

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, Pla14], that [Lecture 20 on Hybrid Systems & Games](#) started. [Lecture 20](#) saw the introduction of differential game logic with a focus on identifying and highlighting the new dynamical aspect of adversarial dynamics. The meaning of hybrid games in differential game logic had been left informal, based on the intuition one relates to interactive gameplay and decisions in trees. While it is possible to turn such a tree-type semantics into an operational semantics for hybrid games [Pla14], the resulting development is technically rather involved. Even if such an operational semantics is interesting and touches on interesting concepts from descriptive set theory, it is unnecessarily complicated compared.

This lecture will, thus, be devoted to developing a much simpler yet rigorous semantics, a denotational semantics of hybrid games. [Lecture 20](#) already highlighted subtleties how never-ending game play ruins determinacy, simply because there never is a state in which the winner would be declared. Especially the aspect of repetition and its interplay with differential equations will need careful attention. The denotational semantics will make this subtle aspect crystal clear.

These lecture notes are based on [Pla13, Pla14], where more information can be found on logic and hybrid games. The most important learning goals of this lecture are:

Modeling and Control: We further our understanding of the core principles behind CPS for the adversarial dynamics resulting from multiple agents with possibly conflicting actions that occur in many CPS applications. This time, we devote attention to the nuances of their semantics.

Computational Thinking: This lecture follows fundamental principles from computational thinking to capture the semantics of the new phenomenon of adversarial

dynamics in CPS models. We leverage core ideas from programming languages by extending syntax and semantics of program models and specification and verification logics with the complementary operator of duality to incorporate adversariality in a modular way into the realm of hybrid systems models. This leads to a compositional model of hybrid games with compositional operators. Modularity makes it possible to generalize our rigorous reasoning principles for CPS to hybrid games while simultaneously taming their complexity. This lecture introduces the semantics of *differential game logic* dGL [Pla13, Pla14], which adds adversarial dynamics to differential dynamic logic, which has been used as the specification and verification language for CPS in the other parts of this course. This lecture provides a perspective on advanced models of computation with alternating choices. The lecture will also encourage us to reflect on the relationship of denotational and operational semantics.

CPS Skills: This lecture focuses on developing and understanding the semantics of CPS models with adversarial dynamics corresponding to how a system changes state over time as multiple agents react to each other. This understanding is crucial for developing an intuition for the operational effects of multi-agent CPS. The presence of adversarial dynamics will cause us to reconsider the semantics of CPS models to incorporate the effects of multiple agents and their mutual reactions. This generalization, while crucial for understanding adversarial dynamics in CPS, also shines a helpful complementary light on the semantics of hybrid systems without adversariality by causing us to reflect on choices. The semantics of hybrid games properly generalizes the semantics of hybrid systems from earlier lectures.

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 dGL formulas ϕ that are used in the postconditions of dGL 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 $\llbracket \phi \rrbracket^I$, which defines the semantics of ϕ . The I in that notation is a reminder that the semantics depends on the interpretation of predicate symbols as defined in interpretation I . Thus, when we used to write $\nu \models \phi$ to indicate that dL formula ϕ is true in state ν , we will now write $\nu \in \llbracket \phi \rrbracket^I$, instead, to say that state ν is among the set of states in which ϕ is true. Working with the set of states $\llbracket \phi \rrbracket^I$ in which a formula ϕ is true will come in handy for defining a semantics of hybrid games.

The logic dGL has a denotational semantics. The dGL semantics defines, for each formula ϕ , the set $\llbracket \phi \rrbracket^I$ of states in which ϕ is true. For each hybrid game α and each set of winning states X , the dGL 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^c . 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 $\llbracket \theta \rrbracket_\nu$. The denotational semantics of dGL 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, because dGL 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 later as well.

