# The simplex algorithm and the Hirsch conjecture: Lecture 3 

Thomas Dueholm Hansen

MADALGO \& CTIC Summer School
August 11, 2011
madaLGO =--ュ-ュ
CENTER FOR MASSIVE DATA ALGORITHMICS


## Overview

- Lecture 1:
- Introduction to linear programming and the simplex algorithm.
- Pivoting rules.
- The RandomFacet pivoting rule.
- Lecture 2:
- The Hirsch conjecture.
- Introduction to Markov decision processes (MDPs).
- Upper bound for the LargestCoefficient pivoting rule for MDPs.
- Lecture 3 :
- Lower bounds for pivoting rules utilizing MDPs. Example: Bland's rule.
- Lower bound for the RandomEdge pivoting rule.
- Abstractions and related problems.


## Discounted Markov decision processes



$$
J=\left[\begin{array}{lll}
1 & 0 & 0 \\
1 & 0 & 0 \\
0 & 1 & 0 \\
0 & 1 & 0 \\
0 & 0 & 1 \\
0 & 0 & 1
\end{array}\right]
$$

$$
P=\left[\begin{array}{ccc}
0 & \frac{1}{2} & \frac{1}{2} \\
1 & 0 & 0 \\
1 & 0 & 0 \\
\frac{1}{2} & \frac{1}{4} & \frac{1}{4} \\
0 & 1 & 0 \\
0 & \frac{1}{3} & \frac{2}{3}
\end{array}\right]
$$

$$
c=\left[\begin{array}{c}
7 \\
3 \\
-4 \\
2 \\
5 \\
-10
\end{array}\right]
$$

- A discounted MDP with $n$ states and a total of $m$ actions can be represented by:
- A discount factor $\gamma<1$.
- A zero-one matrix $J \in\{0,1\}^{m \times n}$, with $J_{a, i}=1$ iff $a \in A_{i}$.
- A stochastic matrix $P \in \mathbb{R}^{m \times n}$.
- A reward vector $c \in \mathbb{R}^{m}$.


## Discounted Markov decision processes



$$
I=\left[\begin{array}{lll}
1 & 0 & 0 \\
& & \\
0 & 1 & 0 \\
0 & 0 & 1
\end{array}\right]
$$

$$
P_{\pi}=\left[\begin{array}{ccc}
0 & \frac{1}{2} & \frac{1}{2} \\
& & \\
\frac{1}{2} & \frac{1}{4} & \frac{1}{4} \\
0 & 1 & 0
\end{array}\right]
$$

$$
c_{\pi}=\left[\begin{array}{l}
7 \\
2 \\
5
\end{array}\right]
$$

- A discounted MDP with $n$ states and a total of $m$ actions can be represented by:
- A discount factor $\gamma<1$.
- A zero-one matrix $J \in\{0,1\}^{m \times n}$, with $J_{a, i}=1$ iff $a \in A_{i}$.
- A stochastic matrix $P \in \mathbb{R}^{m \times n}$.
- A reward vector $c \in \mathbb{R}^{m}$.
- A policy $\pi$ is a choice of an action from each state. $\pi$ defines a Markov chain with rewards $\left(P_{\pi}, c_{\pi}\right)$.


## The stopping condition



- The discount factor $\gamma<1$ was introduced because the expected total reward $\sum_{k=0}^{\infty} b^{T} P_{\pi}^{k} c_{\pi}$, where $b$ is some initial distribution, may not converge.
- For every action a, $(1-\gamma)$ may be interpreted as the probability of moving to a terminal state $t$.


## The stopping condition



- The discount factor $\gamma<1$ was introduced because the expected total reward $\sum_{k=0}^{\infty} b^{T} P_{\pi}^{k} c_{\pi}$, where $b$ is some initial distribution, may not converge.
- For every action $a,(1-\gamma)$ may be interpreted as the probability of moving to a terminal state $t$.
- To ensure convergence it is enough to satisfy the following condition:
- Stopping condition: The terminal state is eventually reached with probability 1 from all states.
- Let $P_{\pi} \in \mathbb{R}^{n \times n}$ be a matrix with non-negative entries such that each row sums to at most 1 .
- The difference between 1 and the sum of the a'th row is the probability of moving to the terminal state when using action a.
- Note that $\gamma P$, where $P$ is an $n \times n$ stochastic matrix, is a special case.


## The stopping condition

- Let $P_{\pi} \in \mathbb{R}^{n \times n}$ be a matrix with non-negative entries such that each row sums to at most 1 .
- The difference between 1 and the sum of the a'th row is the probability of moving to the terminal state when using action a.
- Note that $\gamma P$, where $P$ is an $n \times n$ stochastic matrix, is a special case.
- If the stopping condition is satisfied, each row of $P_{\pi}^{n}$ sums to less than 1 , and $P_{\pi}^{k} \rightarrow 0$ for $k \rightarrow \infty$.
- It is again not difficult to show that:

$$
\left(I-P_{\pi}\right)^{-1}=\sum_{k=0}^{\infty} P_{\pi}^{k}
$$

## The stopping condition

- Let $P_{\pi} \in \mathbb{R}^{n \times n}$ be a matrix with non-negative entries such that each row sums to at most 1 .
- The difference between 1 and the sum of the a'th row is the probability of moving to the terminal state when using action a.
- Note that $\gamma P$, where $P$ is an $n \times n$ stochastic matrix, is a special case.
- If the stopping condition is satisfied, each row of $P_{\pi}^{n}$ sums to less than 1 , and $P_{\pi}^{k} \rightarrow 0$ for $k \rightarrow \infty$.
- It is again not difficult to show that:

$$
\left(I-P_{\pi}\right)^{-1}=\sum_{k=0}^{\infty} P_{\pi}^{k}
$$

- Everything that was said in lecture 2 about discounted Markov chains with rewards remain true if $\gamma P$ is replaced by $P$, where $P$ satisfies the stopping condition.


## Markov decision processes



$$
J=\left[\begin{array}{lll}
1 & 0 & 0 \\
1 & 0 & 0 \\
0 & 1 & 0 \\
0 & 1 & 0 \\
0 & 0 & 1 \\
0 & 0 & 1
\end{array}\right]
$$

$$
P=\left[\begin{array}{ccc}
0 & \frac{1}{2} & \frac{1}{2} \\
0 & 1 & 0 \\
\frac{2}{3} & 0 & 0 \\
\frac{1}{2} & \frac{1}{4} & 0 \\
0 & 1 & 0 \\
0 & 0 & 0
\end{array}\right]
$$

$$
c=\left[\begin{array}{c}
7 \\
3 \\
-4 \\
2 \\
5 \\
0
\end{array}\right]
$$

- An MDP with $n$ states and a total of $m$ actions can be represented by:
- A zero-one matrix $J \in\{0,1\}^{m \times n}$, with $J_{a, i}=1$ iff $a \in A_{i}$.
- A matrix $P \in \mathbb{R}^{m \times n}$ with non-negative entries and each row summing to at most 1.
- A reward vector $c \in \mathbb{R}^{m}$.


## Markov decision processes



$$
I=\left[\begin{array}{lll}
1 & 0 & 0 \\
& & \\
0 & 1 & 0 \\
0 & 0 & 1
\end{array}\right]
$$

$$
P_{\pi}=\left[\begin{array}{ccc}
0 & \frac{1}{2} & \frac{1}{2} \\
& & \\
\frac{1}{2} & \frac{1}{4} & 0 \\
0 & 1 & 0
\end{array}\right]
$$

[^0]- An MDP with $n$ states and a total of $m$ actions can be represented by:
- A zero-one matrix $J \in\{0,1\}^{m \times n}$, with $J_{a, i}=1$ iff $a \in A_{i}$.
- A matrix $P \in \mathbb{R}^{m \times n}$ with non-negative entries and each row summing to at most 1.
- A reward vector $c \in \mathbb{R}^{m}$.
- An MDP satisfies the stopping condition if all policies $\pi$ satisfy the stopping condition. For simplicity, we will generally assume that MDPs satisfy the stopping condition.


