

Pushing the limits

A concept of a parallel microsimulation framework

Working Paper**Author(s):**

Charypar, David; Horni, Andreas; [Axhausen, Kay W.](#) 

Publication date:

2010-08

Permanent link:

<https://doi.org/10.3929/ethz-b-000156442>

Rights / license:

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

Originally published in:

Arbeitsberichte Verkehrs- und Raumplanung 640

1 **Pushing the Limits: A Concept of a Parallel Microsimulation Framework**

2 Date of submission: 2010-08-01

3 David Charypar

Institute for Transport Planning and Systems, ETH Zurich, CH-8093 Zurich

phone: +41-44-633 35 62

fax: +41-44-633 10 57

charypar@ivt.baug.ethz.ch

4 Andreas Horni

Institute for Transport Planning and Systems, ETH Zurich, CH-8093 Zurich

phone: +41-44-633 31 51

fax: +41-44-633 10 57

horni@ivt.baug.ethz.ch

5 Kay W. Axhausen

Institute for Transport Planning and Systems, ETH Zurich, CH-8093 Zurich

phone: +41-44-633 39 43

fax: +41-44-633 10 57

axhausen@ivt.baug.ethz.ch

6 Words: 6447 + 3 Figures

ABSTRACT

7 This paper describes the concept of a versatile transport simulation framework to be used for
8 the development of integrated transport simulation models. It is designed to improve modular-
9 ization and thereby simplify the collaboration between scientists and engineers from different
10 fields. Modularization is believed to be important due to the increasing complexity of the im-
11 plementation task if models from different areas are to be integrated. In performance critical
12 software projects this complexity is often further increased by the desire or the need for an
13 implementation that can be run in parallel on multiple processors. Furthermore, developing an
14 efficient parallel simulation is not trivial. In addition to the complexity of the modeling task as
15 such one has to deal with communication delays and data availability issues. The idea of the
16 presented framework is to handle this complexity by defining simple rules according to which
17 user developed modules must act. These rules include certain minimum delays between the
18 observation of a change in the system and the triggered reactions, limited vision, and limited
19 traveling speed. To illustrate how the framework is to be used a simple modeling scenario
20 is created and a possible implementation employing the framework is sketched. The example
21 scenario consists of the integration of a pedestrian simulation with a load estimation module
22 and a travel time estimator. Finally, an outlook on the next steps in the development of the
23 framework is given.

INTRODUCTION

24 In recent years microscopic, dynamic demand and traffic simulations modeling has caught in-
25 creased attention and found broader use in the field of traffic forecasting and transport planning.
26 One advantage such approaches is that the necessary behavioral models are relatively simple.
27 The reason for this is that they only model directly the behavior of one individual or a small
28 group of individuals (e. g. households). The rest of the behavioral richness is assumed to emerge
29 from the interaction of thousands of such individuals that in the end form a complex system.
30 Another advantage is and that during analysis it is possible to follow the line of influence down
31 to the individuals, which makes the interpretation of the results more intuitive.

32 However, microscopic simulations have (at least) one important and well known drawback:
33 for any reasonable sized scenario they are computationally demanding, as each of the individ-
34 uals naturally must be represented separately.

35 This in turn makes them relatively inefficient for solving low resolution scenarios (for in-
36 stance based on coarse zonal data). The reason for this is that during the necessary disaggrega-
37 tion process, a lot of random detail is generated and afterwards simulated. Clearly, this detail
38 does not contribute to the explanatory power of the output and must be interpreted as overhead
39 of such simulation approaches.

40 On the other hand, if high resolution data is available microsimulations become comparably
41 efficient as, to provide the same precision, any aggregated model needs to refine the zoning to
42 the degree where it fits the level of detail of the input data. Consequently, the resulting OD-
43 matrices become huge as their size grows quadratically with the number of zones.

44 Overall, it seems to be worthwhile to develop and use microscopic transport models, and
45 they have already been successfully applied to large real world problems. Often, the average
46 work day is under investigation—a problem that can be addressed with equilibrium models.
47 Such modeling challenges can already be handled with available models (e. g. (1) even though
48 they require a lot of resources both in time and computing power. One way to reduce the
49 computation time is through parallel execution of simulation programs. While this is becoming
50 more common through multi-core processors that are available nowadays it is still a challenging
51 task to *develop* parallel software. At the same time it is necessary to follow the path of paral-
52 lelization in future microsimulation software developments to increase the range of problems
53 that can be addressed.

54 There are, however, different modeling task that cannot be solved using equilibrium models.
55 One important examples is the simulation of unpredictable events like accidents, emergencies,
56 or disasters. Another example is the simulation of longer periods than a single day. Multi-day
57 simulations increase the computation time of microscopic equilibrium models in two distinct
58 ways: First, since the period of interest is longer the simulation of this period naturally also
59 takes longer. Basically doubling the simulated time doubles the computation time. In agent-
60 based microsimulations finding the equilibrium is often approximated by a learning loop of the
61 agents (e. g. in MATSim (2) or in TRANSSIMS (3)). This loop represents an iterative algorithm
62 converging towards the desired result. The number of iterations necessary to achieve a result
63 of a certain precision naturally depends on the complexity of the solution and hence also on the
64 size of the solution vector: More complex solutions need more iterations to be found. It is clear
65 that finding the equilibrium for a 7-day period is substantially more complex than finding it for
66 a single day.

