AGENT-BASED MODEL (ABM) FOR CITY-SCALE TRAFFIC SIMULATION: A CASE STUDY ON SAN FRANCISCO

B. Zhao^{1}, K. Kumar², G. Casey^{1,3} and K. Soga⁴*

¹ *University of Cambridge, Cambridge, UK*

² *University of Texas at Austin, Austin, USA*

³ *Arup, London, UK*

⁴ *University of California, Berkeley, USA*

** Corresponding author*

ABSTRACT Agent-Based Model (ABM) is a promising tool for city-scale traffic simulation to understand the complex behaviour of the entire urban transportation system under different scenarios. In the ABM, traffic is intuitively simulated as movements and interactions between large numbers of agents, each capable of finding the route for an individual traveller or vehicle. This paper presents such an ABM development to reproduce the traffic patterns of the city of San Francisco. The model features a detailed road network and hour-long simulation time step to capture realistic variations in traffic conditions. Agent speed is determined according to a simplified volume-delay macroscopic relationship, which is more efficient than applying microscopic rules (e.g., car following) for evaluating city-scale traffic conditions. Two particular challenges of building such an ABM are addressed in this paper: data availability and computational cost. The key inputs to the ABM are sourced from standard and publicly available datasets, including the travel demand surveys published by local transport authorities and the road network data from the OpenStreetMap. In addition, an efficient priority-queue based Dijkstra algorithm is implemented to overcome the computational bottleneck of agent routing. The ABM is designed to run on High Performance Computing (HPC) clusters, thereby improving the computational speed significantly. Preliminary validation of the ABM is conducted by comparing its results with a published model. Overall, the ABM has been demonstrated to run efficiently and produce reliable results.

1. Introduction

Traffic modelling provides an alternative means to forecast distributions and behaviour of traffic under different situations or policy scenarios especially on busy urban roads, where it is impractical to carry out physical experiments. As a result, they are widely used by the transport community to study various short- and long-term dynamics of the urban traffic systems, from the effect of signal control at a single intersection (Behrisch et al., 2011) to the flow and congestion distributions in a regional network (Cetin et al., 2003).

Modelling traffic in a road (network) has often been seen as analogous to modelling fluid flow in a pipe (network) (Lighthill & Whitham, 1955). However, unlike the many assumptions in the fluid mechanics (conservation of mass, energy and momentum), the only physical law that governs the traffic flow is the conservation of vehicles (Papageorgiou, 1998). This itself is not sufficient to solve for the two sets of variables in a traffic model: speed and density (or headway). As a result, an empirical behaviour rule is introduced to offer additional constraints to the traffic modelling. Based on the level of detail of this behaviour rule, the traffic models can be classified as macroscopic models and microscopic models. Macroscopic models introduce behaviour constraints on aggregated traffic variables, such as the fundamental diagrams where space-mean-speed adjusts instantaneously to density (Lighthill & Whitham, 1955) or higher order models considering the acceleration time for speed to adjust to

density (Payen, 1979). While in microscopic models, behaviour rules are targeted at individual cars, such as to keep a safe distance (Pipes, 1953). Microscopic models provide a detailed representation of the traffic process, which makes them most suitable for evaluation of complicated traffic facilities. However, they are not suitable for modelling large networks. On the other hand, a macroscopic model capture traffic dynamics in lesser detail, but are faster and easier to apply and calibrate than microscopic models. We explore a mesoscopic traffic simulation model that captures the traffic flow dynamics at a city-scale through a combination of individual agent level travel demand/ route choice and macroscopic road link dynamics with volume-delay relationship.

In this paper, an agent-based macroscopic traffic simulation model is built for the city of San Francisco. It serves in the broader vision of creating a city-scale infrastructure resiliency tool, enabling real-time decisions in response to natural disasters or disruptive events. The traffic simulation is designed to balance the abstractions and details in modelling with the goal of achieving efficient city-scale analysis. The simulation is conducted under the agent-based modelling (ABM) framework, where traffic is intuitively simulated as the movements and interactions between large numbers of agents, each representing an individual vehicle. The ABM allows individual characters to be incorporated, enabling the inclusion of complex human behaviour observed in the real world (Dia, 2002). However, agent mobility is based on the

simplified assumption of volume-delay relationship between macroscopic variables (flow and average speed). This simplification allows simulation to be carried out more efficiently to suit the needs of real-time modelling, forecast and decision making. Travel demand in this study is trip-based. However, with the detailed network representation and the fast simulation speed, the model can also be adapted for activity-based simulations in the future (Bowman & Ben-Akiva, 2001).

