

Computer Science Methods for Effective and Sustainable Simulation Studies

Wentong Cai*¹, Christopher Carothers*², David M. Nicol*³, and Adelinde M. Uhrmacher*⁴

1 Nanyang TU – Singapore, SG. aswtcai@ntu.edu.sg

2 Rensselaer Polytechnic Institute – Troy, US. chrisc@cs.rpi.edu

3 University of Illinois – Urbana Champaign, US. dmmicol@illinois.edu

4 Universität Rostock, DE. adelinde.uhrmacher@uni-rostock.de

Abstract

This report documents the program and the (preliminary) outcomes of Dagstuhl Seminar 22401 “Computer Science Methods for Effective and Sustainable Simulation Studies”. The seminar has been dedicated to addressing central methodological challenges in conducting effective and sustainable simulation studies. Lightning talks provided the opportunity for participants to present their current research and ideas to advance methodological research in modeling and simulation. However, the lion’s share of the seminar was dedicated to working groups. One working group investigated how machine learning and modeling and simulation can be effectively integrated (Intelligent Modeling and Simulation Lifecycle). Another working group focused on methodological challenges to support policy via simulation (Policy by simulation: seeing is believing for interactive model co-creation and effective intervention). A third working group identified 4 challenges closely tied to the quest for sustainable simulation studies (Context, composition, automation, and communication – towards sustainable simulation studies) thereby, focusing on the role of model-based approaches and related methods.

Seminar October 3–7, 2022 – <http://www.dagstuhl.de/22401>

2012 ACM Subject Classification Computing methodologies

Keywords and phrases Modeling, simulation, high performance computing, machine learning, visual analytics

Digital Object Identifier 10.4230/DagRep.12.10.1

1 Executive Summary

Adelinde M. Uhrmacher (Universität Rostock, DE)

Wentong Cai (Nanyang TU – Singapore, SG)

Christopher Carothers (Rensselaer Polytechnic Institute – Troy, US)

David M. Nicol (University of Illinois – Urbana Champaign, US)

License © Creative Commons BY 4.0 International license

© Adelinde M. Uhrmacher, Wentong Cai, Christopher Carothers, and David M. Nicol

Motivation. Simulation becomes more and more important in application areas, establishing itself as the third way of science in addition to theory and (real) experiments. To answer research questions, simulation studies form increasingly intricate processes that intertwine the design and execution of various, often calculation-intensive simulation experiments, the generation and refinement of simulation models, and steps of analysis.

* Editor / Organizer



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Computer Science Methods for Effective and Sustainable Simulation Studies, *Dagstuhl Reports*, Vol. 12, Issue 10, pp. 1–60

Editors: Wentong Cai, Christopher Carothers, David M. Nicol, and Adelinde M. Uhrmacher



DAGSTUHL Dagstuhl Reports

REPORTS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

	Monday, 10/03	Tuesday, 10/04	Wednesday, 10/05	Thursday, 10/06	Friday, 10/07
08 a.m.		Breakfast	Breakfast	Breakfast	Breakfast
09 a.m.		Introduction incl. identifying challenges	Working groups	Intermediate results	Summary/Work.G.
10 a.m.		30 min. Coffee	30 min. Coffee	30 min. Coffee	30 min. Coffee
11 a.m.		Model-based approaches for Modeling and Simulation	Working groups	Working groups	Final summary
12 noon		Lunch	Lunch	Lunch	Lunch
01 p.m.		Visual analytics, experiment design, statistical model checking	Lightning talks	Working groups	Departure
02 p.m.		10 min. Coffee			
03 p.m.		High Performance Computing: exploiting new architectures			
04 p.m.	Arrival	30 min. Coffee	30 min. Coffee	30 min. Coffee	
05 p.m.		Collecting working group ideas	Hiking	Working groups	
06 p.m.		Dinner	Dinner	Dinner	

■ **Figure 1** Schedule of the Dagstuhl Seminar *Computer Science Methods for Effective and Sustainable Simulation Studies*.

The Dagstuhl Seminar has been dedicated to addressing central methodological challenges in supporting the conduction of effective and sustainable simulation studies. Thereby, the seminar focused on problems and solutions related to improving:

- Effectiveness: the usage of resources, including computing infrastructure and data, and the assistance of humans throughout a simulation study.
- Sustainability: continuing a simulation study into the future through support for reusing or building upon its central products, such as simulation model, data, and processes as well as the software used.

The last decades have seen a wide range of methodological developments in computer science that are likely to be instrumental in achieving effective and sustainable simulation studies. However, those efforts are scattered across different computer science fields that include high-performance computing, (modeling) language design, operations research, visual analytics, workflows, provenance, and machine learning, as well as modeling and simulation. The seminar brought participants with diverse computer science backgrounds together to enhance the methodological basis for conducting simulation studies.

Organization and results. Being one day shorter than typical seminars, the seminar started on Tuesday with a short round of introduction and continued with collecting ideas about achievements and challenges of modeling and simulation from the participants on 2 pinboards (see Figure 1). 3 talks and partly extensive discussions followed, one focusing on modeling and model-based approaches applied to simulation studies, one on high-performance computing for simulation, and one on analysis and experiment designs. In the late afternoon, the information gathered on the pinboards was revisited. In the end 3 working groups formed to work towards state-of-the-art and open-challenges papers on the following topics:

- Intelligent Modeling and Simulation Lifecycle
- Policy by simulation: seeing is believing for interactive model co-creation and effective intervention
- Context, composition, automation, and communication: towards sustainable simulation studies

Among the application fields as diverse as cell biological systems, traffic systems, or computer networks, one application dominated the discussions, i.e., Covid 19 simulation. The Covid pandemic showed the importance of modeling and simulation studies being conducted in an efficient, reliable manner, and, accordingly, of comprehensive, intelligent computer support for these studies, it revealed limitations, including those referring to communicating effectively modeling and simulation studies and their results to decision-makers. The results of the working groups are included as short summaries in this report. Wednesday afternoon, the participants presented their current research work and ideas in a series of lightning talks whose abstracts are also included in the report. However, most of the time was dedicated to the working groups. Plenary sessions on Thursday and finally on Friday allowed the participants to catch up with ideas and the progress made in the different working groups.

2 Table of Contents

Executive Summary

Adeline M. Uhrmacher, Wentong Cai, Christopher Carothers, and David M. Nicol 1

Overview of Talks

Towards Differentiable Agent-Based Simulation <i>Philipp Andelfinger</i>	6
Parametric verification of stochastic model using stochastic variational inference <i>Luca Bortolussi</i>	7
Beyond DDDAS and symbiotic simulation <i>Wentong Cai</i>	8
Towards a new facility for model-based design and evaluation of sustainable complex systems <i>Rodrigo Castro</i>	10
Challenges for Sustainable Twinning <i>Joachim Denil and Stijn Bellis</i>	10
The ASTRÉE Analyzer <i>Jérôme Feret</i>	12
Fighting COVID-19 with Simulation <i>Peter Frazier</i>	13
Virtual Time Integration of Emulation and Simulation Systems for Smart Grid Application Testing and Evaluation <i>Dong (Kevin) Jin</i>	13
Towards an Open Repository for Reproducible Performance Comparison of Parallel and Distributed Discrete-Event Simulators <i>Till Köster, Philipp Andelfinger, and Adeline M. Uhrmacher</i>	15
Simulation Based Analysis of Social Systems – Models, Data and Policy <i>Michael Lees</i>	16
Parallel Simulation – What Worked and What Not <i>Jason Liu</i>	17
Simulation at the Edge <i>Margaret Loper</i>	19
Bottlenecks of using simulation for policy making <i>Fabian Lorig</i>	21
Supporting Transparent Simulation Studies: The Role of Provenance Information <i>Bertram Ludäscher</i>	22
Interactive Visual Analysis for Simulation <i>Kresimir Matkovic</i>	24
A logic-based approach to reason about large-scale spatially-distributed systems <i>Laura Nenzi</i>	26
On the Attractiveness of Speculative PDES: Challenges and Pitfalls <i>Alessandro Pellegrini</i>	26

Methods for Integrated Simulation – from Data Acquisition to Decision Support <i>Niki Popper</i>	29
Using the Adaptable I/O System (ADIOS) for Effective and Sustainable Simulation Studies <i>Caitlin Ross</i>	32
A Simulation Architecture to Study Diffusion Processes in Multiplex Networks <i>Cristina Ruiz-Martin</i>	34
Model-based Software and Systems Engineering for Digital Twins <i>Bernhard Rumpe</i>	36
Data Farming: Better Decisions Via Inferential Big Data <i>Susan Sanchez</i>	37
Decision Making using Reinforcement Learning in Contested and Dynamic Environments <i>Claudia Szabo</i>	38
Simulation-based Inference for Automatic Model Construction <i>Wen Jun Tan</i>	39
Models and Specifications within the Modeling and Simulation Life Cycle <i>Adelinde M. Uhrmacher and Claudia Szabo</i>	40
A discrete-event approach for the study of sustainability in buildings <i>Gabriel A. Wainer</i>	40
Automatically Generating Simulation Experiments based on Provenance <i>Pia Wilsdorf</i>	43
Working groups	
Intelligent Modeling and Simulation Lifecycle <i>Wentong Cai, Philipp Andelfinger, Luca Bortolussi, Christopher Carothers, Dong (Kevin) Jin, Till Köster, Michael Lees, Jason Liu, Margaret Loper, Alessandro Pellegrini, Wen Jun Tan, and Verena Wolf</i>	44
Policy by simulation: seeing is believing for interactive model co-creation and effective intervention <i>Rodrigo Castro, Joachim Denil, Jérôme Feret, Kresimir Matkovic, Niki Popper, Susan Sanchez, and Peter Slood</i>	47
Context, composition, automation and communication: towards sustainable simulation studies <i>Adelinde M. Uhrmacher, Peter Frazier, Reiner Hähnle, Franziska Klügl, Fabian Lorig, Bertram Ludäscher, Laura Nenzi, Cristina Ruiz-Martin, Bernhard Rumpe, Claudia Szabo, Gabriel A. Wainer, and Pia Wilsdorf</i>	53
Participants	60

3 Overview of Talks

3.1 Towards Differentiable Agent-Based Simulation

Philipp Andelfinger (Universität Rostock, DE)

License © Creative Commons BY 4.0 International license

© Philipp Andelfinger

Main reference Philipp Andelfinger: “Towards Differentiable Agent-Based Simulation”, ACM Trans. Model. Comput. Simul., Vol. 32(4), Association for Computing Machinery, 2023.

URL <https://doi.org/10.1145/3565810>

Agent-based simulation models aiding the understanding, design, and optimization of systems reside on a spectrum with two extremes [5, 8]: *mechanistic models* are constructed manually by formalizing domain knowledge about the structure and behavior of the system under study, whereas *data-driven models* are generated by fitting generic parametric models against empirical observations of the system. While mechanistic models are typically also parameterized, adjusting the parameters does not alter the fundamental model logic.

The parameter synthesis for model calibration or optimization takes different forms depending on the model category: data-driven models usually permit the computation of the partial derivatives of the model output with respect to the parameters using automatic differentiation algorithms such as backpropagation [9]. Hence, gradient-based methods can be used to steer the parameter combination towards local optima in the model’s response surface. In contrast, mechanistic agent-based models tend to incorporate discrete decision-making logic [5], which can lead to response surfaces dominated by discontinuities and plateaus, thus largely preventing the fruitful use of gradient-based methods. As a consequence, most simulation optimization efforts using agent-based models employ metaheuristics [6] such as genetic algorithms or metamodeling approaches [4, 3], which generate an approximative surrogate of the original model.

We explore methods to make agent-based models involving discontinuous building blocks amenable to the automatic computation of gradients, under the hypothesis that the directed local search afforded by gradient-based methods may exhibit better convergence behavior than the existing approaches. Further, our aim is to enable the integration of the domain knowledge encoded in mechanistic agent-based models with the flexibility of data-driven models. By capturing the behavior of such a combined model, the computed gradients can serve to swiftly identify high-quality solutions for problems in calibration, optimization, control, and reinforcement learning.

Previously, we showed that when weighting the effects of different branches in the logic of agent-based models using a smoothing function, gradients computed using automatic differentiation can be used to accelerate the progress in traffic light control problems using microscopic traffic simulations [1, 2]. Moreover, we showed that the integration of the differentiable simulation model with a neural network enables the gradient-based training of a neural traffic light controller. The weighted execution of branches can be regarded as an approximation of an exact probabilistic program semantics, which is prohibitively expensive in practical cases [7]. However, using simplifications, the computational costs can be reduced to an acceptable level.

In the future, to avoid the need for modelers to manually apply smoothing to their agent-based models, languages or APIs are required that allow simulation models to be executed in their original or in a smoothed form. As models must be expected to vary severely in their suitability for smoothing and their potential for improvements in optimization progress, e.g., depending on the presence of continuous model elements, a categorization of models according to such properties would be beneficial.

By providing a natural way to unify mechanistic and data-driven models and the gradient-based methods used for optimization, we hope for this work to reduce the gap between the communities focused on mechanistic and data-driven modeling.

References

- 1 Philipp Andelfinger. Differentiable agent-based simulation for gradient-guided simulation-based optimization. In *Proceedings of the 2021 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, pages 27–38, 2021.
- 2 Philipp Andelfinger. Towards differentiable agent-based simulation. *ACM Trans. Model. Comput. Simul.*, 2022.
- 3 Russell R Barton. Tutorial: metamodeling for simulation. In *2020 Winter Simulation Conference (WSC)*, pages 1102–1116. IEEE, 2020.
- 4 Atharv Bhosekar and Marianthi Ierapetritou. Advances in surrogate based modeling, feasibility analysis, and optimization: A review. *Computers & Chemical Engineering*, 108:250–267, 2018.
- 5 Eric Bonabeau. Agent-based modeling: methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(suppl 3):7280–7287, 2002.
- 6 Benoît Calvez and Guillaume Hutzler. Automatic tuning of agent-based models using genetic algorithms. In *International Workshop on Multi-Agent Systems and Agent-Based Simulation*, pages 41–57. Springer, 2005.
- 7 Swarat Chaudhuri and Armando Solar-Lezama. Smoothing a program soundly and robustly. In *International Conference on Computer Aided Verification*, pages 277–292. Springer, 2011.
- 8 Hamdi Kavak, Jose J Padilla, Christopher J Lynch, and Saikou Y Diallo. Big data, agents, and machine learning: towards a data-driven agent-based modeling approach. In *Proceedings of the Annual Simulation Symposium*, pages 1–12, 2018.
- 9 David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.

3.2 Parametric verification of stochastic model using stochastic variational inference

Luca Bortolussi (University of Trieste, IT)

License © Creative Commons BY 4.0 International license
© Luca Bortolussi

Joint work of Luca Bortolussi, Francesca Cairolì, Ginevra Carbone, Paolo Pulcini

Main reference Luca Bortolussi, Francesca Cairolì, Ginevra Carbone, Paolo Pulcini: “Scalable Stochastic Parametric Verification with Stochastic Variational Smoothed Model Checking”, arXiv, 2022.

URL <https://doi.org/10.48550/ARXIV.2205.05398>

Parametric verification for stochastic models can be expressed as checking the satisfaction probability of a certain property as a function of the parameters of the model. Smoothed model checking (smMC) [1] aims at inferring the satisfaction function over the entire parameter space from a limited set of observations obtained via simulation. As observations are costly and noisy, smMC is framed as a Bayesian inference problem so that the estimates have an additional quantification of the uncertainty. In [1] the authors use Gaussian Processes (GP), inferred by means of the Expectation Propagation algorithm. This approach provides accurate reconstructions with statistically sound quantification of the uncertainty. However, it inherits the well-known scalability issues of GP. In this paper, we exploit recent advances in probabilistic machine learning to push this limitation forward, making Bayesian inference of smMC scalable to larger datasets, enabling its application to models with high dimensional

parameter spaces. We propose Stochastic Variational Smoothed Model Checking (SV-smMC), a solution that exploits stochastic variational inference (SVI) to approximate the posterior distribution of the smMC problem. The strength and flexibility of SVI make SV-smMC applicable to two alternative probabilistic models: Gaussian Processes (GP) and Bayesian Neural Networks (BNN). The core ingredient of SVI is a stochastic gradient-based optimization that makes inference easily parallelizable and it enables GPU acceleration. In this paper, we compare the performances of smMC [1] against those of SV-smMC by looking at the scalability, the computational efficiency and at the accuracy of the reconstructed satisfaction function.

References

- 1 Bortolussi, L., Milios, D., Sanguinetti, G.: Smoothed model checking for uncertain continuous-time markov chains. *Information and Computation* 247, 235–253 (2016)

3.3 Beyond DDDAS and symbiotic simulation


Wentong Cai (Nanyang TU – Singapore, SG)

License  Creative Commons BY 4.0 International license
© Wentong Cai

Darema in 2000 as a paradigm where measurement data from an operational system is dynamically incorporated into an execution model of that system, and computational results from the model are then used to guide the measurement process. Independently, Symbiotic Simulation [2] paradigm proposed in Dagstuhl Seminar on Grand Challenges for Modelling and Simulation in 2002 solves the what-if problem by having the simulation system and the physical system interact in a mutually beneficial manner. Since their inception, many techniques have been developed to support DDDAS and symbiotic simulations. Digital Twin, originally proposed by Michael Grieves in 2002 and popularized in recent years due to the rise of IoT and AI, is “a real-time virtual representation of a real-world physical system or process that serves as the indistinguishable digital counterpart of it for practical purposes, such as system simulation, integration, testing, monitoring, and maintenance” [3]. Digital twin is a concept. DDDAS and symbiotic simulation techniques form the basis for the realization of digital twins. They emphasize on the interaction between the virtual and physical systems.

Moving beyond DDDAS and symbiotic simulation, data-driven and machine learning (ML) techniques should be integrated in various stages of M&S. For instance, in addition to use sensor data from physical system to calibrate simulation models, ML techniques can be used to extract useful knowledge and insight from the data to facilitate model development. Data analytics and ML techniques can also be used to manipulate or steer simulation experiments on the fly (see Figure 2).

Some examples of our recent works along this direction include: Our PADS’21 paper [4] is about how to use ML approach to create a car-following model (instead of using traditional physics-based model) and how to dynamically calibrate the model using online data. Our recent research work focuses more on using data-driven approach to improve performance of simulation execution and simulation-based optimization. Our PADS’20 & PADS’22 papers [5, 6] are about dynamically analyzing simulation state to determine level of details to be used in the model of a simulation entity during simulation execution. The objective is to reduce the simulation runtime while maintaining accuracy of the simulation results. We applied the approach to a semi-conductor manufacturing simulation. And our WSC’18



■ **Figure 2** Beyond DDDAS – Integration of machine learning.

paper [7] uses an approach to dynamically predict the usefulness of a simulation run. If the results of a simulation run won't contribute to the overall optimization objective, then the simulation run can be terminated early. In this way, the total number of simulation runs required in a simulation-based optimization process will be reduced.

References

- 1 Dynamic Data Driven Applications System. (2022), https://en.wikipedia.org/wiki/Dynamic_Data_Driven_Applications_System, [Online; accessed 23-Sept-2022]
- 2 R. Fujimoto, D. Lunceford, E. Page, and A. Uhrmacher. Grand challenges for modeling and simulation. *Schloss Dagstuhl*. 350 (2002)
- 3 Digital twin. (2022), https://en.wikipedia.org/wiki/Digital_twin, [Online; accessed 23-Sept-2022]
- 4 Htet Naing, Wentong Cai, Nan Hu, Tiantian Wu, and Liang Yu. Data-driven microscopic traffic modelling and simulation using dynamic lstm. *Proceedings Of The 2021 ACM SIGSIM Conference On Principles Of Advanced Discrete Simulation*. pp. 1-12 (2021)
- 5 Moon Gi Seok, Chew Wye Chan, Wentong Cai, Daejin Park, and Hessam S. Sarjoughian. Adaptive abstraction-level conversion framework for accelerated discrete-event simulation in smart semiconductor manufacturing. *IEEE Access*. 8 pp. 165247-165262 (2020)
- 6 Moon Gi Soon, Wen Jun Tan, and Wentong Cai. Hyperparameter Tuning in Simulation-based Optimization for Adaptive Digital-Twin Abstraction Control of Smart Manufacturing System. *Proceedings Of The 2022 ACM SIGSIM Conference On Principles Of Advanced Discrete Simulation*. pp. 61-68 (2022)
- 7 Philipp Andelfinger, Sajeev Udayakumar, David Eckhoff, Wentong Cai, and Alois Knoll. A. Model preemption based on dynamic analysis of simulation data to accelerate traffic light timing optimisation. *2018 Winter Simulation Conference (WSC)*. pp. 652-663 (2018)

3.4 Towards a new facility for model-based design and evaluation of sustainable complex systems

Rodrigo Castro (*University of Buenos Aires, AR*)

License © Creative Commons BY 4.0 International license
© Rodrigo Castro

Problems involving societies and their interactions with cybernetic systems in the context of physical restrictions represent a paradigmatic case of Complex Adaptive Systems (CAS) involving emergent behavior and micro-macro loops.

