



HAL
open science

Games Where You Can Play Optimally with Arena-Independent Finite Memory

Patricia Bouyer, Stéphane Le Roux, Youssouf Oualhadj, Mickael Randour,
Pierre Vandenhove

► **To cite this version:**

Patricia Bouyer, Stéphane Le Roux, Youssouf Oualhadj, Mickael Randour, Pierre Vandenhove. Games Where You Can Play Optimally with Arena-Independent Finite Memory. 31st International Conference on Concurrency Theory (CONCUR'20), Sep 2020, Vienna, Austria. 10.4230/LIPIcs.CONCUR.2020.18 . hal-02917552

HAL Id: hal-02917552

<https://hal.science/hal-02917552>

Submitted on 17 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Games Where You Can Play Optimally with Arena-Independent Finite Memory

Patricia Bouyer

LSV – CNRS & ENS Paris-Saclay, Université Paris-Saclay, France

Stéphane Le Roux

LSV – CNRS & ENS Paris-Saclay, Université Paris-Saclay, France

Youssef Oualhadj

LACL – UPEC, France

Mickael Randour

F.R.S.-FNRS & UMONS – Université de Mons, Belgium

Pierre Vandenhover

F.R.S.-FNRS & UMONS – Université de Mons, Belgium

LSV – CNRS & ENS Paris-Saclay, Université Paris-Saclay, France

Abstract

For decades, two-player (antagonistic) games on graphs have been a framework of choice for many important problems in theoretical computer science. A notorious one is controller synthesis, which can be rephrased through the game-theoretic metaphor as the quest for a winning strategy of the system in a game against its antagonistic environment. Depending on the specification, *optimal* strategies might be simple or quite complex, for example having to use (possibly infinite) memory. Hence, research strives to understand which settings allow for simple strategies.

In 2005, Gimbert and Zielonka [27] provided a complete characterization of preference relations (a formal framework to model specifications and game objectives) that admit *memoryless* optimal strategies for both players. In the last fifteen years however, practical applications have driven the community toward games with complex or multiple objectives, where memory — finite or infinite — is almost always required. Despite much effort, the exact frontiers of the class of preference relations that admit *finite-memory* optimal strategies still elude us.

In this work, we establish a complete characterization of preference relations that admit optimal strategies using *arena-independent* finite memory, generalizing the work of Gimbert and Zielonka to the finite-memory case. We also prove an equivalent to their celebrated corollary of great practical interest: if both players have optimal (arena-independent-)finite-memory strategies in all one-player games, then it is also the case in all two-player games. Finally, we pinpoint the boundaries of our results with regard to the literature: our work completely covers the case of arena-independent memory (e.g., multiple parity objectives, lower- and upper-bounded energy objectives), and paves the way to the arena-dependent case (e.g., multiple lower-bounded energy objectives).

2012 ACM Subject Classification Theory of computation → Formal languages and automata theory

Keywords and phrases two-player games on graphs, finite-memory determinacy, optimal strategies

Digital Object Identifier [10.4230/LIPIcs.CONCUR.2020.18](https://doi.org/10.4230/LIPIcs.CONCUR.2020.18)

Related Version Full paper on arXiv: <https://arxiv.org/abs/2001.03894>

Funding 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 Vandenhover is an F.R.S.-FNRS Research Fellow.

Acknowledgements We extend our warmest thanks to Mathieu Sassolas, for inspiring discussions that were essential in starting this work.



© Patricia Bouyer, Stéphane Le Roux, Youssef Oualhadj, Mickael Randour, and Pierre Vandenhover; licensed under Creative Commons License CC-BY

31st International Conference on Concurrency Theory (CONCUR 2020).

Editors: Igor Konnov and Laura Kovács; Article No. 18; pp. 18:1–18:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

45 **1** Introduction

46 **Controller synthesis through the game-theoretic metaphor.** Two-player games on (finite)
 47 graphs are studied extensively, in particular for their application to controller synthesis for
 48 reactive systems (see, e.g., [29, 37, 7, 2]). The seminal model is *antagonistic* (i.e., *zero-sum* if
 49 one chooses a quantitative view): player 1 (\mathcal{P}_1) is seen as the system to control, player 2 (\mathcal{P}_2)
 50 as its antagonistic environment, and the game models their interaction. Each vertex of the
 51 game graph (called *arena*) models a *state* of the system and belongs to one of the players.
 52 Players take turns moving a pebble from state to state according to the edges, each player
 53 choosing the destination whenever the pebble is on one of his states. These choices are made
 54 according to the *strategy* of the player, which, in general, might use memory (bounded or
 55 not) of the past moves to prescribe the next action.

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

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

64 Winning strategies are *formal blueprints* for controllers to implement in applications.
 65 Therefore, their complexity is of tremendous importance: the simpler the strategy, the easier
 66 and cheaper it will be to build the controller and maintain it. This explains why a lot of
 67 research effort is constantly put in identifying the exact complexity (in terms of memory and/or
 68 randomness) of strategies needed to play *optimally* (i.e., to the best of the player's ability)
 69 for each specific class of games and objectives (e.g., [27, 17, 14, 42, 22, 10, 1, 4, 41, 5, 11]).
 70 Alongside the *practical interest* lies the *theoretical puzzle*: understanding the underlying
 71 mechanisms and implicit properties of games that lead to "simple" strategies being sufficient.
 72 Given the numerous connections between games and various branches of mathematics and
 73 computer science, this fundamental question has interest in its own right.

74 **Preference relations.** There are two prominent ways to formalize a game objective in the
 75 literature. The first one, dubbed *quantitative* and inspired by games in economics, is to use
 76 *payoff functions* mapping plays to numerical values, and to see \mathcal{P}_1 as a maximizer player.
 77 This is for example the case of mean-payoff games [19]. The second one, called *qualitative*, is
 78 to define a *set of winning plays* — called *winning condition* — induced by some property, as
 79 in, e.g., parity games [20, 43]. The two formalisms are strongly linked: the classical decision
 80 problem for quantitative games is to fix a payoff threshold and ask if \mathcal{P}_1 has a strategy to
 81 guarantee it, transforming the problem into a qualitative one (where the winning plays are
 82 all those with a payoff at least equal to the threshold). To define payoff functions or winning
 83 conditions, one often uses weights, priorities, colors, etc, on states or edges of the arena.

84 In this work, we walk in the footsteps of Gimbert and Zielonka [27]: we associate a *color*
 85 to each edge of our arenas, and we adopt the abstract formalism of *preference relations* over
 86 infinite sequences of colors (induced by plays). This general formalism permits to encode
 87 virtually all classical game objectives, both qualitative and quantitative, and lets us reason
 88 in a well-founded framework under minimal assumptions.

89 **Memoryless optimal strategies.** Remarkably, several canonical classes of games that have
 90 been around for decades and proved their usefulness over and over — e.g., mean-payoff [19],
 91 parity [20, 43], or energy games [12] — share a desirable property: they all admit *memoryless*
 92 *optimal strategies for both players*. That is, for every strategy σ_i of \mathcal{P}_i , there is a strategy
 93 σ_i^{ML} which is *at least as good* (i.e., wins whenever σ_i wins or ensures at least the same payoff)
 94 and that uses no memory at all. Such a memoryless strategy always picks the same edge
 95 when in the same state, regardless of what happened before in the game.

96 Memoryless strategies are the simplest kind of strategies one can use. Therefore, it
 97 is quite interesting that they suffice for objectives as rich as the ones we just discussed.
 98 Following this observation, a lot of effort has been put in understanding which games
 99 admit memoryless optimal strategies, and in identifying the exact frontiers of *memoryless*
 100 *determinacy*. Let us mention, non-exhaustively, works by Gimbert and Zielonka [26, 27]
 101 (culminating in a complete characterization), Aminof and Rubin [1] (through the prism of
 102 first-cycle games), and Kopczyński [33] (half-positional determinacy). All these advances
 103 were built by identifying the common underlying mechanisms in ad hoc proofs for specific
 104 classes of games, and generalizing them to wide classes (e.g., the first-cycle games of [1] are
 105 inspired by the seminal paper of Ehrenfeucht and Mycielski on mean-payoff games [19]).

106 **Gimbert and Zielonka’s approach.** Arguably, the most important result in this direction is
 107 the *complete characterization* of preference relations admitting memoryless optimal strategies,
 108 established in [27], fifteen years ago. By complete characterization, we mean sufficient *and*
 109 necessary conditions on the preference relations.


110 It can be stated as follows: a preference relation admits memoryless optimal strategies
 111 for both players on all arenas if and only if the relation (used by \mathcal{P}_1) and its inverse (used
 112 by \mathcal{P}_2) are *monotone* and *selective*. These concepts will be defined formally in Section 3,
 113 but let us give an intuition here. Roughly, a preference relation is monotone if it is stable
 114 under *prefix addition*: given two sequences of colors such that one is *strictly* preferred to the
 115 other, it is impossible to reverse this order of preference by adding the same prefix to both
 116 sequences. Selectivity is similarly defined with regard to *cycle mixing*: if a preference relation
 117 is selective, then, starting from two sequences of colors, it is impossible to create a third one
 118 by mixing the first two in such a way that the third one is strictly preferred to the first two.
 119 These elegant notions coincide with the natural intuition that memoryless strategies suffice if
 120 there is no interest in behaving differently in a state depending on what happened before.

121 In addition to this complete characterization, Gimbert and Zielonka proved another great
 122 result, of high interest in practice [27, Corollary 7]: as a by-product of their approach, they
 123 obtain that if memoryless strategies suffice in all one-player games of \mathcal{P}_1 and all one-player
 124 games of \mathcal{P}_2 , they also suffice in all two-player games. Such a *lifting corollary* provides a
 125 neat and easy way to prove that a preference relation admits memoryless optimal strategies
 126 *without proving monotony and selectivity at all*: proving it in the two one-player subcases,
 127 which is generally much easier as it boils down to graph reasoning, and then lifting the result
 128 to the general two-player case through the corollary.

129 **The rise of memory.** The need to model complex specifications has shifted research toward
 130 games where multiple (quantitative and qualitative) objectives co-exist and interact, requiring
 131 the analysis of *interplay* and *trade-offs* between several objectives. Hence, a lot of effort
 132 is put in studying games where objectives are actually conjunctions of objectives, or even
 133 richer Boolean combinations. See for example [16] for combinations of parity, [13, 17, 32] for
 134 combinations of energy and parity, [42] for combinations of mean-payoff, [5, 4] for combinations

18:4 Games Where You Can Play Optimally with Arena-Independent Finite Memory

135 of energy and average-energy, [11] for combinations of energy and mean-payoff, [14] for
 136 combinations of total-payoff, or [14, 10, 8] for combinations of window objectives.



■ **Figure 1** \mathcal{P}_1 (circle) needs infinite memory to win.

137 When considering such rich objectives, *memoryless strategies usually do not suffice*, and
 138 one has to use an amount of memory that can quickly become an obstacle to implementation
 139 (e.g., exponential memory) or that can prevent it completely (infinite memory). Establishing
 140 precise memory bounds for such general combinations of objectives is tricky and sometimes
 141 counterintuitive. For example, while energy games and mean-payoff games are inter-reducible
 142 in the single-objective setting, exponential-memory strategies are both sufficient and necessary
 143 for conjunctions of energy objectives [17, 32] while *infinite-memory* strategies are required
 144 for conjunctions of mean-payoff ones [42].

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

152 Having said that, one has to hope that the following corollary can be established: “if
 153 *finite-memory* strategies suffice in all one-player games of \mathcal{P}_1 and all one-player games of \mathcal{P}_2 ,
 154 they also suffice in all two-player games.” Unfortunately, this hope is but a delusion.

155 **Lifting corollary: a counterexample.** Consider games where the colors are integers, and
 156 the objective of \mathcal{P}_1 is to create a play such that (a) the running sum of weights grows up to
 157 infinity (e.g., consider its \liminf to define it properly), *or* (b) this running sum of weights
 158 takes value zero infinitely often. As this defines a qualitative objective, the corresponding
 159 preference relation induces only two equivalence classes: *winning* and *losing* plays. The
 160 inverse relation, used by \mathcal{P}_2 , is trivial to obtain. It is fairly easy to prove that \mathcal{P}_1 always
 161 has finite-memory optimal strategies in his one-player games (i.e., games where \mathcal{P}_2 has no
 162 choice), and so does \mathcal{P}_2 in his one-player games.

