

Synthesizing Optimally Resilient Controllers^{*}

Daniel Neider¹, Alexander Weinert² and Martin Zimmermann²

¹ Max Planck Institute for Software Systems, 67663 Kaiserslautern, Germany
`neider@mpi-sws.org`

² Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany
`{weinert,zimmermann}@react.uni-saarland.de`

Abstract. Recently, Dallal, Neider, and Tabuada studied a generalization of the classical game-theoretic model used in program synthesis, which additionally accounts for unmodeled intermittent disturbances. In this extended framework, one is interested in computing optimally resilient strategies, i.e., strategies that are resilient against as many disturbances as possible. Dallal, Neider, and Tabuada showed how to compute such strategies for safety specifications.

In this work, we compute optimally resilient strategies for a much wider range of winning conditions and show that they do not require more memory than winning strategies in the classical model. Our algorithms only have a polynomial overhead in comparison to the ones computing winning strategies. In particular, for parity conditions optimally resilient strategies are positional and can be computed in quasipolynomial time.

1 Introduction

Reactive synthesis is an exciting and promising approach to solving a crucial problem, whose importance is ever-increasing due to ubiquitous deployment of embedded systems: obtaining correct and verified controllers for safety-critical systems. Instead of an engineer programming a controller by hand and then verifying it against a formal specification, synthesis automatically constructs a correct-by-design controller from the given specification (or reports that no such controller exists).

Typically, reactive synthesis is modeled as a two-player zero-sum game on a finite graph that is played between the system, which seeks to satisfy the specification, and its environment, which seeks to violate it. Although this model is well understood, there are still multiple obstacles to overcome before synthesis can be realistically applied in practice. These obstacles include not only the high computational complexity of the problem, but also more fundamental ones. Among the most prohibitive issues in this regard is the need for a complete model of the interaction between the system and its environment, including an accurate model of the environment, the actions available to both antagonists, as well as the effects of these actions.

This modeling task often places an insurmountable burden on engineers as the environments in which real-life controllers are intended to operate tend to be highly complex or not fully known at design time. Moreover, when a controller is deployed in the real world, a common source of errors is a mismatch between the controller’s intended result of an action and the result that actually manifests. Such situations arise, for instance, in the presence of disturbances, when the effect of an action is not precisely known, or when the intended control action of the controller cannot be executed, e.g., when an actuator malfunctions. By a slight abuse of notation from control theory, we subsume all such errors under the generic term *disturbance*.

To obtain controllers that can handle disturbances, one has to yield control over their occurrence to the environment. However, due to the antagonistic setting of the two-player zero-sum game, this would allow the environment to violate the specification by causing disturbances at will. Overcoming this requires the engineer to develop a realistic disturbance model, which is a highly complex task, as such disturbances are assumed to be rare events. Also, incorporating such a model into the game leads to a severe blowup in the size of the game, which can lead to intractability due to the high computational complexity of synthesis.

To overcome these difficulties, Dallal, Neider, and Tabuada [9] proposed a conceptually simple, yet powerful extension of infinite games termed “games with unmodeled intermittent disturbances”. Such

^{*} Supported by the project “TriCS” (ZI 1516/1-1) of the German Research Foundation (DFG).




Fig. 1: A (max-) parity game with disturbances. Disturbance edges are drawn as dashed arrows. Vertices are labeled with both a name and a color. Vertices under control of Player 0 are drawn as circles, while vertices under control of Player 1 are drawn as rectangles.

games are played similarly to classical infinite games: two players, called Player 0 and Player 1, move a token through a finite graph, whose vertices are partitioned into vertices under the control of Player 0 and Player 1, respectively; the winner is declared based on the resulting play. In contrast to classical games, however, the graph is augmented with additional *disturbance edges* that originate in vertices of Player 0 and may lead to any other vertex. Moreover, the mechanics of how Player 0 moves is modified: whenever Player 0 moves the token, her move might be overridden, and the token instead moves along a disturbance edge. This change in outcome implicitly models the occurrence of a disturbance—the intended result of the controller and the actual result differ—but it is not considered to be antagonistic. Instead, the occurrence of a disturbance is treated as a rare event without any assumptions on frequency, distribution, etc.

This non-antagonistic nature of disturbances is different from existing approaches in the literature and causes many interesting phenomena that do not occur in the classical theory of infinite games. Some of these already manifest in the parity game shown in Figure 1, in which vertices are labeled with non-negative integers, so-called colors, and Player 0 wins if the highest color seen infinitely often is even. Consider, for instance, vertex v_2 . In the classical setting without disturbances, Player 0 wins every play reaching v_2 by simply looping in this vertex forever (since the highest color seen infinitely often is even). However, this is no longer true in the presence of disturbances: a disturbance in v_2 causes a play to proceed to vertex v_1 , from which Player 0 can no longer win. In vertex v_7 , Player 0 is in a similar, yet less severe situation: she wins every play with finitely many disturbances but loses if infinitely many disturbances occur. Finally, vertex v_9 falls into a third category: from this vertex, Player 0 wins every play even if infinitely many disturbances occur. In fact, disturbances partition the set of vertices from which Player 0 can guarantee to win into three disjoint regions (indicated as shaded boxes in Figure 1): (A) vertices from which she can win if at most a fixed finite number of disturbance occur, (B) vertices from which she can win if any finite number of disturbances occurs but not if infinitely many occur, and (C) vertices from which she can win even if infinitely many disturbances occur.

The observation above gives rise to a question that is both theoretically interesting and practically important: if Player 0 can tolerate different numbers of disturbances from different vertices, how should she play to be *resilient*³ to as many disturbances as possible, i.e., to tolerate as many disturbances as possible while still winning? Put slightly differently, disturbances induce an order on the space of winning strategies (“a winning strategy is better if it is more resilient”), and the natural problem is to compute optimally resilient winning strategies, yielding optimally resilient controllers. Note that this is in stark contrast to the classical theory of infinite games, where the space of winning strategies is unstructured.

Dallal, Neider, and Tabuada [9] have already solved the problem of computing optimally resilient winning strategies for safety games. Their approach exploits the existence of maximally permissive winning strategies in safety games, which allows Player 0 to avoid “harmful” disturbance edges during a play. In games with more expressive winning conditions, however, this is no longer possible, as witnessed by vertex v_4 in the example of Figure 1: although Player 0 can avoid a disturbance edge by looping in v_4 forever, she needs to move to v_2 eventually in order to see an even color (otherwise she loses), thereby

³ We have deliberately chosen the term *resilience* so as to avoid confusion with the already highly ambiguous notions of *robustness* and *fault tolerance*.

potentially risking to lose if a disturbance occurs. In fact, the problem of constructing optimally resilient winning strategies for games other than safety is still open.

In this paper, we solve this problem for an extensive class of infinite games, including parity games and Muller games. In particular, our contributions are as follows:

- We introduce a novel concept, termed *resilience*, which captures for each vertex how many disturbances need to occur for Player 0 to lose. This concept generalizes the notion of determinacy and allows us to derive optimally resilient winning strategies.
- We present an algorithm for computing the resilience of vertices and optimally resilient winning strategies. Our algorithm uses solvers for the underlying game without disturbances as a subroutine, which it invokes a linear number of times on various subgames. For many winning conditions, the time complexity of our algorithm thus falls into the same complexity class as solving the original game without disturbances. In particular, we obtain a quasipolynomial algorithm for parity games with disturbances, which matches the currently best known upper bound for classical parity games.
- In addition to natural assumptions on the winning condition, e.g., that games are determined and effectively solvable, our algorithm requires the winning condition to be prefix-independent. However, we show that the classical notion of game reduction carries over to the setting of games with disturbances. As a consequence, our algorithm can be applied to an extensive class of infinite games (using a reduction from prefix-dependent games to prefix-independent ones if necessary), including all ω -regular games.
- Finally, we discuss various further phenomena that arise in the presence of disturbances. Amongst others, we illustrate how the additional goal of avoiding disturbances whenever possible affects the memory requirements of strategies. Moreover, we raise the question of how benevolent disturbances can be leveraged to recover from losing a play. However, an in-depth investigation of these phenomena is outside the scope of this paper and left for future work.

This paper is structured as follows: after setting up definitions and notations in Section 2, we present our algorithm for computing optimally resilient strategies in Section 3. In Section 4, we discuss the necessary assumptions on the winning condition in detail and show that the notion of game reduction carries over to games with disturbances. In Section 5, we identify further interesting research questions arising in the context of disturbances. Finally, we discuss related work in Section 6.

2 Preliminaries

For notational convenience, we employ some ordinal notation à la von Neumann: the non-negative integers are defined inductively as $0 = \emptyset$ and $n+1 = n \cup \{n\}$. Now, the first limit ordinal is $\omega = \{0, 1, 2, \dots\}$, the set of the non-negative integers. The next two successor ordinals are $\omega+1 = \omega \cup \{\omega\}$ and $\omega+2 = \omega+1 \cup \{\omega+1\}$. These ordinals are ordered by set inclusion, i.e., we have $0 < 1 < 2 < \dots < \omega < \omega+1 < \omega+2$. For convenience of notation, we also denote the cardinality of ω by ω .

2.1 Infinite Games with Disturbances

An arena (with unmodeled disturbances) $\mathcal{A} = (V, V_0, V_1, E, D)$ consists of a finite directed graph (V, E) , a partition $\{V_0, V_1\}$ of V into the set of vertices V_0 of Player 0 (denoted by circles) and the set of vertices of Player 1 (denoted by squares), and a set $D \subseteq V_0 \times V$ of disturbance edges (denoted by dashed arrows). Note that only vertices of Player 0 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 play in \mathcal{A} is an infinite sequence $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots \in (V \times \{0, 1\})^\omega$ 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 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 $(v_0, b_0) \cdots (v_j, b_j)$ is defined similarly and ends in v_j . The number of disturbances in a play $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ 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 disturbances) $\mathcal{G} = (\mathcal{A}, \text{Win})$ consists of an arena $\mathcal{A} = (V, V_0, V_1, E, D)$ and a winning condition $\text{Win} \subseteq V^\omega$. A play $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ is winning for Player 0, if

$v_0v_1v_2\cdots \in \text{Win}$, otherwise it is winning for Player 1. Hence, winning is oblivious to occurrences of disturbances. A winning condition Win is prefix-independent if for all $\rho \in V^\omega$ and all $w \in V^*$ we have $\rho \in \text{Win}$ if and only if $w\rho \in \text{Win}$.

