Workshop on
Formal Verification and Synthesis for Hybrid Systems

The University of Texas at Austin
AFOSR Contract # FA9550-17-1-0229
End of the Performance Period: August 26, 2017

Prepared by:

Dr. Ufuk Topcu
The University of Texas at Austin
utopcu@utexas.edu
September 24, 2017

Abstract

This report aims to reflect on the outcomes of the Workshop on Formal Verification and Synthesis of Hybrid Systems held on June 1-2, 2017 at the University of Texas at Austin.

Contents

1 Objectives **1**

2 Accomplishments **2**

3 Personnel Supported **9**

4 Technical Publications **9**

5 Interactions/Transitions **9**

6 Patent Disclosures **10**

7 Honors **10**

1 Objectives

Models and specifications for modern systems are typically hybrid in nature. Autonomous systems, for example, are governed by not only the underlying physics but also their computing and communication platforms. They are constrained by not only low-level actuation limitations and stability requirements but also high-level mission specifications.

It has long been recognized that hybrid systems cut across multiple disciplines. However, majority, if not all, of synthesis and verification methods proposed for hybrid systems ironically fail to be hybrid. Each discipline attempts to fit the underlying problems to its toolset through unnatural transformations (e.g., finite-state discretization) and restrictions. Principles and tricks that have been successfully utilized in one domain are naively applied on modeling and specification languages for which they were never intended. The outcome of this monolithic approach has been methods mostly with limited applicability and/or scalability.

The objective of the proposed workshop is to challenge the conventional, monolithic view by arguing that verification and synthesis methods for hybrid systems should be hybrid themselves. The questions to be explored include, but are not limited to, the following: How can we cope with the scalability barrier and develop techniques that can handle systems of arbitrary size? How can we fuse geometric, algebraic and computational approaches toward verification and synthesis of hybrid systems without explicit discretization of continuous dynamics? How can we develop techniques that provide rigorously provable guarantees on system properties and their sensitivities with respect to aspects that are ignored in the abstraction used for verification or synthesis? How can we develop methods that account for the computational resources available and the interpretability and determinism of the outcomes? What are fundamental advances in mathematics, control theory and computer science that promise novel approaches?

2 Accomplishments

The workshop was held with the participation of the following individuals: Jeremy Avigad, Georgios Fainekos, Meeko Oishi, Ufuk Topcu, Danielle Tarraf, Andy Teel, Brandon Hencey, Fred Leve, Melkior Ornik, Min Wen, Tony Bloch, Chris Petersen, Ivan Papusha, Laura Humphrey, Raphael Jungers, Sean Gao, Ufuk Topcu, Will Curtis, Amir Ali Ahmadi, Geir Dullerud, Stanley Bak, Hasan Poonawala, Jared Culbertson, Mohamadreza Ahmadi, Tristan Nguyen and Stathis Bakolas.

The agenda of the workshop is shown below.

Workshop on Verification and Synthesis of Hybrid Systems
University of Texas at Austin
Peter O'Donnell Jr. Building (POB)

June 1st

9-9:15 -- Welcome and introduction (Leve, Sanfelice, Topcu)
9:15-10:15 -- Overviews (15 min each + 5 min discussion)
Controls – Teel + Oishi
Applied math – Avigad + Bloch
Formal methods – Fainekos + Topcu
10:15-10:45 -- Break
10:45-12 -- Lightning talks: 3 minutes and 1-2 slides per participant focusing the following.
12-1:30 -- Lunch
1:30-2:15 -- Challenge the other side (10 min each + 5 min discussion)
Controls – Dullerud + Jungers
Applied math – Lerman + Ames
Formal methods -- Gao + Mitra
2:15-2:30 -- Introduction of the breakout sessions
Break-out 1 – Lead: Tarraf
Break-out 2 – Lead: Sanfelice
Break-out 3 – Lead: Ahmadi
2:30-2:45 -- Break
2:45-4 -- Break-out sessions
4-5 -- Summary of the break-out session discussions and discussion (Tarraf, Sanfelice, Ahmadi)
6 – No-host dinner (Location: TBD)

June 2nd

9-10:15 -- AF talks (75/number of AF participants – 1 min for each presenter)
Bak, Cambie, Culbertson, Curtis, Hencey, Humphrey, Leve, Nguyen, Petersen
10:15-10:45 -- Break
10:45-12 -- Non-AF talks (75/number of AF participants – 1 min for each presenter)
All non-AF participants are welcome
12-1 -- Working lunch in three groups: Write your own objective and ambitious solicitation
Solicitation group 1 – Lead: Fainekos
Solicitation group 2 – Lead: Oishi
Solicitation group 3 – Lead: Avigad
1-1:30 -- Break
1:30-2:15 -- Discuss and merge the solicitations (Fainekos, Oishi, Avigad)
2:15-2:30 -- Wrap-up

Figure 1: The agenda of the workshop.