This paper addresses two major challenges of building the city-scale macroscopic ABM: data availability and computational cost. First of all, in Section 2, it is discussed of how to obtain a cleaned network from the OpenStreetMap and the agent-level disaggregate travel demand from travel surveys with aggregated data. Next, in Section 3, the traffic modelling framework is presented, as well as the computational performance after implementing an efficient priority-queue based Dijkstra algorithm. Section 4 presents the simulation results and validations against official models. The conclusions are summarised in Section 5.

2. Data

Two types of data are required for the agent-based macroscopic traffic simulations: the network properties (topology, capacity, speed limit, ...) and the travel demand (origin, destination, departure time, ...). This section details the process of data collection from openly available data sources as well as the assumptions involved in data processing.

2.1 Network supply from the OpenStreetMap (OSM)

OpenStreetMap (OSM) is a popular service that offers free editable digital map of the world. Map data on the OSM come from public domain mapsets, licensed aerial imagery or GPS tracelogs uploaded by volunteers (OSM Wiki, 2017). In the US, where the study area is located, the OSM road network was initially populated by the public domain TIGER maps in 2007/08 and has been gradually updated by the community over the years. The good level of completion and standard data format have made the OSM a useful network dataset used in many previous transportation studies, including the microscopic traffic simulation package SUMO (Behrisch *et al.*, 2011) and OSMnx, a comprehensive Python tool for downloading, cleaning, analysing and visualising street networks (Boeing, 2017).

Downloading the OSM road network

The OSM road network can be conveniently filtered and downloaded with the Overpass API (Overpass API Wiki, 2018). Figure 1 shows the Overpass Query Language (QL) script used to download the road network for San Francisco. Map features in the OSM are denoted by tags, such as “building”, “railway”, “power”, “waterway”, etc. In particular, the “highway” tag identifies all roads and footpaths, which not only includes the real highways but also residential roads and pedestrian-only paths (OSM Wiki, 2018). The “way[highway]” in the QL script selects all the

roads in the given bounding box. Appendix A.1 provides an example of the downloaded road network data in the JSON (JavaScript Object Notation) output format.

Figure 1 Overpass QL script to download the San Francisco road network.

Data cleaning for the OSM road network

Data downloaded from the OSM contain useful information on the topology and attributes of a traffic network, e.g., speed limits and lane counts. However, the raw data also include redundant and/or missing details that should be handled. Overall, the OSM network needs to be converted into a concise, directed graph so that subsequent traffic simulations (e.g., the shortest path finding algorithm) can be run efficiently. In this section, the pipeline for processing the OSM data to a directed graph is explained. It consists of 3 steps: (1) removing redundant nodes and constructing a road network graph from the OSM data; (2) adding directionalities; (3) populating graph properties.

Unlike, an edge of a graph that connects exactly two nodes, the OSM “ways” string together multiple nodes (≥ 2). Consider the OSM way element that represents the famous Lombard Street in SF (Figure 2). It contains 150 nodes (called “geometry points” shown as dots in Figure 2), but only the first and last nodes are meaningful intersections to adjacent roads. The rest 148 nodes only serve to depict the geometry of the hairpins. This 150-node way element is eventually replaced by a direct link between the start and end nodes, with the new link inheriting the properties of the original way element. On the other hand, if some node in the middle of a way element is an intersection that connects with other roads, the way element is split into multiple links at these locations. The second type of redundant nodes to be removed is “fake intersections”. These are different from the “geometry points” which are in the middle of one-way elements, but “fake intersections” are nodes shared by two-way elements. For instance, the Fell Street is represented as two connected way elements in the block shown in Figure 3. However, travellers starting from A are allowed no other choice at B but to go to C. Thus, the presence of node B is not necessary and is omitted in the road network graph. Instead, a new link A-C is created containing the aggregated properties of the original A-B and B-C. After removing the redundant nodes and fake intersections step, an undirected graph of the SF road network is created.

The directionality of an edge in an OSM road network is encoded using a tag called “oneway” and not merely by the order of the “nodes” forming an edge. A way element with an “oneway” affirmative tag value of “yes”, “true” or “1”, it indicates the first node is indeed the entrance to this road and

the last node is the exit. If the value is “reverse” or “-1”, then the directionality of the edge is defined with the order of the associated nodes reversed. When the value of the “oneway” tag is “no” or missing (using the default value), the edge is considered two-way accessible. In this case, the original edge is replaced by two new edges with the opposite node orders (Figure 4). This step creates a directed graph of the SF road network, with each edge representing a specific direction explicitly.

Figure 2 Curvature nodes on Lombard Street.

Figure 3 Fake Intersections on Fell Street.

Figure 4 Adding directionalities to two-way roads.