In examples, we often use the parity condition, the canonical ω -regular winning condition. Let $\Omega: V \rightarrow \omega$ be a coloring of a set V of vertices. The (max-) parity condition

$$\text{Parity}(\Omega) = \{v_0v_1v_2\cdots \in V^\omega \mid \limsup \Omega(v_0)\Omega(v_1)\Omega(v_2)\cdots \text{ is even}\}$$

requires the maximal color occurring infinitely often during a play to be even. A game $(\mathcal{A}, \text{Win})$ is a parity game, if $\text{Win} = \text{Parity}(\Omega)$ for some coloring Ω of the vertices of \mathcal{A} . In figures, we label vertices of a parity game by a pair v/c where v is the name of the vertex and c its color.

Furthermore, in our proofs we make use of the safety condition

$$\text{Safety}(U) = \{v_0v_1v_2\cdots \in V^\omega \mid v_j \notin U \text{ for every } j \in \omega\}$$

for a given set $U \subseteq V$ of unsafe vertices. It requires Player 0 to only visit safe vertices, i.e., Player 1 wins a play if it visits at least one unsafe vertex.

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., 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 $v_0\cdots v_j \in V^*V_i$.

Remark 1. Note that 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.

2.2 Infinite Games without Disturbances

We characterize the classical notion of infinite games, i.e., those without disturbances, (see, e.g., [15]) 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 $v \in V$ such that Player i has a winning strategy from v . Thus, the winning regions of \mathcal{G} are independent of the disturbance edges, i.e., we obtain the classical notion of infinite games. We say that Player i wins \mathcal{G} from v , if $v \in \mathcal{W}_i(\mathcal{G})$. Solving a game amounts to determining its winning regions. Note that every game has disjoint winning regions. In contrast, a game is determined, if every vertex is in either winning region.

2.3 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. Note that *every* strategy is 0-resilient, as no play has strictly less than zero disturbances. Furthermore, a strategy is 1-resilient from v if and only if it is winning for Player 0 from v .

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. Nevertheless, due to the remarks above, resilient strategies generalize winning strategies.

Remark 2. Let \mathcal{G} be a determined game. Then, $r_{\mathcal{G}}(v) > 0$ if and only if $v \in \mathcal{W}_0(\mathcal{G})$.

A strategy σ is optimally resilient, if it is $r_{\mathcal{G}}(v)$ -resilient from every vertex v . Every such strategy is a *uniform* winning strategy for Player 0, i.e., a strategy that is winning from every vertex in her winning region. Hence, positional optimally resilient strategies can only exist in games which have uniform positional winning strategies for Player 0.

Our goal is to determine the mapping $r_{\mathcal{G}}$ and to compute an optimally resilient strategy.

3 Computing Optimally Resilient Strategies

To compute optimally resilient strategies, we first characterize the vertices of finite resilience in Subsection 3.1. All other vertices either have resilience ω or $\omega + 1$. To distinguish between these possibilities, we show how to determine the vertices with resilience $\omega + 1$ in Subsection 3.2. In Subsection 3.3, we show how to compute optimally resilient strategies using the results of the first two subsections. We only consider prefix-independent winning conditions in Subsections 3.1 and 3.3. In Section 4, we show how to overcome this restriction.

3.1 Characterizing Vertices of Finite Resilience

Our goal in this subsection is to characterize vertices with finite resilience in a game with prefix-independent winning condition, i.e., those vertices from which Player 0 can win even under $k - 1$ disturbances, but not under k disturbances, for some $k \in \omega$. To illustrate our approach, consider the parity game in Figure 1 (on Page 2). The winning region of Player 1 only contains the vertex v_1 . Thus, by Remark 2, v_1 is the only vertex with resilience zero, every other vertex has a larger resilience.

Now, consider the vertex v_2 , which has a disturbance edge leading into the winning region of Player 1. Due to this edge, v_2 has resilience one. From v_1 , a disturbance-free play violating the winning condition starts that is consistent with every strategy for Player 0. Due to prefix-independence, prepending the disturbance edge does not change the winner and consistency with every strategy for her. Hence, this play witnesses that v_2 has resilience at most one, while v_2 being in Player 0's winning region yields the matching lower bound.

However, v_2 is the only vertex to which this reasoning applies. Now, consider v_3 : from here, Player 1 can force a play to visit v_2 using a standard edge. From this property, one can argue that v_3 has resilience one as well. Again, this is the only vertex to which this reasoning is applicable.

In particular, from v_4 Player 0 can avoid reaching the vertices for which we have determined the resilience by using the self loop. However, this comes at a steep price for her: doing so results in a losing play, as the color of v_4 is odd. Thus, if she wants to have a chance at winning, she has to take a risk by moving to v_2 , from which she has a 1-resilient strategy that is winning, if no more disturbances occur. For this reason, v_4 has resilience one as well. The same reasoning applies to v_6 : Player 1 can force the play to v_4 and from there Player 0 has to take a risk by moving to v_2 .

The vertices v_3 , v_4 , and v_6 share the property that Player 1 can either enforce a play violating the winning condition or to reach a vertex with already determined finite resilience. These three vertices are the only ones currently satisfying this property. They all have resilience one since Player 1 can enforce to reach a vertex of resilience one, but he cannot enforce to reach a vertex of resilience zero. Now, we can also determine the resilience of v_5 . The disturbance edge from v_5 to v_3 witnesses that it is two.

Afterwards, these two arguments no longer apply to new vertices: no disturbance edge leads from a $v \in \{v_7, \dots, v_{10}\}$ to a vertex whose resilience is already determined and Player 0 has a winning strategy from each of these vertices that additionally avoids vertices whose resilience is already determined. Thus, our reasoning cannot determine their resilience. This is consistent with our goal, as all four vertices have non-finite resilience, i.e., v_7 and v_8 have resilience ω and v_9 and v_{10} have resilience $\omega + 1$. Note that our reasoning here cannot distinguish these two values. We solve this problem in Subsection 3.2.

In this subsection, we formalize the reasoning described above: starting from the vertices in Player 1's winning region having resilience zero, we use a disturbance update and a risk update to determine all vertices of finite resilience. To simplify our proofs, we describe both as monotone operators updating partial rankings mapping vertices to ω , which might update already defined values. We show that alternatingly applying these updates eventually yields a stable ranking that indeed characterizes the vertices of finite resilience.

Throughout this section, we fix a game $\mathcal{G} = (\mathcal{A}, \text{Win})$ with $\mathcal{A} = (V, V_0, V_1, E, D)$ and with prefix-independent $\text{Win} \subseteq V^\omega$ satisfying the following condition: the game $(\mathcal{A}, \text{Win} \cap \text{Safety}(U))$ is determined for every $U \subseteq V$. We discuss this requirement in Section 4.

A ranking for \mathcal{G} is a partial mapping $r: V \rightarrow \omega$. The domain of r is denoted by $\text{dom}(r)$, its image by $\text{im}(r)$. Let r and r' be two rankings. We say that r' refines r if $\text{dom}(r') \supseteq \text{dom}(r)$ and $r'(v) \leq r(v)$ for all $v \in \text{dom}(r)$. A ranking r is sound, if we have $r(v) = 0$ if and only if $v \in \mathcal{W}_1(\mathcal{G})$ (cf. Remark 2).

Let r be a ranking for \mathcal{G} . We define the ranking r' as

$$r'(v) = \min(\{r(v)\} \cup \{r(v') + 1 \mid v' \in \text{dom}(r) \text{ and } (v, v') \in D\}),$$

with $\{r(v)\} = \emptyset$, if $r(v) \notin \text{dom}(r)$. We call r' the disturbance update of r .

Lemma 1. *The disturbance update r' of a sound ranking r is sound and refines r .*

Proof. As the minimization defining $r'(v)$ ranges over a superset of $\{r(v)\}$, we have $r'(v) \leq r(v)$ for every $v \in \text{dom}(r)$. This immediately implies refinement. From this inequality, we also obtain $r'(v) = 0$ for every $v \in \mathcal{W}_1(\mathcal{G})$, due to soundness of r . Finally, consider some $v \in \mathcal{W}_0(\mathcal{G})$. Then, $r(v) > 0$ by soundness of r . Thus, $r'(v) > 0$ as well, as both $r(v)$ and $r(v') + 1$ are greater than zero. Altogether, r' is sound as well. \square

Again, let r be a ranking for \mathcal{G} . For every $k \in \text{im}(r)$ let

$$A_k = \mathcal{W}_1(\mathcal{A}, \text{Win} \cap \text{Safety}(\{v \in \text{dom}(r) \mid r(v) \leq k\}))$$

the winning region of Player 1 in the game where he either wins by reaching a vertex v with $r(v) \leq k$ or by violating the winning condition. Now, define $r'(v) = \min\{k \mid v \in A_k\}$, where $r'(v)$ is undefined if v is in none of the A_k . We call r' the risk update of r .

Lemma 2. *The risk update r' of a sound ranking r is sound and refines r .*

Proof. We will show $r'(v) \leq r(v)$ for every $v \in \text{dom}(r)$, which implies both refinement and $r'(v) = 0$ for every $v \in \mathcal{W}_1(\mathcal{G})$, as argued in the proof of Lemma 1.

Thus, let $v \in \text{dom}(r)$. Trivially, $v \in \{v' \in \text{dom}(r) \mid r(v') \leq r(v)\}$. Thus, Player 1 wins the game $(\mathcal{A}, \text{Win} \cap \text{Safety}(\{v' \in \text{dom}(r) \mid r(v') \leq r(v)\}))$ from v by violating the safety condition right away. Hence, $v \in A_{r(v)}$ and thus $r'(v) \leq r(v)$.

To complete the proof of soundness of r' , we just have to show $r'(v) > 0$ for every $v \in \mathcal{W}_0(\mathcal{G})$. Towards a contradiction, assume $r'(v) = 0$, i.e., $v \in A_0$. Thus, Player 1 has a strategy τ from v that ensures that either the winning condition is violated or that a vertex v' with $r(v') = 0$ is reached, i.e., $v' \in \mathcal{W}_1(\mathcal{G})$ by soundness of r . Hence, Player 1 has a winning strategy $\tau_{v'}$ for \mathcal{G} from v' . This implies that he also has a winning strategy from v : play according to τ until a vertex v' with $r(v') = 0$ is reached. From there, mimic $\tau_{v'}$ when starting from v' . Every resulting disturbance-free play has a suffix that violates Win. Thus, by prefix-independence, the whole play violates Win as well, i.e., it is winning for Player 1. Thus, $v \in \mathcal{W}_1(\mathcal{G})$, which yields the desired contradiction, as winning regions are always disjoint. \square

Let r_0 be the unique sound ranking with domain $\mathcal{W}_1(\mathcal{G})$, i.e., r_0 maps exactly the vertices in Player 1's winning region to zero. Starting with r_0 , we inductively define a sequence of rankings $(r_j)_{j \in \omega}$ such that r_j for an odd (even) $j > 0$ is the disturbance update (the risk update) of r_{j-1} , i.e., we alternate between disturbance and risk updates.