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

1. $\llbracket p(\theta_1, \dots, \theta_k) \rrbracket^I = \{\nu \in \mathcal{S} : (\llbracket \theta_1 \rrbracket_\nu, \dots, \llbracket \theta_k \rrbracket_\nu) \in I(p)\}$
That is, the set of states in which a predicate $p(\theta_1, \dots, \theta_k)$ is true is the set of states ν in which the tuple $(\llbracket \theta_1 \rrbracket_\nu, \dots, \llbracket \theta_k \rrbracket_\nu)$ of values of the terms θ_i in ν is in the relation $I(p)$ associated to predicate symbol p .
2. $\llbracket \theta_1 \geq \theta_2 \rrbracket^I = \{\nu \in \mathcal{S} : \llbracket \theta_1 \rrbracket_\nu \geq \llbracket \theta_2 \rrbracket_\nu\}$
That is, the set of states in which $\theta_1 \geq \theta_2$ is true is the set in which the value of θ_1 is greater than or equal to the value θ_2 .
3. $\llbracket \neg \phi \rrbracket^I = (\llbracket \phi \rrbracket^I)^c$
That is, the set of states in which $\neg \phi$ is true is the complement of the set of states in which ϕ is true.
4. $\llbracket \phi \wedge \psi \rrbracket^I = \llbracket \phi \rrbracket^I \cap \llbracket \psi \rrbracket^I$
That is, the set of states in which $\phi \wedge \psi$ is true is the intersection of the states in which ϕ is true with the set of states in which ψ is true.
5. $\llbracket \exists x \phi \rrbracket^I = \{\nu \in \mathcal{S} : \nu_x^r \in \llbracket \phi \rrbracket^I \text{ for some } r \in \mathbb{R}\}$
That is, the states in which $\exists x \phi$ is true are those which only differ in the value of x from a state in which ϕ is true.
6. $\llbracket \langle \alpha \rangle \phi \rrbracket^I = \varsigma_\alpha(\llbracket \phi \rrbracket^I)$
That is, the set of states in which $\langle \alpha \rangle \phi$ is true is Angel's winning region to achieve $\llbracket \phi \rrbracket^I$ in hybrid game α , i.e. the set of states from which Angel has a winning strategy in hybrid game α to reach a state where ϕ holds.
7. $\llbracket [\alpha] \phi \rrbracket^I = \delta_\alpha(\llbracket \phi \rrbracket^I)$
That is, the set of states in which $[\alpha] \phi$ is true is Demon's winning region to achieve $\llbracket \phi \rrbracket^I$ in hybrid game α , i.e. the set of states from which Demon has a winning strategy in hybrid game α to reach a state where ϕ holds.

A dGL formula ϕ is *valid in I* , written $I \models \phi$, iff $\llbracket \phi \rrbracket^I = \mathcal{S}$. Formula ϕ is *valid*, $\models \phi$, iff $I \models \phi$ for all interpretations I .

The semantics $\varsigma_\alpha(X)$ and $\delta_\alpha(X)$ of Angel's and Demon's winning regions still needs to be defined, which is the next goal.

Note that the semantics of $\langle \alpha \rangle \phi$ cannot be defined as it would in dL via

$$\llbracket \langle \alpha \rangle \phi \rrbracket^I = \left\{ \nu \in \mathcal{S} : \omega \in \llbracket \phi \rrbracket^I \text{ for some } \omega \text{ with } (\nu, \omega) \in \rho(\alpha) \right\}$$

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 $\omega \in \llbracket \phi \rrbracket^I$ 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 $\varsigma_\alpha(\llbracket \phi \rrbracket^I)$ is supposed to capture. It captures the set of states from which Angel has a winning strategy in hybrid game α to achieve one of the states in which ϕ holds true. What a beneficial coincidence that the semantics of dG \mathcal{L} formulas was already defined in terms of the set of states in which they are true.

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.¹

¹The 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. For brevity, the following informal explanations sometimes say “win the game” when really they mean “have a winning strategy to win the game”.