The study of dynamic CAS lack analytical solutions thus requiring simulation as the only means for quantitative research.

If we add a layer of goal-seeking governance to inform real-world policy-making, legitimately contradictory worldviews must be factored in. This paves the way to reaching what has been termed as Wicked Problems, those that do not accept “correct” definitions nor “optimal” solutions, but rather require discursive processes to reach a consensus about the models themselves.

We postulate Planning for Sustainable Egalitarian Development as an embracing, flagship case study that pushes the envelope of sustainable simulation, thus challenging the state of the art and practice of several computer science disciplines (HPC, model checking, validation and verification, visualization, to name a few)

The proposal includes building on control theory to include simulation models in the loop of decision-making processes, creating a simulation-assisted arena to experiment with interventions, thus obtaining “feasible future developments” and analyzing them with advanced visual analytics.

We envision a framework within which planners (human intelligence) and algorithms (artificial intelligence) inform each other to obtain better strategies to use the simulation models as demonstrators of feasible paths of development.

An advanced facility, such as an interactive and immersive visualization room for complex data and reactive simulations shall integrate and boost capabilities for participatory model-based design and evaluation of sustainable complex systems.

3.5 Challenges for Sustainable Twinning

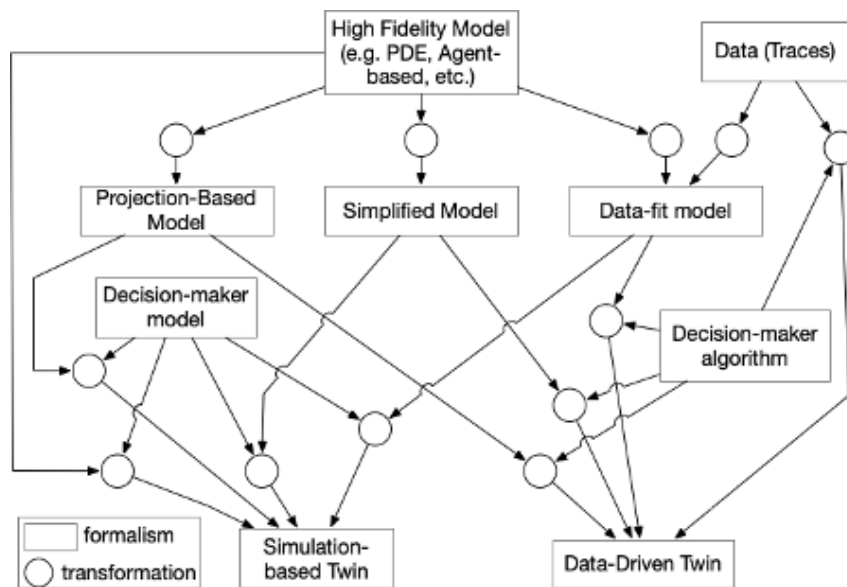
Joachim Denil (*University of Antwerp, BE*) and Stijn Bellis

License © Creative Commons BY 4.0 International license
© Joachim Denil and Stijn Bellis

Main reference Stijn Bellis, Joachim Denil: “Challenges and possible approaches for sustainable digital twinning”, in Proc. of the 25th International Conference on Model Driven Engineering Languages and Systems: Companion Proceedings, MODELS 2022, Montreal, Quebec, Canada, October 23-28, 2022, pp. 643–648, ACM, 2022.

URL <https://doi.org/10.1145/3550356.3561551>

Advances in digital twin technology are creating value for many companies. We consider, from a sustainability perspective, how these digital twins can be better developed. At first glance, the energy consumption of the twin during its life cycle can be described as follows: $E_{total} = E_{design} + E_{local} + E_{networking} + E_{cloud} + E_{update}$. Decomposing this formula allows us to see several challenges in the design and operation of a digital twin.



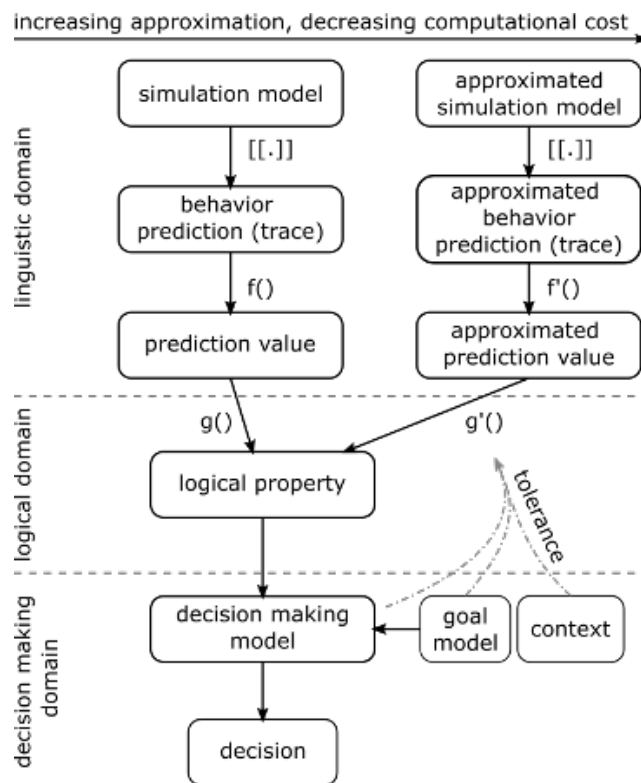
■ **Figure 3** Transformation of Models for a Digital Twin Architecture.

The choice of formalism. When developing a twin from detailed physics models (typically modelled using partial differential equations), different choices concerning the formalism(s) must be made. In Figure 3, we show some simplified paths to obtaining a digital twin. Each operation takes a certain amount of energy while running a simulation using a specific formalism during the twin lifecycle also takes energy. This results in trade-offs that need to be managed. For example, creating a lumped model takes a lot of engineering time while training a neural net takes a lot of energy.

The value proposition of the twin. From the perspective of simulation engineering, the purpose has a huge influence on the engineering of the simulation model. Figure 4 shows how simulation is used to decide on properties of interest. It also shows how approximated models provide value for their users. Estimating how much uncertainty the decision maker can tolerate for gaining enough value from the digital twin is a difficult problem. Furthermore, once we know the allowable tolerance, we still need to include this in the formalism selection process.

The evolution of the system and the twin. To allow for the long and continuous operation of the digital twin, we need insight into the range of validity of the model in combination with insights into the system's evolution. Including the dimensions of evolution within the formalism selection process is needed. Having a good estimate of these evolutions in frequency and severity helps determine the needed boundary conditions and validity of the model. If not considered, a new model needs to be used, possibly a model that consumes more energy.

The deployment of the twin. The final challenge is to reason about the deployment choices related to the deployment architecture used for the digital twin. Most choices are impacted by the system's requirements, which in turn depend on the value proposition of the twin. Some examples of choices include (a) where to run (parts of) the twin: local, edge or cloud? (b) networking and telemetry choices, (c) cloud infrastructure and storage options.



■ **Figure 4** Approximations and their Effect on Validity.

3.6 The ASTRÉE Analyzer

Jérôme Feret (ENS – Paris, FR)

License © Creative Commons BY 4.0 International license
© Jérôme Feret

Joint work of Bruno Blanchet, Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, Antoine Miné, David Monniaux, Xavier Rival

Main reference Bruno Blanchet, Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, Antoine Miné, David Monniaux, Xavier Rival: "A static analyzer for large safety-critical software", in Proc. of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation 2003, San Diego, California, USA, June 9-11, 2003, pp. 196–207, ACM, 2003.

URL <https://doi.org/10.1145/781131.781153>

The ASTRÉE analyzer [2, 3, 4] aims at statically proving the absence of Run-Time errors in critical embedded software. Following the abstract interpretation framework [1], it is based on a formal semantics of C. Then various abstract domains are available to abstract this semantics. The correctness of the approach is proven by construction. ASTRÉE has been used successfully to certify the absence of run-time errors in the primary flight control software of the A340 (2003), the primary flight control software of the A380 (2005), and a C version of the animatic docking software of the Jules Vernes ATV (for the International Space Station). Since 2009, it is commercialized by Absint Angewandte Informatik (Saarbrücken) and is used mainly in automotive software, but also, in avionics and nuclear software.

References

- 1 Cousot, P., Cousot, R.: Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In: Graham, R.M., Harrison, M.A., Sethi, R. (eds.) POPL'77: Proceedings of the 4th Symposium on Principles of Programming Languages. pp. 238–252. ACM (1977)

- 2 Blanchet, B., Cousot, P., Cousot, R., Feret, J., Mauborgne, L., Miné, A., Monniaux, D., Rival, X.: Design and implementation of a special-purpose static program analyzer for safety-critical real-time embedded software. In: Mogensen, T.Æ., Schmidt, D.A., Sudborough, I.H. (eds.) *The Essence of Computation, Complexity, Analysis, Transformation. Essays Dedicated to Neil D. Jones [on occasion of his 60th birthday]*. Lecture Notes in Computer Science, vol. 2566, pp. 85–108. Springer (2002). https://doi.org/10.1007/3-540-36377-7_5
- 3 Blanchet, B., Cousot, P., Cousot, R., Feret, J., Mauborgne, L., Miné, A., Monniaux, D., Rival, X.: A static analyzer for large safety-critical software. In: Cytron, R., Gupta, R. (eds.) *Proceedings of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation 2003, San Diego, California, USA, June 9-11, 2003*. pp. 196–207. ACM (2003). <https://doi.org/10.1145/781131.781153>
- 4 Cousot, P., Cousot, R., Feret, J., Mauborgne, L., Miné, A., Monniaux, D., Rival, X.: The astreé analyzer. In: Sagiv, S. (ed.) *Programming Languages and Systems, 14th European Symposium on Programming, ESOP 2005, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings*. Lecture Notes in Computer Science, vol. 3444, pp. 21–30. Springer (2005). https://doi.org/10.1007/978-3-540-31987-0_3

3.7 Fighting COVID-19 with Simulation

Peter Frazier (Cornell University – Ithaca, US)

License © Creative Commons BY 4.0 International license
© Peter Frazier

Joint work of Peter I. Frazier, J. Massey Cashore, Ning Duan, Shane G. Henderson, Alyf Janmohamed, Brian Liu, David B. Shmoys, Jiayue Wan, Yujia Zhang
Main reference Peter I. Frazier, J. Massey Cashore, Ning Duan, Shane G. Henderson, Alyf Janmohamed, Brian Liu, David B. Shmoys, Jiayue Wan, Yujia Zhang: “Modeling for COVID-19 college reopening decisions: Cornell, a case study”, *Proceedings of the National Academy of Sciences*, Vol. 119(2), p. e2112532119, 2022.

URL <https://doi.org/10.1073/pnas.2112532119>

Universities faced a difficult decision in summer 2020: whether to reopen for in-person instruction despite the pandemic and how to protect students, employees, and the surrounding community if they did. Simulation was critical to this decision at Cornell University in the USA, which successfully reopened for in-person instruction in fall 2020 under the protection of an asymptomatic screening that tested all undergraduate students twice per week. This talk discusses key factors that helped simulation modelers earn credibility and influence: transparency, designing a testing strategy that was robust to the unknown, explaining analysis in a simple clear way, understanding stakeholders’ incentives, responsiveness to stakeholder questions, and a focus on providing value.

3.8 Virtual Time Integration of Emulation and Simulation Systems for Smart Grid Application Testing and Evaluation

Dong (Kevin) Jin (University of Arkansas – Fayetteville, US)

License © Creative Commons BY 4.0 International license
© Dong (Kevin) Jin

Modern energy systems are increasingly adopting Internet technology to boost control efficiency, which unfortunately opens up a new security frontier. As a result, extensive applications have been proposed to enhance the cyber resilience and security of those critical

infrastructures. Incorporation of new technologies in such systems is very challenging because of strong real-time requirements, continuous system availability, and many resource-constrained legacy devices. Therefore, a testing platform targeting such cyber-physical systems is strongly needed for the research community to evaluate the new application and system designs and their impact on the power systems before the real deployment.

We develop a unique testbed, DSSnet [1], combining container-based network emulation and power system simulation using a novel Linux-kernel-based virtual time system. DSSnet enables the modeling of a modern power distribution system and simulates the Intelligent Electrical Devices that make it up. DSSnet also enables high-fidelity analysis by allowing real networking applications to run in the network emulator and interact with the power simulator. DSSnet is composed of the following main components: (1) a container-based network emulator, Mininet [2] that allows execution of real software and communication network applications, (2) an electrical power distribution system simulator, OpenDSS that enables power flow simulation studies [3], (3) a unique Linux-kernel-based virtual time system [4] for synchronization of the two sub-systems, which significantly enhances the temporal fidelity issues in ordinary co-simulation or hardware-in-the-loop testbeds; (4) two coordinators for interfacing with the cyber- and physical-side modules and the virtual time system; and (5) a distributed software-defined networking (SDN) control environment, ONOS [5] that provides high-level abstractions and APIs for power grid control applications to manage, monitor and program the emulated communication network.

One key challenge is synchronizing the execution of the power simulator and the container-based emulator. This is because all the processes in the emulator execute real programs and use the system clock to advance experiments, while the simulator executes models to advance experiments with respect to its simulation virtual clock. To address this issue, we developed an independent and lightweight middleware in the Linux kernel to support virtual time for Linux container [4]. Our system transparently provides the virtual time to processes inside the containers, while returning the ordinary system time to other processes. No change is required in applications. Next, we expanded the capability of the testbed with a distributed virtual time system [6] that enables processes and their clocks to be paused, resumed, and dilated across embedded Linux devices through the use of hardware interrupts and a common kernel module. The distributed system architecture uniquely consists of a common virtual time Linux kernel module and three communication channels, one for virtual time synchronization using general-purpose-input-and-output (GPIO) hardware interrupts, one for connecting the embedded Linux devices, and one for interfacing with the physical system simulation that performs an offline computation. Additionally, we modeled and analyzed the temporal error during non-CPU operations, such as disk I/O, network I/O, and GPU computation, and developed a barrier-based time compensation mechanism to enable accurate virtual time advancement with precise I/O time measurement and compensation [7].

In summary, we present DSSnet, a testing platform that combines an electrical power system simulator and a communication network emulator using a virtual time system. DSSnet can be used to model and simulate power flows, communication networks, and smart grid control applications, and to test and evaluate the effect of network applications on the smart grid.

References

- 1 Christopher Hannon, Jiaqi Yan, and Dong Jin. DSSnet: A smart grid modeling platform combining electrical power distribution system simulation and software-defined networking emulation. ACM SIGSIM-PADS, 2016.
- 2 Mininet: An Instant Virtual Network on your Laptop (or other PC). <http://mininet.org/>

- 3 OpenDSS: An Electrical Power System Simulation Tool for Distribution Systems, *Elect. Power Res. Inst.* <http://smartgrid.epri.com/SimulationTool.aspx>
- 4 Jiaqi Yan and Dong Jin. A Virtual Time System for Linux-container-based Emulation of Software-defined Networks. *ACM SIGSIM-PADS*, 2015.
- 5 ONOS: Open Network Operating System. <https://opennetworking.org/onos/>
- 6 Christopher Hannon, Jiaqi Yan, Dong Jin, and Yuan-An Liu. A Distributed Virtual Time System on Embedded Linux for Evaluating Cyber-Physical Systems. *ACM SIGSIM-PADS*, 2019.
- 7 Gong Chen, Zheng Hu, and Dong Jin. Integrating I/O Time to Virtual Time System for High Fidelity Container-based Network Emulation. *ACM SIGSIM-PADS 2022*.

3.9 Towards an Open Repository for Reproducible Performance Comparison of Parallel and Distributed Discrete-Event Simulators

Till Köster (Universität Rostock, DE), Philipp Andelfinger (Universität Rostock, DE), and Adelinde M. Uhrmacher (Universität Rostock, DE)

License © Creative Commons BY 4.0 International license

© Till Köster, Philipp Andelfinger, and Adelinde M. Uhrmacher

Main reference Till Köster, Adelinde M. Uhrmacher, Philipp Andelfinger: “Towards an Open Repository for Reproducible Performance Comparison of Parallel and Distributed Discrete-Event Simulators”, in *Proc. of the SIGSIM-PADS ’22: SIGSIM Conference on Principles of Advanced Discrete Simulation*, Atlanta, GA, USA, June 8 – 10, 2022, pp. 31–32, ACM, 2022.

URL <https://doi.org/10.1145/3518997.3534989>

Performance is one of the core motivations in the field of parallel and distributed simulation. Contributions for new methods and optimizations frequently rely on custom models, parametrizations, and baseline implementations. This makes a direct comparison between methods and approaches difficult. We present our vision and initial steps towards COMPADS, a benchmark model and repository for reproducibly comparing the performance of parallel and distributed simulators and their respective algorithms. COMPADS[1] is short for COMparing Parallel And Distributed Simulators. The first results include a novel deterministic-by-design synthetic benchmark model inspired by PHOLD and La-pdes. The benchmark output is a checksum that attests to the correctness of an implementation and its execution. So far, implementations exist for the simulators ROOT-Sim and ROSS.

References

- 1 Till Köster, Adelinde Uhrmacher, and Philipp Andelfinger. 2022. Towards an Open Repository for Reproducible Performance Comparison of Parallel and Distributed Discrete-Event Simulators. In *Proceedings of the 2022 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (SIGSIM-PADS ’22)*. Association for Computing Machinery, New York, NY, USA, 31–32. <https://doi.org/10.1145/3518997.3534989>

3.10 Simulation Based Analysis of Social Systems – Models, Data and Policy

Michael Lees (University of Amsterdam, NL)

License © Creative Commons BY 4.0 International license
© Michael Lees

Joint work of Michael Lees, Eric Dignum, Andreas Flache, Willem Boterman

Main reference Eric Dignum, Efi Athieniti, Willem Boterman, Andreas Flache, Michael Lees: “Mechanisms for increased school segregation relative to residential segregation: a model-based analysis”, *Comput. Environ. Urban Syst.*, Vol. 93, p. 101772, 2022.

URL <https://doi.org/10.1016/j.compenvurbsys.2022.101772>

The role of modelling and simulation in the scientific domains has had a long history and *computational X* is now a well-established area in Physics, Chemistry, Biology and more. During the last decade with the increase in detailed data, and the development of novel modelling techniques, the application of modelling and simulation has become more commonplace in the social sciences. In some cases, these models provide scientists and policymakers with unique ways to reason about sociological challenges (e.g., Polarization, Segregation and Inequality).

In this talk, I present a sample of current work [1] in which we develop models to understand the process of primary school choice and school segregation with the Municipality of Amsterdam. In the talk, I present an agent-based model where households face residential decisions depending on neighbourhood compositions and make school choices based on distance and school compositions. Using a global sensitivity analysis we demonstrate that the observed excess (the level of school segregation compared to residential segregation) segregation in schools occurs for a wide range of parameters and demonstrate that asymmetric preferences (residential vs. school selection) are not a requirement for excess school segregation.

Using this case study I highlight a number of challenges and opportunities for modelling and simulation within the social sciences. Firstly, the use of models for theory building and testing in social sciences, in particular how simple models with clear assumptions can demonstrate potential causes for social phenomena. Secondly, using techniques from simulation-based inference I demonstrate how novel calibration methods offer promising solutions to calibrate city-scale models of social dynamics using microdata. Finally, I present some real-world cases where a “digital twin” of the city can be used to answer important policy questions for the municipality of Amsterdam and help them statistically estimate the likely outcomes for different interventions.

I conclude the talk by highlighting a number of initiatives [2] within Amsterdam and the Netherlands where new computational infrastructure presents unique opportunities to conduct computational social science and modelling simulation at a city and country-wide scale.

References

- 1 Dignum, E., Athieniti, E., Boterman, W., Flache, A., & Lees, M. (2022). Mechanisms for increased school segregation relative to residential segregation: A model-based analysis. *Computers, Environment and Urban Systems*, 93, 101772.
- 2 Emery, T., Braukmann, R., Wittenberg, M., van Ossenbruggen, J., Siebes, R., & van de Meer, L. (2020). The ODISSEI Portal: Linking Survey and Administrative Data.

3.11 Parallel Simulation – What Worked and What Not

Jason Liu (Florida International University – Miami, US)

License © Creative Commons BY 4.0 International license
© Jason Liu

Parallel discrete-event simulation (PDES) refers to the class of techniques and tools for efficiently running discrete-event simulation on parallel and distributed computing platforms. Applications of PDES include performance modeling and simulation of large systems. Simulation of large systems exacts high computational demand, and successful PDES efforts must be able to cope with both the scale and complexity of the target systems on modern parallel computing architectures. In this talk, we summarize some of our research efforts for developing and applying parallel simulation of various systems.

