

Simulation driven design of the German toll system – evaluation and enhancement of simulation performance

Tommy Baumann*, Bernd Pfitzinger^{†‡}, Thomas Jestädt[†]

* Andato GmbH & Co. KG, Ehrenbergstraße 11, 98693 Ilmenau, Germany. Email: tommy.baumann@andato.com

[†] Toll Collect GmbH, Linkstraße 4, 10785 Berlin, Germany. Email: {bernd.pfitzinger|thomas.jestaedt}@toll-collect.de

[‡] FOM Hochschule für Oekonomie & Management, Zeltnerstraße 19, 90443 Nürnberg, Germany

Abstract—We study the performance issues of a realistic simulation of the German toll system for heavy trucks using a discrete-event system (DES) simulation. The article first introduces the German toll system and the simulation framework developed to analyze the systems' behavior. A number of performance limitations of several commercial and non-commercial DES simulation kernels are described, measured and benchmarked. The application-level performance of a DES implementation of the German toll system is then compared using two commercial DES tools and several optimizations are introduced both on the simulation model and kernel level to achieve the necessary performance for a detailed and realistic simulation of a fleet of 700 000 trucks.

I. INTRODUCTION

SIMULATIONS are a vital step in the design of systems or the assessment of planned changes [1], [2] - reducing the inherent risk of ongoing system development and allowing for a faster system deployment. Additionally simulations predict the dynamic system behavior which can become highly non-linear or chaotic even for simple systems [3].

To specify and evaluate the German toll system, a simulation-driven design approach has been selected [4]. The approach is characterized by applying modeling and simulation technologies in the early design stages, i.e. at a time when most of the important design decisions have to be made. As a result both the systems and processes are specified in the form of executable models. The approach allows to validate and optimize the overall system architecture already in the specification phase – avoiding expensive integration issues in the subsequent implementation and integration phases.

Consequently, specification speed and quality is considerably increased while the system and product uncertainty is decreased. It is noteworthy that simulation-driven design not only refers to the system under design but also includes the surrounding design process, i.e. the process is also captured as an executable specification which allows automating design steps like architecture optimization, validation against operational scenarios, and tracking of design decisions.

The outline of the article is as follows: Section II gives an overview of the automatic German toll system, the corresponding simulation model and typical simulation results. Section III introduces the performance properties of discrete event simulations and discusses appropriate performance metrics

and benchmarking models. This is followed in section IV by a discussion of the simulation performance and scaling of several discrete-event simulation tools for basic simulation operations. Using an existing microscopic holistic execution specification of the automatic German toll system [5] we describe several performance optimizations both on level of the simulation model and the simulation kernel in section V.

II. EXECUTABLE SPECIFICATION OF THE GERMAN TOLL SYSTEM

Toll Collect GmbH is the provider of the German electronic toll for heavy goods vehicles (HGVs). The system automatically collects the toll fees on federal motorways using an on-board-unit (OBU) installed in most of the trucks¹. Currently there are more than 700 000 OBUs deployed, each determining the toll fees according to an up-to-date map of the chargeable roads using a GNSS receiver coupled to the vehicles' speed and directional data and communicating via GSM with the Toll Collect data center. In total, the HGVs drove $26.7 \cdot 10^9$ km on the chargeable federal motorways in 2011, [6]).

Following the idea of simulation-driven design we use executable models to analyze and evaluate the behavior of the IT systems of Toll Collect GmbH. Due to the necessity to express dynamically changing system states (over time) we selected the execution domain Discrete Event Simulation (DES). The simulation model of the Toll Collect system is used to predict the behavior of the current system as well as effects of changes to the system, especially to maintain the high level of accuracy (with an error rate of less than 1 in 1 000, [7]) needed due to service-level-agreements. Changes to the Toll Collect system occur every day – in the past four years more than 15 major changes (releases) and more than 1 500 medium-sized changes were implemented.

A. Modeling of the Toll Collect system

The simulation model of the Toll Collect system consists of three blocks as shown in fig. 1. The model execution is controlled by a discrete event scheduler (in our example either MLDesigner [8] or MSArchitect [9]), responsible to initialize the vehicle fleet and to run the simulation. The vehicle fleet

¹An alternate mode of operations is available which offers the ability of manual booking.




Fig. 1. High-Level simulation model of the Toll Collect system.

treats each OBU as an individual with a distinct configuration and internal state. This state changes according to the simulation and the externally pre-calculated (statistically realistic) driving pattern of the OBUs containing their configuration (e.g. hardware and software versions), their power cycles and the toll charging instants.

Starting with the OBU and its driving pattern the model simulates the automatic communication between the OBU and the central systems either to transmit the tolls collected or to update the OBU state, map data or software. Due to the arbitrary power cycles of the OBU and various resource restrictions (e.g. limited bandwidth, high latency, intentionally limited number of parallel connections for the central systems) it is common for data transmissions to be interrupted and subsequently recovered by application-level protocols.

