15-424: Foundations of Cyber-Physical Systems

Lecture Notes on Winning & Proving Hybrid Games

André Platzer

Carnegie Mellon University
Lecture 22

1 Introduction

This lecture continues the study of hybrid games and their logic, differential game logic [Pla13], that Lecture 20 on Hybrid Systems & Games and Lecture 21 on Winning Strategies & Regions started.

These lecture notes are based on [Pla13], where more information can be found on logic and hybrid games.

2 Deficiencies of the ω-Strategic Semantics

Lecture 21 on Winning Strategies & Regions raised the question whether the semantics of repetition could be defined by the *ω-strategic semantics*

$$\varsigma_{\alpha^*}(X) \stackrel{?}{=} \varsigma_\alpha^\omega(X) = \bigcup_{n<\omega} \varsigma_\alpha^n(X)$$

For winning condition $X \subseteq \mathcal{S}$ the iterated winning region of α is defined inductively:

$$\varsigma_\alpha^0(X) \stackrel{\text{def}}{=} X$$
$$\varsigma_\alpha^{\kappa+1}(X) \stackrel{\text{def}}{=} X \cup \varsigma_\alpha(\varsigma_\alpha^\kappa(X))$$

Does this give the right semantics for repetition of hybrid games? Does it match the existence of winning strategies that we were hoping to define?

Would the following dGL formula be valid in the ω-strategic semantics?

$$\langle (x := 1; x' = 1^d \cup x := x - 1)^* \rangle (0 \leq x < 1) \tag{1}$$

Before you read on, see if you can find the answer for yourself.

Abbreviate

$$\langle(\underbrace{\underbrace{x := 1; x' = 1^d}_{\beta} \cup \underbrace{x := x - 1}_{\gamma}}_{\alpha})^*\rangle\, (0 \le x < 1)$$

It is easy to see that $\varsigma^\omega_\alpha([0,1)) = [0,\infty)$, because $\varsigma^n_\alpha([0,1)) = [0,n)$ for all $n \in \mathbb{N}$ by a simple inductive proof (recall $\alpha \equiv \beta \cup \gamma$):

$$\begin{aligned}
&\varsigma^1_{\beta\cup\gamma}([0,1)) = [0,1)\\
&\varsigma^{n+1}_{\beta\cup\gamma}([0,1)) = [0,1) \cup \varsigma_{\beta\cup\gamma}(\varsigma^n_{\beta\cup\gamma}([0,1))) \overset{\text{IH}}{=} [0,1) \cup \varsigma_{\beta\cup\gamma}([0,n))\\
&= [0,1) \cup \varsigma_{\beta\cup\gamma}([0,n)) \cup \varsigma_\beta([0,n)) = [0,1) \cup \emptyset \cup [1,n+1) = [0,n+1)
\end{aligned}$$

Consequently,

$$\varsigma^\omega_\alpha([0,1)) = \bigcup_{n<\omega} \varsigma^n_\alpha([0,1)) = \bigcup_{n<\omega} [0,n) = [0,\infty)$$

Hence, the $\omega-$semantics would indicate that the hybrid game (1) can exactly be won from all initial states in $[0,\infty)$, that is, for all initial states that satisfy $0 \le x$.

Unfortunately, this is quite some nonsense. Indeed, the hybrid game in dGL formula (1) can be won from all initial states that satisfy $0 \le x$. But it can also be won from other initial states! So the ω-strategic semantics $\varsigma^\omega_\alpha([0,1))$ misses out on winning states. It is way too small for a winning region. There are cases, where the ω-semantics is minuscule compared to the true winning region and arbitrarily far away from the truth [Pla13].

In (1), this ω-level of iteration of the strategy function for winning regions misses out on Angel's perfectly reasonable winning strategy "first choose $x := 1; x' = 1^d$ and then always choose $x := x - 1$ until stopping at $0 \le x < 1$". This winning strategy wins from every initial state in $\mathbb{R}$, which is a much bigger set than $\varsigma^\omega_\alpha([0,1)) = [0,\infty)$.

