Looking at Mean-Payoff through Foggy Windows

Paul Hunter, Guillermo Pérez, Jean-François Raskin

Université Libre de Bruxelles

Cassting meeting (Oct'15) @ Cachan

Outline

1. Mean-payoff games with PO
2. Window objectives
   - Defs. and nice properties
   - Partial-Observation generalizations
3. Results
   - Lower bounds
   - Upper bounds
4. Conclusions

MPGs with partial-observation: example

- $(Q, q_I, \quad\quad\quad\quad)$

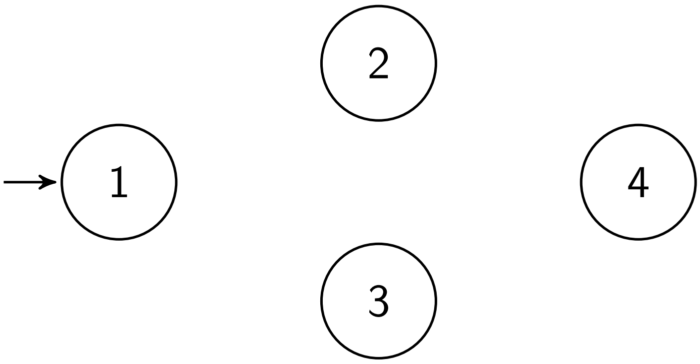

MPGs with partial-observation: example

- $(Q, q_I, \Sigma, \qquad\qquad)$

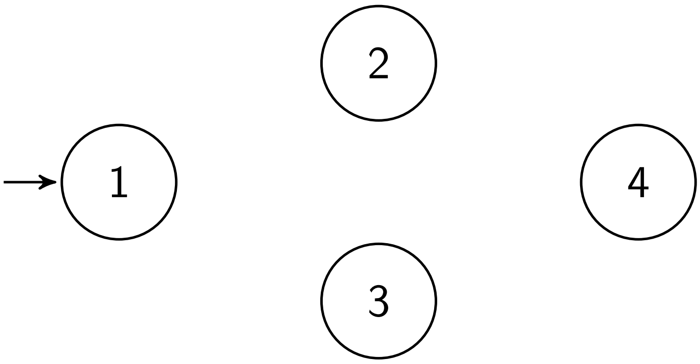

MPGs with partial-observation: example

- $(Q, q_I, \Sigma, \Delta, w, \quad)$
- we assume Δ is total

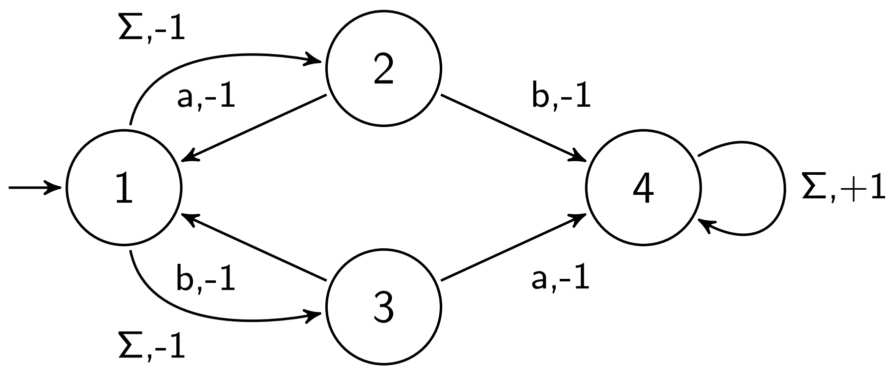

MPGs with partial-observation: example

- $(Q, q_I, \Sigma, \Delta, w, Obs)$
- we assume Δ is total

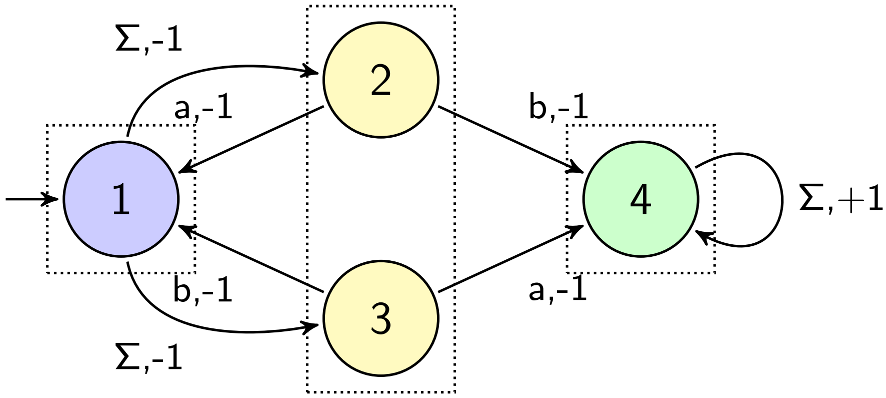

MPGs with partial-observation: example

- $(Q, q_I, \Sigma, \Delta, w, Obs)$
- we assume Δ is total
- each turn: $\exists$ve chooses an action, $\forall$dam resolves non-determinism (repeat to ∞)

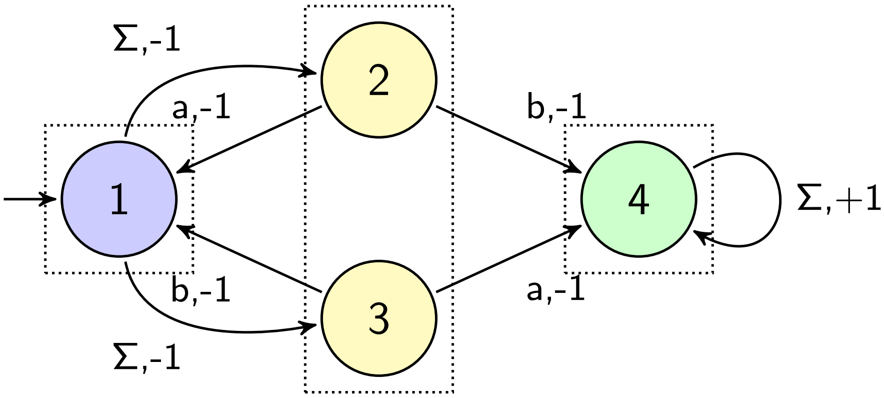

MPGs with partial-observation: example

- $(Q, q_I, \Sigma, \Delta, w, Obs)$
- we assume Δ is total
- each turn: ∃ve chooses an action, ∀dam resolves non-determinism (repeat to ∞)
- ∃ve only sees colors, ∀dam sees everything

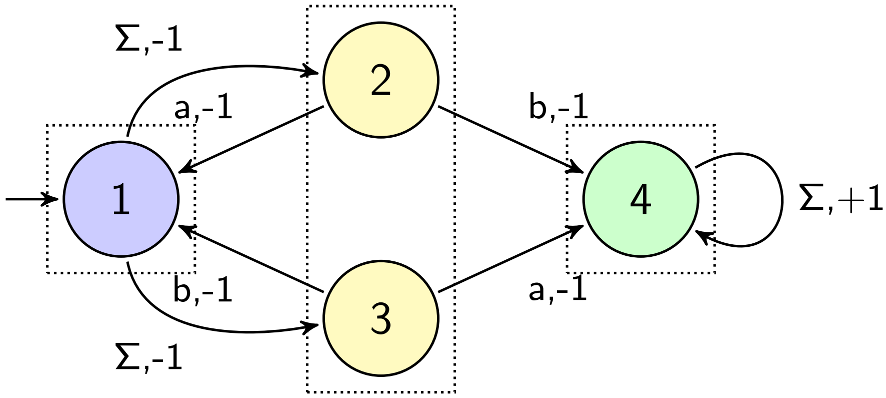

MPGs with partial-observation: example

- $(Q, q_I, \Sigma, \Delta, w, Obs)$
- we assume Δ is total
- each turn: $\exists$ve chooses an action, $\forall$dam resolves non-determinism (repeat to ∞)
- $\exists$ve only sees colors, $\forall$dam sees everything

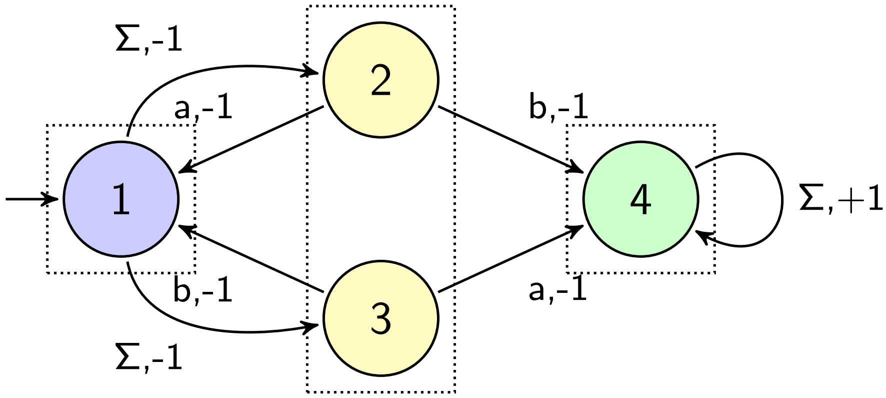

MPGs with partial-observation: example

- $(Q, q_I, \Sigma, \Delta, w, Obs)$
- $\exists$ve sees an abstract play, e.g. $o_0 \sigma_0 o_1 \sigma_1 \ldots$

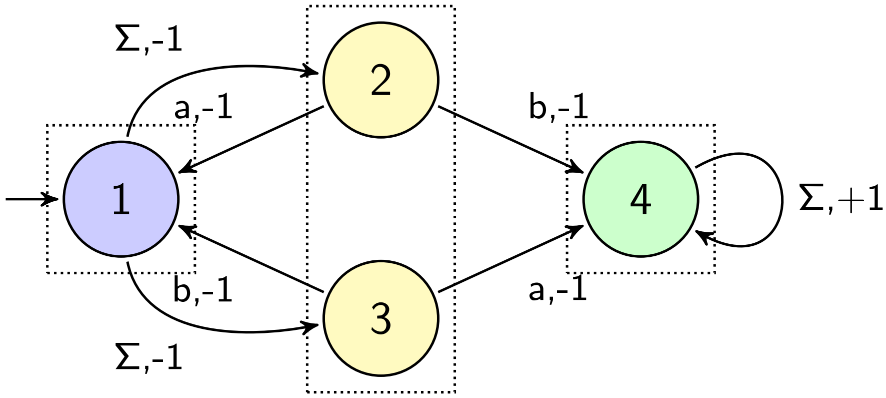

MPGs with partial-observation: example

- $(Q, q_I, \Sigma, \Delta, w, Obs)$
- $\exists$ve sees an abstract play, e.g. $o_0\sigma_0 o_1\sigma_1 \ldots$
- there may be several concrete plays corr. to one abstract play

