# An exponential lower bound for Cunningham's rule 

David Avis and Oliver Friedmann

the date of receipt and acceptance should be inserted later


#### Abstract

In this paper we give an exponential lower bound for Cunningham's least recently considered (round-robin) rule as applied to parity games, Markhov decision processes and linear programs. This improves a recent subexponential bound of Friedmann for this rule on these problems. The round-robin rule fixes a cyclical order of the variables and chooses the next pivot variable starting from the previously chosen variable and proceeding in the given circular order. It is perhaps the simplest example from the class of history-based pivot rules.

Our results are based on a new lower bound construction for parity games. Due to the nature of the construction we are also able to obtain an exponential lower bound for the round-robin rule applied to acyclic unique sink orientations of hypercubes (AUSOs). Furthermore these AUSOs are realizable as polytopes. We believe these are the first such results for history based rules for AUSOs, realizable or not.

The paper is self-contained and requires no previous knowledge of parity games.


Keywords simplex method, Cunningham's rule, parity games, acyclic unique sink orientations, Markhov decision processes

Mathematics Subject Classification (2000) 90C05

Oliver Friedmann
University of Munich
E-mail: me@oliverfriedmann.de
David Avis
School of Informatics, Kyoto University, Kyoto, Japan and
School of Computer Science, McGill University, Montréal, Québec, Canada
E-mail: avis@cs.mcgill.ca

## 1 Introduction

The search for a polynomial time pivoting rule for the simplex method is as old as the method itself. Klee and Minty showed in 1970 [16] that Dantzig's original rule was exponential and similar results were soon found for most of the other known rules. All such lower bound constructions were based on variations of the deformed hypercube that appeared in Klee and Minty's original paper. The constructions have the property that some variables pivoted only a very few times - sometimes only once - in the exponential pivot path. This motivated Cunningham [5], Zadeh [25] and others to consider so-called history based rules. See [1] for a formal description of these and several other history based rules.

Cunningham's least recently consider rule (round-robin rule) assigns a cyclic order to the variables and remembers the last variable to enter the basis. The next entering variable is chosen to be the first allowable candidate starting from the last chosen variable and following the given circular order. Zadeh's least entered rule chooses the entering variable to be the candidate that has entered the basis least often. History based rules defeat the deformed hypercube constructions because they tend to average out how many times a variable pivots. This pseudo-random behaviour held out the possibility that they might be at worst subexponential, if not polynomial, since the random facet rule [15] [17] is subexponential.

Friedmann gave the first evidence that history-based rules can sometimes be non-polynomial by showing the least entered rule [9] and the round-robin rule [10] are superpolynomial in the worst case. These results were obtained by first constructing a certain two person game, known as a parity game, for which the players (zero and one) follow a superpolynomial number of moves. These games are then related to Markhov decision processes (MDPs) and finally to linear programs (LPs). It is shown that each strategy in the parity game corresponds to a vertex in the derived LP and improving from one strategy to the next corresponds to a pivot step in the LP.

In this paper we first obtain a new lower bound construction for parity games. We define a strategy improvement rule for player zero that corresponds to the least recently considered rule and show an exponential lower bound on the number of strategy improvements made to complete the game. Using the earlier transformations this shows an exponential lower bound for MDPs and LPs using this rule.

However the nature of the new construction allows us to do more. An acyclic unique sink orientation of a hypercube (AUSO) [23] is an orientation of the hypercube's edges so that the resulting directed graph has no cycles and each face of the hypercube has a unique sink (vertex of outdegree zero). AUSOs appear in many applications and are an abstraction of linear programming itself [13]. LP pivot rules have natural analogues on AUSOs and their analysis has been the subject of several papers. For example in [18] an exponential lower bound is given for the random edge pivot selection rule. We are able
to show an exponential lower bound for the least recently considered rule on AUSOs, which are realizable as LPs.

The paper is organized as follows. In the next section we begin by defining parity games and policy iteration, and present the new lower bound construction. The longest part of the paper is a proof that this game requires an exponential number of moves before terminating. The section is concluded with some applications of this result to other types of games. In Section 3 we review the known connection between parity games and Markhov decision processes. This gives rise to an exponential lower bound on MDPs that use an analogue of Cunningham's rule. In Section 4 we again exploit a known connection to obtain explicit LPs from the MDP examples. The least recently considered pivot rule on these LPs is shown to require an exponential number of steps. We turn to AUSOs in Section 5. Here we show that the parity game exhibited in Section 2 gives a natural acyclic orientation of a hypercube built on player zero's strategies. Cunningham's rule on this AUSO follows an exponentially long path. Furthermore we show that the each AUSO can be realized as a polytope. In fact the polytope is the one that arises from the same parity game after transforming it to a MDP and then to an LP. The paper concludes with some open problems for future research.

## 2 Parity Game Policy Iteration Lower Bound

This section is organized as follows. We first define parity games and how the general strategy improvement algorithms operate on them. Since some readers may not be familiar with this material we will illustrate them on an example that will later be generalized for the lower bound results. We then describe the lower bound construction and prove it correct. Finally, we show how to extend the results to related game classes.

### 2.1 Parity Games

A parity game is a tuple $G=\left(V, V_{0}, V_{1}, E, \Omega\right)$ where $(V, E)$ forms a directed graph whose node set is partitioned into $V=V_{0} \cup V_{1}$ with $V_{0} \cap V_{1}=\emptyset$, and $\Omega: V \rightarrow \mathbb{N}$ is the priority function that assigns to each node a natural number called the priority of the node. We assume the graph to be total, i.e. for every $v \in V$ there is a $w \in V$ s.t. $(v, w) \in E$.

We depict parity games as directed graphs where nodes owned by player 0 are drawn as circles and nodes owned by player 1 are drawn as rectangles; all nodes are labelled with their respective name and priority. An example of such a graph is shown in Figure 1. For each node the name, $a_{2}, F_{1}, t, \cdots$, is on top and the priority is underneath. For the moment we ignore the colours on the edges and that some edges are shown dashed. (For monochromatic figures, we ignore that some edges are bold and some are dashed.) A very important property of this game is that the out-degree of each node belonging to player 0 is two. We call a game with this property a binary game.


Fig. 1 Parity Game Lower Bound Graph $G_{3}$

We use infix notation $v E w$ instead of $(v, w) \in E$ and define the set of all successors of $v$ as $v E:=\{w \mid v E w\}$. The size $|G|$ of a parity game $G=$ $\left(V, V_{0}, V_{1}, E, \Omega\right)$ is defined to be the cardinality of $E$, i.e. $|G|:=|E|$; since we assume parity games to be total w.r.t. $E$, this is a reasonable way to measure the size. The example has size 36 and $F_{2} E:=\left\{h_{2}, e_{2}, d_{2}\right\}$.

The game is played between two players called 0 and 1: starting at a node $v_{0} \in V$, they construct an infinite path through the graph as follows. If the construction so far has yielded a finite sequence $v_{0} \ldots v_{n}$ and $v_{n} \in V_{i}$ then
player $i$ selects a $w \in v_{n} E$ and the play continues with $v_{0} \ldots v_{n} w$. In the example a game may have started as the sequence $a_{2} g_{2} F_{2} d_{2}$ ending at a node owned by player 0 . She can choose between $F_{2}$ and $g_{1}$ and could continue by appending either node to the sequence.

Every play has a unique winner given by the parity of the greatest priority that occurs infinitely often. The winner of the play $v_{0} v_{1} v_{2} \ldots$ is player $i$ iff $\max \left\{p \mid \forall j \in \mathbb{N} \exists k \geq j: \Omega\left(v_{k}\right)=p\right\} \equiv_{2} i^{1}$. That is, player 0 tries to make an even priority occur infinitely often without any greater odd priorities occurring infinitely often, player 1 attempts the converse. In the example we may consider the infinite path $a_{2} g_{2} F_{2} d_{2} F_{2} d_{2} \ldots$. In this case 6 is the largest priority that occurs infinitely often and player 0 wins since this number is even.

A strategy for player $i$ is a - possibly partial - function $\sigma: V^{*} V_{i} \rightarrow V$, s.t. for all sequences $v_{0} \ldots v_{n}$ with $v_{j+1} \in v_{j} E$ for all $j=0, \ldots, n-1$, and all $v_{n} \in V_{i}$ we have: $\sigma\left(v_{0} \ldots v_{n}\right) \in v_{n} E$. A play $v_{0} v_{1} \ldots$ conforms to a strategy $\sigma$ for player $i$ if for all $j \in \mathbb{N}$ we have: if $v_{j} \in V_{i}$ then $v_{j+1}=\sigma\left(v_{0} \ldots v_{j}\right)$. Intuitively, conforming to a strategy means to always make those choices that are prescribed by the strategy. A strategy $\sigma$ for player $i$ is a winning strategy in node $v$ if player $i$ wins every play that begins in $v$ and conforms to $\sigma$.

A strategy $\sigma$ for player $i$ is called positional if for all $v_{0} \ldots v_{n} \in V^{*} V_{i}$ and all $w_{0} \ldots w_{m} \in V^{*} V_{i}$ we have: if $v_{n}=w_{m}$ then $\sigma\left(v_{0} \ldots v_{n}\right)=\sigma\left(w_{0} \ldots w_{m}\right)$. That is, the choice of the strategy on a finite path only depends on the last node on that path. So in this case we need only specify $\sigma(v)$ for each $v \in V$. The set of positional strategies for player $i$ is denoted by $\mathcal{S}_{i}(G)$. In the example a partial positional strategy could consist of $\sigma\left(a_{2}\right)=g_{2}, \sigma\left(g_{2}\right)=F_{2}, \sigma\left(F_{2}\right)=d_{2}$, $\sigma\left(d_{2}\right)=F_{2}$ and player 0 wins. Note that we do not need to give a strategy for out-degree one nodes, such as $g_{2}$, and will omit these in the sequel.

Recall that a binary game is one where each node belonging to player 0 has out-degree two. In this case player 0's positional strategy has a very simple representation. Suppose she owns $n$ nodes and labels the out-edges for each of them 0 or 1 in any arbitrary way. Then her positional strategy can be represented as a binary $n$-vector specifying for each node which edge is chosen in the strategy.

With $G$ we associate two sets $W_{0}, W_{1} \subseteq V$, where $W_{i}$ is the set of all nodes $v$ where player $i$ wins the game $G$ starting at $v$. Here we may restrict ourselves to positional strategies because it is well-known that a player has a (general) winning strategy if and only if she has a positional winning strategy for a given game. In fact, parity games enjoy positional determinacy meaning that for every node $v$ in the game either $v \in W_{0}$ or $v \in W_{1}$ [7]. Furthermore, it is not difficult to show that, whenever player $i$ has winning strategies $\sigma_{v}$ for all $v \in U$ for some $U \subseteq V$, then there is also a single strategy $\sigma$ that is winning for player $i$ from every node in $U$.

The problem of solving a parity game is to compute $W_{0}$ and $W_{1}$ as well as corresponding winning strategies $\sigma_{0}$ and $\sigma_{1}$ for the players on their respective

[^0]winning regions. In the example $F_{3} \in W_{1}$ since player 1 can set $\sigma\left(F_{1}\right)=h_{3}$ and end up in the infinite loop on $t$ with priority one which is odd.

### 2.2 Strategy Improvement

We describe here the basic definitions of the strategy improvement algorithm. For a given parity game $G=\left(V, V_{0}, V_{1}, E, \Omega\right)$, the reward of node $v$ is defined as follows: $\operatorname{rew}(v):=\Omega(v)$ if $\Omega(v) \equiv_{2} 0$ and $\operatorname{rew}(v):=-\Omega(v)$ otherwise. The set of profitable nodes for player 0 resp. 1 is defined to be $V_{\oplus}:=\{v \in V \mid$ $\left.\Omega(v) \equiv_{2} 0\right\}$ resp. $V_{\ominus}:=\left\{v \in V \mid \Omega(v) \equiv_{2} 1\right\}$.

The relevance ordering $<$ on $V$ is induced by $\Omega: v<u: \Longleftrightarrow \Omega(v)<\Omega(u)$. Additionally one defines the reward ordering $\prec$ on $V$ by $v \prec u: \Longleftrightarrow \operatorname{rew}(v)<$ $\operatorname{rew}(u)$. In our construction, although priorities are not unique, they are unique on each cycle. Therefore on each cycle both orderings are total.

Let $v$ be a node, $\sigma$ be a positional player 0 strategy and $\tau$ be a positional player 1 strategy. Starting at $v$, there is exactly one path $\pi_{\sigma, \tau, v}$ that conforms to $\sigma$ and $\tau$. Since $\sigma$ and $\tau$ are positional strategies, this path can be uniquely written as follows.

$$
\pi_{\sigma, \tau, v}=v_{1} \ldots v_{k}\left(w_{1} \ldots w_{l}\right)^{\omega}
$$

The superscript $\omega$ denotes an infinite cycle on the given vertex or vertices. Here $v_{1} \ldots v_{k}$ is a (possibly empty) non-repeating set of vertices and $w_{1} \ldots w_{l}$ is an infinite cycle with $w_{1}>w_{j}$ for all $1<j \leq l$. If $k \geq 1$ then $v=v_{1}$ otherwise $v$ is a member of the cycle. Note that the uniqueness follows from the fact that all nodes on the cycle have different priorities.

Discrete strategy improvement relies on a more abstract description of such a play $\pi_{\sigma, \tau, v}$. In fact, we only consider the dominating cycle node $w_{1}$, the set of more relevant nodes - i.e. all $v_{i}>w_{1}-$ on the path to the cycle node, and the length $k$ of the path leading to the cycle node.

The node valuation of $v$ w.r.t. $\sigma$ and $\tau$ is defined as follows.

$$
\vartheta_{\sigma, \tau, v}:=\left(w_{1},\left\{v_{i}>w_{1} \mid 1 \leq i \leq k\right\}, k\right)
$$

Given a node valuation $\vartheta$, we refer to $w_{1}$ as the cycle component, to $\left\{v_{i}>w_{1} \mid\right.$ $1 \leq i \leq k\}$ as the path component, and to $k$ as the length component of $\vartheta$.

