Adding freight traffic to MATSim

Michael Zilske* Stefan Schröder† Kai Nagel* Gernot Liedtke†

August 1, 2012

This paper explores the integration of freight traffic into an agent-oriented simulation system for large scale traffic simulations. An agent type for freight operators is described, and it is implemented with a minimal interface to the existing simulation system. Simulating freight traffic with passenger traffic requires to consider different temporal planning horizons, as well as a traffic flow model with heterogeneous vehicle fleets.

Keywords: Agent-based freight transportation modeling; Large scale simulation of agent-based microscopic traffic models

1 Introduction

The MATSim Multi-Agent Transport Simulation toolkit (Balmer et al. [2009]) is mainly designed for large-scale simulations of personal vehicle traffic in urban areas. It consists of a travel demand model of individual travellers and their choices, and a traffic flow simulation, which can be understood as the physical layer where a population of users meets a capacity-constrained transport system. Commercial vehicle traffic has usually been modeled as a background network load without much adaptive behavior and without taking into account its distinct physical characteristics such as lower speed and higher road capacity consumption. This paper explores the integration of freight transport into MATSim. It focuses on freight-related adaptations to the software and on behavioral aspects in relation to the freight agents.

2 From people with activity plans to freight operators with schedules

2.1 Passengers

The travel demand model implemented in MATSim consists of a set of agents representing individual users of a traffic system. Every agent is equipped with a plan, which describes locations, times and types of all the activities the agent will conduct, with legs connecting each physical activity location to the next. Each leg is travelled using a specified transport mode and, depending on the transport mode, along a specified route through the transport system. All agents simultaneously execute their plans in a concurrent simulation of the transport system. The simulation picks up congestion effects, missed public transit connections, and delayed arrivals at activity locations. The results of the simulation are fed back to the agent as observations, and they are used to evaluate the plan using a scoring function which can be personalized for each individual, for example by depending on their age or income. In

---

*Technische Universitat Berlin, Fachgebiet Verkehrssystemplanung und Verkehrstelematik

†Karlsruher Institut fur Technologie, Institut fur Wirtschaftspolitik und Wirtschaftsforschung

reaction to the traffic flow simulation, some agents modify their plan using one of several re-planning modules, which correspond to the choice dimensions available to the agent. This includes choosing a new route, switching transport modes, and choosing new times for activities. Simulation, scoring and re-planning are iterated until a relaxed state resembling a dynamic user equilibrium is reached. The planning and re-planning model employed here is tailored to the daily schedules of private passengers who use the transport system to get from home to work and to leisure or shopping locations: Agents might re-schedule shopping trips to avoid rush-hour congestion, or switch to public transport for commuting. Additionally, plans wrap around over night: It is assumed that the last activity of the day is the same as the first (usually: being at home), and in consequence, the starting time of the first activity is the same as the starting time of the last activity of the day. The advantage is that one does not need to make special assumptions for the temporal boundaries of the system, similar to periodic spatial boundary conditions in materials science simulations.

2.2 Carriers

Up to now, real-world scenarios set up with MATSim have modeled the freight traffic share of the demand by using a set of plans with activities at the depot and at pick-up and delivery locations, but with no variability in any choice dimension except route choice. We improve on this situation by modeling freight vehicles as non- autonomous agents employed by and serving the interests of freight operators. The missing choice dimensions of freight vehicle drivers are then realized as logistics decisions made by the freight operators who employ them. In the freight transport sector, decisions are distributed among actors with different roles. Freight transport decisions include lot-size choice, path-searches in logistical networks, vehicle choice and tour planning. The planning problem of a freight operator is therefore quite different from its passenger counterpart:

1. Firstly, the success of freight transport plans is not determined by the utility of time spent at activity locations, but rather by their commercial success. They have to meet the requirements of customers, like meeting time windows and providing enough capacity at a reasonable cost.
2. Secondly, freight operators often do not operate one single vehicle but several, and their options include rescheduling deliveries from one vehicle to another or even changing the number of vehicles which are used at all.

Therefore, a new software layer populated by *carrier agents* was introduced into the simulation. Each carrier agent represents a firm with a vehicle fleet, depots and contracts. Contracts determine type and quantity of goods to be carried. A contract contains the respective origin and destination as well as pick-up and delivery time windows. The plan of a carrier agent contains a schedule of a tour for each of the vehicles in the fleet. The schedule contains planned pick-up, delivery or arrival times at customer locations and a route through the physical network. In our basic model, all vehicle schedules of a carrier begin and end at one of its depots. When a simulation scenario is initialized, the carrier agents build a schedule for each of their vehicles, including a route through the transport network, with pick-up and delivery activities corresponding to their contracts. At the interface between the freight operator plans and the mobility simulation, the set of routed vehicles from each carrier plan is injected into the traffic demand as individual driver agents, where they move through the traffic system along with passenger vehicles. While executing their plans, the freight driver agents report their shipment-related activities back to the carrier. When all plans have been executed, agents evaluate the success of their plan. The carrier agents use a custom utility function that captures their economic success. Their cost is calculated as a sum of vehicle-dependent distance and time costs incurred by their scheduled vehicles, and some individual fixed costs, plus penalties incurred by missed time windows. Finally, carrier agents create new plans to improve their performance in the next iteration. For instance, a time dependent vehicle routing heuristic can be plugged in to replan vehicle schedules. Shipments can be switched between vehicles, or even an entire vehicle added or removed. During repeated executions of their plans, passengers as well as carriers collect experience from the transport system. The carriers pick up congestion and other disturbances in the traffic system when they incur a higher cost through longer vehicle usage, or by penalizing missed pick-up and delivery times.