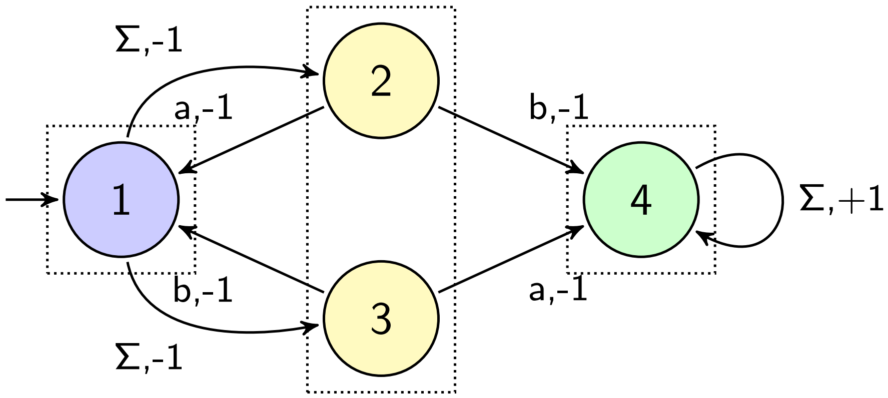

MPGs with partial-observation: example

- $(Q, q_I, \Sigma, \Delta, w, Obs)$
- $\exists$ve sees an abstract play, e.g. $o_0\sigma_0 o_1\sigma_1 \ldots$
- there may be several concrete plays corr. to one abstract play
- for a concrete play, $\forall$dam pays $\exists$ve the limit average of the sequence of weights

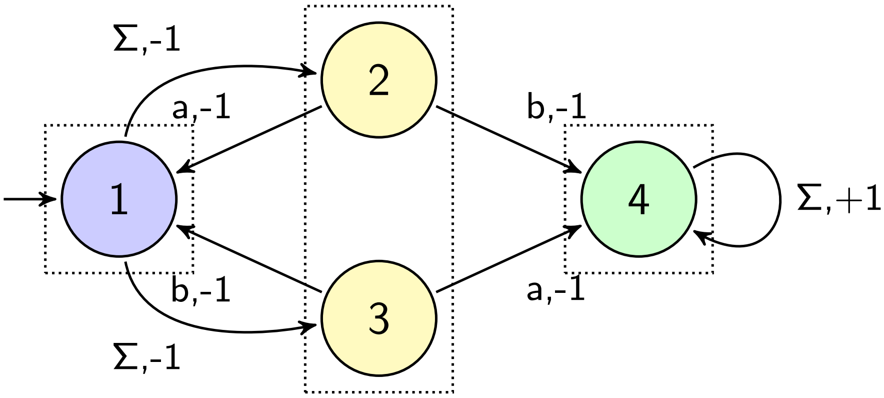

MPGs with partial-observation: example

- $(Q, q_I, \Sigma, \Delta, w, Obs)$
- $\exists$ve sees an abstract play, e.g. $o_0\sigma_0 o_1\sigma_1 \ldots$
- there may be several concrete plays corr. to one abstract play
- for a concrete play, $\forall$dam pays $\exists$ve the limit average of the sequence of weights

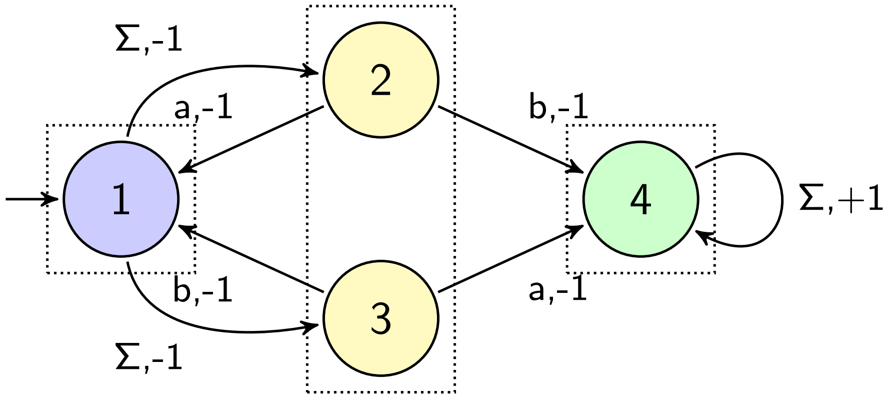

Some formal definitions

Definition (Strategies for ∃ve)

An observation-based strategy for ∃ve is a function from finite sequences $(Obs \cdot \Sigma)^* Obs$ to the next action.

Some formal definitions

Definition (Strategies for ∃ve)

An observation-based strategy for ∃ve is a function from finite sequences $(Obs \cdot \Sigma)^* Obs$ to the next action.

Definition (MP value)

Given a concrete play $\pi = q_0 \sigma_0 q_1 \ldots \in (Q \cdot \Sigma)^\omega$ and weight function $w : \Delta \to \mathbb{Z}$, the MP value is $\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} w(q_i, \sigma_i, q_{i+1})$ $(\overline{MP}, \underline{MP})$.

Some formal definitions

Definition (Strategies for $\exists$ve)

An observation-based strategy for $\exists$ve is a function from finite sequences $(Obs \cdot \Sigma)^* Obs$ to the next action.

Definition (MP value)

Given a concrete play $\pi = q_0\sigma_0 q_1 \ldots \in (Q \cdot \Sigma)^\omega$ and weight function $w : \Delta \to \mathbb{Z}$, the MP value is $\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} w(q_i, \sigma_i, q_{i+1})$ ($\overline{MP}$, $\underline{MP}$).

Problem (The threshold problem)

Given a threshold $\nu \in \mathbb{N}$, the MPG is won by $\exists$ve iff she can ensure $MP \geq \nu$ for all concretizations of her witnessed play.[a]

[a] W.l.o.g. we assume $\nu = 0$.

Motivation

Full-observation (i.e. each state a unique color)

Worst-case analysis of[a]

- online metrical task systems,
- finite window online string matching,
- selection with limited storage.

[a] given in [ZP96]

Motivation

Full-observation (i.e. each state a unique color)

Worst-case analysis of[a]

- online metrical task systems,
- finite window online string matching,
- selection with limited storage.

[a] given in [ZP96]

In the context of reactive synthesis

- MPGs are natural models for systems where we want to optimize the limit-average usage of a resource.
- Partial-observation arises from the fact that most systems have a limited amount of sensors and input data.

What is known?

Theorem (Full-observation [EM79, ZP96])

- *MPGs are determined, i.e. if $\exists$ve doesn't have a winning strategy then $\forall$dam does (and vice versa).*
- *Positional strategies suffice for either $\forall$dam or $\exists$ve to win a MPG.*
- *The threshold problem is in* NP $\cap$ coNP.

What is known?

Theorem (Full-observation [EM79, ZP96])

- *MPGs are determined, i.e. if ∃ve doesn't have a winning strategy then ∀dam does (and vice versa).*
- *Positional strategies suffice for either ∀dam or ∃ve to win a MPG.*
- *The threshold problem is in* $\mathsf{NP} \cap \mathsf{coNP}$.

Theorem (Partial-observation [DDG+10] and [HPR14])

- *∃ve learns about the state of the game by using (possibly infinite) memory.*
- *The limit does not always exist, i.e.* $\overline{MP} \neq \underline{MP}$.
- *The threshold problem is undecidable.*

Outline

UNIVERSITÉ LIBRE DE BRUXELLES

Window objectives

Introduced in [CDRR13]

Original motivation:

- Quantitative objectives that specify timing requirements
- Stronger than MP

UNIVERSITÉ LIBRE DE BRUXELLES

Window objectives

Introduced in [CDRR13]

Original motivation:

- Quantitative objectives that specify timing requirements
- Stronger than MP

Key Idea (main construction block)

A good window from turn i...

- Window of fixed size sliding along a play
- ∃ve must see a local positive MP before hitting the end of the window

