Games Where You Can Play Optimally with Arena-Independent Finite Memory*: an extended abstract

Patricia Bouyer[1], Stéphane Le Roux[1], Youssouf Oualhadj[2], Mickael Randour[3], and Pierre Vandenhove[1,3]

[1] LSV – CNRS & ENS Paris-Saclay, Université Paris-Saclay, France
[2] LACL – UPEC, France
[3] F.R.S.-FNRS & UMONS – Université de Mons, Belgium

In this document, we provide a comprehensive abstract of the article [6], which is currently a preprint available on arXiv.[4] We discuss the context in which our work takes place, its contributions, related work, and we sketch the content of the upcoming student presentation at MOVEP 2020.

Controller synthesis through the game-theoretic metaphor. Two-player games on (finite) graphs are studied extensively, in particular for their application to controller synthesis for reactive systems (see, e.g., [26,30,7,2]). The seminal model is *antagonistic* (i.e., *zero-sum* if one chooses a quantitative view): player 1 ($\mathcal{P}_1$) is seen as the system to control, player 2 ($\mathcal{P}_2$) as its antagonistic environment, and the game models their interaction. Each vertex of the game graph (called *arena*) models a *state* of the system and belongs to one of the players. Players take turns moving a pebble from state to state according to the edges, each player choosing the destination whenever the pebble is on one of his states. These choices are made according to the *strategy* of the player, which, in general, might use memory (bounded or not) of the past moves to prescribe the next action.

The resulting infinite sequence of states, called *play*, represents the execution of the system. The objective of $\mathcal{P}_1$ is to enforce a given *specification*, often encoded as a *winning condition* (i.e., a set of winning plays) or as a *payoff function* to maximize (i.e., a quantitative performance to optimize). This paradigm focuses on the *worst-case* performance of the system, hence $\mathcal{P}_2$'s goal is to prevent $\mathcal{P}_1$ from achieving his objective.

The goal of *synthesis* is thus to decide if $\mathcal{P}_1$ has a *winning strategy*, i.e., one ensuring a given winning condition or guaranteeing a given payoff threshold, against all possible strategies of $\mathcal{P}_2$, and to build such a strategy efficiently if it exists.

Winning strategies are essentially *formal blueprints* for controllers to implement in practical applications. Therefore, the complexity of these strategies is of tremendous importance: the simpler the strategy, the easier and cheaper it will be to build the corresponding controller and maintain it. This explains why a lot of research effort is constantly put in identifying the exact complexity (in terms of memory and/or randomness) of strategies needed to play *optimally* (i.e., to the best of the player's ability) for each specific class of games and objectives (e.g., [24,16,14,33,19,10,1,4,32,5,11]). Alongside the *practical interest* of this question lies the *theoretical puzzle*: understanding the underlying mechanisms and implicit properties of games that lead to "simple" strategies being sufficient. Given the numerous connections between two-player games and various branches of mathematics and computer science, this fundamental question has interest in its own right.

Memoryless optimal strategies. Remarkably, several canonical classes of games that have been around for decades and proved their usefulness over and over — e.g., mean-payoff [17], parity [18,34], or energy games [12] — share a desirable property: they all admit *memoryless optimal strategies for both players*. That is, for every strategy σ_i of $\mathcal{P}_i$, there is a strategy σ_i^{ML} which is *at least as good* (i.e., wins whenever σ_i wins or ensures at least the same payoff) and that uses no memory at all. Such a

* Research supported by F.R.S.-FNRS under Grant n° F.4520.18 (ManySynth), F.R.S.-FNRS mobility funding for scientific missions (Y. Oualhadj in UMONS, 2018), and ENS Paris-Saclay visiting professorship (M. Randour, 2019). Mickael Randour is an F.R.S.-FNRS Research Associate and Pierre Vandenhove is an F.R.S.-FNRS Research Fellow.

[4] Link to the paper: https://arxiv.org/abs/2001.03894.

memoryless strategy always picks the same edge when in the same state, regardless of what happened before in the game.

Memoryless strategies are the simplest kind of strategies one can use in a turn-based game on a graph. Therefore, it is quite interesting that they suffice for objectives as rich as the ones we just discussed. Following this observation, a lot of effort has been put in understanding which games admit memoryless optimal strategies, and in identifying the exact frontiers of *memoryless determinacy*. Let us mention, non-exhaustively, works by Gimbert and Zielonka [23,24] (culminating in a complete characterization), Aminof and Rubin [1] (through the prism of first-cycle games), and Kopczynski [29] (half-positional determinacy). All these advances were built by identifying the common underlying mechanisms in ad hoc proofs for specific classes of games, and generalizing them to wide classes (e.g., the first-cycle games of Aminof and Rubin are inspired by the seminal paper of Ehrenfeucht and Mycielski on mean-payoff games [17]).

