Omega-Regular Half-Positional Winning Conditions

Eryk Kopczyński

Warsaw University

June 9, 2008

Outline

example

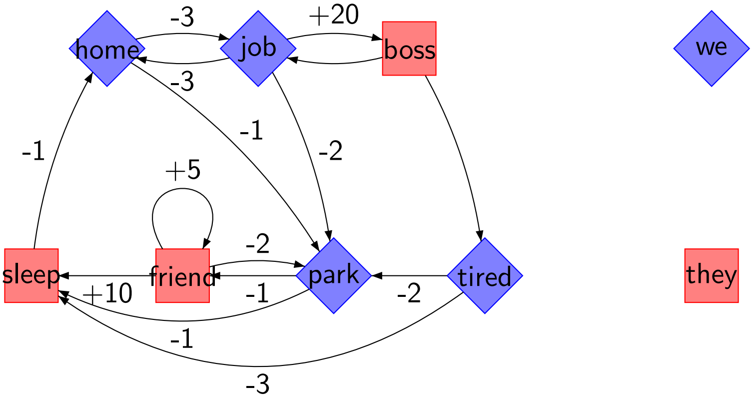
-3
+20
home
job
boss
we
-3
-1
-1
-2
+5
-2
sleep
friend
park
tired
they
+10
-1
-2
-1
-3

example: parity

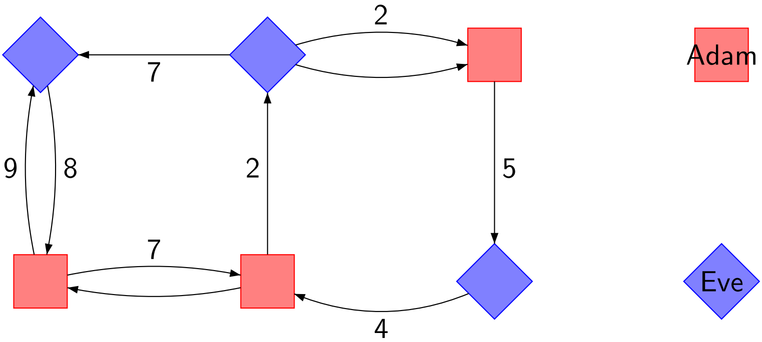

Eve wins iff the greatest number appearing infinitely often during an infinite play is even.

example: parity...

We create a computer program for our business. Good or bad things could happen...

the program uses too much resources

the program hangs

the program works as it should

we lose some money

we break our moral rules

we earn some money

we become rich

newspapers write about us

we go to jail

example: parity...

We create a computer program for our business. Good or bad things could happen...

the program uses too much resources	1
the program hangs	1
the program works as it should	2
we lose some money	3
we break our moral rules	3
we earn some money	4
we become rich	4
newspapers write about us	4
we go to jail	5

Games

C = set of colors

Game = arena + winning condition

Arena:
$G = (\mathrm{Pos}_A, \mathrm{Pos}_E, \mathrm{Mov})$ where $\mathrm{Pos} = \mathrm{Pos}_A \cup \mathrm{Pos}_E$,
$\mathrm{Mov} \subseteq \mathrm{Pos} \times \mathrm{Pos} \times (C \cup \{\epsilon\})$

Winning condition:
Subset $W \subseteq C^\omega$; we assume that it is prefix independent, i.e.
$u \in W \iff cu \in W$

Plays and strategies

A play π is a sequence of moves such that $\mathrm{source}(\pi_{n+1}) = \mathrm{target}(\pi_n)$.

A strategy for Eve (Adam) is a partial function $s : \mathrm{Pos} \cup \mathrm{Mov}^* \to \mathrm{Mov}$ which tells Eve (Adam) what they should do in a given situation (the current position, history so far).

A strategy s is winning for X if each play consistent with s is winning for X.

A strategy s is positional if $s(\pi)$ depends only on $\mathrm{target}(\pi)$.

Determinacy

Definition

A game (G, W) is determined if for each starting position one of players has a winning strategy. (Not all games are determined.) If the game is determined, we have $\mathrm{Pos} = \mathrm{Win}_E \cup \mathrm{Win}_A$ and strategies s_E and s_A such that each play π with $\mathrm{source}(\pi) \in \mathrm{Win}_X$ and consistent with s_X is winning for X.

Determinacy types

Definition

A determinacy type is given by three parameters:

- admissible strategies for Eve: positional or arbitrary
- admissible strategies for Adam: positional or arbitrary
- admissible arenas: finite or infinite

Definition

A winning condition W is (α, β, γ)-determined, if for each γ-arena G the game (G, W) is (α, β)-determined, i.e. for each starting position either Eve has a winning α-strategy or Adam has a winning β-strategy.

Half-positional conditions

For short, we call (positional, arbitrary, infinite)-determined conditions half-positional, and (positional, arbitrary, finite)-determined conditions finitely half-positional.

We will focus on half-positional and finitely half-positional winning conditions.

Büchi and co-Büchi conditions

Definition

Büchi condition: $WB_S = C^*(SC)^\omega$
Eve wants colors from S to appear infinitely often.
co-Büchi condition: $WB'_S = C^*(C - S)^\omega$
Eve wants colors from S to appear finitely often.