UNIVERSITÉ LIBRE DE BRUXELLES

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

UNIVERSITÉ LIBRE DE BRUXELLES

[1] Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

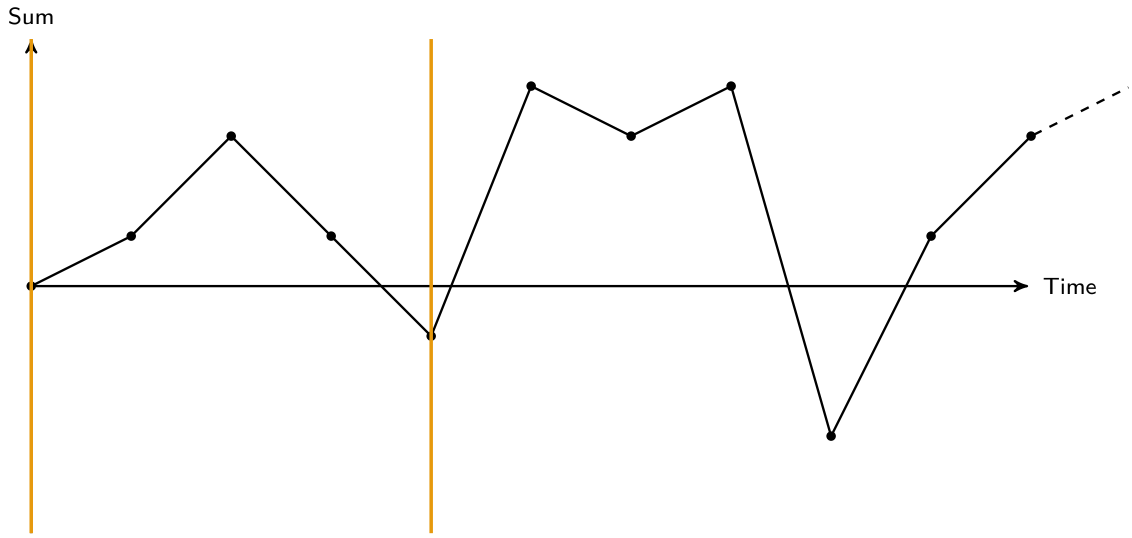

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

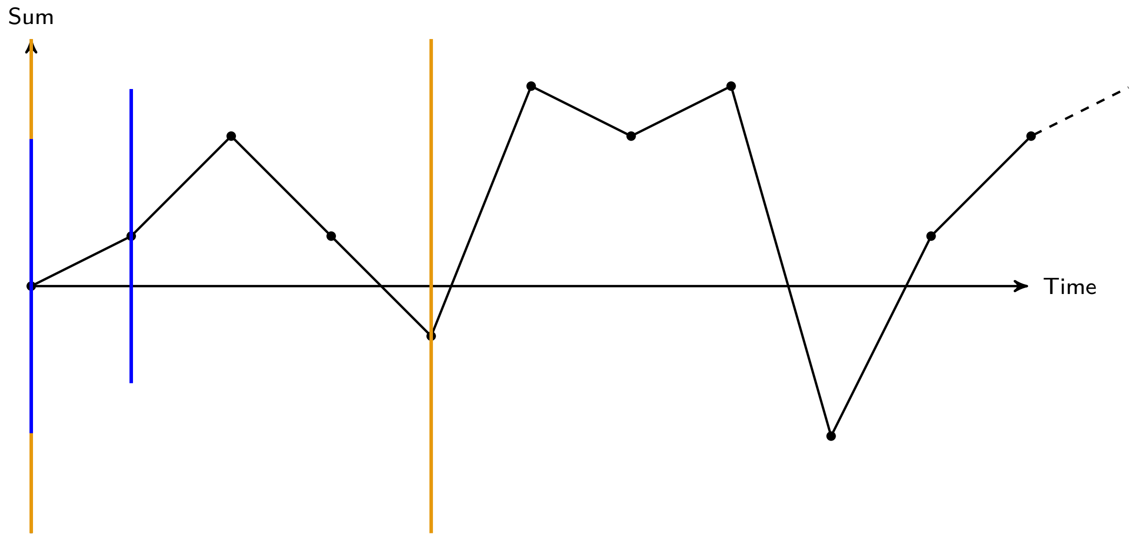

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

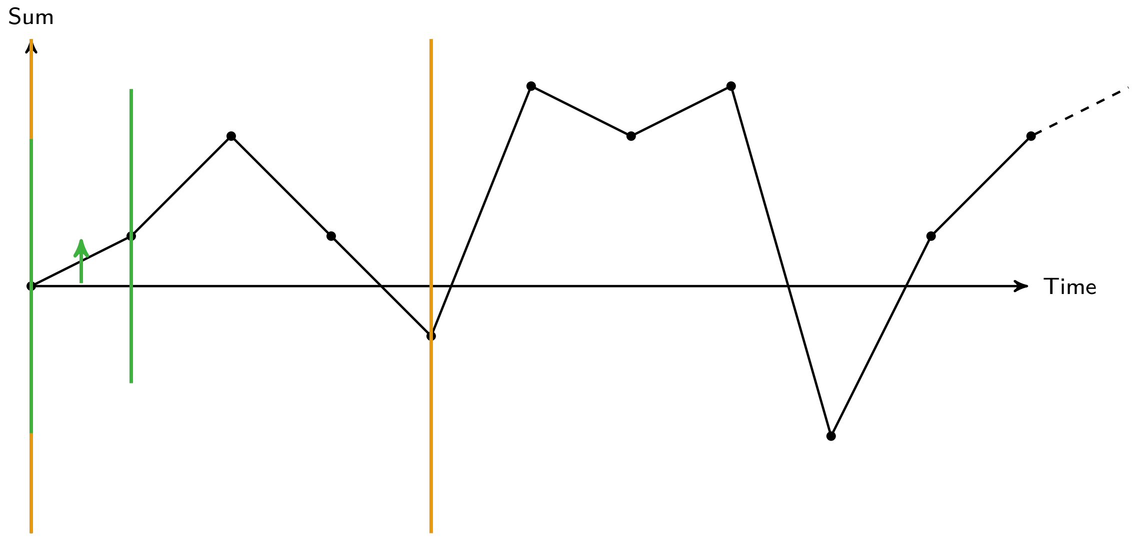

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

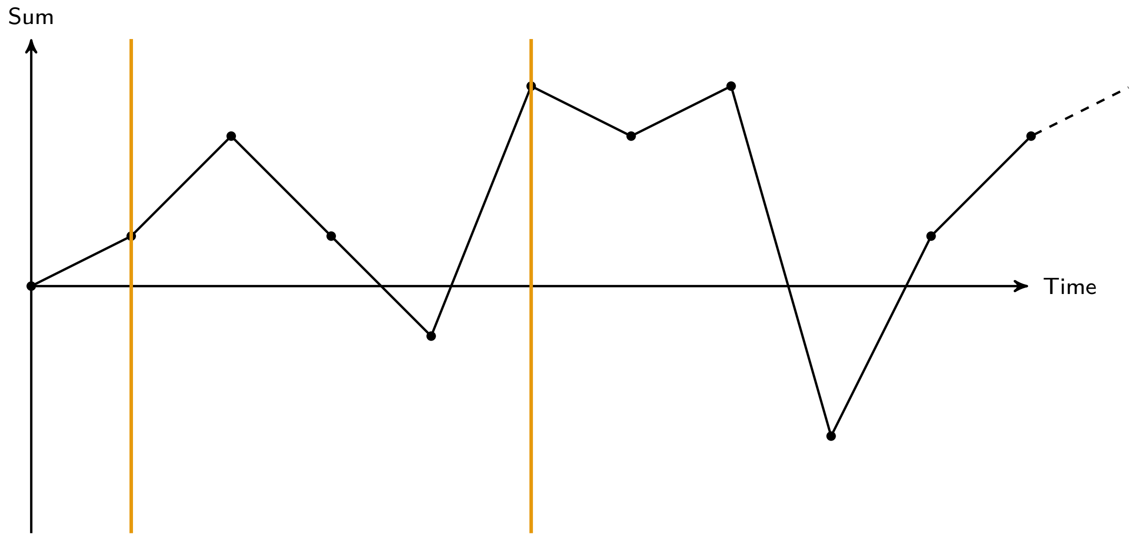

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

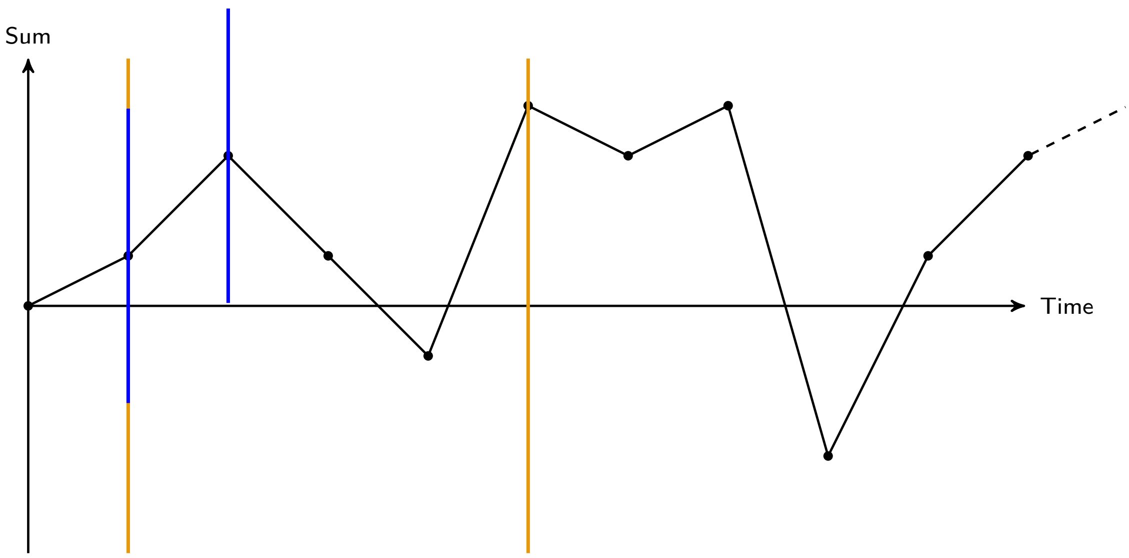

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\textbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