The mobile data network includes provider specific transmission properties (e.g. bandwidth and latency) and resource constraints (e.g. actively managed number of simultaneous connections allowed). On the network layer the simulation includes the bandwidths and latencies observed for the various OBU hardware platforms and mobile network operators. The simulation includes the GPRS connection handling, the authentication handshake and IP address handling but does not include the IP network layer.

The block “Central System” includes the typical systems required to authenticate, receive and validate data transmission (e.g. firewalls, proxy servers, load balancers, database and application servers) each with their individual resource constraints. From a service management perspective the system is a sizeable service value chain spread across several service providers [10–13].

To achieve realistic simulation results the model tries to include as many details as possible. Accordingly the vehicle fleet should include as many individual OBUs as in the real Toll Collect system (more than 700 000) with statistically realistic driving patterns for several consecutive months. In that way it is possible to simulate long-term behavior (e.g. a software update of the whole fleet) without resorting to scaling. The behavior of each OBU is implemented at a high-level of detail up to including the original source code of the OBU in the handling of internal state transitions.

B. Application and results

The simulation model is used to determine the effects of the systems’ configuration, e.g. on the progress of software updates or on the return to normal operations after outages of the central systems. This can be extended to determine the optimal system configuration e.g. considering the trade-off between operational costs and the cost of financing (of not yet processed toll fees) [5].




Fig. 2. Weekly driving pattern of a vehicle fleet of 140 000 HGVs. Each point represents the chargeable kilometers driven within a 30 minute period.

The execution of the Toll Collect model using MLDesigner takes about 6 hours of single-core CPU time (Intel Xeon X5670 at 2.93 GHz) to simulate 16 weeks with a fleet size of 140 000 HGVs (corresponding to a 1:5 scaling). In the example each HGV drives on average 18 493 km per year on German toll roads (compared to an average of 18 872 km [6]) at an average speed of 80 km/h (corresponding to the speed limit for HGVs). The pre-calculated driving pattern changes according to the day of week (see fig. 2) and is based on a statistical analysis of the driving pattern over a 15 week period (in early 2011). The weekend and Sunday truck ban on German highways is clearly visible in fig. 2.

Additional data from the Toll Collect test fleet (> 2 000 HGVs) is used to parametrize the number and duration of power cycles. The example uses an average of 1112 power cycles per OBU and year (with a minimum duration of one minute per power cycle).

This microscopic simulation model is also used to determine macroscopic effects, e.g. the periodic update of map and tariff data on the OBU. The update process is initiated by the OBU which periodically checks for the availability of updates and schedules the download of new updates randomly prior to the start of the validity period of the new data. Fig. 3 shows the result of two consecutive updates of the map data, where each OBU has one version of the data installed and possibly either knows about the existence of a new version or has it already downloaded (but not yet activated). With the start of the validity period of the new version OBUs that had it previously downloaded will immediately switch to the new version (provided the OBU is powered on). OBUs that are unable to download the new data in time (e.g. OBUs staying outside of the German mobile data network coverage) will try to retrieve the update as soon as the power restored to the OBU and the OBU is within reach of the German mobile data network. Across the whole vehicle fleet we observe that about 10% of all OBUs do not connect to the data center within a given 15 week period.

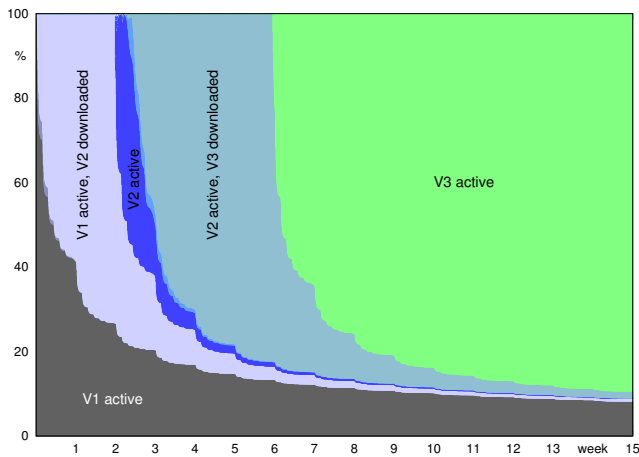


Fig. 3. Two consecutive map data updates of a fleet of 140 000 OBUs over a 15 week period, the update is downloaded and activated only after the start of the validity period.

C. Simulation performance

To achieve realistic simulation results we decided against the use of a simplified simulation model (as compared to the real-world system) and aim for a 1:1 scale, i.e. more than 700 000 individual OBUs within the simulation and a realistic behavior on the network layer. Therefore the typical time-scales within the simulation are on the order of 100 ms. However, the business processes of interest have a typical time-scale of one to two months: e.g. map and software updates are intentionally spread over many weeks to be able to reach HGVs that are operating outside of the German mobile network coverage.

