



One-To-Two-Player Lifting for Mildly Growing Memory

Alexander Kozachinskiy  

Steklov Mathematical Institute of Russian Academy of Sciences, Moscow, Russia

Abstract

We investigate a phenomenon of “one-to-two-player lifting” in infinite-duration two-player games on graphs with zero-sum objectives. More specifically, let \mathcal{C} be a class of strategies. It turns out that in many cases, to show that all two-player games on graphs with a given payoff function are determined in \mathcal{C} , it is sufficient to do so for one-player games. That is, in many cases the determinacy in \mathcal{C} can be “lifted” from one-player games to two-player games. Namely, Gimbert and Zielonka (CONCUR 2005) have shown this for the class of positional strategies. Recently, Bouyer et al. (CONCUR 2020) have extended this to the classes of arena-independent finite-memory strategies. Informally, these are finite-memory strategies that use the same way of storing memory in all game graphs.

In this paper, we put the lifting technique into the context of memory complexity. The memory complexity of a payoff function measures, how many states of memory we need to play optimally in game graphs with up to n nodes, depending on n . We address the following question. Assume that we know the memory complexity of our payoff function in one-player games. Then what can be said about its memory complexity in two-player games? In particular, when is it finite?

In this paper, we answer this questions for strategies with “chromatic” memory. These are strategies that only accumulate sequences of colors of edges in their memory. We obtain the following results.

- Assume that the chromatic memory complexity in one-player games is sublinear in n on some infinite subsequence. Then the chromatic memory complexity in two-player games is finite.
- We provide an example in which (a) the chromatic memory complexity in one-player games is linear in n ; (b) the memory complexity in two-player games is infinite.

Thus, we obtain the exact barrier for the one-to-two-player lifting theorems in the setting of chromatic finite-memory strategies. Previous results only cover payoff functions with constant chromatic memory complexity.

2012 ACM Subject Classification Theory of computation → Logic

Keywords and phrases Games on graphs, one-to-two-player lifting, strategy complexity, positional determinacy, finite-memory determinacy

Digital Object Identifier 10.4230/LIPIcs.STACS.2022.43

Related Version *Full Version:* [arXiv:2104.13888](https://arxiv.org/abs/2104.13888) [17]

Funding This work was performed at the Steklov International Mathematical Center and supported by the Ministry of Science and Higher Education of the Russian Federation (agreement no. 075-15-2019-1614).

1 Introduction

We study two-player infinite-duration games on graphs. These games are of interest in many areas of computer science, ranging from purely theoretical disciplines, such as decidability of logical theories [22, 23], to more practically-oriented ones, such as controller synthesis [15].

These games are played as follows. There is a finite directed graph with a token. We will call this graph *arena*. Initially, the token is placed in one of the nodes of the arena. In each turn, one of the two players takes the token and moves it to some other node. A restriction is that there must be an edge to the new location of the token. For each node of the arena, it



© Alexander Kozachinskiy;

licensed under Creative Commons License CC-BY 4.0

39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022).

Editors: Petra Berenbrink and Benjamin Monmege; Article No. 43; pp. 43:1–43:21

Leibniz International Proceedings in Informatics



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



is fixed in advance which of the players is the one to move the token in this node. The game proceeds for infinitely many turns. The outcome of the game is decided by the resulting trajectory of the token (it forms an infinite path in the arena).

We restrict ourselves to *zero-sum* games. Correspondingly, the players will be called Max and Min from now on. In a zero-sum game, objectives of the players are defined through a *payoff function* – a function of the form $\varphi: C^\omega \rightarrow \mathcal{W}$, where C is a set of *colors*, and (\mathcal{W}, \leq) is an arbitrary linearly ordered set. Next, we assume that arenas are edge-colored by elements of C . To compute the outcome of a play (which will be an element of \mathcal{W}), we take the trajectory of the token in this play, then consider the infinite sequence of colors $\gamma \in C^\omega$ written on the edges of the trajectory, and, finally, apply φ to γ . The aim of Max is to maximize $\varphi(\gamma)$, while the aim of Min is to minimize it (with respect to the ordering of \mathcal{W}).

As usually, a pair of strategies of the players in which the first strategy is the best response to the second one, and vice versa, is called an *equilibrium*. Next, a strategy which belongs to some equilibrium is called *optimal*. Now, a payoff function is called *determined* if in every arena there exists an equilibrium with respect to this payoff function.

We will study determinacy with respect to restricted classes of strategies. Namely, if \mathcal{C} is a class of strategies, then we say that a payoff function is determined in \mathcal{C} if the following holds: in every arena there is an equilibrium for this payoff function in which both strategies are from \mathcal{C} . The smaller is \mathcal{C} , the stronger is this requirement.

One of the main research directions in the area of games of graphs is *strategy complexity*. Its goal, broadly speaking, is to find out, for a payoff function φ of our interest, what is the “simplest” class of strategies \mathcal{C} in which φ is determined. This is highly relevant when our task is to actually implement in practice one of the optimal strategies for φ . For instance, this is the case when we want to produce a device whose performance is measured by φ . If this device is meant to act in the environment, then the execution of this device can be modeled as a game – between the controller of the device and the environment. In this framework, the controller realizes one of the strategies in this game. Ideally, we want an optimal performance w.r.t. φ at the lowest cost (in terms of the resources we need to implement the controller). The lower is strategy complexity of φ , the easier is this task.

Classically, there are two classes of strategies that are often considered in this context. One is the class of *positional* strategies and the other is the class of *finite-memory* strategies.

Let us first consider positional strategies. A strategy is positional if, for every node v of the arena, it always makes the same move when the token is in v , no matter what was the path of the token to this node. Sometimes these strategies are called *memory-less* – they do not need to “remember” anything about the previous development of the game. For brevity, we call payoff functions that are determined in the class of positional strategies *positionally determined*. Classical examples of games with positionally determined payoff functions are Parity Games, Mean-Payoff Games and Discounted Games [21, 10, 24].

These games, especially Parity Games, had a tremendous impact on such areas as verification, model checking and program analysis [11, 12, 1]. However, say, in controller synthesis, it is often required to consider more complex games, namely, those for which positional strategies do not suffice. This brings us to a more general class of strategies – the class of finite-memory strategies.

Unlike positional strategies, finite-memory strategies can store some information about the previous development of the game. The point is that during the whole play, which is infinitely long, the amount of this information should never exceed some constant.

The storage of information in finite-memory strategies is carried out by *memory skeletons*. A memory skeleton \mathcal{M} is a deterministic finite automaton whose input alphabet is the set of colors. Now, an \mathcal{M} -strategy is a strategy which, informally, stores information according

to the memory skeleton \mathcal{M} . To understand how it works, imagine that during the game, each time the token is shifted along some edge, the color of this edge is fed to \mathcal{M} . Then, at every moment, the current state of \mathcal{M} represents the current content of the memory. Correspondingly, the moves of an \mathcal{M} -strategy depend solely on the current state of \mathcal{M} and the current node with the token.

A strategy is finite-memory if it is an \mathcal{M} -strategy for some memory skeleton \mathcal{M} . For brevity, we call payoff functions that are determined in the class of finite-memory strategies *finite-memory determined*.

► **Remark 1.** Finite-memory strategies as defined above are sometimes called “chromatic”. This is because one can consider a more general definition. Namely, one can allow memory skeletons to take the whole edge as an input, not only its color. However, as shown by Le Roux [18], determinacy in general finite-memory strategies is equivalent to determinacy in chromatic finite-memory strategies. In this paper, we work only with chromatic finite-memory strategies.

1.1 One-to-two-player lifting

One of the techniques in the area of strategy complexity is called *one-to-two-player lifting*. Our paper is devoted to this technique. It relies on the notion of *one-player arenas*. An arena is called one-player if for one of the players the following holds: all the nodes of the arena from which this player is the one to move have exactly one out-going edge. This means that one of the players is given no choice and has only one way of playing. Correspondingly, there are two types of one-player arenas – those in which Max has no choice and those in which Min has no choice.

It turns out that to study determinacy in some class of strategies \mathcal{C} , it is sometimes sufficient to consider only one-player arenas. As was shown by Gimbert and Zielonka [13], this applies to the class of positional strategies. More specifically, their result states the following. Assume that a payoff function is such that all *one-player* arenas have an equilibrium of two positional strategies¹ with respect to this payoff function. Then *all* arenas, not only one-player ones, have an equilibrium of two positional strategies with respect to this payoff function. That is, then this payoff function is positionally determined. In a way, this means the positional determinacy of one-player games can always be “lifted” to two-player games.

This result has fundamental significance for studying the positional determinacy. This is because often one-player arenas are considerably easier to analyze than two-player ones. Indeed, assume we have an arena in which, say, Min has no choice. A question of whether such an arena has a positional equilibrium reduces to the following question. Is there a “lasso” (a simple path to a simple cycle over which we rotate infinitely many times) which maximizes our payoff function over all infinite paths? Often this can be figured out with a simple graph reasoning. For instance, this is fairly easy for Parity Games and Mean Payoff Games. Thus, through the lifting theorem of Gimbert and Zielonka one gets simple proofs of positional determinacy of these games. In turn, proofs that existed prior to the paper of Gimbert and Zielonka were highly non-trivial.

Given such a success in the case of positional strategies, it is tempting to extend this to larger classes of strategies. This was recently investigated for the class of finite-memory strategies by Bouyer et al. in [4]. It turns out that the situation is quite different for this

¹ Note that in one-player arenas, one of the players has just one strategy (and this strategy is positional). So this requirement means that the other player has a positional strategy which is at least as good against the unique strategy of the opponent as any other strategy.

class. More specifically, Bouyer et al. have constructed a payoff function such that **(a)** all one-player arenas have an equilibrium of two finite-memory strategies with respect to this payoff function **(b)** there is an arena (in fact, with just 2 nodes) which is not one-player and which has no equilibrium of two finite-memory strategies with respect to this payoff function.

Thus, the class of positional strategies admits one-to-two-player lifting and the class of finite-memory strategies does not. Bouyer et al. suggested to study intermediate classes. Namely, their approach was as follows. By definition, the class of finite-memory strategies is the union of the classes of \mathcal{M} -strategies over all memory skeletons \mathcal{M} . Let us now fix a memory skeleton \mathcal{M} and consider the class of \mathcal{M} -strategies for this specific \mathcal{M} . Bouyer et al. show that for every \mathcal{M} this class admits one-to-two-player lifting.

