

The 3rd International Workshop on Agent-based Mobility, Traffic and Transportation Models,
Methodologies and Applications (ABMTRANS)

Heterogeneous tolls and values of time in multi-agent transport simulation

Kai Nagel^{a,b,*}, Benjamin Kickhöfer^b, Johan W. Joubert^a

^aIndustrial and Systems Engineering, University of Pretoria, Pretoria, South Africa

^bTransport Systems Planning and Transport Telematics, TU Berlin, Berlin, Germany

Abstract

In evolutionary algorithms, agents' *genotypes* are often generated by more or less random mutation, followed by selection based on the fitness of their *phenotypes*. This paper shows that elements of this principle can be applied in multi-agent transport simulations, in the sense that a router, when faced with complex interactions between heterogeneous toll levels and heterogeneous values of time, can resort to some amount of randomness rather than being able to compute the exact best solution in every situation. The computational illustrations are based on a real world case study in the province of Gauteng, South Africa.

© 2014 Published by Elsevier B.V. Open access under [CC BY-NC-ND license](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Selection and Peer-review under responsibility of the Program Chairs.

Keywords: transport simulation; multi-agent simulation; evolutionary algorithm; value of time

1. Introduction: Consistency between plans and scoring, or between genotype and fitness function

A co-evolutionary algorithm^{1,8,9,19}, mapped towards travel behavior, could be a variant of the following:^{3,18}

1. **Initiation:** Generate at least one plan for every person, and select one plan per person.
2. **Iterations:** Repeat the following steps many times:
 - (a) **Scoring:** Obtain a score for every person's selected plan, by executing all selected plans, one for each person, simultaneously in a simulation (synthetic reality), and deriving a performance measure from this.
 - (b) **Selection:** Decrease the occurrences of "bad" solutions.
 - (c) **Construction of new plans/innovation/choice set generation:** For some of the persons, generate new plans, for example as "best replies" (e.g. a fastest path based on the last iteration's link travel times) or as mutations of existing plans (e.g. small departure time changes).
 - (d) **Choice:** Select, for every person, one of the plans.

* Corresponding author. Tel.: +49-30-314-23308 ; fax: +49-30-314-26269.

E-mail address: nagel@vsp.tu-berlin.de

An important feature of these algorithms is that the innovation is separated from the scoring. Scoring is based on a simulation of the system (“synthetic reality”) whereas innovation can be based on arbitrary (mental) models of that system. For example, a traveler could consider a new route by looking at a map, but when actually trying it out, it may have too many potholes or too many stop signs. The separation of innovation and scoring is explicitly recommended⁷, since it modularizes the software system: The designers of the synthetic reality can concentrate on a realistic model, and the designers of the innovative modules can concentrate on their algorithmic aspects. There may, in fact, be many different innovative algorithms, e.g. suggesting a new path, a new time structure, a new mode, a new location for an activity, etc. Each of these modules has different modeling and algorithmic challenges.

The overall approach is derived from Arthur (1994)¹, which in turn has learned from Hrabner et al. (1994)⁹, Palmer et al. (1994)¹⁹, Holland (1992)⁸, and others. In the sense of Russel and Norvig (2010)²⁰, the MATSim agents can be classified as utility-based* learning agents; they possess, however, very little reactivity while they move in the traffic flow simulation: all adaptation is moved *between* repeated runs of the traffic flow simulation. The interaction between agents is indirect: Agents react to delays caused by congestion. The agents have knowledge (mostly: their score) based on their own experience and nothing else, but they get “proposals” for better plans by external modules, for example a router which suggests routes which may be better than the ones that the agent already knows.

An important question with such approaches concerns the consistency between the innovative modules and the scoring. For example, if an innovative module *never* makes a suggestion which is good for the synthetic person, then the synthetic person is also unable to select a good option. If, however, an innovative module sometimes makes suggestions which are good for the person and sometimes not, then the person can keep the good suggestions in his or her choice set, and reject the other ones.

This paper will look at the problem in the specific context of heterogeneous tolls and heterogeneous values of time. The question is how well the innovative module, here the toll-sensitive router, needs to be adapted to the scoring.