Both of these classes winning conditions are positional.

. . . and their union

Example

$C = \{a, b, c\}$, $W = WB'_{\{a\}} \cup WB'_{\{b\}}$. Eve wants at least one of a and b to appear only finitely often.

. . . and their union

Example

$C = \{a, b, c\}$, $W = WB'_{\{a\}} \cup WB'_{\{b\}}$. Eve wants at least one of a and b to appear only finitely often.

Why half-positional: If Eve can win, then ultimately she always can avoid one of the letters. This can be done with a positional strategy.

. . . and their union

Example

$C = \{a, b, c\}$, $W = WB'_{\{a\}} \cup WB'_{\{b\}}$. Eve wants at least one of a and b to appear only finitely often.

Why half-positional: If Eve can win, then ultimately she always can avoid one of the letters. This can be done with a positional strategy.

Witness for non-positionality: One Adam's position with two moves, a and b.

n letters a in a row

Example

$C = \{a, b\}$, Eve wants a^n to appear only finitely often.

n letters a in a row

Example

$C = \{a, b\}$, Eve wants a^n to appear only finitely often.

Why half-positional: Eve can always assume the worst case (the greatest possible number of a so far). This can be done with a positional strategy.

n letters a in a row

Example

$C = \{a, b\}$, Eve wants a^n to appear only finitely often.

Why half-positional: Eve can always assume the worst case (the greatest possible number of a so far). This can be done with a positional strategy.

Witness for non-positionality: One Adam's position with two moves, $a^{n-1}b$ and ba^{n-1}.

Applications

- automata theory (automata on infinite structures)
- modal μ-calculus
- model checking
- interactive systems

Facts about determinacy

Theorem (Martin, 1975)

All Borel winning conditions are determined.

Facts about determinacy

Theorem (Martin, 1975)

All Borel winning conditions are determined.

Theorem (Emerson-Jutla, Mostowski 1991)

The parity condition

$$WP_n = \{w \in \{0, \ldots, n\}^\omega : 2 | \limsup_{n\to\infty} w_n\} \tag{1}$$

is positionally determined.

Facts about determinacy

Theorem (Ehrenfeucht, Mycielski 1979)

The mean payoff games are finitely positionally determined.

Facts about determinacy

Theorem (Ehrenfeucht, Mycielski 1979)

The mean payoff games are finitely positionally determined.

Theorem (Klarlund 1992)

The Rabin condition is half-positional.

Facts about determinacy

Any Borel winning condition	(arbitrary, arbitrary, infinite)
Parity condition	(positional, positional, infinite)
Mean payoff	(positional, positional, finite)
Rabin condition	(positional, arbitrary, infinite)

Remark

If W is (α, β, γ)-determined, then $C^\omega - W$ is (β, α, γ)-determined.

Three types of arenas

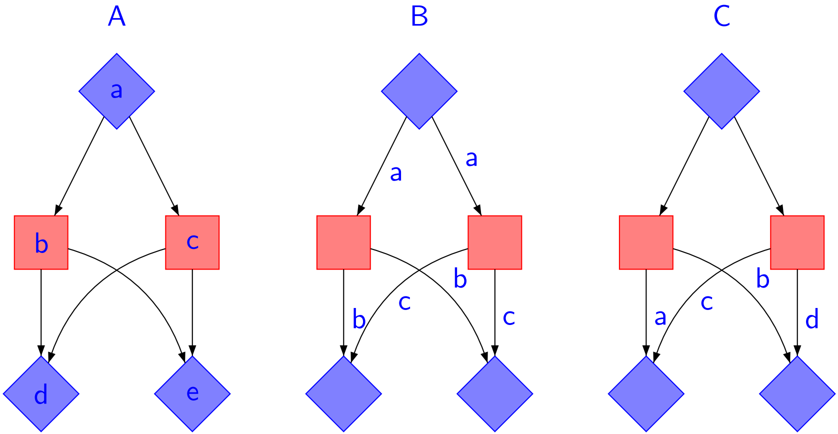

Positional determinacy characterizations

Theorem (Colcombet, Niwiński 2006)

A (prefix independent) winning condition $W \subseteq C^\omega$ is positional iff it is a generalized parity condition, i.e. there is a mapping $h : C \to \{0, 1, \ldots, n\}$ such that $u \in W$ iff $h(u) \in WP_n$.

Note: this theorem requires edge-colored (B) or epsilon-arenas (C) — if we restrict to position-colored arenas (A) there are more positional winning conditions, for example $C^*(ab)^*$ or min-parity.

Positional determinacy characterizations

Theorem (Colcombet, Niwiński 2006)

A (prefix independent) winning condition $W \subseteq C^\omega$ is positional iff it is a generalized parity condition, i.e. there is a mapping $h : C \to \{0, 1, \ldots, n\}$ such that $u \in W$ iff $h(u) \in WP_n$.

Note: this theorem requires edge-colored (B) or epsilon-arenas (C) — if we restrict to position-colored arenas (A) there are more positional winning conditions, for example $C^*(ab)^*$ or min-parity.

