


Design and implementation of a parallel queue-based traffic flow simulation

Working Paper**Author(s):**

Dobler, Christoph; Axhausen, Kay W. 

Publication date:

2011

Permanent link:

<https://doi.org/10.3929/ethz-b-000040273>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

Arbeitsberichte Verkehrs- und Raumplanung 732

1 **Design and Implementation of a Parallel Queue-Based Traffic Flow Simulation**

2 Date of submission: 2010-jul-28

3 Christoph Dobler
IVT, ETH Zurich, CH-8093 Zurich
phone: +41-44-633 65 29
fax: +41-44-633 10 57
dobler@ivt.baug.ethz.ch

4 Kay W. Axhausen
IVT, ETH Zurich, CH-8093 Zurich
phone: +41-44-633 39 43
fax: +41-44-633 10 57
axhausen@ivt.baug.ethz.ch

5 Words: 6032, Figures: 5

ABSTRACT

6 Today, agent based micro-simulations are widely used in the field of transport planning and
7 traffic management. One important requirement is the ability to simulate large scale scenarios
8 in reasonable time. An obvious approach to reduce the computation time of such scenarios is
9 to use multiple CPU cores.

10 This paper presents the implementation of a parallel queue simulation for MATSim written
11 in Java. Existing parallel traffic micro-simulations are reviewed concerning their paralleliza-
12 tion approaches as well as the reached performance gains. Various concepts how to model
13 the progress of time and how to distribute computational workload among multiple CPU cores
14 are discussed. Based on an analysis of the MATSim framework regarding its structure, per-
15 formance and extensibility the concepts for the parallel queue simulation are selected and im-
16 plemented. Performance tests with different sized scenarios are conducted. An analysis of the
17 results shows that especially for large scale scenarios a significant performance gain is reach-
18 able.

INTRODUCTION AND RELATED WORK

19 Today, agent based micro-simulations are widely used in the field of transport planning and
20 traffic management. One important requirement is the ability to simulate large scale scenarios
21 in reasonable time. Until the end of the last millennium, the main focus in CPU development
22 was to increase the computing power of a single CPU core. As a result, a simulation could be
23 simply sped up by using a faster CPU.

24 Within the last years, the development focus has changed dramatically. Today, it can be
25 assumed that in the near future computers with multi core CPUs will become state of the art.
26 Increasing the computing power of a CPU is mainly based on the usage of multiple cores where
27 each core for itself will not have a significantly better performance than an old single core CPU.

28 As a result, existing program code has to be adapted to be able to benefit from this new
29 multi-core architecture. Typically, this makes considerable changes in the program structure
30 necessary because the program logic has to be switched from sequential to parallel. This paper
31 presents the implementation of a parallel queue simulation which results in a major speedup
32 and therefore reduces the simulation time of large scale scenarios significantly.

33 Here, existing work in the field parallel transport simulations is determined. This includes
34 an overview of existing parallel simulations tools as well as the techniques which they use for
35 the parallelization. Additionally commonly used approaches to model the progress of time
36 within a simulation are described and analyzed regarding their suitability for a parallel im-
37 plementation. Moreover it is discussed how the computation effort of a simulation can be
38 distributed among multiple CPU cores. Subsequently, MATSim, a framework for iterative,
39 agent-based micro-simulations, is described with a special focus on its simulation modules.
40 Based on the findings from the previous sections, in the implementation section the selection
41 and implementation of a parallelization approach is described. Afterwards the performance of
42 the implementation is measured and evaluated using various real world scenarios. The paper
43 closes with some conclusions and the outlook on further work.

44 **Parallel Traffic Flow Micro-Simulations**

45 In this section we present a selection of previous work related to parallel traffic flow micro-
46 simulations. An overview on traffic flow simulations in general is for example given by (1).
47 (2, 3, 4, 5) give a detailed overview on the topics of (multi-)agent-systems and simulations.

48 Various existing micro-simulations have been ported to parallel computers. As described
49 by (6) AIMSUN2 uses a shared memory approach based on a parallel threads. Each of these
50 threads is a sequence of instructions executed within the context of a process. If a process hosts
51 multiple threads they can access the same data at the same time which may lead to inconsisten-
52 cies or deadlocks. Therefore, it has to be ensured that changes on the data are only allowed by
53 one thread at a time. The distribution of the calculation effort is done by introducing a system
54 with so called blocks and layers. A block contains objects which interact with each other in a
55 simulation step. A layer groups blocks which do not influence each other and therefore can be
56 simulated simultaneously. By using 8 parallel threads, they reach performance gains up to a
57 factor of 3.5.

58 A mesoscopic traffic simulation model is implemented by DYNEMO (7). A parallel imple-
59 mentation for distributed multiprocessor systems with distributed memory has been developed.
60 The parallelization is based on the usage of subnetworks which are created by splitting the net-
61 work along intersections. As a result the split intersections are duplicated and exist in multiple

62 subnetworks. To simulate traffic between the subnetworks so called transit-storage links are
63 introduced. They accumulate cars which want to proceed to links belonging to another subnet-
64 work. After each simulation step the subnetworks exchange the cars on those transit-storage
65 links. A speed-up of factor 15 using 19 processors is reported.

66 The parallel implementation of TRANSIMS uses a cellular automata approach (8). The
67 calculation effort is distributed among parallel distributed processors by splitting the network
68 into domains. The cuts are performed in the middle of links. Each of the divided links is
69 fully represented in both domains. The consistency between different processors is maintained
70 by exchanging information about the divided links. TRANSIMS uses an iterative simulation
71 approach to do adaptive load balancing. During each iteration the calculation times for all
72 intersections and links are measured. By using this data the load balance is optimized from
73 iteration to iteration.

74 An event-driven parallel queue-based micro-simulation for MATSim is introduced by
75 Charypar *et al.* (17). In contrast to the other described parallel simulations it can be run on
76 shared memory computers utilizing multiple CPUs. The workload is distributed by an adaptive
77 domain decomposition approach. A small test scenario is sped up by a factor of 53 when using
78 64 CPU cores.