As a consequence the simulation performance must allow to simulate at least three consecutive months of a realistic driving pattern with a full-size vehicle fleet. Using the simulation as part of the design process or to validate changes to the systems' configuration necessitates that a typical simulation run delivers results within the business day. Unfortunately, the tools used do not yet allow the automatic distribution of the simulation across several CPUs (or even CPU cores).

The initial implementation of the simulation model with MLDesigner led to various performance bottlenecks due to the large number of OBU objects and scheduled events within the simulation. The extraction of the OBU logic from the simulation model to conventional C++ classes alleviated the performance degradation and the memory usage. However, the simulation model is still effectively limited to less than a 50% fleet size while completely losing the ability to model (rather than to write C++ classes). As a consequence we evaluated different discrete event simulation tools with respect to their performance characteristics.

III. EVALUATION OF DISCRETE EVENT SIMULATION PERFORMANCE

A. Importance of performance

It is well known [14], [15] that software-intensive systems evolve towards ever increasing complexity – fulfilling more user requirements, interfacing with additional other systems and of course requiring ever more lines of code. Modeling and simulation methodologies and technologies [4] can be applied to design, analyze, evaluate, validate, and optimize such systems – far in advance of the actual implementation. Executable models are created as blueprints of the new system and are used as functional (“virtual”) prototype. At an early stage of the design process these virtual prototypes give insight into the systems' behavior e.g. regarding the scaling properties, the advantages and disadvantages of the system topology. At any time simulations can be used to explore operational scenarios (especially those exceeding the systems' specification) and the inherent risks (operational and procedural).

In this context simulation performance needs to keep up with the enormous complexity increase of executable models, which in turn follows the complexity increase of systems and processes. In addition, executable models should include a high level of detail. Together with systems including a large number of active components (e.g. users, machine-to-machine networks) this results in a complex simulation model – both from a static and dynamic perspective.

In a business context, the simulation is often part of an optimization process, i.e. the optimal solution is determined by iterative optimization loops. In that case many steps consisting of a complete simulation run (possibly including test batteries) need to be evaluated to determine the optimal solution. Of course this approach is beneficial only if the simulation results are both reliable and available well in advance of traditional software engineering approaches. Hence simulation performance in terms of speed and memory consumption and its benchmarking became a critical aspect in system design.

B. Performance benchmarking

There exist several approaches for benchmarking simulation performance, especially kernel benchmarks and application benchmarks are common [16], [17]. A kernel benchmark consists of several, typically smaller test cases where each test case stresses a single elementary function of the simulation kernel (see fig. 4). Therefore kernel benchmarks are useful to analyze the built-in performance of low-level mechanisms. The results are typically weighted according to their importance for a given application domain – however, the predictive power of kernel benchmarks for real-world application performance is limited. To compare application performance, the benchmark measurements include a number of real-world examples from the application domain. The selected applications should exhibit different characteristics and represent typical challenging workloads.

In the Toll Collect example we started with an existing simulation model using a given simulation tool (MLDesigner). To




Fig. 4. Basic test model for kernel-level benchmarks. Elementary functions are used to test e.g. the performance of the pseudo-random number generator or the future event list.

achieve the necessary performance (as outlined in section II-C) both kinds of benchmarks were used: A low-level analysis of the simulation kernel allows to identify performance bottlenecks in the existing simulation model and tool. Additionally, the kernel benchmark is easily adopted toward different simulation tools. Section IV-B gives a description of the kernel benchmarks used to benchmark a total of five different DES simulation tools, followed by a comparison and discussion of the kernel benchmark test results.

As a consequence of the kernel benchmark results the Toll Collect simulation model was ported to a second simulation tool – requiring considerable effort and expertise (both of the simulation model and tool). Having the same simulation model implemented for two different simulation tools allows for direct comparison and benchmark at the application-level (as shown in section V-A).

IV. KERNEL-LEVEL BENCHMARKS

DES simulations are typically split into the simulation environment, and the simulation model. The simulation environment itself is used to create models (using an interactive and graphical user interface), to execute existing simulation models and possibly also to visualize the progress and results of a simulation run. The simulation model itself contains all static model entities, their relationships and methods to handle events during the model execution. Thus the simulation model determines the dynamic properties of a simulation run, e.g. the number of entities present during the model execution and the number of events created.

A. Simulation kernel benchmark tests

Since the execution control resides with the simulation kernel, the implementation of event handling (especially the future event list (FEL) and its update mechanism), the data and memory handling (e.g. pass-by-reference vs. pass-by-value, garbage collection) and the use of caches determines the simulation kernel performance. Following [18] we include the following elementary factors in our kernel-level benchmarks:

- FEL management: The event scheduling mechanism is the core of any discrete event simulation determining the dynamic behavior of the simulation. At any given time the model entities create new events scheduled to take place in the future and sometimes cancel existing future events as well. The crucial performance factor of a DES simulation kernel is therefore the handling of the future event list. Its management can be more time consuming than the actual data manipulation.

