Mean-Payoff Games with ω-Regular Specifications

Extended Abstract

Thomas Steeples
Department of Computer Science
University of Oxford
Oxford, United Kingdom
thomas.steeples@cs.ox.ac.uk

Julian Gutierrez
Faculty of Information Technology
Monash University
Melbourne, Australia
julian.gutierrez@monash.edu

Michael Wooldridge
Department of Computer Science
University of Oxford
Oxford, United Kingdom
michael.wooldridge@cs.ox.ac.uk

ABSTRACT

Multi-player mean-payoff games are a natural formalism for modelling the behaviour of concurrent and multi-agent systems with self-interested players. Players in such a game traverse a graph, while trying to maximise a mean-payoff function that depends on the plays so generated. As with all games, the equilibria that could arise may have undesirable properties. However, as system designers, we typically wish to ensure that equilibria in such systems correspond to desirable system behaviours, for example, satisfying certain safety or liveness properties. One natural way to do this would be to specify such desirable properties using temporal logic. Unfortunately, the use of temporal logic specifications causes game theoretic verification problems to have very high computational complexity. To this end, we consider ω-regular specifications, which offer a concise and intuitive way of specifying desirable behaviours of a system. The main results of this work are characterisation and complexity bounds for the problem of determining if there are equilibria that satisfy a given ω-regular specification in a multi-player mean-payoff game in a number of computationally relevant game-theoretic settings.

CCS CONCEPTS

• **Theory of computation** → Logic & verification; • **Computing methodologies** → Multi-agent systems.

KEYWORDS

Multi-player games, Mean-payoff games, Automated verification, Temporal logic, Game theory, Equilibria, Multi-agent systems

ACM Reference Format:
Thomas Steeples, Julian Gutierrez, and Michael Wooldridge. 2021. Mean-Payoff Games with ω-Regular Specifications: Extended Abstract. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 3 pages.

1 INTRODUCTION

Modelling concurrent and multi-agent systems as games in which players interact by taking actions in pursuit of their preferences is an increasingly common approach in both formal verification and artificial intelligence [1, 2, 12]. One widely adopted semantic framework for modelling such systems is that of concurrent game structures [2]. On top of this framework, we can impose additional structure to represent each player's preferences over the possible runs of the system. In this paper, we assign a weight to every state of the game, and then consider each player's *mean-payoff* over generated runs: a player prefers runs that give them a greater mean-payoff [6, 16, 18]. These games are effective in modelling resource-bounded reactive systems, as well as any scenario with multiple agents and quantitative features. Under the assumption that each agent in the system is acting rationally, concepts from game theory offer a natural framework for understanding its possible behaviours [14]. This approach is expressive enough to capture applications of interest, and has been receiving increasing attention recently [5]. As such, equilibria for multi-player games with mean-payoff objectives are well studied, and the computation of Nash equilibria in such games has been shown to be NP-complete [16].

However, a given game-theoretic equilibrium may have undesirable computational properties from the point of view of a system designer. An equilibrium may visit dangerous states, or get stuck in a deadlock. Thus, one may also want to check if there exist equilibria which satisfy some additional desirable properties associated with the game. This problem — determining whether a given formal specification is satisfied on some/every equilibrium of a multi-agent system — is known as *Rational Verification* [8, 17].

Previous approaches to rational verification have borrowed their methodology from temporal logic model checking. However, since rational verification subsumes automated synthesis, the use of temporal logic specifications introduces high computational complexity [15]. To mitigate this problem, one might use fragments of temporal logic with lower complexity (*e.g.*, GR(1) [4, 11]), but in this work we adopt a different approach. Taking inspiration from automata theory, and in particular from [3], we consider system specifications given by a formal language for expressing *ω-regular specifications.* With this approach, the complexity of the main game-theoretic decision problems is considerably lower than in the case with temporal logic specifications.

In this paper, we offer the following main contributions: we introduce a syntax for ω-regular specifications and demonstrate they are a natural construct for reasoning qualitatively about concurrent games. We then study multi-player mean-payoff games with ω-regular specifications in the non-cooperative setting [14], and consider the natural decision problems relating to these games and their Nash equilibria. Following this, we take inspiration from cooperative game theory and look at equivalent decision problems with respect to a cooperative solution concept derived from the core [10, 14]. Finally, we look at reactive module games [9] as a way of inducing succinctness in our system representations, and look at how this affects our established complexity results.

