AFRL-AFOSR-VA-TR-2019-0140

Computation-based Probabilistic Validation of Complex Cyber-controlled Systems

Geir Dullerud
UNIVERSITY OF ILLINOIS

05/09/2019
Final Report

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
06-28-2018	Final	04-01-2015 to 03-31-2018

4. TITLE AND SUBTITLE
Computation-based Probabilistic Validation of
Complex Cyber-controlled Systems

5a. CONTRACT NUMBER
FA9550-15-1-0059

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Geir E. Dullerud (PI), Mahesh Viswanathan, and Matthew West

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Illinois
1206 W Green Street
Urbana, IL 61801

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
875 N Randolph Street STE 325 RM 3112
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Distribution A

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The project research was concerned with verification of stochastic hybrid systems with temporal logic performance specifications. The major technical theme of the research was the leveraging of the computational tools from both probabilistic physical system verification, and those of discrete system verification, to create a computationally effective and automated verification methodology. A project emphasis was on developing methods that can scale gracefully to large systems. The goals were pursued both on a theory level, developing approaches that reduce hybrid verification problems with temporal logic specifications to more tractable computations that still provide safety and performance guarantees, as well as by developing new computational and simulation algorithms that are principle-based and provide improved efficiency over those currently available. These methods include counter-example guided schemes. Also investigated were fundamental questions about the existence of verification algorithms for certain hybrid system classes. As well, part of the research involved development of new tools for the analysis of hybrid systems that specifically contain the sampling of signals, and the research resulted in significant improvements over existing methods for verifying

15. SUBJECT TERMS
stochastic hybrid systems, Mori-Zwanzig model reduction, statistical model checking, temporal logic, MITL, PCTL, lattice Markov chain simulation, sampled-data nonlinear systems, switched systems.

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Geir E. Dullerud
					19b. TELEPHONE NUMBER *(Include area code)* 217-265-5078

Reset

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18
Adobe Professional 7.0

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/ monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

Computation-based Probabilistic Validation of Complex Cyber-controlled Systems: FINAL REPORT

FA9550-15-1-0059

Geir E. Dullerud (PI)
Mechanical Science and Engineering

Mahesh Viswanathan
Computer Science

Matthew West
Mechanical Science and Engineering

University of Illinois at Urbana-Champaign

Prepared for: Dr. Frederick A. Leve, Program Officer
Program: Dynamics and Control (RTA), AFOSR

Contents

1 Executive Summary

This 3-year project started on April 1, 2015 and has resulted in a substantial number of important technical accomplishments, creating new predictive capability in the verification and simulation of hybrid systems that have sophisticated performance specifications given in terms of various temporal logics suitable for hybrid systems, as well as other types of performance metrics. This research has resulted in a sequence of research papers published, or currently under review, in top conferences and archival journals. The PIs have also been active in disseminating the project research by giving plenary and keynote lectures at conferences and workshops, as well as giving seminars at universities and other technical institutions. Additionally, the project has also trained several graduate students on the specific methods of the project, and more generally in the area of verification of hybrid systems. The technical achievements have both been on the theory of hybrid systems, as well as explicit computational methods for deploying the theory, and significant effort has been applied to creating computational verification tools, and applying them to bench mark cases. Recently, we have also been developing new distributed codes that can be run on large-scale computing platforms. This overall research effort has resulted in 21 articles appearing in the archived proceedings of major conferences or top disciplinary research journals; further, 2 additional papers have been submitted and are under review. The research in all of these publications is aimed at computational approaches for either analysis, verification or synthesis for controlled hybrid systems. The major theme of our research was to leverage the computational tools from both probabilistic physical system verification, and those of discrete system verification, to create a computationally effective and automated verification methodology. *An emphasis of this work has been on developing methods that can scale gracefully to large systems.* We have pursued this goal both on a theory level, developing approaches that reduce hybrid verification problems with temporal logic specifications to more tractable computations that still provide safety and performance guarantees, as well as by developing new computational and simulation algorithms that are principle-based and provide improved efficiency over those currently available. These methods include counter-example guided schemes. We have also investigated fundamental questions about the existence of verification algorithms for certain hybrid system classes. Finally, we have developed new tools for the analysis of hybrid systems that specifically contain the sampling of signals, and developed significantly improved results for verifying these. In the final stages of the program we combined the newly developed methods for large-scale system simulation with the verification methodology created.

1.0.1 Project Overview

The main goal of the program was the development of automated verification tools that can be used for large-scale systems that contain both complex physics-based and digitally-based dynamics, and was based on the widely-held vision that the battlespace of the future will include heterogeneous teams of UASs and other agents performing multiple missions autonomously, and sharing airspace with human-piloted vehicles. And the recognition that despite this vision, and promising steps, there remained (and still remains) a formidable barrier to achieving this operational capability, with a key problem being that although it is possible to engineer and physically construct high-performance systems that can "run" without human assistance, predicting how they will behave when interacting autonomously with their environment and other vehicles with any degree of confidence remains a major challenge. This uncertainty arises because the control algorithms and software for such systems are necessarily incredibly complex. A major obstacle to achieving the vision on both single- and multi-agent scales is the challenge of verifying and validating algorithms that control these inherently hybrid systems; a further complicating factor is the complex specifications for such systems are

not the much studied traditional notions of stability and reachability. This collaborative program developed such automated techniques and tools for verification and validation of hybrid systems with temporal logic specifications, and had a specific emphasis on the use of large-scale computation.

In the program we developed a computational approach to verification of the hybrid mathematical models that are formed when combining physics-based models with discrete-transition models, such as those which model software algorithms; namely, the types of models that arise when physical processes are interconnected with digital hardware. The technical foundation for the research was to use a Markov process formulation, simulation, and probabilistic model checking to achieve a unified approach to verification of systems that contain both continuous- and discrete-state dynamics. A significant advantage is gained in this probabilistic setting because the models used in these individual settings are already Markov processes, and thus the approach enabled leveraging the two large bodies of work on probabilistic analysis of purely physics-based models, and probabilistic model checking for purely discrete-transition systems. In the project we first developed a new systematic framework for considering automated analysis. A central aspect of the research was the development of probabilistic model checking algorithms and sophisticated abstraction techniques. We introduced new design logics and algorithms to specifically reason about spatio-temporal system properties. In summary the fulfilled technical objectives were:

- *Constructed a general framework for automated analysis based on Markov chain models of ODEs, PDEs, SDEs and SPDEs, coupled with spatio-temporal logics.*

- *Employed dynamical systems methods for model simplification and abstraction.* Combine simulation-based methods with set-oriented system representations to produce a unified model abstraction framework. Created extended, *weaker* notions of simulation. Investigated and characterized the appropriate relationship between the Markov models and the ODE/PDE models, and studied the class of logical properties that are *reflected*, and *preserved* by this reduction process.

- *Designed logics and algorithms to reason about spatio-temporal properties of systems.* Automated verification techniques for stochastic systems have been primarily developed for specifications that require reasoning about the measures of executions satisfying specified properties. However, the process of transforming a physical model described by ODEs and PDEs into a Markov chain essentially consists of characterizing how probability distributions are evolved by the model. The correspondence between the constructed Markov chain and the original ODE/PDE only applies to the way the Markov chain transforms distributions. There are no efficient algorithms to reason about the types of extended logics needed for specifications in this context. In the project developed a probabilistic model checking approach to reason about such specification properties. The primary idea will be the use of samples to obtain systematic estimates about distributions, and then translate formulas to automata on which temporal reasoning can be completed. The error estimation will be extended for the statistical approach to the project framework.

- Developed explicit verification and simulation codes; illustrate and validate on simulation-based platforms.

1.0.2 Synopsis of Achievements

A summary of the technical successes of the project are given below:

Scalable hybrid modeling and abstraction framework with temporal logic specifications [1,8,20]. Given a hybrid system with temporal logic specifications, an abstraction method was developed that led to a reduced model given in terms of a finite state Markov chain. This was accomplished by using the Mori-Zwanzig model reduction method from physical sciences. It was shown that this Markov chain system abstraction is approximately equivalent to the original system in a distributional sense. This is an important property as the approximate equivalence of the stochastic hybrid system and its Markov chain abstraction means that analyzing the Markov chain with respect to a strengthened property, allows one to conclude whether or not the stochastic hybrid system meets its temporal specification.

