**Doctoral School of Information and Biomedical Technologies**
**Polish Academy of Sciences (TIB PAN)**

---

## SUBJECT:

Model Checking Parametric Timed Strategic Abilities

## SUPERVISORS:

Wojciech Penczek
Institute of Computer Science, PAS, Warsaw
penczek@ipipan.waw.pl

Laure Petrucci
LIPN, CNRS UMR 7030, Université Sorbonne Paris Nord Villetaneuse, France
Laure.Petrucci@lipn.fr

## DESCRIPTION:

Nowadays, socio-technical systems are ubiquitous and they become more and more complex and critical. The rapid spreading of Artificial Intelligence (AI) methods to decision-making in everyday-life autonomous systems renders the security and safety aspects especially important, and compromising them can lead to dramatic consequences such as casualties or expensive recalls. Since modern systems are open, they are subject to security breaches as well. Typical examples of such systems range from new technologies (e.g. autonomous vehicles, medical robots), or interaction of software technology with human behaviour (e.g. electronic voting protocols [26, 32, 15], interactive multimedia systems [41, 39]). Application domains include space mission planning and air traffic control [17, 21], defense and security [20, 38], logistics and production planning [22, 23], etc. (see [34] for a survey). Let us illustrate our objective with a small example of autonomous cars. Up to now, the management of a road on which autonomous cars circulate essentially depends on a centralised controller, i.e. a dedicated unit that controls the traffic. Even though this solution can be satisfactory at a small scale, with the advent of today's use of autonomous cars, a distributed control is required, where each car can communicate and collaborate with others. Thus, it should be possible to synthesise strategies for each car to reach its destination safely, within a time frame, respecting the autonomy of vehicles, and taking into account different traffic conditions depending on the time and day of circulation, so as to offer the most appropriate route. Such systems require handling individual actors and capturing temporal constraints. Autonomous agents systems provide a powerful paradigm for modelling and analyzing socio-technical systems. They embed networks of communicating agents taking autonomous decisions based on AI methods. Modelling strategic behaviours in a real-time context is a key aspect for guaranteeing safety and security of agents systems. Current approaches such as Asynchronous Multi-Agents Systems (AMAS) [25] and timed automata allow for tackling strategic abilities and time constraints in isolation only. Moreover, most works consider uncertainty due to the environment of the system only. Uncertainty in the design should also be taken into account, for instance, to consider situations where the total number of agents or the time needed for an event to occur are unknown. In our experience, such uncertainty is better captured by parametric models. This turns the analysis into a synthesis problem, so as to find the strategies that guarantee the crucial properties. Model checking tools provide means to analyse models of multi-agent systems (MOCHA [3, 5], MCK [19, 40], MCMAS [33], STV [29, 30]), or of (parametric) timed automata (IMITATOR [1], UPPAAL [2]). However, they lack the ability to capture these aspects of interest altogether and focus only on specific features.

### PhD objectives

The PhD aims at leveraging several limitations of the current research, provide more flexibility, with a framework for models and logics, and associated model checking algorithms. The systems targeted

exhibit different characteristics that include agent's strategic abilities and timing constraints. These aspects are usually studied separately in the literature. The aim is to consider them all together.

**Timed agents models and logics**

A first challenge is then to design a framework that encompasses both the agents and the time paradigms. Although the models in the state-of-the-art are different variants of automata, a simple extension, such as a superset of all existing models, might seem appropriate at a syntactic level, but has fatal flaws at the semantic level. Indeed, the new framework should also properly deal with the intertwinement between both aspects. Hence, defining the semantics of the new framework is not a trivial task. The type of synchronisation between the different parts or agents of the system is also an important issue. Asynchronicity is often considered in real-time systems, whereas agents models are mostly synchronous, i.e. agents take simultaneous steps. This problem is alleviated with the recent introduction of AMAS [25], which will serve as a basis. Designing the model of a system is not sufficient: its expected properties should also be defined so as to be checked. Several variants of temporal logics exist, that allow for reasoning about the order of events in time, the causality of events and provide a timing flavour [13]. In order to specify temporal and strategic properties of such systems, we intend to use, in addition to standard temporal logics such as LTL and CTL*, the strategy logics ATL and ATL* [4, 37], their timed versions TATL and TATL* [27], and recently
defined STCTL [8]. We may also consider to exploit subsets of Strategy Logic (SL), such as SL[SG] and SL[1G] [9]. Alternating-time temporal logic (ATL) allows for reasoning about strategic interactions in such systems, by extending the framework of temporal logic with the notion of strategic ability. In order to be able to express properties related to cognitive aspects such as knowledge and common knowledge, the logics have to include also an epistemic component.
Most of the tools and algorithmic solutions focus on agents with perfect information [3, 5, 33], i.e. agents who at any time know exactly the global state of the system, which is clearly unrealistic in all but the simplest multi-agent scenarios. The use of imperfect information semantics of TATL and STCTL, or strategy logic [24, 14, 12, 10, 18, 11], which is much more natural, but it does not admit an alternation-free fixpoint characterisation. This makes incremental synthesis of strategies
impossible, or at least difficult to achieve. Second, the semantics of strategic logics are almost exclusively based on synchronous concurrent game models. That is, one implicitly assumes the existence of a global clock that triggers subsequent global events in the system. It is worth noticing that many real-life systems (e.g. e-voting protocols) are inherently asynchronous, and do not operate on a global clock that perfectly synchronises the atomic steps of all the components.