More precisely, the lifting theorem of Bouyer et al. states that for any memory skeleton \mathcal{M} the following holds. Assume that a payoff function is such that all one-player arenas have an equilibrium of two \mathcal{M} -strategies. Then the same holds for all arenas, with exactly this memory skeleton \mathcal{M} . That is, then this payoff function is determined in \mathcal{M} -strategies.

Observe that positional strategies are exactly \mathcal{M} -strategies if the memory skeleton \mathcal{M} has just one state. Thus, the lifting theorem Bouyer et al. includes the lifting theorem of Gimbert and Zielonka as a special case.

Bouyer et al. call payoff functions to which one can apply their lifting theorem *arena-independent finite-memory determined*. That is, a payoff function φ is arena-independent finite-memory determined if there exists a memory skeleton \mathcal{M} such that φ is determined in \mathcal{M} -strategies.

In the literature there is a number of games with arena-independent finite-memory determined payoff functions. For example, one can list games with ω -regular winning conditions [7] and bounded multidimensional energy games [3]. In turn, unbounded multidimensional energy games are finite-memory determined but not arena-independently [9].

1.2 Our results

The aim of this work is to extend the lifting technique beyond the class of arena-independent finite-memory determined payoff functions.

For payoff functions beyond this class, there is no single memory \mathcal{M} skeleton which suffices for all arenas (here “suffices” means the existence of an equilibrium of two \mathcal{M} -strategies). Instead, larger arenas require larger memory skeletons. This motivates a notion of the *memory complexity* of a payoff function. It can be defined as follows. For every n consider the minimal memory skeleton which is sufficient for all arenas with up to n nodes (w.r.t. our payoff function). Let the size of this memory skeleton (that is, the number of its states) be S_n . Then we call the function $n \mapsto S_n$ the memory complexity of our payoff function. Observe that arena-independent finite-memory determined payoff functions have memory complexity $O(1)$.

The memory complexity is the decisive factor in practice – if it grows too quickly, we might have no resources to implement optimal strategies for our payoff function. This complexity measure was studied for a number of payoff functions in [8, 9]

We initiate the study of the memory complexity in the context of one-to-two-player lifting. More specifically, we address the following question. Assume that we know the memory complexity of our payoff function in one-player arenas. Then what can be said about its memory complexity in all arenas? Thus, our approach differs from the approach of Bouyer et al. in the following regard. Instead of lifting determinacy in some fixed class of strategies from one-player arenas to all arenas, we lift bounds on the memory complexity.

To formulate our results, we introduce the following notation. Let \mathbb{Z}^+ denote the set of positive integers, and let $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ be a function. Then by $\text{FMD}(f)$ we denote the class of all payoff functions φ such that for all $n \in \mathbb{Z}^+$ there exists a memory skeleton \mathcal{M} with at most $f(n)$ states such that every arena with at most n nodes has an equilibrium of two \mathcal{M} -strategies with respect to φ . In other words, $\text{FMD}(f)$ is the class of all payoff function with memory complexity at most f . We also introduce similar notation for one-player arenas. Namely, we let $\text{1playerFMD}(f)$ be the class of all payoff functions φ such that for all $n \in \mathbb{Z}^+$ there exists a memory skeleton \mathcal{M} with at most $f(n)$ states such that every *one-player* arena with at most n nodes has an equilibrium of two \mathcal{M} -strategies with respect to φ . Again, $\text{1playerFMD}(f)$ is the class of payoff functions whose memory complexity in one-player arenas is at most f . Obviously, $\text{FMD}(f) \subseteq \text{1playerFMD}(f)$. Additionally, we let FMD stand for the class of all finite-memory determined payoffs. Finally, let 1playerFMD be the class of all payoff functions φ such that every one-player arena has an equilibrium of two finite-memory strategies w.r.t. φ .

In this notation, the question we address in this paper can be formulated as follows: for which functions f and g do we have $\text{1playerFMD}(f) \subseteq \text{FMD}(g)$?

► **Remark 2.** One could consider an alternative definition of $\text{FMD}(f)$, in which different arenas of size up to n may be mapped to different memory skeletons of size $f(n)$. Unfortunately, it is not clear how to extend results of this paper to this setting.

Before presenting our results, let us express previous ones in this notation. For technical convenience, we assume from now on that the set C of colors is finite. This is not an essential restriction, as any arena involves only finitely many colors. Hence, if C is infinite, one can study, separately all finite subsets $C' \subseteq C$, arenas that involve colors only from C' .

First, let us understand what payoff functions are included² in $\text{FMD}(1)$. By definition, these are payoff functions such that for every n there is a memory skeleton \mathcal{M} with 1 state such that all arenas with up to n nodes are determined in \mathcal{M} -strategies – or, equivalently, in positional strategies. Thus, $\text{FMD}(1)$ is exactly the class of positionally determined payoff functions. Observe then that the lifting theorem of Gimbert and Zielonka can be stated as the equality $\text{1playerFMD}(1) = \text{FMD}(1)$.

In fact, the lifting theorem of Bouyer et al. asserts that, more generally, for any constant $k \in \mathbb{Z}^+$ we have $\text{1playerFMD}(k) = \text{FMD}(k)$. Indeed, take any $\varphi \in \text{1playerFMD}(k)$. Our goal is to show that $\varphi \in \text{FMD}(k)$. By definition, for every n there exists a memory skeleton \mathcal{M} with at most k states such that all one-player arenas with at most n nodes have an equilibrium of two \mathcal{M} -strategies w.r.t. φ . A problem is that these \mathcal{M} may be different for different n . However, since the set C of colors is finite, there are only finitely many memory skeletons with up to k states. One of them works for infinitely many n – and, hence, for all one-player arenas. Due to the lifting theorem of Bouyer et al., the same memory skeleton works for all arenas. Thus, since this memory skeleton has at most k states, we have $\varphi \in \text{FMD}(k)$.

Let us note that the class of arena-independent finite-memory determined payoffs is the class $\text{FMD}(O(1)) = \bigcup_{k \in \mathbb{Z}^+} \text{FMD}(k)$.

² Here, formally, by $\text{FMD}(1)$ we mean $\text{FMD}(f)$ for the function $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ such that $f(n) = 1$ for all $n \in \mathbb{Z}^+$. More generally, if there is some expression in n defining a function $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$, we will use FMD of this expression instead of $\text{FMD}(f)$. For example, if $f(n) = 2n^2 + 2$ for all $n \in \mathbb{Z}^+$, then we will write $\text{FMD}(2n^2 + 2)$ instead of $\text{FMD}(f)$.

Finally, since lifting does not hold for the whole class of finite-memory strategies, we have $\text{1playerFMD} \neq \text{FMD}$. In fact, this means that for some function f we have $\text{1playerFMD}(f) \not\subseteq \text{FMD}$. This is because

$$\text{FMD} = \bigcup_f \text{FMD}(f), \quad \text{1playerFMD} = \bigcup_f \text{1playerFMD}(f)$$

over all $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$. Why is it so? For example, let us show this for FMD. We have to show that for any $\varphi \in \text{FMD}$ and for every n there exists a memory skeleton \mathcal{M} such that all arenas with up to n nodes have an equilibrium of two \mathcal{M} -strategies (w.r.t φ). A point is that, since C is finite, for every n the number of such arenas is also finite (w.l.o.g. we may assume that between each pair of nodes there are at most $|C|$ edges). In each of these arenas, fix a pair of finite-memory strategies forming an equilibrium (this is possible since $\varphi \in \text{FMD}$). This gives a finite set of finite-memory strategies such that every arena with up to n nodes is determined in strategies from this set. It remains to set \mathcal{M} to be the product of the memory skeletons of these strategies. Then all these strategies will be \mathcal{M} -strategies.

We proceed to our main result. Let $\Omega(n)$ denote the set of functions $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ for which there exists $C > 0$ such that $f(n) \geq Cn$ for all $n \in \mathbb{Z}^+$. We obtain the following lifting theorem:

► **Theorem 3.** *Consider any function $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$, $f \notin \Omega(n)$. Define $g: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$, $g(n) = f\left(\min\left\{m \mid \frac{f(m)}{m+1} \leq \frac{1}{2n}\right\}\right)$. Then $\text{1playerFMD}(f) \subseteq \text{FMD}(g)$.*

First, why is the function g well-defined? Since $f \notin \Omega(n)$, the fraction $f(m)/m$ gets arbitrarily close to 0 for some m . Hence, the minimum in the definition of g is always over a non-empty set.

Now consider the case when, as in the lifting theorem of Bouyer et al., the function f is constant, that is $f(n) = k$ for some constant $k \in \mathbb{Z}^+$ and for all $n \in \mathbb{Z}^+$. Then we have $g(n) = k$ for all $n \in \mathbb{Z}^+$ as well. That is, our main results implies the equality $\text{1playerFMD}(k) = \text{FMD}(k)$, and this equality is the lifting theorem of Bouyer et al.

It is instructive to consider an example when $f \notin \Omega(n)$ and is super-constant. Say, assume that $f(n) = O(n^\gamma)$ for some $\gamma < 1$. It is easy to see that then $g(n) = O(n^{\gamma/(1-\gamma)})$. Now there is a gap between memory complexity in one-player arenas and in all arenas. The closer γ is to 1, the larger is this gap.

When γ gets equal to 1, Theorem 3 becomes inapplicable. We demonstrate that this is not due to the weakness of our technique.

► **Theorem 4.** $\text{1playerFMD}(2n + 2) \not\subseteq \text{FMD}$.

This result shows the sharpness of Theorem 3. Namely, in order to obtain at least some bound on the memory complexity in all arenas, the memory complexity in one-player arenas should be a function not from $\Omega(n)$. In other words, it should be sublinear on some infinite subsequence. In turn, when it is already just linear, we might have no finite-memory determinacy.

Thus, our paper pushes the technique of one-to-two-player lifting to its limit. Unfortunately, this limit turns out to be very low. We are not aware of a payoff function which has been considered in the literature and to which one can apply Theorem 3, but which is not arena-independent finite-memory determined. For example, let us consider unbounded multidimensional energy games – as we have indicated, they are finite-memory determined but not arena-independently. As shown in [16], these games are in $\text{FMD}(n^{O(1)})$. Here the constant in $O(1)$ depends on the dimension and the maximum of the norms of the weights. In any case, this bound is not sufficient for Theorem 3.