Along with the agenda, the information in the following figure was shared with the participants in order to facilitate the discussions during the workshop.

Workshop on Verification and Synthesis of Hybrid Systems
University of Texas at Austin
Peter O'Donnell Jr. Building (POB)

June 1st

9-9:15 -- Welcome and introduction (Leve, Sanfelice, Topcu)
9:15-10:15 -- Overviews (15 min each + 5 min discussion)
Controls – Teel + Oishi
Applied math – Avigad + Bloch
Formal methods – Fainekos + Topcu
10:15-10:45 -- Break
10:45-12 -- Lightning talks: 3 minutes and 1-2 slides per participant focusing the following.
12-1:30 -- Lunch
1:30-2:15 -- Challenge the other side (10 min each + 5 min discussion)
Controls – Dullerud + Jungers
Applied math – Lerman + Ames
Formal methods -- Gao + Mitra
2:15-2:30 -- Introduction of the breakout sessions
Break-out 1 – Lead: Tarraf
Break-out 2 – Lead: Sanfelice
Break-out 3 – Lead: Ahmadi
2:30-2:45 -- Break
2:45-4 -- Break-out sessions
4-5 -- Summary of the break-out session discussions and discussion (Tarraf, Sanfelice, Ahmadi)
6 – No-host dinner (Location: TBD)

June 2nd

9-10:15 -- AF talks (75/number of AF participants – 1 min for each presenter)
Bak, Cambie, Culbertson, Curtis, Hencey, Humphrey, Leve, Nguyen, Petersen
10:15-10:45 -- Break
10:45-12 -- Non-AF talks (75/number of AF participants – 1 min for each presenter)
All non-AF participants are welcome
12-1 -- Working lunch in three groups: Write your own objective and ambitious solicitation
Solicitation group 1 – Lead: Fainekos
Solicitation group 2 – Lead: Oishi
Solicitation group 3 – Lead: Avigad
1-1:30 -- Break
1:30-2:15 -- Discuss and merge the solicitations (Fainekos, Oishi, Avigad)
2:15-2:30 -- Wrap-up

Figure 2: Supplementary information on the sessions of the workshop.

The main outcome of the workshop has been the reports on the challenges and opportunities in the verification and synthesis for hybrid systems. These reports are products of discussions among multiple subgroups. The discussion leads and the workshop chair also produced a consolidates report.

Consolidated report
Background:

Hybrid systems are used to model systems with discrete and continuous components, but their application to the design of controllers for autonomous systems poses special challenges. Such systems are often deployed in scenarios with features and behaviors that cannot easily be predicted in advance. Moreover, conventional methods typically do not scale well to complex systems and environments, including multiagent systems. Continuous dynamics are often discretized to facilitate the analysis, but these discretizations can mask important qualitative features and lead to a blow-up in the size of the representation. Conversely, symbolic methods are particularly susceptible to scaling problems, and face uncomputability barriers, even for the simplest nontrivial models.

Progress towards the analysis of autonomous systems and the synthesis of effective controllers will require refining mathematical models to model software, controllers, and their environments more faithfully; adapting methods from the analysis of complex hardware and software systems to hybrid autonomous systems; and developing new computational methods for the analysis, synthesis, and verification of controller design. Several fundamentally different approaches have been proposed in the hybrid systems community, and it is far from clear which are the most effective, or what specific features of an application determine which techniques are most relevant.

Challenges:

While considerable progress has been made in the development of theory and tools for problems in reachability, viability, invariance, formal specifications, and stability, these approaches cannot accommodate the scale or complexity of realistic hybrid systems.

Uncertainty in hybrid systems is pervasive, however few approaches exist to rigorously address the effect of uncertainty on verification and synthesis. Assurances of robustness are further complicated by the need for scalability and eventual real-time operation in environments with uncertainty.

While the lack of full knowledge of the uncertainty may unavoidably lead to designs for worst-case situations, statistical/probabilistic methods should be exploited to reduce conservativeness. Further, uncertainty may be present not only in modeling, but also in probabilistic performance or safety specifications.

Probabilistic and statistical methods can provide quantitative answers whose confidence-level can be chosen by using the appropriate computational resources available; this is in contrast with other methods that can provide boolean ?yes? or ?no? answers, but provide no information if the computational resources are insufficient.

