A Cure for Stuttering Parity Games

Sjoerd Cranen, Jeroen J.A. Keiren, and Tim A.C. Willemse

Department of Computer Science and Mathematics
Eindhoven University of Technology
PO Box 513, 5600MB Eindhoven, The Netherlands

Abstract. We define governed stuttering bisimulation for parity games, weakening stuttering bisimulation by taking the ownership of vertices into account only when this might lead to observably different games. We show that governed stuttering bisimilarity is an equivalence for parity games and allows for a natural quotienting operation. Moreover, we prove that all pairs of vertices related by governed stuttering bisimilarity are won by the same player in the parity game. Thus, our equivalence can be used as a preprocessing step when solving parity games. Governed stuttering bisimilarity can be decided in $\mathcal{O}(n^2m)$ time for parity games with n vertices and m edges. Our experiments indicate that governed stuttering bisimilarity is mostly competitive with stuttering equivalence on parity games encoding typical verification problems.

1 Introduction

Parity games [11,22,27] are played by two players (represented by $\diamond$ and $\square$) on a directed graph in which every vertex is owned by one of the players, and vertices are assigned a priority. The game is played by moving a single token along the edges in the graph; the choice where to move next is dictated by the player owning the vertex on which the token currently resides. Both players try to play such that the resulting infinite path is *winning* for them, and a vertex is won by the player that can play such that, however the opponent plays, every path from that vertex is won by her. The winner of a vertex is uniquely determined [11,22,27] and partitioning the graph in vertices that are won by player $\diamond$ and those won by player $\square$ is referred to as *solving* the parity game.

The parity game framework is a key instrument in solving practical verification and synthesis problems, see [11,2]. Its practical significance is mirrored by its role in searching for the true complexity of model checking: modal μ-calculus model checking is polynomially reducible to parity game solving, and *vice versa* [25]. Despite the apparent simplicity of the latter problem, the precise complexity of solving parity games is still open: the problem is known to be in NP $\cap$ coNP, and more specifically in UP $\cap$ coUP [17], suggesting there just might exist a polynomial time algorithm. Indeed, non-trivial classes of parity games have been identified that admit polynomial time solving algorithms, see *e.g.* [4,23].

In the past decade, several advanced algorithms for solving parity games have been designed. These include algorithms exponential in the number of priorities, such as Jurdziński's *small progress measures* algorithm [18] and Schewe's *bigstep* algorithm [24], as well as the sub-exponential algorithm due to Jurdziński *et al.* [19]. Orthogonally to

the algorithmic improvements, heuristics have been devised that may speed up solving parity games that occur in practice [12]. Such heuristics work particularly well for verification problems, which give rise to games with only few different priorities.

The heuristic that we consider in this paper, following, *e.g.*, Fritz and Wilke's study of *delayed simulation* [14], is based on the use of fine-grained equivalence relations that approximate the solution to a parity game. The idea is to recast the solving problem as the problem of deciding *winner equivalence* between vertices: two vertices in a parity game are equivalent whenever they are won by the same player. Finding equivalence relations that refine winner equivalence and that are decidable in polynomial time yields a preprocessing step that can be used to reduce games prior to solving.

From a practical viewpoint, we are particularly interested in those simulation and equivalence relations that strike a favourable balance between their power to compress the game graph and their computational complexity. Stuttering bisimulation [7] for Kripke Structures is among a select number of candidates worth considering, with an $\mathcal{O}(nm)$ time complexity (n being the number of vertices and m the number of edges). Observe that stuttering bisimulation only preserves a fragment of the μ-calculus when applied to Kripke Structures. It may therefore be surprising that it does preserve the winner of parity games, including those that stem from encodings of arbitrary μ-calculus model checking problems. As earlier experiments [10] indicate, off-the-shelf stuttering bisimulation reduction algorithms can be competitive when compared to modern available parity game solvers. Stuttering bisimulation, however, is inadequate when faced with alternations between players along the possible plays: it cannot relate vertices belonging to different players. Controller synthesis problems *e.g.* [2], and constructs such as $\Box\Diamond\phi$ and $\Diamond\Box\phi$ in μ-calculus verification, give rise to such parity games.

A natural question is, therefore, whether stuttering bisimulation can at all be modified so that it is able to relate vertices that belong to different players. We answer this question in this paper by defining a relation, which we dub *governed stuttering bisimulation* (reflecting that a player's ruling capabilities are taken as primitive), which we show to be strictly weaker than stuttering bisimilarity. In addition, we prove that governed stuttering bisimilarity:

- is an equivalence relation on parity games.
- refines winner equivalence.
- is decidable in $\mathcal{O}(n^2m)$ time using a partition refinement algorithm.

The time complexity for deciding governed stuttering bisimilarity is a factor n worse than that for stuttering bisimilarity; this is due to finding a splitter, for which our algorithm requires $\mathcal{O}(mn)$ rather than $\mathcal{O}(m)$ time. Our experiments, however, indicate that this factor does not manifest itself in practice; in fact, our algorithm is mostly competitive with the one for stuttering bisimilarity.

Structure of the paper. Section 2 briefly introduces the parity game framework. We recall the definition of stuttering bisimulation and we define governed stuttering bisimulation in Section 3. In Section 4, we show that governed stuttering bisimulation is an equivalence relation, we show it refines winner equivalence, and we address its decidability. We discuss our experiments with a prototype implementation of our algorithm for deciding governed stuttering bisimulation in Section 5. Related work is discussed in

Section 6, and future work is described in Section 7. Note that, due to space restrictions, details of the proofs have been omitted. Detailed proofs are provided in [9].

2   Preliminaries

A parity game is a two-player graph game, played by two players on a directed graph. The game is formally defined as follows.