Road attributes, such as lane counts and speed limits, if missing, are set to the OSM default values (OSM Wiki, 2017). Road capacities are calculated according to the Highway Capacity Manual 2010 (Colak, Lima and Gonzalez, 2016). The capacity is determined according to the link speed limit and the number of lanes. It should be noted that the capacity calculation does not involve the detailed geometry of the road links, nor the presence of features such as a hard shoulder, as such information is hard to obtain from open data sources. The free flow travel time of an edge is set according to the speed limit, with additional delays at the intersections (Figure 5). Figure 6 shows a cleaned network for the study area, with 9643 nodes and 26893 edges.

Figure 5 Calculating edge attributes.

Figure 6 (a) Simplified SF road network. (b) Red: nodes in raw OSM data; black: nodes in the simplified network.

2.2 Intra-city travel demand from aggregated data

The OSM data offers a highly accurate representation of the traffic network. However, travel demand data of matching levels of details are hard to obtain. One possibility is to generate a synthetic population based on the demographic background as in activity-based models (Galli et al., 2009; Ory 2017); another method is to extrapolate zonal-level travel surveys to obtain node-to-node travel demand. The latter approach is adopted in this study to take advantage of existing datasets.

Zonal-level travel demand: Uber/Lyft pick-ups and drop-offs

Researchers at the San Francisco County Transportation Authority (SFCTA) and the Northeastern University collected, analysed and released data on passenger pick-ups and drop-offs by the Transportation Network Companies (TNCs, including Uber and Lyft) in San Francisco (SFCTA, 2017). Through their analysis, they reported that the TNCs account for about 15% of all intra-San Francisco vehicle trips, making it moderately representative of the overall traffic pattern. Despite the rather short data collection period, this dataset is utilised to inform the intra-city travel demand due to its representativeness (the ratio of 15% intra-city trips of the TNC dataset is higher than the response rate a typical travel survey) and high spatiotemporal resolutions (Table 1). Apart from the limitations listed in Table 1, compared with other demand generation approaches (such as demographic based approaches), the method used in this study rely heavily on the existence of the particular TNC dataset in the study area and does not relate the spatio-temporal variations in travel demands based on tangible demographic factors.

Table 1 Pros and Cons of SFCTA TNCs data

	Spatial data	Temporal data
Pros:	Origins & destinations associated with 981 Traffic Analysis Zones (TAZs), each containing 4-8 blocks.	Origins & destinations associated with a specific hour of the week.
Cons:	Only trips starts & ends in SF considered.	6 weeks: mid-Nov to mid-Dec, 2016.
high resolution		
limited coverage		

Generating disaggregated nodal-level travel demand

In the United States, Traffic Analysis Zones (TAZs) is the most widely used spatial unit in travel demand models for transportation planning (Zhao and Zhao, 2017). It is also the unit that the TNC pick-up and drop-off data are aggregated to. Figure 7(a) shows the TNC passenger pick-ups for each TAZ on a typical Monday at 9 AM. The rest of the dataset is similar to Figure 7(a), except for the time (24 hours in a seven-day week) and events (pick-ups or drop-offs).

The hourly TNC pick-ups and dropoffs are scaled to obtain the all intra-SF vehicular travel demand for each hour. The all vehicular travel demand include trips made by not only the TNCs, but also taxis, private autos and public transit vehicles (SFCTA, 2017). Figure 7(b) shows the estimated ratios between the numbers of TNC and of all vehicular trips by time periods and supervisorial districts (SFCTA, 2017). Contrary to the general traffic behaviour, the proportions of TNC trips are lower during peak hours, which may suggest that commuter trips during peak hours are still primarily taken by private cars. Also, the TNC trip ratios are higher in downtown, indicating better coverage of TNC services in these areas. TNC pick-ups and drop-offs as shown in Figure 7(a) are scaled by the TNC trip ratios as shown in Figure 7(b) to obtain the origin and destination counts by the hour and TAZ for all vehicular traffic. For example, on average there are 11.8 TNC pick-ups in TAZ 621 in downtown SF on a typical Monday at 9 AM (circled in Figure 7(a)). Since it belongs to supervisorial district 6 (circled in Figure 7(b)), the ratio between TNC traffic and all vehicular traffic is estimated to be 32% at that hour. So the number of all vehicular trips leaving TAZ 621 on Monday at 9 AM is calculated as $11.8/0.32 = 36.9$.