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

1. $\varsigma_{x:=\theta}(X) = \{\nu \in \mathcal{S} : \nu_x^{\llbracket \theta \rrbracket} \in X\}$
That is, an assignment $x := \theta$ wins a game into X from any state whose modification $\nu_x^{\llbracket \theta \rrbracket}$ after the change $x := \theta$ is in X .
2. $\varsigma_{x'=\theta \& H}(X) = \{\varphi(0) \in \mathcal{S} : \varphi(r) \in X \text{ for some } r \in \mathbb{R}_{\geq 0} \text{ and (differentiable) } \varphi : [0, r] \rightarrow \mathcal{S} \text{ such that } \varphi(\zeta) \in \llbracket H \rrbracket^I \text{ and } \frac{d\varphi(t)(x)}{dt}(\zeta) = \llbracket \theta \rrbracket_{\varphi(\zeta)} \text{ for all } 0 \leq \zeta \leq r\}$
That is, Angel wins the differential equation game $x' = \theta \& H$ into X from any state $\varphi(0)$ from which there is a solution φ of $x' = \theta$ of any duration r that remains within H all the time and leads to a state $\varphi(r) \in X$ in the end.
3. $\varsigma_{?H}(X) = \llbracket H \rrbracket^I \cap X$
That is, Angel wins into X for a challenge $?H$ from the states which satisfy H to pass the challenge and are already in X , because challenges $?H$ do not change the state.
4. $\varsigma_{\alpha \cup \beta}(X) = \varsigma_\alpha(X) \cup \varsigma_\beta(X)$
That is, Angel wins a game of choice $\alpha \cup \beta$ into X whenever she wins α into X or wins β into X (by choosing a subgame she has a winning strategy for).
5. $\varsigma_{\alpha; \beta}(X) = \varsigma_\alpha(\varsigma_\beta(X))$
That is Angel wins a sequential game $\alpha; \beta$ into X whenever she has a winning strategy in game α to achieve $\varsigma_\beta(X)$, i.e. to make it to one of the states from which she has a winning strategy in game β to achieve X .
6. $\varsigma_{\alpha^*}(X)$ will be defined later.
7. $\varsigma_{\alpha^d}(X) = (\varsigma_\alpha(X^{\complement}))^{\complement}$
That is, Angel wins α^d to achieve X in exactly the states in which she does not have a winning strategy in game α to achieve the opposite X^{\complement} .

Demon's winning regions are defined accordingly (Def. 3).

Definition 3 (Semantics of hybrid games, continued). 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. $\delta_{x:=\theta}(X) = \{\nu \in \mathcal{S} : \nu_x^{\llbracket \theta \rrbracket} \in X\}$
That is, an assignment $x := \theta$ wins a game into X from any state whose modification $\nu_x^{\llbracket \theta \rrbracket}$ after the change $x := \theta$ is in X .
2. $\delta_{x'=\theta \& H}(X) = \{\varphi(0) \in \mathcal{S} : \varphi(r) \in X \text{ for all } r \in \mathbb{R}_{\geq 0} \text{ and (differentiable) } \varphi : [0, r] \rightarrow \mathcal{S} \text{ such that } \varphi(\zeta) \in \llbracket H \rrbracket^I \text{ and } \frac{d\varphi(t)(x)}{dt}(\zeta) = \llbracket \theta \rrbracket_{\varphi(\zeta)} \text{ for all } 0 \leq \zeta \leq r\}$
That is, Demon wins the differential equation game $x' = \theta \& H$ into X from any state $\varphi(0)$ from which all solutions φ of $x' = \theta$ of any duration r that remain within H all the time lead to states $\varphi(r) \in X$ in the end.
3. $\delta_{?H}(X) = (\llbracket H \rrbracket^I)^c \cup X$
That is, Demon wins into X for a challenge $?H$ from the states which violate H so that Angel fails her challenge $?H$ or that are already in X , because challenges $?H$ do not change the state.
4. $\delta_{\alpha \cup \beta}(X) = \delta_\alpha(X) \cap \delta_\beta(X)$
That is, Demon wins a game of choice $\alpha \cup \beta$ into X whenever he wins α into X and also wins β into X (because Angel might choose either subgame).
5. $\delta_{\alpha; \beta}(X) = \delta_\alpha(\delta_\beta(X))$
That is Demon wins a sequential game $\alpha; \beta$ into X whenever he has a winning strategy in game α to achieve $\delta_\beta(X)$, i.e. to make it to one of the states from which he has a winning strategy in game β to achieve X .
6. $\delta_{\alpha^*}(X)$ will be defined later.
7. $\delta_{\alpha^d}(X) = (\delta_\alpha(X^c))^c$
That is, Demon wins α^d to achieve X in exactly the states in which he does not have a winning strategy in game α to achieve the opposite X^c .

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 dGL semantics, because it is based on the existence of winning strategies, not on the strategies themselves. The winning regions for Angel are illustrated in Fig. 1.