Still, we provide an example of a payoff function to which our lifting theorem is applicable and the lifting theorem of Bouyer et al. is not.

► **Theorem 5.** *There exists a function $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$, $f \notin \Omega(n)$ and a payoff function from $1\text{playerFMD}(f)$ which is not arena-independent finite-memory determined.*

1.3 Other related works and concluding remarks

First, the exact analogs of the theorems of Gimbert and Zielonka and Bouyer et al. for *stochastic* games were obtained in other works of these authors [14, 5]. We find it plausible that our result can be lifted to stochastic games as well. Le Roux and Pauly [19] obtained a *two-to-many-players* lifting theorem. Namely, they show that, under some conditions, two-player finite-memory determinacy implies that all multiplayer games have finite-memory Nash equilibrium. A different approach to study finite-memory determinacy can be found in [20].

A natural open question is to extend lifting theorems to strategies with non-chromatic finite memory. As we mentioned, Le Roux [18] has shown that non-chromatic finite memory can always be replaced by the chromatic one. Unfortunately, this transformation is rather costly – the size of the memory grows exponentially in the number of nodes. So even the following modest question seems to be open: is there a payoff function which has constant non-chromatic memory complexity in one-player games but is not finite-memory determined in two-player games?

Organization of the paper. In Section 2 we give preliminaries. In Section 3 we give brief overviews of the proofs of our results. The full proof of Theorem 3 is given in the Appendix B. The full proofs of Theorems 4 and 5 can be found in the arXiv version of this paper [17].

2 Preliminaries

Notation. We denote the set of positive integer numbers by \mathbb{Z}^+ . Given a set A , by A^* and A^ω we denote the sets of finite and, respectively, infinite sequences of elements of A . The length of a sequence $x \in A^* \cup A^\omega$ is denoted by $|x|$. We write $A = B \sqcup C$ for three sets A, B, C if $A = B \cup C$ and $B \cap C = \emptyset$. Function composition is denoted by \circ .

2.1 Arenas

Following previous papers [13, 14, 4, 5], we call graphs on which our games are played *arenas*. We start with some notation regarding arenas. First, take an arbitrary finite set C . We will refer to the elements of C as *colors*. Informally, an arena is just a directed graph with edges colored by elements of C and with nodes partitioned into two sets.

► **Definition 6.** *A tuple $\mathcal{A} = \langle V, V_{\text{Max}}, V_{\text{Min}}, E, \text{source}, \text{target}, \text{col} \rangle$, where*

- $V, V_{\text{Max}}, V_{\text{Min}}, E$ are four finite sets with $V = V_{\text{Max}} \sqcup V_{\text{Min}}$;
- $\text{source}, \text{target}, \text{col}$ are functions of the form $\text{source}: E \rightarrow V, \text{target}: E \rightarrow V, \text{col}: E \rightarrow C$;

*is called an **arena** if for every $v \in V$ there exists $e \in E$ with $v = \text{source}(e)$.*

Elements of V will be called **nodes** of \mathcal{A} and elements of E will be called **edges** of \mathcal{A} . We understand $e \in E$ as a directed edge from the node $\text{source}(e)$ to the node $\text{target}(e)$. There might be parallel edges and loops. Additionally, every edge e of \mathcal{A} is labeled by the color $\text{col}(e) \in C$. Nodes from V_{Max} will be called nodes of Max and nodes from V_{Min} will be called

nodes of Min. The out-degree of a node $v \in V$ is $|\{e \in E \mid \text{source}(e) = v\}|$. By definition, every node in every arena has positive out-degree. An arena is called **one-player** if either all nodes of Max have out-degree 1 or all nodes of Min have out-degree 1.

Fix an arena $\mathcal{A} = \langle V, V_{\text{Max}}, V_{\text{Min}}, E, \text{source}, \text{target}, \text{col} \rangle$. We extend the function col (which determines the coloring of the edges) to arbitrary sequences of edges by setting: $\text{col}(e_1 e_2 e_3 \dots) = \text{col}(e_1) \text{col}(e_2) \text{col}(e_3) \dots$ for $e_1, e_2, e_3, \dots \in E$.

A non-empty sequence of edges $h = e_1 e_2 e_3 \dots \in E^* \cup E^\omega$ is called a **path** if for every $1 \leq n < |h|$ we have $\text{target}(e_n) = \text{source}(e_{n+1})$. We define $\text{source}(h) = \text{source}(e_1)$. When h is finite, we define $\text{target}(h) = \text{target}(e_{|h|})$. In addition, for every $v \in V$ we consider a 0-length path λ_v identified with v , for which we set $\text{source}(\lambda_v) = \text{target}(\lambda_v) = v$. For every $v \in V$ we define $\text{col}(\lambda_v)$ as the empty string.

2.2 Infinite-duration games on arenas

An arena $\mathcal{A} = \langle V, V_{\text{Max}}, V_{\text{Min}}, E, \text{source}, \text{target}, \text{col} \rangle$ induces an infinite-duration two-player game in the following way. First, we call players of this game Max and Min. Informally, Max and Min interact by gradually constructing a longer and longer path in \mathcal{A} . In each turn one of the players extends a current path by some edge from its endpoint. Which of the two players is the one to move is determined by whether this endpoint belongs to V_{Max} or to V_{Min} .

Formally, positions in the game are finite paths in \mathcal{A} . By definition, $\text{target}(h) \in V_{\text{Max}}$ for a finite path h means that Max is the one to move in the position h ; respectively, $\text{target}(h) \in V_{\text{Min}}$ means that Min is the one to move in the position h . A set of moves available in a position h is the set $\{e \in E \mid \text{source}(e) = \text{target}(h)\}$. Making a move $e \in E$ in a position $h = e_1 e_2 \dots e_{|h|}$ brings to a position $he = e_1 e_2 \dots e_{|h|} e$.

We stress that no position is designated as the initial one. We assume that the game can start in any position of the form $\lambda_v, v \in V$, at our choice.

Next we proceed to a notion of strategies. Namely, a strategy of Max is a function $\sigma: \{h \mid h \text{ is a finite path in } \mathcal{A} \text{ with } \text{target}(h) \in V_{\text{Max}}\} \rightarrow E$ such that for every h from the domain of σ we have $\text{source}(\sigma(h)) = \text{target}(h)$. Respectively, a strategy of Min is a function $\tau: \{h \mid h \text{ is a finite path in } \mathcal{A} \text{ with } \text{target}(h) \in V_{\text{Min}}\} \rightarrow E$ such that for every h from the domain of τ we have $\text{source}(\tau(h)) = \text{target}(h)$.

Observe that if \mathcal{A} is one-player, then one of the players has exactly one strategy. For technical consistency we assume that even when one of the players owns all the nodes of \mathcal{A} , the other player still has one “empty” strategy.

A strategy induces a set of positions *consistent* with it (those that can be reached in a play against this strategy). Formally, a finite path $h = e_1 e_2 \dots e_{|h|}$ is consistent with a strategy σ of Max if the following conditions hold:

- $\text{source}(h) \in V_{\text{Max}} \implies \sigma(\lambda_{\text{source}(h)}) = e_1$;
- for every $1 \leq i < |h|$ we have $\text{target}(e_1 e_2 \dots e_i) \in V_{\text{Max}} \implies \sigma(e_1 e_2 \dots e_i) = e_{i+1}$.

Consistency with the strategies of Min is defined similarly. Further, the notion of consistency can be extended to infinite paths. Namely, given a strategy, an infinite path is consistent with it if all finite prefixes of this path are.

For $v \in V$ and for a strategy \mathcal{S} of one of the players $\text{Cons}(v, \mathcal{S})$ denotes the set of all finite and infinite paths that start at v and are consistent with \mathcal{S} . For any strategy σ of Max, strategy τ of Min and $v \in V$, there is a unique infinite path in the intersection $\text{Cons}(v, \sigma) \cap \text{Cons}(v, \tau)$. We denote this path by $h(v, \sigma, \tau)$ and call it *the play* of σ and τ from v .

2.3 Payoff functions and equilibria

We consider only zero-sum games; correspondingly, in our framework objectives of the players are always given by a *payoff function*. A payoff function is any function of the form $\varphi: C^\omega \rightarrow \mathcal{W}$, where (\mathcal{W}, \leq) is a linearly ordered set. Informally, the aim of Max is to play in a way which maximizes the payoff function (with respect to the ordering of \mathcal{W}) while the aim of Min is the opposite one. Technically, to get the value of the payoff function on a play (which is an infinite path in the underlying arena) we first apply the function col to this play; this gives us an infinite sequence of colors; in conclusion, we apply φ to the sequence of colors.

Any payoff function in a standard way induces a notion of an *equilibrium* of two strategies of the players (with respect to this payoff function). Let us first introduce a notion of an *optimal response*. Namely, take a strategy σ of Max and a strategy τ of Min. We say that σ is a **uniformly optimal response** to τ if for all $v \in V$ and for all infinite $h \in \text{Cons}(v, \tau)$ we have $\varphi \circ \text{col}(h(v, \sigma, \tau)) \geq \varphi \circ \text{col}(h)$. The inequality here, of course, is with respect to the ordering of \mathcal{W} . Similarly, we call τ a **uniformly optimal response** to σ if for all $v \in V$ and for all infinite $h \in \text{Cons}(v, \sigma)$ we have $\varphi \circ \text{col}(h(v, \sigma, \tau)) \leq \varphi \circ \text{col}(h)$. Next, we call a pair (σ, τ) a **uniform equilibrium** if σ and τ are uniformly optimal responses to each other.

► **Lemma 7.** *For any arena \mathcal{A} and for any payoff function φ , the set uniform equilibria in \mathcal{A} w.r.t. φ is a Cartesian product.*

Proof. See Appendix A. ◀

Strategies which belong to some uniform equilibrium will be called **uniformly optimal**.

► **Remark 8.** Each payoff function induces a total preorder on C^ω . Two payoff functions that induce the same preorder have the same set of equilibria. Due to this reason, previous papers in this line of work [13, 14, 4, 5] do not consider payoff functions at all. Instead, they directly consider total preorders on C^ω , to which they refer as *preference relations*. We prefer to use a terminology of payoff functions, as it is more standard. Of course, this does not make our results less general – any preference relation is induced by some payoff function.

