



Swansea University
Prifysgol Abertawe



Cronfa - Swansea University Open Access Repository

This is an author produced version of a paper published in:

Cronfa URL for this paper:

<http://cronfa.swan.ac.uk/Record/cronfa43771>

Conference contribution :

Basset, N., Jecker, I., Pauly, A., Raskin, J. & van den Bogaard, M. *Beyond Admissibility: Dominance Between Chains of Strategies*. Computer Science Logic (CSL).

<http://dx.doi.org/10.4230/LIPIcs.CSL.2018.10>

Licensed under Creative Commons License CC-BY

This item is brought to you by Swansea University. Any person downloading material is agreeing to abide by the terms of the repository licence. Copies of full text items may be used or reproduced in any format or medium, without prior permission for personal research or study, educational or non-commercial purposes only. The copyright for any work remains with the original author unless otherwise specified. The full-text must not be sold in any format or medium without the formal permission of the copyright holder.

Permission for multiple reproductions should be obtained from the original author.

Authors are personally responsible for adhering to copyright and publisher restrictions when uploading content to the repository.

<http://www.swansea.ac.uk/library/researchsupport/ris-support/>

Beyond Admissibility: Dominance Between Chains of Strategies


Nicolas Basset

Université Grenoble Alpes
Grenoble, France
bassetni@univ-grenoble-alpes.fr

Ismaël Jecker

Department of Computer Science, Université Libre de Bruxelles
Brussels, Belgium
ismael.jecker@ulb.ac.be

Arno Pauly

Department of Computer Science, Swansea University
Swansea, UK
arno.m.pauly@gmail.com
 <https://orcid.org/0000-0002-0173-3295>

Jean-François Raskin

Department of Computer Science, Université Libre de Bruxelles
Brussels, Belgium
jraskin@ulb.ac.be

Marie Van den Bogaard

Department of Computer Science, Université Libre de Bruxelles
Brussels, Belgium
marie.van.den.bogaard@ulb.ac.be

Abstract

Admissible strategies, i.e. those that are not dominated by any other strategy, are a typical rationality notion in game theory. In many classes of games this is justified by results showing that any strategy is admissible or dominated by an admissible strategy. However, in games played on finite graphs with quantitative objectives (as used for reactive synthesis), this is not the case.

We consider increasing chains of strategies instead to recover a satisfactory rationality notion based on dominance in such games. We start with some order-theoretic considerations establishing sufficient criteria for this to work. We then turn our attention to generalised safety/reachability games as a particular application. We propose the notion of maximal uniform chain as the desired dominance-based rationality concept in these games. Decidability of some fundamental questions about uniform chains is established.

2012 ACM Subject Classification Theory of computation → Solution concepts in game theory, Theory of computation → Automata extensions

Keywords and phrases dominated strategies, admissible strategies, games played on finite graphs, reactive synthesis, reachability games, safety games, cofinal, order theory

Digital Object Identifier 10.4230/LIPIcs.CSL.2018.10

Related Version A full version of the paper is available at [3], <https://arxiv.org/abs/1805.11608>.



© Nicolas Basset, Ismaël Jecker, Arno Pauly, Jean-François Raskin, and Marie Van den Bogaard; licensed under Creative Commons License CC-BY

27th EACSL Annual Conference on Computer Science Logic (CSL 2018).

Editors: Dan Ghica and Achim Jung; Article No. 10; pp. 10:1–10:22

Leibniz International Proceedings in Informatics



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

Funding This work was partially supported by the ERC inVEST (279499), the ARC project “Non-Zero Sum Game Graphs: Applications to Reactive Synthesis and Beyond” (Fédération Wallonie-Bruxelles), the EOS project “Verifying Learning Artificial Intelligence Systems” (FNRS-FWO), and the Belgian FNRS PDR project “Subgame perfection in graph games” (T.0088.18). J.-F. Raskin is Professeur Francqui de Recherche funded by the Francqui foundation.

Acknowledgements We thank Gilles Geeraerts and Guillermo A. Pérez for several fruitful discussions.

1 Introduction

The canonical model to formalize the reactive synthesis problem are two-player win/lose perfect information games played on finite (directed) graphs [22, 1]. In recent years, more general objectives and multiplayer games have been studied (see e.g. [17] or [7] and additional references therein). When moving beyond two-player win/lose games, the traditional solution concept of a *winning strategy* needs to be updated by another notion. The game-theoretic literature offers a variety of concepts of *rationality* to be considered as candidates.

The notion we focus on here is *admissibility*: roughly speaking, judging strategies according to this criterion allows to deem rational only strategies that are not *worse* than any other strategy (ie, that are not *dominated*). In this sense, admissible strategies represent maximal elements in the whole set of strategies available to a player. One attractive feature of admissibility, or more generally, dominance based rationality notions is that they work on the level of an individual agent. Unlike e.g. to justify Nash equilibria, no common rationality, shared knowledge or any other assumptions on the other players are needed to explain why a specific agent would avoid dominated strategies.

The study of admissibility in the context of games played on graphs was initiated by Berwanger in [4] and subsequently became an active research topic (e.g. [12, 9, 2, 8, 10], see related work below). In [4], Berwanger established in the context of perfect-information games with boolean objectives that admissibility is the *good* criterion for rationality: every strategy is either admissible or dominated by an admissible strategy.

Unfortunately, this fundamental property does not hold when one considers quantitative objectives. Indeed, as soon as there are three different possible payoffs, one can find instances of games where a strategy is neither dominated by an admissible strategy, nor admissible itself (see Example 1). This third payoff actually allows for the existence of infinite domination sequences of strategies, where each element of the sequence dominates its predecessor and is dominated by its successor in the chain. Consequently, no strategy in such a chain is admissible. However, it can be the case that no admissible strategy dominates the elements of the chain. In the absence of a *maximal element* above these strategies, one may ask why they should be discarded in the quest of a rational choice. They may indeed represent a type of behaviour that is rational but not captured by the admissibility criterion.

Our contributions. To formalize this behaviour, we study increasing chains of strategies (Definition 3). A chain is weakly dominated by some other chain, if every strategy in the first is below some strategy in the second. The question then arises whether every chain is below a maximal chain. Based on purely order-theoretic argument, a sufficient criterion is given in Theorem 11. However, Corollary 17 shows that our sufficient criterion does not apply to all games of interests. We can avoid the issue by restricting to some countable class of strategies, e.g. just the regular, computable or hyperarithmetical ones (Corollary 19).

We test the abstract notion in the concrete setting of generalised safety/reachability games (Definition 21). Based on the observation that the crucial behaviour captured by chains of strategies, but not by single strategies is *Repeat this action a large but finite number of times*, we introduce the notion of a parameterized automaton (Definition 28), which essentially has just this ability over the standard finite automata. We then show that any finite memory strategy is below a maximal chain or strategy realized by a parameterized automaton (Theorem 31).

Finally, we consider some algorithmic properties of chains and parameterized automata in generalised safety/reachability games. It is decidable in PTIME whether a parameterized automaton realizes a chain of strategies (Theorem 35). It is also decidable in PTIME whether the chain realized by one parameterized automaton dominates the chain realized by another (Theorem 36).

Most proofs are omitted in the paper due to space restrictions. The appendix contains a selection of those. For the full account, we refer to the arXiv version [3].

Related work. As mentioned above, the study of dominance and admissibility for games played on graphs was initiated by Berwanger in [4]. Faella analyzed several criteria for how a player should play a win/lose game on a finite graph that she cannot win, eventually settling on the notion of admissible strategy [15].

Admissibility in quantitative perfect-information sequential games played on graphs was studied in [9]. Concurrent games were considered in [2]. In [8], games with imperfect information, but boolean objectives were explored. The study of decision problems related to admissibility (as we do in Subsection 4.3) was advanced in [12]. The complexity of decision problems related to dominance in normal form games has received attention, see [21] for an overview. For the role of admissibility for synthesis, we refer to [10].

Our Subsection 3.1 involves an investigation of cofinal chains in certain partially ordered sets. A similar theme (but with a different focus) is present in [25].

2 Background

2.1 Games on finite graphs

A *turn-based multiplayer game* \mathcal{G} on a finite graph G is a tuple $\mathcal{G} = \langle P, G, (p_i)_{i \in P} \rangle$ where:

- P is the non-empty finite set of players of the game,
- $G = \langle V, E \rangle$ where the finite set V of vertices of G is equipped with a $|P|$ -partition $\uplus_{i \in P} V_i$, and $E \subseteq V \times V$ is the edge relation of G ,
- for each player i in P , p_i is a *payoff function* that associates to every infinite path in G a payoff in \mathbb{R} .

Outcomes and histories. An *outcome* ρ of \mathcal{G} is an infinite path in G , that is, an infinite sequence of vertices $\rho = (\rho_k)_{k \in \mathbb{N}} \in V^\omega$, where for all $k \in \mathbb{N}$, $(\rho_k, \rho_{k+1}) \in E$. The set of all possible outcomes in \mathcal{G} is denoted $\mathbf{Out}(\mathcal{G})$. A finite prefix of an outcome is called a *history*. The set of all histories in \mathcal{G} is denoted $\mathbf{Hist}(\mathcal{G})$. For an outcome $\rho = (\rho_k)_{k \in \mathbb{N}}$ and an integer ℓ , we denote by $\rho_{\leq \ell}$ the history $(\rho_k)_{0 \leq k \leq \ell}$. The *length* of the history $\rho_{\leq \ell}$, denoted $|\rho_{\leq \ell}|$ is $\ell + 1$. Given an outcome or a history ρ and a history h , we write $h \subseteq_{\text{pref}} \rho$ if h is a prefix of ρ , and we denote by $h^{-1} \cdot \rho$ the unique outcome (or history) such that $\rho = h.(h^{-1} \cdot \rho)$. Given an outcome ρ or a history h and $k \in \mathbb{N}$ (respectively $k < |h|$), we denote by ρ_k (respectively h_k) the $k + 1$ -th vertex of ρ (respectively of h). For a history h , we define the last vertex of h to be $\text{last}(h) := h_{|h|-1}$ and its first vertex $\text{first}(h) := h_0$. For a vertex $v \in V$, its set of *successors* is $E_v = \{v' \in V \mid (v, v') \in E\}$.


Strategy profiles and payoffs. A *strategy* of a player i is a function σ_i that associates to each history h such that $\text{last}(h) \in V_i$, a successor state $v \in E_{\text{last}(h)}$. A tuple of strategies $(\sigma_i)_{i \in P'}$ where $P' \subseteq P$, one for each player in P' is called a *profile* of strategies. Usually, we focus on a particular player i , thus, given a profile $(\sigma_i)_{i \in P}$, we write σ_{-i} to designate the collection of strategies of players in $P \setminus \{i\}$, and the complete profile is written (σ_i, σ_{-i}) . The set of all strategies of player i is denoted $\Sigma_i(\mathcal{G})$, while $\Sigma(\mathcal{G}) = \prod_{i \in P} \Sigma_i(\mathcal{G})$ is the set of all profiles of strategies in the game \mathcal{G} and $\Sigma_{-i}(\mathcal{G})$ is the set of all profiles of all players except Player i . As we consider games with perfect information and deterministic transitions, any complete profile $\sigma_P = (\sigma_i)_{i \in P}$ yields, from any history h , a unique outcome, denoted $\mathbf{Out}_h(\mathcal{G}, \sigma_P)$. Formally, $\mathbf{Out}_h(\mathcal{G}, \sigma_P)$ is the outcome ρ such that $\rho_{\leq |h|-1} = h$ and for all $k \geq |h| - 1$, for all $i \in P$, it holds that $\rho_{k+1} = \sigma_i(\rho_{\leq k})$ if $\rho_k \in V_i$. The set of outcomes (resp. histories) *compatible* with a strategy σ of player i after a history h is $\mathbf{Out}_h(\mathcal{G}, \sigma_i) = \{\rho \in \mathbf{Out}(\mathcal{G}) \mid \exists \sigma_{-i} \in \Sigma_{-i}(\mathcal{G}) \text{ such that } \rho = \mathbf{Out}_h(\mathcal{G}, (\sigma_i, \sigma_{-i}))\}$ (resp. $\mathbf{Hist}_h(\sigma) = \{h \in \mathbf{Hist}(\mathcal{G}) \mid \exists \rho \in \mathbf{Out}_h(\mathcal{G}, \sigma_i), n \in \mathbb{N} \text{ such that } h = \rho_{\leq n}\}$). Each outcome ρ yields a *payoff* $p_i(\rho)$ for each Player i . We denote with $p_i(h, \sigma, \tau)$ the payoff of a profile of strategies (σ, τ) after a history h .