79 There is a multiplicity of other parallel agent-based traffic micro-simulations that are not
80 discussed here because they employ similar concepts (e.g. 9, 10).

81 As can be seen, a major part of those parallel micro-simulations use the concept of dis-
82 tributed computation. Distributed systems consist of multiple computers which are loosely
83 coupled—e.g. through a computer network—where interactions between the computers are
84 relatively slow. Parallel computing in contrast means parallel execution of calculations on
85 multi-processor (and / or multi-core) computing platforms. Interactions between different pro-
86 cessors are significantly faster than on distributed computers (11).

87 When most of those micro-simulations were written distributed computation was a com-
88 monly used technology. Multi-processor systems were expensive and possible scenarios sizes
89 limited by the available amount of memory. However, today the situations has changed dramati-
90 cally. Even typical workstations use multi-core CPUs and several GB of memory. Therefore
91 parallel computing has gained an enormous amount of attractiveness. Especially large scale
92 scenarios—as they are frequently used today—may profit from such a paradigm shift. E.g.
93 applying a domain decomposition approach to a high resolution network will create a huge
94 amount of shared links and / or nodes which again will result in many interactions between
95 different processors. While those interactions can be handled in reasonable time by a parallel
96 computing implementation, they may significantly slow down an approach based on distributed
97 computing.

98 **Modelling the Progress of Time**

99 A common criterion to classify simulations is grouping them by the way they model the
100 progress of time. The two mainly used approaches in the field of traffic flow simulations are
101 time step based and event based.

102 A simple method to model the progress of time is to divide the simulated period into equal
103 sized time slices (*time slice*, *time bin* and *time step* are used synonymous in this context).
104 For each of these time slices the state of the simulated system has to be evaluated—which
105 is one major drawback of this approach. Even if nothing happens between two time steps—

106 and, therefore, the system state does not change—the state of the system has to be calculated.
107 Another problem is determining the size of the time steps. On the one hand, using too short
108 time slices results in unnecessary long calculation times. Too long time bins, on the other hand,
109 may lead to poor or even wrong simulation results. In many simulated systems, the number of
110 events occurring during a time step varies significantly. Thus, it is necessary to choose the time
111 step size according to the peak times.

112 Think of a road where on average every 60 seconds a car is driving along. During the peak
113 hour, significantly more vehicles may pass that road, e.g. one every 10 seconds. Having a time
114 step of 10 seconds seems to be appropriate when looking at the average flow rate but clearly is
115 too large with respect to the rate during the peak hour. One obvious solution for this problem
116 is to adapt the size of the time slices during the simulation, which can be done dynamically
117 depending on the results of previous time steps or based on predefined rules resulting from
118 existing knowledge (e.g. the peak hours of the call center are known). However, a problem
119 that cannot be solved by adapting the size of the time bins is load balancing. Again, this can
120 be illustrated with the simulation of roads. If not a single road but an entire road network is
121 simulated, it is obvious that the traffic flow rate differs depending on the location of a road. As
122 a result, the time step size has to be small during the whole simulation, which again results in
123 a high computational effort.

124 Typically, simulation software based on a time step approach can be parallelized quite sim-
125 ply. The main requirement is that the events which occur within a time step can be separated
126 into groups that are independent from each other. In the road network example, this could be
127 e.g. a group for every road in the network. In a parallel implementation, each of those groups
128 could be handled by a separate simulation thread which synchronizes its data with the other
129 threads at the end of each time step.

130 Another even more intuitive possibility to simulate time is an event driven approach. In a
131 discrete-event simulation, the operations within a simulated system are represented as a chrono-
132 logically ordered list of events. Each event occurs at a given point in time and causes a change
133 of the system state (12). Using again a road as example, every car entering or leaving a road
134 would create such an event.

135 Classic event driven simulation modules use internally a list of events which have to be
136 processed at their scheduled future point in time (13, 14). During the simulation the events are
137 processed in chronological order—when the list is empty the simulation ends. The scheduled
138 events can be predefined before the simulation starts and / or be created during a running simu-
139 lation. In a distributed event driven simulation employing multiple threads leads to a situation
140 where each simulation thread uses its own simulation clock. This clock is not linked to the
141 clocks of the other threads. Combined with the different calculation efforts of the threads, this
142 results in varying current simulation times. Thus, situations will occur where the events are not
143 processed in a chronological order anymore.

144 Several solutions to solve this problem have been proposed which can be divided into opti-
145 mistic and conservative approaches (e.g. 13, 15, 14, 11, 16). Optimistic approaches assume that
146 such timing problems will not occur very often. Therefore, the threads can process the events
147 without checking whether other events should be handled before. However, if a timing problem
148 appears, a roll back procedure has to be performed which turns the multiple simulation clocks
149 back to the point in time where the events can be processed in the correct order. Another pos-
150 sible solution is a conservative approach where each thread has to check whether its simulation
151 time can be advanced or not. Doing so ensures on one hand that no roll back procedures have

152 to be performed but on the other hand causes additional calculation effort for the consistency
153 checks.

154 Depending on the simulated problem the one or the other approach performs better. How-
155 ever, experiments with real world scenarios in the field of traffic flow simulations show that
156 event driven approaches tend to perform better than time step based ones. In such scenarios
157 the traffic volumes and their distribution in the network varies significantly in space and time
158 which results in a large computational overhead for time step based simulations (17).

159 **Distribution of the Workload**

160 The computational power of parallel computers can be utilized by one of two fundamentally
161 different approaches. On the one hand, new simulation software can be developed which in-
162 corporates algorithms that are designed to be run on parallel computers. On the other hand,
163 existing software can be adapted to be able to do parallel simulations (e.g. 18). Regardless of
164 the approach pursued a parallel simulation has to split up the total computational effort into
165 small packages which can be handled by the parallel executed modules of the simulation.

166 In a *functional decomposition*, the tasks that have to be performed are assigned to different
167 simulation modules. In a traffic flow simulation, one module could do the routing while an-
168 other one could execute the movement of the vehicles. Such a decomposition is often easy to
169 implement but the achievable speed-up is limited by the number of tasks that can be performed
170 simultaneously (8).