2.4 Positional strategies and finite-memory strategies

Positional strategies. A strategy \mathcal{S} of one of the players is called positional if for any two positions h_1, h_2 from its domain we have $\text{target}(h_1) = \text{target}(h_2) \implies \mathcal{S}(h_1) = \mathcal{S}(h_2)$. In other words, $\mathcal{S}(h)$ depends solely on $\text{target}(h)$. It makes convenient to consider positional strategies as functions on the set of nodes of the corresponding players (rather than on the set of the positions of this player). I.e., positional strategies of Max can be identified with functions of the form $\sigma: V_{\text{Max}} \rightarrow E$ such that $\text{source}(\sigma(v)) = v$ for all $v \in V_{\text{Max}}$. Similarly, positional strategies of Min can be identified with functions of the form $\tau: V_{\text{Min}} \rightarrow E$ such that $\text{source}(\tau(v)) = v$ for all $v \in V_{\text{Min}}$.

Let us fix some notation regarding positional strategies. First, every edge $e \in E$ is a path (of length 1) and hence also a position in the game induced by \mathcal{A} . If \mathcal{S} is a positional strategy of one of the players, we let $E_{\mathcal{S}}$ be the set of edges that are consistent with \mathcal{S} . Observe the following feature of positional strategies: the set of paths (positions) that are consistent with a positional strategy \mathcal{S} is exactly the set of paths that consist only of edges from $E_{\mathcal{S}}$.

Given a positional strategy \mathcal{S} of one of the players, by $\mathcal{A}_{\mathcal{S}}$ we denote the arena $\mathcal{A}_{\mathcal{S}} = \langle V, V_{\text{Max}}, V_{\text{Min}}, E_{\mathcal{S}}, \text{source}, \text{target}, \text{col} \rangle$. That is, $\mathcal{A}_{\mathcal{S}}$ is obtained from \mathcal{A} by deleting all edges that are inconsistent with \mathcal{S} . Observe that the arena $\mathcal{A}_{\mathcal{S}}$ is one-player; each node of the player who plays \mathcal{S} has exactly one out-going edge in $\mathcal{A}_{\mathcal{S}}$.

Instead of saying “an equilibrium of two positional strategies” we will simply say “a positional equilibrium”.

Finite-memory strategies. A memory skeleton is a deterministic finite automaton $\mathcal{M} = \langle M, m_{init} \in M, \delta: M \times C \rightarrow M \rangle$ whose input alphabet is the set C of colors. Here M is the set of states of \mathcal{M} , the state $m_{init} \in M$ is a designated initial state, and δ is the transition function of \mathcal{M} . By $|\mathcal{M}|$ we denote the number of states of a memory skeleton \mathcal{M} . Given $m \in M$, we extend $\delta(m, \cdot)$ to finite sequences of elements of C in a standard way. Now, a strategy \mathcal{S} of one of the players is called an \mathcal{M} -strategy if for any two positions h_1 and h_2 from the domain of \mathcal{S} it holds that

$$[\text{target}(h_1) = \text{target}(h_2) \text{ and } \delta(m_{init}, \text{col}(h_1)) = \delta(m_{init}, \text{col}(h_2))] \implies \mathcal{S}(h_1) = \mathcal{S}(h_2).$$

In other words, $\mathcal{S}(h)$ depends solely on $\text{target}(h)$ (the node with the token in the position h) and $\delta(m_{init}, \text{col}(h))$ (the state into which \mathcal{M} comes after reading the sequence of colors along h).

A strategy \mathcal{S} of one of the players is called a finite-memory strategy if it is an \mathcal{M} -strategy for some memory skeleton \mathcal{M} . Instead of saying “an equilibrium of two finite-memory strategies” or “an equilibrium of two \mathcal{M} -strategies” we will simply say “a finite-memory equilibrium” and “an \mathcal{M} -strategy equilibrium”.

2.5 Determinacy and memory complexity

- **Definition 9.** Let \mathcal{C} be a class of strategies. We say that a payoff function φ is determined in \mathcal{C} if every arena has a uniform equilibrium of two strategies from \mathcal{C} w.r.t. φ . In particular,
- if \mathcal{C} is the class of positional strategies, then we call φ **positionally determined**.
 - if \mathcal{C} is the class of finite-memory strategies, then we call φ **finite-memory determined**.
 - if \mathcal{C} is the class of \mathcal{M} -strategies for some memory skeleton \mathcal{M} , then we call φ **arena-independent finite-memory determined**.

For our results it is important that we require equilibria to be uniform in these definitions. That is, it is important to have a single pair of strategies from \mathcal{C} which is an equilibrium no matter in which node the game starts. As far as we know, this is the case for all positionally and finite-memory determined payoff functions that have been considered in the literature.

Next we provide definitions regarding the memory complexity.

- **Definition 10.** Let FMD denote the class of functions $\varphi: C^\omega \rightarrow \mathcal{W}$ such that C is a finite set, \mathcal{W} is linearly ordered and φ is finite-memory determined. Let 1playerFMD denote the class of functions $\varphi: C^\omega \rightarrow \mathcal{W}$ such that C is a finite set, \mathcal{W} is a linearly ordered set and such that the following holds: every one-player arena (with edges colored by elements of C) has a uniform finite-memory equilibrium w.r.t. φ .

Next, consider any function $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$. Let $\text{FMD}(f)$ denote the class of functions $\varphi: C^\omega \rightarrow \mathcal{W}$ such that C is a finite set, \mathcal{W} is a linearly ordered set and such that the following holds: for all $n \in \mathbb{Z}^+$ there exists a memory skeleton \mathcal{M} over the set C with $|\mathcal{M}| \leq f(n)$ such that all arenas (with edges colored by elements of C) with at most n nodes have a uniform \mathcal{M} -strategy equilibrium w.r.t. φ . Similarly, let $\text{1playerFMD}(f)$ denote the class of functions $\varphi: C^\omega \rightarrow \mathcal{W}$ such that C is a finite set, \mathcal{W} is a linearly ordered set and such that the following holds: for all $n \in \mathbb{Z}^+$ there exists a memory skeleton \mathcal{M} over the set C with $|\mathcal{M}| \leq f(n)$ such that all one-player arenas (with edges colored by elements of C) with at most n nodes have a uniform \mathcal{M} -strategy equilibrium w.r.t. φ .

3 Overviews of the Proofs

3.1 Theorem 3

First, let us give the exact statement of the lifting theorem of Bouyer et al.

► **Theorem 11** ([4]). *For any payoff function φ and for any memory skeleton \mathcal{M} the following holds. Assume that all one-player arenas have a uniform \mathcal{M} -strategy equilibrium w.r.t. φ . Then all arenas have a uniform \mathcal{M} -strategy equilibrium w.r.t. φ .*

Our main technical contribution is the following strengthening of Theorem 11.

► **Theorem 12.** *For any payoff function φ and for any $n \in \mathbb{Z}^+$ the following holds. Let \mathcal{M} be a memory skeleton such that all one-player arenas with at most $2n \cdot |\mathcal{M}| - 1$ nodes have a uniform \mathcal{M} -strategy equilibrium w.r.t. φ . Then all arenas with at most n nodes have a uniform \mathcal{M} -strategy equilibrium w.r.t. φ .*

Derivation of Theorem 3 from Theorem 12. Take any $\varphi \in \text{1playerFMD}(f)$. Our goal is to show that $\varphi \in \text{FMD}(g)$, where g is as in Theorem 3. That is, our goal is to establish for every $n \in \mathbb{Z}^+$ a memory skeleton \mathcal{M} with at most $g(n)$ states such that all arenas with at most n nodes have a uniform \mathcal{M} -strategy equilibrium.

Take any $n \in \mathbb{Z}^+$. By definition, $g(n) = f(m)$ for some $m \in \mathbb{Z}$ such that $\frac{f(m)}{m+1} \leq \frac{1}{2n}$. Since $\varphi \in \text{1playerFMD}(f)$, there exists a memory skeleton \mathcal{M} with at most $f(m)$ states such that all one-player arenas with at most m nodes have a uniform \mathcal{M} -strategy equilibrium. Now, since $\frac{f(m)}{m+1} \leq \frac{1}{2n}$, we have $m \geq 2n \cdot f(m) - 1 \geq 2n \cdot |\mathcal{M}| - 1$. By Theorem 12, this means that all arenas with at most n nodes have a uniform \mathcal{M} -strategy equilibrium. Since \mathcal{M} has at most $f(m) = g(n)$ states, we are done. ◀

Before discussing our technique, let us briefly overview how Bouyer et al. establish Theorem 11. They start by defining “ \mathcal{M} -monotone payoff functions” and “ \mathcal{M} -selective payoff functions”. Then they show that any payoff function which is \mathcal{M} -monotone and \mathcal{M} -selective is determined in \mathcal{M} -strategies. Finally, they show that for any non- \mathcal{M} -monotone and for any non- \mathcal{M} -selective payoff function there exists a one-player arena which has no uniform \mathcal{M} -strategy equilibrium w.r.t. this payoff function. This also gives a *characterization* of \mathcal{M} -determinacy: a payoff function is determined in \mathcal{M} -strategies if and only if it is \mathcal{M} -monotone and \mathcal{M} -selective.

In this paper, we obtain Theorem 12 (and, thus, Theorem 11) more directly. For the sake of simplicity, in Section 4 we prove it in a special case when \mathcal{M} is a memory skeleton with just one state. In this special case, \mathcal{M} -strategies are positional strategies.

► **Proposition 13** (Special case of Theorem 12). *For any payoff function φ and for any $N \in \mathbb{Z}^+$ the following holds. Assume that all one-player arenas with at most $2N - 1$ nodes have a uniform positional equilibrium w.r.t. φ . Then all arenas with at most N nodes have a uniform positional equilibrium w.r.t. φ .*

As all papers in this line of works, we build upon the inductive technique first invented by Gimbert and Zielonka [13]. Our contribution here is a more direct exposition of this technique, with the emphasis on the size of arenas.

We extend Proposition 13 to all memory skeletons³ in two steps. We first prove an analogue of Proposition 13 for so-called \mathcal{M} -trivial arenas. Informally, these are arenas where states of \mathcal{M} are “hardwired” into nodes. In such arenas, \mathcal{M} -strategies degenerate to positional strategies. We show that Proposition 13 is true even when only \mathcal{M} -trivial arenas are taken into account (in the assumption and in the conclusion).