Just as the semantics dL, the semantics of dGL is *compositional*, i.e. the semantics of a compound dGL 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 dGL 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

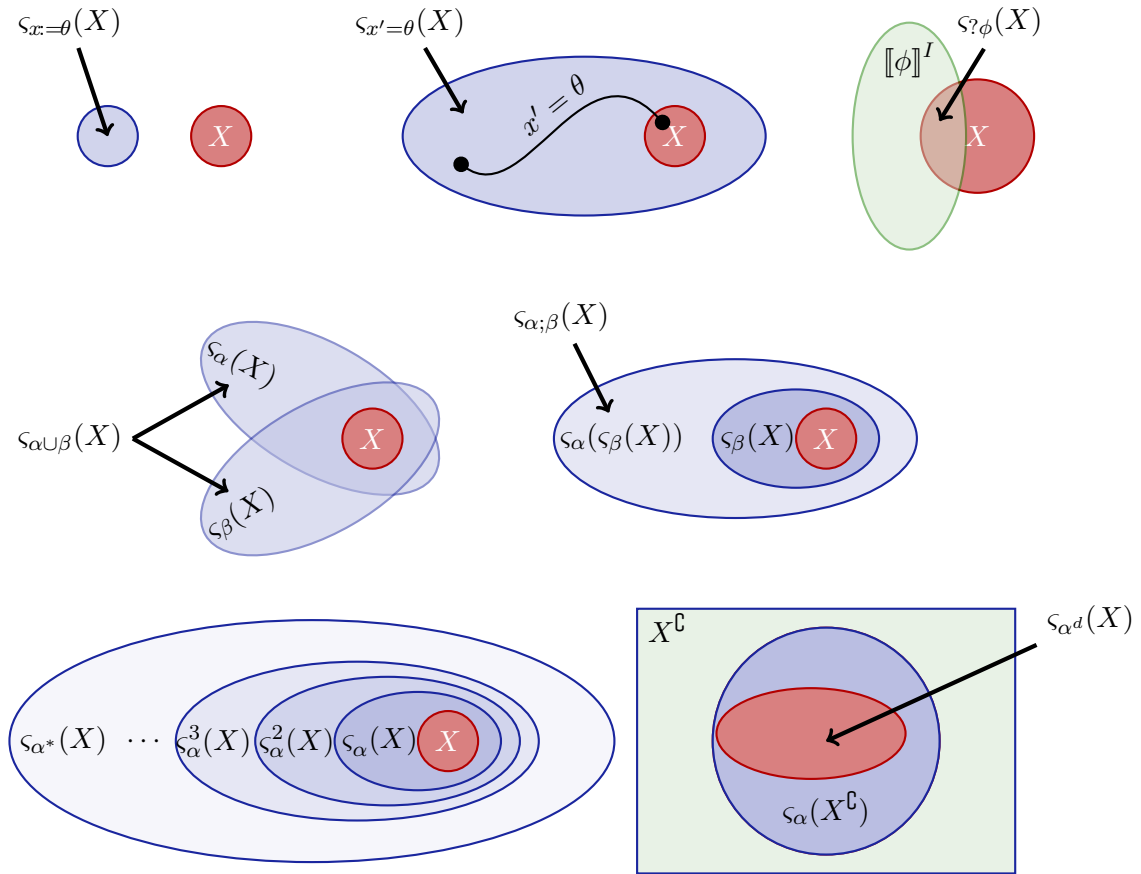


Figure 1: Illustration of denotational semantics of hybrid games as winning regions

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, because it is easier to win into larger sets of winning states.

Lemma 4 (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 $\varsigma_{\alpha^d}(X) = (\varsigma_{\alpha}(X^{\complement}))^{\complement} \subseteq (\varsigma_{\alpha}(Y^{\complement}))^{\complement} = \varsigma_{\alpha^d}(Y)$. \square

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

4 Advance Notice Repetitions

Def. 2 is still missing a definition for the semantics of repetition in hybrid games. 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 $\alpha^0 \equiv ?true$.