Due to refinement, the r_j eventually stabilize, i.e., there is some j_0 such that $r_j = r_{j_0}$ for all $j \geq j_0$. Define $r^* = r_{j_0}$. Due to r_0 being sound and by Lemma 1 and Lemma 2, each r_j , and r^* in particular, is sound. If $v \in \text{dom}(r^*)$, let j_v be the minimal j with $v \in \text{dom}(r_j)$; otherwise, j_v is undefined.

Lemma 3. *If $v \in \text{dom}(r^*)$, then $r_{j_v}(v) = r_j(v)$ for all $j \geq j_v$.*

Proof. We show the following stronger result for every $v \in \text{dom}(r^*)$:

- If j_v is odd, then $r_j(v) = \frac{j_v+1}{2}$ for every $j \geq j_v$.
- If j_v is even, then $r_j(v) = \frac{j_v}{2}$ for every $j \geq j_v$.

The disturbance update increases the maximal rank by at most one and the risk update does not increase the maximal rank at all. Furthermore, due to refinement, the rank of v is set and then only decreases. Hence, we obtain $r_j(v) \leq \frac{jv+1}{2}$ and $r_j(v) \leq \frac{jv}{2}$ for odd and even j , respectively. In the remainder of the proof, we show a matching lower bound.

We say that a vertex v is updated to $k \in \omega$ in r_j if $r_j(v) = k$ and either $v \notin \text{dom}(r_{j-1})$ or both $v \in \text{dom}(r_{j-1})$ and $r_{j-1}(v) \neq k$ (here, r_{-1} is the unique ranking with empty domain). Now, we show the following by induction over j , which implies the matching lower bound.

- If j is odd, then no v is updated in r_j to some $k < \frac{j+1}{2}$.
- If j is even, then no v is updated in r_j to some $k < \frac{j}{2}$.

For $j = 0$, we have $\frac{j}{2} = 0$, and clearly, no vertex is assigned a negative rank by r_0 . For $j = 1$ and $j' = 2$, we obtain $\frac{j+1}{2} = \frac{j'}{2} = 1$. As r_0, r_1 , and r_2 are sound, neither r_1 nor r_2 update some v to zero.

Now let $j > 2$ and first consider the case where j is odd. Towards a contradiction, assume that $v \in V$ is updated in r_j to some value less than $\frac{j+1}{2}$. Since j is odd, r_j is the disturbance update of r_{j-1} . Further, as v is updated in r_j , there exists some disturbance edge $(v, v') \in D$ such that $r_j(v) = r_{j-1}(v') + 1$. Thus, $r_{j-1}(v') < r_j(v) < \frac{j+1}{2}$, i.e., $r_{j-1}(v') \leq \frac{j+1}{2} - 2 = \frac{j-3}{2}$. First, we show $r_{j-3}(v') = r_{j-2}(v') = r_{j-1}(v')$, i.e., the rank of v' is stable during the last two updates.

First assume towards a contradiction $r_{j-2}(v') \neq r_{j-1}(v')$. Then, v' is updated in r_{j-1} to some rank of at most $\frac{j-3}{2}$, which is in turn smaller than $\frac{j-1}{2}$, violating the induction hypothesis for $j-1$. Hence, $r_{j-2}(v') = r_{j-1}(v')$. The same reasoning yields a contradiction to the assumption $r_{j-3}(v') \neq r_{j-2}(v')$. Thus, we indeed obtain $r_{j-3}(v') = r_{j-2}(v') = r_{j-1}(v')$.

Since r_{j-2} is the disturbance update of r_{j-3} , we obtain $r_{j-2}(v) \leq r_{j-3}(v') + 1 = r_{j-1}(v') + 1 = r_j(v)$. Due to refinement, we obtain $r_{j-2}(v) \geq r_j(v)$, i.e., altogether $r_{j-2}(v) = r_{j-1}(v) = r_j(v)$. The latter equality contradicts our initial assumption, namely v being updated in r_j to $r_j(v)$.

Now consider the case where j is even. Again, assume towards a contradiction that $v \in V$ is updated in r_j to some value less than $\frac{j}{2}$. Since j is even, r_j is the risk update of r_{j-1} . Further, as v is updated in r_j , Player 1 wins the game $(\mathcal{A}, \text{Win} \cap \text{Safety}(U))$ from v , where $U = \{v' \in \text{dom}(r_{j-1}) \mid r_{j-1}(v') \leq r_j(v)\}$. Hence, he has a strategy τ such that every play starting in v and consistent with τ either violates Win or eventually visits some vertex v' with $r_{j-1}(v') \leq r_j(v)$. We claim $r_{j-2}(v') = r_{j-1}(v')$ for all $v' \in U$.

Towards a contradiction, assume $r_{j-2}(v') \neq r_{j-1}(v')$ for some $v' \in U$. Note that we have $r_{j-1}(v') \leq r_j(v) < \frac{j}{2}$. Thus, v' is updated in r_{j-1} to some value strictly less than $\frac{j}{2}$, which contradicts the induction hypothesis for $j-1$. Hence, we indeed obtain $r_{j-2}(v') = r_{j-1}(v')$ for all $v' \in U$.

Thus, there are two types of vertices v' in U : those for which $r_{j-3}(v')$ is defined, which implies $r_{j-3}(v') = r_{j-1}(v')$ due to the induction hypothesis and refinement, and those where $r_{j-3}(v')$ is undefined, which implies $r_{j-2}(v') = r_{j-1}(v')$ due to the claim above.

We claim that Player 1 wins $(\mathcal{A}, \text{Win} \cap \text{Safety}(\{v'' \in \text{dom}(r_{j-3}) \mid r_{j-3}(v'') \leq r_j(v)\}))$ from v , which implies $r_{j-2}(v) = r_j(v)$. This contradicts v being updated in r_j , our initial assumption.

To this end, we construct a strategy τ' from v that either violates Win or reaches a vertex v'' with $r_{j-3}(v'') \leq r_j(v)$ as follows. From v , τ' mimics τ until a vertex v' in U is reached (if it is at all). If v' is of the first type, then we have $r_{j-3}(v') = r_{j-1}(v') \leq r_j(v)$. If v' is of the second type, then v' is updated in r_{j-2} to some rank $r_{j-2}(v') = r_{j-1}(v') \leq r_j(v)$. As r_{j-2} is the risk update of r_{j-3} , Player 1 has a strategy $\tau_{v'}$ from v' that either violates Win or reaches a vertex v'' with $r_{j-3}(v'') \leq r_{j-2}(v') \leq r_j(v)$. Thus, starting in v' , τ' mimics $\tau_{v'}$ from v' until such a vertex is reached (if it is reached at all). Thus, every play that starts in v and is consistent with τ' either violates Win or reaches a vertex v'' with $r_{j-3}(v'') \leq r_j(v)$, which proves our claim. \square

From the proof of Lemma 3, we obtain an upper bound on the maximal rank of r^* . This in turn implies that the r_j stabilize quickly, as $r_j = r_{j+1} = r_{j+2}$ implies $r_j = r^*$.

Corollary 1.

1. We have $\text{im}(r^*) = \{0, 1, \dots, n\}$ for some $n < |V|$.
2. We have $r^* = r_{2|V|}$.

Lemma 3 also shows that an algorithm computing the r_j does not need to implement the definition of the two updates as presented above, but can be optimized by taking into account that a rank is never updated once set.

The main result of this section shows that r^* characterizes the resilience of vertices of finite resilience.

Lemma 4. *Let r^* be defined for \mathcal{G} as described above, and let $v \in V$.*

1. *If $v \in \text{dom}(r^*)$, then $r_{\mathcal{G}}(v) = r^*(v)$.*
2. *If $v \notin \text{dom}(r^*)$, then $r_{\mathcal{G}}(v) \in \{\omega, \omega + 1\}$.*

Proof. 1.) Let $v \in \text{dom}(r^*)$. We prove $r_{\mathcal{G}}(v) \leq r^*(v)$ and $r_{\mathcal{G}}(v) \geq r^*(v)$.

“ $r_{\mathcal{G}}(v) \leq r^*(v)$ ”: An α -resilient strategy from v is also α' -resilient from v for every $\alpha' \leq \alpha$. Thus, to prove

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

we just have to show that Player 0 has no $(r^*(v) + 1)$ -resilient strategy from v . By definition, for every strategy σ for Player 0, we have to show that there is a play ρ starting in v and consistent with σ that has at most $r^*(v)$ disturbances and is winning for Player 1. So, fix an arbitrary strategy σ .

We define a play with the desired properties by constructing longer and longer finite prefixes before finally appending an infinite suffix. During the construction, we ensure that each such prefix ends in $\text{dom}(r^*)$ in order to be able to proceed with our construction.

The first prefix just contains the starting vertex $(v, 0)$, i.e., the prefix does indeed end in $\text{dom}(r^*)$. Now, assume we have produced a prefix $w(v', b')$ ending in some vertex $v' \in \text{dom}(r^*)$, which implies that $j_{v'}$ is defined. We consider three cases:

- If $j_{v'} = 0$, then $v' \in \mathcal{W}_1(\mathcal{G})$ by definition of r_0 , i.e., Player 1 has a winning strategy τ from v' . Thus, we extend $w(v', b')$ by the unique disturbance-free play that starts in v' and is consistent with σ and τ , without its first vertex. In that case, the construction of the infinite play is complete.
- If $j_{v'} > 0$ is odd, then v' received its rank $r^*(v')$ during a disturbance update. Hence, there is some v'' such that $(v', v'') \in D$ with $r^*(v') - 1 = r^*(v'')$. In this case, we extend $w(v', b')$ by such a vertex v'' to obtain the new prefix $w(v', b')(v'', 1)$, which satisfies the invariant, as v'' is in $\text{dom}(r^*)$. Further, we have $j_{v''} < j_{v'}$ as the rank of v'' had to be defined in order to be considered during the disturbance update assigning a rank to v' .