Definition 1 (Parity game). *A parity game is a directed graph* $(V, \rightarrow, \Omega, \mathcal{P})$, *where*

- *V is a finite set of vertices,*
- *$\rightarrow \subseteq V \times V$ is a total edge relation* (i.e., *for each $v \in V$ there is at least one $w \in V$ such that $(v, w) \in \rightarrow$),*
- *$\Omega : V \rightarrow \mathbb{N}$ is a priority function that assigns priorities to vertices,*
- *$\mathcal{P} : V \rightarrow \{\diamond, \square\}$ is a function assigning vertices to players.*

If i is a player, then $\neg i$ denotes the opponent of i, *i.e.*, $\neg \diamond = \square$ and $\neg \square = \diamond$. A sequence of vertices $v_1, \dots, v_n$ for which $v_m \rightarrow v_{m+1}$ for all $1 \leq m < n$ is a *path*, and may be denoted using angular brackets: $\langle v_1 \dots v_n \rangle$. The concatenation $p \cdot q$ of paths p and q is again a path. Infinite paths are defined in a similar manner. We use p_n to denote the n^{th} vertex in a path p.

A game starts by placing a token on vertex $v \in V$. Players move the token indefinitely according to the following simple rule: if the token is on some vertex v, player $\mathcal{P}(v)$ moves the token to some vertex w such that $v \rightarrow w$. The result is an infinite path p in the game graph. The *parity* of the lowest priority that occurs infinitely often on p defines the *winner* of the path. If this priority is even, then player $\diamond$ wins, otherwise player $\square$ wins.

A *strategy* for player i is a partial function $\sigma : V^* \rightarrow V$, that is defined only for paths ending in a vertex owned by player i and determines the next vertex to be played onto. The set of strategies for player i in a game $\mathcal{G}$ is denoted $\mathbb{S}^*_{\mathcal{G},i}$, or simply $\mathbb{S}^*_i$ if $\mathcal{G}$ is clear from the context. If a strategy yields the same vertex for every pair of paths that end in the same vertex, then the strategy is said to be *memoryless*. The set of memoryless strategies for player i in a game $\mathcal{G}$ is denoted $\mathbb{S}_{\mathcal{G},i}$, abbreviated to $\mathbb{S}_i$ when $\mathcal{G}$ is clear from the context. A memoryless strategy is usually given as a partial function $\sigma : V \rightarrow V$.

A path p of length n is *consistent* with a strategy $\sigma \in \mathbb{S}^*_i$, denoted $\sigma \Vdash p$, if and only if for all $1 \leq j < n$ it is the case that if σ is defined for $\langle p_1 \dots p_j \rangle$, then $p_{j+1} = \sigma(\langle p_1 \dots p_j \rangle)$. The definition of consistency is extended to infinite paths in the obvious manner. A strategy $\sigma \in \mathbb{S}^*_i$ is said to be a *winning strategy* from a vertex v if and only if i is the winner of every path consistent with σ. A vertex is won by i if i has a winning strategy from that vertex. Parity games are memoryless determined [11], *i.e.* each vertex in the game is won by exactly one player, and it suffices to play a memoryless strategy.

In this paper, we are concerned with relations partitioning the vertices in a parity game such that all related vertices are won by the same player. Let R be a relation over a set V. For $v, w \in V$ we write $v \ R \ w$ for $(v, w) \in R$. For an equivalence relation R, and vertex $v \in V$ we define $[v]_R$, the equivalence class of v under R, as $\{v' \in V \mid v \ R \ v'\}$. The set of equivalence classes of V under R is denoted $V_{/R}$. A

collection $\{B_i \mid i \in I\}$, with $\emptyset \neq B_i \subseteq V$, is called a *partition* of V if $\bigcup_{i \in I} B_i = V$ and for $i \neq j : B_i \cap B_j = \emptyset$. An element B_i of a partition is called a *block*. If P and Q are partitions of V then Q *refines* P if $\forall B_i \in Q : \exists B_j \in P : B_i \subseteq B_j$. We use the notions of equivalence relation and partition interchangeably, and occasionally write $v \, P \, v'$ rather than $v, v' \in B$ for some $B \in P$.

Determinacy of parity games effectively induces a partition on the set of vertices V in those vertices won by player $\diamond$ and those vertices won by player $\square$. This partition is the natural equivalence relation on V.

Definition 2 (Winner equivalence). *Let* $(V, \rightarrow, \Omega, \mathcal{P})$ *be a parity game. Vertices* $v, w \in V$ *are* winner equivalent*, denoted* $v \sim w$ *iff* v *and* w *are won by the same player.*

3   Governed Stuttering Bisimulation

In [10] we introduced *stuttering bisimulation* for parity games. Informally, stuttering bisimulation compresses subsequences of "identical" vertices, *i.e.* vertices with the same priority, owned by the same player, along a path p in a parity game, such that the path retains the essentials of the graph's branching structure.

Before we give the formal definition of stuttering bisimulation, we first introduce some notation. Let $(V, \rightarrow, \Omega, \mathcal{P})$ be a parity game. In the following, let $U \subseteq V$ be arbitrary sets of vertices; we write $v \rightarrow U$ if there exists a $u \in U$ such that $v \rightarrow u$.

Let $R \subseteq V \times V$ be a relation on the set of vertices. The generalised transition relation $v \mapsto_R U$, defined below, formalises that U is eventually reached from v by some computation path through R-related nodes. Likewise, $v \mapsto_R$ expresses that v is the start of an infinite computation path along vertices related through R.

$$v \mapsto_R U \stackrel{\mu}{=} \exists u : v \rightarrow u \wedge (u \in U \vee (v \, R \, u \wedge u \mapsto_R U))$$
$$v \mapsto_R \stackrel{\nu}{=} \exists u : v \rightarrow u \wedge v \, R \, u \wedge u \mapsto_R$$

We next formalise the notion of stuttering bisimulation, deviating notationally from [10]; the definitions, however, are easily seen to coincide and the modifications are standard. Our main reason for deviating from [10] is that the presented definition facilitates explaining the intuition of its generalisation to governed stuttering bisimulation.