171 Another approach commonly used in the field of parallel traffic flow simulation (e.g. 9, 10,
172 8, 7) is *domain decomposition*. The aim is to divide the simulation problem into pieces with
173 approximately equal computational effort. Each of those pieces is handled by one CPU core.
174 Such an approach performs best if the domains do not interact with each other. Then, almost
175 linear performance gains can be realized. However, in typical traffic simulation interactions
176 between domains occur quite frequently. Thus, their influence on the overall performance
177 cannot be ignored. Each time such interactions take place a certain amount of calculation
178 overhead—overhead in this context are calculations that would not be necessary in a non-
179 parallel simulation—is created. Depending on that overhead the reachable performance gain is
180 limited.

181 Another factor that can significantly influence performance is the load balance between the
182 domains. The domain with the highest calculation effort affects the total duration of a simula-
183 tion run. Depending on the kind of simulated problem various solutions can be used to reach
184 an approximately even balance. Using static domains is often sufficient for simple and well
185 known problems where the calculation effort can be predicted with high accuracy. If the prob-
186 lem gets more complex adapting the size of the domains dynamically is an obvious possibility
187 to keep the calculation effort balanced. Various different dynamic load balancing strategies are
188 for example discussed by (19). However, adapting the domain sizes again produces additional
189 overhead. Hence, accepting a certain amount of imbalance between the calculation efforts may
190 be preferable.

191 When applying *domain decomposition* to a traffic flow simulation it is feasible to create the
192 domains based on the simulated network structure. The infrastructure objects like links, nodes
193 and traffic lights are assigned to the domains. The agents are dynamically assigned to the thread
194 that handles the infrastructure object on which they are physically present.

195 Selecting the objects that belong to a domain again can be done in different ways. Using

196 a random assignment typically results in a good load balance between the domains and there-
 197 fore no further mechanisms to check and adapt the balance are needed. Another advantage
 198 of a random approach is that it is simple to implement and no problem specific knowledge is
 199 necessary. A clear drawback is that the amount of interactions between different domains is
 200 extremely high.

201 Another approach is to create the domains based on the network structure. Areas with high
 202 connectivity are consolidated into domains and domain borders are placed in areas with only
 203 low connectivity. A significant advantage of such an approach is that the level of communi-
 204 cation between different domains is minimized because most simulated actions only involve
 205 objects which belong to the same domain. However, creating such domains with comparable
 206 computational workloads is very complicated for typical simulation problems. In a traffic flow
 207 simulation, the computational effort typically depends more on the traffic volume than on the
 208 number of network links. Accordingly, the domain sizes should be chosen based on traffic
 209 flow information. Yet, the load balance may fluctuate significantly during a simulation—e.g. a
 210 domain that contains only housing zones has high traffic volumes in the morning and evening
 211 but only low ones in between.


MATSIM

212 Overview

213 MATSim (Multi Agent Transport Simulation) is a framework for iterative, agent-based micro-
 214 simulations of transport systems that is currently developed by teams at ETH Zurich and TU
 215 Berlin. It consists of several modules that can be used independently or as part of the frame-
 216 work. It is also possible to extend the modules or replace them with new implementations.
 217 Balmer (20) and Balmer *et al.* (21) give a detailed description of the framework, its capabilities
 218 and its structure. Because of its agent-based approach, every person in the system is modeled
 219 as an individual agent in the simulated scenario. Each agent has personalized parameters such
 220 as age, sex, available transport modes and scheduled activities. Due to the modular structure
 221 of the simulation framework, the agent's parameterset can be easily extended my new parame-
 222 ters, for example for the routing strategy that should be used or areas of the road network that
 223 the agent knows. The application of MATSim to a large scale scenario of Switzerland (over
 224 6 million agents simulated on a high resolution network with 1 million links) is presented by
 225 Meister *et al.* (22).

226 Figure 1 shows the structure of a typical, iterative MATSim simulation run. After the cre-
 227 ation of the initial demand, the plans of the agents are modified and optimized in an iterative
 228 process until a relaxed state of the system has been found. The analysis of the results can be
 229 performed afterwards.

FIGURE 1 Iterative MATSim Loop



230 The loop contains the elements *execution (simulation)*, *scoring* and *replanning*. Within
231 the simulation module, the plans of the agents are executed. Afterwards, the scoring module
232 uses a utility function to calculate the quality of the executed plans. The utility function for
233 MATSim is described by Charypar and Nagel (23). Based on the results by scoring module,
234 the replanning module creates new plans by varying start times and durations of activities as
235 well as the routes to travel from one activity to another. Replanning modules currently under
236 development will additionally allow to change order of the planned activities (24) as well as
237 the locations where they are performed (25).

238 Simulation of the traffic behavior is also part of the iterative loop. Currently, four different
239 simulation modules are available. Their task is to execute the plans of the agents within the
240 simulated scenario. The following section describes these four simulation modules.

241 **Simulation Modules**

242 *QueueSimulation*

243 The *QueueSimulation* is a deterministic, Java based re-implementation of Cetin's *SQSim* (26,
244 20). The simulation is based on a queue model and uses a time step based approach with a
245 step size of one second. Within each time step, the state of the queues is considered. As a
246 result the duration of a simulation run increases proportionally to the number of links in the
247 network and is independent of the number of simulated agents. A major disadvantage of the
248 *QueueSimulation* is its single core architecture. While other tasks in an iteration of MATSim
249 can be executed in parallel threads (for example the replanning), the *QueueSimulation* still only
250 uses one CPU core. The *QueueSimulation* offers some benefits like well documented code and
251 its simulation listener concept which allows additional modules to interact with the simulation
252 while it is running.

253 *QSim*

254 Basically the *QSim* can be described as an extended version of the *QueueSimulation*. It con-
255 tains several additional recently developed features like traffic signals (27) or simulated public
256 transport (28). While the *QueueSimulation* can be seen as a default implementation of a traffic
257 simulation module with a stable state, the *QSim* is still under development. Some new features
258 like a redesigned *Within Day Replanning Framework* (based on 29) will be fully implemented
259 in the near future.

