Large-Scale Agent-Based Combined Traffic Simulation of Private Cars and Commercial Vehicles

Journal Article**Author(s):**

Joubert, Johan W.; Fourie, Pieter J.; Axhausen, Kay W.

Publication date:

2010-01

Permanent link:

<https://doi.org/10.3929/ethz-b-000158743>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

Transportation Research Record 2168(1), <https://doi.org/10.3141/2168-04>

Large-Scale Agent-Based Combined Traffic Simulation of Private Cars and Commercial Vehicles

Johan W. Joubert, Pieter J. Fourie, and Kay W. Axhausen

The number of independent and interdependent freight actors (firms), the complex supply chain structures between them, and the sensitivity of shipment data are some reasons that the modeling of freight traffic is lagging its public and private transit counterparts. An agent-based approach was used to reconstruct commercial activity chains and simulate them, along with private vehicles, for a large-scale scenario in Gauteng, South Africa. The simulated activities were compared with the actual observed activities of 5,196 vehicles that were inferred from Global Positioning System logs covering approximately 6 months. The results show that the reconstructed activity chains are both spatially and temporally accurate, especially in areas of high activity density. Freight vehicles are a major contributor to traffic congestion and emissions, and the new approach bridges the gap between the person and the state of the art in commercial transport modeling.

Although freight plays an important role in economic development, freight demand modeling has received much less research attention than its passenger counterpart (1). Possible explanations include the many actors (firms) involved, the complex supply chain structures that result in independent and interdependent logistic networks between the firms (2), the numerous (and variety of) commodities concerned, and the confidentiality and reliability of data gathering.

Rodrigue notes that freight and passengers have different spatial dynamics (3), yet they share the same road space. The number of freight vehicles in urban areas may be small compared with total vehicular traffic, but each vehicle has a disproportionate contribution to certain emissions and traffic congestion.

Decision making about freight is usually viewed in literature from the perspective of the firm (4): the emphasis is on achieving optimal or near-optimal performance with regard to its economic efficiency and service quality. At a strategic level, for example, decisions are made about the location of main facilities (production or distribution) and the procurement of resources such as trucks and rolling stock. Tactical decisions may include the design of the service network through route choice and geographic customer segmentation. Oper-

ational decisions are typically highly dynamic and involve decisions about driver and crew scheduling and vehicle dispatching. Companies most often make these decisions within constrained environments.

Transport planning for central decision-making stakeholders, such as government, takes on a different and aggregated perspective. Their decisions need to reflect the simultaneous actions of many transport players. Friedrich et al. identify senders (consignors), freight forwarders, carriers, drivers, and receivers at the various decision-making levels in freight transport (5). They discuss the diverse process steps and decisions that are taken at the various levels and show how choice sets differ for long-haul and regional freight transport.

This paper addresses the problem of simulating disaggregate commercial vehicle chains effectively in both space and time. The foundation for understanding commercial vehicle movement and activities was laid in an earlier work (6). The present paper builds on that foundation, generating intraprovincial commercial vehicle chains for vehicles performing at least 90% of their activities within the study region in Gauteng, South Africa. Spear et al. refer to these chains as internal distribution trips and distinguish them from through trips and internal-external trips (7). This paper simulates the chains in a multiagent transport simulator (MATSim) tool kit along with private vehicle traffic obtained from Fourie and Joubert (8). The majority of road users were accounted for, the exception being the multitude of small paratransit vehicles and through traffic.

To evaluate the time and space accuracy of the simulated activities, an hour-by-hour comparison with the actual observed activity densities was performed for each of the 457 mesozones as described by the geospatial analysis platform (GAP) developed by CSIR Built Environment (9).

The presented approach to generate commercial vehicle activity chains is shown to produce very accurate results, especially in areas of high activity density.

This paper is significant in several respects. The detailed movements of commercial vehicles are accurately model without the need to model the complex and interdependent relationships among commercial actors such as production facilities (shippers), carriers, and receivers. From a wider transport modeling perspective, significant improvement to private car models is shown with the introduction of commercial vehicles.

FREIGHT TRANSPORT MODELING

Most freight transport models are derived from the classical four-step model originally developed for passenger transport. A detailed review of freight transport models, all derived from the classical approach, is provided by De Jong et al. (10).