## Markov decision processes

- Every policy $\pi$ defines value and flux vectors:

$$
v_{\pi}=\left(I-P_{\pi}\right)^{-1} c_{\pi} \quad x_{\pi}^{T}=e^{T}\left(I-P_{\pi}\right)^{-1}
$$

where $\left(I-P_{\pi}\right)^{-1}=\sum_{k=0}^{\infty} P_{\pi}^{k}$.

- The value of state $i,\left(v_{\pi}\right)_{i}$, is the expected total reward accumulated when starting there.


## Markov decision processes

- Every policy $\pi$ defines value and flux vectors:

$$
v_{\pi}=\left(I-P_{\pi}\right)^{-1} c_{\pi} \quad x_{\pi}^{T}=e^{T}\left(I-P_{\pi}\right)^{-1}
$$

where $\left(I-P_{\pi}\right)^{-1}=\sum_{k=0}^{\infty} P_{\pi}^{k}$.

- The value of state $i,\left(v_{\pi}\right)_{i}$, is the expected total reward accumulated when starting there.
- A policy $\pi^{*}$ is optimal if it maximizes the values of all states: $v_{\pi^{*}} \geq v_{\pi}$ for all $\pi$.


## Markov decision processes

- An optimal policy can be found by solving a linear program:

$$
\text { (P) } \begin{array}{rr}
\max & c^{T} x \\
\text { s.t. } & (J-P)^{T} x \\
& \\
x & \geq 0
\end{array} \quad(D) \quad \begin{gathered}
\min \\
\text { s.t. }
\end{gathered} \quad(J-P) y \geq c
$$

## Markov decision processes

- An optimal policy can be found by solving a linear program:

$$
\text { (P) } \begin{align*}
& \max  \tag{C}\\
& \text { s.t. } \\
&
\end{align*} \quad(J-P)^{T_{x} x}=e \quad \text { (D) } \begin{gathered}
c^{T} x \\
x \geq 0
\end{gathered} \quad \begin{gathered}
e^{T} y \\
\text { s.t. }
\end{gathered} \quad(J-P) y \geq
$$

- There is a one-to-one correspondence between policies and basic feasible solutions of the primal LP $(P)$.


## Markov decision processes

- An optimal policy can be found by solving a linear program:

$$
\text { (P) } \begin{align*}
& \max  \tag{c}\\
& \text { s.t. } \\
&
\end{align*} \quad(J-P)^{T_{x} x}=e \quad \text { (D) } \begin{gathered}
\min ^{T} \begin{array}{c}
e^{T} y \\
\text { s.t. } \\
x
\end{array} \quad(J-P) y \geq
\end{gathered}
$$

- There is a one-to-one correspondence between policies and basic feasible solutions of the primal LP (P).
- The reduced cost vector, i.e. the coefficients of a tableau, $\bar{c}^{\pi} \in \mathbb{R}^{m}$ for a policy $\pi$ is defined by:

$$
\forall i \in S, \forall a \in A_{i}: \quad \bar{c}_{a}^{\pi}=\left(c_{a}+P_{a} v_{\pi}\right)-\left(v_{\pi}\right)_{i}
$$

- $\bar{c}_{a}^{\pi}$ is the improvement over the current value by using a for one step w.r.t. $v_{\pi}$.


## Markov decision processes

- If $\bar{c}_{a}^{\pi}>0$ we say that $a$ is an improving switch w.r.t. $\pi$. l.e., $a \in A_{i}$ is an improving switch iff:

$$
\left(v_{\pi}\right)_{i}<c_{a}+P_{a} v_{\pi}
$$

## Lemma (Howard (1960))

Let $\pi^{\prime}$ be obtained from $\pi$ by jointly performing any non-empty set of improving switches. Then $v_{\pi^{\prime}} \geq v_{\pi}$ and $v_{\pi^{\prime}} \neq v_{\pi}$.

Lemma (Howard (1960))
A policy $\pi$ is optimal iff there are no improving switches.

## Markov decision processes

## Function PolicyIteration ( $\pi$ )

while $\exists$ improving switch w.r.t. $\pi$ do
Update $\pi$ by performing improving switches
return $\pi$

- The simplex algorithm applied to the primal LP $(P)$ is a special case of PolicyIteration.


## Example: A simple MDP

- Notation for graphical representation:
- Circles are states.
- Diamond-shaped vertices are rewards.
- Squares are randomization vertices.
- A policy $\pi$ is shown as bold blue arrows.



## Example: A simple MDP

- Notation for graphical representation:
- Circles are states.
- Diamond-shaped vertices are rewards.
- Squares are randomization vertices.
- A policy $\pi$ is shown as bold blue arrows.
- States and randomization vertices are labelled by corresponding values of $\pi$.



## Example: A simple MDP

- Notation for graphical representation:
- Circles are states.
- Diamond-shaped vertices are rewards.
- Squares are randomization vertices.
- A policy $\pi$ is shown as bold blue arrows.
- States and randomization vertices are labelled by corresponding values of $\pi$.
- Improving switches are indicated
 by red arrows.


## Example: A simple MDP

- Notation for graphical representation:
- Circles are states.
- Diamond-shaped vertices are rewards.
- Squares are randomization vertices.
- A policy $\pi$ is shown as bold blue arrows.
- States and randomization vertices are labelled by corresponding values of $\pi$.
- Improving switches are indicated
 by red arrows.


## Example: A simple MDP

- Notation for graphical representation:
- Circles are states.
- Diamond-shaped vertices are rewards.
- Squares are randomization vertices.
- A policy $\pi$ is shown as bold blue arrows.
- States and randomization vertices are labelled by corresponding values of $\pi$.
- Improving switches are indicated
 by red arrows.


## Example: A simple MDP

- Notation for graphical representation:
- Circles are states.
- Diamond-shaped vertices are rewards.
- Squares are randomization vertices.
- A policy $\pi$ is shown as bold blue arrows.
- States and randomization vertices are labelled by corresponding values of $\pi$.
- Improving switches are indicated
 by red arrows.


## From MDP to LP


$\max -1+2 x_{1}-2 x_{3}-x_{5}$
s.t. $x_{2}=1-\frac{1}{3} x_{1}+\frac{2}{3} x_{3}+\frac{2}{3} x_{5}$

$$
x_{4}=2-x_{3}-x_{5}
$$

$$
x_{6}=1-x_{5}
$$

$$
x_{1}, x_{2}, x_{3}, x_{4}, x_{5}, x_{6} \geq 0
$$



## From MDP to LP


$\max 5-6 x_{2}+2 x_{3}+3 x_{5}$
s.t. $x_{1}=3-3 x_{2}+2 x_{3}+2 x_{5}$
$x_{4}=2-x_{3}-x_{5}$
$x_{6}=1-x_{5}$
$x_{1}, x_{2}, x_{3}, x_{4}, x_{5}, x_{6} \geq 0$


## From MDP to LP


$\max 9-6 x_{2}-2 x_{4}+x_{5}$
s.t. $\quad x_{1}=7-3 x_{2}-2 x_{4}$
$x_{3}=2-x_{4}-x_{5}$
$x_{6}=1-x_{5}$
$x_{1}, x_{2}, x_{3}, x_{4}, x_{5}, x_{6} \geq 0$
$x^{T}=(7,0,2,0,0,1)$

## From MDP to LP



$$
\begin{array}{ll}
\max & 10-6 x_{2}-2 x_{4}-x_{6} \\
\text { s.t. } & x_{1}=7-3 x_{2}-2 x_{4} \\
& x_{3}=1-x_{4}+x_{6} \\
& x_{5}=1-x_{6} \\
& x_{1}, x_{2}, x_{3}, x_{4}, x_{5}, x_{6} \geq 0
\end{array}
$$

$$
x^{T}=(7,0,1,0,1,0)
$$

## Lower bounds for pivoting rules utilizing MDPs

- To prove lower bounds for a pivoting rule for the simplex algorithm, we can prove lower bounds for the corresponding PolicyIteration algorithm for MDPs:
- There is a one-to-one correspondence between policies and basic feasible solutions of the primal LP $(P)$ for MDPs.
- The simplex algorithm for the primal LP $(P)$ is the special case of PolicyIteration, where only single improving switches are performed.