- Memory and data type management: The allocation and maintenance of tokens and memory for dynamic model entities is an important issue. The event handling will inevitably deal with the creation and deletion of a large number of events, events passed between model entities usually need to transport additional (application-level) data between the entities, possibly necessitating the casting between data types (incurring an additional overhead). The efficient storage of the information will directly affect the simulation performance. A *pass-by-value* approach will incur additional overhead (due to the necessity of duplicating the data). A *pass-by-reference* implementation of the FEL management algorithms processing the tokens representing an events should yield better performance – especially if the simulation entities are only referenced from the event tokens. Dealing with memory allocations can be improved by the use of caching mechanisms.
- Pseudo-random number generator performance: A basic requirement of DES simulation execution is the ability to use random numbers to achieve a “non-deterministic” behavior. A typical DES tool includes generators for several different random number distributions. It is critical to be able to use large streams of pseudo random numbers.
- Arithmetic operations: The actual data manipulation is given by arithmetic operations either in an imperative or functional language. This programming language needs to be executed at runtime and can become a performance bottleneck if the chosen programming language does not allow compilation to the underlying CPU architecture.

Additionally, the ability to generate reports or to export reporting data is a basic requirement for any DES simulation. Creating the reports and the underlying data can incur considerable additional computational expense. However, the reporting requirements are typically driven by the application domain. Therefore we do not include reporting in our kernel-level benchmark.

B. Simulation kernel performance tests

We present five different test models for DES simulation kernel benchmarks, addressing all elementary factors presented in section IV-A. These models are applied later to investigate and compare DES kernel performance. The models have been kept simple in order to assure universality regarding different kernels/tools and to avoid possible side effects.

1) *Simulation scaling*: A simulation model with several hierarchy levels (fig. 5) is simulated in a sequence increasing the total number of events, while the size of the future event list remains fixed. This test determines whether the FEL performance is affected by the FEL size. The test used a clock interval of one and the number of events processed increased from $0.1 \cdot 10^6$ to $10\,000 \cdot 10^6$ events.

2) *FEL size scaling*: The second test uses the generic simulation model of fig. 4 with a delay-function. The delay is used to easily configure the (average) number of events waiting in the future events list with minimal variance, while the total number of events processed remains constant. This test




Fig. 5. Test model for simulation scaling




Fig. 6. Test model for FEL adaption with additional single events (SE)




Fig. 7. Test model for memory and data type management for sequential (top) and parallel (bottom) processing

examines the overall performance of the FEL algorithm and data management. We used a uniform distribution of events in the FEL list. Of course, the test can be extended toward non-uniform events distributions, in order to check adaptability of the FEL algorithm on different events densities.

For this test we use a clock interval of one, a fixed number of processed events ($300 \cdot 10^6$) and configure the delay-function to produce a given size of the future events list (between 10^6 to 10^7 events, constant over single experiment).

3) *FEL adaption*: This test extends the future events list size during one test run by changing the parameter of the delay-function dynamically during the simulation execution. The test extends the previous test model by additional single events used to change the parameter of the delay-function (see fig. 6). As a consequence the size of the FEL changes during the test run forcing the simulation kernel to adapt the FEL size (e.g. allocating and deallocating memory) during the simulation run.

The test uses a clock interval of one and a dynamic delay-function parametrized to give a dynamic FEL size of $1000 - 10^6 - 100 - 10^7 - 10$ events during the simulation run. In total one test run consists of $200 \cdot 10^6$ processed events.

4) *Memory and data type management*: The test creates large data arrays of different sizes and passes the data through the simulation model in sequential or parallel order as depicted in figure 7. When executing the model the memory management of the simulation kernel should recognize the passing of unmodified data and use references to this data. Ideally only one datum should be created and send as reference through the model. As long as the delays are set to zero, no difference between serial and parallel passing should be recognizable. The test uses a fixed number of nodes (delay blocks) either in a parallel or serial configuration and data arrays with $1, 0.5 \dots 2 \cdot 10^6$ entries.

5) *Random-number generator performance*: A large number of random values is generated using different distributions. The model uses a constant function as a reference to measure

relative performance of the built-in pseudo-random-number generators. The test computes $20 \cdot 10^6$ random numbers of different random number distributions (normal distribution, Poisson distribution and exponential distribution)

C. Evaluation of simulation kernel

Using the simulation kernel performance test models of the previous section we evaluate several different simulation kernels and try to determine the most efficient one. Currently there are more than 80 simulation tools listed for DES [19]. We selected five, namely PtolemyII [20], OMNeT++ [21], AnyLogic [22], MLDesigner [8] and MSArchitect [9]. Criteria for selecting these tools were their degree of proliferation, their availability, and the flexibility of model acquisition. All of the selected simulation tools provide an actor oriented visual modeling and simulation environment. Additionally they are capable to include external libraries and provide runtime control. Running the test models as outlined in the previous section we collect the following results:

1) *Simulation scaling*: Fig. 8 gives the results of the simulation scaling test. The upper chart shows the event processing performance for different simulation lengths and the bottom chart the private memory consumed during simulation. The tests show that neither the memory consumption nor the event processing performance is affected by increasing the simulation runtime. Looking at the sensitivity of running the tests with additional hierarchy levels we find that only OMNeT++ is sensitive to the additional hierarchy levels.

