

A Modular Approach for Parallel Simulation of Virtual Testbed Applications

Arthur Wahl, Ralf Waspe, Michael Schluse and Juergen Rossmann

Institute for Man-Machine Interaction, RWTH Aachen University, Ahornstrasse 55, Aachen, Germany

Keywords: Parallel Simulation, Virtual Testbeds, Synchronization, Scheduling.

Abstract: Virtual Testbeds are advanced 3D simulation environments that model all relevant aspects of complex technical systems, to enable their systematic evaluation by engineers from various disciplines. Due to the high complexity of the resulting simulation, real-time capabilities are very hard to achieve without applying multi-threading strategies. Therefore, we present a novel, simulation architecture that facilitates a modular approach to perform parallel simulations of arbitrary environments without further effort. Specifically, no explicit knowledge of the underlying simulation algorithms or model partitioning is needed. As a result, engineers can simply distribute simulation components such as rigid-body dynamics, kinematics, renderer, controllers etc. among different threads without being concerned about the specific technical realization. We achieve this by managing (partial) copies of the state data underlying the simulation models. Each copy acts as a self-contained, independent entity and is bound to one thread.

1 INTRODUCTION

Virtual Testbeds (VT) greatly enhance the scope of the classical simulation approach. While typical simulation applications examine only specific aspects of an application, a VT enables development engineers to examine the entire system in its environment and provides a holistic view of the dynamic overall system including internal inter-dependencies in a user-defined granularity. In contrast to the classical bottom-up-strategy this can be seen as a top-down-approach. For classical fields of application of robotics, e.g. in production plants with a well-defined environment, a bottom-up-strategy in the development of simulation models is the approved method, because it allows very detailed insights into the analyzed subsystems. On the other hand, unpredictable effects of the interaction of multiple subsystems may easily be overseen. In particular, when it comes to building larger applications like for example a space robot mission as shown in Figure 1, planetary landing or exploration, autonomous working machines, advanced vehicle assistant systems, forestry application etc. as presented in the Virtual Robotics Testbed (Rossmann and Schluse, 2011), they are hard to describe in an analytical way in order to integrate them into an analytical simulation model and the demand for highly realistic and collaborative simulation environments increases.

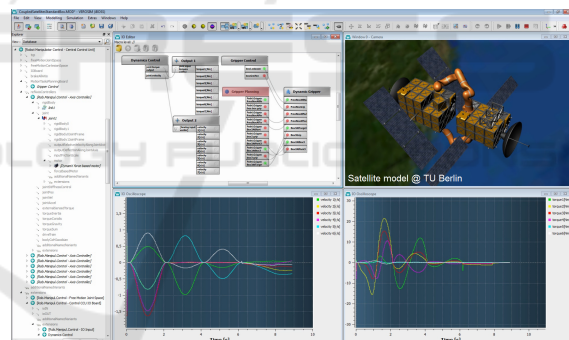


Figure 1: Virtual Robotic Testbed for space environments.

VTs pose new challenges in the field of parallel simulation compared to the classical simulation approach. They demand the development of a parallel simulation architecture that encompasses arbitrary application fields of VTs. Therefore, the concurrent execution of multiple simulation algorithms with different degrees of granularity must be supported. This will enable the simulation to incorporate various aspects of an application with different degrees of detail at the same time. Furthermore, the partitioning scheme provided to distribute the simulation among different threads must enable an engineer to generate partitions that are scalable and adaptable during the development or creation phase of a VT application e.g. an engineer can add new functionality or components to an existing simulation model to inves-

tigate additional aspects of an application. A partition must provide the possibility to include these new extensions. As a consequence, the integration of new simulation algorithms must also be supported. Regarding the amount of different simulation algorithms provided by a VT, an explicit parallelization of each simulation algorithm would require specific knowledge about the algorithms and their internal dependencies, provide no scalability and therefore present an inappropriate approach.