Traditional PDES research has been largely focused on examining efficient parallel synchronization algorithms and incorporating those in simulation tools for general applications. Previously, we developed DaSSF [1], a simulator implemented in C and C++ that incorporates a composite parallel synchronization algorithm [2]. The algorithm was extended to run on both shared and distributed memory architectures and was implemented in a simulator called MiniSSF [3]. We have also explored the use of scripting languages (including Python, LUA, and Javascript) to simplify model development. The simulator, called Simian, was shown to be able to achieve good, and in some cases, even superior performance by taking advantage of just-in-time compilation [4]. One latest effort was the development of a simulator with full-fledged support for the process-oriented simulation world-view in Python for fast development cycle [5]. Abraham Maslow once said: *“If you have a hammer, everything looks like a nail.”* These and other simulators have been used for simulations of large-scale computer systems and networks, including the Internet [6], mobile ad-hoc networks [7], high-performance computing interconnection networks [8], and parallel files systems [9].

PDES not only enables large-scale simulation of complex systems, but can also be incorporated with real-time and interactive simulation of large systems due to its superior simulation performance. We previously designed and implemented a real-time network simulator to run on parallel and distributed platforms [10]. With real-time simulation, simulated network protocols, such as TCP, can seamlessly interact with real network entities in a hybrid simulation and emulation setting [11]. One can also control and steer the network experiments in real time – for example, by injecting network events and observing the results via a remote dashboard during the live experiment [12].

Many large complex systems feature a huge number of components and processes that may inter-operate across multiple layers of the system hierarchy and at different time granularity. A fine-grained simulation may not be able to scale up to the required size even if PDES could achieve linear speedup. Solving problems in many cases does not require brute force. George Box once said: *“All models are wrong but some are useful.”* In this case, one must be able to use multi-resolution models to balance the trade-off between simulation performance and accuracy. A case in point is the hybrid network traffic modeling, which combines fluid traffic models (e.g., based on ordinary differential equations) and packet-oriented simulation to achieve faster-than-real-time performance for large-scale network simulation [13, 14]. Another example is network simulation and emulation symbiosis. To allow high-fidelity high-performance large-scale network experimentation, one can run a full-scale detailed network model on high-performance parallel computing platforms, and an emulation system, which executes unmodified applications in a virtual machine environment configured to represent the target system. Both systems need to represent the same traffic behavior. We

applied model reduction techniques to scale down the model complexity both in emulation to improve its performance and in data exchange between the two systems to reduce the synchronization and communication overhead [15, 16].

We observe PDES research has evolved by leaps and bounds over the last three decades. Many PDES techniques and tools have matured in various domains, although the community continues to discover new techniques, tools, and applications, many coinciding with the emergence of big data systems, machine learning techniques, and data-driven applications. We predict PDES will continue to play an important, and sometimes indispensable, role in modeling dynamic and complex systems, in many cases though combining with new techniques in order to provide its unique capability in solving problems.

References

- 1 J. Liu, D. Nicol, B. J. Premore, and A. L. Poplawski, "Performance prediction of a parallel simulator," in *Proceedings of the thirteenth workshop on Parallel and distributed simulation*. IEEE Computer Society, 1999, pp. 156–164.
- 2 D. M. Nicol and J. Liu, "Composite synchronization in parallel discrete-event simulation," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 13, no. 5, pp. 433–446, 2002.
- 3 J. Liu and R. Rong, "Hierarchical composite synchronization," in *Principles of Advanced and Distributed Simulation (PADS), 2012 ACM/IEEE/SCS 26th Workshop on*. IEEE, 2012, pp. 3–12.
- 4 N. Santhi, S. Eidenbenz, and J. Liu, "The simian concept: parallel discrete event simulation with interpreted languages and just-in-time compilation," in *Proceedings of the 2015 Winter Simulation Conference*. IEEE Press, 2015, pp. 3013–3024.
- 5 J. Liu, "Simulus: easy breezy simulation in python," in *2020 Winter Simulation Conference (WSC)*. IEEE, 2020, pp. 2329–2340.
- 6 D. M. Nicol, J. Liu, M. Liljenstam, and G. Yan, "Simulation of large scale networks using ssf," in *Simulation Conference, 2003. Proceedings of the 2003 Winter*, vol. 1. IEEE, 2003, pp. 650–657.
- 7 J. Liu, Y. Yuan, D. M. Nicol, R. S. Gray, C. C. Newport, D. Kotz, and L. F. Perrone, "Simulation validation using direct execution of wireless ad-hoc routing protocols," in *Parallel and Distributed Simulation, 2004. PADS 2004. 18th Workshop on*. IEEE, 2004, pp. 7–16.
- 8 K. Ahmed, J. Liu, S. Eidenbenz, and J. Zerr, "Scalable interconnection network models for rapid performance prediction of HPC applications," in *2016 IEEE 18th International Conference on High Performance Computing and Communications (HPCC)*, Dec 2016, pp. 1069–1078.
- 9 M. Erazo, T. Li, J. Liu, S. Eidenbenz *et al.*, "Toward comprehensive and accurate simulation performance prediction of parallel file systems," in *Dependable Systems and Networks (DSN), 2012 42nd Annual IEEE/IFIP International Conference on*. IEEE, 2012, pp. 1–12.
- 10 J. Liu, "A primer for real-time simulation of large-scale networks," in *Proceedings of the 41st Annual Simulation Symposium (ANSS'08)*, 2008, pp. 85–94.
- 11 M. A. Erazo, Y. Li, and J. Liu, "SVEET! a scalable virtualized evaluation environment for TCP," in *Proceedings of the 5th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities and Workshops (TRIDENTCOM'09)*, 2009, pp. 1–10.
- 12 N. Van Vorst, M. Erazo, and J. Liu, "PrimoGENI: Integrating real-time network simulation and emulation in GENI," in *Proceedings of the 2011 IEEE Workshop on Principles of Advanced and Distributed Simulation (PADS)*, 2011, pp. 1–9.
- 13 J. Liu, "Packet-level integration of fluid TCP models in real-time network simulation," in *Proceedings of the 2006 Winter Simulation Conference (WSC'06)*, December 2006, pp. 2162–2169.

- 14 J. Liu, “Parallel simulation of hybrid network traffic models,” in *Proceedings of the 21st Workshop on Principles of Advanced and Distributed Simulation (PADS’07)*, June 2007, pp. 141–151.
- 15 M. A. Erazo, R. Rong, and J. Liu, “Symbiotic network simulation and emulation,” *ACM Trans. Model. Comput. Simul.*, vol. 26, no. 1, pp. 2:1–2:25, 2015.
- 16 R. Rong and J. Liu, “Distributed mininet with symbiosis,” in *2017 IEEE International Conference on Communications (ICC)*, 2017, pp. 1–6.

3.12 Simulation at the Edge

Margaret Loper (Georgia Institute of Technology – Atlanta, US)

License © Creative Commons BY 4.0 International license
© Margaret Loper

Mobile networks evolve roughly every ten years, each generation bringing new services and capabilities. By 2030, 6G will bring more capable, intelligent, reliable, scalable, and power-efficient communications. This will enable new applications such as holographic telepresence, collaborative autonomous driving and personalized body area networks [3]. The 6G era will help realize the hyper-connected world of people, data and things – the Internet of Everything (IoE). It will also enable a focus on small data, generated by the plethora of edge devices embedded in our everyday world. The IoE will create a need for efficient processing of massive amounts of small data and edge intelligence, a process where data is collected, analyzed, and insights produced near the end user. The promise of 6G and the emergence of edge intelligence will provide users with actionable real-time information. An extension of this concept is to also provide users with actionable real-time predictions. The intersection of edge computing, sensor networks, artificial intelligence and online predictive simulations enable a new vision called Simulation at the Edge.

Simulation Paradigms. Paradigms for sensor-driven simulations first emerged in the mid-2000s with concepts like symbiotic, ad hoc and data driven adaptive simulation systems. Symbiotic simulation is a paradigm which refers to a close relationship between a simulation system and a physical system. It was defined at the Dagstuhl Seminar on Grand Challenges for Modeling and Simulation in 2002 [5]. The simulation system benefits from real-time measurements about the physical system which are provided by sensors, and the physical system may benefit from decisions made by the simulation [1]. An important concept in symbiotic simulation is the “what-if” analysis, where multiple experiments investigate alternative scenarios based on data from the physical system. Symbiotic simulation does not refer to a specific type of simulation (e.g., real-time, discrete event), it is an umbrella term which refers to independent simulations employed to analyze alternative scenarios regarding a physical system.

The Dynamic Data-Driven Application Systems (DDDAS) concept utilizes online data to drive simulation computations, and the results are then used to optimize the system or adapt the measurement process [2]. For example, live sensor data and analytics can be used to construct or infer the current state of a system and faster-than-real-time simulation can then be used to project the system’s future state. Also, simulation can be used to control an operational system, e.g., data from a real system are fed directly into the simulation model which analyzes alternate options and produces recommended courses of action.

An ad hoc distributed simulation is a collection of autonomous on-line simulations, each modeling some portion of a larger physical system, that are brought together to predict future states of the overall system [4]. In a conventional distributed simulation, the system being modeled is partitioned into non-overlapping elements (e.g., geographic regions) in a top-down fashion. By contrast, ad hoc simulations are constructed bottom-up, resulting in multiple simulators modeling common, overlapping portions of the physical system. In other words, it is constructed in an “ad hoc” fashion, in much the same way a collection of mobile radios join together to form an ad hoc wireless network. Ad hoc simulations are on-line simulation programs, meaning they are able to capture the current state of the system through measurement, and then execute forward as rapidly as possible to project a future state of the system.

Edge Intelligence and Simulation. These simulation paradigms have some similarities in their approach, but different levels of success in accomplishing their vision. For example, they are all data-driven using real-time sensor input and the simulations run faster-than-real-time to predict future state, look at “what if” scenarios or steer measurements. Where they differ include their interaction or feedback with the physical system, as well as the number of simulations in use. For example, DDDAS can be tightly coupled with the real system, steering the measurement process, where symbiotic and ad hoc do not. Further, ad hoc uses more than one simulation to model different parts of a system, whereas symbiotic and DDDAS are focused on a single, central simulation. DDDAS has been quite successful modeling large scale structural systems, urban water systems and transportation systems; and symbiotic simulations have had success in semiconductor manufacturing, business process optimization, and control of unmanned aerial vehicles. While research into ad hoc simulation has been limited to transportation and queueing systems [6], it has so much more potential in the era of 6G and IoE.

The proliferation of mobile computing and IoE, edge computing is an emerging paradigm that pushes computing tasks and services from the network core to the network edge. Edge computing is widely recognized as a promising solution for processing the “zillion” bytes of data generated by IoE devices [7]. It has also attracted attention for its promise to reduce latency, save bandwidth, improve availability, and to keep data secure. At the same time, a proliferation of AI algorithms and models which accelerate the deployment of intelligence in edge devices has emerged. These trends, called Edge Intelligence, can power the evolution of ad hoc simulation to Simulation at the Edge.

Embedding simulations within edge intelligence brings the simulation closer to the data, lessening the need to aggregate sensor data in order to reduce communication bandwidth requirements. It also has the potential to be more resilient to failures, especially communication failures, as portions of the system could be managed under local control. The application of intelligent edge devices embedded with predictive simulations are varied and diverse. They could be used to monitor transportation systems, or rerouting vehicle traffic after a severe accident; track the spread of wild fires, floods, and pollution; optimize emergency responses, such as evacuations during floods or tornadoes; provide self-optimizing communication networks, by reconfiguring the physical network to improve performance and avoid bottlenecks; or respond to breakdowns within a manufacturing system.

Despite the decades of research in symbiotic, DDDAS and ad hoc simulation, research challenges remain. These include: compact representation of system state, fault tolerant and robust systems, multi-resolution modeling, and automatic validation [5]. Like its predecessor, distributed Simulation at the Edge raise a number of intriguing questions. Can such a distributed simulation make sufficiently accurate predictions of future system states to be useful? Can they incorporate new information and revise projections more rapidly and/or effectively than conventional approaches, e.g., global, centralized simulations? How would Simulation at the Edge be organized and operate? The power of edge intelligence and 6G could be the catalyst to help answer these questions.

References

- 1 Heiko Aydt, Stephen John Turner, Wentong Cai, and Malcolm Yoke Hean Low. Research issues in symbiotic simulation. In *Proceedings of the 2009 winter simulation conference (WSC)*, pages 1213–1222. IEEE, 2009.
- 2 F Darema, M Rotea, M Goldberg, DH Newlon, JC Cherniavsky, JE Figueroa, JE Hudson, C Friedman, P Lyster, and R Bohn. Dddas: dynamic data driven applications systems. URL: <http://www.nsf.gov/pubs/2005/nsf05570/nsf05570.htm> (Accessed 20 September 2013), 2005.
- 3 Chamitha De Alwis, Anshuman Kalla, Quoc-Viet Pham, Pardeep Kumar, Kapal Dev, Won-Joo Hwang, and Madhusanka Liyanage. Survey on 6g frontiers: Trends, applications, requirements, technologies and future research. *IEEE Open Journal of the Communications Society*, 2:836–886, 2021.
- 4 Richard Fujimoto, Michael Hunter, Jason Sirichoke, Mahesh Palekar, Hoe Kim, and Wonho Suh. Ad hoc distributed simulations. In *21st International Workshop on Principles of Advanced and Distributed Simulation (PADS'07)*, pages 15–24. IEEE, 2007.
- 5 Richard Fujimoto, WH Lunceford Jr, Ernst H Page, and Adelinde Uhrmacher. Grand challenges for modelling and simulation (dagstuhl seminar 02351). Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2002.
- 6 Ya-Lin Huang, Christos Alexopoulos, Michael Hunter, and Richard M Fujimoto. Ad hoc distributed simulation methodology for open queueing networks. *Simulation*, 88(7):784–800, 2012.
- 7 Zhi Zhou, Xu Chen, En Li, Liekang Zeng, Ke Luo, and Junshan Zhang. Edge intelligence: Paving the last mile of artificial intelligence with edge computing. *Proceedings of the IEEE*, 107(8):1738–1762, 2019.


3.13 Bottlenecks of using simulation for policy making

Fabian Lorig (Malmö University, SE)

License © Creative Commons BY 4.0 International license
© Fabian Lorig

Computer simulation has become an established method for assessing uncertainty by conducting what-if analyses across a variety of disciplines and domains. It allows us to observe the behavior of a system under different circumstances and to investigate the potential consequences of different actions or interventions without actually interfering with the system we want to analyze. During the Covid-19 pandemic, for instance, we could see that a great number of models was developed already in the first months after the initial outbreak of the crisis [1]. Figure 5 shows that models were developed for all different kinds of potential interventions and to investigate how they possibly might affect transmission processes and the overall dynamics of the pandemic.

In practice, however, many researchers experienced that their models and the results generated by their simulation studies were not considered by policy makers when deciding upon which interventions to pursue and implement. In addition to this, other shortcomings in the development and application of simulation models and the conducting of simulation studies could be observed. Most researchers, for instance, did not reuse existing simulations and instead have chosen to start developing new models and to conduct new studies, which affects responsiveness in crisis situations. Potential reasons might be the limited availability of models, ambiguous documentations of the models and previously conducted studies, or



■ **Figure 5** Number of models that support the simulation of non-pharmaceutical and pharmaceutical interventions [1].

a general lack of trust in simulations studies conducted by others. This leads to a series of questions regarding why simulations sometimes fail to support what they were actually intended for namely to provide valuable insights and to facilitate decision making processes.

In this regard, when we discuss how to improve the effectiveness and sustainability of simulation studies, it seems that there might be different perspectives on this very issue. A simulation that we as developers and modellers consider highly insightful, comprehensive, and trustworthy might not be convincing, helpful, or plausible for the decision makers the study was intended for. Therefore, when discussing the effectiveness and sustainability of simulation studies, we should not forget about the stakeholders that might be involved in the process of a simulation study. How can we make sure to develop and generate what decision makers actually need? Why do we prefer to conduct own (new) studies instead of reusing the models and results of others? And how can we improve the rigor and trustworthiness of our studies?

References

- 1 Lorig, F., Johansson, E., & Davidsson, P. *Agent-based Social Simulation of the COVID-19 Pandemic: A Systematic Review*. JASSS: Journal of Artificial Societies and Social Simulation, 24(3).

3.14 Supporting Transparent Simulation Studies: The Role of Provenance Information

Bertram Ludäscher (University of Illinois at Urbana-Champaign, US)

License © Creative Commons BY 4.0 International license
© Bertram Ludäscher

Computational science experiments and simulation studies often require significant computational and human resources, making it impractical or impossible to repeat (i.e., reimplement and/or re-execute) these studies for reproducibility purposes. Transparency, on the other

hand, is arguably always desirable [8] and can be achieved by revealing and laying open the assumptions, general approach, computational methods, codes, parameter settings, runtime environments, and – last not least – the data sources used to conduct a study. Regarding the latter, precisely identifying the correct subsets of data that were used in a computational experiment is already a difficult challenge [3, 12].

Provenance (a.k.a. lineage) information captures the origin and processing history of data products [5, 7, 6], i.e., it is a form of metadata that provides the technical means to support transparency [10]. While the existing W3C standard PROV [11] provides a minimal baseline for exchanging provenance information, domain-specific extensions need to be developed by and for the community to capture more of the application context, the semantics of data and parameters, and assumptions of simulation models and scientific workflows [13, 16].

Transparent research objects [1, 15, 4, 10, 9, 14], with their data, computational, and provenance artifacts, and research papers (which tell the “science story” of a paper) will continue to co-evolve and should ultimately converge towards open, transparent, and reproducible research tales [2]. – *Declaremus et calculemus!*


References

- 1 Bechhofer, S., De Roure, D., Gamble, M., Goble, C., Buchan, I.: Research Objects: Towards Exchange and Reuse of Digital Knowledge. *Nature Precedings* (Jul 2010). <https://doi.org/10.1038/npre.2010.4626.1>
- 2 Brinckman, A., Chard, K., Gaffney, N., Hategan, M., Jones, M.B., Kowalik, K., Kulasekaran, S., Ludäscher, B., Mecum, B.D., Nabrzyski, J., Stodden, V., Taylor, I.J., Turk, M.J., Turner, K.: Computing environments for reproducibility: Capturing the “Whole Tale”. *Future Gener. Comput. Syst.* **94**, 854–867 (2019). <https://doi.org/10.1016/j.future.2017.12.029>
- 3 Buneman, P., Christie, G., Davies, J.A., Dimitrellou, R., Harding, S.D., Pawson, A.J., Sharman, J.L., Wu, Y.: Why data citation isn’t working, and what to do about it. *Database* (Jan 2020). <https://doi.org/10.1093/databa/baaa022>
- 4 Chard, K., Gaffney, N., Jones, M.B., Kowalik, K., Ludäscher, B., McPhillips, T., Nabrzyski, J., Stodden, V., Taylor, I., Thelen, T., Turk, M.J., Willis, C.: Application of BagIt-Serialized Research Object Bundles for Packaging and Re-Execution of Computational Analyses. *15th Intl. Conf. on eScience* pp. 514–521 (Sep 2019). <https://doi.org/10.1109/eScience.2019.00068>
- 5 Freire, J., Koop, D., Santos, E., Silva, C.T.: Provenance for Computational Tasks: A Survey. *Computing in Science Engineering* **10**(3), 11–21 (May 2008). <https://doi.org/10.1109/MCSE.2008.79>
- 6 Herschel, M., Diestelkämper, R., Lahmar, H.B.: A survey on provenance: What for? What form? What from? *The VLDB Journal* **26**(6), 881–906 (Dec 2017). <https://doi.org/10.1007/s00778-017-0486-1>
- 7 Ludäscher, B.: A Brief Tour Through Provenance in Scientific Workflows and Databases. In: Lemieux, V.L. (ed.) *Building Trust in Information*, pp. 103–126. Springer Proceedings in Business and Economics, Springer International Publishing (2016). https://doi.org/10.1007/978-3-319-40226-0_7
- 8 McPhillips, T., Ludäscher, B., Goble, C., Willis, C., Bowers, S.: Workshop Report – T7 Workshop on Provenance for Transparent Research. Tech. Rep., Zenodo, Provenance Week 2021 (Aug 2021). <https://doi.org/10.5281/zenodo.5301583>
- 9 McPhillips, T.M., Thelen, T., Willis, C., Kowalik, K., Jones, M.B., Ludäscher, B.: CPR-A Comprehensible Provenance Record for Verification Workflows in Whole Tale. *Provenance and Annotation of Data and Processes*. LNCS, Springer International Publishing, (2021). https://doi.org/10.1007/978-3-030-80960-7_23

- 10 McPhillips, T.M., Willis, C., Gryk, M.R., Corrales, S.N., Ludäscher, B.: Reproducibility by other means: Transparent research objects. 15th Intl. Conf. on eScience pp. 502–509. (Sep 2019). <https://doi.org/10.1109/eScience.2019.00066>
- 11 Moreau, L., Groth, P., Cheney, J., Lebo, T., Miles, S.: The rationale of PROV. Web Semantics: Science, Services and Agents on the World Wide Web **35**(Part 4), 235–257 (Dec 2015). <https://doi.org/10.1016/j.websem.2015.04.001>
- 12 Rauber, A., Gößwein, B., Zwölf, C.M., Schubert, C., Wörister, F., Duncan, J., Flicker, K., Zettsu, K., Meixner, K., McIntosh, L.D., Jenkyns, R., Pröll, S., Miksa, T., Parsons, M.A.: Precisely and Persistently Identifying and Citing Arbitrary Subsets of Dynamic Data. Harvard Data Science Review **3**(4) (Nov 2021). <https://doi.org/10.1162/99608f92.be565013>
- 13 Ruschinski, A., Wilsdorf, P., Dombrowsky, M., Uhrmacher, A.M.: Capturing and Reporting Provenance Information of Simulation Studies Based on an Artifact-Based Workflow Approach.: Proc. ACM SIGSIM Conference on Principles of Advanced Discrete Simulation. pp. 185–196. (2019). <https://doi.org/10.1145/3316480.3325514>
- 14 Soiland-Reyes, S., Sefton, P., Crosas, M., Castro, L.J., Coppens, F., Fernández, J.M., Garijo, D., Grüning, B., La Rosa, M., Leo, S., Ó Carragáin, E., Portier, M., Trisovic, A., RO-Crate Community, Groth, P., Goble, C.: Packaging research artefacts with RO-Crate. Data Science **5**(2), 97–138 (Jan 2022). <https://doi.org/10.3233/DS-210053>
- 15 Ton That, D.H., Fils, G., Yuan, Z., Malik, T.: Sciunits: Reusable Research Objects. 13th Intl. Conf. on E-Science (eScience). pp. 374–383 (Oct 2017). <https://doi.org/10.1109/eScience.2017.51>
- 16 Wilsdorf, P., Wolpers, A., Hilton, J., Haack, F., Uhrmacher, A.M.: Automatic Reuse, Adaptation, and Execution of Simulation Experiments via Provenance Patterns. ACM Transactions on Modeling and Computer Simulation (Sep 2022). <https://doi.org/10.1145/3564928>

3.15 Interactive Visual Analysis for Simulation

Kresimir Matkovic (VRVis – Wien, AT)

License  Creative Commons BY 4.0 International license
© Kresimir Matkovic

Visualization and interactive visual analysis have been used to explore and analyze simulation data for a long time. At the beginning, results of single run simulation have been visualized. With advancement of storage and computing technologies, the models became more and more complex and, at the same time, ensemble summation or simulation experiments – multiple simulation runs using variations of the same model – became possible. In case of computational fluid dynamics simulation the ensembles contain a relatively small number of members. Due to large and complex spatial-temporal data and long computing time it is not feasible to compute many models. Wang et al. [6] provide a recent survey on visualization of such ensemble simulations. In cases where simulation can be computed fast it is possible to compute hundreds, thousands or even tens of thousands of simulation runs. Matković et al. [4] provide an overview of such approaches.