Statistical model checking of Markov chains [7,10,12]. The project developed the first statistical model checking algorithms to verify finite state Markov chains against correctness properties expressed in linear inequality linear temporal logic (iLTL), and metric interval temporal logic (MITL). This is an important basic algorithm that has application beyond the hybrid systems focus of the project.

New accelerated methods for simulation of discrete-time Markov chains [17,21]. Developed was a new algorithm for simulation of an important class of chains; namely, the class of countable-state, discrete-time Markov chains driven by additive Poisson noise—called, lattice discrete-time Markov chains. In particular, this class can approximately simulate continuous-time Markov chains by employing tau-leaping. The general approach was based on negatively correlating trajectories. Numerical results for example systems, chosen with different levels of complexity and nonlinearity, included two to four orders of magnitude reduction of mean-square error. Analytical results were also obtained showing provable algorithm properties under different special cases. The algorithm is highly implementable on most pre-existing simulation codes.

Exact, variance-reduced, simulation of lattice continuous-time Markov chains [3,17]. New algorithms (based on binomial properties of the Poisson distribution) to reduce the variance of Monte Carlo simulation for lattice continuous-time Markov chains, or lattice CTMCs. This is a very broad class of systems including all processes that can be represented using the so-called random-time-change representation. Numerical experiments demonstrated at least order-of-magnitude performance improvement measured in terms of mean-square error. Exact, analytical expressions were proved for properties/quantities in the algorithm.

Model checking hybrid systems using stratified sampling [2]. In this part of the program we combine work on model checking Markov chains with our new algorithmic approach to simulating them. Our research and numerical examples show that the variance reducing methods developed earlier in the program can improve performance of model checking.

Explicit benchmark testing [10,15]. We have applied the statistical model checking methods developed in the program to a research community challenge bench (high-dimensional power train model). Using our method we were the first to be able to provide statistical guarantees on the system satisfying its performance specification.

Stochastic control [17, 21]. Our results on anti-correlated simulation were combined with an MPC-like framework to produce a strategy for closed stochastic control. This general architecture, numerical simulation combined with iterative control optimization, has significant potential for use on uncertain high-dimensional systems.

Sampled-data and switched systems [4,5,6,14,19]. Considered were hybrid sampled-data systems with arbitrary nonlinear jump maps, and developed was a new methodology for determining their combined stability and performance. The conditions developed are currently the least conservative in the research literature. Similarly, in the context of switched systems, achieved new exact results for analysis and synthesis.

Real-time logics and robustness [7,8,12]. Have studied real-time system logics and developed new decision procedures (also showing that 20-year-old standard procedures are incorrect). Developed a robustness methodology for hybrid systems in terms of syntactic perturbations.

Counterexample guided abstraction refinement [16,18,24]. Developed a new CEGAR technique, which several fundamentally important properties. Created the software tool HARE.

2 Publications

Listed below are the published archival articles (21) of the project, together with 2 additional articles that are under review. Additional articles may also be submitted in the near future.

[1] Wang, Y., N. Roohi, M. West, M. Viswanathan, and G.E. Dullerud, "Verifying Stochastic Hybrid Systems with Temporal Logic Specifications via Mori-Zwanzig Model Reduction," submitted to IEEE Transactions on Automatic Control, 2018.

[2] Wang, Y., N. Roohi, M. West, M. Viswanathan, and G.E. Dullerud, "Statistical Verification of PCTL Using Stratified Samples," to appear in Proceedings IFAC Conference on Analysis and Design of Hybrid Systems (ADHS), 2018.

[3] Maginnis, P. A., M. West, and G.E. Dullerud, "Exact variance-reduced simulation of lattice continuous-time Markov chains with applications in reaction networks," submitted to Bulletin of Mathematical Biology, 2018.

[4] Naghnaeian, M., P.G. Voulgaris, and G.E. Dullerud, "Lp Analysis and Synthesis of Linear Switched Systems: A Unified Framework," to appear in SIAM Journal of Control and Optimization, 2018.

[5] Jansch-Porto, J.P. and G.E. Dullerud, "Decentralized Control of Switched-Systems with Path-Dependent l2-induced Bounds," Proceedings of American Control Conference (ACC), 2018.

[6] Essick, R. and G.E. Dullerud, "Application of a Message-Passing Decomposition of Sparsely-Coupled Linear Programming Problems to the Uniform Stabilization of Positive Switched Linear Systems" Proceedings of American Control Conference (ACC), 2018.

[7] Roohi, N. and M. Viswanathan, "Revisiting MITL to Fix Decision Procedures", Proceedings of the International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI), 2018.

[8] Roohi, N., P. Prabhakar, and Mahesh Viswanathan, "Relating Syntactic and Semantic Perturbations of Hybrid Automata", Proceedings of the International Conference on Concurrency Theory (CONCUR), 2018 (to appear).

[9] Roohi, N., Y. Wang, M. West, G.E. Dullerud, and M. Viswanathan, "Statistical Verification of the Toyota Powertrain Control Verification Benchmark", Proceedings of Hybrid Systems: Control and Computation (HSCC), 2017.

[10] Jansch-Porto, J.P. and G.E. Dullerud, "Decentralized Control with Moving-Horizon Linear Switched Systems: Synthesis and Testbed Implementation," Proceedings of American Control Conference (ACC), 2017.

[11] N. Roohi, P. Prabhakar, M. Viswanathan, "Robust Model Checking of Timed Automata under Clock Drifts", in Proceedings of the International Conference on Hybrid Systems: Computation and Control (ICCPS), 2017.

[12] Wang, Y., S. Mitra, and G.E. Dullerud, "Differential Privacy and Minimum-Variance Unbiased Estimation in Multi-agent Control Systems," Proceedings of IFAC World Congress, 2017.

[13] Strijbosch, N., G.E. Dullerud, A.R. Teel, and W.P.M.H. Heemels, "L2-gain Analysis of Periodic Event-triggered and Self-triggered Control Systems with Delays using Lifting Techniques," Proceedings of IEEE Conference on Decision and Control, 2017.

[14] Buccafusca, L., C.L. Beck, and G.E. Dullerud, "Modeling and Maximizing Power in Wind Turbine Arrays," Proceedings of IEEE Conference on Control Technologies and Applications, 2017.

[15] N. Roohi, P. Prabhakar, M. Viswanathan, "HARE: A Hybrid Abstraction Refinement Engine for Verifying Non-Linear Hybrid Automata", in Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems, 2017.

[16] Maginnis, P.A., M. West, and G.E. Dullerud, "Model Predictive Control of Markov Jump Processes with Anticorrelated Variance Reduced Monte Carlo Estimation," Proceedings of American Control Conference, 2016.

[17] Roohi, N., P. Prabhakar, M. Viswanathan."Hybridization based CEGAR for Hybrid Automata with Affine Dynamics", Proceedings of International Conference on Tools and Analysis for the Construction and Analysis of Systems, 2016.

[18] Heemels, W.P.M.H., G. Dullerud, and A.R. Teel, "L2-gain Analysis for a Class of Hybrid Systems with Applications to Reset and Event-triggered Control: A Lifting Approach," IEEE Transactions on Automatic Control, 2016.

[19] Wang, Y., N. Roohi, M. Viswanathan, M. West, and G. Dullerud, "Verifying Continuous-time Stochastic Hybrid Systems via Mori-Zwanzig Model Reduction", Proceedings of IEEE Conference on Decision and Control (CDC), 2016.

[20] Maginnis, P. A., M. West, and G. Dullerud, "Variance-reduced Tau-leaping using Anticorrelated Sample Paths," Journal of Computational Physics, 2016.

[21] Wang, Y., Z. Huang, S. Mitra, and G.E. Dullerud, "Differential Privacy and Entropy-minimizing Mechanisms in Feedback Systems," IEEE Transactions on Control of Network Systems, 2016.

[22] Wang, Y., N. Roohi, M. Viswanathan, and G. Dullerud, "Stability of Linear Autonomous Systems Under Regular Switching Sequences," IEEE Transactions on Automatic Control, 2016.

[23] P.S. Duggirala, M. Viswanathan, "Parsimonious, Simulation Based Verification of Linear Systems", in Proceedings of the International Conference on Computer-Aided Verification, 2016.