We present a novel parallel simulation architecture for arbitrary VT applications that fulfills the aforementioned requirements. The architecture generates low synchronization overhead during parallel execution and facilitates a modular distribution of a simulation among different threads without further effort. We follow a top-down-approach like the VT approach itself and propose the application of a functional partitioning scheme that distributes simulation algorithms among different threads. As shown in Figure 2, simulation algorithms or components of a VT such as rigid-body dynamics, kinematics, renderer, controllers etc. can simply be assigned by an engineer to a specific thread and executed concurrently. Hereby, no explicit knowledge of the underlying algorithms is required or model partitioning is needed to perform parallel executions of a simulation. The modular distribution provides scalability since parallel execution scales with the number of components provided by a simulation model and partitions can be extended during the development phase by newly included components.

Execution of the simulation components is handled by a scheduler. Each thread owns a scheduler, where components can register and be executed correspondingly to their execution rate. The execution rate can be chosen individually for each com-

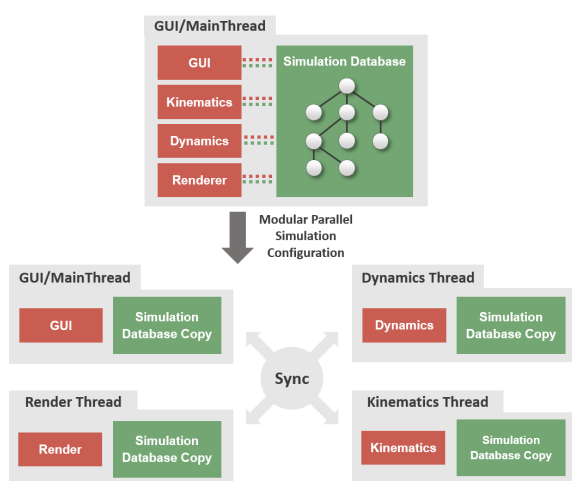


Figure 2: Example distribution of simulation components.

ponent, similar to the approach presented in (Friedmann et al., 2008) for the simulation of autonomous mobile-robots teams. Components can be executed safely in parallel among different threads without interfering with each other. We achieve this by managing (partial) copies of the state data underlying the simulation models. The state data is organized in an active object oriented graph database and represents the current state of the simulation model at a given time-step. Each copy of the database acts as a self-contained, independent entity and is bound to one thread. All components assigned to one thread, work on the same copy of the database and are executed sequentially. Components cannot access or manipulate (no read/write access) the database of different threads and therefore can be executed safely in parallel without interfering with each other. Interaction between components among different threads is handled during a synchronization step. Our approach introduces a conservative synchronization algorithm with variable lookahead. Each scheduler can operate under a different lookahead (synchronization interval), depending on the execution rate of its slowest component. A change detection mechanism keeps track of changes that have occurred during a time-step. These changes represent the progress of each component and are used to update all databases to attain a consistent global simulation state.

2 RELATED WORK

In the context of simulation systems, replication has been exploited in order to evaluate simulation outcomes or to find optimal parameter settings to improve the accuracy of simulation results (Glasserman et al., 1996; Mota et al., 2000; Forlin et al., 2010). The approach is to run multiple independent copies of the same simulation program in parallel with different input parameters. At the end of the runs the results are averaged. This approach has been named Multiple Replication in Parallel (MRIP) (Pawlikowski et al., 1994). Furthermore, MRIP has been applied in stochastic simulations. Stochastic simulations require the execution of multiple replicated simulation runs in order to build confidence intervals for their results (Passerat-Palmbach et al., 2015). MRIP has been applied to accelerate the building process (Eickhoff, 2007; Ballarini et al., 2009; Ewald et al., 2009). This type of concurrent replication is not aimed at increasing the execution speed of each single run, but is aimed at efficiently providing a set of output samples by the concurrent execution of multiple independent differently parametrized sequential simu-

lation runs. Bononi et. al. introduced an approach that combines the concurrent execution of replicas with a parallel discrete event simulation by merging the execution tasks of more than one parallel simulation replica (Bononi et al., 2005). Hereby, blocking (idle CPU time) can be avoided during the synchronization barrier phase of a conservative parallel simulation (processes wait on the completion of more computational intensive processes) by switching to the execution of other replicas which already completed their synchronization phase. Parallel simulation cloning, another approach aimed at introducing replication in the context of parallel simulation was first employed by Hybinette and Fujimoto as a concurrent evaluation mechanism for alternate simulation scenarios (Hybinette and Fujimoto, 1997). The aim of this approach is to allow fast exploration of multiple execution paths, due to sharing of portions of the computation on different paths. Parts of the simulation (logical processes) can be cloned during the simulation at predefined decision points. From then on, the original process and the clones execute along different execution paths in parallel, hence execution paths before cloning are shared.

