

Simple Nash Equilibria for Qualitative Multiplayer Games

Mona Alluwaym 

University of Liverpool, UK

James C. A. Main 

University of Oxford, UK

Sven Schewe 

University of Liverpool, UK

Abstract

We investigate memory requirements for Nash and subgame-perfect equilibria in turn-based deterministic games with ω -regular objectives. We prove that memoryless randomised (i.e., stationary) subgame-perfect equilibria always exist in games with reachability, safety, and 0-2 Muller objectives (i.e., Muller objectives for which accepting sets are either up- or downward closed), and any combination of these objectives. We provide an algorithm to construct such an equilibrium. We also show that randomisation may be required to construct memoryless equilibria in games with reachability or Büchi as well as safety or CoBüchi objectives, and that memoryless equilibria need not exist for any other class of Muller objectives (with respect to the Mostowski hierarchy).

2012 ACM Subject Classification Theory of computation $\rightarrow$ Solution concepts in game theory

Keywords and phrases games on graphs, multiplayer games, Nash equilibria, subgame-perfect equilibria, memoryless strategies

Digital Object Identifier 10.4230/LIPIcs.MFCS.2026.85

Funding This work has been supported by the EPSRC projects EP/Z536179/1, EP/X03688X/1, and EP/X042596/1.

1 Introduction

Games on graphs. We study turn-based games played on finite graphs (e.g., [24, 20]). The structure on which the game is played is called a turn-based deterministic arena (or simply an arena); it is described by a directed graph and a partition of the vertices among the players. At the beginning of a play, a token is placed on an initial vertex and, in each round, the player in control of the current vertex moves the token along an outgoing edge. The players continue this interaction for infinitely many rounds and construct a play, i.e., an infinite path of the graph. The goal of a player is formalised by an objective, i.e., a subset of plays. Players follow strategies describing how to select moves. Formally, strategies are functions assigning to any play history the decisions to be made. In full generality, strategies may use memory and randomisation. A strategy is called pure if it deterministically assigns moves to histories, memoryless (or stationary) if its decisions depend only on the current vertex and not the full history, and positional if it is both memoryless and pure.

Games on graphs are notably useful to solve the controller synthesis problem (e.g., [2]). Given a reactive system, the goal of controller synthesis is to automatically construct a controller for the system that enforces some specification no matter the behaviour of the environment of the system. It can be solved by modelling the interaction of the system and its environment as a zero-sum game on a graph: the goal of the system player is to enforce their objective (which formalises the specification) and the environment player aims to falsify this objective. Strategies of the system player correspond to controllers; thus the goal is



© Mona Alluwaym, James C. A. Main, and Sven Schewe;
licensed under Creative Commons License CC-BY 4.0

51st International Symposium on Mathematical Foundations of Computer Science (MFCS 2026).

Editors: Michal Koucký and Daniela Petrişan; Article No. 85; pp. 85:1–85:17

Leibniz International Proceedings in Informatics



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

to find a winning strategy, i.e., a strategy enforcing the objective of the system against all strategies of its adversary. In general, we are interested in *simple winning strategies* (e.g., using limited memory), as they yield simpler controllers (see [37] and references therein).

Multiplayer games. For systems consisting of several components interacting together, treating each component as competing with the others may be too restrictive (e.g., it precludes cooperation between components). Instead, we consider games in which each component is a player. Each player has their own objective and these are not necessarily incompatible. Such games are called multiplayer *non-zero-sum games*.


Nash equilibria [34] (NE) are the most classical formalisation of rational behaviour in such games. An NE is a strategy profile, i.e., a tuple of one strategy per player, such that no player can increase the probability that their objective holds by unilaterally changing their strategy. *Subgame-perfect equilibria* [39] (SPE) are another well-studied notion of rationality. Intuitively, SPEs are a refinement of NEs requiring that from all histories of the game (even those that are not reachable with the considered strategy profile), the players follow an NE. This stronger requirement excludes NEs based on non-credible threats, i.e., where a player threatens to make a move that is detrimental to the satisfaction of their objective after some history to discourage other players from moving to some vertex.

We focus on the complexity of equilibria in games with ω -regular objectives. The class of ω -regular languages is fundamental in the fields of synthesis and verification [1], and includes, notably, linear temporal logic [36]. In games with ω -regular objectives, there exist pure SPEs such that all strategies are *finite-memory* [41] – see also [10, 19] for a proof of existence of NE with little memory (for a class of objectives subsuming ω -regular objectives).

We focus on three classes of ω -regular objectives. First, *reachability objectives* (also known as reach-a-set objectives), the goal of which is to reach a target. Second, *safety objectives* (also known as stay-in-a-set objectives), the goal of which is to remain within a subset of vertices. For general ω -regular objectives, we consider *Muller objectives*. A Muller objective is described by a family of subsets of vertices; a play satisfies a Muller objective if the set of vertices occurring infinitely often is in the family defining the Muller objective. Muller objectives generalise Büchi, CoBüchi and parity objectives.

When is memory needed? We investigate the existence of memoryless equilibria in games on turn-based deterministic arenas. This question has been settled for more general arenas. We briefly survey known results regarding the existence of (memoryless) equilibria in games on concurrent (stochastic) arenas and turn-based stochastic arenas. In a concurrent arena, the players select actions simultaneously in each vertex and the next vertex is selected following a distribution depending only on the current vertex and chosen actions. A turn-based stochastic arena can be seen as a concurrent arena in which only at most one player has more than one choice of action at each vertex.

For discounted concurrent stochastic games, memoryless randomised SPE are known to exist [21, 40]. This property does not generalise to concurrent games with reachability, safety or Muller objectives. For safety objectives, while NEs exist in concurrent games [38], there need not exist a memoryless NE even on turn-based stochastic arenas [25]. For reachability objectives, NEs need not exist in the concurrent case [22], although memoryless ε -NE have been shown to exist [16] (i.e., strategy profiles in which no unilateral deviation increases the probability of the objective of a player by ε) for all $\varepsilon > 0$. In turn-based stochastic games with ω -regular objectives, pure NE are known to exist [16]. However, there need not exist (randomised) memoryless NE already in turn-based reachability games with absorbing



■ **Figure 1** A two-player reachability game. Circles represent $\mathcal{P}_1$ and hexagons represent $\mathcal{P}_2$. Vertex t_i is the target of $\mathcal{P}_i$ (b) illustrates a memoryless strategy profile that is an SPE from t_2 but not an ENE. (c) illustrates a memoryless strategy profile, which is both an ENE and an SPE.

targets [29, 43]. This negative example for reachability can also be seen as a game with a Muller objective. For games on turn-based deterministic arenas, we highlight the example in [3] of a turn-based deterministic game with no memoryless NE, where the goal of the players is to optimise some payoff attributed to some absorbing vertices and the unique cycle of the arena. Finally, for complex Muller objectives (in the sense of the Mostowski hierarchy), there need not exist memoryless randomised NEs in turn-based deterministic zero-sum games [15].

All of the negative results above require either a stochastic game, payoff functions or complex Muller objectives. It follows that they do not settle the problem of existence of memoryless equilibria in reachability and safety games, and in games with simple Muller objectives on turn-based deterministic arenas.

