


Optimally Resilient Strategies in Pushdown Safety Games

Daniel Neider

Max Planck Institute for Software Systems, 67663 Kaiserslautern, Germany
neider@mpi-sws.org

Patrick Totzke 

University of Liverpool, Liverpool L69 3BX, United Kingdom
totzke@liverpool.ac.uk

Martin Zimmermann 

University of Liverpool, Liverpool L69 3BX, United Kingdom
martin.zimmermann@liverpool.ac.uk

Abstract

Infinite-duration games with disturbances extend the classical framework of infinite-duration games, which captures the reactive synthesis problem, with a discrete measure of resilience against non-antagonistic external influence. This concerns events where the observed system behavior differs from the intended one prescribed by the controller. For games played on finite arenas it is known that computing optimally resilient strategies only incurs a polynomial overhead over solving classical games.

This paper studies safety games with disturbances played on infinite arenas induced by pushdown systems. We show how to compute optimally resilient strategies in triply-exponential time. For the subclass of safety games played on one-counter configuration graphs, we show that determining the degree of resilience of the initial configuration is PSPACE-complete and that optimally resilient strategies can be computed in doubly-exponential time.

2012 ACM Subject Classification Theory of Computation → Automata over infinite objects

Keywords and phrases Controller Synthesis, Infinite Games, Resilient Strategies, Pushdown Games

1 Introduction

Infinite games on finite arenas are a popular approach to the synthesis of reactive controllers from logical specifications. Originally proposed by Büchi and Landweber in 1969 [7], many variations of this classical framework have been studied, including stochastic games [12], games with partial information [14], games with delays [20], and games over infinite arenas such as pushdown graphs [41] and automatic structures [27, 28]. Other variations of this framework stem from the desire to synthesize controllers that exhibit certain user-desired properties. Examples of such properties range from controllers that need to achieve their task, e.g., reaching a goal, as quickly as possible [8] to controllers that are “robust” or “resilient” with respect to the environment in which they are deployed [3, 26, 4, 21, 25, 36, 37, 38]. Furthermore, infinite games have a plethora of applications in logic, automata theory and verification beyond the synthesis of reactive controllers. In this paper, we are concerned with the synthesis application and study infinite games with so-called *unmodeled intermittent disturbances* [13] played on configuration graphs of pushdown machines (pushdown graphs).

Pushdown graphs are finitely represented infinite graphs, typically the simplest class of such graphs one studies. Despite being conceptually simple, they have natural applications in program analysis, static code analysis, and compiler optimization [30, 31] due to their ability to capture recursion, e.g., the call stack of a procedural program. Furthermore, pushdown graphs are known to be well-behaved, and many problems on pushdown graphs are decidable (see, e.g., [5, 32, 33, 35]). In particular, Walukiewicz showed that solving parity games played

on pushdown graphs is EXPTIME-complete [41], paving the way for effective synthesis of *recursive* controllers. Also, Walukiewicz’s result started a long and fruitful line of work on games on pushdown graphs [8, 9, 10, 24, 34]. Of particular interest is the special case of games on configuration graphs of one-counter machines, i.e., pushdown machines with a single stack symbol, which is known to be PSPACE-complete [34, 23].

Games with unmodeled intermittent disturbances were originally introduced by Dallal, Neider, and Tabuada [13] to synthesize resilient controllers. The observation underlying this type of infinite game is that modeling the real-world environment of a controller in sufficiently great detail is often extremely challenging, either because parts of the environment are unknown or because simulating the environment is costly. Moreover, even if a high-resolution model of the environment is available, the resulting games often become prohibitively large. To alleviate this serious obstacle, Dallal, Neider, and Tabuada proposed to augment classical games with what they call unmodeled intermittent disturbances (in the following just called *disturbances* for the sake of brevity). Intuitively, such disturbances modify the outcome of a control action, thus modeling that the intended action of the controller did not have the desired consequences. Note, however, that disturbances are not under the control of the environment and, thus, are not antagonistic. Similarly, one does not consider the occurrence of disturbances as random events, as coming up with an appropriate stochastic error model is typically hard. Instead, the reader should understand them as rare events, such as a robot arm failing to grab an object due a physical phenomenon that has not been fully modeled.

The original work of Dallal, Neider, and Tabuada [13] provides a method to compute *optimally resilient strategies* for safety games over finite arenas, which intuitively are winning strategies that can tolerate as many disturbances as possible. In follow-up work, Neider, Weinert, and Zimmermann [29] have shown that computing optimally resilient strategies in finite arenas only incurs a polynomial overhead over solving classical games (under some mild assumptions on the winning condition), i.e., whenever a class of games is solvable without disturbances, then it is also solvable with disturbances. In particular, they have developed an algorithm that is effective for all standard winning conditions such as Rabin, Muller, and parity. Note, however, that both approaches crucially rely on the arena being finite.

The natural question, which we address here, is how to compute optimally resilient strategies for games on infinite arenas. As this is a very ambitious goal in its full generality, we restrict ourselves here to the setting of *safety games* played on *pushdown graphs*.¹

As argued before, pushdown games are a natural starting point for investigating effective algorithms for games on infinite graphs, and safety specifications are a fundamental class of specifications in practice [15]. While this setting might seem restrictive, recall that both the EXPTIME-hardness of solving pushdown games [41] and the PSPACE-hardness of solving one-counter games [34] already hold for the safety condition. Thus, the complexity of solving pushdown games stems from the transition from finite to infinite graphs, not from the expressiveness of the winning condition. The setting we consider here is still expressive enough to model interesting applications such as reasoning about exception handling in recursive programs. Here, one is interested in determining how many exceptions the program can tolerate while still satisfying a given specification.

To capture the optimization aspect of the problem at hand, we re-use Neider, Weinert,

¹ Some of our results do carry over to other winning conditions, such as reachability and parity, or do not require the underlying arena to be a pushdown graph. If this is the case, we present our arguments and state our results as general as possible. Also, we discuss the additional challenges one has to overcome to generalize all our results to reachability and parity conditions.

and Zimmermann’s notion of *resilience values* [29], which assigns to every vertex v of the arena an ordinal $r_{\mathcal{G}}(v) \leq \omega + 1$, where \mathcal{G} denotes the game in question and ω is the first infinite ordinal. Intuitively, $r_{\mathcal{G}}(v)$ denotes how many disturbances can be tolerated by an optimally resilient strategy from v . This value can be $k \in \omega$ ($k - 1$ disturbances can be tolerated, but not k), ω (finitely many disturbances can be tolerated, but not infinitely many), or $\omega + 1$ (infinitely many disturbances can be tolerated). When moving from finite to infinite arenas, however, various conceptual and technical complications arise, which make computing the resilience values of vertices and, by extension, resilient strategies challenging.

For instance, safety games over infinite arenas no longer guarantee the existence of optimally resilient strategies, i.e., in an infinite arena, one does not necessarily have a strategy that can tolerate an arbitrary finite number of disturbances from a vertex with resilience ω . Instead one has, for every $k \in \omega$, a strategy that can tolerate k disturbances, but not $k + 1$.

Another complication is the fact that it is no longer possible to globally bound the finite resilience values in infinite arenas. In contrast, in the case of finite arenas, the number of vertices is a trivial bound on the finite resilience values [29]. Hence, fixed-point algorithms like the ones devised for finite arenas [13, 29] and algorithms based on exhaustive search do not necessarily terminate.

Our Contributions In the rest of this paper, we study resilience in pushdown safety games, which we introduce in Section 2.

First, we show in Section 3 that no vertex of a finitely branching safety game (which covers pushdown games in particular) can have resilience ω . As a corollary, we show that Player 0 has positional optimally resilient strategies in finitely branching safety games. In contrast, we show that Player 0 does not necessarily have an optimally resilient strategy in infinitely branching safety games, for the reasons explained earlier.

In Sections 4 to 6, we consider the problem of determining the resilience of the initial vertex of a given pushdown safety game. First, we show in Section 4 how to characterize resilience values using classical games (without disturbances): While the notion of resilience is not defined via strategies of the antagonist, we show that one can nevertheless give control over disturbances to the antagonist, if one additionally adjusts the winning condition to control the number of occurrences of disturbances. For certain resilience values, but not all, this adjustment leads to a polynomial time reduction to solving classical games on pushdown games. The values that can be characterized in safety games are fixed finite values k and $\omega + 1$, but not ω .

We then prove that the resilience value of the initial vertex in pushdown safety games can be determined in triply-exponential time (Sections 5) and that of the initial vertex in one-counter safety games in polynomial space (Section 6). The latter result is tight, as associated decision problems are shown to be PSPACE-complete. To show membership, we use the following approach: We prove the existence of an upper bound on the resilience value of the initial vertex in case it is finite. With such an upper bound b , we can use the characterizations developed in Section 4 to perform an exhaustive search on the finite search space (the resilience is either in $\{0, 1, \dots, b\}$ or $\omega + 1$, as we have ruled out ω). For general pushdown games, this search can be implemented in triply-exponential time, as the bound b is doubly-exponential. However, relying on the simplicity of configuration graphs of one-counter systems and on the fact that the bound b is only exponential in this case, we are able to show that the search can be implemented in polynomial space for one-counter safety games. Proving the last result requires the combination of a wide range of techniques, including results from the theory of quantitative pushdown games [17], positional determinacy for quantitative pushdown games, and specifically tailored “hill-cutting” [5, 39] and “summarization” arguments [30, 19], which

we generalize from individual paths in pushdown systems to strategies. Also, we show that a strategy that is optimally-resilient from the initial vertex can be computed in exponential space (triplly-exponential time) for one-counter safety games (pushdown safety games).

Section 7 concludes and discusses directions for future work. Finally, in Appendix A, we present an application of our results, namely, a connection between optimally resilient strategies in pushdown safety games and optimal strategies (in the number of steps to the target) in pushdown reachability games [8, 10]. There, we also discuss which of our results obtained here carry over to pushdown reachability games and discuss the obstacles preventing us from generalizing the other results from safety to reachability.

All proofs omitted due to space restrictions can be found in the Appendix B.

Related Work Resilience, and closely related notions like fault-tolerance and robustness, are not a novel concept in the context of reactive systems synthesis, with numerous formalizations having been proposed. So as to not clutter this paper too much, we refer the reader to Dallal, Neider, and Tabuada [13] as well as Neider, Weinert, and Zimmermann [29] for a comprehensive discussion of how these notions are related to the concept of unmodeled intermittent disturbances. Other notions of resilience against environmental impacts not discussed there include an approach based on imperfect information games that quantifies the resilience of controllers to noise in the input signal [2, 40] (see also the references).

Finally, let us mention that one can implement the characterization of finite resilience values presented in Section 4 by energy conditions [6, 11]. However, solving energy games on pushdown graphs is undecidable [1] and so we do not pursue this approach here. Similarly unfeasible are stochastic methods to quantify resilience in pushdown games. Indeed, checking even the most basic, almost-sure reachability conditions for stochastic games on pushdown graphs is undecidable already for single state systems or single-player games [16].

2 Preliminaries

We use the ordinals $0 < 1 < 2 < \dots < \omega < \omega + 1 < \omega + 2$ to define resilience values. For convenience of notation, we also denote the cardinality of ω by ω .

Infinite Games with Disturbances An *arena* (with unmodeled intermittent disturbances) $\mathcal{A} = (V, V_0, V_1, E, D)$ consists of a countable directed graph (V, E) , a partition $\{V_0, V_1\}$ of V into the set of vertices V_0 of Player 0 and the set of vertices V_1 of Player 1, and a set $D \subseteq V_0 \times V$ of disturbance edges. Note that only vertices of Player 0 may have outgoing disturbance edges. We require that every vertex $v \in V$ has a successor v' with $(v, v') \in E$ to avoid finite plays. A vertex $v \in V$ is a *sink* if it has a single outgoing edge $(v, v) \in E$ leading back to itself but no outgoing disturbance edges.

A *play* in \mathcal{A} is an infinite sequence $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \dots \in (V \times \{0, 1\})^\omega$ such that $b_0 = 0$ and for all $j > 0$: $b_j = 0$ implies $(v_{j-1}, v_j) \in E$, and $b_j = 1$ implies $(v_{j-1}, v_j) \in D$. Hence, the additional bits b_j for $j > 0$ denote whether a standard edge or a disturbance edge has been taken to move from v_{j-1} to v_j . We say ρ starts in v_0 . A play prefix $(v_0, b_0) \dots (v_j, b_j)$ is defined similarly and ends in v_j . The number of disturbances in a play $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \dots$ is defined as $\#_D(\rho) = |\{j \in \omega \mid b_j = 1\}|$, which is either some $k \in \omega$ (if there are finitely many disturbances, namely k) or it is equal to ω (if there are infinitely many). A play ρ is disturbance-free, if $\#_D(\rho) = 0$.

A *game* (with unmodeled intermittent disturbances) $\mathcal{G} = (\mathcal{A}, \text{Win})$ consists of an arena with set V of vertices and a winning condition $\text{Win} \subseteq V^\omega$. A play $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \dots$ is winning for Player 0 if $v_0 v_1 v_2 \dots \in \text{Win}$, otherwise it is winning for Player 1. Hence, winning is oblivious to occurrences of disturbances.

In this work, we focus on safety conditions, but also use the Büchi condition in proofs. Both are induced by a subset F of the set V of vertices.

- $\text{Safety}(F)$ containing the sequences $v_0v_1v_2\cdots \in V^\omega$ with $v_j \notin F$ for every $j \in \omega$ denotes the safety condition induced by F , which requires to avoid F .
- $\text{Büchi}(F)$ containing the sequences $v_0v_1v_2\cdots \in V^\omega$ with $v_j \in F$ for infinitely many $j \in \omega$ denotes the Büchi condition induced by F , which requires to visit F infinitely often.

A game $(\mathcal{A}, \text{Win})$ is a *safety game* if $\text{Win} = \text{Safety}(F)$ for some subset F of the vertices of \mathcal{A} .

A *strategy* for Player $i \in \{0, 1\}$ is a function $\sigma: V^*V_i \rightarrow V$ such that $(v_j, \sigma(v_0 \cdots v_j)) \in E$ for every $v_0 \cdots v_j \in V^*V_i$. A play $(v_0, b_0)(v_1, b_1)(v_2, b_2)\cdots$ is *consistent with* σ if $v_{j+1} = \sigma(v_0 \cdots v_j)$ for every j with $v_j \in V_i$ and $b_{j+1} = 0$, i.e., if the next vertex is the one prescribed by the strategy unless a disturbance edge is used. A strategy σ is *positional*, if $\sigma(v_0 \cdots v_j) = \sigma(v_j)$ for all $v_0 \cdots v_j \in V^*V_i$.

Pushdown Games A *pushdown system* (PDS) $\mathcal{P} = (Q, \Gamma, \mathcal{E}, q_I)$ consists of a finite set Q of states with an initial state $q_I \in Q$, a stack alphabet Γ with a designated stack bottom symbol $\perp \notin \Gamma$, and a transition relation $\mathcal{E} \subseteq Q \times \Gamma_\perp \times Q \times \Gamma_\perp^{\leq 2}$, where $\Gamma_\perp = \Gamma \cup \{\perp\}$ and $\Gamma_\perp^{\leq 2} = \{w \in \Gamma_\perp^* \mid |w| \leq 2\}$. We require \mathcal{E} to neither write nor delete \perp from the stack. Also, we assume every PDS to be *deadlock-free*, i.e., for every $q \in Q$ and $A \in \Gamma_\perp$ there exist $q' \in Q$ and $w \in \Gamma_\perp^{\leq 2}$ such that $(q, A, q', w) \in \mathcal{E}$. Finally, \mathcal{P} is a *one-counter system* (OCS) if $|\Gamma| = 1$.

A stack content is a word in $\Gamma^*\perp$ where the leftmost symbol is assumed to be the top of the stack. A *configuration* of \mathcal{P} is a pair (q, γ) consisting of a state $q \in Q$ and a stack content $\gamma \in \Gamma^*\perp$. The stack height of a configuration (q, γ) is defined by $\text{sh}(q, \gamma) = |\gamma| - 1$. Given two configurations (q, γ) and (q', γ') we write $(q, \gamma) \vdash_{\mathcal{E}} (q', \gamma')$ if there exists a transition $(q, \gamma_0, q', w) \in \mathcal{E}$ with $\gamma' = w\gamma_1 \cdots \gamma_{|\gamma|-1}$.