In the example if we have $\sigma\left(a_{3}\right)=g_{3}, \tau\left(F_{3}\right)=d_{3}$ and $\sigma\left(d_{3}\right)=F_{3}$ then $\pi_{\sigma, \tau, h_{2}}=h_{2} a_{3} g_{3}\left(F_{3} d_{3}\right)^{\omega}$ and node valuation $\vartheta_{\sigma, \tau, h_{2}}:=\left(F_{3},\left\{h_{2}, g_{3}\right\}, 3\right)$. $\left\{F_{3}, d_{3}\right\}$ is the cycle component with dominating node $F_{3},\left\{h_{2}, a_{3}, g_{3}\right\}$ is the path to the cycle and has length $3,\left\{h_{2}, g_{3}\right\}$ are the more relevant nodes.

In order to compare node valuations with each other, we introduce a partial ordering $\prec$ on the set of node valuations. For this we first define a partial ordering on the path components. For a set of nodes $M$, let $\Omega(M)$ denote the priority occurence mapping:

$$
\Omega(M): p \mapsto|\{v \in M \mid \Omega(v)=p\}|
$$

Let $M$ and $N$ be two distinct sets of vertices. If $\Omega(M)=\Omega(N)$ then $M$ and $N$ are not comparable and we write this $M \sim N$. Otherwise let $p$ be the highest
priority s.t. $\Omega(M)(p) \neq \Omega(N)(p)$. Then $M \prec N$ if $\Omega(M)(p)>\Omega(N)(p)$ and $p \equiv_{2} 1$, or $\Omega(M)(p)<\Omega(N)(p)$ and $p \equiv_{2} 0$. Otherwise $N \prec M$. We observe that in a game with all priorities unique $\prec$ defines a total order on the subsets of vertices.

In the example, if we compare $M=\left\{c_{3}, c_{2}, g_{2}\right\}$ with $N=\left\{h_{1}, a_{2}, g_{2}\right\}$, we have $\Omega(M)=\{3 \mapsto 2,11 \mapsto 1\}$ and $\Omega(N)=\{10 \mapsto 1,3 \mapsto 1,11 \mapsto 1\}$. The priority with maximum value in which both sets differ is 10 . Since this priority is even we have $M \prec N$.

Now we are able to extend the partial ordering on sets of nodes to node valuations. If $u \prec v$ then $(u, M, e) \prec(v, N, f)$. Otherwise $\operatorname{rew}(u)=\operatorname{rew}(v)$ and

$$
(u, M, e) \prec(v, N, f) \Longleftrightarrow\left\{\begin{array}{l}
M \prec N \\
M \sim N, e<f \text { and } u \in V_{\ominus} \\
M \sim N, e>f \text { and } u \in V_{\oplus}
\end{array}\right.
$$

We write $(u, M, e) \sim(v, N, f)$ iff neither $(u, M, e) \prec(v, N, f)$ nor $(v, N, f) \prec$ $(u, M, e)$. We write $(u, M, e) \preceq(v, N, f)$ to abbreviate $(u, M, e) \prec(v, N, f)$ or $(u, M, e) \sim(v, N, f)$.

We observe that if all priorities are unique then we have a total order on node valuations. For in this case if $\operatorname{rew}(u)=\operatorname{rew}(v)$ then $u=v$ and if $M \sim N$ then $M=N$. We cannot have $e=f$ for otherwise $(u, M, e)=(v, N, f)$.

The motivation behind the above ordering is a lexicographic measurement of the profitability of a positional play w.r.t. player 0 : the most prominent part of a positional play is the cycle in which the plays eventually stays, and here it is the reward ordering on the dominating cycle node that defines the profitability for player 0 . The second important part is the loopless path that leads to the dominating cycle node. Here, we measure the profitability of a loopless path by a lexicographic ordering on the relevancy of the nodes on path, applying the reward ordering on each component in the lexicographic ordering. Finally, we consider the length, and the intuition behind the definition is that, assuming we have an even-priority dominating cycle node, it is better to reach the cycle fast whereas it is better to stay as long as possible out of the cycle otherwise.

In the example suppose $\sigma$ and $\tau$ give rise to the paths $c_{3} c_{2} g_{2}\left(F_{2} d_{2}\right)^{\omega}$ and $h_{1} a_{2} g_{2}\left(F_{2} d_{2}\right)^{\omega}$. We have node evaluations $\vartheta_{\sigma, \tau, h_{1}}:=\left(F_{2},\left\{h_{1}, g_{2}\right\}, 3\right), \vartheta_{\sigma, \tau, c_{3}}:=$ $\left(F_{2},\left\{g_{2}\right\}, 3\right)$ and $\vartheta_{\sigma, \tau, c_{2}}:=\left(F_{2},\left\{g_{2}\right\}, 2\right)$. We have

$$
\left(F_{2},\left\{g_{2}\right\}, 2\right) \prec\left(F_{2},\left\{g_{2}\right\}, 3\right) \prec\left(F_{2},\left\{h_{1}, g_{2}\right\}, 3\right)
$$

The first $\prec$ is due to the fact that the length of the path to the cycle node is smaller in the first node evaluation. The second $\prec$ is because the symmetric difference of the path components is $h_{1}$ with priority 10 which is even.

Given a player 0 strategy $\sigma$, it is player 1's goal to find a best response counter-strategy $\tau$ that minimizes the associated node valuations. A strategy $\tau$ is an optimal counter-strategy w.r.t. $\sigma$ iff for every opponent strategy $\tau^{\prime}$ and for every node $v$ we have: $\vartheta_{\sigma, \tau, v} \preceq \vartheta_{\sigma, \tau^{\prime}, v}$.

Is is well-known that an optimal counter-strategy always exists and that it is efficiently computable.

Lemma 1 ([24]) Let $G$ be a parity game and $\sigma$ be a player 0 strategy. An optimal counter-strategy for player 1 w.r.t. $\sigma$ exists and can be computed in polynomial time.

A fixed but arbitrary optimal counter-strategy will be denoted by $\tau_{\sigma}$ from now on. The associated game valuation $\Xi_{\sigma}$ is a map that assigns to each node the node valuation w.r.t. $\sigma$ and $\tau_{\sigma}$ :

$$
\Xi_{\sigma}: v \mapsto \vartheta_{\sigma, \tau_{\sigma}, v}
$$

In the example suppose player 0 has played to obtain the path $a_{3} g_{3}\left(F_{3} d_{3}\right)^{\omega}$. It is easy to see that $\tau_{\sigma}\left(F_{3}\right)=h_{3}$ is an optimal counter-strategy and $\vartheta_{\sigma, \tau_{\sigma}, a_{3}}=$ $\left(t,\left\{a_{3}, g_{3}, F_{3}, h_{3}\right\}, 4\right)$. We have $\Xi_{\sigma}\left(a_{3}\right)=\left(t,\left\{a_{3}, g_{3}, F_{3}, h_{3}\right\}, 4\right)$ and player 1 wins.

Game valuations are used to measure the performance of a strategy of player 0 . For a fixed strategy $\sigma$ of player 0 and a node $v$, the associated valuation essentially states which is the worst cycle that can be reached from $v$ conforming to $\sigma$ as well as the worst loopless path leading to that cycle (also conforming to $\sigma$ ). We also write $v \prec_{\sigma} u$ to compare the $\Xi_{\sigma}$-valuations of two nodes, i.e. to abbreviate $\Xi_{\sigma}(v) \prec \Xi_{\sigma}(u)$.

A run of the strategy improvement algorithm can be expressed by a sequence of improving game valuations; a partial ordering on game valuations is quite naturally defined as follows:

$$
\Xi \triangleleft \Xi^{\prime}: \Longleftrightarrow\left(\Xi(v) \preceq \Xi^{\prime}(v) \text { for all } v \in V\right) \text { and }\left(\Xi \neq \Xi^{\prime}\right)
$$

Let $\sigma$ be a strategy, $v \in V_{0}$ and $w \in v E$. We say that $(v, w)$ is a $\sigma$-improving switch iff $\sigma(v) \prec_{\sigma} w$. If $w \prec_{\sigma} \sigma(v)$, we say that $(v, w)$ is a $\sigma$-degradable switch. We say that $\sigma$ is improvable iff $\sigma$ has an improving switch. We write $I_{\sigma}$ to denote the set of improving switches and write $I_{\sigma}(v)=\left\{w \mid(v, w) \in I_{\sigma}\right\}$.

Again things become very simple for binary games. For any node $v$ we can write $v E=\{\sigma(v), w\}$ and an improving switch selects edge $(v, w)$ if $\sigma(v) \prec_{\sigma} w$. Clearly the notion of improving switch requires that the node valuations of $\sigma(v)$ and $w$ are ordered by $\prec_{\sigma}$. The binary games we construct do not have unique priorities and $\prec$ defines only a partial order on the node valuations.

Recall that for binary games player 0's current strategy can be represented by a binary $n$-vector. As improving switch just flips one of the bits in this vector. The connection with paths on hypercubes now becomes apparent.

In the example consider strategies of the two players leading to the infinite path $d_{2} g_{1} F_{1} h_{1} a_{2} a_{3} t^{\omega}$ with node valuation

$$
\vartheta_{\sigma, \tau, d_{2}}:=\left(t,\left\{d_{2}, g_{1}, F_{1}, h_{1}, a_{2}, a_{3}\right\}, 6\right) .
$$

Then $\left(d_{2}, F_{2}\right)$ is an improving switch for player 0 since it leads to the infinite path $\left(d_{2} F_{2}\right)^{\omega}$ with node valuation $\vartheta_{\sigma, \tau, d_{2}}:=\left(F_{2}, \emptyset, 0\right)$ which is better for her since $t \prec F_{2}$.

The improvement step from one strategy to the next is carried out by an improvement rule. It is a map $\mathcal{I}_{G}: \mathcal{S}_{0}(G) \rightarrow \mathcal{S}_{0}(G)$ s.t. $\Xi_{\sigma} \triangleleft \Xi_{\mathcal{I}_{G}(\sigma)}$ for every $\sigma$ and additionally $\Xi_{\sigma} \unlhd \Xi_{\mathcal{I}_{G}(\sigma)}$ if $\sigma$ is improvable. We say that a function $\mathcal{I}_{G}: \mathcal{S}_{0}(G) \rightarrow \mathcal{S}_{0}(G)$ is a standard improvement rule iff it only selects improving switches for finding a successor strategy, i.e.

1. For every node $v \in V_{0}$ it holds that $\sigma(v) \preceq_{\sigma} \mathcal{I}_{G}(\sigma)(v)$.
2. If $\sigma$ is improvable then there is a node $v \in V_{0}$ s.t. $\sigma(v) \prec_{\sigma} \mathcal{I}_{G}(\sigma)(v)$.

Jurdziński and Vöge [24] showed that an improving switch always exists for any non-optimal strategy $\sigma$. They also showed that improving $\sigma$ by an arbitrary, non-empty selection of improving switches can only result in strategies with valuations strictly better than the valuation of $\sigma$.

Theorem 1 ([24]) Let $G$ be a parity game, $\sigma$ be an improvable strategy and $\mathcal{I}_{G}$ be a standard improvement rule. Then $\Xi_{\sigma} \triangleleft \Xi_{\mathcal{I}_{G}(\sigma)}$.

If a strategy is not improvable, the strategy improvement procedure comes to an end. The game has been solved. The winning sets for both players as well as associated winning strategies can be easily derived from the given valuation.

Theorem 2 ([24]) Let $G$ be a parity game and $\sigma$ be a non-improvable strategy. Then the following holds:

1. $W_{0}=\left\{v \mid \Xi_{\sigma}(v)=\left(w,,_{-}\right)\right.$and $\left.w \in V_{\oplus}\right\}$
2. $W_{1}=\left\{v \mid \Xi_{\sigma}(v)=\left(w,_{-},\right)\right.$and $\left.w \in V_{\ominus}\right\}$
3. $\sigma$ is a winning strategy for player 0 on $W_{0}$
4. $\tau_{\sigma}$ is a winning strategy for player 1 on $W_{1}$
5. $\sigma$ is $\unlhd$-optimal

Strategy improvement starts with an initial strategy $\sigma$ and runs for a given improvement rule $\mathcal{I}$ as follows and returns an optimal player 0 strategy as outlined in the pseudo-code of Algorithm 1.

```
Algorithm 1 Strategy Improvement
    procedure StandardStratit( \(\mathcal{I}, G, \sigma)\)
        while \(\sigma\) is improvable do
            \(\sigma \leftarrow \mathcal{I}_{G}(\sigma)\)
        end while
        return \(\sigma\).
    end procedure
```

Given an initial strategy $\sigma$, a game $G$ and a rule $\mathcal{I}$, the unique execution trace, called run, of strategy improvement is the sequence of strategies $\sigma_{1}, \ldots$, $\sigma_{k}$ s.t. $\sigma_{1}=\sigma, \sigma_{i+1}=\mathcal{I}_{G}\left(\sigma_{i}\right)$ for all $i<k, \sigma_{k}$ optimal and $\sigma_{i}$ improvable for all $i<k$. The length of the run is denoted by $k$ and we say that strategy improvement requires $k+1$ iterations to find the optimal strategy.

We call a parity game $G$ (in combination with an initial strategy $\theta$ ) a sink game iff the following two properties hold:

1. Sink Existence: there is a node $v^{*}$ (called the sink of $G$ ) with $v^{*} E v^{*}$ and $\Omega\left(v^{*}\right)=1$ reachable from all nodes; also, there is no other node $w$ with $\Omega(w) \leq \Omega\left(v^{*}\right)$.
2. Sink Seeking: for each player 0 strategy $\sigma$ with $\Xi_{\theta} \unlhd \Xi_{\sigma}$ and each node $w$ it holds that the cycle component of $\Xi_{\sigma}(w)$ equals $v^{*}$.

At this point the reader may wish to verify that the example is a sink game with $v^{*}=t$. Obviously, a sink game is won by player 1. Note that comparing node valuations in a sink game can be reduced to comparing the path components of the respective node valuations, for two reasons. First, the cycle component remains constant. Second, the path-length component equals the cardinality of the path component, because all nodes except the sink node are more relevant than the cycle node itself. In the case of a sink game, we will therefore identify node valuations with their path component.