**Model checking algorithms for real-time agents**

Due to the above mentioned complexity in both the models and the properties of interest, designing algorithms for model-checking becomes also very challenging in the proposed framework. To achieve scalability, abstraction, compositionality, symbolic approaches, and parallelisation of model checking can be used. With the aim of preserving the intended behaviour, the application of those techniques must be carefully studied. After having defined the basic elements of the real-time multi-agents models, the PhD will design new algorithms for their analysis. Broadly speaking, there are two classes of problems:
1. is the system safe/secure/reaches its goal under all circumstances?
2. what is the best (winning or optimal) strategy for an agent?
Unfortunately, these problems are known to be undecidable in full generality [6]. Efficient algorithms will be designed and implemented. Model checking algorithms rely on exhaustive graph search. One trade-off is depth-first exploration (often leading to more efficient algorithms) versus breadth-first exploration (which is easier to parallelise, so to scale the algorithms for more interesting cases). Methods and heuristics to reduce the size of the state space will be designed, based on symmetry reduction, on-the-fly generation, and partial-order reduction [25, 36, 31].
Developing algorithms that can handle the combination of multiple agents and time is very challenging. The PhD will study various heuristics to find exact or approximate answers. In particular, he will combine memoryless approximations [27], bounded model checking [16, 28], and search order

heuristics [35, 7]. The algorithms developed in this part will be implemented in an open source state of-the-art tool. This tool will be evaluated on several benchmarks which will allow us to find better heuristics and algorithmic optimisations.

## BIBLIOGRAPHY:


1. IMITATOR: Parameter Synthesis for Real-time Systems. https://www.imitator.fr.

2. UPPAAL: An Integrated Tool Environment for Modeling, Simulation and Verification of Real-time Systems. https://uppaal.org.

3. Rajeev Alur, Luca de Alfaro, Radu Grosu, Thomas A. Henzinger, M. Kang, Christoph M. Kirsch, Rupak Majumdar, Freddy Y. C. Mang, and Bow-Yaw Wang. JMOCHA: A model checking tool that exploits design structure. In Hausi A. Müller, Mary Jean Harrold, and Wilhelm Schäfer, editors, Proceedings of the 23rd International Conference on Software Engineering, ICSE 2001, 12-19 May 2001, Toronto, Ontario, Canada, pages 835–836. IEEE Computer Society, 2001.

4. Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. J. ACM, 49(5):672–713, 2002.

5. Rajeev Alur, Thomas A. Henzinger, Freddy Y. C. Mang, Shaz Qadeer, Sriram K. Rajamani, and Serdar Tasiran. MOCHA: modularity in model checking. In Alan J. Hu and Moshe Y. Vardi, editors, Computer Aided Verification, 10th International Conference, CAV '98, Vancouver, BC, Canada, June 28 - July 2, 1998, Proceedings, volume 1427 of Lecture Notes in Computer Science, pages 521–525. Springer, 1998.

6. Rajeev Alur, Thomas A. Henzinger, and Moshe Y. Vardi. Parametric real-time reasoning. In S. Rao Kosaraju, David S. Johnson, and Alok Aggarwal, editors, Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA, pages 592–601. ACM, 1993.

7. Étienne André, Jaime Arias, Laure Petrucci, and Jaco van de Pol. Iterative bounded synthesis for efficient cycle detection in parametric timed automata. In Jan Friso Groote and Kim Guldstrand Larsen, editors, Tools and Algorithms for the Construction and Analysis of Systems - 27th International Conference, TACAS 2021, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2021, Luxembourg City, Luxembourg, March 27 - April 1, 2021, Proceedings, Part I, volume 12651 of Lecture Notes in Computer Science, pages 311–329. Springer, 2021.

8. J. Arias,W. Jamroga,W. Penczek, L. Petrucci, and T. Sidoruk. Strategic (timed) computation tree logic. CoRR, abs/2302.13405, 2023.

9. Francesco Belardinelli, Wojciech Jamroga, Damian Kurpiewski, Vadim Malvone, and Aniello Murano. Strategy logic with simple goals: Tractable reasoning about strategies. In Sarit Kraus, editor, Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019, pages 88–94. ijcai.org, 2019.