Fix a PDS $\mathcal{P} = (Q, \Gamma, \mathcal{E}, q_I)$, a partition $\{Q_0, Q_1\}$ of Q and an additional transition relation $\Delta \subseteq Q_0 \times \Gamma_\perp \times Q \times \Gamma_\perp^{\leq 2}$, which is also required to neither write nor delete \perp from the stack. These induce the (pushdown) arena (V, V_0, V_1, E, D) with


- $V = \{(q, \gamma) \mid q \in Q, \gamma \in \Gamma^*\perp\}$ is the set of configurations of \mathcal{P} ,
- $V_i = \{(q, \gamma) \in V \mid q \in Q_i\}$ for $i \in \{0, 1\}$ is the set of configurations whose state is in Q_i ,
- $E = \{(v, v') \mid v \vdash_{\mathcal{E}} v'\}$ is the set of edges, induced by the transition relation \mathcal{E} , and
- $D = \{(v, v') \mid v \vdash_{\Delta} v'\}$ is the set of disturbance edges, which is induced by the transition relation Δ , where \vdash_{Δ} is defined analogously to $\vdash_{\mathcal{E}}$.

Typically, we are interested in the initial vertex of the arena, which is defined as (q_I, \perp) .

A *pushdown safety game* is a safety game whose arena is induced by a pushdown system \mathcal{P} and whose winning condition is induced by a subset of \mathcal{P} 's states, i.e., $F \subseteq Q$ induces the set $\{(q, \gamma) \in V \mid q \in F\}$ of vertices. One-counter safety games are defined analogously.

When using a pushdown game as an input for an algorithm, we represent it by the underlying PDS, the partition of its states, the additional transition relation for the disturbance edges, and a subset of the states inducing the winning condition. We define the size of the input as $|Q| + |\Gamma|$, as all these objects can be represented in polynomial size in the number of states and stack symbols of the underlying PDS.

Infinite Games without Disturbances For technical convenience, we characterize the classical notion of infinite games, i.e., those without disturbances, (see, e.g., [18]) as a special case of games with disturbances. Let \mathcal{G} be a game with vertex set V . A strategy σ for Player i in \mathcal{G} is said to be a *winning strategy* for her from $v \in V$, if every disturbance-free play that starts in v and that is consistent with σ is winning for Player i . The winning region $\mathcal{W}_i(\mathcal{G})$ of Player i in \mathcal{G} contains those vertices from which Player i has a winning strategy. Thus, the winning regions of \mathcal{G} are independent of the disturbance edges, i.e., we obtain the classical notion of infinite games. Player i wins \mathcal{G} from v , if $v \in \mathcal{W}_i(\mathcal{G})$.



■ **Figure 1** A one-counter arena, restricted to vertices reachable from the initial vertex (q_I, \perp) . All vertices are in V_0 , disturbance edges are drawn as dashed arrows, and doubly-lined vertices are in F .

Resilient Strategies Let \mathcal{G} be a game with vertex set V and let $\alpha \in \omega + 2$. A strategy σ for Player 0 in \mathcal{G} is α -resilient from $v \in V$ if every play ρ that starts in v , that is consistent with σ , and with $\#_D(\rho) < \alpha$, is winning for Player 0. Thus, a k -resilient strategy with $k \in \omega$ is winning even under at most $k - 1$ disturbances, an ω -resilient strategy is winning even under any finite number of disturbances, and an $(\omega + 1)$ -resilient strategy is winning even under infinitely many disturbances.

We define the *resilience* of a vertex v of \mathcal{G} as

$$r_{\mathcal{G}}(v) = \sup\{\alpha \in \omega + 2 \mid \text{Player 0 has an } \alpha\text{-resilient strategy for } \mathcal{G} \text{ from } v\}.$$

Note that the definition is not antagonistic, i.e., it is not defined via strategies of Player 1. A strategy σ is *optimally resilient* if it is $r_{\mathcal{G}}(v)$ -resilient from every vertex v .

► **Example 1.** Consider the game $\mathcal{G} = (\mathcal{A}, \text{Safety}(F))$ where \mathcal{A} is the arena from Figure 1 and $\text{Safety}(F)$ is the safety condition induced by $F = \{q_2\}$.

We have that $r_{\mathcal{G}}(q_I, A^n \perp) = \omega + 1$, $r_{\mathcal{G}}(q_1, A^n \perp) = n$ for all $n \in \omega$, and $r_{\mathcal{G}}(q_2, \perp) = 0$. Furthermore, the strategy that indefinitely stays in state q_I is optimally resilient.

3 Resilience in Infinite Safety Games

Player 0 has optimally resilient strategies in every safety game played in a finite arena [13]. In this section, we show that this result also holds for pushdown safety games, but fails for safety games in arbitrary infinite arenas. We start by observing that in safety games in infinite arenas, vertices with resilience ω may exist, unlike in safety games in finite arenas [13].

► **Example 2.** Consider the one-counter arena presented in Figure 1 with the safety condition induced by $F = \{q_2\}$, i.e., Player 0 wins if she avoids visiting a vertex with state q_2 . As argued in Example 1, the resulting game \mathcal{G} has vertices of resilience $\omega + 1$ and k , for each $k \in \omega$, i.e., all values but ω are assumed.

Let us add a vertex $v \in V_0$ to \mathcal{G} with outgoing edges to all vertices of the form $(q_1, A^n \perp)$ to obtain the game \mathcal{G}' (which is infinitely branching and therefore no longer a pushdown arena). Let σ_k , for $k > 0$, be a strategy that moves from v to $(q_1, A^k \perp)$. We have that $r_{\mathcal{G}'}(v) \geq \omega$, as σ_k is k -resilient from v . Consider an arbitrary strategy σ : From v , it moves to some $(q_1, A^k \perp)$ from which k disturbances force the play into the losing sink. Hence, σ is not $(k + 1)$ -resilient and therefore not ω -resilient. Thus, there is no optimally resilient strategy in \mathcal{G}' .

The underlying issue is that $r_{\mathcal{G}}(v) \geq \omega$ can be witnessed either (a) by the existence of a strategy that is ω -resilient from v , or (b) by the existence of a family $(\sigma_k)_{k \in \omega}$ of strategies where each σ_k is k -resilient from v , but not ω -resilient from v . The second case only exists

as ω is a limit ordinal (the only one we consider). For all $\alpha \neq \omega$, we have that $r_{\mathcal{G}}(v) = \alpha$ if and only if Player 0 has an α -resilient strategy from v . The games studied in previous work [13, 29] only exhibited the former case, as these only considered finite arenas. As witnessed in Example 2, this is no longer true in games in infinite arenas.

Note that there is a change of quantifiers between these two cases: by definition, an ω -resilient strategy is k -resilient for every $k \in \omega$, i.e., in the former case there is a uniform strategy that is k -resilient for every $k \in \omega$. In the latter case, for every $k \in \omega$, there is a strategy that is k -resilient, but not ω -resilient. Hence, in the following, we distinguish between these two cases. We say that a vertex v of a game \mathcal{G} with $r_{\mathcal{G}}(v) = \omega$ has a uniform witness², if there is an ω -resilient strategy from v . A game with a vertex of resilience ω without a uniform witness has no optimally resilient strategy by definition.

For safety games in infinite arenas, the existence of optimally resilient strategies depends on the branching of the arena. We say that an arena (V, V_0, V_1, E, D) is *finitely branching* if the set $\{v' \mid (v, v') \in E\}$ of successors of v is finite for every $v \in V$. Otherwise, if there is a vertex with infinitely many successors, then the arena is *infinitely branching*. Note that pushdown arenas are finitely branching.

The following theorem shows that the games presented in Example 2 already exhibit all possible resilience values in safety games, and that infinite branching is necessary to obtain a vertex of resilience ω . We formulate the result for arbitrary infinite arenas, as the proof technique we use here does not rely on the arena being a pushdown arena.

► **Lemma 3.** *Let \mathcal{G} be a safety game with vertex set V .*

1. *There is no $v \in V$ with $r_{\mathcal{G}}(v) = \omega$ that has a uniform witness.*
2. *If \mathcal{A} is finitely branching, then there is no $v \in V$ with $r_{\mathcal{G}}(v) = \omega$.*

Finally, the main result of this section shows that optimally resilient strategies exist in all finitely branching safety games, i.e., in particular in pushdown safety games.

► **Theorem 4.** *Player 0 has positional optimally resilient strategies in finitely branching safety games.*

4 Characterizing Resilience Values via Classical Games

In this section, we characterize the existence of α -resilient strategies by games without disturbances. This generalizes a characterization for $\alpha = \omega + 1$ in finite arenas [29] to infinite arenas and all $\alpha \in \omega + 2$.

The main idea is to give Player 1 control over the disturbances and to restrict the number of their occurrences using the winning condition. Intuitively, when it is Player 0's turn at a vertex v , we let Player 1 first decide whether to simulate a disturbance edge from D or whether to allow Player 0 to pick a standard edge from E . To this end, we add v to Player 1's vertices and he can either move to some vertex v' such that the disturbance edge (v, v') exists. By doing his, he has to visit the fresh vertex (v, v') , which allows to keep track of the number of simulated disturbances. This vertex has exactly one outgoing edge leading to v' . On the other hand, if he does not simulate a disturbance edge, he moves from v to a fresh copy \bar{v} of v from which Player 0 has edges leading to the successors of v . Finally, the moves

² Note that uniformity here refers to having a single strategy σ that is k -resilient from v for every k . It is *not* related to the concept of uniform winning strategies, i.e., strategies that are winning from every vertex in a winning region.

at Player 1's original vertices are unchanged, but we subdivide the edge so that a play in the extended arena always alternates between vertices from V and auxiliary vertices.

Formally, given an arena $\mathcal{A} = (V, V_0, V_1, E, D)$, we define the rigged arena $\mathcal{A}_{\text{rig}} = (V', V'_0, V'_1, E', D')$ with $V' = V \cup A$ for the set

$$A = \{\bar{v} \mid v \in V_0\} \cup D \cup \{(v, v') \in E \mid v \in V_1\}$$

of auxiliary vertices, $V'_0 = \{\bar{v} \mid v \in V_0\}$, $V'_1 = V' \setminus V'_0$, $D' = \emptyset$, and E' is the union of the following sets of edges:

- $\{(v, (v, v')), ((v, v'), v') \mid (v, v') \in D\}$: Player 1 simulates a disturbance edge $(v, v') \in D$ by moving from v to v' via the auxiliary vertex (v, v') that signifies that a disturbance is simulated.
- $\{(v, \bar{v}) \mid v \in V_0\}$: Player 1 does not simulate a disturbance edge and instead gives control to Player 0 by moving to the auxiliary vertex \bar{v} .
- $\{(\bar{v}, v') \mid v \in V_0 \text{ and } (v, v') \in E\}$: Player 0 has control at the auxiliary vertex \bar{v} and simulates a standard move from $v \in V_0$ to v' .
- $\{(v, (v, v')), ((v, v'), v') \mid (v, v') \in E \text{ and } v \in V_1\}$: Player 1 simulates a standard move from $v \in V_1$ to v' by moving via the auxiliary vertex (v, v') .

Let $R_{\geq k}$ denote the set of sequences $v_0 v_1 v_2 \cdots \in (V')^\omega$ such that $|\{j \mid v_j \in D\}| \geq k$, i.e., those plays in which Player 1 simulates at least k disturbances. Finally, given a winning condition $\text{Win} \subseteq V^\omega$ for \mathcal{A} , we define the rigged winning condition

$$\text{Win}_{\text{rig}} = \{v_0 v_1 v_2 \cdots \in (V')^\omega \mid v_0 \in V \text{ and } v_0 v_2 v_4 \cdots \in \text{Win}\},$$

which contains all plays in \mathcal{A}_{rig} that start in V and are in Win after removing the auxiliary vertices. Note that $\text{Büchi}(D)$ contains those plays that simulate infinitely many disturbances.

► **Lemma 5.** *Let $\mathcal{G} = (\mathcal{A}, \text{Win})$ be a game, let v be a vertex of \mathcal{G} , and let $k \in \omega$.*

1. *Player 0 has an $(\omega + 1)$ -resilient strategy for \mathcal{G} from v if and only if $v \in \mathcal{W}_0(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}})$.*
2. *Player 0 has an ω -resilient strategy for \mathcal{G} from v if and only if $v \in \mathcal{W}_0(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}} \cup \text{Büchi}(D))$.*
3. *Player 0 has a k -resilient strategy for \mathcal{G} from v if and only if $v \in \mathcal{W}_0(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}} \cup R_{\geq k})$.*

5 Resilience in Pushdown Safety Games

The goal of this section is to develop an algorithm that determines the resilience of the initial vertex of a pushdown safety game. To this end, we rely on the characterizations presented in the previous section as well as an upper bound on the possible finite resilience values that can be realized by the initial vertex of such a game. We begin by showing that the first two characterizations presented in Lemma 5 (for $\omega + 1$ and ω) are effective for pushdown games. Intuitively, we prove that a pushdown machine \mathcal{P} inducing an arena \mathcal{A} can in polynomial time be turned into a pushdown machine \mathcal{P}_{rig} inducing the arena \mathcal{A}_{rig} .

We state the result for parity conditions (see, e.g., [18] for a definition of parity conditions), which subsume safety conditions.

► **Lemma 6.** *The following problem is EXPTIME-complete (and PSPACE-complete if inputs are restricted to one-counter games): “Given a pushdown parity game \mathcal{G} with initial vertex v_I and $\alpha \in \{\omega, \omega + 1\}$, does Player 0 have an α -resilient strategy for \mathcal{G} from v_I ?”. If yes, such a strategy can be computed in exponential time.*

Both EXPTIME-hardness and PSPACE-hardness already hold for pushdown safety games and one-counter safety games, respectively. The third characterization of Lemma 5 (for $k \in \omega$) is effective as well (even for parity games). Here the running time depends on k .

► **Lemma 7.** *The following problem is in 2EXPTIME (in EXPSPACE if the input is one-counter): “Given a pushdown parity game \mathcal{G} with initial vertex v_I and $k \in \omega$ (encoded in binary), does Player 0 have a k -resilient strategy for \mathcal{G} from v_I ?”. If yes, such a strategy can be computed in doubly-exponential time.*

There are no vertices of resilience ω in pushdown safety games (Lemma 3.2). Thus, the effective characterizations we have presented so far suffice to determine the resilience of the initial vertex in such a game: First, check whether it is $\omega + 1$; if not, then it has to be finite. Hence, for increasing k , check whether the resilience is greater than k . As the resilience is finite, this algorithm will eventually terminate and report the resilience correctly. However, without an upper bound on the possible finite resilience values of the initial vertex, there is no bound on the running time, just a termination guarantee. In the remainder of this section, we present a tight doubly-exponential upper bound $b(\mathcal{P})$ on the resilience of the *initial vertex* in pushdown safety games in the case the resilience is finite. That is, if $r_{\mathcal{G}}(v_I) \in \omega$ then $r_{\mathcal{G}}(v_I) < b(\mathcal{P})$. Note that any proof of the upper bound has to depend on the vertex under consideration being initial, as we have shown that there is in general no upper bound on finite resilience values assumed in pushdown safety games (cf. Example 2). The bound $b(\mathcal{P})$ only depends on the pushdown system \mathcal{P} inducing the game and yields an effective algorithm to determine the resilience of the initial vertex v_I , presented as Algorithm 1.

Algorithm 1 Computing the resilience of the initial vertex v_I of a pushdown safety game $\mathcal{G} = (\mathcal{A}, \text{Safety}(F))$ induced by a PDS \mathcal{P}

```

1: if  $v_I \in \mathcal{W}_0(\mathcal{A}_{\text{rig}}, \text{Safety}(F)_{\text{rig}})$  then
2:   return  $\omega + 1$ 
3: for  $k = 1$  to  $b(\mathcal{P})$  do
4:   if  $v_I \in \mathcal{W}_1(\mathcal{A}_{\text{rig}}, \text{Safety}(F)_{\text{rig}} \cup R_{\geq k})$  then
5:     return  $k - 1$ 

```

Given a PDS \mathcal{P} with set Q of states and set Γ of stack symbols let \mathcal{P}_{rig} be the PDS obtained from \mathcal{P} by implementing the transformation from an arena to the rigged arena. The cardinality of the set Q' of states of \mathcal{P}_{rig} is bounded quadratically in $|Q|$ and the set of stack symbols used by \mathcal{P}_{rig} is still Γ . We define $b(\mathcal{P}) = |Q'| \cdot h(\mathcal{P}) \cdot |\Gamma|^{h(\mathcal{P})}$, where $h(\mathcal{P}) = |Q'| \cdot |\Gamma| \cdot 2^{|Q'|+1} + 1$. Note that $b(\mathcal{P}) \in 2^{2^{\mathcal{O}(|\mathcal{P}|^2)}}$ and $b(\mathcal{P}) \in 2^{\mathcal{O}(|\mathcal{P}|^2)}$ if \mathcal{P} is an OCS.

► **Lemma 8.** *Let \mathcal{G} be a pushdown safety game with initial vertex v_I . If $r_{\mathcal{G}}(v_I) \neq \omega + 1$, then $r_{\mathcal{G}}(v_I) < b(\mathcal{P})$, where \mathcal{P} is the PDS underlying \mathcal{G} .*

This upper bound immediately implies correctness of Algorithm 1, which determines the resilience of the initial vertex of a pushdown safety game.