3 Technical Accomplishments and Scientific Progress

Provided below are more detailed technical descriptions of the research accomplishments of the program.

3.1 Verification of hybrid systems with temporal logic formulas via Mori-Zwanzig model abstraction

In the project we developed a scalable approach to verification of an important class of stochastic hybrid systems, whose specifications were given in terms of temporal logic, using a dynamical systems model reduction approach. This work treats the dynamical system and its approximate model, and shows that our reduced models are approximately equivalent to the original hybrid system. This work is now described in some detail.

Verification problems of hybrid systems with temporal logic specifications are provably hard problems, even for relatively simple models. And the fundamental difficulty of the verification problem largely arises from the fact that the state space of such systems has uncountably many states. The computational challenge posed by the verification problem is often addressed by constructing a simpler finite state model of the system, and then analyzing the finite state model. The finite state model is typically a so-called *abstraction* or a conservative over-approximation of the original system, i.e., every behavior of the system is exhibited by the finite state model, but the finite state model may have additional behaviors that are not system behaviors. For such abstractions, if the finite state model is safe then so is the original system. However, if the finite state model is unsafe, then not much can be concluded about the safety of the original system because the finite state model is an over-approximation.

Our work developed under this project provides a scalable approach to verification of stochastic hybrid system that relies on constructing a finite state approximation that is "equivalent" to the original system. The advantage of using a reduction that is approximately equivalent to the original system is that analyzing the finite state model not only allows us to conclude the safety of the hybrid stochastic system, but also its non-safety. We construct the finite state Markov chain reduction by using the Mori-Zwanzig model reduction method. In order to explain the relationship between the Markov chain we construct and the stochastic hybrid system, *it is useful to recall that there are two broad approaches to defining the semantics of a stochastic process.* One approach is to view a stochastic system as defining a measure space on the collection of executions; by execution here we mean a sequence, or trajectory, of states that the system may possibly go through. The other approach is to view the stochastic system as defining a transformation on distributions; in such a view, the behavior of the stochastic model is captured by a sequence of distributions, starting from some initial distribution. It has been observed that with respect to the first semantics (of measures on executions) it is not possible to construct a finite state Markov chain that is "equivalent" to an infinite state system. In contrast, in our work we show that the Mori-Zwanzig reduction method constructs a finite state Markov chain that *is* approximately equivalent to a stochastic hybrid system with respect to the second semantics. That is, we show that the distribution on states of the Markov chain at any time, is close to the distribution at the same time defined by the stochastic hybrid system.

Having proved that our reduced Markov model is approximately equivalent to the original stochastic hybrid system, we can exploit this to verify stochastic hybrid systems. Approximate equivalence ensures that analyzing the reduced model with respect to a suitably strengthened property, allows us to determine whether the initial stochastic hybrid system meets or violates its requirements. Therefore, a scalable verification approach can be obtained by developing algorithms to verify finite state Markov chains. Since the reduced system, even though finite state, is likely to have a large number of states, we use a statistical approach to verification as opposed to a symbolic one. In statistical model checking, the model being verified is simulated multiple times, and the drawn simulations are analyzed to see if they constitute a statistical evidence for the correctness of the model. Statistical model checking algorithms have been developed for logics that reason about measures of executions. However, since our reduced Markov chain is only close to the stochastic hybrid system in a distributional sense, we could not leverage these existing algorithms. We developed new statistical model checking algorithms for temporal logics (over both discrete and continuous time) that reason about sequences of distributions.

We believe our new approach to verifying stochastic hybrid systems is scalable. Our initial experimental evaluation supports this claim. Using this approach, we were the first to successfully verify a highly nonlinear model including lookup tables of a powertrain control system that was proposed as a community challenge problem for verification tools by Toyota engineers.

Continuous-time Stochastic Hybrid System Model

We define our hybrid model using a Fokker-Planck formulation and interpretation. We denote the continuous and discrete states by $x \in \mathbb{R}^d$ and $q \in \mathcal{Q}$ respectively, where $\mathcal{Q} = \{q_1, \ldots, q_m\}$ is a finite set. We call the combination (q, x) the state of the system, and the product set $\mathbb{X} = \mathcal{Q} \times \mathbb{R}^d$ the state space.

The state space $\mathbb{X}$ of the system is divided into two regions: a flow set $\mathbb{A}$ and a jump set $\mathbb{B} = \mathbb{X} \backslash \mathbb{A}$. Define $\mathbb{A}_q = \left\{x \in \mathbb{R}^d \mid (q, x) \in \mathbb{A}\right\}$, and $\mathbb{B}_q$ similarly. We assume that each $\mathbb{A}_q$ is compact, and the boundaries $\partial\mathbb{A}_q$ are second-order continuously differentiable in x. On the flow set, the state x of the system evolves by a stochastic differential equation

$$\mathrm{d}\mathbf{x} = f(\mathbf{q}, \mathbf{x})\mathrm{d}t + g(\mathbf{q}, \mathbf{x})\mathrm{d}B_t, \tag{1}$$

where $\mathbf{q}$ and $\mathbf{x}$ are random processes describing the stochastic evolution of the continuous and discrete states, and B_t is the standard n-dimensional Brownian motion. The vector-valued function f specifies the drift of the state, and the matrix-valued function g describes the intensity of the diffusion. Meanwhile, the system jumps spontaneously by a non-negative integrable rate function $r_{\mathbb{A}}(q, x)$. The probability distribution of the jumping target is given by a non-negative integrable target distribution $h_{\mathbb{A}}(q', x', q, x)$. When the state of the system falls onto the jump set $\mathbb{B}$, the system is forced to jump. The probability distribution of the jumping target is given by a non-negative integrable target distribution $h_{\mathbb{B}}(q', x', q, x)$. The two target distributions $h_{\mathbb{A}}$ and $h_{\mathbb{B}}$ defined on two disjoint sets $\mathbb{A}$ and $\mathbb{B}$ are combined into one target transition h defined on the state space $\mathbb{X}$ of the system and satisfying

$$\sum_{q \in \mathcal{Q}} \int_{\mathbb{R}^d} h(q', x', q, x)\mathrm{d}x' = 1. \tag{2}$$

The probability distribution $F(t, q, x)$ of the state of the system is determined by the standard Fokker-

Planck equation

$$\begin{aligned}
&\frac{\partial F(t,q,x)}{\partial t} = L(F(t,q,x)) \\
&= \underbrace{-\sum_{a=1}^{d} \frac{\partial}{\partial x_a}(f_a(q,x)F(t,q,x))}_{\text{drift}} \\
&\underbrace{+\sum_{a=1}^{d}\sum_{b=1}^{d} \frac{\partial^2}{\partial x_a \partial x_b} \sum_{c=1}^{d} \frac{g_{ac}(q,x)g_{cb}(q,x)F(t,q,x)}{2}}_{\text{diffusion}} \\
&\underbrace{-r(q,x)F(t,q,x)}_{\text{jump-out}} \\
&\underbrace{+\sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^d} h(q,x,q',x')r(q',x')F(t,q',x')\mathrm{d}x,}_{\text{jump-in}}
\end{aligned} \tag{3}$$

where L is the Fokker-Planck operator for the system. We write symbolically that $F(t,q,x) = e^{tL}F(0,q,x)$. In (3), the four terms on the right hand side describe "drift", "diffusion", "jump-out" and "jump-in", respectively.

On the other hand, a Fokker-Planck equation with proper boundary condition that gives unique solution defines a stochastic differential equation with jump and diffusion, and we therefore make the following assumption which ensures the system equations have well-defined solutions.

Assumption 1. *A unique solution exists for the Fokker-Planck equation corresponding to the system.*

A key component of the Mori-Zwanzig model reduction method is the invariant distribution, and we assume that the continuous-time stochastic hybrid system has an invariant distribution with probability distribution function $F_{\text{inv}}(q,x)$ such that

$$L(F_{\text{inv}}(q,x)) = 0. \tag{4}$$

And, for any initial state, the probability distribution function $F(t,q,x)$ converges to the invariant distribution function $F_{\text{inv}}(q,x)$.