Theorem (Gimbert, Zielonka 2005)

A winning condition $W \subseteq C^\omega$ is finitely positional iff the winner can win positionally for all arenas where all positions belong to the same player.

Examples of half-positional winning conditions

Concavity

Definition

A winning condition W is convex if for all sequences of words (u_n), $u_n \in C^*$, if

- $u_1 u_3 u_5 u_7 \ldots \in W$,
- $u_2 u_4 u_6 u_8 \ldots \in W$,

then $u_1 u_2 u_3 u_4 \ldots \in W$.

A winning condition is concave if its complement is convex.

Concavity

Theorem

Concave winning conditions are finitely half-positional.

Concavity

Theorem

Concave winning conditions are finitely half-positional.

Example

The parity conditions are both concave and convex.

Concavity

Theorem

Concave winning conditions are finitely half-positional.

Example

The parity conditions are both concave and convex.

Not all half-positional conditions are concave.

Weakening the assumptions

Theorem (Gimbert, Zielonka)

A winning condition that is both weakly convex and weakly concave is finitely positional.

Definition

A winning condition W is weakly convex if for all sequences of words (u_n), $u_n \in C^*$, if

- $u_1 u_3 u_5 u_7 \ldots \in W$ and $u_2 u_4 u_6 u_8 \ldots \in W$,
- for all i we have $(u_i)^\omega \in W$,

then $u_1 u_2 u_3 u_4 \ldots \in W$.

Weakening the assumptions

Theorem (Gimbert, Zielonka)

A winning condition that is both weakly convex and weakly concave is finitely positional.

Definition

A winning condition W is weakly convex if for all sequences of words (u_n), $u_n \in C^*$, if

- $u_1 u_3 u_5 u_7 \ldots \in W$ and $u_2 u_4 u_6 u_8 \ldots \in W$,
- for all i we have $(u_i)^\omega \in W$,

then $u_1 u_2 u_3 u_4 \ldots \in W$.

However, weak concavity is not a sufficient condition for half-positional determinacy.

Geometrical conditions

Let $C = [0,1]^d$.
For $u \in C^+$, let $P(u)$ be the average color of u.
For $w \in C^\omega$, let $P_n(w) = P(w_{|n})$.

w_n	0	1	0	1	0	1	0	1	0	...
$P_n(w)$	0	1/2	1/3	2/4	2/5	3/6	3/7	4/8	4/9	...

Geometrical conditions

Let $C = [0,1]^d$.
For $u \in C^+$, let $P(u)$ be the average color of u.
For $w \in C^\omega$, let $P_n(w) = P(w_{|n})$.

w_n	0	1	0	1	0	1	0	1	0	...
$P_n(w)$	0	1/2	1/3	2/4	2/5	3/6	3/7	4/8	4/9	...

Let A be a subset of C. Let $WF(A) \subset C^\omega$ be a set of w such that each cluster point of $(P_n(w))$ is an element of A, and $WF'(A)$ be a set of w such that at least one cluster point of $(P_n(w))$ is an element of A.

Half-positional determinacy vs geometry

For which A's $WF(A)$ and $WF'(A)$ are half-positional?

Half-positional determinacy vs geometry

For which A's $WF(A)$ and $WF'(A)$ are half-positional?

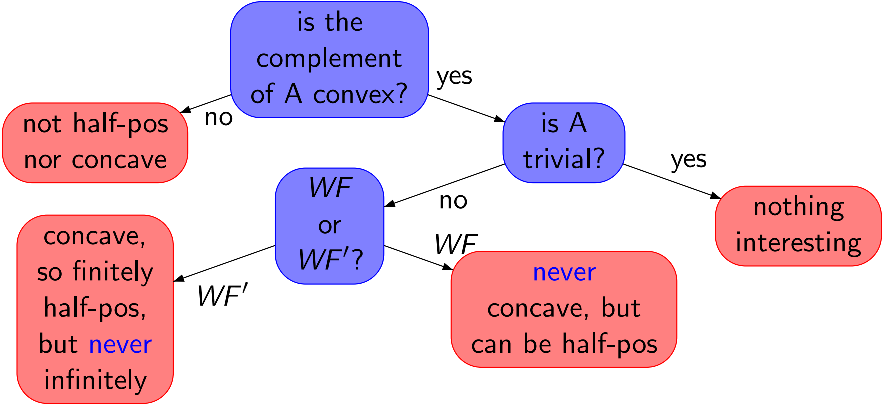

Half-positional determinacy vs geometry

For which A's $WF(A)$ and $WF'(A)$ are half-positional?