Given a parity game $G$ the sink existence property can be verified by standard graph algorithms. Given an initial strategy $\theta$ the sink seeking property can also be easily checked, as shown by the following lemma.

Lemma 2 ([8]) Let $G$ be a parity game with initial strategy $\theta$ fulfilling the sink existence property w.r.t. $v^{*} . G$ is a sink game iff $G$ is completely won by player 1 (i.e. $W_{1}=V$ ) and for each node $w$ it holds that the cycle component of $\Xi_{\theta}(w)$ equals $v^{*}$.

Let $G$ be a sink game and $v, r \in V_{G}$. We define $\Xi_{\sigma}^{>r}(v)$ to be the path component of $\Xi_{\sigma}(v)$ by filtering the nodes which are more relevant than $r$, i.e.

$$
\Xi_{\sigma}^{>r}(v)=\left\{u \in \Xi_{\sigma}^{>r}(v) \mid \Omega(u)>\Omega(r)\right\}
$$

It is easy to see that $\Xi_{\sigma}^{>r}(v) \prec \Xi_{\sigma}^{>r}(u)$ implies $\Xi_{\sigma}(v) \prec \Xi_{\sigma}(u)$. We assume from now on that every game we consider is a sink game.

Cunningham's rule [5] is a deterministic history based pivot rule for selecting entering variables in the network simplex method. It fixes an initial ordering on all variables and then selects the entering variables in a roundrobin fashion starting from the last entering variable selected. The history is simply to remember this variable. The rule can be adapted in a straightforward manner to other local improvement algorithms.

We describe Cunningham's pivoting rule in the context of parity games. We assume that we are given a total ordering $\prec$ on the player 0 edges of the parity game. The history is simply to record the last edge that has been applied.

Given a non-empty subset of player 0 edges $\emptyset \neq F \subseteq E_{0}$ and a player 0 edge $e \in E_{0}$, we define a successor operator as follows:

$$
\operatorname{succ}_{\prec}(e, F):= \begin{cases}\min _{\prec}\left\{e^{\prime} \in F \mid e \preceq e^{\prime}\right\} & \text { if }\left\{e^{\prime} \in F \mid e \preceq e^{\prime}\right\} \neq \emptyset \\ \min _{\prec}\left\{e^{\prime} \in F \mid e^{\prime} \preceq e\right\} & \text { otherwise }\end{cases}
$$

See Algorithm 2 for a pseudo-code specification of the Cunningham's rule for applied to parity games.

```
Algorithm 2 Cunningham's Improvement Algorithm
    procedure \(\operatorname{RoundRobin}(G, \sigma, \prec, e)\)
        while \(\sigma\) is improvable do
            \(e \leftarrow \operatorname{succ} \prec\left(e, I_{\sigma}\right)\)
            \(\sigma \leftarrow \sigma[e]\)
        end while
        return \(\sigma\).
    end procedure
```

Let $\left(\sigma_{1}, e_{1}\right), \ldots,\left(\sigma_{n}, e_{n}\right)$ be a trace of the algorithm w.r.t. some selection ordering $\prec$. We write $(\sigma, e) \leadsto \prec\left(\sigma^{\prime}, e^{\prime}\right)$ iff there are $i<j$ s.t. $(\sigma, e)=\left(\sigma_{i}, e_{i}\right)$ and $\left(\sigma^{\prime}, e^{\prime}\right)=\left(\sigma_{j}, e_{j}\right)$.

In the original specification of Cunningham's rule [5] it is assumed that the ordering on the edges and the initial edge $e$ is given as part of the input. In fact, we know that the asymptotic behavior of Cunningham's improvement rule highly depends on the ordering used, at least in the world of parity games and strategy improvement for games in general. We have the following theorem which is easy to verify (the idea is that there is at least one improving switch towards the optimal strategy in each step).

Theorem 3 Let $G$ be a parity game with $n$ nodes and $\sigma_{0}$ be a strategy. There is a sequence policies $\sigma_{0}, \sigma_{1}, \ldots, \sigma_{N}$ and a sequence of different switches $e_{1}, e_{2}, \ldots, e_{N}$ with $N \leq n$ s.t. $\sigma_{N-1}$ is optimal, $\sigma_{i+1}=\sigma_{i}\left[e_{i+1}\right]$ and $e_{i+1}$ is an $\sigma_{i}$-improving switch.

Since all switches are different in the sequence, it follows immediately that there is always a way to select an ordering that results in a linear number of pivoting steps to solve a parity game with Cunningham's improvement rule. However, there is no obvious method to efficiently find such an ordering. In order to derive a lower bound we are entitled to give both the input graph and the ordering to be used.

### 2.3 Lower Bound Construction

Our lower bound construction is a natural generalization of the parity game $G_{3}$, shown in Figure 1, that we used throughout the last subsection. For each $n \geq 3$ we define the underlying graph $G_{n}=\left(V_{0}, V_{1}, E, \Omega\right)$ as follows.

$$
\begin{aligned}
V_{0} & :=\left\{a_{i}, c_{i}, d_{i} \mid 1<i \leq n\right\} \cup\left\{b_{i} \mid 1<i<n\right\} \cup\left\{e_{i} \mid 1 \leq i \leq n\right\} \\
V_{1} & :=\left\{F_{i} \mid 1 \leq i \leq n\right\} \cup\left\{g_{i}, h_{i} \mid 1 \leq i \leq n\right\} \cup\{s, t\}
\end{aligned}
$$

Figure 2 defines the edge sets and the priorities of $G_{n}$. For convenience of notation, we identify the node names $a_{n+1}$ with $t, b_{1}$ with $g_{1}$, and $c_{1}$ with $g_{1}$. Explicit constructions of $G_{n}$ for small $n$ are available online [11].

Lemma 3 For every n, the game $G_{n}$ is a binary sink parity game.

| Node | Successors | Priority | Node | Successors | Priority |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $a_{i}$ | $g_{i}, a_{i+1}$ | 3 | $F_{i}$ | $h_{i}, d_{i>1}, e_{i}$ | 6 |
| $b_{i}$ | $g_{i}, b_{i-1}$ | 3 | $g_{i}$ | $F_{i}$ | $2 \cdot i+7$ |
| $c_{i}$ | $g_{i}, c_{i-1}$ | 3 | $h_{i}$ | $a_{i+1}$ | $2 \cdot i+8$ |
| $d_{i}$ | $F_{i}, b_{i-1}$ | 5 | $s$ | $c_{n}$ | 8 |
| $e_{i}$ | $F_{i}, s$ | 5 | $t$ | $t$ | 1 |

Fig. 2 Parity Game Lower Bound Graph

To avoid special cases we assume $n \geq 3$. It is easy to verify that the total number of nodes is $G_{n}$ is $8 n-3$, the total number of edges is $15 n-9$, the number of different priorities is $2 n+5$ and the highest priority is $2 n+8$. Therefore we have $\left|G_{n}\right| \in \mathcal{O}(n)$.

We refer to the edges of player 0 by the names given in Figure 3. For convenience of notation, we write $\sigma\left(a_{i}\right)=j$ to indicate that $a_{i}^{j} \in \sigma$ etc.

| Name | Node | Successor | Name | Node | Successor |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $a_{i}^{1}$ | $a_{i}$ | $g_{i}$ | $a_{i}^{0}$ | $a_{i}$ | $a_{i+1}$ |
| $b_{i}^{1}$ | $b_{i}$ | $g_{i}$ | $b_{i}^{0}$ | $b_{i}$ | $b_{i-1}$ |
| $c_{i}^{1}$ | $c_{i}$ | $g_{i}$ | $c_{i}^{0}$ | $c_{i}$ | $c_{i-1}$ |
| $d_{i}^{1}$ | $d_{i}$ | $F_{i}$ | $d_{i}^{0}$ | $d_{i}$ | $b_{i-1}$ |
| $e_{i}^{1}$ | $e_{i}$ | $F_{i}$ | $e_{i}^{0}$ | $e_{i}$ | $s$ |

Fig. 3 Parity Game Lower Bound Player 0 Edges

From a high-level point of view, a run of the strategy improvement algorithm mimics the counting process of a binary counter, yielding an exponential number of steps. Obviously, the specifics of the run depend on our choice of the edge-ordering. Every strategy that we obtain during a run corresponds to the state of the binary counter. However, a single increment step of the binary counter corresponds to several consecutive improvements in strategy improvement. These intermediate steps can be partitioned into well-defined phases.

Before describing the phases in detail, consider the layout of the game graph. It is separated into uniform layers that correspond to the different bits of the binary counter. The first and the last bit have less nodes than all the other bits. We could include the additional nodes in the game graph but they would be of no use for the counting process. Every (disregarding the first and the last) such layer contains five player 0 nodes $a_{i}, b_{i}, c_{i}, d_{i}, e_{i}$, and three player 1 nodes $F_{i}, g_{i}, h_{i}$.

The general construction extends Figure 1 in a natural way. We use terms such as up, down, left, right, etc. based on the layout used in the figure. The $a_{i}$-nodes build up a ladder-like structure that connect layers with each other, starting from the least-significant to the most-significant bit. For every layer, player 0 has to the choice to either enter the layer or to directly pass on to the next layer. By making $g_{i}$ highly unprofitable and $h_{i}$ highly profitable, it
follows that it will only be profitable for player 0 to enter a layer, if player 1 moves from $F_{i}$ to $h_{i}$. This corresponds to a set bit.

Player 0 can force player 1 to move from $F_{i}$ to $h_{i}$ by moving from both $d_{i}$ and $e_{i}$ to $F_{i}$. This is due to the fact the game is a sink game and so the optimal counter-response to a strategy cannot be moving into any other cycle than $t$.

The other two remaining nodes of player $0, b_{i}$ and $c_{i}$, build up ladder-like structures as well that connect layers with each other, but this time starting from the most-significant to the least-significant bit. The nodes $d_{i}$ and $e_{i}$ have direct access to these additional two ladders which will allow them to get reset, corresponding to unsetting a set bit in the binary counter.

Next, we explain, from a high-level point of view, how the intermediate phases contribute to incrementing the counter. At beginning of phase 1 the switches have the following settings depending on the value in the binary counter. All $a_{i}$ and $e_{i}$ point upwards if and only if bit $i$ is zero, all $d_{i}$ point to the left, and both ladder $b$ and $c$ move down to the least significant set bit in the counter. We initiate the binary counter at $00 \cdots 001$ which corresponds to the to the initial strategy $\left\{a_{*}^{0}, b_{*}^{0}, c_{*}^{0}, d_{*}^{1}, e_{1}^{1}, e_{*>1}^{0}\right\}$.

1. We apply improving switches in this phase s.t. all $e_{i}$ point left, exactly all those $d_{i}$ point to the left that correspond to a set bit in the current or in the next counter state, and update the ladder $b$ s.t. it moves down to least significant set bit in the next counter state.
2. In phase 2, we update the ladder $c$ s.t. it moves down to least significant set bit in the next counter state.
3. In phase 3 , we apply improving switches s.t. exactly all those $e_{i}$ point left that correspond to a set bit in the next counter state.
4. In phase 4 , we apply improving switches s.t. all $d_{i}$ point left.
5. In phase 5 , we update the $a_{i}$ ladder to point to the right at exactly those bits which are set in the next counter state.

After completing phase 5, the counter has been incremented by one and we can start with phase 1 again.

Phases 1-5 finish when the counter reaches all ones. The $a_{i}$ nodes point to the right, the $d_{i}$ and $e_{i}$ nodes point to the left, the $c$ chain leads down to $g_{2}$, the $b$ chain leads down to $g_{1}$ and all $F_{i}$ nodes point vertically up. This corresponds to the strategy $\left\{a_{*}^{1}, b_{*}^{0}, c_{*}^{0}, d_{*}^{1}, e_{*}^{1}\right\}$ with the exception that $c_{2}^{1}$ is chosen instead of $c_{2}^{0}$. In fact replacing $c_{2}^{1}$ by $c_{2}^{0}$ is an improving switch and the only such switch. Therefore it will be chosen by Algorithm 2. At this point Player 0 has no improving switches and loses the game. We call $\left\{a_{*}^{1}, b_{*}^{0}, c_{*}^{0}, d_{*}^{1}, e_{*}^{1}\right\}$ the terminal strategy of the game.

Next we specify a total ordering of player 0 edges to be used in Algorithm 2.

$$
\underbrace{\left\{b_{*}^{*}, d_{*}^{0}, e_{*}^{1}\right\}}_{\text {Phase 1 }} \prec \underbrace{\left\{c_{*}^{*}\right\}}_{\text {Phase 2 }} \prec \underbrace{\left\{e_{*}^{0}\right\}}_{\text {Phase 3 }} \prec \underbrace{\left\{d_{*}^{1}\right\}}_{\text {Phase 4 }} \prec \underbrace{\left\{a_{*}^{*}\right\}}_{\text {Phase 5 }}
$$

The detailed ordering for every phase is as follows:
Phase 1 : $e_{1}^{1} \prec d_{2}^{0} \prec e_{2}^{1} \prec b_{2}^{1} \prec b_{2}^{0} \prec \ldots \prec d_{n-1}^{0} \prec e_{n-1}^{1} \prec b_{n-1}^{1} \prec b_{n-1}^{0} \prec d_{n}^{0} \prec e_{n}^{1}$
Phases 2,3: $c_{2}^{0} \prec c_{2}^{1} \prec \ldots \prec c_{n}^{0} \prec c_{n}^{1}, \quad e_{1}^{0} \prec e_{2}^{0} \prec \ldots \prec e_{n}^{0}$
Phases 4,5: $d_{2}^{1} \prec d_{3}^{1} \prec \ldots \prec d_{n}^{1}, \quad a_{n}^{1} \prec a_{n}^{0} \prec \ldots \prec a_{2}^{1} \prec a_{2}^{0}$
For each $n \geq 3$ we define $P_{n}$ to be the sequence of improving switches generated by Algorithm 2 following this edge ordering, starting at the initial strategy and ending at the terminal strategy. Our goal will be to show that the sequence $P_{n}$ has exponential length in $n$.

