Research Statement

Ritam Raha

My research vision is to **apply formal methods and verification techniques to enhance the safety, security, and reliability of cyber-physical systems**. In the past decade, there has been an unprecedented rise in the incorporation of cyber-physical systems being used in performing complex tasks. Verifying these systems is necessary to ensure their safety and reliability, especially in safety-critical scenarios. Formal methods have been proven to be time-tested to systematically verify these systems. In particular, it provides mathematically sound techniques to systematically design and verify the system dynamics against a formal requirement.

Despite the success of formal methods in verifying cyber-physical systems, they encounter two main technical challenges: (i) While they offer a rigorous approach to system verification, their scalability becomes a significant concern when applied to large-scale cyber-physical systems. The state space explosion problem, where the number of possible system states grows exponentially with the system's size, can hinder the feasibility of exhaustive formal verification. (ii) They encounter challenges when dealing with these systems' dynamic and continuous aspects. Traditional formal verification techniques excel at exhaustively analyzing discrete state transitions but can struggle with the intricate real-time behaviours inherent to controllers. On the contrary, control theory offers powerful tools for designing controllers that optimize system behaviours. Still, it often assumes idealized models that can lead to performance degradation in the presence of uncertain or nonlinear dynamics. Existing techniques in controller synthesis for real-world systems often fall short of providing formal guarantees of correctness, as synthesizing safe control is undecidable for most real-world systems. By bridging the gap between these two disciplines, I aim to create a synergy that overcomes their individual limitations. Integrating formal methods with control theory has the potential to provide a comprehensive autonomous framework for designing and verifying cyber-physical systems, leading to safer, more resilient systems that can adapt to real-world conditions while assuring correctness and stability.
My Master's thesis and Doctoral dissertation research contribute towards these aspects and challenges. The subsequent sections focus on utilizing formal methods and techniques across three stages of system design and verification. These sections explain my research contributions, future vision, and goals for each aspect.

Learning Specifications

Formal specifications are models to formally analyze a system's behaviour and requirements and assist in its design by verifying essential properties of interest. Such specifications are 'formal' in that they possess a well-defined syntax, their semantics fall within a specific domain, and they can infer useful information. The process of manually writing specifications is widely recognized as tedious and error-prone. A significant challenge in formal verification lies in synthesizing functional, correct, and easily interpretable specifications and accurately capturing the design requirements. Towards this, temporal logics (e.g., Linear Temporal Logic (LTL), Metric Temporal Logic (MTL), and Signal Temporal Logic (STL)) have gained widespread interest as they are close to natural

language, human-interpretable and have efficient verification algorithms. Over the past few years, learning temporal logic has been identified as an essential goal in robotics, specification mining, and artificial intelligence. The fundamental problem is to build a specification in the form of a temporal logic formula from a set of execution traces of a system that are classified as positive and negative such that (i) the formula is consistent with the sample and (ii) the formula can be efficiently used to apply in formal verification of the system.
Several approaches exist to learn Linear Temporal Logic (LTL) formulas from observed system traces using SAT solvers, Bayesian inference, and decision trees. However, these existing approaches often suffer from the limitations of scaling. Indeed, theoretical studies have shown that finding a concise LTL formula is NP-hard already for very small fragments, explaining the difficulties found in practice. To address these, we devised an approximation algorithm in [RRFN22]. On a high-level overview, our algorithm searches for LTL formulas in a particular shape (Directed LTL) using dynamic programming and searches for the separating formulas using Boolean combinations of the same using a greedy approach. To analyze and assess our algorithms against the previous algorithms, we have implemented a prototype of our algorithm in Python 3 in a tool named SCARLET and showcased its scalability against existing techniques. Next, we focus on automatically learning Metric Temporal Logic (MTL) and Signal Temporal Logic (STL) specifications. These logical formalisms extend LTL to reason about the continuous evolution of a CPS over time and are hence popularly employed in verifying and monitoring complex and hybrid controllers used in aviation, drones, robots, etc. Despite being extensively used in verification, to the best of our knowledge, a limited number of works focus on learning MTL specifications for verification. Towards this, we present a novel algorithm in [RRF$^+$23] to automatically learn MTL specifications that are concise and efficient for monitoring. We also implemented our algorithm in a Python 3-based tool named TEAL and ran it over several benchmarks to establish the efficiency of our algorithm. Our techniques are general enough to extend it to learn specifications in more general formalisms like STL. This research direction has a lot of potential to generalize these techniques in two ways - (i) learn specifications being used in industries with a more expressive nature (RE, PSL), and (ii) devise more efficient learning algorithms that will be scalable for industry applications.