## Lower bounds for pivoting rules utilizing MDPs

- To prove lower bounds for a pivoting rule for the simplex algorithm, we can prove lower bounds for the corresponding PolicyIteration algorithm for MDPs:
- There is a one-to-one correspondence between policies and basic feasible solutions of the primal LP $(P)$ for MDPs.
- The simplex algorithm for the primal $\mathrm{LP}(P)$ is the special case of PolicyIteration, where only single improving switches are performed.
- We next construct an exponential lower bound for BLAND'S RULE as a warmup before sketching the $2^{\Omega\left(n^{1 / 4}\right)}$ lower bound for RandomEdge by Friedmann, Hansen and Zwick (2011).


## Bland's Rule

- Bland's Rule, Bland (1977)
- Always pick the available variable with the smallest index, both for entering and leaving the basis.


## Bland's Rule

- Bland's Rule, Bland (1977)
- Always pick the available variable with the smallest index, both for entering and leaving the basis.
- Bland's rule for MDPs
- Perform the improving switch a with the smallest index.
- There is always only one action that can be exchanged with $a$, namely the current action that originates from the same state as $a$.


## Bland's Rule

- Bland's Rule, Bland (1977)
- Always pick the available variable with the smallest index, both for entering and leaving the basis.
- Bland's rule for MDPs
- Perform the improving switch a with the smallest index.
- There is always only one action that can be exchanged with $a$, namely the current action that originates from the same state as $a$.
- When constructing a lower bound, we may pick a worst-case ordering of the indices.


## Lower bound construction

- We define a family of lower bound MDPs $G_{n}$ such that Bland's rule, and later RandomEdge, simulates an $n$-bit binary counter.


## Lower bound construction

- We define a family of lower bound MDPs $G_{n}$ such that Bland's Rule, and later RandomEdge, simulates an n-bit binary counter.
- We make use of exponentially growing rewards (and penalties): To get a higher reward the MDP is willing to sacrifice everything that has been built up so far.


## Lower bound construction

- We define a family of lower bound MDPs $G_{n}$ such that Bland's Rule, and later RandomEdge, simulates an n-bit binary counter.
- We make use of exponentially growing rewards (and penalties): To get a higher reward the MDP is willing to sacrifice everything that has been built up so far.
- Notation: Integer priority $p$ corresponds to reward $(-N)^{p}$, where $N=3 n+1$.

$$
\ldots<5<3<1<2<4<6<\ldots
$$



## Lower bound for Bland's RUle







## Lower bound for Bland's RULE



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}$

## Lower bound for Bland's RULE



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, \underline{a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}}$

## Lower bound for Bland's RUle



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}$

## Lower bound for Bland's RUle



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}$

## Lower bound for Bland's RUle



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}$

## Lower bound for Bland's RUle



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}$

## Lower bound for Bland's RUle



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}$

## Lower bound for Bland's RULE



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, \underline{a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}}$

## Lower bound for Bland's RUle



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, \underline{a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}}$

## Lower bound for Bland's RUle



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, \underline{a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}}$

## Lower bound for Bland's RUle



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}$

## Lower bound for Bland's RUle



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, \underline{a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}}$

## Lower bound for Bland's RUle



Order: $u_{1}^{1}, u_{1}^{2}, u_{2}^{1}, u_{2}^{2}, w_{1}^{1}, w_{1}^{2}, w_{2}^{1}, w_{2}^{2}, b_{1}^{0}, b_{2}^{0}, a_{1}^{0}, a_{2}^{0}, a_{1}^{1}, a_{2}^{1}, b_{1}^{1}, b_{2}^{1}$

## Lower bound for Bland's rule

- Let $k$ be the lowest unset bit. Incrementing the counter happened roughly through five phases:
(1) Make $b_{k}=1$.
(2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(3) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
(4) Make $a_{k}=1$.
(5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.
- Only the last part of the ordering, involving $a_{i}^{1}$ and $b_{i}^{1}$ edges, was important.
- To implement a lower bound for RandomEdge we start out with the same construction. We need a gadget to delay improving switches like $a_{i}^{1}$ and $b_{i}^{1}$, however.


## Delaying events



- By replacing a vertex by a chain of vertices, a specific sequence of improving switches has to be performed to get the same effect as performing one improving switch originally.
- RandomEdge performs uniformly random improving switches, and a longer sequence therefore gives a longer delay.


## Delaying events



- By replacing a vertex by a chain of vertices, a specific sequence of improving switches has to be performed to get the same effect as performing one improving switch originally.
- RandomEdge performs uniformly random improving switches, and a longer sequence therefore gives a longer delay.


## Delaying events



- By replacing a vertex by a chain of vertices, a specific sequence of improving switches has to be performed to get the same effect as performing one improving switch originally.
- RandomEdge performs uniformly random improving switches, and a longer sequence therefore gives a longer delay.


## Delaying events



- By replacing a vertex by a chain of vertices, a specific sequence of improving switches has to be performed to get the same effect as performing one improving switch originally.
- RandomEdge performs uniformly random improving switches, and a longer sequence therefore gives a longer delay.


## RandomEdge lower bound, first step

- Let $k$ be the lowest unset bit. Incrementing the counter happens through five phases:
$\Rightarrow$ (1) Make $b_{k}=1$.
(2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(3) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
(9) Make $a_{k}=1$.
(5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.



## RandomEdge lower bound, first step

- Let $k$ be the lowest unset bit. Incrementing the counter happens through five phases:
(1) Make $b_{k}=1$.
$\Rightarrow$ (2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(3) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
(4) Make $a_{k}=1$.
(5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.



## RandomEdge lower bound, first step

- Let $k$ be the lowest unset bit. Incrementing the counter happens through five phases:
(1) Make $b_{k}=1$.
$\Rightarrow$ (2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(3) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
(4) Make $a_{k}=1$.
(5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.



## RandomEdge lower bound, first step

- Let $k$ be the lowest unset bit. Incrementing the counter happens through five phases:
(1) Make $b_{k}=1$.
(2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
$\Rightarrow(3)$ Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
(4) Make $a_{k}=1$.
(5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.



## Fast resetting



- Moving in the other directions happens much faster since all actions are improving switches simultaneously.


## RandomEdge lower bound, first step

- Let $k$ be the lowest unset bit. Incrementing the counter happens through five phases:
(1) Make $b_{k}=1$.
(2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
$\Rightarrow$ (3) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
(9) Make $a_{k}=1$.
(5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.



## RandomEdge lower bound, first step

- Let $k$ be the lowest unset bit. Incrementing the counter happens through five phases:
(1) Make $b_{k}=1$.
(2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
$\Rightarrow$ (3) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
(9) Make $a_{k}=1$.
(5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.



## RandomEdge lower bound, first step

- Let $k$ be the lowest unset bit. Incrementing the counter happens through five phases:
(1) Make $b_{k}=1$.
(2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(3) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
$\Rightarrow$ (9) Make $a_{k}=1$.
(5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.



## RandomEdge lower bound, first step

- Let $k$ be the lowest unset bit. Incrementing the counter happens through five phases:
(1) Make $b_{k}=1$.
(2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(3) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
(4) Make $a_{k}=1$.
$\Rightarrow$ (5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.



