Beyond admissibility:
Dominance between chains of strategies

Nicolas Basset
Université Grenoble Alpes, Grenoble, France
nicolas.basset1@univ-grenoble-alpes.fr

Ismaël Jecker
Université Libre de Bruxelles, Brussels, Belgium
ismael.jecker@ulb.ac.be

Arno Pauly
Swansea University, Swansea, UK
Arno.M.Pauly@gmail.com

Jean-François Raskin
Université Libre de Bruxelles, Brussels, Belgium
jraskin@ulb.ac.be

Marie Van den Bogaard
Université Libre de Bruxelles, Brussels, Belgium
marie.van.den.bogaard@ulb.ac.be

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.

1 Introduction

The canonical model to formalize the reactive synthesis problem are two-player win/lose perfect information games played on finite (directed) graphs [20, 1]. In recent years, more general objectives and multiplayer games have been studied (see e.g. [16] 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, 11], 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 10. However, Corollary 16 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 hyperarithmetic ones (Corollary 18).

We test the abstract notion in the concrete setting of generalised safety/reachability games (Definition 20). 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 27), 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 30).

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 31). It is also decidable in PTime whether the chain realized by one parameterized automaton dominates the chain realized by another (Theorem 32).

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 [14]. 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 [19] for an overview. For the role of admissibility for synthesis, we refer to [11]. Our Subsection 3.1 involves an investigation of cofinal chains in certain quasi-ordered sets. A similar theme (but with a different focus) is present in [21].

Full version A full version including the proofs omitted here is available on the arXiv [3].

2 Background

2.1 Games on finite graphs

We presume familiarity with the basic notions related to games played on finite graphs. Our games have payoff functions (denoted by p_i). The set of strategies of player i is denoted by Σ_i. Given some strategy profile (σ, τ) the induced infinite path through the graph is denoted by $\mathbf{Out}(\sigma, \tau)$.

Dominance relation. In order to compare the "efficiency" of different strategies 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 $\sigma \preceq_{q_0} \sigma'$, where q_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 $\sigma \prec_{q_0} \sigma'$, where q_0 is the initial state of $\mathcal{G}$. Strategies that are not dominated by any other strategies are called *admissible*:

Antagonistic and Cooperative Values In order to compare 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 best payoff Player i can ensure with σ from h in the worst-case scenario (ie, the best possible payoff to ensure against the other players playing *antagonistically*), as well as the best possible payoff Player i can hope for with σ from h in the best scenario (ie, in case the other players are playing *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 $\mathrm{aVal}_i(\mathcal{G}, h, \sigma) := \inf_{\tau \in \Sigma_{-i}} p_i(\mathbf{Out}_h(\sigma, \tau))$ and $\mathrm{cVal}_i(\mathcal{G}, h, \sigma) := \sup_{\tau \in \Sigma_{-i}} p_i(\mathbf{Out}_h(\sigma, \tau))$.

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 $\mathrm{aVal}_i(\mathcal{G}, h) := \sup_{\sigma \in \Sigma_i} \mathrm{aVal}_i(\mathcal{G}, h, \sigma)$ and the *cooperative value of* h for Player i as $\mathrm{cVal}_i(\mathcal{G}, h) := \sup_{\sigma \in \Sigma_i} \mathrm{cVal}_i(\mathcal{G}, h, \sigma)$. 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 $acVal_i(\mathcal{G}, h) := \sup\{\mathrm{cVal}_i(\mathcal{G}, h, \sigma) \mid \sigma \in \Sigma_i \text{ and } \mathrm{aVal}_i(\mathcal{G}, h, \sigma) \geq \mathrm{aVal}_i(\mathcal{G}, h)\}$. From now on, we will omit to precise $\mathcal{G}$ when it is clear from the context.

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 $\mathrm{aVal}_i(h, \sigma) = \mathrm{aVal}(h)$, *and* a strategy $\sigma' \in \Sigma_i$ such that $\mathrm{cVal}_i(h, \sigma') = \mathrm{cVal}(h)$. In other words, at every history h, Player i has a strategy that ensures the payoff $\mathrm{aVal}_i(h)$, and a strategy that allows the other players to cooperate to yield a payoff of $cVa_i(h)$. In [10], authors have shown that Player i has admissible strategies in every game that is well-formed for Player i.