Our contribution. We prove the existence of *memoryless randomised SPE* in games where all players have an objective from three classes: safety objectives, reachability objectives, or simple Muller objectives, where all winning sets of vertices are up- or downwards closed (such as parity objectives with colours 0, 1, and 2). Our proof is constructive and can be used to effectively compute strategy profiles that are *everywhere NEs (ENEs)* – strategy profiles that are NEs from every vertex of the game, rather than from a fixed initial vertex. This is strictly stronger than NEs from a single initial vertex. For this class of games, every ENE is an SPE, although this implication does not hold in general. Figure 1 illustrates this distinction: the memoryless strategy profile in Figure 1b is an SPE from the designated initial vertex but fails to be an ENE, whereas the strategy profile depicted in Figure 1c is both an ENE and an SPE. Moreover, building these equilibria is simple and elegant, and our algorithm runs in polynomial time.

We complement our existence results with negative examples, which show that we cannot expect much more: when we have players with reachability (or Büchi) and safety (or CoBüchi) objectives, there may not be positional SPEs. For Muller objectives that do not fall into the class we treat, the situation is worse: even if we ask for the dual class, where all losing sets of vertices are up- or downwards closed, these games may not have any memoryless NE at all.

Related work. We discuss three related directions: memory requirements in games, decision problems for equilibria and the interplay of memory and randomisation in strategy complexity.

Memory requirements in games on graphs have been thoroughly studied for zero-sum games. In turn-based deterministic arenas, Gimbert and Zielonka provide a characterisation of objectives for which there exist optimal memoryless strategies for both players [23]. In the finite-memory case, there exists an analogous characterisation for objectives for which the memory structure for winning strategies of both players can be chosen independently of the arena [6]. Characterisations of objectives for which memoryless strategies and finite-memory strategies suffice to win are given in [35] and [13] respectively using universal graphs. For multiplayer games, [30] provides sufficient conditions for the existence of finite-memory NEs based on the existence of finite-memory winning strategies in related zero-sum games.

From an algorithmic perspective, deciding the existence of an NE or an SPE is trivial in ω -regular games on turn-based deterministic arenas, as they always exist. A more relevant problem is the constrained equilibrium existence problem. In pure strategies, it asks whether there exists an equilibrium from a given initial vertex whose unique outcome is winning for (at least) at given subset of players. In randomised strategies, this translates to lower bounds on the satisfaction probability of each objective by the equilibrium. The main motivation is that, in general, one is not interested in any equilibrium, but instead desires an equilibrium that satisfies some objectives. The constrained (pure) NE existence problem is **NP**-complete for reachability [9], safety [18] and co-Büchi [42] games, and it can be solved in polynomial time for Büchi games. The constrained (pure) SPE existence problem is **PSPACE**-complete for reachability and safety games [8]. Solutions to the constrained equilibrium existence problem require memory in general; see [12] for a discussion of the role of memory in solutions to the constrained equilibrium existence problem in reachability games and [31] for upper bounds on memory requirements for solutions to the constrained NE existence problem in reachability, safety, Büchi and co-Büchi games. The constrained equilibrium existence problem has also been studied in more general models, e.g., in concurrent games [5, 7] and timed games [4, 11].

Finally, we discuss the interplay of memory and randomisation. In our construction, we use randomisation to ensure that all vertices of some parts of the arena are visited infinitely often instead of doing so with memory. Several works study trade-offs between memory and randomisation. For instance, it is shown in [15] that in zero-sum Muller games with upward-closed accepting sets, memory is necessary in general for pure winning strategies but not when using randomised strategies. The same property holds in games where the goal is to enforce a conjunction of a parity and mean-payoff objective [17]. More generally, one can lessen the amount of memory needed by considering randomised finite-memory strategies; see, e.g., [14, 26] for zero-sum Muller games. Memory bounds are sensitive to the chosen definition of a randomised finite-memory strategy; not all models are equally compact, nor are they equally expressive. The relative expressiveness of these models is discussed in [32].

Outline. We provide required definitions in Section 2. In Section 3, we present examples that establish the need of randomisation to obtain memoryless SPEs in the class of games we study, and show that memoryless NEs need not exist for Muller objectives outside of the class we treat. In Section 4, we prove our existence result: we first illustrate our construction then present an algorithm to construct memoryless SPEs.

2 Preliminaries

Notation. For any $n \geq 1$, we let $\llbracket n \rrbracket = \{1, \dots, n\}$. For any finite set $A \neq \emptyset$, we let $\mathcal{D}(A)$ denote the set of distributions over A , i.e., of functions $\mu: A \rightarrow [0, 1]$ such that $\sum_{a \in A} \mu(a) = 1$.

Arenas. Let $n \geq 1$. For each $i \in \llbracket n \rrbracket$, we denote by $\mathcal{P}_i$ the i -th player. Let $G = (V, E)$ be a directed graph where V is a finite set of vertices and $E \subseteq V \times V$ is a set of edges. We assume that the graph is deadlock-free, i.e., for all $v \in V$, there exists $v' \in V$ such that $(v, v') \in E$. An n -player arena on G is a tuple $A = ((V_i)_{i \in \llbracket n \rrbracket}, E)$, where $(V_i)_{i \in \llbracket n \rrbracket}$ is a partition of V .

Players take turns choosing the next vertex depending on which player controls the current vertex. This induces an infinite path in the arena, called a play. Formally, a play in A is an infinite sequence of vertices $\rho = \rho_0 \rho_1 \dots \in V^\omega$ such that for all $\ell \in \mathbb{N}$, $(\rho_\ell, \rho_{\ell+1}) \in E$. A history is a finite prefix of a play $h = h_0 \dots h_\ell \in V^*$; we denote h_ℓ by $\text{last}(h)$. The set of plays (resp. histories) is denoted by Plays (resp. Hist), and we let $\text{Hist}_i = \text{Hist} \cap V^* V_i$.

Let C be a strongly connected component (SCC) of G . We say that C is a *leaf* SCC if there are no edges from C to $V \setminus C$, i.e., once a play enters C , it remains in C forever.

Strategies. Strategies describe how players select edges. They can use both memory and randomisation in general. Let $i \in \llbracket n \rrbracket$. Formally, a (randomised) strategy of $\mathcal{P}_i$ is a function $\sigma_i : \text{Hist}_i \rightarrow \mathcal{D}(V)$ such that for all $h \in \text{Hist}_i$ and all $v \in V$, if $\sigma_i(h)(v) > 0$, then $(\text{last}(h), v) \in E$. A strategy σ_i is pure if it uses no randomisation, i.e., if for all $h \in \text{Hist}_i$, $\sigma_i(h)$ is a Dirac distribution. A strategy σ_i is memoryless if for all histories $h, h' \in \text{Hist}_i$, $\text{last}(h) = \text{last}(h')$ implies $\sigma_i(h) = \sigma_i(h')$. We view memoryless strategies as functions $V_i \rightarrow \mathcal{D}(V)$. Strategies that are both memoryless and pure are called positional. We say that a play ρ is played according to σ_i if, whenever $\rho_\ell \in V_i$, we have that $\sigma_i(\rho_0 \dots \rho_\ell)(\rho_{\ell+1}) > 0$.