The next step is to connect the origin zones to destination zones, producing a list of origin-destination zone pairs. Like many other travel surveys, the TNC dataset only tells the total number of trips starting (ending) in each zone, but not their corresponding destinations (origins). To match the origins to destinations, a random sampling scheme is adopted. The aim of the sampling scheme is not to reduce the total trips in the traffic simulation, but rather to build a list of origin destination zone pairs by picking the zones where the trip starts and ends from the separate zonal origin counts and destination counts. For each hour, the total hourly intra-SF vehicular trip count is taken as the average of the “total origin counts” and “total destination counts” in all TAZs (these two

values are different, but very close). The TAZ-level origins/destinations are normalised to represent the probability of a trip starting/ending in each of the 981 TAZs. A list of TAZ pairs is sampled based on these probabilities. A trip with distance of 1.6-3.2 km (1-2 miles) is considered walkable or bikeable (Riggs, 2014). So if the distance between the centroids of a sampled TAZ pair is less than 2.5 km (about 30 minutes walking), the TAZ pair will not be included as it is assumed that this trip will be taken by walking or other non-vehicular means. In the end, zonal-level origin-destinations of trips are produced for each hour during the study period of a typical week (Figure 8).

Figure 7 (a) TNC pick-ups on Monday at 9 AM. (b) Percentages of TNC trips in all vehicle trips by time periods and supervisorial districts (reproduced from SFCTA, 2017).

Figure 8 Generating zone-to-zone travel demand through sampling for a typical Monday at 9 AM: (a) Sampling an origin TAZ based on probability defined by zonal-level origin counts; (b) Sampling a destination TAZ based on probability defined by zonal-level destination counts.

In the last step, nodal-level origin-destination pairs are generated by selecting a random node within the starting or ending TAZ for each zonal-level origin-destination pair. This leads to a list of node-to-node travel demand (Figure 9). It can be seen that for the all-vehicular travel demand scaled up from the TNC dataset, the morning peak and evening peak are clearly visible for workdays, with the magnitude and duration of Friday's evening peak being the highest and longest, extending over the midnight. For Saturday and Sunday, only a late afternoon/evening peak presents. This, together with the intercity demand in the next section, will be used to seed agents in the traffic simulation.

Figure 9 Generated node-to-node travel demand by hour and day of week.

2.3 Intercity travel demand

San Francisco is connected to nearby cities through four “gates”: (1) Golden Gate Bridge (to the North); (2) Bay Bridge (to cities on the East); (3) California State Route 1 (SR 1, to the South) and (4) San Francisco International Airport (SFO, air traffic). As the intra-SF traffic was collected during mid-November to mid-December, intercity traffic through the above four “gates” is also obtained for the same period (Appendix A.2). It is further assumed that vehicles enter/leave SF via these four “gates” are through-traffic, i.e., they do not start or stop within the city. This assumption can be relaxed in the future by incorporating regional travel surveys (e.g., the California Household Travel Survey, 2010-12) to decide the origins and destinations inside SF for intercity trips.

Table 2 Daily intercity travel demand in SF through four “gates”.

Entrance \ Exit	GGB ^a	BB ^b	SR 1 ^c	SFO ^d	Total entrances ^e
GGB ^a	-	35,038	19,656	13,895	68,589 55k
BB ^b	35,466	-	64,187	44,703	144,356 180k
SR 1 ^c	19,950	63,433	-	24,974	108,357 100k
SFO ^d	13,926	44,885	24,846	-	83,657 70k
Total exits ^e	69,342 55k	143,356 180k	108,689 100k	83,572 70k	404,959 405k

^a GGB: Golden Gate Bridge;

^b BB: Bay Bridge;

^c SR 1: California State Route 1;

^d SFO: San Francisco International Airport;

^e Total entrances/exits: trip counts used in the simulation are given above the dash line, while trip counts published on the websites of the four “gates” are given below the dash line and in bold.

Table 2 shows the daily intercity traffic volumes between the four “gates” in the form of an origin-destination (OD) matrix. The numbers in bold (below the dash line in the last row and

column) in Table 2 are the approximate daily entrances/exits at each of the four “gates” published on their websites (Appendix A.2). For the Golden Gate Bridge and the Bay Bridge, the daily traffic volumes are not provided by direction, so it is assumed that the traffic is equally split between the entering and exiting directions at these two “gates”. For the SR 1 and SFO, daily traffic volumes for each direction are known. However, as the differences in traffic volumes between the two directions are close, a rounded average number is taken as the total daily entrances as well as exits for each of them. A sampling process similar to the one described in Section 2.2 (inter-SF traffic) is utilised to obtain a list of sampled origin-destination pairs based on the observed total entrances and exits (bold numbers in Table 2). The probability of being selected as an entrance or an exit point is equal to the normalised observed daily traffic volume. Despite the large sampled population of a total of 405,000 intercity trips, the sampled distribution (numbers above the dash line in the last row and column in Table 2) is still not exactly the same as the theoretical probability.