J. W. Joubert, Industrial and Systems Engineering, University of Pretoria, Pretoria, 0002, South Africa; CSIR Built Environment, P.O. Box 395, Pretoria, 0001, South Africa. P. J. Fourie, CSIR Built Environment, P.O. Box 395, Pretoria, 0001, South Africa. K. W. Axhausen, Institute for Transport Planning and Systems, HIL F 31.3, ETH Zurich, Wolfgang-Pauli-Strasse 15, CH-8093 Zurich, Switzerland. Corresponding author: J. W. Joubert, johan.joubert@up.ac.za; jjoubert@csir.co.za.

Transportation Research Record: Journal of the Transportation Research Board, No. 2168, Transportation Research Board of the National Academies, Washington, D.C., 2010, pp. 24–32.
DOI: 10.3141/2168-04

The present work is cast against the framework proposed by Tavasszy (11), illustrated in Figure 1, which extends the four-step modeling approach to account for decisions and issues relevant to freight modeling.

Production and consumption volumes (in tons) are generated for each location, typically a zone, by using land use and transport interaction, trip generation, facility location, and various freight–economy coupling approaches. Tapio warns that traditional coupling of freight traffic volumes and gross domestic product should be considered with caution (12). Trade values are converted to volume to establish commodity flows, linking locations. Logistics involve choices regarding inventory location and supply chain management. Transportation is challenged with modal choice, intermodal, and transshipment decisions. Discrete choice techniques and trip conversion factors are often used. At this stage, tons are often converted to ton kilometers, or even to the individual consignment level. Kveiborg and Fosgerau suggest a distinction between freight traffic when vehicle kilometer is the unit of measure and freight transport for ton kilometer (13). The most disaggregate level, network and routing, often seeks decision support regarding congestion, tour planning, and city access. Here, chosen techniques include network assignment models and simulation.

Although fairly sophisticated models can be used throughout, many of the approaches handle flows only at an aggregate level (zones) and the detail movement of the freight carriers is not considered (14). Urban tours, in which vehicles make multiple stops, are usually

ignored. Friedrich et al. report on recent models in which activity chain generation is addressed, but none models the behavior at the individual carrier level (5). Wisetjindawat et al. consider commodity flow following a top-down approach that explains commodity movement through the interaction among several freight agents in the supply chain (15). The spatial discrete choice model distinguishes between shippers, receivers, and the relationships among them. Their model addresses the commodity movement (generation and distribution), shipment sizing, carrier choice, and the routing and traffic assignment of the carrier vehicles. Hunt and Stefan present a tour-based microsimulation model of urban commercial movements by using data from an extensive survey of 37,000 activity chains in Calgary, Alberta, Canada (16).

Agent-based modeling techniques, as opposed to four-step variants, allow one to incorporate and embed decision making across the trade, logistics, transportation, and network levels of the framework. Instead of a top-down approach, the individual stakeholders (vehicles, firms, commodities, or industries) have their own autonomous decision-making mandate. All agents are simulated simultaneously, and their interactions with one another and with the environment allows for emergent phenomena that is otherwise lost if a top-down system description approach is followed. Liedtke notes three advantages that agent-based simulation models present in the freight context: (a) through statistical modeling one is able to represent the heterogeneous nature of freight transport actors and objects; (b) normative optimization engines can be incorporated into the individual agents'

FIGURE 1 Conceptual framework of freight transport modeling: (a) production and consumption, (b) trade (scales and sourcing), (c) logistics, (d) transportation, and (e) network and routing (11).

behavioral models, allowing observation of agent reactions to policy measures on different time-horizons; and (c) micro model results can be aggregated *ex post* as desired to support decisions of various nature (17).

Liedtke argues in his agent-based approach that shippers are generally more influential in decision making than are the carriers (17). Whereas the movement pattern of a person (participating in home, work, education, and leisure activities) closely resembles the movement pattern of a (private) vehicle, the situation in freight transportation is more complex. Shippers and recipients make decisions about production, warehousing, and ordering. Once orders are placed, carriers (either in-house or outsourced) deal with the tour-planning decisions. Other actors, such as freight forwarders, may get involved to build transport chains, especially in a multinational context.