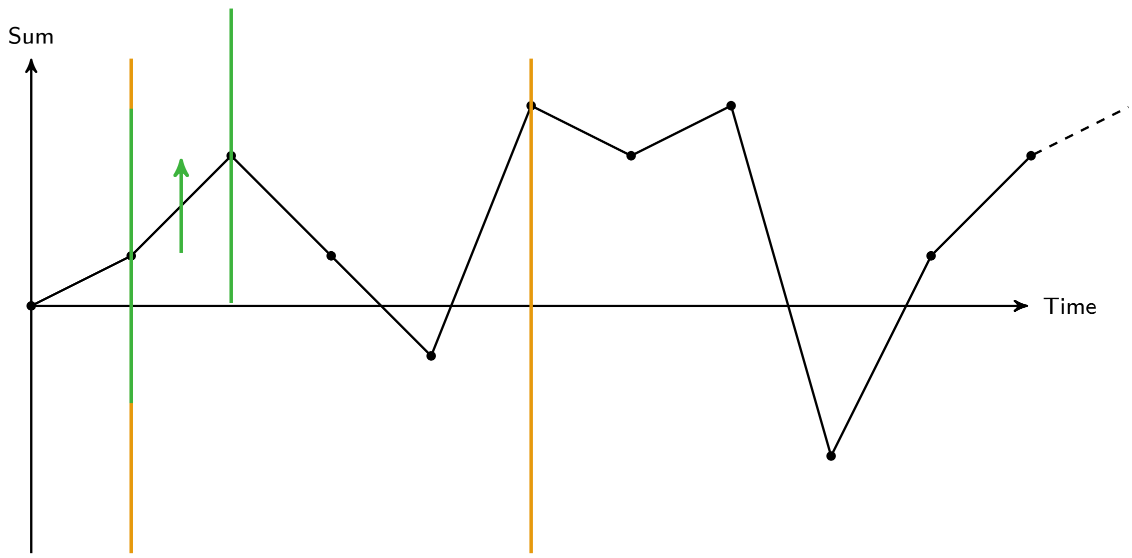

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

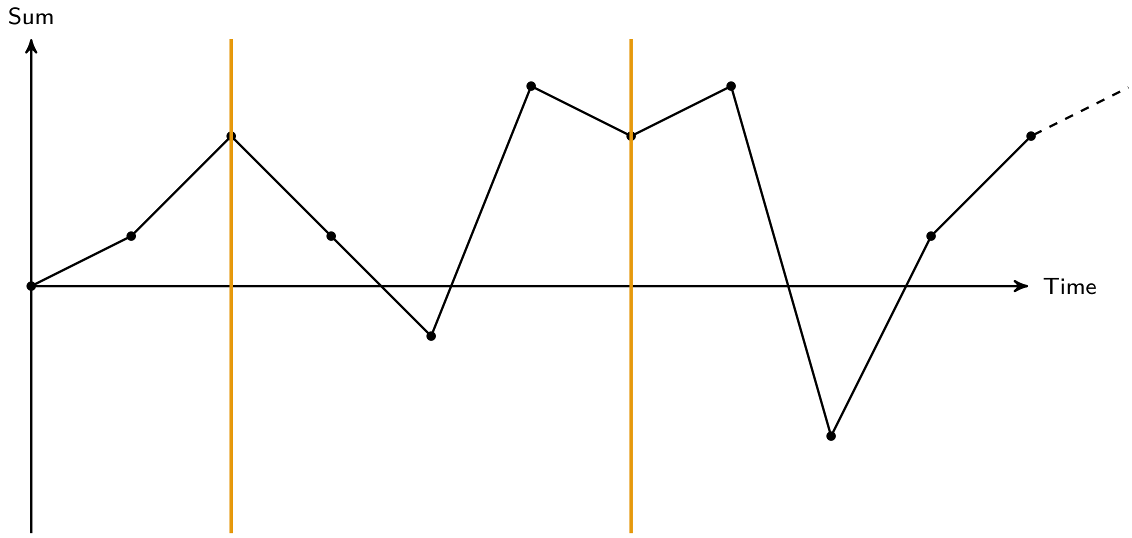

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

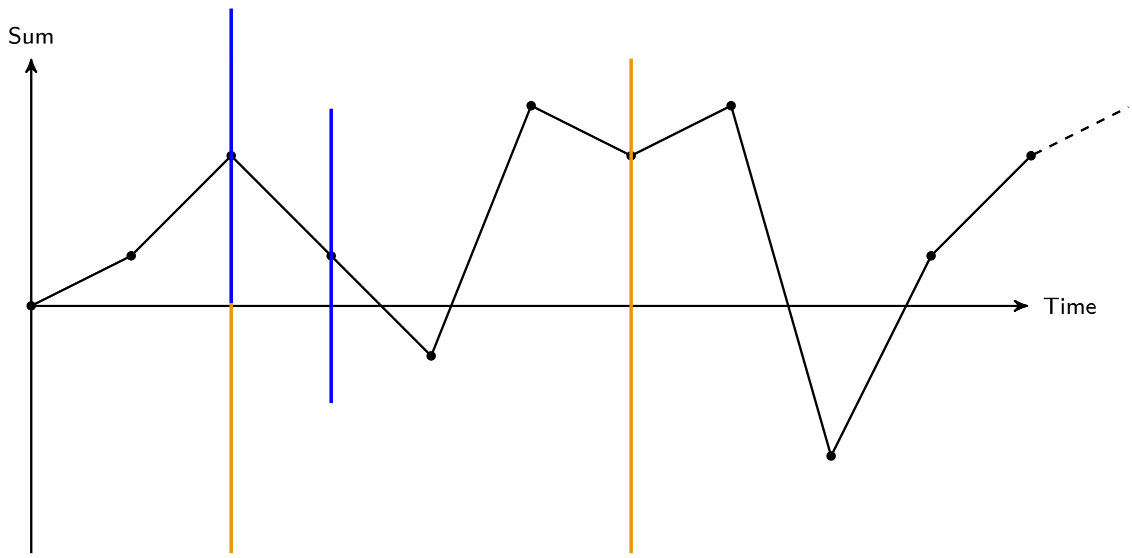

[1] Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

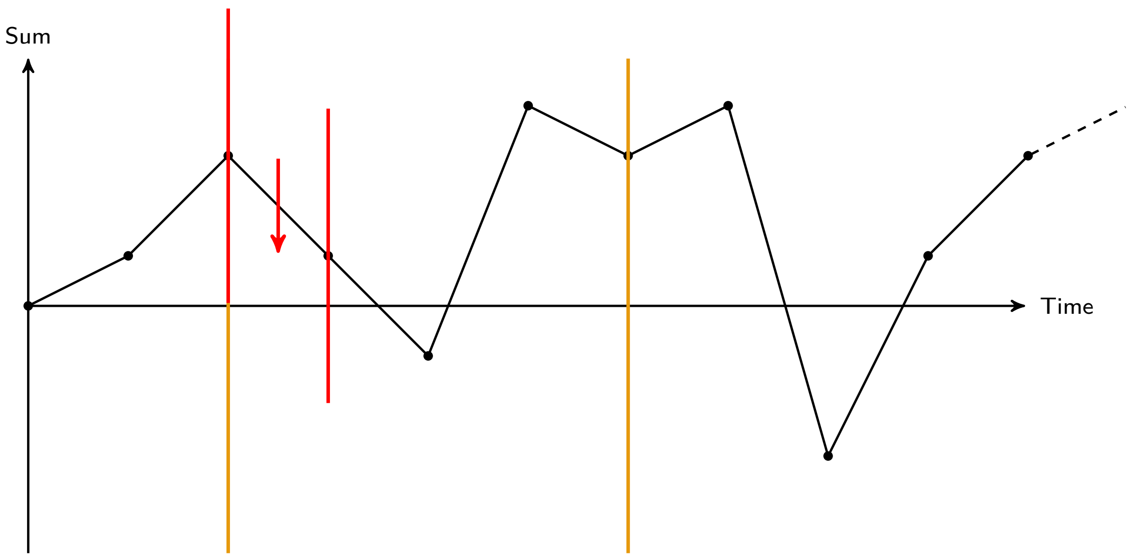

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

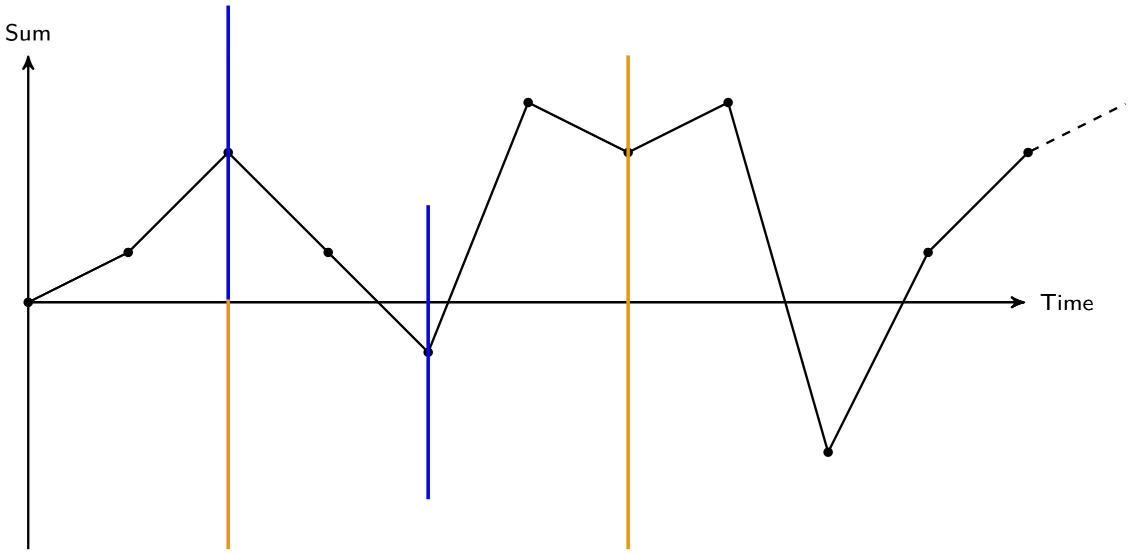

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

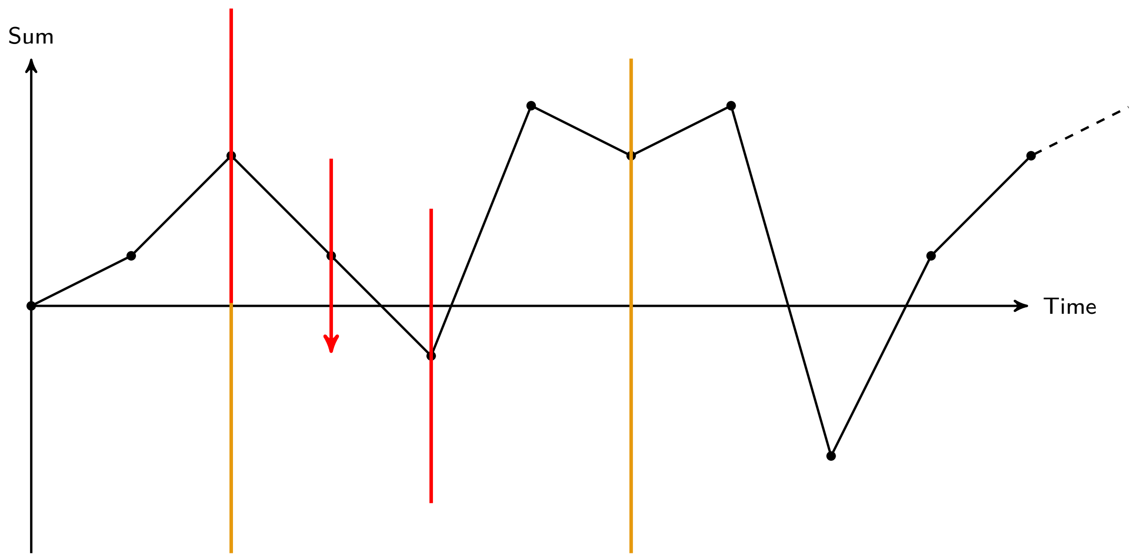

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

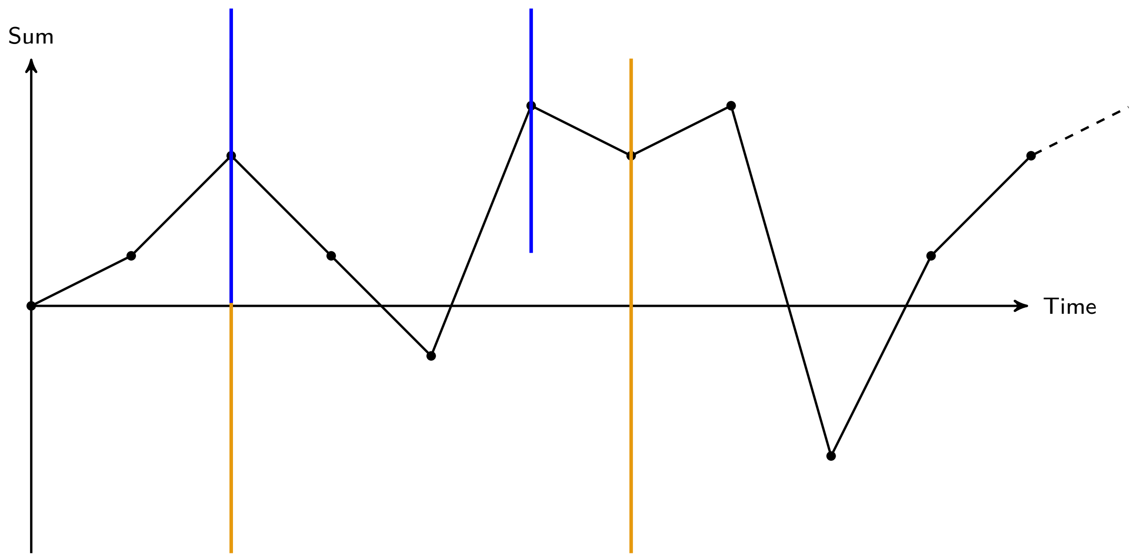

[1]Slide by M. Randour

Window MP, threshold zero, maximal window = 4

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Consider a fixed play and its running weight per turn

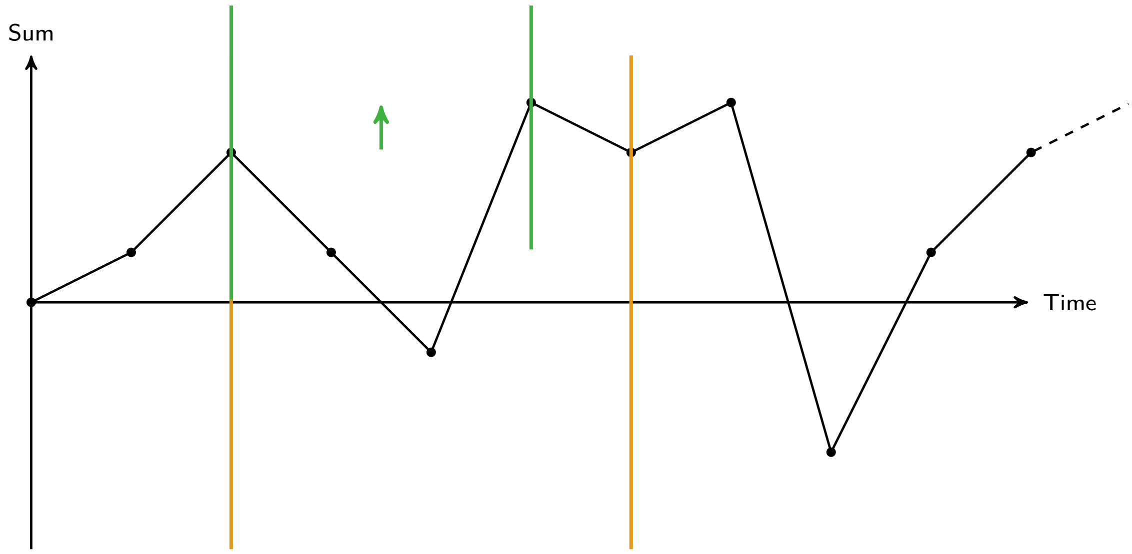

[1]Slide by M. Randour

Multiple variants

Several window objectives were defined in [CDRR13]:

fixed window objectives (FWMP)

using **GW** and a given $\ell_{\max}$ of polynomial size

- a safety-like WMP objective
- a co-Büchi-like WMP objective

UNIVERSITÉ LIBRE DE BRUXELLES

Multiple variants

Several window objectives were defined in [CDRR13]:

fixed window objectives (FWMP)

using **GW** and a given $\ell_{\max}$ of polynomial size

- a safety-like WMP objective
- a co-Büchi-like WMP objective

bounded window objectives (BWMP)

quantifying $\ell_{\max}$ existentially

- safety and co-Büchi WMP objectives

ULB UNIVERSITÉ LIBRE DE BRUXELLES

Properties

Theorem (Conservative approximations [CDRR13])

If ∃ve wins any window objective, she wins the MPG for $M \geq 0$. If she wins the MPG for $MP > 0$, $\exists \ell_{max}$ with which she wins all FWMP.

[1] ℓ_{max} is assumed to be polynomial w.r.t. the input

Properties

Theorem (Conservative approximations [CDRR13])

If ∃ve wins any window objective, she wins the MPG for $M \geq 0$. If she wins the MPG for $MP > 0$, $\exists \ell_{max}$ with which she wins all FWMP.

Complexity[1]

	1-dim	k-dims
FWMP	P-c	$\in$ EXP, PSPACE-h
BWMP	$\in$ NP $\cap$ coNP, MPG-h	NPR-h

UNIVERSITÉ LIBRE DE BRUXELLES

[1] ℓ_{max} is assumed to be polynomial w.r.t. the input

Outline

UNIVERSITÉ LIBRE DE BRUXELLES

Window objectives

Idea/Motivation

Use the **GW** property of concrete plays to define:

- quantitative objectives that specify timing requirements
- which are stronger than MP.

UNIVERSITÉ LIBRE DE BRUXELLES

Window objectives

Idea/Motivation

Use the **GW** property of concrete plays to define:

- quantitative objectives that specify timing requirements
- which are stronger than MP.

Main issue

- Where to quantify over concretizations? e.g. for a bounded window objective, should all concretizations have the same $\ell_{\max}$?
- We consider all distinct cases.

PO variants (definitions)

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Using **GW** and a given $\ell_{\max}$ of polynomial size

Direct Fixed Window (i.e. safety-like)

$\mathbf{DirFix}(\ell_{\max}) \equiv \forall \pi, \Box, \mathbf{GW}(\ell_{\max}) \equiv \Box, \forall \pi, \mathbf{GW}(\ell_{\max})$

PO variants (definitions)

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Using **GW** and a given $\ell_{\max}$ of polynomial size