► **Theorem 9.** *The following problem can be solved in triply-exponential time: “Given a pushdown safety game \mathcal{G} with initial vertex v_I , determine $r_{\mathcal{G}}(v_I)$ ”. If yes, an $r_{\mathcal{G}}(v_I)$ -resilient strategy can be computed in triply-exponential time.*

Note that there is a gap between the triply-exponential upper bound and the exponential lower bound obtained for the related decision problems for ω and $\omega + 1$ (Lemma 6).

The complexity for the special case of one-counter safety games is much smaller, i.e., the resilience of the initial vertex can be computed in exponential space, as the winner of one-counter safety games can be computed in polynomial space [34] and the upper bound on finite resilience values of the initial vertex is only exponential. Furthermore, a witnessing strategy can be computed in doubly-exponential time using Lemma 7. In the next section, we prove that one can do even better by exploiting the simple structure of one-counter arenas.

To conclude this section, we claim that the bound $b(\mathcal{P})$ on the resilience of an initial vertex in a pushdown safety game with finite resilience is tight: There is an exponential lower bound for the one-counter case and a doubly-exponential lower bound for the pushdown case. Both constructions are generalizations of constructions that appeared in the literature previously [10]. To simplify our notation, let p_j denote the j -th prime number and define the primorial $p_k\# = \prod_{j=1}^k p_j$ to be the product of the first k prime numbers. We have $p_k\# \geq 2^k$.

► **Lemma 10.** *Let $k \in \omega$.*

1. *There is a one-counter safety game \mathcal{G}_k with initial state v_I such that $r_{\mathcal{G}}(v_I) = p_k\#$ and the underlying OCS has polynomially many states in k .*
2. *There is a pushdown safety game \mathcal{G}'_k with initial state v_I such that $r_{\mathcal{G}}(v_I) = 2^{p_k\#} - 1$ and the underlying PDS has polynomially many states in k and two stack symbols.*

6 Resilience in One-counter Safety Games

In this section, we show that one can compute the resilience of the initial vertex in a one-counter safety game in polynomial space, significantly improving the exponential space requirement derived in the previous section.

► **Theorem 11.** *The following problem can be solved in polynomial space: “Given a one-counter safety game \mathcal{G} with initial vertex v_I , determine $r_{\mathcal{G}}(v_I)$ ”.*

To prove this result, we show that one can implement Algorithm 1 in polynomial space if the underlying pushdown system is one-counter. In this case, one can run the check “ $v_I \in \mathcal{W}_0(\mathcal{A}_{\text{rig}}, \text{Safety}(F)_{\text{rig}})$ ” in Line 1 in polynomial space due to Lemma 6, and can implement the counter in Line 3 in polynomial space, as the upper bound $b(\mathcal{P})$ is exponential (see the definition on Page 9). It remains to show that one can check in polynomial space, for a given $k \leq b(\mathcal{P})$, if $v_I \in \mathcal{W}_1(\mathcal{A}_{\text{rig}}, \text{Safety}(F)_{\text{rig}} \cup R_{\geq k})$ holds. In the rest of this section we show that this is indeed possible.

Fix, the rigged game $\mathcal{G}_k = (\mathcal{A}_{\text{rig}}, \text{Safety}(F)_{\text{rig}} \cup R_{\geq k})$ for some $k \leq b(\mathcal{P})$ with $\mathcal{A}_{\text{rig}} = (V', V'_0, V'_1, E', \emptyset)$, with initial vertex v_I , where \mathcal{P} is the OCS underlying the original game \mathcal{G} that induces \mathcal{G}_k . We show that the existence of winning strategies for Player 1 in \mathcal{G}_k can be witnessed by a finite graph structure, as follows.

A *strategy graph* for \mathcal{G}_k is a tuple $(V^\circ, E^\circ, \mu_r^\circ, \mu_d^\circ)$ with $\mu_r^\circ: V^\circ \rightarrow \{0, \dots, k-1\}$ and $\mu_d^\circ: V^\circ \rightarrow \{0, \dots, |V^\circ|\}$ such that the following properties are satisfied:

1. (V°, E°) is a directed graph with $V^\circ \subseteq V'$, $E^\circ \subseteq E'$, $v_I \in V^\circ$, and $\text{sh}(v) \leq (2k)^{|Q|^2}$ for all $v \in V^\circ$. Note that $(2k)^{|Q|^2}$ is exponential in the size of the pushdown system \mathcal{P} underlying \mathcal{G} , even though $k \leq b(\mathcal{P})$ may itself be exponential.
2. For all $v \in (V^\circ \cap V'_0) \setminus F$ and all $(v, v') \in E'$, we have $(v, v') \in E^\circ$.
3. For all $v \in (V^\circ \cap V'_1) \setminus F$ there is a unique outgoing edge $(v, v') \in E'$ with $(v, v') \in E^\circ$.
4. For all $(v, v') \in E^\circ$, we have $\mu_r^\circ(v) \geq \mu_r^\circ(v')$ with strict inequality if $v \in D$.
5. For all $(v, v') \in E^\circ$, we have $\mu_d^\circ(v) > \mu_d^\circ(v')$.

► **Lemma 12.** *Player 1 wins \mathcal{G}_k from v_I if and only if there exists a strategy graph for \mathcal{G}_k .*

To simplify the proof, we transform \mathcal{G}_k into a game \mathcal{G}'_k where all reachable vertices in F are sinks of stack height zero. To do this, we replace all outgoing (standard and disturbance) edges of vertices $(q, A^n \perp) \in F$ with $n > 0$ by an edge to $(q, A^{n-1} \perp)$ (which is also in F) and the all outgoing (standard and disturbance) edges of vertices $(q, \perp) \in F$ by an edge to a sink vertex (q_f, \perp) , where q_f is a fresh state. Then, \mathcal{G}'_k is the game played in the modified arena with winning condition $\text{Safety}(\{q_f\})_{\text{rig}} \cup R_{\geq k}$. Intuitively, once a vertex in F is reached in the modified arena, the players no longer have strategic choices; instead, the stack is emptied (without simulating any disturbances) and the unsafe sink vertex (q_f, \perp) is reached.

It is straightforward to verify that we have $v \in \mathcal{W}_i(\mathcal{G}_k)$ if and only if $v \in \mathcal{W}_i(\mathcal{G}'_k)$ for every vertex of \mathcal{A}_{rig} and $i \in \{0, 1\}$ by transferring winning strategies between the games. So, in the following, we assume without loss of generality, that the only vertices of \mathcal{G}_k in F that are reachable from the initial vertex are sinks of stack height zero. In this situation, a play can no longer simulate a disturbance edge once a vertex in F has been reached.

To prove Lemma 12, we show that if Player 1 wins \mathcal{G}_k with some arbitrary winning strategy, then also with a winning strategy that can be turned into a strategy graph. To simplify our notation, given a strategy τ , let $\text{maxSh}(\tau) = \sup_v \text{sh}(v)$, where v ranges over all vertices reachable by a play prefix starting in v_I that is consistent with τ , i.e., $\text{maxSh}(\tau)$ is the maximal stack height visited by a play that is starting in the initial vertex and consistent with τ . Using this, we show that Player 1 wins \mathcal{G}_k from v_I if and only if he has a positional winning strategy from v_I with $\text{maxSh}(\tau) \leq (2k)^{|\mathcal{Q}|^2}$. The latter can then be transformed into a strategy graph.

We only have to consider the implication from left to right, as the other one is trivial. Let Player 1 win \mathcal{G}_k from v_I , i.e., he has a winning strategy τ for \mathcal{G}_k from v_I . We proceed in two steps: First, We turn τ in a positional winning strategy τ' from v_I (Lemma 13). Then, we turn τ' into a positional winning strategy τ'' with $\text{maxSh}(\tau'') \leq (2k)^{|\mathcal{Q}|^2}$ (Lemma 14).

For the first step, we generalize a standard argument for turning an arbitrary, not necessarily positional, winning strategy τ in a reachability game into a positional one: At a vertex $v \notin F$, consider all play prefixes that are consistent with τ and end in v , and mimic the move τ prescribes for a longest one (call it $\text{rep}(v)$). The resulting strategy τ' is obviously positional and winning as every play consistent with τ' and ending in some $v \notin F$ can be shown to be at most as long as the play $\text{rep}(v)$ whose moves are mimicked to define $\tau'(v)$. Here, we have to refine this argument to ensure that the resulting strategy τ' still simulates at most $k - 1$ disturbances during each play.

► **Lemma 13.** *If Player 1 wins \mathcal{G}_k from v_I then he has a positional winning strategy for \mathcal{G}_k from v_I .*

The second step of our construction is to bound the stack height reached by plays consistent with the winning strategy (while preserving positionality). To this end, we generalize a classical argument for pushdown safety games: In such games, Player 1, who has a reachability objective, has a positional winning strategy τ from v_I with exponentially bounded $\text{maxSh}(\tau)$, if he wins at all from v_I . This is typically proven by a “hill-cutting” argument [5, 39] showing that a winning strategy exceeding this bound can be turned into one of smaller maximal stack height by removing infixes of plays that increase the stack without reaching states that have not been reached at smaller stack height already. Here, we again have to generalize this argument to additionally ensure that the number of simulated disturbances remains bounded by $k - 1$. This is done using “summarizations” of paths in pushdown systems (see e.g. [30, 19]) that take the number of disturbances into account.

► **Lemma 14.** *If Player 1 wins \mathcal{G}_k from v_I then he has a positional winning strategy from v_I with $\max\text{Sh}(\tau) \leq (2k)^{|\mathcal{Q}|^2}$.*

A positional strategy as in Lemma 14 is essentially a strategy graph. So, we have proven Lemma 12: The existence of strategy graphs for \mathcal{G}_k captures Player 1 winning \mathcal{G}_k . Hence, it remains to prove that we can decide the existence of strategy graphs in polynomial space. Here, we use the fact that k is at most $b(\mathcal{P}) \in \mathcal{O}(2^{|\mathcal{P}|^2})$, where \mathcal{P} is the pushdown system underlying the game inducing \mathcal{G}_k , to guess and verify a strategy graph in polynomial space.

► **Lemma 15.** *The following problem is in PSPACE: “Given a one-counter safety game \mathcal{G} induced by a PDS \mathcal{P} and $k \leq b(\mathcal{P})$ (encoded in binary), is there a strategy graph for \mathcal{G}_k ?”*

While we consider one-counter systems with unit updates, i.e., each transition changes the counter value by at most one, our results are also applicable to one-counter systems where each transition updates the counter by some integer (encoded in binary). Such *binary updates* can be simulated by unit updates, albeit with an exponential blowup. Hence, the algorithm above computes the resilience of the initial vertex of a one-counter safety game with binary updates in exponential space. A matching lower bound follows from the EXPSpace hardness of solving disturbance-free one-counter safety games with binary updates [22].

7 Conclusion

In this work, we have investigated pushdown safety games with disturbances, thereby extending the theory of games with disturbances from finite to infinite arenas. In particular, we have determined the possible resilience values in safety games, presented effective characterizations for all possible values, and presented algorithms that determine the resilience of the initial vertex (and a witnessing strategy) in one-counter and pushdown safety games. As an application of our results, we obtain a polynomial space algorithm for computing optimal winning strategies for one-counter reachability games (see Appendix A.2). This is, to the best of our knowledge, the first improvement over the general doubly-exponential time algorithm for pushdown reachability games due to Carayol and Hague [10].

The algorithm computing the resilience in one-counter safety games runs in polynomial space, which is optimal, as the corresponding decision problems are PSPACE-complete. However, the algorithm for pushdown games has triply-exponential running time. Here, there is a gap, as some of the corresponding decision problems are EXPTIME-complete (e.g., those for resilience $\omega + 1$ and ω) while the complexity of others is open (e.g., that for finite resilience values). In future work, we aim to close this gap. An interesting first step in that direction would be to determine the complexity of checking whether the resilience of the initial vertex is at least k , where k is part of the input and encoded in binary. Here, one has to keep in mind that algorithms for computing the resilience also yield algorithms computing optimal strategies in reachability games. The latter problem also has a complexity gap between the currently best algorithms and known lower bounds. Finally, another obvious open problem is to consider more general winning conditions, e.g., reachability (see Appendix A.1) or parity.

The main obstacle is that one either has to develop an effective characterization of vertices with resilience ω without a uniform witness, or to obtain an upper bound on the finite resilience value an initial vertex can assume. The first option is challenging due to the quantifier change discussed in Section 5. Hence, the more promising route seems to be the second option. The main challenge here is to bound the number of disturbances that are necessary to prevent Player 0 from ever reaching the target states, i.e., Player 1 now has a safety objective in conjunction with a limited number of disturbances at his disposal.

References

- 1 Parosh Aziz Abdulla, Mohamed Faouzi Atig, Piotr Hofman, Richard Mayr, K. Narayan Kumar, and Patrick Totzke. Infinite-state energy games. In *CSL-LICS 2014*, pages 7:1–7:10. ACM, 2014. URL: <https://doi.org/10.1145/2603088.2603100>.
- 2 Shaull Almagor and Orna Kupferman. Latticed-LTL synthesis in the presence of noisy inputs. *Discrete Event Dynamic Systems*, 27(3):547–572, 2017. doi:10.1007/s10626-017-0242-0.
- 3 Roderick Bloem, Krishnendu Chatterjee, Karin Greimel, Thomas A. Henzinger, Georg Hofferek, Barbara Jobstmann, Bettina Könighofer, and Robert Könighofer. Synthesizing robust systems. *Acta Inf.*, 51(3-4):193–220, 2014. doi:10.1007/s00236-013-0191-5.
- 4 Roderick Bloem, Krishnendu Chatterjee, Thomas A. Henzinger, and Barbara Jobstmann. Better quality in synthesis through quantitative objectives. In *CAV 2009*, volume 5643 of *LNCS*, pages 140–156. Springer, 2009. doi:10.1007/978-3-642-02658-4_14.
- 5 Stanislav Böhm, Stefan Göller, and Petr Jancar. Bisimulation equivalence and regularity for real-time one-counter automata. *J. Comput. Syst. Sci.*, 80(4):720–743, 2014. doi:10.1016/j.jcss.2013.11.003.
- 6 Patricia Bouyer, Ulrich Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Jiri Srba. Infinite runs in weighted timed automata with energy constraints. In *FORMATS 2008*, volume 5215 of *LNCS*, pages 33–47. Springer, 2008.
- 7 J. Richard Büchi and Lawrence H. Landweber. Solving sequential conditions by finite-state strategies. *Trans. Amer. Math. Soc.*, 138:295–311, 1969.
- 8 Thierry Cachat. Symbolic strategy synthesis for games on pushdown graphs. In *ICALP 2002*, volume 2380 of *LNCS*, pages 704–715. Springer, 2002. doi:10.1007/3-540-45465-9_60.
- 9 Thierry Cachat. Higher order pushdown automata, the caucal hierarchy of graphs and parity games. In *ICALP 2003*, volume 2719 of *LNCS*, pages 556–569. Springer, 2003. URL: https://doi.org/10.1007/3-540-45061-0_45.
- 10 Arnaud Carayol and Matthew Hague. Optimal strategies in pushdown reachability games. In *MFCS 2018*, volume 117 of *LIPICs*, pages 42:1–42:14. Schloss Dagstuhl - LZI, 2018. doi:10.4230/LIPICs.MFCS.2018.42.
- 11 Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga. Resource interfaces. In *EMSOFT 2003*, volume 2855 of *LNCS*, pages 117–133. Springer, 2003.
- 12 Anne Condon. On algorithms for simple stochastic games. In *Advances in Computational Complexity Theory*, pages 51–73. American Mathematical Society, 1993.
- 13 Eric Dallal, Daniel Neider, and Paulo Tabuada. Synthesis of safety controllers robust to unmodeled intermittent disturbances. In *CDC 2016*, pages 7425–7430. IEEE, 2016. doi:10.1109/CDC.2016.7799416.
- 14 L. Doyen and J.-F. Raskin. Games with imperfect information: Theory and algorithms. In *Lectures in Game Theory for Computer Scientists*, pages 185–212. Cambridge University Press, 2011.
- 15 Matthew B. Dwyer, George S. Avrunin, and James C. Corbett. Patterns in property specifications for finite-state verification. In *ICSE 1999*, pages 411–420. ACM, 1999. doi:10.1145/302405.302672.
- 16 Kousha Etessami and Mihalis Yannakakis. Recursive markov decision processes and recursive stochastic games. *J. ACM*, 62(2):11:1–11:69, 2015. doi:10.1145/2699431.
- 17 Wladimir Fridman and Martin Zimmermann. Playing pushdown parity games in a hurry. In *GandALF 2012*, volume 96 of *EPTCS*, pages 183–196, 2012. doi:10.4204/EPTCS.96.14.
- 18 Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *LNCS*. Springer, 2002. doi:10.1007/3-540-36387-4.
- 19 Lukás Holík, Roland Meyer, and Sebastian Muskalla. Summaries for context-free games. In *FSTTCS*, volume 65 of *LIPICs*, pages 41:1–41:16. Schloss Dagstuhl - LZI, 2016. URL: <https://doi.org/10.4230/LIPICs.FSTTCS.2016.41>.