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 $\mathrm{aVal}_i, \mathrm{cVal}_i, acVal_i, p_i$, etc..

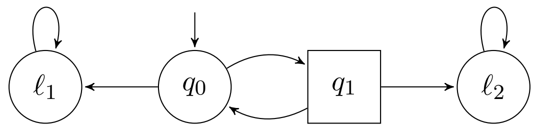

Figure 1: The *Help-me?*-game

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

$$p(\rho) = \begin{cases} 0 \text{ if } \rho = (q_0q_1)^\omega, \\ 1 \text{ if } \rho = (q_0q_1)^n q_0 \ell_1^\omega \text{ where } n \in \mathbb{N}, \\ 2 \text{ if } \rho = (q_0q_1)^n \ell_2^\omega \text{ where } n \in \mathbb{N}. \end{cases}$$

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 q_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 q_1 as the successor vertex at the first k visits of q_0, and ℓ_1 at the $k+1$-th visit, while s_ω denotes the strategy that prescribes q_1 at every visit of q_0. It holds that $s_k \prec s_{k+1}$: Indeed, for all $\tau \in \Sigma_{-i}$, if $p(s_k, \tau) = 2$, then there exists $j \leq k$ such that $\tau((q_0q_1)^j) = \ell_2$. As s_k and s_{k+1} agree up to $(q_0q_1)^k q_0$, we have that $\mathbf{Out}(s_{k+1}, \tau) = (q_0q_1)^j \ell_2^\omega = \mathbf{Out}(s_k, \tau)$, thus $p(s_{k+1}, \tau) = 2$ as well. Furthermore, consider a strategy τ such that $\tau((q_0q_1)^j) = q_0$ for all $j \leq k$ and $\tau((q_0q_1)^{k+1}) = \ell_2$. Then $p(s_k, \tau) = 1$ while $p(s_{k+1}, \tau) = 2$. Finally, consider the strategy τ such that $\tau((q_0q_1)^k) = q_0$ 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 q_1 yields a payoff of 1 against strategy s_k but 2 against strategy s_ω. Thus, $s_\omega \not\preceq s_k$ for any $k \in \mathbb{N}$.

It is also easy to see that $s_k \not\preceq s_\omega$ for any $k \in \mathbb{N}$: Let $\tau \in \Sigma_{-i}$ be such that $\tau((q_0q_1)^k) = q_0$ for all $k \in \mathbb{N}$. Then $p(s_k, \tau) = 1 > 0 = p(s_\omega, \tau) = 0$. 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. [18]).

A *linear order* is a total, transitive and antisymmetric relation. A linearly ordered set $(R, \prec)$ is a *well-order*, if every subset of R has a minimal element w.r.t. $\prec$. 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 *quasi order* is a transitive and reflexive relation. Let $(X, \preceq)$ be a quasi-ordered set. 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 quasi-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. [18]).

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 quasi-order of increasing chains in a given quasiorder $(X, \preceq)$. 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:

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

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 $(\text{IC}, \sqsubseteq)$ is a quasiorder 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 $\mathrm{cof}((y_\gamma)_{\gamma<\delta})$ 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. $\mathrm{cof}((y_\gamma)_{\gamma<\delta})$ 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 $(x_n)_{n<\omega}$.

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

Lemma 8. The following are equivalent:

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

Let us illustrate the problem of extending Lemma 8 by an example:

Example 9 ([18, 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 $x^\gamma_n = (\gamma, n)$, this corresponds to the case $\alpha = \omega$, $\delta = \omega_1$ in Lemma 8. 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 10. If all increasing chains of elements in IC (ie, increasing chains of increasing chains of elements of $(X, \preceq)$) have a countable number of elements, then for every $A \in \mathrm{IC}$ there exists a maximal $B \in \mathrm{IC}$ with $A \sqsubseteq B$.

Proof. We first argue that Condition 2 in Lemma 8 is vacuously true. As all increasing chains in IC are countable, the only possible value $\delta > 1$ for $\delta = \mathrm{cof}(((x^\gamma_\beta)_{\beta<\alpha})_{\gamma<\delta})$ 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 8, 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 10 with the simpler requirement that $(X, \preceq)$ has only countable increasing chains:

Example 11. Let $X = \omega_1 \times \omega$, and let $(\alpha, n) \prec (\beta, m)$ iff $\alpha \leq \beta$ and $n < m$. Then $(X, \preceq)$ has only countable increasing chains, but IC still has the chain of chains given by $x^\gamma_n = (\gamma, n)$ as in Example 9.