67 When combining both considerations above it can be seen that the computational burden

68 of the described iterated approach really becomes an issue. The computation time increases
69 disproportionately with the length of the study period. Consequently, there is a relatively short
70 (computational) limit of what time periods can be investigated using agent-based models with
71 an iteration-type learning loop.

72 Further more, there exists another argument against iterative learning and against modeling
73 long periods as an equilibrium: In the agent-based context, an (admittedly simplified) interpre-
74 tation of an equilibrium is that all agents have considered (all) possible choices and found the
75 sequence of actions that maximizes their utility in the given environment. It can be doubted if
76 real people really plan their weeks, months, or even years completely in advance.

77 Another problem is the assumption that two consecutive runs of the same simulation with
78 only minimal perturbations yield the same result when the list of planned actions is kept the
79 same. One has to remember that complex systems (as the transport infrastructure) tend to
80 amplify disturbances, and this can lead to completely different results at the end. Such effects
81 have been observed in at least one implementation of a microscopic integrated demand model
82 (4).

83 We believe that it would make much more sense to model longer, multi-day periods as
84 a continuously evolving scenario, where the modeled persons (agents) constantly make deci-
85 sions on the following time frame. see e. g. Märki *et al.* (5) However, this makes it necessary
86 to make current information about the state of the system available to all objects in the simu-
87 lation. The online estimation of state variables during a running simulation and consequently
88 the propagation thereof represents a substantial increase in complexity of the simulation. In
89 iterative frameworks this information exchanges happens at the end of the iteration where the
90 generated output is analyzed, processed and made available as static information to agents for
91 re-planning.

92 In the proposed framework online information processing and spreading across the simula-
93 tion is provided as a service. This makes it possible that entities in the simulation (e. g. agents)
94 can use them for their continuous planning process.

95 Based on the above line of reasoning, the development of a framework for large continuous
96 spacial microsimulation was started. The purpose of this tool is to encapsulate all necessary
97 complexity for parallelization and spacial information interchange and hide it from the user.
98 Consequently, a module employing the framework will be comparably simple as it will be able
99 to rely on the framework's functionality provided through a clear interface.

100 The remainder of this paper is structured as follows: The next section discusses related
101 work, after that the proposed framework is specified and an implementation of three basic
102 modules of an integrated agent-based microsimulation is sketched as it could be done using
103 the described framework. Finally an outlook on planned future steps is given, and a discussion
104 concludes the paper.

RELATED WORK

105 The following is a short overview about other work that was performed either in the field of
106 parallel microsimulations or in the design of frameworks for transport or urban modeling.

107 **Parallel Traffic Simulations**

108 There are numerous examples of parallel implementations of traffic simulations. Barceló *et al.*
109 (6) showed a parallel implementation of their microsimulator AIMSUN achieving a parallel
110 speedup of 3.5 when run on 8 processors. The parallelization concept was to make all data
111 globally accessible. PTV's VISSIM traffic microsimulator also has the capability to run in par-
112 allel using a multi-threaded concept. (7)

113 Nagel and Rickert (8, 3) showed a parallel version of a cellular automaton used for traf-
114 fic flow simulation in TRANSIMS (9). They used message passing between processors and
115 achieved a speedup of 10 with 32 processors. They reported latency problems due to Ethernet
116 data communications.

117 There has been some work on parallel queue-based models (e. g. 10, 11, 12, 13) Using
118 message passing between cluster nodes, the queue-based model presented in (10, 11) achieved
119 a speed-up of 32 using 64 CPUs when simulating a peak period. In (13) the authors report a
120 parallel speedup of 53 when using 64 processors for simulating a large scenario.

121 A number parallel implementations of mesoscopic transport models have been presented
122 in the past. METROPOLIS(14, 15) is able to simulate large scenarios efficiently by using
123 a parallel implementation based on up to 16 threads. DynaMIT(16, 17) does not parallelize
124 the traffic flow simulation itself but uses task parallelization i.e. different modules are run in
125 parallel. Unfortunately this limits the number of usable processors to the number of modules.
126 DYNEMO(18, 19) was run in parallel (19) by using a message passing technique on 19 CPUs
127 for simulating small scenarios. Larger numbers of CPUs were reported to be inefficient.

128 **Frameworks for Integrated Modeling**

129 Ferreira *et al.* (20) present a framework (MAS-T2er) for integrated multi-agent systems. Their
130 focus is on control strategies and intelligent transport systems. The intended use of their soft-
131 ware is "...for cooperative design, visualization and engineering, allowing for the cooperative
132 decision-making by different traffic and transport experts". Their framework is designed to run
133 on distributed systems. It is still under development.

134 The goal of UrbanSim(21) is to model and simulate urban development by modeling the
135 interactions of many different actors that make decisions in the markets for land, housing, non-
136 residential space and transportation.