2. Toll in multi-agent simulations of traffic flow and travel behavior

2.1. The charging of toll and a scoring function including toll

The agent-based simulation needs to make sure that heterogeneous toll levels and heterogeneous values of time are correctly implemented. For toll, this is rather straightforward: The simulation can be set up such that every time a vehicle enters a tolled road link, the toll level is looked up in a database, and charged to the driver. Toll levels might vary by time of day,¹⁵

or by vehicle type and traffic state.¹⁴

For the values of time, it is more involved, since scoring full days rather than single trips involves an opportunity cost of time for activities, or marginal utility of time as a resource.¹⁰ Since the computational experiments are done with MATSim⁴, the following discussion concentrates on the MATSim scoring function; alternative approaches will often have a similar structure. Scoring in MATSim works as follows: At the end of the simulated day, each synthetic traveler adds up all scoring function results belonging to the different elements of its daily activities:⁵

- (i) Each **activity** i adds $\beta^{prf} \cdot t_{typ(i)} \cdot \ln \frac{t_i}{t_{0(i)}}$, where β^{prf} is an approximation to the marginal opportunity cost of time (see below), $t_{typ(i)}$ the typical duration of activity i , t_i its actual duration in the simulation, and $t_{0(i)}$ corresponds to an alternative-specific constant per activity type, but has no influence as long as the chain of activities is fixed.
- (ii) Each **trip** from one activity i to the following activity $i + 1$ adds a typically negative value $\beta_{mode(i)}^{trv} t_i^{trv}$, where t_i^{trv} is the trip’s duration, and $\beta_{mode(i)}^{trv}$ is a mode-dependent parameter whose relation to the utility of travel time savings will be explained below.
- (iii) The (dis)utility of the **toll** payments τ is added as $-\beta^{mon} \tau$, where β^{mon} is the marginal utility of money.

* Note that ‘score’ is the technical term in most MATSim-related publications. ‘Utility’ is the common expression in economics. In this paper, both terms refer to the same absolute value.

2.2. Understanding utility functions for full days

The marginal opportunity cost of time (in units of utility), or **marginal utility of time as a resource**, of activity i can be computed from considering small variations of its actual duration:

$$mUoT = \frac{\partial}{\partial t_i} \left(\beta^{prf} \cdot t_{typ(i)} \cdot \ln \frac{t_i}{t_{0(i)}} \right) = \beta^{prf} \cdot t_{typ(i)} \cdot \frac{1}{t_i}. \quad (1)$$

When travel times between activity locations are constant, then the optimal allocation of time between activities is reached when all marginal utilities of time are the same:[†] $\beta^{prf} \cdot t_{typ(i)} \cdot \frac{1}{t_i} \stackrel{!}{=} \beta^{prf} \cdot t_{typ(j)} \cdot \frac{1}{t_j} \ (\forall i, j)$. Rearranging yields $\frac{t_i}{t_j} = \frac{t_{typ(i)}}{t_{typ(j)}}$, that is, at the agent's optimum, the ratio of two activity times is the same as the ratio of their typical durations. This provides an intuitive way to provide a first set of plausible parameters by setting typical durations of activities. Also, if a synthetic person is under time pressure, t_i will be smaller than $t_{typ(i)}$, and thus the marginal utility of time larger than β^{prf} . This makes much intuitive sense.

Now let us assume that a trip can be made faster by ϵ , and the time gain is given to the subsequent activity. The marginal utility change (= **marginal utility of travel time savings**) is

$$mUTTS = \frac{\partial}{\partial \epsilon} \left(\beta^{trv}_{mode(i)} (t_i^{trv} - \epsilon) + \beta^{prf} \cdot t_{typ(i+1)} \cdot \ln \frac{t_{i+1} + \epsilon}{t_{0(i+1)}} \right) \stackrel{\epsilon \rightarrow 0}{=} -\beta^{trv}_{mode(i)} + \beta^{prf} \cdot \frac{t_{typ(i+1)}}{t_{i+1}}. \quad (2)$$

That is, the marginal disutility of traveling is composed of two contributions:

- the base marginal utility (or opportunity cost) of time as a resource. This is the same value for all modes, but it depends on time pressure.
- a mode-dependent marginal offset to that base marginal utility. If traveling is more convenient than “doing nothing”, then $\beta^{trv}_{mode(i)}$ is positive, otherwise negative. Different modes will generally have different values.