In many applications, the state of the system is only partially observable, and in our model we define system observables by

$$\begin{aligned}
y(t) &= \mathbb{E}[y(q(t),x(t))] \\
&= \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^d} \gamma(q,x)F(t,q,x)\mathrm{d}x,
\end{aligned} \tag{5}$$

where $\gamma(q,x)$ is a weight function on $\mathbb{X}$, which is integrable in x for each $q \in \mathcal{Q}$.

Temporal Logic Specifications

Our focus was on verifying temporal properties of continuous-time stochastic hybrid systems. These properties are be specified in the following way: the atomic propositions are inequalities $y \sim c$ ($c \in \mathbb{Q}$,

$\sim \in \{<, \leq, \geq, >\}$) on the observables of the system; and they are concatenated by the syntax of Metric Interval Temporal Logic (MITL). This type of logic is also referred to as Signal Temporal Logic (STL) in the literature. The syntax of MITL is given in Definition 1.

Definition 1 (MITL Syntax)**.** *Let $\mathbb{I}_{\geq 0}$ be the set of intervals on $\mathbb{R}$ that are both non-negative and non-singleton. An MITL formula is defined using the following BNF form:*

$$\varphi ::= \bot \mid \top \mid p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \mathcal{U}_I \varphi \mid \varphi \mathcal{R}_I \varphi,$$

where $p \in \mathtt{AP}$ and $I \in \mathbb{I}_{\geq 0}$.

We note that this syntax does not contain negation ($\neg$), since $\{<, \leq, \geq, >\}$ is closed under negation. For a standard MITL formula, negation on non-atomic formulas can always be pushed inside as part of the atomic propositions. For example, $\neg(y > 0)$ is equivalent to $y \leq 0$, $\neg(\varphi_1 \vee \varphi_2)$ is equivalent to $(\neg\varphi_1) \wedge (\neg\varphi_2)$, and $\neg(\varphi \mathcal{U}_I \psi)$ is equivalent to $(\neg\varphi)\mathcal{R}_I(\neg\psi)$.

Given a set of atomic propositions, the continuous-time stochastic hybrid system induces a signal $f : \mathbb{R}_{\geq 0} \to 2^{\mathtt{AP}}$, in which $f(t)$ is the set of atomic proposition that holds on the system at time t. The semantics of MITL are defined with respect to the function $f(t)$ as follows.

Definition 2 (MITL Semantics)**.** *Let φ be an MITL formula and f be a signal $f : \mathbb{R}_{\geq 0} \to 2^{\mathtt{AP}}$. The satisfaction relation $\models$ between f and φ is defined according to the following inductive rules:*

$$\begin{array}{lll}
f \models \bot & \text{iff} & \textit{always false} \\
f \models \top & \text{iff} & \textit{always true} \\
f \models y \sim c & \text{iff} & (y \sim c) \in f(0) \\
f \models \varphi \wedge \psi & \text{iff} & (f \models \varphi) \wedge (f \models \psi) \\
f \models \varphi \vee \psi & \text{iff} & (f \models \varphi) \vee (f \models \psi) \\
f \models \varphi \mathcal{U}_I \psi & \text{iff} & \exists t \in I, (f^t \models \psi) \\
 & & \wedge \forall t' \in (0, t), f^{t'} \models \varphi \\
f \models \varphi \mathcal{R}_I \psi & \text{iff} & \forall t \in I, (f^t \models \psi) \quad \textit{or} \\
 & & \exists t \in \mathbb{R}_{>0}, (f^t \models \varphi) \\
 & & \quad \wedge \forall t' \in [0, t_1] \cap I, f^{t'} \models \psi
\end{array}$$

where f^r is a signal that maps t to $f(t + r)$. We define $[\![\varphi]\!]$ to be the set of signals that satisfy φ.

Satisfiability and model checking problems for MITL with *abstract* atomic propositions are known to be EXPSPACE-complete. The corresponding decision procedure has a close connection with timed automata.

Definition 3. *Timed automaton A is a tuple $(\mathtt{Q}, \mathtt{X}, \Sigma, \mathtt{L}, \mathtt{I}, \mathtt{E}, \mathtt{Q}^{\mathsf{init}}, \mathtt{Q}^{\mathtt{final}})$ where*

- $\mathtt{Q}$ *is a finite non-empty set of* locations.
- $\mathtt{X}$ *is a finite set of* clocks.
- Σ *is a finite* alphabet.
- $\mathtt{L} \in \mathtt{Q} \to \Sigma$ *maps each location to the* label *of that location.*
- $\mathtt{I} \in \mathtt{Q} \to (\mathtt{X} \to \mathbb{I}_{\geq 0})$ *maps each location to its* invariant *which is the set of possible values of variables in that location.*
- $\mathtt{E} \subseteq \mathtt{Q} \times \mathtt{Q} \times 2^{\mathtt{X}}$ *is a finite set of* edges *of the form (s, d, j), where $s = \mathtt{S}e$ is* source *of the edge; $d = \mathtt{D}e$ is* destination *of the edge; and $j = \mathtt{J}e$ is the set of clocks that are* reset *by the edge.*
- $\mathtt{Q}^{\mathsf{init}} \subseteq \mathtt{Q}$ *is the set of* initial locations.

- $\mathtt{Q}^{\mathtt{final}} \subseteq \mathtt{Q}$ *is the set of* final locations.

A *run* of the timed automaton A is a sequence of tuples $(\rho, \tau, \eta) \in \mathtt{Q} \times \mathbb{I}_{\geq 0} \times \mathtt{E}$ with the following conditions holds:

(i) $\rho_0 \in \mathtt{Q}^{\mathsf{init}}$, *i.e.*, ρ starts from an initial location $\mathtt{Q}^{\mathsf{init}}$;

(ii) $(\mathtt{S}\eta_n = \rho_n) \wedge (\mathtt{D}\eta_n = \rho_{n+1})$, *i.e.*, the source and destination of edges η_n are ρ_n and ρ_{n+1};

(iii) $\tau_0, \tau_1, \ldots$ is an ordered and disjoint partition of the time horizon $\mathbb{R}_{\geq 0}$; and

(iv) $\forall t \in \tau_n, x \in \mathtt{X}$, we have $\varrho_n(x) + t - \underline{\tau}_n \in \mathtt{I}(\varrho_n, x)$, where $\varrho_{n+1}(x)$ is defined inductively by

$$\varrho_{n+1}(x) = \begin{cases} 0, & \text{if } x \in \mathtt{J}\eta_n \\ \varrho_n(x) + \overline{\tau}_n - \underline{\tau}_n, & \text{otherwise} \end{cases}$$

i.e., the clock times must satisfy the invariant of the current location.

Here, $\underline{\tau}$ and $\overline{\tau}$ are the lower and upper bound of the interval.

A run satisfying the condition $\mathtt{inf}(\rho) \cap \mathtt{Q}^{\mathtt{final}} \neq \emptyset$, *i.e.*, some location from $\mathtt{Q}^{\mathtt{final}}$ has been visited infinitely many times by ρ, is called an *accepting run* of A. Note that every run of A induces a function f of type $\mathbb{R}_{\geq 0} \to \Sigma$ that maps t to $\mathtt{L}(\rho_n)$, where n is uniquely determined by the condition $t \in \tau_n$. We define the *language* of A, denoted by $\mathtt{Lang}(A)$, to be the set of all functions that are induced by accepting runs of A.

Abstraction of Continuous-time Hybrid Systems

To implement the Mori-Zwanzig model reduction method for continuous-time stochastic systems, we divide the continuous state space into finitely many partitions $\mathbb{S} = \{s_1, \ldots, s_n\}$, and treat each of them as a discrete state. We assume that for each s_i, there exists $q \in \mathcal{Q}$ such that $s_i \subseteq \{q\} \times \mathbb{A}_q$, and denote its measure by $\mu(s_i)$. Let $m(\mathbb{X})$ and $m(\mathbb{S})$ be set of probability distribution functions on $\mathbb{X}$ and $\mathbb{S}$, respectively. Then we can define a projection $P : m(\mathbb{X}) \to m(\mathbb{S})$ and an injection $R : m(\mathbb{S}) \to m(\mathbb{X})$ between $m(\mathbb{X})$ and $m(\mathbb{S})$ by

$$p_j = (PF(q,x))_j = \int_{s_j} F(q,x)\mathrm{d}x, \tag{6}$$

where p_j is the jth element of p, and

$$Rp = \sum_{j=1}^{n} p_j \mathbf{U}_{s_j}, \tag{7}$$

where $\mathbf{U}_{s_j}$ is the uniform distribution on s_j:

$$\mathbf{U}_{s_j}(x) = \begin{cases} \frac{1}{\mu(s_j)}, & \text{if } x \in s_j \\ 0, & \text{otherwise.} \end{cases} \tag{8}$$

Here the projection P and the injection R are defined for probability distributions, and extend naturally to L_1 functions on $\mathbb{X}$ and $\mathbb{S}$ respectively. The projection P is the left inverse of the injection R but not *vice versa*, namely $PR = I$ but $RP \neq I$.