137 The Multi-Agent Transport Simulation Toolkit (MATSim-T) (2) is a simulation framework
138 for modular development of an integrated transport simulation for large-scale applications. Sev-
139 eral modules that were written for/in MATSim run in parallel, e. g. certain versions of the
140 traffic flow simulator(e. g. 13), the processing of simulation events, and the activity planning
141 module. However, the program code of the framework itself is executed sequentially.

FRAMEWORK

142 In this section first, the problematics that necessitate the creation of the described programming
143 framework are discussed, second, the model concept is derived from these problematics, and
144 third, the design of the software is elaborated on.

145 **Motivation**

146 The context of this work is the microsimulation of different aspects of travel. Microsimulation
147 can be a very powerful tool to gain insight into travel behavior, emerging dynamics of systems
148 with human actors, and effects resulting from external measures. Unfortunately, microsimula-
149 tion models are computationally demanding, which makes it hard to apply them to sufficiently
150 large problems. Consequently, a tool that accelerates the development and execution of mi-
151 crosimulation models would increase their range of application.

152 Integrating different models in one more complex microsimulation also widens the range
153 applications: It enables researchers to investigate interactions between different aspects of the
154 modeled scenario. For instance integrating a traffic simulator and a routing module with a lo-
155 cation choice module makes it possible to study the effect of congestion on location choice. To
156 keep the modeling task as simple as possible it is desirable to have a high level of modulariza-
157 tion and to have the development process of modules as isolated as possible. This also helps to
158 control code complexity.

159 In many research projects emerging phenomena (e. g. urban gridlock) are of special interest.
160 Such emergence naturally only occurs in sufficiently large systems, and hence researchers must
161 be able to cope with such systems computationally. Here, being able to efficiently use parallel
162 computers might make the difference if interesting effects might be investigated or not.

163 Unfortunately, until now developing parallel programs tremendously increases the com-
164 plexity of a coding project and makes it hard to handle.

165 The classical modeling approach in transport planning is to compute a user equilibrium
166 to investigate long term effects of changes. Recently the immediate response to unexpected
167 events is catching more and more interest. To be able to model such effects, users of the system
168 (i. e. agents) must have access to estimates of the current state of the system. When develop-
169 ing a simulation software, including online state estimation corresponds to integrating another
170 module. Obviously this further increases software complexity.

171 Based on the above considerations it seems that a framework solving the described prob-
172 lems with code complexity while at the same time simplifying parallel and modular program-
173 ming would be very useful. It would facilitate further research in the field of integrated transport
174 microsimulations.

175 **Concept**

176 The main objective for the described framework is to create a tool that simplifies the modular-
177 ization of a complex transport modeling/simulation project and to reduce the code complexity
178 of the modules at the same time. This is achieved through taking over the tasks of code paral-
179 lelization and distribution of the workload on different computers/processors, and information
180 distribution and interchange between different entities of the simulation.

181 To enable the collaboration between user-developed modules and the framework, respec-
182 tively it is necessary to specify the cut line between their individual fields of responsibility. One
183 of the key aspects of the presented framework is to limit the vision and motion capabilities of
184 object in the simulation. This is done to avoid a problem that otherwise often occurs when a
185 simulation program is later enhanced to run on a parallel computer. The following paragraph
186 should illustrate that problem

187 When developing a simple simulation program (i. e. at the beginning of the coding project),

188 initially the visibility of information about spatially distributed objects is usually assumed to
189 be global, for the sake of simplicity. For example when developing a car following model the
190 speed and position of all cars on a road might be stored in an array representing the state of the
191 system, and this state is stored at one specific spot in computer memory. As the model grows,
192 larger simulation are being performed and the desire for a parallel implementation arises. Often
193 distribution across multiple threads is tried here first, unfortunately with limited speedup of the
194 simulation. The problem is that all objects in the simulation access the same data set (the array
195 described above) and also make changes there. As a result this part of the simulation becomes
196 a bottleneck, essentially slowing down the simulation to single-CPU speed.

197 The underlying problem is, that all data is visible globally and instantly in the simulation.
198 As a result, when a data point is changed by one processor this new information must be
199 propagated to all other processors before they can continue with their individual tasks. Since
200 communication speed between processors is physically limited the simulation essentially stalls
201 until the message has been propagated.