A strategy profile is a tuple $\sigma = (\sigma_1, \dots, \sigma_n)$ where, for all $i \in \llbracket n \rrbracket$, σ_i is a strategy of $\mathcal{P}_i$. We write $\mathbb{P}_v^\sigma$ for the distribution over plays induced by a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ from an initial vertex $v \in V$, defined in the usual way (see, e.g., [32]). A strategy profile σ is called memoryless (resp. positional) if all strategies in σ are memoryless (resp. positional).


Objectives and games. An objective is a measurable subset of Plays. We say that a play ρ satisfies an objective Ω if $\rho \in \Omega$. An n -player game $\mathcal{G} = (A, (\Omega_i)_{i \in \llbracket n \rrbracket})$ is an arena equipped with a profile of objectives $(\Omega_i)_{i \in \llbracket n \rrbracket}$.

We focus on the following three classes of objectives. A reachability objective is defined by a target set $T \subseteq V$; a play ρ satisfies the reachability objective for T if it visits T . Formally, we let $\text{Reach}(T) = \{\rho \in \text{Plays} \mid \exists \ell \geq 0, \rho_\ell \in T\}$. A safety objective is defined by a set $S \subseteq V$; a play ρ satisfies the safety objective for S if it never leaves S . Formally, we let $\text{Safe}(S) = \{\rho \in \text{Plays} \mid \forall \ell \geq 0, \rho_\ell \in S\}$.


For any play ρ , let $\text{Inf}(\rho)$ be the set of vertices occurring infinitely often in ρ . A Muller objective is defined by a family $\mathcal{F} \subseteq 2^V$. A play ρ satisfies the Muller objective described by $\mathcal{F}$, denoted by $\text{Muller}(\mathcal{F})$, if $\text{Inf}(\rho) \in \mathcal{F}$. An objective $\text{Muller}(\mathcal{F})$ is a 0-2 Muller objective if for all $F \in \mathcal{F}$, either all subsets of F are in $\mathcal{F}$ (F is downward closed in $\mathcal{F}$) or all supersets of F are in $\mathcal{F}$ (F is upward closed in $\mathcal{F}$), and $\text{Muller}(\mathcal{F})$ is a 1-3 Muller objective if its complement $\text{Muller}(2^V \setminus \mathcal{F})$ is a 0-2 Muller objective.

Nash equilibria. Let $\mathcal{G} = (A, (\Omega_i)_{i \in \llbracket n \rrbracket})$ be a game and $v \in V$ and let σ be a strategy profile. A strategy σ'_i of $\mathcal{P}_i$ is a profitable deviation for $\mathcal{P}_i$ from v with respect to σ if $\mathbb{P}_v^{(\sigma_{-i}, \sigma'_i)}(\Omega_i) > \mathbb{P}_v^\sigma(\Omega_i)$, where (σ_{-i}, σ'_i) denotes the profile obtained by replacing the strategy of $\mathcal{P}_i$ in σ by σ'_i . A strategy profile σ is a Nash equilibrium (NE) from $v \in V$ if there are no profitable deviations with respect to σ from v . A strategy profile σ is an everywhere NE (ENE) if for every vertex $v \in V$, σ is an NE from v .

Subgame-perfect equilibria. Given a possibly empty history $h \in \{\varepsilon\} \cup \text{Hist}$, we define the subgame of $\mathcal{G}$ rooted at h as $\mathcal{G}_h = (A, (\Omega_h^i)_{i \in \llbracket n \rrbracket})$, where for all $i \in \llbracket n \rrbracket$, we have $\Omega_h^i = \{\rho \in \text{Plays} \mid h\rho \in \Omega^i\}$. Given a strategy profile σ and a history h , we let σ_h denote the strategy profile defined by $(\tau_1, \dots, \tau_n)$ such that, for all $i \in \llbracket n \rrbracket$ and all histories $h' \in \text{Hist}_i$ such that hh' is a history, $\tau_i(h') = \sigma_i(hh')$. A strategy profile σ is a subgame-perfect equilibrium (SPE) from $v \in V$ in $\mathcal{G}$ if, for all histories $hv' \in \text{Hist}$ starting in v , σ^h is an NE in $\mathcal{G}_h$ from v' . In particular, every SPE is an NE.



■ **Figure 2** A six-player game: circles denote reachability players, hexagons safety players. (a) Game graph. (b) Fixing the direction of the play for reachability players. (c) The play unstable under unilateral deviations (Case (1), of Example 1).



■ **Figure 3** Case (2) of Example 1: fixing f counterclockwise forces d to play clockwise, yielding an unstable profile under unilateral deviations.

3 Impossibility Results


Before we introduce the construction of randomised SPEs for any combination of players with reachability, safety, or 0-2 Muller objective, we first show that randomisation is required even if we only consider reachability and safety objectives, while 1-3 Muller objectives may not have any memoryless NEs at all.

3.1 Randomisation

We illustrate the need of randomisation for reachability (or Büchi) and safety (or CoBüchi) on a simple example.

► **Example 1.** We now consider a game where six players are standing in a circle, and in each move they hand the token to a neighbour (see Fig. 2a). Players a , c , and e have the individual reachability objective that the token shall leave their respective neighbourhood (their own nodes + the two neighbouring nodes), or the Büchi objective that it shall do so infinitely often. The other three players, b , d , and f , have the individual safety objective never to leave their respective neighbourhood, or the CoBüchi objective to eventually always stay in this neighbourhood.

We now examine whether a positional ENE exists for this example. We start with the strategies of the reachability/Büchi players and observe that, modulo renaming, there are only two such strategies:



■ **Figure 4** The Mostowski hierarchy. For all levels below the separator, there exists a memoryless (randomised) ENE and SPE.

1. All reachability/Büchi players pass the token in the same direction (w.l.o.g. counter-clockwise).

Figure 2b shows this case together with the only optimal counter-strategy of the safety/-CoBüchi players to always play back/clockwise. (All other positional strategies of the individual safety/CoBüchi players can be improved upon and are therefore not stable). But for this strategy of the safety/CoBüchi players, none of the strategies of the reachability/Büchi players is stable: they lose from their vertices, but would win if they played clockwise. (violates the NE condition).

2. Two reachability/Büchi players pass the token toward the same player; w.l.o.g. we consider the case where a and c play to b , while e plays to f . (All other such strategies are equivalent to this one modulo renaming.)

Player f now has only one stable policy (to move back to e), and in an ENE, it has to take it. We therefore fix this decision.

If d also plays to e (Figure 3a), then the respective reachability/Büchi player can win from both a and c by changing their policy. But for either positional strategy played from b , either the reachability/Büchi player owning a or the reachability/Büchi player owning c loses from that vertex. Her choice is therefore not stable, so this cannot happen in an ENE.

If d instead plays to c , (Figure 3b), then e can improve her outcome by turning to d instead (Figure 3c), so this also cannot result in an ENE.


We can extend the example from ENEs to SPEs by simply adding a seventh vertex for a seventh player. This vertex is initial, safe for all safety players and not a target for either reachability player, and has all six vertices as successors, then the same argument works – it is just more difficult to draw and the additional initial vertex distracts from the core argument.

► **Theorem 2.** *Games with reachability or Büchi objectives as well as safety or CoBüchi objectives may neither allow for positional ENE nor for positional SPE.*

3.2 1-3 Muller games

We now show that Muller games where all rejecting infinity sets are upwards *or* downwards closed may not have *any* memoryless NE.