Now this is the final answer for the winning region of (1). In particular, the dGL formula (1) is valid. Yet, is there a direct way to see that $\varsigma^\omega_\alpha([0,1)) = [0,\infty)$ is not the final answer for (1) without having to put the winning region computations aside and constructing a separate ingenious winning strategy?

Before you read on, see if you can find the answer for yourself.

The crucial observation is the following. The fact $\varsigma_\alpha^\omega([0,1)) = [0,\infty)$ shows that the hybrid game in (1) can be won from all nonnegative initial values with at most ω ("first countably infinitely many") steps. Let's recall how the proof worked, which showed $\varsigma_\alpha^n([0,1)) = [0,n)$ for all $n \in \mathbb{N}$. Its inductive step basically showed that if, for whatever reason (by inductive hypothesis really), $[0,n)$ is in the winning region, then $[0,n+1)$ also is in the winning region by simply applying $\varsigma_\alpha(\cdot)$ to $[0,n)$.

How about doing exactly that again? For whatever reason (i.e. by the above argument), $[0,\infty)$ is in the winning region. Doesn't that mean that $\varsigma_\alpha([0,\infty))$ should again be in the winning region by exactly the same inductive argument above?

Before you read on, see if you can find the answer for yourself.

Note 1. *Whenever a set Y is in the winning region $\varsigma_{\alpha^*}(X)$ of repetition, then $\varsigma_\alpha(Y)$ also should be in the winning region $\varsigma_{\alpha^*}(X)$, because it is just one step away from Y and α^* could simply repeat once more.*

Thus, the winning region $\varsigma_{(\beta\cup\gamma)^*}([0,\infty))$ should also contain

$$\varsigma_{\beta\cup\gamma}([0,\infty)) = \varsigma_\beta([0,\infty)) \cup \varsigma_\gamma([0,\infty)) = \mathbb{R} \cup [0,\infty) = \mathbb{R}$$

Beyond that, the winning region cannot contain anything else, because $\mathbb{R}$ is the whole state space. And, indeed, trying to use the winning region construction once more on $\mathbb{R}$ does not change the result:

$$\varsigma_{\beta\cup\gamma}(\mathbb{R}) = \varsigma_\beta(\mathbb{R}) \cup \varsigma_\gamma(\mathbb{R}) = \mathbb{R} \cup [0,\infty) = \mathbb{R}$$

This result, then coincides with what the ingenious winning strategy above told us as well: formula (1) is valid, because there is a winning strategy for Angel from every initial state. Except that the repeated $\varsigma_{\beta\cup\gamma}(\cdot)$ winning region construction seems more systematic than an ingenious guess of a smart winning strategy. So it gives a more constructive and explicit semantics.

Let's recap. In order to find the winning region of the hybrid game described in (1), it took us not just infinitely many steps, but more than that. After ω many iterations to arrive at $\varsigma_\alpha^\omega([0,1)) = [0,\infty)$, it took us one more step to arrive at

$$\varsigma_{(\beta\cup\gamma)^*}([0,1)) = \varsigma_\alpha^{\omega+1}([0,1)) = \mathbb{R}$$

where we denote the number of steps we took overall by $\omega+1$, since it was one more step than (first countable) infinitely many (i.e. ω many); see Fig. 1 for an illustration. More than infinitely many steps to get somewhere are plenty. Even worse: there are cases where even $\omega+1$ has not been enough of iteration to get to the repetition. The number of iterations needed to find $\varsigma_{\alpha^*}(X)$ could in general by much larger [Pla13].

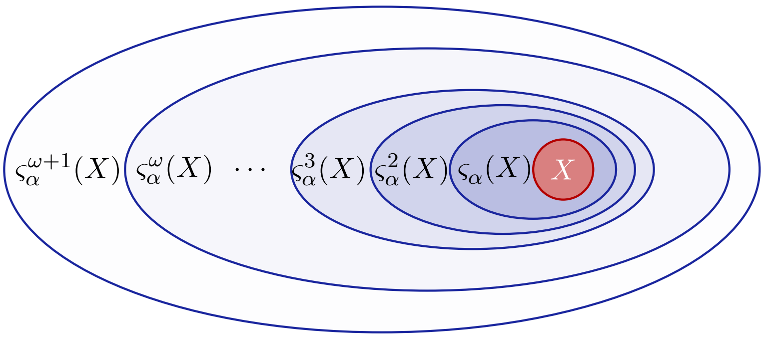

Figure 1: Iteration $\varsigma_\alpha^{\omega+1}(X)$ of $\varsigma_\alpha(\cdot)$ from winning condition $X = [0,1)$ stops when applying $\varsigma_\alpha(\cdot)$ to the ωth infinite iteration $\varsigma_\alpha^\omega(X)$.