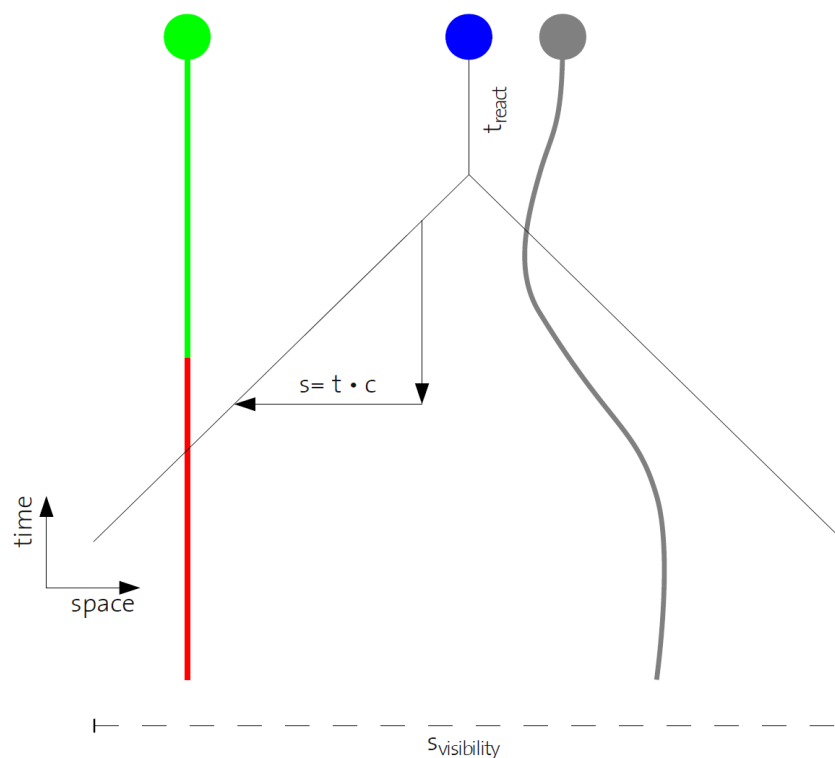
202 Our approach to this problem is to limit the assumed visibility of information and the speed
203 of its propagation and hence give the processors more time to synchronize the state of their
204 memory. This avoids stalling the CPUs and hence improves the simulation speed.

205 The first such constraint is *limited visibility of information*. To reduce the amount of data
206 that must be held readily available to an observer, we assume a maximum radius of vision. This
207 radius can be chosen by the user at the beginning of the simulation run. It will depend to a great
208 extent on the problem at hand. In case of a pedestrian simulation for evacuations of buildings
209 it might be set to e. g. 20 meters. In another example, where we want to simulate freeway car
210 traffic, some 500 meters might be more appropriate. After having specified such a visibility
211 radius it is assumed that no module will request or need information from farther away and on
212 the other hand that data provided by the framework is complete inside this radius.

213 The second introduced constraint is *delayed perception*. One can imagine this as a reaction
214 time. Virtually any system shows delayed reaction to external information. In the case of a car
215 following simulation this might be set to half a second. This is the time from the moment an
216 information can be perceived by some entity to the moment this object can take some action
217 based on it.

218 The third and maybe most important constraint is *limited speed of motion and information*
219 *propagation*. If at one point in the simulation there is some change to the state of the system,
220 this change cannot be observed instantly in the whole area around this point. Rather, the infor-
221 mation has to travel (at a certain speed) through the system much like a sound wave travels from
222 the source. Observers will not take notice of the change until this *information front* hits them.
223 Similarly, moving objects are not allowed to move faster than a certain maximum speed. While
224 this last constraint can be observed in reality (objects and information cannot travel faster than
225 light), it might seem odd to limit the speed in a simulation somewhat arbitrarily to a relatively
226 low value.

227 The final set of constraints is illustrated in Figure 1, where the perception of an object (blue)
228 is shown in a space time diagram. $s_{visibility}$ is the visibility range, t_{react} is the reaction time,
229 c is the speed of information propagation, t represents time, and s space. The thin black line
230 represents the front of information available to the blue object. The green object is stationary
231 and changed its state from red to green at a certain point in the past. Since the blue object is
232 relatively far away this change can not yet be perceived and hence the red state is still relevant
233 to the blue object. The gray object is moving and the position where its path intersects the

FIGURE 1 Limited Perception of Information Based on Introduced Constraints

234 information front is the latest position visible to blue.