Definition 3 (Stuttering bisimulation for parity games [10]). *Let* $(V, \rightarrow, \Omega, \mathcal{P})$ *be a parity game. Let* $R \subseteq V \times V$ *be an equivalence relation on vertices;* R *is a* stuttering bisimulation *if* $v \, R \, v'$ *implies*

a) $\Omega(v) = \Omega(v')$ *and* $\mathcal{P}(v) = \mathcal{P}(v')$;
b) $v \rightarrow \mathcal{C}$ *implies* $v' \mapsto_R \mathcal{C}$, *for all* $\mathcal{C} \in V_{/R} \setminus \{[v]_R\}$.
c) $v \mapsto_R$ *iff* $v' \mapsto_R$;

Two states v *and* v' *are said to be* stuttering bisimilar*, denoted* $v \simeq v'$ *iff there is a stuttering bisimulation relation* R*, such that* $v \, R \, v'$.

Our objective is to weaken stuttering bisimulation so that it will be able to relate vertices of different players. However, we cannot simply weaken clause *a)* to $\Omega(v) = \Omega(v')$ without modifying the remaining clauses, as this would enable us to relate vertices won by different players, as the following parity game demonstrates:
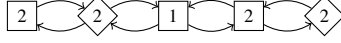
The suggested weakening of clause *a)* would allow us to relate all vertices with priority 2; the two left vertices, however are won by player $\diamond$, whereas the other vertices are won by player $\square$.

The problem in the above example is that the computation paths that appear in clauses *b)* and *c)* may consist of vertices owned by different players. This means that a fixed player is at the mercy of her opponent to stay on a computation path: the opponent may simply choose an alternative next vertex if that would better suit her. We are therefore forced to reason about computation trees, taking all the opponent's choices into account. Effectively, clause *b)* must be strengthened to ensure that a player eventually reaches class $\mathcal{C}$ along some computation tree, and clause *c)* must be strengthened to ensure that a player can construct an infinite computation tree not leaving its own class.

We first extend our notation to facilitate reasoning about computation trees rather than computation paths. Given a memoryless strategy σ for some player, the ability to move from vertex v to another vertex u depends on this strategy.

$$
v \,_\sigma\!\to u = \begin{cases} v \to u \wedge \sigma(v) = u, & \text{if } \sigma(v) \text{ is defined} \\ v \to u, & \text{otherwise} \end{cases}
$$

From the viewpoint of a fixed player and her memoryless strategy σ, a token may be moved along the edges of a computation tree that only branches at vertices owned by her opponent. This notation $v \,_\sigma\!\mapsto_R U$, defined below, formalises that all plays allowed by σ eventually reach the set of vertices U immediately when they follow an edge to a vertex that is no longer related under relation R. The notation $v \,_\sigma\!\mapsto_R$ is dual; it expresses that all plays allowed by σ can reach only vertices related under R to the previous vertex in that play:

$$
v \,_\sigma\!\mapsto_R U \stackrel{\mu}{=} \forall u : v \,_\sigma\!\to u \implies u \in U \vee (v \, R \, u \wedge u \,_\sigma\!\mapsto_R U)
$$
$$
v \,_\sigma\!\mapsto_R \stackrel{\nu}{=} \forall u : v \,_\sigma\!\to u \implies v \, R \, u \wedge u \,_\sigma\!\mapsto_R
$$

If the strategy is unimportant to the purpose at hand, we abstract from the specific strategy that is used and reason only in terms of a player i having a strategy with the capability of forcing a play to a set of vertices U, and, dually, for i to be able to force the play to *diverge* within a class of R:

$$
x \,_i\!\mapsto_R U = \exists \sigma \in \mathbb{S}_i : x \,_\sigma\!\mapsto_R U
$$
$$
x \,_i\!\mapsto_R = \exists \sigma \in \mathbb{S}_i : x \,_\sigma\!\mapsto_R
$$

We omit R if it is the relation $V \times V$. Note that $v \,_i\!\mapsto_R \emptyset$ never holds. On the other hand, $v \,_i\!\mapsto_R V$ and $v \,_i\!\mapsto$ are trivially true. We write $v \,_i\!\not\mapsto_R U$ for $\neg(v \,_i\!\mapsto_R U)$; likewise for all other arrows. If $\mathcal{U} \subseteq V_{/R}$, then we write $v \,_i\!\mapsto_R \mathcal{U}$ to denote $v \,_i\!\mapsto_R \bigcup_{\mathcal{C} \in \mathcal{U}} \mathcal{C}$.

Definition 4 (Governed stuttering bisimulation). *Let* $(V, \to, \Omega, \mathcal{P})$ *be a parity game. Let* $R \subseteq V \times V$ *be an equivalence relation. Then* R *is a* governed stuttering bisimulation *if* $v \, R \, v'$ *implies*

a) $\Omega(v) = \Omega(v')$;

b) $v \to \mathcal{C}$ implies $v' \, _{\mathcal{P}(v)}\!\mapsto_R \mathcal{C}$, for all $\mathcal{C} \in V_{/R} \setminus \{[v]_R\}$.

c) $v \, _i\!\mapsto_R$ iff $v' \, _i\!\mapsto_R$ for $i \in \{\Diamond, \Box\}$.

Vertices v and v' are governed stuttering bisimilar, *denoted $v \approxeq v'$, iff a governed stuttering bisimulation R exists such that $v \, R \, v'$.*

If we additionally require that $\mathcal{P}(v) = \mathcal{P}(v')$, we find that $v \mapsto_R U$ iff $v \, _{\mathcal{P}(v)}\!\mapsto_R U$, and, likewise, $v \mapsto_R$ iff $v \, _{\mathcal{P}(v)}\!\mapsto_R$. This is the basis for the following proposition.

Proposition 1. *Let $R \subseteq V \times V$ be a governed stuttering bisimulation, such that $v \, R \, v'$ implies $\mathcal{P}(v) = \mathcal{P}(v')$. Then R is a stuttering bisimulation.*

Example 1. Consider the parity game depicted in Figure 1a. The equivalence relation that relates vertices with equal priorities is a governed stuttering bisimulation. Stuttering bisimulation does not relate any of the vertices.
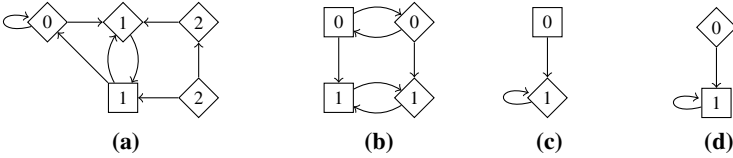

Fig. 1. All vertices in (a) with the same priorities can be related using governed stuttering bisimilarity. Both (c) and (d) are minimal representations of (b).