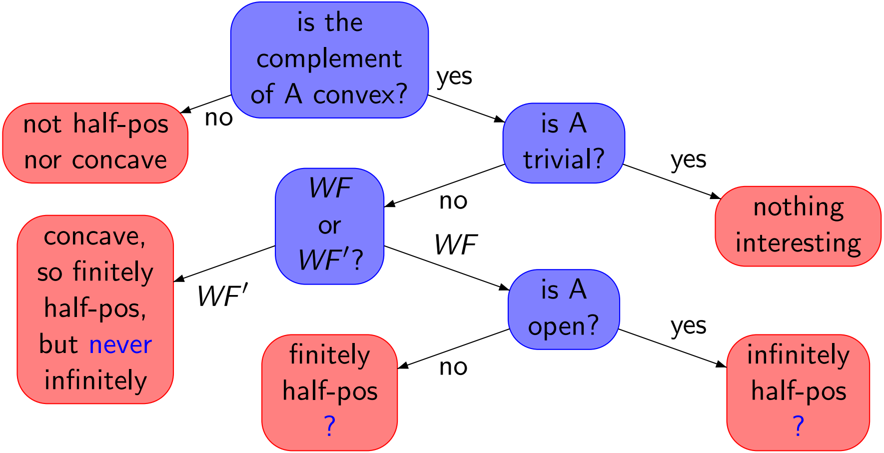

Half-positional determinacy vs geometry

For which A's $WF(A)$ and $WF'(A)$ are half-positional?

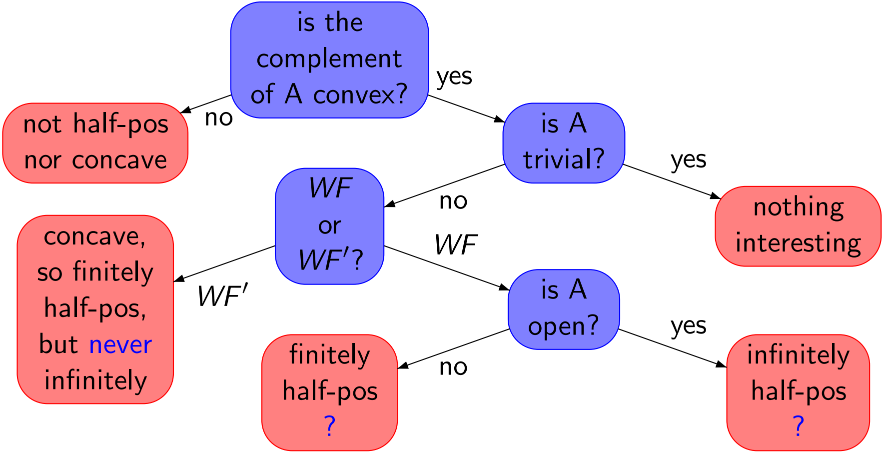

If A is an open half-space, then $WF(A)$ is half-positional.

Monotonic automata

Consider the languages: $C^* a^n C^*$, $C^* a^{n-1} b C^*$, $C^* b a^{n-1} C^*$ over $C = \{a, b, c\}$.

Monotonic automata

Consider the languages: $C^*a^nC^*$, $C^*a^{n-1}bC^*$, $C^*ba^{n-1}C^*$ over $C = \{a, b, c\}$.

They can be recognized by deterministic finite automata

Monotonic automata

Consider the languages: $C^*a^nC^*$, $C^*a^{n-1}bC^*$, $C^*ba^{n-1}C^*$ over $C = \{a, b, c\}$.

They can be recognized by deterministic finite automata satisfying the following special conditions:

- The set of states $Q = \{0, \ldots, n\}$;
- 0 is the initial state, n is the only accepting state;
- The transition function σ is monotonic, i.e. $q \geq q'$ implies $\sigma(q, c) \geq \sigma(q', c)$.

We call such an automaton a monotonic automaton $A = (n, \sigma)$ over C.

Monotonic automata

Let A be a monotonic automaton. We call the set $WM_A = C^\omega - L_A^\omega$ a monotonic condition.

Monotonic automata

Let A be a monotonic automaton. We call the set $WM_A = C^\omega - L_A^\omega$ a monotonic condition.

Theorem

Monotonic conditions are half-positional.

Monotonic automata

Let A be a monotonic automaton. We call the set $WM_A = C^\omega - L_A^\omega$ a monotonic condition.

Theorem

Monotonic conditions are half-positional.

For $L_A = C^* a^2 C^*$ the resulting WM_A is not concave: $(babab)^\omega$ is a combination of $(bbbaa)^\omega$ and $(aabbb)^\omega$. (However, monotonic conditions are weakly concave.)

Proof (an idea)

Let (G, WM_A) be a game with winning condition WM_A. We construct a new arena G' where $\mathrm{Pos}' = \mathrm{Pos} \times Q$. Eve automatically loses when a game reaches a position (v, q) where $q = n$. We can implement this as a parity condition, so we can find a positional strategy s' for Eve in some set $\mathrm{Win}'_E \subseteq \mathrm{Pos}'$. If Win'_E is empty, it ensures that Adam can win from any position. Otherwise, we project s' to $\mathrm{Win}_E = \{v : (v, 0) \in \mathrm{Win}'_E\}$ by $s(v) = \pi s'(v, q(v))$ where $q(v)$ is the greatest q such that (v, q) is in Win_E. This gives a positional strategy for Eve in some subset. We remove this subset from arena and repeat.

General properties of half-positional winning conditions

Basic tools

Let D be a determinacy type.