We then derive Theorem 12 from this using the product arena construction [2, Chapter 2]. Take any (two-player) arena \mathcal{A} with up to n nodes. We have to derive the existence of an \mathcal{M} -strategy equilibrium in \mathcal{A} from the assumption of Theorem 12. It is a classical observation that \mathcal{M} -strategies in \mathcal{A} can be viewed as positional strategies in the product arena $\mathcal{M} \times \mathcal{A}$. This product arena is obtained by first pairing states of \mathcal{M} with nodes of \mathcal{A} , and then by drawing edges of \mathcal{A} in all possible ways that are consistent with the transition function of \mathcal{M} . Now we only have to establish a positional equilibrium in $\mathcal{M} \times \mathcal{A}$. This arena is \mathcal{M} -trivial, so we use Proposition 13 for \mathcal{M} -trivial arenas and $N = n \cdot |\mathcal{M}|$. The size of $\mathcal{M} \times \mathcal{A}$ is the product of the sizes of \mathcal{M} and \mathcal{A} , so it does not exceed N . It remains to show that all one-player \mathcal{M} -trivial arenas with up to $2N - 1 = 2n \cdot |\mathcal{M}| - 1$ nodes have a positional equilibrium. Indeed, by the assumption of Theorem 12, all one-player arenas (not only \mathcal{M} -trivial) of this size have an \mathcal{M} -strategy equilibrium. But in \mathcal{M} -trivial arenas these \mathcal{M} -strategy equilibria are automatically positional.

The full proof of Theorem 12 is given in Appendix B.

3.2 Theorem 4

Let the set of colors be $C = \{-1, 1\}$. We define a payoff function $\psi: C^\omega \rightarrow \{0, 1\}$ as follows. We set $\psi(c_1 c_2 c_3 \dots) = 1$ if and only if either $(\lim_{n \rightarrow \infty} \sum_{i=1}^n c_i = +\infty)$ or $(\sum_{i=1}^n c_i = 0 \text{ for infinitely many } n)$. We assume the standard ordering on $\{0, 1\} = \psi(C^\omega)$, so that 1 is interpreted as victory of Max and 0 is interpreted as victory of Min.

We show that $\psi \in \text{1playerFMD}(2n+2) \setminus \text{FMD}$. In fact, this payoff function was defined by Bouyer et al. in [4, Section 3.4]. They have shown that this payoff function is finite-memory determined in one-player arenas but not in two-player arenas. So our contribution here is an upper bound $\psi \in \text{1playerFMD}(2n+2)$ on its memory complexity in one-player arenas. In other words, for every n we provide a memory skeleton \mathcal{M}_n with $2n+2$ states such that every one-player arena \mathcal{A} with up to n nodes has a uniform \mathcal{M}_n -strategy equilibrium. Let us describe the main ideas needed to obtain this upper bound. In this overview, we only consider those one-player arenas where all nodes of Min have out-degree 1. We use similar ideas for one-player arenas of the opposite type (but they require a bit more care).

It will be more convenient to refer to the elements of C as *weights* rather than as colors. Correspondingly, by the weight of a path we will mean the sum of the weights of its edges. Further, we will call a path *positive* if its weight is positive. We define negative and zero paths similarly.

Take an arena with up to n nodes where all nodes of Min have out-degree 1 (that is, essentially Max is the one to move everywhere). First, we can remove all the nodes from where one can reach a positive cycle. Indeed, Max has a positional winning strategy from these nodes (Max can go to the closest simple positive cycle, and then start rotating over it forever). Here it is important that our arena is one-player. Two-player arenas might have positive cycles, but Max might be unable to stay on them.

³ Our technique in this part is rather similar to a technique from a recent paper of Bouyer et al. [5] (see the arXiv version [6] of their paper for more details). In this paper, they give a direct proof of an analogue of Theorem 11 for stochastic games.

Now the only way Max can win is by making the sum of the weights equal to 0 infinitely many times. As a first attempt, consider an “illegal” memory skeleton \mathcal{M}_∞ , which simply stores the sum of the weights along the current play. It is illegal since the sum of the weights can be arbitrarily large, so \mathcal{M}_∞ will have infinitely many states. Still, our winning condition for Max can be reformulated in terms of \mathcal{M}_∞ . Indeed, Max just has to bring \mathcal{M}_∞ into a state “the current sum is 0” infinitely many times. Notice that this is a parity condition in the product of our initial arena and the memory skeleton \mathcal{M}_∞ . Since parity games are positionally determined [23], we have a uniform positional equilibrium in the product arena, and this gives a uniform \mathcal{M}_∞ -strategy equilibrium in the initial arena.

To turn this idea into a proof, we “truncate” \mathcal{M}_∞ . For arenas with up to n nodes we consider a memory skeleton \mathcal{M}_n , which stores the current sum of the weights while its absolute value is at most n ; if it exceeds n , our memory skeleton comes into a special invalid state. Observe that such memory skeleton requires just $2n + 2$ states.

We now make use of the fact that w.l.o.g our arena has no positive cycles. Since our weights are ± 1 , there is no path of weight larger than n . Indeed, any path can be decomposed into cycles and a simple path. The contribution of cycles is non-negative, and the contribution of a simple path is at most n , just because its length is at most n . So the current sum of the weights can never become larger than n . It can become smaller than $-n$, and in this case Max loses (he can never make it equal to 0 again). So the goal of Max is, first, to avoid a state “the current sum exceeded n in the absolute value”, and second, to reach a state “the current sum is 0” infinitely many times. This is a parity condition in the product of our initial arena and the memory skeleton \mathcal{M}_n . Therefore, we get a uniform \mathcal{M}_n -strategy equilibrium in our initial arena.

3.3 Theorem 5

Let the set of colors be $C = \{0, 1\}$. Fix a set $T \subseteq \mathbb{Z}^+$. Define a payoff function $\varphi: \{0, 1\}^\omega \rightarrow \{0, 1\}$ by setting $\varphi(\alpha) = 1$ for $\alpha = \alpha_1\alpha_2\alpha_3\dots \in \{0, 1\}^\omega$ if and only if at least one of the following two conditions holds:

- α contains only finitely many 0’s;
- for some $t \in T$, the sequence α contains the word 01^t0 .

We show that, under some condition on T , the payoff function φ is not arena-independent finite-memory determined, but belongs to $\mathbf{1playerFMD}(f)$ for some $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+, f \notin \Omega(n)$. This condition is called *isolation*. Roughly speaking, it requires that there are infinitely many elements in T such that far to the left and to the right of them there are no other elements of T . More precisely, $T \subseteq \mathbb{Z}^+$ is isolated if there are infinitely many $k \in T$ such that $l \notin T$ for all $k/2 < l < k^4, l \neq k$. We call such k *isolated elements* of T .

From now on, we fix any isolated set T , for example, $T = \{2^{4^n} \mid n \in \mathbb{Z}^+\}$. To show that $\varphi \in \mathbf{1playerFMD}(f)$ for some $f: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+, f \notin \Omega(n)$, we construct, for every k , the following memory skeleton \mathcal{M}_k . It simply counts the number of 1’s after the last 0. If this number exceeds k , it stops counting (it just remembers a fact that there are more than k ones after the last 0). Now, when our memory skeleton receives a 0, there are two cases. If the current value of the counter is some number from $T \cap [1, k]$, then \mathcal{M}_k transits into a special “winning state”, and stays in it forever. Otherwise, it resets the counter to 0 and starts counting again.

Note that \mathcal{M}_k can be realized with $k + O(1)$ states. We show that if k is an isolated element of T , then all arenas (even two-player) with up to k^2 nodes have a uniform \mathcal{M}_k -strategy equilibrium. This will show that $\varphi \in \mathbf{FMD}(f)$ for some function f such that $f(n) \leq 2\sqrt{n}$ for infinitely many n .

Consider any arena with up to k^2 nodes. We define an auxiliary game in which Max wins if either \mathcal{M}_k was brought to the “winning state” or there were just finitely many 0’s in the play. Note that if Max wins in the auxiliary game, then Max wins w.r.t. φ . The auxiliary game, however, is not entirely equivalent to φ , because a play can be winning for Min in the auxiliary game but loosing for Min w.r.t. φ (if this play contains 01^t0 for some $t \in T, t > k$). Still, it holds that if Min can win in the auxiliary game, then Min can also win w.r.t. φ . To prove this claim, we notice that the auxiliary game is a parity game in the product of our initial arena and the memory skeleton \mathcal{M}_k . So if Min can win in it, then Min can do so via some positional strategy τ in the product arena. We observe that τ is also winning w.r.t. φ . Indeed, otherwise there is a play against τ which contains 01^t0 for some $t \in T, t > k$. Since k is an isolated element of T , we have $t \geq k^4$. Therefore, as the size of the product arena is $(k + O(1)) \cdot k^2 < k^4$, there must be a cycle which is consistent with τ and which consists entirely of 1’s. But then Max can win against τ in the auxiliary game, contradiction.

As we pointed out, the auxiliary game is a parity game in the product of our arena with \mathcal{M}_k . Thus, it has a positional equilibrium there. This positional equilibrium translates into an \mathcal{M}_k -strategy equilibrium in the initial arena. Finally, as shown in the previous paragraph, any equilibrium in the auxiliary game is also an equilibrium w.r.t. φ .

Showing that φ is not arena-independent finite-memory determined is much easier. Take an isolated element $k \in T$. The idea is to construct an arena with a node which “cuts” the word 01^k0 in $\Omega(k)$ different ways near the middle. Due to isolation, the only way for Max to win in this arena is to go through one of the cuts. However, Min can choose any of the cuts, so Max needs $\Omega(k)$ states to distinguish between different cuts. Since k can be arbitrarily large, this shows that no single memory skeleton can be sufficient for φ in all arenas.

4 Proof of Proposition 13

The proof is by induction on the number of edges of an arena. More precisely, we are proving by induction on m the following claim: for every m every arena with m edges and at most N nodes has a uniform positional equilibrium.