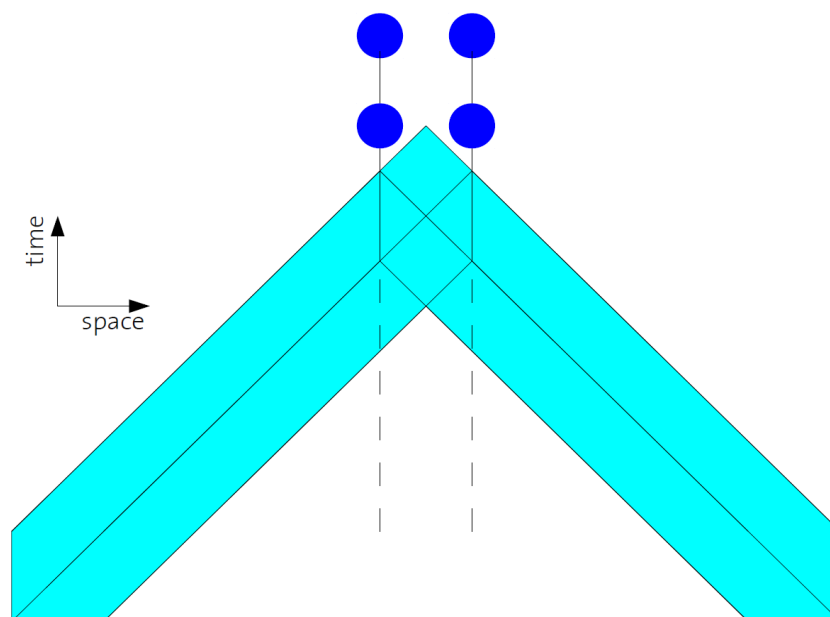
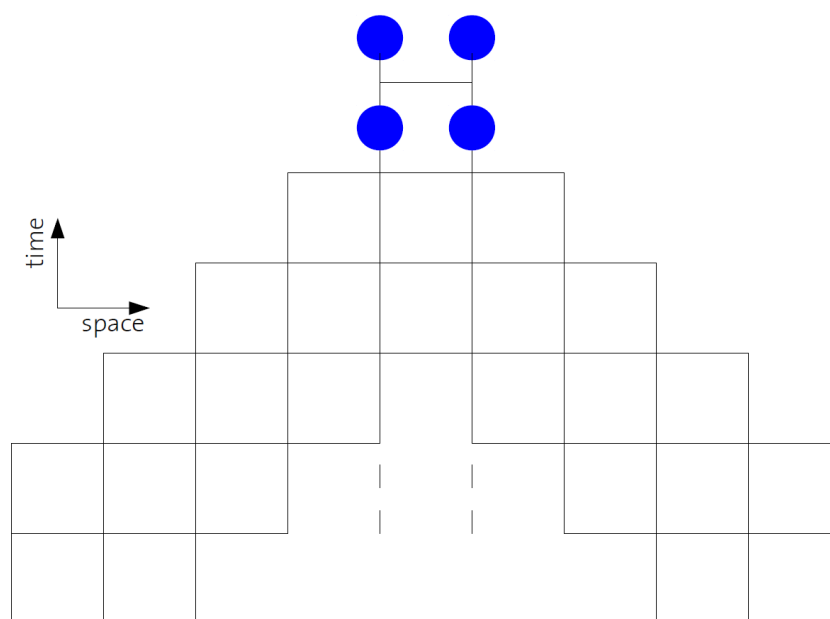
235 One basic concept of our simulation is to map the spatial ordering of the simulated area
 236 to the processors and hence to computer memory. In a sense, if information travels through
 237 the virtual domain of the simulation, it travels from processor to processor and from memory
 238 bank to memory bank involved in the computation. Since we effectively limit the speed of
 239 information propagation in the virtual world, the data also travels at limited speed between
 240 the involved computers. This in turn increases the achievable computing speed as the reduced
 241 requirements are more easily satisfied.

242 Design

243 In our framework the simulation domain is subdivided using a uniform grid with cells of side
 244 length $s = t_{react} \cdot c$, which is defined during the configuration phase of the framework. This size
 245 of the cells was selected as it simplifies the information exchange across processor boundaries.

246 As a result, the maximum number of processors that can be used is equal to the number
 247 of cells used to subdivide the domain. However, it is possible to join multiple cells to use
 248 on single CPU, eliminating the need for physical communication between these parts of the
 249 simulation. This is desirable if scenarios show differently loaded cells. Cells with comparably
 250 little work to do should be joined and assigned to a single CPU while heavily loaded cells
 251 should be simulated exclusively on a separate processor.

252 Internally, the framework holds a comprehensive list of replicated cells for each real cell,
 253 that is simulated. These lists form a discretized form of the information front described in
 254 Figure 1. Further more it is expanded to represent not only the information for one single point

FIGURE 2 Information Potentially Needed in a Cell During a Certain Time Period**FIGURE 3 Discretized Information Domain in a Cell**

255 but from the union of all data that might be needed by any entity situated in the cell during
256 one time period as long as the reaction time. A graphical representation of this can be found in
257 Figure 2. Further more, to simplify the process of exchanging information between adjacent
258 cells it is useful to discretize the information domain in a similar way as the simulation domain
259 is discretized into square cells. This is illustrated in Figure 3 which shows a discretized form
260 of Figure 2.

261 The goal of the framework described in this paper is to encapsulate as much complex-
 262 ity associated with information interchange and parallelization, and hide it from application
 263 programmers developing a module that is part of a larger integrated simulation. The basic
 264 approach is to employ a client-server architecture where both, framework and user modules
 265 take a dual role. On the one hand the framework is the server and modules are clients during
 266 phases where information about the state of the environment is requested by the modules. On
 267 the other hand the framework becomes the client after computations have been completed by
 268 user modules and the resulting new information needs to be published to the system which is
 269 when modules are in the server role.

270 It seems to be straightforward to define two interfaces here. One defining how information
 271 about the state of the system can be gathered by user-developed modules (the *gathering inter-*
 272 *face*) and a second interface for the opposite direction of data flow, the *publishing interface*. The
 273 gathering interface consists basically of a function called with the following arguments:

- **cell_index:** Index of the cell of interest for the current query.
- **start_time:** Start time of the query period.
- **end_time:** End time of the query period.
- **predicate:** A mathematical predicate. True for relevant information.