The main idea here is to deploy coordinated multiple views that visualize multidimensional parameter space and complex simulation results at once. Interaction is used to support exploration and analysis. The user can brush, i.e. interactively select a subset of data in any view, and the items correspond to the brushed subset will be highlighted in all views. The idea functions if the number of parameters does not exceed half a dozen, or so. A

parameter space of a higher dimensionality requires many runs in order to be sufficiently covered. Interactive simulating ensemble steering can be used in such cases. A kind of automatic analysis can be integrated in order to guide the expert.

In the case of multi-model simulation, there is little available support from the visualization community. Simultaneous exploration of simulation results computed with models of different levels of detail remains a challenge for the visualization. Large data, hierarchical data structure, and a need for fluent switching of context depending on level (and maybe even providing the overall context across the levels at the same time) is subject of future research. However, based on successful deployment of interactive visualization in the simulation community up to know, it is plausible to reason that interactive visualization can become a key interface in navigating in complicated multi-level models and simulation results, and a great support in comprehending underlying phenomena for different stake-holders. Scalability often represents a serious problem in visualization for simulation. In case of multi-model simulation, scalability will also represent a challenge. Finally, provenance tracking will gain in importance, as user actions across multiple levels need to be stored and recalled on demand. Finally, Dimara and Stasko [1] recently reported on the missing link between user tasks in visualisation taxonomies (e.g., sensemaking) and the high-level task of decision-making. One of the key causes of this mismatch is a lack of interdisciplinary approaches. A close collaboration with simulation experts represents a valid approach to establishing the missing link.


In order to support future requirements from the simulation community we need further advances in interactive visualization. We expect the novel approaches to be inspired by several directions of interactive visualization, depending on the simulation methods and corresponding models, as well as on the identified explorative tasks. Besides coordinated multiple views, the promising pillars of future research are certainly focus and context techniques [3], comparative visualization [2], tree [5], networks and graphs visualization, as well as plethora of existing visualization for simulation methods and provenance tracking.

References

- 1 Evanthia Dimara and John Stasko. A critical reflection on visualization research: Where do decision making tasks hide? *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1128–1138, 2022.
- 2 Michael Gleicher, Danielle Albers, Rick Walker, Ilir Jusufi, Charles D. Hansen, and Jonathan C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, oct 2011.
- 3 Helwig Hauser. Generalizing focus+context visualization. In Georges-Pierre Bonneau, Thomas Ertl, and Gregory M. Nielson, editors, *Scientific Visualization: The Visual Extraction of Knowledge from Data*, pages 305–327, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- 4 Krešimir Matković, Denis Gračanin, and Helwig Hauser. Visual analytics for simulation ensembles. In *Proceedings of the 2018 Winter Simulation Conference, WSC '18*, page 321–335. IEEE Press, 2018.
- 5 Tamara Munzner, François Guimbretière, Serdar Tasiran, Li Zhang, and Yunhong Zhou. Treejuxtaposer: Scalable tree comparison using focus+context with guaranteed visibility. *ACM Trans. Graph.*, 22(3):453–462, jul 2003.
- 6 Junpeng Wang, Subhashis Hazarika, Cheng Li, and Han-Wei Shen. Visualization and visual analysis of ensemble data: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 25(9):2853–2872, 2019.

3.16 A logic-based approach to reason about large-scale spatially-distributed systems

Laura Nenzi (University of Trieste, IT)

License  Creative Commons BY 4.0 International license
© Laura Nenzi

From the reliability in a wireless sensor network to the formation of traffic jams, spatio-temporal patterns are key in understanding how complex behaviors can emerge in a network of locally interacting dynamical systems. One of the most important and intriguing questions is how to describe such behaviors in a formal and human-understandable specification language. A possible approach consists in using formal methods and in particular logic languages. In this talk, we show how a logic specification can be used to specify and analyse complex behaviours of large-scale spatially-distributed systems. Furthermore, we briefly show how we can use the logic for parameter synthesis, anomaly detection and the automatic feature extraction from spatio-temporal data.

3.17 On the Attractiveness of Speculative PDES: Challenges and Pitfalls

Alessandro Pellegrini (University of Rome “Tor Vergata”, IT)

License  Creative Commons BY 4.0 International license
© Alessandro Pellegrini

Introduction: The last 40 years have witnessed an explosion of methodologies and techniques related to speculative PDES [8]. Looking at the history of this research line, while performance aspects have always been paramount, each decade has seen its hot points. Roughly speaking, in the first decade, algorithmic solutions were proposed that enabled significant speedups with relatively little effort (see, e.g., [12]). The second decade began to propose solutions for the adaptivity of particular simulation aspects (see, e.g., [6, 10]). The third decade focused on effectively supporting models’ programming by hiding the complexity of Speculative PDES (see, e.g., [23]). The fourth decade showed the technique’s robustness even on new hardware/software architectures [2, 13, 21].

We can conclude that Speculative PDES brings together effective solutions and methodologies for executing many classes of models, offering significant speedups and making tractable problems that otherwise would not be.

However, these results do not appear to be generally exploited. Market penetration of these techniques is severely limited and often, with some exceptions, confined to the academic sphere. The research community has already shown that there are various scenarios in which exploiting speculative PDES can bring benefits (e.g. Agent-Based Models [1], Spiking Neural Networks, SNN [20, 16], Traffic [11] and Hardware Architecture Simulations [25]). However, the most widely used simulation software does not consider this methodology at all. To give a few examples, in traffic simulation, SUMO [3] is single-threaded; in SNN simulations, most simulators are sequential (Brian [22]) or, if parallel, are based on conservative/time stepped algorithms (NEST [9]). Architecture simulators, such as gem5 [4], are also sequential.

Over the past few years, we have been wondering what the reasons might be for this poor uptake of methodologies that the scientific community has proven to be effective in many fields. From these motivations, we have tried to identify some challenges that we believe are relevant to make Speculative PDES more attractive to other fields of research and industry.

Challenges. In general, using speculative PDES to support model execution is difficult. However long and deep the problem has been studied, it is still not wholly possible to hide the complexity of algorithms such as Time Warp from model developers. While some solutions allow the technical complexities to be hidden, it is still true that a model not explicitly developed for Time Warp may provide inadequate performance. Work on self-tuning and self-optimization has shown that it is possible to improve overall performance by optimizing specific parts of the algorithm (e.g., the GVT computation). However, a significant challenge could be to build an optimization approach capable of mixing different, more or less optimistic algorithms that can provide no worse performance than a sequential simulation.

Attractiveness to other domains must necessarily come through real-world models, which has also recently been recognized as vital [19]. The scientific community has often used synthetic benchmarks (see, for example, [7, 17]) to study Time Warp behaviour. However, scientists in other fields or industrial settings may be unable to map these results to their use cases. An important challenge could be to design and implement a real-world benchmark suite to show how effective speculative PDES can be in these areas. Approaches of this kind have already been followed in other research fields (see [15]) and could be successful for more widespread adoption of speculative PDES.

Another problem that makes using Speculative PDES difficult is that among the various implementations built by research groups (see, e.g., [5, 18, 14]), there is no uniformity of interface. PDES researchers have done much work on abstract model representations (e.g., [24]), but an agreement on interfaces among developers of runtime environments would also be desirable. Moreover, this would greatly help in the cross-fertilization of approaches used in the runtimes and would more easily allow experimental studies to be carried out, also bringing benefit to the scientific community dealing with PDES.

Another problem we encountered is related to the quality of support libraries for PDES runtime environments. Many tools rely on MPI interfaces for message exchange, which is a correct choice to support deployment on supercomputers. Unfortunately, it has been observed that the most modern MPI features are not widely used in the High Performance Computing world, especially those related to asynchronous execution. The result is that the most common implementations often suffer from correctness bugs¹. Moreover, MPI libraries do not consider highly concurrent and asynchronous usage patterns. Thus, even if the implementations are correct, they do not hold up to the message rate, especially if the runtime is highly optimised. This makes the work done on runtime environments unusable by researchers in other fields. The experience gained in developing runtime environments for speculative simulations could be used to create extremely fast message exchange libraries geared toward speculative PDES simulations. In this way, the usability of research results could be significantly improved.

References

- 1 Sameera Abar, Georgios K Theodoropoulos, Pierre Lemarinier, and Gregory M P O'Hare. Agent based modelling and simulation tools: A review of the state-of-art software. *Computer Science Review*, 24:13–33, 2017.
- 2 Peter D Barnes, Christopher D Carothers, David R Jefferson, and Justin M LaPre. Warp speed: executing time warp on 1,966,080 cores. In *Proceedings of the 1st ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, SIGSIM PADS '13, pages 327–336, New York, NY, USA, May 2013. ACM.

¹ Some issues that we identified while developing a high-performance PDES runtime [18] can be found following the links [here](#), [here](#), and [here](#).

- 3 Michael Behrisch, Laura Bieker, Jakob Erdmann, and Daniel Krajzewicz. SUMO—simulation of urban mobility: an overview. In *Proceedings of SIMUL 2011, The Third International Conference on Advances in System Simulation*. elib.dlr.de, 2011.
- 4 Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R Hower, Tushar Krishna, Somayeh Sardashti, Rathijit Sen, Korey Sewell, Muhammad Shoaib, Nilay Vaish, Mark D Hill, and David A Wood. The gem5 simulator. *SIGARCH Comput. Archit. News*, 39(2):1–7, August 2011.
- 5 Christopher D Carothers, David Bauer, and Shawn Pearce. ROSS: A high-performance, low-memory, modular time warp system. *Journal of parallel and distributed computing*, 62(11):1648–1669, November 2002.
- 6 Josef Fleischmann and Philip A Wilsey. Comparative analysis of periodic state saving techniques in time warp simulators. In *Proceedings of the 9th workshop on Parallel and Distributed Simulation*, PADS '95, pages 50–58, Piscataway, NJ, USA, July 1995. IEEE Computer Society.
- 7 Richard M Fujimoto. Performance of time warp under synthetic workloads. In David Nicol, editor, *Proceedings of the SCS Multiconference on Distributed Simulation*, pages 23–28, San Diego, CA, USA, 1990. Society for Computer Simulation International.
- 8 Richard M Fujimoto, Rajive Bagrodia, Randal E Bryant, K Mani Chandy, David Jefferson, Jayadev Misra, David Nicol, and Brian Unger. Parallel discrete event simulation: The making of a field. In *2017 Winter Simulation Conference (WSC)*, pages 262–291, December 2017.
- 9 Marc-Oliver Gewaltig and Markus Diesmann. *NEST (NEural Simulation Tool)*, volume 2, chapter 4. Scholarpedia, 2007.
- 10 D W Glazer and Carl Tropper. On process migration and load balancing in time warp. *IEEE Transactions on Parallel and Distributed Systems*, 4:318–327, 1993.
- 11 Masatoshi Hanai, Toyotaro Suzumura, Georgios Theodoropoulos, and Kalyan S Perumalla. Exact-Differential Large-Scale traffic simulation. In *Proceedings of the 3rd ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, SIGSIM PADS '15, pages 271–280, New York, NY, USA, June 2015. Association for Computing Machinery.
- 12 David R Jefferson. Virtual time. *ACM Transactions on Programming Languages and Systems*, 7(3):404–425, July 1985.
- 13 Xihu Liu and Philipp Andelfinger. Time warp on the GPU: Design and assessment. In *Proceedings of the 2017 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, SIGSIM-PADS '17, pages 109–120, New York, NY, USA, May 2017. ACM.
- 14 D E Martin, T J McBrayer, and P A Wilsey. WARPED: a time warp simulation kernel for analysis and application development. In *Proceedings of the 29th Hawaii International Conference on System Sciences*, volume 1 of *HICSS*, pages 383–386 vol.1, Piscataway, NJ, USA, January 1996. IEEE Computer Society.
- 15 Chi Cao Minh, Jaewoong Chung, Christos Kozyrakis, and Kunle Olukotun. STAMP: Stanford transactional applications for Multi-Processing. In *2008 IEEE International Symposium on Workload Characterization*, pages 35–46, September 2008.
- 16 Quang Anh Pham Nguyen, Philipp Andelfinger, Wentong Cai, and Alois Knoll. Transitioning spiking neural network simulators to heterogeneous hardware. In *Proceedings of the 2019 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, SIGSIM-PADS, pages 115–126, New York, NY, USA, May 2019. ACM.
- 17 Eunjung Park, Stephan Eidenbenz, Nandakishore Santhi, Guillaume Chapuis, and Bradley Settlemyer. Parameterized benchmarking of parallel discrete event simulation systems: Communication, computation, and memory. In *2015 Winter Simulation Conference (WSC)*, pages 2836–2847. ieeexplore.ieee.org, December 2015.

- 18 Alessandro Pellegrini, Roberto Vitali, and Francesco Quaglia. The ROME OpTimistic simulator: Core internals and programming model. In *Proceedings of the 4th International ICST Conference on Simulation Tools and Techniques*, SIMUTOOLS, pages 96–98, Brussels, Belgium, April 2012. ICST.
- 19 Kalyan Perumalla, Maximilian Bremer, Kevin Brown, Cy Chan, Stephan Eidenbenz, K Scott Hemmert, Adolfo Hoisie, Benjamin Newton, James Nutaro, Tomas Opielstrup, and Others. Computer science research needs for parallel discrete event simulation (PDES). Technical report, USDOE Office of Science (SC)(United States), 2022.
- 20 Adriano Pimpini, Andrea Piccione, Bruno Ciciani, and Alessandro Pellegrini. Speculative distributed simulation of very large spiking neural networks. In *Proceedings of the 2022 SIGSIM Conference on Principles of Advanced Discrete Simulation*, SIGSIM PADS, New York, NY, USA, 2022. ACM.
- 21 Shafiqur Rahman, Nael Abu-Ghazaleh, and Walid Najjar. PDES-A: a parallel discrete event simulation accelerator for FPGAs. In *Proceedings of the 2017 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, SIGSIM-PADS '17, pages 133–144, New York, NY, USA, May 2017. Association for Computing Machinery.
- 22 Marcel Stimberg, Romain Brette, and Dan F M Goodman. Brian 2, an intuitive and efficient neural simulator. *eLife*, 8(e47314):e47314, August 2019.
- 23 Roberto Toccaceli and Francesco Quaglia. DyMeLoR: Dynamic memory logger and restorer library for optimistic simulation objects with generic memory layout. In *Proceedings of the 22nd Workshop on Principles of Advanced and Distributed Simulation*, PADS, pages 163–172, Piscataway, NJ, USA, June 2008. IEEE Computer Society.
- 24 Bernard P Zeigler, Tag Gon Kim, and Herbert Praehofer. *Theory of Modeling and Simulation*. Academic Press, January 2000.
- 25 L Zhu, G Chen, B K Szymanski, C Tropper, and T Zhang. Parallel logic simulation of million-gate VLSI circuits. In *13th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, pages 521–524. ieeexplore.ieee.org, September 2005.

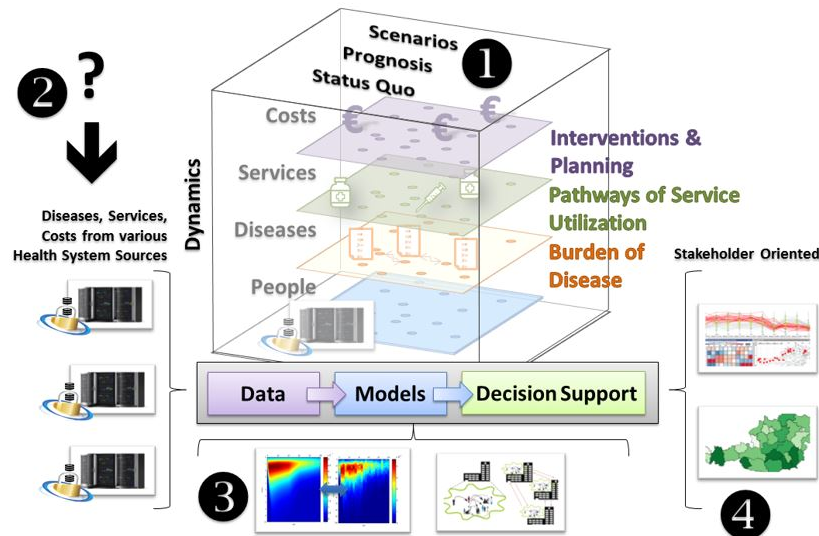
3.18 Methods for Integrated Simulation – from Data Acquisition to Decision Support

Niki Popper (*Technische Universität Wien, AT*)

License © Creative Commons BY 4.0 International license
© Niki Popper

Diversity and Heterogeneity of man-made systems increase rapidly and so do the costs spent for them. Measuring efficiency and effectiveness becomes more and more elaborate but is an urgent need. Development of new methods, models and technologies is needed to support analysing, planning and controlling. The quantity and quality of available data strongly increase and therefore facilitate the description and analysis of systems like health care. Bringing together necessary technologies is an enormous challenge.

Data-based Demographic models have to be combined with models for the spread of diseases. Dynamic modelling concepts must be parametrised with complex data sets from various sources. For system simulation, an important aspect is the possibility to implement changes inside the system, like interventions within the computer model, and analysing their effects. As a recent example see Covid-19 Modelling at TU Wien [1]).



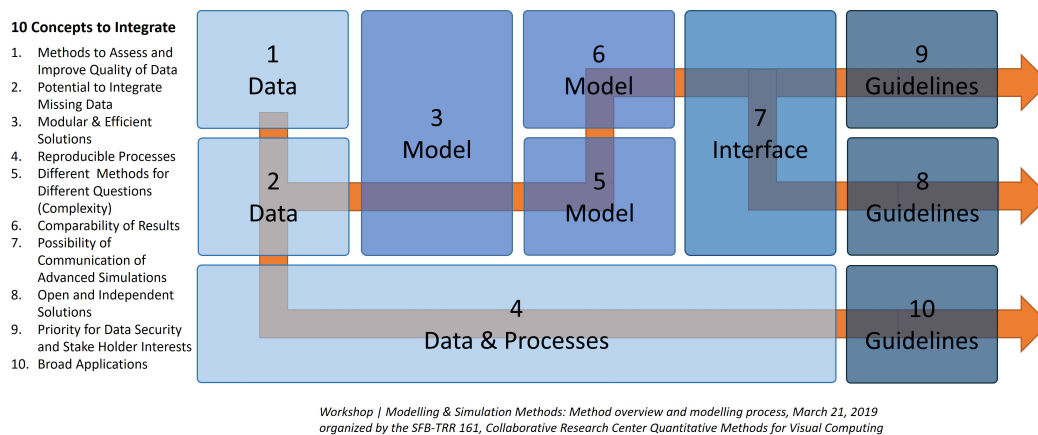
■ **Figure 6** Schematic overview of DEXHELPP infrastructure and process.