The induction base ($m = 1$) is trivial (any arena with one edge is one-player and has exactly one node, so we can just refer to the assumption of the lemma). We proceed to the induction step. Take an arena $\mathcal{A} = \langle V, V_{\text{Max}}, V_{\text{Min}}, E, \text{source}, \text{target}, \text{col} \rangle$ with at most N nodes and assume that all arenas with at most N nodes and with fewer edges than \mathcal{A} have a uniform positional equilibrium. We prove the same for \mathcal{A} . Since the set of uniform equilibria is a Cartesian product by Lemma 7, it is enough to establish the following two claims:

- **(a)** in \mathcal{A} there exists a uniform equilibrium including a positional strategy of Max;
- **(b)** in \mathcal{A} there exists a uniform equilibrium including a positional strategy of Min.

We only show **(a)**, a proof of **(b)** is similar.

We may assume that \mathcal{A} is not one-player (otherwise we are done due to the assumptions of the lemma). Hence there exists a node $w \in V_{\text{Max}}$ with out-degree at least 2. Partition the set $E(w) = \{e \in E \mid \text{source}(e) = w\}$ into two non-empty disjoint subsets $E_1(w)$ and $E_2(w)$. Define two new arenas \mathcal{A}_1 and \mathcal{A}_2 . The arena \mathcal{A}_1 is obtained from \mathcal{A} by deleting edges from the set $E_2(w)$. Similarly, the arena \mathcal{A}_2 is obtained from \mathcal{A} by deleting edges from the set $E_1(w)$. So in \mathcal{A}_i for $i = 1, 2$ the set of edges with the source in w is $E_i(w)$.


Both \mathcal{A}_1 and \mathcal{A}_2 have fewer edges than \mathcal{A} . So both these arenas have a uniform positional equilibrium. Let (σ_i, τ_i) be a uniform positional equilibrium in \mathcal{A}_i for $i = 1, 2$. We will first define two auxiliary strategies τ_{12} and τ_{21} of Min; then we will show that either (σ_1, τ_{12}) or (σ_2, τ_{21}) is a uniform equilibrium in \mathcal{A} . After that **(a)** will be proved.

Strategies τ_{12} and τ_{21} will not be positional. In a sense, they are combinations of τ_1 and τ_2 . In both strategies Min has a counter I which can only take two values, 1 and 2. The counter I indicates to Min which of the strategies τ_1 or τ_2 to use. I.e., whenever Min should make a move from a node $v \in V_{\text{Min}}$, he uses an edge $\tau_I(v)$. The value of I changes each time in the node w Max uses an edge not from a set $E_I(w)$. It only remains to specify the initial value of I . There are two ways to do this, one will give us strategy τ_{12} , and the other will give τ_{21} . More specifically, in τ_{12} the initial value of I is 1 and in τ_{21} the initial value of I is 2.

It is not hard to see that τ_{12} is a uniformly optimal response to σ_1 and τ_{21} is a uniformly optimal response to σ_2 . For instance, let us show this for τ_{12} and σ_1 . By definition, τ_1 is a uniformly optimal response to σ_1 in the arena \mathcal{A}_1 , and hence also in the arena \mathcal{A} (because any play against σ takes place inside \mathcal{A}_1). It remains to notice that τ_{12} plays exactly as τ_1 against σ_1 . Indeed, σ_1 never uses edges from $E_2(w)$, so the counter I always equals 1 against σ_1 .

It remains to show that either σ_1 is a uniformly optimal response to τ_{12} or σ_2 is a uniformly optimal response to τ_{21} (in the arena \mathcal{A}). We derive it from the assumption of the lemma applied to an auxiliary one-player arena \mathcal{B} with at most $2N - 1$ nodes.

Namely, we define \mathcal{B} as follows. Recall that in our notation $(\mathcal{A}_1)_{\tau_1}$ and $(\mathcal{A}_2)_{\tau_2}$ stand for two arenas obtained from, respectively, \mathcal{A}_1 and \mathcal{A}_2 by throwing away edges that are inconsistent with, respectively, τ_1 and τ_2 . Consider an arena consisting of two “independent” parts one of which coincides with $(\mathcal{A}_1)_{\tau_1}$ and the other with $(\mathcal{A}_2)_{\tau_2}$ (“independent” means that there are no edges between the parts). From each part take a node corresponding to the node w . Then merge these two nodes into a single one. The resulting arena with $2|V| - 1 \leq 2N - 1$ nodes will be \mathcal{B} .



■ **Figure 1** Arena \mathcal{B} .

For each node of \mathcal{A} there are two “copies” of it in \mathcal{B} – one from $(\mathcal{A}_1)_{\tau_1}$ and the other from $(\mathcal{A}_2)_{\tau_2}$. We will call copies of the first kind *left copies* and copies of the second kind *right copies*. Note that the left and the right copy of w is the same node in \mathcal{B} . Any other node of \mathcal{A} has two distinct copies. Now, by the *prototype* of a node v' of \mathcal{B} we mean a node v of \mathcal{A} of which v' is a copy.

Note that in \mathcal{B} all nodes of Min have out-degree 1 (because they do so inside $(\mathcal{A}_1)_{\tau_1}$ and $(\mathcal{A}_2)_{\tau_2}$, and the only node of \mathcal{B} which was obtained by merging two nodes is a node of Max). Thus, \mathcal{B} is a one-player arena.

An important feature of \mathcal{B} is that it can “emulate” any play against τ_{12} and τ_{21} in \mathcal{A} . Formally,

► **Lemma 14.** *For any infinite path h in \mathcal{A} which is consistent with τ_{12} there exists an infinite path h' in \mathcal{B} with $\text{col}(h') = \text{col}(h)$ and with the source in the left copy of $\text{source}(h)$. Similarly, for any infinite path h in \mathcal{A} which is consistent with τ_{21} there exists an infinite path h' in \mathcal{B} with $\text{col}(h') = \text{col}(h)$ and with the source in the right copy of $\text{source}(h)$.*

Proof. We only give an argument for τ_{12} , the argument for τ_{21} is similar. We construct h' from the left copy of $\text{source}(h)$ by always moving in the same “local direction” as h . There will be no problem with that for the nodes of Max because they have the same set of out-going edges in \mathcal{B} as their prototypes have in \mathcal{A} . Now, for the nodes of Min we should be more accurate. The path h is consistent with τ_{12} , so from the nodes of Min it applies either τ_1 or τ_2 . Now, in \mathcal{B} strategy τ_1 is available only in the left ellipse of Figure 1, and τ_2 is available only in the right ellipse. So each time h wants to apply τ_1 , the path h' should be in the left ellipse. Similarly, each time h wants to apply τ_2 , the path h' should be in the right ellipse. Initially, until its counter changes, τ_{12} applies τ_1 , and correspondingly h' starts in the left ellipse. Now, each time τ_{12} switches to τ_2 , it does so because Max used an edge from $E_2(w)$ in w . Correspondingly, h' enters the right ellipse at this moment. Similarly, whenever τ_{12} switches back to τ_1 , the path h' returns to the left ellipse. ◀

Note that in \mathcal{B} Min has exactly one strategy. We denote it by T . The arena \mathcal{B} is one-player and has at most $2N - 1$ nodes, so by the assumption of the lemma there is a uniform positional equilibrium $(\widehat{\Sigma}, T)$ in it. We claim the following:

- if $\widehat{\Sigma}$ applies an edge from $E_1(w)$ in w , then σ_1 is a uniformly optimal response to τ_{12} in \mathcal{A} ;
- if $\widehat{\Sigma}$ applies an edge from $E_2(w)$ in w , then σ_2 is a uniformly optimal response to τ_{21} in \mathcal{A} .

We only show the first claim, the proof of the second one is analogous. Consider a restriction of $\widehat{\Sigma}$ to the left ellipse of \mathcal{B} . This defines a positional strategy $\widehat{\sigma}$ of Max in \mathcal{A} . Note that in each node of \mathcal{A} the strategy σ_1 is at least as good against τ_{12} as $\widehat{\sigma}$. Indeed, $\sigma_1(w), \widehat{\sigma}(w) \in E_1(w)$. Hence $\sigma_1, \widehat{\sigma}$ are strategies in the arena \mathcal{A}_1 , where σ_1 is a uniformly optimal response to τ_1 . It remains to notice that τ_{12} plays exactly as τ_1 against σ_1 and $\widehat{\sigma}$ since these two strategies of Max never use edges from $E_2(w)$.

Therefore, it is enough to show that $\widehat{\sigma}$ is a uniformly optimal response to τ_{12} in \mathcal{A} . Take any node $v \in V$ and any play h against τ_{12} from v . Our goal is to show that the play of $\widehat{\sigma}$ and τ_{12} from v is at least as good from the Max’s perspective as h . Now, by Lemma 14 some infinite path h' from the left copy of v is colored exactly as h . On the other hand, the play of $\widehat{\Sigma}$ and T from the left copy of v is at least as good for Max as h' (and hence as h). This is because h' is consistent with T (as there are simply no other strategies of Min in \mathcal{B}) and because $(\widehat{\Sigma}, T)$ is an equilibrium. It remains to note that the play of $\widehat{\Sigma}$ and T from the left copy of v is colored exactly as the play of $\widehat{\sigma}$ and τ_{12} from v . Indeed, as we have already observed, τ_{12} plays exactly as τ_1 against $\widehat{\sigma}$. On the other hand, the play of $\widehat{\Sigma}$ and T can never leave the left ellipse as $\widehat{\Sigma}$ points to the left in w . Moreover, restrictions of these strategies to the left ellipse coincide with $\widehat{\sigma}$ and τ_1 ; for $\widehat{\Sigma}$ this is just by definition and for T this is because the left ellipse coincides with the arena $(\mathcal{A}_1)_{\tau_1}$.

References

- 1 Paolo Baldan, Barbara König, Christina Mika-Michalski, and Tommaso Padoan. Fixpoint games on continuous lattices. *Proceedings of the ACM on Programming Languages*, 3(POPL):1–29, 2019.
- 2 Mikołaj Bojańczyk and Wojciech Czerwiński. An automata toolbox. A book of lecture notes, available at <https://www.mimuw.edu.pl/~bojan/upload/reduced-may-25.pdf>, 2018.
- 3 Patricia Bouyer, Uli Fahrenberg, Kim G Larsen, Nicolas Markey, and Jiří Srba. Infinite runs in weighted timed automata with energy constraints. In *International Conference on Formal Modeling and Analysis of Timed Systems*, pages 33–47. Springer, 2008.