163 Now, consider the very simple two-player game depicted in Figure 1. Player \mathcal{P}_1 (circle)
 164 has an *infinite-memory* strategy to win: keeping track of the running sum of weights (which
 165 is unbounded, hence the need for infinite memory) and looping in s_1 up to the point where
 166 this sum hits zero, then going to s_2 . This strategy ensures victory because either \mathcal{P}_2 always
 167 goes back to s_1 , in which case (b) is satisfied; or \mathcal{P}_2 eventually loops forever on s_2 , in which
 168 case (a) is satisfied. It remains to argue that \mathcal{P}_1 has no *finite-memory* winning strategy in
 169 this game. This can be done using a standard argument: whatever the amount of memory
 170 used by \mathcal{P}_1 , \mathcal{P}_2 may loop in s_2 long enough as to exceed the bound up to which \mathcal{P}_1 can track
 171 the sum accurately; thus dooming \mathcal{P}_1 to fail to reset the sum to zero in s_1 infinitely often.

172 This modest example proves that *Gimbert and Zielonka’s approach cannot work in full*
 173 *generality in the finite-memory case*, and for good reasons. Informally, in this case, the
 174 corollary breaks down because of (the absence of some sort of) monotony. In the case of

175 *memoryless* strategies, as in [27], \mathcal{P}_1 is already doomed in one-player games in the absence of
 176 monotony: two prefixes to distinguish — in order to play optimally — can be hardcoded as
 177 different paths leading to the same state in a game arena, as if they were chosen by \mathcal{P}_2 in a
 178 two-player game. In the case of *finite-memory* strategies, however, the situation is different.
 179 In one-player games, the number of such paths that can be hardcoded in an arena is always
 180 bounded, hence finite memory might suffice to react, i.e., to keep track of which prefix is the
 181 current one and how to behave accordingly. However, in two-player games, \mathcal{P}_2 might create
 182 an *infinite* number of prefixes to distinguish (using a cycle), thus requiring \mathcal{P}_1 to use infinite
 183 memory to be able to do so. This is exactly what happens in the example above: in any
 184 one-player game, the largest sum that \mathcal{P}_1 has to track is bounded, whereas \mathcal{P}_2 can make this
 185 sum as large as he wants in two-player games.

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

- 191 ■ All memoryless-determined relations — studied in [27] — use arena-independent memory:
 192 the memory required, *none*, is the same for all arenas.
- 193 ■ Combinations of parity objectives use arena-independent memory [16]: the memory only
 194 depends on the number of objectives and the number of priorities — both parameters of
 195 the preference relation, not on the size of the arena.
- 196 ■ Lower- and upper-bounded energy objectives also use arena-independent memory (see,
 197 e.g., [3, 5, 4]): the memory only depends on the bounds and the weights — parameters
 198 of the preference relation, not on the size of the arena.
- 199 ■ On the contrary, combinations of lower-bounded energy objectives (with no upper bound)
 200 require *arena-dependent* memory [17, 32]: it depends on the size of the arena in addition
 201 to the weights used in it. Such a setting falls outside the scope of our results.

202 This informal concept of arena-independent memory is transparent in our work: in all
 203 our results, we use *memory skeletons* — essentially Mealy machines without a next-action
 204 function (Section 2) — that suffice for *all* arenas, and that are at the basis of the strategies
 205 we build. A quick look at our main concepts (Section 3) and results (Section 4) suffices to
 206 grasp the formalism behind this intuition.

207 This restriction to arena-independent memory is natural given the counterexample to a
 208 general approach presented above. It is also important to note that it is not as restrictive as
 209 it may seem, as hinted by the examples above: we are *not restricted to constant memory*
 210 but to memory only depending on the *parameters of the preference relation* (or equivalently,
 211 objective), and not of the arena. This framework thus already encompasses many objectives
 212 from the literature — e.g., [19, 20, 43, 12, 5, 21, 16, 10, 14, 3, 5, 4], as well as possible
 213 extensions. More details in Appendix A.

214 The *arena-independent* case, which we solve here, is an exact equivalent to Gimbert and
 215 Zielonka’s results in the finite-memory case: the memoryless case is de facto arena-independent.
 216 Therefore, this paper strictly generalizes [27] by allowing to study any arena-independent
 217 memory skeleton instead of the unique trivial one corresponding to memoryless strategies.

218 **Outline of our contributions.** Informally, our characterization can be stated as follows:
 219 given a preference relation and a memory skeleton \mathcal{M} , both players have optimal finite-
 220 memory strategies based on skeleton \mathcal{M} in all games if and only if the relation and its inverse

221 are \mathcal{M} -monotone and \mathcal{M} -selective.

222 These last two concepts are keys to our approach. Intuitively, they correspond to Gimbert
 223 and Zielonka’s monotony and selectivity, *modulo a memory skeleton*. Recall that monotony
 224 and selectivity are related to stability of the preference relation with regard to prefix addition
 225 and cycle mixing, respectively. Our more general concepts of \mathcal{M} -*monotony* and \mathcal{M} -*selectivity*
 226 serve the same purpose, but they only compare sequences of colors that are deemed equivalent
 227 by the memory skeleton. For the sake of illustration, take selectivity: it implies that one has
 228 no interest in mixing different cycles of the game arena. For its generalization, the memory
 229 skeleton is taken into account: \mathcal{M} -selectivity implies that one has no interest in mixing cycles
 230 of the game arena *that are read as cycles on the same memory state in the skeleton \mathcal{M}* .

231 Let us give a quick breakdown of our paper. Due to space constraints, we only provide
 232 an intuitive overview of our results and technical approach in this conference version: formal
 233 details and proofs, along with additional results, can be found in the full article [6].

234 In Section 2, we introduce some basic notions, including the memory skeletons, and we
 235 establish several technical results. We also discuss *optimal strategies* and *Nash equilibria*.
 236 In Section 3, we introduce \mathcal{M} -monotony and \mathcal{M} -selectivity, cornerstones of our work. We
 237 also present two essential tools: *prefix-covers* and *cyclic-covers* of arenas. Section 4 states
 238 formally our characterization (Theorem 9), as well as the corresponding lifting corollary
 239 (Corollary 10), from one-player to two-player games. We show an example of application
 240 in Section 5. Finally, we give an overview of the technical highlights of our approach in
 241 Section 6 — its details are broken down in several intermediate results in our full paper [6].

242 In a nutshell, the proof of the characterization (Theorem 9) is split in two. We first
 243 establish that (the sufficiency of) finite memory based on \mathcal{M} implies \mathcal{M} -monotony and
 244 \mathcal{M} -selectivity of the preference relation. The crux is to build game arenas based on automata
 245 recognizing the languages involved in the two concepts, and to use the existence of finite-
 246 memory optimal strategies in these arenas to prove that \mathcal{M} -monotony and \mathcal{M} -selectivity hold.
 247 To prove the converse implication, we proceed in two steps, first establishing the existence of
 248 *memoryless* optimal strategies in “covered” arenas, and then building on it to obtain the
 249 existence of *finite-memory* optimal strategies in general arenas. The main technical tools we
 250 use are Nash equilibria and the aforementioned notions of prefix-covers and cyclic-covers.

251 Alongside the technical details, we analyze our characterization in Appendix A: we
 252 highlight some limitations and interesting features, compare our techniques with Gimbert and
 253 Zielonka’s, discuss our place in the research landscape, and sketch directions for future work.
 254 Let us just stress already that our result — relating a memory skeleton \mathcal{M} and preference
 255 relations for which this skeleton suffices — cannot be obtained by simply considering product
 256 arenas and invoking Gimbert and Zielonka’s result on memoryless determinacy [27]. While,
 257 of course, memoryless strategies on product arenas correspond to memoryfull strategies on
 258 original arenas (as we will formally establish in Lemma 1), invoking [27] requires to be able
 259 to quantify on *all* arenas, not only *product* arenas. Filling this gap is exactly the goal of this
 260 paper, and it is made possible through the new concepts we sketched above.

261 **2 Preliminaries**

262 We only give here the notions and results necessary to understand this overview. Necessary
 263 notions and results — some of them interesting in their own right — to understand the
 264 technical details of the approach are found in the full paper [6].

265 **Automata and languages of colors.** Let C be an arbitrary set of *colors*. We assume knowl-
 266 edge of *non-deterministic finite-state automata (NFA)*, which recognize *regular languages*.
 267 For any finite subset $B \subseteq C$, we denote by $\text{Reg}(B)$ the set of all regular languages over B .
 268 Let $\mathcal{R}(C) = \bigcup_{B \subseteq C, |B| < \infty} \text{Reg}(B)$, that is, all the regular languages built over C .

269 Let $K \subseteq C^*$ be a language of *finite* words. We write $\text{Prefs}(K)$ for the set of all *prefixes* of
 270 words in K . We define the set of *infinite* words $[K] = \{w = c_1 \dots \in C^\omega \mid \forall n \geq 1, c_1 \dots c_n \in$
 271 $\text{Prefs}(K)\}$, which contains all infinite words for which every finite prefix is a prefix of a word
 272 in K . Intuitively, if K is regular, $[K]$ is the language of infinite words that correspond to
 273 infinite paths that can always branch and reach a final state, on an automaton for K . Given
 274 a finite word $w \in C^*$ and a language $K \subseteq C^*$, we write wK for their concatenation, i.e., the
 275 language $wK = \{w' = ww'' \mid w'' \in K\} \subseteq C^*$.

276 **Arenas.** We consider two players: player 1 (\mathcal{P}_1) and player 2 (\mathcal{P}_2). An *arena* is a tuple
 277 $\mathcal{A} = (S_1, S_2, E)$ such that $S = S_1 \uplus S_2$ (disjoint union) is a finite set of *states* partitioned into
 278 states of \mathcal{P}_1 (S_1) and \mathcal{P}_2 (S_2), and $E \subseteq S \times C \times S$ is a finite set of *edges*. Let $\text{col}: E \rightarrow C$ be
 279 the projection of edges to *colors* and $\widehat{\text{col}}$ its natural extension to *sequences* of edges. For an
 280 edge $e \in E$, we use $\text{in}(e)$ and $\text{out}(e)$ to denote its starting state and arrival state respectively,
 281 i.e., $e = (\text{in}(e), \text{col}(e), \text{out}(e))$. We assume all arenas to be non-blocking, i.e., for all $s \in S$,
 282 there exists $e \in E$ such that $\text{in}(e) = s$. For $i \in \{1, 2\}$, we call an arena $\mathcal{A} = (S_1, S_2, E)$ a \mathcal{P}_i 's
 283 *one-player arena* if for all $s \in S_{3-i}$, $|\{e \in E \mid \text{in}(e) = s\}| = 1$ — that is, \mathcal{P}_{3-i} has no choice.

284 Let $\text{Hists}(\mathcal{A}, s)$ denote the *histories* in \mathcal{A} from $s \in S$, i.e., finite sequences of edges
 285 $\rho = e_1 \dots e_n \in E^+$ such that $\text{in}(e_1) = s$ and for all i , $1 \leq i < n$, $\text{out}(e_i) = \text{in}(e_{i+1})$. Let
 286 $\text{Plays}(\mathcal{A}, s)$ denote the *plays* in \mathcal{A} from $s \in S$, i.e., infinite sequences $\pi = e_1 e_2 \dots \in E^\omega$ such
 287 that $\text{in}(e_1) = s$ and for all $i \geq 1$, $\text{out}(e_i) = \text{in}(e_{i+1})$. We write $\text{Hists}(\mathcal{A}, S')$ and $\text{Plays}(\mathcal{A}, S')$
 288 for unions over $S' \subseteq S$, and write $\text{Hists}(\mathcal{A})$ and $\text{Plays}(\mathcal{A})$ for the unions over all states of \mathcal{A} .

289 Let $\rho = e_1 \dots e_n \in \text{Hists}(\mathcal{A})$ (resp. $\pi = e_1 e_2 \dots \in \text{Plays}(\mathcal{A})$): we extend the operator in
 290 to histories (resp. plays) by identifying $\text{in}(\rho)$ (resp. $\text{in}(\pi)$) to $\text{in}(e_1)$. We proceed similarly
 291 for out and histories: $\text{out}(\rho) = \text{out}(e_n)$. For the sake of convenience, we consider that any
 292 set $\text{Hists}(\mathcal{A}, s)$ contains the *empty history* λ_s such that $\text{in}(\lambda_s) = \text{out}(\lambda_s) = s$. We write
 293 $\text{Hists}_i(\mathcal{A}, s)$ and $\text{Hists}_i(\mathcal{A})$ for the subsets of histories ρ such that $\text{out}(\rho) \in S_i$, $i \in \{1, 2\}$, i.e.,
 294 histories whose last state belongs to \mathcal{P}_i . For any set $H \subseteq \text{Hists}(\mathcal{A})$, we write $\widehat{\text{col}}(H)$ for its
 295 projection to colors, i.e., $\widehat{\text{col}}(H) = \{\widehat{\text{col}}(\rho) \mid \rho \in H\}$. We do the same for sets of plays.