For this, we first very briefly outline the Mostowski hierarchy [33] for Muller objectives (see Figure 3.2). The 0-2 class is those objectives, where all winning sets are up- or downward closed, such as parity objectives with colours 0,1,2. The 1-3 class is the reverse: all losing sets are up- or downward closed, such as parity objectives with colours 1,2,3.



■ **Figure 5** A 1-3 Muller game with no memoryless NEs. Circles represent $\mathcal{P}_1$ and hexagons represent $\mathcal{P}_2$.

While we will treat the 0-2 class in the algorithm we introduce in the next section, we now provide an example, where two players with Muller objectives from the 1-3 class do not have any memoryless NE.

► **Example 3.** Consider a game with two players (circle and hexagon), where each player has two outgoing transitions to the other vertex, as depicted in Figure 5. The players have almost opposing objectives, except that both players lose if all four middle vertices are seen infinitely often. If three are seen infinitely often, the player for whom only one of her successor vertices is seen infinitely often wins. If two of the middle vertices are seen infinitely often, $\mathcal{P}_1$ wins for vertices a, c and b, d , while $\mathcal{P}_2$ wins for vertices a, d and b, c .

This way, if both players use randomisation, then both benefit from deviating, so this is no NE. If both players play pure, the losing player can win by changing to her other pure choice, so this cannot be an NE. If one player plays randomised while the other plays pure, the randomised player loses, but can change to a winning pure choice, so this cannot be an NE either.

We note that the losing sets with three and four of the middle vertices are upwards closed for either player, while the losing sets with two vertices are downwards closed for either player, so that these are 1-3 Muller objectives.

► **Theorem 4.** *Games with 1-3 Muller objectives may not have NEs, even if we restrict the games to mutually exclusive 1-3 Muller objectives and two players.*

4 Memoryless Nash Equilibria in Games with Reachability, Safety, and 0-2 Muller Objectives

In this section, we study the existence of memoryless NE for games with reachability, safety, and 0-2 Muller objectives. We first illustrate the principles and ideas of our construction for memoryless randomised NE for reachability, safety, and 0-2 Muller objectives. We then give the proof of the main result in Section 4.4.

4.1 Reachability Objectives

► **Example 5.** Consider the four-player game with reachability objectives played on the arena of Figure 6. The main idea is to incrementally define memoryless strategies in a way that allows as many players as possible to win. A vertex is considered evaluated once a move is assigned to it. When evaluating a vertex, we record the set of players who will win when the game starts at this vertex.

For reachability games, desirable strategies are particularly easy to identify in leaf SCCs: a *maximally mixed* strategy (i.e., such that all outgoing transitions of every vertex are played uniformly at random) will almost-surely visit all vertices in the leaf component infinitely often, thus in particular almost-surely satisfy any reachability objective with a target in the leaf. At the same time, players with no target in the leaf component cannot win at all.

Let C be the leaf SCC in the green area of Figure 6. Playing uniformly at random constitutes an NE from every vertex in C , we fix this strategy and can mark t_2 , v_3 , and t_4 as winning for $\mathcal{P}_2$ and $\mathcal{P}_4$, but not for the other two players.

Once we have evaluated vertices, we incrementally construct a strategy profile. As part of this process, we remove edges to evaluated vertices that are not marked as winning for the player controlling the outgoing vertex. Intuitively, such edges are not interesting for this player. To keep track of these modifications, we maintain a reduced arena containing exactly the edges that have not been removed.

In order to do this, we consider, for each unevaluated vertex $v \in V_i$ its successors.

1. If there exists a successor v' of v that is marked as winning for $\mathcal{P}_i$, we fix the edge $v \rightarrow v'$ in the strategy we build. This is safe to do because we have already established that $\mathcal{P}_i$ almost-surely wins from v' , thus $\mathcal{P}_i$ will almost-surely win from v this way; no other successor can provide a better result, and the decision to move to v will therefore be stable. We then mark v as winning for all players who are winning from v' , as well as for those whose target is v . For example, consider v_5 owned by $\mathcal{P}_4$, with successors v_3 and t_4 , both are marked as winning for her. We fix one of them (here, $v_5 \rightarrow t_4$) and mark v_5 as winning for $\mathcal{P}_2$ and $\mathcal{P}_4$.
2. If v has more than one successor, and among them is an evaluated successor v' that is marked as losing for $\mathcal{P}_i$, we remove the edge from $v \rightarrow v'$. This is safe to do because we have already established that we lose from that successor; so whichever of the remaining successors will be chosen, moving instead to v' cannot be advantageous for $\mathcal{P}_i$; not considering this edge in the future cannot lead to instability.
For example, consider v_2 owned by $\mathcal{P}_3$, with successors v_1 , v_3 , and v_5 . Since $\mathcal{P}_3$ does not win from the evaluated vertices v_3 and v_5 , we can successively remove the edges $v_2 \rightarrow v_3$ and $v_2 \rightarrow v_5$. This leaves v_2 with a single successor. While we cannot evaluate v_2 yet (as v_1 is not evaluated), we have simplified the game.
3. If v has only one successor v' and that successor is evaluated, we fix the edge $v \rightarrow v'$ and evaluate v as in Rule 1. This is safe to do as there is no other edge left to consider.


To continue with the example, we could next evaluate t_1 , owned by $\mathcal{P}_2$, fix the edge to t_4 (i.e., $t_1 \rightarrow t_4$), and mark t_1 as winning for $\mathcal{P}_1$, $\mathcal{P}_2$, and $\mathcal{P}_4$ by Rule 1. We continue by evaluating t_3 , v_1 , v_2 , v_4 , and v_0 . The resulting strategy profile is depicted in Figure 6, where solid arrows represent the selected edges (with associated probabilities).

The vertices may also be evaluated in a different order, i.e., the construction is non-deterministic. For instance, Rule 1 may offer different successor vertices to choose from while Rule 2 may offer different edges to drop. Therefore, the approach may yield several strategy profiles.


If unevaluated vertices remain and none of Rules 1-3 apply, all unevaluated vertices have only unevaluated successors. Due to the edge deletions, there is a new leaf SCC consisting of unevaluated vertices; we extend the strategy by letting the players play uniformly at random within this leaf SCC. This phenomenon occurs in the game of Figure 7: after evaluating the leaf component $\{v_3, t_1, t_4\}$ and v_5 , and removing the edges from v_2 and t_1 to the aforementioned vertices (their targets cannot be reached using them), a new leaf emerges in the reduced arena (Figure 8).

4.2 Safety Objectives

While playing maximally mixed in leaf SCCs is optimal for players with reachability objectives, it is the worst opponent for safety players – which curiously provides a different form of stability.



■ **Figure 6** A four-player reachability game where $\mathcal{P}_1$ is represented by circles, $\mathcal{P}_2$ by squares, $\mathcal{P}_3$ by diamonds, and $\mathcal{P}_4$ by hexagons. For each i , t_i denotes the target of player i . Edges that receive no probability weight in our example run of the algorithm are shown as dotted.



■ **Figure 7** A four-player reachability game, where t_i denotes a target of $\mathcal{P}_i$ and t_{ij} denotes a target shared by $\mathcal{P}_i$ and $\mathcal{P}_j$.

► **Example 6.** Let us consider the four-player game played on the arena shown in Figure 9. Each player $\mathcal{P}_i$ is associated with a set of safe vertices $S_i \subseteq V$, as depicted in the figure.