- 4 Patricia Bouyer, Stéphane Le Roux, Youssouf Oualhadj, Mickael Randour, and Pierre Vandenhove. Games where you can play optimally with arena-independent finite memory. In *31st International Conference on Concurrency Theory (CONCUR 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- 5 Patricia Bouyer, Youssouf Oualhadj, Mickael Randour, and Pierre Vandenhove. Arena-Independent Finite-Memory Determinacy in Stochastic Games. In Serge Haddad and Daniele Varacca, editors, *32nd International Conference on Concurrency Theory (CONCUR 2021)*, volume 203 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:18, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CONCUR.2021.26.
- 6 Patricia Bouyer, Youssouf Oualhadj, Mickael Randour, and Pierre Vandenhove. Arena-independent finite-memory determinacy in stochastic games. *arXiv preprint*, 2021. arXiv:2102.10104.
- 7 J Richard Büchi and Lawrence H Landweber. Solving sequential conditions by finite-state strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.
- 8 Krishnendu Chatterjee and Laurent Doyen. Energy parity games. *Theoretical Computer Science*, 458:49–60, 2012.
- 9 Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin. Strategy synthesis for multi-dimensional quantitative objectives. *Acta informatica*, 51(3-4):129–163, 2014.
- 10 Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979.
- 11 E Allen Emerson and Charanjit S Jutla. Tree automata, mu-calculus and determinacy. In *FoCS*, volume 91, pages 368–377. Citeseer, 1991.
- 12 E Allen Emerson, Charanjit S Jutla, and A Prasad Sistla. On model-checking for fragments of μ -calculus. In *International Conference on Computer Aided Verification*, pages 385–396. Springer, 1993.
- 13 Hugo Gimbert and Wiesław Zielonka. Games where you can play optimally without any memory. In *International Conference on Concurrency Theory*, pages 428–442. Springer, 2005.
- 14 Hugo Gimbert and Wiesław Zielonka. Pure and stationary optimal strategies in perfect-information stochastic games with global preferences. *arXiv preprint*, 2016. arXiv:1611.08487.
- 15 Erich Gradel and Wolfgang Thomas. *Automata, logics, and infinite games: a guide to current research*, volume 2500. Springer Science & Business Media, 2002.
- 16 Marcin Jurdziński, Ranko Lazić, and Sylvain Schmitz. Fixed-dimensional energy games are in pseudo-polynomial time. In *International Colloquium on Automata, Languages, and Programming*, pages 260–272. Springer, 2015.
- 17 Alexander Kozachinskiy. One-to-two-player lifting for mildly growing memory. *arXiv preprint*, 2021. arXiv:2104.13888.
- 18 Stéphane Le Roux. Time-aware uniformization of winning strategies. In *Conference on Computability in Europe*, pages 193–204. Springer, 2020.
- 19 Stéphane Le Roux and Arno Pauly. Extending finite-memory determinacy to multi-player games. *Information and Computation*, 261:676–694, 2018.
- 20 Stéphane Le Roux, Arno Pauly, and Mickael Randour. Extending finite-memory determinacy by boolean combination of winning conditions. In *38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, 2018.
- 21 A Mostowski. Games with forbidden positions. Technical report, Preprint No. 78, Uniwersytet Gdanski, Instytut Matematyki, 1991.
- 22 An A Muchnik. Games on infinite trees and automata with dead-ends: a new proof for the decidability of the monadic second order theory of two successors. *Bulletin-European Association For Theoretical Computer Science*, 48:219–219, 1992.
- 23 Wiesław Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1-2):135–183, 1998.
- 24 Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996.

A

 Uniform equilibria and T -wise equilibria

For the proof of Theorem 12 we need to generalize the notion of a uniform equilibrium. Take any arena $\mathcal{A} = \langle V, V_{\text{Max}}, V_{\text{Min}}, E, \text{source}, \text{target}, \text{col} \rangle$ and any payoff function $\varphi: C^\omega \rightarrow \mathcal{W}$. Fix a subset $T \subseteq V$, a strategy σ of Max and a strategy τ of Min. We say that σ is a **T -wise optimal response** to τ w.r.t. φ if for all $v \in T$ and for all infinite $h \in \text{Cons}(v, \tau)$ we have $\varphi \circ \text{col}(h(v, \sigma, \tau)) \geq \varphi \circ \text{col}(h)$. Similarly, we call τ a **T -wise optimal response** to σ w.r.t. φ if for all $v \in T$ and for all infinite $h \in \text{Cons}(v, \sigma)$ we have $\varphi \circ \text{col}(h(v, \sigma, \tau)) \leq \varphi \circ \text{col}(h)$. Finally, we call a pair (σ, τ) a **T -wise equilibrium** w.r.t. φ if σ and τ are T -wise optimal responses to each other.

When $T = V$ is the whole set of nodes, then T -wise equilibria are uniform equilibria, and vice versa. Thus, the following lemma generalizes Lemma 7.

► **Lemma 15.** *For any arena $\mathcal{A} = \langle V, V_{\text{Max}}, V_{\text{Min}}, E, \text{source}, \text{target}, \text{col} \rangle$, for any payoff function φ , and for any subset $T \subseteq V$, the set of T -wise equilibria in \mathcal{A} w.r.t. φ is a Cartesian product.*

Proof. It is sufficient to show the following: if (σ_1, τ_1) and (σ_2, τ_2) are T -wise equilibria, then so is (σ_1, τ_2) . That is, our goal is to show that σ_1 is a T -wise optimal response to τ_2 , and that τ_2 is a T -wise optimal response to σ_1 . We only prove the first claim, the second one can be proved similarly. Take any $v \in T$ and any infinite $h \in \text{Cons}(v, \tau_2)$. We have to show that $\varphi \circ \text{col}(h(v, \sigma_1, \tau_2)) \geq \varphi \circ \text{col}(h)$. We first show that $\varphi \circ \text{col}(h(v, \sigma_1, \tau_2)) = \varphi \circ \text{col}(h(v, \sigma_1, \tau_1)) = \varphi \circ \text{col}(h(v, \sigma_2, \tau_2))$. Indeed,

$$\begin{aligned} \varphi \circ \text{col}(h(v, \sigma_1, \tau_1)) &\geq \varphi \circ \text{col}(h(v, \sigma_2, \tau_1)) \geq \varphi \circ \text{col}(h(v, \sigma_2, \tau_2)) \\ &\geq \varphi \circ \text{col}(h(v, \sigma_1, \tau_2)) \geq \varphi \circ \text{col}(h(v, \sigma_1, \tau_1)). \end{aligned}$$

The first inequality here holds because σ_1 is a T -wise optimal response to τ_1 . The second inequality here holds because τ_2 is a T -wise optimal response to σ_2 . The third inequality here holds because σ_2 is a T -wise optimal response to τ_2 . The fourth inequality here holds because τ_1 is a T -wise optimal response to σ_1 .

As we have shown, $\varphi \circ \text{col}(h(v, \sigma_1, \tau_2)) = \varphi \circ \text{col}(h(v, \sigma_2, \tau_2))$. In turn, since $h \in \text{Cons}(v, \tau_2)$, and since σ_2 is a T -wise optimal response to τ_2 , we have that $\varphi \circ \text{col}(h(v, \sigma_2, \tau_2)) \geq \varphi \circ \text{col}(h)$. Therefore, we get $\varphi \circ \text{col}(h(v, \sigma_1, \tau_2)) \geq \varphi \circ \text{col}(h)$. ◀

B

 Proof of Theorem 12

We reduce Theorem 12 to a statement about positional strategies (namely, to Lemma 20 below). First we need a classical concept of *product arenas*.

► **Definition 16** (Product arenas). *Let $\mathcal{M} = \langle M, m_{\text{init}}, \delta: M \times C \rightarrow M \rangle$ be a memory skeleton and $\mathcal{A} = \langle V, V_{\text{Max}}, V_{\text{Min}}, E, \text{source}, \text{target}, \text{col} \rangle$ be an arena. Then $\mathcal{M} \times \mathcal{A}$ stands for an arena, where*

- the set of nodes is $M \times V$;
- the set of Max's nodes is $M \times V_{\text{Max}}$;
- the set of Min's nodes is $M \times V_{\text{Min}}$;
- the set of edges is $M \times E$;
- the source function is defined as follows: $\text{source}((m, e)) = (m, \text{source}(e))$;
- the target function is defined as follows: $\text{target}((m, e)) = (\delta(m, \text{col}(e)), \text{target}(e))$;
- the coloring function is defined as follows: $\text{col}((m, e)) = \text{col}(e)$.

The following is a standard observation that product arenas reduce finite-memory determinacy to positional determinacy.

► **Observation 17.** *Let $\mathcal{M} = \langle M, m_{init}, \delta: M \times C \rightarrow M \rangle$ be a memory skeleton and $\mathcal{A} = \langle V, V_{Max}, V_{Min}, E, source, target, col \rangle$ be an arena. Then for every $S \subseteq V$ the following holds: if $\mathcal{M} \times \mathcal{A}$ has an $(\{m_{init}\} \times S)$ -wise positional equilibrium, then \mathcal{A} has an S -wise \mathcal{M} -strategy equilibrium.*

Its full proof can be found in the arXiv version of this paper [17].

Next we introduce one more concept which we need for the reduction, namely, one of \mathcal{M} -triviality.

► **Definition 18.** *Let $\mathcal{M} = \langle M, m_{init}, \delta: M \times C \rightarrow M \rangle$ be a memory skeleton. A pair (\mathcal{A}, f) of an arena $\mathcal{A} = \langle V, V_{Max}, V_{Min}, E, source, target, col \rangle$ and a function $f: V \rightarrow M$ is called **\mathcal{M} -trivial** if for every $e \in E$ it holds that $\delta(f(source(e)), col(e)) = f(target(e))$.*

Informally, f is a mapping from \mathcal{A} to the transition graph of \mathcal{M} which takes into account the colors of the edges. Of course, there are arenas that belong to no \mathcal{M} -trivial pair. We observe that \mathcal{M} -strategies, in a sense, degenerate to positional ones in \mathcal{M} -trivial pairs.