The (marginal) **value of travel time savings** is obtained by dividing mUTTS by the marginal utility of money, β^{mon} :

$$mVTTS = \frac{mUTTS}{\beta^{mon}} = \frac{-\beta^{trv}_{mode(i)} + \beta^{prf} \cdot \frac{t_{typ(i+1)}}{t_{i+1}}}{\beta^{mon}}. \quad (3)$$

2.3. A toll-sensitive router

The MATSim router will generate a good route r according to an objective function $f(r)$. $f(r)$ can consider minimization of a single attribute, such as (expected) travel time t_r , or of a combination of attributes, such as

$$f(r) = \alpha \cdot t_r + \beta \cdot \tau_r, \quad (4)$$

where τ_r is the (expected) toll along route r . The parameters α and β need to be determined correctly from a linearization of the scoring function. According to Eq. (2) and adding toll, the correct form is

$$f(r) = mUTTS \cdot t_r + \beta^{mon} \cdot \tau_r = \left(-\beta^{trv}_{mode(i)} + \beta^{prf} \cdot \frac{t_{typ(i+1)}}{t_{i+1}} \right) t_r + \beta^{mon} \cdot \tau_r. \quad (5)$$

There are several challenges to get this right: (1) There needs to be a universal understanding that the above is the correct mathematical form. (2) The information needs to be correctly shared between the scoring function and the router module. (3) It may well happen that at least β^{mon} and possibly other parameters are person-specific.¹³ (4) The expected duration of the following activity, t_{i+1} , needs to be known to the router. All of these issues can be resolved by diligent conceptual work, followed by diligent software engineering. It is, however, our experience that such an approach is not very robust. Maybe more importantly, it goes against the design idea of the co-evolutionary algorithm: As explained in Sec. 1, it is in fact *desirable* that the mental modules, and the router is one of them, are conceptually decoupled from the scoring. In consequence, it makes sense to consider alternative approaches to the problem.

[†] Obtained e.g. from optimization under the constraint that the length of the day is fixed; in micro-economics, it corresponds to the statement that a budget is optimally allocated if the marginal utilities of all products with respect to changes in budget allocation are the same.

2.4. A diversity generating router

An approach that comes to mind from an evolutionary computing perspective is to increase the variability of the solutions that the router generates^{17,21}, e.g. by varying the ratio between α and β in Eq. (4) each time the router is called. In this way, effectively a different mVTTS is used every time an agent obtains a new route. If that mVTTS is consistent with the synthetic person's true mVTTS, then the route will probably be a good one for the synthetic person, it will obtain a high score, and be retained in the synthetic person's memory.

Technically, α is left fixed at $-\beta_{mode(i)}^{trv} + \beta^{prf}$, while β is drawn from a log-normal distribution realized as $\beta = \beta^{mon} \cdot \exp(\sigma \cdot Z) / \exp(\sigma^2/2)$, where Z is drawn from a Gaussian distribution with mean zero and variance one using the `nextGaussian()` method in Java.²² The log-normal distribution was selected since it generates values between 0 and ∞ with a strong focus around one. σ is a width parameter; the denominator $\exp(\sigma^2/2)$ brings the expectation value of the expression to β^{mon} . The approach can be seen as a randomized version of finding Pareto-optimal paths^{6,16}, and thus an approach to generate alternative routes².

2.5. Scoring for freight

In the study considered here, freight follows the same scoring (= utility) function as private transport, i.e. items (i) to (iii) above. More realistic freight-specific scoring functions need to be considered in future work. For the present study, the simulation is set up such that only re-routing is switched on as choice dimension, and freight activities are fixed by their end times, which are the same as the departure times for the following trips. Thus, travel time savings are converted into score improvements by multiplying them with Eq. (2).