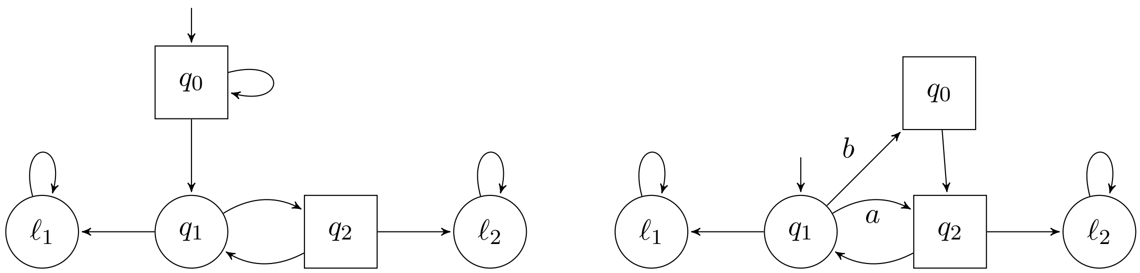

(a) A variant of the *Help-me?* game with an extra loop

(b) A variant of the *Help-me?* game with two paths from q_0 to q_1

Figure 2: Variants of the *Help-me?* game

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 10. In fact, we will see that the chain of chains of uncountable length we construct is not below any maximal chain.

Example 12. 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 $f : \mathbb{N} \to \mathbb{N} \cup \{\infty\}$ describing how often the protagonist is willing to repeat the second loop (between q_1 and q_2) given the number of repetitions the antagonist made in the first loop (at q_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 13. Let $\mathbb{N}^\mathbb{N}$ denote the set of functions $f : \mathbb{N} \to \mathbb{N}$. 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 14. There is an embedding of $(\mathbb{N}^\mathbb{N}, \leq)$ into the strategies of the game in Example 12 ordered by dominance such that no strategy in the range of embedding is dominated by a strategy outside the range of the embedding.

Proposition 15 ([1])**.** For every chain $(f_n)_{n\in\mathbb{N}}$ in $(\mathbb{N}^\mathbb{N}, \leq)$ there exists a chain of chains $((f^\alpha_n)_{n<\omega})_{\alpha<\omega_1}$ of length ω_1 with $(f^0_n)_{n<\omega} \sqsupseteq (f_n)_{n<\omega}$.

Corollary 16. The game in Example 12 has uncountably long chains of chains not below any maximal chains.

Proof. Combine Observation 14 and Proposition 15. □

3.3 Chains over countable quasiorders $(X, \preceq)$

Our proof of Proposition 15 crucially relied on functions of type $f : \mathbb{N} \to \mathbb{N}$ 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

[1]This result is adapted from an answer by user *Deedlit* on `math.stackexchange.org` [15].

computable strategies often makes sense. Many games played on finite graphs are even finite-memory determined (see [17] 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 15. The reason is that the cardinality of the length of a chain of chains cannot exceed that of the underlying quasiorder $(X, \preceq)$:

Proposition 17. For any increasing chain $((x^\gamma_\beta)_{\beta<\alpha})_{\gamma<\delta}$ in $\mathrm{IC}(X, \preceq)$ we find that $|\delta| \leq |X|$.

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

Corollary 18. If $(X, \preceq)$ is countable, then any increasing chain is maximal or below a maximal chain.

Proof. Proposition 17 shows that Theorem 10 applies. □

Example 19. We return to the *Help-me?* game (Example 1, Figure 1). 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 18 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 20. 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 [17] (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 $\mathrm{cVal}(h, \sigma') = \mathrm{cVal}(h, \sigma)$ and $\mathrm{aVal}(h, \sigma') = \mathrm{aVal}(h, \sigma)$. We shall thus restrict our attention to finite memory strategies, of which there are only countably many. We then obtain immediately from Corollary 18:

Corollary 21. 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 22. 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 an 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 $\text{cVal}(h, \sigma_1) > \text{aVal}(h, \sigma_2)$.

