A Verification Framework for Concurrent Systems based on Higher Dimensional Automata

Keywords

Formal Methods, Automata Theory, Model checking, Petri Nets, Extended Automata

Profile and skills required

The goal of this thesis is develop a framework for working with higher dimensional automata and use them in different modeling and verification contexts, in particular for Petri Nets and extended automata in a new and hopefully more efficient way. This goal should be reflected in the profile and skills of the candidate:

- M2 in computer science or mathematics
- Solid background in formal methods and automata theory
- Some experience in tool develop and advanced C++ techniques
- Affinity for conceiving efficient data structures and algorithms

Summary of thesis project

Modeling and verifying concurrent systems is a corner stone in the formal methods community and has attracted a tremendous amount of research over several decades.

In recent years higher dimensional automata (HDA) have become an increasingly active field of research within the concurrent systems community, see [1, 2, 3, 6, 7, 9] and references therein. This is in part due to their intuitive interpretation as they form an extension of standard automata, in part to their expressiveness, which allows for the modeling of true, non-interleaving concurrency.

Indeed it was shown in [7] that HDA is a strictly more expressive model than most standard concurrency models. That is, any instance of all these other models can in theory be translated to an HDA representing an equivalent behavior. Despite this expressiveness, the promising properties and the significant theoretical advancements in recent years, practical applications and implementations for HDAs are lacking.

The goal of this thesis is to push forward practical applications for HDA and devise new and efficient algorithms to solve classical and new problems related to model checking. These practical advances will also provide guidance for which aspects of the theory are

to be developed further or if our model needs additional restrictions or features. To keep up with this rapidly evolving theoretical landscape, a flexible yet efficient platform for working with HDAs is necessary. It has to be modular enough to accommodate intermediate abstractions or changing the internal representation while allowing to fully take advantage of the abstract representation of concurrency HDAs provide.

The thesis can be divided into a theoretical and a practical part, with each of them having some well-defined sub-tasks. However they are interconnected as the advances and discovered demands of one impact the other.

Envisaged theoretical contributions

(I) Comprehensive complexity analysis of different algorithms on HDAs, including potential improvements to classical algorithms
(II) Define the concepts of determinism and history-determinism for HDA
(II) Detailed comparison between the different Petri Net semantics, extended automata and the corresponding version of HDA
(II) Define properties and languages of extended HDA
(III) Extend HDA to $\omega-$HDA working on infinite words
(III) Study the translation from LTL (or LTLf) like specifications to HDAs

Envisaged practical contributions

(I) Devise and compare data structures and representation for the different building blocks of HDAs and their languages (Cells, iipomsets, etc.)
(I) Devise classic algorithms for HDAs (determinisation, complementation, etc.)
(II) Petri Net and extended automata to HDA translation
(II) Model checking algorithms for HDAs
(II) Competing at the annual model checking contest, possibly in multiple categories
(III) Translation of LTLf (LTL)-like formulas to (ω)-HDAs

The thesis is mainly composed of two parts, a first part for which no big surprises from the theoretical side are expected and that concerns mainly the comparison of different alternatives and the adaptation of classic algorithms to HDAs. There currently exist two implementations to represent higher dimensional automata, one open source project with links to process graphs (pg2hda) and a student project at Epita with links to Petri Nets. However neither of these implementations takes advantage of sparse representations for HDAs and are very limited in their extensibility and usability. The second part of the thesis in more challenging on a practical and especially on the theoretical side. A comprehensive comparison of different Petri Net semantics, extended automata and HDA is an extension to the work in [7] and needs to be englobe colored and possibly timed Petri Nets. This work is crucial to successfully participate at the model checking competition, which regroups industrial and academic examples of Petri Nets and their verification. The investigation into extended HDA, which could be used in settings like the one described in [8], might reveal substantial differences to *classical* extended automata, due to the true concurrency of the HDA model. The connections between LTL (LTLf)-like formulae and (ω-)HDAs are also of great theoretical and practical interest,

see [4]. On the practical side as possibly deriving an algorithm to translate LTL-like specifications to $\omega-$HDAs could significantly reduce the complexity of LTL-based model checking. On the theoretical side it has been shown that non-deterministic Büchi-HDAs do not possess the same expressiveness as non-deterministic Muller-HDAs, in contrast to standard automata theory, opening up interesting research directions.