Opposed to replication and cloning, our approach uses multiple independent copies of the state data (organized in an active object oriented graph database) underlying the simulation model that all participate in the parallel execution of one simulation run. Through the utilization of state data copies, simulation components can be executed safely in parallel among different threads without interfering with each other. OpenSG (an open source project for portable scenegraph systems) uses a similar multi buffer approach, where partial copies of the scenegraph structure are used to introduce multi-threading (Voß et al., 2002; Roth et al., 2004). The highly modular nature of such an architecture allows for the accommodation of multiple projects with varying requirements and presents the ideal basis for the development of arbitrary parallel VT applications.

3 ACTIVE REAL-TIME SIMULATION DATABASE

A key component for our Virtual Testbed is the versatile simulation database (VSD), which is an active object-oriented, self-reflecting graph database, as introduced in (Roßmann et al., 2013). The VSD represents the core of our simulation architecture. The function of the core is to store the structure and state of the simulation upon which integrated simulation components will act on.

Besides information about the logical and spatial arrangement of a 3D scenes, as given by a scenegraph, our database approach also incorporates functionality in form of integrated simulation components, flexibly adopts to multiple data schemes, is independent from the type of simulation (continuous or discrete) and provides all the data needed for simulating VT applications. In addition, each database copy offers an efficient change detection mechanism that collects all the changes that have happened to a database during one simulation time-step. Synchronization (restoring a global consistent simulation state among all copies), thereby is reduced to the distribution of such changes among all database copies.

4 MODULAR PARALLEL SIMULATION OF VIRTUAL TESTBEDS

We introduce a parallel simulation approach that meets the challenges presented by VT architectures. The approach must be applicable to arbitrary application fields of VT, facilitate the concurrent execution of multiple simulation algorithms with different degrees of granularity and allow for scalability as well as the integration of new algorithms. VTs are composed of a variety of different integrated simulation components such as rigid-body dynamics, kinematics, sensors, actuators, renderer etc. which can be regarded as logical, independent processes. We propose a partitioning scheme that utilizes the modular VT structure and partition applications along their simulation components. We follow the top-down-approach and distribute simulation components among simulation threads for parallel execution.

4.1 Simulation Thread

A simulation thread owns a scheduler, a time controller, a (partial) copy of the main VSD and a unique context id e.g. "simulation thread 1-n". Components can be assigned by a development engineer to a simulation thread by choosing a thread specific context id and hereby activating their execution through the corresponding scheduler. A scheduler executes only components with a matching context id. The execution rate can be chosen individually for each component. Components assigned to one simulation thread are executed sequentially. The execution order is given by the execution rate of the individual components. In consequence, the execution order of a set generated by sequential simulation run is also main-

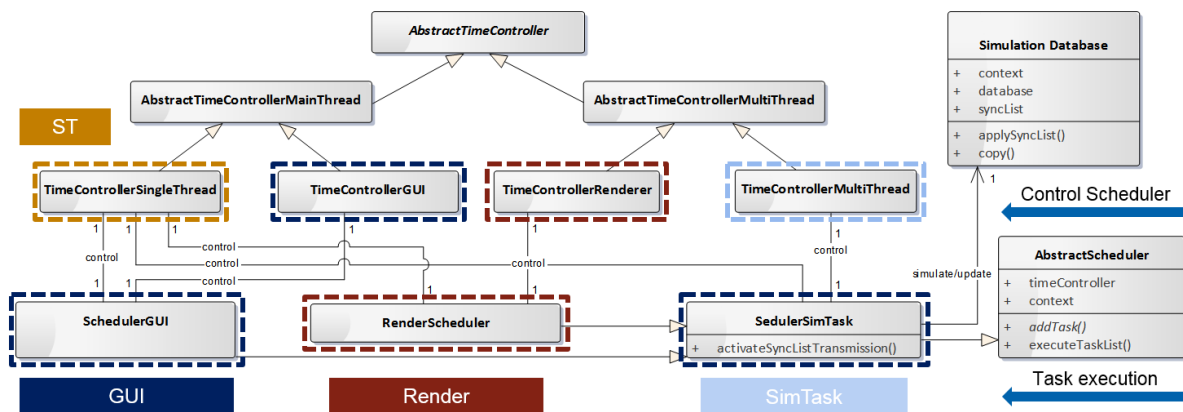


Figure 3: Simulation architecture.