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, $\text{cVal}(h, \sigma_1)$) is less than the one σ_2 can ensure (that is, $\text{aVal}(h, \sigma_2)$), 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 [10].

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 23. 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 $\text{cVal}(h, \sigma_2) > \text{aVal}(h, \sigma_1)$, $\text{cVal}(h, \sigma_3) > \text{aVal}(h, \sigma_2)$, $\text{cVal}(h, \sigma_1) \leq \text{aVal}(h, \sigma_2)$ and $\text{cVal}(h, \sigma_2) \leq \text{aVal}(h, \sigma_3)$. It follows that $\text{aVal}(h, \sigma_1) < \text{aVal}(h, \sigma_3)$ and $\text{cVal}(h, \sigma_1) < \text{cVal}(h, \sigma_3)$. 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. □

Similarly, we can also extract witnesses for a strategy to be non-maximal (non-admissible or strictly dominated). 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:

Lemma 24. 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 $\mathrm{aVal}(h,\sigma) \leq \mathrm{cVal}(h,\sigma) \leq \mathrm{aVal}(h) \leq acVal(h)$ where at least one inequality is strict.

Definition 25. Call a history h as in Lemma 24 a non-admissibility witness for σ. Call σ preadmissible, if for every non-admissibility witness hv of σ we find that $h = h'vh''$ with $\mathrm{aVal}(h'v,\sigma) = \mathrm{aVal}(h'v)$ and $\mathrm{cVal}(h'v,\sigma) = acVal(h'v)$.

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 26. 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

Definition 27. Let a *parameterized automaton* be a Mealy automaton that in addition can access a single counter. The special counter-access-states have an outgoing *green* transitions, which can only be taken if the counter value is positive, and decrement the counter. They also have *red* transitions which are taken if the counter value is 0. The other transitions will be called *black* transitions. The parameterized automata are parameterized by the initial value of the counter.

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 28. Let a chain $(\sigma_n)_{n\in\mathbb{N}}$ of regular strategies be called a *uniform* chain if there is a parameterized automaton M that realizes σ_n if the counter is initialized with the value n.

Example 29. 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 edge corresponds to the transition to take when the counter value is greater than 0 and should be decremented, while the red 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:

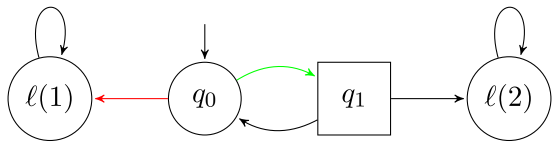

Figure 3: Product of the *Help-me?* game with parameterized automaton with a single memory state realizing $(s_k)_{k\in\mathbb{N}}$

Theorem 30. In a generalised safety/reachability game, every dominated finite memory strategy is dominated by a maximal finite memory strategy or by a maximal[2] uniform chain.

Proof. By Lemma 26 it suffices to prove the claim for preadmissible strategies (Definition 25). We thus start with a preadmissible finite memory strategy σ.

By the prefix-independence of generalised safety/reachability games, for any combination of vertex v in the game and state s in the automaton realizing σ, either a history ending in v and state s is a witness for non-admissibility of σ or not. Let N be the set of such pairs corresponding to non-admissibility witnesses. By the definition of preadmissibility, we cannot reach any $(v, s) \in N$ without first passing through some (v, s^i_v) with $\mathrm{aVal}(v, s^i_v) = \mathrm{aVal}(v)$ and $\mathrm{cVal}(v, s^i_v) = acVal(v, s^i_v)$. By expanding the automaton if necessary (to remember where we were when first encountering some vertex), we can assume that for any $(v, s) \in N$ there is canonic choice of prior (v, s^i_v).