This projection P and injection R can reduce the Fokker-Planck operator to a transition rate matrix on $\mathbb{S}$, and hence reduce the continuous-time stochastic hybrid system into a continuous-time Markov chain.

Theorem 1. *Let $\mathbb{S} = \{s_1, s_2, \ldots, s_n\}$ be a partition of the continuous state space $\mathbb{X}$ and P, R be the corresponding projection and injection defined in* (6)-(8)*. The Fokker-Planck operator given in* (3) *reduces to the transition rate matrix A of a continuous-time Markov chain on $\mathbb{S}$ by*

$$A = PLR \tag{9}$$

where the transition rate from state s_i to s_j at time t is given by

$$\begin{aligned} A_{ij} = &\int_{\partial s_i \cap \partial s_j} f(q,x)\mathrm{d}x \\ &+ \frac{1}{\mu(s_i)} \int_{s_i} r(q,x)\mathbf{I}_{h(q,x)\in s_j}\mathrm{d}x \end{aligned} \tag{10}$$

for $a, b = 1, \ldots, n$, where $\mathbf{I}_{h(q,x)\in s_j} = 1$ when $h(q,x) \in s_j$, and 0 otherwise.

One can informally interpret the transition rate between two partitions in the same location is the flux of $f(q,x)$ across the boundary and the transition rate between two different locations is the flux of $r(q,x)$.

Reducing MITL Specifications to checkable formulas on finite Markov chains

The observables on the continuous-time stochastic hybrid system reduce to the corresponding continuous-time Markov chain using the projection P. Let y be an observable on the continuous-time stochastic hybrid system with weight function $\gamma(q,x)$. We define a corresponding observable y' on the continuous-time Markov chain that derives from the model reduction procedure by

$$y'(t) = \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^d} \gamma(q,x) PF(t,q,x)\mathrm{d}x, \tag{11}$$

where $\mathcal{S}$ is a trajectory of the continuous-time Markov chain that obeys the distribution $p(t)$. We denote the corresponding observable on the CTMC by y' for any observable y on the continuous-time stochastic hybrid system.

For a given observable y with weight function $\gamma(q,x)$, the error of the projection P with respect to the observable y is defined by the maximal possible difference between y and y',

$$\Delta_y = \Big| \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^d} \gamma(q,x)(F(0,q,x) - RPF(0,q,x))\mathrm{d}x \Big|. \tag{12}$$

Remark 1. *When refining the partition of $\mathbb{X}$, $RP \to I$ in the weak operator topology, thus $\Delta_y \to 0$ for any given y.*

By the definition of Δ_y, we know that, at the initial time, the atomic propositions on the continuous-time stochastic hybrid system and the CTMC have the relations

$$y(0) > c \Longrightarrow y'(0) > c - \Delta_y, \tag{13}$$
$$y(0) < c \Longrightarrow y'(0) < c + \Delta_y, \tag{14}$$

and similarly,

$$y'(0) > c + \Delta_y \Longrightarrow y(0) > c, \tag{15}$$
$$y'(0) < c - \Delta_y \Longrightarrow y(0) < c. \tag{16}$$

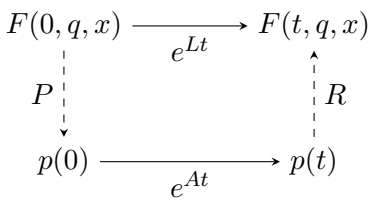

Figure 1: Diagram(noncommutative) for reduction error.

To derive the relations of the observables between the continuous-time stochastic hybrid system and the CTMC at any time, we define the reduction error of the observable y at time t due to the model reduction process by

$$\begin{aligned}\Theta_y(t) &= |y(t) - y'(t)| \\ &= \Big| \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^d} \gamma(q,x)(e^{Lt} - Re^{At}P)F(0,q,x)\mathrm{d}x \Big|,\end{aligned} \tag{17}$$

where $F(0,q,x)$ is an initial distribution of the continuous-time stochastic hybrid system and $y'(t)$ is the corresponding observable of $y(t)$ on the CTMC. This reduction error is illustrated in Fig. 1. Note that the diagram is *not* commutative; actually the difference between going along the two paths is related to the reduction error.

In general, the reduction error $\Theta(t)$ may not be bounded as $t \to \infty$. To find a sufficient condition for boundedness, we define the reduction error of the Fokker-Planck operator L by

$$\delta(t,q,x) = (L - RPL)e^{tRPL}F(0,q,x). \tag{18}$$

Accordingly, we define the integration of $\delta(t,q,x)$ with respect to the weight function $\gamma(q,x)$ by

$$\Lambda_y = \sup_{t\geq 0} \Big| \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^d} \gamma(q,x)(L - RPL)e^{tRPL}F(0,q,x)\mathrm{d}x \Big|, \tag{19}$$

which captures the maximal change of the time derivative of observable y.

A sufficient condition to find a uniform bound over time is that the reduction error of the Fokker-Planck operator $\delta(f(q,x))$ converges exponentially in time for any $f(q,x) \in m(\mathbb{X})$.

Definition 4. *For $\alpha > 0$, $\beta \geq 1$ and a given observable y, the continuous-time stochastic hybrid system is α-contractive with respect to y, if for any initial distribution function $F(0,q,x)$ on the state space, we have*

$$\begin{aligned}&\Big| \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^d} \gamma(q,x)e^{tL}\delta(t,q,x)\mathrm{d}x \Big| \\ &\qquad \leq \beta e^{-\alpha t} \Big| \sum_{q\in\mathcal{Q}} \int_{\mathbb{R}^d} \gamma(q,x)\delta(t,q,x)\mathrm{d}x \Big|.\end{aligned} \tag{20}$$

where $\delta(t,q,x)$ is given by (18)*.*

This contractivity condition, though it seems restrictive, is valid for a relatively wide range of systems including asymptotically stable systems. It is a commonly-used sufficient condition to guarantee the existence and uniqueness of an invariant measure for general dynamical systems, and the contractivity factor α is usually derived case-by-case. Using Definition 4, we derive the following theorem.

Theorem 2. *If the continuous-time stochastic hybrid system is α-contractive, then for any $t \geq 0$, the reduction error $\Theta_y(t)$ for an observable y satisfies*

$$\Theta_y(t) \leq \frac{\beta \Lambda_y}{\alpha} + \Delta_y. \tag{21}$$

Theorem 2 implies the following relations between the atomic propositions on the continuous-time stochastic hybrid system and the CTMC.

Theorem 3. *If the continuous-time stochastic hybrid system is α-contractive, then we have*

$$y(t) > c \Longrightarrow y'(t) > c - \Big(\frac{\beta \Lambda_y}{\alpha} + \Delta_y\Big), \tag{22}$$

$$y(t) < c \Longrightarrow y'(t) < c + \Big(\frac{\beta \Lambda_y}{\alpha} + \Delta_y\Big), \tag{23}$$

and similarly,

$$y'(t) > c + \Big(\frac{\beta \Lambda_y}{\alpha} + \Delta_y\Big) \Longrightarrow y(t) > c, \tag{24}$$

$$y'(t) < c - \Big(\frac{\beta \Lambda_y}{\alpha} + \Delta_y\Big) \Longrightarrow y(t) < c. \tag{25}$$

The above theorem gives the following result.

Theorem 4. *Given a MITL specification φ on the continuous-time stochastic hybrid system that is α-contractive, it can be strengthened to ψ by replacing the atomic propositions according to* (24)-(25)*. If ψ is true on the corresponding CTMC, then φ is true on the continuous-time stochastic hybrid system.*

We now turn to model checking of the MITL formula on an CTMC; our work provides the first such algorithm.