3. Model

The intra-SF and inter-city travel demand is loaded/assigned onto the SF road network graph to calculate the traffic distributions during the modelling phase. Specifically, a travel agent is created for each OD pair and traverses the network through a series of graph edges that form the optimal route between the origin and destination nodes. Temporally, the traffic ABM progresses by one-hour time steps. Inside each time step, trips are dispatched in 10 batches (iterations) and link-level travel times are updated after each batch. The framework for traffic simulation is shown in Figure 10.

3.1 The outer loop: temporal evolution

The outer loop of the traffic simulation controls the progress of time. The body of the outer loop specifies how traffic is simulated in each one-hour time slice and will be explained in detail in Section 3.2. For each day simulated, the outer loop is executed for 24 times, starting from 3 AM in the morning (when the traffic is the lightest (Ragland *et al.*, 1992)) till 3 AM in the next day, so the traffic simulation in the first time step of each day can be assumed to run with all link-level travel times equal to their free-flow travel times. In this study, traffic is simulated for the seven typical days in a week.

3.2 The inner loop: traffic assignment in each hour

The inner loop of the traffic simulation implements a routing algorithm to find the route/path for each of the agent that is scheduled to travel in the specific hour (set by the outer loop). A route or path is a sequence of graph edges that connect the origin node to the destination node. It is assumed that an agent will seek the optimal (fastest) route when travelling, ignoring other factors such as route distance, toll cost or habitual route. This is equivalent to (1) setting the time to traverse each road link as the edge weight in the road network graph; (2) finding the path that has the smallest total weight (shortest travel time); (3) updating the link-level travel time based on the simulated traffic flow.

Figure 10 Framework for the traffic simulation.

Link-level travel time

Link-level travel time is set according to the well-known Bureau of Public Roads (BPR, 1964) volume-delay curves. It has the following form (Colak et al., 2016):

$$t_t = t_f \times (1 + \alpha(\text{volume}/\text{capacity})^\beta) \times f_p \quad (1)$$

where t_t is the link-level travel time with traffic on the link; t_f is the free flow travel time of the link; *volume* is the amount of traffic on the link; *capacity* is the link capacity; α, β are calibration parameters (set to 0.6 and 4) and f_p is the city-specific correction factor (set to 1.3 for SF, (Colak et al., 2016)). t_f and *capacity* are calculated in Section 2.1. *volume* is initialised with the link-level traffic from the previous time step, or zero if in the first time step. Within one time step, the *volume* (so consequently t_t) of the network is updated 10 times during the iterative assignment process, which will be explained shortly.

Shortest-path finding

A priority-queue based Dijkstra shortest path algorithm is implemented to find the fastest route for each agent. Dijkstra algorithm is an efficient way to find the shortest path between two nodes (vertices), or to produce the shortest path tree from a “source” node to all other nodes in a weighted graph (Dijkstra, 1959). In the ABM, the Dijkstra algorithm is applied to the road network graph weighted by link-level travel time and finds the weighted shortest path (i.e., fastest route to travel) from the agent’s origin node to its destination node. Dijkstra algorithm can be implemented with a min-priority queue data structure. The priority queue stores all unvisited graph vertices (whose shortest distances to the

origin vertex have not been decided), prioritised according to their current (not necessarily the shortest) distances to the origin vertex. This implementation allows the Dijkstra algorithm to take advantage of the fast operations of the priority queue, including inserting a node, extracting the node with the lowest priority and decreasing the priority of a node (Montanaro, 2013).

Update of link-level travel time

Link-level travel time at the beginning of each time step is set according to the simulated traffic from the previous hour. This is updated at the end as well as in the middle of the current time step through a process called “iterative assignment” (NPTEL, 2007). In the iterative assignment of this study, the hourly travel demand is randomly partitioned into 10 chunks of equal size. For example, there are around 87,000 agents travelling on a typical Monday at 8 AM (63,000 intra-SF and 24,000 intercity). So 10% of the total hourly demand, e.g., 8,700 agents, are assigned (the fastest) routes in the first batch. Link-level travel time is updated for those links that were used in this batch but not updated for other links (since the traffic volume did not increase). This “traffic assignment and graph weight update” process is repeated for 10 times until routes are assigned for all the 87,000 agents that travel on Monday at 8 AM (Figure 10). The iterative assignment can produce more realistic traffic patterns than the single-step assignment, where link-level travel time is fixed during each time step. It is also more efficient than the “user equilibrium” assignment, wherein the trip assignment is iteratively adjusted until no driver can unilaterally reduce his/her travel cost by shifting to another route (Miller, 2014; NPTEL, 2007).