We now construct either a parameterized automaton from σ that either realizes a single maximal strategy, or a maximal uniform chain. If N is empty, we are done. Otherwise, consider $(v, s) \in N$ and the corresponding (v, s^i_v), and compare the associated values: Since the antagonist can reach (v, s) from (v, s^i_v), it has to hold that $\mathrm{aVal}(v, s^i_v) \leq \mathrm{aVal}(v, s) \leq \mathrm{cVal}(v, s) \leq \mathrm{cVal}(v, s^i_v)$. By choice of (v, s^i_v), we have $\mathrm{aVal}(v, s) \leq \mathrm{aVal}(v) = \mathrm{aVal}(v, s^i_v)$, and thus $\mathrm{aVal}(v, s^i_v) = \mathrm{aVal}(v, s)$. Since $(v, s) \in N$, we see that even $\mathrm{aVal}(v, s^i_v) = \mathrm{aVal}(v, s) = \mathrm{cVal}(v, s) < \mathrm{cVal}(v, s^i_v)$ holds by Lemma 24.

If $\mathrm{aVal}(v, s^i_v) \leq 0$, we modify the automaton to act in (v, s) as it does in (v, s^i_v). If $\mathrm{aVal}(v, s^i_v)$, then we add green edges to let the automaton act in (v, s) as in (v, s^i_v), and red edges to act as it would do originally. The comparison of the values lets us conclude via Lemma 22 that the parameterized automaton $\mathcal{M}$ either realizes a single strategy dominated σ, or a uniform chain dominating σ.

It remains to argue that the strategy/uniform chain realized by $\mathcal{M}$ is maximal. Due to reasons of space, we omit that reasoning here. □

4.3 Algorithmic properties

Theorem 31. Given a generalised safety/reachability game and a parameterized automaton, we can decide whether the automaton realizes a uniform chain of strategies, and whether it realizes an increasing chain of strategies.

[2]There are two different potential meanings of *maximal uniform chain*: It could be a uniform chain not dominated by another uniform chain, or a uniform chain that is not dominated by any chain comprised of finite memory strategies. The latter is the stronger property, and our proof establishes it.

Theorem 32. Given a generalised safety/reachability game and two parameterized automata realizing uniform chains of strategies, we can decide whether the chain realized by the first is dominated by the one from the second.

The proofs of the preceding theorems have a common structure: Proposition 34 allows us to reduce questions about chains to questions about single strategies, which can be decided by applying Lemma 33.

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 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 22, an 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, corresponding to $\text{cVal}(h, \sigma)$, and a, corresponding to $\text{aVal}(h, \sigma')$.
- 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.
- 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.

□

The other main ingredient of our algorithms are bounds on the parameters:

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 $N_{\preceq} = |G||\mathcal{S}|$.
   Then $(S_n)_{n\in\mathbb{N}}$ is a chain if and only if $T_i \preceq T_{i+1}$ for every $1 \leq i \leq N_{\preceq}$.
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 $M \not\preceq T_{N_T}$.
3. Let $N_{\mathcal{S}} = |G||\mathcal{S}|(2|\mathcal{T}| + 1)$, 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 $S_{N_{\mathcal{S}}} \not\preceq (T_n)_{n\in\mathbb{N}}$.

5 Conclusions and outlook

We have observed that admissibility is lacking as a rationality criterion for infinite sequential games with quantitative payoffs. Our primary counterexample 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 in a TCS setting. A more in-depth exploration of the game-theoretic merits of such a notion of rationality based on chains of strategies is left for the future.

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. Moreover, the generalization of our results from generalized safety/reachability games to games with ω-regular objectives seems achievable - our proofs make only very limited use of the special features of the former. Both these endeavours could benefit from a better understanding of parameterized automata in general.

References

[1] Martín Abadi, Leslie Lamport & Pierre Wolper (1989): *Realizable and Unrealizable Specifications of Reactive Systems*. In: *ICALP'89*, *LNCS* 372, Springer.

[2] Nicolas Basset, Gilles Geeraerts, Jean-François Raskin & Ocan Sankur (2017): *Admissibility in Concurrent Games*. *CoRR* abs/1702.06439. Available at `http://arxiv.org/abs/1702.06439`.

[3] Nicolas Basset, Ismaël Jecker, Arno Pauly, Jean-François Raskin & Marie Van den Boogard (2018). *Beyond admissibility: Dominance between chains of strategies*. arXiv 1805.11608. Available at `https://arxiv.org/abs/1805.11608`.

[4] Dietmar Berwanger (2007): *Admissibility in Infinite Games*. In: *STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science, Aachen, Germany, February 22-24, 2007, Proceedings*, pp. 188–199. Available at `https://doi.org/10.1007/978-3-540-70918-3_17`.