tained by a simulation thread during parallel execution. A distinction must be made between sequences of components which have to be executed in a predefined order and sequences which are not dependent from the execution order. Sequences which have to be executed in a predefined order can only be assigned to one thread, to preserve that order. A time controller of a simulation thread provides multi-thread safe functionality to start, stop and pause a scheduler. The controller keeps track of how the simulation-time progresses compared to the wall-time. The execution of the next scheduling interval can be paused by the time controller, if the simulation time progresses faster than the wall-time. In return, if simulation time progresses slower, the time controller provides the possibility to skip rendering as well as execution cycles of components to catch up with the wall-time. Our partitioning scheme facilitates the possibility to spawn an individual simulation thread for each sequence, thus providing the scalability to exploit the maximum number of cores on a CPU.

4.2 Simulation Architecture

We propose the following architecture, see Figure 3. Four different types of controllers can be assigned to a simulation thread: a time controller for multi-thread execution, a time controller for rendering, a time controller for the GUI update of the simulation system and a time controller for sequential execution. Each time controller controls the execution of a corresponding scheduler. Both time controller types for parallel and sequential execution are derived from abstract classes that provide functionality to start, stop and pause a scheduler as well as to monitor how simulation-time progresses compared to wall-time. A simulation task scheduler can either be controlled by a corresponding time controller during parallel or sequential execution. Tasks with rendering content are

executed by a render scheduler. A render scheduler determines the render update rate and is responsible for updating all transformation matrices of the applications geometry before rendering the scene. The GUI update of the simulation system is handled by the GUI scheduler. Responsiveness of the application can be achieved by solely handling GUI event updates in the main thread. All scheduler classes are derived from an abstract base class. The abstract base class provides functionality to manage the execution of simulation tasks, handle synchronization and holds a reference to the corresponding time controller. The architecture presented here is capable of concurrently executing distributed render and non-render simulation components among multiple simulation threads.

4.3 Configuration of Simulation Threads

The creation of multiple independent simulation databases, that each facilitates all the data and functionality needed for simulation, represents the core idea of our parallel simulation approach. Each simulation thread has the ability to execute an independent, autonomous part of the whole simulation. All components assigned to one thread, work on the same copy of the VSD. This ensures that access to the VSD of a simulation thread only happens successively by components that share the same context. Typical problems like simultaneous data access, race conditions etc. that come along with multi-threading are being avoided because components from different threads do not share the same VSD but hold their own copy.

The approach presented here, provides the ability to create complex configurations of simulation threads. An example configuration for parallel simulations of on-orbit satellite servicing applications, as presented in the Virtual Space Robotics Testbed (Rossmann et al., 2016) is shown in Figure 4.

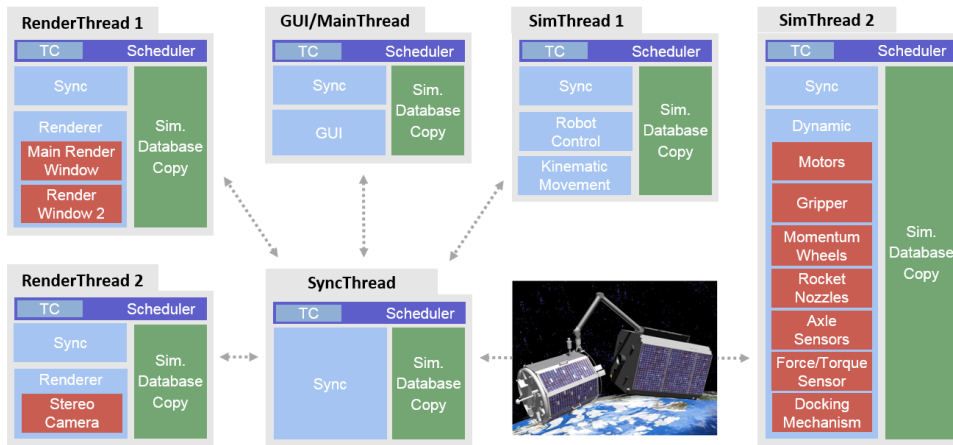


Figure 4: Modular parallel simulation of an on-orbit satellite servicing scenario.