The existence of the above winning strategy is only found at the level $\varsigma_\alpha^{\omega+1}([0,1)) = \varsigma_\alpha([0,\infty)) = \mathbb{R}$. Even though any particular use of the winning strategy in any game play uses only some finite number of repetitions of the loop, the argument why it will always work requires $> \omega$ many iterations of $\varsigma_\alpha(\cdot)$, because Demon can change x to an arbitrarily big value, so that ω many iterations of $\varsigma_\alpha(\cdot)$ are needed to conclude that Angel has a winning strategy for any positive value of x. There is no smaller upper bound on the number of iterations it takes Angel to win, in particular Angel cannot promise ω as a bound on the repetition count, which is what the ω-semantics would effectively require her to do. But strategies do converge after $\omega + 1$ iterations.

Note 2. *The ω-semantics is inappropriate, because it can be arbitrarily far away from characterizing the winning region of hybrid games.*

3 Characterizing Winning Repetitions

Is there a more immediate way of characterizing the winning region $\varsigma_{\alpha^*}(X)$ of repetition?

Whenever a set Y is in the winning region $\varsigma_{\alpha^*}(X)$ of repetition, then $\varsigma_\alpha(Y)$ also should be in the winning region $\varsigma_{\alpha^*}(X)$, because it is just one step away from Y and α^* could simply repeat once more. Thus,

$$Y \subseteq \varsigma_{\alpha^*}(X) \;\Rightarrow\; \varsigma_\alpha(Y) \subseteq \varsigma_{\alpha^*}(X)$$

In particular, the set $Y \stackrel{\text{def}}{=} \varsigma_{\alpha^*}(X)$ itself is expected to satisfy

$$\varsigma_\alpha(\varsigma_{\alpha^*}(X)) \subseteq \varsigma_{\alpha^*}(X) \tag{2}$$

because repeating α once more from the winning region $\varsigma_{\alpha^*}(X)$ of repetition of α should not give us any states that did not already have a winning strategy in α^*. Consequently, a set $Z \subseteq \mathcal{S}$ only qualifies as a candidate for being the winning region $\varsigma_{\alpha^*}(X)$ of repetition if

$$\varsigma_\alpha(Z) \subseteq Z \tag{3}$$

That is, strategyzing along α from Z does not give anything that Z would not already know about.

So what is this set Z? Is there only one choice? Or multiple? If there are multiple choices, which Z is it? Does such a Z always exist, even?

Before you read on, see if you can find the answer for yourself.

One such Z always exist, even though it may be rather boring. The empty set $Z \stackrel{\text{def}}{=} \emptyset$ certainly satisfies $\varsigma_\alpha(\emptyset) = \emptyset$, because it is rather hard to win a game that requires Angel to enter the empty set of states to win.

But the empty set is maybe a bit small. The winning region $\varsigma_{\alpha^*}(X)$ of repetition of α should at least contain the winning condition X, because the winning condition X is particularly easy to reach from states in X that have already let Angel won by simply suggesting Angel to repeat zero times. Consequently, the only Z that qualify as a candidate for being $\varsigma_{\alpha^*}(X)$ should satisfy (3) and

$$X \subseteq Z \tag{4}$$

Both conditions (3) and (4) together can be summarized in a single condition as follows:

Note 3 (Prefixpoint)**.** *Every candidate Z for the winning region $\varsigma_{\alpha^*}(X)$ satisfies:*

$$X \cup \varsigma_\alpha(Z) \subseteq Z \tag{5}$$

Again: what is this set Z that satisfies (5)? Is there only one choice? Or multiple? If there are multiple choices, which Z is it? Does such a Z always exist, even?