Similarly to Example 5, we start from the leaf SCC. $\mathcal{P}_4$ is not winning under uniform random play, but can win unilaterally from s_{124} by fixing her strategy to move to s_{24} ; we fix this edge and remove the edge to s_3 . This is safe to do because it guarantees that $\mathcal{P}_4$ is winning regardless of how the other players play, as she wins surely against a maximally mixed strategy. (This is the stability mentioned earlier: if the other players narrow down their options later, she will still win, so that these choices remain stable.) We do *not* evaluate at this point, we only fix the strategy of $\mathcal{P}_4$.

This strategy for $\mathcal{P}_4$ changes the structure of the game, with the new (only) leaf consisting of s_{24} and s_{124} . In this leaf component, there is no player who loses when all vertices in the leaf component are visited infinitely often, but can win unilaterally anywhere in this leaf component if all other players play this way. We now fix this strategy for the complete leaf, knowing that it is stable, and evaluate it by marking s_{24} and s_{124} as winning for $\mathcal{P}_2$ and $\mathcal{P}_4$.


We then select winning successors according to Rule 1, discard losing successors (where alternatives remain) according to Case 2, and evaluate vertices where the only successor is evaluated according to (Rule 3), much like for the reachability case.

In this example, we could evaluate s_3 next, as it has only one successor. Vertex s_3 is only safe for $\mathcal{P}_3$, and its only successor is marked as winning for $\mathcal{P}_2$ and $\mathcal{P}_4$, so that no player wins from s_3 . The remaining vertices can be evaluated in the following order: v_3 , s_{12} , v_2, s_{134} , v_3 , s_{14} , and v_0 , as in Example 5. We depict a memoryless strategy profile σ obtained in this way by solid arrows, while dotted arrows denote unused edges. This profile is an ENE and an SPE from all vertices.

Like in the reachability case, new leaves can occur during evaluation, and we turn to their evaluation when none of the Rules 1-3 apply.



■ **Figure 8** The reduced arena of (Figure 7) after removing edges (removed edges are shown as dotted lines) from evaluated vertices; a new leaf SCC emerges (highlighted in gray).



■ **Figure 9** A four-player safety game. Vertices s_i , s_{ij} , and s_{ijk} are safe for $\mathcal{P}_i, \mathcal{P}_j$ and $\mathcal{P}_k$ respectively, and v_i is safe for all players. Edges that receive no probability weight in our example run of the algorithm are shown as dotted.

4.3 0-2 Muller Objectives

While we have seen that 1-3 Muller objectives may not have any memoryless NE, 0-2 Muller objectives share enough properties with reachability that we can treat them similarly.

Maximally mixing strategies in leaves are at the same time optimal for upwards closed winning sets (much like for reachability objectives) and the worst opponent for downwards closed winning sets (much like for safety objectives). As a result, we can treat them like safety games: if a player almost-surely loses in a leaf when all players play uniformly at random, but can almost surely win unilaterally on a component within the leaf, they can do so using a positional strategy and we can fix it. This is because this leaf cannot contain an upwards closed winning region of that player – for if it would, then the leaf itself would be a winning region and all players playing uniformly at random would be almost surely winning for her. Consequently, she can only win if she can force the game to eventually stay in some downward closed winning set. This is a safety objective; in particular, if other players later remove some of their edges, the region visited can only shrink, and the downward closedness ensures that the resulting infinity set is winning for her.

4.4 Constructing Nash Equilibria

We now generalise the ideas presented earlier in this section to prove the existence of memoryless SPE in games with players that may have reachability objectives, safety objectives, and 0-2 Muller objectives. Formally, the main result of this section is the following.

► **Theorem 7.** *Let $\mathcal{G} = (A, (\Omega_i)_{i \in \llbracket n \rrbracket})$ be a game where, for all $i \in \llbracket n \rrbracket$, Ω_i is either a reachability, safety or a 0-2 Muller objective (i.e., defined by a family $\mathcal{F}$ such that, for all $F \in \mathcal{F}$, either all subsets of F are in $\mathcal{F}$ or all supersets of F are in $\mathcal{F}$). Then there exists a strategy profile σ that is a memoryless randomised SPE from any vertex such that for all $i \in \llbracket n \rrbracket$ and all $v \in V$, $\mathbb{P}_v^\sigma(\Omega_i) \in \{0, 1\}$.*

Let $\mathcal{G} = (A, (\Omega_i)_{i \in \llbracket n \rrbracket})$ be such a game. We provide an algorithm to construct this SPE. In the following, we view a memoryless strategy profile as a function $\sigma: V \rightarrow \mathcal{D}(V)$.

We construct the strategy profile incrementally as a sequence $(\sigma_k)_{k \geq 0}$ of partially-defined strategy profiles, where each σ_{k+1} extends σ_k . We label the vertices v in the domain of σ_k by the set of players W_v whose objectives are almost-surely satisfied from v when playing according to σ_k (this set is independent of k). In some steps of the algorithm, we may not extend σ_k and instead remove edges that do not lead to a winning scenario for the player who controls the vertex.

Initially, we let σ_0 be defined over the empty set. We say that a vertex $v \in V$ is evaluated (at step k) if $\sigma_k(v)$ is defined, and say that v is unevaluated otherwise. The algorithm iteratively applies the following rules in arbitrary order until the domain of σ_k is V .

1. If there are no unevaluated vertices with evaluated successors, pick a leaf SCC C . We distinguish the following rules:
 - a. There exists a player $\mathcal{P}_i$ whose objective is almost-surely falsified from C when all players play uniformly at random, but who can win with positive probability from some vertex in C when the other players play uniformly at random. In this case, the objective of $\mathcal{P}_i$ is either a safety objective or a Muller objective given by $\mathcal{F}_i$ whose satisfaction is witnessed by a set $F \in \mathcal{F}_i$ whose subsets are all in $\mathcal{F}_i$. In either case, there exists a subset $C' \subseteq C$ and a positional strategy τ_i such that, for all strategy profiles τ_{-i} of the other players that do not exit C , all plays played according to (τ_i, τ_{-i}) starting in C' remain in C' and satisfy Ω_i (i.e., τ_i is a winning strategy in a zero-sum game and C' is the set of vertices from which $\mathcal{P}_i$ has a sure winning strategy). We fix such a pure memoryless winning strategy τ_i for $\mathcal{P}_i$ over C' and remove outgoing edges of C' inconsistent with τ_i . We let $\sigma_{k+1} = \sigma_k$.
 - b. Otherwise, we extend σ_k to σ_{k+1} by choosing successors uniformly at random from all vertices in C . For all $v \in C$, we define W_v to be the set of players whose objectives are satisfied when playing according to σ_{k+1} (we explain how to determine this set below).
2. If there exists $i \in \llbracket n \rrbracket$ and an unevaluated vertex $v \in V_i$ such that one of its successors v' is evaluated and satisfies $i \in W_{v'}$, we define $\sigma_{k+1}(v) = v'$. We define W_v as above and remove all other outgoing edges of v .
3. If there exists a vertex v that is unevaluated with a unique evaluated successor v' , we let $\sigma_{k+1}(v) = v'$. We define W_v as above.
4. Finally, if there exists $i \in \llbracket n \rrbracket$ and an unevaluated vertex $v \in V_i$ that has at least two successors, one of these successors v' is evaluated and satisfies $i \notin W_{v'}$, we remove the edge from v to v' and let $\sigma_k = \sigma_{k+1}$.