However, from the test results it is already obvious that the different tools vary in event processing performance by an order of magnitude: MSArchitect provides the highest speed. MLDesigner, AnyLogic, and OMNeT++ provide half the speed. Ptolemy II is ten times slower. Looking at the memory usage during the simulation the difference between

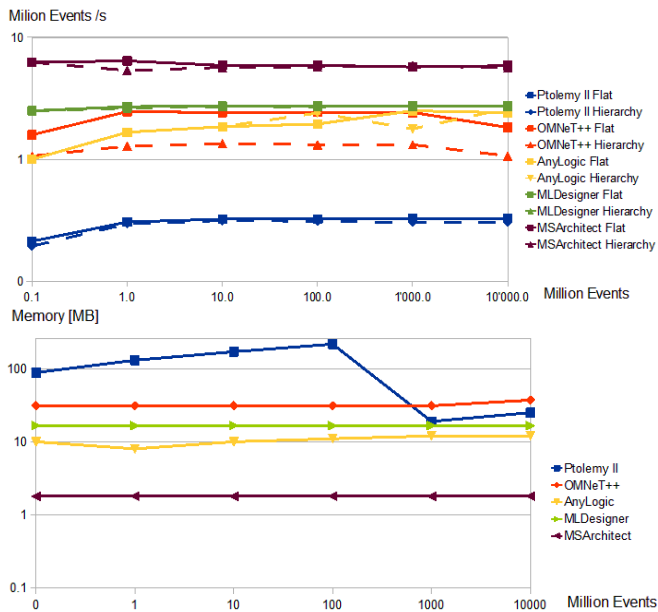


Fig. 8. Runtime performance (top) and memory usage (bottom) scaling of different DES tools

the tools is again more than an order of magnitude – the slowest tool using the most memory and the fastest tool using the least.

Two of the tools (Ptolemy II and AnyLogic) are based on the Java programming language, where explicit memory deallocation is not possible. Apparently the Ptolemy II test run triggers the JVM garbage collection during the simulation run and is able to free 90% of its memory. As a result Ptolemy II memory consumption is then comparable to the next three simulation tools. The AnyLogic test run starts already with a much lower memory consumption than Ptolemy II and no effect of JVM garbage collection is visible.

2) *FEL size scaling*: Fig. 9 gives the results of the FEL size scaling test. As expected, a systematic performance reduction can be observed with increasing FEL size, due to the increasing overhead for FEL management. Most of the tools tested initially start with “constant” performance (on a log-log scale). With increasing FEL size three of the five tools develop drastic performance degradation. This coincides with a rapid growth of memory consumption with increasing FEL size. We propose that the performance reduction is correlated with increased FEL memory usage due to a performance penalty of calendar queue based schedulers for large queue sizes.

Again, Ptolemy II has the lowest performance in this test. OMNeT++ is nearly not affected in the considered FEL size interval. In absolute numbers, MSArchitect has the best test performance and the lowest memory usage.

3) *FEL adaption*: In this test the simulation kernel is subjected to a varying demand to its future events list. Beyond the runtime needed for the test the main result is the memory consumption during the test run as given in fig. 10. The simulation took 3 122 s with MLDesigner, 673 s with AnyLogic, 425

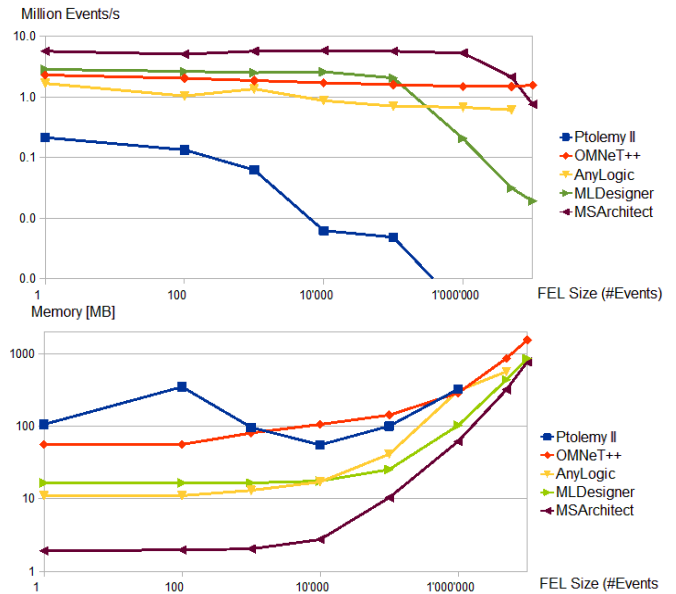


Fig. 9. Future event list size scaling test results for runtime performance (top) and memory usage (bottom) scaling of different DES tools




Fig. 10. Test results for the future events list adaption to varying demand

s with MSArchitect, 376 s with OMNeT++ and over 2 hours with Ptolemy II. During that time the dynamically changing memory usage varies widely between the different tools. Most tools tend to allocate memory in chunks visible as steps in fig. 10.

