Multi-Resolution Traffic Simulation

for Connected Car Applications using VIRES VTD

By **AUDI AG:** Andreas Kern
Technical University of Munich: Manuel Schiller
Institute of Transportation Systems: Daniel Krajzewicz
VIRES, part of Hexagon: Marius Dupuis

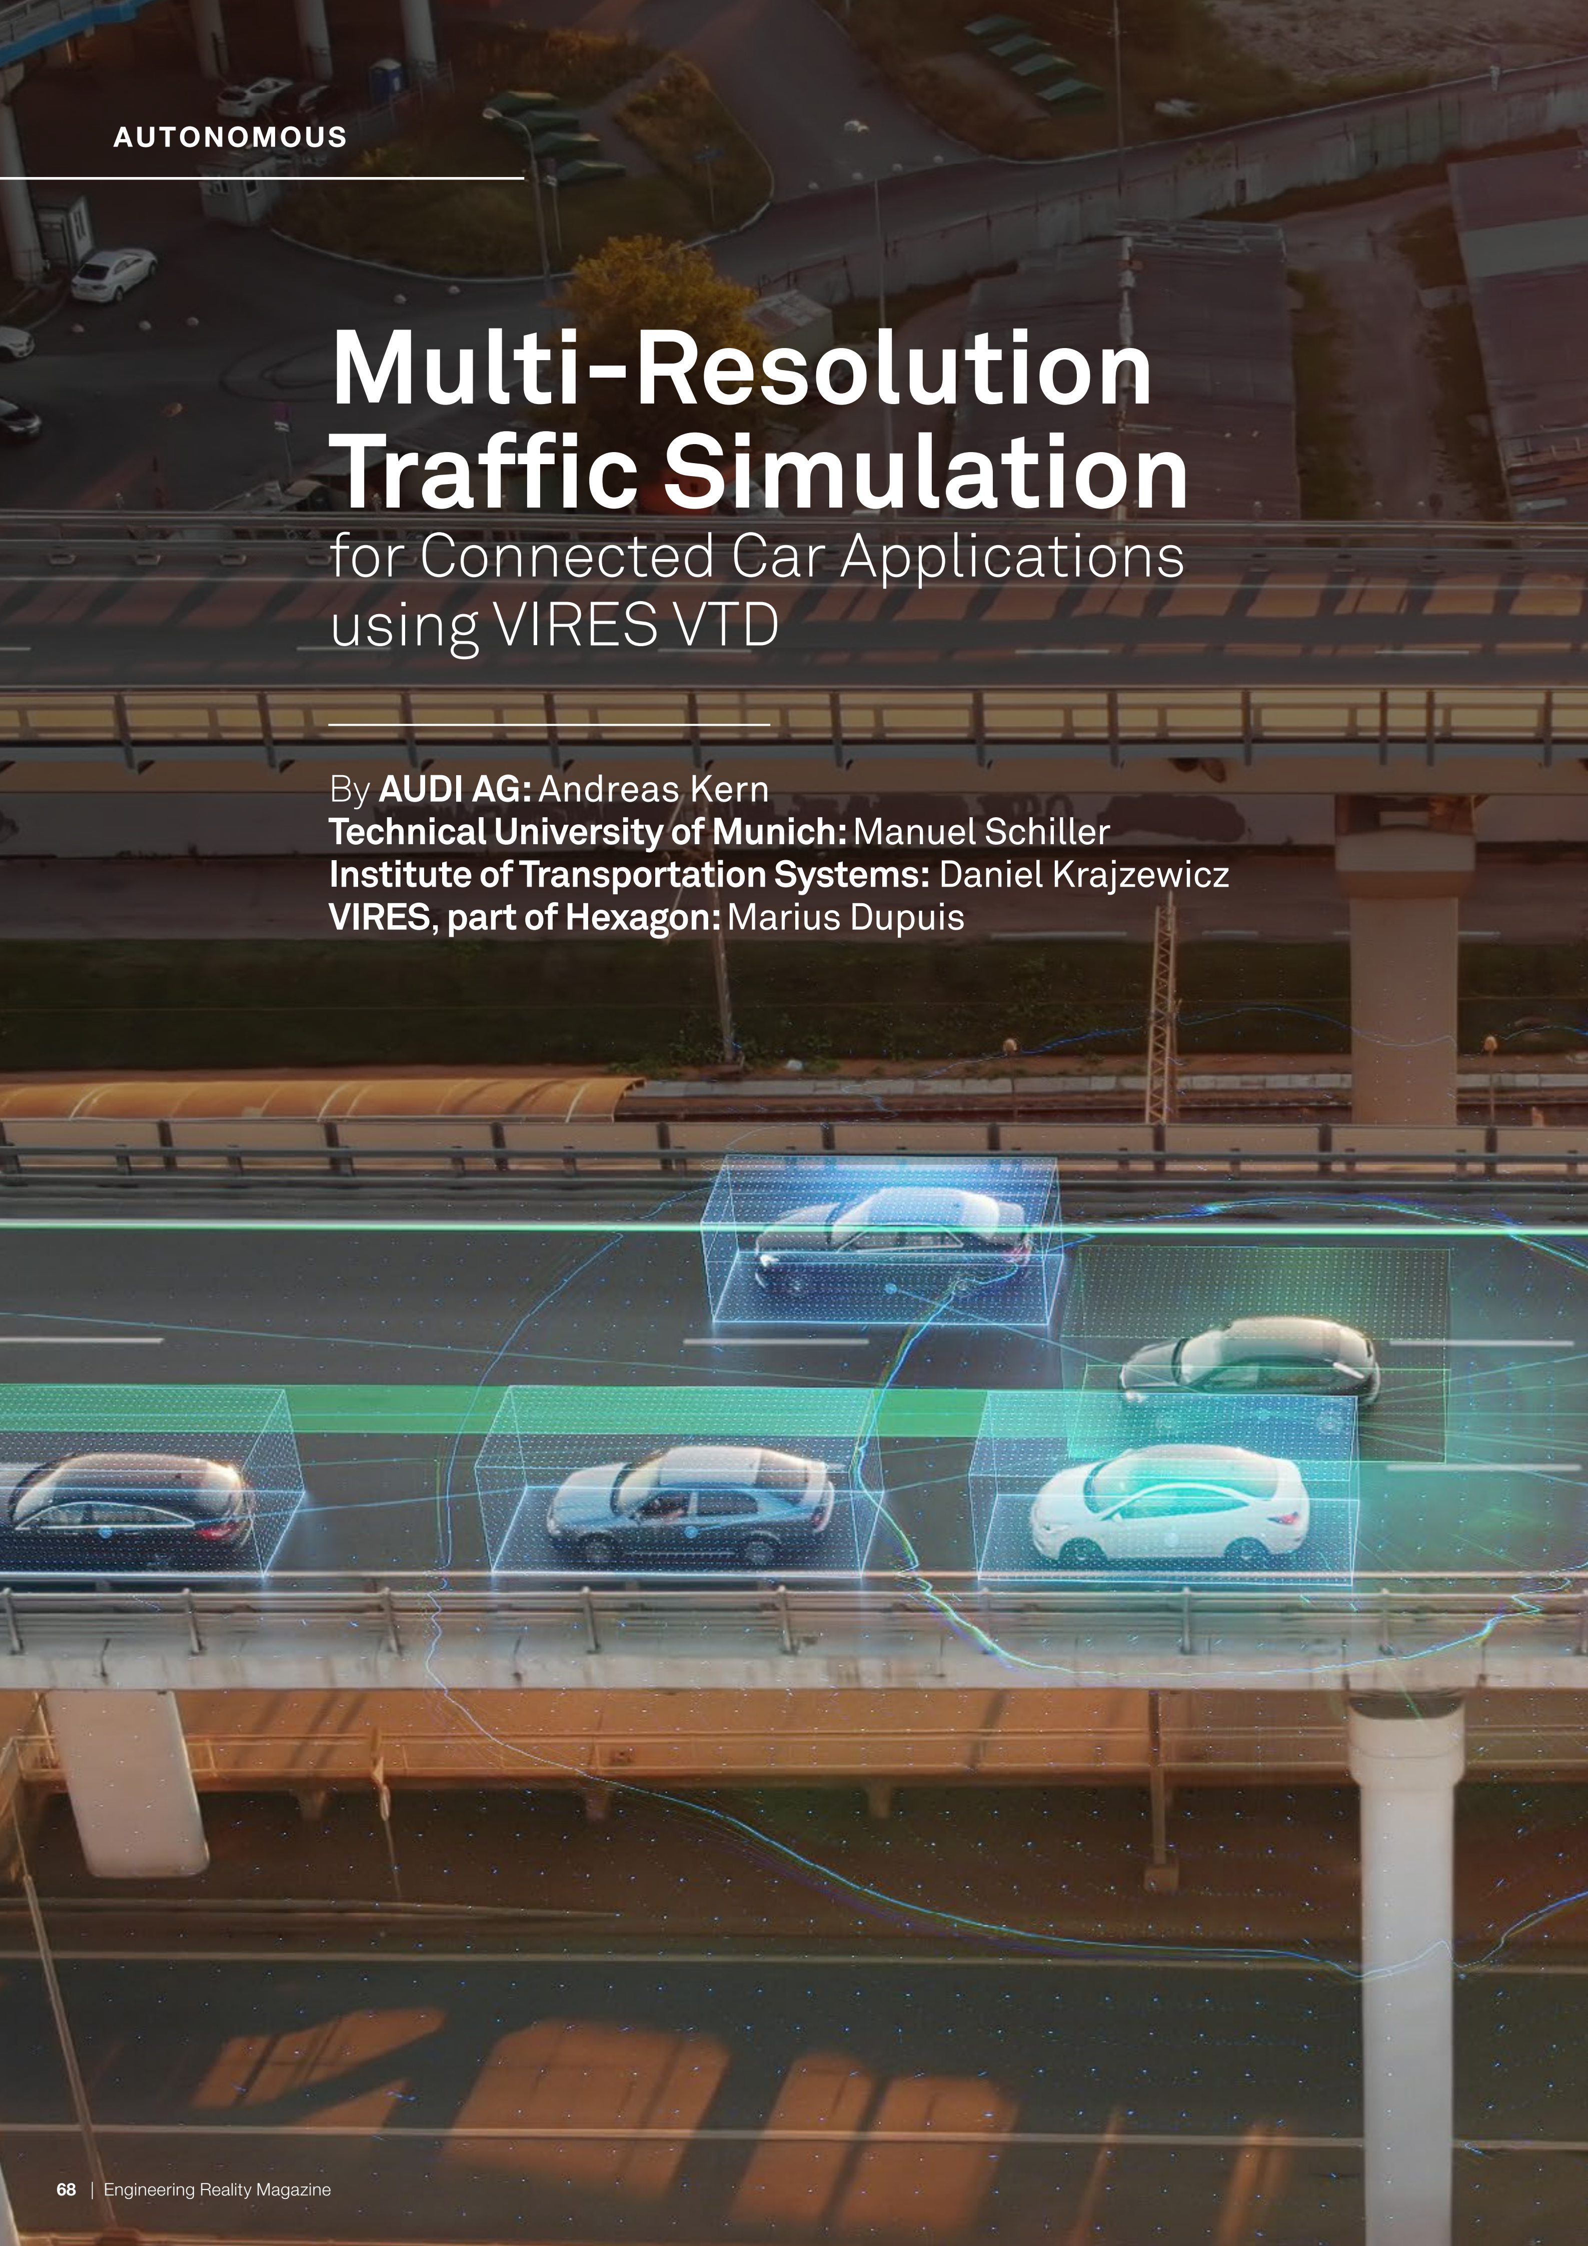

Vehicular ad hoc networks (VANETs) have attracted a lot of research attention over the last few years because they have the potential to improve traffic safety, efficiency and driver comfort. In fact, several ADAS (Advanced Driver Assistance Systems) applications, such as cooperative driving and subsequently automated driving, can only be achieved through wireless communication between the vehicles on the road.

Since those systems often exhibit safety-critical features, rigorous testing and validation must be completed before their mass adoption. Although real road tests using physical prototype vehicles offer the highest degree of realism, the large amount of resources needed to perform large-scale and extensive testing of vehicular networks renders their use impossible. Simulations are essential to validate the performance of such solutions in large-scale virtual environments. Furthermore, simulation-based evaluation techniques are invaluable for testing those complex systems in a wide variety of dangerous and critical scenarios without putting humans at risk.