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

Problem	Complexity
Memoryless Nash Membership	P upper bound
Memoryless E-Nash	NP-complete
E-Nash	NP-complete
Memoryless Core Membership	co-NP complete
Memoryless E-Core	Σ_2^P upper bound
E-Core	Σ_2^P-hard
WRMG E-Nash	NEXPTIME upper bound and EXPTIME-hard

Table 1: Summary of main results

2 GAMES, SPECIFICATIONS AND DECISION PROBLEMS

As mentioned previously, we use *concurrent game structures* [2] as our model for games. Informally, these consist of a set of agents, each with a set of actions, a set of states and a transition function which given a state, along with an action for each player, provides a new state for the game. These games are played out by repeatedly having each player choose an action, and moving the game to a new state as dictated by the transition function. Additionally, each state of the game has a vector of weights attached to it, one for each player. We define the *mean-payoff* of a player to be the lim inf of the average of the weights attached to the states visited; players prefer runs which give them a larger mean-payoff. Finally, we model player behaviour using *strategies* — arbitrary mathematical functions mapping histories of states to actions. In addition to being interested in functions in their full mathematical generality, we also consider *memoryless* and *finite memory* strategies. With our game model set, we are interested in what outcomes will emerge when our players act rationally. One natural way of doing this is to use *solution concepts* from game theory, and we are particularly interested in two of them: the Nash equilibrium [13] and the core [10]. The Nash equilibrium consists of those strategy profiles which are invariant to unilateral deviations, and the core consists of those strategy profiles which are invariant to multilateral deviations, under the assumption that the remaining players could also change their action when faced with a deviation. Nash equilibria are well-known in non-cooperative game theory, but the core is a concept borrowed from cooperative game theory. Both of these solution concepts capture a notion of 'stable' behaviour.

One problem is that even if a strategy profile is a Nash equilibrium, or it lies in the core, despite its 'stability', it may still have socially undesirable properties. Thus, we are not solely interested in the Nash equilibria or the core of the game — we are interested in those members of them which satisfy certain desirable properties. This problem is called *Rational Verification* [8, 17]. Traditionally, temporal logic specifications are used to describe the desired executions of a system. One problem with this is that it has a high computational overhead — computing if a mean-payoff game has a Nash equilibrium is NP-complete [16], whilst determining if a mean-payoff game has a Nash equilibrium which models a given LTL specification is PSPACE-complete [11]. Thus, we look for a more computationally amenable way of specifying system behaviours. We do this by introducing ω-regular specifications, borrowing the syntax of [3]. These consist of Boolean combinations of atoms of the form $\mathrm{Inf}(F)$, where F is some set of states. Semantically, these specifications describe the sets of states that are visited infinitely (or finitely) often in a natural way. For instance, in a game with the states $\{s_1, s_2, s_3\}$, the following ω-regular specification,

$$\mathrm{Inf}(\{s_1, s_2\}) \vee \mathrm{Fin}(\{s_2, s_3\}),$$

describes those runs of the game which either visits s_1 and s_2 infinitely often or visits s_2 and s_3 finitely often. Using this syntax (and potentially using a product construction), we can succinctly and naturally describe all ω-regular behaviours of a system, from Büchi to Muller to Parity conditions. We also note that our ω-regular specifications are equivalent to Emerson-Lei conditions [7], albeit with a different syntax.

With this machinery in place, there are two key decision problems of interest. The first problem asks, given a game, a strategy profile, a solution concept, and an ω-regular specification, is the given strategy profile a member of the given solution concept, and does it satisfy the specification? We call this problem the Membership problem, and we can consider finite-memory strategies, or restrict our attention to memoryless strategies. The second problem asks, given a game, a solution concept and an ω-regular specification, whether there exists some member of the solution concept which models the specification. We call this the E-Nash(/E-Core) problem. We analysed this problems in the setting of concurrent game structures, and well as in the succinctly represented model of weighted reactive module games (WRMGs) [9]. Our results are as in Table 1. In addition to this, we have also established a number of results which offer characterisations of these games and offer insight into the sort of behaviours they can display.

3 CONCLUDING REMARKS

In this paper we introduced ω-regular specifications as a natural, expressive, computationally tractable way of reasoning qualitatively about concurrent games. In particular, we establish several results within the rational verification framework, the most important of which are the complexity bounds for the **E-NASH**, **E-CORE** and **WRMG-E-NASH** problems. Moving forward, we would like to determine some of the bounds missing from Table 1, and following this, there are other directions that appear to be fruitful; for example, introducing both imperfect information and nondeterminism offers a closer approximation to real-world systems and we are also interested in using ω-regular specifications to understand ω-regular games in a unified, principled way.

ACKNOWLEDGMENTS

T. Steeples gratefully acknowledges the support of the EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines and Systems EP/L015897/1 and the Ian Palmer Memorial Scholarship. M. Wooldridge was supported by JP Morgan and the Alan Turing Institute.