260 *DEQSim*

261 Another implementation is the *DEQSim*, which implements an extended queue model and is
262 described in detail by Charypar *et al.* (1) and Charypar *et al.* (17). In addition to the FIFO
263 (first in, first out) behavior of the queues, a gap is simulated that moves backwards through
264 the queues which allows to simulate congestion more realistically. Two major attributes of this
265 implementation are its multi-threaded architecture and its event based approach. As a result
266 the calculation effort scales with the number of agents. Compared with the time step based
267 approach of the *QueueSimulation* the event based implementation of the *DEQSim* achieves sig-
268 nificantly shorter calculation time. A disadvantage of the *DEQSim* is that it is implemented in
269 C++ whereas MATSim is written in Java. Therefore the communication between them is done
270 using a time consuming file input/output interface which produces noticeable longer computa-
271 tion times.

272 *JDEQSim*

273 The *JDEQSim* is the fourth simulation module currently available in MATSim. It is a re-
274 designed re-implementation of the *DEQSim* in Java that is described in detail by Waraich *et al.*
275 (30). Due to conceptual differences between C++ and Java it was not possible to reach per-
276 formance gains by implementing the multi-threaded architecture of the *DEQSim*. Therefore,
277 the *JDEQSim* uses only a single CPU core. However, due to its event based approach the
278 calculation effort is significant lower compared to the *QueueSimulation* and the *QSim*.

IMPLEMENTATION

279 **General Conditions of the Parallelization Approach**

280 The first decision that has to be made is whether a new simulation module should be written
281 from scratch or if an existing one should be adapted. As presented, already multiple different
282 simulation modules for MATSim are available. They offer a wide range of functionality and
283 have already been used for various projects. Additionally, several simulations are documented
284 which can be used for performance comparisons. Therefore, reusing one of those simulations
285 is preferred.

286 As second step, it has to be decided whether the parallel simulation should base on a dis-
287 tributed or a parallel computing approach. Using the second one is preferred for two reasons.
288 On the one hand, the implementation in the existing MATSim framework should possible with
289 less effort and higher performance due to the fast data exchange between multiple threads. On
290 the other hand, the performance of desktop computers has increased significantly within the
291 last years—concerning computational power as well as available memory.

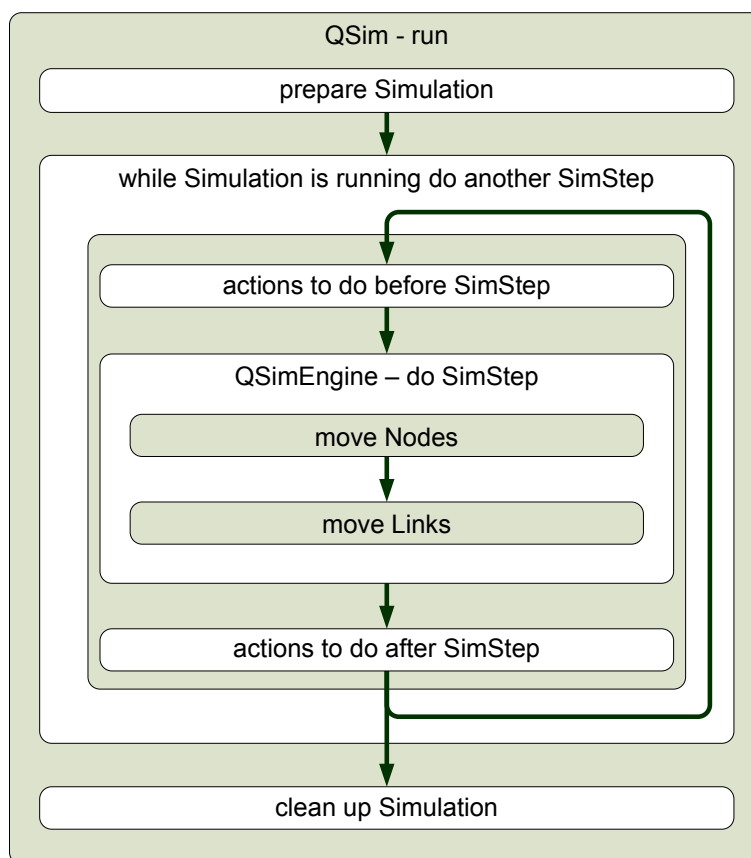
292 The next decision to be made concerns the workload distribution. A first implementation
293 of a functional decomposition for the MATSim simulation modules has already been presented
294 (30). By handling the events occurring in a separated thread, a remarkable reduction of compu-
295 tation times is reached. However, the remaining computational effort of the simulation cannot
296 be divided into further functional blocks. Thus, implementing a domain decomposition ap-
297 proach is necessary to reach further performance improvements.

298 Finally, it has to be decided whether a simulation with a time step based or a event driven
299 approach should be used for the implementation. While event driven approaches tend to per-
300 form better in the field of traffic flow simulations time step based approaches seem to be easier
301 to parallelize. There, the time steps can be used as fixed synchronization points which should
302 reduce the communication overhead between multiple threads dramatically. The consideration
303 of above factors and additional analysis of the source codes leads to the conclusion that the
304 time step based *QSim* is best basis for the implementation of a parallel micro-simulation.

305 **Analysis of QSim**

306 A simplified picture of the structure of *QSim* is shown in Figure 2. At first the simulation mod-
307 ule has to be prepared, for example to create the simulated agents. Afterwards the simulation
308 itself is started. In a loop the state of the simulated scenario is calculated for each time step.
309 When no further time steps have to be simulated, some data structures, which were only used
310 by *QSim*, are removed from memory.

311 A performance analysis shows that the *doSimStep* method in the *QSimEngine* consumes
312 over 90% of the computation time of a simulation run. In this context, only the computation

FIGURE 2 Simplified structure of QSim

313 time of the simulation itself is considered, efforts for the scoring and replanning modules are
 314 ignored. Thus, the main focus is on the parallelization of that method. Within *doSimStep* two
 315 methods with comparable computational effort are called—*moveNodes* and *moveLinks*.

