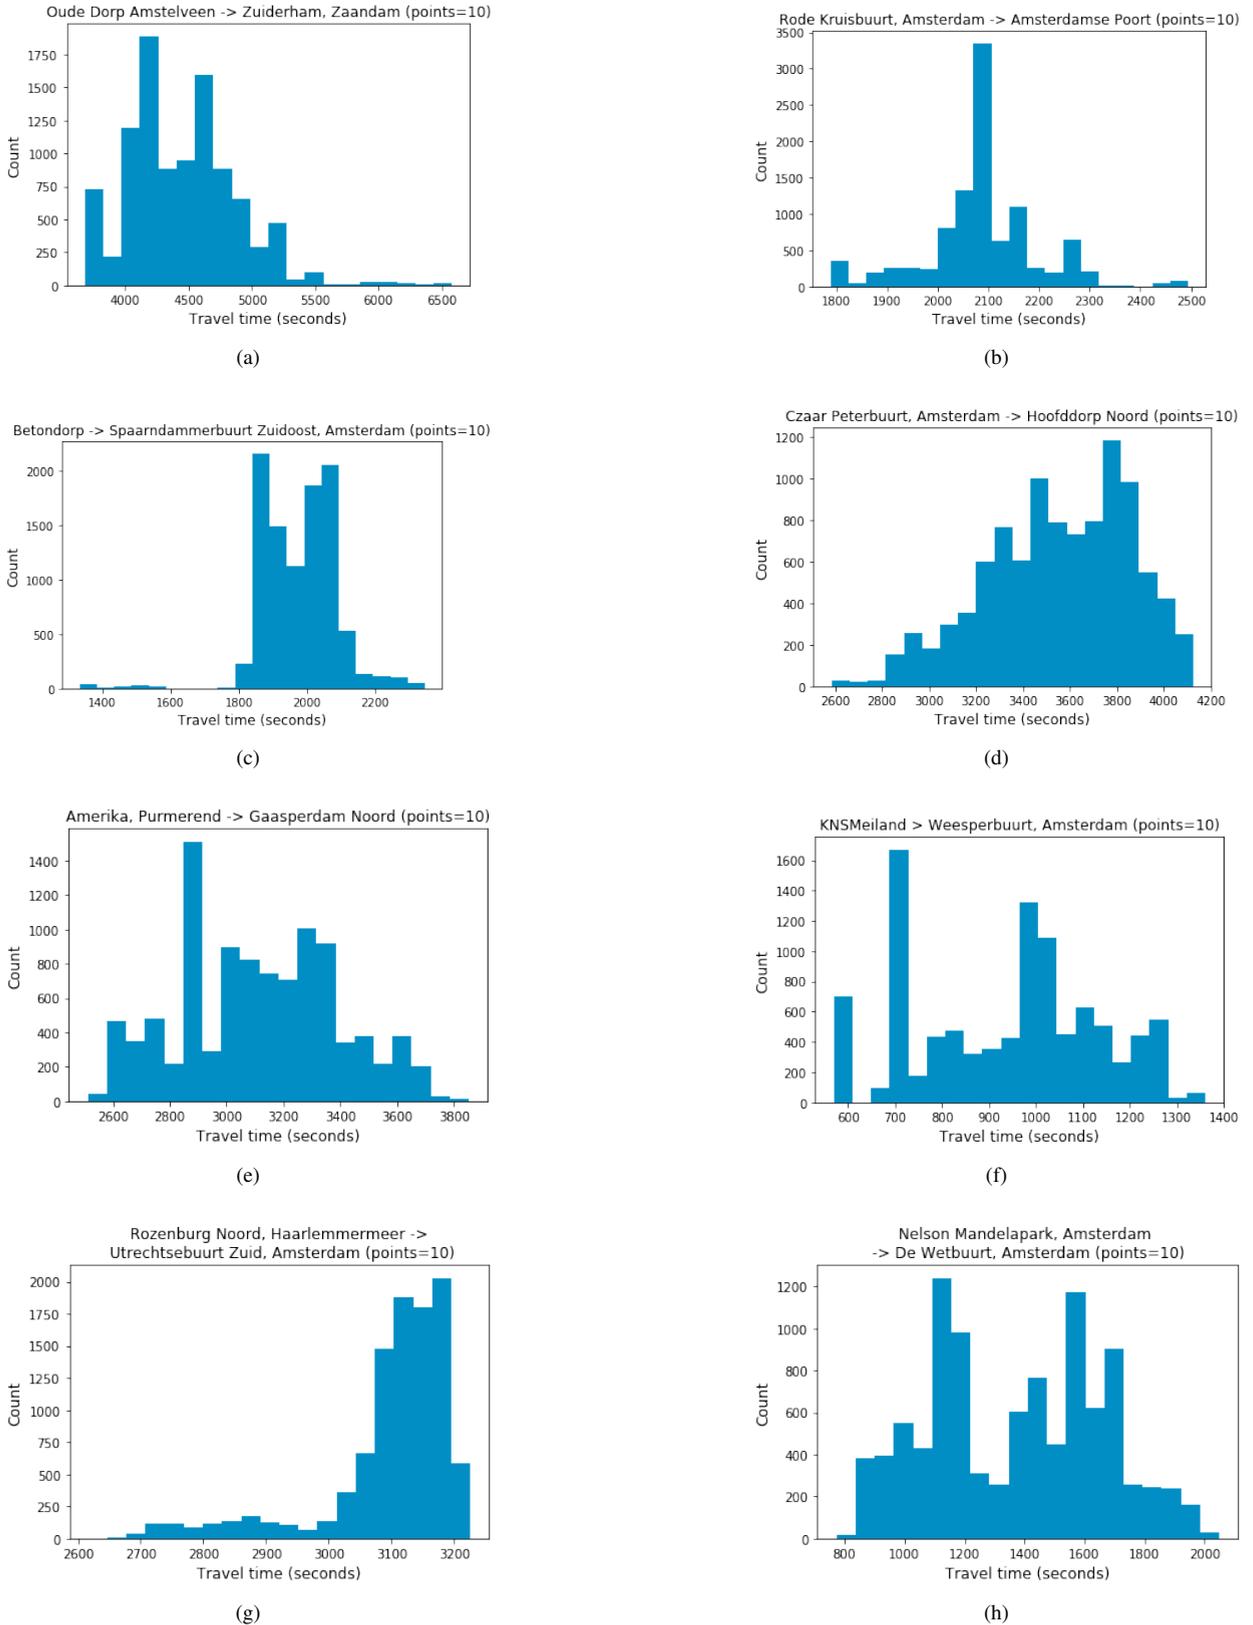
Fig. 2: Histograms of travel times door-to-door using public transit between 8 pairs of origin and destination zones in and around Amsterdam. Travel times were calculated 100 times between 10 sampled points in the origin zone and 