Before you read on, see if you can find the answer for yourself.

One such Z certainly exists. The empty set does not qualify unless $X = \emptyset$. The set X itself is too small unless the game has no incentive to start repeating, because $\varsigma_\alpha(X) \subseteq X$. But the full space $Z \stackrel{\text{def}}{=} \mathcal{S}$ always satisfies (5) trivially. Now, the whole space is a little big to call it Angel's winning region independently of the hybrid game α. Even if the full space may very well be the winning region for some particularly Demonophobic Angel-friendly hybrid games like (1), it is hardly the right winning region for any arbitrary α^*. For example for Demon's favorite game where he always wins, $\varsigma_{\alpha^*}(X)$ had better be $\emptyset$, not $\mathcal{S}$. Thus, the largest solution Z of (5) hardly qualifies.

So which solution Z of (5) should be the definition of $\varsigma_{\alpha^*}(X)$ now?

Before you read on, see if you can find the answer for yourself.

Among the many Z that solve (5), the largest one is not informative, because the largest Z simply degrades to $\mathcal{S}$. So smaller solutions Z are preferable. How do multiple solutions relate at all? Suppose Y, Z are both solutions of (5). That is

$$X \cup \varsigma_\alpha(Y) \subseteq Y \tag{6}$$
$$X \cup \varsigma_\alpha(Z) \subseteq Z \tag{7}$$

Then, by monotonicity lemma, Lemma 3:

$$X \cup \varsigma_\alpha(Y \cap Z) \overset{\text{mon}}{\subseteq} X \cup \varsigma_\alpha(Y) \cap \varsigma_\alpha(Z) \overset{(6),(7)}{\subseteq} Y \cap Z \tag{8}$$

Hence, by (8), the intersection $Y \cap Z$ of solutions Y and Z of (5) also is a solution of (5):

Lemma 1 (Intersection closure)**.** *Whenever there are two solutions Y, Z of* (5)*, a (possibly) smaller solution of* (5) *can be obtained by intersection $Y \cap Z$.*

So whenever there are two solutions Z_1, Z_2 of (5), their intersection $Y_1 \cap Z_2$ solves (5). When there's yet another solution Z_3 of (5), their intersection $Y_1 \cap Y_2 \cap Y_3$ also solves (5). Similarly for *any* larger family of solutions. If we keep on intersecting solutions, we will arrive at smaller solutions until, some fine day, there's not going to be a smaller one. This yields the smallest solution Z of (5) which can be characterized directly.

Among the many Z that solve (5), the smallest Z that solves (5) is informative and can be used to define $\varsigma_{\alpha^*}(X)$:

$$\varsigma_{\alpha^*}(X) = \bigcap \{Z \subseteq \mathcal{S} \,:\, X \cup \varsigma_\alpha(Z) \subseteq Z\} \tag{9}$$

The set on the right-hand side of (9) is an intersection of solutions, thus, a solution by Lemma 1 (or its counterpart for families of solutions). Hence $\varsigma_{\alpha^*}(X)$ itself satisfies (5):

$$X \cup \varsigma_\alpha(\varsigma_{\alpha^*}(X)) \subseteq \varsigma_{\alpha^*}(X) \tag{10}$$

Also compare this with what we argued earlier in (2). Could it be the case that the inclusion in (10) is strict, i.e. not equals? No this cannot happen, because $\varsigma_{\alpha^*}(X)$ is the smallest. In detail, by (10), the set $Z \overset{\text{def}}{=} X \cup \varsigma_\alpha(\varsigma_{\alpha^*}(X))$ satisfies $Z \subseteq \varsigma_{\alpha^*}(X)$ and, thus, by Lemma 3:

$$X \cup \varsigma_\alpha(Z) \overset{\text{mon}}{\subseteq} X \cup \varsigma_\alpha(\varsigma_{\alpha^*}(X)) = Z$$

Consequently, both inclusions hold, so $\varsigma_{\alpha^*}(X)$ satisfies

$$X \cup \varsigma_\alpha(\varsigma_{\alpha^*}(X)) = \varsigma_{\alpha^*}(X) \tag{11}$$

That is, $\varsigma_{\alpha^*}(X)$ is even a *fixpoint* solving the equation

$$X \cup \varsigma_\alpha(Z) = Z \tag{12}$$

and it is the *least fixpoint*, i.e. the smallest Z solving the equation (12).