316 The *moveNodes* method handles vehicles that leave one link and enter another one. Typi-
 317 cally a *Random* object is used to select in which order the ingoing links are handled (If all nodes
 318 are controlled by light signals, the *Random* object is never used). Therefore the result of a sim-
 319 ulation is influenced by the order in which the nodes are processed. This can be avoided by
 320 assigning a *Random* object to each node. Doing so will create deterministic simulation results
 321 that will slightly differ from results calculated with *QSim* because other sets of random num-
 322 bers will be used. When using multiple *Random* objects on parallel threads it is necessary to
 323 guarantee that the random numbers are independent from each other. This for example would
 324 not be the case if each *Random* object is initialized with the same initial value.

325 *moveLinks* simulates the actions (e.g. agents which start and end activities) on the links as
 326 each link can be treated independently from the other ones. Therefore the links can be simu-
 327 lated on multiple threads without concerning about race conditions with one exception (Race
 328 conditions occur in situations where the result of an operation depends on the timing of events
 329 that are created concurrently on parallel running threads, which leads to an indeterministic be-
 330 havior of the system). *QSim* can teleport vehicles from one link to another one. If within one
 331 time step multiple vehicles are teleported from different threads to one link, their order may
 332 vary. In that case they have to be ordered by their agent Id to ensure that the simulation result

333 is deterministic.

334 At some points within *moveNodes* and *moveLinks* calls to methods in global objects (*QSim*,
335 *QSimEngine* and *Simulation*) are executed. If multiple threads performs such method calls
336 concurrently this may result in an unpredictable behavior of the simulation. This can be avoided
337 by using one of two strategies. A simple but slow approach is to allow only one thread at a time
338 to call such a method. Especially if many concurrent calls from multiple threads occur this
339 will be a performance bottleneck. The second strategy is more complex and requires more
340 changes in the code but results in better performance. The method is moved from the global
341 object to one which exists once per parallel thread. Additionally it may be necessary to create
342 an additional, supervising method that is executed from the main thread.

343 This can be illustrated with a simple example. Links that do not contain active vehicles are
344 deactivated by *QSim* to reduce the calculation effort. When a vehicles enters the link, the link
345 has to be reactivated which is done by calling a method in the *QSimEngine*. There the link
346 is added to a list which is processed at a later point in time. In a parallel *QSim* each thread
347 could contain such a list. Finally the additional supervising method can instruct all threads
348 concurrently to reactivate the links which they have marked before.

349 **Structure of ParallelQSim**

350 If the current design of *QSim* that contains a single *QSimEngine* would be used in a parallel im-
351 plementation, a lot of method calls in the *QSimEngine* would have to be synchronized to avoid
352 problems with indeterministic behavior. This would result in a poor performance. This prob-
353 lem can be avoided by using one *QSimEngine* per thread. The *ParallelQSim* introduces a code
354 structure where a single *MultiThreadQSimEngine* manages an array of *QSimEngineThreads*
355 that extend the Java *Thread* class and can act as *QSimEngines*. This results in the structure
356 shown in Figure 3(a).

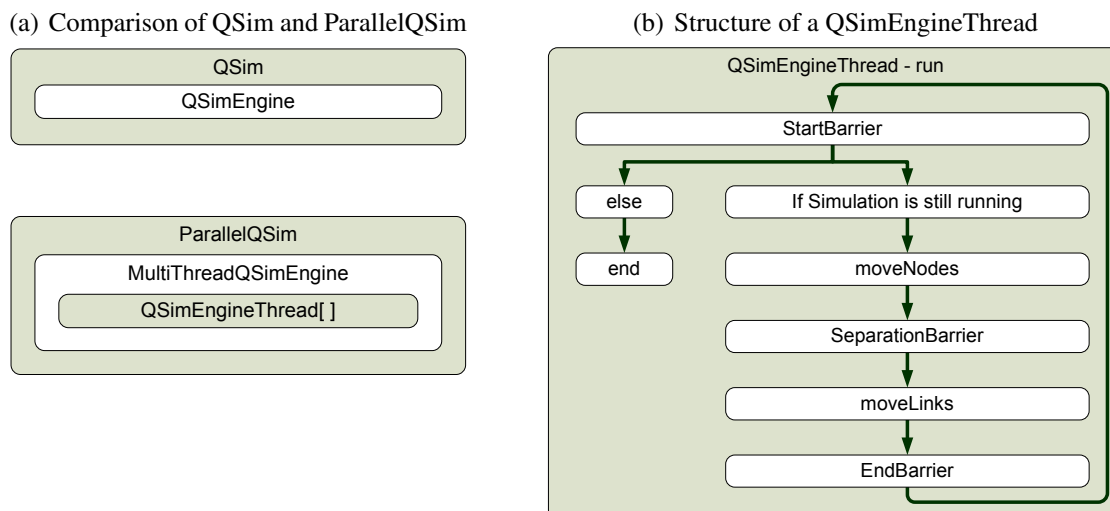
357 The *MultiThreadQSimEngine* is a wrapper class that manages the communication between
358 the *ParallelQSim* and the *QSimEngineThreads*. As a result of that structure, the *ParallelQSim*
359 sees only the *MultiThreadQSimEngine* and is not involved in the handling of the threads—it
360 does not even recognize that there are multiple threads involved in the simulation.

361 The *QSimEngineThreads* are created once per iteration of the simulation and reused in every
362 sim step which is considerably faster than creating new threads in each sim step. As shown in
363 Figure 3(b) this is realized by two *CyclicBarriers* (*StartBarrier* and *EndBarrier*) that are part of
364 the Java *concurrent* package. A third *CyclicBarrier* (*SeparationBarrier*) is used to synchronize
365 the *moveNodes* and *moveLinks* actions. The threads must have handled all their nodes before
366 they can continue with the links. When the *doSimStep* method of the *MultiThreadQSimEngine*
367 is called, it starts the threads by triggering the *StartBarrier* and then waits until all threads have
368 reached the *EndBarrier*.