Again, Ptolemy II is the slowest tool in the comparison and also requires more memory than any other tool in the benchmark. The memory usage of OMNeT++ indicates the ability to dynamically free already allocated memory. However, this simulation tool also allocates considerably more memory than any of the other tools for a brief period of time during the test run.

4) *Memory and data type management*: This test passes large arrays of data through the simulation running either in a parallel or serial configuration. The memory consumption during the test run is shown in fig. 11. Most tools handle serial and parallel passing of token data in a different way, which can be recognized by the gap in memory consumption between both serial and parallel versions. Ptolemy II and MSArchitect do not show a difference between the parallel and




Fig. 11. FEL adaption test results

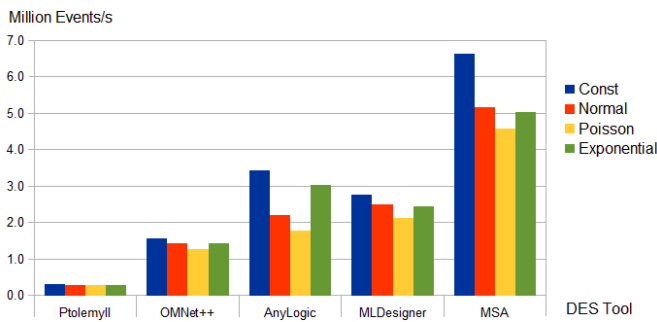


Fig. 12. Performance of the pseudo-random number generators in different DES tools

serial version, only references are passed when only delays are used. However, Ptolemy II requires more memory and shows a different behavior according to the memory allocation: a large portion of the memory is allocated at initialization time with standard modeling elements.

5) *Random-number generator performance*: The last test is the PRNG performance test. As depicted in fig. 12 the performance does not depend on the type of the generated distribution, since there are only minor differences to the generation of constant numbers. Again, Ptolemy II is an order of magnitude slower than OMNeT++, AnyLogic, and MLDesigner in this test. MSArchitect gives the highest performance in this comparison.

It is not known whether the PRNG algorithm differs between the five simulation tools or if the PRNG performance is adversely affected by event management overhead. Since this test relies on the correct implementation of the PRNG i.e. we do not check the statistical quality of the random numbers generated, the test results might not be fair if one of the tools were to use low-quality but high-speed generators. The fastest tool (MSArchitect) is known to use the PRNG of the standard C++ libraries.

It can be concluded, that Ptolemy II is inferior in all simulation kernel benchmarks performed. MLDesigner is equal to or better than AnyLogic in all categories but FEL adoption. Due to the utilization of the JVM, AnyLogic requires more memory in equivalent models and therefore scales worse with increasing FEL size. Both, MSArchitect as well as OMNeT++ show the best performance in some categories. MSArchitect

is the fastest simulation kernel in most categories and requires least memory for data handling. OMNeT++ is better at adapting to different and dynamic FEL size requirements.

V. ENHANCING SIMULATION PERFORMANCE OF THE TOLL SYSTEM MODEL

While section IV focused on the performance evaluation of DES kernels we now focus on the application-level performance, i.e. how to specify *efficient* executable DES models.

A. Evaluation of model architecture

A simulation model can be thought of as a (simplified) copy of an existing or imaginary system, created for a certain purpose. The model and the process of creating the model are a key to learning and communicating about the system itself. This implies that the right level of abstraction needs to be found so as to include only the system behavior relevant to the models' purpose. Bearing this in mind the most important rule in designing efficient models can be derived: *The level of detail always follows the model purpose*. For instance it makes a huge difference for choosing the appropriate level of detail when designing a data transmission model compared to modeling a rather abstract business process. Of course, any useful model must be connected in some way to the reality. A second point directly connected to that rule is to focus on measurable system behavior. Otherwise the model could become worthless when using it for analyzes and optimization.

Designing the simulation model directly affects the runtime properties (e.g. performance and memory consumption). From a technical point of view the performance of a model can be improved considerably by addressing several issues:

- the number of simulation entities present during the simulation run,
- the data transport between model components to reduce the number of DES events and
- the execution time and memory consumption of the model components.

Hence whenever possible, highly interacting model components should be merged together to avoid the time consuming data exchange via the simulation kernel. In addition data should always be transmitted in form of references/pointers (*pass-by-reference*). References are values that enable indirect access to a particular datum, indistinct from the data itself. They are used to efficiently pass large or mutable data. In that way the time-consuming and unnecessary copying of data is avoided.