We let σ denote an output of the above algorithm (it need not be unique). The sets W_v can be computed in practice as follows. When playing uniformly at random in a leaf SCC C due to Rule 1, the set W_v for any $v \in C$ contains the players with a reachability objective with a target that intersects C , those with a safety objective for whom all vertices of C are safe and those with a Muller objective for which C is in the family defining the objective. As all vertices of a leaf SCC C are visited infinitely often almost-surely when playing at random, we remark that for all $v \in C$ and $i \notin W_v$, $\mathbb{P}_v^\sigma(\Omega_i) = 0$. For a vertex v such that $\sigma(v)$ is given

by Rule 2 or 3, for all $i \in \llbracket n \rrbracket$, whether $i \in W_v$ depends on the objective of $\mathcal{P}_i$: if $\mathcal{P}_i$ has a reachability objective, then $i \in W_v$ if, and only if, $v \in W_{\sigma(v)}$ or v is in the target of $\mathcal{P}_i$; if $\mathcal{P}_i$ has a safety objective, then $i \in W_v$ if, and only if, $v \in W_{\sigma(v)}$ and v is safe for $\mathcal{P}_i$; if $\mathcal{P}_i$ has a Muller objective, then $W_v = W_{\sigma(v)}$. We obtain the following property of σ by the above.

► **Lemma 8.** *For all $i \in \llbracket n \rrbracket$ and all $v \in V$, $\mathbb{P}_v^\sigma(\Omega_i) \in \{0, 1\}$.*

We now prove that σ is an ENE. Our proof is by induction: we show that, for all $k \in \mathbb{N}$, any strategy σ' that extends σ_k is an NE from the vertices in the domain of σ_k .

► **Lemma 9.** *Any strategy profile σ constructed by the above algorithm is an ENE.*

Proof. We prove the result by induction. In the base case, σ_0 is defined over $\emptyset$ and therefore the claim is vacuously true.

We now assume by induction that any memoryless strategy profile extending σ_k is an NE from every vertex the domain of σ_k . Let σ_{k+1} be the extension of σ_k obtained after one additional iteration, and let σ' be a (fully defined) memoryless profile extending σ_{k+1} . We show that it is an NE from every vertex in the domain of σ_{k+1} . It suffices to show that no player has a profitable deviation with respect to σ' from vertices over which σ_{k+1} is defined and σ_k is not. We consider three rules.

First, assume that σ_{k+1} is obtained from σ_k by playing uniformly at random over some set of vertices C by Rule 1 (C is a leaf SCC that emerged during the algorithm). We show that σ' is an NE from all $v \in C$. We remark that $W_v = W_{v'}$ for all $v, v' \in C$. Let $i \in \llbracket n \rrbracket$. If, for all $v \in C$, $i \in W_v$ (i.e., $\mathbb{P}_v^{\sigma'}(\Omega_i) = 1$), then $\mathcal{P}_i$ has no profitable deviation from any $v \in C$ with respect to σ' . Assume therefore that $\mathbb{P}_v^{\sigma_{k+1}}(\Omega_i) = 0$ (cf. Lemma 8). By Rule 1(a), $\mathcal{P}_i$ has no profitable deviation that remains in C . If all outgoing edges of $V_i \cap C$ are in C , this shows that $\mathcal{P}_i$ has no profitable deviation. We now show that $\mathcal{P}_i$ cannot improve their gain by using an edge from $V_i \cap C$ to $V \setminus C$. Assume towards contradiction that this is the case, and let (v, v') be such an edge. This edge has been removed by Rule 4, and thus $i \notin W_{v'}$. It follows that $\mathcal{P}_i$ has a profitable deviation from v' with respect to σ' . This contradicts the induction assumption that σ' is an NE from v' .

Second, we assume that σ_{k+1} is obtained from σ_k by Rule 2 and let v be the vertex that is added to the domain of σ_k . Let $i \in \llbracket n \rrbracket$. First, assume that $v \in V_i$. If $\mathcal{P}_i$ has a reachability or Muller objective, we have $i \in W_v$ and thus $\mathcal{P}_i$ has no profitable deviation. If $\mathcal{P}_i$ has a safety objective, then $i \in W_v$ if and only if v is safe for $\mathcal{P}_i$; $\mathcal{P}_i$ has no profitable deviation from v regardless of the safety of v . Assume now that $\mathcal{P}_i$ does not control v . If $\mathcal{P}_i$ has a profitable deviation from v , then $\mathcal{P}_i$ must have a profitable deviation from $\sigma_{k+1}(v)$. By induction, there are no such profitable deviations.

Finally, we assume that σ_{k+1} is obtained from σ_k by Rule 3. Let v be the vertex that is added to the domain of σ_k . Let $i \in \llbracket n \rrbracket$. We first consider the case $v \in V_i$. We observe that no outgoing edges of v could have been removed by Rule 1(a); otherwise, we would have $i \in W_v$. It follows that all successors v' of v satisfy $i \notin W_{v'}$ by Rule 4. If $\mathcal{P}_i$ has a profitable deviation from v , it must be the case that $\mathcal{P}_i$ has a profitable deviation from one of the successors, and there are no such deviations by induction. We conclude similarly as in the previous case for players such that $v \notin V_i$.

We have shown that σ' is an NE from all vertices in the domain of σ_{k+1} . This shows that σ is an ENE. ◀

We now prove Theorem 7 by using Lemma 9.

Proof of Theorem 7. Let $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ be a strategy profile constructed by the algorithm. By Lemma 8, we have $\mathbb{P}_v^\sigma(\Omega_i) \in \{0, 1\}$ for all $i \in \llbracket n \rrbracket$ and all $v \in V$.

It remains to prove that σ is an SPE from all vertices. We prove it by contradiction. Assume that there exists a history hv (with h a history and $v \in V$), a player $\mathcal{P}_i$ and a profitable deviation σ'_i for $\mathcal{P}_i$ with respect to σ from v in the subgame $\mathcal{G}_h$. Then σ'_i is also a profitable deviation from v with respect to σ in $\mathcal{G}$ (due to the objectives we consider), contradicting Lemma 9. This shows that σ is an SPE from all initial vertices. ◀

We now highlight a restricted class of games with reachability and safety objectives in which we can show the existence of pure memoryless strategy profiles that are SPEs from all vertices by adapting our construction. To identify this class, we reflect on when randomisation is used in our construction, and when it is absent or can be omitted. Randomised decisions are only made in leaf SCCs that emerge during the algorithm. Its purpose is to enforce stability by visiting all targets within the SCC for the reachability players and by violating the objectives of the safety players for whom there is an unsafe vertex in the SCC.

We identify two situations in which randomisation is not used or is not helpful in achieving stability when our algorithm builds a strategy profile. First, if all vertices in a leaf SCC have a single successor, then all strategies of this SCC are pure by default. In particular, in absorbing vertices (i.e., vertices who are their own unique successor), randomisation is not used. Second, if no targets and no unsafe states occur in a leaf SCC, fixing any pure strategy on this SCC leads to a behaviour that is as stable as a maximally mixed strategy. On the one hand, for the safety players, plays that stay in this leaf SCC satisfy their safety objective, therefore they have no profitable deviation. On the other hand, for the reachability players, they cannot win within the SCC and none of the deleted edges that exit the leaf SCC can lead to a profitable deviation by construction (as explained in the proof of Lemma 9).