Gimbert and Zielonka's approach. Arguably, the most important result in this direction is the *complete characterization* of preference relations admitting memoryless optimal strategies, established in [24], fifteen years ago. This result can be stated as follows: a preference relation admits memoryless optimal strategies for both players on all arenas if and only if the relation (used by $\mathcal{P}_1$) and its inverse (used by $\mathcal{P}_2$) satisfy two properties (namely, *monotony* and *selectivity*).

As a by-product of their approach, Gimbert and Zielonka proved another great result, of high interest in practice [24, Corollary 7]: if memoryless strategies suffice in all one-player games of $\mathcal{P}_1$ and all one-player games of $\mathcal{P}_2$, they also suffice in all two-player games. Such a *lifting corollary* provides a neat and easy way to prove that a preference relation admits memoryless optimal strategies *without proving monotony and selectivity at all*: proving it in the two one-player subcases, which is generally much easier as it boils down to graph reasoning, and then lifting the result to the general two-player case through the corollary.

The rise of memory. Over the last decade, the increasing need to model complex specifications has shifted research toward games where multiple (quantitative and qualitative) objectives co-exist and interact, requiring the analysis of *interplay* and *trade-offs* between several objectives. Hence, a lot of effort is put in studying games where objectives are actually conjunctions of objectives, or even richer Boolean combinations. See for example [15] for combinations of parity, [13,16,28] for combinations of energy and parity, [33] for combinations of mean-payoff, [5,4] for combinations of energy and average-energy, [11] for combinations of energy and mean-payoff, [14] for combinations of total-payoff, or [14,10,8] for combinations of window objectives.

When considering such rich objectives, *memoryless strategies usually do not suffice*, and one has to use an amount of memory which can quickly become an obstacle to implementation (e.g., exponential memory) or which can prevent it completely (infinite memory). Establishing precise memory bounds for such general combinations of objectives is tricky and sometimes counterintuitive. For example, while energy games and mean-payoff games are inter-reducible in the single-objective setting, exponential-memory strategies are both sufficient and necessary for conjunctions of energy objectives [16,28] while *infinite-memory* strategies are required for conjunctions of mean-payoff ones [33].

A natural question arises: *which preference relations do admit finite-memory optimal strategies?* Surprisingly, whether an equivalent to Gimbert and Zielonka's characterization could be obtained in the finite-memory case or not has remained an open question up to now. It is worth noticing that such an equivalent could be of tremendous help in practice, especially if a *lifting corollary* also holds: see for example [5,4,11], where proving that finite-memory strategies suffice in one-player games was fairly easy, in contrast to the high complexity of the two-player case — a lifting corollary could grant the two-player case for free!

Having said that, one has to hope that the following result can be established: "if *finite-memory* strategies suffice in all one-player games of $\mathcal{P}_1$ and all one-player games of $\mathcal{P}_2$, they also suffice in all two-player games." Unfortunately, we will show during the talk an example illustrating that such a result does not hold in practice.

Our approach. We generalize Gimbert and Zielonka's results — characterization *and* lifting corollary — to the case of *arena-independent* finite memory. That is, we encompass *all* situations where the amount of memory needed by the two players is *solely dependent on the preference relation* (e.g., colors, dimensions of weight vectors), and *not* on the game arena (i.e., number of edges/states). Let us take some classical examples to illustrate this notion.

– All memoryless-determined relations — studied in [24] — use arena-independent memory: the amount of memory required, *none*, is the same for all arenas.
– Combinations of parity objectives use arena-independent memory [15]: the amount of memory only depends on the number of objectives and the number of priorities — both parameters of the preference relation, not on the size of the arena.
– Lower- and upper-bounded energy objectives also use arena-independent memory (see, e.g., [3,5,4]): the memory only depends on the bounds — parameters of the preference relation, not on the size of the arena, nor the weights used in it.
– Combinations of lower-bounded energy objectives (with no upper bound) require *arena-dependent* memory [16,28]: the memory depends on the size of the arena in addition to the weights used in it. Such a setting falls outside the scope of our results.

This restriction to arena-independent memory is natural given that there are counterexamples to a general approach. It is also important to note that it is not as restrictive as it may seem, as hinted by the examples above: we are *not restricted to constant memory* but to memory only depending on the *parameters of the preference relation* (or equivalently, objective), and not of the arena. This framework thus already encompasses many objectives from the literature, as well as possible extensions.