The simulation kernel benchmarks in the previous section identified the future events list as a key factor in the kernel performance. The tests were designed to continuously create new events leading to different FEL sizes. Creating, scheduling and passing events is certainly the key feature of DES simulation kernels. However, a simulation model sometimes needs to be able to cancel scheduled future events before they are executed. Many simulation tools lack a good implementation for canceling events from the FEL, possibly needing to traverse the whole FEL in the search of the canceled event and possibly

fleet size	Runtime [s]	
	MLDesigner	MSArchitect
70 000	0.25 M	900
700 000	49.00 M	8 700

Fig. 13. Simulation performance for the Toll Collect example with a simulated time period of three months

triggering memory reorder after removing the canceled event from the FEL. Obviously, simulation runs with a large FEL are more affected.

The performance of simulation models can be improved by transferring part of the simulation model to existing standardized model components or even extending the simulation tools' existing catalog of standard components.

B. Performance enhancement to our solution

As the project of modeling the German toll system was launched our team had no clear picture of the coming performance issues: Existing simulation models of HGV tolling systems both at Toll Collect and in the literature were limited to a few thousand simulated HGVs [23, 24] and reaching 500 000 HGVs over a 4 week period [25]. Our model aimed to include a more detailed behavior and a vehicle fleet almost two orders of magnitude larger (comparable but still larger than simulations of metropolitan car traffic, e.g. [26] using 200 000 drivers with a shorter simulated time frame).

After putting together and validating the basic DES model in MLDesigner, including the dynamic behavior of the vehicle fleet, mobile providers and central system we tried to scale up to the real world situation. This meant to run simulations of vehicle sizes of up to 700 000 HGVs over a simulated time period of at least 3 months. The disappointing simulation performance results are shown in the second column of table 13. The desired scenario took about 49 million seconds, over 6 times *slower than reality*, an unacceptable result.

By transferring the model from the system design tool MLDesigner to MSArchitect the simulation performance could be increased dramatically. On the one hand the throughput of events is about 3.5 to 4 times higher in MSArchitect (as confirmed by the kernel benchmarks in section IV). On the other hand we recognized huge performance issues in the management of complex data structures in MLDesigner. We analyzed the differences by comparing the simulation performance of using MLDesigner data structures versus using pointers to external C++ classes for data transport (the default in MSArchitect). In total the transfer of our simulation model from MLDesigner to MSArchitect brought a 120-fold speed increase.

On top we redesigned our model architecture. First of all removing all "cancel event" operations from the model – being rather expensive operations in both simulation tools. The canceling of events was replaced by introducing an additional boolean tag to store whether the next receiving event is ignored or not. By doing so the overall amount of events in the FEL is increased and more memory is needed but time consuming cancel operations can be avoided.

Next we removed several retry processes between vehicle fleet and mobile data network providers and merged heavily interacting model components to minimize data transport across the simulation kernel. In addition we switched to the data structure mechanism of MSArchitect which automatically uses references when sending or receiving unchanged data tokens.

In total a further significant performance increase could be achieved. The rightmost column of table 13 shows the results of two different scenarios executed with MSArchitect. Simulating the scenario stated above (700 000 HGVs over a three months period) took about 8 700 seconds. Thus the simulation speed could be increased by a factor of 5 630 compared to the initial runs using MLDesigner. It is noteworthy that the model used with MSArchitect also includes additional functionality of the German toll system.

VI. SUMMARY AND OUTLOOK

The paper showed how to analyze and evaluate the performance behavior of simulation tools in general and specifically the German toll system using a simulation-driven design approach. For this purpose an executable model of the toll system has been created and validated using historical data. Afterwards different possibilities to improve simulation performance were evaluated and applied. In practical terms, five simulation tools and several design rules in creating efficient model architectures have been investigated. As a result an increase in simulation performance by a factor of 5 630 could be achieved.

The improved performance facilitates new areas of application for the toll system model. On the one hand the model is now used as part of the system engineering process to evaluate design decisions before putting them into reality (test batteries). The operational risk inherent to configuration updates of the toll system can be investigated and therefore decreased enormously. Additionally the performance achieved allows us to iterate over large amounts of functional, architectural and environmental (operational scenarios) alternatives for optimization purposes.

In the future we plan to further improve simulation performance by using parallel DES (PDES) in MSArchitect. In this context the model architecture needs to be adapted to facilitate the automatic distribution of the model execution across multiple CPUs.

REFERENCES

- [1] J. Banks and B. Nelson, *Discrete-Event System Simulation*. Prentice Hall, 2010.
- [2] T. Baumann, *Automatisierung der frühen Entwurfsphasen verteilter Systeme*. Saarbrücken, Germany: Südwestdeutscher Verlag für Hochschulschriften, 2009.
- [3] E. Mosekilde, *Topics in Nonlinear Dynamics: Applications to Physics, Biology and Economic Systems*. Singapore: World Scientific Pub Co Inc, 1996.
- [4] T. Baumann, "Simulation-driven design of distributed systems," *SAE Technical Paper*, no. 2011-01-0458, 2011.
- [5] B. Pfitzinger, T. Baumann, and T. Jestädt, "Analysis and evaluation of the german toll system using a holistic executable specification," *Hawaii International Conference on System Sciences*, vol. 0, pp. 5632–5638, 2012.