[5] Endre Boros, Khaled Elbassioni, Vladimir Gurvich & Kazuhisa Makino (2012): *On Nash equilibria and improvement cycles in pure positional strategies for Chess-like and Backgammon-like n-person games*. *Discrete Mathematics* 312(4), pp. 772 – 788.

[6] Endre Boros, Vladimir Gurvich & Emre Yamangil (2013): *Chess-Like Games May Have No Uniform Nash Equilibria Even in Mixed Strategies*. *Game Theory* .

[7] Romain Brenguier, Lorenzo Clemente, Paul Hunter, Guillermo A. Pérez, Mickael Randour, Jean-François Raskin, Ocan Sankur & Mathieu Sassolas (2016): *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*, *Lecture Notes in Computer Science* 9618, Springer, pp. 3–23.

[8] Romain Brenguier, Arno Pauly, Jean-François Raskin & Ocan Sankur (2017): *Admissibility in Games with Imperfect Information (Invited Talk)*. In: Roland Meyer & Uwe Nestmann, editors: *28th International Conference on Concurrency Theory (CONCUR 2017)*, *LIPIcs* 85, Schloss Dagstuhl, pp. 2:1–2:23. Available at `http://drops.dagstuhl.de/opus/volltexte/2017/7806`.

[9] Romain Brenguier, Guillermo A. Pérez, Jean-François Raskin & Ocan Sankur (2016): *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*, *LIPIcs* 65, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 42:1–42:14.

[10] Romain Brenguier, Guillermo A. Pérez, Jean-François Raskin & Ocan Sankur (2016): *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*, pp. 42:1–42:14. Available at `https://doi.org/10.4230/LIPIcs.FSTTCS.2016.42`.

[11] Romain Brenguier, Jean-François Raskin & Ocan Sankur (2017): *Assume-admissible synthesis*. *Acta Inf.* 54(1), pp. 41–83. Available at `https://doi.org/10.1007/s00236-016-0273-2`.

[12] Romain Brenguier, Jean-François Raskin & Mathieu Sassolas (2014): *The complexity of admissibility in Omega-regular games*. In: *CSL-LICS '14, 2014*, ACM. Available at `http://doi.acm.org/10.1145/2603088.2603143`.

[13] Thomas Brihaye, Veroniqué Bruyère & Julie De Pril (2010): *Equilibria in quantitative reachability games*. In: *Proc. of CSR*, *LNCS* 6072, Springer.

[14] Marco Faella (2009): *Admissible Strategies in Infinite Games over Graphs*. In: *MFCS 2009*, *Lecture Notes in Computer Science* 5734, Springer, pp. 307–318.

[15] Deedlit (https://math.stackexchange.com/users/21465/deedlit). *Height of a certain order*. Mathematics Stack Exchange. Available at `https://math.stackexchange.com/q/2436983`. URL:https://math.stackexchange.com/q/2436983 (version: 2017-09-20).

[16] Orna Kupferman (2016): *On High-Quality Synthesis*. In: S. Alexander Kulikov & J. Gerhard Woeginger, editors: *11th International Computer Science Symposium in Russia, CSR 2016*, Springer International Publishing, pp. 1–15.

[17] Stéphane Le Roux & Arno Pauly (201X): *Extending finite memory determinacy*. *Information and Computation* Available at `https://www.sciencedirect.com/science/article/pii/S0890540118300270`.

[18] George Markowsky (1976): *Chain-complete posets and directed sets with applications*. *Algebra Universalis* 6, pp. 53–68. Available at `http://dx.doi.org/10.1007/BF02485815`.

[19] Arno Pauly (2016): *The Computational Complexity of Iterated Elimination of Dominated Strategies*. *Theory of Computing Systems* , pp. 52–75.

[20] Amir Pnueli & Roni Rosner (1989): *On the Synthesis of a Reactive Module*. In: *POPL*, pp. 179–190. Available at `http://doi.acm.org/10.1145/75277.75293`.

[21] Wang Shang-Zhi & Li Bo-Yu (1984): *On the minimal cofinal subsets of a directed quasi-ordered set*. *Discrete Mathematics* 48(2), pp. 289 – 306. Available at `http://www.sciencedirect.com/science/article/pii/0012365X84901894`.