Before proceeding with the proof let us apply Algorithm 2 to the example in Figure 1. The edge ordering is

$$
e_{1}^{1}, d_{2}^{0}, e_{2}^{1}, b_{2}^{1}, b_{2}^{0}, d_{3}^{0}, e_{3}^{1}, c_{2}^{0}, c_{2}^{1}, c_{3}^{0}, c_{3}^{1}, e_{1}^{0}, e_{2}^{0}, e_{3}^{0}, d_{2}^{1}, d_{3}^{1}, a_{3}^{1}, a_{3}^{0}, a_{2}^{1}, a_{2}^{0}
$$

Figure 1 corresponds to the initial state 001 of the binary counter, The current strategy for player 0 is shown by the blues edges, and that for player 1 by the red edges. Improving edges for player 0 are shown in dotted green and the other non-strategy edges are shown in dotted black. (Coloured edges show in bold on monochromatic printing.)

Note that since bits 2 and 3 of the counter (reading from right to left) are set to zero the strategy edges for $a_{2}, a_{3}, e_{2}$ and $e_{3}$ all point upwards. Each $d_{i}$ edge points to $F_{i}$, and the $b$ and $c$ ladders point to the first bit, which is the least bit set. The improving edges for player zero are $e_{2}^{1}, e_{3}^{1}$ and $e_{2}^{1}$ is chosen as it comes first in order. Since player 0 stands to win on the infinite loop ( $e_{2} F_{2}$ ) player 1 counters by changing the strategy on $F_{2}$ to point to $h_{2}$. The first five phases involve nine player 0 moves and are given in Figure 4. Note in some cases player 1 does not make a response as the position is still winning for him. At the end of the nine moves the counter has moved to 010 and we are back to the settings required to initiate Phase 1.

| Phase | Improving <br> Edges | Selected <br> Edge | Player 1 <br> Response |
| :---: | :---: | :---: | :---: |
| 1 | $e_{2}^{1}, e_{3}^{1}$ | $e_{2}^{1}$ | $F_{2} h_{2}$ |
| 1 | $a_{2}^{1}, b_{2}^{1}, c_{2}^{1}, e_{3}^{1}$ | $b_{2}^{1}$ |  |
| 1 | $a_{2}^{1}, c_{2}^{1}, d_{3}^{0}, e_{3}^{1}$ | $d_{3}^{0}$ |  |
| 1 | $a_{2}^{1}, c_{2}^{1}, e_{3}^{1}$ | $e_{3}^{1}$ | $F_{3} d_{3}$ |
| 2 | $a_{2}^{1}, c_{2}^{1}, d_{3}^{1}$ | $c_{2}^{1}$ |  |
| 3 | $a_{2}^{1}, d_{3}^{1}, e_{1}^{0}, e_{3}^{0}$ | $e_{3}^{0}$ |  |
| 3 | $a_{2}^{1}, d_{3}^{1}, e_{1}^{0}$ | $e_{1}^{0}$ |  |
| 4 | $a_{2}^{1}, d_{3}^{1}$ | $d_{3}^{1}$ | $F_{3} e_{3}$ |
| 5 | $a_{2}^{1}, e_{3}^{1}$ | $a_{2}^{1}$ | $F_{1} e_{1}$ |

Fig. 4 Binary counter moving from 001 to 010

Continuing in this fashion for a total of 36 improving switches by player 0 we arrive at the terminating position where she loses from each node and has no further improving edges. A complete simulation is given on the website [11].

### 2.4 Lower Bound Proof

In this section we prove the fundamental result of the paper.
Theorem 4 The sequence $P_{n}$ of improving switches followed by Algorithm 2 in $G_{n}$ from the initial strategy to the terminal strategy using the ordering (1) has length at least $2^{n}$ where $G_{n}$ has size $O(n)$.

First, we introduce notation to succinctly describe binary counters. It will be convenient for us to consider counter configurations with an infinite tape, where unused bits are zero. The set of $n$-bit configurations is formally defined as $\mathcal{B}_{n}=\left\{\mathfrak{b} \in\{0,1\}^{\infty} \mid \forall i>n: \mathfrak{b}_{i}=0\right\}$.

We start with index one, i.e. $\mathfrak{b} \in \mathcal{B}_{n}$ is essentially a tuple $\left(\mathfrak{b}_{n}, \ldots, \mathfrak{b}_{1}\right)$, with $\mathfrak{b}_{1}$ being the least and $\mathfrak{b}_{n}$ being the most significant bit. By $\mathbf{0}$, we denote the configuration in which all bits are zero, and by $\mathbf{1}_{n}$, we denote the configuration in which the first $n$ bits are one. We write $\mathcal{B}=\bigcup_{n>0} \mathcal{B}_{n}$ to denote the set of all counter configurations.

The integer value of a $\mathfrak{b} \in \mathcal{B}$ is defined as usual, i.e. $|\mathfrak{b}|:=\sum_{i>0} \mathfrak{b}_{i} \cdot 2^{i-1}<$ $\infty$. For two $\mathfrak{b}, \mathfrak{b}^{\prime} \in \mathcal{B}$, we induce the lexicographic linear ordering $\mathfrak{b}<\mathfrak{b}^{\prime}$ by $|\mathfrak{b}|<\left|\mathfrak{b}^{\prime}\right|$. It is well-known that $\mathfrak{b} \in \mathcal{B} \mapsto|\mathfrak{b}| \in \mathbb{N}$ is a bijection. For $\mathfrak{b} \in \mathcal{B}$ let $\mathfrak{b}^{\oplus}$ denote the unique $\mathfrak{b}^{\prime}$ s.t. $\left|\mathfrak{b}^{\prime}\right|=|\mathfrak{b}|+1$.

Given a configuration $\mathfrak{b}$, we access the least unset bit by $\mu(\mathfrak{b})=\min \{j \mid$ $\left.\mathfrak{b}_{j}=0\right\}$ and the least set bit by $\nu(\mathfrak{b})=\min \left\{j \mid \mathfrak{b}_{j}=1\right\}$. Let $\mathfrak{b}^{\mu}$ denote $\mathfrak{b}[\mu(\mathfrak{b}) \mapsto 1]$.

We are now ready to formulate the conditions for strategies that fulfill one of the five phases along with the improving edges. See Figure 5 for a complete description (with respect to a bit configuration $\mathfrak{b}$ ). We say that a strategy $\sigma$ is a phase $p$ strategy with configuration $\mathfrak{b}$ iff every node is mapped by $\sigma$ to a choice included in the respective cell of the table.

| Phase | $1: \exists j \geq \mu(\mathfrak{b}), k \in\{1,2,3\}$ | 2 | 3 | 4 | 5: $\exists j \leq \mu(\mathfrak{b})$ |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $\sigma\left(a_{i}\right)$ | $\mathfrak{b}_{i}$ | $\mathfrak{b}_{i}$ | $\mathfrak{b}_{i}$ | $\mathfrak{b}_{i}$ | $\begin{cases}\mathfrak{b}_{i} & \text { if } i \leq j \\ \mathfrak{b}_{i}^{\oplus} & \text { if } i>j\end{cases}$ |
| $\sigma\left(b_{i}\right)$ | $\begin{cases}\mathbb{1}_{i=\mu(\mathfrak{b})} & \text { if } j>i \\ \mathbb{1}_{i=\nu(\mathfrak{b})} & \text { otherwise }\end{cases}$ | $\mathbb{1}_{i=\mu(\mathfrak{b})}$ | $\mathbb{1}_{i=\mu(\mathfrak{b})}$ | $\mathbb{1}_{i=\mu(\mathfrak{b})}$ | $\mathbb{1}_{i=\mu(\mathfrak{b})}$ |
| $\sigma\left(c_{i}\right)$ | $\mathbb{1}_{i=\nu(\mathfrak{b})}$ | $\mathbb{1}_{i=\nu(\mathfrak{b})}$ | $\mathbb{1}_{i=\mu(\mathfrak{b})}$ | $\mathbb{1}_{i=\mu(\mathfrak{b})}$ | $\mathbb{1}_{i=\mu(\mathfrak{b})}$ |
| $\sigma\left(d_{i}\right)$ | $\begin{cases}\mathfrak{b}_{i}^{\mu} & \text { if } j>i \vee(j=i \wedge k>1) \\ 1 & \text { otherwise }\end{cases}$ | $\mathfrak{b}_{i}^{\mu}$ | $\mathfrak{b}_{i}^{\mu}$ | $1, \mathfrak{b}_{i}^{\mu}$ | 1 |
| $\sigma\left(e_{i}\right)$ | $\begin{cases}1 & \text { if } j>i \vee(j=i \wedge k>2) \\ \mathfrak{b}_{i} & \text { otherwise }\end{cases}$ | 1 | $1, \mathfrak{b}_{i}^{\oplus}$ | $\mathfrak{b}_{i}^{\oplus}$ | $\mathfrak{b}_{i}^{\oplus}$ |

Fig. 5 Policy Phases

The following lemma computes an optimal counter-strategy along with the associated valuations of the nodes $F_{i}$ given a strategy $\sigma$ belonging to one of the phases.

Lemma 4 Let $n \geq 3$ and $\sigma$ be a strategy belonging to one of the phases w.r.t. $\mathfrak{b}$. Let $k=\max \left(\left\{i \mid \sigma\left(a_{i}\right) \neq \mathfrak{b}_{i}^{\oplus}\right\} \cup\{1\}\right)$. Then the following holds: $\left(\tau_{\sigma}\left(F_{i}\right), \Xi_{\sigma}^{>6}\left(F_{i}\right)\right)=$
$\begin{cases}\left(h_{i},\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}\right) & \text { if } \sigma\left(d_{i}\right)=\sigma\left(e_{i}\right)=1 \\ \left(e_{i},\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1\right\} \cup\{s\}\right) & \text { if } \sigma\left(e_{i}\right)=0 \wedge P=1 \\ \left(d_{i},\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\}\right) & \text { if } \sigma\left(d_{i}\right)=0 \wedge\left(\sigma\left(e_{i}\right)=1 \vee P \neq 1\right) \\ \left(e_{i},\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} \cup\{s\}\right) & \text { if } \sigma\left(e_{i}\right)=0 \wedge \sigma\left(d_{i}\right)=1 \wedge P \neq 1 \wedge i \geq k \\ \left(h_{i},\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}\right) & \text { if } \sigma\left(e_{i}\right)=0 \wedge \sigma\left(d_{i}\right)=1 \wedge P \neq 1 \wedge i<k=\mu(\mathfrak{b}) \\ \left(h_{i},\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1 \vee i<j<k\right\} \cup\left\{s, h_{i}, g_{k}\right\}\right) & \text { if } \sigma\left(e_{i}\right)=0 \wedge \sigma\left(d_{i}\right)=1 \wedge P \neq 1 \wedge i<k<\mu(\mathfrak{b})\end{cases}$

Proof Let $n \geq 3$ and $\sigma$ be a strategy belonging to one of the phases w.r.t. $\mathfrak{b}$.
Let $k=\max \left(\left\{i \mid \sigma\left(a_{i}\right) \neq \mathfrak{b}_{i}^{\oplus}\right\} \cup\{1\}\right)$.

- Phase 1: First, if $\sigma\left(d_{i}\right)=\sigma\left(e_{i}\right)=1$, it follows that $\tau_{\sigma}\left(F_{i}\right)=h_{i}$ since we have a sink game. Furthermore it follows by definition of phase 1 that

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}
$$

Otherwise, $i \geq \mu(\mathfrak{b})$. It follows by definition of phase 1 that

$$
\begin{aligned}
\Xi_{\sigma}^{>6}(s) & =\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1\right\} \cup\{s\} \\
\Xi_{\sigma}^{>6}\left(h_{i}\right) & =\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}
\end{aligned}
$$

and if $\sigma\left(d_{i}\right)=0$ that

$$
\Xi_{\sigma}^{>6}(s)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\}
$$

Second, if $\sigma\left(d_{i}\right)=1$ and $\sigma\left(e_{i}\right)=0$, it follows that $e_{i} \prec_{\sigma} h_{i}$ and therefore $\tau_{\sigma}\left(F_{i}\right)=e_{i}$ and

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1\right\} \cup\{s\}
$$

Third, if $\sigma\left(d_{i}\right)=0$ and $\sigma\left(e_{i}\right)=1$, it follows that $d_{i} \prec_{\sigma} h_{i}$ and therefore $\tau_{\sigma}\left(F_{i}\right)=d_{i}$ and

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\}
$$

Fourth, if $\sigma\left(d_{i}\right)=\sigma\left(e_{i}\right)=0$, it follows that $e_{i} \prec_{\sigma} d_{i} \prec_{\sigma} h_{i}$ and therefore $\tau_{\sigma}\left(F_{i}\right)=e_{i}$ and

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1\right\} \cup\{s\}
$$

- Phase 2: First, if $\mathfrak{b}_{i}^{\mu}=1$, it follows that $\sigma\left(d_{i}\right)=\sigma\left(e_{i}\right)=1$ by definition of phase 2. Since we have a sink game, it follows that $\tau_{\sigma}\left(F_{i}\right)=h_{i}$. Furthermore it follows by definition of phase 2 that

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}
$$

Second, if $\mathfrak{b}_{i}^{\mu}=0$, if follows that $i>\mu(\mathfrak{b})$ and from definition of phase 2 that $\sigma\left(d_{i}\right)=0, \sigma\left(e_{i}\right)=1$, and

$$
\begin{aligned}
\Xi_{\sigma}^{>6}\left(h_{i}\right) & =\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\} \\
\Xi_{\sigma}^{>6}\left(b_{i-1}\right) & =\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\}
\end{aligned}
$$

Hence, $b_{i-1} \prec_{\sigma} h_{i}$ and therefore $\tau_{\sigma}\left(F_{i}\right)=d_{i}$ and

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\}
$$

- Phase 3 \& 4: First, if $\mathfrak{b}_{i}^{\oplus}=1$, it follows that $\sigma\left(d_{i}\right)=\sigma\left(e_{i}\right)=1$ by definition of phase 3 and phase 4 . Since we have a sink game, it follows that $\tau_{\sigma}\left(F_{i}\right)=h_{i}$. Furthermore it follows by definition of phase 3 and phase 4 that

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}
$$

Second, if $\mathfrak{b}_{i}^{\mu}=0$, it follows that $i>\mu(\mathfrak{b})$. By definition of phase 3 and phase 4, we have

$$
\begin{aligned}
\Xi_{\sigma}^{>6}\left(h_{i}\right) & =\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\} \\
\Xi_{\sigma}^{>6}\left(b_{i-1}\right) & =\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} \\
\Xi_{\sigma}^{>6}(s) & =\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} \cup\{s\}
\end{aligned}
$$

It follows that $b_{i-1} \prec_{\sigma} s \prec_{\sigma} h_{i}$ and hence $\tau_{\sigma}\left(F_{i}\right)=d_{i}$ and

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\}
$$

if $\sigma\left(d_{i}\right)=0$, and $\tau_{\sigma}\left(F_{i}\right)=e_{i}$ and

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} \cup\{s\}
$$

if $\sigma\left(d_{i}\right)=1$.
Third, if $i<\mu(\mathfrak{b})=k$, we show by backward induction on $i$ that $\tau_{\sigma}\left(F_{i}\right)=$ $h_{i}$, which implies that

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}
$$

By induction hypothesis or by considering the base case $i=k-1$ directly, we have

$$
\Xi_{\sigma}^{>6}\left(h_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}
$$

and therefore $h_{i} \prec_{\sigma} s$, implying $\tau_{\sigma}\left(F_{i}\right)=h_{i}$.