4   Properties of Governed Stuttering Bisimulation

We next study three key properties of governed stuttering bisimulation, *viz.*, governed stuttering bisimilarity is an equivalence on parity games, it refines winner equivalence and it is decidable in polynomial time.

4.1   Governed Stuttering Bisimilarity is an Equivalence

Proving that $\approxeq$ is an equivalence relation on parity games is far from straightforward: transitivity no longer bows to the standard proof strategies that work for stuttering bisimilarity and branching bisimilarity [26]. As a result of the asymmetry in the use of two different transition relations in clause *b)* of Definition 4, proving that the equivalence closure of the union of two governed stuttering bisimulation relations is again a governed stuttering bisimulation relation is equally problematic.

The strategy we pursue is as follows. We characterise governed stuttering bisimulation, in two steps, by a set of symmetric requirements. The obtained alternative characterisation is then used in our equivalence proof. These alternative characterisations do not facilitate the reuse of standard proof strategies, but they are instrumental in the technically involved proof that the equivalence closure of two governed stuttering bisimulation relations is again a governed stuttering bisimulation relation. Apart from being

convenient technically, the characterisations offer more insight into the nature of governed stuttering bisimilarity. Hence, instead of providing the details of our equivalence proof, we focus on the alternative characterisations of governed stuttering bisimulation.

Our result below states that we can rephrase condition *b)* of governed stuttering bisimulation by requiring that a fixed player must have the same power to force the play from any pair of related vertices to reach an arbitrary class. Thus, we abstract from the player that takes the initiative to leave its class in one step.

Theorem 1. *Let $R \subseteq V \times V$ and $v, v' \in V$. Then R is a governed stuttering bisimulation iff R is an equivalence relation and $v \mathrel{R} v'$ implies:*

a) $\Omega(v) = \Omega(v')$;
b) $v \mathrel{}_i\!\!\mapsto_R \mathcal{C}$ iff $v' \mathrel{}_i\!\!\mapsto_R \mathcal{C}$ for all $i \in \{\diamond, \square\}, \mathcal{C} \in V_{/R} \setminus \{[v]_R\}$;
c) $v \mathrel{}_i\!\!\mapsto_R$ iff $v' \mathrel{}_i\!\!\mapsto_R$ for all $i \in \{\diamond, \square\}$.

While the above alternative characterisation of governed stuttering bisimulation is now fully symmetric, the restriction on the class $\mathcal{C}$ that is considered in clause *b)* turns out to be too strong to facilitate our proof that $\approx$ is an equivalence relation. We therefore generalise this clause once more to reason about sets of classes. A perhaps surprising side-result of this generalisation is that the divergence requirement of clause *c)* becomes superfluous. Note that this last generalisation is not trivial, as $v \mathrel{}_i\!\!\mapsto_R \{\mathcal{C}_1, \mathcal{C}_2\}$ is in general neither equivalent to saying that $v \mathrel{}_i\!\!\mapsto_R \mathcal{C}_1$ and $v \mathrel{}_i\!\!\mapsto_R \mathcal{C}_2$, nor $v \mathrel{}_i\!\!\mapsto_R \mathcal{C}_1$ or $v \mathrel{}_i\!\!\mapsto_R \mathcal{C}_2$.

Theorem 2. *Let $R \subseteq V \times V$ and $v, v' \in V$. Then R is a governed stuttering bisimulation iff R is an equivalence relation and $v \mathrel{R} v'$ implies:*

a) $\Omega(v) = \Omega(v')$;
b) $v \mathrel{}_i\!\!\mapsto_R \mathcal{U}$ iff $v' \mathrel{}_i\!\!\mapsto_R \mathcal{U}$ for all $i \in \{\diamond, \square\}, \mathcal{U} \subseteq V_{/R} \setminus \{[v]_R\}$.

Note that the divergence requirement $v \mathrel{}_i\!\!\mapsto_R$ iff $v' \mathrel{}_i\!\!\mapsto_R$ can be recovered by instantiating set $\mathcal{U}$ by $V_{/R} \setminus \{[v]_R\}$ for player $\neg i$ in the above theorem. The last characterisation enables us to prove the following theorem.

Theorem 3. $\approx$ *is an equivalence relation.*

As a side-result of the proof of Theorem 3, we find that the *equivalence closure* of the union of two governed stuttering bisimulations is again a governed stuttering bisimulation. The union of *all* governed stuttering bisimulations is again a governed stuttering bisimulation, which coincides with governed stuttering bisimilarity.

4.2   Quotienting

The main reason for studying equivalence relations for parity games is that they may offer the prospect of minimising the parity game by collapsing vertices that are considered equivalent. The resulting minimised structure is referred to as the quotient. However, not all equivalence relations admit such a quotienting operation; in particular, the delayed simulation [14] for parity games fails to have a natural quotienting operation.

Quotienting for governed stuttering bisimulation can be done efficiently. Due to the nature of governed stuttering bisimulation, we have some freedom in the definition of the quotient, in particular when assigning vertices to players. We therefore first define a notion of minimality, and we subsequently define the quotient in terms of that notion.

Definition 5 (Minimality). *A $\bumpeq$-minimal representation of a parity game $(V, \to, \Omega, \mathcal{P})$ is defined as a game $(V_m, \to_m, \Omega_m, \mathcal{P}_m)$, that satisfies the following conditions (where $c, c', c'' \in V_m$):*

$$V_m = \{ [v]_{\bumpeq} \mid v \in V \}$$
$$\Omega_m(c) = \Omega(v) \text{ for all } v \in c$$
$$\mathcal{P}_m(c) = i, \text{ if for all } v \in c, \text{ and some } c' \neq c \text{ we have } v \,{}_i{\mapsto}_{\bumpeq} c' \text{ and } v \,{}_{\neg i}{\not\mapsto}_{\bumpeq} V \setminus c'$$
$$c \to_m c \text{ iff } v \,{}_i{\mapsto}_{\bumpeq} \text{ for all } v \in c \text{ for some player } i$$
$$c \to_m c' \text{ iff } v \,{}_i{\mapsto}_{\bumpeq} c' \text{ for all } v \in c \text{ for some player } i \text{ and } c' \neq c$$

Observe that for the third clause above, if from some vertex v the play could be forced to c' by i without $\neg i$ having the opportunity to diverge, player i is in charge of the game when the play arrives in c. This requires the representative in the quotient to be owned by player i.