- 20 Frederick A. Hosch and Lawrence H. Landweber. Finite delay solutions for sequential conditions. In *ICALP 1972*, pages 45–60. North-Holland, Amsterdam, 1972.
- 21 Chung-Hao Huang, Doron A. Peled, Sven Schewe, and Farn Wang. A game-theoretic foundation for the maximum software resilience against dense errors. *IEEE Trans. Software Eng.*, 42(7):605–622, 2016. doi:10.1109/TSE.2015.2510001.
- 22 Paul Hunter. Reachability in succinct one-counter games. In Mikołaj Bojańczyk, Sławomir Lasota, and Igor Potapov, editors, *RP 2015*, volume 9328 of *LNCS*, pages 37–49. Springer, 2015. doi:10.1007/978-3-319-24537-9_5.
- 23 Petr Jancar and Zdenek Sawa. A note on emptiness for alternating finite automata with a one-letter alphabet. *Inf. Process. Lett.*, 104(5):164–167, 2007. doi:10.1016/j.ipl.2007.06.006.
- 24 Orna Kupferman and Moshe Y. Vardi. An automata-theoretic approach to reasoning about infinite-state systems. In *CAV 2000*, volume 1855 of *LNCS*, pages 36–52. Springer, 2000. doi:10.1007/10722167_7.
- 25 Rupak Majumdar, Elaine Render, and Paulo Tabuada. A theory of robust omega-regular software synthesis. *ACM Trans. Embedded Comput. Syst.*, 13(3):48:1–48:27, 2013. doi:10.1145/2539036.2539044.
- 26 David E. Muller and Paul E. Schupp. The theory of ends, pushdown automata, and second-order logic. *Theor. Comput. Sci.*, 37:51–75, 1985. doi:10.1016/0304-3975(85)90087-8.
- 27 Daniel Neider. Reachability games on automatic graphs. In *CIAA 2010*, volume 6482 of *LNCS*, pages 222–230. Springer, 2010. doi:10.1007/978-3-642-18098-9_24.
- 28 Daniel Neider and Ufuk Topcu. An automaton learning approach to solving safety games over infinite graphs. In *TACAS 2016*, volume 9636 of *LNCS*, pages 204–221. Springer, 2016. doi:10.1007/978-3-662-49674-9_12.
- 29 Daniel Neider, Alexander Weinert, and Martin Zimmermann. Synthesizing optimally resilient controllers. In *CSL 2018*, volume 119 of *LIPICs*, pages 34:1–34:17. Schloss Dagstuhl - LZI, 2018. doi:10.4230/LIPICs.CSL.2018.34.
- 30 Thomas W. Reps, Susan Horwitz, and Shmuel Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *POPL 1995*, pages 49–61. ACM Press, 1995. doi:10.1145/199448.199462.
- 31 Thomas W. Reps, Akash Lal, and Nicholas Kidd. Program analysis using weighted pushdown systems. In *FSTTCS 2007*, volume 4855 of *LNCS*, pages 23–51. Springer, 2007. doi:10.1007/978-3-540-77050-3_4.
- 32 Gérard Sénizergues. $L(a)=l(b)$? decidability results from complete formal systems. *Theor. Comput. Sci.*, 251(1-2):1–166, 2001. doi:10.1016/S0304-3975(00)00285-1.
- 33 Gérard Sénizergues. The bisimulation problem for equational graphs of finite out-degree. *SIAM J. Comput.*, 34(5):1025–1106, 2005. doi:10.1137/S0097539700377256.
- 34 Olivier Serre. Parity games played on transition graphs of one-counter processes. In *FOSSACS 2006*, volume 3921 of *LNCS*, pages 337–351. Springer, 2006. doi:10.1007/11690634_23.
- 35 Jiří Srba. Roadmap of infinite results. In Gheorghe Paun, Grzegorz Rozenberg, and Arto Salomaa, editors, *Current Trends in Theoretical Computer Science*, pages 337–350. World Scientific, 2004. URL: https://doi.org/10.1142/9789812562494_0054.
- 36 Paulo Tabuada, Sina Yamac Caliskan, Matthias Rungger, and Rupak Majumdar. Towards robustness for cyber-physical systems. *IEEE Trans. Automat. Contr.*, 59(12):3151–3163, 2014. doi:10.1109/TAC.2014.2351632.
- 37 Paulo Tabuada and Daniel Neider. Robust linear temporal logic. In *CSL 2016*, volume 62 of *LIPICs*, pages 10:1–10:21. Schloss Dagstuhl - LZI, 2016. doi:10.4230/LIPICs.CSL.2016.10.
- 38 Ufuk Topcu, Necmiye Ozay, Jun Liu, and Richard M. Murray. On synthesizing robust discrete controllers under modeling uncertainty. In *HSCC 2012*, pages 85–94. ACM, 2012. doi:10.1145/2185632.2185648.
- 39 Leslie G. Valiant. *Decision Procedures for Families of Deterministic Pushdown Automata*. PhD thesis, University of Warwick, 1973.

- 40 Yaron Velner and Alexander Rabinovich. Church synthesis problem for noisy input. In *FoSSaCS 2011*, volume 6604 of *LNCS*, pages 275–289. Springer, 2011. URL: https://doi.org/10.1007/978-3-642-19805-2_19.
- 41 Igor Walukiewicz. Pushdown processes: Games and model-checking. *Inf. Comput.*, 164(2):234–263, 2001. doi:10.1006/inco.2000.2894.

APPENDIX

In the appendix, we consider reachability games with disturbances (Appendix A) and present the proofs omitted in the main part (Appendix B).

A Beyond Safety: Reachability Games with Disturbances

Up to now, we were concerned with pushdown safety games with disturbances and have shown that they provide a rich model with interesting properties that go beyond the rather straightforward case of finite safety games with disturbances. Nevertheless, there are many other winning conditions that can be studied in pushdown games with disturbances. Probably the simplest class of winning condition besides safety conditions are reachability conditions: Given a set $F \subseteq V$ of vertices, the reachability condition $\text{Reach}(F) = \{v_0 v_1 v_2 \dots \mid v_j \in F \text{ for some } j \in \omega\}$ requires to visit F at least once. While safety and reachability conditions are dual, games with disturbances are asymmetric. Thus, we cannot directly transfer results for safety to reachability games and vice versa.


A.1 Resilience in Pushdown Reachability Games

Many of our results trivially carry over to reachability conditions while others can be recovered with some more effort. We begin by claiming that, in contrast to safety games, all possible resilience values can be realized in reachability games.

► **Lemma 16.** *All possible resilience values $\alpha \in \omega + 2$ can be realized in a one-counter reachability game that has vertices of resilience ω with and without a uniform witness.*

Proof. Consider the one-counter reachability game \mathcal{G} presented in Figure 2 where the reachability condition is induced by the doubly-lined vertices, i.e., Player 0 wins if and only if a doubly-lined vertex is visited. For every $\alpha \in \omega + 2$, there is a vertex v with $r_{\mathcal{G}}(v) = \alpha$, where the lower vertex of resilience ω has a uniform witness while the upper row of vertices does not, for reasons that are analogous to the ones presented in Example 2: Essentially, the upper row of vertices implements the fresh vertex v to obtain \mathcal{G}' . ◀

Thus, let us consider the problem of determining the resilience of the initial vertex of a pushdown reachability game. Lemma 6 and Lemma 7 are formulated for parity games, and therefore hold in particular for reachability games, as these are subsumed by parity



■ **Figure 2** A one-counter reachability game with all possible resilience values (depicted as labels below vertices). All vertices belong to V_0 and disturbance edges are dashed.

games. Thus, we can determine whether the initial vertex of a pushdown reachability game has resilience $\omega + 1$, ω with a uniform witness, or k , for a fixed k . However, as the value could also be ω without a uniform witness (for which we have no characterization) and as we have no upper bound on possible finite values, we do not obtain a complete algorithm determining the resilience of the initial vertex of every game.

Such an upper bound would immediately yield a complete algorithm similar to Algorithm 1: one just has to add a line checking whether the resilience is ω (and returning that result) and if the resilience is not equal to some k below the upper bound, then the algorithm returns “ ω without uniform witness”.

Finally, the examples witnessing the lower bounds presented in Lemma 10 can be turned into reachability games. Intuitively, one uses disturbances to push content on the stack, i.e., replaces the moves of Player 1 by disturbance edges.

A.2 Optimal Strategies in One-counter Reachability Games

In this subsection, we build a bridge between resilience in pushdown safety games and a classical problem in the theory of infinite games on infinite arenas: computing optimal strategies in pushdown reachability games, i.e., winning strategies that reach a fixed set F of target states in the least number of steps possible. This problem has been first studied by Cachat [8] and recently been revisited by Carayol and Hague [10]. Here, we first present a connection between both problems and then prove a, to the best of our knowledge, novel result about optimal strategies in one-counter games.

Fix a reachability game $\mathcal{G} = (\mathcal{A}, \text{Reach}(F))$ with a disturbance-free finitely branching arena $\mathcal{A} = (V, V_0, V_1, E, \emptyset)$. Given a play $\rho = (v_0, 0)(v_1, 0)(v_2, 0) \cdots$, let

$$\text{val}_{\mathcal{G}}(\rho) = \min\{j \in \omega \mid v_j \in F\}$$

where we define $\min \emptyset = \omega + 1$ for technical convenience. Hence, $\text{val}_{\mathcal{G}}(\rho)$ is the minimal position in F , if $v_0 v_1 v_2 \cdots \in \text{Reach}(F)$, and $\omega + 1$ if $v_0 v_1 v_2 \cdots \notin \text{Reach}(F)$. Given a strategy σ for Player 0, we define $\text{val}_{\mathcal{G}}(v, \sigma) = \sup_{\rho} \text{val}_{\mathcal{G}}(\rho)$, where ρ ranges over all plays starting in v that are consistent with σ . Using König’s Lemma shows that σ is a winning strategy for Player 0 in $(\mathcal{A}, \text{Reach}(F))$ from v if and only if $\text{val}_{\mathcal{G}}(v, \sigma) < \omega$ (here we use the assumption that \mathcal{A} is finitely branching). Otherwise, i.e., if σ is not winning from v , then we have $\text{val}_{\mathcal{G}}(v, \sigma) = \omega + 1$. A strategy σ for Player 0 is reachability optimal if $\text{val}_{\mathcal{G}}(v, \sigma) \leq \text{val}_{\mathcal{G}}(v, \sigma')$ for every $v \in V$ and every strategy σ' for Player 0. The existence of reachability optimal strategies follows straightforwardly from the correctness proof of the attractor construction for reachability games [18].

► **Proposition 17.** *Every reachability game in a finitely branching arena has a reachability optimal strategy.*

Now, define $\mathcal{A}' = (V \cup E, V_1 \cup E, V_0, E', D)$ with

- $E' = \{(v, (v, v')), ((v, v'), (v, v')) \mid (v, v') \in E\}$ and
- $D = \{((v, v'), v') \mid (v, v') \in E\}$,

i.e., we split every edge $(v, v') \in E$ into a sequence of a standard edge from v to the newly introduced vertex (v, v') , a standard self-loop at (v, v') , and a disturbance edge from (v, v') to v' . Intuitively, it takes one disturbance in \mathcal{A}' to simulate a move in \mathcal{A} while the overall structure of the arena, including the strategic choices for the players, is preserved. Note that we flip the positions of Player 0 and 1 when turning \mathcal{A} into \mathcal{A}' . Hence, we also dualize the winning condition and turn $\text{Reach}(F)$ into the dual safety condition $\text{Safety}(F) = V^\omega \setminus \text{Reach}(F)$ and

define $\mathcal{G}' = (\mathcal{A}', \text{Safety}(F))$. If \mathcal{G} is a pushdown (one-counter) reachability game, then \mathcal{G}' is a pushdown (one-counter) safety game and the blow-up of the transformation is polynomial.

Now, we relate the resilience of vertices in \mathcal{G}' with the value of optimal strategies in \mathcal{G} .

► **Lemma 18.** *A reachability optimal strategy σ satisfies $\text{val}_{\mathcal{G}}(v, \sigma) = r_{\mathcal{G}'}(v)$ for every $v \in V$.*

Proof. Fix a reachability optimal strategy σ_{opt} for Player 0 in \mathcal{G} . We sketch the proof ideas, but leave the straightforward details to the reader.

$\text{val}_{\mathcal{G}}(v, \sigma_{\text{opt}}) \leq r_{\mathcal{G}'}(v)$: The statement is trivial if $r_{\mathcal{G}'}(v) = \omega + 1$. Hence, assume $r_{\mathcal{G}'}(v) \in \omega$, say $r_{\mathcal{G}'}(v) = k$. Then, Player 1 has a winning strategy τ for $(\mathcal{A}'_{\text{rig}}, \text{Safety}(F)_{\text{rig}} \cup R_{\geq k+1})$ from v due Lemma 5.3. This strategy can be turned into a strategy σ for Player 0 in \mathcal{A} that mimics τ while ignoring the auxiliary vertices in $\mathcal{A}'_{\text{rig}}$ that are not in \mathcal{A} . An induction shows that F is reached within k moves when starting in v and playing according to σ . Thus, $\text{val}(v, \sigma_{\text{opt}}) \leq \text{val}(v, \sigma) \leq k = r_{\mathcal{G}'}(v)$.

$r_{\mathcal{G}'}(v) \leq \text{val}_{\mathcal{G}}(v, \sigma_{\text{opt}})$: The statement is trivial, if $\text{val}_{\mathcal{G}}(v, \sigma_{\text{opt}}) = \omega + 1$. Hence, assume $\text{val}_{\mathcal{G}}(v, \sigma_{\text{opt}}) \in \omega$, say $\text{val}_{\mathcal{G}}(v, \sigma_{\text{opt}}) = k$. Then, by definition, F is reached within k moves when starting in v and playing according to σ_{opt} . The strategy σ_{opt} for Player 0 in \mathcal{G} can be turned into a strategy τ for Player 1 in $\mathcal{A}'_{\text{rig}}$ that simulates the moves of σ_{opt} and, as long as F has not been visited, simulate a disturbance whenever possible (there is a unique disturbance edge at every vertex with outgoing disturbance edges). After visiting F , no more disturbances are simulated. An induction shows that F is reached and at most k disturbances are simulated when starting in v and playing according to τ . Hence, Player 1 wins $(\mathcal{A}'_{\text{rig}}, \text{Safety}(F)_{\text{rig}} \cup R_{\geq k+1})$, which implies $r_{\mathcal{G}'}(v) \leq k = \text{val}_{\mathcal{G}}(v, \sigma_{\text{opt}})$ by Lemma 5.3. ◀

Using the connection between resilience and values of reachability optimal strategies allows us to compute the value of a reachability optimal strategy in the initial vertex of a pushdown game. In particular, for one-counter systems, we obtain an algorithm with polynomial space requirements. Thereby, we close a gap in our knowledge about reachability optimal strategies in pushdown games.

► **Theorem 19.** *The following problem can be solved in polynomial space: “Given a one-counter reachability game \mathcal{G} with initial vertex v_I , determine $\text{val}_{\mathcal{G}}(v_I, \sigma)$ for a reachability optimal strategy σ ”.*

Note that the approach via a reduction to computing the resilience presented here is not the simplest one: One could simplify the constructions presented in Section 6 and obtain a direct algorithm.

B Proofs Omitted in the Main Part

We begin by stating some simple properties of resilient strategies and resilience values that are useful throughout the appendix.

► **Remark 20.** A strategy σ does not have access to the bits indicating whether a disturbance occurred or not. However, this is not a restriction: let $(v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ be a play with $b_j = 1$ for some $j > 0$. We say that this disturbance is consequential (w.r.t. σ), if $v_j \neq \sigma(v_0 \cdots v_{j-1})$, i.e., if the disturbance transition (v_{j-1}, v_j) traversed by the play did not lead to the vertex the strategy prescribed. Such consequential disturbances can be detected by comparing the actual vertex v_j to σ 's output $\sigma(v_0 \cdots v_{j-1})$. On the other hand, inconsequential disturbances will just be ignored. In particular, the number of consequential disturbances is always at most the number of disturbances.

The following remark lists some simple consequences of the definition of resilience.

► **Remark 21.** The following hold for every vertex v .