Usually, we consider games instances such that players start to play at a fixed vertex. Thus, we call an *initialized game* a pair (\mathcal{G}, v_0) of a game \mathcal{G} and a vertex $v_0 \in V$. When the initial vertex v_0 is clear from context, we speak directly from \mathcal{G} , $\mathbf{Out}(\mathcal{G}, \sigma_P)$ and $p_i(\sigma_P)$ instead of (\mathcal{G}, v_0) , $\mathbf{Out}_{v_0}(\mathcal{G}, \sigma_P)$ and $p_i(v_0, \sigma_P)$.

Dominance relation. In order to compare different strategies of a player i in terms of payoffs, we rely on the notion of dominance between strategies: A strategy $\sigma \in \Sigma_i$ is *weakly dominated* by a strategy $\sigma' \in \Sigma_i$ at a history h compatible with σ and σ' , denoted $\sigma \preceq_h \sigma'$, if for every $\tau \in \Sigma_{-i}$, we have $p_i(h, \sigma, \tau) \leq p_i(h, \sigma', \tau)$. We say that σ is weakly dominated by σ' , denoted $\sigma \preceq \sigma'$ if $\sigma \preceq_{v_0} \sigma'$, where v_0 is the initial state of \mathcal{G} . A strategy $\sigma \in \Sigma_i$ is *dominated* by a strategy $\sigma' \in \Sigma_i$, at a history h compatible with σ and σ' , denoted $\sigma \prec_h \sigma'$, if $\sigma \preceq_h \sigma'$ and there exists $\tau \in \Sigma_{-i}$, such that $p_i(h, \sigma, \tau) < p_i(h, \sigma', \tau)$. We say that σ is dominated by σ' , denoted $\sigma \prec \sigma'$ if $\sigma \prec_{v_0} \sigma'$, where v_0 is the initial state of \mathcal{G} . Strategies that are not dominated by any other strategies are called *admissible*: A strategy $\sigma \in \Sigma_i$ is *admissible* (respectively *from* h) if $\sigma \not\prec \sigma'$ (resp. $\sigma \not\prec_h \sigma'$) for every $\sigma' \in \Sigma_i$.

Antagonistic and Cooperative Values. To study the rationality of different behaviours in a game \mathcal{G} , it is useful to be able to know, for a player i , a fixed strategy $\sigma \in \Sigma_i$ and any history h , the worst possible payoff Player i can obtain with σ from h (i.e., the payoff he will obtain assuming the other players play *antagonistically*), as well as the best possible payoff Player i can hope for with σ from h (i.e., the payoff he will obtain assuming the other players play *cooperatively*). The first value is called the *antagonistic value of the strategy* σ of Player i at history h in \mathcal{G} and the second value is called the *cooperative value of the strategy* σ of Player i at history h in \mathcal{G} . They are formally defined as $\text{aVal}_i(\mathcal{G}, h, \sigma) := \inf_{\tau \in \Sigma_{-i}} p_i(\mathbf{Out}_h(\sigma, \tau))$ and $\text{cVal}_i(\mathcal{G}, h, \sigma) := \sup_{\tau \in \Sigma_{-i}} p_i(\mathbf{Out}_h(\sigma, \tau))$.

Prior to any choice of strategy of Player i , we can define, for any history h , the *antagonistic value of h* for Player i as $\text{aVal}_i(\mathcal{G}, h) := \sup_{\sigma \in \Sigma_i} \text{aVal}_i(\mathcal{G}, h, \sigma)$ and the *cooperative value of h* for Player i as $\text{cVal}_i(\mathcal{G}, h) := \sup_{\sigma \in \Sigma_i} \text{cVal}_i(\mathcal{G}, h, \sigma)$. Furthermore, one can ask, from a history h , what is the maximal payoff one can obtain *while* ensuring the antagonistic value of h . Thus, we define the *antagonistic-cooperative value of h* for Player i as $\text{acVal}_i(\mathcal{G}, h) := \sup\{\text{cVal}_i(\mathcal{G}, h, \sigma) \mid \sigma \in \Sigma_i \text{ and } \text{aVal}_i(\mathcal{G}, h, \sigma) \geq \text{aVal}_i(\mathcal{G}, h)\}$. From now on, we will omit to precise \mathcal{G} when it is clear from the context.



■ **Figure 1** The *Help-me?*-game.

An initialized game (\mathcal{G}, v_0) is *well-formed for Player i* if, for every history $h \in \mathbf{Hist}_{v_0}(\mathcal{G})$, there exists a strategy $\sigma \in \Sigma_i$ such that $\text{aVal}_i(h, \sigma) = \text{aVal}(h)$, and a strategy $\sigma' \in \Sigma_i$ such that $\text{cVal}_i(h, \sigma') = \text{cVal}(h)$. In other words, at every history h , Player i has a strategy that ensures the payoff $\text{aVal}_i(h)$, and a strategy that allows the other players to cooperate to yield a payoff of $\text{cVal}_i(h)$.

In the following, we will always focus on the point of view of one player i , thus we will sometimes refer to him as the *protagonist* and assume it is the first player, while the other players $-i$ can be seen as a coalition and abstracted to a single player, that we will call the *antagonist*. Furthermore, we will omit the subscript i to refer to the protagonist when we use the notations $\text{aVal}_i, \text{cVal}_i, \text{acVal}_i, p_i$, etc..

► **Example 1.** Consider the game depicted in Figure 1. The protagonist owns the circle vertices. The payoffs are defined as follows for the protagonist :

$$p(\rho) = \begin{cases} 0 & \text{if } \rho = (v_0v_1)^\omega, \\ 1 & \text{if } \rho = (v_0v_1)^n v_0 \ell_1^\omega \text{ where } n \in \mathbb{N}, \\ 2 & \text{if } \rho = (v_0v_1)^n \ell_2^\omega \text{ where } n \in \mathbb{N}. \end{cases}$$

Let us first look at the possible behaviours of the protagonist in this game, when he makes no assumption on the payoff function of the antagonist. He can choose to be “optimistic” and opt to try (at least for some time, or forever) to go to v_1 in the hope that the antagonist will cooperate to bring him to ℓ_2 , or settle from the start and go directly to ℓ_1 , not counting on any help from the antagonist. We denote by s_k the strategy that prescribes to choose v_1 as the successor vertex at the first k visits of v_0 , and ℓ_1 at the $k + 1$ -th visit, while s_ω denotes the strategy that prescribes v_1 at every visit of v_0 .

Fix $k \in \mathbb{N}$. Then, $s_k \prec s_{k+1}$: Indeed, for all $\tau \in \Sigma_{-i}$, if $p(s_k, \tau) = 2$, then there exists $j \leq k$ such that $\tau((v_0v_1)^j) = \ell_2$. As s_k and s_{k+1} agree up to $(v_0v_1)^k q_0$, we have that $\mathbf{Out}(s_{k+1}, \tau) = (v_0v_1)^j \ell_2^\omega = \mathbf{Out}(s_k, \tau)$, thus $p(s_{k+1}, \tau) = 2$ as well. Furthermore, consider a strategy τ such that $\tau((v_0v_1)^j) = v_0$ for all $j \leq k$ and $\tau((v_0v_1)^{k+1}) = \ell_2$. Then $p(s_k, \tau) = 1$ while $p(s_{k+1}, \tau) = 2$. Finally, consider the strategy τ such that $\tau((v_0v_1)^k) = v_0$ for all $k \in \mathbb{N}$. Then $p(s_k, \tau) = 1 = p(s_{k+1}, \tau)$. Hence, $s_k \prec s_{k+1}$. In addition, we observe that s_ω is admissible: for any strategy s_k , the strategy τ of the antagonist that moves to ℓ_2 at the $k + 1$ -th visit of v_1 yields a payoff of 1 against strategy s_k but 2 against strategy s_ω . Thus, $s_\omega \not\prec s_k$ for any $k \in \mathbb{N}$.

Quantitative vs Boolean setting. Remark that in the boolean variant of the *Help-me?* game considered in Example 1, where the payoff associated with the vertex ℓ_1 is 0 and the payoff associated with the vertex ℓ_2 is 1, every strategy s_k for $k \in \mathbb{N}$ is in fact dominated by s_ω , as s_k and s_ω both yield payoff 0 against τ such that $\tau((v_0v_1)^k) = v_0$ for all $k \in \mathbb{N}$. In fact, Berwanger in [4], showed that boolean games with ω -regular objectives enjoy the following fundamental property: every strategy is either admissible, or dominated by an admissible strategy. The existence of an admissible strategy in any such game follows as an immediate corollary.

Let us now illustrate how admissibility fails to capture fully the notion of rational behaviour in the quantitative case. Firstly, recall that the existence of admissible strategies is not guaranteed in this setting (see for instance the examples given in [9]). In [9], the authors identified a class of games for which the existence of admissible strategies (for Player i) is guaranteed: well-formed games (for Player i). However, even in such games, the desirable fundamental property that holds for boolean games is not assured to hold anymore. In fact, this is already true for quantitative well-formed games with only three different payoffs and really simple payoff functions. Indeed, consider again the *Help-me?* game in Figure 1. Remark that it is a well-formed game for the protagonist. We already showed that any strategy s_k is dominated by the strategy s_{k+1} . Thus, none of them is admissible. The only admissible strategy is s_ω . It is easy to see that $s_k \not\preceq s_\omega$ for any $k \in \mathbb{N}$: Let $\tau \in \Sigma_{-i}$ be such that $\tau((v_0 v_1)^k) = v_0$ for all $k \in \mathbb{N}$. Then $p(s_k, \tau) = 1 > 0 = p(s_\omega, \tau)$. To sum up, we see that there exists an infinite sequence $(s_k)_{k \in \mathbb{N}}$ of strategies such that none of its elements is dominated by the only admissible strategy s_ω . However, the sequence $(s_k)_{k \in \mathbb{N}}$ is totally ordered by the dominance relation. Based on these observations, we take the approach to not only consider single strategies, but also such ordered sequences of strategies, that can represent a type of rational behaviour not captured by the admissibility concept.

2.2 Order theory

In this paragraph we recall the standard results from order theory that we need (see e.g. [19]).