The design of emerging autonomous systems require combinations and compositions of heterogeneous modules for which different levels of information are available, including physics, computation, and networking. For traditional kinds of modules, a complete mathematical model may be available, but for some, the mathematical model may be available but too big or complicated to be manipulated computationally, and other modules may be only available as an executable black-box without a mathematical description.

One very foundational limitation of the existing verification methods is their reliance to accurate system models which are known in advance. Therefore, even though offline model based verification approaches can tremendously help in system design, there is also a need for methods which are online and/or can analyze the system safety and performance without relying on an accurate fixed model of the system (both in physics and computation). The development of analysis algorithms that can exploit available model information (even when it is quite limited), in combination with system identification and learning approaches for obtaining integrated identification-verification algorithms, are also important. New theories

are needed for analyzing the efficiency and optimality of these approaches, as the traditional computational resources (CPU time and memory) have to be augmented with new constraints such as data availability, sample efficiency, and sensor resolution.

In the case of autonomous systems utilizing uncertain sensor modalities, we need to develop new control frameworks which not only receive data from the sensors, but also enable active sensing by providing feedback to the sensors. In such an active sensing framework it is expected that perception errors will be reduced substantially, increasing thus the runtime performance and safety of the system.

In order to enable true fully automated and scalable verification, the new methods should be cleanly implemented into software tools without hidden parameters so that the implementation can be independently trusted and certified. Additionally, demonstrations that show application of the new methods to industrially relevant cases studies should be developed.

Research Concentration Areas:

Suggested research directions can include, but are not limited to, the following: 1) developing models and specification languages to describe autonomous systems and controllers, providing hierarchical descriptions and abstractions that support decompositional analysis and facilitate an analysis of continuous aspects of a model together with their discretizations; 2) extending deductive verification to richer models, for example, by refining Lyapunov methods; 3) developing decomposition methods for analyzing such systems in modular terms; 4) improving scalability of optimization-based techniques for formal verification; 5) introducing and taking advantage of stochastic components through probabilistic specifications and verification and synthesis methods; 6) developing methods for ensuring correctness of control systems by design, and taking advantage of user interaction and intervention in the design process; 7) developing new mathematical and computational methods for on-the-fly verification of hybrid systems; 8) abstractions and methods that scalably account for the limitations in sensory data and 9) developing software tools and benchmarks to compare different approaches.

Report 1

Lead: Georgios Fainekos

Contributors: Danielle Tarraf, Andy Teel, Brandon Hencey, Fred Leve, Melkior Ornik, Min Wen, Tony Bloch

Autonomous systems cannot become a widely adopted technology without a minimum level of confidence and guarantees on the system behavior and performance. Currently, the verification and validation methods which are able to analyze system behavior with respect to some safety requirements have several limitations. The most important limitation is that verification methods cannot handle systems with more than 10 dimensions in a fully automated way. Typically, extensive computational resources along with deep knowledge about the system and the verification methodology are required in order to be able to scale the verification methods to more complex systems similar to what is published in the literature. But even in these special cases, the handled systems are not of industrial size and complexity.

In order to enable true fully automated and scalable verification, we need new methods which:

- generalize existing successful specialized case studies to frameworks which can be applied to complex systems,
- smartly discretize the problem by automatically exploiting system structure and by allowing adaptive and hierarchical discretization,

- can relate models at different levels of fidelity so that analysis results on simpler models can be propagated to more complex models, and
- can be cleanly implemented into software tools without hidden parameters and black-box components so that the implementation can be independently trusted and certified.

In addition, one very foundational limitation of the existing verification methods is their reliance to accurate system models which are known in advance. Therefore, even though offline model based verification approaches can tremendously help in system design, there is also a need for methods which are online and/or can analyze the system safety and performance without relying on an accurate fixed model of the system (both in physics and computation). As an example, consider a space satellite whose physical behavior is changed by a minor collision with another space object and whose computer has been affected by radiation effects. In such problems, hard bounds on the uncertainty lead to extremely conservative results which may not be practically interesting, or they may even overconstrain the problem rendering it thus infeasible.

In order to enable the deployment of autonomous systems, we need to:

- develop online verification methods which check the system correctness at runtime while taking the computational resources of the system into account, and
- move away from worst-case system analysis and prove system properties within probabilistic frameworks reasoning about the expected system behavior.

Finally, in the special case of autonomous systems utilizing uncertain sensor modalities, we need to develop new control frameworks which not only receive data from the sensors, but also enable active sensing by providing feedback to the sensors. In such an active sensing framework it is expected that perception errors will be reduced substantially, increasing thus the runtime performance and safety of the system.