296 **Memory skeletons.** A *memory skeleton* is a tuple $\mathcal{M} = (M, m_{\text{init}}, \alpha_{\text{upd}})$ where M is a *finite*
 297 set of *states*, $m_{\text{init}} \in M$ is a fixed *initial* state and $\alpha_{\text{upd}}: M \times C \rightarrow M$ is an *update function*.
 298 We write $\widehat{\alpha}_{\text{upd}}$ for the natural extension of α_{upd} to sequences in C^* . Memory skeletons are
 299 deterministic and might have an *infinite* number of transitions, in contrast to NFA. We define
 300 the *trivial memory skeleton* as $\mathcal{M}_{\text{triv}} = (M = \{m_{\text{init}}\}, m_{\text{init}}, \alpha_{\text{upd}}: \{m_{\text{init}}\} \times C \rightarrow \{m_{\text{init}}\})$: it
 301 permits to formalize memoryless strategies [27] in our framework.

302 Let $\mathcal{M} = (M, m_{\text{init}}, \alpha_{\text{upd}})$ be a skeleton. For $m, m' \in M$, we define the language $L_{m, m'} =$
 303 $\{w \in C^* \mid \widehat{\alpha}_{\text{upd}}(m, w) = m'\}$ that contains all words that can be read from m to m' in \mathcal{M} .

304 Let $\mathcal{M}^1 = (M^1, m_{\text{init}}^1, \alpha_{\text{upd}}^1)$ and $\mathcal{M}^2 = (M^2, m_{\text{init}}^2, \alpha_{\text{upd}}^2)$. Their *product* $\mathcal{M}^1 \otimes \mathcal{M}^2$ is the
 305 memory skeleton $\mathcal{M} = (M, m_{\text{init}}, \alpha_{\text{upd}})$ where $M = M^1 \times M^2$, $m_{\text{init}} = (m_{\text{init}}^1, m_{\text{init}}^2)$, and, for
 306 all $m^1 \in M^1$, $m^2 \in M^2$, $c \in C$, $\alpha_{\text{upd}}((m^1, m^2), c) = (\alpha_{\text{upd}}^1(m^1, c), \alpha_{\text{upd}}^2(m^2, c))$. That is, the
 307 memories are updated in parallel when a color is read.

308 **Product arenas.** Let $\mathcal{A} = (S_1, S_2, E)$ be an arena and $\mathcal{M} = (M, m_{\text{init}}, \alpha_{\text{upd}})$ be a skeleton.
 309 Their *product* $\mathcal{A} \times \mathcal{M}$ is the arena (S'_1, S'_2, E') where $S'_1 = S_1 \times M$, $S'_2 = S_2 \times M$, and

310 $E' \subseteq S' \times C \times S'$, with $S' = S'_1 \uplus S'_2$, is such that $((s_1, m_1), c, (s_2, m_2)) \in E'$ if and only if
 311 $(s_1, c, s_2) \in E$ and $\alpha_{\text{upd}}(m_1, c) = m_2$. That is, the memory is updated according to the colors
 312 of the edges in E . Though \mathcal{M} might contain an infinite number of transitions, $\mathcal{A} \times \mathcal{M}$ is
 313 always finite, as E is finite. Since we assume \mathcal{A} is non-blocking, it is also the case of $\mathcal{A} \times \mathcal{M}$.

314 **Strategies.** A strategy σ_i for \mathcal{P}_i , $i \in \{1, 2\}$, on arena $\mathcal{A} = (S_1, S_2, E)$, is a function
 315 $\sigma_i: \text{Hists}_i(\mathcal{A}) \rightarrow E$ such that for all $\rho \in \text{Hists}_i(\mathcal{A})$, $\text{in}(\sigma_i(\rho)) = \text{out}(\rho)$. Let $\Sigma_i(\mathcal{A})$ be the set
 316 of all strategies of \mathcal{P}_i on \mathcal{A} . A *finite-memory strategy* σ_i can be encoded as a *Mealy machine*,
 317 i.e., a skeleton $\mathcal{M} = (M, m_{\text{init}}, \alpha_{\text{upd}})$ with transitions over a *finite subset of colors* $B \subseteq C$,
 318 enriched with a *next-action function* $\alpha_{\text{next}}: M \times S_i \rightarrow E$ such that for all $m \in M$, $s \in S_i$,
 319 $\text{in}(\alpha_{\text{next}}(m, s)) = s$. Given a Mealy machine $\Gamma_{\sigma_i} = (\mathcal{M}, \alpha_{\text{next}})$, strategy σ_i is defined as follows:

- 320 ■ $\forall s \in S_i, \sigma_i(\lambda_s) = \alpha_{\text{next}}(m_{\text{init}}, s)$,
- 321 ■ $\forall \rho \cdot e \in \text{Hists}_i(\mathcal{A}), e \in E, \sigma_i(\rho \cdot e) = \alpha_{\text{next}}(\widehat{\alpha_{\text{upd}}}(m_{\text{init}}, \widehat{\text{col}}(\rho \cdot e)), \text{out}(e))$.

322 Let $\Sigma_i^{\text{FM}}(\mathcal{A})$ be the finite-memory strategies of \mathcal{P}_i on \mathcal{A} . We say that a strategy $\sigma_i \in \Sigma_i^{\text{FM}}(\mathcal{A})$
 323 is *based on memory skeleton* \mathcal{M} if it can be encoded as a Mealy machine $\Gamma_{\sigma_i} = (\mathcal{M}, \alpha_{\text{next}})$, as
 324 above. We implicitly assume that strategies of $\Sigma_i^{\text{FM}}(\mathcal{A})$ are built by restricting the transitions
 325 of their skeleton \mathcal{M} to the actual subset of colors appearing in \mathcal{A} . A strategy σ_i is *memoryless*
 326 if it is a function $\sigma_i: S_i \rightarrow E$, or equivalently, if it is based on $\mathcal{M}_{\text{triv}}$.

327 We write $\text{Plays}(\mathcal{A}, s, \sigma_i)$ for the plays *consistent* with a strategy σ_i of \mathcal{P}_i from a state s , i.e.,
 328 plays $\pi = e_1 e_2 \dots \in \text{Plays}(\mathcal{A}, s)$ such that for all $\rho = e_1 \dots e_n$, $\text{out}(\rho) \in S_i \implies \sigma_i(\rho) = e_{n+1}$.
 329 We write $\text{Plays}(\mathcal{A}, s, \sigma_1, \sigma_2)$ for the singleton set containing the unique play consistent with a
 330 couple of strategies for the two players. We use similar notations for histories.

331 **Preference relations.** Let \sqsubseteq be a total preorder on C^ω , called *preference relation*. We
 332 consider *antagonistic* games, where the objective of \mathcal{P}_1 is to create the best possible play
 333 with regard to \sqsubseteq whereas the objective of \mathcal{P}_2 is the opposite. That is, \mathcal{P}_2 uses the inverse
 334 relation \sqsubseteq^{-1} . This corresponds to *zero-sum* games when using a quantitative framework.

335 Given $w, w' \in C^\omega$, we write $w \sqsubset w'$ if we have $\neg(w' \sqsubseteq w)$ since the preorder is total. We
 336 extend \sqsubseteq to subsets: for $W, W' \subseteq C^\omega$, $W \sqsubseteq W' \iff \forall w \in W, \exists w' \in W', w \sqsubseteq w'$. We also
 337 write $W \sqsubset W' \iff \exists w' \in W', \forall w \in W, w \sqsubset w'$. Note that $W \sqsubset W' \iff \neg(W' \sqsubseteq W)$.

338 To compare a word $w \in C^\omega$ with a language $K \subseteq C^\omega$, we simply identify it to $\{w\}$.

339 **Games.** A (deterministic turn-based two-player) *game* is a tuple $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$ where \mathcal{A} is
 340 an arena and \sqsubseteq is a preference relation. All classical objectives from the literature (both
 341 qualitative and quantitative) can be expressed in the general framework of preference relations
 342 (see Example 3 in [6]). For $i \in \{1, 2\}$, a \mathcal{P}_i 's *one-player game* is a game $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$ such that
 343 \mathcal{A} is a \mathcal{P}_i 's one-player arena.

344 **Optimal strategies.** Let $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$ be a game on arena $\mathcal{A} = (S_1, S_2, E)$. Given a \mathcal{P}_i -
 345 strategy $\sigma_i \in \Sigma_i(\mathcal{A})$ and a state $s \in S$, we define

$$346 \quad \text{UCol}_{\sqsubseteq}(\mathcal{A}, s, \sigma_i) = \{w \in C^\omega \mid \exists \sigma_{3-i} \in \Sigma_{3-i}(\mathcal{A}), \widehat{\text{col}}(\text{Plays}(\mathcal{A}, s, \sigma_1, \sigma_2)) \sqsubseteq w\},$$

$$347 \quad \text{DCol}_{\sqsubseteq}(\mathcal{A}, s, \sigma_i) = \{w \in C^\omega \mid \exists \sigma_{3-i} \in \Sigma_{3-i}(\mathcal{A}), w \sqsubseteq \widehat{\text{col}}(\text{Plays}(\mathcal{A}, s, \sigma_1, \sigma_2))\}.$$

349 Note that $\text{DCol}_{\sqsubseteq}(\mathcal{A}, s, \sigma_i) = \text{UCol}_{\sqsubseteq^{-1}}(\mathcal{A}, s, \sigma_i)$. Intuitively, $\text{UCol}_{\sqsubseteq}$ and $\text{DCol}_{\sqsubseteq}$ represent the
 350 *upward* and *downward closures* of sequences of colors (consistent with a strategy) with respect
 351 to the preference relation.

352 Taking the standpoint of \mathcal{P}_1 , we say that $\sigma_1 \in \Sigma_1(\mathcal{A})$ is *at least as good as* $\sigma'_1 \in \Sigma_1(\mathcal{A})$
 353 from $s \in S$ if $\text{UCol}_{\sqsubseteq}(\mathcal{A}, s, \sigma_1) \subseteq \text{UCol}_{\sqsubseteq}(\mathcal{A}, s, \sigma'_1)$. Intuitively, σ_1 is at least as good as σ'_1 if
 354 the “worst-case” plays consistent with σ_1 are at least as good as the ones consistent with σ'_1 .
 355 The UCol operator is useful to define this notion properly even in the case where there is no
 356 “worst-case” play (i.e., if the infimum used in the classical quantitative setting is not reached).
 357 Similar notions have been used before, e.g., in [38]. Symmetrically, for \mathcal{P}_2 , we say that
 358 $\sigma_2 \in \Sigma_2(\mathcal{A})$ is at least as good as $\sigma'_2 \in \Sigma_2(\mathcal{A})$ from $s \in S$ if $\text{DCol}_{\sqsubseteq}(\mathcal{A}, s, \sigma_2) \subseteq \text{DCol}_{\sqsubseteq}(\mathcal{A}, s, \sigma'_2)$.

359 Now, we say that $\sigma_i \in \Sigma_i(\mathcal{A})$ of \mathcal{P}_i is *optimal* from $s \in S$, aka *s-optimal*, if it is at least
 360 as good as every other $\sigma'_i \in \Sigma_i(\mathcal{A})$ from s . We extend this notation to subsets of states in
 361 the natural way, and we say that a strategy σ_i is *uniformly-optimal* if it is *S-optimal*.

362 Our goal is to characterize the preference relations that admit *uniformly-optimal finite-*
 363 *memory (UFM) strategies based on a given skeleton* \mathcal{M} in all arenas. We also discuss the
 364 simpler case of *uniformly-optimal memoryless (UML) strategies*, which corresponds to the
 365 subcase studied by Gimbert and Zielonka [27], using the trivial skeleton $\mathcal{M}_{\text{triv}}$.