The fact that $\varsigma_{\alpha^*}(X)$ is defined as the least of the fixpoints makes sure that Angel only wins games by a well-founded number of repetitions. That is, she only wins a repetition if she ultimately stops repeating, not by postponing termination forever. See [Pla13] for more details.

It is also worth noting that it would still have been possible to make the iteration of winning region constructions work out using the seminal fixpoint theorem of Knaster-Tarski. Yet, this requires the iterated winning region constructions to go significantly transfinite [Pla13] way beyond ω.

4 Semantics of Hybrid Games

The semantics of differential game logic from Lecture 21 was still pending a definition of the winning regions $\varsigma_\alpha(\cdot)$ and $\delta_\alpha(\cdot)$ for Angel and Demon, respectively, in the hybrid game α. Rather than taking a detour for understanding those by operational game semantics (as in Lecture 20), the winning regions of hybrid games can be defined directly, giving a denotational semantics to hybrid games.

The only difference compared to the definition in Lecture 21 is the new case of repetition α^*.

Definition 2 (Semantics of hybrid games)**.** The *semantics of a hybrid game* α is a function $\varsigma_\alpha(\cdot)$ that, for each interpretation I and each set of Angel's winning states $X \subseteq \mathcal{S}$, gives the *winning region*, i.e. the set of states $\varsigma_\alpha(X)$ from which Angel has a winning strategy to achieve X (whatever strategy Demon chooses). It is defined inductively as follows[a]

1. $\varsigma_{x:=\theta}(X) = \{\nu \in \mathcal{S} \;:\; \nu_x^{[\![\theta]\!]_\nu} \in X\}$
2. $\varsigma_{x'=\theta \,\&\, H}(X) = \{\varphi(0) \in \mathcal{S} \;:\; \varphi(r) \in X$ for some $r \in \mathbb{R}_{\geq 0}$ and (differentiable) $\varphi : [0, r] \to \mathcal{S}$ such that $\varphi(\zeta) \in [\![H]\!]^I$ and $\frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ for all $0 \leq \zeta \leq r\}$
3. $\varsigma_{?H}(X) = [\![H]\!]^I \cap X$
4. $\varsigma_{\alpha\cup\beta}(X) = \varsigma_\alpha(X) \cup \varsigma_\beta(X)$
5. $\varsigma_{\alpha;\beta}(X) = \varsigma_\alpha(\varsigma_\beta(X))$
6. $\varsigma_{\alpha^*}(X) = \bigcap\{Z \subseteq \mathcal{S} \;:\; X \cup \varsigma_\alpha(Z) \subseteq Z\}$
7. $\varsigma_{\alpha^d}(X) = (\varsigma_\alpha(X^\complement))^\complement$

The *winning region* of Demon, i.e. the set of states $\delta_\alpha(X)$ from which Demon has a winning strategy to achieve X (whatever strategy Angel chooses) is defined inductively as follows

1. $\delta_{x:=\theta}(X) = \{\nu \in \mathcal{S} \;:\; \nu_x^{[\![\theta]\!]_\nu} \in X\}$
2. $\delta_{x'=\theta \,\&\, H}(X) = \{\varphi(0) \in \mathcal{S} \;:\; \varphi(r) \in X$ for all $r \in \mathbb{R}_{\geq 0}$ and (differentiable) $\varphi : [0, r] \to \mathcal{S}$ such that $\varphi(\zeta) \in [\![H]\!]^I$ and $\frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ for all $0 \leq \zeta \leq r\}$
3. $\delta_{?H}(X) = ([\![H]\!]^I)^\complement \cup X$
4. $\delta_{\alpha\cup\beta}(X) = \delta_\alpha(X) \cap \delta_\beta(X)$
5. $\delta_{\alpha;\beta}(X) = \delta_\alpha(\delta_\beta(X))$
6. $\delta_{\alpha^*}(X) = \bigcup\{Z \subseteq \mathcal{S} \;:\; Z \subseteq X \cap \delta_\alpha(Z)\}$
7. $\delta_{\alpha^d}(X) = (\delta_\alpha(X^\complement))^\complement$

[a] The semantics of a hybrid game is not merely a reachability relation between states as for hybrid systems [Pla12], because the adversarial dynamic interactions and nested choices of the players have to be taken into account.