Liedtke's INTERLOG model prototype simulates the auction of transport contracts among independent carriers. A simulated population of actors is generated by using economic activity, number of employees, and location from national statistical descriptors. Shippers and customers are then matched, and an exchange of goods is set, expressed at a microscopic level in metric tons per year. The market-interaction module maps how shippers derive shipments and how these shipments are allocated to carriers. Through an experience-based learning process, actors follow behavior-related models to pursue cost minimization, allowing the simulation to settle into a dynamic equilibrium.

Although the autonomous decision-making and learning aspects of Liedtke are attractive (17), these require extensive information that frequently is not available, such as detailed sectoral data about companies, intersectoral commodity exchanges for different transport markets, and the sensitivity of the transport market to price and distance. Also, the INTERLOG model addresses the behavior of only a very small portion—fixed contracts—of total commercial traffic.

The MATSim tool kit is another agent-based approach with experienced-based learning that until now had been applied only to private vehicles (18). MATSim has been applied successfully in Gauteng, South Africa, where access to information is limited (8). This paper aims to extend the existing agent-based simulation capability of the MATSim tool kit to include freight and aims to extend the existing South African implementation.

MULTIAGENT TRANSPORT SIMULATION

The MATSim tool kit allows for agent-based simulation of large-scale transport scenarios. Figure 2 illustrates the controlling procedure of the tool kit. The model is initialized with a synthesized population of agents, each having at least one possible plan. A plan consists of a set of sequential activities, each with a location that is associated with a given network. Sequential activities in the plan are connected with a travel leg, described as a route through the network and an intended mode of transport.

When the model is executed, each agent selects a single plan from its set of plans, and all agents execute their chosen plans simultane-

ously in a mobility simulation. Simulation information is recorded in the form of discrete events such as the start and end of an activity, entering and leaving a link in the network, or waiting to access a link. Simulation events are then interpreted to derive a score for each agent's selected plan (19). During the replanning step, agent plans can be adapted by changing the start time and location of the activities or the mode and routing of the travel legs. Adapted plans are added to the database of plans, and poorly performing plans are discarded. Each agent selects a plan, and the simultaneous execution of plans is repeated (20).

After a specified number of iterations, the simulation terminates and presents statistics such as count comparisons with actual observed counting station data. The higher the number of iterations, the more stable the set of selected plans becomes.

DEMAND GENERATION

For the purpose of this paper, both the private vehicle plans and the network description were taken from the work of Fourie and Joubert (8). Private persons' activity chains are limited to home–work–home to account for the majority of trips during peak periods of the day. This section describes the procedure followed to generate an initial demand for commercial vehicles.

From the original 31,041 vehicles used by Joubert and Axhausen (6) to describe the activity chains of commercial vehicles, those that spend 90% or more of their activities within the province of Gauteng were selected. For those 5,196 vehicles, activities and activity chains were extracted. Activities were classified as either major, lasting more than 5 h, or minor, lasting less than 5 h. Major activities can be thought of as depots or home locations for commercial vehicles. Minor activities include any (and all) activities for which vehicles are switched off for periods of less than 5 h. When the vehicle is switched on again, the activity is considered finished. In the activity extraction, false starts are eliminated, and all consecutive activities at the same location are merged. Examples of minor activities are vehicle refueling, stopping to perform a service, and loading and off-loading. Figures 3b and 3c present the kernel density estimation of the major and minor activities in Gauteng, respectively. Conceptually, kernel density estimation fits a smooth surface over the observed activities. For the estimations in Figure 3, a grid size of 500 by 500 m and a search radius of 5,000 m were used. This was achieved by using the ESRI ArcGIS implementation of a quadratic kernel function, popularized by Silverman (21).

The rich activity set was used to build conditional probabilities for four chain parameters: the start location, denoted by M^l ; start time, M^t ; chain duration, M^d ; and number of activities, M^n . To create a plan for a commercial vehicle i , the first (major) activity is created and assigned a location sampled from the kernel density estimation of Figure 3b, M_i^l . Although it is known (6) that commercial vehicles, albeit the minority, may start and end their chains at different locations, it is assumed in this paper that chains start and end at the same location. The chain start time given M_i^l was sampled; the conditional

FIGURE 2 MATSim controller (18).