REFERENCES

[1] Luca de Alfaro and Thomas A. Henzinger. 2000. Concurrent Omega-Regular Games. In *15^{th} Annual IEEE Symposium on Logic in Computer Science, Santa Barbara, California, USA, June 26-29, 2000.* IEEE Computer Society, 141–154. https://doi.org/10.1109/LICS.2000.855763

[2] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *J. ACM* 49, 5 (2002), 672–713. https://doi.org/10.1145/585265.585270

[3] Tomás Babiak, Frantisek Blahoudek, Alexandre Duret-Lutz, Joachim Klein, Jan Kretínský, David Müller, David Parker, and Jan Strejcek. 2015. The Hanoi Omega-Automata Format. In *Computer Aided Verification - 27^{th} International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 9206)*, Daniel Kroening and Corina S. Pasareanu (Eds.). Springer, 479–486. https://doi.org/10.1007/978-3-319-21690-4_31

[4] Roderick Bloem, Barbara Jobstmann, Nir Piterman, Amir Pnueli, and Yaniv Sa'ar. 2012. Synthesis of Reactive(1) designs. *J. Comput. Syst. Sci.* 78, 3 (2012), 911–938. https://doi.org/10.1016/j.jcss.2011.08.007

[5] Krishnendu Chatterjee and Thomas A. Henzinger. 2012. A survey of stochastic \(ømega\)-regular games. *J. Comput. Syst. Sci.* 78, 2 (2012), 394–413. https://doi.org/10.1016/j.jcss.2011.05.002

[6] Andrzej Ehrenfeucht and Jan Mycielski. 1979. Positional strategies for mean payoff games. *International Journal of Game Theory* 8, 2 (1979), 109–113.

[7] E. Allen Emerson and Chin-Laung Lei. 1987. Modalities for Model Checking: Branching Time Logic Strikes Back. 8, 3 (1987), 275–306. https://doi.org/10.1016/0167-6423(87)90036-0

[8] Dana Fisman, Orna Kupferman, and Yoad Lustig. 2009. Rational Synthesis. *CoRR* abs/0907.3019 (2009). http://arxiv.org/abs/0907.3019 _eprint: 0907.3019.

[9] Julian Gutierrez, Paul Harrenstein, and Michael J. Wooldridge. 2017. From model checking to equilibrium checking: Reactive modules for rational verification. *Artif. Intell.* 248 (2017), 123–157. https://doi.org/10.1016/j.artint.2017.04.003

[10] Julian Gutierrez, Sarit Kraus, and Michael J. Wooldridge. 2019. Cooperative Concurrent Games. In *Proceedings of the 18^{th} International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, Edith Elkind, Manuela Veloso, Noa Agmon, and Matthew E. Taylor (Eds.). International Foundation for Autonomous Agents and Multiagent Systems, 1198–1206. http://dl.acm.org/citation.cfm?id=3331822

[11] Julian Gutierrez, Muhammad Najib, Giuseppe Perelli, and Michael J. Wooldridge. 2019. On Computational Tractability for Rational Verification. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, Sarit Kraus (Ed.). ijcai.org, 329–335. https://doi.org/10.24963/ijcai.2019/47

[12] Thomas A. Henzinger. 2005. Games in system design and verification. In *Proceedings of the 10^{th} Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2005), Singapore, June 10-12, 2005*, Ron van der Meyden (Ed.). National University of Singapore, 1–4. https://dl.acm.org/citation.cfm?id=1089935

[13] John F Nash et al. 1950. Equilibrium points in n-person games. *Proceedings of the national academy of sciences* 36, 1 (1950), 48–49.

[14] Martin J Osborne and Ariel Rubinstein. 1994. *A course in game theory.* MIT press.

[15] Amir Pnueli and Roni Rosner. 1989. On the Synthesis of an Asynchronous Reactive Module. In *Automata, Languages and Programming, 16^{th} International Colloquium, ICALP89, Stresa, Italy, July 11-15, 1989, Proceedings (Lecture Notes in Computer Science, Vol. 372)*, Giorgio Ausiello, Mariangiola Dezani-Ciancaglini, and Simona Ronchi Della Rocca (Eds.). Springer, 652–671. https://doi.org/10.1007/BFb0035790

[16] Michael Ummels and Dominik Wojtczak. 2011. The Complexity of Nash Equilibria in Limit-Average Games. *CoRR* abs/1109.6220 (2011). http://arxiv.org/abs/1109.6220

[17] Michael J. Wooldridge, Julian Gutierrez, Paul Harrenstein, Enrico Marchioni, Giuseppe Perelli, and Alexis Toumi. 2016. Rational Verification: From Model Checking to Equilibrium Checking. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, Dale Schuurmans and Michael P. Wellman (Eds.). AAAI Press, 4184–4191. http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12268

[18] Uri Zwick and Mike Paterson. 1996. The complexity of mean payoff games on graphs. *Theoretical Computer Science* 158, 1-2 (1996), 343–359.