275 The function returns a list of information objects that all are relevant in the given cell during
 276 the given time period and satisfy the query predicate. Information objects represent published
 277 information about simulation entities that might be relevant to other entities. For example, in
 278 a car following model this might be information about the current position, speed, and accel-
 279 eration of a car, but it would usually not contain the route, the destination, or its desired speed
 280 of the car. The querying module can assume that no information objects are forgotten and that
 281 now new relevant information may become available during the processing of the current time
 282 period. It is clear that this is only possible if the current simulation time t_{now} , the reaction time
 283 t_{react} , the query start time t_{start} , and end time t_{end} satisfy certain condition:

$$284 \quad t_{\text{end}} - t_{\text{start}} \leq t_{\text{react}} \quad (1)$$

$$285 \quad t_{\text{end}} \leq t_{\text{now}} \quad (2)$$

286 In the other direction of information exchange the interface looks very similar. Each of the
 287 simulated entities, and hence the implemented modules must provide a function that can be
 288 called by the framework to obtain all information objects. At the end of a simulation time step
 289 the framework goes through all cells and for each entity the calls the data providing function
 290 with the following arguments:

- **start_time:** Start time of the query period.
- **end_time:** End time of the query period.

292 The object should react with a list of information objects that describe all publicly available
 293 information about the object at hand.

EXAMPLE MODULES

294 Using a small example, we would like to sketch how the described framework would be em-
 295 ployed to solve a real modeling task, involving the implementation of a couple of modules. The
 296 scene of interest is a music festival with many visitors that in general spend many hours on the
 297 festival venue. During their stay they have to get something to food and drink from time to

298 time. For this purpose there is a number of food stands distributed on the periphery of the area
299 while the main stage can be found near the center. The modeling task is now to simulate (and
300 maybe predict) the movements conducted by pedestrians looking for available food stands. It
301 would not be realistic to assume that the one food stand closest to the stage would have the
302 capacity to serve all visitors in a reasonable time. For the modeling task at hand it is of partic-
303 ular interest how visitors would use more distant and hence less crowded stands to get served
304 in shorter time.

305 The following is a relatively straight forward example of how the above modeling task
306 might be solved. Certainly it is not very sophisticated and leaves a lot of room for improve-
307 ments. The proposed realization is meant to be illustrative rather than comprehensive and
308 should give an idea how the framework would be employed.

309 Clearly, one necessary module is a pedestrian simulation that models how people walk
310 around based on their direct surrounding. One possible choice for the underlying simulation
311 model might be (22) as it already implements attraction through other objects that might be
312 used to model how visitors generally want to get as close to the central stage as possible. A
313 second module would have to be added with the task of estimating the current loading of all
314 food stands. Finally, a third module for estimating pedestrian densities would provide a way of
315 estimating travel times to different locations.

316 For each of the modules it must be now decided what are the simulated entities, what is the
317 information these entities needed about the environment for proper operation, and what is the
318 information generated. For the generated information it is especially instructive to think about
319 for what the information is to be used and what new information will be generated from it in
320 turn. In the case of the pedestrian simulation it is relatively clear that the simulated entities
321 are pedestrians and that they need information about surrounding persons to be able to act an
322 react. Symmetrically, the information pedestrians need to publish is their position, velocity,
323 plus current activity, namely if the are waiting for food or not. This part will be important
324 later in this section. Now, when a pedestrian receives the positions and velocities of all other
325 pedestrians near by (within a range of $s_{visibility}$), it has sufficient information to take its next
326 actions. (E. g. decelerating, avoiding a collision, or changing direction).

327 When it comes to the pedestrian density estimator a simple grid based approach is used
328 in this example. The simulated entities are nodes an a square lattice with a spacing of one
329 visibility range of a pedestrian. Each node uses the gathering interface to get the positions of
330 all pedestrians near by. From this the node can easily compute a local density estimate using a
331 kernel method for instance. It is clear that at this point in time the node cannot know anything
332 about pedestrian densities outside the visibility range. For this reason it is essential that this
333 local estimate is published by the node. In the next turn, each node not only gets the information
334 about pedestrians near by but also the (local) density estimates of neighboring nodes. By storing
335 them, the node can construct a density map of an area larger than the visibility range. In the
336 next turn, this whole density map is published through the framework, providing information
337 to nodes even farther away By iterating this procedure, very soon each node will possess an
338 estimate of the pedestrian densities in the whole simulated domain. This is actually a density
339 map and can be used e. g. for routing and travel time estimation.

340 The food stand loading estimator (FSLE) has as similar task as the density estimator. For
341 this reason its design is similar and it also operates in a similar way. The FSLE is also designed
342 using nodes. Each node holds a list of all food stands in the domain and how many visitors
343 are currently near them waiting for food. This list is initially empty and filled with information

344 as the simulation progresses. In a first step the loading of local (to the node) food stands is
345 estimated from the pedestrians' positions and their current activity. The value is stored in the
346 list. In the next step, each node takes its updated list and publishes it using the communication
347 framework. Then the new information is collected once again through the gathering interface
348 and the estimates from neighboring FSLE nodes are merged into the current list of food stand
349 loads. At this point the list already contains more information. In each turn, the content of
350 the food stand loading list grows, until each FSLE node has a complete list of food stand load
351 estimates.