10. Francesco Belardinelli and Alessio Lomuscio. Agent-based abstractions for verifying Alternating-Time Temporal Logic with imperfect information. In Kate Larson, Michael Winikoff, Sanmay Das, and Edmund H. Durfee, editors, Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017, pages 1259–1267. ACM, 2017.

11. Francesco Belardinelli, Alessio Lomuscio, and Vadim Malvone. An abstraction-based method for verifying strategic properties in multi-agent systems with imperfect information. In The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019, pages 6030–6037. AAAI Press, 2019.

12. Francesco Belardinelli, Alessio Lomuscio, Aniello Murano, and Sasha Rubin. Verification of multi-agent systems with imperfect information and public actions. In Kate Larson, Michael Winikoff, Sanmay Das, and Edmund H. Durfee, editors, Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017, pages 1268–1276. ACM, 2017.

13. Béatrice Bérard, Michel Bidoit, Alain Finkel, François Laroussinie, Antoine Petit, Laure Petrucci, Philippe Schnoebelen, and Pierre McKenzie. Systems and Software Verification, Model-Checking Techniques and Tools. Springer, 2001.

14. Raphaël Berthon, Bastien Maubert, Aniello Murano, Sasha Rubin, and Moshe Y. Vardi. Strategy logic with imperfect information. In 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017, pages 1–12. IEEE Computer Society, 2017.

15. David Chaum, Richard Carback, Jeremy Clark, Aleksander Essex, Stefan Popoveniuc, Ronald L. Rivest, Peter Y. A. Ryan, Emily Shen, Alan T. Sherman, and Poorvi L. Vora. Corrections to scantegrity II: end-to-end verifiability by voters of optical scan elections through confirmation codes. IEEE Trans. Inf. Forensics Secur., 5(1):194, 2010.

16. Edmund M. Clarke, Armin Biere, Richard Raimi, and Yunshan Zhu. Bounded model checking using satisfiability solving. Formal Methods Syst. Des., 19(1):7–34, 2001.

17. Mitchell K. Colby, Logan Michael Yliniemi, and Kagan Tumer. Autonomous multiagent space exploration with high-level human feedback. J. Aerosp. Inf. Syst., 13(8):301–315, 2016.

18. Catalin Dima, Constantin Enea, and Dimitar P. Guelev. Model-checking an alternating-time temporal logic with knowledge, imperfect information, perfect recall and communicating coalitions. In Angelo Montanari, Margherita Napoli, and Mimmo Parente, editors, Proceedings First Symposium on Games, Automata, Logic, and Formal Verification, GANDALF 2010, Minori (Amalfi Coast), Italy, 17-18th June 2010, volume 25 of EPTCS, pages 103–117, 2010.

19. Peter Gammie and Ron van der Meyden. MCK: model checking the logic of knowledge. In Rajeev Alur and Doron A. Peled, editors, Computer Aided Verification, 16th International Conference, CAV 2004, Boston, MA, USA, July 13-17, 2004, Proceedings, volume 3114 of Lecture Notes in Computer Science, pages 479–483. Springer, 2004.

20. José Manuel Gascueña and Antonio Fernández-Caballero. On the use of agent technology in intelligent, multisensory and distributed surveillance. Knowl. Eng. Rev., 26(2):191–208, 2011.

21. Sergio Alejandro Gómez, Anca Goron, Adrian Groza, and Ioan Alfred Letia. Assuring safety in air traffic control systems with argumentation and model checking. Expert Syst. Appl., 44:367–385, 2016.

22. Christoph Greulich, Stefan Edelkamp, and Niels Eicke. Cyber-physical multiagentsimulation

in production logistics. In Jörg P. Müller, Wolf Ketter, Gal A. Kaminka, Gerd Wagner, and Nils Bulling, editors, Multiagent System Technologies - 13th German Conference, MATES 2015, Cottbus, Germany, September 28-30, 2015, Revised Selected Papers, volume 9433 of Lecture Notes in Computer Science, pages 119–136. Springer, 2015.

23. Naihui He, David Zhang, and Qiang Li. Agent-based hierarchical production planning and scheduling in make-to-order manufacturing system. International Journal of Production Economics, 149:117–130, 03 2014.

24. Wojciech Jamroga, Michal Knapik, Damian Kurpiewski, and Lukasz Mikulski. Approximate verification of strategic abilities under imperfect information. Artif. Intell., 277, 2019.

25. Wojciech Jamroga, Wojciech Penczek, Teofil Sidoruk, Piotr Dembinski, and Antoni W. Mazurkiewicz. Towards partial order reductions for strategic ability. J. Artif. Intell. Res., 68:817–850, 2020.