As the framework needs to encompass all aspects from handling formulae, manipulating the automata and hosting a large spectrum of different algorithms, we will draw some inspiration and the lessons learned from spot [5].

The list of envisaged results also defines a rough schedule for the different steps to be taken during this thesis. Tasks marked by (I) are expected to by accomplished in large parts by the end of the first year. Whereas as the ones marked by (II) and (III) are expected to start within the second and third year respectively.

Context

Formal methods are mathematically rigorous techniques to ensure that a system design meets the given specification. There are numerous mathematical models, answering to different problems. These approaches reach from simple finite state machines to formal proof systems. Models for concurrent systems, that aim to represent and analyze systems that can execute multiple actions at the same time, or out of order, have been of particular interest.

Petri Nets stand out as one of the most well-established models for concurrency among the various existing frameworks. They can model numerous semantics of concurrence by using a built-in notion of resources (tokens). Moreover, Petri Nets are an established standard in academia and industry as they have a nice graphical representation while retaining a high expressiveness. A more recent framework for modeling concurrency are Higher Dimensional Automata (HDA), which are an automaton based model representing both interleaving and non-interleaving concurrency.

As an example, consider Fig. 1 showing Petri Nets and HDA models for a system with two events, labelled a and b. The Petri Net and HDA on the left side model the (mutually exclusive) interleaving of a and b as either $a.b$ or $b.a$; those to the right model concurrent execution of a and b. In the HDA, this "non-interleaving independence" is indicated by a filled-in square.

Another interesting use case for HDA is to represent the asynchronous product of extended automata. Indeed, traditional automata are the same as an HDA of dimension at most 1 and therefore the product of N such automata results in an at most N-dimensional HDA. Expanding our HDA model to englobe extended automata with synchronizations might allow to scale up approaches such as the one described in [8], used to detect races and faults in communication networks.

In recent years, the automata theoretical foundations for HDAs have been layed by a relatively small but very active and tightly connected group of international researchers (main researchers are Georg Struth, Sheffield, UK; Krzysztof Ziemiański, Warsaw, PL; Christian Johansen, Oslo, NO; Uli Fahrenberg, Paris, FR). The model has shown very

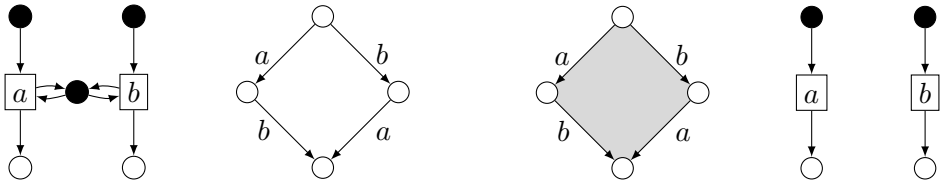

Figure 1: In Petri Nets, there are two types of "states": places (circles) and transitions (squares). A transition can be activated if all incoming places have a resource available (black, filled-in circle), destroying the incoming resource and producing a new resource on the outgoing places. On the left the Petri Net and HDA model for interleaving concurrency with the executions $a.b + b.a$ are shown. That is a and b can happen in any order, however, if we take take the actions to be noninstantaneous, they can never happen at the same time or overlap. On the right we have the models for "true", non-interleaving concurrency. Here a and b can, in addition to one after another, also happen at the same time or partially overlap, denoted as $a \parallel b$.

promising *theoretical* properties and might therefore be able to alleviate some of the inherent complexity of concurrent systems and put to scale the verification of such systems.

The quick progress on the theoretical side is in stark contrast to an almost non-existing evolution on the practical side, which is however crucial to further motivate the interest in this model and guide further fundamental research. The only practical use of HDA is the tool PG2HDA by Thomas Kahl, which takes a process given in PROMELA and translates it into an HDA and allows to following subsequent topological analysis via the CHomP tool. However this tool does not take advantage of a sparse representation for HDA and is not modular or easily modifiable. We therefore think it is time to capitalize upon this progress and develop an open source framework that allows for efficient HDA manipulation and usage. Learning the lessons from developing spot, we wish to construct the framework as modular as possible using modern C++ techniques such that exchanging each component is possible. For instance allowing swapping out different internal representation of the HDA or its alphabet is easily possible. Putting HDAs to practice in this fashion will also guide and necessitate further theoretical advances, depending on the findings.