On basis of experiences of the Austrian DEXHELPP Competence Centre for Decision Support in Health Policy and Planning [2], which started in 2014 a concept was developed how large, interdisciplinary teams can handle these complex processes in the future and what are similarities and differences between health systems and other complex man-made DEXHELPP developed. DEXHELPP developed an innovative research infrastructure with (1) a flexible virtualised health system, (2) methods to cope with data, (3) an adaptive analysis and simulation methods pool and (4) stakeholder-oriented interfaces to enable researchers and other stakeholders to share data and methods for research and decision making. (see Figure 6).

“Ten Concepts to Integrate” were identified and first presented at Invited Talks at the University of Rostock [3] and University of Stuttgart [4]. The ten concepts are (1) Methods to Assess and Improve Quality of Data, (2) Potential to Integrate Missing Data, (3) Modular & Efficient Solutions, (4) Reproducible Processes, (5) Different Methods for Different Research Questions, (6) Comparability of Results, (7) Possibility of Communication of Advanced Simulations, (8) Open and Independent Solutions, (9) Priority for Data Security and Stake Holder Interests and (10) Broad Applications like Health System, Energy, Industry, Energy, Mobility and Infrastructure Planning and Usage (see Figure 7).

“Methods to Assess and Improve data Quality of Data” and “Potential to Integrate Missing Data” (summarized as “Data” in Figure 7) offer the possibility to find wrong data and correct them, ideally also during the simulation process. For this purpose, methods of interactive visualisation and statistics are used, among others, to preprocess data collected unilaterally, e.g. sensor data or reimbursement data or to link data that are unstructured or have different structures without direct linkage, e.g. [5].

To develop “Modular & Efficient Solutions” using “Different Methods for Different Research Questions” and maintain “Comparability of Results” (summarised as “Model” in Figure 7) includes sustainable, modular models that can be quickly adapted to new problems and concepts for comparing, combining and linking models (qualitatively and quantitatively) in order to demonstrate the benefits and limitations of the concrete model. As an example



■ **Figure 7** 10 Concepts to integrate into simulation processes.

the DEXHELPP Population Model GEPOC (Generic Population Concept) was used for Covid-19 Modelling in order to guarantee rapid and quality-assured implementation, e.g. [6] also to compare it in the ECDC Covid-19 Scenario Hub [7].

“Reproducible Processes” and the “Possibility of Communication of Advanced Simulations” (Blocks 4 and 7 in Figure 7) are essential to guarantee the credibility and usability of the models. We need tools to manage and share data (e.g. [8]) and models and to communicate not only the simulation results, but also the modelling process and model construction.

Last but not least, guidelines, standards that go beyond the concrete implementation are crucial (Blocks 8-10 in Figure 7). Here, the concepts of “Open and Independent Solutions”, “Priority for Data Security and Stake Holder Interests” and “Broad Applications” are crucial. The possibility of publication is limited, for example, by (justified) economic or data protection interests, which, however, leads to a lack of comparability of different models and thus jeopardises quality. This requires fundamental regulations such as those addressed in the General Data Protection Regulation and Data Governance Act. And clear and transparent processes are necessary for every project (even before the start of a simulation development) as well as the reuse of models is necessary to ensure quality and sustainability over time.

It is planned to publish the process description and examples as a White Paper within the European Umbrella Organisation of Simulation Societies EUROSIM [9].


References

- 1 Covid-19 Simulation in Austria, <http://www.dexhelpp.at/en/appliedprojects/tmp/covid-19/>, [Online; accessed 4.11.2022]
- 2 DEXHELPP, <http://www.dexhelpp.at/en/project-description/state-of-the-art/>, [Online; accessed 4.11.2022]
- 3 Abstract Kolloquiumsvortrag 14.6.2018 “Sharing Data, Methods, and Simulation Models – New Opportunities for Digital Health Care”, <https://www.informatik.uni-rostock.de/veranstaltungen/detailansicht-des-events/n/kolloquiumsvortrag-von-nikolas-popper-director-dexhelpp-wien-coordinator-centre-for-computational-complex-systems-tu-wien-46337/>, [Online; Accessed 4.11.2022]
- 4 Abstract Talk 25.3.2019 “Integrated Processes for Modelling & Simulation”, https://www.sfbtrr161.de/newsandpress/events_sfbtrr161/pastevents/, [Online; Accessed 4.11.2022]

- 5 N. Popper, B. Glock, F. Endel and G. Endel. Deterministic Record Linkage of Health Data as Preparatory Work in Modelling and Simulation-Use Case: Hospitalizations in Austria. *Proceedings of the 6th International Workshop on Innovative Simulation for Health Care*. pp. 44-49 (2017)
- 6 M. Bicher, C. Rippinger, C. Urach, D. Brunmeir, U. Siebert, N. Popper. Evaluation of Contact-Tracing Policies against the Spread of SARS-CoV-2 in Austria: An Agent-Based Simulation. *Medical Decision Making*, 41-8. pp. 1017-1032 (2021)
- 7 ECDC Covid-19 Scenario Hub, <https://covid19scenariohub.eu/>, [Online; accessed 4.11.2022]
- 8 N. Popper, M. Zechmeister, D. Brunmeir, C. Rippinger, N. Weibrecht, C. Urach, M. Bicher, G. Schneckenreither and A. Rauber. Synthetic Reproduction and Augmentation of COVID-19 Case Reporting Data by Agent-Based Simulation. *Data Science Journal*, 20(1). pp. 16ff (2021)
- 9 EUROSIM Technical Committee “Data Driven System Simulation”, <https://www.eurosim.info/tcs/tc-ddss/>, [Online; accessed 4.11.2022]

3.19 Using the Adaptable I/O System (ADIOS) for Effective and Sustainable Simulation Studies

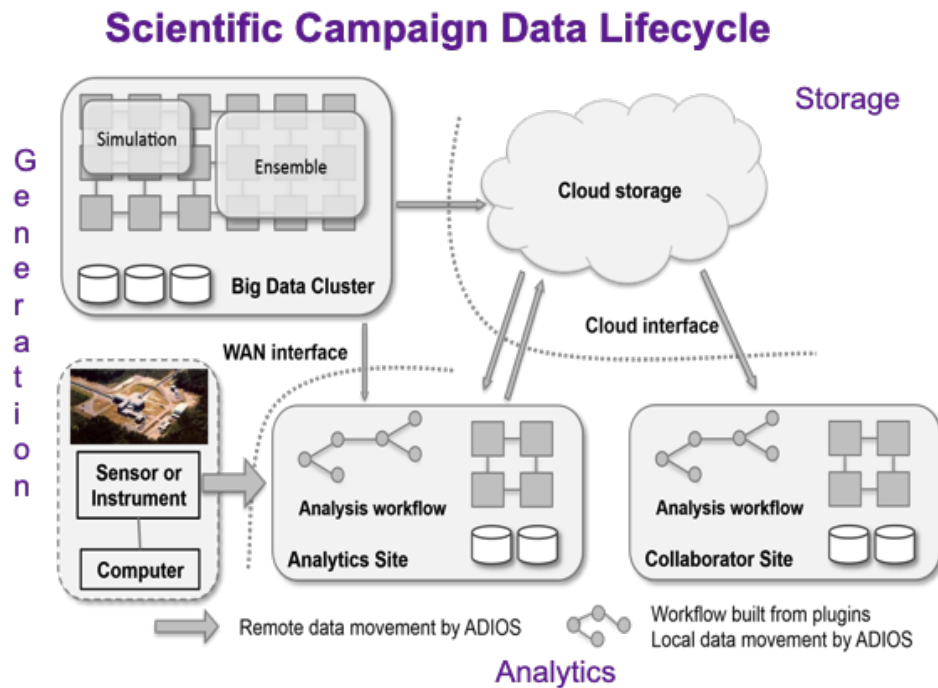
Caitlin Ross (Kitware – Clifton Park, US)

License  Creative Commons BY 4.0 International license
© Caitlin Ross

With the large amount of data written by large-scale scientific simulations, efficient parallel I/O is critical. However, leadership class systems have a number of factors that make developing applications difficult, from complex storage hierarchies to different parallel filesystems with different configurations. The Adaptable I/O System (ADIOS) [1] aims to simplify I/O, while achieving high performance on leadership class supercomputers and providing a simple API for developers. ADIOS is an open source C++ library that also provides bindings in C, Fortran, and Python. ADIOS is used by applications in various scientific fields, such as plasma physics, earth science, and aerospace engineering. Figure 8 shows the data lifecycle of a large-scale scientific campaign. ADIOS can be used to transfer data from simulations or sensors to persistent storage, couple simulations to exchange data, provide checkpoint/restart capabilities, or transfer data to other systems for analysis workflows.

ADIOS has a single, simple API to perform I/O, whether data is to be written to disk or streamed over a network. ADIOS provides this capability through different engines that can be selected at runtime. File engines include a self-describing format called binary packed (BP), as well as a HDF5 engine. There are a number of engines that stream data over a network, such as the Sustainable Staging Transport (SST) engine which can be used for exchanging data between loosely coupled jobs running on either the same compute cluster or over a wide area network. There is also an Inline engine that provides in process coupling of simulation and analysis/visualization for in situ analysis. Regardless of the engine chosen, it is selected in an XML configuration file and ingested by ADIOS at runtime, so data can be handled in different ways without needing to change the source code or using different API calls.

Recent work with ADIOS has involved integrating it into visualization tools such as ParaView [2] and the Visualization Toolkit (VTK) [3]. The Fides library was developed for this purpose and provides a data model schema which enables users to describe their data in a simple JSON format mapping ADIOS variables to the mesh and field characteristics of a



■ **Figure 8** The data generation, transfer, storage, and processing stages of a scientific campaign. ADIOS can be used for efficient transfer of data locally in an analysis workflow, as well as remotely over wide area network interfaces.

high-level data model. With ADIOS and Fides, simulations can use ParaView out of the box to visualize their data with post hoc, in situ, or in transit methods, without needing to write specialized adaptors or having a deep understanding of the VTK data model.

To the best of our knowledge, ADIOS is not currently being used in parallel discrete event simulations (PDES), but it could be useful in several ways for PDES. For instance, a recent US Department of Energy roundtable report [4] discusses the need for interoperability of multiple simulators that is portable and efficient on emerging platforms. ADIOS can be used for coupling simulations and provides portability and efficient parallel I/O.

The roundtable report also discusses incorporating machine learning (ML) frameworks into PDES engines/models. Some future work planned by the ADIOS team is to extend Fides metadata and provide data transformation utilities for ML data, which would enable ADIOS codes to integrate ML methods with their simulations. In addition to extending Fides metadata for ML data, we would also like to extend it for non-spatial data such as performance data and provide Python bindings. This would enable the use of other visualization and analysis tools besides scientific visualization tools like ParaView.

References

- 1 W. Godoy, N. Podhorszki, R. Wang, C. Atkins, G. Eisenhauer, J. Gu, P. Davis, J. Y. Choi, K. Germaschewski, K. Huck, A. Huebl, M. Kim, J. Kress, T. Kurc, Q. Liu, J. Logan, K. Mehta, G. Ostrouchov, M. Parashar, and S. Klasky, "ADIOS 2: The adaptable input output system. A framework for high-performance data management," *SoftwareX*, vol. 12, p. 100561, 07 2020.

- 2 U. Ayachit, The ParaView Guide. Kitware, Inc., 2015, 978-1-930934-30-6. [Online]. Available: <http://www.paraview.org>
- 3 Kitware Inc., “The Visualization Tool Kit (VTK),” <http://www.vtk.org/>, September 2021, [Online; accessed 08-October-2021].
- 4 K. Perumalla, M. Bremer, K. Brown, C. Chan, S. Eidenbenz, K. S. Hemmert, A. Hoisie, B. Newton, J. Nutaro, T. Opielstrup et al., “Computer science research needs for parallel discrete event simulation (PDES),” USDOE Office of Science (SC)(United States), Tech. Rep., 2022.

3.20 A Simulation Architecture to Study Diffusion Processes in Multiplex Networks


Cristina Ruiz-Martin (Carleton University – Ottawa, CA)

License © Creative Commons BY 4.0 International license
© Cristina Ruiz-Martin

Main reference Cristina Ruiz Martin, Gabriel A. Wainer, Adolfo López-Paredes: “Discrete-Event Modeling and Simulation of Diffusion Processes in Multiplex Networks”, *ACM Trans. Model. Comput. Simul.*, Vol. 31(1), pp. 6:1–6:32, 2021.

URL <https://doi.org/10.1145/3434490>

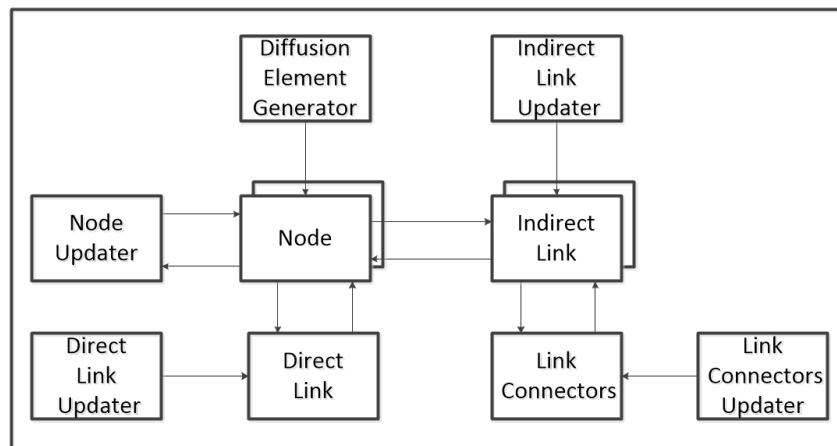
A variety of phenomena (such as the spread of diseases, pollution in rivers, etc.) can be studied as diffusion processes over networks (i.e. the diffusion of the phenomenon over a set of interconnected entities). There are different methods for studying diffusion processes, but most are based on various kinds of entities that are interconnected (networks). Two main approaches have been used to model the diffusion process over the network: Differential Equations (DE) and Agent-Based Modeling (ABM). The main advantage of nonlinear DE is that they can include a wide range of feedback effects (i.e. how the current value of a parameter of the system affects its future value, as in closed-loop systems). However, when they are used to study diffusion processes, one typically needs to aggregate nodes into fewer states or categories. Instead, ABM uses different attributes in each category, and different nodes in the same category may have different behavior. The network structure is clearly defined, and the behavior of each node is modeled individually at an increased computational cost. Likewise, neither of these two approaches provides well-established modeling and simulation (M&S) mechanisms for incorporating diffusion algorithms into multiplex dynamic networks and running simulations. The limitations of DE and ABM to studying diffusion processes pose the following question: how can we study diffusion processes in Multiplex Dynamic Networks to overcome the limitations of DE and ABM? Is there a framework that allows us to do this? To answer these questions, we introduced a method to study diffusion processes in multiplex dynamic networks, maintaining separations of concerns in all phases of modeling, implementation, and experimentation. The results include: a) an Architecture to study Diffusion Processes in Multiplex dynamic networks (ADPM); b) a systematic Process to define, implement and simulate diffusion processes over such networks. We use Network Theory formal specifications to define the topology of the diffusion process, Agent-Based Modeling (ABM) to define the behavior of the entities involved, and a formal specification of both for simulation modeling. This research uses the Discrete Event System Specification formalism (DEVS) to define the formal simulation model. The Architecture (and development process) to simulate Diffusion Processes in Multiplex dynamic networks (ADPM), is presented in Figure 9. The architecture is generic and can be used to study several types of diffusion processes.



■ **Figure 9** ADPM organization and workflow.

ADPM includes the Requirements and Assumptions Document of the problem; a Network model of the relations among components; an Agent-Based model of behavior; the Diffusion Abstract Model (DAM) (Figure 10), a formal representation based on the Network and Agent models; a Diffusion Computerized Model (DCM) of the DAM; the Simulation Logs and a Results Analysis Report.

The DAM is defined using a formal specification, DEVS, in our case. This solves some of the limitations of Network Theory, such as the lack of a formally verified simulator to simulate a diffusion process over the network. By including ABM, we can also define the behavior of both the nodes and links in the network. The formal DAM confers ABM rigor while separating modeling, verification, and experimentation. The main advantage of combining Network Theory, ABM, and DEVS is that we can use the most appropriate method to model the various aspects of the problem. Network Theory is well suited to model the relations between components; ABM is better suited to model behavior. DEVS provides a formal specification to define the whole model as components with hierarchical and modular specifications, which can be executed using well-established abstract simulation algorithms, which are proven to execute models correctly. This combination allows us to separate concerns and clearly differentiate each part of the problem, as well as separating models from simulation engines and experiments that are independent software entities.



■ **Figure 10** The DAM structure.

3.21 Model-based Software and Systems Engineering for Digital Twins

Bernhard Rumpe (RWTH Aachen, DE)

License © Creative Commons BY 4.0 International license
© Bernhard Rumpe

Main reference Manuela Dalibor, Malte Heithoff, Judith Michael, Lukas Netz, Jérôme Pfeiffer, Bernhard Rumpe, Simon Varga, Andreas Wortmann: “Generating customized low-code development platforms for digital twins”, *J. Comput. Lang.*, Vol. 70, p. 101117, 2022.

URL <https://doi.org/10.1016/j.cola.2022.101117>

Modeling is an important technique in many engineering disciplines. Modern modeling languages and tools allow developers to early concentrate on key aspects of the product and thus frontload quality assurance e.g. through simulation.

Equally important, explicit models of system requirements, technology independent function and architecture models, as well compact software and system component models enable reuse and variability management. Composition, tracing and refactoring assist evolutionary development and simulative quality assurance in a way that greatly reduces development cost for products of all kinds.

Using explicit models in appropriate modeling languages, like SysML or even domain specific languages, and an integrated, highly automated tool chain for construction, analysis and simulation is key to integrate all forms of system components. A holistic development approach needs a clear decomposition and decoupling and thus also well defined integration techniques.

Furthermore, models are a good basis for the construction of a digital twin, because a digital twin shares a lot of characteristics with a model. A digital twin is like a real twin: it is an active instance that interacts with the real system, allows to share real operative data, but also to simulate what the real twin would do and thus predict a systems behavior.

We examine the current state and problems of modeling for cyberphysical systems. In particular, we discuss how to make use of models in large development projects, where a set of heterogeneous models of different languages needs is developed and needs to fit together. A model based development process (both with UML/SysML as well as a domain specific modeling language) heavily relies on modeling core parts individually and composing those through generators to early and repeatedly cut simulative and productive code as well as a digital twin infrastructure from these models.

3.22 Data Farming: Better Decisions Via Inferential Big Data

Susan Sanchez (Naval Postgrad. School – Monterey, US)

License © Creative Commons BY 4.0 International license
© Susan Sanchez

Recently, the ready availability of “big data” has led to the adoption of data mining methods by organizations around the globe, as they seek to sift through massive volumes of data to find interesting patterns that are, in turn, transformed into actionable information. Yet a key drawback to the big data paradigm is that it relies on observational data, limiting the types of insights that can be gained.

The simulation world is different. Simulation models are integral to modern scientific research, national defense, industry and manufacturing, and in public policy debates. These models tend to be extremely complex, often with hundreds of thousands of potential factors (inputs or embedded parameters) and many sources of uncertainty. A “data farming” metaphor captures the notion of purposeful data generation from simulation models using efficient designed experiments.

Efficient design of experiments are required because it is literally impossible to examine even moderate numbers of factors by brute force. When the first supercomputer with petaflop performance (a quadrillion operations per second) was launched in 2008, the New York Times stated that machines such as this had “*the potential to fundamentally alter science and engineering*” by letting researchers “*ask questions and receive answers virtually interactively*” and “*perform experiments that would previously have been impractical*” [3]. Fourteen years later, the “Frontier” leads the world as the first exaflop machine [8], capable of over 10^{18} or a quintillion operations per second. But let us take a closer look at the practicality of a brute-force approach. Suppose a simulation has 100 factors, each factor has two levels (low and high) of interest, and we decide to look at all combinations of these 100 factors. A single replication of this brute-force experiment would take over 40 millenia on the Frontier, even if each of the 2^{100} simulation runs consisted of a single machine instruction! Efficient design of experiments can break this curse of dimensionality where expensive hardware cannot; experiments involving 100 factors can be completed in hours to weeks even for simulations with runtimes of minutes or hours [7].