2.3 Implementation

MATSim has a modular and extensible architecture, allowing users to plug in custom reporting tools as well as specialized re-planning modules. The traffic flow simulation delivers its output as a stream of events which is used as input by other core parts of MATSim as well as by user-provided code. While implementing the freight extension, we took care to avoid tight couplings to the rest of the system as well as changes to the core code. The interface between MATSim and the freight extension could be kept conveniently small:

1. The carrier agent layer produces a population of freight driver agents which are injected into the mobility simulation upon start.
2. It listens to simulation events, keeping track of the distance travelled by the freight drivers and of the experienced pick-up and delivery times.
3. After each iteration of the mobility simulation, any number of user-provided scoring and re-planning modules for the carriers are called. They can make use of the time- dependent link travel times from the mobility simulation. Re-planning modules could, for example, employ dynamic vehicle routing algorithms to optimize tours based on the experienced travel times. Another re-planning module might be a transport market model, which would redistribute contracts among carriers. In principle, anything which affects the plans or contracts of a carrier can be plugged in (Schröder et al. [2011]).

It was not required to make any changes to the mobility simulation or even to introduce custom code for the freight driver agents. As far as the mobility simulation is concerned, freight drivers execute routes through the network, just like private vehicle operators. Changes to the mobility simulation would indeed be required for custom within-day behavior as opposed to plan-following. For example, the current implementation does not allow for a freight vehicle to wait if the goods it is supposed to pick up are unexpectedly not available yet due to traffic conditions on the upstream leg. It is assumed that the penalty of being late in one leg of a multi-leg delivery is avoided by the re-planning, and that in a relaxed system state, the planned pick-up time for a shipment is no earlier than the time at which it becomes available.

3 Simulating sub-populations with different planning horizons

Simulating long-haul freight traffic alongside a population of private car users raises the problem of different planning horizons. Personal traffic demand is modeled as a collection of plans for a day, and the agent which re-plans its schedule considers its options regarding daily activities. Trips which cannot be understood as being part of a daily routine, such as long-distance travel to a vacation spot, often do not appear in an activity-based demand model. In any case, one day is the time frame considered in replanning, and it is also the time period which is simulated in each iteration of the mobility simulation. On the other hand, freight traffic contains a substantial share of long-haul traffic with tours spanning several days, or shipments taking several days. There are two ways of solving this problem:

1. Leaving the simulated time period at one day. This requires that multi-day trips are

broken down (Joubert et al. [2009]) into segments starting at 00:00 and ending at 23:59.

2. Switching the simulated time period to something longer. Passenger plans are either repeated every 24 hours, or also switched to something longer.

If the purpose of the simulation is modeling decisions of freight operators, which includes weighing long tours versus short tours, the second option is desirable. As before, it is advantageous to have some periodicity in the simulated time period, since this reduces the problems with the boundaries. The next longer period which comes to mind is the week. With respect to freight, a week is plausible since many freight companies attempt to have their drivers at home over the weekend, even if this is just because of weekend restrictions for freight vehicles. In order to get from a population with daily plans to a population of weekly plans, a first step could be to unwind the plans, replicating them to fill a week. This is straight- forward except for the question of what to do in early iterations of the mobility simulation, where many vehicles may be severely delayed up to the point where they do not get home by the end of the day. Two options come to mind:

1. Treat the weekly plan the same way daily plans are treated. This means that delays are accumulated, and in the extreme case, an agent could be caught in congestion for several days.
2. Reset all agents to their home locations at a time where almost everyone plans to be at home (such as 03:00) and have them start their day afresh.

The re-planning step also requires some consideration for weekly plans. The initial demand is generated from data for one typical workday. This leads to the idea that re- planning should be continually performed on a daily plan which is then unwound in each iteration, so the population changes their daily routine for the whole week at once. If passenger data specifically for Saturdays and Sundays were available, three classes of days could be considered and re-planned separately: workdays, Saturdays and Sundays. In each iteration, each of the three day templates would be re-planned and a new weekly plan would be built from five copies of the workday plan and one copy of the Saturday and Sunday plan. However, if the unwound weekly plan is re-planned as a whole, most probably every workday in a plan will end up differently. The variability in the resulting plans may only represent uncertainty, but it could also suggest that agents have reasons for following different routines on different days of the week, for example by splitting their weekly work-hours unevenly over the days to avoid peak-hour congestion. Another option worth exploring and which is expected to speed up the relaxation process is to warm-start the passengers by using an already relaxed set of daily plans which was simulated without the freight share, in a simulation of one day. These plans would then be unwound to a week, the freight share would be added, and the weekly simulation with freight traffic would be started from there.