10 sampled points in the the destination zone.
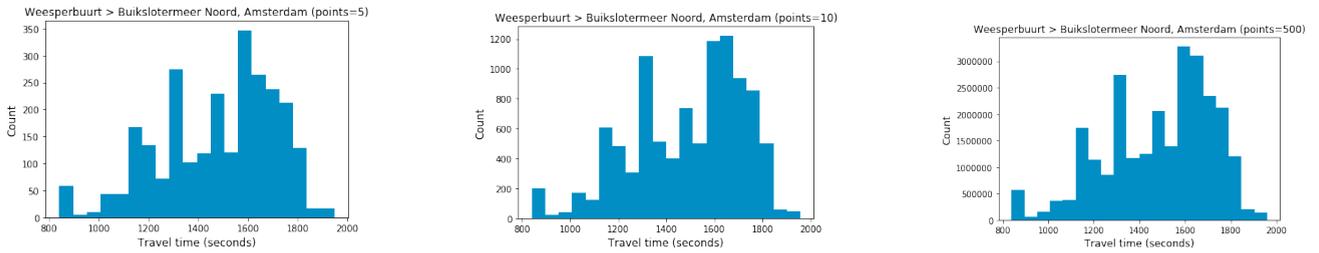
Fig. 3: Histograms of travel times door to door using public transit for *(a)* 5 *(b)* 10 *(c)* 500 points in both origin and destination zone
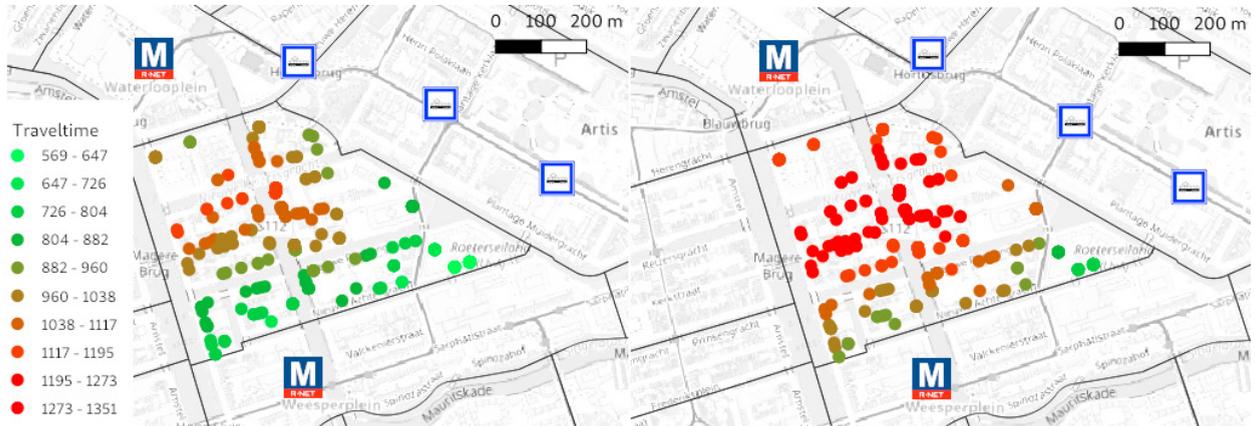