352 At this point all parts of our simulation are ready to be used. We would like to show now the
353 steps taken if a pedestrian in the simulation becomes hungry and hence wants to find an avail-
354 able food stand: In the first step it gets the last version of the food stand loading list published
355 by the nearest FSLE node. This functionality is made available by the data gathering interface.
356 Now, for each food stand that is not overcrowded, a travel time is estimated using the last pub-
357 lished pedestrian density map. This data is also received through the data gathering interface
358 of our framework. Finally, the pedestrian selects the best of the available choices, consider-
359 ing expected waiting time at the food stand, travel times, and travel distances. The pedestrian
360 starts to walk towards the selected food stand, thereby reacting to all other pedestrians that he
361 encounters.

FUTURE WORK

362 The described framework is still in a relatively early state of development. The implementa-
363 tions must be tagged as prototypes and the interfaces are not yet finalized. One of the aspects
364 we are still working on is how information objects are transported from one simulation cell to
365 the others. There are different paradigms that one can follow here. In general one has to trade
366 off between communication and processing overhead. One extreme is the very sophisticated
367 selection of only the bare minimum of information that needs to be transferred between proces-
368 sors to guarantee correct results. This obviously comes at the expense of spending a lot of time
369 evaluating the necessity of a data point. At the other end of the scale stands the preference for
370 quick checks selecting a relatively large volume of data for transfer to other cells. Hence, using
371 this design, processor loads will be relatively low while communication demands will be high.

372 The next steps will be to finalize the communication paradigm based on performance and
373 complexity considerations, finalize the interfaces to user modules, and then create a first pub-
374 lished version of our framework. Since the presented software is meant to ease collaboration
375 on integrated transport modeling projects we seek cooperation with other interested researchers
376 that would like to implement modules using the framework. Consequently, the software is
377 meant to be released to public domain.

378 We are currently working on one first project employing our parallelization framework. In
379 that project we are aiming to simulate periods of more than 30 days in an integrated agent-based
380 environment. A special focus is the activity planning process, especially the resulting weekly
381 rhythms and effects of business-holidays on infrastructure usage. The activity planning mod-
382 ule is currently under development and shows first promising results (5). Apart from activity
383 planning there are other modules necessary for the functioning of this integrated simulation:
384 For the adaptive creation of routes there is a on-line travel time estimator under development.
385 Furthermore, we envisage a location choice module based on current load factors to represent
386 the flexible choice of shopping and leisure locations.

387 Apart from getting interesting insights into modeling and simulation of multi-day periods,

388 we plan to measure the effectiveness of our framework by testing different scenario size with
389 various numbers of processors for parallelization. If our approach proves to be right, we should
390 be able to demonstrate good scaling of performance with the number of processors.

DISCUSSION AND CONCLUSION

391 The concept of a parallel framework to be used in the development of integrated transport mi-
392 crosimulations was presented. It should increase and ease the cooperation between researchers
393 and engineers from different fields by allowing for strict modularization of different model
394 parts.

395 The framework takes care of the complex tasks of parallelization and information exchange
396 between modules. In the experience of the authors these are often critical parts of integrated
397 simulations which often lead to problems in performance, reproducibility of results, and stabil-
398 ity of the over all software package. By assuming a two dimensional domain for all modules
399 (this can be easily extended to three dimensions) and by setting explicit limits on what actions
400 can be performed (maximum speed of motion), how quickly they can be perceived (reaction
401 time), and how far away they are visible (visibility radius) it becomes possible to confine the
402 tasks of parallelization and information exchange in the described framework.

403 It must be noted that the introduced limits might also produce problems in certain cases.
404 Many models implicitly assume global availability of information and it is not clear in advance
405 if and how they can be fit into the described framework. Also some modules might not have
406 an obvious spatial interpretation. However, the authors believe that mapping everything to 2D
407 space and assuming a certain delay in the propagation of information does not represent a real
408 problem in all but very few cases.

409 In the example shown it was demonstrated how the framework would be used for a simu-
410 lation of intelligent pedestrians performing food stand location choice based on load estimates
411 and route calculation through crowded areas. Employing the framework in this example was
412 shown to be straightforward.

413 The current prototype implementations of the framework show reasonable parallel perfor-
414 mance. However, future test will have to show the performance when simulating real integrated
415 models. There certainly will be an overhead involved with our framework, as there always is
416 when introducing a layer of abstraction. However, we believe that the benefit through im-
417 proved efficiency in the development of integrated simulations will by far exceed the moderate
418 simulation overheads.