1. Every strategy is 0-resilient from v .
2. A strategy is 1-resilient from v if and only if it is winning for Player 0 from v .
3. If a strategy is α -resilient from v and $\alpha > \alpha'$ then it is also α' -resilient from v .

Note that every game has disjoint winning regions. A game is determined, if every vertex is in either winning region. The previous remark implies that resilience refines winning regions.

► **Lemma 22.** *Let \mathcal{G} be a game and v a vertex of \mathcal{G} .*

1. $r_{\mathcal{G}}(v) > 0$ if and only if $v \in \mathcal{W}_0(\mathcal{G})$.
2. If \mathcal{G} is determined, then $r_{\mathcal{G}}(v) = 0$ if and only if $v \in \mathcal{W}_1(\mathcal{G})$.

Proof. 1.) The resilience of v is greater than zero if and only if Player 0 has a 1-resilient strategy from v due to Item 3 of Remark 21. The latter condition is equivalent to Player 0 having a winning strategy for \mathcal{G} from v , i.e., equivalent to $v \in \mathcal{W}_0(\mathcal{G})$, due to Item 2 of Remark 21.

2.) Due to Items 3 and 2 of Remark 21, the resilience of v is zero if and only if Player 0 has no winning strategy for \mathcal{G} from v , i.e., $v \notin \mathcal{W}_0(\mathcal{G})$. Due to determinacy, this is equivalent to $v \in \mathcal{W}_1(\mathcal{G})$. ◀

Note that determinacy is a necessary condition for Item 2. In an undetermined game, the vertices that are in neither winning region have resilience zero, due to Item 1, but are in particular not in $\mathcal{W}_1(\mathcal{G})$.

For determined disturbance-free games, i.e., those without disturbance edges in the arena, we obtain a tighter connection between resilience and winning regions: There are only two possible resilience values and they characterize the winning regions.

► **Remark 23.** Let \mathcal{G} be a determined disturbance-free game and v a vertex of \mathcal{G} .

1. $r_{\mathcal{G}}(v) = \omega + 1$ if and only if $v \in \mathcal{W}_0(\mathcal{G})$.
2. $r_{\mathcal{G}}(v) = 0$ if and only if $v \in \mathcal{W}_1(\mathcal{G})$.

To conclude, notice that one can turn a safety game into a parity game. Formally, given a coloring $\Omega: V \rightarrow \omega$, which is required to have a finite range $\Omega(V)$, $\text{Parity}(\Omega)$ containing the sequences $v_0v_1v_2 \cdots \in V^\omega$ with even

$$\limsup \Omega(v_0)\Omega(v_1)\Omega(v_2) \cdots$$

denotes the (max-) parity condition induced by Ω , which requires the maximal color occurring infinitely often during a play to be even. As the range of Ω is finite, every play has a maximal color occurring infinitely often.

Now, given a safety game $\mathcal{G} = (\mathcal{A}, \text{Safety}(F))$ with vertex set V we turn all vertices in F into sinks, obtaining the arena \mathcal{A}' with vertex set V . Then, we have $r_{\mathcal{G}}(v) = r_{\mathcal{G}'}(v)$, where $\mathcal{G}' = (\mathcal{A}', \text{Parity}(\Omega))$ for the coloring Ω mapping vertices in F to 1 and all other vertices to 2.

B.1 Proofs Omitted in Section 3

We begin by giving a characterization of the resilience values in finitely branching safety games that will be the basis of both the proof of Lemma 3.2 and the proof of Theorem 4. The characterization is a generalization of a similar one for safety games in finite arenas [13].

Fix a finitely branching safety game $\mathcal{G} = (\mathcal{A}, \text{Safety}(F))$ with $\mathcal{A} = (V, V_0, V_1, E, D)$. First, we recall the attractor construction for Player 1. Fix a set $X \subseteq V$. Let $A_0 = X$ and define, for every $j \geq 0$, A_{j+1} as follows.

$$A_{j+1} = A_j \cup \{v \in V_0 \mid \text{for all } (v, v') \in E: v' \in A_j\} \cup \\ \{v \in V_1 \mid \text{there exists } (v, v') \in E \text{ with } v' \in A_j\}$$

We call $\text{Att}_1(X) = \bigcup_{j \in \omega} A_j$ the 1-attractor of X in \mathcal{A} .

By construction, Player 1 has a positional strategy τ such that every disturbance-free play starting in $\text{Att}_1(X)$ and being consistent with τ visits X at least once. Dually, Player 0 has a positional strategy σ such that every disturbance-free play starting in $V \setminus \text{Att}_1(X)$ and being consistent with σ never visits X . We refer to τ and σ as the attractor and trap strategy associated to $\text{Att}_1(X)$. Finally, we call

$$\text{Bnd}_D(X) = \{v \in V_0 \setminus X \mid \text{there exists } (v, v') \in D \text{ with } v' \in X\}$$

the D -boundary of X , which contains all vertices $v \notin X$ from which a disturbance edge leads into X .

In the following, we alternately apply the attractor and the boundary operation starting with the set F of vertices that Player 0 has to avoid in order to win. Then, we show that every vertex in the limit has finite resilience while every other vertex has resilience $\omega + 1$, which completes the proof.

Formally, let $S_0 = \text{Att}_1(F)$ be the 1-attractor of F , $S_{j+1} = \text{Att}_1(S_j \cup \text{Bnd}_D(S_j))$ for every $j \in \omega$, and define $S = \bigcup_{j \in \omega} S_j$. Now, for $v \in S$, let $r(v) = \min\{j \mid v \in S_j\}$ be the index at which v is added to S .

B.1.1 Proof of Lemma 3

► **Lemma 3.** *Let \mathcal{G} be a safety game with vertex set V .*

1. *There is no $v \in V$ with $r_{\mathcal{G}}(v) = \omega$ that has a uniform witness.*
2. *If \mathcal{A} is finitely branching, then there is no $v \in V$ with $r_{\mathcal{G}}(v) = \omega$.*

Proof. 1.) Let $\mathcal{G} = (\mathcal{A}, \text{Safety}(F))$. Towards a contradiction assume that there is a vertex $v \in V$ with $r_{\mathcal{G}}(v) = \omega$ and that there is a strategy σ that is ω -resilient from v . Due to $r_{\mathcal{G}}(v) < \omega + 1$, σ is not $(\omega + 1)$ -resilient from v . Thus, there is a play $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ that starts in v , is consistent with σ , satisfies $\#_D(\rho) < \omega + 1$ (which is a tautology), and such that $v_0 v_1 v_2 \cdots \notin \text{Safety}(F)$, i.e., there is a j such that $v_j \in F$. Consider a play of the form $\rho' = (v_0, b_0) \cdots (v_j, b_j) \rho''$ that is consistent with σ and such that $(v_j, b_j) \rho''$ is disturbance-free. Such a play exists, as each vertex in V_0 has a non-disturbance successor. The play ρ' starts in v , is consistent with σ , satisfies $\#_D(\rho') \leq j$, as disturbances can only occur in the prefix $(v_0, b_0) \cdots (v_j, b_j)$, but violates the safety condition, as $v_j \in F$ is visited by ρ' . Therefore, σ is not $(j + 1)$ -resilient from v , and in particular not ω -resilient from v , which contradicts our assumption.

2.) Let $\mathcal{G} = (\mathcal{A}, \text{Safety}(F))$ with finitely branching $\mathcal{A} = (V, V_0, V_1, E, D)$ and let the corresponding values $r(v)$ and the set S be defined as on Page 20. We claim $r_{\mathcal{G}}(v) \leq r(v)$ for every $v \in S$ and $r_{\mathcal{G}}(v) = \omega + 1$ for every $v \notin S$, which proves our claim.

Fix a vertex $v \in S$. To show $r_{\mathcal{G}}(v) \leq r(v)$, we need to show for every strategy σ for Player 0 that there is a play that starts in v , is consistent with σ , has at most $r(v)$ disturbances, and is losing for Player 0, i.e., it visits F at least once. We fix any strategy σ and construct such a play inductively starting with the play prefix $(v_0, b_0) = (v, 0)$. During the construction, we ensure that the prefix constructed thus far is consistent with σ and that it ends in S . Thus, assume we have constructed a play prefix $w = (v_0, b_0) \cdots (v_j, b_j)$ satisfying the invariant. To extend it, we distinguish two cases:

1. Assume $r(v_j) = 0$, i.e., $v_j \in S_0 = \text{Att}_1(F)$. Then, consider the unique disturbance-free play $(v_j, 0)\rho$ consistent with σ and the attractor strategy for Player 1 associated with $\text{Att}_1(F)$. We extend w by ρ to complete the construction of the desired play. The resulting play $w\rho$ is consistent with σ due to our invariant and the choice of ρ , and contains a vertex from F .
2. Assume $r(v_j) > 0$, i.e., $v_j \in S_{r(v_j)} = \text{Att}_1(S_{r(v_j)-1} \cup \text{Bnd}_D(S_{r(v_j)-1}))$. Consider the unique disturbance-free play $(v_j, 0)\rho$ consistent with σ and the attractor strategy for Player 1 associated with

$$\text{Att}_1(S_{r(v_j)-1} \cup \text{Bnd}_D(S_{r(v_j)-1})).$$

Let $(v_j, 0)(v_{j+1}, 0) \cdots (v_{j+j'}, 0)$ be the minimal prefix of $(v_j, 0)\rho$ such that $v_{j+j'} \in S_{r(v_j)-1} \cup \text{Bnd}_D(S_{r(v_j)-1})$. If $v_{j+j'} \in S_{r(v_j)-1}$ (which implies $j' > 0$ due to $v_j \notin S_{r(v_j)-1}$) then we extend w to $w(v_{j+1}, 0) \cdots (v_{j+j'}, 0)$ to obtain the next prefix in our inductive construction. If $v_{j+j'} \in \text{Bnd}_D(S_{r(v_j)-1})$, then there is a vertex $v_{j+j'+1} \in S_{r(v_j)-1}$ and $(v_{j+j'}, v_{j+j'+1}) \in D$ due to the definition of the D -boundary. Thus, we extend w to $w(v_{j+1}, 0) \cdots (v_{j+j'}, 0)(v_{j+j'+1}, 1)$ to obtain the next prefix in our inductive construction. The resulting prefix is consistent with σ and its last vertex is in $S_{r(v_j)-1} \subseteq S$, i.e., our invariant is satisfied.

Now, let $v_{j_0}, v_{j_1}, v_{j_2}, \dots$ be the sequence of last vertices of the prefixes obtained during the construction. In particular, $v_{j_0} = v$. By construction, we have $r(v_{j_0}) > r(v_{j_1}) > r(v_{j_2}) \cdots$. Hence, we apply the second case at most $r(v_{j_0})$ many times and then have to apply the first case. Hence, we indeed obtain an infinite play ρ starting in v , which is consistent with σ due to our invariant, and which visits F , as the first case is eventually applied. Finally, ρ has at most $r(v_{j_0}) = r(v)$ many disturbances, as each application of the second case adds at most one disturbance edge and the first case adds none. Thus, ρ witnesses that σ is not $(r(v) + 1)$ -resilient from v . As we have picked σ arbitrarily, we conclude $r_{\mathcal{G}}(v) \leq r(v)$ as desired.

It remains to show $r_{\mathcal{G}}(v) = \omega + 1$ for every $v \notin S$. We start by listing some properties of such vertices:

1. $v \notin F$, as $F \subseteq \text{Att}_1(F) = S_0 \subseteq S$.
2. If $v \in V_0$, then there is a v' with $(v, v') \in E$ and $v' \notin S$. Towards a contradiction, assume there is no such v' . Then, all successors of v are in S . As v has only finitely many successors by assumption on \mathcal{A} , there is a j such that all these successors are in S_j . Hence, $v \in \text{Att}_1(S_j) \subseteq S_{j+1} \subseteq S$, which contradicts $v \notin S$.
3. If $v \in V_1$, then all v' with $(v, v') \in E$ satisfy $v' \notin S$. Towards a contradiction, assume there is a successor of v in S . Then, v' is in some S_j and $v \in \text{Att}_1(S_j) \subseteq S_{j+1} \subseteq S$, which contradicts $v \notin S$.
4. If $v \in V_0$ and $(v, v') \in D$, then $v' \notin S$. Again, towards a contradiction assume there is a disturbance edge leading from v to v' in S . Then, v' is in some S_j and $v \in \text{Bnd}_D(S_j) \subseteq S_{j+1} \subseteq S$, which contradicts $v \notin S$.

Thus, due to Property 2, Player 0 must have a positional strategy σ that moves from any vertex $v \notin S$ to some successor $v' \notin S$. Now, consider a play ρ that starts in a vertex $v \notin S$, is consistent with σ , and has an arbitrary number of disturbances. It starts outside of S , Player 0 does not move into S by definition of σ , Player 1 cannot due to Property 3, and disturbances do not lead into S due to Property 4. Hence, ρ never visits S and thus also avoids F , due to Property 1. Hence, ρ is winning for Player 0. As v and ρ are arbitrary, we have shown $r_{\mathcal{G}}(v) = \omega + 1$ for every $v \notin S$. \blacktriangleleft

B.1.2 Proof of Theorem 4

► **Theorem 4.** *Player 0 has positional optimally resilient strategies in finitely branching safety games.*

Proof. Let $\mathcal{G} = (\mathcal{A}, \text{Safety}(F))$ with finitely branching $\mathcal{A} = (V, V_0, V_1, E, D)$, and let the values $r(v)$ and the set S be defined as on Page 20. We have shown $r_{\mathcal{G}}(v) \leq r(v)$ for every $v \in S$ and $r_{\mathcal{G}}(v) = \omega + 1$ for every $v \notin S$ in the proof of Lemma 3.2. We now show $r_{\mathcal{G}}(v) \geq r(v)$ for every $v \in S$.

To simplify our notation, let $X_0 = F$ and $X_{j+1} = S_j \cup \text{Bnd}_D(S_j)$, i.e., $S_j = \text{Att}_1(X_j)$ for every j . Now, for every $j \in \omega$, let σ_j be the trap strategy for Player 0 associated with $S_j = \text{Att}_1(X_j)$, i.e., every disturbance-free play that starts in $V \setminus S_j$ and is consistent with σ_j never visits X_j . Recall that we defined $r(v) = \min\{j \mid v \in S_j\}$ for all $v \in S$. Thus, if $r(v) > 0$, then $v \notin S_{j-1}$.

We define a positional strategy σ for Player 0 as follows:

- If $v \in V_0 \cap S$ with $r(v) > 0$ then $\sigma(v) = \sigma_{r(v)-1}(v)$.
- If $v \in V_0 \cap S$ with $r(v) = 0$ then $\sigma(v) = v'$ for some arbitrary successor v' of v .
- If $v \in V_0 \setminus S$ then $\sigma(v) = v'$ for some successor v' of v with $v' \notin S$. We have argued in the proof of Lemma 3.2, that such a successor always exists if $v \notin S$.

Fix some $v \in S$ and consider a play $\rho = (\rho_0, b_0)(\rho_1, b_1)(\rho_2, b_2) \cdots$ starting in $v \in S$, consistent with σ , and with $k < r(v)$ disturbances. A straightforward induction on $j \geq 0$ shows that $r(\rho_j) \geq r(v) - \#_D((\rho_0, b_0) \cdots (\rho_j, b_j))$ for every j . Thus, $r(\rho_j) \geq r(v) - k > 0$, which implies $\rho_j \notin F \subseteq S_0$, i.e., ρ is winning for Player 0.

Therefore, σ is $r(v)$ -resilient from every $v \in S$. Conversely, in the proof of Lemma 3.2, we have shown $r_{\mathcal{G}}(v) \leq r(v)$. Hence, $r(v) = r_{\mathcal{G}}(v)$, i.e., σ is $r_{\mathcal{G}}(v)$ -resilient from every $v \in S$. Furthermore, the arguments presented in the proof of Lemma 3.2 for vertices $v \notin S$ show that σ is $(\omega + 1)$ -resilient from every $v \notin S$.

Altogether, σ is optimally resilient. \blacktriangleleft


B.2 Proofs Omitted in Section 4

We illustrate the definition of the construction of the rigged arena in Figure 3.

B.2.1 Proof of Lemma 5

► **Lemma 5.** *Let $\mathcal{G} = (\mathcal{A}, \text{Win})$ be a game, let v be a vertex of \mathcal{G} , and let $k \in \omega$.*

1. *Player 0 has an $(\omega + 1)$ -resilient strategy for \mathcal{G} from v if and only if $v \in \mathcal{W}_0(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}})$.*
2. *Player 0 has an ω -resilient strategy for \mathcal{G} from v if and only if $v \in \mathcal{W}_0(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}} \cup \text{Büchi}(D))$.*
3. *Player 0 has a k -resilient strategy for \mathcal{G} from v if and only if $v \in \mathcal{W}_0(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}} \cup R_{\geq k})$.*



■ **Figure 3** The rigged arena \mathcal{A}_{rig} for the arena \mathcal{A} presented in Figure 1, restricted to vertices reachable from the initial vertex (q_I, \perp) . Round vertices are in V_0 , square ones in V_1 , and a gray vertex indicates that a disturbance has been simulated.