- If $j_{v'} > 0$ is even, then v' received its rank $r^*(v')$ during a risk update. We claim that Player 1 has a strategy $\tau_{v'}$ that guarantees one of the following outcomes from v' : either the resulting play violates Win or it encounters a vertex $v'' \neq v'$ that satisfies $r^*(v'') \leq r^*(v')$ and $j_{v''} < j_{v'}$.

In that case, consider the unique disturbance-free play ρ' that starts in v' and is consistent with σ and the strategy $\tau_{v'}$ as above. If ρ' violates Win, then we extend $w(v', b')$ by ρ' without its first vertex. In that case, the construction of the infinite play is complete.

If ρ' does not violate Win, then we extend $w(v', b')$ by the prefix of ρ' without its first vertex and up to (and including) the first occurrence of a vertex v'' in ρ' satisfying the properties described above. Note that this again satisfies the invariant.

It remains to argue our claim: v' was assigned its rank $r^*(v') = r_{j_{v'}}(v')$ because it is in Player 1's winning region in the game with winning condition $\text{Win} \cap \text{Safety}(U)$, for

$$U = \{v'' \in \text{dom}(r_{j_{v'}-1}) \mid r_{j_{v'}-1}(v'') \leq r_{j_{v'}}(v')\}.$$

Hence, from v' , Player 1 has a strategy to either violate the winning condition or to reach U . Thus, $r_{j_{v'}-1}(v'') = r^*(v'')$ for every $v'' \in \text{dom}(r_{j_{v'}-1})$ yields $r^*(v'') \leq r^*(v')$. Finally, we have $j_{v''} < j_{v'}$, as the rank of v' was assigned due to the vertices in U already having ranks.

Note that only in two cases, we extend the prefix to an infinite play. In the other two cases, we just extend the prefix to a longer finite one. Thus, we first show that this construction always results in an infinite play. To this end, let $w_0(v_0, b_0)$ and $w_1(v_1, b_1)$ two of the prefixes constructed above such that $w_1(v_1, b_1)$ is an extension of $w_0(v_0, b_0)$. A simple induction proves $j_{v_1} < j_{v_0}$. Hence, as the value can only decrease finitely often, at some point an infinite suffix is added. Thus, we indeed construct an infinite play.

Finally, we have to show that the resulting play has the desired properties: by construction, the play starts in v and is consistent with σ . Furthermore, by construction, it has a disturbance-free suffix that violates Win. Thus, by prefix-independence, the whole play also violates Win. It remains to show that it has at most $r^*(v)$ disturbances. To this end, let $w_0(v_0, b_0)$ and $w_1(v_1, b_1)$ two of the prefixes such that $w_1(v_1, b_1)$ is obtained by extending $w_0(v_0, b_0)$ once. If the extension consists of taking the disturbance edge $(v_0, v_1) \in D$, then we have $r^*(v_1) = r^*(v_0) + 1$. The only other possibility is the extension

consisting of a finite play prefix that is consistent with the strategy τ_{v_0} . Then, by construction, we obtain $r^*(v_1) \leq r^*(v_0)$.

Thus, there are at most $r^*(v)$ many disturbances in the play, as the current rank decreases with every disturbance edge and does not increase with the other type of extension, but is always non-negative.

“ $r_{\mathcal{G}}(v) \geq r^*(v)$ ”: Here, we construct a strategy σ_f for Player 0 that is $r^*(v)$ -resilient from every $v \in \text{dom}(r^*)$, i.e., from v , σ_f has to be winning even under $r^*(v) - 1$ disturbances. As every strategy is 0-resilient, we only have to consider those v with $r^*(v) > 0$.

The proof is based on the fact that r^* is both stable under the disturbance and under the risk update, i.e., the disturbance update and the risk update of r^* are r^* , which yields the following properties. Let $(v, v') \in D$ be a disturbance edge such that $r^*(v) > 0$. Then, we have $r^*(v') \geq r^*(v) - 1$. Also, for every $v \in \text{dom}(r^*)$ with $r^*(v) > 0$, Player 0 has a winning strategy σ_v from v for the game $\mathcal{G}_v = (\mathcal{A}, \text{Win} \cap \text{Safety}(\{v' \in \text{dom}(r^*) \mid r^*(v') < r^*(v)\}))$ (note the strict inequality). Here, we apply the determinacy of the game \mathcal{G}_v , as the risk update is formulated in terms of Player 1’s winning region.

Now, we define σ_f as follows: it always mimics a strategy $\sigma_{v_{\text{cur}}}$ for some $v_{\text{cur}} \in \text{dom}(r^*)$, which is initialized by the starting vertex. The strategy $\sigma_{v_{\text{cur}}}$ is mimicked until a consequential (w.r.t. $\sigma_{v_{\text{cur}}}$) disturbance edge is taken, say by reaching the vertex v' . In that case, the strategy σ_f discards the history of the play constructed so far, updates v_{cur} to v' , and begins mimicking $\sigma_{v'}$. This is repeated ad infinitum.

Now, consider a play that starts in $\text{dom}(r^*)$, is consistent with σ_f , and has less than $r^*(v)$ disturbances. The part up to the first consequential disturbance edge (if it exists at all) is consistent with σ_v . Now, let (v_0, v'_0) be the corresponding disturbance edge. Then, we have $r^*(v_0) \geq r^*(v)$, as σ_v being a winning strategy for the safety condition never visits vertices with a rank smaller than $r^*(v)$. Thus, we conclude $r^*(v'_0) \geq r^*(v_0) - 1 \geq r^*(v) - 1$. Similarly, the part between the first and the second consequential disturbance edge (if it exists at all) is consistent with $\sigma_{v'_0}$. Again, if (v_1, v'_1) is the corresponding disturbance edge, then we have $r^*(v'_1) \geq r^*(v_1) - 1 \geq r^*(v) - 2$. Continuing this reasoning shows that less than $r^*(v)$ (consequential) disturbance edge lead to a vertex v' with $r^*(v') > 0$, as the rank is decreased by at most one for every disturbance edge. The suffix starting in this vertex is disturbance-free and consistent with $\sigma_{v'}$. Hence, the suffix satisfies Win, i.e., by prefix-independence, the whole play satisfies Win as well. Thus, σ_f is indeed $r^*(v)$ -resilient from every $v \in \text{dom}(r^*)$.

2.) Let $X = V \setminus \text{dom}(r^*)$. The disturbance update of r^* being r^* implies that every disturbance edge starting in X leads back to X . Similarly, the risk update of r^* being r^* implies $X = \mathcal{W}_0(\mathcal{G}_X)$ for $\mathcal{G}_X = (\mathcal{A}, \text{Win} \cap \text{Safety}(V \setminus X))$. Thus, from every $v \in X$, Player 0 has a strategy σ_v such that every disturbance-free play that starts in v and is consistent with σ_v satisfies the winning condition Win and never leaves X . Using these properties, we construct a strategy σ_ω that is ω -resilient from every $v \in X$, which implies $r_{\mathcal{G}}(v) \in \{\omega, \omega + 1\}$.

The definition of the strategy σ_ω here is similar to the one above yielding the lower bound on the resilience. Again, σ_ω always mimics a strategy $\sigma_{v_{\text{cur}}}$ for some $v_{\text{cur}} \in X$, which is initialized by the starting vertex. The strategy $\sigma_{v_{\text{cur}}}$ is mimicked until a consequential (w.r.t. $\sigma_{v_{\text{cur}}}$) disturbance edge is taken, say by reaching the vertex v' . In that case, the strategy σ_ω discards the history of the play constructed so far, updates v_{cur} to v' , and begins mimicking $\sigma_{v'}$. This is repeated ad infinitum.

Due to the properties of the disturbance edges and the strategies σ_v , such a play never leaves X , even if disturbances occur. Furthermore, if only finitely many disturbances occur, then the resulting play has a disturbance-free suffix that starts in some $v' \in X$ and is consistent with $\sigma_{v'}$. As $\sigma_{v'}$ is winning from v' in \mathcal{G}_X , this suffix satisfies Win. Hence, by prefix-independence of Win, the whole play also satisfies Win. Thus, σ_ω is indeed an ω -resilient strategy from every $v \in X$. \square

Combining Corollary 1 and Lemma 4, we obtain an upper bound on the resilience of vertices with finite resilience.

Corollary 2. *We have $r_{\mathcal{G}}(V) \cap \omega = \{0, 1, \dots, n\}$ for some $n < |V|$.*

3.2 Characterizing Vertices of Resilience $\omega + 1$

Our goal in this subsection is to determine the vertices of resilience $\omega + 1$, i.e., those from which Player 0 can win even under an infinite number of disturbances. Intuitively, in this setting, we give Player 1 control over the occurrence of disturbances, as he cannot execute more than infinitely many disturbances during a play. To this end, consider again the parity game in Figure 1 (on Page 2). From v_9 Player 0 wins even if

Player 1 controls whether the disturbance edge is taken from v_9 , as both v_9 and v_{10} have color zero. On the other hand, giving Player 1 control over the disturbance edges implies that he wins from v_7 , as he can use the one incident to v_7 infinitely often to move to v_8 , which has color one.




Fig. 2: The rigged game induced by the game shown in Figure 1 (on Page 2).

In the following, we prove this intuition to be correct. To this end, we transform the arena of the game so that at a Player 0 vertex, first Player 1 gets to choose whether he wants to take one of the disturbance edges and, if not, gives control to Player 0, who is then able to use a standard edge. Figure 2 depicts the game corresponding to the one from Figure 1. And indeed, Player 0 wins this so-called rigged game, which still has a parity condition, from v_9 and v_{10} , but not from any other vertex. These are the only vertices of resilience $\omega + 1$ in the original game.

Given a game $\mathcal{G} = (\mathcal{A}, \text{Win})$ with $\mathcal{A} = (V, V_0, V_1, E, D)$, define the rigged game $\mathcal{G}_{\text{rig}} = (\mathcal{A}', \text{Win}')$ with $\mathcal{A}' = (V', V'_0, V'_1, E', D')$ such that $V' = V'_0 \cup V'_1$ with $V'_0 = \{\bar{v} \mid v \in V_0\}$ and $V'_1 = V$ and $D' = \emptyset$. The set E' of edges is the union of the following sets:

- D : Player 1 uses a disturbance edge.
- $\{(v, \bar{v}) \mid v \in V_0\}$: Player 1 does not use a disturbance edge and yields control to Player 0.
- $\{(\bar{v}, v') \mid (v, v') \in E \text{ and } v \in V_0\}$: Player 0 has control and picks a standard edge.
- $\{(v, v') \mid (v, v') \in E \text{ and } v \in V_1\}$: Player 1 takes a standard edge.