- Phase 5: First, if $\mathfrak{b}_{i}^{\oplus}=1$, it follows that $\sigma\left(d_{i}\right)=\sigma\left(e_{i}\right)=1$ by definition of phase 5 . Since we have a sink game, it follows that $\tau_{\sigma}\left(F_{i}\right)=h_{i}$. Furthermore it follows by definition of phase 5 that

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}
$$

Second, if $\mathfrak{b}_{i}^{\oplus}=0$ and $i \geq k$, it follows by definition of phase 5 that

$$
\begin{aligned}
\Xi_{\sigma}^{>6}\left(h_{i}\right) & =\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1, j>i\right\} \cup\left\{h_{i}\right\} \\
\Xi_{\sigma}^{>6}(s) & =\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} \cup\{s\}
\end{aligned}
$$

It follows that $s \prec_{\sigma} h_{i}$ and hence $\tau_{\sigma}\left(F_{i}\right)=e_{i}$ and

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} \cup\{s\}
$$

Third, if $i<\mu(\mathfrak{b})=k$, we show by backward induction on $i$ that $\tau_{\sigma}\left(F_{i}\right)=$ $h_{i}$, which implies that

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}
$$

By induction hypothesis or by considering the base case $i=k-1$ directly, we have

$$
\Xi_{\sigma}^{>6}\left(h_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\}
$$

and therefore $h_{i} \prec_{\sigma} s$, implying $\tau_{\sigma}\left(F_{i}\right)=h_{i}$.
Fourth, if $i<k<\mu(\mathfrak{b})$, we show by backward induction on $i$ that $\tau_{\sigma}\left(F_{i}\right)=$ $h_{i}$, which implies that

$$
\Xi_{\sigma}^{>6}\left(F_{i}\right)=\Xi_{\sigma}^{>6}(s) \cup\left\{g_{j+1}, h_{j} \mid i \leq j<k\right\}
$$

By induction hypothesis or by considering the base case $i=k-1$ directly, we have

$$
\Xi_{\sigma}^{>6}\left(h_{i}\right)=\Xi_{\sigma}^{>6}(s) \cup\left\{g_{j+1}, h_{j} \mid i<j<k\right\} \cup\left\{h_{i}, g_{i+1}\right\}
$$

and therefore $h_{i} \prec_{\sigma} s$, implying $\tau_{\sigma}\left(F_{i}\right)=h_{i}$.

Figure 6 specifies the sets of improving switches for each phase $p$. It should be read as follows: an edge $e$ is included in the set of improving switches $I(\sigma)$ iff $e \notin \sigma$ and the condition holds that is specified in the respective cell. If a cell contains a question mark, we do not specify whether the edge is included in the set.

| Phase | 1 | 2 | 3 | 4 | 5 |
| :--- | :---: | :---: | :---: | :---: | :---: |
| $a_{i}^{1} \in I(\sigma)$ | $?$ | $?$ | $?$ | $?$ | $i=\mu(\mathfrak{b})$ |
| $a_{i}^{0} \in I(\sigma)$ | $?$ | $?$ | $?$ | $?$ | $i<\mu(\mathfrak{b}) \wedge \sigma\left(a_{i+1}\right)=\mathfrak{b}_{i}^{\oplus}$ |
| $b_{i}^{1} \in I(\sigma)$ | $i=\mu(\mathfrak{b}) \wedge \sigma\left(e_{i}\right)=1$ | $?$ | $?$ | $?$ | $?$ |
| $b_{i}^{0} \in I(\sigma)$ | $i>\mu(\mathfrak{b}) \wedge \sigma\left(e_{\mu(\mathfrak{b})}\right)=1$ | $?$ | $?$ | $?$ | $?$ |
| $c_{i}^{1} \in I(\sigma)$ | $?$ | $i=\mu(\mathfrak{b})$ | $?$ | $?$ | $?$ |
| $c_{i}^{0} \in I(\sigma)$ | $?$ | $i \neq \mu(\mathfrak{b})$ | $?$ | $?$ | $?$ |
| $d_{i}^{1} \in I(\sigma)$ | $?$ | $?$ | $?$ | Yes | $?$ |
| $d_{i}^{0} \in I(\sigma)$ | $\sigma\left(b_{\mu(\mathfrak{b})}\right)=1 \wedge \sigma\left(b_{\nu(\mathfrak{b})}\right)=0$ | $?$ | $?$ | $?$ | $?$ |
| $e_{i}^{1} \in I(\sigma)$ | $Y e s$ | $?$ | $?$ | $?$ | $?$ |
| $e_{i}^{0} \in I(\sigma)$ | $?$ | $?$ | $\mathfrak{b}_{i}^{\oplus}=0$ | $?$ | $?$ |

Fig. 6 Improving Switches

We finally arrive at the following main lemma describing the improving switches.
Lemma 5 Let $n \geq 3$. The improving switches from policies that belong to the phases in Figure 5 are as specified in Figure 6.

Proof Let $n \geq 3$ and $\sigma$ be a strategy belonging to one of the phases w.r.t. $\mathfrak{b}$.

- Phase 1: It follows from Lemma 4 that

$$
\Xi_{\sigma}^{>6}\left(g_{i}\right)= \begin{cases}\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{g_{i}, h_{i}\right\} & \text { if } \sigma\left(d_{i}\right)=\sigma\left(e_{i}\right)=1 \\ \left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1\right\} \cup\left\{g_{i}, s\right\} & \text { if } \sigma\left(e_{i}\right)=0 \\ \left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} \cup\left\{g_{i}\right\} & \text { if } \sigma\left(d_{i}\right)=0 \wedge \sigma\left(e_{i}\right)=1\end{cases}
$$

Since $\tau_{\sigma}\left(F_{i}\right)=e_{i}$ if $\sigma\left(e_{i}\right)=0$, it immediately follows that $e_{i}^{1}$ is an improving switch. It can furthermore be easily observed that $d_{i}^{0}$ is an improving switch for $\mathfrak{b}_{i}=0, i>\nu(\mathfrak{b})$ iff $\sigma\left(e_{\nu(\mathfrak{b})}\right)=1$ and $\sigma\left(b_{\nu(\mathfrak{b})}\right)=1$.

- Phase 2: It follows from Lemma 4 that

$$
\Xi_{\sigma}^{>6}\left(g_{i}\right)= \begin{cases}\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{g_{i}, h_{i}\right\} & \text { if } \mathfrak{b}_{i}^{\mu}=1 \\ \left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} \cup\left\{g_{i}\right\} & \text { if } \mathfrak{b}_{i}^{\mu}=0\end{cases}
$$

If $\mathfrak{b}_{1}=0$ and $\mathfrak{b}_{2}=1$, it follows that

$$
\Xi_{\sigma}^{>6}\left(c_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1\right\}
$$

It immediately follows that $c_{2}^{0}$ is the only improving switch w.r.t. $c$.
Otherwise, if $\mathfrak{b}_{1}=0$ and $\mathfrak{b}_{2}=0$, it follows that

$$
\Xi_{\sigma}^{>6}\left(c_{i}\right)= \begin{cases}\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1\right\} & \text { if } i \geq \nu(\mathfrak{b}) \\ \left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}\right\} & \text { if } i<\nu(\mathfrak{b})\end{cases}
$$

It immediately follows that $c_{\nu(\mathfrak{b})}^{0}$ is the only improving switch w.r.t. $c$.
Otherwise, if $\mathfrak{b}_{1}=1$, it follows that

$$
\Xi_{\sigma}^{>6}\left(c_{i}\right)=\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1\right\}
$$

Hence, it follows that $c_{\mu(b i t)}^{1}$ is the only improving switch w.r.t. $c$.

- Phase 3: Let $i$ s.t. $\sigma\left(e_{i}\right)=1$. It follows from Lemma 4 that

$$
\begin{aligned}
\Xi_{\sigma}^{>6}\left(F_{i}\right) & = \begin{cases}\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{h_{i}\right\} & \text { if } \sigma\left(d_{i}\right)=1 \\
\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} & \text { if } \sigma\left(d_{i}\right)=0\end{cases} \\
\Xi_{\sigma}^{>6}(s) & =\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} \cup\{s\}
\end{aligned}
$$

Hence, we have $F_{i} \prec_{\sigma} s$.

- Phase 4: Let $i$ s.t. $\sigma\left(d_{i}\right)=0$. It follows from Lemma 4 that $\tau_{\sigma}\left(F_{i}\right)=d_{i}$ and hence $\Xi_{\sigma}\left(F_{i}\right)=\left\{F_{i}\right\} \cup \Xi_{\sigma}\left(d_{i}\right)$, i.e. $\sigma\left(d_{i}\right) \prec_{\sigma} F_{i}$.
- Phase 5: Let $k=\max \left(\left\{i \mid \sigma\left(a_{i}\right) \neq \mathfrak{b}_{i}^{\oplus}\right\} \cup\{1\}\right)$. It follows from Lemma 4 that

$$
\Xi_{\sigma}^{>6}\left(g_{i}\right)= \begin{cases}\left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}=1, j>i\right\} \cup\left\{g_{i}, h_{i}\right\} & \text { if } \mathfrak{b}_{i}^{\oplus}=1 \vee(i<k=\mu(\mathfrak{b})) \\ \left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1\right\} \cup\left\{g_{i}, s\right\} & \text { if } \mathfrak{b}_{i}^{\oplus}=0 \wedge i \geq k \\ \left\{g_{j}, h_{j} \mid \mathfrak{b}_{j}^{\oplus}=1 \vee i<j<k\right\} \cup\left\{s, g_{k}, g_{i}, h_{i}\right\} & \text { if } \mathfrak{b}_{i}^{\oplus}=0 \wedge i<k<\mu(\mathfrak{b})\end{cases}
$$

Hence, it follows that if $k>1$, the only improving switch is either $a_{k}^{1}$ or $a_{k}^{0}$ w.r.t. $a$.

We are now ready to formulate our main lemma describing the transitioning from an initial phase 1 strategy corresponding to $\mathfrak{b}$ to a successor initial phase 1 strategy corresponding to $\mathfrak{b}^{\oplus}$, complying with the given ordering selection.

Lemma 6 Let $\sigma$ be a phase 1 strategy with configuration $\mathbf{0}_{n}<\mathfrak{b}<\mathbf{1}_{n}$. Let $z$ be an edge with $z \preceq e_{1}^{1}$ or $a_{n}^{1} \preceq e$. Then, there is a phase 1 strategy $\sigma^{\prime}$ with configuration $\mathfrak{b}^{\oplus}$ and an edge $z^{\prime}$ with $z^{\prime} \preceq e_{1}^{1}$ or $a_{n}^{1} \preceq z^{\prime}$ s.t. $(\sigma, z) \leadsto \prec\left(\sigma^{\prime}, z^{\prime}\right)$.

Proof The proof of the lemma is ultimately based on the five phases described in Figure 5, the corresponding improving switches given in Figure 6 (proven correct in Lemma 5) and the introduced selection ordering.

We prove the lemma by outlining the complete sequence of switches that are applied to $\sigma$ in order to obtain $\sigma^{\prime}$ (we do not explicitly describe the intermediate strategies which can be derived by applying all mentioned switches up to that point).

Let $i_{1}, \ldots, i_{k}$ be the complete sequence of ascending indices s.t. $\mathfrak{b}_{i_{j}}=0$ for $1 \leq j \leq k$. The following holds:

$$
\begin{aligned}
& \stackrel{\mathrm{P} 1}{\sim} e_{i_{1}}^{1} \leadsto b_{i_{1}}^{1} \leadsto d_{i_{2}}^{0} \leadsto e_{i_{2}}^{1} \leadsto d_{i_{3}}^{0} \leadsto e_{i_{3}}^{1} \leadsto \ldots \leadsto d_{i_{k}}^{0} \leadsto e_{i_{k}}^{m} 1 \\
& \stackrel{\text { P2 }}{\sim} c_{i_{1}}^{1} \\
& \stackrel{\text { P3 }}{\sim}\left\{e_{i}^{0} \mid \mathfrak{b}^{\oplus}=0\right\} \\
& \stackrel{\text { P4 }}{\sim}\left\{d_{i}^{1} \mid \mathfrak{b}^{\nu}=0\right\} \\
& \stackrel{\text { P5 }}{\sim} a_{i_{1}}^{*} \leadsto a_{i_{1}-1}^{*} \leadsto \ldots \leadsto a_{2}^{*}
\end{aligned}
$$

It follows immediately that the parity game provided here indeed simulate a binary counter by starting with the designated initial strategy and the $\prec-$ minimal edge. This completes the proof of Theorem 4.

### 2.5 Application to Other Games

Theorem 4 can be applied to a variety of other games using known connections between these games. For completeness we give these results here. However we will not give definitions of the games concerned, referring instead to the relevant literature, as they will not be needed in the rest of the paper.

Parity games can be reduced to mean payoff games [21], mean payoff games to discounted payoff games, and the latter ones to turn-based stochastic games [26]. Friedmann showed [8] that the strategy improvement algorithm for payoff and stochastic games, when applied to the reduced game graphs of sink parity games, behaves exactly the same on the reduced games. In other words, the strategies in both the original game and the reduced game graphs coincide as well as the associated sets of improving switches.