These observations suggest that if all targets and all unsafe states are absorbing, we can construct pure memoryless strategy profiles that are SPEs from all vertices by changing Rule 1(b) to fix an arbitrary outgoing edge in all vertices of leaf SCCs. This yields the following result.

► **Theorem 10.** *Let $\mathcal{G} = (A, (\Omega_i)_{i \in \llbracket n \rrbracket})$ be a game where, for all $i \in \llbracket n \rrbracket$, Ω_i is either a reachability or a safety objective such that all unsafe vertices and target vertices are absorbing. Then there exists a strategy profile σ that is a memoryless pure SPE from any vertex.*

Theorem 10 slightly generalises the main result of [28], which states this existence result for NE. The proof approach used in [28] is not algorithmic: the main idea is to consider a discounted version of the game and then transform a memoryless equilibrium into a positional one. While the existence of memoryless equilibria in discounted games is known [21], the proof is not constructive and relies on Kakutani's fixed point theorem [27]. In contrast, the adaptation of our algorithm outlined above provides a direct construction of an equilibrium.

5 Conclusion

In this paper, we have constructed memoryless randomised ENEs in turn-based games where all players have reachability, safety, or simple Muller objectives whose winning sets are up- or downward closed (the 0-2 class in the Mostowski hierarchy). We have first shown that randomisation is necessary even for simple combinations of reachability and safety objectives, and that memoryless NEs may fail to exist for 1-3 Muller objectives. Our algorithm constructs an ENE in polynomial time. We proved that every ENE is an SPE in this setting, establishing the existence of memoryless randomised SPEs for this class of games. It remains open whether games in which all players have only reachability objectives, or only safety objectives, admit positional NEs, ENEs, or SPEs.

References

- 1 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- 2 Roderick Bloem, Krishnendu Chatterjee, and Barbara Jobstmann. Graph games and reactive synthesis. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors, *Handbook of Model Checking*, pages 921–962. Springer, 2018. doi:10.1007/978-3-319-10575-8_27.
- 3 Endre Boros, Vladimir Gurvich, Martin Milanic, Vladimir Oudalov, and Jernej Vicić. A three-person deterministic graphical game without Nash equilibria. *Discrete Applied Mathematics*, 243:21–38, 2018. doi:10.1016/J.DAM.2018.01.008.
- 4 Patricia Bouyer, Romain Brenguier, and Nicolas Markey. Nash equilibria for reachability objectives in multi-player timed games. In Paul Gastin and François Laroussinie, editors, *Proceedings of the 21st International Conference on Concurrency Theory, CONCUR 2010, Paris, France, August 31–September 3, 2010.*, Lecture Notes in Computer Science, pages 192–206. Springer, 2010. doi:10.1007/978-3-642-15375-4_14.
- 5 Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. Nash equilibria in concurrent games with büchi objectives. In Supratik Chakraborty and Amit Kumar, editors, *Proceedings of the 31st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2011, Mumbai, India, December 12–14, 2011*, LIPIcs, pages 375–386. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2011. doi:10.4230/LIPIcs.FSTTCS.2011.375.
- 6 Patricia Bouyer, Stéphane Le Roux, Youssouf Oualhadj, Mickael Randour, and Pierre Vandenhove. Games where you can play optimally with arena-independent finite memory. *Logical Methods in Computer Science*, 18(1), 2022. doi:10.46298/lmcs-18(1:11)2022.
- 7 Romain Brenguier. *Équilibres de Nash dans les jeux concurrents : application aux jeux temporisés. (Nash equilibria in concurrent games : application to timed games)*. PhD thesis, École normale supérieure de Cachan, France, 2012. URL: <https://tel.archives-ouvertes.fr/tel-00827027>.
- 8 Thomas Brihaye, Véronique Bruyère, Aline Goeminne, and Jean-François Raskin. Constrained existence problem for weak subgame perfect equilibria with ω -regular boolean objectives. *Information and Computation*, 278:104594, 2021. doi:10.1016/J.IC.2020.104594.
- 9 Thomas Brihaye, Véronique Bruyère, Aline Goeminne, and Nathan Thomasset. On relevant equilibria in reachability games. *Journal of Computer and System Sciences*, 119:211–230, 2021. doi:10.1016/j.jcss.2021.02.009.
- 10 Thomas Brihaye, Julie De Pril, and Sven Schewe. Multiplayer cost games with simple Nash equilibria. In Sergei N. Artëmov and Anil Nerode, editors, *Proceedings of the International Symposium on Logical Foundations of Computer Science, LFCS 2013, San Diego, CA, USA, January 6–8, 2013*, volume 7734 of *Lecture Notes in Computer Science*, pages 59–73. Springer, 2013. doi:10.1007/978-3-642-35722-0_5.
- 11 Thomas Brihaye and Aline Goeminne. On subgame perfect equilibria in turn-based reachability timed games. In Nathalie Bertrand and Nils Jansen, editors, *Proceedings of the 18th International Conference on Formal Modeling and Analysis of Timed Systems, FORMATS 2020, Vienna, Austria, September 1–3, 2020*, Lecture Notes in Computer Science, pages 94–110. Springer, 2020. doi:10.1007/978-3-030-57628-8_6.
- 12 Thomas Brihaye, Aline Goeminne, James C. A. Main, and Mickael Randour. Reachability games and friends: A journey through the lens of memory and complexity (invited talk). In Patricia Bouyer and Srikanth Srinivasan, editors, *Proceedings of the 43rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2023, IIT Hyderabad, Telangana, India, December 18–20, 2023*, volume 284 of *LIPIcs*, pages 1:1–1:26. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.FSTTCS.2023.1.
- 13 Antonio Casares and Pierre Ohlmann. Characterising memory in infinite games. *Logical Methods in Computer Science*, 21(1), 2025. doi:10.46298/LMCS-21(1:28)2025.