The obvious counterpart for the semantics of repetition in hybrid games would, thus, be

$$\varsigma_{\alpha^*}(X) \stackrel{?}{=} \bigcup_{n < \omega} \varsigma_{\alpha^n}(X) \quad (1)$$

where ω is the first infinite ordinal (if you have never seen ordinals before, just read $n < \omega$ as n in natural numbers, i.e. as $n \in \mathbb{N}$). 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 according to (1) 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 α^\times . This advance notice makes it easier for Demon to win loops α^* and easier for Angel to win loops α^\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 dGL (see Fig. 2), but would not be valid in the advance notice semantics:

$$x = 1 \wedge a = 1 \rightarrow \langle ((x := a; a := 0) \cap x := 0)^* \rangle x \neq 1 \quad (2)$$

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 dGL 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 (2) indicated in Fig. 2(left) shows that this dGL formula is valid.

Since the advance notice semantics misses out on the existence of perfectly reasonable winning strategies, dGL does not choose this semantics. Nevertheless, the advance notice semantics can be a useful semantics to consider for other purposes [QP12]. But it is not interactive enough for proper hybrid game play.

5 ω -Strategic Semantics

The trouble with the semantics in Sect. 4 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

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

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

The only states from which a repetition can win without actually repeating are the ones that start at the goal X already ($\varsigma_{\alpha}^0(X) = X$). And the states from which a repetition can win into X with $\kappa + 1$ repetitions are those that start in X as well as all the states for which there is a winning strategy in the hybrid game α to achieve a state in $\varsigma_{\alpha}^{\kappa}(X)$.

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. 3 for an illustration.




Figure 3: 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 ω th iteration of α .

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

$$\langle (x := 1; x' = 1^d \cup x := x - 1)^* \rangle (0 \leq x < 1) \quad (3)$$

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

Abbreviate

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

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

$$\begin{aligned} \varsigma_{\beta \cup \gamma}^0([0, 1]) &= [0, 1) \\ \varsigma_{\beta \cup \gamma}^{n+1}([0, 1]) &= [0, 1) \cup \varsigma_{\beta \cup \gamma}(\varsigma_{\beta \cup \gamma}^n([0, 1])) \stackrel{\text{IH}}{=} [0, 1) \cup \varsigma_{\beta \cup \gamma}([0, n + 1)) \\ &= [0, 1) \cup \varsigma_{\beta}([0, n + 1)) \cup \varsigma_{\gamma}([0, n)) = [0, 1) \cup \emptyset \cup [1, n + 2) = [0, n + 1 + 1) \end{aligned}$$

Consequently,

$$\varsigma_{\alpha}^{\omega}([0, 1]) = \bigcup_{n < \omega} \varsigma_{\alpha}^n([0, 1]) = \bigcup_{n < \omega} [0, n + 1) = [0, \infty)$$

Hence, the ω -semantics would indicate that the hybrid game (3) can exactly be won from all initial states in $[0, \infty)$, that is, for all initial states that satisfy $0 \leq x$.

Unfortunately, this is quite some nonsense. Indeed, the hybrid game in dGL formula (3) can be won from all initial states that satisfy $0 \leq x$. But it can also be won from other initial states! So the ω -strategic semantics $\varsigma_{\alpha}^{\omega}([0, 1])$ 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 (3), 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 $0 \leq x < 1$ ". This winning strategy wins from every initial state in \mathbb{R} , which is a much bigger set than $\varsigma_{\alpha}^{\omega}([0, 1]) = [0, \infty)$.

Now this is the final answer for the winning region of (3). In particular, the dGL formula (3) is valid. Yet, is there a direct way to see that $\varsigma_{\alpha}^{\omega}([0, 1]) = [0, \infty)$ is not the final answer for (3) 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 (3) 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 5 (+1 argument). 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. That is

$$Y \subseteq \varsigma_{\alpha^*}(X) \text{ then } \varsigma_{\alpha}(Y) \subseteq \varsigma_{\alpha^*}(X)$$

Applying Note 5 to the situation at hand works as follows. The above inductive proof showed $\varsigma_{\alpha}^{\omega}([0, 1)) = [0, \infty)$, which explains that at least $[0, \infty) \subseteq \varsigma_{(\beta \cup \gamma)^*}([0, 1))$ is in the winning region of repetition. By Note 5, the winning region $\varsigma_{(\beta \cup \gamma)^*}([0, 1))$ should, thus, also contain its one-step winning region $\varsigma_{\beta \cup \gamma}([0, \infty)) \subseteq \varsigma_{(\beta \cup \gamma)^*}([0, 1))$. Computing what that is gives

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

Beyond that, the winning region cannot contain anything else, because \mathbb{R} is the whole state space already and it is kind of hard to add anything to that. 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 (3) 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 (3), 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. 4 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 be much larger [Pla13].