REFERENCES

- 419 1. Balmer, M., A. Horni, K. Meister, F. Ciari, D. Charypar and K. W. Axhausen (2009)
420 Wirkungen der Westumfahrung Zürich: Eine Analyse mit einer Agenten-basierten
421 Mikrosimulation, *Final Report*, Baudirektion Kanton Zurich, IVT, ETH Zurich, Zurich,
422 February 2009.
- 423 2. Balmer, M., M. Rieser, K. Meister, D. Charypar, N. Lefebvre and K. Nagel (2009)
424 MATSim-T: Architecture and simulation times, in A. L. C. Bazzan and F. Klügl (eds.)
425 *Multi-Agent Systems for Traffic and Transportation Engineering*, 57–78, Information Sci-
426 ence Reference, Hershey.
- 427 3. Rickert, M. and K. Nagel (2001) Dynamic traffic assignment on parallel computers in

- 428 TRANSIMS, *Future Generation Computer Systems*, **17** (5) 637–648.
- 429 4. Rieser, M. and K. Nagel (2008) Network breakdown “at the edge of chaos” in multi-agent
430 traffic simulations, *The European Physical Journal B - Condensed Matter and Complex*
431 *Systems*, **63** (3) 321–327.
- 432 5. Märki, F., D. Charypar and K. W. Axhausen (forthcoming) Continuous activity planning
433 for a continuous traffic simulation, paper presented at the *90th Annual Meeting of the*
434 *Transportation Research Board*, Washington, D.C.
- 435 6. Barceló, J., J. L. Ferrer, D. Garcia, M. Florian and E. Le Saux (1998) Microscopic traffic
436 simulation, in P. Marcotte and S. Nguyen (eds.) *Equilibrium and Advanced Transportation*
437 *Modelling*, chap. 1, 1–26, Kluwer, Dordrecht.
- 438 7. PTV America (2010) PTV America, webpage, <http://www.ptvamerica.com>.
- 439 8. Nagel, K. and M. Rickert (2001) Parallel implementation of the TRANSIMS micro-
440 simulation, *Parallel Computing*, **58** (2) 1611–1639.
- 441 9. Nagel, K., D. E. Wolf, P. Wagner and P. M. Simon (1998) Two-lane traffic rules for cellular
442 automata: A systematic approach, *Physical Review E*, **58** (2) 1611–1639.
- 443 10. Cetin, N. (2005) Large-scale parallel graph-based simulations, Ph.D. Thesis, ETH Zurich,
444 Zurich.
- 445 11. Cetin, N., A. Burri and K. Nagel (2003) A large-scale multi-agent traffic microsimulation
446 based on queue model, paper presented at the *3rd Swiss Transport Research Conference*,
447 Ascona, March 2003.
- 448 12. Charypar, D., K. W. Axhausen and K. Nagel (2007) An event-driven queue-based traffic
449 flow microsimulation, *Transportation Research Record*, **2003**, 35–40.
- 450 13. Charypar, D., M. Balmer and K. W. Axhausen (2009) High-performance traffic flow mi-
451 crosimulation for large problems, paper presented at the *88th Annual Meeting of the Trans-*
452 *portation Research Board*, Washington, D.C., January 2009.
- 453 14. Marchal, F. (2001) Contribution to dynamic transportation models, Ph.D. Thesis, Univer-
454 sity of Cergy-Pontoise, Cergy-Pontoise.
- 455 15. de Palma, A. and F. Marchal (2002) Real cases applications of the fully dynamic
456 METROPOLIS tool-box: An advocacy for large-scale mesoscopic transportation systems,
457 *Networks and Spatial Economics*, **2** (4) 347–369.
- 458 16. Ben-Akiva, M. E., M. Bierlaire, H. Koutsopoulos and R. Mishalani (1998) DynaMIT: A
459 simulation-based system for traffic prediction, paper presented at the *DACCORS Short*
460 *Term Forecasting Workshop*.
- 461 17. DynaMIT (2006) Intelligent transportation system program, webpage, [http://mit.](http://mit.edu/its/dynamit.html)
462 [edu/its/dynamit.html](http://mit.edu/its/dynamit.html).
- 463 18. Schwerdtfeger, T. (1984) DYNEMO: A model for the simulation of traffic flow in motor-
464 way networks, in J. Volmuller and R. Hamerslag (eds.) *Proceedings of the Ninth Inter-*
465 *national Symposium on Transportation and Traffic Theory*, chap. 4, 65–87, VNU Science
466 Press, Utrecht.

- 467 19. Nökel, K. and M. Schmidt (2002) Parallel DYNEMO: Meso-scopic traffic flow simulation
468 on large networks, *Networks and Spatial Economics*, **2** (4) 387–403.
- 469 20. Ferreira, P. A. F., E. F. Esteves, R. J. F. Rossetti and E. C. Oliveira (2008) A cooperative
470 simulation framework for traffic and transportation engineering, in L. Yuhua (ed.) *Coop-*
471 *erative Design, Visualization, and Engineering*, vol. 5220 of *Lecture Notes in Computer*
472 *Science*, 89–97, Springer, Berlin.
- 473 21. Waddell, P. (2002) Urbansim: Modeling urban development for land use, transportation,
474 and environmental planning, *Journal of the American Planning Association*, **68** (3) 297–
475 314.
- 476 22. Helbing, D. and P. Molnár (1995) Social force model for pedestrian dynamics, *Physical*
477 *Review E*, **51** (5) 4282–4286.