PERFORMANCE MEASUREMENTS

369 **Hardware**

370 The experiments employed to compare the performance of the *ParallelQSim* with the existing
371 *QSim* are run on a computer with two quad core CPUs (each a AMD Opteron 2380) and 24 GB
372 of shared memory. A maximum of 7 cores is used for the *ParallelQSim*. The remaining core is
373 used for (parallel-)events handling and some background processes.

FIGURE 3 Structure of the Implementation

374 Scenarios

375 As a first scenario, a 1% example of Berlin is used which is a basic example scenario used by
 376 MATSim. It contains about 16K agents who perform 28K trips and is simulated on a network
 377 with about 11k nodes and 28K links. During a simulation run 1M events are created.

378 For the second and third scenario a model of Canton Zurich is used—once as 25% sample
 379 with 400K agents and 1.3M performed trips and once as 100% sample with 1.6M agents and
 380 5.1M performed trips. The network contains 73K nodes and 163K links. A simulation run
 381 creates 47M and 158M events, respectively.

382 These are real world scenarios that are typically simulated with MATSim. It is assumed, that
 383 the results of the performance measurements can be reached on other, comparable, scenarios
 384 as well.

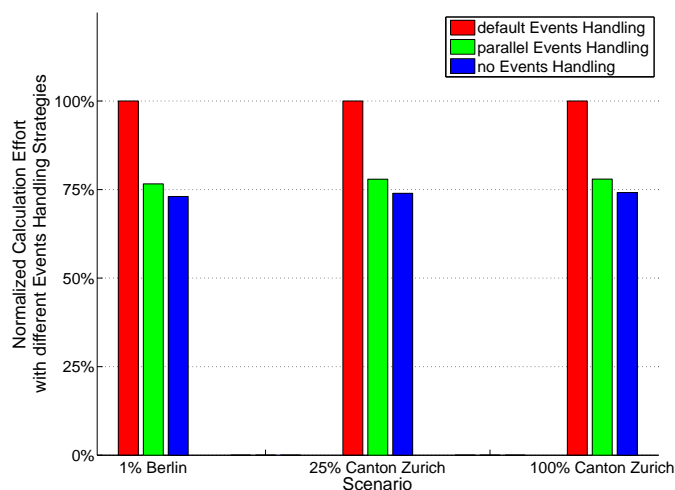
385 Results

386 The *ParallelQSim* uses the same simulation logic as *QSim*. However, simulation results pro-
 387 duced by the *ParallelQSim* are slightly different from the ones created by *QSim*, which is a
 388 result of using multiple Random objects instead of a single one. From a traffic planning point
 389 of view the results are absolutely comparable and therefore the results of the simulations runs
 390 in this section are only analyzed regarding the performance of the used simulation setup (queue
 391 simulation and events handling strategy). Conclusions concerning the results from a traffic
 392 planning view have already been drawn (21, 22).

393 Figure 4 shows the calculation effort for the events handling in the three scenarios. The
 394 results show that the effort constitutes 25% of the total calculation effort and is not influenced
 395 by the size of the scenario. According to Amdahl's Law (31), which describes the maximum
 396 achievable speedup of a programm with partially parallelized code, this affects significantly the
 397 performance gain reachable. The influence of the non-parallel code can be illustrated with a
 398 simple example. If code that consumes 5% of the computation time of a program cannot be
 399 parallelized, the total calculation time cannot be reduced by more than a factor twenty—even

400 if the remaining code could be handled in zero seconds. As a result, the events handling limits
 401 the possible speed gain to a factor four of the calculation time of *QSim* with non-parallel events
 402 handling. Relative to the runs of *QSim* with parallel events handling, a performance gain of
 403 factor three is possible.

FIGURE 4 Performance of different Events Handling Strategies



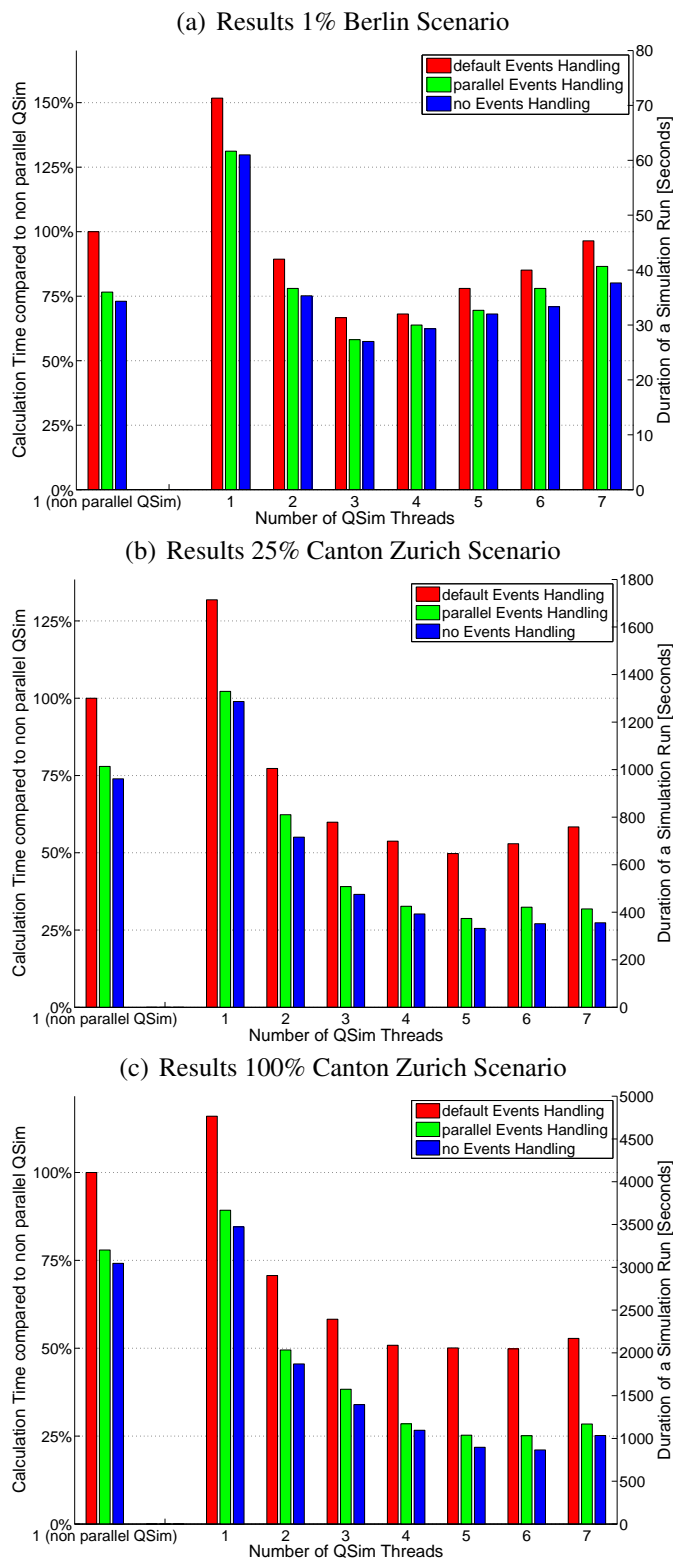
404 Another important finding, that is also depicted in Figure 4, is the high efficiency of *ParallelEventsManager* in combination with *QSim*. Almost the entire calculation effort of the events
 405 handling is moved from the main thread to a separate thread. As a result, the simulation is
 406 nearly as fast as it would be without any events handling.
 407