FIGURE 3 Demarcation and activity densities for Gauteng, South Africa: (a) geospatial analysis platform mesozones, (b) kernel density estimation of major activities, and (c) kernel density estimation of minor activities.

sampling is denoted $M_i^l | M_i^l$, the number of activities $M_i^n | (M_i^l, M_i^l)$, and the chain duration $M_i^d | (M_i^l, M_i^l, M_i^n)$.

To fill chain i , M_i^n minor activities were sampled from the kernel density estimation shown in Figure 3c and were sequenced randomly, and their start times were spaced equally between M_i^l and $M_i^l + M_i^d$. The chain was completed by adding another major activity with the same location as M_i^l but with a start time of $M_i^l + M_i^d$.

Since, per Joubert and Axhausen (6), nearly 50% of commercial vehicle activity chains observed has duration in excess of 24 h, the generated chains in this study had to be adapted to fit within the 24-h simulation time imposed on the simulation scenario. Figure 4 depicts the process followed in this paper. The procedure can handle any time window denoted by $[t^{\min}, t^{\max}]$. For this paper, the time window is specified in seconds for a complete 24-h day as $[0; 86400]$.

FIGURE 4 Procedure to wrap activity chain artificially around given time window, with independent chains as a result, each starting and finishing within time window: (a) original activity chain exceeding specified time window, (b) introduction of dummy major activity at end of time window (dummy shares activity type of first portion of chain and ending time and location with first activity of subsequent portion) and (c) resulting two activity chains that depict two artificial and independent commercial vehicles.

Consider the activity chain in Figure 4a starting with a major activity, M_1 , within the first 24-h window. The activity chain contains n minor activities, denoted by $m_1, \dots, m_n$. The activity chain is completed with major activity M_2 .

Say for the purpose of this example that the 24-h window ends somewhere between minor activities m_e and m_{e+1} . The activity chain is split between minor activities m_e and m_{e+1} with the introduction of a dummy major activity denoted d (Figure 4b). The dummy shares its activity type with the first activity, M_1 . The location of d is the same as the location of m_{e+1} . The major activity is added at the end of the first segment of the chain as d with start time 24:00:00 (or 86,400 s). This has the result that traveling from activity m_e to activity m_{e+1} is forced to occur at the end of the first segment of the chain. It is assumed that the impact of whether traveling occurs as the last leg of the first portion or the first leg of the second portion is negligible.

Once the dummy major activity is added, all the activity start times following the dummy activity are adjusted by simply subtracting 24 h, resulting in the two independent chains illustrated in Figure 4c, where the starred version of the activities denote the activities with adjusted start and end times. Each of the resulting chains is treated as an independent agent without weighing of the chains. Should the remaining portion of the chain still exceed the time window, $t^{\max}$, the procedure is repeated on the remainder.

SIMULATION RESULTS

To account for the variation as a result of the random sampling of activities from the kernel density estimation, six populations of freight agents were generated, each having a different random seed. Each population was added to the same population of private car agents to derive six unique scenarios. The scenarios were each simulated for 100 iterations. This section reports on the simulated results as the average over the six scenarios—they were not significantly different. To measure the repeatability of the six scenarios, the trip durations for all agents were compared across 13 5-min intervals for each scenario, with the standard deviation calculated during each

interval. The average standard deviation across all intervals was then calculated as $\sigma = 1.52$.

The private and commercial vehicle plans were combined and simulated for 100 iterations. To evaluate the activity location and timing of the commercial activities, the discrete events produced during the 100th iteration were analyzed. Activity start times and the location of all minor activities were of interest.

To determine if the simulated commercial activities occur at the right geographic location at the right time of day, activities were geographically aggregated to the GAP mesozone level (Figure 3a) so as to be comparable with the observed activities reported by Joubert and Axhausen (6). Temporally, the events were aggregated to the hour of the day in which they occurred. Figure 5 shows, as an example, a geographic comparison of the actual versus simulated activity counts for all activities in the ninth hour, that is, between 08:00:00 and 08:59:59. The actual observed activities span approximately 180 days, so both the actual and the simulated activity densities were normalized to the highest activity count found in the respective sets. As an example, consider the process of normalizing the simulated data set. The reference is the highest activity count found in any zone at any time of the day. This was found to be 1,167 for Zone 4917 in the 10th hour. If then 245 activities were simulated to start in Zone 4570 in the ninth hour, the normalized density for Zone 4570 in the ninth hour would be $245/1167 \approx 21\%$.