Let us highlight however that the *arena-independent* case, which we solve here, is an exact equivalent to Gimbert and Zielonka's results in the finite-memory case: the memoryless case is de facto arena-independent. Therefore, this paper strictly generalizes [24] and captures lots of widely-studied classes of games.

Related work. We have already discussed many important related papers, notably [24]. Following the same motivation as our work — the need to characterize (combinations of) objectives admitting finite-memory optimal strategies, Le Roux et al. [32] take another road: whereas our work permits to lift results from one-player games to two-player games, they provide a lifting from the single-objective case to the multi-objective one.

Our work focuses on *deterministic turn-based* two-player games. Sufficient conditions have been published for stochastic models but to the best of our knowledge, no complete characterization, even for the simplest case of Markov decision processes (e.g., [21]). Two unpublished articles contain interesting results on stochastic games [25,22], including an extension of Gimbert and Zielonka's original work, by the same authors [25]. Whether part of our approach can be useful to tackle the finite-memory case in this context, or in richer contexts mixing games and stochastic models (e.g., [9]) is a question for future research. Some sufficient criteria, orthogonal to our approach, were studied for *concurrent* games in [31].

Outline of the talk. In the talk, we will recall Gimbert and Zielonka's results [24], and explain the main challenges to extend them from *memoryless determinacy* to *finite-memory determinacy*. We will justify why the case of *arena-independent* finite memory is relevant, both because it is necessary to obtain a natural generalization from [24] and because of its applicability. We will also discuss how it paves the way to the *arena-dependent* case, and difficulties to generalize the results in this direction.

References

1. Benjamin Aminof and Sasha Rubin. First-cycle games. *Inf. Comput.*, 254:195–216, 2017.
2. Roderick Bloem, Krishnendu Chatterjee, and Barbara Jobstmann. Graph games and reactive synthesis. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors, *Handbook of Model Checking*, pages 921–962. Springer, 2018.

3. Patricia Bouyer, Ulrich Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Jirí Srba. Infinite runs in weighted timed automata with energy constraints. In Franck Cassez and Claude Jard, editors, *Formal Modeling and Analysis of Timed Systems, 6th International Conference, FORMATS 2008, Saint Malo, France, September 15-17, 2008. Proceedings*, volume 5215 of *Lecture Notes in Computer Science*, pages 33–47. Springer, 2008.
4. Patricia Bouyer, Piotr Hofman, Nicolas Markey, Mickael Randour, and Martin Zimmermann. Bounding average-energy games. In Javier Esparza and Andrzej S. Murawski, editors, *Foundations of Software Science and Computation Structures - 20th International Conference, FOSSACS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*, volume 10203 of *Lecture Notes in Computer Science*, pages 179–195, 2017.
5. Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim G. Larsen, and Simon Laursen. Average-energy games. *Acta Inf.*, 55(2):91–127, 2018.
6. Patricia Bouyer, Stéphane Le Roux, Youssouf Oualhadj, Mickael Randour, and Pierre Vandenhove. Games where you can play optimally with finite memory. *CoRR*, abs/2001.03894, 2020.
7. Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour, Jean-François Raskin, Ocan Sankur, and Mathieu Sassolas. Non-zero sum games for reactive synthesis. In Adrian-Horia Dediu, Jan Janousek, Carlos Martín-Vide, and Bianca Truthe, editors, *Language and Automata Theory and Applications - 10th International Conference, LATA 2016, Prague, Czech Republic, March 14-18, 2016, Proceedings*, volume 9618 of *Lecture Notes in Computer Science*, pages 3–23. Springer, 2016.
8. Thomas Brihaye, Florent Delgrange, Youssouf Oualhadj, and Mickael Randour. Life is random, time is not: Markov decision processes with window objectives. In Fokkink and van Glabbeek [20], pages 8:1–8:18.
9. Véronique Bruyère, Emmanuel Filiot, Mickael Randour, and Jean-François Raskin. Meet your expectations with guarantees: Beyond worst-case synthesis in quantitative games. *Inf. Comput.*, 254:259–295, 2017.
10. Véronique Bruyère, Quentin Hautem, and Mickael Randour. Window parity games: an alternative approach toward parity games with time bounds. In Domenico Cantone and Giorgio Delzanno, editors, *Proceedings of the Seventh International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2016, Catania, Italy, 14-16 September 2016*, volume 226 of *EPTCS*, pages 135–148, 2016.
11. Véronique Bruyère, Quentin Hautem, Mickael Randour, and Jean-François Raskin. Energy mean-payoff games. In Fokkink and van Glabbeek [20], pages 21:1–21:17.
12. Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga. Resource interfaces. In Rajeev Alur and Insup Lee, editors, *EMSOFT*, volume 2855 of *Lecture Notes in Computer Science*, pages 117–133. Springer, 2003.
13. Krishnendu Chatterjee and Laurent Doyen. Energy parity games. *Theor. Comput. Sci.*, 458:49–60, 2012.
14. Krishnendu Chatterjee, Laurent Doyen, Mickael Randour, and Jean-François Raskin. Looking at mean-payoff and total-payoff through windows. *Inf. Comput.*, 242:25–52, 2015.
15. Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Generalized parity games. In Helmut Seidl, editor, *Foundations of Software Science and Computational Structures, 10th International Conference, FOSSACS 2007, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2007, Braga, Portugal, March 24-April 1, 2007, Proceedings*, volume 4423 of *Lecture Notes in Computer Science*, pages 153–167. Springer, 2007.
16. Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin. Strategy synthesis for multi-dimensional quantitative objectives. *Acta Inf.*, 51(3-4):129–163, 2014.
17. Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *Int. Journal of Game Theory*, 8(2):109–113, 1979.
18. E. Allen Emerson and Charanjit S. Jutla. The complexity of tree automata and logics of programs. In *FOCS*, pages 328–337. IEEE Computer Society, 1988.
19. Nathanaël Fijalkow, Florian Horn, Denis Kuperberg, and Michal Skrzypczak. Trading bounds for memory in games with counters. In Halldórsson et al. [27], pages 197–208.
20. Wan Fokkink and Rob van Glabbeek, editors. *30th International Conference on Concurrency Theory, CONCUR 2019, August 27-30, 2019, Amsterdam, the Netherlands*, volume 140 of *LIPIcs*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
21. Hugo Gimbert. Pure stationary optimal strategies in Markov decision processes. In Wolfgang Thomas and Pascal Weil, editors, *STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science, Aachen, Germany, February 22-24, 2007, Proceedings*, volume 4393 of *Lecture Notes in Computer Science*, pages 200–211. Springer, 2007.
22. Hugo Gimbert and Edon Kelmendi. Two-player perfect-information shift-invariant submixing stochastic games are half-positional. Unpublished, 2014.