Proof. We begin by introducing translations between plays that are useful in all three cases.

First, we translate a play prefix w in \mathcal{A} into a play prefix $t'(w)$ in \mathcal{A}_{rig} satisfying the following invariant:

$$t'((v_0, b_0) \cdots (v_j, b_j))$$

starts in v_0 and ends in v_j . We proceed by induction starting with $t'(v_0, b_0) = (v_0, 0)$. For the induction step, we have to consider a play prefix $(v_0, b_0) \cdots (v_j, b_j)(v_{j+1}, b_{j+1})$ such that $t'((v_0, b_0) \cdots (v_j, b_j))$ is already defined, which ends in v_j due to our invariant. We consider several cases:

- If $b_{j+1} = 1$, then (v_j, v_{j+1}) is a disturbance edge, which is simulated in \mathcal{A}_{rig} by Player 1 taking control at v_j , moving to (v_j, v_{j+1}) and then to v_{j+1} . Hence, we define

$$t'((v_0, b_0) \cdots (v_j, b_j)(v_{j+1}, b_{j+1})) = t'((v_0, b_0) \cdots (v_j, b_j)) \cdot ((v_j, v_{j+1}), 0)(v_{j+1}, 0).$$

- If $b_{j+1} = 0$ and $v_j \in V_0$, then (v_j, v_{j+1}) is a non-disturbance edge picked by Player 0, which is simulated in \mathcal{A}_{rig} by Player 1 ceding control at v_j to Player 0 by moving \bar{v}_j , from where Player 0 can then move to v_{j+1} . Hence, we define

$$t'((v_0, b_0) \cdots (v_j, b_j)(v_{j+1}, b_{j+1})) = t'((v_0, b_0) \cdots (v_j, b_j)) \cdot (\bar{v}_j, 0)(v_{j+1}, 0).$$

- If $b_{j+1} = 0$ and $v_j \in V_1$, then (v_j, v_{j+1}) is a non-disturbance edge picked by Player 1, which is simulated in \mathcal{A}_{rig} by Player 1 directly moving to v_{j+1} . Hence, we define

$$t'((v_0, b_0) \cdots (v_j, b_j)(v_{j+1}, b_{j+1})) = t'((v_0, b_0) \cdots (v_j, b_j)) \cdot ((v_j, v_{j+1}), 0)(v_{j+1}, 0).$$

In each case, the invariant is satisfied and

$$t'((v_0, b_0) \cdots (v_j, b_j)(v_{j+1}, b_{j+1}))$$

is indeed a play prefix due to $t'((v_0, b_0) \cdots (v_j, b_j))$ ending in v_j .

Furthermore, we extend t' to infinite plays by defining $t'((v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots)$ to be the unique play ρ' in \mathcal{A}_{rig} such that $t'((v_0, b_0) \cdots (v_j, b_j))$ is a prefix of ρ' for every $j \in \omega$. Let $\rho = (v_0, 0)(v_1, 0)(v_2, 0) \cdots$ be a play in \mathcal{A} . Then we have $t'(\rho) = (v_0, 0)(a_0, 0)(v_1, 0)(a_1, 0)(v_2, 0)(a_2, 0) \cdots$ for auxiliary vertices $a_0 a_1 a_2 \cdots$ and $\#_D(\rho) = |\{j \mid a_j \in D\}|$, i.e., the number of disturbances during a play ρ in \mathcal{A} is equal to the number of vertices from $D \subseteq A$ occurring in $t'(\rho)$.

Finally, we can use the translation t' to transform a strategy σ' for Player 0 in \mathcal{A}_{rig} to a strategy σ for her in \mathcal{A} . To this end, let b^- denote the homomorphism from $(V' \times \{0, 1\})^*$ to $(V')^*$ that removes the second component. Then, we define

$$\sigma(v_0 \cdots v_j) = \sigma'(b^-(t'((v_0, b_0) \cdots (v_j, b_j)))) \cdot \bar{v}_j$$

where $b_0 = 0$ and for every $0 < j' \leq j$, $b_{j'} = 1$ if and only if $v_{j'-1} \in V_0$ and $v_{j'} \neq \sigma(v_0 \cdots v_{j'-1})$, i.e., we reconstruct the consequential disturbances with respect to σ as defined thus far. A simple induction shows that a play ρ in \mathcal{A} being consistent with σ implies that $t'(\rho)$ in \mathcal{A}_{rig} is consistent with σ' .

Now, we consider the other direction and translate a play prefix w in \mathcal{A}_{rig} into a play prefix $t(w)$ in \mathcal{A} . Here, we only consider play prefixes w starting and ending in a vertex from $V' \setminus A$, i.e., only play prefixes that do not start or end in one of the auxiliary vertices. This satisfies the following invariant: $t((v_0, 0) \cdots (v_j, 0))$ starts in v_0 and ends in v_j (recall that \mathcal{A}_{rig} has no disturbance edges, which implies that all bits b_j in w are equal to zero). Again, we proceed by induction and start with $t(v_0, 0) = (v_0, 0)$. For the induction step, consider a play prefix $(v_0, 0) \cdots (v_j, 0)(a_j, 0)(v_{j+1}, 0)$ such that $t((v_0, b_0) \cdots (v_j, b_j))$ is already defined, which ends in v_j due to our invariant.

- If the prefix is of the form

$$(v_0, 0) \cdots (v_j, 0)((v_j, v_{j+1}), 0)(v_{j+1}, 0)$$

with $v_j \in V_0$, then the last move simulated during the play prefix is the disturbance edge $(v_j, v_{j+1}) \in D$. Hence, we define

$$t((v_0, 0) \cdots (v_j, 0)((v_j, v_{j+1}), 0)(v_{j+1}, 0)) = t((v_0, 0) \cdots (v_j, 0)) \cdot (v_{j+1}, 1).$$

- If the prefix is of the form

$$(v_0, 0) \cdots (v_j, 0)(\bar{v}_j, 0)(v_{j+1}, 0),$$

then the last move simulated during the play prefix is the non-disturbance edge $(v_j, v_{j+1}) \in E$ with $v_j \in V_0$. Hence, we define

$$t((v_0, 0) \cdots (v_j, 0)(\bar{v}_j, 0)(v_{j+1}, 0)) = t((v_0, 0) \cdots (v_j, 0)) \cdot (v_{j+1}, 0).$$

- If the prefix is of the form

$$(v_0, 0) \cdots (v_j, 0)((v_j, v_{j+1}), 0)(v_{j+1}, 0)$$

with $v_j \in V_1$, then the last move simulated during the play prefix is the non-disturbance edge $(v_j, v_{j+1}) \in E$. Hence, we define

$$t((v_0, 0) \cdots (v_j, 0)((v_j, v_{j+1}), 0)(v_{j+1}, 0)) = t((v_0, 0) \cdots (v_j, 0)) \cdot (v_{j+1}, 0).$$

In each case, the invariant is satisfied and the extension is indeed a play prefix due to $t((v_0, 0) \cdots (v_j, 0))$ ending in v_j .

Again, we extend the function t to infinite plays by defining $t((v_0, 0)(v_1, 0)(v_2, 0) \cdots)$ to be the unique play ρ in \mathcal{A} such that $t((v_0, 0) \cdots (v_j, 0))$ is a prefix of ρ for every $j \in \omega$. Let $\rho' = (v_0, 0)(a_0, 0)(v_1, 0)(a_1, 0)(v_2, 0)(a_2, 0) \cdots$ be a play in \mathcal{A}_{rig} starting in V . Hence, $t(\rho') = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ for some bits b_j , and $|\{j \mid a_j \in D\}| = \#_D(t(\rho'))$, i.e., the number of vertices from $D \subseteq A$ occurring in ρ' is equal to the number of disturbances during the play $t(\rho')$ in \mathcal{A} .

To conclude, we again show that we can use the translation t to transform a strategy σ for Player 0 in \mathcal{A} to a strategy σ' for her in \mathcal{A} . Here, let b^- denote the homomorphism from $(V \times \{0, 1\})^*$ to V^* that removes the second component in each letter. Now, we define

$$\sigma'(v_0 \cdots v_j \bar{v}_j) = \sigma(b^-(t((v_0, 0) \cdots (v_j, 0)))).$$

Finally, a simple induction shows that a play ρ' in \mathcal{A}_{rig} being consistent with σ' implies that $t(\rho')$ in \mathcal{A} is consistent with σ .

After these preparations, the proof of the three characterizations is straightforward employing the transformation of strategies described above.

1.) Let $v \in \mathcal{W}_0(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}})$, i.e., Player 0 has a winning strategy σ' from v . Let the strategy σ for Player 0 in \mathcal{A} be obtained from σ' as described above. We claim that σ is $(\omega + 1)$ -resilient from v . To this end, let $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ be a play in \mathcal{G} that starts in v , is consistent with σ , and has an arbitrary number of disturbances. We need to show that ρ is winning for Player 0, i.e., $v_0 v_1 v_2 \cdots \in \text{Win}$.

As argued above, the play $t'(\rho)$ in \mathcal{A}_{rig} is of the form

$$(v_0, 0)(a_0, 0)(v_1, 0)(a_1, 0)(v_2, 0)(a_2, 0) \cdots,$$

starts in v , and is consistent with σ' . This implies $t'(\rho) \in \text{Win}_{\text{rig}}$. So, by definition of Win_{rig} , we have indeed $v_0 v_1 v_2 \cdots \in \text{Win}$.

Now, assume Player 0 has an $(\omega + 1)$ -resilient strategy σ for \mathcal{G} from v . Let the strategy σ' for Player 0 in \mathcal{A}_{rig} be obtained from σ as described above. We claim that σ' is a winning strategy from v in the game $(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}})$. To this end, let $\rho' = (v_0, 0)(a_0, 0)(v_1, 0)(a_1, 0)(v_2, 0)(a_2, 0) \cdots$ be a play in \mathcal{A}_{rig} starting in v and consistent with σ' . We need to show that ρ' is winning for Player 0.

As argued above, the play $t(\rho') = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ in \mathcal{A} starts in v and is consistent with σ . Since σ is $(\omega + 1)$ -resilient from v , $t(\rho')$ is winning for Player 0, as it has at most ω disturbances. Thus, $v_0 v_1 v_2 \cdots \in \text{Win}$. Hence, $\rho' \in \text{Win}_{\text{rig}}$ by definition of Win_{rig} , i.e., ρ' is indeed winning for Player 0.

2.) As this proof is a refinement of the previous one, we only sketch the differences.

First, let $v \in \mathcal{W}_0(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}} \cup \text{Büchi}(D))$, i.e., Player 0 has a winning strategy σ' from v which induces a strategy σ for her in \mathcal{A} . We show that σ is ω -resilient from v . To this end, let $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ be a play in \mathcal{G} that starts in v , is consistent with σ , and has a finite number of disturbances. We need to show that ρ is winning for Player 0.

Again, the play $t'(\rho)$ in \mathcal{A}_{rig} starts in v and is consistent with σ' . Now, we additionally have that $t'(\rho)$ visits vertices in D only finitely often, as the number of these visits is equal to the number of disturbances in ρ , as argued above. Hence, $t'(\rho)$ is not in $\text{Büchi}(D)$, which implies $t'(\rho) \in \text{Win}_{\text{rig}}$, as $t'(\rho)$ is consistent with the winning strategy σ' . This allows us, as before, to conclude that ρ is indeed winning for Player 0.

Now, assume Player 0 has an ω -resilient strategy σ for \mathcal{G} from v and let σ' be the induced strategy for her in \mathcal{A}_{rig} . We show that σ' is winning from v in the game $(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}})$, i.e.,

every play $\rho' = (v_0, 0)(a_0, 0)(v_1, 0)(a_1, 0)(v_2, 0)(a_2, 0) \cdots$ in \mathcal{A}_{rig} starting in v and consistent with σ' is winning for Player 0.

If $v_0 a_0 v_1 a_1 v_2 a_2 \cdots$ is in $\text{Büchi}(D)$, then ρ' is winning for Player 0. Thus, assume it is not. Then, consider the play $t(\rho') = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ in \mathcal{A} . It starts in v , is consistent with σ , and has the same finite number of disturbances as ρ' has visits to vertices in D . Hence, as σ is ω -resilient from v , $t(\rho')$ is winning for Player 0. From this we can conclude, as before, that ρ' is indeed winning for Player 0.

3.) Analogously to the previous one arguing about “less than k disturbances” instead of “finitely many disturbances”. ◀

B.3 Proofs Omitted in Section 5

B.3.1 Proof of Lemma 6

► **Lemma 6.** *The following problem is EXPTIME-complete (and PSPACE-complete if inputs are restricted to one-counter games): “Given a pushdown parity game \mathcal{G} with initial vertex v_I and $\alpha \in \{\omega, \omega + 1\}$, does Player 0 have an α -resilient strategy for \mathcal{G} from v_I ?”. If yes, such a strategy can be computed in exponential time.*

Proof. Given a pushdown arena \mathcal{A} induced by a PDS \mathcal{P} with set Q of states, a partition $\{Q_0, Q_1\}$ of Q , and a transition relation Δ inducing the disturbance edges, a PDS \mathcal{P}' with set Q' of states and a partition $\{Q'_0, Q'_1\}$ of Q' inducing \mathcal{A}_{rig} can be computed in linear time. If \mathcal{P} is one-counter, then so is \mathcal{P}' . Further, given a coloring Ω of Q , one can determine

- a coloring Ω' of Q' such that $\text{Parity}(\Omega') = \text{Parity}(\Omega)_{\text{rig}}$, and
- a coloring Ω'' of Q' such that $\text{Parity}(\Omega'') = \text{Parity}(\Omega)_{\text{rig}} \cup \text{Büchi}(D)$, where D is the set of disturbance edges of \mathcal{A} .

In Ω' , all vertices in V inherit their colors from Ω and auxiliary vertices are colored by zero, which makes them irrelevant, while in Ω'' , all vertices in V inherit their colors from Ω , all vertices in D are assigned an even color that is larger than all colors in Ω 's range, and all other auxiliary vertices are colored by zero.

Hence, the games characterizing the existence of $(\omega + 1)$ -resilient and ω -resilient strategies are pushdown (one-counter) parity games that can be efficiently constructed. Finally, checking whether Player 0 wins a pushdown parity game from the initial vertex is EXPTIME-complete [41] while checking whether Player 0 wins a one-counter parity game from the initial vertex is PSPACE-complete [23, 34]. Furthermore, the first algorithm directly yields winning strategies for the rigged games, which can easily be turned into $(\omega + 1)$ -resilient or ω -resilient strategies for the original game.

The lower bounds hold already for determining the winner of a disturbance-free pushdown (one-counter) safety game, which is hard for EXPTIME [41] (PSPACE [23]³). ◀

B.3.2 Proof of Lemma 7

► **Lemma 7.** *The following problem is in 2EXPTIME (in EXPSPACE if the input is one-counter): “Given a pushdown parity game \mathcal{G} with initial vertex v_I and $k \in \omega$ (encoded in binary), does Player 0 have a k -resilient strategy for \mathcal{G} from v_I ?”. If yes, such a strategy can be computed in doubly-exponential time.*

³ The result cited pertains to emptiness of alternating word automata over a singleton alphabet. However it is easy to see that this problem can be reduced to solving one-counter safety games.

Proof. Assume the input $\mathcal{G} = (\mathcal{A}, \text{Parity}(\Omega))$ is induced by a PDS \mathcal{P} with set Q of states, a partition $\{Q_0, Q_1\}$ of Q , and a coloring Ω of Q . Then, we construct a PDS \mathcal{P}' with set Q' of states and a partition $\{Q'_0, Q'_1\}$ of Q' inducing \mathcal{A}_{rig} as for the proof of Lemma 5. Now, we turn \mathcal{P}' into a PDS \mathcal{P}'_k with set $Q' \times \{0, \dots, k\}$ of states which uses the additional component to keep track of the number of simulated disturbances, up to k . Further, we use the partition

$$\{Q'_0 \times \{0, \dots, k\}, Q'_1 \times \{0, \dots, k\}\}$$

and define the coloring Ω' such that $\Omega'(q, k') = \Omega(q)$ for $k' < k$ and $\Omega'(q, k) = 1$.