A *linear order* is a total, transitive and antisymmetric relation. A linearly ordered set $(R, <)$ is a *well-order*, if every subset of R has a minimal element w.r.t. $<$. The ordinals are the canonical examples of well-orders, in as far as any well-order is order-isomorphic to an ordinal. The ordinals themselves are well-ordered by the relation $<$ where $\alpha \leq \beta$ iff α order-embeds into β . The first infinite ordinal is denoted by ω , and the first uncountable ordinal by ω_1 .

A *partial order* is a transitive and reflexive relation. Let (X, \preceq) be a partially ordered set (*poset* for short). A *chain* in (X, \preceq) is a subset of X that is totally ordered by \preceq . An *increasing chain* is an ordinal-indexed family $(x_\beta)_{\beta < \alpha}$ of elements of X such that $\beta < \gamma < \alpha \Rightarrow x_\beta \prec x_\gamma$. If we only have that $\beta < \gamma$ implies $x_\beta \preceq x_\gamma$, we speak of a *weakly increasing chain*. We are mostly interested in (weakly) increasing chains in this paper, and will thus occasionally suppress the words *weakly increasing* and only speak about *chains*.

A subset Y of a partially ordered set (X, \preceq) is called *cofinal*, if for every $x \in X$ there is a $y \in Y$ with $x \preceq y$. A consequence of the axiom of choice is that every chain contains a cofinal increasing chain, which is one reason for our focus on increasing chains. It is obvious that having multiple maximal elements prevents the existence of a cofinal chain, but even a lattice can fail to admit a cofinal chain. An example we will go back to is $\omega_1 \times \omega$ (cf. [19]).

If (X, \preceq) admits a cofinal chain, then its *cofinality* (denoted by $\text{cof}(X, \preceq)$) is the least ordinal α indexing a cofinal increasing chain in (X, \preceq) . The possible values of the cofinality are 1 or infinite regular cardinals (it is common to identify a cardinal and the least ordinal of that cardinality). In particular, a countable chain can only have cofinality 1 or ω . The first uncountable cardinal \aleph_1 is regular, and $\text{cof}(\omega_1) = \omega_1$.

We will need the probably most-famous result from order theory:

► **Lemma 2 (Zorn's Lemma).** *If every chain in (X, \preceq) has an upper bound, then every element of X is below a maximal element.*

3 Increasing chains of strategies

3.1 Ordering chains

In this subsection, we study the poset of increasing chains in a given poset (X, \preceq) . We denote by $\text{IC}(X, \preceq)$ the set of increasing chains in (X, \preceq) . Our intended application will be that (X, \preceq) is the set of strategies for the protagonist in a game ordered by the dominance relation. However, in this subsection we are not exploiting any properties specific to the game-setting. Instead, our approach is purely order-theoretic.

► **Definition 3.** We introduce an order \sqsubseteq on $\text{IC}(X, \preceq)$ by defining:

$$(x_\beta)_{\beta < \alpha} \sqsubseteq (y_\gamma)_{\gamma < \delta} \quad \text{if} \quad \forall \beta < \alpha \exists \gamma < \delta \quad x_\beta \preceq y_\gamma$$

Note that \sqsubseteq is a partial order. Let \doteq denote the corresponding equivalence relation. We will occasionally write short IC for $(\text{IC}(X, \preceq), \sqsubseteq)$.

Inspired by our application to dominance between strategies in games, we will refer to both \preceq and \sqsubseteq as the *dominance* relation, and might express e.g. $(x_\beta)_{\beta < \alpha} \sqsubseteq (y_\gamma)_{\gamma < \delta}$ as $(x_\beta)_{\beta < \alpha}$ is dominated by $(y_\gamma)_{\gamma < \delta}$, or $(y_\gamma)_{\gamma < \delta}$ dominates $(x_\beta)_{\beta < \alpha}$. There is no risk to confuse whether \preceq or \sqsubseteq is meant, since $x \preceq y$ iff $(x)_{\beta < 1} \sqsubseteq (y)_{\gamma < 1}$. Continuing the identification of $x \in X$ and $(x)_{\beta < 1} \in \text{IC}$, we will later also speak about a single strategy dominating a chain or vice versa.

The central notion we are interested in will be that of a maximal chain:

► **Definition 4.** $A \in \text{IC}$ is called *maximal*, if $A \sqsubseteq B$ for $B \in \text{IC}$ implies $B \sqsubseteq A$.

We desire situations where every chain in IC is either maximal or below a maximal chain. Noting that this goal is precisely the conclusion of Zorn's Lemma (Lemma 2), we are led to study chains of chains; for if every chain of chains is bounded, Zorn's Lemma applies. Since (IC, \sqsubseteq) is a poset just as (X, \preceq) is, notions such as cofinality apply to chains of chains just as they apply to chains. We will gather a number of lemmas we need to clarify when chains of chains are bounded.

In a slight abuse of notation, we write $(x_\beta)_{\beta < \alpha} \subseteq (y_\gamma)_{\gamma < \delta}$ iff $\{x_\beta \mid \beta < \alpha\} \subseteq \{y_\gamma \mid \gamma < \delta\}$. Clearly, $(x_\beta)_{\beta < \alpha} \subseteq (y_\gamma)_{\gamma < \delta}$ implies $(x_\beta)_{\beta < \alpha} \sqsubseteq (y_\gamma)_{\gamma < \delta}$. We can now express cofinality by noting that $(x_\beta)_{\beta < \alpha}$ is cofinal in $(y_\gamma)_{\gamma < \delta}$ iff $(x_\beta)_{\beta < \alpha} \subseteq (y_\gamma)_{\gamma < \delta}$ and $(y_\gamma)_{\gamma < \delta} \sqsubseteq (x_\beta)_{\beta < \alpha}$. We recall that the cofinality of $(y_\gamma)_{\gamma < \delta}$ (denoted by $\text{cof}((y_\gamma)_{\gamma < \delta})$) is the least ordinal α such that there exists some $(x_\beta)_{\beta < \alpha}$ which is cofinal in $(y_\gamma)_{\gamma < \delta}$.

► **Lemma 5.** If $(x_\beta)_{\beta < \alpha} \doteq (y_\gamma)_{\gamma < \delta}$, then there is some $(y'_\lambda)_{\lambda < \alpha'} \subseteq (y_\gamma)_{\gamma < \delta}$ with $\alpha' \leq \alpha$ and $(y'_\lambda)_{\lambda < \alpha'} \doteq (y_\gamma)_{\gamma < \delta}$.

► **Corollary 6.** $\text{cof}((y_\gamma)_{\gamma < \delta})$ is equal to the least ordinal α such that there exists $(x_\beta)_{\beta < \alpha}$ with $(x_\beta)_{\beta < \alpha} \doteq (y_\gamma)_{\gamma < \delta}$.

► **Corollary 7.** For every chain $(y_\gamma)_{\gamma < \delta}$ there exists an equivalent chain $(x_\beta)_{\beta < \alpha}$ such that $\alpha = 1$ or α is an infinite regular cardinal. In particular, if δ is countable, then $(y_\gamma)_{\gamma < \delta}$ is equivalent to a singleton or some chain $(x_n)_{n < \omega}$.

We briefly illustrate the concepts introduced so far in the game setting. Notice that for a game \mathcal{G} and a Player i , the pair $(\Sigma_i(\mathcal{G}), \preceq)$ is indeed a partially ordered set. We can thus consider the set $\text{IC}(\Sigma_i(\mathcal{G}), \preceq)$ of increasing chains of strategies in \mathcal{G} .

► **Example 8.** Recall the *Help-me?* game of Figure 1 and consider the set (Σ_i, \preceq) of strategies of the protagonist partially ordered by the weak dominance relation. Any single strategy is an increasing chain, indexed by the ordinal 1. We already noted that the strategy s_ω is admissible, thus the chain consisting of s_ω is maximal with respect to \sqsubseteq . Furthermore, the sequence of strategies $(s_k)_{k < \omega}$ is an increasing chain. Indeed, we know that for any $k < \omega$, we have $s_k \prec s_{k+1}$. It is a maximal one: in fact, since the set of strategies of the protagonist solely consists of the strategies of this chain and s_ω , and as $s_k \not\preceq s_\omega$ for any $k < \omega$, we get that any chain $(\sigma_\beta)_{\beta < \alpha}$ such that $(s_k)_{k < \omega} \sqsubseteq (\sigma_\beta)_{\beta < \alpha}$ satisfies $(\sigma_\beta)_{\beta < \alpha} \subseteq (s_k)_{k < \omega}$. Thus, $(\sigma_\beta)_{\beta < \alpha} \sqsubseteq (s_k)_{k < \omega}$. Let $(\sigma_\beta)_{\beta < \alpha}$ be an increasing chain indexed by the ordinal α . First, remark that $\alpha \leq \omega$. If $\alpha < \omega$, then the cofinality of $(\sigma_\beta)_{\beta < \alpha}$ is 1 as $(\sigma_\beta)_{\beta < \alpha}$ is equivalent to the strategy $\sigma_{\alpha-1}$: every strategy of $(\sigma_\beta)_{\beta < \alpha}$ is weakly dominated by $\sigma_{\alpha-1}$, and as the strategy $\sigma_{\alpha-1}$ is included in the increasing chain $(\sigma_\beta)_{\beta < \alpha}$, it is weakly dominated by $(\sigma_\beta)_{\beta < \alpha}$. If $\alpha = \omega$, then the cofinality of $(\sigma_\beta)_{\beta < \alpha}$ is ω : As for every finite chain $(\sigma'_{\beta'})_{\beta' < \alpha'}$ with $1 < \alpha' < \omega$, there exists $n < \omega$ such that $(\sigma'_{\beta'})_{\beta' < \alpha'} \sqsubseteq \sigma_n$, and thus $(\sigma_\beta)_{\beta < \alpha}$ is not (weakly) dominated by $(\sigma'_{\beta'})_{\beta' < \alpha'}$. Moreover, we have that $(\sigma_\beta)_{\beta < \alpha} \dot{=} (s_k)_{k < \omega}$ and is thus maximal. Indeed, since $(\sigma_\beta)_{\beta < \alpha}$ is a chain that is not a singleton, we already know that $(\sigma_\beta)_{\beta < \alpha} \subseteq (s_k)_{k < \omega}$, that is $(\sigma_\beta)_{\beta < \alpha} \sqsubseteq (s_k)_{k < \omega}$. Let now $k < \omega$. As $(\sigma_\beta)_{\beta < \alpha}$ is an increasing chain and $\alpha = \omega$, we have that there exists $n < \omega$ and $k' \geq k$ such that $\sigma_n = s_{k'}$. Thus, $s_k \preceq \sigma_n$ since $(s_k)_{k < \omega}$ is an increasing chain. Hence, we also have $(s_k)_{k < \omega} \sqsubseteq (\sigma_\beta)_{\beta < \alpha}$.

Now we are ready to prove the main technical result of this section 3.1, which identifies the potential obstructions for each chain in IC to have an upper bound:

► **Lemma 9.** *The following are equivalent:*

1. *If $((x_\beta^\gamma)_{\beta < \alpha_\gamma})_{\gamma < \delta}$ is an increasing chain in IC, then it has an upper bound in IC.*
2. *If $((x_\beta^\gamma)_{\beta < \alpha})_{\gamma < \delta}$ is an increasing chain in IC with $\alpha \neq \delta$, $\text{cof}((x_\beta^\gamma)_{\beta < \alpha}) = \alpha > 1$ and $\text{cof}(((x_\beta^\gamma)_{\beta < \alpha})_{\gamma < \delta}) = \delta > 1$, then it has an upper bound in IC.*

Let us illustrate the problem of extending Lemma 9 by an example:

► **Example 10** ([19, Example 1]). Let $(X, \preceq) = \omega_1 \times \omega$, i.e. the product order of the first uncountable ordinal and the first infinite ordinal. Consider the chain of chains given by $x_n^\gamma = (\gamma, n)$, this corresponds to the case $\alpha = \omega$, $\delta = \omega_1$ in Lemma 9. If this chain of chains had an upper bound, then $\omega_1 \times \omega$ would need to admit a cofinal chain. However, this is not the case.