Note that a parity game may have multiple $\bumpeq$-minimal representations. It is not hard to verify that every parity game contains at least as many vertices and edges as its $\bumpeq$-minimal representations. Moreover, any parity game is governed stuttering bisimulation equivalent to all its $\bumpeq$-minimal representations. As a result, the governed stuttering bisimulation quotient of a graph can be defined as its least $\bumpeq$-minimal representation, given some arbitrary ordering on parity games. A natural ordering would be one that is induced by an ordering on players, *e.g.*, $\square < \diamond$.

Example 2. Consider the parity game in Figure 1b. Two of its four minimal representations are in Figure 1c and 1d. Observe that the particular player chosen for the 0 and 1 vertices is arbitrary and does not impact the solution to the games.

4.3 Governed Stuttering Bisimilarity Refines Winner Equivalence

In this section, we prove that governed stuttering bisimilarity is strictly finer than winner equivalence. That is, vertices that are won by different players are never related by governed stuttering bisimilarity. In order to prove this result, we must first lift the concept of governed stuttering bisimilarity to paths.

Paths of length 1 are equivalent if the vertices they consist of are equivalent. If paths p and q are equivalent, then $p \cdot \langle v \rangle \simeq q$ iff v is equivalent to the last vertex in q, and $p \cdot \langle v \rangle \simeq q \cdot \langle w \rangle$ iff $v \simeq w$. An infinite path p is equivalent to a path q if for all finite prefixes of p there is an equivalent prefix of q and *vice versa*.

We define $\Pi_\varphi^n(v)$ to be the set of paths of length n that start in v and that are allowed by some strategy φ. $\Pi_\varphi^\omega(v)$ is then the set of all infinite paths allowed by φ, starting in v. In a similar fashion, we also define $\Psi_\varphi^n(v)$, which contains those paths starting in v that are allowed by φ and that consist of exactly n segments in which all vertices in a segment are related by $\bumpeq$, except the last vertex. Also included in $\Psi_\varphi^n(v)$ are infinite paths that stay in the same class forever after n or less such segments.

Definition 6 (Levels). *Formally we define the nth* level *paths* $\Psi_\varphi^n(v)$ *of a strategy* φ *from root vertex* v *for all paths* p *as follows:*

$$p \in \Psi_\varphi^0(v) \text{ iff } p = \langle v \rangle$$
$$p \cdot q \in \Psi_\varphi^{n+1}(v) \text{ iff } p \in \Psi_\varphi^n(v) \ \wedge \varphi \Vdash p \cdot q \ \wedge$$
$$((p \cdot q \approx p \wedge |q| = \infty) \ \vee$$
$$(\exists \bar{q}, v : q = \bar{q} \cdot \langle v \rangle \wedge p \cdot \bar{q} \approx p \wedge p \cdot q \napprox p))$$

Note that $\Pi_\varphi^\omega(v) = \Psi_\varphi^\omega(v)$. The following lemma is the basis for establishing that governed stuttering bisimilarity refines winner equivalence.

Lemma 1. *Given some* $v, w \in V$ *such that* $v \approx w$, *then for every strategy* $\varphi \in \mathbb{S}_i$ *we have a strategy* $\psi \in \mathbb{S}_i^*$ *such that* $\forall n \in \mathbb{N} : \forall p \in \Psi_\psi^n(w) : \exists p' \in \Psi_\varphi^n(v) : p \approx p'$

We are now in a position to prove that governed stuttering bisimilarity refines winner equivalence.

Theorem 4. *Governed stuttering bisimulation strictly refines winner equivalence.*

Proof. Let φ be a strategy for player i that wins from $v \in V$. Without loss of generality assume that φ is memoryless, and let $w \in V$ such that $v \approx w$. By Lemma 1, we know that there is some strategy ψ such that for every path in $\Psi_\psi^\omega(w)$ there is a related path in $\Psi_\varphi^\omega(v)$. As $\Pi_\varphi^\omega(v) = \Psi_\varphi^\omega(v)$, this means that for every path starting in w that is allowed by ψ, we have an equivalent path starting in v that is allowed by φ. Equivalent paths have the same set of infinitely often recurring priorities. Any priority that may be visited infinitely often under ψ could therefore also have been visited infinitely often under φ. Therefore, ψ must be a winning strategy. The strictness of the refinement follows from, *e.g.*, the example in Figure 1.c, in which player $\square$ wins both vertices. $\square$

4.4 Decidability

Our algorithm for deciding governed stuttering bisimilarity is based on Groote and Vaandrager's $\mathcal{O}(nm)$ algorithm for deciding stuttering bisimilarity [16]. Before we provide the details, we introduce the necessary additional concepts.

Our algorithm requires a generalisation of the well-known notion of *attractor sets* [22] along the lines of the generalisation used for the computation of the *Until* in the alternating-time temporal logic ATL [1]. The generalisation introduces a parameter restricting the set of vertices that are considered in the attractor sets.

$$\begin{aligned}
_B Attr_i^0(U) &= U \\
_B Attr_i^{n+1}(U) &= {_B Attr_i^n(U)} \\
&\quad \cup \{v \in B \mid \mathcal{P}(v) = i \wedge \exists v \rightarrow v' : v' \in {_B Attr_i^n(U)}\} \\
&\quad \cup \{v \in B \mid \mathcal{P}(v) \neq i \wedge \forall v \rightarrow v' : v' \in {_B Attr_i^n(U)}\} \\
_B Attr_i(U) &= {_B Attr_i^\omega(U)} \\
Leave_i(B, W) &= {_B Attr_i(W)} \cap B
\end{aligned}$$

The set $Leave_i(B, W)$ captures the subset of B from which player i can force the game to $W \subseteq V$. The formal correspondence between *Leave* and $_i\!\mapsto$ is formalised below; this allows for restating the criteria from Definition 4 in terms of *Leave*.