Theorem

Let $W \subseteq C^\omega$ be a winning condition such that for each nonempty D-arena G over C, there exists a position $v \in G$ such that in the game (G, W) one of the players has a D-strategy winning from v. Then W is D-determined.

Enhancing with Büchi conditions

Theorem

Let $W \subseteq C^\omega$ be a D-determined winning condition, and $S \subseteq C$. Then $W \cup WB_S$ is a D-determined winning condition as well.

Eve wins if she either wins W, or colors from S appear infinitely often.

Enhancing with Büchi conditions

Theorem

Let $W \subseteq C^\omega$ be a D-determined winning condition, and $S \subseteq C$. Then $W \cup WB_S$ is a D-determined winning condition as well.

Eve wins if she either wins W, or colors from S appear infinitely often.
Dually, $W \cap WB'_S$ (Eve wins if she wins W, and also colors from S appear only finitely often) is a D-determined winning condition as well.

Enhancing with Büchi conditions

Theorem

Let $W \subseteq C^\omega$ be a D-determined winning condition, and $S \subseteq C$. Then $W \cup WB_S$ is a D-determined winning condition as well.

Eve wins if she either wins W, or colors from S appear infinitely often.

Dually, $W \cap WB'_S$ (Eve wins if she wins W, and also colors from S appear only finitely often) is a D-determined winning condition as well.

By applying this theorem n times one can get positional determinacy of parity condition.

Closure under union?

If W is half-positional, then $W \cup WB_S$ is also half-positional.

Closure under union?

If W is half-positional, then $W \cup WB_S$ is also half-positional.
Is this a general fact: whenever W_1 and W_2 are half-positional, $W_1 \cup W_2$ is also half-positional?

Closure under union?

If W is half-positional, then $W \cup WB_S$ is also half-positional.
Is this a general fact: whenever W_1 and W_2 are half-positional, $W_1 \cup W_2$ is also half-positional?
Concave conditions are closed under union.
Monotonic conditions are closed under finite union.
A union of a concave and a monotonic condition is also finitely half-positional.

Uncountable union

A union of an uncountable family of half-positional conditions need not be half-positional — even for Büchi conditions.

Uncountable union

A union of an uncountable family of half-positional conditions need not be half-positional — even for Büchi conditions.

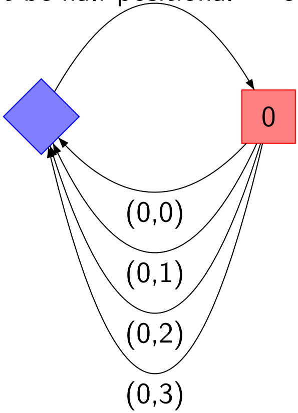

Uncountable union

A union of an uncountable family of half-positional conditions need not be half-positional — even for Büchi conditions.

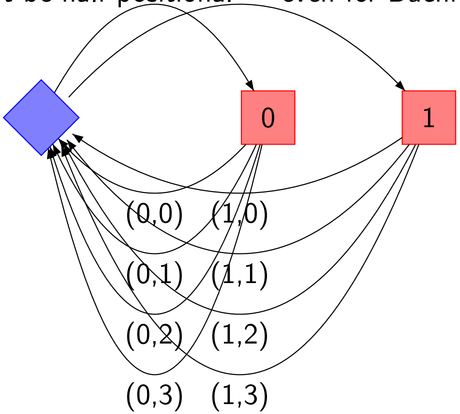

Uncountable union

A union of an uncountable family of half-positional conditions need not be half-positional — even for Büchi conditions.

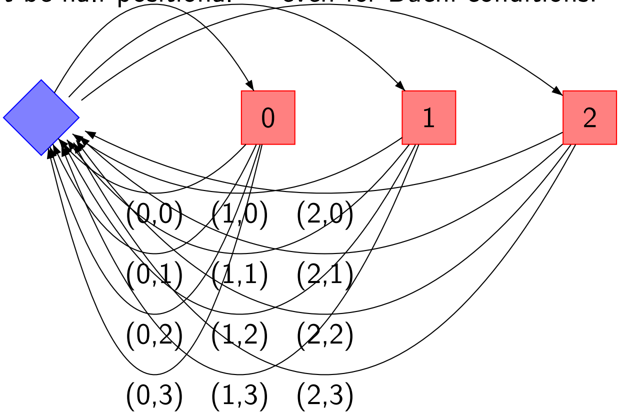

Uncountable union

A union of an uncountable family of half-positional conditions need not be half-positional — even for Büchi conditions.

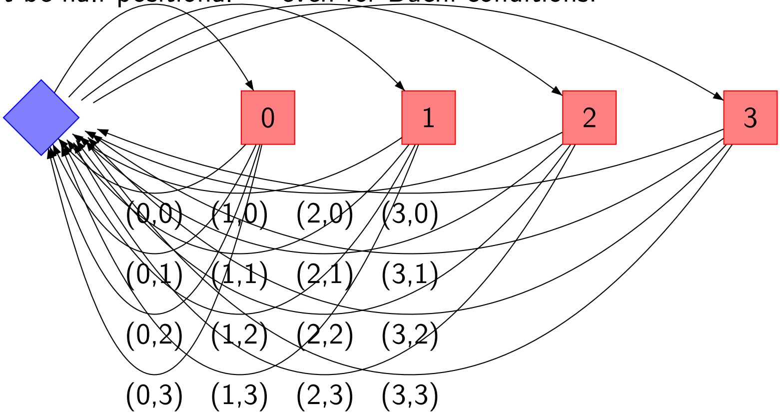

Uncountable union

A union of an uncountable family of half-positional conditions need not be half-positional — even for Büchi conditions.