Direct Fixed Window (i.e. safety-like)

$\mathbf{DirFix}(\ell_{\max}) \equiv \forall\pi, \Box, \mathbf{GW}(\ell_{\max}) \equiv \Box, \forall\pi, \mathbf{GW}(\ell_{\max})$

Fixed Window (i.e. co-Büchi-like)

- $\mathbf{UFix}(\ell_{\max}) \equiv \Diamond, \forall\pi, \Box, \mathbf{GW}(\ell_{\max})$
- $\mathbf{Fix}(\ell_{\max}) \equiv \forall\pi, \Diamond, \Box, \mathbf{GW}(\ell_{\max})$

PO variants (definitions)

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Using **GW** and a given $\ell_{\max}$ of polynomial size

Direct Fixed Window (i.e. safety-like)

$\mathbf{DirFix}(\ell_{\max}) \equiv \forall \pi, \Box, \mathbf{GW}(\ell_{\max}) \equiv \Box, \forall \pi, \mathbf{GW}(\ell_{\max})$

Fixed Window (i.e. co-Büchi-like)

- $\mathbf{UFix}(\ell_{\max}) \equiv \Diamond, \forall \pi, \Box, \mathbf{GW}(\ell_{\max})$
- $\mathbf{Fix}(\ell_{\max}) \equiv \forall \pi, \Diamond, \Box, \mathbf{GW}(\ell_{\max})$

Quantifying $\ell_{\max}$ existentially: *(Direct) Bounded Window* . . .

UNIVERSITÉ LIBRE DE BRUXELLES

Categorizing objectives

$$
\begin{array}{ccccccccccc}
\text{DirFix} & & \subset & & \text{UFix} & \subset & \text{Fix} & & & & \\
\Downarrow & & & & \Downarrow & & \Downarrow & & & & \\
\text{UDirBnd} & \subset & \text{DirBnd} & \subset & \text{UBnd} & \subset & \text{Bnd} & \subset & \text{MPInf} & (\subset \text{MPSup}) &
\end{array}
$$

Figure: Implications among the objectives

Categorizing objectives

$$
\begin{array}{ccccccccccc}
\mathsf{DirFix} & & \subset & & \mathsf{UFix} & \subset & \mathsf{Fix} & & & & \\
\Downarrow & & & & \Downarrow & & \Downarrow & & & & \\
\mathsf{UDirBnd} & \subset & \mathsf{DirBnd} & \subset & \mathsf{UBnd} & \subset & \mathsf{Bnd} & \subset & \mathsf{MPInf} & (\subset \mathsf{MPSup}) &
\end{array}
$$

Figure: Implications among the objectives

Theorem (Conservative approximations)

If $\exists$ve wins any window objective, she wins the MPG for $MP \geq 0$. If she wins the MPG with finite memory for $\overline{MP} = \underline{MP} > 0$, $\exists \ell_{max}$ with which she wins all FWMP.

UNIVERSITÉ LIBRE DE BRUXELLES

Outline

Lower bounds

From k-dims window objectives

We can "simulate" several dimensions using PO, so bounded window objs. are at least NPR-h.

UNIVERSITÉ LIBRE DE BRUXELLES

Lower bounds

From k-dims window objectives

We can "simulate" several dimensions using PO, so bounded window objs. are at least NPR-h.

From safety games with PO

Even safety games with PO is EXP-h [BD08, AG11]. All window objs. can simulate safety games. Hence, they are all EXP-h.

ULB UNIVERSITÉ LIBRE DE BRUXELLES

Lower bounds

From k-dims window objectives

We can "simulate" several dimensions using PO, so bounded window objs. are at least NPR-h.

From safety games with PO

Even safety games with PO is EXP-h [BD08, AG11]. All window objs. can simulate safety games. Hence, they are all EXP-h.

From weighted automata

Using a reduction from the UNIVERSALITY PROBLEM for weighted automata, we show all bounded window objs. are undecidable.

More details on the undecidability. . .

Weighted automata

A weighted finite automaton $(Q, q_I, \Sigma, \Delta, w)$ realizes a function $\Sigma^* \to \mathbb{Z}$ by mapping a word x to the minimal (sum) value of its runs on x.