366 In that respect, the following link is important to observe.

367 ► **Lemma 1.** *Let $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$ be a game on arena $\mathcal{A} = (S_1, S_2, E)$. Let $\mathcal{M} = (M, m_{\text{init}}, \alpha_{\text{upd}})$
 368 be a memory skeleton and let $\sigma_i \in \Sigma_i^{\text{FM}}(\mathcal{A})$ be a finite-memory strategy encoded by the Mealy
 369 machine $\Gamma_{\sigma_i} = (\mathcal{M}, \alpha_{\text{next}})$. Then, σ_i is a UFM strategy in \mathcal{G} if and only if α_{next} corresponds to
 370 an $(S \times \{m_{\text{init}}\})$ -optimal memoryless strategy in $\mathcal{G}' = (\mathcal{A} \times \mathcal{M}, \sqsubseteq)$.*

371 **Nash equilibria.** We use Nash equilibria [36] as tools to establish the existence of optimal
 372 strategies in some of our proofs. Let $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$ be a game on arena $\mathcal{A} = (S_1, S_2, E)$. Formally,
 373 a *Nash equilibrium* (NE) from a state $s \in S$ is a couple of strategies $(\sigma_1, \sigma_2) \in \Sigma_1(\mathcal{A}) \times \Sigma_2(\mathcal{A})$
 374 such that, for all $\sigma'_1 \in \Sigma_1(\mathcal{A})$, $\sigma'_2 \in \Sigma_2(\mathcal{A})$,

$$375 \quad \widehat{\text{col}}(\text{Plays}(\mathcal{A}, s, \sigma'_1, \sigma_2)) \sqsubseteq \widehat{\text{col}}(\text{Plays}(\mathcal{A}, s, \sigma_1, \sigma_2)) \sqsubseteq \widehat{\text{col}}(\text{Plays}(\mathcal{A}, s, \sigma_1, \sigma'_2)). \quad (1)$$

376 Similarly to optimal strategies, we call an NE *uniform* if it is an NE from all states $s \in S$.

377 The connection between *optimal strategies* and *Nash equilibria* in our specific context of
 378 *antagonistic* games is interesting to discuss, especially with respect to Gimbert and Zielonka’s
 379 original work [27]. We defer this discussion to [6] due to space constraints, and only provide
 380 here a brief account of the results one has to know to understand this overview. First, NE
 381 are de facto pairs of optimal strategies. Second, it is possible to mix different NE.

382 ► **Lemma 2.** *Let $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$ be a game on arena $\mathcal{A} = (S_1, S_2, E)$, and let $s \in S$ be a state.
 383 Let (σ_1^a, σ_2^a) and $(\sigma_1^b, \sigma_2^b) \in \Sigma_1(\mathcal{A}) \times \Sigma_2(\mathcal{A})$ be two Nash equilibria from s . Then, (σ_1^a, σ_2^b) is
 384 also a Nash equilibrium from s .*

385 ► **Remark 3.** Lemma 2 crucially relies on the assumption (transparent in our definition of
 386 Nash equilibrium) that we consider *antagonistic* games, that is, \mathcal{P}_2 uses the inverse preference
 387 relation \sqsubseteq^{-1} . ◁

388 3 Concepts

389 **Generalizing monotony and selectivity.** As seen in Section 1, Gimbert and Zielonka’s
 390 characterization [27] relies on *monotony* and *selectivity* of the preference relation. The main
 391 difference between their technical approach and ours is the following. In the memoryless
 392 setting, all the reasoning can be *abstracted away* from the underlying arena and done on
 393 sequences of colors. In the finite-memory one, however, one has to pay attention to how

394 sequences of colors are composed and compared, to maintain consistency with regard to
 395 the memory and the game arena. This need to intertwine abstract reasoning on arbitrary
 396 sequences of colors with concrete tracking of memory updates is the key obstacle to overcome.

397 Much of our effort was thus spent on trying to define concepts that would preserve the
 398 elegance of monotony and selectivity while allowing us to lift the theory to the finite-memory
 399 case. As often the case, the good concepts turned out to be the most natural ones, capturing
 400 the intuitive idea that one needs monotony and selectivity *modulo a memory skeleton*.

401 ► **Definition 4** (\mathcal{M} -monotony). *Let $\mathcal{M} = (M, m_{\text{init}}, \alpha_{\text{upd}})$ be a memory skeleton. A preference
 402 relation \sqsubseteq is \mathcal{M} -monotone if for all $m \in M$, for all $K_1, K_2 \in \mathcal{R}(C)$,*

$$403 \quad (\exists w \in L_{m_{\text{init}}, m}, [wK_1] \sqsubset [wK_2]) \implies (\forall w' \in L_{m_{\text{init}}, m}, [w'K_1] \sqsubseteq [w'K_2]). \quad (2)$$

404 Recall that a skeleton \mathcal{M} has a fixed initial state m_{init} . Intuitively, \mathcal{M} -monotony extends
 405 Gimbert and Zielonka’s monotony by comparing prefixes *belonging to the same language*
 406 $L_{m_{\text{init}}, m}$, that is, prefixes that are deemed equivalent by skeleton \mathcal{M} . This property roughly
 407 captures that \sqsubseteq is *stable with regard to prefix addition*, for memory-equivalent prefixes.

408 The original monotony notion is equivalent to our \mathcal{M} -monotony with \mathcal{M} being the trivial
 409 skeleton $\mathcal{M}_{\text{triv}}$: that is, the memoryless case is naturally a subcase of our framework.

410 ► **Definition 5** (\mathcal{M} -selectivity). *Let $\mathcal{M} = (M, m_{\text{init}}, \alpha_{\text{upd}})$ be a memory skeleton. A preference
 411 relation \sqsubseteq is \mathcal{M} -selective if for all $w \in C^*$, $m = \widehat{\alpha_{\text{upd}}}(m_{\text{init}}, w)$, for all $K_1, K_2 \in \mathcal{R}(C)$ such
 412 that $K_1, K_2 \subseteq L_{m, m}$, for all $K_3 \in \mathcal{R}(C)$,*

$$413 \quad [w(K_1 \cup K_2)^* K_3] \sqsubseteq [wK_1^*] \cup [wK_2^*] \cup [wK_3]. \quad (3)$$

414 Similarly, \mathcal{M} -selectivity extends Gimbert and Zielonka’s selectivity by asking one to
 415 compare sequences of colors *belonging to the same language* $L_{m, m}$, that is, sequences read as
 416 cycles on the memory skeleton. Note also that the memory state m should be consistent
 417 with the prefix w read from the initial memory state m_{init} . This property roughly captures
 418 that \sqsubseteq is *stable with regard to cycle mixing*, for memory-equivalent cycles.

419 Again, the original selectivity notion is exactly equivalent to $\mathcal{M}_{\text{triv}}$ -selectivity.

420 In a nutshell, \mathcal{M} -monotony deals with prefixes up to the first cycle (on memory) and
 421 \mathcal{M} -selectivity deals with the cycles thereafter; we will see that memory skeletons can be built
 422 in a compositional way based on these two orthogonal yet complementary tasks.

423 Our notions respect the natural intuition that access to additional memory should always
 424 be helpful: if a skeleton \mathcal{M} is sufficient to classify sequences of colors in a way that guarantees
 425 \mathcal{M} -monotony and \mathcal{M} -selectivity, then it should also be the case for “more powerful” skeletons.

426 ► **Lemma 6.** *Let \mathcal{M} and \mathcal{M}' be two memory skeletons. If \sqsubseteq is \mathcal{M} -monotone (resp. \mathcal{M} -
 427 selective) then, it is also $(\mathcal{M} \otimes \mathcal{M}')$ -monotone (resp. $(\mathcal{M} \otimes \mathcal{M}')$ -selective).*

428 **Prefix-covers and cyclic-covers.** While the concepts of \mathcal{M} -monotony and \mathcal{M} -selectivity
 429 are the primordial ones for stating the characterization, we still need two additional notions
 430 to prove it. Let us sketch the issue. To prove that monotone and selective preference relations
 431 yield UML strategies, Gimbert and Zielonka deploy an *inductive argument* on the number
 432 of choices in an arena. Intuitively, we want to use a similar approach for UFM strategies,
 433 but because of the unavoidable coupling between the memory skeleton and the arena (e.g.,
 434 Lemma 1), the induction argument breaks, as adding one choice in the arena results in adding
 435 many in the *product arena* (as many as there are memory states), where the reasoning needs
 436 to take place. New insight and techniques are thus needed to patch this scheme.

437 To solve this issue, we decouple the two aspects. We first establish that, on arenas that
 438 inherently share the same good properties as product arenas (i.e., they already “classify”
 439 prefixes and cycles as the memory would), we can deploy the induction argument and obtain
 440 UML strategies. Then, we obtain UFM strategies on *general* arenas as a corollary. The crux
 441 is identifying such “good” arenas: this is done through the following notions.

442 ► **Definition 7** (Prefix-covers and cyclic-covers). Let $\mathcal{M} = (M, m_{\text{init}}, \alpha_{\text{upd}})$ be a memory
 443 skeleton and $\mathcal{A} = (S_1, S_2, E)$ be an arena. Let $S_{\text{cov}} \subseteq S$.

444 We say that \mathcal{M} is a prefix-cover of S_{cov} in \mathcal{A} if for all $s \in S$, there exists $m_s \in M$ such
 445 that, for all $\rho \in \text{Hists}(\mathcal{A})$ such that $\text{in}(\rho) \in S_{\text{cov}}$, $\text{out}(\rho) = s$ and such that for all ρ' proper
 446 prefix of ρ , $\text{out}(\rho') \neq s$, we have $\widehat{\alpha_{\text{upd}}}(m_{\text{init}}, \widehat{\text{col}}(\rho)) = m_s$.

447 We say that \mathcal{M} is a cyclic-cover of S_{cov} in \mathcal{A} if for all $\rho \in \text{Hists}(\mathcal{A})$ such that $\text{in}(\rho) \in S_{\text{cov}}$,
 448 if $s = \text{out}(\rho)$ and $m = \widehat{\alpha_{\text{upd}}}(m_{\text{init}}, \widehat{\text{col}}(\rho))$, for all $\rho' \in \text{Hists}(\mathcal{A})$ such that $\text{in}(\rho') = \text{out}(\rho') = s$,
 449 $\widehat{\alpha_{\text{upd}}}(m, \widehat{\text{col}}(\rho')) = m$.

450 Intuitively, \mathcal{M} is a prefix-cover for a set of states S_{cov} if the histories starting in S_{cov} and
 451 visiting a given state $s \in S$ for the first time are read up to the same memory state in the
 452 memory skeleton. Similarly, \mathcal{M} is a cyclic-cover of \mathcal{A} if the cycles of \mathcal{A} are read as cycles in
 453 the memory skeleton, once the memory has been initialized properly.

454 As hinted above, the canonical example of a prefix- and cyclic-covered arena is a product
 455 arena (but many more may be in this case; it is beneficial to be general with these concepts).

456 ► **Lemma 8.** Let $\mathcal{M} = (M, m_{\text{init}}, \alpha_{\text{upd}})$ be a memory skeleton and $\mathcal{A} = (S_1, S_2, E)$ be an
 457 arena. Then \mathcal{M} is a prefix- and cyclic-cover for $S_{\text{cov}} = S \times \{m_{\text{init}}\}$ in $\mathcal{A} \times \mathcal{M}$.

458 4 Characterization

459 **Equivalence.** We now have the necessary ingredients to state our equivalence result.

460 ► **Theorem 9** (Equivalence). Let \sqsubseteq be a preference relation and let \mathcal{M} be a memory skeleton.
 461 Then, both players have UFM strategies based on memory skeleton \mathcal{M} in all games $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$
 462 if and only if \sqsubseteq and \sqsubseteq^{-1} are \mathcal{M} -monotone and \mathcal{M} -selective.

463 We state this theorem broadly and with a *focus on UFM strategies*. The actual results
 464 we have for each direction of the equivalence — see [6, Section 4 and Section 5] — are a
 465 bit stronger, of wider applicability and/or more interesting, but this statement carries the
 466 take-home message of our work. It is also meant to mirror the seminal result of Gimbert
 467 and Zielonka [27, Theorem 2]: their result can be retrieved from Theorem 9 by taking the
 468 trivial memory skeleton $\mathcal{M}_{\text{triv}}$. As such, our work brings a *strict generalization of Gimbert*
 469 *and Zielonka’s results* [27] to the finite-memory case.

470 **Lifting corollary.** As discussed in Section 1, the work of Gimbert and Zielonka contains
 471 not one, but *two* great results. Alongside the aforementioned equivalence result, Gimbert
 472 and Zielonka provide a corollary of high practical interest [27, Corollary 7]: they essentially
 473 obtain as a by-product of their approach that if memoryless strategies suffice in all one-player
 474 games of \mathcal{P}_1 and all one-player games of \mathcal{P}_2 , they also suffice in all two-player games.

475 This provides an elegant way to prove that a preference relation (equivalently, an objective)
 476 admits memoryless optimal strategies *without proving monotony and selectivity at all*: proving
 477 it in the two one-player subcases, which is generally much easier as it boils down to graph
 478 reasoning, and then lifting the result to the general two-player case through the corollary.
 479 See examples of one-player vs. two-player complexity in [5, 4, 11].