2.6. Heterogeneous values of time

From Eq. (3) one obtains for road $\beta^{mon} = (-\beta_{road(i)}^{trv} + \beta^{prf} \cdot \frac{t_{typ(i+1)}}{t_{i+1}}) / mVTTS \approx (-\beta_{road}^{trv} + \beta^{prf}) / mVTTS$, where the approximation holds when the actual duration, t_i , is similar to the typical duration, $t_{typ(i)}$. This is used to determine utilities of money that reflect externally given values of time:

$$\beta_{class}^{mon} := \frac{-\beta_{road}^{trv} + \beta^{prf}}{mVTTS_{class}}. \quad (6)$$

For the simulations, the MATSim default values $\beta_{road}^{trv} = -6/h$, and $\beta^{prf} = 6/h$ are used, and from discussions with SANRAL[‡] $mVTTS_{class} = \text{ZAR } 440/h$ for class C vehicles/drivers, ZAR 220/h for class B, and ZAR 110/h for all others.[§] The actual mVTTS is now given by inserting Eq. (6) into Eq. (3):

$$mVTTS_{a,i} = \frac{-\beta_{road}^{trv} + \beta^{prf} \cdot \frac{t_{typ(a,i+1)}}{t_{a,i+1}}}{-\beta_{road}^{trv} + \beta^{prf}} \cdot mVTTS_{class(a)}. \quad (7)$$

That is, each trip's actual mVTTS will be different, not only depending on the β_x which are the same for all drivers, or on the base value $mVTTS_{class(a)}$ of the class to which the agent a belongs, but also depending on the type and actual duration of the activity immediately after the trip.

3. Illustrative scenario: Gauteng Freeway Improvement Project

3.1. Context and network

The Gauteng Freeway Improvement Project (GFIP) is a phased upgrade and implementation of a 560km freeway network, financed through a toll collected by an open-road electronic system ("e-toll") using 45 overhead gantries.

[‡] South African National Roads Agency Limited

[§] 1 ZAR (South African Rand) \approx 0.0666 EUR. Exchange rate on 04.02.2014. The resulting VTTS are then approximately 29.00 EUR, 14.50 EUR, and 7.25 EUR, respectively.

The network on which the simulation is run is derived from OpenStreetMap (www.osm.org) and consists of 86 567 nodes and 235 854 links.

3.2. Population

The synthetic population that represents the initial demand consists of four road user groups:

1. 1 776 980 **commuters** traveling by private car, 201 810 coming from outside the province. Only *home* and *work* activities are modeled. These agents are all assumed to be *Class A* light vehicles.
2. 2 090 **public transport buses**. Agents plans represent the vehicle trips, not the individual journeys of the patrons. All buses are assumed to be *Class B* vehicles.
3. 81 330 **paratransit (minibus) vehicles**. All minibuses are assumed to be *Class A*.
4. 100 000 **commercial vehicles** with activity chains.^{11,12} Half of the vehicles are assumed to be *Class B*, and the other half *Class C* heavy vehicles, consistent with vehicle counts from SANRAL.

The synthetic population totals 1.96-million agents. For computational speed reasons, a randomly selected 1% of this population was used, and network capacities were reduced accordingly in the MATSim queue model.

3.3. Toll and discount structure

The e-toll scheme is based on base rates of ZAR 0.70/km, 2.10/km, and 4.20/km for Classes A, B, and C, respectively. A discount of 25% for e-tag holders is offered. For commuters within the province we assume an e-tag penetration of 40%; for bus, 50%; paratransit, 40%; commercial vehicles 60%; and external commuters, 25%. Period discounts are specified on an hour-by-hour basis, and range between 0% and 20% for a weekday. A public transport discount is offered to both buses (55%) and paratransit (30%). Although a frequency discount is also offered, it was not considered in this paper since the current agent-based model only represents a single full day. As stated earlier, toll is calculated whenever an agent enters a tolled link in the network. Both the time of day and the agent ID are known, so one is able to accurately determine the correct base rate. The relevant discounts can then be calculated and applied to the base rate.

4. Results

All simulations start from a base case provided by earlier studies.¹² That scenario came with equilibrated routes; the toll which vehicles would have paid with these routes will be used as a baseline. Fig. 1 shows on the left the cumulative number of commercial vehicles that would pay a certain amount of toll; for example, if the e-toll would not lead to re-routing (red line), then 82'000 commercial vehicles would pay a toll of 500 ZAR or less.[¶] Fig. 1 (right) shows the resulting cumulative revenue. For example, if the e-toll would not lead to rerouting (red line), then the toll income from all vehicles paying 500 ZAR or less would be 12 million ZAR. The curve levels out at the complete toll revenue, which would be roughly 13.5 million ZAR if nobody would re-route because of the toll.