Human in the Loop: Active Learning. Another possible future direction toward learning temporal specification will be to investigate the problem in an active learning setting. This setting assumes the presence of a teacher, where the learner tries to learn the desired temporal specification incrementally by asking the teacher two kinds of queries: membership and equivalence. This setting is interesting as an efficient active learning algorithm allows an engineer to incorporate newly observed system behaviours to update the already learned specification. Also, this setting incorporates a human (a control engineer) in the loop that will guide the process of obtaining desired specifications, which is more practical in real-world scenarios. Although there are substantial works on active learning settings for automata, to the best of our knowledge, there are not many works on a similar setting for temporal logic. Although we can use a passive learning setting to devise an algorithm for active learning, that is far from efficient. The crux of the hardness towards this is twofold- (i) The equivalence query for LTL is computationally hard and more expensive than that of automata, and (ii) there is no canonical minimal LTL formula for a given language. The *Directed LTL* fragment already shows the first steps towards this as this fragment bypasses both the above-mentioned problems. Hence, it will be interesting to investigate this direction, resulting in a learning algorithm more suitable for practical settings.

Design Controller: Control Synthesis

Control synthesis is the process of designing algorithms or controllers that govern the behaviour of CPS and help in the construction of CPS design. The main idea is to devise a control law that fulfils a requirement (qualitative/quantitative) under the presence of uncontrollable physical components. Games on graphs are one of the popular techniques towards designing correct controllers against a particular requirement. They are a fascinating branch of graph theory where vertices represent players, edges signify interactions, and states evolve through gameplay between these two players. The main idea is that the controller aims to devise strategies to fulfil the given requirement in the graph against the environment.

Weighted games are a common way to formally address questions related to resource consumption, production and storage, in which transitions carry positive or negative integers, representing the accumulation or consumption of resources. Various objectives have been considered for such arenas, such as optimizing the total or average amount of resources collected along the play or maintaining the total amount within given bounds. The latter kind of objectives, usually called energy objectives, simulate the dynamic changes in available energy within autonomous systems. In addition to fulfilling their primary tasks, these systems must prioritize regular battery recharging to prevent power depletion. Energy objectives also find application in modelling moulding machines, where they regulate liquid levels in a tank to maintain adequate pressure for injecting molten plastic into moulds, balancing the need for pressure with the valve's service life considerations. In [HMR19, HMR22], we showed that with strong lower and upper bounds, energy objective (quantitative requirements) combined with reachability (qualitative requirement) can be reducible to only energy objective. In the same work, We also devised efficient algorithms when we 'relax' the strong bounds on energy with reachability objectives. These relaxed objectives can be used to design batteries or systems where scores punish violations and do not result in immediate failures. One possible research direction will be to explore similar settings with more intricate objectives like mean-payoff or discounted payoff that help to model the adaptive control strategies by allowing the system to adjust its behaviour based on the changing value of different states and actions. This adaptability is essential for CPS to respond effectively to varying conditions while still achieving their goals. On the other hand, *reactive synthesis* provides the platform to design correct-by-construction control laws from a desired requirement against a dynamic environment. Temporal logic-based reactive synthesis provides assurances of correctness and practical scalability, even in the face of computational complexity challenges. This provides a future direction that allows for combining temporal and quantitative objectives in the games on graphs setting that helps design more robust and reliable controllers against changing environments where the requirements are more practical in real-world scenarios. Existing works in this direction mostly explain the hardness from the theoretical point of view. The possible future direction will be to explore heuristics or to find better objectives that devise scalable algorithms for real-world cases.

Another possible direction is to consider settings where the environment is not fully antagonistic, which is more practical in real-world scenarios, and to devise 'strategy templates' that will be more robust to dynamic system changes. Existing techniques only allow for qualitative objectives towards this direction, but combining quantitative and qualitative objectives in this direction is also provides potential for interesting research direction.

Verification

Formal verification has been proven to be a time-tested method to ensure the safety and reliability of CPS. It involves the mathematically rigorous techniques to establish a specification's validity

within a formal model of the system.