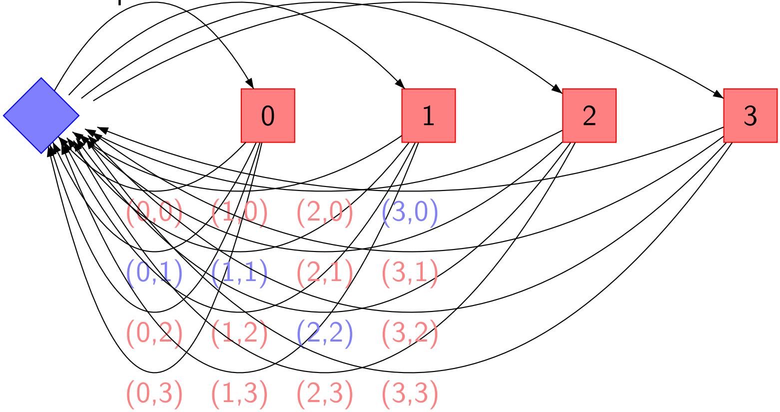

Closure under uncountable union

Arena: One Eve's position E and ω Adam's positions (A_n)
In E Eve chooses n and moves to A_n.
In A_n Adam chooses r and returns to E. This move is colored with (n, r).

For each $f : \omega \to \omega$, W_f is the Büchi condition given by $S_f = \{(n, f(n)) : n \in \omega\}$: Eve wins W_f if Adam uses moves colored with S_f infinitely many times.

Eve can win $\bigcup_{f:\omega\to\omega} W_f$, but only if she uses a non-positional strategy.

Positional/suspendable conditions

Definition

W is a positional/suspendable condition iff for each arena G Eve always has a positional strategy in her winning set, and Adam always has a suspendable strategy in his winning set.

Strategy for Adam is suspendable if from time to time Adam can stop using it (and do something else) and return later and still win (if he did not leave his winning set).

Positional/suspendable conditions: examples

The following winning conditions are positional/suspendable:

- Co-Büchi condition WB'_S.

Positional/suspendable conditions: examples

The following winning conditions are positional/suspendable:

- Co-Büchi condition WB'_S.
- The geometrical condition $WF(A)$, for an open half-space A.
- Monotonic conditions.

Positional/suspendable conditions: examples

The following winning conditions are positional/suspendable:

- Co-Büchi condition WB'_S.
- The geometrical condition $WF(A)$, for an open half-space A.
- Monotonic conditions.
- Countable unions of positional/suspendable conditions.

XPS conditions

Definition

The class of **extended positional/suspendable (XPS)** conditions over C is the smallest set of winning conditions that contains all Büchi and positional/ suspendable conditions, is closed under intersection with co-Büchi conditions, and is closed under finite union.

Theorem

XPS conditions are half-positional.

ω-regular Winning Conditions

A language $L \subseteq C^\omega$ is ω-regular iff it is accepted by a deterministic finite automaton with parity acceptance condition.

ω-regular Winning Conditions

A language $L \subseteq C^\omega$ is ω-regular iff it is accepted by a deterministic finite automaton with parity acceptance condition.

Definition

A DFA with parity acceptance condition is a tuple $A = (Q, q_I, \delta, \mathrm{rank})$

ω-regular Winning Conditions

A language $L \subseteq C^\omega$ is ω-regular iff it is accepted by a deterministic finite automaton with parity acceptance condition.

Definition

A DFA with parity acceptance condition is a tuple $A = (Q, q_I, \delta, \mathrm{rank})$ where:

- Q — set of states
- q_I — initial state
- $\delta : Q \times C \to Q$ — transition function

ω-regular Winning Conditions

A language $L \subseteq C^\omega$ is ω-regular iff it is accepted by a deterministic finite automaton with parity acceptance condition.

Definition

A DFA with parity acceptance condition is a tuple $A = (Q, q_I, \delta, \mathrm{rank})$ where:

- Q — set of states
- q_I — initial state
- $\delta : Q \times C \to Q$ — transition function
- $\mathrm{rank} : Q \to \{0 \ldots d\}$ — rank function

Let q_n be the state after reading first n letters. An infinite word is accepted iff $\limsup \mathrm{rank}\ q_n$ is even.

Simplifying the Witness Arena

Theorem

If W is ω-regular and not finitely half-positional then there is a witness arena (i.e. such that Eve has a winning strategy, but no positional winning strategy)

Simplifying the Witness Arena

Theorem

If W is ω-regular and not finitely half-positional then there is a witness arena (i.e. such that Eve has a winning strategy, but no positional winning strategy) where there is only one Eve's position, and only two moves from this position (no restriction on Adam's positions and moves).