- 14 Krishnendu Chatterjee. Optimal strategy synthesis in stochastic Müller games. In Helmut Seidl, editor, *Proceedings of the 10th International Conference on Foundations of Software Science and Computational Structures, FoSSaCS 2007, Held as Part of ETAPS 2007, Braga, Portugal, March 24 – April 1, 2007*, volume 4423 of *Lecture Notes in Computer Science*, pages 138–152. Springer, 2007. doi:10.1007/978-3-540-71389-0_11.
- 15 Krishnendu Chatterjee, Luca de Alfaro, and Thomas A. Henzinger. Trading memory for randomness. In *Proceedings of the 1st International Conference on Quantitative Evaluation of Systems, QEST 2004, Enschede, The Netherlands, 27–30 September 2004*, pages 206–217. IEEE Computer Society, 2004. doi:10.1109/QEST.2004.1348035.
- 16 Krishnendu Chatterjee, Rupak Majumdar, and Marcin Jurdzinski. On Nash equilibria in stochastic games. In Jerzy Marcinkowski and Andrzej Tarlecki, editors, *Proceedings of the 18th International Workshop on Computer Science Logic, CSL 2004, 13th Annual Conference of the EACSL, Karpacz, Poland, September 20–24, 2004*, Lecture Notes in Computer Science, pages 26–40. Springer, 2004. doi:10.1007/978-3-540-30124-0_6.
- 17 Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin. Strategy synthesis for multi-dimensional quantitative objectives. *Acta Informatica*, 51(3-4):129–163, 2014. doi:10.1007/s00236-013-0182-6.
- 18 Rodica Condurache, Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. The complexity of rational synthesis. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, Rome, Italy, July 11–15, 2016*, volume 55 of *LIPIcs*, pages 121:1–121:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.ICALP.2016.121.
- 19 Julie De Pril. *Equilibria in Multiplayer Cost Games*. PhD thesis, UMONS, 2013.
- 20 Nathanaël Fijalkow, Nathalie Bertrand, Patricia Bouyer-Decitre, Romain Brenguier, Arnaud Carayol, John Fearnley, Hugo Gimbert, Florian Horn, Rasmus Ibsen-Jensen, Nicolas Markey, Benjamin Monmege, Petr Novotný, Mickael Randour, Ocan Sankur, Sylvain Schmitz, Olivier Serre, and Mateusz Skomra. Games on graphs. *CoRR*, abs/2305.10546, 2023. doi:10.48550/arXiv.2305.10546.
- 21 A Fink. Equilibrium points of stochastic noncooperative games. *Journal of Science Hiroshima University Series A*, 28:89–93, 1964.
- 22 János Flesch, Frank Thuijsman, and O. J. Vrieze. Recursive repeated games with absorbing states. *Mathematics of Operations Research*, 21(4):1016–1022, 1996. doi:10.1287/MOOR.21.4.1016.
- 23 Hugo Gimbert and Wiesław Zielonka. Games where you can play optimally without any memory. In Martín Abadi and Luca de Alfaro, editors, *Proceedings of the 16th International Conference on Concurrency Theory, CONCUR 2005, San Francisco, CA, USA, August 23–26, 2005*, volume 3653 of *Lecture Notes in Computer Science*, pages 428–442. Springer, 2005. doi:10.1007/11539452_33.
- 24 Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002. doi:10.1007/3-540-36387-4.
- 25 Kristoffer Arnsfelt Hansen and Mikhail A. Raskin. A stay-in-a-set game without a stationary equilibrium. In Jérôme Leroux and Jean-François Raskin, editors, *Proceedings of the Tenth International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2019, Bordeaux, France, 2–3 September 2019*, EPTCS, pages 83–90, 2019. doi:10.4204/EPTCS.305.6.
- 26 Florian Horn. Random fruits on the Zielonka tree. In Susanne Albers and Jean-Yves Marion, editors, *Proceedings of the 26th International Symposium on Theoretical Aspects of Computer Science, STACS 2009, Freiburg, Germany, February 26–28, 2009*, volume 3 of *LIPIcs*, pages 541–552. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2009. doi:10.4230/LIPIcs.STACS.2009.1848.

- 27 Shizuo Kakutani. A generalization of Brouwer's fixed point theorem. *Duke Mathematical Journal*, 8(3):457–459, 1941. doi:10.1215/S0012-7094-41-00838-4.
- 28 Athanasios Kehagias. A note on the Nash equilibria of some multi-player reachability/safety games. *Journal of Dynamics and Games*, 9(1):117–122, 2022. doi:10.3934/jdg.2021028.
- 29 Jeroen Kuipers, János Flesch, Gijs Schoenmakers, and Koos Vrieze. Pure subgame-perfect equilibria in free transition games. *European Journal of Operational Research*, 199(2):442–447, 2009. doi:10.1016/J.EJOR.2008.11.038.
- 30 Stéphane Le Roux and Arno Pauly. Extending finite-memory determinacy to multi-player games. *Information and Computation*, 261:676–694, 2018. doi:10.1016/j.ic.2018.02.024.
- 31 James C. A. Main. Arena-independent memory bounds for Nash equilibria in reachability games. *Information and Computation*, 309:105425, 2026. doi:10.1016/J.IC.2026.105425.
- 32 James C. A. Main and Mickael Randour. Different strokes in randomised strategies: Revisiting Kuhn's theorem under finite-memory assumptions. *Information and Computation*, 301:105229, 2024. doi:10.1016/J.IC.2024.105229.
- 33 Andrzej Włodzimierz Mostowski. Regular expressions for infinite trees and a standard form of automata. In Andrzej Skowron, editor, *Computation Theory - Fifth Symposium, Zaborów, Poland, December 3-8, 1984, Proceedings*, volume 208 of *Lecture Notes in Computer Science*, pages 157–168. Springer, 1984. doi:10.1007/3-540-16066-3_15.
- 34 John F. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950. doi:10.1073/pnas.36.1.48.
- 35 Pierre Ohlmann. Characterizing positionality in games of infinite duration over infinite graphs. *TheoretiCS*, 2, 2023. doi:10.46298/THEORETICS.23.3.
- 36 Amir Pnueli. The temporal logic of programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science, FOCS 1977, Providence, Rhode Island, USA, October 31 – November 1, 1977*, pages 46–57. IEEE Computer Society, 1977. doi:10.1109/SFCS.1977.32.
- 37 Mickael Randour. Simplicity lies in the eye of the beholder: A strategic perspective on controllers in reactive synthesis. In Pierre Ganty and Alessio Mansutti, editors, *Proceedings of the 19th International Conference on Reachability Problems, RP 2025, Madrid, Spain, October 1–3, 2025*, *Lecture Notes in Computer Science*, pages 31–48. Springer, 2025. doi:10.1007/978-3-032-09524-4_3.
- 38 Piercesare Secchi and William D. Sudderth. Stay-in-a-set games. *International Journal of Game Theory*, 30(4):479–490, 2002. doi:10.1007/S001820200092.
- 39 Reinhard Selten. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit: Teil I: Bestimmung des dynamischen Preisgleichgewichts. *Zeitschrift für die gesamte Staatswissenschaft*, 121(2):301–324, 1965.
- 40 Masayuki Takahashi. Equilibrium points of stochastic non-cooperative n -person games. *Journal of Science of the Hiroshima University, Series AI (Mathematics)*, 28(1):95–99, 1964.
- 41 Michael Ummels. Rational behaviour and strategy construction in infinite multiplayer games. In S. Arun-Kumar and Naveen Garg, editors, *Proceedings of the 26th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FST TCS 2006, Kolkata, India, December 13–15, 2006*, volume 4337 of *Lecture Notes in Computer Science*, pages 212–223. Springer, 2006. doi:10.1007/11944836_21.
- 42 Michael Ummels. The complexity of Nash equilibria in infinite multiplayer games. In Roberto M. Amadio, editor, *Proceedings of the 11th International Conference on Foundations of Software Science and Computational Structures, FoSSaCS 2008, Held as Part of ETAPS 2008, Budapest, Hungary, March 29 – April 6, 2008*, volume 4962 of *Lecture Notes in Computer Science*, pages 20–34. Springer, 2008. doi:10.1007/978-3-540-78499-9_3.
- 43 Michael Ummels. *Stochastic multiplayer games: theory and algorithms*. PhD thesis, RWTH Aachen University, 2011. URL: <http://darwin.bth.rwth-aachen.de/opus3/volltexte/2011/3451/pdf/3451.pdf>.