Shared-memory parallelisation

The Dijkstra shortest path algorithm is executed for each agent in the ABM at every time-step. This quickly becomes the computational bottleneck as the number of agents and graph size increase in a city-scale simulation. However, as the graph weights do not change within a batch of the iterative assignment process (Figure 10), the choice of route for an agent is not affected by the route assignment of other agents in the same batch. As a result, the Dijkstra algorithm can be executed in parallel for agents in the same batch. In this study, the SF model is deployed on a single-node of the Cambridge High Performance Computing (HPC) cluster CSD3 (<https://www.hpc.cam.ac.uk>), utilising 32 cores (parallel processes). The hourly number of agent trips to be simulated varies from 4,111 (Tuesday at 3 AM) to 151,170 (Friday at 7 PM). The number of agents in a batch (10% of hourly agent trips whose optimal routes are calculated in parallel) is between 400 to 15,000. This adds up to 9 million agent trips for the simulation of a typical week. It took ~22 minutes to simulate the traffic of a typical week in the SF road network with 9643 vertices and 26893 edges.

4. Results

A direct result from the ABM simulation is the traffic volume on every road link in each hourly time step. Figure 11 shows

such results on a typical Friday at 6 AM (off-peak) and at 6 PM (evening peak hour). At 6 AM, the traffic is light on most roads in the city except from the highways. While at 6 PM, significantly more traffic is seen on the highways as well as roads in the inner part of the city. However, downtown SF (in the northeast) is not particularly congested. This is probably because of the dense placements of traffic lights in the city centre, making it more delayed and thus less attractive to go through these streets in agents routing.

Figure 11 Hourly link-level traffic volume: (a) on a typical Friday at 6 AM; (b) on a typical Friday at 6 PM.

Figure 12 shows the volume-to-capacity ratio of each road link on a typical Friday at 6 AM and at 6 PM. It is obtained by dividing the traffic volume data in Figure 11 by the capacity of each road as calculated in Section 2.1. At 6 AM, traffic is below capacity on almost every road. While at 6 PM, about 10% of the roads have traffic that exceeds capacity. The critical links with high volume-to-capacity ratio are not exactly the same as the links that have more traffic. For example, the Bay Bridge (the long link in the northeast corner) has almost the heaviest traffic flow in SF at Friday 6 PM. But as it also has a high capacity, its volume is still below the capacity.

Figure 12 Road volume-to-capacity ratio: (a) on a typical Friday at 6 AM; (b) on a typical Friday at 6 PM.

Figure 13 compares the results from the ABM simulation in this study with results from the Bay Area Metropolitan Transport Commission's (MTC) Travel Model One (MTC, 2018). Specifically, the "2015_06_022" scenario in the MTC model result repository is chosen for comparison as its forecast year of 2015 is the closest to the year of 2016, when the travel demand data used in this study were collected. The MTC model utilises a simplified road network for the whole Bay Area but predicts the traffic distributions of different types of vehicles (passenger cars, commercial vehicles of various size, etc.). Temporally, it provides results for five time periods in a typical weekday (early morning, AM peak, midday, PM peak and evening) (MTC, 2018). As weekends are not covered by the MTC model, it is decided that the total weekday traffic from the two models should be compared. Figure 13(a) shows the total weekday traffic (Monday to

Friday, 24 hourly time steps per day) from the ABM simulation in this study. Figure 13(b) is the counterpart from the MTC's model, cropped to the study area.

Figure 13 Total simulated traffic in SF from Monday to Friday: (a) results from this study; (b) results from MTC's Travel Model One.

Figure 14 Weekday link-level traffic against cumulative link length ("mileage").

An initial visual inspection of Figure 13 suggests that similar total weekday traffic is predicted by these two simulation models. However, it is difficult to conduct more quantitative comparisons of the results from Figure 13. To begin with, there is no common field in these two networks to identify corresponding road links, such as street names or the road ID. Besides, the MTC network is sparser geographically and has more simplified road geometries than the OSM data, so the same roads do not align spatially in these two maps. Due to these practical difficulties in aligning the results from these two models, an alternative approach is taken to facilitate the comparison instead. First, the road links in each model are re-ordered, ranked by the total weekday traffic in descending order. Here, the "total weekday traffic" is used as a measure/proxy of link criticality. The more traffic on a link, the higher its criticality and the higher its rank among all links. Based on this criticality rank, a "mileage" number is calculated for each link by taking the cumulative sum of link length (total length for roads ranked before itself). Road links ranked higher receive smaller mileage numbers. Figure 14 shows the scatter plot of the link-level weekday traffic against the link mileage. It can be seen that (1) the total length of the MTC network (1328 km) is about 40% of the OSM network (3137 km), which is expected as the MTC network contains only important roads and leaves out many residential streets, while the OSM network includes all drivable roads; (2) the data series associated with the ABM results in this study (labelled as "SF ABM" in Figure 14) is higher when the road mileage is below 100 km, indicating that the ABM this study

predicts more traffic than the MTC model for the top 100 km of roads ranked in terms of traffic volumes (the heavier traffic regime of the network); (3) the two data series overlap after 100 km of mileage, indicating that similar amounts of traffic are predicted by the two models for road links with mileages between 100-1000 km.