## RandomEdge lower bound, first step

- Let $k$ be the lowest unset bit. Incrementing the counter happens through five phases:
(1) Make $b_{k}=1$.
(2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(3) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
(4) Make $a_{k}=1$.
$\Rightarrow$ (5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.



## RandomEdge lower bound, first step

- Let $k$ be the lowest unset bit. Incrementing the counter happens through five phases:
$\Rightarrow$ (1) Make $b_{k}=1$.
(2) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(3) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$.
(9) Make $a_{k}=1$.
(5) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.



## Resetting partially set higher bits

- Problem: Higher bits get a head start in later counting steps. When resetting lower bits we must also reset higher bits that are partially set.


## Resetting partially set higher bits

- Problem: Higher bits get a head start in later counting steps. When resetting lower bits we must also reset higher bits that are partially set.
- Higher bits must have access to lower bits: Introduce actions moving down to $u_{1}$.


## Resetting partially set higher bits

- Problem: Higher bits get a head start in later counting steps. When resetting lower bits we must also reset higher bits that are partially set.
- Higher bits must have access to lower bits: Introduce actions moving down to $u_{1}$.
- Note: The resulting MDP does not actually satisfy the stopping condition, but this just means the LP us unbounded towards $-\infty$. Alternatively, we can introduce randomization and always move up with an insignificant probability.


## Resetting partially set higher bits

- Problem: Higher bits get a head start in later counting steps. When resetting lower bits we must also reset higher bits that are partially set.
- Higher bits must have access to lower bits: Introduce actions moving down to $u_{1}$.
- Note: The resulting MDP does not actually satisfy the stopping condition, but this just means the LP us unbounded towards $-\infty$. Alternatively, we can introduce randomization and always move up with an insignificant probability.
- No state can have direct access to a large reward: Introduce a stochastic action such that this happens only with an insignificant probability $\epsilon=N^{-(4 n+8)}$.


## Resetting partially set higher bits

- Problem: Higher bits get a head start in later counting steps. When resetting lower bits we must also reset higher bits that are partially set.
- Higher bits must have access to lower bits: Introduce actions moving down to $u_{1}$.
- Note: The resulting MDP does not actually satisfy the stopping condition, but this just means the LP us unbounded towards $-\infty$. Alternatively, we can introduce randomization and always move up with an insignificant probability.
- No state can have direct access to a large reward: Introduce a stochastic action such that this happens only with an insignificant probability $\epsilon=N^{-(4 n+8)}$.
- Resetting higher bits requires alternating behaviour: Introduce an additional chain of $c_{i}$ vertices.


## RandomEdge lower bound, full construction



## RandomEdge lower bound, full construction



## RandomEdge lower bound, full construction



## RandomEdge lower bound, full construction

- Incrementing the counter happens through seven phases:
(1) Make $b_{k}=1$.
(2) Make $c_{k}=1$.
(3) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(4) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$. Reset $b_{i}$ for unset bits $i>k$.
(9) Make $a_{k}=1$.
(6) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.
(1) Reset $c_{i}$ for all unset bits.



## RandomEdge lower bound, full construction

- Incrementing the counter happens through seven phases:
$\Rightarrow$ (1) Make $b_{k}=1$.
(2) Make $c_{k}=1$.
(3) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(4) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$. Reset $b_{i}$ for unset bits $i>k$.
(9) Make $a_{k}=1$.
(6) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.
(7) Reset $c_{i}$ for all unset bits.



## RandomEdge lower bound, full construction

- Incrementing the counter happens through seven phases:
(1) Make $b_{k}=1$.
$\Rightarrow$ (2) Make $c_{k}=1$.
(3) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(4) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$. Reset $b_{i}$ for unset bits $i>k$.
(5) Make $a_{k}=1$.
(6) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.
(7) Reset $c_{i}$ for all unset bits.



## RandomEdge lower bound, full construction

- Incrementing the counter happens through seven phases:
(1) Make $b_{k}=1$.
(2) Make $c_{k}=1$.
$\Rightarrow$ (3) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(4) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$. Reset $b_{i}$ for unset bits $i>k$.
(9) Make $a_{k}=1$.
(6) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.
(1) Reset $c_{i}$ for all unset bits.



## RandomEdge lower bound, full construction

- Incrementing the counter happens through seven phases:
(1) Make $b_{k}=1$.
(2) Make $c_{k}=1$.
(3) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
$\Rightarrow$ (9) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$. Reset $b_{i}$ for unset bits $i>k$.
(9) Make $a_{k}=1$.
(6) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.
(1) Reset $c_{i}$ for all unset bits.



## RandomEdge lower bound, full construction

- Incrementing the counter happens through seven phases:
(1) Make $b_{k}=1$.
(2) Make $c_{k}=1$.
(3) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(4) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$. Reset $b_{i}$ for unset bits $i>k$.
$\Rightarrow$ (6) Make $a_{k}=1$.
(6) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.
(7) Reset $c_{i}$ for all unset bits.



## RandomEdge lower bound, full construction

- Incrementing the counter happens through seven phases:
(1) Make $b_{k}=1$.
(2) Make $c_{k}=1$.
(3) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(4) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$. Reset $b_{i}$ for unset bits $i>k$.
(5) Make $a_{k}=1$.
$\Rightarrow$ (0) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.
(1) Reset $c_{i}$ for all unset bits.



## RandomEdge lower bound, full construction

- Incrementing the counter happens through seven phases:
(1) Make $b_{k}=1$.
(2) Make $c_{k}=1$.
(3) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(4) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$. Reset $b_{i}$ for unset bits $i>k$.
(5) Make $a_{k}=1$.
(6) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.
$\Rightarrow$ (1) Reset $c_{i}$ for all unset bits.



## RandomEdge lower bound, full construction

- Incrementing the counter happens through seven phases:
$\Rightarrow$ (1) Make $b_{k}=1$.
(2) Make $c_{k}=1$.
(3) Make $u_{k}=1$ and $u_{i}=0$, for $i<k$.
(4) Make $b_{i}=0$ and $a_{i}=0$, for $i<k$. Reset $b_{i}$ for unset bits $i>k$.
(5) Make $a_{k}=1$.
(6) Make $w_{k}=1$ and $w_{i}=0$, for $i<k$.
(1) Reset $c_{i}$ for all unset bits.



## Competing chains

- The greatest challenge when setting the parameters is to make sure that the lowest unset bit is incremented next with high probability.
- The situation occurs when two chains $b_{i}$ and $b_{i+1}$ of lengths $\ell_{i}$ and $\ell_{i+1}$ are competing to change from 0 to 1 .
- In both chains there is always exactly one improving switch, which means that the RANDOMEDGE pivoting rule will pick either of them with equal probability.
- We bound the probability of failure with a Chernoff bound, and show that it suffices to set $\ell_{i}=\Theta\left(i^{2} n\right)$.