This approach is valid because GAP Zone 4917 had the highest activity counts in the observed and all the simulated scenarios. Hence, densities were normalized to the same zone.

If geographic location is less important, the simulated results may also be presented as in Figure 6. Each point represents a unique zone–hour combination. The error appears to be greatest in low-activity zones.

Next the accuracy of the simulated activities for different hours of the day is considered. Figure 7 shows box-and-whisker plots for the normalized difference between the simulated number of activity starts during each hour and the observed number of activity starts. Positive values indicate instances in which simulated activity counts were higher than the actual observed activity counts, and negative

FIGURE 5 Comparison of actual versus simulated activities for ninth hour: (a) actual observed activity densities, normalized to highest actual activity count over 24-h period; (b) simulated activity densities, normalized to highest simulated activity count; and (c) absolute difference between normalized actual and normalized simulated activity counts.

FIGURE 6 Comparison of simulated versus actual activity densities for all hours across all GAP zones. From top to bottom, linear bands represent 2:1, 1.5:1, 1:1, 1:1.5, and 1:2 ratios of simulated to observed normalized activity densities.

values indicate where they were lower. All mesozones are accounted for in Figures 7a and 7c and thus the comparisons are referred to as unbiased. Many of the mesozones, however, had no observed activity starts and were nonsimulated. In itself, this is arguably good, but the large number of zero-observed–zero-simulated mesozones obscures the results somewhat. Removing all the zero-observed–zero-simulated mesozones yields a result that is biased toward simulation differences, presented in Figures 7b and 7d.

For both the unbiased and biased cases, the box widths, the inter-quartile ranges (IQRs) are very tight, and thus Figures 7c and 7d zoom in on the range $[-2, 2]$.

In both plots, the whisker ends indicate the lowest datum still within 1.5 times the IQR of the lower quartile and the highest datum within 1.5 times the IQR of the upper quartile. With all the medians located at (or very near) the zero datum and the very narrow box widths, the commercial activities are accurately simulated, at least at a time–space granularity of hours and mesozones. The model overpredicts commercial activities during the period 08:00 to 12:00 and underpredicts from 15:00 to 00:00. Possible assignable causes of overestimating are that only home–work–home trips are modeled for private vehicle users, essentially freeing up a large portion of road capacity during off-peak hours. Also, the effective road network considered by commercial vehicles is probably more limited than the complete network currently made available to all road users, allowing for more accessibility than is realistic.

Two approaches were used to examine the effect of the inclusion of commercial vehicles on the private cars that were modeled in the simulation. First, each road link was associated with a mesozone, and the average speed of a mesozone was calculated as the average speed of private vehicles across all links associated with that mesozone. If no activity occurred during a specific hour on the link, the link's free speed was used. Only when no network link was associated with the

mesozone was no average speed calculated. In the second approach, the total trip time of each private vehicle agent was measured and the average of all agents was calculated. A comparison was performed of these measures when only private vehicles were modeled and the model in which commercial vehicles were included.

The overall impact across the study area is shown in Figure 8. With slower traffic expected in the economically dense areas, the two economic centers (Tshwane in the north and the larger Johannesburg–Ekurhuleni in the south) are distinguishable with their lighter shade, whereas the introduction of commercial vehicles shows only a slight decrease in average speeds. The slight change may be because commercial vehicles, in absolute numbers, represent only about 5% of the total traffic volume.

An hour-by-hour comparison for a single mesozone (Zone 4917) is shown in Figure 9a. The impact that commercial vehicles have on the average speed of private vehicles again is not significant. Except between 13:00 and 14:00, slower average speeds are reported for all hours.

A different view of the same analysis is reported in Figure 9b, where the overall trip times also did not change dramatically. There are fewer short trips and more trips exceeding 20 min. Since only internal distribution trips were simulated, which accounted for only an estimated 35% of Gauteng commercial traffic (6), the impact will become higher when through traffic is incorporated. Still, the inclusion of commercial vehicles in the agent-based simulation model does have the desired (expected) effect: slowing of traffic, and thus increased travel time.