23. Hugo Gimbert and Wieslaw Zielonka. When can you play positionally? In Jirí Fiala, Václav Koubek, and Jan Kratochvíl, editors, *Mathematical Foundations of Computer Science 2004, 29th International Symposium, MFCS 2004, Prague, Czech Republic, August 22-27, 2004, Proceedings*, volume 3153 of *Lecture Notes in Computer Science*, pages 686–697. Springer, 2004.
24. Hugo Gimbert and Wieslaw Zielonka. Games where you can play optimally without any memory. In Martín Abadi and Luca de Alfaro, editors, *CONCUR 2005 - Concurrency Theory, 16th International Conference, CONCUR 2005, San Francisco, CA, USA, August 23-26, 2005, Proceedings*, volume 3653 of *Lecture Notes in Computer Science*, pages 428–442. Springer, 2005.
25. Hugo Gimbert and Wieslaw Zielonka. Pure and Stationary Optimal Strategies in Perfect-Information Stochastic Games with Global Preferences. Unpublished, December 2009.
26. Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
27. Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors. *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, volume 9135 of *Lecture Notes in Computer Science*. Springer, 2015.
28. Marcin Jurdzinski, Ranko Lazic, and Sylvain Schmitz. Fixed-dimensional energy games are in pseudo-polynomial time. In Halldórsson et al. [27], pages 260–272.
29. Eryk Kopczynski. Half-positional determinacy of infinite games. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part II*, volume 4052 of *Lecture Notes in Computer Science*, pages 336–347. Springer, 2006.
30. Mickael Randour. Automated synthesis of reliable and efficient systems through game theory: A case study. In *Proc. of ECCS 2012*, Springer Proceedings in Complexity XVII, pages 731–738. Springer, 2013.
31. Stéphane Le Roux. Concurrent games and semi-random determinacy. In Igor Potapov, Paul G. Spirakis, and James Worrell, editors, *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, August 27-31, 2018, Liverpool, UK*, volume 117 of *LIPIcs*, pages 40:1–40:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.
32. Stéphane Le Roux, Arno Pauly, and Mickael Randour. Extending finite-memory determinacy by Boolean combination of winning conditions. In Sumit Ganguly and Paritosh K. Pandya, editors, *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2018, December 11-13, 2018, Ahmedabad, India*, volume 122 of *LIPIcs*, pages 38:1–38:20. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018.
33. Yaron Velner, Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, Alexander Moshe Rabinovich, and Jean-François Raskin. The complexity of multi-mean-payoff and multi-energy games. *Inf. Comput.*, 241:177–196, 2015.
34. Wieslaw Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theor. Comput. Sci.*, 200(1-2):135–183, 1998.