Contrat Doctoral — ED Galilée

Titre du sujet : Verification of Imperfectly Rational Multi-Agent Systems

- ➢ Unité de recherche : Laboratoire d'Informatique de Paris Nord (LIPN)
- ➢ Discipline : Informatique
- ➢ Direction de thèse : Munyque Mittelmann (CR CNRS), Laure Petrucci (PR USPN)
- ➢ Contact : mittelmann@lipn.univ-paris13.fr
- ➢ Domaine de recherche : Formal Verification
- ➢ Mots-clés : Model Checking, Rationality, Multi-Agent Systems, Strategic Reasoning, Logics

---

Context

Multi-Agent Systems (MAS) are systems composed of several autonomous entities that can interact with each other and aim at achieving private or common objectives, such as teams of cooperative robots, distributed sensor networks, or software agents negotiating in online markets. Strategy logics [2, 7] are formalisms to allow **reasoning about strategies of agents with temporal goals**, incorporating both collaborative and adversarial interactions with other agents. For instance, they can express that a group of agents has a *strategy* to avoid the system entering a failure state. A traditional approach to enable verification of systems is **model checking**. It consists on determining whether a given specification (expressed as a logical formula) holds on to some or all of the system possible executions. In other words, model checking evaluates the correctness of a system with respect to a formal specification.

In Game Theory, we usually consider *rational agents*, which means they act according to their preferences and goals. More precisely, each agent is modelled as having well-defined objectives, whether expressed as one or several goals or as a utility function, and is assumed to act towards achieving the best attainable outcome relative to those objectives (e.g., maximizing the utility function). To evaluate the expected behaviour of a MAS, we can take this assumption to filter the system executions by the ones that occur when agents act rationally. This problem is also known as **rational verification** [1].

The standard assumption in rational verification is that agents are perfect reasoners and can, therefore, compute their rational strategies. This assumption is unreasonable whenever we consider agents' with limited computing capacities or memory. Computing rational strategies is particularly challenging as agents' have to consider the effect of their actions, their knowledge about the state of the system, and the behaviour of other rational agents. Moreover, most of the research in strategy logics abstracts away agents' preferences and focuses on goals, usually Boolean. Capturing notions of partial and bounded rationality as well as agents' preferences is crucial to better evaluate MAS executions.

Recent research has proposed a strategy logic with minimal rationality [4]. This notion of minimal rationality simply requires an agent not to play a strategy that is always outperformed by another of its strategies (i.e., in a sense, the agent does not follow his worse strategies). In [3], it is shown that this minimal dominance-based requirement of rationality is particularly suitable for defining the deontic notion of obligation, namely, what an agent or coalition ought to do.

Subject

The objective of the thesis will be to investigate imperfect rationality in strategy logics. The starting point will be the recent work on minimal rationality [4] and iterated dominance [6]. Notions of iterated rationality will be proposed to capture agents reasoning about other agents' rationality.

Different directions will be envisioned during the thesis:

- Proposing strategy logics that capture iterated and bounded notions of rationality as well as agents' preferences;
- Extend the framework with the notion of group preferences resulting from the aggregation of individual agent preferences;
- Investigate the complexity of decision problems for the proposed logics, with focus on model checking and rational verification
- Proposing and implementing algorithms for solving the target decision problems. The implementation will rely on existing verification tools, such as EVE [1], MCMAS [5], and IMITATOR [8].

Candidate profile We are seeking for candidates having the following profile: (i) highly motivated Master students in computer science or closely related subjects; (ii) good knowledge in mathematics, theory of computation, and logic; (iii) good English writing and presentation skills.

References

[1] Abate, A., Gutierrez, J., Hammond, L., Harrenstein, P., Kwiatkowska, M., Najib, M., Perelli, G., Steeples, T., and Wooldridge, M. J. Rational verification: game-theoretic verification of multi-agent systems. *Appl. Intell. 51*, 9 (2021), 6569–6584.

[2] Alur, R., Henzinger, T. A., and Kupferman, O. Alternating-time temporal logic. *Journal of the ACM 49*, 5 (2002), 672–713.

[3] Horty, J. F. *Agency and Deontic Logic.* Oxford University Press, 2001.

[4] Li, Y., Lorini, E., and Mittelmann, M. Rational Capability in Concurrent Games. In *Proc. of the International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2025* (2025).

[5] Lomuscio, A., and Raimondi, F. Mcmas: A model checker for multi-agent systems. In *Proc. of the Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems* (2006).

[6] Lorini, E. On the epistemic foundation for iterated weak dominance: An analysis in a logic of individual and collective attitudes. *J. Philos. Log. 42*, 6 (2013), 863–904.

[7] Mogavero, F., Murano, A., Perelli, G., and Vardi, M. Y. Reasoning about strategies: On the model-checking problem. *ACM Trans. Comput. Log. 15*, 4 (2014), 34:1–34:47.

[8] Penczek, W., Petrucci, L., and Sidoruk, T. Satisfiability checking for (strategic) timed CTL using IMITATOR. In *Proc. of the Int. Conf. on Agents and Artificial Intelligence* (2025).