In the automotive industry, the use of simulation (Figure 1) is well established in the development process of traditional driver assistance and active safety systems, which primarily focus on the simulation of individual vehicles with a very high level of detail. When investigating and evaluating the performance of ADAS based on vehicular communication, this isolated view of a single vehicle alone or a small number of vehicles in the simulation is not sufficient anymore. Potentially, every vehicle equipped with wireless communication technology could be coupled in a feedback loop with the other road users participating in the vehicular network, and therefore, the number of influencers that need to be considered is drastically increased.

These considerations lead to a trade-off between accuracy in terms of the simulation details for each vehicle and scalability in terms of the number of vehicles that can be simulated with the available computing resources. This article presents an approach to solve this trade-off by coupling multiple resolutions of traffic simulation to get highly accurate simulation results where they are needed, and simultaneously achieve an efficient simulation of large-scale scenarios of the surrounding environment.

The Developed Multi-Resolution Traffic Simulation

Microscopic Traffic Simulator: SUMO

We chose to use Simulation of Urban MObility (SUMO) as the traffic simulator responsible for the simulation of the low-resolution area (LRA). SUMO is a microscopic, space-continuous, and time-discrete simulator. While it is employed in a wide range of research domains, its most notable use is shown in a high number of research papers regarding VANET simulations. SUMO is well known for its high execution speed, as well as for its extensibility. SUMO is ideally suited to simulate a high number of vehicles residing in the LRA due to its efficiency, which is partly achieved through its simplified driver model (which determines the path a vehicle will take).

Nanoscopic Traffic and Vehicle Simulator: VIRES Virtual Test Drive

We employ the nanoscopic traffic and vehicle simulator VIRES Virtual Test Drive (VTD) for the simulation of the high-resolution vehicles. VTD was developed