This notation uses $\varsigma_\alpha(X)$ instead of $\varsigma_\alpha^I(X)$ and $\delta_\alpha(X)$ instead of $\delta_\alpha^I(X)$, because the interpretation I that gives a semantics to predicate symbols in tests and evolution domains is clear from the context. Strategies do not occur explicitly in the d$\mathcal{GL}$ semantics, because it is based on the existence of winning strategies, not on the strategies themselves.

Just as the semantics $\mathsf{d}\mathcal{L}$, the semantics of $\mathsf{dG}\mathcal{L}$ is *compositional*, i.e. the semantics of a compound $\mathsf{dG}\mathcal{L}$ formula is a simple function of the semantics of its pieces, and the semantics of a compound hybrid game is a function of the semantics of its pieces. Furthermore, existence of a strategy in hybrid game α to achieve X is independent of any game and $\mathsf{dG}\mathcal{L}$ formula surrounding α, but just depends on the remaining game α itself and the goal X. By a simple inductive argument, this shows that one can focus on memoryless strategies, because the existence of strategies does not depend on the context, hence, by working bottom up, the strategy itself cannot depend on past states and choices, only the current state, remaining game, and goal. This also follows from a generalization of a classical result by Zermelo. Furthermore, the semantics is monotone, i.e. larger sets of winning states induce larger winning regions.

Monotonicity is what Lecture 21 looked into for the case of hybrid games without repetition. But it continues to hold for general hybrid games.

Lemma 3 (Monotonicity [Pla13])**.** *The semantics is* monotone, *i.e.* $\varsigma_\alpha(X) \subseteq \varsigma_\alpha(Y)$ *and* $\delta_\alpha(X) \subseteq \delta_\alpha(Y)$ *for all* $X \subseteq Y$*.*

Proof. A simple check based on the observation that X only occurs with an even number of negations in the semantics. For example, $\varsigma_{\alpha^*}(X) = \bigcap\{Z \subseteq \mathcal{S} : X \cup \varsigma_\alpha(Z) \subseteq Z\} \subseteq \bigcap\{Z \subseteq \mathcal{S} : Y \cup \varsigma_\alpha(Z) \subseteq Z\} = \varsigma_{\alpha^*}(Y)$ if $X \subseteq Y$. Likewise, $X \subseteq Y$ implies $X^\complement \supseteq Y^\complement$, hence $\varsigma_\alpha(X^\complement) \supseteq \varsigma_\alpha(Y^\complement)$, so $\varsigma_{\alpha^d}(X) = (\varsigma_\alpha(X^\complement))^\complement \subseteq (\varsigma_\alpha(Y^\complement))^\complement = \varsigma_{\alpha^d}(Y)$. □

Monotonicity implies that the least fixpoint in $\varsigma_{\alpha^*}(X)$ and the greatest fixpoint in $\delta_{\alpha^*}(X)$ are well-defined [HKT00, Lemma 1.7]. The semantics of $\varsigma_{\alpha^*}(X)$ is a least fixpoint, which results in a well-founded repetition of α, i.e. Angel can repeat any number of times but she ultimately needs to stop at a state in X in order to win. The semantics of $\delta_{\alpha^*}(X)$ is a greatest fixpoint, instead, for which Demon needs to achieve a state in X after every number of repetitions, because Angel could choose to stop at any time, but Demon still wins if he only postpones $X^\complement$ forever, because Angel ultimately has to stop repeating. Thus, for the formula $\langle\alpha^*\rangle\phi$, Demon already has a winning strategy if he only has a strategy that is not losing by preventing ϕ indefinitely, because Angel eventually has to stop repeating anyhow and will then end up in a state not satisfying ϕ, which makes her lose. The situation for $[\alpha^*]\phi$ is dual.