More details on the undecidability...

Weighted automata

A weighted finite automaton $(Q, q_I, \Sigma, \Delta, w)$ realizes a function $\Sigma^* \to \mathbb{Z}$ by mapping a word x to the minimal (sum) value of its runs on x.

Universality

A weighted automaton is universal if $\forall x \in \Sigma^*$ the value assigned by the automaton to x is negative. Undecidability from [ABK11].

More details on the undecidability...

Weighted automata

A weighted finite automaton $(Q, q_I, \Sigma, \Delta, w)$ realizes a function $\Sigma^* \to \mathbb{Z}$ by mapping a word x to the minimal (sum) value of its runs on x.

Universality

A weighted automaton is universal if $\forall x \in \Sigma^*$ the value assigned by the automaton to x is negative. Undecidability from [ABK11].

With the bounded window objectives, we construct simple gadgets to:

- force $\exists$ve to play infinitely many $\#$ (a fresh new symbol),

ULB UNIVERSITÉ LIBRE DE BRUXELLES

More details on the undecidability. . .

Weighted automata

A weighted finite automaton $(Q, q_I, \Sigma, \Delta, w)$ realizes a function $\Sigma^* \to \mathbb{Z}$ by mapping a word x to the minimal (sum) value of its runs on x.

Universality

A weighted automaton is universal if $\forall x \in \Sigma^*$ the value assigned by the automaton to x is negative. Undecidability from [ABK11].

With the bounded window objectives, we construct simple gadgets to:

- force ∃ve to play infinitely many # (a fresh new symbol),
- force her to play # in intervals of bounded length, and

UNIVERSITÉ LIBRE DE BRUXELLES

More details on the undecidability...

Weighted automata

A weighted finite automaton $(Q, q_I, \Sigma, \Delta, w)$ realizes a function $\Sigma^* \to \mathbb{Z}$ by mapping a word x to the minimal (sum) value of its runs on x.

Universality

A weighted automaton is universal if $\forall x \in \Sigma^*$ the value assigned by the automaton to x is negative. Undecidability from [ABK11].

With the bounded window objectives, we construct simple gadgets to:

- force $\exists$ve to play infinitely many $\#$ (a fresh new symbol),
- force her to play $\#$ in intervals of bounded length, and
- when $\exists$ve plays a word $\#x\#$ with positive value, she closes a good window.

Outline

DirFix games: a game on functions

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Direct Fixed Window:
$\mathbf{DirFix}(\ell_{\max}) \equiv \forall\pi, \Box, \mathbf{GW}(\ell_{\max}) \equiv \Box, \forall\pi, \mathbf{GW}(\ell_{\max})$

DirFix games: $\forall\pi, \Box, \mathbf{GW}(\ell_{\max})$

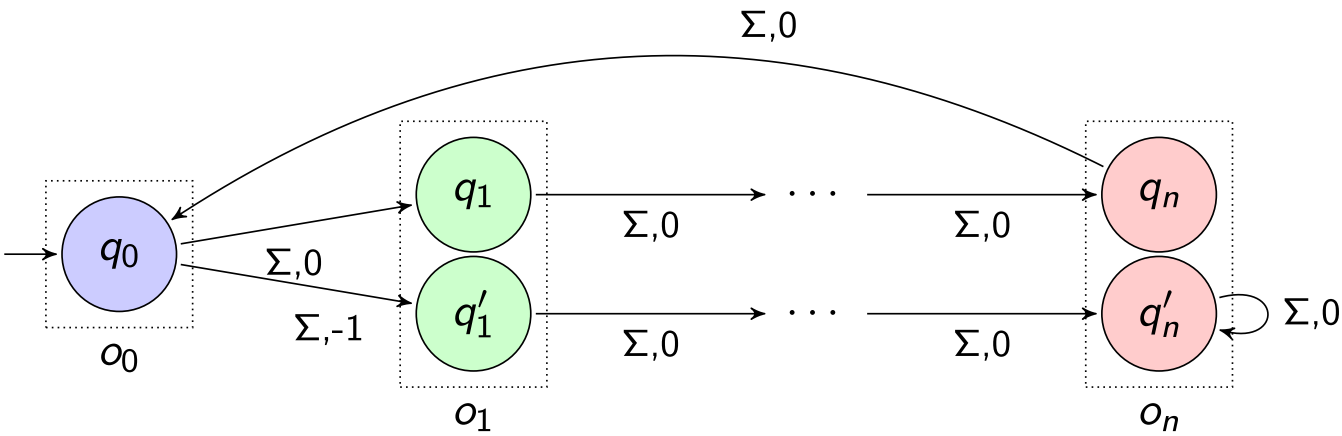

[1] $n = \ell_{\max} + 1$

DirFix games: $\forall\pi, \Box, \mathbf{GW}(\ell_{\max})$

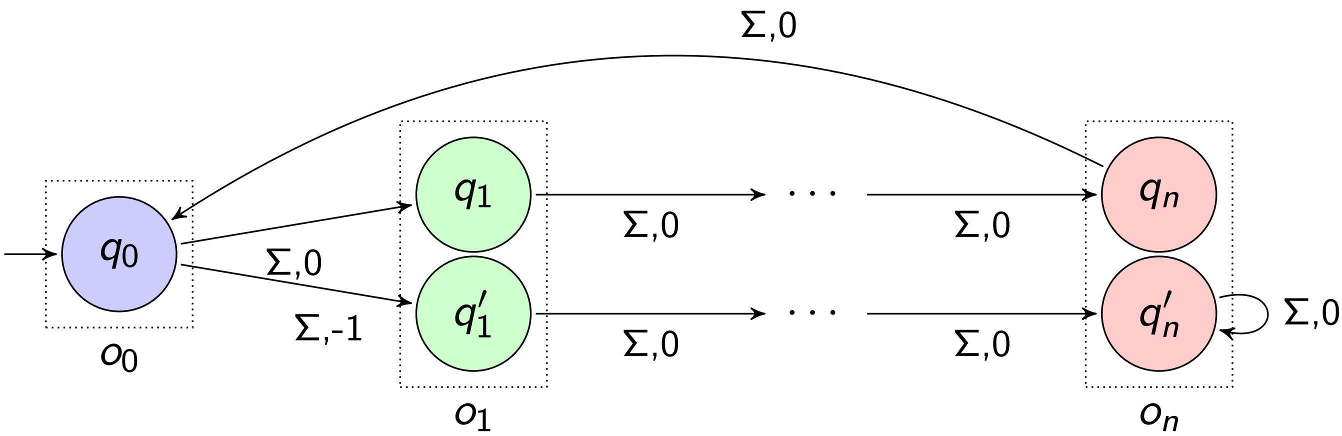

obs:
f. play:
curr. $f \in Q \to (\{1, \ldots, \ell_{\max}\} \to \{-W \cdot \ell_{\max}, \ldots 0\})$

[1] $n = \ell_{\max} + 1$

DirFix games: $\forall\pi, \Box, \mathbf{GW}(\ell_{\max})$

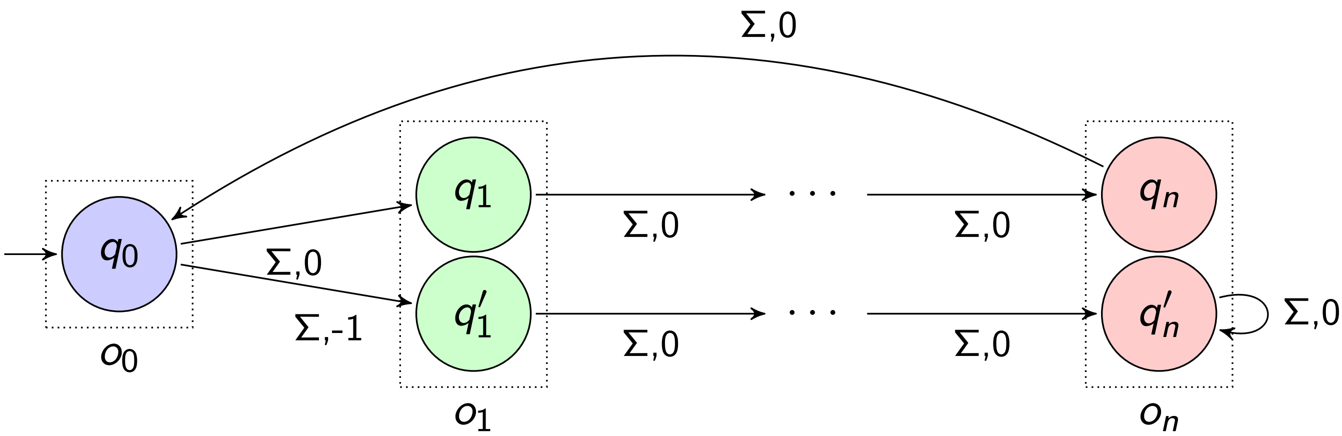

obs: o_0
f. play: f_0
curr. $f \in Q \to (\{1, \ldots, \ell_{\max}\} \to \{-W \cdot \ell_{\max}, \ldots 0\})$
$q_0 \mapsto (i \mapsto 0, \forall i)$

[1] $n = \ell_{\max} + 1$

DirFix games: $\forall\pi, \Box$, **GW**$(\ell_{\max})$

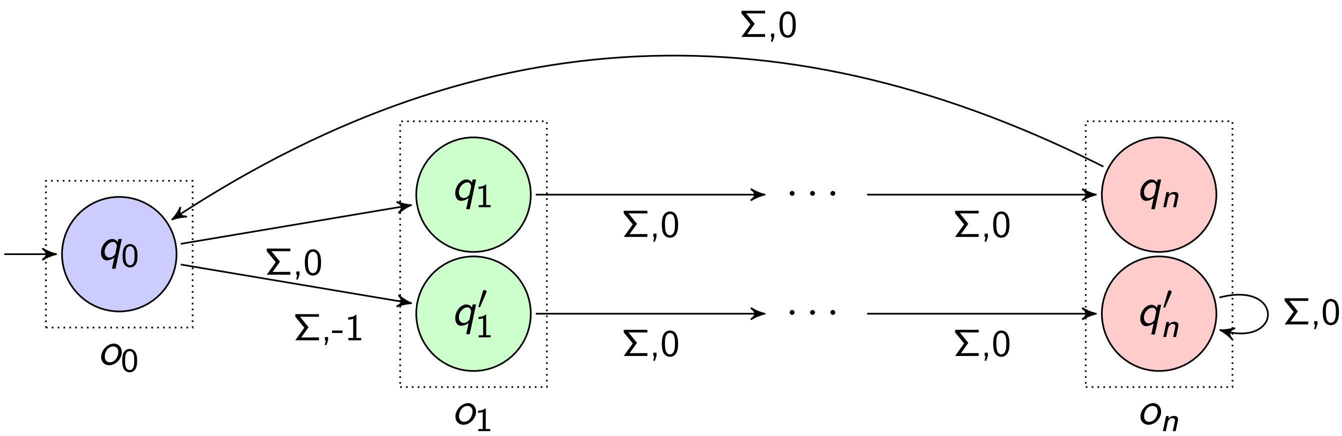

obs: $o_0\sigma_0o_1$

f. play: $f_0\sigma_0f_1$

curr. $f \in Q \to (\{1,\ldots,\ell_{\max}\} \to \{-W\cdot\ell_{\max},\ldots 0\})$