The resulting pushdown game is equivalent to $(\mathcal{A}_{\text{rig}}, \text{Win}_{\text{rig}} \cup R_{\geq k})$ and the winner from the initial vertex $((q_I, 0), \perp)$ can be determined in exponential time in k and the size of \mathcal{P} [41], i.e., in doubly-exponential time in the size of the input, as k is encoded in binary. Due to Lemma 5.3, Player 0 wins from the initial vertex if and only if she has a k -resilient strategy from v_I in \mathcal{G} , i.e., if and only if $r_{\mathcal{G}}(v_I) \geq k$. Furthermore, the algorithm computes winning strategies for Player 0 in doubly-exponential time, if they exist at all. These can easily be turned into k -resilient strategies for the original game.

If the input is one-counter, then the resulting pushdown game is one-counter as well and the winner from the initial vertex can be determined in polynomial space in k and the size of \mathcal{P} [34], i.e., in exponential space in the input. ◀

B.3.3 Proof of Lemma 8

► **Lemma 8.** *Let \mathcal{G} be a pushdown safety game with initial vertex v_I . If $r_{\mathcal{G}}(v_I) \neq \omega + 1$, then $r_{\mathcal{G}}(v_I) < b(\mathcal{P})$, where \mathcal{P} is the PDS underlying \mathcal{G} .*

To prove this result, we apply a result about winning strategies for Player 1 in pushdown safety games (Player 1 has a reachability condition in a safety game: he wins if F is visited at least once). Fix a disturbance-free pushdown safety game $\mathcal{G} = (\mathcal{A}, \text{Safety}(F))$ with initial vertex v_I . We say that a winning strategy τ for Player 1 from v_I *bounds the stack height to* $n \in \omega$ if every play $v_0 v_1 v_2 \dots$ that starts in v_I and is consistent with τ satisfies the following condition for all $j \in \omega$: either there is some $j' \leq j$ with $v_{j'} \in F$ or $\text{sh}(v_j) \leq n$. Thus, such a strategy ensures a visit to F when starting in the initial vertex, and ensures that the stack height n is never exceeded before F is visited for the first time. The next proposition shows that such a strategy always exists for $n = h(\mathcal{P})$, if Player 1 wins from v_I at all.

► **Lemma 24.** *If $v_I \in \mathcal{W}_1(\mathcal{G})$, then Player 1 has a winning strategy τ that bounds the stack height to $h(\mathcal{P})$, where \mathcal{P} is the PDS underlying \mathcal{G} .*

Proof. We transform \mathcal{G} into a parity game as described at the end of Section 2 on Page 19. This transformation can be implemented on the PDS inducing \mathcal{G} without increasing the number of states or the number of stack symbols. Furthermore, the parity condition only uses two colors, say 0 for states outside of F and 1 for states in F , which are sinks. Now, the desired result follows from a result on the existence of strategies in pushdown games that bound the occurrence of undesirable colors (here, the color 0, which is undesirable for Player 1) [17]. Slightly more formally, in the resulting parity game, the *stair score* for the color 0 after a play prefix (see [17] for definitions) is equal to the stack height of the prefix. Now, the main result in the work cited above shows that Player 1 has a strategy that bounds the stair score for 0 by $h(\mathcal{P})$, if he wins at all. Thus, this strategy bounds the stack height to $h(\mathcal{P})$. ◀

Now, we are able to prove the upper bound $b(\mathcal{P})$ on the resilience of the initial vertex of a pushdown safety game induced by \mathcal{P} in case this value is finite.

Proof of Lemma 8. Let $r_{\mathcal{G}}(v_I) \neq \omega + 1$. As pushdown arenas are finitely branching, Lemma 3 yields $r_{\mathcal{G}}(v_I) \in \omega$, say $r_{\mathcal{G}}(v_I) = k$. By definition, Player 0 has a k -resilient strategy for \mathcal{G} from v_I , but no $(k + 1)$ -resilient strategy. Hence, due to Lemma 5.3, Player 1 wins the game

$$(\mathcal{A}_{\text{rig}}, \text{Safety}(F)_{\text{rig}} \cup R_{\geq k+1})$$

from v_I . Thus, he also wins the safety game

$$(\mathcal{A}_{\text{rig}}, \text{Safety}(F)_{\text{rig}})$$

from v_I , as every winning strategy for Player 1 for the former game is also one for the latter. Hence, applying Lemma 24 yields the existence of a winning strategy τ for the latter game from v_I that bounds the stack height by $h(\mathcal{P})$. Note that we can assume τ to be positional (see Lemma 13 on Page 11 for a stronger statement and note that the construction presented in its proof preserves bounds on the stack height).

Now, every play that starts in v_I and is consistent with τ visits each vertex with stack height at most $h(\mathcal{P})$ at most once before reaching F . There are at most $b(\mathcal{P})$ such vertices, i.e., after at most $b(\mathcal{P}) - 1$ moves, F is reached.

Now, we show that Player 0 has no $b(\mathcal{P})$ -resilient strategy from v_I in \mathcal{G} . To this end, we show for that every strategy σ for her, there is a play ρ that starts in v_I , is consistent with σ , has at most $b(\mathcal{P}) - 1$ many disturbances, and visits a vertex in F , i.e., it is losing for Player 0.

Let σ' be the strategy for Player 0 in \mathcal{A}_{rig} obtained by transforming σ as described in the proof of Lemma 5. Now, let ρ' be the unique play of \mathcal{A}_{rig} starting in v_I that is consistent with σ' and τ , which visits F after at most $b(\mathcal{P}) - 1$ many moves. Hence, there are at most $b(\mathcal{P}) - 1$ many simulated disturbances in ρ' before the first visit to F . Now, $t(\rho')$ starts in v , is consistent with σ , and there are at most $b(\mathcal{P}) - 1$ many disturbances in $t(\rho')$ before the first visit to F (which occurs). Now, we just replace the suffix of $t(\rho')$ after the first visit to F by some disturbance-free suffix so that the resulting play ρ is still consistent with σ . We obtain a play ρ starting in v_I , consistent with σ , with at most $b(\mathcal{P}) - 1$ many disturbances that is losing for Player 0. Hence, σ is indeed not $b(\mathcal{P})$ -resilient. As we have picked σ arbitrarily, there is no $b(\mathcal{P})$ -resilient strategy from v_I and therefore $r_{\mathcal{G}}(v_I) < b(\mathcal{P})$. ◀

B.3.4 Proof of Theorem 9


► **Theorem 9.** *The following problem can be solved in triply-exponential time: “Given a pushdown safety game \mathcal{G} with initial vertex v_I , determine $r_{\mathcal{G}}(v_I)$ ”. If yes, an $r_{\mathcal{G}}(v_I)$ -resilient strategy can be computed in triply-exponential time.*

Proof. Algorithm 1 is correct due to Lemma 8. The triply-exponential running time stems from the doubly-exponential bound $b(\mathcal{P})$ presented in Lemma 8, which has to be plugged into Lemma 7 to implement the check in Line 4. The check in Line 1 runs in exponential time (Lemma 6) and the for-loop terminates after at most doubly-exponentially many iterations. ◀

B.3.5 Proof of Lemma 10

► **Lemma 10.** *Let $k \in \omega$.*

1. *There is a one-counter safety game \mathcal{G}_k with initial state v_I such that $r_{\mathcal{G}}(v_I) = p_k\#$ and the underlying OCS has polynomially many states in k .*



■ **Figure 4** The one-counter safety game \mathcal{G}_2 for the proof of Lemma 10.1. Round vertices are in V_0 , square ones in V_1 , and disturbance edges are dashed. Player 0 wins if and only if (s, \perp) is never visited. Vertices in the upper gray rectangle implement a modulo-2 counter while vertices in the lower rectangle implement a modulo-3 counter.

2. There is a pushdown safety game \mathcal{G}'_k with initial state v_I such that $r_G(v_I) = 2^{p_k\#} - 1$ and the underlying PDS has polynomially many states in k and two stack symbols.

Proof. 1.) We show the game \mathcal{G}_2 in Figure 4 and later explain the general case.

The winning condition is defined such that Player 0 wins a play if and only if the state s is never reached. Now, a play starting in the initial vertex of \mathcal{G}_2 proceeds as follows: Player 1 either stays in the state i ad infinitum, and thereby allows Player 0 to win, or he eventually moves to some vertex of the form $(c, A^n \perp)$. Now, Player 0 has three choices, moving to $((2, 0), A^n \perp)$, $((3, 0), A^n \perp)$, or $(d, A^n \perp)$. In the first case, there is only one continuation of the play prefix, which results in a disturbance-free play that is winning for Player 0 if and only if $n \bmod 2 \neq 0$. Similarly, in the second case, there is only one continuation of the play prefix, which results in a disturbance-free play that is winning for Player 0 if and only if $n \bmod 3 \neq 0$. Finally, moving to $(d, A^n \perp)$ means that Player 0 wins if strictly less than n disturbances occur in the continuation of the play prefix, but loses if n disturbances occur.

We claim that the initial vertex has resilience $6 = p_2\#$. A 6-resilient strategy for Player 0 moves from $(c, A^n \perp)$ to $(d, A^n \perp)$ if n is a multiple of 6. Else, it moves to $((p_j, 0), A^n \perp)$ for some $p_j \in \{2, 3\}$ such that $n \bmod p_j \neq 6$, which always exists. Applying the reasoning above implies that every play starting in the initial vertex, consistent with the strategy, and with at most 5 disturbances is winning for Player 0. Thus, the strategy is indeed 6-resilient.

Now, consider an arbitrary strategy σ for Player 0. We show that it is not 7-resilient, which is yields the desired result. To this end, consider the unique play prefix leading to $(c, A^6 \perp)$, which is consistent with σ . If σ prescribes a move to some $((p_j, 0), A^6 \perp)$, then, as argued above, there is disturbance-free play that is consistent with σ , but losing for her. The only other choice for σ is to move to $(d, A^6 \perp)$. Then, as argued above as well, there is a play that is consistent with σ with 6 disturbances that is losing for her. In both cases, we have shown that the strategy is indeed not 7-resilient.

The general case is obtained by having modulo counters in \mathcal{G}_k for the first k prime numbers instead of only the first two as in \mathcal{G}_2 . Using the same reasoning as above for arbitrary k instead of $k = 2$ shows that the initial vertex of \mathcal{G}_k has resilience $p_k\#$.

Finally, the number of states of the one-counter system inducing \mathcal{G}_k is bounded by $\mathcal{O}(k^3)$.

2.) We modify the one-counter safety game \mathcal{G}_k to obtain a pushdown safety game \mathcal{G}'_k . We use the stack alphabet $\{0, 1\}$, which allows us to interpret stack contents as binary encodings of natural numbers, with the least significant bit at the top of the stack. In the following, we give an informal account of the structure of \mathcal{G}'_k and leave the implementation by a pushdown system to the reader. Here, we reuse the modulo-counters of \mathcal{G}_k which forces that Player 1 to reach a stack height that is a multiple of $p_k\#$, as he would lose otherwise.

Thus, Player 1 is initially forced to push a multiple of $p_k\#$ 1's on the stack and then gives control to Player 0. As the stack height is a multiple of $p_k\#$, she can only go to a state d where all 0's are popped from the stack until the first 1 is uncovered (note that initially there is no 0 to pop). If there is no such 1, i.e., if the bottom of the stack is reached by removing 0's, then the play reaches a losing sink for Player 0. Otherwise, if a 1 is uncovered, then Player 0 only has a self-loop that leaves the stack unchanged, but there is also a disturbance edge that removes the topmost 1 by a 0 and hands back control to Player 1. He can now push as many 1's as necessary to again reach a stack height that is a multiple of $p_k\#$.

Now, if Player 1 never exceeds the stack height $p_k\#$, the stack always contains $p_k\#$ bits when Player 0 gains control. Assume now that Player 0 uses a strategy which moves to d in that situation and uses the correct modulo counter to win in all other situations (as described in more detail above for \mathcal{G}_k). Then, the stack contents reached at the positions where Player 0 gains control implement a binary counter with $p_k\#$ bits that is decremented each time Player 0 gains control, starting with the value $1^{p_k\#}$. Hence, as each decrement requires exactly one disturbance (and there are no others), the strategy described above is $(2^{p_k\#} - 1)$ -resilient from the initial vertex.

On the other hand, $2^{p_k\#} - 1$ disturbances suffice to reach a stack containing only 0's at some configuration where Player 0 gains control. Then, the unique continuation of that play is losing for her.

The only other choice for Player 0 is to enter a modulo counter at an “unsuitable” configuration, which also leads to a losing play with less than $2^{p_k\#}$ disturbances. Hence, Player 0 has no $2^{p_k\#}$ -resilient strategy from the initial vertex, i.e., it has indeed resilience $2^{p_k\#} - 1$.

Finally, the number of states of the one-counter system inducing \mathcal{G}_k is bounded by $\mathcal{O}(k^3)$. ◀

B.4 Proofs Omitted in Section 6

B.5 Proof of 12

► **Lemma 12.** *Player 1 wins \mathcal{G}_k from v_I if and only if there exists a strategy graph for \mathcal{G}_k .*

Recall that we split the proof into several steps, which we have to prove first.

► **Lemma 13.** *If Player 1 wins \mathcal{G}_k from v_I then he has a positional winning strategy for \mathcal{G}_k from v_I .*

Proof. Assume a winning strategy τ for Player 1 from v_I . Let us call a play prefix $v_0 \cdots v_j$ *unsettled* if it starts in v_I , is consistent with τ , and no strict prefix contains a vertex in the target F . Notice that there must be a uniform bound $\ell \in \omega$ such that $|w| < \ell$ for every unsettled w . Indeed, if there was no such bound, then it is possible to arrange an infinite set of arbitrarily long play prefixes not visiting F into an infinite finitely branching tree. By

König's Lemma, this tree has an infinite path which corresponds to an infinite play starting in v_I , consistent with τ , but not containing a vertex in F , which contradicts the assumption that τ is winning.

Given an unsettled prefix w , let $\text{val}(w) = d \cdot \ell + |w|$ where d is the number of simulated disturbances during w . Let $U(v)$ for $v \in V'$ denote the set of unsettled play prefixes ending in v . Further, for every $v \in V$ with non-empty $U(v)$ let $\text{rep}(v)$ be an element from $U(v)$ such that $\text{val}(\text{rep}(v)) \geq \text{val}(w)$ for all $w \in U(v)$. Such an element exists, as the $\text{val}(w)$ for $w \in U(v)$ are bounded by $k \cdot \ell - 1$: Each unsettled prefix is consistent with the winning strategy τ , which implies that it simulates at most $k - 1$ disturbance edges, and its length is bounded by $\ell - 1$, as argued above.

Based on this we define the positional strategy τ' via $\tau'(v) = \tau(\text{rep}(v))$ if $\text{rep}(v)$ is defined and $\tau'(v) = v'$ for some arbitrary successor v' of v if $\text{rep}(v)$ is undefined (note that it suffices to define $\tau'(v)$ for $v \in V_1$ to define a positional strategy for Player 1). We claim that τ' is winning from v_I . To this end, let $\rho = v_0 v_1 v_2 \dots$ start in v_I and be consistent with τ' . We need to show that it visits a vertex in F and that it simulates at most $k - 1$ disturbance edges.

A simple induction shows that every length- j prefix $v_0 \dots v_{j-1}$ that does not visit F must satisfy that

$$\text{val}(v_0 \dots v_j) \leq \text{val}(\text{rep}(v_j)) \tag{1}$$

The induction start $j = 0$ is trivial, as we have $v_0 = v_I$ and $v_I \in U(v_I)$, which implies $\text{val}(v_I) \leq \text{val}(\text{rep}(v_I))$ as required.

For the induction step, consider some $j > 0$ such that $v_0 \dots v_{j-1}$ does not visit F . The induction hypothesis yields $\text{val}(v_0 \dots v_{j-1}) \leq \text{val}(\text{rep}(v_{j-1}))$. Let $\text{rep}(v_{j-1}) = wv_{j-1}$, which is consistent with τ . If $v_{j-1} \in V'_0$, then $wv_{j-1}v_j$ is consistent with τ as well, as it is Player 0's turn at v_{j-1} . Similarly, if $v_{j-1} \in V'_1$, then we have

$$v_j = \tau'(v_0 \dots v_{j-1}) = \tau(\text{rep}(v_{j-1})) = \tau(wv_{j-1}).$$

Hence, $wv_{j-1}v_j$ is again consistent with τ . Furthermore, wv_{j-1} does not contain a vertex in F , as v_{j-1} is not in F (recall that vertices in F are sinks). Thus, we conclude that $wv_{j-1}v_j$ is unsettled, which implies $\text{val}(wv_{j-1}v_j) \leq \text{val}(\text{rep}(v_j))$, by our definition of $\text{rep}(v_j)$. To finish the induction step, let $x = 1$ if $v_j \in D$, i.e. a disturbance edge is simulated, and $x = 0$ otherwise. Then, we have

$$\begin{aligned} \text{val}(v_0 \dots v_j) &= \text{val}(v_0 \dots v_{j-1}) + x \cdot \ell + 1 \\ &\leq \text{val}(\text{rep}(v_{j-1})) + x \cdot \ell + 1 \\ &= \text{val}(wv_{j-1}) + x \cdot \ell + 1 \\ &= \text{val}(wv_{j-1}v_j) \leq \text{val}(\text{rep}(v_j)). \end{aligned}$$

Applying Equation 1, we can show that ρ is indeed winning. First, towards a contradiction, assume ρ does not visit a vertex in F . Then, Equation 1 is applicable to every prefix $v_0 \dots v_j$ and we thus obtain for every $j > 0$, that

$$j + 1 = |v_0 \dots v_j| \leq \text{val}(v_0 \dots v_j) \leq \text{val}(\text{rep}(v_j)) \leq k \cdot \ell - 1$$

which is a contradiction as the term on the right is constant.