Theorem 5 ([8]) Let $G$ be a sink parity game. Discrete strategy improvement requires the same number of iterations to solve $G$ as strategy improvement for the induced payoff games as well as turn-based stochastic games to solve the respective game $G^{\prime}$. The game $G^{\prime}$ is induced by applying the standard reduction from $G$ to the respective game class, assuming that the improvement rule solely depends on the combinatorial valuation-ordering of the improving edges.

By this, we can conclude that the exponential lower bound for Cunningham's rule on parity games presented here also applies to payoff and turnbased stochastic games. We will see in the next section that it also applies to Markov decision processes, which can be seen as a one-player version of turn-based stochastic games.

Corollary 1 There is a family of mean payoff games, discounted payoff games resp. turn-based stochastic games on which number of improving steps performed by Algorithm 2 is at least $2^{n}$, where the size of the games are $O(n)$.

## 3 Markov Decision Process Policy Iteration Lower Bound

Markov decision processes (MDPs) provide a mathematical model for sequential decision making under uncertainty. They are employed to model stochastic optimization problems in various areas ranging from operations research, machine learning, artificial intelligence, economics and game theory. For an indepth coverage of MDPs, see the books of Howard [14], Derman [6], Puterman [22] and Bertsekas [3].

### 3.1 Markov Decision Processes and Policy Iteration

Formally, an MDP is defined by its underlying graph $G=\left(V_{0}, V_{R}, E_{0}, E_{R}, r, p\right)$. Here, $V_{0}$ is the set of vertices (states) operated by the controller, also known as player 0 , and $V_{R}$ is a set of randomization vertices corresponding to the probabilistic actions of the MDP. We let $V=V_{0} \cup V_{R}$. The edge set $E_{0} \subseteq$ $V_{0} \times V_{R}$ corresponds to the actions available to the controller. The edge set $E_{R} \subseteq V_{R} \times V_{0}$ corresponds to the probabilistic transitions associated with each action. The function $r: E_{0} \rightarrow \mathbb{R}$ is the immediate reward function. The function $p: E_{R} \rightarrow[0,1]$ specifies the transition probabilities. For every $u \in V_{R}$, we have $\sum_{v:(u, v) \in E_{R}} p(u, v)=1$, i.e., the probabilities of all edges emanating from each vertex of $V_{R}$ sum up to 1 .

A policy $\sigma$ is a function $\sigma: V_{0} \rightarrow V$ that selects for each vertex $u \in V_{0}$ a target node $v$ corresponding to an edge $(u, v) \in E_{0}$, i.e. $(u, \sigma(u)) \in E_{0}$. We assume that each vertex $u \in V_{0}$ has at least one outgoing edge). There are several objectives for MDPs and we consider the expected total reward objective here. The values $\operatorname{val}_{\sigma}(u)$ of the vertices under $\sigma$ are defined as the
unique solutions of the following set of linear equations:

$$
\operatorname{val}_{\sigma}(u)= \begin{cases}\operatorname{val}_{\sigma}(v)+r(u, v) & \text { if } u \in V_{0} \text { and } \sigma(u)=v \\ \sum_{v:(u, v) \in E_{R}} p(u, v) \operatorname{val}_{\sigma}(v) & \text { if } u \in V_{R}\end{cases}
$$

together with the condition that $\operatorname{val}_{\sigma}(u)$ sum up to 0 on each irreducible recurrent class of the Markov chain defined by $\sigma$.

All MDPs considered in this paper satisfy the unichain condition (see [22]). It states that the Markov chain obtained from each policy $\sigma$ has a single irreducible recurrent class.

This condition implies, in particular, that all vertices have the same value. It is not difficult to check that $\operatorname{val}_{\sigma}(u)$ is indeed the expected reward per turn, when the process starts at $u$ and policy $\sigma$ is used. The potentials $\operatorname{pot}_{\sigma}(u)$ represent biases. Loosely speaking, the expected reward after $N$ steps, when starting at $u$ and following $\sigma$, and when $N$ is sufficiently large, is about $N \operatorname{val}_{\sigma}(u)+\operatorname{pot}_{\sigma}(u)$.

Howard's [14] policy iteration algorithm is the most widely used algorithm for solving MDPs. It is closely related to the simplex method, which provides a practical way to solve such problems. Nevertheless in the worst case, Melekopoglou and Condon [19] showed that the simplex method with the smallest index pivot rule needs an exponential number of iterations to compute an optimal policy for a specific MDP problem regardless of discount factors.

Although most MDPs can be solved in polynomial time using the interior point method the search continues for a strongly polynomial time method. Post and Ye [20] have recently made progress in this direction by showing that the simplex method with the greedy pivot rule terminates in at most $O\left(m^{3} n^{2} \log ^{2} m\right)$ pivot steps when discount factors are uniform, or in at most $O\left(m^{5} n^{3} \log ^{2} m\right)$ pivot steps with non-uniform discounts. No such results have been proved for the policy iteration method. As history based pivot rules provide good candidates for subexponential time behaviour it is important to analyze their worst case performance for policy improvement algorithms.

As is the case for parity games, the policy iteration algorithm starts with some initial policy $\sigma_{0}$ and generates an improving sequence $\sigma_{0}, \sigma_{1}, \ldots, \sigma_{N}$ of policies, ending with an optimal policy $\sigma_{N}$. In each iteration the algorithm first evaluates the current policy $\sigma_{i}$, by computing the values $\mathrm{val}_{\sigma_{i}}(u)$ of all vertices. An edge $\left(u, v^{\prime}\right) \in E_{0}$, such that $\sigma_{i}(u) \neq v^{\prime}$ is then said to be an improving switch if and only if either $\operatorname{val}_{\sigma_{i}}\left(v^{\prime}\right)>\operatorname{val}_{\sigma_{i}}(u)$. Given a policy $\sigma$, we again denote the set of improving switches by $I_{\sigma}$.

A crucial property of policy iteration is that $\sigma$ is an optimal policy if and only if there are no improving switches with respect to it (see, e.g., [14], [22]). Furthermore, if $\left(u, v^{\prime}\right) \in I_{\sigma}$ is an improving switch w.r.t. $\sigma$, and $\sigma^{\prime}$ is defined as $\sigma\left[\left(u, v^{\prime}\right)\right]$ (i.e., $\sigma^{\prime}(u)=v^{\prime}$ and $\sigma^{\prime}(w)=\sigma(w)$ for all $w \neq u$ ), then $\sigma^{\prime}$ is strictly better than $\sigma$, in the sense that for every $u \in V_{0}$, we have $\operatorname{val}_{\sigma^{\prime}}(u) \geq \operatorname{val}_{\sigma}(u)$, with a strict inequality for at least one vertex $u \in V_{0}$.

Lemma 7 Let $\sigma$ be a policy and $(v, w)$ be a $\sigma$-improving switch. Let $\sigma^{\prime}=$ $\sigma[v \mapsto w]$. Then the following holds:

1. $\operatorname{val}_{\sigma^{\prime}}(u) \geq \operatorname{val}_{\sigma}(u)$ for all $u \in V$,
2. $\operatorname{val}_{\sigma^{\prime}}(v)>\operatorname{val}_{\sigma}(v)$, and particularly
3. $\sum_{u \in V_{0}} \operatorname{val}_{\sigma^{\prime}}(u)>\sum_{u \in V_{0}} \operatorname{val}_{\sigma}(u)$.

### 3.2 Lower Bound Construction

We relate the description of the lower bound construction for Markov decision processes closely to the construction of the parity games. For that reason, we relax our definition of MDPs such that it corresponds almost directly to parity games.

As defined in the previous subsection, the underlying graph $G$ is bipartite. However one can relax this condition and allow edges from $V_{0}$ to $V_{0}$ and from $V_{R}$ to $V_{R}$. It suffices to subdivide these edges by inserting a node belonging to player 1 or player 0 , respectively, with out-degree 1 and no priority. This leads to the following definition.

A relaxed $M D P$ is a tuple $M=\left(V, V_{0}, V_{R}, E, E_{0}, E_{R}, r, p\right)$, where $V=$ $V_{0} \cup V_{R}, E \subseteq V \times V, E_{0}=E \cap\left(V_{0} \times V\right), E_{R}=E \cap\left(V_{R} \times V\right), r: E \mapsto \mathbb{R}$ and $p: E_{R} \rightarrow[0,1]$ with $\sum_{w \in v E} p(v, w)=1$ for all $v \in V_{R}$.

Relaxed MDPs allow us to show the close relationship between the original parity games and the corresponding MDPs.

1. Edges $(v, w) \in E_{0} \cap V_{0} \times V_{0}$ can be realized by adding a randomization node $(v, w)$ and by replacing the edge $(v, w)$ with new edges $(v,(v, w))$ and $((v, w), w)$. The outgoing edge from $(v, w)$ obviously has probability 1.
2. Edges $(v, w) \in E_{R} \cap V_{R} \times V_{R}$ can be realized by adding a player 0 node $(v, w)$ and by replacing the edge $(v, w)$ with new edges $(v,(v, w))$ and $((v, w), w)$. Note that player 0 does not obtain new choices by adding this node since the out-degree is one.
3. Randomization edges are not allowed to have rewards in MDPs. Hence we insert, for every outgoing edge of a randomization node with reward, a new node of player 0 connected to a new node of the randomizer connected to the original target node. We push the reward to the new player 0 edge.
4. Since we consider the expected total reward here, adding new intermediate nodes does not change the value of the nodes.

For each $n \geq 3$ we define the underlying graph $M_{n}=\left(V, V_{0}, V_{R}, E, E_{0}, E_{R}, r, p\right)$ of a relaxed MDP as shown schematically in Figure 8 which the reader is invited to compare with Figure 1.

More formally:

$$
\begin{aligned}
V_{0} & :=\left\{a_{i}, c_{i}, d_{i} \mid 1<i \leq n\right\} \cup\left\{b_{i} \mid 1<i<n\right\} \cup\left\{e_{i} \mid 1 \leq i \leq n\right\} \\
V_{R} & :=\left\{F_{i} \mid 1 \leq i \leq n\right\} \cup\left\{g_{i}, h_{i} \mid 1 \leq i \leq n\right\} \cup\{s, t\}
\end{aligned}
$$

With $M_{n}$, we associate a large number $N \in \mathbb{N}$ and a small number $0<\varepsilon$. We require $N$ to be at least as large as the number of nodes with priorities, i.e. $N \geq 2 n$ and $\varepsilon^{-1}$ to be significantly larger than the largest occurring priority induced reward, i.e. $\varepsilon \leq \frac{1}{2 n}$.

Some of the vertices are assigned integer priorities. If a vertex $v$ has priority $\Omega(v)$ assigned to it, then a reward of $\langle v\rangle=(-N)^{\Omega(v)}$ is added to all edges emanating from $v$. This idea of using priorities is inspired by the reduction from parity games to mean payoff games, see [21].

Figure 7 defines the edge sets, the probabilities and the priorities of $M_{n}$. For convenience of notation, we identify the node names $a_{n+1}$ with $t, b_{1}$ with $g_{1}$, and $c_{1}$ with $g_{1}$. Explicit constructions for small $n$ are available online [11].

| Node | Successors | Probability | Node | Successors |
| :---: | :---: | :---: | :---: | :---: |
| $F_{i}$ | $h_{i}$ | $\varepsilon$ |  |  |
| $F_{i}$ | $d_{i}$ | $0.5 \cdot(1-\varepsilon)$ | $a_{i}$ | $g_{i}, a_{i+1}$ |
| $F_{i}$ | $e_{i}$ | $0.5 \cdot(1-\varepsilon)$ | $b_{i}$ | $g_{i}, b_{i-1}$ |
| Node | Successors | Priority | $c_{i}$ | $g_{i}, c_{i-1}$ |
| $g_{i}$ | $F_{i}$ | $2 \cdot i-1$ | $e_{i}$ | $F_{i}, b_{i-}$ $F_{i}, s$ |
| $h_{i}$ | $a_{i+1}$ | $2 \cdot i$ | , | , |
| $s$ | $c_{n}$ | 0 |  |  |

Fig. 7 MDP Lower Bound Graph

In comparing Figure 8 with Figure 1 the connections between MDPs and parity games becomes clear. Blue edges show respectively the current policy and current policy for player 0 , and red edges for player 1 . Improving edges for player 0 are shown in dotted green and the other current non-policy and non-strategy edges are shown in dotted black. (Coloured edges show in bold on monochromatic printing.) Whereas Figure 1 shows the initial strategy for parity game $G_{3}$, Figure 8 shows the initial policy for the MDP $M_{3}$. The sequence $P_{n}$ of improving switches we constructed in $G_{n}$ starting at the initial strategy and following Algorithm 2 using the ordering (1) will be shown to correspond to an identical sequence of improving switches in $M_{n}$ using policy iteration and the same ordering.

Lemma 8 The Markov chains obtained by any policy reach the sink $t$ almost surely (i.e. the sink $t$ is the single irreducible recurrent class).

It is not too hard to see that the absolute value of all nodes corresponding to policies are bounded by $\varepsilon^{-1}$. More formally we have:

Lemma 9 Let $P=\left\{s, g_{*}, h_{*}\right\}$ be the set of nodes with priorities. For a subset $S \subseteq P$, let $\sum(S)=\sum_{v \in S}\langle v\rangle$. For non-empty subsets $S \subseteq P$, let $v_{S} \in S$ be the node with the largest priority in $S$.

1. $\left|\sum(S)\right|<N+1$ and $\varepsilon \cdot\left|\sum(S)\right|<1$ for every subset $S \subseteq P$, and
2. $\left|v_{S}\right|<\left|v_{S^{\prime}}\right|$ implies $\left|\sum(S)\right|<\left|\sum\left(S^{\prime}\right)\right|$ for non-empty subsets $S, S^{\prime} \subseteq P$.

### 3.3 Lower Bound Proof

In this section we prove the following theorem. The initial policy, $\left\{a_{*}^{0}, b_{*}^{0}, c_{*}^{0}, d_{*}^{1}, e_{*}^{0}\right\}$, corresponds exactly to the initial strategy for the parity game $G_{n}$.