26. Wojciech Jamroga and Masoud Tabatabaei. Preventing coercion in e-voting: Be open and commit. In Robert Krimmer, Melanie Volkamer, Jordi Barrat, Josh Benaloh, Nicole J. Goodman, Peter Y. A. Ryan, and Vanessa Teague, editors, Electronic Voting - First International Joint Conference, E-Vote-ID 2016, Bregenz, Austria, October 18-21, 2016, Proceedings, volume 10141 of Lecture Notes in Computer Science, pages 1–17. Springer, 2016.

27. Michal Knapik, Étienne André, Laure Petrucci, Wojciech Jamroga, and Wojciech Penczek. Timed ATL: forget memory, just count. J. Artif. Intell. Res., 66:197–223, 2019.

28. Michal Knapik and Wojciech Penczek. Bounded model checking for parametric timed automata. Trans. Petri Nets Other Model. Concurr., 5:141–159, 2012.

29. Damian Kurpiewski, Wojciech Jamroga, and Michal Knapik. STV: model checking for strategies under imperfect information. In Edith Elkind, Manuela Veloso, Noa Agmon, and Matthew E. Taylor, editors, Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019, pages 2372–2374. International Foundation for Autonomous Agents and Multiagent Systems, 2019.

30. Damian Kurpiewski, Witold Pazderski, Wojciech Jamroga, and Yan Kim. Stv+reductions: Towards practical verification of strategic ability using model reductions. In Frank Dignum, Alessio Lomuscio, Ulle Endriss, and Ann Nowé, editors, AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021, pages 1770–1772. ACM, 2021.

31. Alfons Laarman, Elwin Pater, Jaco van de Pol, and Henri Hansen. Guard-based partial-order reduction. Int. J. Softw. Tools Technol. Transf., 18(4):427–448, 2016.

32. Byoungcheon Lee, Colin Boyd, Ed Dawson, Kwangjo Kim, Jeongmo Yang, and Seungjae Yoo. Providing receipt-freeness in mixnet-based voting protocols. In Jong In Lim and Dong Hoon Lee, editors, Information Security and Cryptology - ICISC 2003, 6th International Conference, Seoul, Korea, November 27-28, 2003, Revised Papers, volume 2971 of Lecture Notes in Computer Science, pages 245–258. Springer, 2003.

33. Alessio Lomuscio, Hongyang Qu, and Franco Raimondi. MCMAS: an open-source model checker for the verification of multi-agent systems. Int. J. Softw. Tools Technol. Transf., 19(1):9–30, 2017.

34. Jörg P. Müller and Klaus Fischer. Application impact of multi-agent systems and technologies: A survey. In Onn Shehory and Arnon Sturm, editors, Agent-Oriented Software Engineering

Reflections on Architectures, Methodologies, Languages, and Frameworks, pages 27–53. Springer, 2014.

35. Hoang Gia Nguyen, Laure Petrucci, and Jaco van de Pol. Layered and collecting NDFS with subsumption for parametric timed automata. In 23rd International Conference on Engineering of Complex Computer Systems, ICECCS 2018, Melbourne, Australia, December 12-14, 2018, pages 1–9. IEEE Computer Society, 2018.

36. Huyen T. T. Nguyen, César Rodriguez, Marcelo Sousa, Camille Coti, and Laure Petrucci. Quasi-optimal partial order reduction. In Hana Chockler and Georg Weissenbacher, editors, Computer Aided Verification - 30th International Conference, CAV 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part II, volume 10982 of Lecture Notes in Computer Science, pages 354–371. Springer, 2018.

37. Pierre-Yves Schobbens. Alternating-time logic with imperfect recall. Electron. Notes Theor. Comput. Sci., 85(2):82–93, 2004.

38. Munindar P. Singh. Cybersecurity as an application domain for multiagent systems. In Gerhard Weiss, Pinar Yolum, Rafael H. Bordini, and Edith Elkind, editors, Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2015, Istanbul, Turkey, May 4-8, 2015, pages 1207–1212. ACM, 2015.

39. Kıvanç Tatar and Philippe Pasquier. Musical agents: A typology and state of the art towards musical metacreation. Journal of New Music Research, 47:1–50, 09 2018.

40. Ron van der Meyden. Optimizing epistemic model checking using conditional independence (extended abstract). In Jérôme Lang, editor, Proceedings Sixteenth Conference on Theoretical Aspects of Rationality and Knowledge, TARK 2017, Liverpool, UK, 24-26 July 2017, volume 251 of EPTCS, pages 398–414, 2017.

41. Rodolfo Daniel Wulfhorst, Lauro Nakayama, and Rosa Maria Vicari. A multiagent approach for musical interactive systems. In The Second International Joint Conference on Autonomous Agents & Multiagent Systems, AAMAS 2003, July 14-18, 2003, Melbourne, Victoria, Australia, Proceedings, pages 584–591. ACM, 2003