5 SYNCHRONIZATION

To achieve correctness of parallel executed simulation runs, synchronization across the processors is required. The goal is to ensure that a parallel execution of a simulation produces the same result as an equivalent sequential execution. Therefore, dependencies have to be preserved among partitions and events as well as intermediate results have to be processed in the right order during the computation of the simulation state across the processors to attain a global, consistent simulation state. Assuming that the simulation consists of a collection of logical processes that resulted from a partition scheme and that the logical processes communicate by exchanging timestamped messages or events, this can be achieved during synchronization by preserving the local causality constraint (Fujimoto, 2001). The local causality constraint states that in order to achieve identical simulation results, each logical process or partition must process events in non-decreasing time-stamp order. As a consequence results are also repeatable.

Our approach introduces a conservative synchronization architecture with variable lookahead (synchronization interval). Each scheduler can operate under a different lookahead, depending on the execution of its slowest component. A change detection mechanism keeps track of changes that have occurred during a time-step. These changes represent the progress of each component and are used to update all VSDs.

5.1 Synchronization Architecture

The change detection mechanism of the VSD stores changes inside a list (synclist). The synclist represents the difference between the actual and the last

simulation state. It can be used to update other VSDs. Synchronization is realized through the exchange and application of synclists. A VSD can connect to another VSD, it is interested in receiving updates from, by enabling a synchronization connection. Therefore, we introduce the following communication mechanism, see Figure 5.

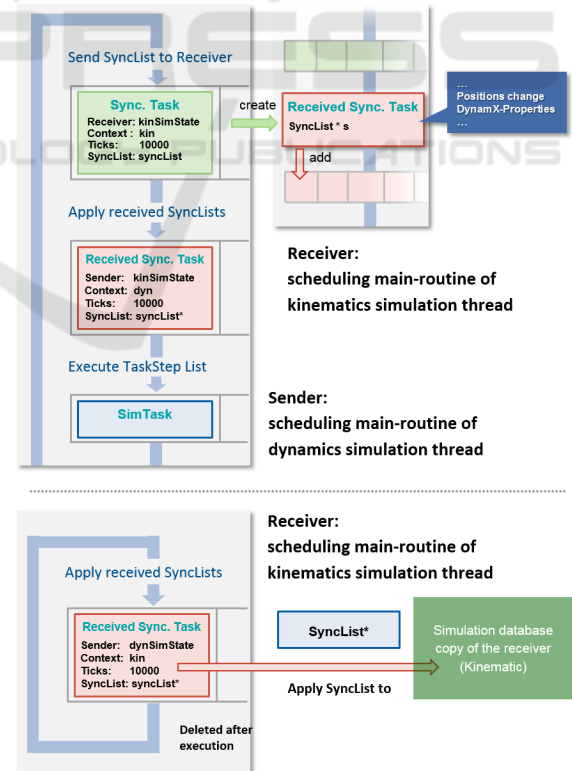


Figure 5: Synchronization architecture.

A synchronization connection is a sender/receiver-relation between two VSDs. Communication is done

via synchronization tasks that function as containers for the synclists to be sent. These tasks are transmitted and executed during the synchronization phase of the scheduling main-routine, before the next simulation time-step starts. After establishing the synchronization connections, each sender scheduler possesses a list of synchronization tasks, containing all receiver references and the necessary information to transmit synclists to them. The list is processed during each scheduling-cycle of the sender scheduler. Hereby, the current synclist of the sender VSD together with the current time-step of the sender scheduler are passed to the receiver scheduler. This is done by creating a new synchronization task containing the aforementioned information and inserting the task into the corresponding task-list of the receiver scheduler. That way, the receiver does not have to halt immediately when receiving a synchronization task. The execution of simulation components is not influenced by the broadcast of synchronization tasks. The receiver continues the execution of simulation components and only processes received synchronization tasks when the scheduling-routine arrives at that processing step. In return, a sender with a lower scheduling time-step than the receiver can perform unimpeded, multiple scheduling cycles and send multiple synclists to the receiver while the receiver may still be processing simulation components until both reach the same time-step. The received synchronization tasks are applied in time-stamp order and the simulation state of the receiver VSD is updated by applying all synclists. Pursuing this strategy allows an independent execution of sender and receiver schedulers with different time-steps until both schedulers reach the same time-step and the synchronization algorithm is applied to ensure that the causality constraint is not violated.