Simplifying the Witness Arena part I

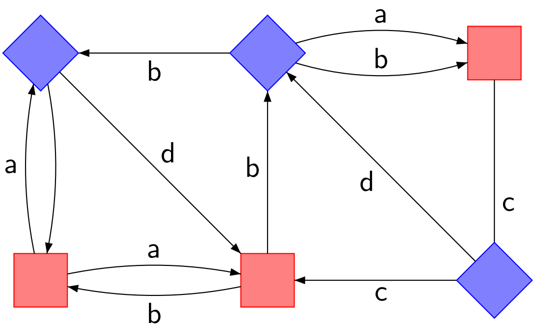

Simplifying the Witness Arena part I

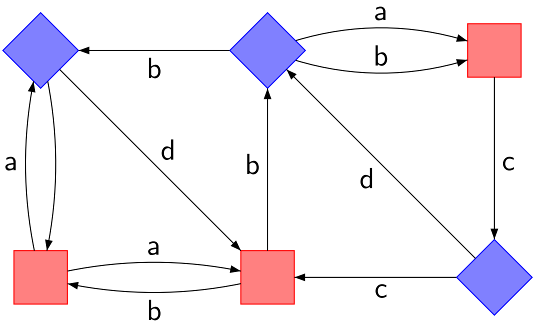

We can assume that for each Eve's position no strategy exists which always uses the same move in this position.

Simplifying the Witness Arena part I

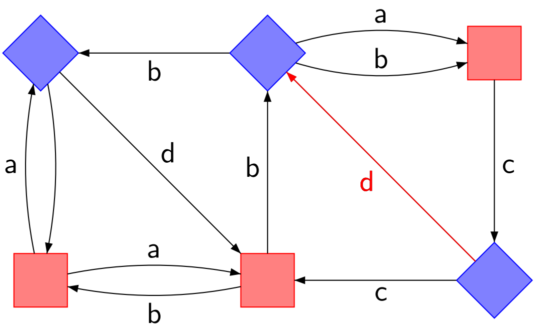

We can assume that for each Eve's position no strategy exists which always uses the same move in this position.

Simplifying the Witness Arena part I

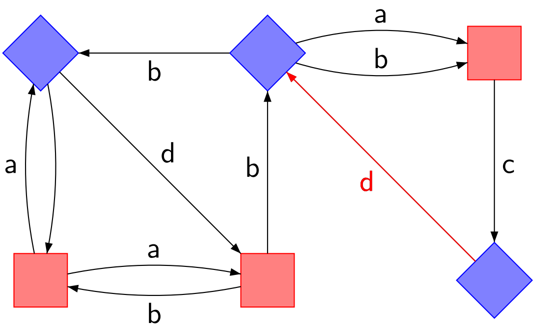

We can assume that for each Eve's position no strategy exists which always uses the same move in this position.

Simplifying the Witness Arena part I

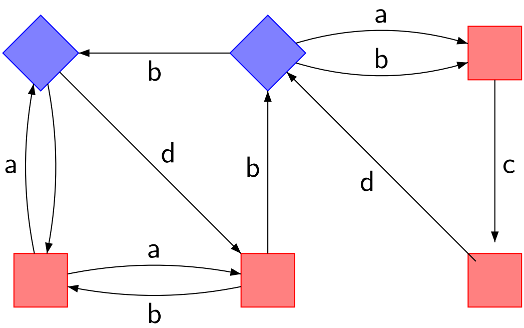

We can assume that for each Eve's position no strategy exists which always uses the same move in this position.

