15-424: Foundations of Cyber-Physical Systems

Lecture Notes on Winning Strategies & Regions

André Platzer

Carnegie Mellon University

Lecture 21

1 Introduction

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

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

2 Semantics

What is the most elegant way of defining a semantics for differential game logic? How could a semantics be defined at all? First of all, the dG$\mathcal{L}$ formulas ϕ that are used in the postconditions of dG$\mathcal{L}$ modal formulas $\langle\alpha\rangle\phi$ and $[\alpha]\phi$ define the winning conditions for the hybrid game α. Thus, when playing the hybrid game α, we need to know the set of states in which the winning condition ϕ is satisfied. That set of states in which ϕ is true is denoted $[\![\phi]\!]^I$, which defines the semantics of ϕ.

The logic dG$\mathcal{L}$ has a denotational semantics. The dG$\mathcal{L}$ semantics defines, for each formula ϕ, the set $[\![\phi]\!]^I$ of states in which ϕ is true. For each hybrid game α and each set of winning states X, the dG$\mathcal{L}$ semantics defines the set $\varsigma_\alpha(X)$ of states from which Angel has a winning strategy to achieve X in hybrid game α, as well as the set $\delta_\alpha(X)$ of states from which Demon has a winning strategy to achieve X in α.

A *state* ν is a mapping from variables to $\mathbb{R}$. An *interpretation* I assigns a relation $I(p) \subseteq \mathbb{R}^k$ to each predicate symbol p of arity k. The interpretation further determines the set of states $\mathcal{S}$, which is isomorphic to a Euclidean space $\mathbb{R}^n$ when n is the number of relevant variables. For a subset $X \subseteq \mathcal{S}$ the complement $\mathcal{S} \setminus X$ is denoted $X^\complement$. Let ν_x^d denote the state that agrees with state ν except for the interpretation of variable x, which is changed to $d \in \mathbb{R}$. The value of term θ in state ν is denoted by $[\![\theta]\!]_\nu$. The denotational semantics of dG$\mathcal{L}$ formulas will be defined in Def. 1 by simultaneous induction along with the denotational semantics, $\varsigma_\alpha(\cdot)$ and $\delta_\alpha(\cdot)$, of hybrid games, defined later in Def. 2, because dG$\mathcal{L}$ formulas are defined by simultaneous induction with hybrid games. The *(denotational) semantics of a hybrid game* α defines for each interpretation I and each set of Angel's winning states $X \subseteq \mathcal{S}$ 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). 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 subsequently in Def. 2 as well.

Definition 1 (dG$\mathcal{L}$ semantics)**.** The *semantics of a* dG$\mathcal{L}$ *formula* ϕ for each interpretation I with a corresponding set of states $\mathcal{S}$ is the subset $[\![\phi]\!]^I \subseteq \mathcal{S}$ of states in which ϕ is true. It is defined inductively as follows

1. $[\![p(\theta_1, \dots, \theta_k)]\!]^I = \{\nu \in \mathcal{S} : ([\![\theta_1]\!]_\nu, \dots, [\![\theta_k]\!]_\nu) \in I(p)\}$
2. $[\![\theta_1 \geq \theta_2]\!]^I = \{\nu \in \mathcal{S} : [\![\theta_1]\!]_\nu \geq [\![\theta_2]\!]_\nu\}$
3. $[\![\neg\phi]\!]^I = ([\![\phi]\!]^I)^\complement$
4. $[\![\phi \wedge \psi]\!]^I = [\![\phi]\!]^I \cap [\![\psi]\!]^I$
5. $[\![\exists x\, \phi]\!]^I = \{\nu \in \mathcal{S} : \nu_x^r \in [\![\phi]\!]^I \text{ for some } r \in \mathbb{R}\}$
6. $[\![\langle\alpha\rangle\phi]\!]^I = \varsigma_\alpha([\![\phi]\!]^I)$
7. $[\![[\alpha]\phi]\!]^I = \delta_\alpha([\![\phi]\!]^I)$

A dG$\mathcal{L}$ formula ϕ is *valid in* I, written $I \models \phi$, iff $[\![\phi]\!]^I = \mathcal{S}$. Formula ϕ is *valid*, $\vDash \phi$, iff $I \models \phi$ for all interpretations I.

Note that the semantics of $\langle\alpha\rangle\phi$ cannot be defined as it would in d$\mathcal{L}$ via

$$[\![\langle\alpha\rangle\phi]\!]^I = \{\nu \in \mathcal{S} : \omega \in [\![\phi]\!]^I \text{ for some } \omega \text{ with } (\nu, \omega) \in \rho(\alpha)$$

First of all, the reachability relation $(\nu, \omega) \in \rho(\alpha)$ is only defined when α is a hybrid program, not when it is a hybrid game. But the deeper reason is that the above shape is too harsh. Criteria of this shape would require Angel to single out a single state ν that satisfies the winning condition $\omega \in [\![\phi]\!]^I$ and then get to that state ω by playing α from ν. Yet all that Demon then has to do to spoil that plan is lead the play into a different state (e.g., one in which Angel would also have won) but which is different from the projected ω. More generally, winning into a single state is really difficult. Winning by leading the play into one of several states that satisfy the winning condition is more feasible. This is what the winning region $\varsigma_\alpha([\![)]\!]^{\phi\phi}$ is supposed to capture.