## RandomEdge lower bound

Theorem (Friedmann, Hansen and Zwick (2011))
The worst-case expected number of pivoting steps performed by RandomEdge on linear programs with $n$ equalities and $2 n$ non-negative variables is $2^{\Omega\left(n^{1 / 4}\right)}$.
maximize

$$
\begin{aligned}
& \sum_{i=1}^{n} \sum_{j=1}^{h}\left((-N)^{4 i-1}+j \epsilon\right) a_{i, j}^{0}+ \\
& \sum_{i=1}^{n}\left((-N)^{4 i+1}+\epsilon(-N)^{i+2}\right)\left(a_{i, 1}^{1}+u_{i}^{1}\right)+ \\
& \sum_{i=1}^{n}\left(\epsilon(-N)^{4 i+2}\right)\left(b_{i, 1}^{1}+c_{i, 1}^{1}\right)+ \\
& \sum_{i=1}^{n} \sum_{j=1}^{\ell_{i}} j \epsilon b_{i, j}^{0}+\sum_{i=1}^{n} \sum_{j=1}^{g} j \epsilon c_{i, j}^{0}
\end{aligned}
$$

subject to

$$
\begin{array}{ll}
\forall 1 \leq i \leq n: & \begin{aligned}
a_{i, h}^{0}+a_{i, h}^{1} & =1+w_{i}^{1} \\
\forall 1 \leq i \leq n, \forall 1 \leq j<h: & a_{i, j}^{0}+a_{i, j}^{1}
\end{aligned}=1+a_{i, j+1}^{1} \\
\forall 1 \leq i \leq n: & b_{i, \ell_{i}}^{0}+b_{i, \ell_{i}}^{1}
\end{array}=1+\frac{1-\epsilon \epsilon}{2}\left(a_{i, 1}^{1}+b_{i, 1}^{1}+c_{i, 1}^{1}+u_{i}^{1}\right) .
$$

```
Function RandomFacet ( \(G, \pi\) )
if \(\pi\) contains all actions then
    return \(\pi\)
else
    Choose unused action a uniformly at random
    \(\pi^{\prime} \leftarrow \operatorname{RANDOMFACET}(G \backslash\{a\}, \pi)\)
    if \(a\) is improving switch w.r.t. \(\pi^{\prime}\) then
    \(\pi^{\prime \prime} \leftarrow \pi^{\prime}[a]\)
    return RandomFacet( \(G, \pi^{\prime \prime}\) )
else
    return \(\pi^{\prime}\)
```

```
Function RandomFacet ( \(G, \pi, \varphi\) )
if \(\pi\) contains all actions then
    | return \(\pi\)
else
    \(a \leftarrow \operatorname{argmin}_{a \in A \backslash \pi} \varphi(a)\)
    \(\pi^{\prime} \leftarrow \operatorname{RANDOMFACET}(G \backslash\{a\}, \pi, \varphi)\)
    if \(a\) is improving switch w.r.t. \(\pi^{\prime}\) then
        \(\pi^{\prime \prime} \leftarrow \pi^{\prime}[\mathrm{a}]\)
    return RandomFacet \(\left(G, \pi^{\prime \prime}, \varphi\right)\)
    else
        return \(\pi^{\prime}\)
```

- In Friedmann, Hansen and Zwick (SODA, 2011) we proved a $2^{\tilde{\Omega}\left(n^{1 / 2}\right)}$ lower bound, for parity games, for the "modified RANDOMFACET algorithm" starting with a uniformly random permutation.
- We incorrectly believed, until less than three weeks ago, that by linearity of expectation the RANDOMFACET algorithm required the same expected number of steps. We now know that this is not true.
- Fortunately, using the same construction with different parameters we have been able to prove a $2^{\tilde{\Omega}\left(n^{1 / 3}\right)}$ lower bound, which will be made public as soon as all the details have been written and verified.
- Looking closer at the "modified RandomFacet algorithm", it turns out to be a dual, recursive variant of the Randomized Bland's rule.
- In Friedmann, Hansen and Zwick (STOC, 2011) we showed a simple transformation of our lower bound parity games to MDPs, thereby getting lower bounds for the simplex algorithm. This transformation remains the same.
- The main result of this paper was the $2^{\Omega\left(n^{1 / 4}\right)}$ lower bound for RandomEdge which is unaffected.


## Simplified RandomFacet lower bound



## Simplified RandomFacet lower bound

- The RandomFacet algorithm picks a random edge, here $b_{2}^{1}$, and removes it, thereby disabling the bit.



## Simplified RandomFacet lower bound

- The RandomFacet algorithm picks a random edge, here $b_{2}^{1}$, and removes it, thereby disabling the bit.
- The MDP is then solved recursively.



## Simplified RandomFacet lower bound

- The RandomFacet algorithm picks a random edge, here $b_{2}^{1}$, and removes it, thereby disabling the bit.
- The MDP is then solved recursively.
- $b_{2}^{1}$ is reintroduced and a switch is made.



## Simplified RandomFacet lower bound

- The RandomFacet algorithm picks a random edge, here $b_{2}^{1}$, and removes it, thereby disabling the bit.
- The MDP is then solved recursively.
- $b_{2}^{1}$ is reintroduced and a switch is made.



## Simplified RandomFacet lower bound

- A new random edge, $a_{2}^{1}$, is removed.



## Simplified RandomFacet lower bound

- A new random edge, $a_{2}^{1}$, is removed.
- The MDP is again solved recursively.



## Simplified RandomFacet lower bound

- A new random edge, $a_{2}^{1}$, is removed.
- The MDP is again solved recursively.
- $a_{2}^{1}$ is reintroduced and a switch is made.



## Simplified RandomFacet lower bound

- A new random edge, $a_{2}^{1}$, is removed.
- The MDP is again solved recursively.
- $a_{2}^{1}$ is reintroduced and a switch is made.



## We simulate a "randomized bitcounter"

Start with $n$ bits with value 0 :

## We simulate a "randomized bitcounter"

Start with $n$ bits with value 0 :
Pick a random bit $i$ and fix it:

00000
00000

## We simulate a "randomized bitcounter"

Start with $n$ bits with value 0 :
Pick a random bit $i$ and fix it:
00000

Count recursively with the remaining $n-1$ bits: 11011

## We simulate a "randomized bitcounter"

Start with $n$ bits with value 0 :00000Pick a random bit $i$ and fix it: ..... 00000
Count recursively with the remaining $n-1$ bits: ..... 11011
Increment the $i$ 'th bit: ..... 11111

## We simulate a "randomized bitcounter"

Start with $n$ bits with value 0 :
Pick a random bit $i$ and fix it:
00000
$00 \underline{0} 00$
Count recursively with the remaining $n-1$ bits: $11 \underline{111}$
Increment the $i$ 'th bit: 11111
Reset the $i-1$ lower bits: 11100

## We simulate a "randomized bitcounter"

Start with $n$ bits with value 0 :00000Pick a random bit $i$ and fix it: ..... 00000
Count recursively with the remaining $n-1$ bits: ..... 11011
Increment the $i$ 'th bit: ..... 11111
Reset the $i-1$ lower bits: ..... 11100
Count recursively with the $i-1$ lower bits: ..... 11100

## We simulate a "randomized bitcounter"

Start with $n$ bits with value 0 :
Pick a random bit $i$ and fix it: 00000
Count recursively with the remaining $n-1$ bits: 00000
Increment the $i$ 'th bit:
11011
11111
Reset the $i-1$ lower bits:
11100
Count recursively with the $i-1$ lower bits:

- Expected number of increments:

$$
\begin{aligned}
& f(0)=0 \\
& f(n)=f(n-1)+1+\frac{1}{n} \sum_{i=0}^{n-1} f(i) \quad \text { for } \quad n>0
\end{aligned}
$$

## We simulate a "randomized bitcounter"

Start with $n$ bits with value 0 :
Pick a random bit $i$ and fix it:
Count recursively with the remaining $n-1$ bits:
00000
11011
Increment the $i$ 'th bit:
11111
Reset the $i-1$ lower bits:
11100
Count recursively with the $i-1$ lower bits:

- Expected number of increments:

$$
\begin{aligned}
& f(0)=0 \\
& f(n)=f(n-1)+1+\frac{1}{n} \sum_{i=0}^{n-1} f(i) \quad \text { for } \quad n>0
\end{aligned}
$$

- Solving the recurrence gives: $f(n)=2^{\Theta(\sqrt{n})}$


## Different challenges

- When constructing lower bounds for the RandomFacet pivoting rule, instead of delaying improving switches, the challenge is to make sure that certain actions are not removed before certain other actions.