480 Again, we are able to lift this corollary to the arena-independent finite-memory case.

481 ► **Corollary 10.** Let \sqsubseteq be a preference relation and $\mathcal{M}_1, \mathcal{M}_2$ be two memory skeletons.
 482 Assume that

- 483 1. for all one-player arenas $\mathcal{A} = (S_1, S_2 = \emptyset, E)$, \mathcal{P}_1 has a UFM strategy $\sigma_1 \in \Sigma_1^{\text{FM}}(\mathcal{A})$ based
 484 on memory skeleton \mathcal{M}_1 in $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$;
- 485 2. for all one-player arenas $\mathcal{A} = (S_1 = \emptyset, S_2, E)$, \mathcal{P}_2 has a UFM strategy $\sigma_2 \in \Sigma_2^{\text{FM}}(\mathcal{A})$ based
 486 on memory skeleton \mathcal{M}_2 in $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$.

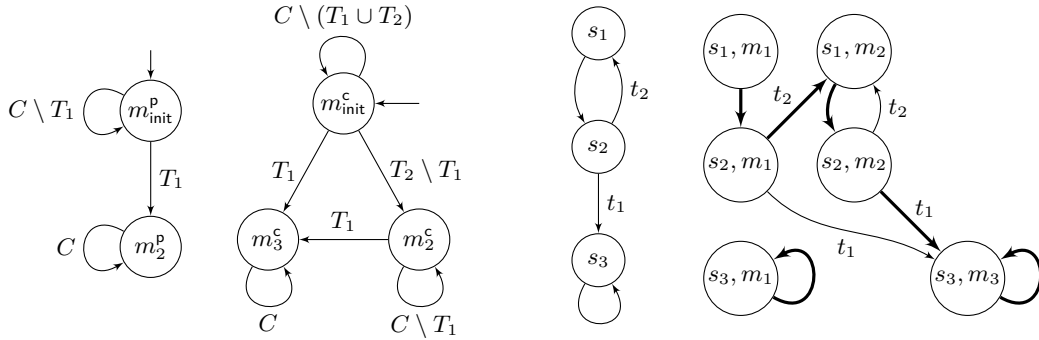
487 Then, for all two-player arenas $\mathcal{A} = (S_1, S_2, E)$, both \mathcal{P}_1 and \mathcal{P}_2 have UFM strategies
 488 $\sigma_i \in \Sigma_i^{\text{FM}}(\mathcal{A})$ based on memory skeleton $\mathcal{M} = \mathcal{M}_1 \otimes \mathcal{M}_2$ in $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$.

489 We highlight the two (possibly different) skeletons of the two players to maintain a
 490 compositional approach, but if the same skeleton \mathcal{M} works in both one-player versions, it
 491 also suffices in the two-player version.

492 5 Example of application

493 We present an illustrative application, thereby proving the existence of UFM strategies for
 494 a specific preference relation: the conjunction of two reachability objectives, a subcase of
 495 *generalized reachability games*, studied extensively in [21]. Let C be an arbitrary set of colors,
 496 and $T_1, T_2 \subseteq C$ be two target sets of colors that have to be visited. Formally, let $W \subseteq C^\omega$
 497 be the set of words $w = c_1 c_2 \dots$ such that $\exists i, j \in \mathbb{N}, c_i \in T_1 \wedge c_j \in T_2$. This winning condition
 498 induces a two-level (i.e., win/lose) preference relation \sqsubseteq .

499 In this example, we will use Theorem 9 directly in order to provide one thorough
 500 illustration of the definitions of \mathcal{M} -monotony and \mathcal{M} -selectivity. However, in practice, using
 501 Corollary 10 is preferable, as it yields a much shorter proof: by exhibiting the right skeletons
 502 for \mathcal{P}_1 and \mathcal{P}_2 , we simply have to show that these skeletons are sufficient to play optimally
 503 on both players' *one-player arenas*, which amounts to graph reasoning.



■ **Figure 2** First and second: memory skeletons \mathcal{M}^P and \mathcal{M}^C for two-target reachability games; third: arena \mathcal{A} ; fourth: product arena $\mathcal{A} \times \mathcal{M}$ (only states reachable from $S \times \{m_{\text{init}}\}$ are depicted). We assume that $T_1 = \{t_1\}$, $T_2 = \{t_2\}$. The $(S \times \{m_{\text{init}}\})$ -optimal memoryless strategy is in bold.

504 We start by showing that this preference relation is not $\mathcal{M}_{\text{triv}}$ -monotone (that is, is not
 505 monotone for [27]). Assume $c_1 \in T_1 \setminus T_2$, $c_2 \in T_2 \setminus T_1$, and $c_3 \notin T_1 \cup T_2$. Take $K_1 = c_1^*$,
 506 $K_2 = c_2^*$. For $w = c_1$, $w' = c_2$, we have $[wK_1] \sqsubseteq [wK_2]$, but $[w'K_2] \sqsubseteq [w'K_1]$. This means
 507 that the preference relation is not stable with regard to prefix addition (at least, without
 508 distinguishing different classes of prefixes). Similarly, it is not $\mathcal{M}_{\text{triv}}$ -selective (take w as the
 509 empty word, $K_1 = c_1^*$, $K_2 = c_2^*$, $K_3 = c_3^*$: to win, K_1 and K_2 need to be mixed).

510 In Figure 2, we exhibit skeletons $\mathcal{M}^P = (M^P, m_{\text{init}}^P, \alpha_{\text{upd}}^P)$ and $\mathcal{M}^C = (M^C, m_{\text{init}}^C, \alpha_{\text{upd}}^C)$ such
 511 that \sqsubseteq is \mathcal{M}^P -monotone and \mathcal{M}^C -selective. Note that such skeletons are obviously not unique.

512 Let us prove that \sqsubseteq is \mathcal{M}^P -monotone. Let $m \in M^P$, $K_1, K_2 \in \mathcal{R}(C)$; we want to show that
 513 Equation (2) is satisfied. We assume that there exists $w \in L_{m_{\text{init}}, m}^P$ such that $[wK_1] \sqsubseteq [wK_2]$;
 514 this means that all words of $[wK_1]$ are losing, and that there exists a winning word in $[wK_2]$.
 515 Let $w' \in L_{m_{\text{init}}, m}^P$; we show that we necessarily have that $[w'K_1] \sqsubseteq [w'K_2]$. Note that if
 516 $[K_1]$ is empty, this always holds; we now assume that $[K_1]$ is non-empty. We study the two
 517 possible values of m separately.

- 518 ■ If $m = m_{\text{init}}^P$, then w and w' do not reach T_1 . If w does not reach T_2 either, as there
 519 is a winning word in $[wK_2]$, then there must be a winning word in $[K_2]$. This word
 520 is still winning after prepending w' to it, so there is a winning word in $[w'K_2]$, and
 521 $[w'K_1] \sqsubseteq [w'K_2]$. If w reaches T_2 , then $[K_1]$ cannot have a word reaching T_1 . As w' does
 522 not reach T_1 either, all words of $[w'K_1]$ are losing, so $[w'K_1] \sqsubseteq [w'K_2]$.
- 523 ■ If $m = m_2^P$, then w and w' reach T_1 . Clearly, w cannot reach T_2 (as $[wK_1]$ would be
 524 winning). This implies that $[K_2]$ must contain a word reaching T_2 ; as w' reaches T_1 , the
 525 concatenation of w' with the word of $[K_2]$ reaching T_2 means that there is a winning
 526 word in $[w'K_2]$, so $[w'K_1] \sqsubseteq [w'K_2]$.

527 Let us now prove that \sqsubseteq is \mathcal{M}^C -selective. Let $w \in C^*$, $m = \widehat{\alpha_{\text{upd}}^C}(m_{\text{init}}^C, w)$, $K_1, K_2 \in \mathcal{R}(C)$
 528 such that $K_1, K_2 \subseteq L_{m, m}$, and $K_3 \in \mathcal{R}(C)$. We show that Equation (3) is satisfied, i.e., that
 529 $[w(K_1 \cup K_2)^*K_3] \sqsubseteq [wK_1^*] \cup [wK_2^*] \cup [wK_3]$. If all words of $[w(K_1 \cup K_2)^*K_3]$ are losing, this
 530 equation trivially holds; we thus assume that this set contains a winning word. We therefore
 531 have to show that there is a winning word in $[wK_1^*]$, $[wK_2^*]$, or $[wK_3]$. We study the three
 532 possible values of m separately.

- 533 ■ If $m = m_{\text{init}}^C$, then w does not reach T_1 nor T_2 , and the same holds for all words of K_1
 534 and K_2 , as $K_1, K_2 \subseteq L_{m_{\text{init}}^C, m_{\text{init}}^C}$. Therefore, if a word of $[w(K_1 \cup K_2)^*K_3]$ is winning, this
 535 must be because a word of $[wK_3]$ is winning.
- 536 ■ If $m = m_2^C$, then w reaches T_2 but not T_1 , and K_1, K_2 do not reach T_1 . Thus, a word of
 537 $[K_3]$ must reach T_1 ; in particular, a word of $[wK_3]$ must reach both T_1 and T_2 .
- 538 ■ If $m = m_3^C$, we distinguish two cases. If w reaches T_2 and T_1 , then $[wK_1^*] \cup [wK_2^*] \cup [wK_3]$
 539 trivially contains only winning words. If w reaches T_1 but not T_2 , then there must be a
 540 word reaching T_2 in $[(K_1 \cup K_2)^*K_3]$. Hence, at least one set among $[K_1^*]$, $[K_2^*]$, and $[K_3]$
 541 must contain a word reaching T_2 , so $[wK_1^*]$, $[wK_2^*]$, or $[wK_3]$ contains a winning word.

542 Similar arguments can be laid out to show that the preference relation \sqsubseteq^{-1} of \mathcal{P}_2 is
 543 \mathcal{M}^P -monotone and $\mathcal{M}_{\text{triv}}$ -selective (where $\mathcal{M}_{\text{triv}}$ is the trivial memory skeleton defined earlier).
 544 Let $\mathcal{M} = \mathcal{M}^P \otimes \mathcal{M}^C \otimes \mathcal{M}_{\text{triv}}$ be the product of all the considered skeletons. Although \mathcal{M}
 545 formally has six states, its only reachable part is isometric to skeleton \mathcal{M}^C , with $m_1 \leftrightarrow m_{\text{init}}^C$
 546 as initial state, $m_2 \leftrightarrow m_2^C$, and $m_3 \leftrightarrow m_3^C$: we thus do not depict it to save space.

547 By Lemma 6, we have that both \sqsubseteq and \sqsubseteq^{-1} are \mathcal{M} -monotone and \mathcal{M} -selective. Using
 548 Theorem 9, we obtain that *both players have UFM strategies based on skeleton \mathcal{M} in all*
 549 *games $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$. Note that memory skeleton \mathcal{M} is minimal (no memory skeleton with two*
 550 *states or less suffices for \mathcal{P}_1 to play optimally in all arenas [21]).*

551 We provide an example of a one-player arena $\mathcal{A} = (S_1, S_2 = \emptyset, E)$ in Figure 2, and
 552 show that there is a UFM strategy for the preference relation \sqsubseteq based on skeleton \mathcal{M} .
 553 To do so, we invoke Lemma 1: we show equivalently that the product $\mathcal{A} \times \mathcal{M}$ admits an
 554 $(S \times \{m_{\text{init}}\})$ -optimal memoryless strategy for \sqsubseteq . Notice that no memoryless strategy suffices
 555 to play optimally in $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$, as when starting in s_2 , \mathcal{P}_1 should first visit s_1 before going
 556 to s_3 . Also, the $(S \times \{m_{\text{init}}\})$ -optimal memoryless strategy for the product arena is only
 557 optimal if the initial state is in $S \times \{m_{\text{init}}\}$; it is for instance not optimal from state (s_2, m_2) .

558 **6** Technical sketch

559 Due to space constraints, we only sketch our proof schemes here: full proofs are in [6].

560 **From finite memory based on \mathcal{M} to \mathcal{M} -monotony and \mathcal{M} -selectivity.** For the left-to-
 561 right implication of Theorem 9, it suffices to consider the weaker assumption involving only
 562 one-player: we establish that if UFM strategies based on \mathcal{M} exist in all *one-player games* of
 563 \mathcal{P}_1 (resp. \mathcal{P}_2), then his preference relation \sqsubseteq (resp. \sqsubseteq^{-1}) is \mathcal{M} -monotone and \mathcal{M} -selective. To
 564 maintain a compositional approach, we consider \mathcal{M} -monotony and \mathcal{M} -selectivity separately.
 565 Details are in [6, Section 4].