5 Hybrid Game Axioms

An axiomatization for differential game logic has been found in previous work [Pla13], where we refer to for more details. The study of proof rules for differential game logic will be deferred to next lecture. But its axioms can be discussed today.

Note 7 (Differential game logic axioms [Pla13]).

$$([\cdot])\ [\alpha]\phi \leftrightarrow \neg\langle\alpha\rangle\neg\phi$$

$$(\langle{:=}\rangle)\ \langle x := \theta\rangle\phi(x) \leftrightarrow \phi(\theta)$$

$$(\langle{'}\rangle)\ \langle x' = \theta\rangle\phi \leftrightarrow \exists t{\geq}0\, \langle x := y(t)\rangle\phi \qquad (y'(t) = \theta)$$

$$(\langle ? \rangle)\ \langle ?H\rangle\phi \leftrightarrow (H \wedge \phi)$$

$$(\langle\cup\rangle)\ \langle\alpha \cup \beta\rangle\phi \leftrightarrow \langle\alpha\rangle\phi \vee \langle\beta\rangle\phi$$

$$(\langle ; \rangle)\ \langle\alpha;\beta\rangle\phi \leftrightarrow \langle\alpha\rangle\langle\beta\rangle\phi$$

$$(\langle ^* \rangle)\ \phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi \rightarrow \langle\alpha^*\rangle\phi$$

$$(\langle ^d \rangle)\ \langle\alpha^d\rangle\phi \leftrightarrow \neg\langle\alpha\rangle\neg\phi$$

6 Determinacy

Every particular game play in a hybrid game is won by exactly one player, because hybrid games are zero-sum and there are no draws. Hybrid games actually satisfy a much stronger property: *determinacy*, i.e. that, from any initial situation, either one of the players always has a winning strategy to force a win, regardless of how the other player chooses to play.

If, from the same initial state, both Angel and Demon had a winning strategy for opposing winning conditions, then something would be terribly inconsistent. It cannot happen that Angel has a winning strategy in hybrid game α to get to a state where $\neg\phi$ and, from the same initial state, Demon supposedly also has a winning strategy in the same hybrid game α to get to a state where ϕ holds. After all, a winning strategy is a strategy that makes that player win no matter what strategy the opponent follows. Hence, for any initial state, at most one player can have a winning strategy for complementary winning conditions. This argues for the validity of $\vDash \neg([\alpha]\phi \wedge \langle\alpha\rangle\neg\phi)$, which can also be proved (Theorem 4).

So it cannot happen that both players have a winning strategy for complementary winning conditions. But it might still happen that no one has a winning strategy, i.e. both players can let the other player win, but cannot win strategically themselves (recall, e.g., the filibuster example from Lecture 20, which first appeared as if no player might have a winning strategy but then turned out to make Demon win). This does not happen for hybrid games, though, because at least one (hence exactly one) player has a winning strategy for complementary winning conditions from any initial state.

Theorem 4 (Consistency & determinacy [Pla13]). *Hybrid games are consistent and determined, i.e.* $\vDash \neg\langle\alpha\rangle\neg\phi \leftrightarrow [\alpha]\phi$.

Proof. The proof shows by induction on the structure of α that $\varsigma_\alpha(X^\complement)^\complement = \delta_\alpha(X)$ for all $X \subseteq \mathcal{S}$ and all I with some set of states $\mathcal{S}$, which implies the validity of $\neg\langle\alpha\rangle\neg\phi \leftrightarrow [\alpha]\phi$ using $X \stackrel{\text{def}}{=} [\![\phi]\!]^I$.