408 Figures 5(a) to 5(c) show the results of the runs with the three test scenarios. Each figure
 409 contains the results of runs employing *ParallelQSim* using one to seven cores and different
 410 event handling strategies. Additionally the same scenarios have been run with the non-parallel
 411 *QSim*.

412 When the computation times of the *QSim* and the *ParallelQSim* using only one thread are
 413 compared, the difference between the calculation times is the overhead caused by the paral-
 414 lelization such as distributing and synchronizing data between the threads. In the Berlin sce-
 415 nario, a significant overhead of over 50% is found. As a result, the *ParallelQSim* cannot reduce
 416 the calculation time significantly compared to the *QSim*. However, the *ParallelQSim* itself per-
 417 forms quite well. The calculation time decreases by 50% if three threads are used instead of
 418 a single one. The Canton Zurich scenarios show that the calculation overhead is less signif-
 419 icant if the scenario gets more complex. The overhead reduces to 30% (25% scenario, using
 420 5 threads) and 20% (100% scenario, using 6 threads). Therefore, the performance gains rise
 421 up to the—according to Amdahl’s Law—highest reachable value of a factor three when using
 422 parallel events handling.

423 The comparison of the results of Canton Zurich runs with parallel events handling and runs
 424 without events handling shows that there is still only a small difference in computation time.
 425 Hence, we can assume that the computation times of the events handling and the *ParallelQSim*
 426 are almost alike in these scenarios. The runs without events handling would have a notice-
 427 able shorter computation time if the events handling had become a bottleneck. However, if
 428 the simulated scenarios get even bigger events handling could clearly become a performance

FIGURE 5 Results of the Sample Scenarios



429 bottleneck.

430 Considering only the results of the *ParallelQSim*, the number of cores used which results

431 in the lowest calculation times rises with the total calculation effort for the scenario. While the
432 Berlin scenario performs best with only three cores, the 25% Canton Zurich scenario should
433 be run with five cores and the 100% Canton Zurich scenario benefits from up to six cores. An
434 important detail is that using too many cores results in increased computation times which is a
435 consequence of the synchronization effort that increases with the number of used cores.

CONCLUSION AND OUTLOOK

436 This paper describes the the development and implementation of a new simulation module in
437 MATSim that reaches short calculation times by using multiple CPU cores. An adaption of
438 the *QSim* was chosen because its time steps can be used to synchronize the data between the
439 parallel calculation threads. Distributing the workload is done by a simple approach where the
440 assignment of the network's links and nodes to the threads is done randomly.

441 The results of the performance tests show that—depending on the scenario size—the cal-
442 culation time can be reduced by a factor of four. Based on Amdahl's Law it is shown that the
443 events handling could become a bottleneck when simulating large scale scenarios. A speedup
444 of more than a factor of four is not possible. However, it is shown that the existing parallel
445 events handling reduces the calculation effort in the main thread very efficiently. Moving the
446 events that are created in the main thread to another thread where they are handled is done
447 within negligible time.

448 Another important results of the performance tests concerns the number of CPU cores used.
449 Depending on the complexity and size of the scenario the number of cores resulting in the best
450 simulation performance varies. Hence, a user should keep the results of the sample scenarios
451 in mind when choosing the number of cores for another scenario. Comparing the scenario with
452 the given samples in this paper should lead to a reasonable choice.

453 Although the results are already very satisfying there are still some further performance
454 optimizations possible and desirable. One major point concerns the synchronization effort
455 between the threads. Especially in smaller scenarios, this performance bottleneck reduces the
456 attractiveness of using the *ParallelQSim*. Another topic for further developments is the events
457 handling. Having a setup where each thread has its own set of events handlers would reduce
458 the amount of synchronized method calls significantly and therefore should results in further
459 performance gains.

REFERENCES

- 460 1. Charypar, D., K. W. Axhausen and K. Nagel (2007) An event-driven queue-based traffic
461 flow microsimulation, *Transportation Research Record*, **2003**, 35–40.
- 462 2. Ferber, J. (1999) *Multi-Agent Systems: An Introduction to Distributed Artificial Intelli-*
463 *gence*, Addison-Wesley, Boston.
- 464 3. Wooldridge, M. (2000) *Reasoning about Rational Agents*, MIT Press, Cambridge.
- 465 4. Klügl, F. (2001) *Multiagentensimulation - Konzepte, Werkzeuge, Anwendung*, Addison-
466 Wesley, Munich.
- 467 5. Eymann, T. (2003) *Digitale Geschäftsagenten - Softwareagenten im Einsatz*, Springer,
468 Berlin.