► **Observation 19.** *Let $\mathcal{M} = \langle M, m_{init}, \delta: M \times C \rightarrow M \rangle$ be a memory skeleton. Then for every \mathcal{M} -trivial pair (\mathcal{A}, f) the following holds: if \mathcal{A} has a uniform \mathcal{M} -strategy equilibrium, then \mathcal{A} has an $f^{-1}(m_{init})$ -wise positional equilibrium.*

Proof. Note that for any finite path h in \mathcal{A} we have:

$$\delta(f(source(h)), col(h)) = f(target(h)).$$

Indeed, this holds by definition as long as h is a single edge; for longer h this can be easily proved by induction on $|h|$.

To show the observation, we simply show that any \mathcal{M} -strategy coincides with some positional one on all plays that start in the nodes of $f^{-1}(m_{init})$. Indeed, a move of an \mathcal{M} -strategy in a position h depends solely on $target(h)$ and $\delta(m_{init}, col(h))$. However, $\delta(m_{init}, col(h)) = \delta(f(source(h)), col(h)) = f(target(h))$ for all h with $source(h) \in f^{-1}(m_{init})$. In other words, for all such h a move of an \mathcal{M} -strategy in h is a function only of $target(h)$, as required. ◀

We are ready to formulate a statement about positional strategies to which we reduce Theorem 12.

► **Lemma 20.** *Let $\mathcal{M} = \langle M, m_{init}, \delta: M \times C \rightarrow M \rangle$ be a memory skeleton. Assume that for every \mathcal{M} -trivial pair (\mathcal{A}, f) such that \mathcal{A} is one-player and has at most $2N - 1$ nodes there exists an $f^{-1}(m_{init})$ -wise positional equilibrium in \mathcal{A} .*

Then for every \mathcal{M} -trivial pair (\mathcal{A}, f) such that \mathcal{A} has at most N nodes there exists an $f^{-1}(m_{init})$ -wise positional equilibrium in \mathcal{A} .

Derivation of Theorem 12 from Lemma 20. Let $\mathcal{A} = \langle V, V_{Max}, V_{Min}, E, source, target, col \rangle$ be an arena with at most n nodes. Our goal is to show that \mathcal{A} has a uniform \mathcal{M} -strategy equilibrium. By Observation 17, it is sufficient to show that the arena $\mathcal{M} \times \mathcal{A}$ has an $\{m_{init}\} \times V$ -wise positional equilibrium. It is easy to see that a pair $(\mathcal{M} \times \mathcal{A}, f)$, where

$$f: M \times V \rightarrow M, \quad f((m, v)) = m,$$

is an \mathcal{M} -trivial pair, by definition of $\mathcal{M} \times \mathcal{A}$. Observe that $\{m_{init}\} \times V = f^{-1}(m_{init})$, so we only have to show that $\mathcal{M} \times \mathcal{A}$ has an $f^{-1}(m_{init})$ -wise positional equilibrium. Since $\mathcal{M} \times \mathcal{A}$ has at most $|\mathcal{M}| \cdot n$ nodes, it remains to explain why the assumption of Lemma 20 holds for $N = |\mathcal{M}| \cdot n$.

By the assumption of Theorem 12 all one-player arenas with at most $2|\mathcal{M}| \cdot n - 1 = 2N - 1$ nodes have a uniform \mathcal{M} -strategy equilibrium. In particular, this applies to any one-player arena \mathcal{A}' with at most $2N - 1$ nodes belonging to some \mathcal{M} -trivial pair (\mathcal{A}', f) . By Observation 19 this means that all such \mathcal{A}' have a $f^{-1}(m_{init})$ -wise positional equilibrium, as required. \blacktriangleleft

Proof of Lemma 20. We use the same technique and terminology as in Section 4. We are now proving by induction on m the following claim: for every m and for every \mathcal{M} -trivial pair (\mathcal{A}, f) such that \mathcal{A} has m edges and at most N nodes there exists an $f^{-1}(m_{init})$ -wise positional equilibrium in \mathcal{A} .

Induction base ($m = 1$) again requires no argument, and we proceed to the induction step. Consider any \mathcal{M} -trivial pair (\mathcal{A}, f) , where $\mathcal{A} = \langle V, V_{\text{Max}}, V_{\text{Min}}, E, \text{source}, \text{target}, \text{col} \rangle$ has at most N nodes. Our goal is to show that \mathcal{A} has an $f^{-1}(m_{init})$ -wise positional equilibrium, provided that an analogous claim is already proved for all \mathcal{M} -trivial pairs (\mathcal{A}', f') in which \mathcal{A}' has at most N nodes and fewer edges than \mathcal{A} . Since the set of $f^{-1}(m_{init})$ -wise equilibria is a Cartesian product by Lemma 15, it is enough to establish the following two claims:

- (a) in \mathcal{A} there exists an $f^{-1}(m_{init})$ -wise equilibrium including a positional strategy of Max;
- (b) in \mathcal{A} there exists an $f^{-1}(m_{init})$ -wise equilibrium including a positional strategy of Min.

We only show (a), a proof of (b) is similar. As before, we may assume that \mathcal{A} is not one-player so that there exists a node $w \in V_{\text{Max}}$ with out-degree at least 2. We partition the set of its out-going edges into two disjoint non-empty sets $E_1(w)$ and $E_2(w)$. Then we define arenas \mathcal{A}_1 and \mathcal{A}_2 exactly as in Section 4. Since (\mathcal{A}, f) is an \mathcal{M} -trivial pair, then so are pairs (\mathcal{A}_1, f) and (\mathcal{A}_2, f) . Indeed, \mathcal{A}_1 and \mathcal{A}_2 were obtained by simply throwing away some edges of \mathcal{A} . The remaining edges satisfy the definition of \mathcal{M} -triviality with respect to f just because they do so inside \mathcal{A} .

Note that \mathcal{A}_1 and \mathcal{A}_2 both have fewer edges than \mathcal{A} and at most as many nodes. So by the induction hypothesis both these arenas have an $f^{-1}(m_{init})$ -wise positional equilibrium. Let (σ_1, τ_1) be an $f^{-1}(m_{init})$ -wise positional equilibrium in \mathcal{A}_1 and (σ_2, τ_2) be an $f^{-1}(m_{init})$ -wise positional equilibrium in \mathcal{A}_2 . Next, we define two auxiliary strategies τ_{12} and τ_{21} of Min exactly as in Section 4. Our goal is to show that either (σ_1, τ_{12}) is an $f^{-1}(m_{init})$ -wise equilibrium in \mathcal{A} or (σ_2, τ_{21}) is an $f^{-1}(m_{init})$ -wise equilibrium in \mathcal{A} .

By the same argument as in Section 4, we have that τ_{12} is an $f^{-1}(m_{init})$ -wise optimal response to σ_1 and τ_{21} is an $f^{-1}(m_{init})$ -wise optimal response to σ_2 . The main challenge is to show the opposite for at least one of the pairs (σ_1, τ_{12}) and (σ_2, τ_{21}) .

For that we define a one-player arena \mathcal{B} exactly as in Section 4 (see Figure 1). It has $2|V| - 1 \leq 2N - 1$ nodes. We will apply the assumption of Lemma 20 to \mathcal{B} . More precisely, this will be done for some \mathcal{M} -trivial pair which includes \mathcal{B} . For that we define the following mapping g from the set of nodes of \mathcal{B} to the set of states of \mathcal{M} . Namely, if v' is a node of \mathcal{B} , we set $g(v') = f(v)$, where v is the prototype of v' . Observe that (\mathcal{B}, g) is an \mathcal{M} -trivial pair. Indeed any edge of \mathcal{B} is between two nodes whose prototypes are connected in \mathcal{A} by an edge of the same color. Thus, by the assumption of Lemma 20, the arena \mathcal{B} has a $g^{-1}(m_{init})$ -wise positional equilibrium $(\hat{\Sigma}, T)$ (as before, in \mathcal{B} there are no strategies of Min other than T). It is sufficient to establish the following two claims:

- if $\widehat{\Sigma}$ applies an edge from $E_1(w)$ in w , then σ_1 is an $f^{-1}(m_{init})$ -wise optimal response to τ_{12} in \mathcal{A} ;
- if $\widehat{\Sigma}$ applies an edge from $E_2(w)$ in w , then σ_2 is an $f^{-1}(m_{init})$ -wise optimal response to τ_{21} in \mathcal{A} .

A key observation here is that $g^{-1}(m_{init})$ is the union of the left copies of the nodes from $f^{-1}(m_{init})$ and the right copies of the nodes of $f^{-1}(m_{init})$. In fact, for a proof of the first claim we only need a fact that $g^{-1}(m_{init})$ includes all the left copies of the nodes from $f^{-1}(m_{init})$. Correspondingly, only the right copies of $f^{-1}(m_{init})$ are relevant for a proof of the second claim.

We only show the first claim, the second one can be proved similarly. As in Section 4, the argument is carried out through a positional strategy $\widehat{\sigma}$ of Max in \mathcal{A} obtained by restricting $\widehat{\Sigma}$ to the left ellipse. First we observe that in any node from $f^{-1}(m_{init})$ the strategy σ_1 is at least as good against τ_{12} as $\widehat{\sigma}$. Indeed, both σ_1 and $\widehat{\sigma}$ are strategies in \mathcal{A}_1 whereas σ_1 is an $f^{-1}(m_{init})$ -wise optimal response to τ_1 in \mathcal{A}_1 by definition. On the other hand, τ_{12} plays against σ_1 and $\widehat{\sigma}$ exactly as τ_1 .

It remains to show that $\widehat{\sigma}$ is an optimal response to τ_{12} in any node from $f^{-1}(m_{init})$. This can be done by exactly the same argument as in the last paragraph of Section 4. A difference is that now we have a weaker assumption about $\widehat{\Sigma}$; namely, we only know that $\widehat{\Sigma}$ is optimal in the nodes from $g^{-1}(m_{init})$ (while before it was optimal everywhere in \mathcal{B}). Correspondingly, we are proving a weaker statement. Namely, instead of proving that $\widehat{\sigma}$ is an optimal response to τ_{12} everywhere in \mathcal{A} , we are only proving this for all $v \in f^{-1}(m_{init})$. It can be checked that in the argument for a specific v we only require optimality of $\widehat{\Sigma}$ in the left copy of v ; so if $v \in f^{-1}(m_{init})$, then its left copy is in $g^{-1}(m_{init})$ so that the argument still works. ◀