Fig. 4: Illustrating the effect of shifting the origin position in a neighbourhood 700 meter. The left figure originates from the east of the KNSM-eiland neighbourhood with a tram service directly to the indicated tram stops. The right figure originates 700 meter further away. The bus line connecting with the metro networks passes both points.

Additionally important to note is that in this study we used straight line distances to model the access and egress to transit stops; we expect that when using actual distances instead the variability of transit travel times would further increase.

4.2. Weighting of points per zone

What we saw with both transit and car accessibility is that heavily weighted points influence the distributions of the travel time. Since we aim to sample plausible origin and destination choices and not just addresses, we applied weights to addresses based on their function and surface. Namely, certain single addresses such as a hospital will function as the destination for thousands of trips per day whereas a single home is likely only the destination for a handful of people. Hence for the points we selected a single address will occur multiple times. This has an effect on the distributions as well, as the repetition of the same destination will increase the number of duplicate travel times. However with transit we see an additional interaction effect as those popular destinations will be made more accessible in well designed transit systems. For example in figure 4 there is a university building in the south of that neighbourhood, that is very well accessible and is selected multiple times in our sampling and forms a portion of that peak of 700 seconds.

4.3. Number of points per zone

What we noticed is that the number of points per zone has little influence on the distribution of the travel times. It might be better to repeat a lower amount of points per zone multiple times instead of running a higher amount of points.

5. Conclusion

Travel time level of service matrices for public transit containing a single value for each TAZ pair are commonly used. They represent point estimates applying to uni-modal travel.

In this research we introduce real multi-modal trips by accounting for first and last mile legs of journeys as well as for mode change and transfers. TAZ to TAZ travel times were computed by sampling locations in the origin and destination zone and by computing sets of Pareto optimal travel options. We show that in many cases the total travel time is not normally distributed but rather has a bi-model or even tri-modal distribution. We also show how this originates from accessibility properties of public transit stops. Elaborated experiments show that the distribution for the travel time can be accurately estimated using ten or a few dozens of locations sampled in each TAZ. Careful sampling based on an address database specifying location types (residence, school, shop, hospital, etc) is required.

6. Future research

Extension of this research aims to replace the point estimates used in classical skim matrices by a specification of the travel time distribution. A possible research path is to discover typical kinds of distributions, classify them and estimate parameters to describe them. Each cell in the skim matrix will then consist of the distribution type identification and its parameters.