Statistical Model Checking of MITL

Here we present our basic algorithm on model checking CTMCs. And show that given a CTMC C and a MITL formula φ with atomic propositions $\mathtt{AP}_\varphi$, we can construct a timed automaton $T_{C,\mathtt{AP}_\varphi}$ by sampling such that reachable locations of this automaton at time t are labeled by the subset of atomic propositions in φ that are true in C at that time. By $[\![C, \mathtt{AP}_\varphi]\!]$ we denote the singleton set containing the unique signal induced by C and φ.

For simplicity, we assume $\mathtt{AP}_\varphi$ is a singleton and focus only on constructing $T_{C,\{P\}}$ for an atomic formula $P = \{y > c\}$. Let $f(t)$ be the set of atomic formulas that y satisfies at time t. Formally, $(y, c) \in f(t)$ iff $y(t) > c$. Also, let $T_{C,\{P\}}(t)$ be the set of reachable locations of $T_{C,\{P\}}$ at time t.

We assume $\delta' > 0$ and y^* (an estimation of the invariant distribution y^{inv}) are given such that for all $(y, c) \in \mathtt{AP}, |y^{\text{inv}} - c| > \delta'$ and $\|y^{\text{inv}} - y^*\|_1 < \frac{\delta'}{3}$. Since f converges to y^{inv} due to contractivity, we know $|y^{\text{inv}} - c| > \delta'$, for large enough t. This can be checked by a statistical algorithm $\mathcal{A}_1(y, y', \alpha, \delta)$ using standard techniques.

Using $\mathcal{A}_1$, we can find time T such that $\forall t \in [T, \infty), \|y_t - y^*\|_1 < \frac{\delta'}{3}$, which immediately give us $\forall t \in [T, \infty), \|y_t - y^{\text{inv}}\|_1 < \frac{2\delta'}{3}$, as given in Algorithm 1. For any $t \in [T, \infty)$ we know $-\frac{2\delta'}{3} \leq y^* - y^{\text{inv}} \leq \frac{2\delta'}{3}$ and hence $y^{\text{inv}} - c - \frac{2\delta'}{3} \leq y^* - c \leq y^{\text{inv}} - c + \frac{2\delta'}{3}$. Therefore, if $y^{\text{inv}} - c > \delta$ then $y^* - c > 0$. Similarly, if $y^{\text{inv}} - c < -\delta$ then $y^* - c < 0$. Note that exactly one of $y^{\text{inv}} - c > \delta$ and $y^{\text{inv}} - c < -\delta$ is true.

Furthermore, $y^* - c > 0$ and $y^* - c < 0$ cannot be both true. Therefore, $y^* - c > 0$ implies $y^{\text{inv}} - c > \delta$, $y^* - c < 0$ implies $y^{\text{inv}} - c < -\delta$, and $y^* - c$ is never zero. Finally, since y^* is precisely known, value of $y^* - c$ is also known precisely. Therefore, there will be no error (probabilistic or otherwise) in determining the value of f after time T.

ALGORITHM 1: Truncating time horizon

Data: CTMC (T, y_0), estimation of invariant distribution y^*, MITL formula φ, parameters α, γ, δ, and δ'
Function `DurationOfSimulation`
  $t \leftarrow 0$
  $\eta \leftarrow \max_{(y,c)\in \mathtt{AP}} \|y\|_1$
  while $\mathcal{A}_1\Big(y_t, y^*, \frac{1}{2}\min\{\alpha,\gamma\}, \frac{\delta'}{3\eta}\Big) = \mathtt{failed}$ **do**
    $t \leftarrow t + 1$;
  end
  return t

For any $\delta_1 \in \mathbb{R}_{>0}$, let $\Delta = \frac{\delta_1}{3\max\{|\dot{y}_i(t)||t\in[0,T]\}}$. Then, for any $t \in [0,T]$ and $t' \in [t-\Delta, t+\Delta] \cap [0,T]$, we have

1. if $y_i(t) - c > \frac{\delta_1}{3}$ then $y_i(t') - c > 0$,
2. if $y_i(t) - c < -\frac{\delta_1}{3}$ then $y_i(t') - c < 0$,
3. if $|y_i(t) - c| \leq \frac{2\delta_1}{3}$ then $|y_i(t') - c| \leq \delta_1$.

We partition $[0,T)$ into at least $\lfloor \frac{T}{2\Delta} \rfloor + 1$ intervals, each of size smaller than 2Δ. Let $[t_1, t_2]$ be one of these intervals. We then run $\mathcal{A}_1$, for $t = \frac{1}{2}(t_1 + t_2)$ as follows:

$$\mathtt{res}_1 = \mathcal{A}_1^{\delta_1/3}\left(y_i(t), c + \frac{\delta_1}{3}, \alpha', \gamma'\right),$$
$$\mathtt{res}_2 = \mathcal{A}_1^{\delta_1/3}\left(y_i(t), c - \frac{\delta_1}{3}, \alpha', \gamma'\right),$$

where $\mathcal{A}_1$ statistically check If $\mathtt{res}_1 = \mathtt{yes}$ then $\forall t' \in [t_1, t_2), (y_i(t') > c)$ holds with bounded error α'. Therefore, we set $T_{C,\{P\}}(t) = \{P\}$. If $\mathtt{res}_2 = \mathtt{no}$ then for any time $t' \in [t_1, t_2]$, we know $y_i(t') < c$ holds with bounded error α'. Therefore, we set $T_{C,\{P\}}(t) = \{\emptyset\}$. Otherwise, for any time t' in the interval, $|y_i(t') - c| \leq \delta_1$ with bounded error $\max(\alpha', \gamma')$. In this case, we set 1. $T_{C,\{P\}}(t) = \{q, q'\}$, 2. $\mathtt{L}(q) = \{P\}$ and $\mathtt{L}(q') = \emptyset$, 3. entry to q or q', and 4. switches between q and q' for arbitrary number of times, while their common invariant permits. The result of the above procedure $\mathtt{res} = \mathcal{A}^{\delta_1,\delta_2}(C, y_0, \varphi, \alpha, \beta)$ satisfies

$$\mathbb{P}[\mathtt{res} = \mathtt{no} \mid C \models \varphi] \leq \alpha \tag{26}$$
$$\mathbb{P}[\mathtt{res} = \mathtt{yes} \mid C \not\models \varphi] \leq \alpha \tag{27}$$

As for the `unknown` output, let $B^{\delta_1}(y)$ be the δ_1-ball centered at y in the L_∞ norm. The algorithm guarantees that

$$\mathbb{P}[\mathtt{res} = \mathtt{unknown}] \leq \alpha + \beta \tag{28}$$

for all $y' \in B^{\delta_1}(y)$.