This is now followed by a sequence of numerical experiments. Each experiment is run for 100 iterations to achieve a relaxed state where agents' plans are adapted to the traffic conditions created. At the end of each iteration, 10% of agents adapt their routes, while the remaining 90% switch between plans according to a process which converges to a logit model.¹⁸ At iteration 80, the "innovative" rerouting module is switched off, and agents thus only switch between their existing plans. This is meant to "cool down" the simulation by fixing each agent's choice sets.

The specific experiments and their most important results are as follows:

1. At first, the e-Toll is switched on and the simulation is run "as is", i.e. with higher tolls for commercial vehicles, but neither adjustment of the scoring function nor adjustment of the router. – This results in a large number of commercial vehicles avoiding the toll (green lines).

[¶] All numbers in the presentation of the results are re-scaled to 100% travel demand.




Fig. 1. Results for commercial vehicles from different runs. LEFT: Cumulative number of commercial vehicles paying toll as a function of their toll payments. RIGHT: Cumulative revenue from commercial vehicles of the different runs as a function of their toll payments.

Table 1. Average scores from executed plans of commercial vehicles for the different numerical experiments.

run	av. executed score
scoring corrected	188.25
+ router with randomness ($\sigma = 1$)	198.40
+ more randomness ($\sigma = 3$)	209.01
Comparison: manually calibrated router	196.51

- Next, the scoring function for the commercial vehicles is adjusted to higher values of time. This reduces the toll avoidance somewhat (blue lines).
- Next, the router is randomized as explained in Sec. 2.4, with $\sigma = 1$. This reduces the toll avoidance even more (pink lines).
- Increasing the level of randomness by setting $\sigma = 3$ leads to even more commercial toll payers (cyan line in Fig. 1 left), and even to higher toll revenues (cyan line in Fig 1 right).
- Finally, for comparison each user's base value of time β_{class}^{mon} is manually inserted into the router. This leads to the brown lines. Clearly, this approach leads to fewer toll revenues than the “randomizing” router.

The question now is why the manually calibrated router predicts considerably less toll revenue than the strongly randomizing router (Fig. 1 right brown vs. cyan), while the total number of toll payers is quite similar (Fig. 1 left). A possible reason is that a router that generates a diversity of options may find “better” solutions than those found by a manually calibrated, deterministic, router. The reason lies in the fact that the latter is operating on approximations: First of the network, e.g. using link travel times aggregated into 15 min time bins. Therefore, even a router which is correctly adapted to the scoring function may not find the best solution. Second, the router assumes $\beta^{prf} - \beta^{trv}$ as marginal disutility of time, as explained in Sec. 2.4. However, as also explained in that section, the true marginal disutilities depend on the time pressure inherent in each agent's plans. That is, even agents with the same scoring function may have different marginal disutilities of time. An even more involved router might be able to pick up on those. The implementation is rather difficult, since one needs to consider both the preceding and the following activity as well as any time constraints such as opening times.

This explanation is corroborated when looking at the average scores from executed plans in Tab.1: the experiment with more randomness produces a larger average score than the one with the manually calibrated router. Overall, the morale is that, for heterogeneous multi-agent simulations, an innovative module that generates diversity significantly reduces both the consistency burden of the software system (the necessity to keep the values of time consistent between different modules) and the level of detail that the innovative module needs to have about the problem (e.g. the necessary time resolution of the routing module).

5. Acknowledgements and disclaimer

KN thanks University of Pretoria for hospitality during a sabbatical. – The results are based on an uncalibrated model to illustrate the issue of different routing procedures; the results, in particular the daily revenues, are not realistic.