566 Let us sketch the proof for \mathcal{M} -monotony and \mathcal{P}_1 . We need to establish Equation (2).
 567 We instantiate the four languages involved in it: $\{w\}$, $\{w'\}$, K_1 and K_2 . We take NFA
 568 recognizing them and build an NFA \mathcal{N} that joins them in such a way that, when \mathcal{N} is
 569 considered as a one-player game arena, its plays correspond exactly to the languages of
 570 infinite words considered in Equation (2). This arena is composed of two chains emulating
 571 the two prefixes w and w' and leading to a state t where \mathcal{P}_1 has to pick a side corresponding
 572 to the two languages $[K_1]$ and $[K_2]$. Now, proving the \mathcal{M} -monotony of \sqsubseteq boils down to
 573 invoking an optimal strategy σ in the corresponding game, the crux being that σ always
 574 picks the same edge in t (i.e., the same side between subarenas corresponding to $[K_1]$ and
 575 $[K_2]$) as both prefixes w and w' are deemed equivalent by the memory skeleton \mathcal{M} .

576 The proof for \mathcal{M} -selectivity is similar. The main difference is that the “joining” state t
 577 can be visited repeatedly in this case — possibly infinitely often. This is because Equation (3)
 578 is about cycles and their languages. Our proof takes that into account.

579 **From \mathcal{M} -monotony and \mathcal{M} -selectivity to finite memory based on \mathcal{M} .** The right-to-left
 580 implication of Theorem 9 is more complex to establish. In this case, we want the result for
 581 *two-player games*, provided both preference relations are \mathcal{M} -monotone and \mathcal{M} -selective. The
 582 general scheme we use is an induction on the number of choices in arenas. The main issue
 583 is dealing with the memory: one additional choice in an arena results in many ones in the
 584 corresponding product arena. To circumvent this obstacle, we proceed in two steps. Details
 585 are in [6, Section 5].

586 We first establish the existence of *UML strategies in (prefix- and cyclic-)covered arenas*.
 587 Let us focus on the induction step we use to prove this result, as an example of the techniques
 588 involved. For an arena $\mathcal{A} = (S_1, S_2, E)$, we write $n_{\mathcal{A}} = |E| - |S|$ for its number of choices.
 589 To simplify, let us say we have a skeleton \mathcal{M} such that \sqsubseteq is \mathcal{M} -monotone and \mathcal{M} -selective,
 590 and that we assume that memoryless — *for the two players* — NE exist from all covered
 591 states in arenas with less than n choices. Then we establish that we can also build an NE in
 592 arenas with n choices, in which \mathcal{P}_1 uses a memoryless strategy — but maybe not \mathcal{P}_2 ! To
 593 prove this, we proceed as follows.

594 Let \mathcal{A} be an arena with $n_{\mathcal{A}} = n$ choices. We identify a state t in which \mathcal{P}_1 has at least two
 595 outgoing edges. By splitting the edges in t in two sets, we obtain two corresponding subarenas
 596 \mathcal{A}_a and \mathcal{A}_b such that $n_{\mathcal{A}_a}, n_{\mathcal{A}_b} < n$, along with the corresponding subgames. The induction
 597 hypothesis gives us two memoryless NE (from covered states) in these subgames: (σ_1^a, σ_2^a)
 598 and (σ_1^b, σ_2^b) . The arguments can then be unfolded as follows. First, using \mathcal{M} -monotony and
 599 \mathcal{M} being a prefix-cover, we identify one subarena (say \mathcal{A}_a) which is clearly at least as good
 600 as the other for \mathcal{P}_1 . Second, we build a strategy profile $(\sigma_1^\#, \sigma_2^\#)$, that we claim to be an NE
 601 in $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$, in the following way: \mathcal{P}_1 uses strategy σ_1^a (the one from the best subarena)
 602 and \mathcal{P}_2 reacts to \mathcal{P}_1 's actions by playing the corresponding best-response strategy. I.e., if

603 \mathcal{P}_1 plays in \mathcal{A}_a , \mathcal{P}_2 plays according to σ_2^a , and otherwise he plays according to σ_2^b . Third, it
 604 remains to prove the two inequalities of Equation (1). The rightmost one is easy, as well as
 605 the leftmost one in the subcase where the unique play $\pi \in \text{Plays}(\mathcal{A}, s, \sigma_1^\#, \sigma_2^\#)$ does not visit
 606 state t : they can both be proved thanks to the induction hypothesis and easy construction
 607 arguments. The crux of the proof is thus in the last step: proving that the leftmost inequality
 608 holds when the play visits t . This can be achieved thanks to \mathcal{M} -selectivity and \mathcal{M} being a
 609 cyclic-cover, properties of the union operator in languages of prefixes, inherent properties of
 610 the preference relation, \mathcal{A}_a being the best subarena thanks \mathcal{M} -monotony, and the induction
 611 hypothesis, in that order.

612 The actual induction step and its proof are more subtle — for example, we use different
 613 skeletons for monotony and selectivity and obtain the result in a compositional way; but
 614 the main intuition is carried here. The same result can be established symmetrically for \mathcal{P}_2 ,
 615 but again the resulting NE is only memoryless for \mathcal{P}_2 . Yet, assuming both \sqsubseteq and \sqsubseteq^{-1} are
 616 \mathcal{M} -monotone and \mathcal{M} -selective, we have two half-memoryless NE that we can mix to obtain a
 617 truly memoryless NE via Lemma 2; thus proving the existence of UML strategies in covered
 618 arenas. Observe this interesting by-product of our approach: we can actually detect *arenas*
 619 *where memory is not needed at all* thanks to our concepts of prefix- and cyclic-covers.

620 The final result — the existence of UFM strategies based on \mathcal{M} in all arenas — can then
 621 be obtained as a corollary, based on the link between memoryless strategies in product arenas
 622 and memoryfull ones in original arenas (Lemma 1). Another nice by-product of our approach,
 623 witnessed in Corollary 10, is that *the product of individual memories from one-player games*
 624 *is sufficient to play optimally in two-player games*, for both players. This is in stark contrast
 625 to the counter-example discussed in Section 1 and it illustrates that our characterization
 626 matches well-behaved preference relations.

627 **Equivalence and lifting corollary.** The equivalence (Theorem 9) is easily obtained by putting
 628 together its two directions. Note that we also establish a similar equivalence in the one-player
 629 case as a by-product.

630 ► **Theorem 11 (One-player equivalence).** *Let \sqsubseteq be a preference relation and let \mathcal{M} be a*
 631 *memory skeleton. Then, \mathcal{P}_1 has UFM strategies based on memory skeleton \mathcal{M} in all his*
 632 *one-player games $\mathcal{G} = (\mathcal{A}, \sqsubseteq)$ if and only if \sqsubseteq is \mathcal{M} -monotone and \mathcal{M} -selective.*

633 Although this looks like a weak version of Theorem 9 at first sight, this is actually a
 634 distinct result as both sides of the equivalence are weaker: on the left side, it only handles
 635 the memory requirements for \mathcal{P}_1 's one-player games; on the right side, it does not assume
 636 anything about the inverse preference relation \sqsubseteq^{-1} .

637 The lifting corollary, Corollary 10, is also a consequence of our approach. As we have
 638 seen, the existence of UFM strategies based on a skeleton \mathcal{M} in *one-player* games suffices to
 639 yield \mathcal{M} -monotony and \mathcal{M} -selectivity of the corresponding preference relation. Hence if both
 640 players have UFM strategies in their one-player games, both relations satisfy these properties
 641 and we can take the other direction of Theorem 9 to ensure that UFM strategies also exist in
 642 *two-player* games. As explained above, this approach can actually be used compositionally.

643 All proofs, as well as the one-player equivalence, are presented in details in [6].

644 — References

- 645 1 Benjamin Aminof and Sasha Rubin. First-cycle games. *Inf. Comput.*, 254:195–216, 2017. URL:
646 <https://doi.org/10.1016/j.ic.2016.10.008>, doi:10.1016/j.ic.2016.10.008.
- 647 2 Roderick Bloem, Krishnendu Chatterjee, and Barbara Jobstmann. Graph games and reactive
648 synthesis. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem,
649 editors, *Handbook of Model Checking*, pages 921–962. Springer, 2018. URL: [https://doi.org/](https://doi.org/10.1007/978-3-319-10575-8)
650 [10.1007/978-3-319-10575-8](https://doi.org/10.1007/978-3-319-10575-8), doi:10.1007/978-3-319-10575-8_27.
- 651 3 Patricia Bouyer, Ulrich Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Jiri Srba.
652 Infinite runs in weighted timed automata with energy constraints. In Franck Cassez
653 and Claude Jard, editors, *Formal Modeling and Analysis of Timed Systems, 6th International*
654 *Conference, FORMATS 2008, Saint Malo, France, September 15-17, 2008. Proceedings*, volume
655 5215 of *Lecture Notes in Computer Science*, pages 33–47. Springer, 2008. URL: [https://doi.org/](https://doi.org/10.1007/978-3-540-85778-5_4)
656 [10.1007/978-3-540-85778-5_4](https://doi.org/10.1007/978-3-540-85778-5_4), doi:10.1007/978-3-540-85778-5_4.
- 657 4 Patricia Bouyer, Piotr Hofman, Nicolas Markey, Mickael Randour, and Martin Zimmer-
658 mann. Bounding average-energy games. In Javier Esparza and Andrzej S. Murawski, ed-
659 itors, *Foundations of Software Science and Computation Structures - 20th International*
660 *Conference, FOSSACS 2017, Held as Part of the European Joint Conferences on Theory*
661 *and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings*,
662 volume 10203 of *Lecture Notes in Computer Science*, pages 179–195, 2017. URL: [https://doi.org/](https://doi.org/10.1007/978-3-662-54458-7_11)
663 [10.1007/978-3-662-54458-7_11](https://doi.org/10.1007/978-3-662-54458-7_11), doi:10.1007/978-3-662-54458-7_11.
- 664 5 Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim G. Larsen, and Simon Laursen.
665 Average-energy games. *Acta Inf.*, 55(2):91–127, 2018. URL: [https://doi.org/10.1007/](https://doi.org/10.1007/s00236-016-0274-1)
666 [s00236-016-0274-1](https://doi.org/10.1007/s00236-016-0274-1), doi:10.1007/s00236-016-0274-1.
- 667 6 Patricia Bouyer, Stéphane Le Roux, Youssef Oualhadj, Mickael Randour, and Pierre Van-
668 denhov. Games where you can play optimally with arena-independent finite memory. *CoRR*,
669 abs/2001.03894, 2020. URL: <https://arxiv.org/abs/2001.03894>, arXiv:2001.03894.
- 670 7 Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour,
671 Jean-François Raskin, Ocan Sankur, and Mathieu Sassolas. Non-zero sum games for reactive
672 synthesis. In Adrian-Horia Dediu, Jan Janousek, Carlos Martín-Vide, and Bianca Truthe,
673 editors, *Language and Automata Theory and Applications - 10th International Conference,*
674 *LATA 2016, Prague, Czech Republic, March 14-18, 2016, Proceedings*, volume 9618 of *Lecture*
675 *Notes in Computer Science*, pages 3–23. Springer, 2016. URL: [https://doi.org/10.1007/](https://doi.org/10.1007/978-3-319-30000-9)
676 [978-3-319-30000-9](https://doi.org/10.1007/978-3-319-30000-9), doi:10.1007/978-3-319-30000-9_1.
- 677 8 Thomas Brihaye, Florent Delgrange, Youssef Oualhadj, and Mickael Randour. Life is
678 random, time is not: Markov decision processes with window objectives. In Fokkink and
679 van Glabbeek [23], pages 8:1–8:18. URL: <https://doi.org/10.4230/LIPIcs.CONCUR.2019.8>,
680 doi:10.4230/LIPIcs.CONCUR.2019.8.
- 681 9 Véronique Bruyère, Emmanuel Filiot, Mickael Randour, and Jean-François Raskin. Meet
682 your expectations with guarantees: Beyond worst-case synthesis in quantitative games. *Inf.*
683 *Comput.*, 254:259–295, 2017. URL: <https://doi.org/10.1016/j.ic.2016.10.011>, doi:10.
684 [1016/j.ic.2016.10.011](https://doi.org/10.1016/j.ic.2016.10.011).
- 685 10 Véronique Bruyère, Quentin Hautem, and Mickael Randour. Window parity games: an
686 alternative approach toward parity games with time bounds. In Domenico Cantone and Giorgio
687 Delzanno, editors, *Proceedings of the Seventh International Symposium on Games, Automata,*
688 *Logics and Formal Verification, GandALF 2016, Catania, Italy, 14-16 September 2016*, volume
689 226 of *EPTCS*, pages 135–148, 2016. URL: <https://doi.org/10.4204/EPTCS.226.10>, doi:
690 [10.4204/EPTCS.226.10](https://doi.org/10.4204/EPTCS.226.10).
- 691 11 Véronique Bruyère, Quentin Hautem, Mickael Randour, and Jean-François Raskin. Energy
692 mean-payoff games. In Fokkink and van Glabbeek [23], pages 21:1–21:17. URL: [https://doi.org/](https://doi.org/10.4230/LIPIcs.CONCUR.2019.21)
693 [10.4230/LIPIcs.CONCUR.2019.21](https://doi.org/10.4230/LIPIcs.CONCUR.2019.21), doi:10.4230/LIPIcs.CONCUR.2019.21.