With a data farming mindset, we can achieve tremendous leaps in the breadth, depth, and timeliness of the insights yielded by simulation. Large-scale experiments let us grow the simulation output efficiently and effectively. Modern statistical and visual analytic methods help us explore massive input spaces, uncover interesting features of complex simulation response surfaces, and explicitly identify cause-and-effect relationships [7, 2, 4]. Yet despite the current benefits offered by data farming, opportunities remain for advancing both the practice and research of simulation studies. A brief list of a few opportunities and challenges follows: see [1, 6, 5] for more.

From the practical standpoint, decision makers in many areas (climate change, economics, public health, to name a few) face increasingly complex problems where computational models are better than simple analytic models at capturing the complexity of the underlying system. At the same time, decision makers are increasingly comfortable with computerized and computer-based decisions based on observational big data, such as machine learning, artificial intelligence, and the use of digital twins. The rapid evolution of data science means that a greater number of simulation and non-simulation professionals are becoming more adept at scripting, modeling, graphical and statistical displays. Decision makers may, similarly, be less likely to shy away from using *model-driven* data to inform their decisions

if they are made aware of the potential benefits – particularly as they seek solutions to complex problems. By varying inputs in carefully chosen ways and exploring or building metamodels of the input/output relationships, we use simulation as an inferential (rather than descriptive) decision support tool. Rather than simply answering ‘what is?’ and ‘what if?’ questions, we can explore ‘what matters?, how?, and why?’

From a research perspective, the portfolio of methods suitable for addressing complex questions needs to be expanded. Further research is needed on multi-objective procedures; exploitation of parallel computing; adaptive sequential design methods; and methods that leverage the structure of inferential big data, rather than observational big data, for analysis and visualization tools. Regarding simulation optimization and other adaptive search techniques, it may be that we should be doing optimization on metamodels, rather than on the simulations themselves – and that we need automated ways of reoptimizing as these metamodels evolve over time. Adaptive sequential procedures are particularly important if a simulation study is viewed as an ongoing process, rather than a terminating one.

References

- 1 Elmegreen, B. E., S. M. Sanchez, and A. S. Szalay. 2014. The Future of Computerized Decision Making. In *Proceedings of the 2014 Winter Simulation Conference*, edited by A. Tolk et al., 943–949. Piscataway, New Jersey: IEEE.
- 2 Feldkamp, N., S. Bergmann, and S. Strassburger. 2020. Knowledge discovery in simulation data. *ACM Transactions on Modeling and Computer Simulation*, 30(4): Article 22, 1–25.
- 3 Markoff, J. 2008. Military Supercomputer Sets Record. *New York Times*, June 9.
- 4 Matković, K., D. Gračanin, and H. Hauser. 2018. Visual analytics for simulation ensembles. In *Proceedings of the 2018 Winter Simulation Conference*, edited by M. Rabe et al., 321–335. Piscataway, New Jersey: IEEE.
- 5 Sanchez, S. M. 2020. Data farming: Methods for the present, opportunities for the future. *ACM Transactions on Modeling and Computer Simulation*, 30(4): Article 22, 1–30.
- 6 Sanchez, S. M., and P. J. Sanchez. 2017. Better big data via data farming experiments. In *Advances in Modeling and Simulation: Seminal Research from 50 Years of Winter Simulation Conferences*, edited by A. Tolk, J. Fowler, G. Shao, and E. Yücesan, 159–179. Cham, Switzerland: Springer International Publishing.
- 7 Sanchez, S. M., P. J. Sanchez, and H. Wan. 2021. Work smarter, not harder: A tutorial on designing and conducting simulation experiments. In *Proceedings of the 2021 Winter Simulation Conference*, edited by S. Kim et al., 15 pages. Piscataway, New Jersey: IEEE.
- 8 TOP500 News Team 2022, May 30. ORNL’s Frontier first to break the exaflop ceiling. <https://www.top500.org/news/ornl-frontier-first-to-break-the-exaflop-ceiling/>.

3.23 Decision Making using Reinforcement Learning in Contested and Dynamic Environments

Claudia Szabo (University of Adelaide, AU)

License  Creative Commons BY 4.0 International license
© Claudia Szabo

Systems operating in military operations and crisis situations usually do so in contested and dynamic environments with poor and unreliable network conditions. Individual nodes within these systems usually have an incomplete, local and changing view of the system and its operating environment, and as such optimizing how nodes communicate in order to improve decision making is critical. Reinforcement learning approaches have been very

successful at solving problems in dynamic environments and in some cases when dealing with incomplete information. In this talk, we discuss the challenges of training and executing reinforcement learning approaches within such environments, in particular when considering communication middlewares.

3.24 Simulation-based Inference for Automatic Model Construction

Wen Jun Tan (Nanyang TU – Singapore, SG)

License © Creative Commons BY 4.0 International license
© Wen Jun Tan

Simulation-based inference aims to address the question of linking simulation models with empirical data by designing statistical inference procedures that can be applied to simulators [1]. Data assimilation is an inference method, in which the observed data are assimilated into the model to produce a time sequence of estimated system states [2]. Data assimilation was initially developed in the field of numerical weather prediction. However, simply inserting point-wise measurements into the dynamic weather models results in large instabilities in the simulation, rendering the forecasts meaningless. Hence, data assimilation is developed to initialize a model using observation data while making sure to maintain stability in the model by iteratively selecting the best system estimation.

The initial concept of data assimilation relies on the initialization of an existing model. However, changes in the observation data may also result from dynamic structural changes in the system. This means that the original simulation model built from the past system will be invalid; initialization of an outdated model will render the simulation results useless. Instead of assimilating the observation data into an existing model, we propose to build the model concurrently with the data assimilation procedure.

Considering a smart factory with Internet-of-Things (IoT) sensors to monitor physical events occurring in the factory. Process mining is frequently used to analyze operational processes based on event logs. In process mining, process discovery aims at constructing a process model as an abstract representation of an event log [3]. The goal is to build a model (e.g., a Petri net) that provides insight into the behavior captured in the log. A petri net is a class of discrete event dynamic system, which can be used to simulate discrete event systems, such as the manufacturing process in a factory.

Due to dynamic changes in the customers' orders, there will be frequent changes to the production operations in the factory. A new approach to combine data assimilation with process mining is proposed to accurately model the dynamic manufacturing processes. By monitoring the sequence of physical events, these events are captured into snapshots of the event lists for each event's arrival. Concurrently for each event, the real system performance is also measured, e.g., factory throughput. Process discovery will be performed on these snapshots to obtain the process models for each snapshot. Process enhancement will be used to enrich these process models by mining additional operational behaviours from the event logs, e.g., queuing or batching operations, to achieve a more realistic discrete event model of the factory. Each process model is simulated to obtain the system performance. Data assimilation will be performed by comparing the simulation performance with the real system performance and selecting the process model with the best estimated system states. By iterating discovering new process models and selecting the most accurate process model, the proposed approach is able to adapt to changes in the system structure while producing a calibrated simulation model at the same time.

References

- 1 Cranmer, K., Brehmer, J. & Louppe, G. The frontier of simulation-based inference. *Proceedings Of The National Academy Of Sciences*. **117**, 30055-30062 (2020)
- 2 Bouttier, F. & Courtier, P. Data assimilation concepts and methods March 1999. *Meteorological Training Course Lecture Series. ECMWF*. **718** pp. 59 (2002)
- 3 Van Dongen, B., Medeiros, A. & Wen, L. Process mining: Overview and outlook of petri net discovery algorithms. *Transactions On Petri Nets And Other Models Of Concurrency II*. pp. 225-242 (2009)

3.25 Models and Specifications within the Modeling and Simulation Life Cycle

Adelinde M. Uhrmacher (Universität Rostock, DE) and Claudia Szabo (University of Adelaide, AU)

License © Creative Commons BY 4.0 International license
© Adelinde M. Uhrmacher and Claudia Szabo

Modeling means organizing knowledge about a system of interest. In the talk, we will extend this citation, which has been attributed to Bernhard Zeigler, to the various artifacts of the modeling and simulation life cycle, as well as the life cycle itself. Thus, we will consider not only simulation models, but also models about simulation experiments, simulation methods, or entire simulation studies. With the subject of modeling the means of modeling varies as well, and includes approaches as diverse as domain-specific languages, formalisms, meta-models, ontologies, and logics. We will exemplify the interrelations between the subject of the model and the modeling approach. The second part of the talk is dedicated to the challenges that developing and applying model-based approaches face in modeling and simulation studies. Examples of such challenges including expressivity trade-offs, computational efficiency, reusability, accessibility, explainability, and evaluation are discussed.

3.26 A discrete-event approach for the study of sustainability in buildings

Gabriel A. Wainer (Carleton University – Ottawa, CA)

License © Creative Commons BY 4.0 International license
© Gabriel A. Wainer


Joint work of Vinu Subashini Rajus, Joseph Boi-Ukeme, Rishabh Sudhir Jiresal, Nicolas Arellano Risopatrón, Pedram Nojedehi, Gabriel A. Wainer, Liam O'Brien, Stephen Fai

Main reference Vinu Subashini Rajus, Joseph Boi-Ukeme, Rishabh Sudhir Jiresal, Nicolas Arellano Risopatrón, Pedram Nojedehi, Gabriel A. Wainer, Liam O'Brien, Stephen Fai: Measured Data Reliability for Building Performance and Maintenance. *IEEE Instrum. Meas. Mag.* 25(1): 55-61 (2022)

URL <https://ieeexplore.ieee.org/document/9693445>

Indoor factors like thermal, visual, acoustic, and chemical exposure all impact occupants' wellbeing. One of the main approaches to improve wellbeing is to have continuous monitoring of the building. Similarly, building design issues or technical flaws in the building software can have negative impact on occupants' health [1]. Modern building systems include sensors, actuators, and control devices including a tight coupling of hardware and software features. Advanced building systems collect data to improve building performance, operation, and maintenance. Building systems today use a variety of state-of-the-art equipment, such as

embedded hardware, wide-area connectivity, and software for decision making [2]. We use cloud computing to increase collaboration among various building stakeholders and to display real-time data from buildings systems for performance and maintenance analysis. Using cloud services, designers can store large Building Information Models (BIM), historical data or simulated data. BIM allows managing the digital representation of building components, their digital geometry, metadata, relationships, and the parametric rules to manipulate them [3]. We propose a system architecture and a workflow called SUSTAIN (Sensor-based Unified Simulation Techniques for Advanced In-building Networks). SUSTAIN integrates fault tolerance models, a text-mining framework and BIM to improve building system reliability. The overall software architecture is shown in Fig.11.




■ **Figure 11** SUSTAIN Architecture and Workflow.

The data collected from SAFE is integrated into BIM through Autodesk Forge, a cloud-based platform. Autodesk Forge allows to store and access various BIM and includes an integrated visualization component to visualize the real-time data through various web services. The BIM displays the sensor locations and ties the timeline data to its location. The aggregated data is represented as a shader to the room volume depending on the data parameters like temperature, humidity, or CO₂.


The VSIM BIM is a replica of the actual building with all the building elements (windows, doors, desks, etc.) and building systems (sensors, HVAC, etc.). We show a case study using temperature sensors with an operating range 0°C – 50°C (with an accuracy of $\pm 2\%$). We conduct simulation studies with different temperature readings within the measuring range were used. To evaluate the SAFE framework, faults were deliberately injected in both simulation and experimentation using different scenarios. The BIM of the lab was used to fabricate a maquette at a scale of 1:20 using laser cutting on acrylic sheets and adding control hardware (Fig. 12).

This physical model mimics the real-world setting of the research lab, which consists of six workstations, a fixed window, an operable blind, automated HVAC, and lighting controls. Fig. 13 shows the integration of the various sensors. Although the physical model does not behave as the real-world setting (e.g., the temperature in the maquette becomes steady faster than in the real building; however, when they are both in steady state, their behavior is similar), the objective here is different. SUSTAIN is focused on the development of the



■ **Figure 12** Physical prototype for experimentation [4].

building control software, the integration with 3D visualization, simulation and support in a cloud environment. The physical model permits conducting a variety of experiments useful during the software development cycle, and this can be done safely and without affecting the actual building operations.



■ **Figure 13** Integration of data with BIM using Autodesk Forge [5].

References

- 1 M. Arif, M. Kafatygiotou, A. Mazroei, A. Kaushik, and E. Elsarrag. *Impact of indoor environmental quality on occupant well-being and comfort: A review of the literature*. International Journal of Sustainable Built Environment, 5(1), pp.1-11, 2016.
- 2 V. Reppa, M.M. Polycarpou, and C. G. Polycarpou. *Sensor Fault Diagnosis*. Foundations and Trends in 2020 Spring Simulation Conference (SpringSim), (virtual) pp.1-12. IEEE, 2020.
- 3 S.A. Viktoros, M.K. Michael, and M.M Polycarpou. *Compact Fault Dictionaries for Efficient Sensor Fault Diagnosis in IoT-enabled CPSs*. In 2020 IEEE International Conference on Smart Internet of Things (SmartIoT), IEEE, 2020.
- 4 J. Boi-Ukeme, and G. Wainer. *A framework for the extension of DEVS with sensor fusion capabilities*. Systems and Control, 3(1-2), pp.1-248, 2016.
- 5 Autodesk Inc. *Dynamo*. <https://dynamobim.org/>. 2022.

3.27 Automatically Generating Simulation Experiments based on Provenance

Pia Wilsdorf (Universität Rostock, DE)

License © Creative Commons BY 4.0 International license
© Pia Wilsdorf

Main reference Pia Wilsdorf, Anja Wolpers, Jason Hilton, Fiete Haack, Adelinde M. Uhrmacher: “Automatic Reuse, Adaption, and Execution of Simulation Experiments via Provenance Patterns”, *ACM Trans. Model. Comput. Simul.*, Association for Computing Machinery, 2022.

URL <https://doi.org/10.1145/3564928>

Simulation experiments play a vital role during a simulation study, be it for calibration, validation, or exploration [1]. Many simulation experiments are used repeatedly throughout a simulation study or across related simulation studies. For instance, models are successively calibrated and validated after each major model refinement. Similarly, if a model is built by extension of a previous study, cross-validation experiments are carried out. And in the case where different models represent alternative hypotheses about a system, or the same model has been implemented for different modeling and simulation platforms, again slightly adapted simulation experiments are used to compare these model alternatives.

The Reuse and Adapt framework for Simulation Experiments (RASE) supports modelers in conducting these repeated experiments in a more systematic, effective, and efficient manner [2]. Based on predefined activity patterns and inference rules, key user activities are automatically detected (e.g., model refinements), whereupon suitable, previously executed simulation experiments (e.g., validation experiments) are selected, adapted, and executed in the context of the new simulation model, possibly reusing other information such as data or requirements. The framework is founded on the notion of provenance, i.e., information about the various entities (the research questions, simulation models, simulation experiments, simulation data, input data, requirements, qualitative model, assumptions, theories, etc.) and how they participated in or were generated by the diverse activities of modeling and simulation [3]. This information needs to be captured during a simulation study and stored as a directed acyclic graph based on the PROV-DM standard [4, 5] to be accessible to the pattern detection mechanism of the framework.

Currently, a new approach for capturing the provenance of simulation studies on-the-fly with minimal user involvement is being developed [6]. However, also the means for formalizing the context of a simulation model (provenance meta-data) need to be enhanced for provenance to be fully machine-interpretable. Here, we might take advantage of existing model-based approaches, such as signal temporal logic for specifying requirements [7], whereas, for the specification of other entities, such as assumptions or theories, no approaches yet exist.

References

- 1 Robert G. Sargent. An introduction to verification and validation of simulation models. In *2013 Winter Simulations Conference (WSC)*, pages 321–327, 2013.
- 2 Pia Wilsdorf, Anja Wolpers, Jason Hilton, Fiete Haack, and Adelinde M. Uhrmacher. Automatic reuse, adaption, and execution of simulation experiments via provenance patterns. *ACM Transactions on Modeling and Computer Simulation*, 2022.
- 3 Kai Budde, Jacob Smith, Pia Wilsdorf, Fiete Haack, and Adelinde M. Uhrmacher. Relating simulation studies by provenance—developing a family of wnt signaling models. *PLoS Computational Biology*, 17(8):1–27, 2021.
- 4 Andreas Ruschewski, Pia Wilsdorf, Marcus Dombrowsky, and Adelinde M. Uhrmacher. Capturing and reporting provenance information of simulation studies based on an artifact-

based workflow approach. In *Proceedings of the 2019 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, SIGSIM-PADS '19, page 185–196, 2019.

- 5 Khalid Belhajjame, Reza B'Far, James Cheney, Sam Coppens, Stephen Cresswell, Yolanda Gil, Paul Groth, Graham Klyne, Timothy Lebo, Jim McCusker, et al. PROV-DM: The PROV data model. *W3C Recommendation*, 14:15–16, 2013.
- 6 Anja Wolpers. *Lightweight Provenance Capturing for Simulation Studies*. PhD thesis, University of Rostock, 2022.
- 7 Laura Nenzi and Luca Bortolussi. Specifying and monitoring properties of stochastic spatio-temporal systems in signal temporal logic. *EAI Endorsed Transactions on Cloud Systems*, 1(4), 2015.

4 Working groups

4.1 Intelligent Modeling and Simulation Lifecycle

Wentong Cai (Nanyang TU – Singapore, SG), Philipp Andelfinger (Universität Rostock, DE), Luca Bortolussi (University of Trieste, IT), Christopher Carothers (Rensselaer Polytechnic Institute – Troy, US), Dong (Kevin) Jin (University of Arkansas – Fayetteville, US), Till Köster (Universität Rostock, DE), Michael Lees (University of Amsterdam, NL), Jason Liu (Florida International University – Miami, US), Margaret Loper (Georgia Institute of Technology – Atlanta, US), Alessandro Pellegrini (University of Rome “Tor Vergata”, IT), Wen Jun Tan (Nanyang TU – Singapore, SG), and Verena Wolf (Universität des Saarlandes – Saarbrücken, DE)

License © Creative Commons BY 4.0 International license
 © Wentong Cai, Philipp Andelfinger, Luca Bortolussi, Christopher Carothers, Dong (Kevin) Jin, Till Köster, Michael Lees, Jason Liu, Margaret Loper, Alessandro Pellegrini, Wen Jun Tan, and Verena Wolf

Modeling and simulation (M&S) has shown to be effective at conducting what-if analyses of complex scenarios. However, the current societal and technical challenges require building increasingly complex models and carrying out larger scale simulation experiments, which call for more efficient and intelligent approaches in all aspects of simulation studies, from model creation to model execution and experimentation.

In his state-of-the-art and open-challenges (STAROC) presentation on “High Performance Computing (HPC): Exploiting New Architectures”, Christopher Carothers introduced an array of new emerging hardware that supports effective and sustainable simulation studies, including ultra-dense chips, fast analog computing, wafer scale processor, low precision computing, and brain inspired chips. He went on to discuss the demand for Artificial Intelligence (AI) / Machine Learning (ML) and how it is driving the development of hardware and applications in the HPC sector. The trend of growing system complexity is also accelerating the development of software tools and hardware platforms specifically developed for AI/ML [1]. In his presentation, Carothers posed the following questions:

- Can AI/ML help M&S be more sustainable?
- Can M&S help improve AI/ML predictions?

To enable efficient utilization of the HPC resources, M&S needs to harness these new emerging computing hardware and platforms developed for AI/ML. In addition, AI/ML also opens up new and interesting opportunities to enhance and enrich M&S [2].

Initially, there were two topics identified for the working group: i) AI/ML + Simulation, and ii) Enabling Models to Run Efficiently on Heterogeneous Hardware. The first topic involves the integration of AI/ML techniques with M&S; and the second topic focuses on the efficient utilization of heterogeneous hardware to execute simulation models. As seen from Carothers' presentation, the current development of HPC systems and applications are "riding the AI/ML wave", i.e. the new emerging hardware developed for AI/ML. As a result, the working group decided to change the questions to incorporate the notion of the AI/ML wave: i) how can the emerging hardware be efficiently utilized for M&S? and ii) can M&S also ride on the AI/ML wave to make it more efficient and sustainable?

In the last decade, research challenges on exploiting emerging computing platforms and using heterogeneous hardware to accelerate the execution of discrete event simulations have been addressed in several research papers:

- 2015 Paper on "Grad Challenges for Modelling and Simulation" [3],
- 2016 Richard Fujimoto's STAROC paper on "Parallel and Distributed Simulation (PADS)" [4],
- 2016 NSF Workshop on "Future Research in Modelling and Simulation" [5],
- 2019 Survey paper on "Agent-based Simulation Using Hardware Accelerators" [6], and
- 2022 DOE Roundtable report on "Computer Science Research Needs for Parallel Discrete Event Simulation (PDES)" [7].

Particularly, Perumalla et al presented a computer science-oriented view of research challenges in PDES and identified priority research opportunities in advanced computing for PDES [7].

Given this prior research, the working group decided to focus the discussion on a merged topic: Intelligent Modeling and Simulation Lifecycle. The group believed the STAROC paper should focus on how modelling and simulation can benefit from recent advances in AI/ML techniques as well as the emerging powerful and pervasive hardware, computing paradigms and systems. In the proposed paper, we will examine the existing and emerging hardware/systems and AI/ML techniques in the context of modeling and simulation lifecycle (e.g., model creation, calibration, execution, and experimentation). We will also identify the major challenges and opportunities faced in M&S applications and outline important research directions for intelligent modeling and simulation, to improve speed, accuracy, and capabilities, by taking advantage of AI/ML and computing paradigms.