Lemma 2. *Let P be a partition of V, and let $B \in P$ be a block. Then for all $u \in B$: $u \ _i{\mapsto}_P$ if and only if $u \notin Leave_{\neg i}(B, V \setminus B)$.*

Lemma 3. *Let P partition V, and let $B, B' \in P$ such that $B \neq B'$. Let $v \in B$ such that $v \to B'$. Then for all $w \in B$ it holds that $w \ _{\mathcal{P}(v)}{\mapsto}_P B'$ if and only if $w \in Leave_{\mathcal{P}(v)}(B, B')$.*

Groote and Vaandrager's algorithm for stuttering bisimulation repeatedly refines a carefully chosen initial partition P_0 using a so-called *splitter*. We apply the same principle, choosing P_0 such that for all $v, v' \in V$, $v \ P_0 \ v'$ if and only if $\Omega(v) = \Omega(v')$ as our initial partition. As our splitter, we define a function *pos* that returns the set of vertices in B from which a given player i can force the play to reach B', or, in case $B = B'$, force the play to diverge:

$$
pos_i(B, B') = \begin{cases} \{v \in B \mid v \ _i{\mapsto}_P\} & \text{if } B = B' \\ \{v \in B \mid v \ _i{\mapsto}_P B'\} & \text{if } B \neq B' \end{cases}
$$

In line with [16], we say that B' is a *splitter* of B if and only if $\emptyset \neq pos_i(B, B') \neq B$ for some player i. A partition P is *stable with respect to a block $B \in P$* if B is not a splitter of any block in P. The partition itself is stable if it is stable with respect to all its blocks.   A high-level description of our algorithm for governed stuttering bisimulation,

---

Algorithm 1. Decision procedure for $\approx$

$n \leftarrow 0$
repeat
  $splitter \leftarrow \bot$
  for each $B \in P_n$ and player i **do** { Find splitter in $\mathcal{O}(nm)$ }
    if there exists $v \in B$ with $v \to B'$ and $\emptyset \neq pos_i(B, B') \neq B$ for $B' \in P_n$ **then**
      $splitter \leftarrow (B, pos_i(B, B'))$
    end if
  end for
  if $splitter = (B, Pos)$ **then** { Refine partition in $\mathcal{O}(m)$ }
    $P_{n+1} \leftarrow (P_n \setminus \{B\}) \cup \{Pos, B \setminus Pos\}$
  end if
  $n \leftarrow n + 1$
until $P_{n-1} = P_n$

---

is given as Algorithm 1. Note that this does not compute the quotient. Correctness of the algorithm follows the same line of reasoning as in [16]. Based on this algorithm, we obtain the following complexity result.

Theorem 5. *Algorithm 1 decides $\approx$ in $\mathcal{O}(n^2 m)$ time for a parity game that contains n vertices and m edges.*

Our time complexity is worse than the $\mathcal{O}(nm)$ achieved by the original algorithm for deciding stuttering bisimulation. The extra factor $\mathcal{O}(n)$ is due to the complexity required to search for a splitter which, in our case, requires $\mathcal{O}(nm)$ time, instead of the original $\mathcal{O}(m)$ time.

5   Experiments

While the running time of our algorithm for governed stuttering bisimilarity is theoretically worse than that of the algorithm for stuttering bisimilarity, we expect that for solving parity games, in practice both are comparable. We test this hypothesis on a set of over 200 real-life model checking problems, part of which was previously used to study the effect of stuttering bisimilarity for parity games, see [10].

Whereas in [10] a signature based approach [5] was used, in the present paper we use the Groote-Vaandrager algorithm for computing stuttering bisimilarity in order to answer our hypothesis. For computing governed stuttering bisimilarity we have modified the implementation of Groote-Vaandrager to include the changes presented in Algorithm 1.

For running our experiments we reuse the setup of [10] for solving parity games, *i.e.*, we use an optimised C++ implementation of the *small progress measures* algorithm [18], and the optimised variants of the small progress measures, recursive [22,27] and *bigstep* algorithms [24] offered by the PGSolver [12] toolset.

All experiments were conducted on a machine consisting of 28 Intel® Xeon© E5520 Processors running at 2.27GHz, with 1TB of shared main memory, running a 64-bit Linux distribution using kernel version 2.6.27. None of our experiments employ multi-core features.

5.1   Test Sets

The parity games that were used for our experiments are clustered into four test sets. We give a brief description of each of the sets below.

Model checking. Our main interest is in the practical implications of governed stuttering bisimilarity reduction on solving model checking problems. To this end, a number of model checking problems have been selected from the literature [21,3,15]. The properties that have been checked include fairness, liveness and safety properties.

Games. The second test set considers a number of turn-based, two player board games. For each of these games, and for each player, we have encoded the property that said player can play the game in such a way that, regardless of the play of the opponent, she can win the game.

PGSolver. The third test set was taken from [12] and consists of the elevator problem and the Hanoi towers problem described in that paper. It also includes alternative encodings of these problems, taken from [10].

Equivalence checking. The last test set consists of a number of equivalence checking problems encoded into parity games as described in [8].

The problems in these test sets are scalable. In every test set, a number of instances of every problem is included. Each problem gives rise to a parity game with at most 4 different priorities, which is typical for practical verification problems.

The model checking, PGSolver and equivalence checking problems were studied before in the setting of stuttering bisimilarity [10]. We extended that test set to include more examples of parity games with alternations between players and priorities. We can expect improved reductions compared to stuttering bisimilarity in these cases.

## 5.2   Measurements: Size and Time

To analyse the performance of our reduction, we measured the difference in the sizes (computed as the sum of the number of vertices and the number of edges) of the stuttering and governed stuttering minimal games. A reduction of 0% means that the governed stuttering bisimilarity reduced game has the same size as the stuttering bisimilarity reduced game.