Furthermore, $\text{Win}' = \{\rho \in (V')^\omega \mid h(\rho) \in \text{Win}\}$ where h is the homomorphism induced by defining $h(v) = v$ and $h(\bar{v}) = \varepsilon$ for every $v \in V$.

We now show that $\mathcal{W}_0(\mathcal{G}_{\text{rig}})$ characterizes the vertices of resilience $\omega + 1$ in \mathcal{G} . Note that we have no assumptions on \mathcal{G} here.

Lemma 5. *Let v be a vertex of a game \mathcal{G} . Then, $v \in \mathcal{W}_0(\mathcal{G}_{\text{rig}})$ if and only if $r_{\mathcal{G}}(v) = \omega + 1$.*

Proof. “ \Rightarrow ”: Let Player 0 win \mathcal{G}_{rig} from v , say with winning strategy σ' . We inductively translate play prefixes w in \mathcal{G} into play prefixes $t'(w)$ in \mathcal{G}_{rig} satisfying the following invariant: $t'((v_0, b_0) \cdots (v_j, b_j))$ starts in v_0 and ends in v_j .

For the induction start, we define $t'(v_0, b_0) = (v_0, 0)$; to define $t'((v_0, b_0) \cdots (v_j, b_j)(v_{j+1}, b_{j+1}))$, we consider several cases:

- If $b_{j+1} = 1$, then $(v_j, v_{j+1}) \in D$, i.e., the play traverses the disturbance edge (v_j, v_{j+1}) . This move is mimicked by defining $t'((v_0, b_0) \cdots (v_j, b_j)(v_{j+1}, b_{j+1})) = t'((v_0, b_0) \cdots (v_j, b_j)) \cdot (v_{j+1}, 0)$.
- If $b_{j+1} = 0$, i.e., $(v_j, v_{j+1}) \in E$, and $v_j \in V_0$, then the play did not traverse a disturbance edge and instead allowed Player 0 to pick a standard edge (v_j, v_{j+1}) to traverse. This move is mimicked by defining $t'((v_0, b_0) \cdots (v_j, b_j)(v_{j+1}, b_{j+1})) = t'((v_0, b_0) \cdots (v_j, b_j)) \cdot (\bar{v}_j, 0) \cdot (v_{j+1}, 0)$.
- If $b_{j+1} = 0$, i.e., $(v_j, v_{j+1}) \in E$, and $v_j \in V_1$, then the play traversed the standard edge (v_j, v_{j+1}) . This move is mimicked by defining $t'((v_0, b_0) \cdots (v_j, b_j)(v_{j+1}, b_{j+1})) = t'((v_0, b_0) \cdots (v_j, b_j)) \cdot (v_{j+1}, 0)$.

Note that our invariant is satisfied in any case. Also, we lift t' to infinite plays by taking limits as usual.

Using this translation, we define a strategy σ for Player 0 in \mathcal{G} via

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

where $b_0 = 0$ and where for every $j' > 0$, $b_{j'} = 1$ if and only if $v_{j'} \neq \sigma(v_0 \cdots v_{j'-1})$, i.e., we reconstruct the consequential disturbances. A straightforward induction shows that for every play $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ in \mathcal{G} that is consistent with σ , the play $t'(\rho)$ is consistent with σ' . Hence, $t'(\rho) \in \text{Win}'$ for every ρ starting in v . Furthermore, we have $h(t'(\rho)) = v_0 v_1 v_2 \cdots \in \text{Win}$, as $t'(\rho) \in \text{Win}'$. Thus, $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ is winning for Player 0. As we have no restriction on the number of disturbances in ρ , σ is $(\omega + 1)$ -resilient from v . Thus, $r_{\mathcal{G}}(v) = \omega + 1$.

“ \Leftarrow ”: Now, let $r_{\mathcal{G}}(v) = \omega + 1$, i.e., Player 0 has an $(\omega + 1)$ -resilient strategy σ from v in \mathcal{G} . This time, we inductively define a translation t of play prefixes in \mathcal{G}_{rig} into play prefixes in \mathcal{G} . Here, it suffices to consider those prefixes that start and end in V'_1 . For these, we satisfy the following invariant: if w starts in v_0 and ends in v_j , then $t(w)$ starts in v_0 and ends in v_j as well. Note that \mathcal{G}_{rig} has no disturbance edges. Hence, thus the bits indicating whether such an edge has been traversed are always zero in plays of \mathcal{G}_{rig} . Thus, we define $t(v_0, 0) = (v_0, 0)$ and consider several cases for the inductive step:

- First, assume we have a prefix of the form $(v_0, 0) \cdots (v_j, 0)(v_{j+1}, 0)$ for some $v_j \in V_0$, i.e., Player 1’s move simulates the disturbance edge $(v_j, v_{j+1}) \in D$. Then, we define $t((v_0, 0) \cdots (v_j, 0)(v_{j+1}, 0)) = t((v_0, 0) \cdots (v_j, 0)) \cdot (v_{j+1}, 1)$.
- Next, assume we have a prefix of the form $(v_0, 0) \cdots (v_j, 0)(v_{j+1}, 0)$ for some $v_j \in V_1$, i.e., Player 1’s move simulates the standard edge $(v_j, v_{j+1}) \in E$. Then, we define $t((v_0, 0) \cdots (v_j, 0)(v_{j+1}, 0)) = t((v_0, 0) \cdots (v_j, 0)) \cdot (v_{j+1}, 0)$.
- Finally, the last case is a prefix of the form $(v_0, 0) \cdots (v_j, 0)(\overline{v_j}, 0)(v_{j+1}, 0)$ for some $v_j \in V_0$, i.e., Player 0’s move simulates the standard edge $(v_j, v_{j+1}) \in E$. Then, we define $t((v_0, 0) \cdots (v_j, 0)(\overline{v_j}, 0)(v_{j+1}, 0)) = t((v_0, 0) \cdots (v_j, 0)) \cdot (v_{j+1}, 0)$.

The invariant is satisfied in any case. Also, we can again lift t to infinite plays via limits.

Now, we define a strategy σ' for Player 0 in \mathcal{G}_{rig} via

$$\sigma'(v_0 \cdots v_j \overline{v_j}) = \sigma(t(v_0 \cdots v_j)).$$

A straightforward induction shows that for every play ρ that is consistent with σ' , the play $t(\rho)$ is consistent with σ . Hence, $t(\rho)$ satisfies the winning condition, if ρ starts in v , as σ is $(\omega + 1)$ -resilient from v . Let $t(\rho) = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$. Then, $v_0 v_1 v_2 \cdots \in \text{Win}$. Now, $h(\rho) = v_0 v_1 v_2 \cdots$ implies $\rho \in \text{Win}'$. Thus, σ' is a winning strategy for Player 0 from v . \square

Consider the first implication proved above. If σ is positional, then σ' is positional as well. Thus, applying both implications yields the following corollary.

Corollary 3. *Assume Player 0 has a positional winning strategy for \mathcal{G}_{rig} from v . Then, Player 0 has an $(\omega + 1)$ -resilient positional strategy from v .*

3.3 Computing Optimally Resilient Strategies

This subsection is concerned with computing the resilience and optimally resilient strategies. Recall that in the proof of Lemma 4, we have constructed strategies σ_f and σ_ω such that σ_f (σ_ω) is $r_{\mathcal{G}}(v)$ -resilient from every v with $r_{\mathcal{G}}(v) \in \omega$ (with $r_{\mathcal{G}}(v) = \omega$). However, even if the strategies σ_v used to construct them are positional, the strategies σ_f and σ_ω are in general not positional, as they have to store the vertex v_{cur} to simulate the strategy $\sigma_{v_{\text{cur}}}$. In the following, we show how to combine such positional strategies σ_v and a positional one for \mathcal{G}_{rig} into a single positional optimally resilient strategy. To this end, we refine the following standard technique that combines positional winning strategies for games with prefix-independent winning conditions.

Assume we have a positional strategy σ_v for every vertex v in some set $W \subseteq V$ such that σ_v is winning from v . Furthermore, let R_v be the set of vertices visited by plays that start in v and are consistent with σ_v . Furthermore, let $m(v) = \min_{\prec} \{v' \in V \mid v \in R_{v'}\}$ for some strict total ordering \prec of W . Then, the positional strategy σ defined by $\sigma(v) = \sigma_{m(v)}(v)$ is winning from each $v \in W$, as along every play that starts in some $v \in W$ and is consistent with σ , the value of the function m only decreases. Thus, after it has stabilized, the remaining suffix is consistent with some strategy σ_v . Hence, the suffix is winning for Player 0 and prefix-independence implies that the whole play is winning for her as well.

Here, we have to adapt this reasoning to respect the resilience of the vertices and to handle disturbance edges. In particular, we have to pay attention to vertices of resilience $\omega + 1$, as plays starting in such vertices have to be winning under infinitely many disturbances.

Recall the requirements from Subsection 3.1 for a game $(\mathcal{A}, \text{Win})$: Win is prefix-independent and the game \mathcal{G}_U is determined for every $U \subseteq V$, where we write \mathcal{G}_U for the game $(\mathcal{A}, \text{Win} \cap \text{Safety}(U))$ for some $U \subseteq V$. To prove the results of this subsection, we need to impose some additional effectiveness requirements: we require that each game \mathcal{G}_U and the rigged game \mathcal{G}_{rig} can be effectively solved. Also, we first assume that Player 0 has positional winning strategies for each of these games, which have to be effectively computable as well. We discuss the severity of these requirements in Section 4.

Theorem 1. *Let \mathcal{G} satisfy all the above requirements. Then, the resilience of \mathcal{G} 's vertices and a positional optimally resilient strategy can be effectively computed.*

Proof. The effective computability of the resilience follows from the effectiveness requirements on \mathcal{G} : to compute the ranking r^* , it suffices to compute the disturbance and risk updates. The former are trivially effective while the effectiveness of the latter ones follows from our assumption. Lemma 4 shows that r^* correctly determines the resilience of all vertices with finite resilience. Finally by solving the rigged game, we also correctly determine the resilience of the remaining vertices (Lemma 5). Again, this game can be solved by our assumption.