- 469 6. Barceló, J., J. L. Ferrer, D. Garcia, M. Florian and E. Le Saux (1998) Microscopic traffic
470 simulation, in P. Marcotte and S. Nguyen (eds.) *Equilibrium and Advanced Transportation*
471 *Modelling*, chap. 1, 1–26, Kluwer, Dordrecht.
- 472 7. Nökel, K. and M. Schmidt (2002) Parallel DYNEMO: Meso-scope traffic flow simulation
473 on large networks, *Networks and Spatial Economics*, **2** (4) 387–403.
- 474 8. Nagel, K. and M. Rickert (2001) Parallel implementation of the TRANSIMS micro-
475 simulation, *Parallel Computing*, **58** (2) 1611–1639.
- 476 9. Niedringhaus, W. P., J. M. Opper, L. Rhodes and B. L. Hughes (1994) Ivhs traffic modeling
477 using parallel computing: Performance results, in H. J. Siegel (ed.) *Proceedings of the*
478 *8th International Symposium on Parallel Processing*, 688–693, IEEE Computer Society,
479 Washington, D.C.
- 480 10. Cameron, G. D. B. and G. I. D. Duncan (1996) Dynamic process and equilibrium in trans-
481 portation network: Towards a unifying theory, *Journal of Supercomputing*, **10** (1) 25–53.
- 482 11. Fujimoto, R. M. (2001) Parallel and distributed simulation systems, in B. A. Peter, J. S.
483 Smith, D. J. Medeiros and M. W. Rohrer (eds.) *WSC '01: Proceedings of the 33rd confer-*
484 *ence on Winter simulation*, 147–157, IEEE Computer Society, Washington, D.C.
- 485 12. Robinson, S. (2004) *Simulation: The Practice of Model Development and Use*, John Wiley
486 & Sons, Chichester.
- 487 13. Hartrum, T. C. and B. J. Donlan (1988) HYPERSIM: Distributed discrete-event simulation
488 on an iPSC, in G. Fox (ed.) *Proceedings of the third conference on Hypercube concurrent*
489 *computers and applications: Architecture, software, computer systems, and general issues*,
490 vol. 1, 745–747, Association for Computing Machinery, New York.
- 491 14. Ferscha, A. (1996) Parallel and distributed simulation of discrete event systems, in A. Y. H.
492 Zomaya (ed.) *Parallel and Distributed Computing Handbook*, 1003–1041, McGraw-Hill,
493 New York.
- 494 15. Fujimoto, R. M. (1989) Parallel discrete event simulation, in E. A. MacNair, K. J. Mus-
495 selman and P. Heidelberger (eds.) *WSC '89: Proceedings of the 21st conference on Winter*
496 *simulation*, 19–28, Association for Computing Machinery, New York.
- 497 16. Logan, B. and G. Theodoropoulos (2001) The distributed simulation of multi-agent sys-
498 tems, *Proceedings of the IEEE*, **89** (2) 174–186.
- 499 17. Charypar, D., K. W. Axhausen and K. Nagel (2007) An event-driven parallel queue-based
500 microsimulation for large scale traffic scenarios, paper presented at the *11th World Confer-*
501 *ence on Transportation Research*, Berkeley, June 2007.
- 502 18. Hanebutte, U. R. and A. M. Tentner (1995) Traffic simulations on parallel computers using
503 domain decomposition techniques, paper presented at the *2nd World Congress on Intelli-*
504 *gent Transport Systems*, Yokohama, November 1995.
- 505 19. Willebeek-LeMair, M. H. and A. P. Reeves (1993) Strategies for dynamic load balancing
506 on highly parallel computers, *IEEE Transactions on Parallel and Distributed Systems*, **4** (9)
507 979–993.

- 508 20. Balmer, M. (2007) Travel demand modeling for multi-agent traffic simulations: Algo-
509 rithms and systems, Ph.D. Thesis, ETH Zurich, Zurich, May 2007.
- 510 21. Balmer, M., M. Rieser, K. Meister, D. Charypar, N. Lefebvre, K. Nagel and K. W. Ax-
511 hausen (2008) MATSim-T: Architektur und Rechenzeiten, paper presented at the *Heureka*
512 '08, Stuttgart, March 2008.
- 513 22. Meister, K., M. Balmer, F. Ciari, A. Horni, M. Rieser, R. A. Waraich and K. W. Axhausen
514 (2010) Large-scale agent-based travel demand optimization applied to Switzerland, in-
515 cluding mode choice, paper presented at the *12th World Conference on Transportation*
516 *Research*, Lisbon, July 2010.
- 517 23. Charypar, D. and K. Nagel (2005) Generating complete all-day activity plans with genetic
518 algorithms, *Transportation*, **32** (4) 369–397.
- 519 24. Feil, M. (2010) Choosing the daily schedule: Expanding activity-based travel demand
520 modelling, Ph.D. Thesis, ETH Zurich, Zurich.
- 521 25. Horni, A., D. M. Scott, M. Balmer and K. W. Axhausen (2009) Location choice modeling
522 for shopping and leisure activities with MATSim: Combining micro-simulation and time
523 geography, *Transportation Research Record*, **2135**, 87–95.
- 524 26. Cetin, N. (2005) Large-scale parallel graph-based simulations, Ph.D. Thesis, ETH Zurich,
525 Zurich.
- 526 27. Neumann, A. (2008) Modellierung und Evaluation von Lichtsignalanlagen in Queue-
527 Simulationen, Master Thesis, Technical University Berlin, Berlin.
- 528 28. Rieser, M. (2010) Adding transit to an agent-based transportation simulation, Ph.D. Thesis,
529 Technical University Berlin, Berlin.
- 530 29. Dobler, C. (2009) Implementations of within day replanning in MATSim-T, *Working Pa-*
531 *per*, **598**, IVT, ETH Zurich, Zurich.
- 532 30. Waraich, R. A., D. Charypar, M. Balmer and K. W. Axhausen (2009) Performance im-
533 provements for large scale traffic simulation in MATSim, paper presented at the *9th Swiss*
534 *Transport Research Conference*, Ascona, September 2009.
- 535 31. Amdahl, G. M. (1967) Validity of the single processor approach to achieving large scale
536 computing capabilities, paper presented at the *Spring Joint Computer Conference*, New
537 York, April 1967.