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.




Figure 4: Iteration $\varsigma_\alpha^{\omega+1}(X)$ of $\varsigma_\alpha(\cdot)$ from winning condition $X = [0, 1)$ stops when applying $\varsigma_\alpha(\cdot)$ to the ω th infinite iteration $\varsigma_\alpha^\omega(X)$.

Note 6. *The ω -semantics is inappropriate, because it can be arbitrarily far away from characterizing the winning region of hybrid games.*

More generally, the semantics of repetition could be defined using

$$\begin{aligned} \varsigma_\alpha^0(X) &\stackrel{\text{def}}{=} X \\ \varsigma_\alpha^{\kappa+1}(X) &\stackrel{\text{def}}{=} X \cup \varsigma_\alpha(\varsigma_\alpha^\kappa(X)) \\ \varsigma_\alpha^\lambda(X) &\stackrel{\text{def}}{=} \bigcup_{\kappa < \lambda} \varsigma_\alpha^\kappa(X) \quad \lambda \neq 0 \text{ a limit ordinal} \end{aligned}$$

where we keep on computing winning regions at limit ordinals λ such as ω as the union of all previous winning regions. The semantics of repetition could then be defined as the union of all winning regions for all ordinals:

$$\varsigma_{\alpha^*}(X) = \varsigma^\infty(\alpha)X = \bigcup_{\kappa \text{ ordinal}} \varsigma_\alpha^\kappa(X)$$

Note 7. *Unfortunately, hybrid games might require rather big infinite ordinals until this inflationary style of computing their winning regions stops [Pla14]. That translates into an infinite amount of work and then some more, infinitely often, to compute the winning region starting from \emptyset . Hardly the sort of thing we would like to wait for until we finally know who wins the game.*

Finally look back at dGL formula (3) and observe what the above argument about the winning region computation terminating at $\omega + 1$ implies about bounds on how long it takes Angel to win the game in (3). Since the winning region only terminates at $\omega + 1$, she could not win with any finite bound $n \in \mathbb{N}$ on the number of repetitions it takes her to win. Even though she will surely win in the end according to her winning strategy, she has no way of saying how long that would take.

Not that Angels would ever do that. But suppose she were to brag to impress Demon by saying she could win within $n \in \mathbb{N}$ repetitions, then it would be hard for her to keep that promise. No matter how big a bound $n \in \mathbb{N}$ she chose, Demon could always spoil it from any negative initial state by evolving his differential equation $x' = 1^d$ for longer than n time units so that it takes Angel more than n rounds to decrease the resulting value down to $[0, 1)$ again.

This illustrates the dual of the discussion on the advance notice semantics in Sect. 4, which showed that Demon could make Angel win faster than she announced just to make her lose in the final round. In (3), Demon can always make Angel win later than she promised even if she ultimately will still win. This is the sense in which $\omega + 1$ is the only bound on the number of rounds it takes Angel to win the hybrid game in (3). This shows that a variation of the advance notice semantics based on Angel announcing to repeat at most $n \in \mathbb{N}$ times (as opposed to exactly $n \in \mathbb{N}$ times) would not capture the semantics of repetition appropriately.