Report 2

Lead: Jeremy Avigad

Contributors: Chris Petersen, Ivan Papusha, Laura Humphrey, Raphael Jungers, Sean Gao, Ufuk Topcu, Will Curtis, Amir Ali Ahmadi

Background: Hybrid systems are used to model systems with discrete and continuous components, but their application to the design of controllers for autonomous systems poses special challenges. Such systems are often deployed in scenarios with features and behaviors that cannot easily be predicted in advance. Moreover, conventional methods typically do not scale well to complex systems and environments, including multiagent systems. Continuous dynamics are often discretized to facilitate the analysis, but these discretizations can mask important qualitative features and lead to a blow-up in the size of the representation. Conversely, symbolic methods are particularly susceptible to scaling problems, and face uncomputability barriers, even for the simplest nontrivial models.

Progress towards the analysis of autonomous systems and the synthesis of effective controllers will require refining mathematical models to model software, controllers, and their environments more faithfully; adapting methods from the analysis of complex hardware and software systems to hybrid autonomous systems; and developing new computational methods for the analysis, synthesis, and verification of controller design. Several fundamentally different approaches have been proposed in the hybrid systems community, and it is far from clear which are the most effective, or what specific features of an application determine which techniques are most relevant.

Objective: The goal of this solicitation is to develop methods to synthesize controllers for complex hybrid systems that scale with the system?s size, the level of uncertainty, and

the richness of the requirements; and to verify their correctness, robustness and performance, especially in unpredictable, mission-critical scenarios.

Research Concentration Areas: Suggested research directions can include, but are not limited to, the following: 1) developing models and specification languages to describe autonomous systems and controllers, providing hierarchical descriptions and abstractions that support decompositional analysis and facilitate an analysis of continuous aspects of a model together with their discretizations; 2) extending deductive verification to richer models, for example, by refining Lyapunov methods; 3) developing decomposition methods for analyzing such systems in modular terms; 4) improving scalability of optimization-based techniques for formal verification; 5) introducing and taking advantage of stochastic components; 6) developing methods for ensuring correctness of control systems by design, and taking advantage of user interaction and intervention in the design process; 7) developing new mathematical and computational methods to address the specific challenges enumerated above; and 8) developing tools and benchmarks to compare different approaches.

Report 3

Lead: Meeko Oishi

Contributors: Geir Dullerud, Stanley Bak, Hasan Poonawala, Jared Culbertson, Mohamadreza Ahmadi, Tristan Nguyen, Stathis Bakolas

There is a clear need for methods and tools for verification and synthesis of hybrid systems, that are rigorous, computationally efficient, and can cope with uncertainty. Approaches are sought which can address ?hard? problems at the intersections of control theory, formal methods, and applied mathematics. While considerable progress has been made in the development of theory and tools for problems in reachability, viability, invariance, formal specifications, and stability, these approaches cannot accommodate the scale or complexity of realistic hybrid systems. Dramatic improvements are necessary in modeling, analysis, and synthesis of algorithms that meet complex specifications. Such approaches should ideally provide scalability, exactness, robustness, and support algorithmic approaches that are computationally feasible.

Uncertainty in hybrid systems is pervasive, however few approaches exist to rigorously address the effect of uncertainty on verification and synthesis. Assurances of robustness are further complicated by the need for scalability and eventual real-time operation in environments with uncertainty. While the lack of full knowledge of the uncertainty may unavoidably lead to designs for worst-case situations, statistical/probabilistic methods should be exploited to reduce conservativeness. Further, uncertainty may be present not only in modeling, but also in probabilistic performance or safety specifications. Methods must be developed that can replace computationally hard, exact formulations with provably good approximations of lower complexity. However, although numerical validation may provide satisfaction of performance requirements and constraints, rigorous, analysis based, mathematical descriptions of the properties attained by the algorithms are paramount.

Probabilistic and statistical methods can provide quantitative answers whose confidence-level can be chosen by using the appropriate computational resources available; this is in contrast with other methods that can provide boolean ?yes? or ?no? answers, but provide no information if the computational resources are insufficient. To advance validation and verification of autonomous systems, new frameworks and methodologies are needed that follow the sophisticated principle-based probabilistic paradigm. The design of emerging autonomous systems require combinations and compositions of heterogeneous modules for which different levels of information are available,

including physics, computation, and networking. For traditional kinds of modules, a complete mathematical model may be available, but for some, the mathematical model may be available but too big or complicated to be manipulated computationally, and other modules may be only available as an executable black-box without a mathematical description. Thus, the development of analysis algorithms that can exploit available model information (even when it is quite limited), in combination with system identification and learning approaches for obtaining integrated identification-verification algorithms, are also important. New theories are needed for analyzing the efficiency and optimality of these approaches, as the traditional computational resources (CPU time and memory) have to be augmented with new constraints such as data availability, sample efficiency, and sensor resolution.