ALGORITHM 2: Constructing the signal for atomic proposition P

$h \leftarrow \max\{|\dot{y}_i(t)| \mid t \in [0,T]\}$, $\Delta \leftarrow \frac{\delta_1}{3h}$, $n \leftarrow |\mathtt{AP}|\lceil \frac{T}{\Delta} \rceil$, $T_{C,\{P\}} \leftarrow$ an empty automaton, $\mathtt{X} \leftarrow \{t\}$, $q_{\text{last}} \leftarrow \perp$
forall $i \leftarrow 0$ *to* $\lfloor \frac{T}{\Delta} \rfloor$ **do**
  $\alpha' \leftarrow \min(\frac{\alpha}{4n}, \frac{\beta}{2n})$, $\beta' \leftarrow \frac{\beta}{n}$
  $\mathtt{res}_1 \leftarrow Algorithm_1^{\delta_1/3}\big(y_i((i+\frac{1}{2})\Delta), c + \frac{\delta_1}{3}, \alpha', \beta'\big)$
  $\mathtt{res}_2 \leftarrow Algorithm_1^{\delta_1/3}\big(y_i((i+\frac{1}{2})\Delta), c - \frac{\delta_1}{3}, \alpha', \beta'\big)$
  add a new location q to $\mathtt{Q}$
  if $\mathtt{res}_1 = \mathtt{yes}$ **then**
    $\mathtt{L}(q) \leftarrow \{P\}$
  else if $\mathtt{res}_2 = \mathtt{no}$ **then**
    $\mathtt{L}(q) \leftarrow \emptyset$
  else
    $\mathtt{L}(q) \leftarrow \mathtt{unknown}$
  $\mathtt{I}(q) \leftarrow 2i\Delta \le t < 2(i+1)\Delta$
  if $q_{\text{last}} \neq \perp$ **then**
    $\mathtt{E} \leftarrow \mathtt{E} \cup \{(q_{\text{last}}, q, \emptyset)\}$
  else
    $\mathtt{Q}^{\mathsf{init}} \leftarrow \{q\}$
  $q_{\text{last}} = q$
end
add a new location q to $\mathtt{Q}$
$\mathtt{I}(q) \leftarrow \mathtt{true}$, $\mathtt{Q}^{\mathtt{final}} \leftarrow \{q\}$
$\mathtt{E} \leftarrow \mathtt{E} \cup \{(q_{\text{last}}, q, \emptyset), (q, q, \emptyset)\}$
if $y^{\text{inv}} > c$ **then**
  $\mathtt{L}(q) \leftarrow \{P\}$
else
  $\mathtt{L}(q) \leftarrow \emptyset$
$T_{C,\{P\}} \leftarrow$ replace any `unknown` location in $\mathtt{Q}$ with q and q' labeled $\{P\}$ and $\emptyset$. Duplicate edges from/to q and q' accordingly
Add $(q, q', \emptyset)$ and $(q', q, \emptyset)$ to $\mathtt{E}$ for every split locations in the previous step.
return $T_{C,\{P\}}$

Definition 5. *For any $\epsilon > 0$ let $y + B_\epsilon$ be the set of observables achieved by slightly perturbing y. Let C_ϵ be any object with observables in the set $y + B_\epsilon$. We say satisfaction relation of CTMC C and MITL formula φ is* ϵ-robust*, if one of the following is true: 1. For all y' induced by C_ϵ we have $y' \models \varphi$, or 2. For all y' induced by C_ϵ we have $y' \not\models \varphi$. We say satisfaction relation is* robust*, if it is ϵ-robust for some $\epsilon > 0$.*

With this definition in hand, we have shown that given any CTMC C and MITL formula φ, if C is robust on φ, iteratively reducing δ_1 using our algorithm guarantees that it will eventually return an answer which is not `unknown` while satisfying conditions (26) and (27).

To summarize, this subsection has presented in some detail the project research techniques, and some results, on verification of hybrid systems. The general approach is believed to be scalable to large systems, and as mentioned in the executive summary we have carried out and published some application-oriented case studies that support this claim. An important future direction is to bring synthesis of closed-loop control and supervisory policies into the framework.

3.2 Lattice continuous-time Markov jump processes

During the project we have made significant advances in developing new techniques for simulation of large-scale stochastic systems. First, we have developed several anticorrelated variance reduction techniques for the simulation and estimation of discrete-time lattice Markov jump processes; such systems approximate continuous-time Markov processes via so-called tau-leaping approximations. These discrete time methods were 10 to 1000 times faster than conventional methods on the spectrum of numerical experiments investigated; the test systems ranged from relatively simple linear models, to complex nonlinear ones. Following on this work, in contrast, we have further developed *exact* continuous time approaches leveraging our general approach from the discrete time setting. These continuous time methods reduce the variance of continuous-time Monte Carlo for Markov jump process systems, and we have rigorously constructed antithetic Poisson processes, and have analytical demonstrations of negative correlation between pairs. Indeed, the primary contribution of this new work is the construction of antithetic Poisson processes and their use in the random time-change representation for the purpose of exact simulation and variance reduced mean estimation of continuous-time processes. Considerable effort has been devoted to developing computer codes that run on large-scale distributed computing platforms such as that of a supercomputer; we have applied these new methods to several initial examples using a distributed platform and seen significant improvements (between 1 and 2 orders of magnitude) in computational efficiency over the standard methods currently used.

3.3 Stratified model checking

A significant project accomplishment has been progress on the problem using stratified samples to check probabilistic computation tree logic formulas on discrete-time Markov chains. Compared to previous statistical verification methods using naive sampling, stratified samples are repellent to each other, con- sequently, their sample mean has higher concentration. By large sample approximation, we have developed a sequential probability ratio testing (SPRT) algorithm to check PCTL formulas, and have showed that this new algorithm reduces of the number of samples needed for a given confidence level on several benchmarks. More generally, statistical verification of probabilistic temporal logic formulas on probabilistic systems has flourished during the past decade, due to its scalability to large-scale real-world problems with complicated dynamics. The general idea is to treat the satisfaction problem of a probabilistic temporal logic formula as a hypothesis testing problem. Therefore, by drawing a sufficient number of samples from the underlying probabilistic system and using proper statistical inference, the satisfaction of the probabilistic temporal logic formula can be determined with given confidence level. Previously, most statistical verification algorithms are based on naive sampling, with the underlying probabilistic system treated as a black-box that generates sample paths and a new sample is drawn from it at each round. Consequently, the samples are independently and identically distributed (i.i.d.). Though this approach has the advantage of being able to handle systems with unknown dynamics, it is usually not efficient, when the underlying probabilistic system is known. In our research, we have developed a method of statistically verifying probabilistic computational tree logic (PCTL) formulas on a finite time horizon. Instead of drawing one sample at each round, m-samples that are repellent to each other are generated simultaneously. This leads to faster exploration of the path space and less statistical error in estimation and inference. Initially we have focused on PCTL formulas, restricted to situations of a simple formula without probabilistic operator. Namely, the correctness of can be determined from a single sample path. PCTL formulas in general form with nested probabilistic operators can also be handled and we are currently considering generalizations.

3.4 Benchmark implementation

During the project we have applied our computational approach to a significant community benchmark challenge problems that captures features of realistic automotive dynamics. We have statistically verified the most complicated of the powertrain control models proposed in these community benchmark problem (provided by Toyota), which includes features like delayed differential and difference equations, look-up tables, and highly non-linear dynamics. Our results show that for at least 98% of the possible initial operating conditions the desired properties hold. These are the first successful verification results for this model using any verification technique, statistical or otherwise. More specifically, the powertrain control problem is one of regulating the air-to-fuel ratio in a automotive engine. A series of models of such controllers, with increasing levels of sophistication and fidelity to real-world designs, have been recently proposed by researchers as challenge problems for today's verification technologies. Ever since these models were proposed 3 years ago, they have served as a high water mark for evaluating verification tools and techniques for cyberphysical systems. In our work, unlike previous unsuccessful approaches to verifying this model class, we developed statistical model checking methods to analyze the system traces. We established formally that the system model behaves correctly for most settings of the initial parameters. Statistical techniques are used to bound sample sizes, and we develop new ideas to enable checking whether a single sample execution satisfies the prescribed STL property. Mathematically, an execution of the benchmark powertrain controller is a function that species for each time instant, the state of the system. We developed conditions on how executions must be sampled (based on the correctness property) and an approximate satisfaction relation that together guarantee that if discretely sampled executions satisfy an STL property, then so does the actual, continuous execution.

3.5 Stochastic control

We have developed a technique for applying our anti-correlated sampling techniques to model predictive control (MPC), focusing on stochastic MPC, which is usually characterized by a combination of stochastic dynamics (possibly characterized by stochastic noise) and/or probabilistic constraints. Specifically, the research focused on techniques that leveraged Monte Carlo simulation (often called scenario-based methods in this context) in service of approximating solutions to the finite-horizon open-loop stochastic optimization component of MPC implementation. We leveraged our previous work on anticorrelated simulation of Markov jump processes (a broad class of potentially nonlinear and non-Gaussian Markov processes on a countable state-space) to reduce the variance of mean estimates of the finite-horizon expected cost. The goal was to allow for a reduced Monte Carlo budget to achieve the same or better performance of stochastic MPC. The main contributions of the work were the presentation of an algorithm for variance-reduced stochastic model predictive control of Markov jump processes, and its demonstration on a non-linear reaction network. The stochastic dynamical setting we considered was the Markov jump process, a continuous-time, countable statespace process that experiences transitions that can be classified by a finite number of reaction channels. This class is a broad collection of systems, and appears commonly in models for chemical reaction systems, gene regulation systems, and atmospheric aerosol simulation. To simulate such models, the Gillespie stochastic simulation algorithm (SSA) is frequently used. However, when the frequency of at least some of the reaction events is relatively large, SSA can quickly become expensive and impractical. In these cases, a discrete-time approximation method known as tau-leaping is often used. The tau-leaping method involves simulating the system at discrete time intervals, where the number of transitions due to each continuous time Poisson process during that interval are approximated by the sampling of an appropriately chosen Poisson random variable. The convergence properties of the tau-leaping algorithm have been rigorously proven

and many variants exist, including implicit tau-leaping and adaptive stepping methods. In our work, we restricted attention to a commonly used explicit, fixed step size approach. In the work we demonstrated that anticorrelated ensembles can be drawn for open loop simulations of the process during MPC to produce reduced error estimates of the true average cost.