5. Summary

This paper presents the development of a macroscopic agent-based model (ABM) for traffic simulation based on a case study set in the city of San Francisco (SF). The simulation model is based entirely on inputs processed from openly available datasets, including (1) a detailed road network graph from the OpenStreetMap, with nearly 10,000 vertices and 27,000 edges; (2) intra-SF travel demand inferred from Uber/Lyft pick-ups and drop-offs; (3) intercity trips that enters and leaves the study area at four “gates” on the boundary of SF. The hourly intra-SF and intercity travel demand add up to trips for nearly 9 million agents during the one-week time scope analysed in this study. To handle the route finding of the large numbers of agents, a priority-queue based Dijkstra algorithm is implemented to accelerate the shortest path computation. The simulation model uses shared-memory parallelisation. Results of hourly link-level traffic volume and volume-to-capacity ratio are presented for a peak hour time step and an off-peak time step. In addition, the simulation results are compared against an official model both spatially and along a “mileage” system that linearises the road network. Overall through this case study, the model has demonstrated good computational performance as well as the ability to produce sensible traffic simulation results.

References

- Behrisch M *et al.* (2011) Sumo—simulation of urban mobility. *The Third International Conference on Advances in System Simulation (SIMUL 2011), Barcelona* (Omerovic A, Simoni DA and Bobashev G (eds)). ThinkMind Digital Library, pp 55-60.
- Boeing G (2017) OSMnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks. *Computers, Environment and Urban Systems* 65: 126-139.
- Bowman JL and Ben-Akiva ME (2001) Activity-based disaggregate travel demand model system with activity schedules. *Transportation Research Part A: Policy and Practice* 35(1): 1-28.
- BPR (Bureau of Public Roads) (1964) Traffic Assignment Manual. Washington, DC.
- Cetin N *et al.* (2003) A large-scale agent-based traffic microsimulation based on queue model. *Proceedings of Swiss Transport Research Conference (STRC), Monte Verita*.
- Çolak S *et al.* (2016) Understanding congested travel in urban areas. *Nature communications* 7: 10793.
- Daganzo CF (1997) A continuum theory of traffic dynamics for freeways with special lanes. *Transportation Research Part B: Methodological* 31(2): 83-102.
- Dijkstra EW (1959) A note on two problems in connexion with graphs. *Numerische mathematik* 1(1): 269-271.
- Dia H (2002) An agent-based approach to modelling driver route choice behaviour under the influence of real-time information. *Transportation Research Part C: Emerging Technologies* 10(5-6): 331-349.
- Galli E *et al.* (2009) ActivitySim: large-scale agent-based activity generation for infrastructure simulation. *Proceedings of the 2009 spring simulation multiconference, San Diego*. Society for Computer Simulation International, San Diego, CA, USA, article no. 16.
- Golden Gate Bridge, Highway and Transportation District (2018). Golden Gate Bridge Southbound Traffic. http://goldengate.org/news/transit/trends_traffic-transit.php#ggbtrends (accessed 01/2019).
- Hoogendoorn SP and Bovy PH (2001) State-of-the-art of vehicular traffic flow modelling. *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering* 215(4): 283-303.
- Lighthill MJ and Whitham GB (1955) On kinematic waves II. A theory of traffic flow on long crowded roads. *Proceedings of the Royal Society A* 229(1178): 317-345.
- Maerivoet S and De Moor B (2005) Transportation planning and traffic flow models. *arXiv preprint physics/0507127*.
- McNelly (2007) The Four-Step Model. *Handbook of Transport Modelling* (Henshner DA and Button KJ (eds)). Emerald Group Publishing Limited, Bingley, UK, pp. 35-53.
- Miller MK (2014) *Seismic risk assessment of complex transportation networks*. PhD thesis, Stanford University, USA.
- Montanaro A (2013) Priority-queues and Dijkstra's algorithm. <https://people.maths.bris.ac.uk/~csxam/teaching/dsa/dijkstra.pdf> (accessed 01/2019).
- MTC (Metropolitan Transportation Commission) (2017). 2016 Traffic Volumes on California State Highways. http://www.dot.ca.gov/trafficops/census/docs/2016_aadt_volumes.pdf (accessed 01/2019).
- MTC (2018). Travel Model Data: Data Repository. <http://data.mtc.ca.gov/data-repository/>
- MTC (2019). Historic Toll-Paid Vehicle Counts and Total Toll Revenues. <https://mtc.ca.gov/about-mtc/what-mtc/mtc-organization/three-agencies-one-bay-area-toll-authority/historic-toll-paid> (accessed 01/2019).
- NPTEL (National Programme on Technology Enhanced Learning) (2007) Traffic Assignment. <https://nptel.ac.in/courses/105101087/downloads/Lec-10.pdf> (accessed 01/2019).
- OpenStreetMap Wiki (2017) About OpenStreetMap. https://wiki.openstreetmap.org/wiki/About_OpenStreetMap (accessed 01/2019).
- OpenStreetMap Wiki (2018). Map Features. https://wiki.openstreetmap.org/wiki/Map_Features#Highway (accessed 01/2019).
- Ory D (2007) Introducing Doppelgänger, Model Lab's open-source population synthesizer.