- 694 12 Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga. Resource
695 interfaces. In Rajeev Alur and Insup Lee, editors, *EMSOFT*, volume 2855 of *Lecture Notes in*
696 *Computer Science*, pages 117–133. Springer, 2003.
- 697 13 Krishnendu Chatterjee and Laurent Doyen. Energy parity games. *Theor. Comput. Sci.*,
698 458:49–60, 2012. doi:10.1016/j.tcs.2012.07.038.
- 699 14 Krishnendu Chatterjee, Laurent Doyen, Mickael Randour, and Jean-François Raskin. Looking
700 at mean-payoff and total-payoff through windows. *Inf. Comput.*, 242:25–52, 2015. URL:
701 <https://doi.org/10.1016/j.ic.2015.03.010>, doi:10.1016/j.ic.2015.03.010.
- 702 15 Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski. Mean-payoff parity
703 games. In *20th IEEE Symposium on Logic in Computer Science (LICS 2005)*, 26–29 June
704 2005, Chicago, IL, USA, *Proceedings*, pages 178–187. IEEE Computer Society, 2005. URL:
705 <https://doi.org/10.1109/LICS.2005.26>, doi:10.1109/LICS.2005.26.
- 706 16 Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Generalized parity
707 games. In Helmut Seidl, editor, *Foundations of Software Science and Computational*
708 *Structures, 10th International Conference, FOSSACS 2007, Held as Part of the Joint Eu-*
709 *ropean Conferences on Theory and Practice of Software, ETAPS 2007, Braga, Portugal,*
710 *March 24–April 1, 2007, Proceedings*, volume 4423 of *Lecture Notes in Computer Science*,
711 pages 153–167. Springer, 2007. URL: https://doi.org/10.1007/978-3-540-71389-0_12,
712 doi:10.1007/978-3-540-71389-0_12.
- 713 17 Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin. Strategy synthesis for
714 multi-dimensional quantitative objectives. *Acta Inf.*, 51(3-4):129–163, 2014. URL: <https://doi.org/10.1007/s00236-013-0182-6>, doi:10.1007/s00236-013-0182-6.
- 715 18 Florent Delgrange, Joost-Pieter Katoen, Tim Quatmann, and Mickael Randour. Simple
716 strategies in multi-objective MDPs. In Armin Biere and David Parker, editors, *Tools and*
717 *Algorithms for the Construction and Analysis of Systems - 26th International Conference,*
718 *TACAS 2020, Held as Part of the European Joint Conferences on Theory and Practice of*
719 *Software, ETAPS 2020, Dublin, Ireland, April 25–30, 2020, Proceedings, Part I*, volume
720 12078 of *Lecture Notes in Computer Science*, pages 346–364. Springer, 2020. URL: https://doi.org/10.1007/978-3-030-45190-5_19, doi:10.1007/978-3-030-45190-5_19.
- 721 19 Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *Int.*
722 *Journal of Game Theory*, 8(2):109–113, 1979.
- 723 20 E. Allen Emerson and Charanjit S. Jutla. The complexity of tree automata and logics of
724 programs. In *FOCS*, pages 328–337. IEEE Computer Society, 1988.
- 725 21 Nathanaël Fijalkow and Florian Horn. The surprising complexity of reachability games. *CoRR*,
726 abs/1010.2420, 2010. URL: <http://arxiv.org/abs/1010.2420>, arXiv:1010.2420.
- 727 22 Nathanaël Fijalkow, Florian Horn, Denis Kuperberg, and Michal Skrzypczak. Trading bounds
728 for memory in games with counters. In Halldórsson et al. [30], pages 197–208. URL: https://doi.org/10.1007/978-3-662-47666-6_16, doi:10.1007/978-3-662-47666-6_16.
- 729 23 Wan Fokkink and Rob van Glabbeek, editors. *30th International Conference on Concurrency*
730 *Theory, CONCUR 2019, August 27–30, 2019, Amsterdam, the Netherlands*, volume 140 of
731 *LIPICs*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. URL: [http://www.dagstuhl.](http://www.dagstuhl.de/dagpub/978-3-95977-121-4)
732 [de/dagpub/978-3-95977-121-4](http://www.dagstuhl.de/dagpub/978-3-95977-121-4).
- 733 24 Hugo Gimbert. Pure stationary optimal strategies in Markov decision processes. In Wolfgang
734 Thomas and Pascal Weil, editors, *STACS 2007, 24th Annual Symposium on Theoretical*
735 *Aspects of Computer Science, Aachen, Germany, February 22–24, 2007, Proceedings*, volume
736 4393 of *Lecture Notes in Computer Science*, pages 200–211. Springer, 2007. URL: https://doi.org/10.1007/978-3-540-70918-3_18, doi:10.1007/978-3-540-70918-3_18.
- 737 25 Hugo Gimbert and Edon Kelmendi. Two-player perfect-information shift-invariant submixing
738 stochastic games are half-positional. Unpublished, 2014.
- 739 26 Hugo Gimbert and Wiesław Zielonka. When can you play positionally? In Jirí Fi-
740 ala, Václav Koubek, and Jan Kratochvíl, editors, *Mathematical Foundations of Com-*
741 *puter Science 2004, 29th International Symposium, MFCS 2004, Prague, Czech Repub-*
742 *lic*

- 746 *lic*, August 22-27, 2004, *Proceedings*, volume 3153 of *Lecture Notes in Computer Science*,
 747 pages 686–697. Springer, 2004. URL: https://doi.org/10.1007/978-3-540-28629-5_53,
 748 [doi:10.1007/978-3-540-28629-5_53](https://doi.org/10.1007/978-3-540-28629-5_53).
- 749 **27** Hugo Gimbert and Wieslaw Zielonka. Games where you can play optimally without any
 750 memory. In Martín Abadi and Luca de Alfaro, editors, *CONCUR 2005 - Concurrency Theory*,
 751 *16th International Conference, CONCUR 2005, San Francisco, CA, USA, August 23-26, 2005*,
 752 *Proceedings*, volume 3653 of *Lecture Notes in Computer Science*, pages 428–442. Springer,
 753 2005. URL: https://doi.org/10.1007/11539452_33, [doi:10.1007/11539452_33](https://doi.org/10.1007/11539452_33).
- 754 **28** Hugo Gimbert and Wieslaw Zielonka. Pure and Stationary Optimal Strategies in Perfect-
 755 Information Stochastic Games with Global Preferences. Unpublished, December 2009. URL:
 756 <https://hal.archives-ouvertes.fr/hal-00438359>.
- 757 **29** Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite*
 758 *Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume
 759 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
- 760 **30** Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, edi-
 761 tors. *Automata, Languages, and Programming - 42nd International Colloquium, ICALP*
 762 *2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, volume 9135 of *Lecture Notes*
 763 *in Computer Science*. Springer, 2015. URL: <https://doi.org/10.1007/978-3-662-47666-6>,
 764 [doi:10.1007/978-3-662-47666-6](https://doi.org/10.1007/978-3-662-47666-6).
- 765 **31** Marcin Jurdzinski. Deciding the winner in parity games is in $UP \cap co-UP$. *Inf. Process.*
 766 *Lett.*, 68(3):119–124, 1998. URL: [https://doi.org/10.1016/S0020-0190\(98\)00150-1](https://doi.org/10.1016/S0020-0190(98)00150-1), [doi:10.1016/S0020-0190\(98\)00150-1](https://doi.org/10.1016/S0020-0190(98)00150-1).
 767
- 768 **32** Marcin Jurdzinski, Ranko Lazic, and Sylvain Schmitz. Fixed-dimensional energy games
 769 are in pseudo-polynomial time. In Halldórsson et al. [30], pages 260–272. URL: <https://doi.org/10.1007/978-3-662-47666-6>, [doi:10.1007/978-3-662-47666-6_21](https://doi.org/10.1007/978-3-662-47666-6_21).
 770
- 771 **33** Eryk Kopczyński. Half-positional determinacy of infinite games. In Michele Bugliesi, Bart Pre-
 772 neel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming*,
 773 *33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings*,
 774 *Part II*, volume 4052 of *Lecture Notes in Computer Science*, pages 336–347. Springer, 2006.
 775 URL: https://doi.org/10.1007/11787006_29, [doi:10.1007/11787006_29](https://doi.org/10.1007/11787006_29).
- 776 **34** Eryk Kopczyński. *Half-positional Determinacy of Infinite Games*. PhD thesis, Warsaw
 777 University, 2008.
- 778 **35** Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.
- 779 **36** Martin J. Osborne and Ariel Rubinstein. *A course in game theory*. The MIT Press, Cambridge,
 780 USA, 1994.
- 781 **37** Mickael Randour. Automated synthesis of reliable and efficient systems through game theory:
 782 A case study. In *Proc. of ECCS 2012*, Springer Proceedings in Complexity XVII, pages
 783 731–738. Springer, 2013. [doi:10.1007/978-3-319-00395-5_90](https://doi.org/10.1007/978-3-319-00395-5_90).
- 784 **38** Stéphane Le Roux. Infinite sequential Nash equilibrium. *Logical Methods in Computer Science*,
 785 9(2), 2013. URL: [https://doi.org/10.2168/LMCS-9\(2:3\)2013](https://doi.org/10.2168/LMCS-9(2:3)2013), [doi:10.2168/LMCS-9\(2:3\)](https://doi.org/10.2168/LMCS-9(2:3)2013)
 786 [2013](https://doi.org/10.2168/LMCS-9(2:3)2013).
- 787 **39** Stéphane Le Roux. Concurrent games and semi-random determinacy. In Igor Potapov,
 788 Paul G. Spirakis, and James Worrell, editors, *43rd International Symposium on Mathematical*
 789 *Foundations of Computer Science, MFCS 2018, August 27-31, 2018, Liverpool, UK*, volume 117
 790 of *LIPICs*, pages 40:1–40:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018. URL:
 791 <https://doi.org/10.4230/LIPICs.MFCS.2018.40>, [doi:10.4230/LIPICs.MFCS.2018.40](https://doi.org/10.4230/LIPICs.MFCS.2018.40).
- 792 **40** Stéphane Le Roux and Arno Pauly. Extending finite-memory determinacy to multi-player
 793 games. *Inf. Comput.*, 261(Part):676–694, 2018. URL: <https://doi.org/10.1016/j.ic.2018.02.024>, [doi:10.1016/j.ic.2018.02.024](https://doi.org/10.1016/j.ic.2018.02.024).
 794
- 795 **41** Stéphane Le Roux, Arno Pauly, and Mickael Randour. Extending finite-memory determinacy by
 796 Boolean combination of winning conditions. In Sumit Ganguly and Paritosh K. Pandya, editors,

- 797 *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Com-*
798 *puter Science, FSTTCS 2018, December 11-13, 2018, Ahmedabad, India*, volume 122 of *LIPIcs*,
799 pages 38:1–38:20. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2018. URL: <https://doi.org/10.4230/LIPIcs.FSTTCS.2018.38>, doi:10.4230/LIPIcs.FSTTCS.2018.38.
800
- 801 42 Yaron Velner, Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, Alexander Moshe
802 Rabinovich, and Jean-François Raskin. The complexity of multi-mean-payoff and multi-energy
803 games. *Inf. Comput.*, 241:177–196, 2015. URL: <https://doi.org/10.1016/j.ic.2015.03.001>,
804 doi:10.1016/j.ic.2015.03.001.
- 805 43 Wieslaw Zielonka. Infinite games on finitely coloured graphs with applications to automata
806 on infinite trees. *Theor. Comput. Sci.*, 200(1-2):135–183, 1998.

807 **A Discussion**

808 We close our paper with a discussion of the assets and limits of our approach, its applicability
 809 with regard to the current research landscape, and the directions we aim to follow in future
 810 work.