Thus, it remains to show how to compute a positional optimally resilient strategy. To this end we compute a positional strategy σ_v for every v satisfying the following properties.

- For every $v \in V$ with $r_{\mathcal{G}}(v) \in \omega \setminus \{0\}$, the strategy σ_v is winning for Player 0 from v for the game $(\mathcal{A}, \text{Win} \cap \text{Safety}(\{v' \in V \mid r_{\mathcal{G}}(v') < r_{\mathcal{G}}(v)\}))$. The existence of such a strategy has been shown in the proof of Item 1 of Lemma 4.
- For every $v \in V$ with $r_{\mathcal{G}}(v) = \omega$, the strategy σ_v is winning for Player 0 from v for the game $(\mathcal{A}, \text{Win} \cap \text{Safety}(\{v' \in V \mid r_{\mathcal{G}}(v') \in \omega\}))$. The existence of such a strategy has been shown in the proof of Item 2 of Lemma 4.
- For every $v \in V$ with $r_{\mathcal{G}}(v) = \omega + 1$, the strategy σ_v is $(\omega + 1)$ -resilient from v . The existence of such a strategy follows from Corollary 3, as we assume Player 0 to win \mathcal{G}_{rig} with positional strategies.
- Finally, for every $v \in V$ with $r_{\mathcal{G}}(v) = 0$, we fix an arbitrary positional strategy σ_v for Player 0.

Furthermore, we fix a strict linear order \prec on V such that $v \prec v'$ implies $r_{\mathcal{G}}(v) \leq r_{\mathcal{G}}(v')$, i.e., we order the vertices by ascending resilience. For a vertex v with $r_{\mathcal{G}}(v) \neq \omega + 1$, let R_v be the set of vertices reachable via disturbance-free plays that start in v and are consistent with σ_v . On the other hand, for a vertex v with $r_{\mathcal{G}}(v) = \omega + 1$, let R_v be the set of vertices reachable via plays with arbitrarily many disturbances that start in v and are consistent with σ_v .

We claim $R_v \subseteq \{v' \in V \mid r_{\mathcal{G}}(v') \geq r_{\mathcal{G}}(v)\}$ for every $v \in V$ (*). For v with $r_{\mathcal{G}}(v) \neq \omega + 1$ this follows immediately from the choice of σ_v . Thus, let us argue the claim for v with $r_{\mathcal{G}}(v) = \omega + 1$. Assume σ_v reaches a vertex v' of resilience $r_{\mathcal{G}}(v') \neq \omega + 1$. Then, there exists a play ρ' starting in v' that is consistent with σ_v , has less than $r_{\mathcal{G}}(v')$ many disturbances and is losing for Player 0. Thus the play obtained by first taking the play prefix to v' and then appending ρ' without its first vertex yields a play starting in v , consistent with σ_v , but losing for Player 0. This play implies that σ_v is not $(\omega + 1)$ -resilient from v , which yields the desired contradiction.

Let $m: V \rightarrow V$ be given as $m(v) = \min_{\prec}\{v' \in V \mid v \in R_{v'}\}$ and define the positional strategy σ as $\sigma(v) = \sigma_{m(v)}(v)$. By our assumptions, σ can be effectively computed. It remains to show that it is optimally resilient.

To this end, we apply the following two properties of edges (v, v') that may appear during a play that is consistent with σ , i.e., we either have $v \in V_0$ and $\sigma(v) = v'$ (which implies $(v, v') \in E$), or $v \in V_1$ and $(v, v') \in E$, or $v \in V_0$ and $(v, v') \in D$:

1. If $(v, v') \in E$, then we have $r_{\mathcal{G}}(v) \leq r_{\mathcal{G}}(v')$ and $m(v) \geq m(v')$. The first property follows from minimality of $m(v)$ and (*) while the second one follows from the definition of R_v .
2. If $(v, v') \in D$, then we have to distinguish several subcases, which all follow immediately from the definition of resilience:
 - If $r_{\mathcal{G}}(v) \in \omega$, then $r_{\mathcal{G}}(v') \geq r_{\mathcal{G}}(v) - 1$.
 - If $r_{\mathcal{G}}(v) = \omega$, then $r_{\mathcal{G}}(v') = \omega$, and

- If $r_{\mathcal{G}}(v) = \omega + 1$, then $r_{\mathcal{G}}(v') = \omega + 1$ and $m(v) \geq m(v')$ (here, the second property follows from the definition of R_v for v with $r_{\mathcal{G}}(v) = \omega + 1$, which takes disturbance edges into account).

Now, consider a play $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ that is consistent with σ . If $r_{\mathcal{G}}(v_0) = 0$ then we have nothing to show, as every strategy is 0-resilient from v .

Now, assume $r_{\mathcal{G}}(v_0) \in \omega \setminus \{0\}$. We have to show that if ρ has less than $r_{\mathcal{G}}(v_0)$ disturbances, then it is winning for Player 0. An inductive application of the above properties shows that in that case the last disturbance edge leads to a vertex of non-zero resilience. Furthermore, as the values $m(v_j)$ are only decreasing afterwards, they have to stabilize at some later point. Hence, there is some suffix of ρ that starts in some v' with non-zero resilience and that is consistent with the strategy $\sigma_{v'}$. Thus, the suffix is winning for Player 0 by the choice of $\sigma_{v'}$ and prefix-independence implies that ρ is winning for her as well.

Next, assume $r_{\mathcal{G}}(v_0) = \omega$. We have to show that if ρ has a finite number of disturbances, then it is winning for Player 0. Again, an inductive application of the above properties shows that in that case the last disturbance edge leads to a vertex of resilience ω or $\omega + 1$. Afterwards, the values $m(v_j)$ stabilize again. Hence, there is some suffix of ρ that starts in some v' with non-zero resilience and that is consistent with the strategy $\sigma_{v'}$. Thus, the suffix is winning for Player 0 by the choice of $\sigma_{v'}$ and prefix-independence implies that ρ is winning for her as well.

Finally, assume $r_{\mathcal{G}}(v_0) = \omega + 1$. Then, the above properties imply that ρ only visits vertices with resilience $\omega + 1$ and that the values $m(v_j)$ eventually stabilize. Hence, there is a suffix of ρ that is consistent with some $(\omega + 1)$ -resilient strategy $\sigma_{v'}$, where v' is the first vertex of the suffix. Hence, the suffix is winning for Player 0, no matter how many disturbances occurred. This again implies that ρ is winning for her as well. \square

Next, we analyze the complexity of the algorithm sketched above in some more detail. The inductive definition of the r_j can be turned into an algorithm computing r^* (using the results of Lemma 3 to optimize the naive implementation), which has to solve $\mathcal{O}(|V|)$ many games (and compute winning strategies for some of them) with winning condition $\text{Win} \cap \text{Safety}(U)$. Furthermore, the rigged game, which is of size $\mathcal{O}(|V|)$, has to be solved and winning strategies have to be determined. Thus, the overall complexity is in general dominated by the complexity of solving these tasks. We explicitly state one complexity result for the important case of parity games, using the fact that each of these games is then a parity game as well. Also, we use one of the recently presented quasipolynomial time algorithms for solving parity games [6,12,17] to solve the \mathcal{G}_U and \mathcal{G}_{rig} .

Theorem 2. *Optimally resilient strategies in parity games are positional and can be computed in quasipolynomial time.*

Using similar arguments, one can also analyze games where positional strategies do not suffice. As above, assume \mathcal{G} satisfies the same determinacy and effectiveness assumptions, but only require that Player 0 has finite-state winning strategies⁴ for each game with winning condition $(\mathcal{A}, \text{Win} \cap \text{Safety}(U))$ and for the rigged game \mathcal{G}_{rig} . Then, one can show that she has a finite-state optimally resilient strategy. In fact, by reusing memory states, one can construct an optimally resilient strategy that it is not larger than any constituent strategy.

4 Discussion

In this section, we discuss the assumptions required to be able to compute (positional) optimally resilient strategies with the algorithm presented in Section 3. To this end, fix a game $\mathcal{G} = (\mathcal{A}, \text{Win})$ with vertex set V and recall that \mathcal{G}_{rig} is the corresponding rigged game and that we defined $\mathcal{G}_U = (\mathcal{A}, \text{Win} \cap \text{Safety}(U))$ for every $U \subseteq V$. Now, the assumptions on \mathcal{G} for Theorem 1 to hold are as follows:⁵

1. Every game \mathcal{G}_U is determined.
2. Player 0 has a positional winning strategy from every vertex in her winning regions in the games \mathcal{G}_U and in the rigged game \mathcal{G}_{rig} .

⁴ A finite state strategy is implemented by a finite automaton that processes play prefixes and outputs vertices to move to. See Section 4 for a formal definition.

⁵ The subsequent reasoning also holds true for games with finite-state strategies.

3. Each game \mathcal{G}_U and the game \mathcal{G}_{rig} can be effectively solved and positional winning strategies can be effectively computed for each such game.
4. Win is prefix-independent.

First, consider the determinacy assumption. It is straightforward to show

$$\mathcal{W}_0(\mathcal{G}_U) = \mathcal{W}_0(\mathcal{A} \setminus W, \text{Win} \cap (V \setminus W)^\omega)$$

with $W = \mathcal{W}_1(\mathcal{A}, \text{Safety}(U))$. Hence, one can first determine and then remove the winning region of Player 1 in the safety game and only then solve the subgame of \mathcal{G} played in Player 0's winning region of the safety game. Thus, all subgames of \mathcal{G} being determined is a sufficient condition for our determinacy requirement being satisfied. The winning conditions one typically studies, e.g., parity and in fact all Borel ones [20], satisfy this property.

The next requirement concerns the existence of positional winning strategies for the games \mathcal{G}_U and \mathcal{G}_{rig} . For the former games, this requirement is satisfied if Player 0 has positional winning strategies for all subgames of \mathcal{G} . As every positional optimally resilient strategy is also a winning strategy in a certain subgame, this condition is necessary. Now, consider the rigged game, whose winning condition can be written as $\text{Win} = \{v_0v_1v_2 \cdots \in (V')^\omega \mid v_0v_2v_4 \cdots \in \text{Win}\}$. The winning conditions one typically studies, e.g., again the Borel ones, are closed under taking such supersequences. Thus, if \mathcal{G} is from a class of winning conditions that allows for positional winning strategies for Player 0, then this class typically also contains \mathcal{G}_{rig} .