3.6 Hybrid sampled-data

We have also investigated the stability and L_2-gain properties of a class of hybrid systems that exhibit linear flow dynamics, periodic time-triggered jumps and arbitrary nonlinear jump maps. This class of hybrid systems is relevant for a broad range of applications including periodic event-triggered control, sampled-data reset control, sampled-data saturated control, and certain networked control systems with scheduling protocols. For this class of continuous-time hybrid systems we have developed new stability and L_2-gain analysis methods. Inspired by ideas from lifting we show that the stability and the contractivity in L_2-sense of the continuous-time hybrid system is equivalent to the stability and the contractivity of an appropriate discrete-time nonlinear system. These new characterizations generalize earlier (more conservative) conditions provided in the literature. We have provided explicit examples of reset control and an event-triggered control applications, for which stability and contractivity in L_2-sense is the same as stability and contractivity in ℓ_2-sense of a discrete-time piecewise linear system, that the new conditions are significantly less conservative than the existing ones in the literature. And we have shown that the prior conditions in the literature can be reinterpreted as a conservative ℓ_2-gain analysis of a discrete time piecewise linear system based on common quadratic storage/Lyapunov functions. These new insights were obtained by the adopted lifting-based perspective on this problem, which leads to computable ℓ_2-gain (and thus L_2-gain) conditions.

3.7 Counter-example guided abstraction-refinement framework.

We have developed the theoretical foundations and software tools to automatically verify hybrid systems using abstractions. Abstractions play an important role in the verification of cyber-physical systems, where complex continuous dynamics are abstracted into simpler dynamics that are amenable to automated analysis. This is because the general problem of safety verification is undecidable even for very simple class of continuous dynamics. The success of the abstraction based method depends on finding the right abstraction, which can be difficult. One approach that tries to address this issue is the counter example guided abstraction refinement (CEGAR) framework that tries to automatically discover the right abstraction through a process of progressive refinement based on analyzing spurious counter examples in abstract models. Over the past three years, we have been developing a CEGAR framework to automatically analyze hybrid systems, identifying algorithms to automatically construct coarse abstractions, and refine the abstractions based on model checking the abstract designs. We have mathematical theorems that characterize progress achieved during each iteration of abstraction-refinement, along with a characterization of termination under special circumstances. We developed a new algorithm for model checking safety problems of hybrid automata with affine dynamics and rectangular constraints in a CEGAR framework. We show that our algorithm is sound and the approach has been built into a software tool called HARE (Hybrid Abstraction-Refinement Engine) that can automatically verify safety properties of hybrid systems with non-linear dynamics. Experimental evaluation of the tool has demonstrated that the approach works well in practice — HARE can analyze designs that cannot be verified by state-of the art verification tools, as well as work much faster on simpler systems. We also compared the performance of our tool with several state-of-the-art tools, which further demonstrated its advantages.

3.8 Syntactic and Semantic Perturbations of Hybrid Automata

Hybrid automata provide a mathematical framework in which to model systems consisting of a digital controller interacting with a continuously evolving physical process, and the models have discrete modes corresponding to phases in the digital controller. Transitions between modes model discrete changes to actuator inputs from the controller based on sensory feedback. Formal models of cyberphysical systems are typically "best effort" descriptions that may not be completely faithful to the actual system. There are several sources of inaccuracies. For instance, environment parameters, like network latency, are estimated based on extensive experimentation. Differential equations governing the behavior of the physical plant maybe imprecise, either because of limitations in our mathematical understanding of the physics, or because of a conscious effort to construct a tractable model by approximating. Sensor and actuator delays might either be unpredictable or have been ignored. Finally, inaccuracies in sensor input to the controller may not have been faithfully modeled.

For these reasons, a system modeled by hybrid automaton $\mathcal{H}$, may, in practice, behave like the automaton $\mathcal{H}^\delta$ which is obtained from $\mathcal{H}$ by syntactically perturbing constants, constraints on mode switches, and flow equations governing continuous evolution, by some $\delta > 0$. A natural question to ask is if the automaton $\mathcal{H}$ and its perturbation $\mathcal{H}^\delta$ are semantically close (in some well defined sense). Can a perturbed automaton $\mathcal{H}^\delta$ be arbitrarily close to $\mathcal{H}^\delta$? The challenge in answering this question lies in the presence of discrete mode changes—small changes to the behavior of a hybrid automaton could result in transitions becoming enabled that yield unexpected behavior in the perturbed automaton.

It is known that general syntactic perturbations can result in models that are not semantically close. Our main result is that for a restricted but fairly general class of hybrid automata, that can include hybrid automata with highly non-linear and non-deterministic dynamics, syntactic perturbations are closely related to semantic perturbations. We consider hybrid automata $\mathcal{H}$ all of whose components, like flows and invariants in modes, and guards and resets on transitions, are described using formulas in first-order logic over reals built from constraints of the form $f \geq 0$, where f is a continuous function, and using conjunction, disjunction, and first order quantification. For a formula φ in this logic, its perturbation δ by, φ^δ, is the formula obtained by replacing all atomic constraints $f \geq 0$ in φ by $f + \delta \geq 0$. Using the notion of perturbation of a constraint, we define $\mathcal{H}^\delta$ to be the hybrid automaton obtained from $\mathcal{H}$ by perturbing all constraints φ appearing $\mathcal{H}$ by δ. Using this definition we are able to prove a rigorous result about the sets of automata such perturbations sweep out. Descriptively, the result can be interpreted as follows. Let us consider the function F which maps the automaton $\mathcal{H}$ to its transition system semantics $F(\mathcal{H})$. Our results can be seen as saying that F is a "continuous" map with respect to the metric on the hybrid automata induced by perturbations δ.

The results of this part of the project on relating syntactic and semantic perturbations have significant implications in satisfiability checking as well as verification. In particular, applications include CEGARs (described above) and decision procedures.

4 Awards and Recognitions

Geir Dullerud:

- Keynote lecture at IEEE Symposium on Complex Systems and Cybernetics, Guangzhou, China, 2017;
- Keynote lecture at Design and Analysis of Robust Systems (DARS) Workshop, Computer Automated Verification (CAV) Conference, 2017;
- Plenary lecture at the International Workshop on Operator Theory and Applications (IWOTA), 2016;
- W. Grafton and Lillian B. Wilkins Endowed Professor of Mechanical Engineering, University of Illinois, Urbana-Champaign, 2016.

Mahesh Viswanathan:

- Plenary lecture at Workshop on Probabilistic Reasoning and Formal Methods, International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS), 2017;
- Keynote lecture at Trends and Challenges in Quantitative Verification, Mysore Park Workshop, 2016;
- Promoted to the rank of full Professor in the Department of Computer Science at the University of Illinois, Urbana-Champaign, 2016.

5 Personnel Supported

Geir E. Dullerud	Professor, University of Illinois at Urbana-Champaign
Raymond Essick	PhD Student, University of Illinois at Urbana-Champaign
Peter Maginnis	PhD Student, University of Illinois at Urbana-Champaign
Joao Porto	PhD Student, University of Illinois at Urbana-Champaign
Nima Roohi	PhD Student, University of Illinois at Urbana-Champaign
Mahesh Viswanathan	Professor, University of Illinois at Urbana-Champaign
Yu Wang	PhD Student, University of Illinois at Urbana-Champaign
Matthew West	Associate Professor, University of Illinois at Urbana-Champaign