Simplifying the Witness Arena part II

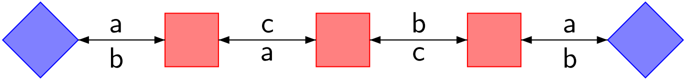

our witness arena

Simplifying the Witness Arena part II

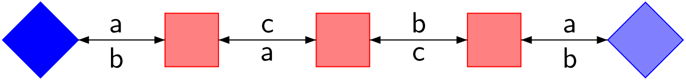

we choose one Eve's position

Simplifying the Witness Arena part II

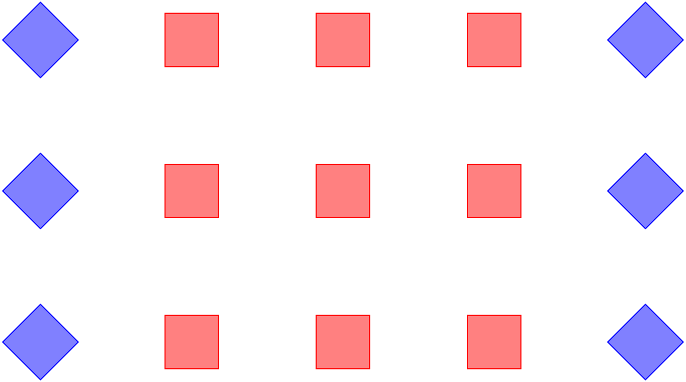

we create an equivalent game on $G \times Q$

Simplifying the Witness Arena part II

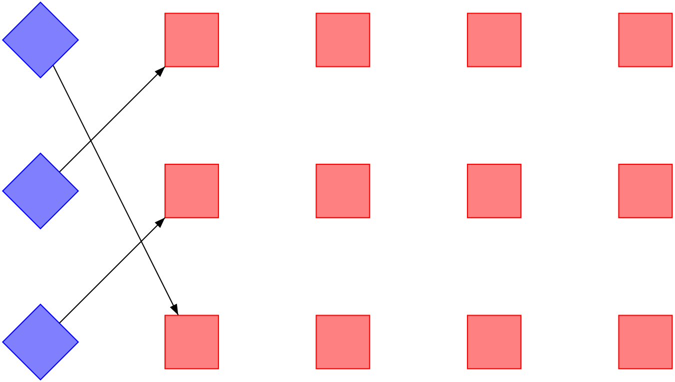

remove unnecessary moves

Simplifying the Witness Arena part II

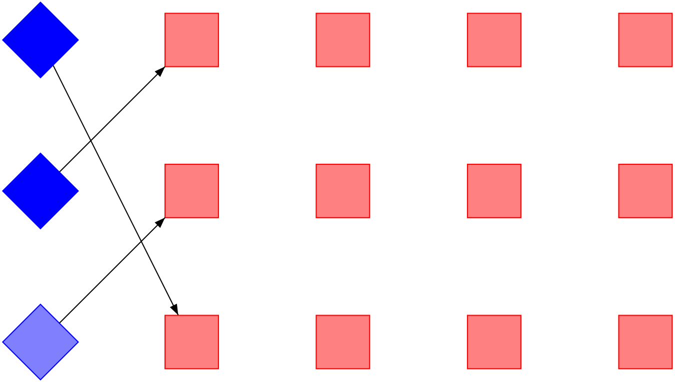

choose two Eve's positions. . .

Simplifying the Witness Arena part II

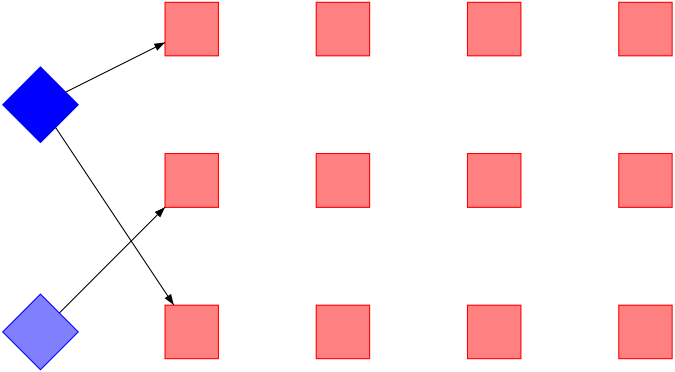

and merge them

Simplifying the Witness Arena part II

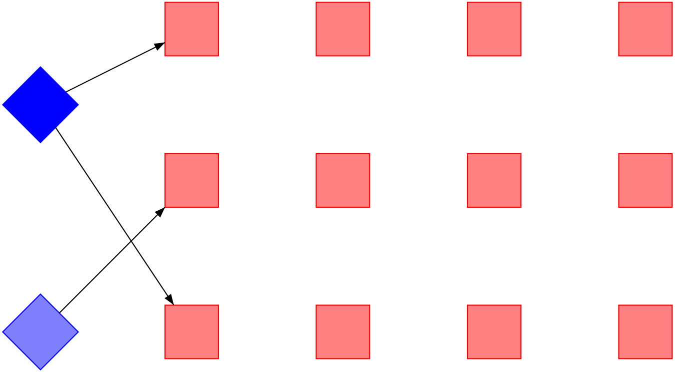

positional → remove the unused move and merge again

Decidability

Theorem

Let W be a (prefix independent) ω-regular winning condition recognized by a DFA with parity acceptance condition with n states. Then finite half-positional determinacy of W is decidable in time $O(n^{n^2})$.

Algorithm idea: Check all possible arenas with only one Eve's position and two Eve's moves.

Conclusion

Future work

Future work:

- more closure properties and examples of interesting half-positional winning conditions?
- Is an union of half-positional winning condition half-positional? Are there any half-positional winning conditions not in XPS?
- What about infinite half-positional determinacy of ω-regular languages? Do the classes of finitely half-positional and half-positional conditions coincide for ω-regular languages?
- Is the algorithm given optimal? Are there simple characterizations of (finitely) half-positional winning conditions?
- What about other geometrical conditions?
- In this work we allow arenas where some moves are colorless. Are there any winning conditions which are half-positional

Future work

How can our results be extended to:

- Finite memory strategies?
- Payoff mappings instead of win-lose?
- Winning conditions which are not prefix independent?
- Position-colored arenas (type A)?
- Stochastic games?

Conclusion

Summary

- infinite games — basic examples and definitions
- examples: concave, geometrical, and monotonic winning conditions
- closure properties
- suspendable strategies
- ω-regular winning conditions, decidability of finite determinacy
- future work

thank you