Fig. 8 (Relaxed) Markov Decision Process Lower Bound Graph

Theorem 6 The sequence $P_{n}$ of improving switches followed by policy iteration in $M_{n}$ from the initial policy to the terminal policy using the ordering (1) has length at least $2^{n}$, where $M_{n}$ has size $O(n)$.

We will show that the sequence $P_{n}$ in Theorem 6 is in one-to-one correspondence with the sequence $P_{n}$ in Theorem 4. We first show that Figure 6 indeed specifies the sets of improving switches for each phase $p$.

Lemma 10 Let $n>1$. The improving switches from policies that belong to the phases in Figure 5 are as specified in Figure 6.
Proof Let $n>1$ and $\sigma$ be a policy belonging to one of the phases w.r.t. $\mathfrak{b}$. Let $\mu=\mu(\mathfrak{b})$. Define $T_{i}=\left\langle h_{i}\right\rangle+\sum_{j>i, \mathfrak{b}_{j}=1}\left(\left\langle g_{j}\right\rangle+\left\langle h_{j}\right\rangle\right)$ and $S_{i}=\left\langle g_{i}\right\rangle+T_{i}$.

- Phase 1: The following holds:

$$
\begin{aligned}
& \operatorname{val}_{\sigma}\left(g_{i}\right)= \begin{cases}S_{i} & \text { if } \sigma\left(d_{i}\right)=1, \sigma\left(e_{i}\right)=1 \\
\varepsilon \cdot T_{i}+\left\langle g_{i}\right\rangle+(1-\varepsilon) \cdot\left(\langle s\rangle+S_{\mu}\right) & \text { if } \sigma\left(d_{i}\right)=1, \sigma\left(e_{i}\right)=0 \\
\varepsilon \cdot T_{i}+\left\langle g_{i}\right\rangle+(1-\varepsilon) \cdot S_{\mu} & \text { if } \sigma\left(d_{i}\right)=0, \sigma\left(e_{i}\right)=1 \\
\varepsilon \cdot T_{i}+\left\langle g_{i}\right\rangle+\frac{1-\varepsilon}{2} \cdot\left(\langle s\rangle+2 \cdot S_{\mu}\right) & \text { if } \sigma\left(d_{i}\right)=0, \sigma\left(e_{i}\right)=0\end{cases} \\
& \operatorname{val}_{\sigma}\left(e_{i}\right)= \begin{cases}\langle s\rangle+S_{\mu} & \text { if } \sigma\left(e_{i}\right)=0\end{cases}
\end{aligned}
$$

It is easy to see that $e_{i}^{1}$ are improving switches as $T_{i}>S_{\mu}$. It can furthermore be easily observed that $d_{i}^{0}$ is an improving switch for $\mathfrak{b}_{i}=0, i>\nu(\mathfrak{b})$ iff $\sigma\left(e_{\nu(\mathfrak{b})}\right)=1$ and $\sigma\left(b_{\nu(\mathfrak{b})}\right)=1$.

- Phase 2: Similar to phase 5.
- Phase 3: Similar to phase 4.
- Phase 4: The following holds:

$$
\begin{aligned}
\operatorname{val}_{\sigma}\left(b_{i}\right) & = \begin{cases}S_{\mu} & \text { if } i \geq \mu \\
S_{1} & \text { otherwise }\end{cases} \\
\operatorname{val}_{\sigma}\left(F_{i}\right) & = \begin{cases}T_{i} & \text { if } \mathfrak{b}_{i}^{\oplus}=1 \\
\varepsilon \cdot \mathcal{O}(1)+\frac{1}{2} \cdot\left(\langle s\rangle+S_{\mu}+\operatorname{val}_{\sigma}\left(b_{i-1}\right)\right) & \text { if } \mathfrak{b}_{i}^{\mu}=0 \wedge \sigma\left(d_{i}\right) \neq F_{i} \\
\varepsilon \cdot \mathcal{O}(1)+\langle s\rangle+S_{\mu} & \text { otherwise }\end{cases}
\end{aligned}
$$

We conclude: $\operatorname{val}_{\sigma}\left(F_{i}\right)-\operatorname{val}_{\sigma}\left(b_{i-1}\right)=$

$$
\begin{cases}T_{i}-\operatorname{val}_{\sigma}\left(b_{i-1}\right) \geq T_{i}-S_{\mu}>0 & \text { if } S_{\mu} \\ \varepsilon \cdot \mathcal{O}(1)+\frac{1}{2} \cdot\left(\langle s\rangle+S_{\mu}-\operatorname{val}_{\sigma}\left(b_{i-1}\right)\right) \geq \varepsilon \cdot \mathcal{O}(1)+\frac{1}{2} \cdot\langle s\rangle>0 & \text { if } \mathfrak{b}_{i}^{\mu}=0 \wedge \sigma\left(d_{i}\right) \neq F_{i} \\ \varepsilon \cdot \mathcal{O}(1)+\langle s\rangle+S_{\mu}-\operatorname{val}_{\sigma}\left(b_{i-1}\right) \geq \varepsilon \cdot \mathcal{O}(1)+\langle s\rangle>0 & \text { otherwise }\end{cases}
$$

- Phase 5: The following holds:

$$
\begin{aligned}
& \operatorname{val}_{\sigma}\left(g_{i}\right)= \begin{cases}S_{i} & \text { if } \mathfrak{b}_{i}^{\oplus}=1 \\
\varepsilon \cdot \mathcal{O}(1)+\left\langle g_{i}\right\rangle+\langle s\rangle+S_{\mu} & \text { otherwise }\end{cases} \\
& \operatorname{val}_{\sigma}\left(a_{i}\right)= \begin{cases}S_{i} & \text { if } \mathfrak{b}_{i}^{\oplus}=1 \wedge(i>\mu \vee \sigma( \\
\operatorname{val}_{\sigma}\left(a_{i+1}\right) & \text { if } \left.i=\mu \wedge \sigma\left(a_{\mu}\right) \neq g_{\mu}\right) \\
S_{\mu} & \text { if } i<\mu \wedge \sigma\left(a_{i}\right) \neq g_{i} \\
\varepsilon \cdot \mathcal{O}(1)+\left\langle g_{i}\right\rangle+\langle s\rangle+S_{\mu} & \text { if } i<\mu \wedge \sigma\left(a_{i}\right)=g_{i}\end{cases}
\end{aligned}
$$

By computing the difference we again see that the improving switches are as described.

Now Lemma 6 becomes applicable again and the proof of Theorem 6 follows.

4 Linear Program Simplex Method Lower Bound
In this section we use the well known transformation from MDPs to linear programs to obtain an exponential lower bound for the simplex method using Cunningham's rule.

### 4.1 Linear Programs and the Simplex Method

We briefly give a few basic definitions and state our notation. For more information the reader is referred to any standard linear programming text, such as [4]. Given an $m$ by $n$ matrix $A$ with $m \leq n$, an $n$-vector $c$ and $m$-vector $b$ we consider the primal linear program in the standard form

$$
\begin{array}{ll}
\max & c^{T} x \\
\text { s.t. } & A x=b \\
& x \geq 0
\end{array}
$$

Let $B$ and $N$ be a partition of the indices $\{1,2, \ldots, n\}$ such that $|B|=m$. We denote by $A_{B}$ the submatrix of $A$ with columns indexed by $B$, and $x_{B}$ the $m$-subvector of $x$ with indices in $B$. We say that $B$ is a feasible basis if $A_{B}$ is non-singular and

$$
x_{B}=A_{B}^{-1} b \geq 0
$$

A corresponding basic feasible solution(BFS) $x$ is obtained if we extend $x_{B}$ to $x$ by setting $x_{N}=0$.

We call $c^{T} x$ the objective function. A pivot from a feasible basis $B$ is defined by a pair of indices $i \in B$ and $j \in N$ for which $B \backslash\{i\} \cup\{j\}$ is also a feasible basis. If the corresponding BFS are $x$ and $x^{\prime}$, then the pivot is improving if $c^{T} x \leq c^{T} x^{\prime}$. If the inequality is strict then we call the pivot step nondegenerate otherwise it is called degenerate. An optimal basis is one for which the corresponding BFS maximizes the objective function. A deterministic pivot rule gives a unique pivot pair for every non-optimal feasible basis.

The simplex method starts from a given feasible basis and applies improving pivots until an optimal basis is obtained. The sequence of pivots depends on the specific pivot rule, and care must be taken to ensure that it does not cycle if there are degenerate pivots.

The corresponding dual LP is written

$$
\begin{aligned}
& \min b^{T} y \\
& \text { s.t. } A^{T} y \geq c .
\end{aligned}
$$

4.2 Markov Decision Processes as LPs

Optimal policies for MDPs that satisfy the unichain condition can be found by solving the following primal linear program (see, e.g., [22].)

$$
\begin{array}{ll}
\max & \sum_{(u, v) \in E_{0}} r(u, v) x(u, v) \\
\text { s.t. } & \sum_{(u, v) \in E} x(u, v)-\sum_{(v, w) \in E_{0},(w, u) \in E_{R}} p(w, u) x(v, w)=1, u \in V_{0}  \tag{P}\\
& x(u, v) \geq 0 \quad, \quad(u, v) \in E_{0}
\end{array}
$$

The variable $x(u, v)$, for $(u, v) \in E_{0}$, stands for the probability (frequency) of using the edge (action) $(u, v)$. The constraints of the linear program are conservation constraints that state that the probability of entering a vertex $u$ is equal to the probability of exiting $u$. It is not difficult to check that the BFS's of ( P ) correspond directly to policies of the MDP. For each policy $\sigma$ we can define a feasible setting of primal variables $x(u, v)$, for $(u, v) \in E_{0}$, such that $x(u, v)>0$ only if $\sigma(u)=(u, v)$. Conversely, for every BFS $x(u, v)$ we can define a corresponding policy $\sigma$. It is well known that the policy corresponding to an optimal BFS of $(\mathrm{P})$ is an optimal policy of the MDP. (See, e.g., [22].)

Lemma 11 Let $\sigma$ be a policy and $x(u, v)$ be a corresponding BFS. Then the following holds:

$$
\sum_{u \in E_{0}} \operatorname{val}_{\sigma}(u)=\sum_{(u, v) \in E_{0}} r(u, v) x(u, v)
$$

The dual linear program (for unichain MDPs) is:

$$
\begin{align*}
& \min \sum_{u \in V} y(u) \\
& \text { s.t. } y(u)-\sum_{(v, w) \in E_{R}} p(v, w) y(w) \geq r(u, v) \quad, \quad(u, v) \in E_{0} \tag{D}
\end{align*}
$$

together with the condition that $y(u)$ sum up to 0 on the single irreducible recurrent class.

If $y^{*}$ is an optimal solution of (D), then $y^{*}(u)$, for every $u \in V_{0}$, is the value of $u$ under an optimal policy. An optimal policy $\sigma^{*}$ can be obtained by letting $\sigma^{*}(u)=(u, v)$, where $(u, v) \in E_{0}$ is an edge for which the inequality constraint in (D) is tight, i.e., $y(u)-\sum_{w:(v, w) \in E_{R}} p(v, w) y(w)=r(u, v)$. Such a tight edge is guaranteed to exist.

### 4.3 Policy Iteration and Simplex Method

A policy iteration algorithm with policy $\sigma$ that perform a single switch at each iteration - like Cunningham's rule - corresponds to a variation of the simplex method where the selection rule behaves like $\sigma$. Indeed $\sigma$ gives rise to a feasible solution $x(u, v)$ of the primal linear program (P). We use $\sigma$ to define a Markov chain and let $x(u, v)$ be the 'steady-state' probability that the edge (action) $(u, v)$ is used. In particular, if $\sigma(u) \neq v$, then $x(u, v)=0$.

We can also view the values corresponding to $\sigma$ as settings of the variables $y(u)$ of the dual linear program (D). By linear programming duality, if
$y(u)$ is feasible then $\sigma$ is an optimal policy. It is easy to check that an edge $\left(u, v^{\prime}\right) \in E_{0}$ is an improving switch if and only if the dual constraint corresponding to $\left(u, v^{\prime}\right)$ is violated. Furthermore, replacing the edge $(u, v)$ by the edge ( $u, v^{\prime}$ ) corresponds to a pivoting step, with a non-negative reduced cost, in which the column corresponding to $\left(u, v^{\prime}\right)$ enters the basis, while the column corresponding to $(u, v)$ leaves the basis.