References

[1] Benenson, I., Martens, K., Rofé, Y., Kwartler, A., 2011. Public transport versus private car gis-based estimation of accessibility applied to the tel aviv metropolitan area. The Annals of Regional Science 47, 499–515.

[2] Blanchard, S.D., Waddell, P., 2017. Urbanaccess: generalized methodology for measuring regional accessibility with an integrated pedestrian and transit network. Transportation research record 2653, 35–44.

[3] Bovy, P.H., Hoogendoorn-Lanser, S., 2005. Modelling route choice behaviour in multi-modal transport networks. Transportation 32, 341–368.

[4] Cats, O., Gkioulou, Z., 2017. Modeling the impacts of public transport reliability and travel information on passengers' waiting-time uncertainty. EURO Journal on Transportation and Logistics 6, 247–270.

[5] CBS, . Wijk- en buurtkaart 2019. https://www.cbs.nl/nl-nl/dossier/nederland-regionaal/geografische-data/wijk-en-buurtkaart-2019. Accessed: 2019-11-29.

[6] Curtis, C., Scheurer, J., 2010. Planning for sustainable accessibility: Developing tools to aid discussion and decision-making. Progress in Planning 74, 53–106.

[7] Delling, D., Pajor, T., Werneck, R.F., 2014. Round-based public transit routing. Transportation Science 49, 591–604.

[8] Delmelle, E.C., Casas, I., 2012. Evaluating the spatial equity of bus rapid transit-based accessibility patterns in a developing country: The case of cali, colombia. Transport Policy 20, 36–46.

[9] Dibbelt, J., Pajor, T., Strasser, B., Wagner, D., 2018. Connection scan algorithm. Journal of Experimental Algorithmics (JEA) 23, 1–7.

[10] Guis, N., Nijënstein, S., 2015. Modelleren van klantvoorkeuren in dienstregelingsstudies, in: Colloquium Vervoersplanologisch Speurwerk. Antwerpen: NS.

[11] Hadas, Y., Ranjitkar, P., 2012. Modeling public-transit connectivity with spatial quality-of-transfer measurements. journal of Transport Geography 22, 137–147.

[12] Horni, A., Nagel, K., Axhausen, K.W., 2016. The multi-agent transport simulation MATSim. Ubiquity Press London.

[13] Kujala, R., Weckström, C., Mladenović, M.N., Saramäki, J., 2018. Travel times and transfers in public transport: Comprehensive accessibility analysis based on pareto-optimal journeys. Computers, Environment and Urban Systems 67, 41–54.

[14] O'Sullivan, D., Morrison, A., Shearer, J., 2000. Using desktop gis for the investigation of accessibility by public transport: an isochrone approach. International Journal of Geographical Information Science 14, 85–104.

[15] OVapi, . Gtfs netherlands. http://gtfs.ovapi.nl/nl/. Accessed: 2019-11-29.

[16] Raveau, S., Guo, Z., Muñoz, J.C., Wilson, N.H., 2014. A behavioural comparison of route choice on metro networks: Time, transfers, crowding, topology and socio-demographics. Transportation Research Part A: Policy and Practice 66, 185–195.

[17] Salonen, M., Toivonen, T., 2013. Modelling travel time in urban networks: comparable measures for private car and public transport. Journal of transport Geography 31, 143–153.

[18] SchweizerischeBundesbahnen, . GitHub matsim-extensions by sbb. https://github.com/SchweizerischeBundesbahnen/matsim-sbb-extensions#skim-matrices. Accessed: 2019-12-01.

[19] Tenkanen, H., Heikinheimo, V., Järv, O., Salonen, M., Toivonen, T., 2016. Open data for accessibility and travel time analyses: Helsinki region travel time and co2 matrix, in: Geospatial data in a changing world: The short papers and poster papers of the 19th agile conference on geographic information science, 14–17 June 2016, Helsinki, Finland.

[20] Tribby, C.P., Zandbergen, P.A., 2012. High-resolution spatio-temporal modeling of public transit accessibility. Applied Geography 34, 345–355.