<https://www.sidewalklabs.com/blog/a-first-step-toward-creating-a-digital-planning-laboratory-is-populating-it/> (accessed 01/2019)

Overpass API Wiki (2018). https://wiki.openstreetmap.org/wiki/Overpass_API (accessed 01/2019).

Papageorgiou M (1998) Some remarks on macroscopic traffic flow modelling. *Transportation Research Part A: Policy and Practice* 32(5): 323-329.

Payne HJ (1979) FREFLO: A macroscopic simulation model of freeway traffic. *Transportation Research Record*, 722.

Pipes LA (1953) An operational analysis of traffic dynamics. *Journal of applied physics* 24(3): 274-281.

Ragland DR *et al.* (1992) Traffic volume and collisions involving transit and non-transit vehicles. *Accident Analysis and Prevention* 24(5): 547-558.

Riggs W (2014) Dealing with parking issues on an urban campus: The case of UC Berkeley. *Case Studies on Transport Policy* 2(3): 168-176.

SFCTA (San Francisco County Transportation Authority) (2017) TNCs Today: A profile of San Francisco Transportation Network Company Activity. <https://www.sfcta.org/emerging-mobility/tncs-today> (accessed 12/2018).

SFO (San Francisco International Airport) (2016). Analysis of Scheduled Airline Traffic Comparative Traffic Report for November, 2016. <http://media.flysfo.com.s3.amazonaws.com/media/sfo/media/air-traffic/as201611.pdf> (accessed 01/2019).

Zhao H and Zhao Y (2017) Traffic analysis zones- how do we move forward? Under review for *AASHTO CTPP and TRB Joint Conference, "Applying Census Data for Transportation", Kansas City, Missouri*. http://onlinepubs.trb.org/onlinepubs/conferences/2017/census-data/TAZ_Paper.pdf (Accessed 03/2019).

Appendix

A.1 Example of downloaded OSM data

```
{
  "version": 0.6,
  "generator": "Overpass API 0.7.55.4 3079d8ea",
  "osm3s": {
    "timestamp_osm_base": "2018-09-10T15:44:03Z",
    "copyright": "The data included in this document is
from www.openstreetmap.org. The data is made available
under ODbL."
  },
  "elements": [

    {
      "type": "node",
      "id": 26117861,
      "lat": 37.6539819,
      "lon": -122.4071324
    },
    ...
    {
      "type": "way",
      "id": 394222581,
      "nodes": [
        704673274,
        3972323851,
        3972323849,
        3972323848,
        3972323847,
        3972323846
      ],
      "tags": {
        "highway": "tertiary",
        "name": "King Drive",
        "oneway": "yes",
        "tiger:cfcc": "A41",
        "tiger:county": "San Mateo, CA",
        "tiger:name_base": "King",
        "tiger:name_type": "Dr"
      }
    },
    ...
  ]
}
```

A.2 Source of Inter-city traffic:

Below are the observed traffic counts for the Golden Gate Bridge. There are 1,654,022 southbound bridge traffic in November 2016 (Golden Gate Bridge, Highway and Transportation District, 2018). Assuming southbound and northbound traffic is equal, this is equal to 55,000 daily traffic per direction.

Next are the observed traffic counts for the Bay Bridge. In the Fiscal year 2016-17, there are 136,813,538 toll-paid vehicles (MTC, 2019). Assuming eastbound and westbound traffic is equal, this is equal to 180,000 daily traffic per direction.

Then, there are the observed traffic counts for the California State Route 1. In 2016, at San Francisco Alemany Boulevard, there southbound and northbound AADT are 103,000 and 96,000, respectively (MTC, 2017). Assume northbound and southbound traffic is equal, this is equal to around 100,000 daily traffic per direction.

At last, there are the observed traffic counts to and from the San Francisco International Airport (SFO). In November 2016, the total enplaned (departure) at SFO is 2,137,473. Total deplaned (arrival) is 2,127,669 (SFO, 2016). Assume every passenger arrives in individual cars, this is equal to 70k daily traffic per direction.