811 **Technical overview.** Naturally, our technical approach is inspired by the one of Gimbert
 812 and Zielonka for the memoryless case [27], which can actually be rediscovered through our
 813 results using a trivial memory skeleton. Two of the most important challenges we had to
 814 overcome were:

- 815 1. establishing natural concepts of *monotony and selectivity modulo memory* that are exactly
 816 as powerful as required to maintain a complete characterization (i.e., sufficient and
 817 necessary conditions) in the finite-memory case;
- 818 2. circumventing the seemingly unavoidable *coupling between the memory skeleton and the*
 819 *arena* in the inductive argument needed to prove the implication from \mathcal{M} -monotony and
 820 \mathcal{M} -selectivity to finite-memory optimal strategies — which we were able to do using our
 821 notions of prefix-covers and cyclic-covers.

822 All along our paper, we highlighted the similarities and discrepancies between our work and
 823 Gimbert and Zielonka’s [27]. As observed through [6, Remark 16], our results are established
 824 using fine-grained assumptions and conclusions, in an effort to push the approach to its
 825 limits. They also preserve *compositionality*, splitting the reasoning for \mathcal{M} -monotony and
 826 \mathcal{M} -selectivity, and for the two players.

827 Alongside \mathcal{M} -monotony and \mathcal{M} -selectivity, we define two other key concepts to solve the
 828 technical issues related to the induction on product arenas: *prefix-covers* and *cyclic-covers*.
 829 These notions are crucial tools to prove the right-to-left implication of Theorem 9.

830 **Some advantages.** The aforementioned concepts of prefix-covers and cyclic-covers also have
 831 benefits from a practical point of view: given a preference relation \sqsubseteq and the corresponding
 832 memory skeleton \mathcal{M} , they let us *identify game arenas where memoryless strategies suffice*
 833 whereas finite memory (based on \mathcal{M}) might be necessary in general. Such arenas are the
 834 ones covered by \mathcal{M} . Hence in practice, this approach permits to obtain UML strategies for
 835 many arenas where a coarser approach would only provide UFM ones.

836 Our approach yields *two methods* to establish that a preference relation (or equivalently
 837 a payoff function or a winning condition) admits UFM strategies. The first one, exhibiting
 838 appropriate memory skeletons and proving \mathcal{M} -monotony and \mathcal{M} -selectivity, is based on
 839 Theorem 9 and can be used *compositionally* through [6, Corollary 25]. The second one follows
 840 the *lifting corollary*, Corollary 10: one only has to study the one-player subcases then invoke
 841 this result to lift the existence of UFM strategies to the two-player case, without checking
 842 for \mathcal{M} -monotony and \mathcal{M} -selectivity at all. Hence this second method is often painless in
 843 practice.

844 Two interesting facts can be seen through Corollary 10. First, there is *no blow-up in the*
 845 *memory* required when going from one-player games to two-player games: the overall memory
 846 simply combines the memory skeletons of the two players. Second, assuming that one has an
 847 algorithm to solve¹ one-player games — say for \mathcal{P}_1 — for a winning condition satisfying our
 848 hypotheses, this lifting corollary also induces a *naive algorithm for the two-player case for*
 849 *free*: thanks to the bounds on memory, one may enumerate the strategies of the adversary, \mathcal{P}_2

¹ I.e., decide who has a winning strategy from a given state.

850 — or guess one if one aims for a non-deterministic algorithm — and solve the corresponding
 851 \mathcal{P}_1 's game(s) where the strategy of \mathcal{P}_2 is fixed. Note that while such a simple algorithm
 852 might not be optimal, it does correspond to the approach giving the best complexity class
 853 known for the renowned family of games in $\text{NP} \cap \text{coNP}$, such as, e.g., parity or mean-payoff
 854 games (e.g., [31]). These last two cases could already be dealt with thanks to Gimbert and
 855 Zielonka's result since they involve memoryless strategies, but now a similar road can be
 856 taken for any objective that admits arena-independent finite-memory optimal strategies, such
 857 as, e.g., generalized parity games.

858 **Applicability.** Let us give a quick tour of some classical (combinations of) objectives —
 859 expressed through winning conditions, payoffs or preference relations — and assess whether
 860 our approach permits to establish the existence of UFM strategies in the corresponding
 861 games.

862 Note that when considering multiple (quantitative) objectives, optimal strategies usually
 863 do not exist, and one has to settle for *Pareto-optimal* ones (e.g., [18]). However, in many
 864 cases, the (decision) problem under study is as follows: given a threshold (vector), define
 865 the winning condition as all the plays achieving at least this threshold, and check for a
 866 winning strategy. Hence multi-objective quantitative games are often de facto reduced
 867 to qualitative win-lose games for this so-called *threshold problem*. Observe that, given a
 868 multi-objective setting, if UFM strategies exist for all threshold problems, then finite-memory
 869 strategies suffice to realize the Pareto front (as each point of this front can be considered
 870 as a threshold). Therefore, *our approach also enables reasoning about the existence of*
 871 *finite-memory Pareto-optimal strategies in multi-objective games.*

872 We start our overview with some game settings that fall under the scope of our approach.
 873 Obviously, *all memoryless-determined objectives* are among them, since we generalize Gimbert
 874 and Zielonka's work [27]: this includes, e.g., mean-payoff [19], parity [20, 43], energy [12]
 875 or average-energy games [5]. As established in Section 1, our results encompass all cases
 876 where *arena-independent* memory suffices. Hence they permit to rediscover the existence of
 877 UFM strategies for games such as, e.g., generalized reachability [21], generalized parity [16],
 878 window parity games [10], some variants of window mean-payoff games [14], or lower- and
 879 upper-bounded (multi-dimension) energy games [3, 5, 4]. Our approach can also be useful to
 880 extend these known results to more general combinations, either via appropriate memory
 881 skeletons or through the lifting corollary (see an application in Section 5).

882 There are many games that do not fit our approach for *good reasons*, as they do not admit
 883 UFM strategies in general: e.g., multi-dimension mean-payoff [42], mean-payoff parity [15],
 884 or energy mean-payoff games [11]. More interesting are games for which finite-memory
 885 strategies exist, but the memory is *arena-dependent*. These notably include games with
 886 multi-dimension lower-bounded energy objectives and no upper bound [17, 32], or other
 887 variants of window mean-payoff games [14]. In such games, the players usually have to keep
 888 track of information such as, e.g., the sum of weights along an acyclic path, which is bounded
 889 for any given arena, but by a value that grows when the arena grows. Hence the need for
 890 memory that grows with the arena parameters. Our results cannot be applied directly to
 891 such cases in order to obtain the existence of finite-memory strategies for all games. An
 892 adaptation of our approach could potentially be used for subclasses of arenas where the
 893 parameters are bounded (in order to regain a skeleton working on all arenas of the class).

894 **Critical analysis.** Let us take a step back and assess the place of our work in its larger line
 895 of research. The natural endgame is characterizing all preference relations admitting finite-

memory optimal strategies, including those using *arena-dependent* memory, and pinpointing the frontiers of application of the lifting corollary — that is, under which conditions is finite-memory determinacy preserved when going from one-player to two-player games?

The road is long from Gimbert and Zielonka’s characterization in the memoryless case [27] to such a general result, and this work is but a first step. We have already established that Gimbert and Zielonka’s approach cannot be fully transposed for finite memory. Our focus on *arena-independent* memory is a way to study the frontiers of this approach while providing an extension of practical interest. While it may seem limited at first, note that our framework already encompasses arguably rich classes of games such as, e.g., generalized parity games and fully-bounded energy games. As argued in Section 1, recall that our result is in no way a simple application of [27] to product arenas.

From a practical point of view, our equivalence result has limitations as it inherently uses the memory skeleton \mathcal{M} . At this point, our approach neither helps in finding an appropriate skeleton, nor in determining the minimal one; two highly interesting questions from a practical standpoint. Nonetheless, to advance toward answering these questions and to be able to find good skeletons automatically, one first has to understand their theoretical characteristics, which we do here as a necessary stepping stone. Focusing on applications, let us note that the equivalence result is often not the most suited tool: this is instead where the lifting corollary shines. As noted before, reasoning on one-player games (i.e., graphs) is generally much easier than in two-player games (e.g., [5, 4, 11]). Hence, a reasonably easy way to tackle practical cases is to find skeletons sufficient for \mathcal{P}_1 and \mathcal{P}_2 in their respective one-player games and to use our constructive result to build a skeleton that suffices for both in two-player games.

Comparison with related work. We already discussed the most important related papers [26, 27, 33, 1, 35, 41] in Section 1. Let us highlight here some works where similar approaches have been considered to establish “meta-theorems” applying to general classes of games, or works that inspire interesting directions of research. First and foremost is the determinacy theorem by Martin that guarantees determinacy (without considering the complexity of strategies) for Borel winning conditions [35].

Aminof and Rubin provide a simpler (but incomplete) approach to memoryless determinacy through the prism of first-cycle games in [1]: a similar take on finite-memory determinacy could be appealing — it could provide sufficient conditions easier to test than \mathcal{M} -monotony and \mathcal{M} -selectivity.

Let us discuss the work of Kopczyński. First, in [33], he establishes sufficient (and relaxed) conditions to ensure the existence of UML strategies *for one player, in two-player games*: it would be interesting to study the corresponding problem in the finite-memory case. Indeed, in many games where infinite memory is needed, it is only the case for one of the players (e.g., [42, 15, 11]) and conditions à la Kopczyński could thus prove useful. Note that this is different from Theorem 11, which gives a sufficient *and necessary* condition but for one-player games only. Second, we recently discovered unpublished content in Kopczyński’s PhD thesis [34]. Kopczyński distinguishes *chromatic* memory (which corresponds to our definition of memory skeleton), and the more powerful *chaotic* memory, where transitions can depend on the actual edges of the arenas, rather than simply on their colors. Chaotic memory is thus intrinsically *arena-dependent*. Our notion of an arena being both prefix- and cyclic-covered by a memory skeleton \mathcal{M} is equivalent to a notion in [34, Definition 8.12], which defines that an arena *adheres to chromatic memory* \mathcal{M} if it is possible to assign a state of \mathcal{M} to every state of the arena such that moving along the edges updates these memory

943 states in a consistent way. Our definitions of *prefix-* and *cyclic-cover* can be seen as two
944 distinct sides of this idea of *adherence*, which when added up, are actually equivalent to it.

945 Following the same motivation as our work — the need to characterize (combinations of)
946 objectives admitting finite-memory optimal strategies, Le Roux et al. [41] take another road:
947 whereas our work permits to lift results from one-player games to two-player games, they
948 provide a lifting from the single-objective case to the multi-objective one. Their techniques, as
949 well as the scope of their results, are somewhat orthogonal to ours. Whether both approaches
950 can be intertwined to obtain results on more general settings remains an open question.

951 Our work focuses on *deterministic turn-based* two-player games. Sufficient conditions
952 have been published for stochastic models but to the best of our knowledge, no complete
953 characterization, even for the simplest case of Markov decision processes (e.g., [24]). Two
954 unpublished articles contain interesting results on stochastic games [28, 25], including an
955 extension of Gimbert and Zielonka’s original work, by the same authors [28]. Whether part
956 of our approach can be useful to tackle the finite-memory case in this context, or in richer
957 contexts mixing games and stochastic models (e.g., [9]) is a question for future research.
958 Some sufficient criteria, orthogonal to our approach, were studied for *concurrent* games
959 in [39].

960 **Limits and future work.** To close this paper, we recall three limits of our approach, and
961 the corresponding open problems.

962 First, as explained throughout the paper, our results cover all cases where *arena-*
963 *independent* memory suffices, and are *limited to these cases*. We have argued that the
964 approach cannot be fully lifted to the general case, for good reasons, as the lifting corollary
965 breaks in some situations (Section 1). Still, we have hope to generalize our approach to some
966 extent to the *arena-dependent* case, through some *function* associating memory skeletons
967 to arenas, as discussed in Section 1. Obtaining a lifting corollary — under well-chosen
968 conditions — in the arena-dependent case would be of tremendous help in practice: see for
969 example [5, 4, 11]. Hence this is clearly the next step in our quest.

970 Second, our result is a characterization *instantiated by a memory skeleton* \mathcal{M} . While
971 the lifting corollary is helpful in applications, it would be fantastic to be able to find an
972 appropriate skeleton automatically, and to be able to determine if a given skeleton is minimal
973 (with regard to a preference relation). This paper is a first step toward these long-term
974 objectives.

975 Lastly, as explained in Remark 3 and [6, Remark 24], most of our arguments carry over
976 to the case of *general Nash equilibria*. That is, when considering not necessarily antagonistic
977 games where the two players use different, not necessarily inverse, preference relations.
978 Whether our approach can be adapted in this case, at the price of an unavoidable blow-up
979 of memory, is an open question worth considering. In particular, we want to study the
980 links between our results (including the lifting from one-player to two-player games) and
981 recent results lifting finite-memory determinacy in two-player games to the existence of
982 finite-memory Nash equilibria in multi-player games [40].