However, we can guarantee the existence of a maximal chain above any chain when there is no uncountable increasing chain of increasing chains.

► **Theorem 11.** *If all increasing chains of elements in IC (i.e., increasing chains of increasing chains of elements of (X, \preceq)) have a countable number of elements, then for every $A \in \text{IC}$ there exists a maximal $B \in \text{IC}$ with $A \sqsubseteq B$.*

Proof. We first argue that Condition 2 in Lemma 9 is vacuously true. As all increasing chains in IC are countable, the only possible value $\delta > 1$ for $\delta = \text{cof}(((x_\beta^\gamma)_{\beta < \alpha})_{\gamma < \delta})$ is $\delta = \omega$. As (X, \preceq) embeds into IC, if all chains in IC are countable, then so are all chains in (X, \preceq) . This tells us that the only possible value for α is $\alpha = \omega$. But then $\alpha \neq \delta$ cannot be satisfied.

By Lemma 9, Condition 1 follows. We can then apply Zorn's Lemma (Lemma 2) to conclude the claim. ◀

A small modification of the example shows that we cannot replace the requirement that IC has only countable increasing chains in Theorem 11 with the simpler requirement that (X, \preceq) has only countable increasing chains:



■ **Figure 2** Two variants of the *Help-me?* game.

► **Example 12.** Let $X = \omega_1 \times \omega$, and let $(\alpha, n) \prec (\beta, m)$ iff $\alpha \leq \beta$ and $n < m$. Then (X, \prec) has only countable increasing chains, but IC still has the chain of chains given by $x_n^\gamma = (\gamma, n)$ as in Example 10.

3.2 Uncountably long chains of chains

Unfortunately, we can design a game such that there exists an uncountable increasing chain of increasing chains. Thus the existence of a maximal element above any chain is not guaranteed by Theorem 11. In fact, we will see that the chain of chains of uncountable length we construct is not below any maximal chain.

► **Example 13.** We consider a variant of the *Help-me?* game (Example 1), depicted in Figure 2a. The strategies of the protagonist in this game can be described by functions $f : \mathbb{N} \rightarrow \mathbb{N} \cup \{\infty\}$ describing how often the protagonist is willing to repeat the second loop (between v_1 and v_2) given the number of repetitions the antagonist made in the first loop (at v_0). With the same reasoning as in Example 1 we find that the strategy corresponding to a function g dominates the strategy corresponding to f iff $\forall n \in \mathbb{N} f(n) = \infty \Leftrightarrow g(n) = \infty$ and $\forall n \in \mathbb{N} f(n) \leq g(n)$.

► **Definition 14.** Let $\mathbb{N}^{\mathbb{N}}$ denote the set of functions $f : \mathbb{N} \rightarrow \mathbb{N}$. For $f, g \in \mathbb{N}^{\mathbb{N}}$, let $f \leq g$ denote that $\forall n \in \mathbb{N} f(n) \leq g(n)$.

► **Observation 15.** There is an embedding of $(\mathbb{N}^{\mathbb{N}}, \leq)$ into the strategies of the game in Example 13 ordered by dominance such that no strategy in the range of embedding is dominated by a strategy outside the range of the embedding.

► **Proposition 16** ⁽¹⁾. For every chain $(f_n)_{n \in \mathbb{N}}$ in $(\mathbb{N}^{\mathbb{N}}, \leq)$ there exists a chain of chains $((f_n^\alpha)_{n < \omega})_{\alpha < \omega_1}$ of length ω_1 with $(f_n^0)_{n < \omega} \supseteq (f_n)_{n < \omega}$.

► **Corollary 17.** *The game in Example 13 has uncountably long chains of chains not below any maximal chains.*

Proof. Combine Observation 15 and Proposition 16. ◀

¹ This result is adapted from an answer by user *Deedlit* on math.stackexchange.org [16].

3.3 Chains over countable posets (X, \preceq)

Our proof of Proposition 16 crucially relied on functions of type $f : \mathbb{N} \rightarrow \mathbb{N}$ with arbitrarily high rate of growth. In concrete applications such functions would typically be unwelcome. In fact, for almost all classes of games of interest in (theoretical) computer science, a countable collection of strategies suffices for the players to attain their attainable goals. Restricting to computable strategies often makes sense. Many games played on finite graphs are even finite-memory determined (see [18] for how this extends to the quantitative case), and thus strategies implementable by finite automata are all that need to be considered.

Restricting consideration to a countable set of strategies indeed circumvents the obstacle presented by Proposition 16. The reason is that the cardinality of the length of a chain of chains cannot exceed that of the underlying partially ordered set (X, \preceq) :

► **Proposition 18.** For any increasing chain $((x_\beta^\gamma)_{\beta < \alpha})_{\gamma < \delta}$ in $\text{IC}(X, \preceq)$ we find that $|\delta| \leq |X|$.

Proof. Let $X_\gamma = \{x \in X \mid \exists \beta < \alpha \ x \preceq x_\beta^\gamma\}$. We find that $X_{\gamma_1} \subsetneq X_{\gamma_2}$ for any $\gamma_1 < \gamma_2 < \delta$ as a direct consequence of $(x_\beta^{\gamma_1})_{\beta < \alpha} \sqsubset (x_\beta^{\gamma_2})_{\beta < \alpha}$. Pick for each $\gamma < \delta$ some $y_\gamma \in X_{\gamma+1} \setminus X_\gamma$. Then $y : \delta \rightarrow X$ is an injection, establishing $|\delta| \leq |X|$. ◀

► **Corollary 19.** If (X, \preceq) is countable, then any increasing chain is maximal or below a maximal chain.

Proof. Proposition 18 shows that Theorem 11 applies. ◀

► **Example 20.** We return to the *Help-me?* game (Example 1, Figure 1). With the analysis done in Example 8, we have seen that any increasing chain C is either maximal or such that $C \sqsubseteq (\sigma_n)_{n < \omega}$, which is maximal. This fact can be derived directly from Corollary 19 as the number of strategies in \mathcal{G} is countable. Note also that the seemingly irrelevant loop we added in Figure 2a has a fundamental impact on the behaviour of chains of strategies!

4 Generalised safety/reachability games

► **Definition 21.** A *generalised safety/reachability game* (for Player i) $\mathcal{G} = \langle P, G, L, (p_i)_{i \in P} \rangle$ is a turn-based multiplayer game on a finite graph such that:

- $L \subseteq V$ is a finite set of *leaves*,
- for each $\ell \in L$, we have that $(\ell, v) \in E$ if, and only if $v = \ell$, that is, each leaf is equipped with a self-loop, and no other outgoing transition,
- for each $\ell \in L$, there exists an associated payoff $n_\ell \in \mathbb{Z}$ such that: for each outcome ρ , we have $p_i(\rho) = \begin{cases} n_\ell & \text{if } \rho \in V^* \ell^\omega, \\ 0 & \text{otherwise.} \end{cases}$

The traditional reachability games can be recovered as the special case where all leaves are associated with the same positive payoff, whereas the traditional safety games are those generalised safety/reachability games with a single negative payoff attached to leaves. This class was studied under the name *chess-like* games in [5, 6].

Generalised safety/reachability games are *well-formed* for Player i . Furthermore, they are prefix-independent, that is, for any outcome ρ and history h , we have that $p_i(h\rho) = p_i(\rho)$. Without loss of generality, we consider that there is either a unique leaf $\ell(n) \in L$ or no leaf for each possible payoff $n \in \mathbb{Z}$.

It follows from the transfer theorem in [18] (in fact, already from the weaker transfer theorem in [13]) that generalised safety/reachability games are finite memory determined. With a slight modification, we see that for any history h and strategy σ , there exists a

finite-memory strategy σ' such that $cVal(h, \sigma') = cVal(h, \sigma)$ and $aVal(h, \sigma') = aVal(h, \sigma)$. We shall thus restrict our attention to finite memory strategies, of which there are only countably many. We then obtain immediately from Corollary 19:

► **Corollary 22.** *In a generalised safety/reachability game, every increasing chain comprised of finite memory strategies is either maximal or dominated by a maximal such chain.*

If our goal is only to obtain a dominance-related notion of rationality, then for generalised safety/reachability games we can be satisfied with maximal chains comprised of finite memory strategies. However, for applications, it would be desirable to have a concrete understanding of these maximal chains. For this, having used Zorn's Lemma in the proof of their existence surely is a bad omen!

After collecting some useful lemmas on dominance in generalised safety/reachability games in Section 4.1, we will introduce the notion of *uniform chains* in Section 4.2. These are realized by automata of a certain kind, and thus sufficiently concrete to be amenable to algorithmic manipulations.

4.1 Dominance in generalised safety/reachability games

Given a generalised safety/reachability game \mathcal{G} and two strategies σ_1 and σ_2 of Player i , we can provide a criterion to show that σ_1 is not dominated by σ_2 :

► **Lemma 23.** *Let σ_1 and σ_2 be two strategies of Player i in a generalised safety/reachability game \mathcal{G} . Then, $\sigma_1 \not\preceq \sigma_2$ if, and only if, there exists a history h compatible with σ_1 and σ_2 such that $\text{last}(h) \in V_i$, $\sigma_1(h) \neq \sigma_2(h)$ and $cVal(h, \sigma_1) > aVal(h, \sigma_2)$.*

Intuitively, if there is no history where the two strategies disagree, they are in fact equivalent, and if, at every history where they disagree, the best payoff σ_1 can achieve (that is, $cVal(h, \sigma_1)$) is less than the one σ_2 can ensure (that is, $aVal(h, \sigma_2)$), then $\sigma_1 \preceq \sigma_2$. On the other hand, if they disagree at a history h and the best payoff σ_1 can achieve is strictly greater than the one σ_2 can ensure, then there exist a strategy of the antagonist that will yield exactly these payoffs against σ_1 and σ_2 respectively, which means that $\sigma_1 \not\preceq \sigma_2$. This result follows from the proof of Theorem 11 in [9]. The proof adapted to our setting can be found in the appendix.

We call such a history h a *non-dominance witness* of σ_1 by σ_2 . The existence of non-dominance witnesses allows us to conclude that in generalised safety/reachability games, all increasing chains are countable (not just those comprised of finite memory strategies).

► **Corollary 24.** *If $(\sigma_\beta)_{\beta < \alpha}$ is an increasing chain in generalised safety/reachability game, then α is countable.*

Proof. Assume that a history h is a witness of non-dominance of σ_2 by σ_1 , and of σ_3 by σ_2 , but not of σ_1 by σ_2 or σ_2 by σ_3 . Then $cVal(h, \sigma_2) > aVal(h, \sigma_1)$, $cVal(h, \sigma_3) > aVal(h, \sigma_2)$, $cVal(h, \sigma_1) \leq aVal(h, \sigma_2)$ and $cVal(h, \sigma_2) \leq aVal(h, \sigma_3)$. It follows that $aVal(h, \sigma_1) < aVal(h, \sigma_3)$ and $cVal(h, \sigma_1) < cVal(h, \sigma_3)$. Thus, if there are k different possible values, then any increasing chain of strategies using h as witness of non-dominance between them can have length at most $2k - 1$.