5.2 Synchronization Algorithm

Simulation threads can be paused during the synchronization phase of the scheduling main-routine. Sender and Receiver schedulers have to wait either for the dispatch or application of synchronization tasks if they progress faster than other schedulers that participate in the synchronization process. Because a synclist only contains references to changed items in the VSD, a sender scheduler always has to wait until all connected receivers with a smaller time-step have finished applying the synclist. We propose the following synchronization sequence, see Figure 6.

A scheduler starts the synchronization process and the main-routine by sending the current synclist containing all changes from the last scheduling time-step ($t_{i-1}^{scheduler}$) to its receivers. Hereby, receivers are wo-

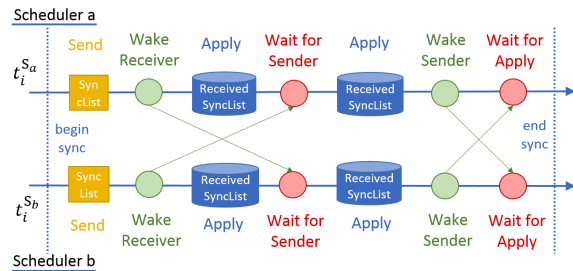


Figure 6: Synchronization sequence.

ken up that have waited on receiving that synclist. Next the scheduler processes the synclists received so far in time-stamp order ($t_i^{synclist} < t_i^{scheduler}$). After the first processing step, the scheduler checks if it has progressed faster than the remaining senders. If the scheduler has progressed faster, it has to wait until all senders have advanced sufficiently and transmitted their synclists up to the current time-step of that scheduler. The scheduler is immediately woken up as soon as a synclist arrives and starts processing it. In the last step of synchronization process the scheduler itself waits on the application of all transmitted synclists and is woken up as soon as the last receiver has applied the transmitted synclists.

Processing synchronization tasks in the sequence presented here, guarantees that synchronization tasks are not processed out of time-stamp order, that synchronization tasks do not contain future VSD changes and that no synchronization tasks are missed during a scheduling-cycle. All changes contained by the received synclists are applied during one time-step in exact the same order to the VSD of a receiver as they happened in the original VSDs of the senders. Hereby, the causality constraint is preserved and repeatable as well as identical results are produced during the parallel execution.

6 RESULTS / VALIDATION

The modular parallel simulation architecture presented here is still under development. First results are very promising regarding the performance. We tested several scenarios of the Virtual Space Robotics Testbed with different simulation component distribution configurations. In all cases the average performance during parallel execution met the desired outcome predicted by our partitioning scheme. The average performance was mainly dependent on the total execution time of the distributed components of the slowest thread. Our approach is capable of achieving performance gains that scale with the number of threads, as shown for the following scenario.

We choose a configuration of three simulation threads (see figure 7) each connected with each other by synchronization connections for the parallel simulation of a modular satellite reconfiguration scenario where a chaser satellite is docked to a satellite composed of individual building blocks (Weise et al., 2012) and utilizes robotic manipulators to exchange malfunctioning components.

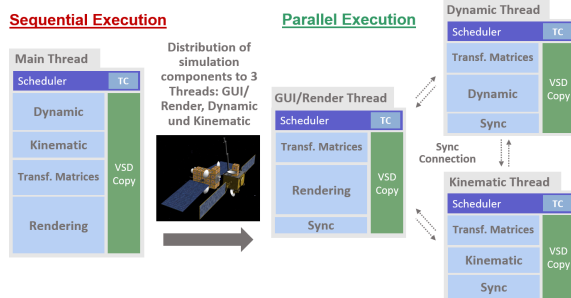


Figure 7: Modular parallel simulation of a satellite reconfiguration scenario.