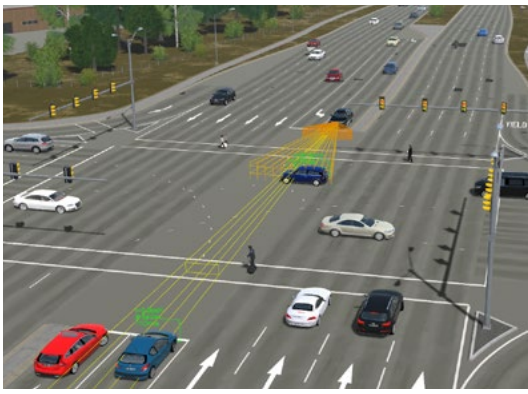

Figure 1: **ADAS simulation**

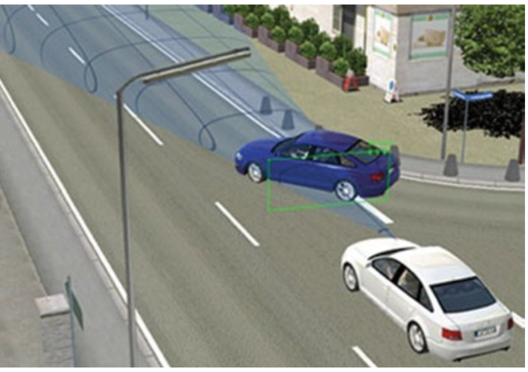

Figure 2: **3D visualization of a simulated RADAR sensor in VIRES VTD**

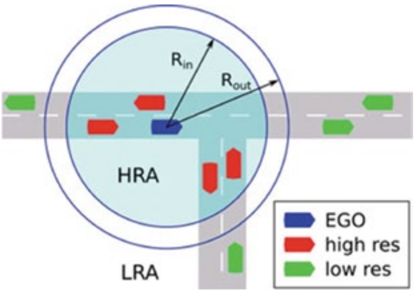

Figure 3: **Dynamic partitioning of the simulation area**

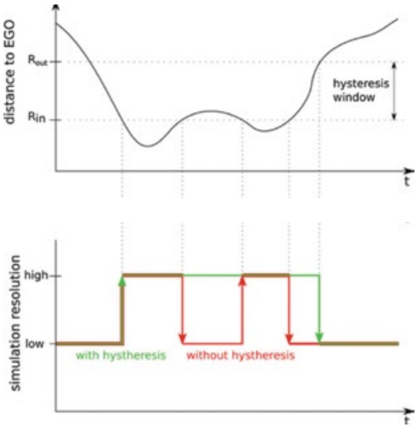

Figure 4: **Comparison of simulation resolution switching**

for the automotive industry as a virtual test environment used for the development of ADAS and Autonomous Vehicles. Its focus lies on the interactive high-realism simulation of driver behaviour, vehicle dynamics, and sensors. VTD is highly modular, so any standard component may be exchanged by a custom and potentially more detailed implementation. Its standard driver model is based on the intelligent driver model; however, an external driver model may be applied if necessary. The same concept applies to the vehicle dynamics simulation, where the standard single-track model can be substituted by an arbitrarily complex vehicle dynamics model adapted for specific vehicles. Each simulated vehicle can be equipped with arbitrary simulated sensors, for example a RADAR sensor, which is shown in Figure 2.

Offline Pre-processing To Enable Coupling

The two simulators rely on different data formats representing the modelled road network. In order to be able to run a co-simulation of both simulators, the underlying data basis must match. VTD uses the OpenDRIVE format to specify the road network. This specification models the road geometry as realistically as possible by using analytical definitions. SUMO on the other hand approximates the road network geometry by line segments. There are also differences in the modelling of intersections and lane geometries. To achieve a matching database, we convert the road network in an offline pre-processing step from OpenDRIVE to the file format SUMO supports.

Online Coupling and Synchronization

The coupling of the simulators at simulation runtime is based on the master-slave principle. Figure 4 shows this sequence of operations during a single simulation step, in which VTD and SUMO can operate with different temporal resolutions without losing synchronization.

S_{VTD} is the length of a time step for the high-resolution area (HRA), whereas S_{SUMO} is the length of a time step for the low-resolution area (LRA). Typically, the nanoscopic simulation is run at a higher frequency than the microscopic one. T_{VTD} and T_{SUMO} respectively denote the local simulation time in each simulator. At the beginning of each simulation step, a new timestep is simulated in VTD. If the next timestep has been reached for SUMO and the condition $T_{VTD} \geq T_{VTD} + S_{SUMO}$ is therefore fulfilled, the state of the high-resolution vehicles is sent to SUMO through a gateway. This triggers the simulation of the next timestep in the low-resolution model, and as a result, the positions of the low-resolution vehicles are passed back. These vehicles are now classified, and, if applicable, the change of resolution is performed for individual vehicles. When an exchange of a vehicle between the simulators happens, the previously mentioned inherent difference in the underlying road network may cause problems if a vehicle cannot be mapped based on its position in a specific lane due to differences in accuracy. This is especially true for complex intersections which are modelled quite differently.

After all the resolution changes have been successfully completed, the simulation is unblocked again and the next timestep can be simulated. This synchronization is very important to ensure reproducible simulation results across multiple simulation runs.

Dynamic Spatial Partitioning of The Simulated Area

Our approach aims to couple traffic simulation models of different resolutions at dynamic regions of interest. Contrary to conventional traffic simulation, we are not interested in investigating a large number of vehicles from a bird's eye perspective, but the focus is rather on a single vehicle (or a limited number of vehicles) which are used to conduct a test drive in the virtual environment. This vehicle of interest has the ADAS system under investigation onboard, and is referred to as the EGO car. The simulated measurements and sensor values are fed into the ADAS, and depending on its type and its use case, the respective ADAS directly or indirectly influences the vehicle's state and behaviour.

Based on this distance criterion, an area of interest is defined that centres around the EGO car, and in which the defined simulative high-fidelity requirements must be fulfilled. Since the EGO car is driving continuously through the virtual environment, this area of interest is likewise being moved along. We therefore partition the global area of the simulation dynamically into a high-resolution area (HRA) and a low-resolution area (LRA). Figure 3 shows a schematic view of the dynamic spatial partitioning. There, the HRA is defined as the area of a circle which is centred around the EGO vehicle. Red vehicles are within that circle and are therefore simulated in high resolution by the involved nanoscopic simulator, whereas the green vehicles are outside of the circle and are consequently simulated in low resolution by a microscopic simulation. All vehicles exist in the microscopic simulation, but in the nanoscopic simulation contains only the high-resolution vehicles, and their movements are applied to their proxy counterparts in the microscopic simulator.