But if there were an uncountably long increasing chain, by the pigeon hole principle it would have an uncountably long subchain where all non-dominance witnesses in the reverse direction are given by the same history. ◀

10:12 Dominance Between Chains of Strategies

As we only handle countable chains, in the following we use the usual notation $(\sigma_n)_{n \in \mathbb{N}}$ to index chains.

The following lemma states that we can also extract witnesses for a strategy to be non-maximal (non-admissible or strictly dominated):

► **Lemma 25.** *Let \mathcal{G} be a generalised safety/reachability game and σ a strategy of Player i . The strategy σ is not admissible if, and only if there exists a history h compatible with σ such that $\text{aVal}(h, \sigma) \leq \text{cVal}(h, \sigma) \leq \text{aVal}(h) \leq \text{acVal}(h)$ where at least one inequality is strict.*

This result is a reformulation of Theorem 11 in [9] catered to our context and with a focus on the non-admissibility rather than on admissibility (see the arXiv version [3] for a proof adapted to our setting).

► **Definition 26.** Call a history h as in Lemma 25 a non-admissibility witness for σ . Call σ preadmissible, if for every non-admissibility witness h of σ we find that $h = h'vh''$ with $\text{aVal}(h'v, \sigma) = \text{aVal}(h'v)$ and $\text{cVal}(h'v, \sigma) = \text{acVal}(h'v)$.

While a preadmissible strategy may fail to be admissible, it is not possible to improve upon it the first time it enters some vertex. Only when returning to a vertex later it may make suboptimal choices. Moreover, before a dominated choice is possible at a vertex, previously both the antagonistic and the antagonistic-cooperative value were realized at that vertex by the preadmissible strategy.

► **Lemma 27.** *In a generalised safety/reachability game, every strategy is either preadmissible or dominated by a preadmissible strategy.*

Proof sketch. Essentially, we can change how a strategy behaves locally on those histories that are an obstacle to it being preadmissible by replacing by a finite memory strategy that realizes the antagonistic and the antagonistic-cooperative value there. ◀


4.2 Parameterized automata and uniform chains

Let a *parameterized automaton* be a Mealy automaton that in addition can access a single counter in the following way: In a counter-access-state, a transition is chosen based on whether the counter value is 0 or not. Furthermore, in these counter-access-states, when the counter value is greater than 0, the counter is decremented by 1, otherwise, it stays at 0. In the remaining states, only one transition is possible and the counter value is not affected.

► **Definition 28.** A *parameterized automaton* for Player $i \in P$ over a game graph $G = (V, E)$ is a tuple $\mathcal{M} = (M, M_C, m_0, V, \mu, \nu)$ where:

- M is a non-empty finite set of memory states and $M_C \subseteq M$ is the set of *counter-access* states,
- m_0 is the initial memory state,
- V is the set of vertices of G ,
- $\mu : M \times V \times \mathbb{N} \rightarrow M \times \mathbb{N}$ is the memory and counter update function,
- $\nu : M \times V_i \times \mathbb{N} \rightarrow V$ is the move choice function for Player i , such that $(v, \nu(m, v, n)) \in E$ for all $m \in M$ and $v \in V_i$ and $n \in \mathbb{N}$.

The memory and counter-update function μ respects the following conditions: for each $m \in M \setminus M_C$, and $v \in V$, there exists $m' \in M$ such that $\mu(m, v, n) = (m', n)$ for all $n \in \mathbb{N}$. for each $m \in M_C$, and $v \in V$, there exists $m' \in M$ such that $\mu(m, v, n) = (m', n - 1)$ for all $n > 0$ and $m'' \in M$ such that $\mu(m, v, 0) = (m'', 0)$. The move choice function ν respects



■ **Figure 3** Product of the *Help-me?* game with parameterized automaton with a single memory state realizing $(s_k)_{k \in \mathbb{N}}$.

the following conditions: for each $m \in M \setminus M_C$, and $v \in V_i$, there exists $v' \in V$ such that $\nu(m, v, n) = v'$ for all $n \in \mathbb{N}$. For each $m \in M_C$, and $v \in V_i$, there exists $v', v'' \in V$ such that $\nu(m, v, n) = v'$ for all $n > 0$ and $\nu(m, v, 0) = v''$.

To ease presentation and understanding, we call transitions that decrement the counter *green* transitions, the transitions only taken when the counter value is 0 *red* transitions, and the ones that do not depend on the counter value *black* transitions. This classification between *green*, *red* and *black* transitions extends naturally to the edges of the product $\mathcal{M} \times G$ (that is, the graph with set of vertices $M \times V$ and edges induced by the functions μ and ν).

Parameterized automata can be seen as a collection of finite Mealy automata, one for each initialization of the counter. Thus, we say that a parameterized automaton \mathcal{M} realizes a *sequence* of finite-memory strategies $(\sigma_n)_{n \in \mathbb{N}}$. In the remainder of the paper, we focus on chains realized by parameterized automata:

► **Definition 29.** Let a chain $(\sigma_n)_{n \in \mathbb{N}}$ of strategies be called a *uniform* chain if there is a parameterized automaton \mathcal{M} that realizes σ_n if the counter is initialized with the value n . If $(\sigma_n)_{n \in \mathbb{N}}$ is maximal for \sqsubseteq amongst the increasing chains comprised of finite memory strategies, we call it a maximal uniform chain.

► **Example 30.** The *Help-me?* game from Figure 1 is clearly a generalised safety/reachability game with two leaves. The chain of strategies $(s_k)_{k \in \mathbb{N}}$ exposed in Example 1 is a uniform chain, as it is realized by the parameterized automaton that loops k times when its counter is initialized with value k . Figure 3 shows the product between this parameterized automaton and the game graph. The green (doubled) edge corresponds to the transition to take when the counter value is greater than 0 and should be decremented, while the red (dashed) edge corresponds to the transition to take when the counter value is 0.

The following theorem shows us that uniform chains indeed suffice to realize any rational behaviour in the sense of maximal chains:

► **Theorem 31.** *In a generalised safety/reachability game, every dominated finite memory strategy is dominated by an admissible finite memory strategy or by a maximal uniform chain.*

Theorem 31 cannot be extended to state that every chain comprised of finite memory strategies is below an admissible strategy or a maximal uniform chain. Note that there are only countably many uniform chains.

► **Example 32.** There is a generalised safety/reachability game where there are uncountably many incomparable maximal chains of finite memory strategies.

Proof. Consider the game depicted in Figure 2b. For any $p \in \{a, b\}^\omega$, define a chain of finite memory strategies by letting the n -strategy be *loop n times while playing the symbols from $p_{\leq n}$, then quit*. For each p , we obtain a different maximal chain. ◀

4.3 Algorithmic properties

In this section, we prove two decidability results concerning parametrized automata.

First, we prove that we can decide whether the sequence of strategies realized by a parameterized automaton is a chain. Note that this decision problem is not trivial: not every parameterized automaton realizes an (increasing) chain of strategies. For instance, if we switch the red and green transitions in the automaton/game graph product of figure 3, the sequence of strategies realized consists of s_ω when the counter is initialized with value 0, and s_0 when it is initialized with any other value. As $s_\omega \not\preceq s_0$, it is not a chain.

Second, we demonstrate that we can compare uniform chains: given two parametrized automata defining chains of strategies, we can decide whether one is dominated by the other. We begin by proving that strategies realized by Mealy automata are comparable.

► **Lemma 33.** *Let \mathcal{G} be a generalised safety/reachability game, let σ and σ' be finite-memory strategies realized by the finite Mealy automata \mathcal{M} and \mathcal{M}' . It is decidable in PTIME whether $\sigma \preceq \sigma'$.*

Proof sketch. We construct the game \mathcal{G}' of perfect information for two players, *Challenger* and *Prover*, such that Prover wins the game if and only if $\sigma \preceq \sigma'$. The goal of Challenger is to show that there exists a non-dominance witness of σ by σ' , that is, according to Lemma 23, a history h compatible with σ and σ' such that $\text{last}(h) \in V_i$, $\sigma(h) \neq \sigma'(h)$ and $\text{cVal}(h, \sigma) > \text{aVal}(h, \sigma')$. The game can be decomposed into the following phases:

- first, Challenger chooses a path \tilde{h} in $\mathcal{M} \times G \times \mathcal{M}'$ such that \tilde{h} has no successor in $\mathcal{M} \times G \times \mathcal{M}'$. This guarantees that h is compatible with σ and σ' , and that $\sigma(h) \neq \sigma'(h)$.
- Challenger then announces two values: c and a , such that $c > a$.
- Prover now can choose to contest either value c or value a .
- If Prover chooses to contest c , the game proceeds to a subgame \mathcal{C} , where Challenger has to find a continuation path in $(\mathcal{M} \times G)$ that yields a payoff c , to prove that $\text{cVal}(h, \sigma) \geq c$.
- If Prover chooses to contest a , the game proceeds to a subgame \mathcal{A} , where Challenger has to find a valid continuation path in $(\mathcal{M}' \times G)$ that yields a payoff a , to prove that $\text{aVal}(h, \sigma') \leq a$.

Informally, if $\sigma \not\preceq \sigma'$, Challenger is able to select correctly a non-dominance witness h of σ by σ' , and the two values $c = \text{cVal}(h, \sigma)$ and $a = \text{aVal}(h, \sigma')$ such that $c > a$. Thus, he can follow in \mathcal{G}' the path \tilde{h} corresponding to h , then continue, depending on the choice of Prover, to follow either a continuation of h that yields a payoff c with strategy σ or a continuation of h that yields a payoff a with strategy σ' . Symmetrically, if $\sigma \preceq \sigma'$, then for any history h compatible with σ and σ' where $\sigma(h) \neq \sigma'(h)$, we have that $\text{cVal}(h, \sigma) \leq \text{aVal}(h, \sigma')$. Thus any choice of pair (c, a) with $c > a$ by Challenger is faulty: either $c > \text{cVal}(h, \sigma)$, in which case Prover can let the game proceed to \mathcal{C} , and Challenger will fail to expose a continuation of h that yields a payoff c with strategy σ , or $a < \text{aVal}(h, \sigma')$, in which case Prover can let the game proceed to \mathcal{A} , and Challenger will fail to expose a continuation of h that yields a payoff a with strategy σ' . As the game graph we construct for this *Prover* game has a size polynomial in the size of the strategy automata and the game graph, and as solving this game amounts to solving a polynomially bounded number of reachability and safety subgames, we obtain that the question whether $\sigma \preceq \sigma'$ is decidable in PTIME. ◀

We now expose equivalences between the decision problems we are interested in, and properties (P₁), (P₂) and (P₃) that can be decided with the use of Lemma 33.