3 Winning Regions

Def. 1 needs 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.

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 $\frac{\mathrm{d}\,\varphi(t)(x)}{\mathrm{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ 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. $\varsigma_{\alpha^d}(X) = (\varsigma_\alpha(X^\complement))^\complement$

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 $\frac{\mathrm{d}\,\varphi(t)(x)}{\mathrm{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ 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. $\delta_{\alpha^d}(X) = (\delta_\alpha(X^\complement))^\complement$

[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 $\varsigma^I_\alpha(X)$ and $\delta_\alpha(X)$ instead of $\delta^I_\alpha(X)$, 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 dG$\mathcal{L}$ semantics, because it is based on the existence of winning strategies, not on the strategies themselves.

Just as the semantics d$\mathcal{L}$, the semantics of dG$\mathcal{L}$ is *compositional*, i.e. the semantics of a compound 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 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.

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, $X \subseteq Y$ implies $X^\complement \supseteq Y^\complement$, hence $\varsigma_\alpha(X^\complement) \supseteq \varsigma_\alpha(Y^\complement)$, so $\varsigma_{\alpha^d}(X) = (\varsigma_\alpha(X^\complement))^\complement \subseteq (\varsigma_\alpha(Y^\complement))^\complement = \varsigma_{\alpha^d}(Y)$. □

Before going any further, however, we need to define a semantics for repetition, which will turn out to be surprisingly difficult.

4 Examples

Consider the following examples and find out whether the formulas are valid or not.

$$\langle (x := x+1; (x' = x^2)^d \cup x := x-1)^* \rangle\, (0 \le x < 1)$$

$$\langle (x := x+1; (x' = x^2)^d \cup (x := x-1 \cap x := x-2))^* \rangle (0 \le x < 1)$$

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

$$\vDash \langle (x := x+1; (x' = x^2)^d \cup x := x-1)^* \rangle\, (0 \le x < 1)$$

$$\nvDash \langle (x := x+1; (x' = x^2)^d \cup (x := x-1 \cap x := x-2))^* \rangle (0 \le x < 1)$$

5 Advance Notice Repetitions

The semantics of repetition in hybrid systems was

$$\rho(\alpha^*) = \bigcup_{n \in \mathbb{N}} \rho(\alpha^n)$$

with $\alpha^{n+1} \equiv \alpha^n; \alpha$ and $\alpha^0 \equiv ?true$.

So the obvious candidate for the semantics of repetition in hybrid games might be

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

where ω is the first infinite ordinal (if you have never seen ordinals before, just read $n < \omega$ as natural numbers $n \in \mathbb{N}$). Would that give the intended meaning to repetition? Is Angel forced to stop in order to win if the game of repetition would be played this way? Yes, she would, because, even though there is no bound on the number of repetitions that she can choose, for each natural number n, the resulting game $\varsigma_{\alpha^n}(X)$ is finite.

Would this definition capture the intended meaning of repeated game play?

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

The issue is that each way of playing a repetition this way would require Angel to choose a natural number $n \in \mathbb{N}$ of repetitions and *expose this number to Demon* when playing α^n so that he would know how often Angel decided to repeat.

That would lead to what is called the *advance notice semantics* for α^*, which requires the players to announce the number of times that game α will be repeated when the loop begins. The advance notice semantics defines $\varsigma_{\alpha^*}(X)$ as $\bigcup_{n<\omega} \varsigma_{\alpha^n}(X)$ where $\alpha^{n+1} \equiv \alpha^n; \alpha$ and $\alpha^0 \equiv ?true$ and defines $\delta_{\alpha^*}(X)$ as $\bigcap_{n<\omega} \delta_{\alpha^n}(X)$. When playing α^*, Angel, thus, announces to Demon how many repetitions n are going to be played when the game α^* begins and Demon announces how often to repeat $\alpha^\times$. This advance notice makes it easier for Demon to win loops α^* and easier for Angel to win loops $\alpha^\times$, because the opponent announces an important feature of their strategy immediately as opposed to revealing whether or not to repeat the game once more one iteration at a time as in Def. 2. Angel announces the number $n < \omega$ of repetitions when α^* starts.

The following formula, for example, turns out to be valid in dG$\mathcal{L}$ (see Fig. 1), but would not be valid in the advance notice semantics:

$$x = 1 \land a = 1 \to \langle ((x := a; a := 0) \cap x := 0)^* \rangle x \neq 1 \tag{1}$$

If, in the advance notice semantics, Angel announces that she has chosen n repetitions of the game, then Demon wins (for $a \neq 0$) by choosing the $x := 0$ option $n - 1$ times followed by one choice of $x := a; a := 0$ in the last repetition. This strategy would not work in the dG$\mathcal{L}$ semantics, because Angel is free to decide whether to repeat α^* after each repetition based on the resulting state of the game. The winning strategy for (1) indicated in Fig. 1(left) shows that this dG$\mathcal{L}$ formula is valid.