$q_1 \mapsto (i \mapsto 0, \forall i), q_1' \mapsto (1 \mapsto -1)$

[1] $n = \ell_{\max} + 1$

DirFix games: $\forall\pi, \Box$, **GW**(ℓ_{max})

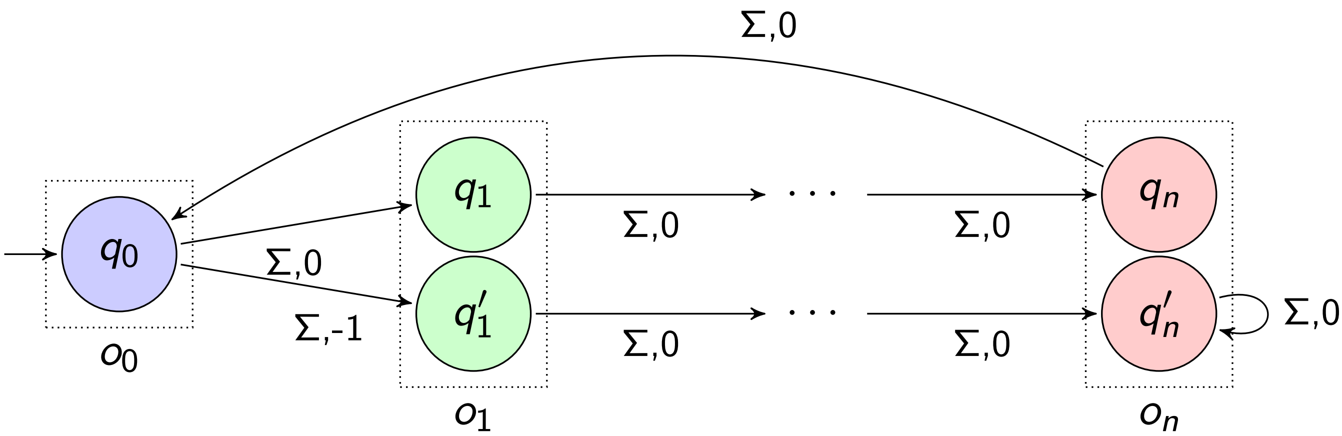

obs: $o_0\sigma_0 o_1\sigma_1 o_2$
f. play: $f_0\sigma_0 f_1\sigma_1 f_2$
curr. $f \in Q \to (\{1,\ldots,\ell_{max}\} \to \{-W\cdot\ell_{max},\ldots 0\})$
$q_2 \mapsto (i \mapsto 0, \forall i), q'_2 \mapsto (2 \mapsto -1)$

[1] $n = \ell_{max} + 1$

DirFix games: $\forall\pi, \square$, **GW**$(\ell_{\max})$

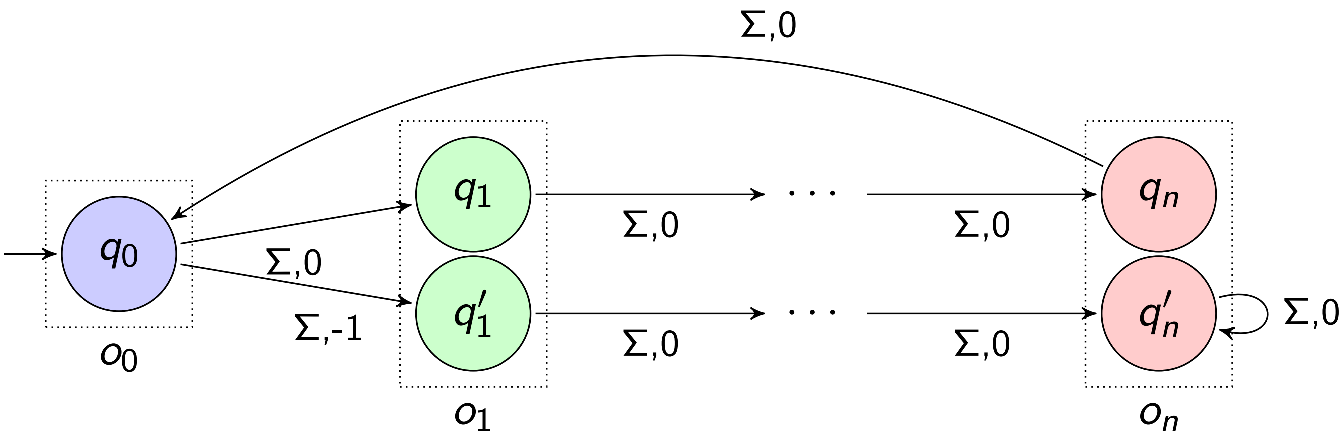

obs: $o_0\sigma_0 o_1\sigma_1 o_2 \ldots \sigma_{n-1} o_n$

f. play: $f_0\sigma_0 f_1\sigma_1 f_2 \ldots \sigma_{n-1} f_n$

curr. $f \in Q \to (\{1, \ldots, \ell_{\max}\} \to \{-W \cdot \ell_{\max}, \ldots 0\})$

$q_n \mapsto (i \mapsto 0, \forall i), q'_n \mapsto (n \mapsto -1)$

a window open for $> \ell_{\max}$ steps in q'_n!

[1] $n = \ell_{\max} + 1$

Fix games: using observers

$$\mathbf{GW}(i, \ell_{\text{max}}) \equiv \exists i \leq \ell_{\text{max}} : w(q_i \ldots q_{i+j}) \geq 0$$

Fixed Window: $\mathbf{Fix}(\ell_{\text{max}}) \equiv \forall \pi, \Diamond, \Box, \mathbf{GW}(\ell_{\text{max}})$

Fix games: using observers

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

Fixed Window: $\mathbf{Fix}(\ell_{\max}) \equiv \forall \pi, \Diamond, \Box, \mathbf{GW}(\ell_{\max})$

The complement of the objective

$\exists \pi, \Box, \Diamond, \neg\mathbf{GW}(\ell_{\max})$, i.e. there is a conretization with ∞-ly many bad windows.

Fix games: $\forall\pi, \Diamond, \Box,$ **GW**$(\ell_{\max})$

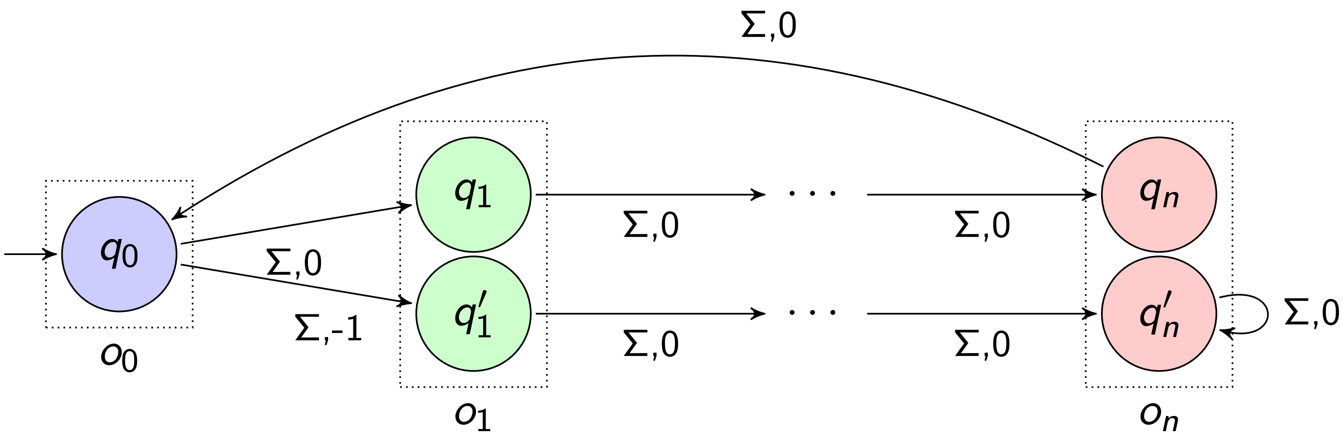

Should we play a co-Büchi game now?

[1] $n = \ell_{\max} + 1$

Fix games: $\forall\pi, \Diamond, \Box$, GW($\ell_{max}$)

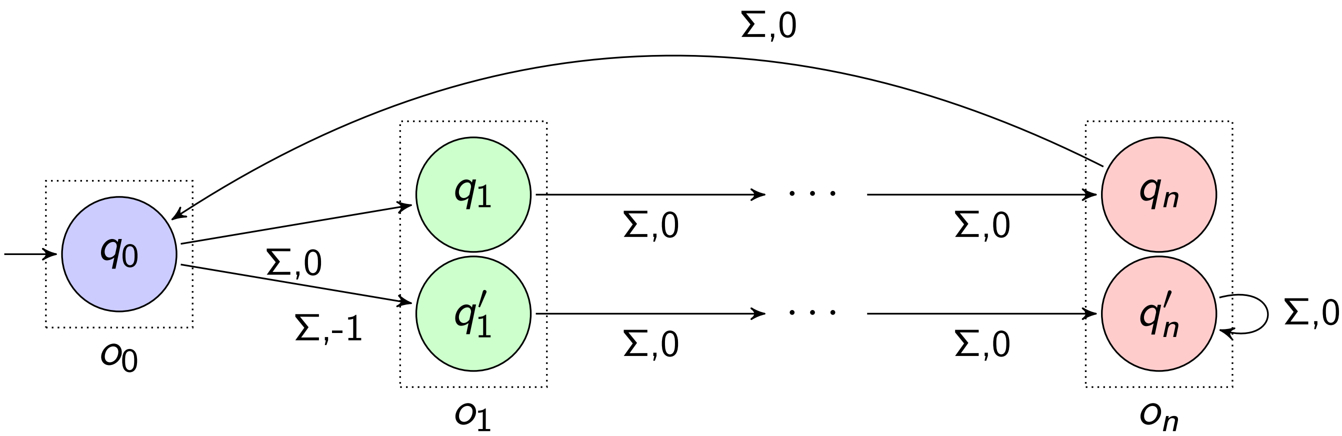

We build a non-det. poly-size Büchi observer:

- e.g. q_1, q'_1 are reachable from q_0 reading (Σ, o_1);
- states now also encode open window information, i.e. functions $\{1, \ldots, \ell_{max}\} \to \{-W \cdot \ell_{max}, \ldots 0\}$);
- states with $f(\ell_{max}) < 0$ are now accepting.