CONCLUSION

The activity chains generated and simulated in this paper were shown to be a geographically and temporally accurate representation of observed vehicle chains. A number of future extensions to this paper are proposed.

The next step is to generate through traffic and trips for which less than 90% of the activities occur in the study area. Once these trips are accounted for, it will be appropriate to compare the simulated commercial traffic against actual traffic counts. This paper shows that commercial vehicle behavior can be modeled for replanning route and activity schedules. The impact on policy evaluation should be clear: a base case can be established by modeling the status quo and validating it by using actual vehicle counts. If we want to evaluate how commercial vehicle movement changes as a result of, say, a road-pricing or lane-reservation policy, we can include those policy measures on the network, resimulate the scenario, and compare the vehicle movement (activity densities, activity timing, average link speed, etc.) against that of the base case. The simulated commercial vehicles will autonomously replan their routes and activity schedules to maximize their expected utilization. In doing so, the simulated agents give a reasonable indication of what can be expected from real autonomous decision-making commercial vehicles in maximization of their utility.

In accordance with Tavasszy (11), this paper suffers from the empirical challenge of disentangling light-goods vehicles from heavier vehicles. Vehicle size does play a role in congestion impact. Activity chains generated in this paper do not distinguish between vehicle sizes or activity types (commercial freight versus commercial services).

The freight trip chains in this paper were based on and sampled from observed trip chains. New chain structures hence cannot emerge from the current model. This paper's approach of randomly drawing attribute values characterizing the commercial chains can be extended.

FIGURE 7 Box-and-whisker plot showing normalized difference between simulated and actual number of activity starts for different hours of day. Positive normalized difference indicates simulated values higher than actual observed values. Whisker ends indicate lowest datum still within 1.5 times IQR of lower quartile and highest datum within 1.5 times IQR of upper quartile. Circles outside whisker extents are considered outliers. (a) Unbiased comparison including all mesozones, (b) biased comparison with zero-observed-zero-simulated mesozones removed, (c) unbiased comparison zoomed in, and (d) biased comparison zoomed in.

FIGURE 8 Comparison of average speed and speed difference for private vehicles aggregated to mesozone level. Speeds are for ninth hour: (a) only private vehicles modeled, (b) private and commercial vehicles modeled, and (c) difference.

For example, the random draw could be replaced with a decision engine capable of simulating and predicting the attribute values based on underlying factors, such as commodity type and the land use of activities. This would require that the underlying factors be adequately described by available data sources.

In the absence of any additional information to the Global Positioning System (GPS) logs, one may be tempted to consider supply chain network structures to be intractable. Yet, with the inferred knowledge of major and minor activity locations, one may be able to analyze the activity chains with the underlying land use data and infer the network structures. Contributions such as those by De Jong and

Ben-Akiva (22), Liedtke (17), and Wisetjindawat et al. (15) may prove useful for inferring network and supply chain structures.

The GPS logs used to infer activities include both service and freight vehicles. Thus it can be argued that since samples are drawn from all observed commercial trip chains, the model results account for empty trips, something often omitted in other modeling approaches. Also, activities may include collections from production locations, deliveries to receivers, and activities at hub locations. Friedrich et al. noted the latter as lacking from network models for freight (5). Empty trips and hub activities, although accounted for, will not be distinguishable within the results here.

FIGURE 9 Comparisons of speed and travel time of private vehicles when only private vehicles are modeled versus when commercial traffic is combined with private vehicles: (a) average speed in single zone (Zone 4917) and (b) overall trip travel time.

ACKNOWLEDGMENT

The South African National Research Foundation partially funded this research.