During the parallel simulation of this scenario, we achieved a constant speed-up factor of 2.8 (see figure 8). Synchronization overhead was very low, on the average around 250 microseconds, while the execution time of the main components was within the range of milliseconds, see Figure 8. Despite the harsh execution times, in the range of a few milliseconds, our approach was capable of achieving a stable and constant speed-up during each simulation time-step, resulting in a linear progression of simulation time. Memory usage increased by 40.5 MB per VSD copy from initially 734 MB (consisting of: 444 MB for the simulation model and 290 MB for the simulation architecture including all components of the Virtual Space Robotics Testbed) to 815 MB after two additional VSD copies were generated for the configuration presented above and the simulation was started. Model geometries and textures were shared among the simulation threads. Due to the processing of synchronization tasks in time-stamp order, the parallel simulation of this scenario produced repeatable and identical results compared to the equivalent sequential execution. The results were generated on a Intel Core i7-3930k multi-core CPU with 6 physical cores at 4.1GHz and 16GB of RAM.

7 CONCLUSION / FUTURE WORK

We introduced a parallel simulation approach that addresses the challenges posed by VTs and developed a

parallel simulation architecture that presents an ideal platform for further research regarding the parallel execution of arbitrary VT applications. Scalability is provided by the functional partitioning scheme that allows a modular distribution of various simulation components among threads. The scheduler based simulation architecture facilitates the concurrent execution of multiple distributed simulations algorithms with different degrees of granularity. Distributed components are executed safely in parallel without interfering with each other, due to the utilization and assignment of independent VSD copies to simulation threads. Each VSD copy facilitates all the data and functionality needed for simulation. The change detection mechanisms provides an efficient distribution method of VSD updates among the simulation threads. In addition, the synchronization architecture allows an independent execution of sender and receiver schedulers. As a result, our approach generates low synchronization overhead and provides speed-ups that are scalable with the number of utilized threads during parallel execution.

Regarding the scalability during parallel execution, our modular approach is limited by the number of simulation components provided by an VT application. In future, we would like to investigate further partitioning schemes for VT applications to increase the scalability of our approach. Our functional partitioning scheme could be extended by a spatial partitioning to generate more complex configurations of distributed VT components. Furthermore we would like to investigate the application of computational load balancing strategies to realize an automatic distribution of components among all simulation threads during parallel execution.

ACKNOWLEDGMENTS

Parts of this work were developed in the context of the research project ViTOS. Supported by the German Aerospace Center (DLR) with funds of the German Federal Ministry of Economics and Technology (BMWi), support code 50 RA 1304.

REFERENCES

- Ballarini, P., Forlin, M., Mazza, T., and Prandi, D. (2009). Efficient parallel statistical model checking of biochemical networks. *arXiv preprint arXiv:0912.2551*.
- Bononi, L., Bracuto, M., D'Angelo, G., and Donatiello, L. (2005). Concurrent replication of parallel and distributed simulations. In *Principles of Advanced and*

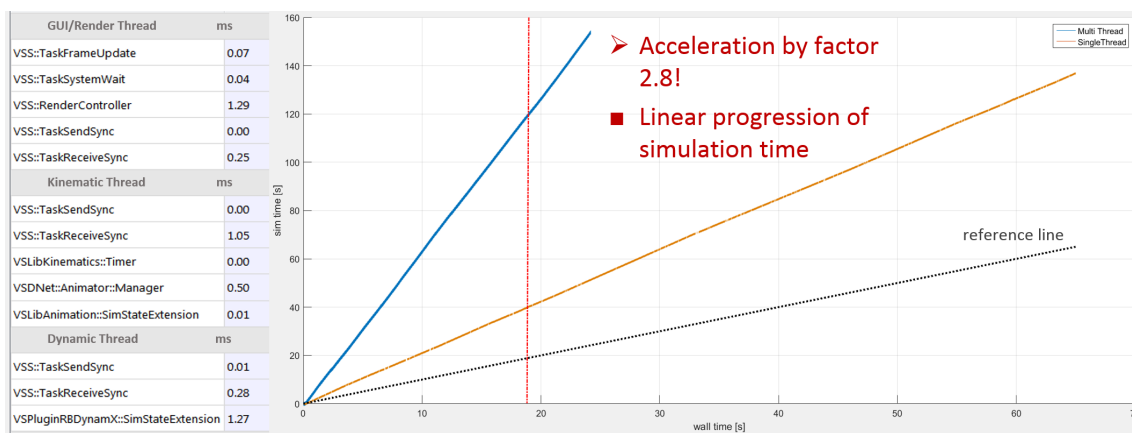


Figure 8: Performance measurements and execution times of the satellite reconfiguration scenario.