4 Traffic flow simulation with vehicles of different speed

The queue model of traffic flow implemented in MATSim (Çetin [2005]) (Gawron [1998]) is designed to be computationally faster than car-following models. It operates on the principle that the time a vehicle spends on a link is split between moving to the end of the link and waiting in a queue. When a vehicle enters a link at time t, its earliest link exit time $t_{exit} = t + \Delta t$ is determined, where Δt is the time it would take a vehicle to pass the link under free flow conditions. Vehicles exit a link according to the following rule: In every time step t, vehicles whose texit has passed are removed from the head of the queue, but only as many as the capacity per time step cap of the link permits, and only if there is enough space on the next link. In (Çetin [2005]), the free flow link travel time of a vehicle is simply $\Delta t = l/v_0$, where l denotes the length of the link and v_0 the free-flow velocity permitted by the link, which is the same for all vehicles. However, the same source already notes that links need not be first-in-first-out queues, but priority queues in which vehicles are sorted by their texit value, and that Δt can, in principle, be any function of the current link condition. If we keep it simple and assign to each link a free-flow velocity $v_0(veh)$ depending on the vehicle type, this already has the following consequences:

1. Faster vehicles pass slower vehicles under non-congested conditions. A link is non-congested if every vehicle leaves the link exactly at the moment when its free-flow travel time has passed. As in the homogeneous case, vehicles do not influence each other.
2. If a vehicle enters a link and hits the end of an already established queue, i.e. if all other vehicles on the link have exceeded their free-flow travel time, it passes no other vehicle but leaves the queue only after all other vehicles have left it.
3. In the general case where some vehicles have exceeded their free-flow travel time and some have not, a faster vehicle will only pass those vehicles which are still in their free-flow phase.

However, even this seemingly simple modification of allowing vehicle-specific free speeds requires the following considerations:

1. The approach would under-estimate the length of the queue. Firstly, in reality, vehicles whose free-flow exit time has expired may still already be part of the queue. As more vehicles become part of the queue, it grows in physical length, so every following vehicle has less space available for free-flow (and being passed). Secondly, one needs to make a decision regarding the queue density. From a traffic flow theoretical perspective, this density is determined by the fundamental diagram of the link, by looking up the congested density corresponding to the outflow of the link. Given that the outflow of the link may change in congested conditions with spillback from downstream links, it might be more parsimonious to just use the maximum density. This typically over-estimates the number of vehicles that are on a link under congested conditions, but it is consistent with the current approach.
2. Accurately computing the position of the end of the queue may use up the computational advantage of the current model. Even with the simple approach, the current simple first-in, first-out data structure needs to be changed into a priority queue. The impact on computational performance will need to be tested.

It remains to be shown to which degree this model can be calibrated to reproduce travel times of a real-world scenario with lane-changing and passing.

5 Conclusion

We discussed some aspects of implementing and adding a freight model to a multi- agent transport simulation which was designed with passenger traffic in mind. Some of them might be of independent interest in other applications: Moving from single autonomous agents to groups of agents which replan as a group with shared resources and obligations is a necessary step to implement shared usage of available cars by families or ride-sharing pools, a feature currently absent from MATSim. Sub-populations with varying planning timeframes can occur whenever the initial demand is generated from different data-sets, for example where weekend travel diaries are available or not available. A traffic flow simulation with heterogeneous vehicle fleets is helpful for scenarios where a large share of road traffic comes from non-car modes, as is common for scenarios in developing countries, but to what degree the interactions between different modes could be captured requires further investigation. The carrier agent class described in the first part of this paper was implemented as the bottom layer of a multi-tier freight market model with the goal of simulating logistics decisions at different levels (Schröder et al. [2011]).

Acknowledgements

Funded in part by the German Research Foundation (DFG) under grant Detaillierte Evaluation verkehrlicher Maßnahmen mit Hilfe von Mikrosimulation.

References

M. Balmer, M. Rieser, K. Meister, D. Charypar, N. Lefebvre, K. Nagel, and K. Axhausen. MATSim-T: Architecture and simulation times. In A. Bazzan and F. Klügl, editors, *Multi-Agent Systems for Traffic and Transportation*, pages 57–78. IGI Global, 2009.

N. Çetin. *Large-scale parallel graph-based simulations.* PhD thesis, Swiss Federal Institute of Technology, Zurich, 2005.

C. Gawron. *Simulation-based traffic assignment.* PhD thesis, Universität zu Köln, 1998.

J. W. Joubert, P. J. Fourie, and K. W. Axhausen. A large scale combined private car and commercial vehicle agent-based traffic simulation. 2009.

S. Schröder, M. Zilske, G. Liedtke, and K. Nagel. Towards a Multi-Agent Logistics and Commercial Transport Model: The Transport Service Provider's View. In *7th International Conference on City Logistics*, Feb. 2011.