## Different challenges

- When constructing lower bounds for the RandomFacet pivoting rule, instead of delaying improving switches, the challenge is to make sure that certain actions are not removed before certain other actions.
- Suppose an action a must not be removed before another action $b$.


## Different challenges

- When constructing lower bounds for the RandomFacet pivoting rule, instead of delaying improving switches, the challenge is to make sure that certain actions are not removed before certain other actions.
- Suppose an action a must not be removed before another action $b$.
- To achieve this with high probability we make use of redundancy: Let $a$ and $b$ be copied $k$ times, in such a way that we only require that at least one copy of $b$ is removed before all copies of a are removed.


## Different challenges

- When constructing lower bounds for the RandomFacet pivoting rule, instead of delaying improving switches, the challenge is to make sure that certain actions are not removed before certain other actions.
- Suppose an action a must not be removed before another action $b$.
- To achieve this with high probability we make use of redundancy: Let $a$ and $b$ be copied $k$ times, in such a way that we only require that at least one copy of $b$ is removed before all copies of a are removed.
- The only bad permutation for the "modified RandomFacet pivoting rule" is then: $a a \ldots a b b \ldots b$


## Different challenges

- When constructing lower bounds for the RandomFacet pivoting rule, instead of delaying improving switches, the challenge is to make sure that certain actions are not removed before certain other actions.
- Suppose an action a must not be removed before another action $b$.
- To achieve this with high probability we make use of redundancy: Let $a$ and $b$ be copied $k$ times, in such a way that we only require that at least one copy of $b$ is removed before all copies of a are removed.
- The only bad permutation for the "modified RandomFacet pivoting rule" is then: $a a \ldots a b b \ldots b$
- The probability of choosing a bad permutation is $\frac{(k!)^{2}}{(2 k)!} \leq \frac{1}{2^{k}}$.


## Different gadgets



- We use different gadgets to ensure that we get the correct behaviour with high probability.
- For the "modified RandomFacet pivoting rule" this only increases the number of states and actions by a polylogarithmic factor.
- For the real RandomFacet pivoting rule the increase needs to be a factor $\tilde{O}(\sqrt{n})$.


## RandomFacet lower bound construction



## Overview

- Lecture 1:
- Introduction to linear programming and the simplex algorithm.
- Pivoting rules.
- The RandomFacet pivoting rule.
- Lecture 2:
- The Hirsch conjecture.
- Introduction to Markov decision processes (MDPs).
- Upper bound for the LargestCoefficient pivoting rule for MDPs.
- Lecture 3 :
- Lower bounds for pivoting rules utilizing MDPs. Example: Bland's rule.
- Lower bound for the RandomEdge pivoting rule.
- Abstractions and related problems.


## 2-player turn-based stochastic games (2TBSGs)



- A 2TBSG is an MDP where the set of states is partitioned into two sets: $S_{1} \cup S_{2}=S$.
- $S_{1}$ is controlled by player 1 , the maximizer.
- $S_{2}$ is controlled by player 2 , the minimizer.
- A strategy $\pi_{1}$ (or $\pi_{2}$ ) for player 1 (or player 2 ) is a choice of an action for each state $i \in S_{1}$ (or $i \in S_{2}$ ).


## 2-player turn-based stochastic games (2TBSGs)

- A strategy profile $\pi=\left(\pi_{1}, \pi_{2}\right)$ is a pair of strategies, defining a Markov chain with rewards.
- The value vector for discounted 2TBSGs is again defined as:

$$
v_{\pi}=\left(I-\gamma P_{\pi}\right)^{-1} c_{\pi}
$$

## 2-player turn-based stochastic games (2TBSGs)

- A strategy profile $\pi=\left(\pi_{1}, \pi_{2}\right)$ is a pair of strategies, defining a Markov chain with rewards.
- The value vector for discounted 2TBSGs is again defined as:

$$
v_{\pi}=\left(I-\gamma P_{\pi}\right)^{-1} c_{\pi}
$$

- For a fixed strategy $\pi_{1}$ for player 1 , a best response from player 2 is a strategy:

$$
\pi_{2}\left(\pi_{1}\right) \in \underset{\pi_{2}}{\operatorname{argmin}} v_{\pi_{1}, \pi_{2}}
$$

## 2-player turn-based stochastic games (2TBSGs)

- A strategy profile $\pi=\left(\pi_{1}, \pi_{2}\right)$ is a pair of strategies, defining a Markov chain with rewards.
- The value vector for discounted 2TBSGs is again defined as:

$$
v_{\pi}=\left(I-\gamma P_{\pi}\right)^{-1} c_{\pi}
$$

- For a fixed strategy $\pi_{1}$ for player 1 , a best response from player 2 is a strategy:

$$
\pi_{2}\left(\pi_{1}\right) \in \underset{\pi_{2}}{\operatorname{argmin}} v_{\pi_{1}, \pi_{2}}
$$

- Note that $\pi_{2}\left(\pi_{1}\right)$ can be computed by solving an MDP.


## 2-player turn-based stochastic games (2TBSGs)

- A strategy profile $\pi=\left(\pi_{1}, \pi_{2}\right)$ is a pair of strategies, defining a Markov chain with rewards.
- The value vector for discounted 2TBSGs is again defined as:

$$
v_{\pi}=\left(I-\gamma P_{\pi}\right)^{-1} c_{\pi}
$$

- For a fixed strategy $\pi_{1}$ for player 1 , a best response from player 2 is a strategy:

$$
\pi_{2}\left(\pi_{1}\right) \in \underset{\pi_{2}}{\operatorname{argmin}} v_{\pi_{1}, \pi_{2}}
$$

- Note that $\pi_{2}\left(\pi_{1}\right)$ can be computed by solving an MDP.
- A best response from player $1, \pi_{1}\left(\pi_{2}\right)$, is defined analogously.


## 2-player turn-based stochastic games (2TBSGs)

- $\pi_{1}^{*}$ and $\pi_{2}^{*}$ are optimal if:

$$
\begin{array}{ll}
\forall \pi_{1}: & v_{\pi_{1}^{*}, \pi_{2}\left(\pi_{1}^{*}\right)} \geq v_{\pi_{1}, \pi_{2}\left(\pi_{1}\right)} \\
\forall \pi_{2}: & v_{\pi_{1}\left(\pi_{2}^{*}\right), \pi_{2}^{*}} \leq v_{\pi_{1}\left(\pi_{2}\right), \pi_{2}}
\end{array}
$$

## 2-player turn-based stochastic games (2TBSGs)

- $\pi_{1}^{*}$ and $\pi_{2}^{*}$ are optimal if:

$$
\begin{array}{ll}
\forall \pi_{1}: & v_{\pi_{1}^{*}, \pi_{2}\left(\pi_{1}^{*}\right)} \geq v_{\pi_{1}, \pi_{2}\left(\pi_{1}\right)} \\
\forall \pi_{2}: & v_{\pi_{1}\left(\pi_{2}^{*}\right), \pi_{2}^{*}} \leq v_{\pi_{1}\left(\pi_{2}\right), \pi_{2}}
\end{array}
$$

- Alternatively, $\pi_{1}^{*}$ and $\pi_{2}^{*}$ are optimal if $\pi_{1}^{*}$ is a best response to $\pi_{2}^{*}$, and $\pi_{2}^{*}$ is a best response to $\pi_{1}^{*}$. We then say that $\left(\pi_{1}^{*}, \pi_{2}^{*}\right)$ is a Nash equilibrium.