Also, the assumption on the effective solvability and computability of positional strategies is obviously necessary, as we solve a more general problem here when determining optimally resilient strategies.




Fig. 3: Counterexample to the correctness of the computation of resilience for games with prefix-dependent winning conditions.

Finally, let us consider prefix-independence. If the winning condition is not prefix-independent, the algorithm presented in Section 3 does not compute the resilience of vertices correctly anymore. As an example, consider the family $\mathcal{G}_k = (\mathcal{A}, \text{Win}_k)$ of games shown in Figure 3. In \mathcal{G}_k , it is Player 0's goal to avoid more than k visits to v . Such a visit only occurs via a disturbance or if the initial vertex is v . Hence, we have $r_{\mathcal{G}_k}(v) = k$ and $r_{\mathcal{G}_k}(v') = k + 1$. Applying the algorithm from Section 3, however, the initial ranking function r_0 has an empty domain, since we have $\mathcal{W}_1(\mathcal{G}_k) = \emptyset$. Thus, the computation of the r_j immediately stabilizes, yielding r^* with empty domain. This is a counterexample to Lemma 4, if the winning condition is prefix-dependent.

To conclude this section, we show that we can, however, still leverage the algorithm from Section 3 in order to compute the resilience of a wide range of games with prefix-dependent winning conditions. To this end, we extend the framework of game reductions to games with disturbances.

Fix a game $(\mathcal{A}, \text{Win})$ with $\mathcal{A} = (V, V_0, V_1, E, D)$. A memory structure for \mathcal{A} is a triple $\mathcal{M} = (M, \text{Init}, \text{Upd})$ where M is a finite set of memory states, $\text{Init}: V \rightarrow M$ is the initialization function, and $\text{Upd}: M \times V \rightarrow M$ is the memory update function.

The update function can be extended to finite play prefixes: $\text{Upd}^+(v) = \text{Init}(v)$ and $\text{Upd}^+(wv) = \text{Upd}(\text{Upd}^+(w), v)$ for $w \in V^+$ and $v \in V$. A next-move function $\text{Nxt}: V_i \times M \rightarrow V$ for Player i has to satisfy $(v, \text{Nxt}(v, m)) \in E$ for all $v \in V_i$ and all $m \in M$. It induces a strategy σ for Player i with memory \mathcal{M} via $\sigma(v_0 \cdots v_j) = \text{Nxt}(v_j, \text{Upd}^+(v_0 \cdots v_j))$. A strategy is called finite-state if it can be implemented by a memory structure.

The arena \mathcal{A} and $\mathcal{M} = (M, \text{Init}, \text{Upd})$ induce the expanded arena $\mathcal{A} \times \mathcal{M} = (V \times M, V_0 \times M, V_1 \times M, E', D')$ where E' is defined via $((v, m), (v', m')) \in E'$ if and only if $(v, v') \in E$ and $\text{Upd}(m, v') = m'$. The disturbance edges D' are defined analogously, i.e., $((v, m), (v', m')) \in D'$ if and only if $(v, v') \in D$ and $\text{Upd}(m, v') = m'$. Every play $(v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$ in \mathcal{A} has a unique extended play $\text{ext}(\rho) = ((v_0, m_0), b_0)((v_1, m_1), b_1)((v_2, m_2), b_2) \cdots$ in $\mathcal{A} \times \mathcal{M}$ defined by $m_0 = \text{Init}(v_0)$ and $m_{j+1} = \text{Upd}(m_j, v_{j+1})$, i.e., $m_j = \text{Upd}^+(v_0 \cdots v_j)$. Play prefixes are translated analogously.

Remark 3. Let ρ be a play in \mathcal{G} . Then, $\#_D(\rho) = \#_D(\text{ext}(\rho))$.

A game $\mathcal{G} = (\mathcal{A}, \text{Win})$ is reducible to $\mathcal{G}' = (\mathcal{A}', \text{Win}')$ via \mathcal{M} , written $\mathcal{G} \leq_{\mathcal{M}} \mathcal{G}'$, if $\mathcal{A}' = \mathcal{A} \times \mathcal{M}$ and every play ρ in \mathcal{G} is won by the same player that wins $\text{ext}(\rho)$ in \mathcal{G}' .

Lemma 6. *Let $\mathcal{G} \leq_{\mathcal{M}} \mathcal{G}'$. Then, $r_{\mathcal{G}}(v) = r_{\mathcal{G}'}(v, \text{Init}(v))$ for all vertices v of \mathcal{G} .*

Proof. We show that Player 0 has an α -resilient strategy σ' for \mathcal{G}' from $(v, \text{Init}(v))$ if and only if she has an α -resilient strategy σ for \mathcal{G} from v , which implies our claim. The translation of the strategies is the same as in the disturbance-free setting (see, e.g., [18]), but here we have to argue about resilience instead of just winning.

“ \Leftarrow ”: Given a strategy σ for \mathcal{G} , we define σ' for \mathcal{G}' via $\sigma'((v_0, m_0) \cdots (v_j, m_j)) = \sigma(v_0 \cdots v_j)$. Consider a play $\rho' = ((v_0, m_0), b_0)((v_1, m_1), b_1)((v_2, m_2), b_2) \cdots$ consistent with σ' . If $m_0 = \text{Init}(v_0)$, then $\rho' = \text{ext}(\rho)$ for $\rho = (v_0, b_0)(v_1, b_1)(v_2, b_2) \cdots$, which is consistent with σ . Hence, ρ' and ρ having the same winner and the same number of faults implies that σ is α -resilient from a vertex v if and only if σ' is α -resilient from $(v, \text{Init}(v))$.

“ \Rightarrow ”: Given a strategy σ' for \mathcal{G}' , we define σ for \mathcal{G} via $\sigma(v_0 \cdots v_j) = v$, if $\sigma'(\text{ext}(v_0 \cdots v_j)) = (v, m)$ for some $m \in M$. A straightforward induction shows that a play in \mathcal{G} is consistent with σ if and only if its extended play in \mathcal{G}' is consistent with σ' . Thus, these plays having the same winner and the same number of faults implies that σ is α -resilient from a vertex v if and only if σ' is α -resilient from $(v, \text{Init}(v))$. \square

As usual for game reductions, we obtain a finite-state strategy for \mathcal{G} when starting with a positional strategy in \mathcal{G}' . To this end, consider the proof of the second implication above. If σ is positional, then the strategy σ' is implemented by \mathcal{M} and the next-move function Nxt given by $\text{Nxt}(v, m) = v'$, if $\sigma(v, m) = (v', m')$ for some $m' \in M$.

A similar construction works in case σ' is finite-state, say implemented by \mathcal{M}' . Then, σ is implemented by the product of \mathcal{M} and \mathcal{M}' , which is defined as expected (we refer to, e.g., [18] for a formal definition). Altogether, we obtain the following result.

Corollary 4. *Let $\mathcal{G} \leq_{\mathcal{M}} \mathcal{G}'$.*

1. *If Player 0 has an α -resilient positional strategy from $(v, \text{Init}(v))$ in \mathcal{G}' , then she has an α -resilient finite-state strategy from v in \mathcal{G} , which is implemented by \mathcal{M} .*
2. *If Player 0 has an α -resilient finite-state strategy from $(v, \text{Init}(v))$ in \mathcal{G}' , say implemented by \mathcal{M}' , then she has an α -resilient finite-state strategy from v in \mathcal{G} , which is implemented by the product of \mathcal{M} and \mathcal{M}' .*

Now, we can formulate the main theorem of this section, which shows that prefix-dependence is not a restriction, as long as the game is reducible to a prefix-independent one. Note that this is in particular true for every ω -regular winning condition (see, e.g., [15]): every such condition is recognized by a deterministic parity automaton, which can be turned into a memory structure which allows to reduce the original game to a parity game.

Theorem 3. *Let $\mathcal{G} \leq_{\mathcal{M}} \mathcal{G}'$ so that \mathcal{G}' has a prefix-independent winning condition, can be effectively computed from \mathcal{G} , and satisfies the assumptions from Section 3 (with finite-state strategies).*

Then, the resilience of \mathcal{G} 's vertices and a finite-state optimally resilient strategy can be effectively computed.

Proof. This is a direct consequence of Lemma 6 and Theorem 1. To obtain an optimally resilient strategy, we apply Corollary 4 for finite-state strategies. \square

Recall the family of games shown in Figure 3 in which Player 0 aims to prevent more than k visits to the vertex v_1 for some parameter $k \in \omega$. Such a game can be reduced to a parity game using a memory structure implementing a counter up to $k + 1$. Thus, we obtain an optimally resilient strategy for Player 0 that is implemented by a memory structure with $k + 1$ states. While this strategy is indeed optimally resilient, it is not minimal: in fact, the unique strategy for Player 0 is positional and optimally resilient. Thus, the approach of computing optimally resilient strategies for games with prefix-dependent winning conditions via reductions to prefix-independent winning conditions is not optimal in that sense, as it may

yield unnecessarily large optimally resilient strategies. In current research, we study how to synthesize minimal optimally resilient strategies for games with prefix-dependent winning conditions.

Moreover, in the case of prefix-dependent winning conditions, the question arises whether or not optimally resilient strategies may be necessarily larger than winning ones. It is easy to construct a game in which Player 0 has a positional winning strategy, but an optimally resilient one requires an infinite amount of memory. One example is a game with a single vertex with a self-loop, from which a disturbance edge leads into a disturbance-free subgame in which Player 0 needs an infinite amount of memory to win. Thus, it is an interesting question for further research whether a similar result to Theorem 1 holds true for prefix-dependent games with positional winning strategies, e.g., weak parity games [7] or bounded parity games [8].⁶

5 Outlook

In this work we have developed a fine-grained view on the quality of strategies: instead of evaluating whether or not a given strategy is winning or not, we evaluate it according to its resilience against intermittent disturbances. While this measure of quality enables the construction of “better” strategies than the distinction between winning and losing strategies, there remain aspects of optimality that are not captured in our notion of resilience. In this section we discuss these aspects and give examples of games in which further analysis yields crucial differences between optimally resilient strategies. In further research, we aim to synthesize optimal strategies with respect to these criteria.




Fig. 4: Intuitively, moving from v_1 to v_3 is preferable for Player 0, as it allows her to possibly “recover” from a first fault with the “help” of a second one.