For every problem in the test set, we compute the reduction as the average reduction over all instances of that problem. We do this in order to measure the reduction rate for the different problems, rather than for the instances. Figure 2a shows the average reduction for problems in each test set, together with the minimal and maximal reduction achieved within that set.

In addition, we measured the times needed to reduce the parity games plus the time needed to solve the reduced game using the fastest solver. That is, the sum of these two is the *total solving time* for a parity game. This way, our results can be compared to those listed in [10]. In Figure 2b, every data point represents a problem instance, of which the total solving time of the stuttering minimal game determines the x-coordinate, and the total solving time for the governed stuttering bisimilarity minimal game determines the y-coordinate.
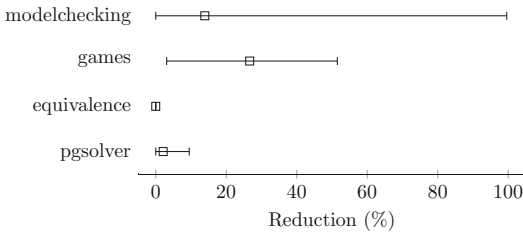
(a) Minimum, maximum and average reduction of parity games using governed stuttering bisimulation reduction, as percentage of the size after stuttering bisimilarity reduction
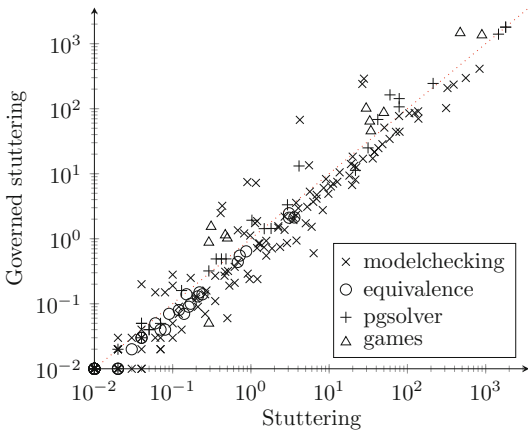
(b) Solving times (=sum of reduction time and subsequent solving) in seconds using governed stuttering bisimilarity set out against the solving times using stuttering bisimilarity. The dotted line is defined as $x = y$ and serves as a reference. Note that the axes are in log scale.

Fig. 2. Comparison of sizes and times of reductions

5.3   Discussion

Figure 2b shows that solving times for governed stuttering bisimilarity are generally comparable to those for stuttering bisimilarity, confirming our hypothesis.

Whether governed stuttering bisimilarity offers additional reductions over stuttering bisimilarity depends largely on the kind of property that is checked, and the resulting structure of the parity game. On average, a modest additional reduction is achieved, and there are practical cases in which the additional reduction is almost 100%.

For several model checking cases, stuttering bisimilarity already reduces the parity game to a graph with one vertex per priority. Obviously, governed stuttering bisimilarity cannot improve on that. However, in one of our problems (model checking a leadership protocol) a reduction of almost 100% is achieved, increasing the average reduction for this test set.

The properties that we considered on two-player games naturally give rise to alternations between players in the parity game. For these type of properties, the reduction achieved using governed stuttering bisimilarity surpasses that of stuttering bisimilarity by about 20% on average. Similar results for parity games obtained for controller synthesis (see *e.g.* [2]) may be obtained as these exhibit similar structures.

For the equivalence cases, stuttering bisimilarity reduction already yields games of a small size and governed stuttering bisimilarity does not reduce any further.[1]

Interestingly, one of the PGSolver cases taken from [12] shows a better reduction using governed stuttering bisimilarity, in contrast to an alternative encoding also used in [10].

Summarising, we conclude that governed stuttering bisimilarity reduces slightly better than stuttering bisimilarity, without noticable loss of performance.

6   Related Work

As observed in Fritz' thesis [13], *direct simulation* for parity games led to disappointing reductions, spurring Fritz and Wilke to investigate a weaker notion, called *delayed simulation* [14] and its induced equivalence. Delayed simulation equivalence is incomparable to governed stuttering bisimilarity. Contrary to governed stuttering bisimilarity, delayed simulation equivalence has the capability to relate vertices with different priorities. On the other hand, governed stuttering bisimilarity can relate vertices with the same priority in cases that delayed simulation equivalence cannot, as illustrated by the two parity games below, in which governed stuttering bisimulation relates all vertices with equal priority whereas delayed simulation equivalence does not:

Contrary to governed stuttering bisimulation, the definition of the simulation relation is entirely in terms of a *simulation game*, *viz.*, a game graph equipped with Büchi winning conditions. The simulation game gives rise to an $\mathcal{O}(d^2 n^3 m)$ algorithm for deciding

---

[1] [10] reports a poor reduction for stuttering equivalence in these cases. This was caused by "optimisations" that were used during generation of the parity games.

delayed simulation (here, n is the number of vertices, m the number of edges, and d the number of different priorities in the game), significantly exceeding our $\mathcal{O}(n^2m)$ complexity for governed stuttering bisimulation.

Apart from delayed simulation, in the setting of Boolean equation systems, the notion of *idempotence-identifying bisimilarity* was defined and investigated [20]. This equivalence relation enables one to relate conjunctive equations to disjunctive equations. In parity games, this translates to being able to relate □ vertices and ◇ vertices, respectively. Idempotence-identifying bisimilarity is much finer than governed stuttering bisimilarity, as the former is based on strong bisimilarity. Interestingly, the complexity of deciding idempotence-identifying bisimilarity is the same as for strong bisimilarity.

7   Concluding Remarks

We have described a non-trivial modification of stuttering bisimulation that allows relating vertices that belong to different players. The resulting relation, dubbed *governed stuttering bisimulation*, is an equivalence relation that can be decided in $\mathcal{O}(n^2m)$ time using a partition refinement algorithm. Although this complexity is worse than the $\mathcal{O}(nm)$ time complexity for deciding stuttering bisimulation, our experiments indicate that this factor does not manifest itself in practice. In fact, the algorithm is largely competitive with the one for stuttering bisimilarity.