### 4.4 Lower Bound Construction

Let $n>1$. The variables of the LP correspond to the edges $E_{0}$ controlled by player 0 , i.e. we have $10(n-1)$ variables

$$
\left\{a_{i}^{1}, a_{i}^{0}, c_{i}^{1}, c_{i}^{0}, d_{i}^{1}, d_{i}^{0} \mid 1<i \leq n\right\} \cup\left\{b_{i}^{1}, b_{i}^{0} \mid 1<i<n\right\} \cup\left\{e_{i}^{1}, e_{i}^{0} \mid 1 \leq i \leq n\right\} .
$$

The LP has $5(n-1)$ constraints, corresponding to the nodes $V_{0}$ controlled by player 0, and labelled

$$
\left\{a_{i}, c_{i}, d_{i} \mid 1<i \leq n\right\} \cup\left\{b_{i} \mid 1<i<n\right\} \cup\left\{e_{i} \mid 1 \leq i \leq n\right\} .
$$

The linear program is defined as follows for each $n \geq 3$ (non-existent variables are assumed to be zero):

$$
\begin{aligned}
& L P_{n}: \\
& \max \sum_{i=1}^{n}\left(\left(a_{i}^{1}+b_{i}^{1}+c_{i}^{1}\right)\left(\Omega\left(g_{i}\right)+\varepsilon \Omega\left(h_{i}\right)\right)+\varepsilon\left(d_{i}^{1}+e_{i}^{1}\right) \Omega\left(h_{i}\right)+e_{i}^{0} \Omega(s)\right)
\end{aligned}
$$

subject to:

```
\(\left(a_{2}\right) \quad a_{2}^{0}+a_{2}^{1}=1+\varepsilon\left(b_{2}^{0}+c_{2}^{0}+d_{2}^{0}+e_{1}^{1}\right)\)
\(\left(a_{i}\right) \quad a_{i}^{0}+a_{i}^{1}=1+a_{i-1}^{0}+\varepsilon\left(a_{i-1}^{1}+b_{i-1}^{1}+c_{i-1}^{1}+d_{i-1}^{1}+e_{i-1}^{1}\right) \quad 3 \leq i \leq n\)
( \(b_{i}\) ) \(b_{i}^{0}+b_{i}^{1}=1+b_{i+1}^{0}+d_{i+1}^{0} \quad 2 \leq i<n\)
( \(c_{i}\) ) \(c_{i}^{0}+c_{i}^{1}=1+c_{i+1}^{0} \quad 2 \leq i<n\)
( \(c_{n}\) ) \(c_{n}^{0}+c_{n}^{1}=1+\sum_{j=1}^{n} e_{j}^{0}\)
\(\left(d_{i}\right) \quad d_{i}^{0}+d_{i}^{1}=1+\frac{1-\varepsilon}{2}\left(a_{i}^{1}+b_{i}^{1}+c_{i}^{1}+d_{i}^{1}+e_{i}^{1}\right) \quad 2 \leq i \leq n\)
\(\left(e_{1}\right) \quad e_{1}^{0}+e_{1}^{1}=1+(1-\varepsilon)\left(b_{2}^{0}+c_{2}^{0}+d_{2}^{0}+e_{1}^{1}\right)\)
\(\left(e_{i}\right) \quad e_{i}^{0}+e_{i}^{1}=1+\frac{1-\varepsilon}{2}\left(a_{i}^{1}+b_{i}^{1}+c_{i}^{1}+d_{i}^{1}+e_{i}^{1}\right) \quad 2 \leq i \leq n\)
```

All variables non-negative
Note that the size of $L P_{n}$ is linear in $n$. Depending on the context we let $L P_{n}$ denote both the linear program and the polytope defined by its constraints. For small values of $n$ explicit constructions of $L P_{n}$ and it dual are available online [11].

We now use the correspondence between the parity game $G_{n}$, the MDP $M_{n}$ and linear program $L P_{n}$ to get a lower bound for the simplex method using Cunningham's rule. The initial strategy for $G_{n}$ and initial policy for $M_{n},\left\{a_{*}^{0}, b_{*}^{0}, c_{*}^{0}, d_{*}^{1}, e_{1}^{1}, e_{*>1}^{0}\right\}$, defines a starting basis for $L P_{n}$. We construct a path on the polytope $L P_{n}$ from this starting basis using the least recently considered rule with ordering (1). Using this construction, Theorem 6 implies that the path generated will be in one-to-one correspondence with the sequence
$P_{n}$ of improving switches generated in $M_{n}$ (and hence $G_{n}$ ). We observe that the objective function strictly increases with each pivot due to Lemma 11 and Lemma 7. Therefore we have the following result.

Theorem 7 The pivot path $P_{n}$ for $L P_{n}$ from the starting basis to the optimum basis followed by the least recently considered rule with ordering (1) has length at least $2^{n}$. The objective function strictly increases with each pivot.

## 5 Acyclic Unique Sink Orientations

Our final result concerns acyclic unique sink orientations (AUSOs), which are abstractions of various optimization problems including linear programming, linear complementarity and binary payoff games. In this section we extend our exponential lower bound to finding the sink of an AUSO using the least considered rule. For background information on AUSOs, see [23,12,13].

### 5.1 Definitions and previous results

AUSOs can be defined on arbitrary polytopes, but here we consider only hypercubes. An AUSO on a $n$-dimensional hypercube is an orientation of its edges that is acyclic and such that every face of the hypercube has a unique sink (vertex of outdegree 0 ). The goal is to find the unique sink of the AUSO. There is at present no known polynomial time algorithm for doing this, nor is it known to be NP-hard.

A natural class of algorithms to find the sink of an AUSO are path following algorithms. Such an algorithm would start at any vertex $v$ of the hypercube and repeatedly choose an outgoing edge according to some rule until the unique sink is located. Each edge of the path corresponds to flipping one bit of the current vertex.

There is a very natural analogy between path following algorithms and pivoting in linear programming. Pivot rules for LPs therefore have natural analogues for AUSOs and a full discussion of this is contained in [1]. In particular the least recently considered rule can be adapted to give a path following algorithm to find the unique sink of an $n$-cube AUSO starting at any given vertex as follows.

We define, for each $i=1,2, \ldots, n$, the variable $v_{i}$ to denote a flip of bit $i$ from 0 to 1 , and variable $v_{n+i}$ to denote a flip of bit $i$ from 1 to 0 . We may now arrange the $2 n$ variables $v_{1}, v_{2}, \ldots, v_{2 n}$ in any cyclic order. For any vertex of the AUSO that is not the sink we find the first allowable flip in this cyclic order starting at the last chosen $v_{i}$.

Suppose we are given a polytope and an objective function that is not constant on any edge of the polytope. We can can then orient each edge in the direction of increasing objective function. The corresponding directed graph on the skeleton of the polytope can be shown to be an AUSO. The converse is not always true. Indeed, we call an AUSO realizable if there is a polytope
and objective function that induces a directed graph on its skeleton which is graph isomorphic to the AUSO. Not all AUSOs are realizable.

The Klee-Minty examples[16] are realizations of AUSOs on hypercubes, so exponential lower bounds for most of the non-history based deterministic LP pivot rules immediately give similar bounds for AUSOs. Gärtner[12] gives an $\exp \left(2 n^{1 / 2}\right)$ lower bound for random facet algorithms and Matousek and Szabó[18] give an $\exp \left(\right.$ const $\left.n^{1 / 3}\right)$ lower bound for the random edge rule on AUSOs. In the next subsection we derive an exponential lower bound for finding the sink of a realizable AUSO using a path following algorithm based on Cunningham's rule.

### 5.2 Lower Bound Construction

As we saw in Section 2 there is a direct relationship between binary parity games and oriented hypercubes. Each vertex $v$ corresponds to a strategy $\sigma$ for player 0 . A partial orientation of the hypercube's edges is given by the notion of improving switches: the orientation goes from $\sigma$ to $\sigma^{\prime}$ iff there is a game edge $e$ belonging to player 0 s.t. $\sigma^{\prime}=\sigma[e]$ and $e$ is $\sigma$-improving. This is only a partial orientation since there are strategies for the parity game that are not used in the lower bound construction and for which the notion of improving is not well defined. For each $n \geq 3$ we denote this partially oriented $n$-cube $H_{n}$.

Our goal is to embed $H_{n}$ into an AUSO $A_{n}$ whose edges orientations are consistent with those already set in $H_{n}$. Note that this is not a trivial operation, as even for $n=3$ it can be readily verified that there are partial acyclic orientations of the 3-cube which are USOs on every complete face but do not embed into an AUSO. We will achieve this embedding via the linear programming formulation of the last section, achieving the stronger result that the AUSO is realizable.

Lemma $12 L P_{n}$ is a realization an AUSO $A_{n}$ which is consistent with the edge orientations of $H_{n}$.

Proof We argued in Section 4.2 that each basic feasible solution of $L P_{n}$ corresponded directly to a policy of the corresponding MDP $M_{n}$, and hence to strategy of the parity game $G_{n}$. Also each edge of $L P_{n}$ corresponded to a switch in $M_{n}$ and hence to an edge owned by player 0 in $G_{n}$. It follows that the vertices and edges of $L P_{n}$ and $H_{n}$ are in one-to-one correspondence. Since $H_{n}$ is an $n$-cube so is the skeleton of $L P_{n}$. By applying symbolic dual perturbation if necessary to resolve ties in the objective function (see, e.g. [4]) $A_{n}$ can be oriented to give an AUSO, which we denote as $A_{n}$. By definition $A_{n}$ is realizable.

According to Theorem 7 the objective function of $L P_{n}$ is non-constant on every edge and increases in the direction corresponding to an improving switch in $M_{n}$. However improving switches in $M_{n}$ correspond to improving edges in $G_{n}$. The edges of $H_{n}$ were directed in the same way as these improving edges. Since dual perturbation will not change the direction of any edges for which
the objective function is strictly increasing, the directed edges in $H_{n}$ maintain their directions in $A_{n}$. The lemma follows.

The lemma implies that not only is $H_{n}$ an AUSO but stronger properties, such as the Holt-Klee condition and the shelling property, also hold (see, e.g. [2]). For small values of $n$ the AUSOs $A_{n}$ are available online [11].

We may now apply the least recently considered rule to $A_{n}$ as described in Section 5.1. The starting basis of $L P_{n}\left\{a_{*}^{0}, b_{*}^{0}, c_{*}^{0}, d_{*}^{1}, e_{*}^{0}\right\}$ defines a starting vertex of $A_{n}$. We construct a path from this vertex using the least recently considered rule with ordering (1). Using this construction, Theorem 7 implies that the path generated will be in one-to-one correspondence with the path $P_{n}$ generated in $L P_{n}$.

Theorem 8 The directed path $P_{n}$ in $A_{n}$ from the starting vertex to the unique sink followed by the least recently considered rule with ordering (1) has length at least $2^{n}$.

## 6 Conclusions

We have shown in this paper that Cunningham's least considered rule can lead to exponential worst case behaviour in parity and other games, Markhov decision processes, linear programs and AUSOs. This appears to be the first such result for a history based rule for AUSOs. However Cunningham's rule was in fact first proposed for the network simplex method. The LPs presented in Section 4 are not networks, but are structurally remarkably close to them. Our first open problem would be to extend our results to network LPs.

As remarked in the introduction, Zadeh's rule has recently been shown to have superpolynomial worst case behavior on linear programs. The parity games behind these LPs were not binary, so it does not immediately follow that Zadeh's rule has similar behaviour on AUSOs. This is a second open problem. More generally it is of interest determine whether all of the history based rules mentioned in [1] have exponential behaviour on AUSOs.

Finally it would be of interest to study the AUSOs presented in Section 5 to see if they have a simple enough structure to be stated explicitly. If so it may be possible to prove the existence of exponentially long paths in a simpler manner than that done in Section 2.

## References

1. Aoshima, Y., Avis, D., Deering, T., Matsumoto, Y., Moriyama, S.: On the existence of hamiltonian paths for history based pivot rules on acyclic unique sink orientations of hypercubes. Discrete Applied Mathematics 160(15), 2104-2115 (2012)
2. Avis, D., Moriyama, S.: On Combinatorial Properties of Linear Programming Digraphs. In: D. Avis, D. Bremner, A. Deza (eds.) Polyhedral Computation, CRM Proceedings and Lecture Notes 48, pp. 1-13. AMS (2009)
3. Bertsekas, D.: Dynamic programming and optimal control, second edn. Athena Scientific (2001)
4. Chvátal, V.: Linear Programmming. W.H. Freeman (1983)
5. Cunningham, W.H.: Theoretical properties of the network simplex method. In: Mathematics of Operations Research, pp. 196-208 (1979)
6. Derman, C.: Finite state Markov decision processes. Academic Press (1972)
7. Emerson, E., Jutla, C.: Tree automata, $\mu$-calculus and determinacy. In: Proc. 32nd Symp. on Foundations of Computer Science, pp. 368-377. IEEE, San Juan (1991)
8. Friedmann, O.: An exponential lower bound for the latest deterministic strategy iteration algorithms. Logical Methods in Computer Science 7(3) (2011)
9. Friedmann, O.: A subexponential lower bound for Zadeh's pivoting rule for solving linear programs and games. In: IPCO, pp. 192-206 (2011)
10. Friedmann, O.: A subexponential lower bound for the least recently considered rule for solving linear programs and games. In: GAMES'2012. Naples, Italy (2012)
11. Friedmann, O.: (2013). http://tcswiki.com/?title=Cunningham's_Rule
12. Gärtner, B.: The random-facet simplex algorithm on combinatorial cubes. Random Structures \& Algorithms 20(3), 353-381 (2002)
13. Gärtner, B., Schurr, I.: Linear programming and unique sink orientations. In: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 749-757 (2006)
14. Howard, R.: Dynamic programming and Markov processes. MIT Press (1960)
15. Kalai, G.: A subexponential randomized simplex algorithm (extended abstract). In STOC, pp. 475-482 (1992)
16. Klee, V., Minty, G.J.: How Good is the Simplex Algorithm? In: O. Shisha (ed.) Inequalities III, pp. 159-175. Academic Press Inc., New York (1972)
17. Matousek, J., Sharir, M., Welzl, E.: A subexponential bound for linear programming. In: Symposium on Computational Geometry, pp. 1-8 (1992)
18. Matousek, J., Szabó, T.: Random edge can be exponential on abstract cubes. In: FOCS, pp. 92-100 (2004)
19. Melekopoglou, M., Condon, A.: On the complexity of the policy improvement algorithm for markov decision processes. INFORMS Journal on Computing 6(2), 188-192 (1994)
20. Post, I., Ye, Y.: The simplex method is strongly polynomial for deterministic markov decision processes. CoRR abs/1208.5083 (2012)
21. Puri, A.: Theory of hybrid systems and discrete event systems. Ph.D. thesis, University of California, Berkeley (1995). URL http://www.eecs.berkeley.edu/Pubs/TechRpts/ 1995/2950.html
22. Puterman, M.: Markov decision processes. Wiley (1994)
23. Szabó, T., Welzl, E.: Unique sink orientations of cubes. In: Proceedings of the 42 th FOCS, pp. 547-555 (2001)
24. Vöge, J., Jurdzinski, M.: A discrete strategy improvement algorithm for solving parity games. In: Proc. 12th Int. Conf. on Computer Aided Verification, CAV'00, LNCS, vol. 1855, pp. 202-215. Springer (2000)
25. Zadeh, N.: What is the worst case behavior of the simplex algorithm. In: Polyhedral Computation, pp. 131-143. American Mathematical Society $(2009,1980)$
26. Zwick, U., Paterson, M.: The complexity of mean payoff games on graphs. Theoretical Computer Science 158(1-2), 343-359 (1996). DOI http://dx.doi.org/10.1016/ 0304-3975(95)00188-3

[^0]:    ${ }^{1} x \equiv_{2} y$ if and only if $x$ and $y$ are congruent $\bmod 2$