REFERENCES

1. Ortúzar, J. d. D., and L. G. Willumsen. *Modelling Transport*, 3rd ed. John Wiley and Sons, Chichester, England, 2001.
2. Hesse, M., and J. Rodrigue. The Transport Geography of Logistics and Freight Distribution. *Journal of Transport Geography*, Vol. 12, No. 3, 2004, pp. 171–184.
3. Rodrigue, J. Transport Geography Should Follow the Freight. *Journal of Transport Geography*, Vol. 14, No. 5, 2006, pp. 386–388.
4. Crainic, T. G., and G. Laporte. Planning Models for Freight Transportation. *European Journal of Operational Research*, Vol. 97, No. 3, 1997, pp. 409–438.
5. Friedrich, M., T. Haupt, and K. Nökel. Freight Modelling: Data Issues, Survey Methods, Demand and Network Models. In *Moving Through Nets: The Physical and Social Dimensions of Travel* (K. W. Axhausen, ed.), Elsevier, Amsterdam, Netherlands, 2007, pp. 249–277.
6. Joubert, J. W., and K. W. Axhausen. Inferring Commercial Vehicle Activities in Gauteng, South Africa. *Journal of Transport Geography* (forthcoming).
7. Spear, B., G. Giaimo, T. R. Curlee, and K. Neels. Freight Modeling: State of the Practice. In *Conference Proceedings 40: Freight Demand Modeling: Tools for Public-Sector Decision Making* (K. L. Hancock, ed.), 2008, Transportation Research Board of the National Academies, Washington, D.C., pp. 21–26.
8. Fourie, P. J., and J. W. Joubert. *Large-Scale Agent-Based Transport Simulation of South African Commuters*. University of Pretoria, South Africa, 2009.
9. *Geospatial Analysis Platform and NSDP Spatial Profiles*. CSIR Built Environment. 2009. <http://www.gapweb.co.za>. Accessed March 13, 2009.
10. De Jong, G., H. Gunn, and W. Walker. National and International Freight Transport Models: An Overview and Ideas for Future Development. *Transport Reviews*, Vol. 24, No. 1, 2004, pp. 103–124.
11. Tavasszy, L. Freight Modeling: An Overview of International Experiences. In *Conference Proceedings 40: Freight Demand Modeling: Tools for Public-Sector Decision Making* (K. L. Hancock, ed.), Transportation Research Board of the National Academies, Washington, D.C., 2008, pp. 47–55.
12. Tapio, P. Towards a Theory of Decoupling: Degrees of Decoupling in the EU and the Case of Road Traffic in Finland Between 1970 and 2001. *Transport Policy*, Vol. 12, No. 2, 2005, pp. 137–151.
13. Kveiborg, O., and M. Fosgerau. Decomposing the Decoupling of Danish Road Freight Traffic Growth and Economic Growth. *Transport Policy*, Vol. 14, No. 1, 2007, pp. 39–48.
14. Fernández, J. E., J. de Cea, and A. Soto. A Multi-Modal Supply-Demand Equilibrium Model for Predicting Intercity Freight Flows. *Transportation Research Part B: Methodological*, Vol. 37, No. 7, 2003, pp. 615–640.
15. Wisetjindawat, W., K. Sano, and S. Matsumoto. Commodity Distribution Model Incorporating Spatial Interactions for Urban Freight Movement. In *Transportation Research Record: Journal of the Transportation Research Board*, No. 1966, Transportation Research Board of the National Academies, Washington, D.C., 2006, pp. 41–50.
16. Hunt, J., and K. Stefan. Tour-Based Microsimulation of Urban Commercial Movements. *Transportation Research Part B: Methodological*, Vol. 41, No. 9, 2007, pp. 981–1013.
17. Liedtke, G. Principles of Micro-Behavior Commodity Transport Modeling. *Transportation Research Part E*, Vol. 45, No. 5, 2009, pp. 795–809.
18. MATSim Project. 2009. <http://www.matsim.org>. Accessed June 24, 2009.
19. Charypar, D., and K. Nagel. Generating Complete All-Day Activity Plans with Genetic Algorithms. *Transportation*, Vol. 32, No. 4, 2005, pp. 369–397.
20. Raney, B., and K. Nagel. An Improved Framework for Large-Scale Multi-Agent Simulations of Travel Behaviour. In *Towards Better Performing European Transportation Systems* (P. Rietveld, B. Jourquin, and K. Westin, eds.), Routledge, London, 2006.
21. Silverman, B. W. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall, London, 1986.
22. De Jong, G., and M. Ben-Akiva. A Micro-Simulation Model of Shipment Size and Transport Chain Choice. *Transportation Research Part B: Methodological*, Vol. 41, No. 9, 2007, pp. 950–965.

The Freight Transportation Planning and Logistics Committee peer-reviewed this paper.