As a first example, consider the parity game shown in Figure 4. Vertices v_0 and v_3 have resilience 1 and $\omega + 1$, respectively, while vertices v_1 , v_2 , and v_2' have resilience 0. Player 0’s only choice consists of choosing to move to v_2 or to v_2' from v_1 . Let σ and σ' be the unique positional strategies for Player 0 that move to v_2 and v_2' from v_1 , respectively. Both strategies are optimally resilient and thus, the algorithm from Section 3 may yield either one, depending on the underlying parity game solver used. Intuitively, however, σ' is preferable for Player 0 over σ , as a play prefix ending in v_2' may proceed to her winning region if a single disturbance occurs. All plays encountering v_2 at some point, however, are losing for her. Hence, another interesting avenue for further research is to study **how to recover from losing**, i.e., how to construct strategies that leverage benevolent disturbances in order to recover from entering Player 1’s winning region. For safety games, this has been addressed by Dallal, Neider, and Tabuada [9].

The previous example shows that Player 0 can still make “meaningful” choices even if the play has moved outside her winning region. The game \mathcal{G} shown in Figure 5 demonstrates that she can do so as well when remaining in vertices of resilience ω . Every vertex in \mathcal{G} has resilience ω , since every play with finitely many disturbances eventually remains in vertices of color 0. Moreover, the only choice to be made by Player 0 is whether to move to vertex v_1 or to vertex v_1' from vertex v_0 . Let σ and σ' be positional strategies that implement the former and the latter choice, respectively.

First consider a scenario in which visiting an odd color models the occurrence of some undesirable event, e.g., that a request has not been answered. In this case, Player 0 should aim to prevent visits to v_3 in \mathcal{G} , the only vertex of odd color. Hence, the strategy σ should be more desirable for her, as it requires two disturbances in direct succession in order to trigger a visit to v_3 . When playing consistently with σ' , however, a single disturbance suffices to make a visit to v_3 unavoidable.

⁶ For both of these conditions, monotonicity arguments allow to transform finite-state optimally resilient strategies into positional ones (similar to the construction in [13, Section 5]).




Fig. 5: Taking the upper path from v_0 allows Player 0 to minimize visits to odd colors, while taking the lower path allows her to minimize the occurrence of disturbances.

On the other hand, consider a setting in which it is Player 0’s goal to avoid the occurrence of disturbances. In that case, σ' is preferable over σ , as it allows for fewer situations in which disturbances may occur, since no disturbances are possible from vertices v_2 and v_3 .

Note that the two goals of minimizing visits to vertices of odd color and minimizing the occurrence of disturbances are not contradictory: if both events are undesirable, it may be optimal for Player 0 to play according to a combination of strategies σ and σ' . In general, it is interesting to study how to **how to best brace for a finite number of disturbances**.




Fig. 6: Additional memory allows Player 0 to remain in the right-hand loop longer and longer, thus decreasing the potential for disturbances.

Recall that, due to Theorem 2, optimally resilient strategies for parity games do not require memory. In contrast, the game shown in Figure 6 demonstrates that additional memory can prove beneficial in order to further improve such strategies. Any strategy for Player 0 that does not stay in v_2 from some point onwards is optimally resilient. However, every visit to v_1 carries with it the possibility of a disturbance occurring, which would lead the play into a losing sink for Player 0. Hence, it is in her best interest to remain in vertex v_2 for as long as possible in order to minimize the possibility for disturbances to occur. This behavior does, however, require memory to implement, as Player 0 needs to track the number of visits to v_2 in order to not remain in that state ad infinitum. Thus, for each optimally resilient strategy σ with finite memory there exists another optimally resilient strategy that uses more memory, but visits v_1 more rarely than σ , reducing the possibilities for disturbances to occur. Hence, it is interesting to study **how to balance the avoidance of disturbances with the satisfaction of the winning condition**. We address all these problems in further research.

6 Related Work

The notion of *unmodeled intermittent disturbances* has recently been formulated by Dallal, Neider, and Tabuada [9]. In this work, the authors also present an algorithm for computing optimally resilient strategies for safety games with disturbances, which is an extension of the classical attractor computation [15]. Due to the relatively simple nature of such games, however, this algorithm cannot easily be extended to handle more expressive winning conditions, and the approach presented in this work relies on fundamentally different ideas.

Resilience is not a novel concept in the context of reactive systems synthesis. It appears, for instance, in the work by Topcu et al. [22] as well as Ehlers and Topcu [11]. A notion of resilience that is very similar to the one considered here has been proposed by Huang et al. [16], where the game graph is augmented with so-called “error edges”. However, this setting differs from the one studied in this work in various aspects. Firstly, Huang et al. work in the framework of concurrent games and model errors as under the control of Player 1. This is in contrast to the setting considered here, in which the players play in alternation and disturbances are seen as rare events rather than antagonistic to Player 0. Secondly, Huang et al. restrict themselves to safety games, whereas we consider a much broader class of infinite games.

Finally, Huang et al. compute resilient strategies with respect to a fixed parameter k , thus requiring to repeat the computation for various values of k to find optimal resilient strategies. By contrast, our approach computes an optimal strategy in a single run.

Related to resilience are various notions of *fault tolerance* [1,10,14] and *robustness* [2,3,4,5,19,21]. In fact, robustness in the area of reactive controller synthesis has attracted considerable interest in the recent years, typically in setting with specifications of the form $\varphi \Rightarrow \psi$ stating that the controller needs to fulfill the guarantee ψ if the environment satisfies the assumption φ . A prominent example of such work is that of Bloem et al. [2], in which the authors understand robustness as the property that “if assumptions are violated temporarily, the system is required to recover to normal operation with as few errors as possible” and consider the synthesis of robust controllers for the GR(1) fragment of Linear Temporal Logic [5]. Other examples include quantitative synthesis [3], where robustness is defined in terms of payoffs, and the synthesis of robust controllers for cyber-physical systems [19,21]. For a more in-depth discussion of related notions of resilience and robustness in reactive synthesis, we refer the interested reader to Dallal, Neider, and Tabuada’s section on related work [9, Section I]. Moreover, a survey of a large body of work dealing with robustness in reactive synthesis has been presented by Bloem et al. [4].

References

1. Attie, P.C., Arora, A., Emerson, E.A.: Synthesis of fault-tolerant concurrent programs. *ACM Trans. Program. Lang. Syst.* 26(1), 125–185 (2004)
2. Bloem, R., Chatterjee, K., Greimel, K., Henzinger, T.A., Hofferek, G., Jobstmann, B., Könighofer, B., Könighofer, R.: Synthesizing robust systems. *Acta Inf.* 51(3-4), 193–220 (2014)
3. Bloem, R., Chatterjee, K., Henzinger, T.A., Jobstmann, B.: Better quality in synthesis through quantitative objectives. In: Bouajjani, A., Maler, O. (eds.) *CAV 2009*. LNCS, vol. 5643, pp. 140–156. Springer (2009)
4. Bloem, R., Ehlers, R., Jacobs, S., Könighofer, R.: How to handle assumptions in synthesis. In: Chatterjee, K., Ehlers, R., Jha, S. (eds.) *SYNT 2014*. EPTCS, vol. 157, pp. 34–50 (2014)
5. Bloem, R., Jobstmann, B., Piterman, N., Pnueli, A., Sa’ar, Y.: Synthesis of Reactive(1) designs. *J. Comput. Syst. Sci.* 78(3), 911–938 (2012)
6. Calude, C.S., Jain, S., Khoussainov, B., Li, W., Stephan, F.: Deciding parity games in quasipolynomial time. In: Hatami, H., McKenzie, P., King, V. (eds.) *STOC 2017*. pp. 252–263. ACM (2017)
7. Chatterjee, K.: Linear time algorithm for weak parity games. *arXiv 0805.1391* (2008)
8. Chatterjee, K., Henzinger, T.A., Horn, F.: Finitary winning in ω -regular games. *ACM Trans. Comput. Log.* 11(1) (2009)
9. Dallal, E., Neider, D., Tabuada, P.: Synthesis of safety controllers robust to unmodeled intermittent disturbances. In: *CDC 2016*. pp. 7425–7430. IEEE (2016)
10. Ebzenasir, A., Kulkarni, S.S., Arora, A.: FTSyn: a framework for automatic synthesis of fault-tolerance. *STTT* 10(5), 455–471 (2008)
11. Ehlers, R., Topcu, U.: Resilience to intermittent assumption violations in reactive synthesis. In: Fränzle, M., Lygeros, J. (eds.) *HSCC 2014*. pp. 203–212. ACM (2014)
12. Fearnley, J., Jain, S., Schewe, S., Stephan, F., Wojtczak, D.: An ordered approach to solving parity games in quasi polynomial time and quasi linear space. In: Erdogmus, H., Havelund, K. (eds.) *SPIN 2017*. pp. 112–121. ACM (2017)
13. Fijalkow, N., Zimmermann, M.: Parity and street games with costs. *Logical Methods in Computer Science* 10(2) (2014)
14. Girault, A., Rutten, É.: Automating the addition of fault tolerance with discrete controller synthesis. *Formal Methods in System Design* 35(2), 190–225 (2009)
15. Grädel, E., Thomas, W., Wilke, T. (eds.): *Automata, Logics, and Infinite Games: A Guide to Current Research*, LNCS, vol. 2500. Springer (2002)
16. Huang, C., Peled, D.A., Schewe, S., Wang, F.: A game-theoretic foundation for the maximum software resilience against dense errors. *IEEE Trans. Software Eng.* 42(7), 605–622 (2016)
17. Jurdzinski, M., Lazic, R.: Succinct progress measures for solving parity games. In: *LICS 2017*. pp. 1–9. IEEE Computer Society (2017)
18. Kaiser, L.: *Logic and Games on Automatic Structures - Playing with Quantifiers and Decompositions*, LNCS, vol. 6810. Springer (2011)
19. Majumdar, R., Render, E., Tabuada, P.: A theory of robust omega-regular software synthesis. *ACM Trans. Embedded Comput. Syst.* 13(3), 48:1–48:27 (2013)
20. Martin, D.A.: Borel determinacy. *Annals of Mathematics* 102, 363–371 (1975)
21. Tabuada, P., Caliskan, S.Y., Rungger, M., Majumdar, R.: Towards robustness for cyber-physical systems. *IEEE Trans. Automat. Contr.* 59(12), 3151–3163 (2014)

22. Topcu, U., Ozay, N., Liu, J., Murray, R.M.: On synthesizing robust discrete controllers under modeling uncertainty. In: Dang, T., Mitchell, I.M. (eds.) HSCC 2012. pp. 85–94. ACM (2012)