After intense discussion and brainstorming, the working group developed the following structure for the proposed STAROC paper:

- The first section of the paper will provide the motivations and introduce the main thesis of this article. We will focus on how advanced AI/ML and existing/emerging computing hardware/platforms will improve M&S at different stages of the M&S lifecycle.
- The second section "M&S Lifecycle, Challenges, and Demands" will provide an introduction to the M&S lifecycle: creation, calibration, execution, simulation experiments. These stages shall be defined clearly with respect to the traditional M&S pipeline: input modeling, concept model & validation, computational model & verification, experiment design, model execution, and output analysis. We will identify major M&S applications, including transportation and mobility, power grid, networking and cybersecurity, HPC, and epidemic modeling, in order to identify challenges and demands at different M&S stages that can be fulfilled by AI/ML and new emerging computing hardware and platforms.

- The third section “Emerging AI/ML-Assisted Approaches for M&S” will review the existing prominent AI/ML techniques and methodologies that have been applied for M&S. This will be achieved by categorizing these approaches to the various stages of the M&S lifecycle, and thereby identifying the gaps that can be explored in future research efforts.
- The fourth section “Heterogeneous and Emerging Computing Platforms for M&S” will categorize the diverse computing systems, platforms and emerging computing hardware that can support the M&S mission. First, we will introduce the existing computing platforms, e.g., HPC, cloud, edge, accelerators (GPUs, FPGAs). Next, we will explore the newer emerging hardware that are specially developed for AI/ML computing, e.g., neuromorphic devices, low-precision AI chips. At the same time, we will also examine the different computational paradigms: serverless computing, micro-services, virtualization, containerization, quantum computing or analog computing.
- The fifth section “Research Challenges and Roadmap for Intelligent M&S Lifecycle” will identify specific research directions and outline a roadmap for applying AI/ML and emerging hardware/computing platforms for the M&S lifecycle. We will create a table to illustrate the gaps or challenges at each stage of the M&S lifecycle and list the corresponding techniques and research directions to be taken to address these challenges. In the subsequent subsections, we will elaborate on these techniques and research directions.
- Finally, we will conclude the article by pointing out that we are focusing on using AI/ML for M&S. However, there are other research directions, such as applying M&S for improving AL/ML (such as explainable AI).

Tentatively, the working group plans to complete and submit the paper by mid-Jan 2023.

References

- 1 Daniel A Reed and Jack Dongarra. Exascale computing and big data. *Communications Of The ACM*. 58(7): 56-68, July 2015.
- 2 Alexander Lavin, Hector Zenil, Brooks Paige, David Krakauer, Justin Gottschlich, Tim Mattson, Anima Anandkumar, et al. Simulation intelligence: Towards a new generation of scientific methods. *ArXiv Preprint ArXiv:2112.03235*. 2021.
- 3 Simon J. E. Taylor, Azam Khan, Katherine L. Morse, Andreas Tolk, Levent Yilmaz, Justyna Zander, and Pieter J. Mosterman. Grand challenges for modeling and simulation: simulation everywhere—from cyberinfrastructure to clouds to citizens. *Simulation: Transactions of the Society for Modeling and Simulation International* . 91(7):648-665, 2015.
- 4 Richard Fujimoto. Research challenges in parallel and distributed simulation. *ACM Transactions On Modeling And Computer Simulation (TOMACS)*. 26(4):22:1 -22:29, May 2016.
- 5 Christopher Carothers, Alois Ferscha, Richard M. Fujimoto, David Jefferson, Margaret Loper, Madhav Marathe, P. Mosterman, Simon Taylor, and H. Vakilzadian. Computational challenges in modeling and simulation, in *Research Challenges In Modeling And Simulation For Engineering Complex Systems*. pp. 45-74, 2017.
- 6 Jiajian Xiao, Philipp Andelfinger, David Eckhoff, Wentong Cai, and Alois Knoll. A survey on agent-based simulation using hardware accelerators. *ACM Computing Surveys (CSUR)*. 51(6):131:1 – 131:35, Nov 2019.
- 7 Kalyan Perumalla, Maximilian Bremer, Kevin Brown, Cy Chan, Stephan Eidenbenz, Scott K. Hemmert, Adolfo Hoisie, Benjamin Newton, James Nutra, Tomas Ooppelstrup, Robert Ross, Markus Schordan, Nathan Urban. Computer Science Research Needs for Parallel Discrete Event Simulation (PDES US Department of Energy, Advanced Scientific Computing Research, Roundtable Report, May 2022. (DOR: 10.2172/1855247)

4.2 Policy by simulation: seeing is believing for interactive model co-creation and effective intervention

Rodrigo Castro (University of Buenos Aires, AR), Joachim Denil (University of Antwerp, BE), Jérôme Feret (ENS – Paris, FR), Kresimir Matkovic (VRVis – Wien, AT), Niki Popper (Technische Universität Wien, AT), Susan Sanchez (Naval Postgrad. School – Monterey, US), and Peter Slood (University of Amsterdam, NL)

License © Creative Commons BY 4.0 International license
© Rodrigo Castro, Joachim Denil, Jérôme Feret, Kresimir Matkovic, Niki Popper, Susan Sanchez, and Peter Slood

Motivation

We understand that policy makers are currently much more willing (than they were 10 years ago, for example) to participate in simulation-assisted planning and decision-making processes. We must therefore be well prepared so we are able to seize the opportunity.

Simulation is particularly useful in decision support when bifurcations can occur, i.e., when a system under study takes qualitatively different courses depending on the input parameters or random events as the system dynamics evolve, a distinctive feature of complex systems. Interventions represent a special example that by definition leads to discrete changes in the dynamics that may not be represented with classical methods alone [13].

Two key aspects of policy by simulation are that the systems being addressed are inevitably complex (they include overlapping facets such as social, physical, biological, cybernetic, etc.) and also that the interventions being applied in the real world must ultimately and inevitably be understood and agreed upon by humans.

As there are usually several stakeholders from different domains, it is essential to provide means which will efficiently support specifications of interventions as well as communication of results which do not rely solely on mathematical-computational models.

In this context, dynamic and reactive visualization emerges as a key element for human understanding of interventions in complex systems.

While the technologies of modeling, simulation and visualization (M&S&V) have made enormous progress during the past several decades, several challenges remain. We need to provide efficient and sustainable mechanisms that engage interdisciplinary teams of decision makers with simulation-assisted processes.

Meanwhile, recent work highlights missing links between user tasks in visualization taxonomies (e.g., sensemaking) and the high-level task of decision-making (Dimara and Stasko [3]). They identified a lack of interdisciplinary approaches as one of the key causes of this mismatch. The latter suggests that a joint approach of modeling, simulation and visualization experts stands as a valid approach to reduce this gap.

We therefore envision “policy by simulation”, a comprehensive co-creation framework that exploits state-of-the-art modeling, simulation, experimentation, visualization and AI to engage policy makers and stakeholders from multiple domains concerned with a common system under study. The goal is an efficient and sustainable framework of interoperable M&S&V-based tools that allow stakeholders to model potential interventions into simulation worlds and deliver timely insights about complex systems for the purpose of planning efficient and effective real-world interventions.

These envisaged capabilities are shared by several initiatives such as the POLDER Simulation Center (POLicy Decision-support and Evidence-based Reasoning) [15], the DEXHELPP Project (DEcision support for HEalth Policy and Planning) [16], the SEED Center (Simulation Experiments and Efficient Designs) for Data Farming [17], and the Simulation and Immersive Visualization Lab [18].

We first present the framework we envision, and then describe some of the open research challenges that must be addressed to fully realize its potential.

A Policy by Simulation Framework

To implement an integrated decision-making process that covers both simulation aspects and stakeholder needs for modeling feasible interventions, we envision an iterative three-level structure. The ultimate aim is to unveil potentials for action and to outline a basic structure as well as minimum requirements without claiming to be exhaustive.

- **Conceptual level.** Generate problem statements, identify mental models, create a narrative definition of goals, identify areas of conflict, and discuss attitudes toward risk. Identify the timeline for making decisions about whether and how to intervene. This will be revisited as external circumstances change, insights arise, or new stakeholders join the team.
- **Logical level.** Determine data availability or specify data model specifications and assumptions. Describe correlations, causations, causal loops, anchor points, levers, and potential interventions. Discuss the key response measures and aggregation levels. Prioritize features to include in the next round of executable models. Delegate responsibility for creating embedded submodels or suitable linkages between stand-alone submodels and the overarching model.
- **Assessment level.** Create, refine, and explore executable models. This can include models at different levels of granularity, from the coarse grain level to a highly detailed level. Types of models could include systems dynamics models (SD), partial differential equations (PDE), agent based models (ABM), discrete event simulation (DES), complex networks (CN), and more. Policy stakeholders will not be involved in the actual model implementation, but are vital for scoping the experimental region(s) of interest. Efficient experiments can be used to grow data that can be jointly and visually explored and assessed, spawning further experimentation as insights are gained and new questions arise.

Visualization techniques shall encompass all stages, providing visual coherence and continuity across all levels of information, connecting input data, simulation models, and simulation results across levels as the study evolves. Visualization is a key enabler for model-based tradeoff discussions among stakeholders with competing priorities and world views.

On the one hand, these visualization techniques should cover all the necessary needs in each stage. On the other hand, they should provide sufficient freedom, especially in early phases, so as not to impose restrictions that could lead to mapping errors and ultimately reduce the quality of the process and the intervention decisions.

Research Challenges: Bridging the Credibility Gap

The overall policy by simulation framework can be implemented using existing tools and methods. However, its full potential will not be realized without additional research in a variety of areas.

Credibility Gap: Even though we can create sophisticated simulation models for policy intervention, we often fall short in communicating their structure and results in transparent and understandable ways. For example, causal loop or process flow diagrams are very understandable, but more detailed models may be difficult for some stakeholders to interpret.

Past experience with model-based approaches (or lack thereof) may also affect stakeholder credibility in different ways.

A priori: The question of which aspects are important depends not only on the subjective view of the stakeholders, but also on the non-assessability of dynamic effects. These are to be investigated either by data analysis or by methods of causal analysis. In this case, the estimation of potentials is often low.

A posteriori: Real-world validation is difficult in the sense of the prevention paradox, interventions that are set cannot be compared with interventions that are not set—at least in real-world settings. Also, for some modeling situations and paradigms the system boundaries must be narrowly chosen in order to be able to create meaningful models.

We describe four threads we feel are key for bridging the credibility gap: model co-creation, visualization, flexibility, and efficiency.

Co-creation

There is a lack of robust methodologies for co-creating a palette of simulation models [12, 6] that can refocus across several levels of model detail, from conceptual high-level models down to very detailed model components. Similar methods are also needed for co-creating and co-exploring disparate models that capture different aspects of the system of interest, regardless of their level of detail, as part of a “many model thinking” paradigm.

Approaches that can support model co-creation, but require further research, include:

- **Immersion.** As simulation professionals, we have to show that we understand the context, the stakeholders and their interests, and how the simulation is embedded in the whole environment. System and simulation boundaries have to be discussed. These can be changed but have to be communicated in transparent ways. Similarly, stakeholders themselves must become immersed in the co-creation process. We need methods for eliciting their inputs and assumptions, engaging them with other stakeholders during the model development process, and engendering their trust in using a model-based framework to gain insights.
- **Formal methods.** Formal methods [4] for model relation and model reduction can help increase stakeholder confidence. They can relate models at different levels of abstraction and automatically derive reduced models while ensuring a formal relationship between the initial and the reduced ones. This can be done either analytically (which eases their application to a wide collection of models) or by language specific approaches (based on the structural properties of the description language). So derived model reduction may be exact [2] (as a reduction of the dimensionality of the model), or approximated [1] (including explicit bounds for numerical errors). Analytic approaches may require unwarranted assumptions and struggle to scale to the complexity of large models. Event graphs are modeling representations that can be implemented in a variety of languages, but are less widely used [11]. The development of language-based approaches is more time consuming since a given tool will work only for a dedicated modeling language. Further research is needed to address these challenges.

Visualization

Although visualization technologies and visual analytics techniques are established as exploratory methods and have been applied to numerous decision making and simulation problems, visualization and interaction with models at different levels of detail in an integrated framework remains an unsolved challenge. Further research on visualization techniques

that leverage the structure of data from designed simulation experiments is also needed. Comparison of different scenarios [5], provenance of information and tracking of interactions [7], efficient visualization for a plethora of stakeholders from different domains [3], all represent significant challenges for visualization research. Finally, as the number of stakeholders increases, there will be more and more scenarios to evaluate. Visualizing many scenarios represents a scalability challenge for visualization [8]. Limited screen space often does not allow to show all possibilities at once. At the same time important information has to be shown.

Flexibility

As demands and goals change over time, we need to be able to zoom in and out in the simulation scenarios across different levels of model detail. We also need the ability to focus on different model components to reveal the worldviews of different stakeholders.

- **Integration.** Integration for Exploration: We must build in flexibility in the form of software hooks from the beginning of the co-creation process. This includes flexibility in setting or modifying inputs and linking them to outputs, so that causal relationships within the model setting can be readily explored and identified. Input links should facilitate switching between real-world data (if available), diverse input data models, or changes driven by a mixture of approaches including interactive visual analytics, single-stage or adaptive experimental designs, or AI/ML-guided searches.
- **Flexible timing.** We need to provide partial insights in a timely way. Since this may involve assessments when the executable models are all at very high levels, input data are limited or unavailable, or only portions of the modeling are complete, best practices should be developed for effectively engaging stakeholders so the overall process provides value. How and when to revisit insights obtained early in the co-creation process, or how to determine the appropriate level for a particular intervention opportunity, remain open challenges.

Efficiency

We take a broad view of efficiency because the overarching goal of policy by simulation is that the stakeholders gain insights about effective interventions in a timely manner. Clearly, modeling efficiency and simulation run-time efficiency play their parts, and are often the focus of benchmark studies. We need metrics that capture the entire path to effective intervention by including person-hours (for modeling, analysis, visualization, documentation, discussions, etc.) as well as the computer cycles needed to arrive at insights. Harder (but extremely important) aspects to capture are the flexibility provided, the quality of the insights obtained, and the effectiveness of the intervention.

- **Sustainability.** There is an unmet need for automatic mechanisms that provide robust reusability of model libraries and experiments so that simulation projects can be properly evolved and adapted across longer periods of time. This includes accommodations for legacy models that continue to be used, as well as sunset clauses to handle models that should be reimplemented or retired as new modeling paradigms or languages are developed. Methods that reduce maintenance expenses for these “living models” are also of interest.
- **Model selection.** As a result of the zoom in/out requirement, we may need to use different methods for detailed modeling producing different outcomes. This imposes

different requirements for the underlying numerical and computational implementations. Over time, “best modeling approaches” might be revealed for certain problem domains. Such guidance would improve the modeling efficiency and so improve the timeliness of the overall process. This is another area that merits further research.

The above presented four threads –co-creation, visualization, flexibility, and efficiency– are interrelated. Even minor advances in any area may lead to dramatic improvements in the process as a whole.

Guiding principles

It is well known that in policy making, choosing not to intervene can be as decisive as enacting a particular intervention. Moreover, there are often windows of opportunity that require making the best possible decision in a given time frame, weighing the best sources of (always imperfect) information available at the time. We believe that the quality of such information can be greatly improved by following processes such as those presented herein, relying on simulation models and visualization techniques for intervention planning.

At early stages it is key to determine what kind of interventions could be feasibly made, identifying the lever points for controllable interventions in the complex system under study. This is a result of a discursive process with multiple stakeholders involved, including those with legitimately contradictory perspectives and worldviews of the overarching problem at hand.

The co-creation of simulation models implies including all necessary subsystems, variables, parameters and data or data models that capture the valid concerns and interests of heterogeneous stakeholders involved in the modeling process.

One important measure of the quality of the resulting models (or family of models) is the degree at which they allow the stakeholders to exercise relevant what-if analysis in an exploratory process. Emergent behavior is a key feature of complex systems [14] (basically anything unexpected can happen) therefore “potentially feasible paths of action” must be explored. While experimenting, the framework shall keep track of all interventions in the models along the exploratory process. To that aim, different alternatives are possible. If the computational time for single runs is long, checkpoints could be recorded enabling the team to “go back and rebranch” if a given path of interventions is rendered inefficient/undesirable. If storage requirements for checkpoints is prohibitive, random number seeds could be recorded enabling the team to initially output limited data (e.g., end-of-run or other summary information) but “go back and restart” to output detailed information or additional performance measures. The exploratory process itself becomes the result of a discussion with a group of “humans in the loop”. Varied data widgets should provide, simultaneously, multiple different views on a same system, both for simulation results and for the structure of model components. Yet, in parallel to human-based analysis, AI/ML-driven algorithms can automatically explore the parametric space (intelligent parameter sweeping) to prune and rule out model configurations that lead to undesirable results (e.g., violate restrictions on state variables). Other automatic analyses that can enrich the knowledge about the underlying model are parameter sensitivity analyses, which can provide a better understanding of the relative potential impacts of lever points. These sensitivity levels can be visually mapped into the model structure to guide the experience in a more efficient way. Large-scale designed experiments (“data farming”) can be used to identify the most impactful model inputs and interactions, uncover other interesting features about the I/O behavior, and assess how uncertainties propagate through a model or hierarchy of models [9, 10]. Such insights can help steer the interactive exploration toward relevant paths.

Related needs

The implementation of the proposed process has implications for other areas of simulation technology -not addressed here- which were dealt with in the Dagstuhl Seminar “Computer Science Methods for Effective and Sustainable Simulation Studies” in October 2022:

- Reduce simulation ecological footprint while providing increased simulation performance
- Provide efficient automatic documentation of models along with their ranges of validity
- Track the provenance of models, simulation experiments, and the associated input and output data
- Create better ways (standards?) to store and reuse model components across several projects with some common parts of their systems under study

Summary

We believe that the “policy by simulation” framework we propose can – and should – be used for identifying and implementing real-world interventions to address many complex problems currently faced by policy makers. If done effectively, this approach will improve the quality, timeliness, and effectiveness of the intervention decisions, will also enhance the insights stakeholders gain from the simulation models, and so improve the credibility of the collections of models used. On a broader level, as policy makers and other stakeholders become more familiar with this framework, we will have more opportunities to create a mindset that values and celebrates the exploration of interventions in the virtual world, rather than bemoaning the inability for a posteriori validation in the real world. We envision policy by simulation as an integrated, interactive, ongoing process, rather than a one-time product. Similarly, the framework itself will evolve as new tools and methods are developed to address the many research challenges identified above. Also, as we CS and simulation professionals, policy makers, and other stakeholders co-create the models and co-learn from analytic and visual exploration of their behavior, we expect other challenges and research opportunities to arise.

References

- 1 Andreea Beica, Jérôme Feret, and Tatjana Petrov. Tropical abstraction of biochemical reaction networks with guarantees. *Electronic Notes in Theoretical Computer Science*, 350:3–32, 2020. Proceedings of SASB 2018, the Ninth International Workshop on Static Analysis and Systems Biology, Freiburg, Germany – August 28th, 2018.
- 2 Vincent Danos, Jérôme Feret, Walter Fontana, Russell Harmer, and Jean Krivine. Abstracting the differential semantics of rule-based models: Exact and automated model reduction. In *2010 25th Annual IEEE Symposium on Logic in Computer Science*, pages 362–381, 2010.
- 3 Evanthia Dimara and John Stasko. A critical reflection on visualization research: Where do decision making tasks hide? *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1128–1138, 2022.
- 4 Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In Robert M. Graham, Michael A. Harrison, and Ravi Sethi, editors, *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages, Los Angeles, California, USA, January 1977*, pages 238–252. ACM, 1977.
- 5 Michael Gleicher, Danielle Albers, Rick Walker, Ilir Jusufi, Charles D. Hansen, and Jonathan C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, Oct 2011.
- 6 Scott E. Page. *The Model Thinker: What You Need to Know to Make Data Work for You*. Basic Books, Inc., USA, 2018.