► **Proposition 34.** Let \mathcal{G} be a generalised safety/reachability game over a graph G . Let \mathcal{M} be a Mealy automaton realizing a finite memory strategy M , and let \mathcal{S} and \mathcal{T} be parameterized automata realizing sequences $(S_n)_{n \in \mathbb{N}}$ and $(T_n)_{n \in \mathbb{N}}$ of finite memory strategies. Then:

1. Let $N_{\prec} = |G||\mathcal{S}|$.
Then $(S_n)_{n \in \mathbb{N}}$ is a chain if and only if (P₁) $S_i \preceq S_{i+1}$ for every $1 \leq i \leq N_{\prec}$.
2. Let $N_T = |G||\mathcal{T}|(|\mathcal{M}| + 1) + 1$, and suppose that $(T_n)_{n \in \mathbb{N}}$ is a chain.
Then $M \not\sqsubseteq (T_n)_{n \in \mathbb{N}}$ if and only if (P₂) $M \not\preceq T_{N_T}$.
3. Let $N_S = |G||\mathcal{S}|(2|\mathcal{T}| + 1)$, and suppose that $(S_n)_{n \in \mathbb{N}}$ and $(T_n)_{n \in \mathbb{N}}$ are chains.
Then $(S_n)_{n \in \mathbb{N}} \not\sqsubseteq (T_n)_{n \in \mathbb{N}}$ if and only if (P₃) $S_{N_S} \not\preceq (T_n)_{n \in \mathbb{N}}$.

Proof sketch. Note that for every item, the backward implication is straightforward. The proof of each forward implication relies on the study of the loops that appear in witnesses of non dominance, whose existence is guaranteed by Lemma 23. For item 1, we prove that, given a witness of non-dominance of T_i by T_{i+1} for any integer $i > N_{\prec}$, we are able to construct a witness of non-dominance of T_j by T_{j+1} for some $j \leq N_{\prec}$ by exposing loops that can be pumped down.

To prove item 2, we show that since $(T_n)_{n \in \mathbb{N}}$ is a chain, $M \not\sqsubseteq (T_n)_{n \in \mathbb{N}}$ if and only if M is not dominated by T_N for arbitrarily large N . If M is dominated by T_{N_T} , we exhibit a loop in a witness of non dominance, which, once pumped, allows us to create witnesses of non dominance of M by T_N for arbitrarily large N , yielding the desired result.

Finally, item 3 is proved as follows. Since $(S_n)_{n \in \mathbb{N}}$ and $(T_n)_{n \in \mathbb{N}}$ are chains, $(S_n)_{n \in \mathbb{N}} \not\sqsubseteq (T_n)_{n \in \mathbb{N}}$ if and only if there exists an integer N such that $S_N \not\preceq (T_n)_{n \in \mathbb{N}}$. Once again, we show that if such an N exists, there is at least one that is smaller than N_S . ◀

Since the property P₁ can be decided in PTIME by applying Lemma 33 with adequately chosen Mealy automata as parameters, we obtain the following theorem.

► **Theorem 35.** *Given a generalised safety/reachability game and a parameterized automaton, we can decide in PTIME whether the automaton realizes a chain of strategies.*

Similarly, the property P₂ can be decided in PTIME by applying Lemma 33 with \mathcal{M} and the Mealy automaton corresponding to the strategy T_{N_T} as parameters. Moreover, by Proposition 34.2, the problem of deciding property P₃ can be reduced in polynomial time to the problem of deciding property P₂. Therefore Proposition 34.3 implies our final decidability result.

► **Theorem 36.** *Given a generalised safety/reachability game and two parameterized automata realizing uniform chains of strategies, we can decide in PTIME whether the chain realized by the first is dominated by the one from the second.*

5 Conclusion and outlook

In quantitative games with more than three possible payoffs, there are strategies that are dominated but not dominated by any admissible strategy. Example 1 suggests that chains of strategies could provide a suitable framework to circumvent this issue. Abstract order-theoretic considerations revealed that in the most general case, this does not work. However, if we restrict to countable collections of strategies, every chain is below a maximal chain. This restriction is very natural, as it covers all computable strategies.

We explored the abstract approach in the concrete setting of generalized safety/reachability games. Here, parameterized automata can give a very concrete meaning to chains of strategies. Several fundamental algorithmic questions are decidable in PTIME. There are

more algorithmic questions to investigate: first and foremost, deciding, given a parameterized automaton, whether the chain realized is maximal or not, is a relevant question left open.

Moreover, our results on this class of games mostly rely on the prefix-independence and finite-range of the payoff function, and on the restriction to finite-memory strategies. Thus, it seems achievable to extend our approach to other classes of games that enjoy these properties, such as quantitative extensions of parity or Muller games, in the sense of [20] and [24]. A more ambitious objective would be to tackle more general classes of games, starting by dropping the finite-range hypothesis to encompass, for instance, mean-payoff games [14].

Finally, in the boolean case, in addition to the fundamental property that a strategy is either admissible or dominated by an admissible strategy, the admissibility notion exhibits other good properties. Indeed, in [4], the author proves that, in games with ω -regular winning conditions on finite graphs, the set of admissible strategies is itself an ω -regular set. Furthermore, as shown in [11], assuming all the players are rational (that is, play admissible strategies) yields robust and resilient solutions for strategy synthesis.

This synthesis problem remains to be investigated in the quantitative setting.

References

- 1 Martín Abadi, Leslie Lamport, and Pierre Wolper. Realizable and unrealizable specifications of reactive systems. In *ICALP'89*, volume 372 of *LNCS*. Springer, 1989.
- 2 Nicolas Basset, Gilles Geeraerts, Jean-François Raskin, and Ocan Sankur. Admissibility in concurrent games. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, volume 80 of *LIPICs*, pages 123:1–123:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/LIPICs.ICALP.2017.123.
- 3 Nicolas Basset, Ismaël Jecker, Arno Pauly, Jean-François Raskin, and Marie Van den Boogaard. Beyond admissibility: Dominance between chains of strategies. arXiv 1805.11608, 2018. URL: <https://arxiv.org/abs/1805.11608>.
- 4 Dietmar Berwanger. Admissibility in infinite games. In *STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science, Aachen, Germany, February 22-24, 2007, Proceedings*, pages 188–199, 2007. doi:10.1007/978-3-540-70918-3_17.
- 5 Endre Boros, Khaled Elbassioni, Vladimir Gurvich, and Kazuhisa Makino. On Nash equilibria and improvement cycles in pure positional strategies for chess-like and backgammon-like n -person games. *Discrete Mathematics*, 312(4):772–788, 2012. doi:10.1016/j.disc.2011.11.011.
- 6 Endre Boros, Vladimir Gurvich, and Emre Yamangil. Chess-like games may have no uniform nash equilibria even in mixed strategies. *Game Theory*, 2013. doi:10.1155/2013/534875.
- 7 Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour, Jean-François Raskin, Ocan Sankur, and Mathieu Sassolas. Non-zero sum games for reactive synthesis. In *Language and Automata Theory and Applications - 10th International Conference, LATA 2016, Prague, Czech Republic, March 14-18, 2016, Proceedings*, volume 9618 of *Lecture Notes in Computer Science*, pages 3–23. Springer, 2016.
- 8 Romain Brenguier, Arno Pauly, Jean-François Raskin, and Ocan Sankur. Admissibility in Games with Imperfect Information (Invited Talk). In Roland Meyer and Uwe Nestmann, editors, *28th International Conference on Concurrency Theory (CONCUR 2017)*, volume 85 of *LIPICs*, pages 2:1–2:23. Schloss Dagstuhl, 2017. doi:10.4230/LIPICs.CONCUR.2017.2.
- 9 Romain Brenguier, Guillermo A. Pérez, Jean-François Raskin, and Ocan Sankur. Admissibility in quantitative graph games. In *36th IARCS Annual Conference on Foundations of*

- Software Technology and Theoretical Computer Science, FSTTCS 2016, December 13-15, 2016, Chennai, India*, pages 42:1–42:14, 2016. doi:10.4230/LIPIcs.FSTTCS.2016.42.
- 10 Romain Brenguier, Jean-François Raskin, and Ocan Sankur. Assume-admissible synthesis. *Acta Inf.*, 54(1):41–83, 2017.
 - 11 Romain Brenguier, Jean-François Raskin, and Ocan Sankur. Assume-admissible synthesis. *Acta Inf.*, 54(1):41–83, 2017.
 - 12 Romain Brenguier, Jean-François Raskin, and Mathieu Sassolas. The complexity of admissibility in omega-regular games. In *CSL-LICS '14, 2014*. ACM, 2014. doi:10.1145/2603088.2603143.
 - 13 Thomas Brihaye, Veronique Bruyère, and Julie De Pril. Equilibria in quantitative reachability games. In *Proc. of CSR*, volume 6072 of *LNCS*. Springer, 2010. doi:10.1007/978-3-642-13182-0_7.
 - 14 A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, Jun 1979. doi:10.1007/BF01768705.
 - 15 Marco Faella. Admissible strategies in infinite games over graphs. In *MFCS 2009*, volume 5734 of *Lecture Notes in Computer Science*, pages 307–318. Springer, 2009.
 - 16 Deedlit (<https://math.stackexchange.com/users/21465/deedlit>). Height of a certain order. Mathematics Stack Exchange. URL:<https://math.stackexchange.com/q/2436983> (version: 2017-09-20). URL: <https://math.stackexchange.com/q/2436983>.
 - 17 Orna Kupferman. On high-quality synthesis. In S. Alexander Kulikov and J. Gerhard Woeginger, editors, *11th International Computer Science Symposium in Russia, CSR 2016*, pages 1–15. Springer International Publishing, 2016. doi:10.1007/978-3-319-34171-2_1.
 - 18 Stéphane Le Roux and Arno Pauly. Extending finite memory determinacy. *Information and Computation*, 201X. doi:10.1016/j.ic.2018.02.024.
 - 19 George Markowsky. Chain-complete posets and directed sets with applications. *Algebra Universalis*, 6:53–68, 1976. doi:10.1007/BF02485815.
 - 20 Soumya Paul and Sunil Easaw Simon. Nash equilibrium in generalised muller games. In *FSTTCS*, volume 4 of *LIPICs*, pages 335–346. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2009.
 - 21 Arno Pauly. The computational complexity of iterated elimination of dominated strategies. *Theory of Computing Systems*, pages 52–75, 2016. doi:10.1007/s00224-015-9637-1.
 - 22 Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190, 1989. doi:10.1145/75277.75293.
 - 23 H.J. Prömel, W. Thumser, and B. Voigt. Fast growing functions based on Ramsey theorems. *Discrete Mathematics*, 95(1):341–358, 1991. doi:0012-365X(91)90346-4.
 - 24 Stéphane Le Roux. From winning strategy to nash equilibrium. *Math. Log. Q.*, 60(4-5):354–371, 2014.
 - 25 Wang Shang-Zhi and Li Bo-Yu. On the minimal cofinal subsets of a directed quasi-ordered set. *Discrete Mathematics*, 48(2):289–306, 1984. doi:10.1016/0012-365X(84)90189-4.

Here we present some, but not all omitted proofs. For the complete account, we refer to the arXiv version at [3].

A Proofs omitted from Section 3

► **Lemma 37.** *If $(x_\beta)_{\beta < \alpha} \sqsubseteq (y_\gamma)_{\gamma < \delta}$ and $\alpha < \text{cof}((y_\gamma)_{\gamma < \delta})$, then there exists $\gamma_0 < \delta$ such that*

$$(x_\beta)_{\beta < \alpha} \sqsubseteq (y_{\gamma_i})_{i < 1}$$

Proof of Lemma 9. It is clear that 2 is a special case of 1. We thus just need to show that any potential obstruction to 1 can be assumed to have the form in 2.