Second, again towards a contradiction, assume that ρ simulates at least k disturbance edges. Then, let j be minimal such that the prefix $v_0 \dots v_j$ simulates exactly k disturbance

edges. As vertices in F are sinks, and therefore have no outgoing edges simulating disturbance edges, Equation 1 is applicable to $v_0 \cdots v_j$ and we obtain that

$$k \cdot \ell \leq \text{val}(v_0 \cdots v_j) \leq \text{val}(\text{rep}(v_j))$$

which is impossible as $\text{val}(\text{rep}(v_j)) \leq k \cdot \ell - 1$. Consequently, ρ visits F and simulates at most $k - 1$ disturbance edges. As ρ was an arbitrary play consistent with τ' , this strategy is indeed winning. \blacktriangleleft

► **Lemma 14.** *If Player 1 wins \mathcal{G}_k from v_I then he has a positional winning strategy from v_I with $\text{maxSh}(\tau) \leq (2k)^{|\mathcal{Q}|^2}$.*

Proof. By Lemma 13 we can pick a positional strategy τ for Player 1 that is winning \mathcal{G}_k from v_I . We show how to turn this into a winning strategy that satisfies the claim.

Notice first that $\text{maxSh}(\tau)$ must be finite. Indeed, if it is unbounded, then for every $n \in \omega$ there is a play prefix w_n starting in v_I , consistent with τ , and ending in a vertex of stack height n . As the stack height is increased by at most one during each move, we have $|w_n| \geq n$. Furthermore, as vertices in F are sinks, these play prefixes can be assumed to not contain a vertex in F . The prefixes w_n can be arranged in an infinite finitely branching tree. By König's Lemma, this tree has an infinite path, which corresponds to an infinite play starting in v_I , consistent with τ , but not visiting a vertex in F . This contradicts τ being a winning strategy.

It suffices to show that if $\text{maxSh}(\tau) > (2k)^{|\mathcal{Q}|^2}$, then τ can be turned into a positional winning strategy τ' from v_I with strictly smaller maximal stack height. For the sake of readability, we will identify a stack content $A^n \perp$ of the one-counter system underlying \mathcal{G}_k by the number $n \in \omega$. Hence, vertices of \mathcal{G}_k are from now on denoted by (q, n) with $n \in \omega$.

Let R denote the set of vertices reachable from v_I via play prefixes that are consistent with τ . For $(q, n) \in R$ with $n > 0$ let $H(q, n)$ be the set of vertices of the form $(q', n - 1)$ reachable from (q, n) via a play prefix $(q, n)(q_1, n_1) \cdots (q_j, n_j)(q', n - 1)$ that is consistent with τ and such that $n_{j'} \geq n$ for every $j' \in \{1, \dots, j\}$, i.e., the last vertex of the play prefix is the first time the stack height along the play prefix is strictly smaller than n . We call such a play prefix a *hill* from (q, n) to $(q', n - 1)$.

For all $n > 0$ define the partial function $h_n: \mathcal{Q} \rightarrow 2^{\mathcal{Q}}$ that maps q to $H(q, n)$ whenever $(q, n) \in R$, and else leaves $h_n(q)$ undefined. Similarly, define the partial function $d_n: \mathcal{Q} \times \mathcal{Q} \rightarrow \{0, \dots, k - 1\}$ by mapping each pair (q, q') with $q' \in H(q, n)$ to the maximal number of disturbances simulated during any hill from (q, n) to $(q', n - 1)$. This value is bounded by $k - 1$, as each hill is part of a play that is consistent with τ . For (q, q') with $q' \notin H(q, n)$, we leave $d_n(q, q')$ undefined.

There are at most $(2k)^{|\mathcal{Q}|^2}$ many different pairs of such functions h_n and d_n . Hence, if R contains a vertex (q, n) with $n > (2k)^{|\mathcal{Q}|^2}$, then there are $0 < n_\ell < n_u$ such that $h_{n_\ell} = h_{n_u}$ and $d_{n_\ell} = d_{n_u}$. Let $s = n_u - n_\ell$. We define the positional strategy τ' via $\tau'(q, n) = \tau(q, n)$, if $n < n_\ell$ and $\tau'(q, n) = \tau(q, n + s)$ if $n \geq n_\ell$ (recall that it suffices to define $\tau'(v)$ for every $v \in V'_1$ to define a positional strategy τ').

We claim that τ' is still winning for Player 1 from v_I in \mathcal{G}_k . To this end, consider an arbitrary play

$$\rho' = (q_0, n_0)(q_1, n_1)(q_2, n_2) \cdots$$

that starts in v_I and is consistent with τ' . We need to show that it visits F and simulates at most $k - 1$ disturbance edges.

If every n_j is strictly smaller than n_ℓ , then ρ' is also consistent with τ , as only the first case of the definition of τ is applied. Hence, it is winning for Player 1, as τ is a winning strategy from v_I .

It remains to consider the case where ρ' reaches stack height n_ℓ . Here, we turn ρ' into a play ρ starting in v_I and consistent with τ , which implies that ρ visits F and simulates at most $k - 1$ disturbance edges. Using the relation between ρ and ρ' , we argue that the ρ' is also winning.

The following remark is useful throughout our argument and follows immediately from the fact that at stack heights $n \geq n_\ell$, τ' mimics the behavior of τ at stack height $n + s$.

► **Remark 25.** Let j and j' be positions of ρ' such that $n_j = n_\ell$ and $n_{j'} \geq n_\ell$ for every $j'' \in \{j + 1, \dots, j'\}$, i.e., the infix between positions j and j' starts at stack height n_ℓ and never reaches a smaller stack height. Then, $(q_j, n_j + s) \cdots (q_{j'+1}, n_{j'+1} + s)$ is consistent with τ (note the $+1!$).

We inductively construct ρ by defining a sequence $(w_m)_{m \in \omega}$ of strictly increasing prefixes whose limit is ρ . To define this sequence, we simultaneously construct a sequence $(j_m)_{m \in \omega}$ of strictly increasing positions of ρ' . During the construction, we satisfy the following invariant: Each w_m is consistent with τ , ends in (q_{j_m}, n_{j_m}) where n_{j_m} is strictly smaller than n_ℓ , and w_m simulates at least as many disturbances as $(q_0, n_0) \cdots (q_{j_m}, n_{j_m})$.

We start with $j_0 = 0$ and $w_0 = (q_0, n_0) = v_I$, which satisfies the invariant due to our choice of n_ℓ being greater than zero. We define w_m and j_m for $m > 0$, based on w_{m-1} and j_{m-1} , as follows. Due to the invariant, w_{m-1} ends in $(q_{j_{m-1}}, n_{j_{m-1}})$ with $n_{j_{m-1}} < n_\ell$ and is consistent with τ .

We consider two cases. In the first, if $(q_{j_{m-1}+1}, n_{j_{m-1}+1})$, the next vertex after $(q_{j_{m-1}}, n_{j_{m-1}})$ in ρ' , satisfies $n_{j_{m-1}+1} < n_\ell$, then define $w_m = w_{m-1}(q_{j_{m-1}+1}, n_{j_{m-1}+1})$ and $j_m = j_{m-1} + 1$. Note that the move from $(q_{j_{m-1}}, n_{j_{m-1}})$ to $(q_{j_{m-1}+1}, n_{j_{m-1}+1})$ is consistent with τ , as it is either Player 0's turn or the first case of the definition of τ' is applied (which mimics τ) due to our invariant. Hence, w_m is again consistent with τ . Similarly, the requirement on the number of simulated disturbances is satisfied as the same edge is used to extend both play prefixes.

In the second case, we have $n_{j_{m-1}+1} \geq n_\ell$, which implies $n_{j_{m-1}+1} = n_\ell$, as the stack height can increase by at most one during every transition.

We claim there is some $j > j_{m-1} + 1$ such that $n_j = n - 1$. Towards a contradiction, assume there is no such j . Then, Remark 25 is applicable to every pair $(j_{m-1} + 1, j)$ with $j > j_{m-1} + 1$. This yields an infinite play ρ_c equal to

$$(q_{j_{m-1}+1}, n_{j_{m-1}+1} + s)(q_{j_{m-1}+2}, n_{j_{m-1}+2} + s)(q_{j_{m-1}+3}, n_{j_{m-1}+3} + s) \cdots$$

that is consistent with τ . The play prefix w_{m-1} starts in v_I , is consistent with τ , and ends in $(q_{j_{m-1}}, n_{j_{m-1}})$. Further, the move from $(q_{j_{m-1}}, n_{j_{m-1}})$ to $(q_{j_{m-1}+1}, n_{j_{m-1}+1})$ in ρ' is consistent with τ' and therefore also with τ , as $n_{j_{m-1}} < n_\ell$ by our invariant. Thus, we have shown $(q_{j_{m-1}+1}, n_{j_{m-1}+1}) \in R$, i.e., there is a play prefix w_c starting in v_I , consistent with τ , and ending in $(q_{j_{m-1}+1}, n_{j_{m-1}+1})$. Altogether, we can combine w_c and ρ_c into an infinite play starting in v_I and consistent with τ that has ρ_c as suffix. Now, ρ_c contains by construction no vertex of stack height zero. As vertices in F are sinks of stack height zero, the combined play can not visit F . This contradicts the assumption that τ is winning from v_I .

Thus, let $j > j_{m-1} + 1$ be minimal such that $n_j = n - 1$. Applying Remark 25 for $j_{m-1} + 1$ and $j - 1$ shows that

$$w = (q_{j_{m-1}+1}, n_{j_{m-1}+1} + s) \cdots (q_j, n_j + s)$$

is a hill from $(q_{j_{m-1}+1}, n_{j_{m-1}+1} + s) = (q_{j_{m-1}+1}, n_u)$ to $(q_j, n_j + s) = (q_j, n_u - 1)$. Hence, by the choice of n_ℓ and n_u there is also a hill w' from $(q_{j_{m-1}+1}, n_\ell)$ to $(q_j, n_\ell - 1)$ that has at least as many simulated disturbances as w .

We obtain w_m from w_{m-1} by appending w' and define $j_m = j$. The requirement on the stack height n_{j_m} is satisfied by our choice of $j_m = j$ while w_m is consistent with τ , as w_{m-1} , the move from $(q_{j_{m-1}}, n_{j_{m-1}})$ (the last vertex of w_{m-1}) to $(q_{j_{m-1}+1}, n_{j_{m-1}+1})$ (the first vertex of w'), and w' are all consistent with τ . The requirement on the number of simulated disturbances is satisfied, as $(q_{j_{m-1}+1}, n_{j_{m-1}+1}) \cdots (q_j, n_j)$ simulates the same number of disturbances as w , which is at most the number of disturbances simulated by w' .

Consider the resulting play ρ , which is by construction winning for Player 1 and consequently simulates at most $k - 1$ disturbances. An inductive application of the invariant above shows that ρ' therefore also simulates at most $k - 1$ disturbances. Furthermore, ρ visits a vertex in F , which has stack height zero. When such a vertex is added during the inductive construction described above, then only in the first case (when $n_{j_{m-1}+1} < n_\ell$) and only because the same vertex appears in ρ' , i.e., ρ' visits F as well. Hence, ρ' is indeed winning for Player 1.

To conclude, we have to show $\max\text{Sh}(\tau') < \max\text{Sh}(\tau)$. An induction on n shows that if (q, n) is reachable from v_I by a play prefix that is consistent with τ' , then:

- If $n \leq n_\ell$, then (q, n) is reachable from v_I by a play prefix that is consistent with τ .
- If $n > n_\ell$, then $(q, n + s)$ is reachable from v_I by a play prefix that is consistent with τ . This implies $\max\text{Sh}(\tau') + s \leq \max\text{Sh}(\tau)$, which yields the desired bound due to $s > 0$. ◀

Having proved the existence of positional winning strategies with exponential maximal stack height, it is straightforward to show that these are essentially strategy graphs, which proves Lemma 12.

Proof of Lemma 12. Let Player 1 win \mathcal{G}_k from v_I . Then, Lemma 14 yields a positional winning strategy τ for him from v_I with $\max\text{Sh}(\tau) \leq (2k)^{|Q|^2}$. We turn τ into a strategy graph for \mathcal{G}_k by defining

- V° to be the set of vertices visited by plays starting in v_I that are consistent with τ ,
- E° to be the set of edges traversed by these plays (ignoring the self-loops at vertices in F),
- $\mu_r^\circ(v)$ to be the maximal number of disturbance edges simulated on plays starting in v that are consistent with τ , and
- $\mu_d^\circ(v)$ to be the maximal length of a play prefix starting in v , being consistent with τ , and the last vertex (but no other) being in F .

It is straightforward to prove that $(V^\circ, E^\circ, \mu_r^\circ, \mu_d^\circ)$ satisfies all properties required of a strategy graph for \mathcal{G}_k .

Conversely, assume there is a strategy graph $(V^\circ, E^\circ, \mu_r^\circ, \mu_d^\circ)$ for \mathcal{G}_k . We turn it into a positional winning strategy τ for Player 1 from v_I . Let $v \in V_1'$. If $v \in V^\circ \setminus F$, then there is a unique outgoing edge $(v, v') \in E^\circ \cap E'$ due to Property 3 of the strategy graph definition. Then, we define $\tau(v) = v'$. Otherwise, i.e., if $v \notin V^\circ \setminus F$, then define $\tau(v)$ to be an arbitrary successor of v in \mathcal{A}_{rig} . We claim that τ is indeed a winning strategy for Player 1 for \mathcal{G}_k from v_I .

To this end, let $\rho = (v_0, 0)(v_1, 0)(v_2, 0) \cdots$ be a play starting in v_I that is consistent with τ . We need to show that ρ is winning for Player 1, i.e., that it visits F and contains at most $k - 1$ simulated disturbance edges.

An induction applying the definition of τ and Property 2 of the strategy graph definition shows that if $v_0 \cdots v_{j-1}$ does not contain a vertex from F , then $v_0 \cdots v_j$ is a path through the graph (V°, E°) . Hence, we have $\mu_d^\circ(v_0) > \mu_d^\circ(v_1) > \cdots > \mu_d^\circ(v_{j-1}) > \mu_d^\circ(v_j)$ by Property 5. As the range of μ_d° is finite, this yields an upper bound on the length of prefixes of ρ that do not visit F , which implies that ρ contains a vertex of F . Hence, let j be the minimal position of ρ with $v_j \in F$. As vertices in F are sinks, no disturbance edges are simulated in ρ after position j . Due to Property 4, we have $\mu_r^\circ(v_0) \geq \mu_d^\circ(v_0) \geq \cdots \geq \mu_r^\circ(v_{j-1}) \geq \mu_r^\circ(v_j)$ with strict inequality whenever a disturbance edge is simulated, as $v_0 \cdots v_j$ is a path through (V°, E°) as argued above. Hence, as the range of μ_r° has at most k elements, there are at most $k - 1$ simulated disturbances in ρ before position j and none afterwards, as argued above. Altogether, ρ visits F and contains at most $k - 1$ simulated disturbance edges, i.e., it is indeed winning for Player 1 in \mathcal{G}_k . ◀

B.5.1 Proof of Lemma 15

► **Lemma 15.** *The following problem is in PSPACE: “Given a one-counter safety game \mathcal{G} induced by a PDS \mathcal{P} and $k \leq b(\mathcal{P})$ (encoded in binary), is there a strategy graph for \mathcal{G}_k ?”*

Proof. Notice that all defining conditions of strategy graphs are local, and can be verified for a vertex $v = (q, n)$ if the values of $\mu_r^\circ(v')$ and $\mu_d^\circ(v')$ are known for all direct neighbors, which have the form (q', n') with $n' \in \{n - 1, n, n + 1\}$. A strategy graph can therefore be guessed and verified on the fly, keeping in memory these values for vertices in $Q \times \{n, n + 1, n + 2\}$ while incrementing n from 0 to $(2k)^{|Q|^2}$. This requires polynomial space, both for the labeling of the vertices (as the numbers are at most exponential in the size of the input) and for the counter. ◀