Expected Results

- Publications in well-known conferences and journals on automata theory and practice
- Open source framework for HDA manipulation
- Participation at the model checking contest

Thématique

Automata Theory, Petri Nets, Verification

Domaine

Theoretical Computer Science

Objectifs

- Contributions to the theory of higher dimensional automata
  - Comprehensive complexity analysis of different algorithms on HDAs
  - Links between HDA and well-known formalisms such as Petri Nets and extended automata
  - Connections between HDA and logics
- A framework for HDA manipulation
  - Devise and compare data structures for HDA
  - Create a range of algorithms for HDA reflecting their counterparts for regular automata
  - Design it in a modular way such that cells and letters (pomsets) are easily modifiable without having to rewrite all algorithms

Supervision and organization

The thesis is a co-tutelle between Télécom SudParis - Samovar (IP Paris) and Epita - LRE (EDITE). It will mainly be located at Télécom SudParis either at the campus in Palaiseau (Place Marguerite Perey, Palaiseau) or Evry (Rue Charles Fourier, Courcouronnes) with frequent visits to Epita (Rue Voltaire, Le Kremlin-Bicêtre). For any further information, or if you wish to candidate for this position, please write to philipp.schlehuber-caissier@telecom-sudparis.eu before submitting. It is partially financed via an AMX grant from IP Paris, partially by Epita.

A corresponding M2 internship is also available, see here.

- Natalia KUSHIK, HDR, Télécom SudParis, Samovar - ACMES
- Philipp SCHLEHUBER-CAISSIER, Télécom SudParis, Samovar - NeSS
- Uli FAHRENBERG, HDR, Epita, LRE - Automata and Applications
- Amazigh AMRANE, Epita, LRE - Automata and Applications

Collaborations envisagées

The supervisors are actively collaborating with a number of international and national partners. Some topics of this thesis could be carried out with these partners or lead to further collaborations.

- Computation research group, The University of Sheffield, United Kingdom
- Institute of Mathematics, University of Warsaw, Poland

- Systems Security Research Group, NTNU, Norway
- IRIF, Université Paris Cité, France

References

[1] Amazigh Amrane, Hugo Bazille, Emily Clement, and Uli Fahrenberg. Languages of higher-dimensional timed automata. In *International Conference on Applications and Theory of Petri Nets and Concurrency*, pages 197–219. Springer, 2024.

[2] Amazigh Amrane, Hugo Bazille, Uli Fahrenberg, and Marie Fortin. Logic and languages of higher-dimensional automata. In *International Conference on Developments in Language Theory*, pages 51–67. Springer, 2024.

[3] Amazigh Amrane, Hugo Bazille, Uli Fahrenberg, and Krzysztof Ziemiański. Closure and decision properties for higher-dimensional automata. In Erika Ábrahám, Clemens Dubslaff, and Silvia Lizeth Tapia Tarifa, editors, *Theoretical Aspects of Computing – ICTAC 2023*, pages 295–312, Cham, 2023. Springer Nature Switzerland.

[4] Emily Clement, Enzo Erlich, and Jérémy Ledent. Expressivity of linear temporal logic for pomset languages of higher dimensional automata, 2024.

[5] Alexandre Duret-Lutz, Etienne Renault, Maximilien Colange, Florian Renkin, Alexandre Gbaguidi Aisse, Philipp Schlehuber-Caissier, Thomas Medioni, Antoine Martin, Jérôme Dubois, Clément Gillard, et al. From spot 2.0 to spot 2.10: what's new? In *International Conference on Computer Aided Verification*, pages 174–187. Springer, 2022.

[6] Uli Fahrenberg, Christian Johansen, Georg Struth, and Krzysztof Ziemiański. Posets with interfaces as a model for concurrency. *Information and Computation*, 285:104914, 2022.

[7] Rob J van Glabbeek. On the expressiveness of higher dimensional automata. *Theoretical computer science*, 356(3):265–290, 2006.

[8] Evgenii M Vinarskii, Natalia Kushik, Nina Yevtushenko, Jorge López, and Djamal Zeghlache. Races in extended input/ouput automata, their compositions and related reactive systems. In *ENASE*, pages 727–734, 2024.

[9] Safa Zouari, Krzysztof Ziemianski, and Uli Fahrenberg. Bisimulations and logics for higher-dimensional automata. *Theoretical Aspects of Computing - ICTAC 2024*, page to appear, 2024.