10:18 Dominance Between Chains of Strategies

By replacing each $(x_\beta^\gamma)_{\beta < \alpha_\gamma}$ with some suitable cofinal increasing chain if necessary, we can assume that $\text{cof}((x_\beta^\gamma)_{\beta < \alpha_\gamma}) = \alpha_\gamma$ for all $\gamma < \delta$.

Consider $\{(x_\beta^\gamma)_{\beta < \alpha_\gamma} \mid \exists \gamma' > \gamma \ \alpha_\gamma < \alpha_{\gamma'}\}$. If this set is cofinal in $((x_\beta^\gamma)_{\beta < \alpha_\gamma})_{\gamma < \delta}$, then for each γ inside that set pick some witness γ' , and let y_γ be the witness obtained from Lemma 37. Now $\{y_\gamma \mid \exists \gamma' > \gamma \ \alpha_\gamma < \alpha_{\gamma'}\}$ is the desired upper bound.

If the set from the paragraph above is not cofinal, then there exists some $\delta' < \delta$ such that for $\delta' \leq \gamma < \gamma' < \delta$ we always have that $\alpha_\gamma \geq \alpha_{\gamma'}$. As the α_γ are ordinals, decreases can happen only finitely many times. Thus, by moving to a suitable cofinal subset we can safely assume that all α_γ are equal to some fixed α .

Again by moving to a suitable cofinal subset, we can assume that $\text{cof}(((x_\beta^\gamma)_{\beta < \alpha})_{\gamma < \delta}) = \delta$. If $\delta = 1$, the statement is trivial. If $\alpha = 1$, then $(x_0^\gamma)_{\gamma < \delta}$ is the desired upper bound. It remains to handle the case $\alpha = \delta > 1$.

We construct some function $f : \alpha \rightarrow \alpha$, such that the desired upper bound $(y_\epsilon)_{\epsilon < \alpha}$ is of the form $y_\epsilon = x_{f(\epsilon)}^\epsilon$. We proceed as follows: Set $f(0) = 0$. Once $f(\zeta)$ has been defined for all $\zeta < \epsilon$, pick for each $\zeta < \epsilon$ some $g(\zeta)$ such that $x_{f(\zeta)}^\zeta \preceq x_{g(\zeta)}^\epsilon$ and $x_\epsilon^\zeta \preceq x_{g(\zeta)}^\epsilon$. As $\epsilon < \alpha$, it cannot be that $\{x_{g(\zeta)}^\epsilon \mid \zeta < \epsilon\}$ is cofinal in $\{x_\beta^\epsilon \mid \beta < \alpha\}$. Thus, it has some upper bound, and we define $f(\epsilon)$ such that $x_{f(\epsilon)}^\epsilon$ is such an upper bound. ◀

Proof of Proposition 16. For each countable limit ordinal α , we fix² some fundamental sequence $(\alpha[m])_{m < \omega}$ of ordinals with $\alpha[m] < \alpha$ and $\sup_{m \in \omega} \alpha[m] = \alpha$.

Let $f_n^0(k) = \max\{f(k), k\}$. Let $f_n^{\alpha+1}(k) = \max_{j \leq k} (f_{n+j}^\alpha)(k) + 1$, and for limit ordinals α , let $f_n^\alpha(k) = \max_{m \leq n+k} f_n^{\alpha[m]}(k)$.

Claim: If $\alpha \leq \beta$, then $(f_n^\alpha)_{n < \omega} \sqsubseteq (f_m^\beta)_{m < \omega}$.

Proof. It suffices to show that if $\alpha \leq \beta$, then $f_n^\alpha \leq f_n^\beta$ for all n greater than some t . If $\beta = \alpha + 1$, this is immediate already for $t = 0$. For β a limit ordinal, we note that $f_n^{\beta[m]} \leq f_n^\beta$ for $n \geq m$.

The claim then follows by induction over β . Recall that if β is a limit ordinal and $\alpha < \beta$, then there is some $m \in \omega$ with $\alpha \leq \beta[m]$. Since for any given α, β , the ordinals γ between α and β we will need to inspect in the induction form a decreasing chain, there are only finitely many such ordinals. In particular, the maximum of all thresholds t we encounter is well-defined. ◀

Claim: If $\alpha > \beta$, then $(f_n^\alpha)_{n < \omega} \not\sqsubseteq (f_m^\beta)_{m < \omega}$.

Proof. Due to transitivity of \sqsubseteq and the previous claim, it suffices to show that $(f_m^{\alpha+1})_{m < \omega} \not\sqsubseteq (f_n^\alpha)_{n < \omega}$. Write $g_n = f_n^\alpha$. Assume the contrary, i.e. that for all $n < \omega$ there exists some $m < \omega$ such that for all $k \in \mathbb{N}$ and for all $j \leq k$ we have that $g_{n+j}(k) + 1 \leq g_m(k)$. In particular, for $n = 0$ we would have that $\forall k \in \mathbb{N} \ \forall j \leq k \ g_j(k) + 1 \leq g_m(k)$, and then setting $k = j = m$, that $g_m(m) + 1 \leq g_m(m)$, which is a contradiction. ◀

² We have no computability or other uniformity requirements to satisfy, and can thus just invoke the axiom of choice. Otherwise, as discussed e.g. in [23, Section 3.1] this approach would fail.

B

 Proofs omitted from Subsection 4.1

Proof of Lemma 23.

\implies Suppose that for every history h compatible with σ_1 and σ_2 such that $\text{last}(h) \in V_i$ and $\sigma_1(h) \neq \sigma_2(h)$, we have that $\text{cVal}(h, \sigma_1) \leq \text{aVal}(h, \sigma_2)$. We show that $\sigma_1 \preceq \sigma_2$. Let τ be a strategy of Player $-i$. Consider $\rho_1 = \mathbf{Out}(\sigma_1, \tau)$ and $\rho_2 = \mathbf{Out}(\sigma_2, \tau)$. If for all prefixes $h' \subseteq_{\text{pref}} \rho_1$ such that $\text{last}(h') \in V_i$, it holds that $\sigma_1(h') = \sigma_2(h')$, then in fact $\rho_1 = \rho_2$ and $p_i(\sigma_1, \tau) = p_i(\sigma_2, \tau)$. Otherwise, let h be the least common prefix of ρ_1 and ρ_2 such that $\text{last}(h) \in V_i$ and $\sigma_1(h) \neq \sigma_2(h)$. We know that $p_i(\rho_1) \leq \text{cVal}(h, \sigma_1)$ and $p_i(\rho_2) \geq \text{aVal}(h, \sigma_2)$ since $h \subseteq_{\text{pref}} \rho_1$ and $h \subseteq_{\text{pref}} \rho_2$. As $\text{cVal}(h, \sigma_1) \leq \text{aVal}(h, \sigma_2)$, we have that $p_i(\sigma_1, \tau) \leq p_i(\sigma_2, \tau)$. Thus, for every $\tau \in \Sigma_{-i}$, it holds that $p_i(\sigma_1, \tau) \leq p_i(\sigma_2, \tau)$, that is, $\sigma_1 \preceq \sigma_2$.

\impliedby Let h be a history compatible with σ_1 and σ_2 such that $\text{last}(h) \in V_i$, $\sigma_1(h) \neq \sigma_2(h)$ and $\text{cVal}(h, \sigma_1) > \text{aVal}(h, \sigma_2)$. Then, there exists two strategies τ_1 and τ_2 of player $-i$ such that $p_i(h, \sigma_1, \tau_1) = \text{cVal}(h, \sigma_1)$ and $p_i(h, \sigma_2, \tau_2) = \text{aVal}(h, \sigma_2)$. Let τ be a strategy

of player $-i$ compatible with h , and define $\tau'(h') = \begin{cases} \tau_1(h') & \text{if } h\sigma_1(h) \subseteq_{\text{pref}} h', \\ \tau_2(h') & \text{if } h\sigma_2(h) \subseteq_{\text{pref}} h', \\ \tau(h') & \text{otherwise.} \end{cases}$

The strategy τ' is well defined, as $\sigma_1(h) \neq \sigma_2(h)$. Furthermore, we have that $p_i(\sigma_1, \tau') = p_i(h, \sigma_1, \tau_1) = \text{cVal}(h, \sigma_1) > \text{aVal}(h, \sigma_2) = p_i(h, \sigma_2, \tau_2) = p_i(\sigma_2, \tau')$, since generalised safety/reachability games are prefix-independent. Thus, $\sigma_1 \not\preceq \sigma_2$. \blacktriangleleft

Proof of Lemma 27. For each vertex v in the game, we fix a finite memory strategy τ^v that realizes $\text{aVal}(v)$ and $\text{acVal}(v)$. Note that since generalised safety/reachability games are prefix independent, values depend only on the current vertex, but not on the entire history.

We start with a finite memory strategy σ . If it is not already preadmissible, then it has witnesses of non-admissibility violating the desired property. Whether a history h is a witness of non-admissibility for a finite memory strategy σ depends only on the last vertex of h and the current state of σ . We now modify σ such that whenever σ is in a combination of vertex v and state s corresponding to a problematic witness of non-admissibility, the new strategy σ' moves to playing τ^v instead. The choices of v , s and τ^k ensure that σ' dominates σ .

The new strategy σ' may fail to be preadmissible, again, and we repeat the construction. Now any problematic history in σ' needs to enter the automaton for some τ^v at some point. By choice of τ^v , the history where τ^v has just been entered cannot be a witness of non-admissibility. It follows that a problematic history entering τ^v cannot end in v . Repeating the updating process for at most as many times as there are vertices in the game graph will yield a preadmissible finite memory strategy dominating σ . \blacktriangleleft

C

 Proofs omitted from Subsection 4.2

To complete the proof of Theorem 31, we need the following intermediary results:

► Lemma 38. *If h is not a witness of non-admissibility of σ , and not a witness of non-dominance of σ by τ , then h is not a witness of non-dominance of τ by σ .*

► Lemma 39. *Given an initialized game with initial vertex v_0 , the following holds: If for two strategies σ and τ it holds that for any maximal history h compatible with both, there is a prefix h' with $\text{aVal}(h', \sigma) = \text{aVal}(h', \tau)$ and $\text{cVal}(h', \sigma) = \text{cVal}(h', \tau)$, then $\text{aVal}(v_0, \sigma) = \text{aVal}(v_0, \tau)$ and $\text{cVal}(v_0, \sigma) = \text{cVal}(v_0, \tau)$.*

► **Lemma 40.** *Given an initialized game with initial vertex v_0 , the following holds: If σ is preadmissible and $\sigma \preceq \tau$, then $\text{aVal}(v_0, \sigma) = \text{aVal}(v_0, \tau)$ and $\text{cVal}(v_0, \sigma) = \text{cVal}(v_0, \tau)$.*

Proof. We show that the conditions of Lemma 39 are satisfied, which will imply our desired conclusion. Consider a maximal history h compatible with both σ and τ . First, assume that h is not a witness of non-admissibility of σ . Since $\sigma \preceq \tau$, by Lemma 23 h cannot be a witness of non-dominance of σ by τ , i.e. $\text{cVal}(h, \sigma) \leq \text{aVal}(h, \tau)$. By Lemma 38, it follows that h is not a witness of non-dominance of τ by σ either, i.e. $\text{cVal}(h, \tau) \leq \text{aVal}(h, \sigma)$. Put together, we have $\text{aVal}(h, \sigma) = \text{cVal}(h, \sigma) = \text{aVal}(h, \tau) = \text{cVal}(h, \tau)$.