## 2-player turn-based stochastic games (2TBSGs)

- $\pi_{1}^{*}$ and $\pi_{2}^{*}$ are optimal if:

$$
\begin{array}{ll}
\forall \pi_{1}: & v_{\pi_{1}^{*}, \pi_{2}\left(\pi_{1}^{*}\right)} \geq v_{\pi_{1}, \pi_{2}\left(\pi_{1}\right)} \\
\forall \pi_{2}: & v_{\pi_{1}\left(\pi_{2}^{*}\right), \pi_{2}^{*}} \leq v_{\pi_{1}\left(\pi_{2}\right), \pi_{2}}
\end{array}
$$

- Alternatively, $\pi_{1}^{*}$ and $\pi_{2}^{*}$ are optimal if $\pi_{1}^{*}$ is a best response to $\pi_{2}^{*}$, and $\pi_{2}^{*}$ is a best response to $\pi_{1}^{*}$. We then say that $\left(\pi_{1}^{*}, \pi_{2}^{*}\right)$ is a Nash equilibrium.
- Shapley (1953): Optimal strategies always exist.


## 2-player turn-based stochastic games (2TBSGs)

- $\pi_{1}^{*}$ and $\pi_{2}^{*}$ are optimal if:

$$
\begin{array}{ll}
\forall \pi_{1}: & v_{\pi_{1}^{*}, \pi_{2}\left(\pi_{1}^{*}\right) \geq} v_{\pi_{1}, \pi_{2}\left(\pi_{1}\right)} \\
\forall \pi_{2}: & v_{\pi_{1}\left(\pi_{2}^{*}\right), \pi_{2}^{*}} \leq v_{\pi_{1}\left(\pi_{2}\right), \pi_{2}}
\end{array}
$$

- Alternatively, $\pi_{1}^{*}$ and $\pi_{2}^{*}$ are optimal if $\pi_{1}^{*}$ is a best response to $\pi_{2}^{*}$, and $\pi_{2}^{*}$ is a best response to $\pi_{1}^{*}$. We then say that $\left(\pi_{1}^{*}, \pi_{2}^{*}\right)$ is a Nash equilibrium.
- Shapley (1953): Optimal strategies always exist.
- Solving a 2TBSG means finding an optimal strategy profile.


## 2-player turn-based stochastic games (2TBSGs)

- $\pi_{1}^{*}$ and $\pi_{2}^{*}$ are optimal if:

$$
\begin{array}{ll}
\forall \pi_{1}: & v_{\pi_{1}^{*}, \pi_{2}\left(\pi_{1}^{*}\right)} \geq v_{\pi_{1}, \pi_{2}\left(\pi_{1}\right)} \\
\forall \pi_{2}: & v_{\pi_{1}\left(\pi_{2}^{*}\right), \pi_{2}^{*}} \leq v_{\pi_{1}\left(\pi_{2}\right), \pi_{2}}
\end{array}
$$

- Alternatively, $\pi_{1}^{*}$ and $\pi_{2}^{*}$ are optimal if $\pi_{1}^{*}$ is a best response to $\pi_{2}^{*}$, and $\pi_{2}^{*}$ is a best response to $\pi_{1}^{*}$. We then say that $\left(\pi_{1}^{*}, \pi_{2}^{*}\right)$ is a Nash equilibrium.
- Shapley (1953): Optimal strategies always exist.
- Solving a 2TBSG means finding an optimal strategy profile.
- Note that the decision problem corresponding to solving 2TBSGs is in NP $\cap$ coNP, since an optimal strategy profile is a witness for both yes and no answers. The problem is not known to be in $\mathbf{P}$.


## 2-player turn-based stochastic games (2TBSGs)

- We again say that $a \in A_{i}$, for $i \in S_{1}$, is an improving switch for player 1 w.r.t. $\pi$ iff:

$$
\left(v_{\pi}\right)_{i}<c_{a}+\gamma P_{a} v_{\pi}
$$

- Similarly, $a \in A_{i}$, for $i \in S_{2}$, is an improving switch for player 2 w.r.t. $\pi$ iff:

$$
\left(v_{\pi}\right)_{i}>c_{a}+\gamma P_{a} v_{\pi}
$$

## 2-player turn-based stochastic games (2TBSGs)

- We again say that $a \in A_{i}$, for $i \in S_{1}$, is an improving switch for player 1 w.r.t. $\pi$ iff:

$$
\left(v_{\pi}\right)_{i}<c_{a}+\gamma P_{a} v_{\pi}
$$

- Similarly, $a \in A_{i}$, for $i \in S_{2}$, is an improving switch for player 2 w.r.t. $\pi$ iff:

$$
\left(v_{\pi}\right)_{i}>c_{a}+\gamma P_{a} v_{\pi}
$$

- The vector of reduced costs for a strategy profile $\pi$ is again defined as:

$$
\bar{c}^{\pi}=c-(J-\gamma P) v_{\pi}
$$

## 2-player turn-based stochastic games (2TBSGs)

- We again say that $a \in A_{i}$, for $i \in S_{1}$, is an improving switch for player 1 w.r.t. $\pi$ iff:

$$
\left(v_{\pi}\right)_{i}<c_{a}+\gamma P_{a} v_{\pi}
$$

- Similarly, $a \in A_{i}$, for $i \in S_{2}$, is an improving switch for player 2 w.r.t. $\pi$ iff:

$$
\left(v_{\pi}\right)_{i}>c_{a}+\gamma P_{a} v_{\pi}
$$

- The vector of reduced costs for a strategy profile $\pi$ is again defined as:

$$
\bar{c}^{\pi}=c-(J-\gamma P) v_{\pi}
$$

- Note that $\left(\pi_{1}^{*}, \pi_{2}^{*}\right)$ is a Nash equilibrium iff there are no improving switches.


## StrategyIteration

Function StrategyIteration ( $\pi_{1}$ )
while $\exists$ improving switch w.r.t. $\left(\pi_{1}, \pi_{2}\left(\pi_{1}\right)\right)$ do
Update $\pi_{1}$ by performing improving switches
return $\left(\pi_{1}, \pi_{2}\left(\pi_{1}\right)\right)$

## StrategyIteration

## Function StrategyIteration $\left(\pi_{1}\right)$

while $\exists$ improving switch w.r.t. $\left(\pi_{1}, \pi_{2}\left(\pi_{1}\right)\right)$ do Update $\pi_{1}$ by performing improving switches return $\left(\pi_{1}, \pi_{2}\left(\pi_{1}\right)\right)$

- Howard's algorithm can be naturally extended to 2TBSGs by choosing:

$$
\forall i \in S_{1}: \quad \pi_{1}(i) \leftarrow \underset{a \in A_{i}}{\operatorname{argmax}} \bar{c}_{a}^{\pi_{1}, \pi_{2}\left(\pi_{1}\right)}
$$

## Non-discounted MDPs and 2TBSGs

- We have already seen that discounted MDPs are a special case of MDPs satisfying the stopping condition, and the same is true for 2TBSGs.
- Liggett and Lippman (1969) showed that for any 2TBSG G there exists a discount factor $\gamma_{G}<1$, such that the same strategies are optimal for all discount factors $\gamma^{\prime} \in\left[\gamma_{G}, 1\right)$.


## Non-discounted MDPs and 2TBSGs

- We have already seen that discounted MDPs are a special case of MDPs satisfying the stopping condition, and the same is true for 2TBSGs.
- Liggett and Lippman (1969) showed that for any 2TBSG G there exists a discount factor $\gamma_{G}<1$, such that the same strategies are optimal for all discount factors $\gamma^{\prime} \in\left[\gamma_{G}, 1\right)$.
- Andersson and Miltersen (2009) showed that $\gamma_{G}$ can be described with a number of bits that is polynomial in the bit complexity of $G$.


## Non-discounted MDPs and 2TBSGs

- We have already seen that discounted MDPs are a special case of MDPs satisfying the stopping condition, and the same is true for 2TBSGs.
- Liggett and Lippman (1969) showed that for any 2TBSG G there exists a discount factor $\gamma_{G}<1$, such that the same strategies are optimal for all discount factors $\gamma^{\prime} \in\left[\gamma_{G}, 1\right)$.
- Andersson and Miltersen (2009) showed that $\gamma_{G}$ can be described with a number of bits that is polynomial in the bit complexity of $G$.
- A 2TBSG $G$ is called non-discounted if it is implicitly using discount factor $\gamma_{G}$.