[1] $n = \ell_{max} + 1$

Fix games: $\forall\pi, \Diamond, \Box$, **GW**$(\ell_{\mathsf{max}})$

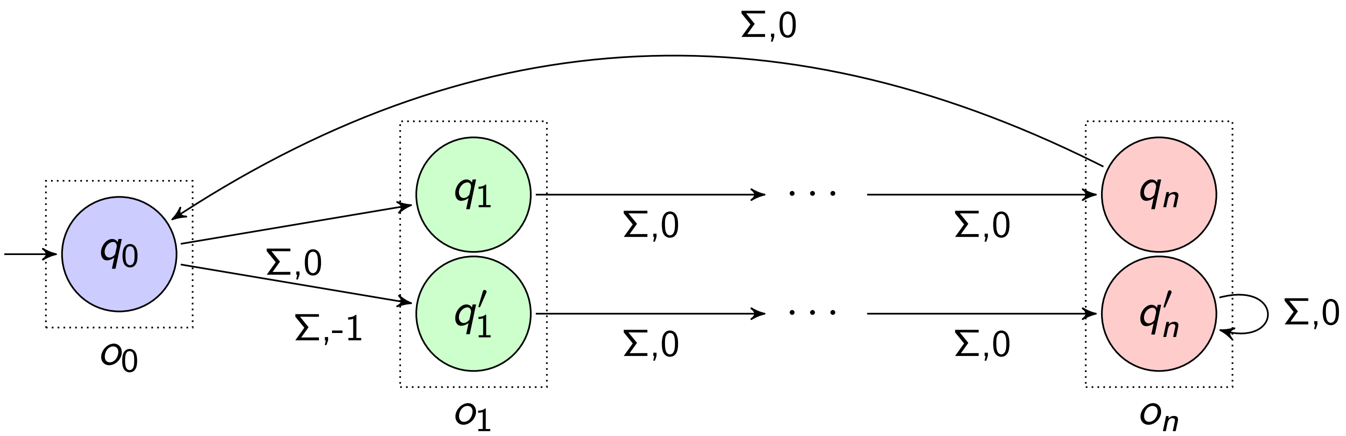

E.g., $q_0, (i \mapsto 0, \forall i)$ non-deterministically transitions to

- $q'_1, (1 \mapsto -1)$ or
- $q_1, (i \mapsto 0, \forall i)$

with input (σ, o_1).

[1] $n = \ell_{\mathsf{max}} + 1$

UFix games: uniformity

$$\mathbf{GW}(i, \ell_{\max}) \equiv \exists i \leq \ell_{\max} : w(q_i \ldots q_{i+j}) \geq 0$$

(Uniform) Fixed Window: $\mathbf{UFix}(\ell_{\max}) \equiv \Diamond, \forall\pi, \Box, \mathbf{GW}(\ell_{\max})$

The complement of the objective

$\Box, \exists\pi, \Diamond, \neg\mathbf{GW}(\ell_{\max})$, i.e. ∞-ly many bad windows.

UNIVERSITÉ LIBRE DE BRUXELLES

UFix: uniformity (Ex.)

Does ∃ve win for **Fix**? for **UFix**?

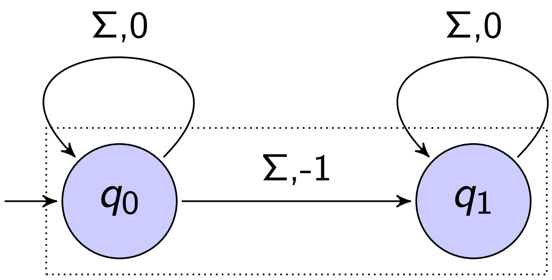

UFix: uniformity (merging trick)

Lemma

An abstract play is not in **UFix**(ℓ_{max}) *iff it has a concretization which merges with infinitely many violating concrete plays.*

UFix: uniformity (merging trick)

Lemma

An abstract play is not in **UFix**(ℓ_{max}) *iff it has a concretization which merges with infinitely many violating concrete plays.*

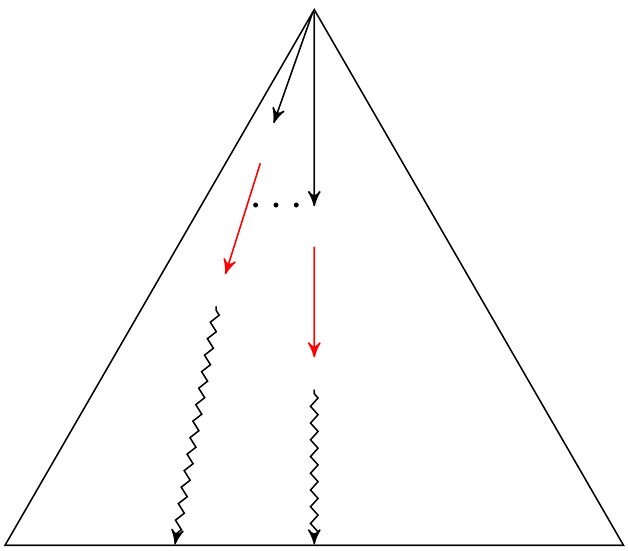

UFix: uniformity (merging trick)

Lemma

An abstract play is not in **UFix**(ℓ_{max}) *iff it has a concretization which merges with infinitely many violating concrete plays.*

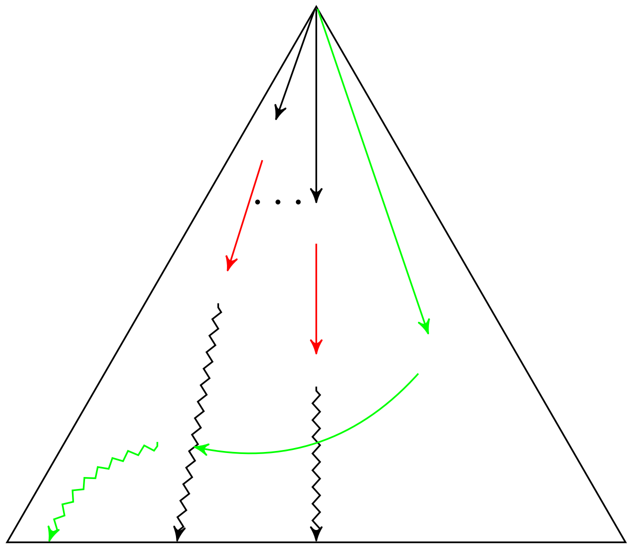

UFix: uniformity (merging trick)

Lemma

An abstract play is not in **UFix**(ℓ_{max}) *iff it has a concretization which merges with infinitely many violating concrete plays.*

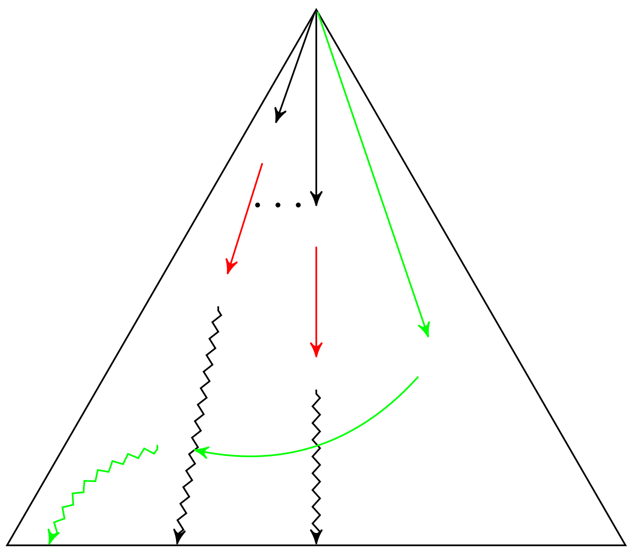

Build another observer to recognize such concretizations!

Conclusions

To summarize:

- we extended window objectives to the partial-observation setting,
- obtained a decidable conservative approximation of MPGs (i.e. the FWMP objectives) with **DirFix** games even allowing for an order on the state-space, and
- established tight complexity bounds for their threshold problem.

DirFix	EXP-c
Fix	EXP-c
UFix	EXP-c
Bounded	Undecidable

Still open: full-observation k-dimension BWMP

UNIVERSITÉ LIBRE DE BRUXELLES

References I

- Shaull Almagor, Udi Boker, and Orna Kupferman.
  What's decidable about weighted automata?
  In ATVA, pages 482–491. Springer, 2011.
- Krzysztof R. Apt and Erich Grädel.
  Lectures in game theory for computer scientists.
  Cambridge University Press, 2011.
- Dietmar Berwanger and Laurent Doyen.
  On the power of imperfect information.
  In FSTTCS, pages 73–82, 2008.
- Krishnendu Chatterjee, Laurent Doyen, Mickaël Randour, and Jean-François Raskin.
  Looking at mean-payoff and total-payoff through windows.
  In ATVA, pages 118–132. Springer, 2013.

References II

- Aldric Degorre, Laurent Doyen, Raffaella Gentilini, Jean-François Raskin, and Szymon Toruńczyk.
  Energy and mean-payoff games with imperfect information.
  In CSL, pages 260–274, 2010.
- A. Ehrenfeucht and J. Mycielski.
  Positional strategies for mean payoff games.
  International Journal of Game Theory, 8:109–113, 1979.
- Paul Hunter, Guillermo A. Pérez, and Jean-François Raskin.
  Mean-payoff games with partial-observation.
  In RP, LNCS. Springer, 2014.
- Uri Zwick and Mike Paterson.
  The complexity of mean payoff games on graphs.
  TCS, 158(1):343–359, 1996.