An obvious question is whether elements of Fritz and Wilke's delayed simulation [14] and governed stuttering bisimulation can be combined. Given the complexity of the proofs of most of our results for governed stuttering bisimulation and our attempts to weakening governed stuttering bisimulation along these lines, we are rather sceptical about the chances of success. Even if one would manage to define such a relation, it would likely have little practical significance due to the prohibitive complexity of delayed simulation.

An interesting extension of our work could be to generalise the concepts of governed stuttering bisimilarity to games with other payoff functions that are insensitive to stuttering. We expect such a generalisation to be reasonably straightforward.

Finally, we observe that stuttering bisimulation (also known as *branching bisimulation* in labelled transition systems) underlies several confluence reduction techniques for syntactic system descriptions, see [6]. Such reductions partly side-step the statespace explosion. We believe that our study offers the required foundations for bringing similar-spirited confluence reduction techniques to a setting of symbolic representations of parity games.

References

1. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. J. ACM 49(5), 672–713 (2002)
2. Arnold, A., Vincent, A., Walukiewicz, I.: Games for synthesis of controllers with partial observation. TCS 303(1), 7–34 (2003)
3. Badban, B., Fokkink, W., Groote, J.F., Pang, J., van de Pol, J.: Verification of a sliding window protocol in μCRL and PVS. FAC 17, 342–388 (2005)

4. Berwanger, D., Dawar, A., Hunter, P., Kreutzer, S.: DAG-Width and Parity Games. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 524–536. Springer, Heidelberg (2006)

5. Blom, S.C.C., Orzan, S.: Distributed branching bisimulation reduction of state spaces. Electronic Notes in Theoretical Computer Science 89(1), 99–113 (2003)

6. Blom, S., van de Pol, J.: State Space Reduction by Proving Confluence. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 596–694. Springer, Heidelberg (2002)

7. Browne, M.C., Clarke, E.M., Grumberg, O.: Characterizing finite Kripke structures in propositional temporal logic. TCS 59, 115–131 (1988)

8. Chen, T., Ploeger, B., van de Pol, J., Willemse, T.A.C.: Equivalence Checking for Infinite Systems Using Parameterized Boolean Equation Systems. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007. LNCS, vol. 4703, pp. 120–135. Springer, Heidelberg (2007)

9. Cranen, S., Keiren, J.J.A., Willemse, T.A.C.: A cure for stuttering parity games. Technical Report 12-05, Eindhoven University of Technology, Eindhoven (2012), http://alexandria.tue.nl/repository/books/732149.pdf

10. Cranen, S., Keiren, J.J.A., Willemse, T.A.C.: Stuttering Mostly Speeds Up Solving Parity Games. In: Bobaru, M., Havelund, K., Holzmann, G.J., Joshi, R. (eds.) NFM 2011. LNCS, vol. 6617, pp. 207–221. Springer, Heidelberg (2011)

11. Emerson, E.A., Jutla, C.S.: Tree automata, mu-calculus and determinacy. In: FOCS 1991, pp. 368–377. IEEE Computer Society, Washington, DC (1991)

12. Friedmann, O., Lange, M.: Solving Parity Games in Practice. In: Liu, Z., Ravn, A.P. (eds.) ATVA 2009. LNCS, vol. 5799, pp. 182–196. Springer, Heidelberg (2009)

13. Fritz, C.: Simulation-Based Simplification of omega-Automata. PhD thesis, Christian-Albrechts-Universität zu Kiel (2005)

14. Fritz, C., Wilke, T.: Simulation Relations for Alternating Parity Automata and Parity Games. In: Ibarra, O.H., Dang, Z. (eds.) DLT 2006. LNCS, vol. 4036, pp. 59–70. Springer, Heidelberg (2006)

15. Groote, J.F., Pang, J., Wouters, A.G.: Analysis of a distributed system for lifting trucks. In: JLAP, vol. 55, pp. 21–56. Elsevier (2003)

16. Groote, J.F., Vaandrager, F.W.: An Efficient Algorithm for Branching Bisimulation and Stuttering Equivalence. In: Paterson, M. (ed.) ICALP 1990. LNCS, vol. 443, pp. 626–638. Springer, Heidelberg (1990)

17. Jurdziński, M.: Deciding the winner in parity games is in UP ∩ co-UP. IPL 68(3), 119–124 (1998)

18. Jurdziński, M.: Small Progress Measures for Solving Parity Games. In: Reichel, H., Tison, S. (eds.) STACS 2000. LNCS, vol. 1770, pp. 290–301. Springer, Heidelberg (2000)

19. Jurdziński, M., Paterson, M., Zwick, U.: A Deterministic Subexponential Algorithm for Solving Parity Games. In: SODA 2006, pp. 117–123. ACM/SIAM (2006)

20. Keiren, J.J.A., Willemse, T.A.C.: Bisimulation Minimisations for Boolean Equation Systems. In: Namjoshi, K., Zeller, A., Ziv, A. (eds.) HVC 2009. LNCS, vol. 6405, pp. 102–116. Springer, Heidelberg (2011)

21. Luttik, S.P.: Description and formal specification of the link layer of P1394. In: Workshop on Applied Formal Methods in System Design, pp. 43–56 (1997)

22. McNaughton, R.: Infinite games played on finite graphs. APAL 65(2), 149–184 (1993)

23. Obdržálek, J.: Clique-Width and Parity Games. In: CSL, pp. 54–68 (2007)

24. Schewe, S.: Solving Parity Games in Big Steps. In: Arvind, V., Prasad, S. (eds.) FSTTCS 2007. LNCS, vol. 4855, pp. 449–460. Springer, Heidelberg (2007)

25. Stirling, C.: Bisimulation, Model Checking and Other Games. In: Notes for Mathfit Workshop on Finite Model Theory, University of Wales Swansea (1996)

26. van Glabbeek, R.J., Weijland, W.P.: Branching time and abstraction in bisimulation semantics. J. ACM 43(3), 555–600 (1996)

27. Zielonka, W.: Infinite games on finitely coloured graphs with applications to automata on infinite trees. TCS 200(1-2), 135–183 (1998)