## Non-discounted MDPs and 2TBSGs

- We have already seen that discounted MDPs are a special case of MDPs satisfying the stopping condition, and the same is true for 2TBSGs.
- Liggett and Lippman (1969) showed that for any 2TBSG G there exists a discount factor $\gamma_{G}<1$, such that the same strategies are optimal for all discount factors $\gamma^{\prime} \in\left[\gamma_{G}, 1\right)$.
- Andersson and Miltersen (2009) showed that $\gamma_{G}$ can be described with a number of bits that is polynomial in the bit complexity of $G$.
- A 2TBSG $G$ is called non-discounted if it is implicitly using discount factor $\gamma_{G}$.
- MDPs (and 2TBSGs) satisfying the stopping condition are a special case of non-discounted MDPs (and 2TBSGs). See, e.g., Puterman (1994).


## Special cases of 2TBSGs

- A non-discounted 2TBSG whose actions are all deterministic is called a mean payoff game.


## Special cases of 2TBSGs

- A non-discounted 2TBSG whose actions are all deterministic is called a mean payoff game.
- An $n$-state mean payoff game where the reward of every action $a$ is described by an integer priority $p_{a}$, such that $c_{a}=(-n)^{p_{a}}$, and where all actions leaving the same state have the same priority, is called a parity game.
- Note that the mean of the rewards of a cycle is positive iff the parity of the largest priority is even.


## Special cases of 2TBSGs

- A non-discounted 2TBSG whose actions are all deterministic is called a mean payoff game.
- An $n$-state mean payoff game where the reward of every action $a$ is described by an integer priority $p_{a}$, such that $c_{a}=(-n)^{p_{a}}$, and where all actions leaving the same state have the same priority, is called a parity game.
- Note that the mean of the rewards of a cycle is positive iff the parity of the largest priority is even.
- There is no known polynomial time algorithm for solving parity games.


## A few results about StrategyIteration

- Friedmann (2009) showed that Howard's algorithm requires exponentially many iterations to solve parity games.
- Fearnley (2010) transformed Friedmann's construction to MDPs.


## A few results about StrategyIteration

- Friedmann (2009) showed that Howard's algorithm requires exponentially many iterations to solve parity games.
- Fearnley (2010) transformed Friedmann's construction to MDPs.
- These lower bounds are precursors for the lower bounds for RandomEdge, RandomFacet and Randomized Bland's rule by Friedmann, Hansen and Zwick (2011), and LeastEntered by Friedmann (2011). All of which were also first obtained for parity games.


## A few results about StrategyIteration

- Friedmann (2009) showed that Howard's algorithm requires exponentially many iterations to solve parity games.
- Fearnley (2010) transformed Friedmann's construction to MDPs.
- These lower bounds are precursors for the lower bounds for RandomEdge, RandomFacet and Randomized Bland's rule by Friedmann, Hansen and Zwick (2011), and LeastEntered by Friedmann (2011). All of which were also first obtained for parity games.
- Ye (2010): $O\left(\frac{m n}{1-\gamma} \log \frac{n}{1-\gamma}\right)$ iterations for discounted MDPs with $n$ states and $m$ actions.


## A few results about StrategyIteration

- Friedmann (2009) showed that Howard's algorithm requires exponentially many iterations to solve parity games.
- Fearnley (2010) transformed Friedmann's construction to MDPs.
- These lower bounds are precursors for the lower bounds for RandomEdge, RandomFacet and Randomized Bland's rule by Friedmann, Hansen and Zwick (2011), and LeastEntered by Friedmann (2011). All of which were also first obtained for parity games.
- Ye (2010): $O\left(\frac{m n}{1-\gamma} \log \frac{n}{1-\gamma}\right)$ iterations for discounted MDPs with $n$ states and $m$ actions.
- Hansen, Miltersen and Zwick (2011): $O\left(\frac{m}{1-\gamma} \log \frac{n}{1-\gamma}\right)$ iterations for discounted 2 TBSGs with $n$ states and $m$ actions.


## 2TBSGs are LP-type problems

- Ludwig (1995), Halman (2007): 2TBSGs are LP-type problems.
- Let $H$ be the set of actions for player 1 , and let $\omega$ map a subgame to the sum of its optimal values. Bases are strategies for player 1.
- Monotonicity: More available actions only increases the value.
- Locality: If $F \subseteq G \subseteq H$ and $-\infty<\omega(F)=\omega(G)$, then $F$ and $G$ share optimal strategies, and if an added action $h \in H$ is an improving switch for one then it also is for the other.


## 2TBSGs are LP-type problems

- Ludwig (1995), Halman (2007): 2TBSGs are LP-type problems.
- Let $H$ be the set of actions for player 1 , and let $\omega$ map a subgame to the sum of its optimal values. Bases are strategies for player 1 .
- Monotonicity: More available actions only increases the value.
- Locality: If $F \subseteq G \subseteq H$ and $-\infty<\omega(F)=\omega(G)$, then $F$ and $G$ share optimal strategies, and if an added action $h \in H$ is an improving switch for one then it also is for the other.
- Hence, the dual RandomFacet algorithm can be used to solve 2 TBSGs, and, in fact, it gives the best known bound for solving the non-discounted problem.


## Unique sink orientations of cubes

- MDPs and 2TBSGs with two actions per state can be described abstractly by acyclic unique sink orientations (AUSOs) of hypercubes:
- Strategies for player 1 map to vertices of the cube, and improving switches define an orientation of the edges such that in each face there is a unique sink.
- An algorithm can evaluate vertices of the cube to learn the orientation of the adjacent edges, and the goal is to find the unique sink of the entire cube.


## Unique sink orientations of cubes

- MDPs and 2TBSGs with two actions per state can be described abstractly by acyclic unique sink orientations (AUSOs) of hypercubes:
- Strategies for player 1 map to vertices of the cube, and improving switches define an orientation of the edges such that in each face there is a unique sink.
- An algorithm can evaluate vertices of the cube to learn the orientation of the adjacent edges, and the goal is to find the unique sink of the entire cube.
- Szabó and Welzl (2001) introduced the FibonacciSeesaw algorithm for solving $n$ dimensional unique sink orientations with $F_{n+2}$ vertex evaluations.


## More extensions of 2 TBSGs

- There is no known way to formulate 2TBSGs as linear programs.
- Gärtner and Rüst (2005), and Jurdziński and Ravani (2008) showed how to formulate 2TBSGs with two actions per state as P -matrix linear complementarity problems.


## More extensions of 2 TBSGs

- There is no known way to formulate 2TBSGs as linear programs.
- Gärtner and Rüst (2005), and Jurdziński and Ravani (2008) showed how to formulate 2TBSGs with two actions per state as P-matrix linear complementarity problems.
- P-matrix linear complementarity problems are also generalized by USOs, but not by AUSOs.


## More extensions of 2TBSGs

- There is no known way to formulate 2TBSGs as linear programs.
- Gärtner and Rüst (2005), and Jurdziński and Ravani (2008) showed how to formulate 2TBSGs with two actions per state as P-matrix linear complementarity problems.
- P-matrix linear complementarity problems are also generalized by USOs, but not by AUSOs.
- Solving P-matrix linear complementarity problems, as well as 2TBSGs, is known to be in PPAD $\cap$ PLS. Daskalakis and Papadimitriou (2011) suggested a new complexity class CLS (continuous local search) for capturing these and other problems.


[^0]:    $\left[\begin{array}{l}7 \\ \\ 2 \\ 5\end{array}\right]$