Fundamental questions remain in the following broad areas.

Robustness for hybrid systems

- What robustness guarantees are possible with reachability analysis? What are relevant mathematical properties of robust or probabilistic reachable and viable sets?
- Are there fundamental differences in analysis precision when different types of modelling error are considered (input uncertainties, dynamics uncertainties, measurement noise, poor characterization of uncertainties)?
- What are appropriate metrics to characterize the effect of disturbances in hybrid systems and to allow for rigorous comparison to the behavior in nominal conditions? Are practical properties acceptable?
- How does the propagation of uncertainty affect stability, reachability, viability, controllability, and robustness of hybrid systems? What are the benefits of a robust approach (as opposed to a probabilistic or statistical approach), particularly when considering nonparametric models?

Scalability in reachability, viability, and invariance calculations

- What are the limiting factors in reachability analysis, both in terms of the size of the continuous and discrete space? Can analysis leverage multiple levels of abstraction within the same analysis run (not just the original system and a discrete abstraction)? Can this abstraction hierarchy be automated?
- Can hierarchical abstractions be generated through reachability analysis (for specific classes of systems, sets, or properties) by e.g., exploiting branch-and-bound algorithms?
- How can safety verification methods, such as those based on compositional barrier certificates, exploit topology and other model structure to improve scalability (and even parallelization)?
- Can we formulate ?safety? properties of an autonomous system in terms of control-theoretic properties of a hybrid system? If the answer is negative, what class of ?safety? properties can be described completely by control-theoretic concepts?
- Discretization of probability spaces and measures for numerical approximations introduces limitations, and often combinatorial explosion of states. How can we verify probabilistic or statistical hybrid systems in a manner that is independent of the fidelity of the discretization, or even avoids discretization altogether?

Constructive verification

- How can verification methods be designed to use invalidating error traces to correct models and controllers?
- Can reachability analysis be performed with models that are partially symbolic, so that results can re-used later for further analysis? Upon reaching an unsafe state, is it possible to correct the model or suggest possible minimal model corrections, with some notion of minimal change?

- How can two error traces be compared or classified? Given an error trace, can we efficiently find all other error traces that are similar according to some metric? Given a set of similar errors, can we find the smallest modification of the controller or model that eliminates all the errors in the set simultaneously?

  Computational feasibility
- How can computational limitations be considered at the design stage, either in the models or as robustness margins, and factored into the certificates needed to guarantee properties of interest (stability, invariance, reachability, robustness)? How can we construct anytime algorithms that also provide assurances of convergence? How can verification and synthesis algorithms be designed to support hardware parallelization, by construction?
- What reachability schemes are amenable to worst-case execution-time analysis? Can methods be designed with bounded error in bounded time? Can metrics be designed which compare two online reachability analysis approaches for a given system? Can GPUs or FPGAs or other hardware components be leveraged to greatly speed up online and offline reachability analysis?

  Modeling constraints and systems of systems
- How can modeling and analysis frameworks be designed to support automated model generation that is amenable to verification and synthesis?
- How can verification be performed in hierarchical systems that involve mathematical as well as black-box or data-driven models?
- Complex, autonomous systems are typically designed and constructed from modular subsystems. In this case, verification and analysis are often carried out individually for each modular component. When, if at all, can properties of subsystems carry over to the entire system?

  Synthesis
- How can reachability analysis be leveraged to solve controller synthesis problems? How do we integrate performance objectives (e.g., minimizing fuel consumption or control effort), in addition to safety objectives and temporal specifications, into reachability and viability calculations, as well as into feedback controller synthesis?
- How can we design algorithms that guarantee the closed-loop system meets the given specifications under uncertainty? Can recovery modes be designed to cope with unexpected situations? How could autonomous learning from ongoing measurements be used to guarantee safety in such situations?
- How can counter-example generation be performed for nonlinear or possibly hybrid models?

3 Personnel Supported

Not applicable.

4 Technical Publications

Not applicable

5 Interactions/Transitions

A number of AFRL researchers from RQ, RV and RW participated the workshop.

6 Patent Disclosures

Not applicable.

7 Honors

Not applicable.