Distributed Simulation, 2005. PADS 2005. Workshop on, pages 234–243. IEEE.

Eickhoff, M. (2007). *Sequential Analysis of Quantiles and Probability Distributions by Replicated Simulations*. PhD thesis, University of Canterbury.

Ewald, R., Leye, S., and Uhrmacher, A. M. (2009). An efficient and adaptive mechanism for parallel simulation replication. In *Principles of Advanced and Distributed Simulation, 2009. PADS'09. ACM/IEEE/SCS 23rd Workshop on*, pages 104–113. IEEE.

Forlin, M., Mazza, T., and Prandi, D. (2010). Predicting the effects of parameters changes in stochastic models through parallel synthetic experiments and multivariate analysis. In *Parallel and Distributed Methods in Verification, 2010 Ninth International Workshop on, and High Performance Computational Systems Biology, Second International Workshop on*, pages 105–115. IEEE.

Friedmann, M., Petersen, K., and von Stryk, O. (2008). Simulation of multi-robot teams with flexible level of detail. In *International Conference on Simulation, Modeling, and Programming for Autonomous Robots*, pages 29–40. Springer.

Fujimoto, R. M. (2001). Parallel simulation: parallel and distributed simulation systems. In *Proceedings of the 33rd conference on Winter simulation*, pages 147–157. IEEE Computer Society.

Glasserman, P., Heidelberger, P., Shahabuddin, P., and Zajic, T. (1996). Splitting for rare event simulation: Analysis of simple cases. In *Proceedings of the 28th Conference on Winter Simulation, WSC '96*, pages 302–308, Washington, DC, USA. IEEE Computer Society.

Hybinette, M. and Fujimoto, R. (1997). Cloning: a novel method for interactive parallel simulation. In *Proceedings of the 29th conference on Winter simulation*, pages 444–451. IEEE Computer Society.

Mota, E., Wolisz, A., and Pawlikowski, K. (2000). A perspective of batching methods in a simulation environment of multiple replications in parallel. In *Proceedings of the 32nd Conference on Winter Simulation, WSC '00*, pages 761–766, San Diego, CA, USA. Society for Computer Simulation International.

Passerat-Palmbach, J., Caux, J., Siregar, P., Mazel, C., and Hill, D. (2015). Warp-level parallelism: Enabling multiple replications in parallel on gpu. *arXiv preprint arXiv:1501.01405*.

Pawlikowski, K., Yau, V. W., and McNickle, D. (1994). Distributed stochastic discrete-event simulation in parallel time streams. In *Proceedings of the 26th conference on Winter simulation*, pages 723–730. Society for Computer Simulation International.

Rossmann, J. and Schluse, M. (2011). Virtual robotic testbeds: A foundation for e-robotics in space, in industry-and in the woods. In *Developments in E-systems Engineering (DeSE), 2011*, pages 496–501. IEEE.

Rossmann, J., Schluse, M., Rast, M., and Atorf, L. (2016). eRobotics combining electronic media and simulation technology to develop (not only) robotics applications. In Kadry, S. and El Hami, A., editors, *E-Systems for the 21st Century – Concept, Developments, and Applications*, volume 2, chapter 10. Apple Academic Press. ISBN: 978-1-77188-255-2.

Roßmann, J., Schluse, M., and Waspe, R. (2013). Combining supervisory control, object-oriented petri-nets and 3d simulation for hybrid simulation systems using a flexible meta data approach. In *SIMULTECH*, pages 15–23.

Roth, M., Voss, G., and Reiners, D. (2004). Multi-threading and clustering for scene graph systems. *Computers & Graphics*, 28(1):63–66.

Voß, G., Behr, J., Reiners, D., and Roth, M. (2002). A multi-thread safe foundation for scene graphs and its extension to clusters. *EGPGV*, 2:33–37.

Weise, J., Briess, K., Adomeit, A., Reimerdes, H.-G., Gller, M., and Dillmann, R. (2012). An intelligent building blocks concept for on-orbit-satellite servicing. In *Proc. International Symposium on Artificial Intelligence Robotics and Automation in Space (iSAIRAS)*, Turin, Italy.