References

1. ARTHUR, B. Inductive reasoning, bounded rationality, and the bar problem. *American Economic Review (Papers and Proceedings)* 84 (1994), 406–411.
2. BADER, R., DEES, J., GEISBERGER, R., AND SANDERS, P. Alternative route graphs in road networks. In *Theory and Practice of Algorithms in (Computer) Systems*, A. Marchetti-Spaccamela and M. Segal, Eds., vol. 6595 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2011, pp. 21–32.
3. BALMER, M. *Travel demand modeling for multi-agent transport simulations: Algorithms and systems*. PhD thesis, Swiss Federal Institute of Technology (ETH) Zürich, Switzerland, 2007.
4. BALMER, M., RIESER, M., MEISTER, K., CHARYPAR, D., LEFEBVRE, N., NAGEL, K., AND AXHAUSEN, K. MATSim-T: Architecture and simulation times. In *Multi-Agent Systems for Traffic and Transportation*, A. Bazzan and F. Klügl, Eds. IGI Global, 2009, pp. 57–78.
5. CHARYPAR, D., AND NAGEL, K. Generating complete all-day activity plans with genetic algorithms. *Transportation* 32, 4 (2005), 369–397.
6. DELL'OLMO, P., GENTILI, M., AND SCOZZARI, A. On finding dissimilar Pareto-optimal paths. *European Journal of Operational Research* 162 (2005), 70–82.
7. FERBER, J. *Multi-agent systems. An Introduction to distributed artificial intelligence*. Addison-Wesley, 1999.
8. HOLLAND, J. *Adaptation in Natural and Artificial Systems*. Bradford Books, 1992. Reprint edition.
9. HRABER, P., JONES, T., AND FORREST, S. The ecology of Echo. *Artificial Life* 3, 3 (1994), 165–190.
10. JARA-DÍAZ, S. *Transport Economic Theory*. Emerald Group Publishing Limited, 2007.
11. JOUBERT, J. W., AND AXHAUSEN, K. W. Inferring commercial vehicle activities in Gauteng, South Africa. *Journal of Transport Geography* 1 (2011), 115–124.
12. JOUBERT, J. W., FOURIE, P., AND AXHAUSEN, K. W. Large-scale agent-based combined traffic simulation of private cars and commercial vehicles. *Transportation Research Record* 2168 (2010), 24–32.
13. KICKHÖFER, B., GREYER, D., AND NAGEL, K. Income-contingent user preferences in policy evaluation: application and discussion based on multi-agent transport simulations. *Transportation* 38 (2011), 849–870.
14. KICKHÖFER, B., AND NAGEL, K. Towards high-resolution first-best air pollution tolls. *Networks and Spatial Economics* (2013), 1–24.
15. KICKHÖFER, B., ZILSKE, M., AND NAGEL, K. Income dependent economic evaluation and public acceptance of road user pricing. In *Proceedings of the Kuhmo Nectar Conference on Transportation Economics* (2010). Also VSP WP 10-03, see www.vsp.tu-berlin.de/publications.
16. MAHMASSANI, H., ZHOU, X., AND LU, C.-C. Toll pricing and heterogeneous users: Approximation algorithms for finding bicriterion time-dependent efficient paths in large-scale traffic networks. *Transportation Research Record: Journal of the Transportation Research Board* 1923 (2005), 28–36.
17. MAULDIN, M. L. Maintaining diversity in genetic search. In *AAAI* (1984), pp. 247–250.
18. NAGEL, K., AND FLÖTTERÖD, G. Agent-based traffic assignment: Going from trips to behavioural travelers. In *Travel Behaviour Research in an Evolving World – Selected papers from the 12th international conference on travel behaviour research*, R. Pendyala and C. Bhat, Eds. International Association for Travel Behaviour Research, 2012, ch. 12, pp. 261–294.
19. PALMER, R., ARTHUR, W., HOLLAND, J. H., LeBARON, B., AND TAYLER, P. Artificial economic life: a simple model of a stockmarket. *Physica D* 75 (1994), 264–274.
20. RUSSEL, S., AND NORVIG, P. *Artificial Intelligence – A Modern Approach*, 3 ed. Pearson Education, Upper Saddle River, New Jersey 07458, 2010.
21. TALBI, E.-G. *Metaheuristics: from design to implementation*. John Wiley & Sons, 2009.
22. WIKIPEDIA. Generating log-normally distributed random variates. http://en.wikipedia.org/wiki/Log-normal_distribution/#Generating_log-normally_distributed_random_variates, accessed 6-feb-2014.