Due to the dynamic nature of road traffic, the EGO car, the high-resolution vehicles as well as the low-resolution vehicles are permitted to move continuously. The classification of the assigned resolution mode is therefore performed after each time step of the simulation. Vehicles for which the classification has led to a change in resolution are transferred to the appropriate simulator. This change of resolution is possible in both directions at every time step. However, since the HRA is defined to be centred around the EGO car, it is always simulated in high resolution.

In order to prevent vehicles which are close to the boundary between HRA and LRA from oscillating very frequently between the two resolution areas, a hysteresis controller as depicted in Figure 5 is applied in the classification process. As shown in Figure 3, the two thresholds R_{in} and R_{out} are defined. A vehicle is transferred into the high-resolution simulation only if its distance to the EGO car falls below the value of R_{in}. The exchange back to the low-resolution simulation is carried not out until the threshold R_{out} is exceeded.

Simulation & Evaluation

Scenario and Simulation Setup

A synthetic scenario was created for testing the coupling concept and evaluating its performance. It consists of a single straight road running west to east with a length of 50 km and two lanes, one for each direction. Each lane is configured to have a constant inlet of 1,000 vehicles per hour heading either east or west. The EGO car is located near the start of the road. It drives from west to east and is followed by a traffic flow, heading towards the oncoming traffic flow. This artificial road was first modelled in the OpenDRIVE format and was then converted to the SUMO road network format.

> **VTD is highly modular, so any standard component may be exchanged by a custom and potentially more detailed implementation.**

We performed two series of experiments. In the first series, the nanoscopic traffic simulator VTD was applied to the whole simulated area. In the second series, we used the described multi-resolution concept to partition the simulation area between VTD and SUMO. We chose a timestep of S_{VTD} 20 ms for the high-resolution area in VTD and a timestep of S_{SUMO} 1 s for the low-resolution area in SUMO. The hysteresis thresholds which define the dynamic area of interest were set to R_{in} 500 m and R_{out} 550 m.

Performance Evaluation

We measured the duration it takes to perform each simulation step over the simulation period of 1,800 s, while the number of vehicles was constantly being increased. Each series consists of five separate simulation runs to account for fluctuations in the measured execution times. To illustrate the trends of the measurements more clearly, the moving average is also displayed in the following figures.

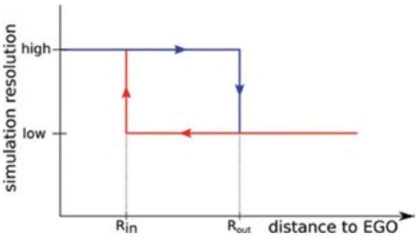

Figure 5: **Hysteresis control of the simulation resolution**

Figure 6 shows the performance development of the nanoscopic simulation while increasing the simulated vehicle count over the simulation period. The duration of each simulation step is almost constant up to a count of 70 vehicles. Until then, the duration is around 12 ms, which is less than the timestep length of 20 ms and therefore yet fulfils the real-time constraint. At around 150 vehicles, the duration is beyond these 20 ms and real-time simulation is not possible anymore. With increasing vehicle count, the duration for each timestep also considerably increases and reaches 180 ms at the end of the simulation period. This results in a factor 15 computation time increase compared to the amount of at the beginning of the simulation. The overall simulation took over 120 min to complete, which is four times more than the simulated time.