It remains the case where h is a witness of non-admissibility of σ . Then by preadmissibility of σ , h has some prefix h' with $\text{aVal}(h', \sigma) = \text{aVal}(h')$ and $\text{cVal}(h', \sigma) = \text{acVal}(h')$. Since $\sigma \preceq \tau$, we must have $\text{aVal}(h', \sigma) \leq \text{aVal}(h', \tau)$, so it follows that $\text{aVal}(h', \sigma) = \text{aVal}(h', \tau)$, and then that $\text{cVal}(h', \tau) \leq \text{acVal}(h') = \text{cVal}(h', \sigma) \leq \text{cVal}(h', \tau)$, i.e. $\text{cVal}(h', \sigma) = \text{cVal}(h', \tau)$. ◀

Proof of Theorem 31. By Lemma 27 it suffices to prove the claim for preadmissible strategies (Definition 26). We thus start with a preadmissible finite memory strategy σ .

Preliminaries. Since we are working with prefix-independent outcomes and strategies realized by automata, we see that any of the values of σ at some history h depends only on the final vertex v of h and the state s the strategy σ is in after reading h . We can thus overload our notation to write $\text{aVal}(v, s)$ for $\text{aVal}(h, \sigma)$ and $\text{aVal}(v)$ for $\text{aVal}(h)$, and so on. In particular, whether some history h is a witness of non-admissibility of σ or not depends only on the final vertex v of h and the state s that σ is in after reading h . Let WNA be the set of such pairs (v, s) corresponding to non-admissibility witnesses. By the definition of preadmissibility, we cannot reach any $(v, s) \in \text{WNA}$ without first passing through some (v, s_v) with $\text{aVal}(v, s_v) = \text{aVal}(v)$ and $\text{cVal}(v, s_v) = \text{acVal}(v, s_v)$. By expanding the automaton if necessary (to remember where we were when first encountering some vertex), we can assume that for any $(v, s) \in \text{WNA}$ there is canonic choice of prior (v, s_v) .

► **Lemma 41.** *For any $(v, s) \in \text{WNA}$ and corresponding (v, s_v) we find that $\text{aVal}(v, s_v) = \text{aVal}(v, s) = \text{cVal}(v, s) < \text{cVal}(v, s_v)$.*

The construction. We now construct a parameterized automaton \mathcal{M} from σ that either realizes a single maximal strategy, or a maximal uniform chain. The parameterized automaton is identical to the one realizing σ everywhere except at the $(v, s) \in \text{WNA}$. In particular, if $\text{WNA} = \emptyset$, we are done. Otherwise, for each $(v, s) \in \text{WNA}$ we make the following modifications: If $\text{aVal}(v, s_v) \leq 0$, we modify the automaton to act in (v, s) as it does in (v, s_v) . If $\text{aVal}(v, s_v) > 0$, then we add green edges to let the automaton act in (v, s) as in (v, s_v) , and red edges to act as it would do originally.

Correctness. The comparison of the values lets us conclude via Lemma 23 that the parameterized automaton \mathcal{M} either realizes a single strategy dominating σ , or a uniform chain dominating σ .

It remains to argue that the strategy/uniform chain realized by \mathcal{M} is maximal. Let σ_n be the strategy where \mathcal{M} is initialized with $n \in \mathbb{N}$. Assume that $\tau \succ \sigma_n$, and let h be a witness of $\tau \not\preceq \sigma_n$ according to Lemma 23, i.e. satisfying $\text{cVal}(h, \tau) > \text{aVal}(h, \sigma_n)$. Since $\sigma_n \preceq \tau$, we have $\text{cVal}(h, \sigma_n) \leq \text{aVal}(h, \tau)$, so $\text{aVal}(h, \sigma_n) \leq \text{cVal}(h, \sigma_n) \leq \text{aVal}(h, \tau) \leq \text{cVal}(h, \tau)$ with one inequality being strict. In particular, h is a witness of non-admissibility of σ_n . By construction of \mathcal{M} the next move after h must be given by a red edge. This already implies that if \mathcal{M} realizes a single strategy, then that strategy is maximal.

Let m be the size of the parameterized automaton \mathcal{M} , let t be the size of the automaton realizing τ , and $N = mt + 1$.

► **Lemma 42.** *At any maximal history compatible with σ_N and τ , σ_N will follow a green or black edge next.*

Proof. Assume there were such a history hv compatible with both σ_N and τ where σ_N is about to apply a red edge, being in state s . If the combination (v, s) has been reached more than t times during hv , then it has to hold that on histories extending hv , τ always acts at v as \mathcal{M} does following the green edge at (v, s) , for τ cannot count up to $t + 1$ (in particular, h is maximal for being compatible with τ and σ_N). It follows that $\text{aVal}(hv, \tau) \leq 0$. Let $h'v$ be a prefix of this form of hv compatible with σ_n not ending in a red edge (this exists, since $n > m$). Then $\text{aVal}(h'v, \tau) \leq 0$, and since $\tau \succeq \sigma_n$, $\text{aVal}(h'v, \sigma_n) = \text{aVal}(h'v, \tau) \leq 0$. But then when constructing \mathcal{M} , we would not have placed red and green edges at (v, s_v) , leading to a contradiction. Thus, at any maximal history compatible with σ_N and τ , σ_N will follow a green or black edge next.

If the combination (v, s) has been visited at most t times during hv , then there has to be some other pair of counter access state s' and vertex v' which was reached more often than t times during hv by the pigeon hole principle (for since σ_N is about to follow a red edge, it has reached a counter access state at least $N = mt + 1$ many times), with σ_N taking the green edge there. Again, by the same reasoning as above, τ always follows the green edge at the corresponding histories, leading to the conclusion that the antagonistic value obtained by τ there is 0, and ultimately a contradiction to s' being created as a counter access state when constructing \mathcal{M} . ◀

If τ is part of a chain $(\tau_i)_{i \in \mathbb{N}}$ with $(\sigma_i)_{i \in \mathbb{N}} \sqsubseteq (\tau_i)_{i \in \mathbb{N}}$, then τ and σ_N have a common upper bound τ' . We proceed to show that this suffices to conclude $\tau \preceq \sigma_N$. This completes our argument, since by induction it follows that if $(\sigma_n)_{n \in \mathbb{N}} \sqsubseteq (\tau_n)_{n \in \mathbb{N}}$, then also $(\tau_n)_{n \in \mathbb{N}} \sqsubseteq (\sigma_n)_{n \in \mathbb{N}}$.

► **Lemma 43.** *If τ and σ_N have common upper bound τ' , then $\tau \preceq \sigma_N$.*

Proof. We proceed by ruling out all candidates for witnesses of non-dominance of τ by σ_N , and conclude our claim by Lemma 23. Any candidate is a maximal history h compatible with both σ_N and τ .

Case 1. Either h is not compatible with τ' , or $\tau'(h) \neq \sigma_N(h)$.

If h is not compatible with τ' , then h has a longest prefix h' compatible with τ' . If h is compatible with τ' , but $\tau'(h) \neq \sigma_N(h)$, we set $h' = h$. By Lemma 42, h' cannot be a witness of non-admissibility of σ_N , and by Lemma 23 it cannot be a witness of non-dominance of σ_N by τ' , since $\sigma_N \preceq \tau'$. Lemma 38 then gives us that h' is not a witness of non-dominance of τ' by σ_N , i.e. $\text{cVal}(h', \tau') \leq \text{aVal}(h', \sigma_N)$. Together with $\sigma_N \preceq \tau'$ we get that $\text{aVal}(h', \sigma_N) = \text{cVal}(h', \sigma_N)$. Since h is compatible with σ_N and extends h' , it follows that $\text{aVal}(h', \sigma_N) = \text{aVal}(h, \sigma_N) = \text{cVal}(h, \sigma_N)$. Since $\tau \preceq \tau'$, it follows that $\text{cVal}(h', \tau) \leq \text{cVal}(h', \tau') = \text{aVal}(h', \sigma_N)$. Since h is compatible with τ and extends h' , it follows that $\text{cVal}(h, \tau) \leq \text{cVal}(h', \tau) \leq \text{aVal}(h', \sigma_N) = \text{aVal}(h, \sigma_N)$, i.e. that h is not a witness of non-dominance of τ by σ_N .

Case 2. h is compatible with τ' and $\tau'(h) = \sigma_N(h)$.

Consider the subgame starting after that move. Since we have chosen N sufficiently big, in this subgame it is impossible for σ_N to pass through a red edge without previously passing through a green edge at the same vertex. By construction, this ensures that σ_N is still preadmissible in this subgame. Since reaching the subgame is compatible

10:22 Dominance Between Chains of Strategies

with τ' and σ_N , restricting to this subgame, we still have that $\sigma_N \preceq \tau'$. Thus, we can apply Lemma 40 to the subgame, and conclude that $\text{aVal}(h, \tau') = \text{aVal}(h, \sigma_N)$ and $\text{cVal}(h, \tau') = \text{cVal}(h, \sigma_N)$. Since h cannot be a witness of non-dominance of τ by τ' , it holds that $\text{cVal}(h, \tau) \leq \text{aVal}(h, \tau') = \text{aVal}(h, \sigma_N)$. Thus, h is not a witness of non-dominance of τ by σ_N either. \blacktriangleleft

D Proofs omitted from Subsection 4.3

The proof of Proposition 34.1 is based on the following auxiliary Lemma, whose demonstration relies on the study of the loops that appear in witnesses of non dominance.

► **Lemma 44.** *Let \mathcal{G} be a generalised safety/reachability game, let \mathcal{M} be a parametrized automaton over the game graph of \mathcal{G} , and let $(T_n)_{n \in \mathbb{N}}$ be the sequence of finite-memory strategies realized by \mathcal{M} . Then for every pair of integers $n_1, n_2 > |G||\mathcal{M}|$ satisfying $T_{n_1} \not\preceq T_{n_2}$, there exists $0 < k \leq |G||\mathcal{M}|$ such that for every $i \in \mathbb{N}$, $T_{n_1+(i-1)k} \not\preceq T_{n_2+(i-1)k}$.*

Proof of Proposition 34.1. Let \mathcal{G} be a generalised safety/reachability game, and let \mathcal{S} be a parametrized automaton over the game graph of \mathcal{G} . We denote by $(S_n)_{n \in \mathbb{N}}$ the sequence of finite-memory strategies realized by \mathcal{S} . Let $N_{\preceq} = |G||\mathcal{S}|$.

Let $U_{\mathcal{S}}$ denote the set composed of the integers n satisfying $S_n \not\preceq S_{n+1}$. It is clear that if $U_{\mathcal{S}}$ is not empty, then $(S_n)_{n \in \mathbb{N}}$ is not a chain. Conversely, if $U_{\mathcal{S}}$ is empty, then $(S_n)_{n \in \mathbb{N}}$ is a chain, since for every pair of integers $n_1 < n_2$, we have $S_{n_1} \preceq S_{n_1+1} \preceq \dots \preceq S_{n_2}$. Let us suppose, towards building a contradiction, that the minimal element m of $U_{\mathcal{S}}$ is strictly greater than N_{\preceq} . Then, we obtain from Lemma 44 that there exists an integer $k > 0$ such that $S_{m-k} \not\preceq S_{m-k+1}$ by setting $i = 0$. This contradicts the minimality of m . As a consequence, $m \leq N_{\preceq}$. This proves that $(S_n)_{n \in \mathbb{N}}$ is a chain if and only if $S_i \preceq S_{i+1}$ for every $1 \leq i \leq N_{\preceq}$. \blacktriangleleft