- 7 Eric D. Ragan, Alex Endert, Jibonananda Sanyal, and Jian Chen. Characterizing provenance in visualization and data analysis: An organizational framework of provenance types and purposes. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):31–40, 2016.
- 8 Gaëlle Richer, Alexis Pister, Moataz Abdelaal, Jean-Daniel Fekete, Michael Sedlmair, and Daniel Weiskopf. Scalability in visualization, arxiv.org/abs/2210.06562, 2022.
- 9 Susan M. Sanchez. Data farming: Methods for the present, opportunities for the future. *ACM Transactions on Modeling and Computer Simulation*, 30(4): Article 22.
- 10 Susan M. Sanchez, Paul J. Sanchez, Hong Wan. Work smarter, not harder: A tutorial on designing and conducting simulation experiments. In *Proceedings of the 2021 Winter Simulation Conference*, 15 pages, 2021.
- 11 Lee W. Schruben. Simulation modeling with event graphs. *Communications of the ACM*, 28(11):957-963, 2013.
- 12 Bernard P. Zeigler. Why should we develop simulation models in pairs? In *2017 Winter Simulation Conference (WSC)*, pages 2–2, 2017.
- 13 Rodrigo Castro. Chapter 24 – Open Research Problems: Systems Dynamics, Complex Systems. *Theory Of Modeling And Simulation (Third Edition)*. pp. 641-658 (2019)
- 14 Foguelman, D., Henning, P., Uhrmacher, A. and Castro, R. EB-DEVS: A formal framework for modeling and simulation of emergent behavior in dynamic complex systems. *Journal Of Computational Science*. **53** pp. 101387 (2021)
- 15 POLDER Simulation Center (POLicy Decision-support and Evidence-based Reasoning) <https://polder.center/>
- 16 DEXHELPP Project (DEcision support for HEalth Policy and Planning) <http://www.dexhelpp.at/>
- 17 SEED Center (Simulation Experiments and Efficient Designs) for Data Farming <https://harvest.nps.edu/>
- 18 Simulation and Immersive Visualization Lab <https://modsimu.exp.dc.uba.ar/>

4.3 Context, composition, automation and communication: towards sustainable simulation studies

Adelinde M. Uhrmacher (Universität Rostock, DE), Peter Frazier (Cornell University – Ithaca, US), Reiner Hähnle (TU Darmstadt, DE), Franziska Klügl (University of Örebro, SE), Fabian Lorig (Malmö University, SE), Bertram Ludäscher (University of Illinois at Urbana-Champaign, US), Laura Nenzi (University of Trieste, IT), Cristina Ruiz-Martin (Carleton University – Ottawa, CA), Bernhard Rumpe (RWTH Aachen, DE), Claudia Szabo (University of Adelaide, AU), Gabriel A. Wainer (Carleton University – Ottawa, CA), and Pia Wilsdorf (Universität Rostock, DE)

License © Creative Commons BY 4.0 International license

© Adelinde M. Uhrmacher, Peter Frazier, Reiner Hähnle, Franziska Klügl, Fabian Lorig, Bertram Ludäscher, Laura Nenzi, Cristina Ruiz-Martin, Bernhard Rumpe, Claudia Szabo, Gabriel A. Wainer, and Pia Wilsdorf

Motivation. Simulation has become a sine qua non in many areas. Particularly, the COVID-19 crisis has revealed the importance of simulation studies [6], but also limitations in terms of how fast we can develop useful models, how to interpret and build on results, and how to communicate the results to decision-makers. Therefore, new and better support for conducting simulation studies is needed, in particular

- to help simulation analysts who conduct or use simulation studies in building on their and others' results, improving the quality of their analyses, and making it easier for them to correctly interpret and reuse results.
- to facilitate the appropriate use of simulation by domain experts and decision-makers (e.g., policymakers in government, industry leaders), as a way to improve the quality of the decisions that they make.

Our goals. For comprehensive support of simulation studies, model-based approaches, in terms of formalisms, language-based approaches, logic, and meta-models play a central role, as they allow to make knowledge about a system explicit, computationally accessible, and interpretable. To move ahead, it is necessary to analyze possibilities already offered by current approaches and to identify challenges for future methodological research in model-based approaches to:

- ensure that simulation studies come with context. Context is important for helping technical experts correctly interpret results and for ensuring the quality of analyses based on the results of others. It is also important to correctly and confidently explain simulation results to decision-makers.
- improve composition and model re-use. Model re-use avoids building models from scratch. It thus saves time and, in addition, may improve analysis quality. Saving time may also increase access to simulation, increasing the number of decisions that it supports.
- support simulation automation. This would again save analyst time and may contribute to the analysis quality since automation facilitates the application of methods and thus may reduce errors in their application and, even, broaden the scope of analysis done with simulation models.
- facilitate communication both between simulation analysts and domain experts and between simulation analysts and decision-makers. We see communication as one of several limiting factors in the use of simulation for decisions. Better communication would also reduce the time required to achieve an impact.

A systematic application of model-based approaches for modeling and simulation should also allow us to harvest synergies with other areas of computer science such as language design, human-computer interaction, high-performance computing, visualization, or machine learning more effectively.

Starting point – the simulation study. Simulation studies that are aimed at developing and applying simulation models are intricate processes that combine different activities, such as model building, model refinement, model composition, calibration, analysis, and validation, and involve different types of entities, such as theories or data, and possibly products from other simulation studies. Conceptual models, requirements, assumptions, simulation models, qualitative models, simulation experiment specifications, and simulation data belong to the primary products of a simulation study, although not all of this information is explicitly represented (or documented). These ingredients can be distilled from work on modeling and simulation life cycles [22] and documentation guidelines for simulation studies [7]. For several of those, different modeling approaches have already been developed and are being applied (see Table 1).

In Table 1, simulation models refer to the class of discrete-event systems. We define requirements as expectations referring to the produced simulation outputs, which are sometimes stated in terms of constraints or inequalities. Requirements are often also given in terms of data (and a distance measure) whether or not the simulation outputs are sufficiently close to the data, or as formal expressions in a (spatial-)temporal logic which can be used for

■ **Table 1** Exemplary model-based approaches for specifying simulation models, simulation experiments, and requirements.

	Simulation Model	Simulation Experiment	Requirements
Formalisms	DEVS [31], Stochastic Petri Nets [1], Process algebras [9]	DEVS [4]	–
Programming languages, in particular DSLs	APIs [26], NetLogo [25], BioNetGen [2]	SESSL [27], NLRX [24]	FITS [17]
Logics	Ontologies for composing and reuse [18]	Ontologies of simulation algorithms [3]	Temporal logic [19]
Active objects	ABS [12, 13]	ABS runtime [30]	event-driven, cooperating processes
Metamodels	ATom3 for multi-formalism modeling [14]	meta-models for simulation experiments [29]	–

statistical model checking or runtime verification [19]. Ontologies are widely used, e.g., in the context of model composition so to map variables of different models consistently [18], in the context of simulation experiments, to classify the simulation algorithms used (and assess the approximation) [3], and referring to requirements again ontologies can be used to clarify what is meant with the variables used within a requirement. Metamodels are used for transformations to compose models defined in different modeling approaches and generate an executable model [14] or to generate simulation experiments for different modeling approaches as diverse as finite element analysis and stochastic discrete event simulation to be executed in different tools [29]. Whereas requirements refer to the model’s output, assumptions are some form of model input (as they influence how the simulation model looks like), and as such are essential for providing context and for semantically consistent model composition [18]. However, so far textual representations prevail, and little research has been done to model assumptions or theories. Not only the ingredients but also the simulation studies themselves have been subject to modeling, e.g., via workflows. Workflows have been used to support documenting and executing (also in terms of guidance) simulation studies [21, 23].

Towards Sustainable Simulation Studies. Referring to sustainability we adopted the definition given in the seminar, i.e., sustainability: continuing a simulation study into the future through support for reusing or building upon its central products, such as simulation model, data, and processes as well as the software used. We emphasized that this also applies to a simulation study currently being conducted (thus how to assist the modeler during a simulation study based on what has already been done and information gathered, see automation) and that the building upon is not restricted to modelers and domain experts but also needs to include decision-makers.

Context: The context of a simulation study is all information that helps in interpreting and reproducing the results achieved in a simulation study and makes up a large part of its documentation according to standardized guidelines [7] or its conceptual model [22, 28]. A recent empirical study shows that efforts invested in the reproducibility of a simulation study enhance its impact [10].

Effective adoption of provenance standards opens up new possibilities as essential processes, and the various sources that contributed to the results (also intermediate ones) become explicit.

Thereby, the interpretability, consistency, and reproduction of simulation studies can be enhanced, and the stored information can be used for user guidance and automation [11]. New research challenges arise, such as an unobtrusive collection of provenance information or on-the-fly abstraction to cater to the needs of different types of users. It should be noted, that the adoption of provenance standards does not alleviate the problem of how to unambiguously specify the different ingredients that make up the documentation, conceptual model, or provenance graph. However, the promise to have more comprehensive, automatic support for conducting simulation studies increases the incentive to tackle the problem.

Composition and reuse: Modeling languages should support composition in one way or the other to facilitate the development of the model and its reuse. The composition of simulation models is supported in formalisms such as DEVS by black-box composition via input and output ports, or process algebras inherently by the parallel composition of processes. Solutions that lie in between the rigidity of DEVS and the rather fluent composition of processes, and allow to flexibly define interaction points and the interface to evolve on demand, require further attention. In addition to these syntactical considerations, semantically valid reuse of simulation models still provides many challenges in modeling and simulation, although various approaches have been developed since its being stated as a grand challenge in 2002 [20]. A central question here is how the consistency of the respective context of the simulation models that shall be composed can be guaranteed. This leads us back to the question of how to represent the different ingredients that make up the context of a simulation study respectively simulation model in an unambiguous and (ideally) computationally accessible manner. A different aspect of reuse concerns variability: simulation models come in a plethora of related variants that are distinguished by differing assumptions, parameters, submodels, etc. It is essential to manage the commonality and variability embodied in these variants in a systematic manner.

Automation: Automation is one means to increase efficiency. As simulation studies are knowledge-intensive processes [5], any effort aimed at automation has to look at means for representing and evaluating the required knowledge computationally [11, 16]. Thus, model-based approaches are a central step towards at least (partly) automating simulation studies. There are various challenges referring to knowledge engineering, e.g., what information is needed to automate various activities such as model building, model refinement, calibrating, validating, or analyzing simulation models, which need to be addressed by the modeling and simulation community. Thereby, also new possibilities that are offered by applying, combining, and possibly revising methods from logic-based reasoning and machine learning have to be taken into account to automate simulation studies, e.g., in selecting methods to conduct specific experiments such as sensitivity analysis [15].

Communication: Much of the work on model-based approaches are aimed at more effective communication with different users. Despite the plethora of formalisms, domain-specific languages, etc., coming up with abstractions that are coherent to the modeling metaphors of domains, adequate for the problem, and even match the mental models of individuals remains a challenge. The research is also hampered by the difficulty to measure the effectiveness or potential impact of new methods and thus the progress made. New developments in visual analytics open up new possibilities to communicate the results

and context of simulation studies [8]. Discussing the various methods already used, their potential, limitations, and open challenges for visualization is the subject of another working group of the seminar. We expect challenges of mapping and adapting model representations of entities and processes involved in conducting a simulation study to different users and their respective needs to carry over to visualization. However, knowledge about the simulation study made explicit by model-based approaches will also further more effective visualization methods.

Conclusion. During the seminar, we identified four crucial areas of research for enhancing the sustainability of simulation studies, i.e., documenting the context of simulation studies, composition and reuse of simulation results, automation for conducting simulation studies more systematically, and communicating the results and processes of simulation studies with domain-experts and decision-makers. Central discussions revolved around whether and how concepts of software and programming languages can be or are already adopted in the field of modeling and simulation. There, the focus has been on composition, reuse, abstraction, and variability. We have still to hone in on the diversity of model-based approaches, the role they already play or might play in this endeavor, and the concrete methodological (and community) challenges associated.

References

- 1 Gianfranco Balbo. Introduction to stochastic petri nets. In *School organized by the European Educational Forum*, pages 84–155. Springer, 2000.
- 2 Michael L Blinov, James R Faeder, Byron Goldstein, and William S Hlavacek. Bionetgen: software for rule-based modeling of signal transduction based on the interactions of molecular domains. *Bioinformatics*, 20(17):3289–3291, 2004.
- 3 Mélanie Courtot, Nick Juty, Christian Knüpfer, Dagmar Waltemath, Anna Zhukova, Andreas Dräger, Michel Dumontier, Andrew Finney, Martin Golebiewski, Janna Hastings, et al. Controlled vocabularies and semantics in systems biology. *Molecular systems biology*, 7(1):543, 2011.
- 4 Joachim Denil, Stefan Klikovits, Pieter J Mosterman, Antonio Vallecillo, and Hans Vangheluwe. The experiment model and validity frame in m&s. In *Proceedings of the Symposium on Theory of Modeling & Simulation*, pages 1–12, 2017.
- 5 Claudio Di Ciccio, Andrea Marrella, and Alessandro Russo. Knowledge-intensive processes: characteristics, requirements and analysis of contemporary approaches. *Journal on Data Semantics*, 4(1):29–57, 2015.
- 6 Peter I Frazier, J Massey Cashore, Ning Duan, Shane G Henderson, Alyf Janmohamed, Brian Liu, David B Shmoys, Jiayue Wan, and Yujia Zhang. Modeling for covid-19 college reopening decisions: Cornell, a case study. *Proceedings of the National Academy of Sciences*, 119(2):e2112532119, 2022.
- 7 Volker Grimm, Jacqueline Augusiak, Andreas Focks, Béatrice M Frank, Faten Gabsi, Alice SA Johnston, Chun Liu, Benjamin T Martin, Mattia Meli, Viktoriia Radchuk, et al. Towards better modelling and decision support: documenting model development, testing, and analysis using trace. *Ecological modelling*, 280:129–139, 2014.
- 8 Jussi Hakanen, Kaisa Miettinen, and Krešimir Matković. Task-based visual analytics for interactive multiobjective optimization. *Journal of the Operational Research Society*, 72(9):2073–2090, 2021.
- 9 Jane Hillston. Process algebras for quantitative analysis. In *20th Annual IEEE Symposium on Logic in Computer Science (LICS'05)*, pages 239–248. IEEE, 2005.
- 10 Sebastian Höpfl, Jürgen Pleiss, and Nicole Radde. Bayesian hypothesis testing reveals that reproducible models in systems biology get more citations. 2022.

- 11 Pia Wilsdorf, Anja Wolpers, Jason Hilton, Fiete Haack, and Adelinde M. Uhrmacher. Automatic reuse, adaption, and execution of simulation experiments via provenance patterns. *ACM Transactions on Modeling and Computer Simulation*, 2022.
- 12 Einar Broch Johnsen, Reiner Hähnle, Jan Schäfer, Rudolf Schlatte, and Martin Steffen. ABS: A core language for abstract behavioral specification. In Bernhard K. Aichernig, Frank de Boer, and Marcello M. Bonsangue, editors, *Proc. 9th Intl. Symp. on Formal Methods for Components and Objects (FMCO 2010)*, volume 6957 of *LNCS*, pages 142–164. Springer, 2011.
- 13 Eduard Kamburjan, Stefan Mitsch, and Reiner Hähnle. A hybrid programming language for formal modeling and verification of hybrid systems. *Leibniz Transactions on Embedded Systems*, 8(1), 2022. Special Issue on Distributed Hybrid Systems.
- 14 Juan de Lara and Hans Vangheluwe. Atom 3: A tool for multi-formalism and meta-modelling. In *International Conference on Fundamental Approaches to Software Engineering*, pages 174–188. Springer, 2002.
- 15 Stefan Leye, Roland Ewald, and Adelinde M Uhrmacher. Composing problem solvers for simulation experimentation: a case study on steady state estimation. *PloS one*, 9(4):e91948, 2014.
- 16 Fabian Lorig. *Hypothesis-Driven Simulation Studies: Assistance for the Systematic Design and Conducting of Computer Simulation Experiments*. Springer Vieweg, Wiesbaden, 2019.
- 17 Fabian Lorig, Colja A Becker, and Ingo J Timm. Formal specification of hypotheses for assisting computer simulation studies. In *Proceedings of the Symposium on Theory of Modeling & Simulation*, pages 1–12, 2017.
- 18 Maxwell Lewis Neal, Matthias König, David Nickerson, Göksel Mısırlı, Reza Kalbasi, Andreas Dräger, Koray Atalag, Vijayalakshmi Chelliah, Michael T Cooling, Daniel L Cook, et al. Harmonizing semantic annotations for computational models in biology. *Briefings in bioinformatics*, 20(2):540–550, 2019.
- 19 Laura Nenzi, Luca Bortolussi, Vincenzo Ciancia, Michele Loreti, and Mieke Massink. Qualitative and quantitative monitoring of spatio-temporal properties. In *Runtime Verification*, pages 21–37. Springer, 2015.
- 20 C Michael Overstreet, Richard E Nance, and Osman Balci. Issues in enhancing model reuse. In *International Conference on Grand Challenges for Modeling and Simulation. San Antonio, Texas, USA*, 2002.
- 21 Judicaël Ribault and Gabriel Wainer. Using workflows and web services to manage simulation studies (wip). In *Proceedings of the 2012 Symposium on Theory of Modeling and Simulation-DEVS Integrative M&S Symposium*, pages 1–6, 2012.
- 22 Stewart Robinson. Conceptual modeling for simulation: issues and research requirements. In *Proceedings of the 2006 winter simulation conference*, pages 792–800. IEEE, 2006.
- 23 Andreas Ruschinski, Tom Warnke, and Adelinde M Uhrmacher. Artifact-based workflows for supporting simulation studies. *IEEE Transactions on Knowledge and Data Engineering*, 32(6):1064–1078, 2019.
- 24 Jan Salecker, Marco Sciaini, Katrin M Meyer, and Kerstin Wiegand. The nlrx r package: A next-generation framework for reproducible netlogo model analyses. *Methods in Ecology and Evolution*, 10(11):1854–1863, 2019.
- 25 Seth Tisue and Uri Wilensky. Netlogo: A simple environment for modeling complexity. In *International conference on complex systems*, volume 21, pages 16–21. Boston, MA, 2004.
- 26 Yentl Van Tendeloo and Hans Vangheluwe. The modular architecture of the python (p) devs simulation kernel. In *Proceedings of the 2014 Symposium on Theory of Modeling and Simulation-DEVS*, pages 387–392, 2014.
- 27 Tom Warnke and Adelinde M Uhrmacher. Complex simulation experiments made easy. In *2018 Winter Simulation Conference (WSC)*, pages 410–424. IEEE, 2018.

- 28 Pia Wilsdorf, Fiete Haack, and Adelinde M Uhrmacher. Conceptual models in simulation studies: Making it explicit. In *2020 Winter Simulation Conference (WSC)*, pages 2353–2364. IEEE, 2020.
- 29 Pia Wilsdorf, Jakob Heller, Kai Budde, Julius Zimmermann, Tom Warnke, Christian Haubelt, Dirk Timmermann, Ursula van Rienen, and Adelinde M Uhrmacher. A model-driven approach for conducting simulation experiments. *Applied Sciences*, 12(16):7977, 2022.
- 30 Peter Y. H. Wong, Elvira Albert, Radu Muschevici, José Proença, Jan Schäfer, and Rudolf Schlatte. The ABS tool suite: modelling, executing and analysing distributed adaptable object-oriented systems. *STTT*, 14(5):567–588, 2012.
- 31 Bernard P Zeigler, Alexandre Muzy, and Ernesto Kofman. *Theory of modeling and simulation: discrete event & iterative system computational foundations*. Academic press, 2018.

Participants

- Philipp Andelfinger
Universität Rostock, DE
- Luca Bortolussi
University of Trieste, IT
- Wentong Cai
Nanyang TU – Singapore, SG
- Christopher Carothers
Rensselaer Polytechnic Institute –
Troy, US
- Rodrigo Castro
University of Buenos Aires, AR
- Joachim Denil
University of Antwerp, BE
- Jérôme Feret
ENS – Paris, FR
- Peter Frazier
Cornell University – Ithaca, US
- Reiner Hähnle
TU Darmstadt, DE
- Dong (Kevin) Jin
University of Arkansas –
Fayetteville, US
- Franziska Klügl
University of Örebro, SE
- Till Köster
Universität Rostock, DE
- Michael Lees
University of Amsterdam, NL
- Jason Liu
Florida International University –
Miami, US
- Margaret Loper
Georgia Institute of Technology –
Atlanta, US
- Fabian Lorig
Malmö University, SE
- Bertram Ludäscher
University of Illinois at
Urbana-Champaign, US
- Kresimir Matkovic
VRVis – Wien, AT
- Laura Nenzi
University of Trieste, IT
- David M. Nicol
University of Illinois –
Urbana Champaign, US
- Alessandro Pellegrini
University of Rome “Tor
Vergata”, IT
- Niki Popper
Technische Universität Wien, AT
- Caitlin Ross
Kitware – Clifton Park, US
- Cristina Ruiz-Martin
Carleton University –
Ottawa, CA
- Bernhard Rumpe
RWTH Aachen, DE
- Susan Sanchez
Naval Postgrad. School –
Monterey, US
- Nadja Schlungbaum
Universität Rostock, DE
- Peter Sloot
University of Amsterdam, NL
- Claudia Szabo
University of Adelaide, AU
- Wen Jun Tan
Nanyang TU – Singapore, SG
- Adelinde M. Uhrmacher
Universität Rostock, DE
- Gabriel A. Wainer
Carleton University –
Ottawa, CA
- Pia Wilsdorf
Universität Rostock, DE
- Verena Wolf
Universität des Saarlandes –
Saarbrücken, DE