Figure 7 shows the performance development of the multi-resolution simulation in the same simulation scenario over the same simulation period. While the total vehicle count is increased the same way as in the pure nanoscopic simulation, the separately plotted nanoscopic vehicle count illustrates the number of cars

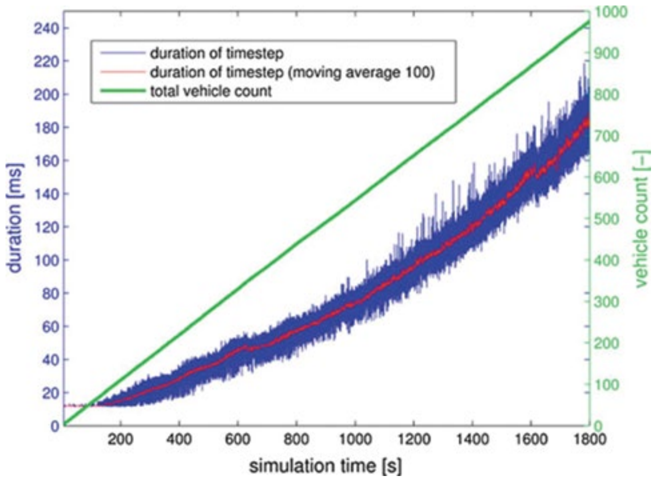

Figure 6: **Simulation performance—nanoscopic simulation only**

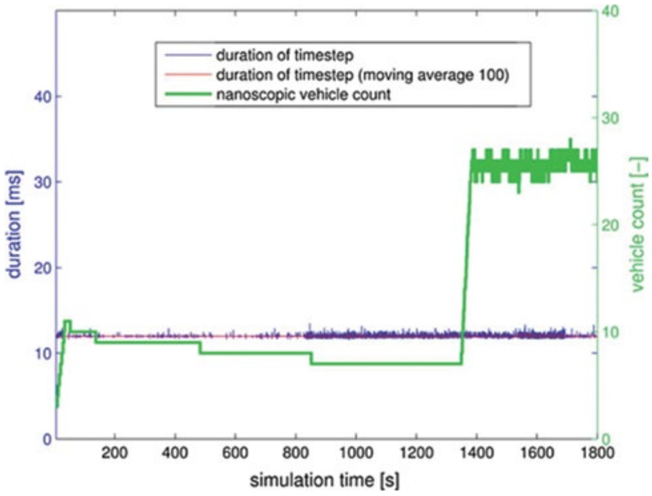

Figure 7: **Simulation performance: multi-resolution simulation**

which are within the high-resolution area. It shows that reducing the nanoscopic model's area of interest fulfils the aim of reducing the overall simulation time. After a local maximum of 11 nanoscopic cars is reached, this count decreases slowly since slower vehicles are left behind the faster moving EGO car. At around simulation time 1,350 s, the two traffic flows from each end of the road meet in the middle of the road, which then increases the nanoscopic vehicle count. However, because the extent of the HRA is limited, the nanoscopic vehicle count does not exceed a certain limit, which for the given configuration is at around 27 vehicles. The duration for the timesteps stays on average constant around 12 ms, so the overhead resulting from the coupling of the two simulators is negligible. The execution time of the microscopic simulator is also shown to be negligible due to its less detailed, yet much more efficient, simulation model. The overall simulation took less than 18 min to complete, so the simulation was faster than real time by factor 1.66 and the real-time constraint was fulfilled throughout the whole simulation period.

Conclusion

In this article, we proposed a concept for coupling traffic simulators of different simulation resolutions to achieve a multi-resolution traffic simulation which focuses on a dynamically-determined area of interest. The presented methodology partitions the simulation area into a variable, highly detailed region of interest represented by a nanoscopic model, with VIRES Virtual Test Drive (VTD), and the surrounding area simulated at low resolution by a microscopic model. The evaluation shows a dramatic reduction of computation time in comparison with a pure nanoscopic simulation of the same simulation dimensions, which even makes real-time simulation possible. This divide-and-conquer strategy enables accurate, realistic, and large-scale testing and validation of real implementations of driver assistance systems based on vehicular networks in a virtual environment. As the next step, we are investigating the application of the multi-resolution simulation methodology for the other domains relevant for the simulation of vehicular networks, namely network simulation and application emulation, to model the whole system across all domains efficiently at high fidelity.

Reference

"Multi-resolution Traffic Simulation for Large-Scale High-Fidelity Evaluation of VANET Applications", Manuel Schiller, Marius Dupuis, Daniel Krajzewicz, Andreas Kern and Alois Knoll, © 3rd SUMO Conference 2015 Berlin, Germany

This divide-and-conquer strategy enables accurate, realistic, and large-scale testing and validation of real implementations of driver assistance systems based on vehicular networks in a virtual environment.