- [6] Bundesamt für Güterverkehr, "Maut-Jahresstatistik 2011/2010," 2012, [accessed 10-March-2012]. [Online]. Available: http://www.bag.bund.de/SharedDocs/Downloads/DE/Statistik/Lkw-Maut/Jahrestab_11_10.pdf?__blob=publicationFile
- [7] Toll Collect GmbH, "Truck toll system proven effective," 2012, [accessed 11-March-2012]. [Online]. Available: <http://www.toll-collect.de/en/company/news/truck-toll-system-proven-effective.html>
- [8] "MLDesigner," 2012, [accessed 10-April-2012]. [Online]. Available: <http://www.mldesigner.com/>
- [9] "MSArchitect," 2012, [accessed 10-April-2012]. [Online]. Available: <http://www.andato.com/>
- [10] B. Pfitzinger, H. Bley, and T. Jestädt, "Service catalogue and service sourcing," in *Informatik 2010, Business Process and Service Science—Proceedings of ISSS and BPSC*, W. Abramowicz, R. Alt, K.-P. Fähnrich, B. Franczyk, and L. A. Maciaszek, Eds., vol. 177, 2010, pp. 55 – 62.
- [11] B. Pfitzinger, T. Gründer, and T. Jestädt, "Sourcing decisions and IT service management," in *Practitioner Track International Symposium on Services Science (ISSS'09)*, R. Alt, K.-P. Fähnrich, and B. Franczyk, Eds., vol. 16, 2009, pp. 71–80.
- [12] F. Opitz, B. Pfitzinger, and T. Jestädt, "Service levels of a cost center organization," in *Practitioner Track International Symposium on Services Science (ISSS'09)*, R. Alt, K.-P. Fähnrich, and B. Franczyk, Eds., vol. 16, 2009, pp. 81 – 86.
- [13] B. Pfitzinger, T. Jestädt, W. Helmers, and S. Kosterski, "Best practices im contract lifecycle," in *Best Practices im Outsourcing*, M. Auerbach, C. Oecking, R. Jahnke, F. Strecker, and M. Weber, Eds. Bitkom, 2010, pp. 185 – 200.
- [14] M. Lehman, "Programs, life cycles, and laws of software evolution," *Proceedings of the IEEE*, vol. 68, no. 9, pp. 1060 – 1076, sept. 1980.
- [15] M. M. Lehman, "The role and impact of assumptions in software development, maintenance and evolution," *IEEE International Workshop on Software Evolvability*, pp. 3–14, 2005.
- [16] T. Teweldeberhan, A. Verbraeck, E. Valentin, and G. Bardounet, "An evaluation and selection methodology for discrete-event simulation software," in *Simulation Conference, 2002. Proceedings of the Winter*, vol. 1, dec. 2002, pp. 67 – 75.
- [17] R. Ronngren, L. Barriga, and R. Ayani, "An incremental benchmark suite for performance tuning of parallel discrete event simulation," in *Proceedings of the Twenty-Ninth Hawaii International Conference on System Sciences*, vol. 1, jan 1996, pp. 373 –382.
- [18] G. S. Fishman, *Discrete-Event Simulation: Modeling, Programming and Analysis*. Berlin: Springer, 2001.
- [19] M. C. Albrecht, "Introduction to discrete event simulation," 2012, [10-April-2012]. [Online]. Available: <http://www.albrechts.com/mike/DES/Introduction%20to%20DES.pdf>
- [20] "Ptolemy Project," 2012, [accessed 10-April-2012]. [Online]. Available: <http://ptolemy.eecs.berkeley.edu/>
- [21] "OMNet++," 2012, [accessed 10-April-2012]. [Online]. Available: <http://www.omnetpp.org/>
- [22] "AnyLogic," 2012, [accessed 10-April-2012]. [Online]. Available: <http://www.xjtek.com/anylogic/>
- [23] M. Hericko, M. Hericko, and A. Zivkovic, "An evaluation of different functional solutions for satellite-based tolling in europe," *Hawaii International Conference on System Sciences*, pp. 1–10, 2011.
- [24] K. Lunde and L. Kieble, "Simulating communication within a satellite-based automated toll collection system," *Proceedings of the 55th International Scientific Colloquium*, pp. 318 – 323, 2010.
- [25] K. Lunde, L. Kieble, and M.-A. Funk, "Prediction strategies in a service level granting prefetching cache for version-controlled gis data," *ISAST Transactions on Computers and Intelligent Systems*, vol. 2, no. 2, pp. 46–51, 2010.
- [26] G. Flötteröd, "Traffic state estimation with multi-agent simulations," Ph.D. dissertation, TU Berlin, 2008.