1. $\varsigma_{x:=\theta}(X^\complement)^\complement = \{\nu \in \mathcal{S} : \nu_x^{[\![\theta]\!]_\nu} \notin X\}^\complement = \varsigma_{x:=\theta}(X) = \delta_{x:=\theta}(X)$
2. $\varsigma_{x'=\theta\,\&\,H}(X^\complement)^\complement = \{\varphi(0) \in \mathcal{S} : \varphi(r) \notin X$ for some $0 \leq r \in \mathbb{R}$ and some (differentiable) $\varphi : [0,r] \to \mathcal{S}$ such that $\frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ and $\varphi(\zeta) \in [\![H]\!]^I$ for all $0 \leq \zeta \leq r\}^\complement = \delta_{x'=\theta\,\&\,H}(X)$, because the set of states from which there is no winning strategy for Angel to reach a state in $X^\complement$ prior to leaving $[\![H]\!]^I$ along $x' = \theta\,\&\,H$ is exactly the set of states from which $x' = \theta\,\&\,H$ always stays in X (until leaving $[\![H]\!]^I$ in case that ever happens).
3. $\varsigma_{?H}(X^\complement)^\complement = ([\![H]\!]^I \cap X^\complement)^\complement = ([\![H]\!]^I)^\complement \cup (X^\complement)^\complement = \delta_{?H}(X)$
4. $\varsigma_{\alpha\cup\beta}(X^\complement)^\complement = (\varsigma_\alpha(X^\complement) \cup \varsigma_\beta(X^\complement))^\complement = \varsigma_\alpha(X^\complement)^\complement \cap \varsigma_\beta(X^\complement)^\complement = \delta_\alpha(X) \cap \delta_\beta(X) = \delta_{\alpha\cup\beta}(X)$
5. $\varsigma_{\alpha;\beta}(X^\complement)^\complement = \varsigma_\alpha(\varsigma_\beta(X^\complement))^\complement = \varsigma_\alpha(\delta_\beta(X)^\complement)^\complement = \delta_\alpha(\delta_\beta(X)) = \delta_{\alpha;\beta}(X)$
6. $\varsigma_{\alpha^*}(X^\complement)^\complement = \left(\bigcap\{Z \subseteq \mathcal{S} : X^\complement \cup \varsigma_\alpha(Z) \subseteq Z\}\right)^\complement = \left(\bigcap\{Z \subseteq \mathcal{S} : (X \cap \varsigma_\alpha(Z)^\complement)^\complement \subseteq Z\}\right)^\complement$ $= \left(\bigcap\{Z \subseteq \mathcal{S} : (X \cap \delta_\alpha(Z^\complement))^\complement \subseteq Z\}\right)^\complement = \bigcup\{Z \subseteq \mathcal{S} : Z \subseteq X \cap \delta_\alpha(Z)\} = \delta_{\alpha^*}(X).$ [1]
7. $\varsigma_{\alpha^d}(X^\complement)^\complement = (\varsigma_\alpha((X^\complement)^\complement)^\complement)^\complement = \delta_\alpha(X^\complement)^\complement = \delta_{\alpha^d}(X)$ $\square$

Exercises

Exercise 1. Explain how often you will have to repeat the winning region construction to show that the following d$\mathcal{GL}$ formula is valid:

$$\langle(x := x + 1; x' = 1^d \cup x := x - 1)^*\rangle\,(0 \leq x < 1)$$

Exercise 2. Can you find d$\mathcal{GL}$ formulas for which the winning region construction takes even longer to terminate? How far can you push this?

Exercise 3. Carefully identify how determinacy relates to the two possible understandings of the filibuster example discussed in an earlier lecture.

[1]The penultimate equation follows from the μ-calculus equivalence $\nu Z.\Upsilon(Z) \equiv \neg\mu Z.\neg\Upsilon(\neg Z)$ and the fact that least pre-fixpoints are fixpoints and that greatest post-fixpoints are fixpoints for monotone functions.

Exercise 4. Prove the elided cases of Lemma 3.

Exercise 5. Find the appropriate soundness notion for the axioms of $\mathsf{dG}\mathcal{L}$ and prove that the axioms are sound.

Exercise 6. Write down a valid formula that characterizes an interesting game between two robots.

References

[HKT00] David Harel, Dexter Kozen, and Jerzy Tiuryn. *Dynamic logic*. MIT Press, 2000.

[Pla12] André Platzer. The complete proof theory of hybrid systems. In *LICS*, pages 541–550. IEEE, 2012. doi:10.1109/LICS.2012.64.

[Pla13] André Platzer. A complete axiomatization of differential game logic for hybrid games. Technical Report CMU-CS-13-100R, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, January, Revised and extended in July 2013.