Since the advance notice semantics misses out on the existence of perfectly reasonable winning strategies, dG$\mathcal{L}$ does not choose this semantics. Nevertheless, the advance notice semantics can be a useful semantics to consider for other purposes [QP12].

6 ω-Strategic Semantics

The trouble with the semantics in Sect. 5 is that Angel's move for the repetition reveals too much to Demon, because Demon can inspect the remaining game α^n to find out once and for all how long the game will be played before he has to do his first move.

Let's try to undo this. Another alternative choice for the semantics would have been to allow only arbitrary finite iterations of the strategy function for computing the winning region by using the *ω-strategic semantics*, which defines

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

along with a corresponding definition for $\delta_{\alpha^*}(X)$. All we need to do for this is define what it means to nest the winning region construction. For any winning condition

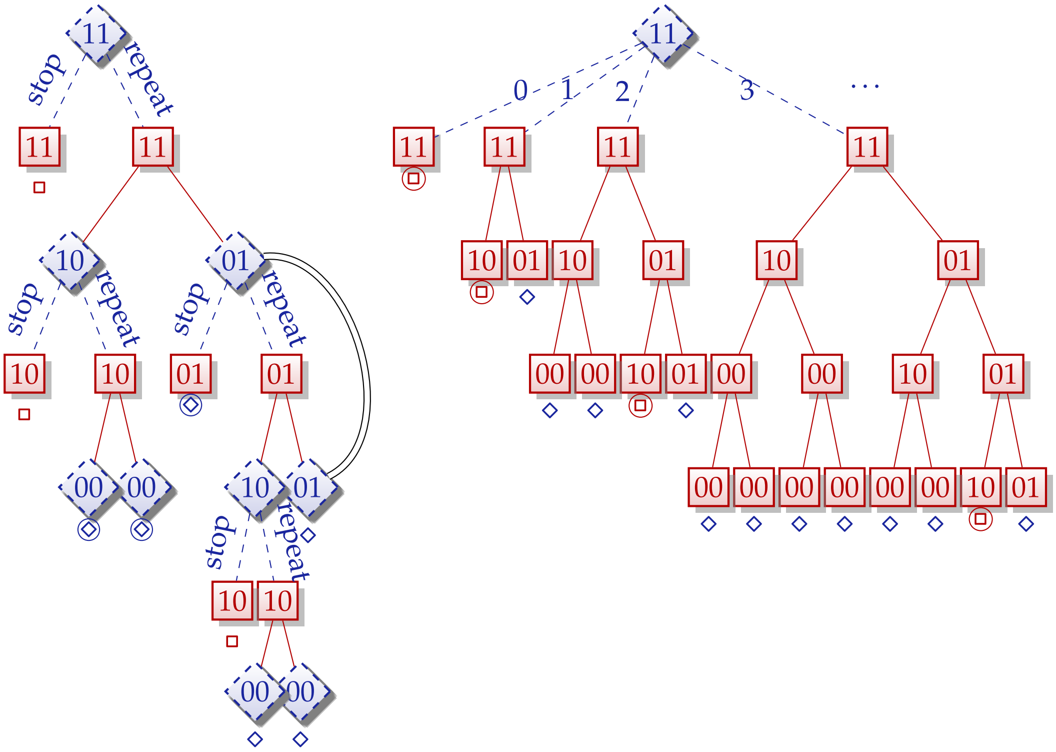

Figure 1: Game trees for $x = 1 \wedge a = 1 \to \langle \alpha^* \rangle x \neq 1$ with game $\alpha \equiv (x := a; a := 0) \cap x := 0$ (notation: x, a). **(left)** valid in dG$\mathcal{L}$ by strategy "repeat once and repeat once more if $x = 1$, then stop" **(right)** false in advance notice semantics by the strategy "$n-1$ choices of $x := 0$ followed by $x := a; a := 0$ once", where n is the number of repetitions Angel announced

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

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

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? See Fig. 2 for an illustration.

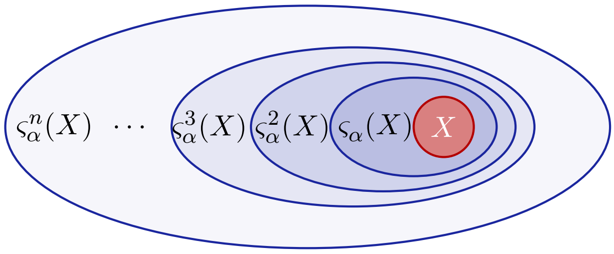

Figure 2: Iteration $\varsigma_\alpha^n(X)$ of $\varsigma_\alpha(\cdot)$ from winning condition X.

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

The surprising answer is *no* for a very subtle but also very fundamental reason. The existence of winning strategies for α^* does not coincide with the ωth iteration of α. This will be investigated further in the next lecture.

References

[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.

[QP12] Jan-David Quesel and André Platzer. Playing hybrid games with KeYmaera. In Bernhard Gramlich, Dale Miller, and Ulrike Sattler, editors, *IJCAR*, volume 7364 of *LNCS*, pages 439–453. Springer, 2012. doi:10.1007/978-3-642-31365-3_34.