Discrete Models. Finite-state machines are widely used as formal models as they can formally encode the discrete behavioural changes in a system. Counter machines are computational models equipped with one or more counters that can be incremented, decremented, and tested. They extend the expressive power of finite state machines by incorporating counters to track and enforce quantitative properties. This makes them essential in formal verification for reasoning about quantitative aspects of system behaviour. It is well-known that counter machines with two counters are computationally equivalent to a Turing machine, making them inefficient to deploy in practice. To address this, several restrictions of the model have been proposed. In [PR22], we worked on a restricted model with one counter, called One-counter automata with parameters. In this model, we allow one counter that can be updated or tested with natural numbers or a set of parameters along its transitions. We show that the parameter synthesis problems for this model are decidable via an encoding to Presburger Arithmetic with Divisibility (PAD). The parameter synthesis problem for an LTL specification is a natural question in formal verification that asks whether a valuation of the parameters exists, such that all runs satisfy a given LTL specification. We also show that these problems are in PSPACE when we restrict the model to only parametric tests, giving us a more implementable algorithm for solving this problem. From a theoretical point of view, there are many interesting open questions in this area. Although the relationship with PAD seems to establish that the problems are inherently hard, no lower bounds are known for these problems beyond $\mathrm{NP}^{\mathrm{NP}}$. On the other hand, finding restriction or approximations of the general model (e.g. Continuous One-Counter Automata, One-counter nets) to get efficient algorithms for these problems are also a growing research area. Allowing all kinds of tests (upper and lower bound) in one-counter automata with parameters corresponds to the famous long-standing open problem: Ibarra's simple programs.

Continuous and Hybrid Models. The hardness of verification increases significantly when we consider the continuous or hybrid behaviour of CPS over time. Hybrid automata is a mathematical model that is general and rich enough to cover and integrate both continuous and discrete dynamics of a system, and therefore, their mathematical model must be general and rich enough to cover and integrate both aspects of the behaviour. Consequently, the reachability question in Hybrid automata is already undecidable (decidable for a very restricted fragment). This leads to using approximations and restrictions in the model or the techniques to tackle the verification questions in these models. Given any controller in a CPS, one of the most fundamental questions is to measure how efficient or cost-effective the controller is compared to the other possible controllers. In [BDK+23], we show that this problem is undecidable in full generality but devised an efficient algorithm using reachability approximation and cloning the behaviour of the given controller in a Deep Neural Network (DNN). This algorithm introduces a baseline for the problem and opens the scope for improved algorithms to the point where it scales for interesting classes of hybrid systems. Another direction is to tackle the parametric verification questions of hybrid automata, as discussed in the discrete cases. From a theoretical point of view, these problems will still be undecidable but that opens up the scope to tackle them using similar approximation approaches and using machine learning techniques that will be relevant in real-world scenarios.

As a formal methods researcher, I am excited to explore new and interesting applications of Formal Methods in interdisciplinary fields that will popularize these techniques to be deployed even more in practical applications. One promising direction is to use these methods in Control Theory and establish harmony between the two fields. I strongly believe that my expertise and research vision

will lead to significant progress in these fields.

Publications

[BDK+23] Stijn Bellis, Joachim Denil, Ramesh Krishnamurthy, Tim Leys, Guillermo Peréz, and Ritam Raha. A framework for the competitive analysis of model predictive controllers. In *17th International Conference on Reachability Problems (RP 2023)*, 2023. (To appear).

[HMR19] Loïc Hélouët, Nicolas Markey, and Ritam Raha. Reachability games with relaxed energy constraints. In Jérôme Leroux and Jean-François Raskin, editors, *Proceedings Tenth International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2019, Bordeaux, France, 2-3rd September 2019*, volume 305 of *EPTCS*, pages 17–33, 2019.

[HMR22] Loïc Hélouët, Nicolas Markey, and Ritam Raha. Reachability games with relaxed energy constraints. *Information and Computation*, 2022.

[PR22] Guillermo Perez and Ritam Raha. Revisiting Parameter Synthesis for One-Counter Automata. In Florin Manea and Alex Simpson, editors, *30th EACSL Annual Conference on Computer Science Logic (CSL 2022)*, volume 216 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 33:1–33:18, Dagstuhl, Germany, 2022.

[RRF+23] Ritam Raha, Rajarshi Roy, Nathanaël Fijalkow, Daniel Neider, and Guillermo Peréz. Synthesizing efficiently monitorable formulas in metric temporal logic. In Submission, 2023.

[RRFN22] Ritam Raha, Rajarshi Roy, Nathanaël Fijalkow, and Daniel Neider. Scalable anytime algorithms for learning fragments of linear temporal logic. In Dana Fisman and Grigore Rosu, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 263–280, Cham, 2022. Springer International Publishing.