Expedition 1 (Ordinal arithmetic). Ordinals extend natural numbers. Natural numbers are inductively defined as the (smallest) set \mathbb{N} containing 0 and the successor $n + 1$ of every number $n \in \mathbb{N}$ that is in the set. Natural numbers are totally ordered. Given any two different natural numbers, one number is going to be strictly smaller than the other one. For every finite set of natural numbers there is a smallest natural number that's bigger than all of them. Ordinals extend this beyond infinity. They just refuse to stop after all natural numbers have been written down:

$$0 < 1 < 2 < 3 < \dots$$

Taking all those (countably infinitely many) natural numbers $\{0, 1, 2, 3, \dots\}$, there is a smallest ordinal that's bigger than all of them. This ordinal is ω , the first infinite ordinal.^a

$$0 < 1 < 2 < 3 < \dots < \omega$$

Unlike the ordinals $1, 2, 3, \dots$ from the natural numbers, the ordinal ω is a *limit ordinal*, because it is not the successor of any other ordinal. The ordinals $1, 2, 3, \dots$ are *successor ordinals*, because each of them is the successor $n + 1$ of another ordinal n . The ordinal 0 is special, because it is not a successor ordinal of any ordinal or natural number.

Now, since ordinals are keen on satisfying that every ordinal has a successor, or that every set of ordinals has an ordinal that is bigger, ω must have a successor as well. Its successor is the successor ordinal $\omega + 1$, the successor of which is $\omega + 2$ and so on:

$$0 < 1 < 2 < \dots < \omega < \omega + 1 < \omega + 2 < \dots$$

Of course, in ordinal land, there ought to be an ordinal that's bigger than even all of those ordinals as well. It's the limit ordinal $\omega + \omega = \omega \cdot 2$, at which point we have counted to countable infinity twice already and will keep on finding bigger

ordinals, because even $\omega \cdot 2$ will have a successor, namely $\omega \cdot 2 + 1$:

$$0 < 1 < 2 < \dots < \omega < \omega + 1 < \omega + 2 < \dots \omega \cdot 2 < \omega \cdot 2 + 1 < \omega \cdot 2 + 2 < \dots$$

Now the set of all these will have a bigger ordinal $\omega \cdot 2 + \omega = \omega \cdot 3$, which again has successors and so on. That happens infinitely often so that $\omega \cdot n$ will be an ordinal for any natural number $n \in \mathbb{N}$. All those infinitely many ordinals will also have a limit ordinal that's bigger than all of them, which is $\omega \cdot \omega = \omega^2$. That one again has a successor $\omega^2 + 1$ and so on, also see Fig. 5:

$$0 < 1 < 2 < \dots \omega < \omega + 1 < \omega + 2 < \dots \omega \cdot 2 < \omega \cdot 2 + 1 < \dots \omega \cdot 3 < \omega \cdot 3 + 1 < \dots \\ \omega^2 < \omega^2 + 1 < \dots \omega^2 + \omega < \omega^2 + \omega + 1 < \dots \omega^\omega < \dots \omega^{\omega^\omega} < \dots \omega_1^{\text{CK}} < \dots \omega_1 < \dots$$

The first infinite ordinal is ω , the Church-Kleene ordinal ω_1^{CK} , i.e. the first nonrecursive ordinal, and ω_1 is the first uncountable ordinal. Every ordinal κ is either a successor ordinal, i.e. the smallest ordinal $\kappa = \iota + 1$ greater than some ordinal ι , or a limit ordinal, i.e. the supremum of all smaller ordinals. Depending on the context, 0 is considered a limit ordinal or separate.

Ordinals support (non-commutative) addition, multiplication, and exponentiation, which can be defined by induction on its second argument:

$$\begin{aligned} \iota + 0 &= \iota \\ \iota + (\kappa + 1) &= (\iota + \kappa) + 1 && \text{for successor ordinals } \kappa + 1 \\ \iota + \lambda &= \bigsqcup_{\kappa < \lambda} \iota + \kappa && \text{for limit ordinals } \lambda \\ \iota \cdot 0 &= 0 \\ \iota \cdot (\kappa + 1) &= (\iota \cdot \kappa) + \iota && \text{for successor ordinals } \kappa + 1 \\ \iota \cdot \lambda &= \bigsqcup_{\kappa < \lambda} \iota \cdot \kappa && \text{for limit ordinals } \lambda \\ \iota^0 &= 1 \\ \iota^{\kappa+1} &= \iota^\kappa \cdot \iota && \text{for successor ordinals } \kappa + 1 \\ \iota^\lambda &= \bigsqcup_{\kappa < \lambda} \iota^\kappa && \text{for limit ordinals } \lambda \end{aligned}$$

where \bigsqcup denotes the supremum or least-upper bound. Carefully note ordinal oddities like the noncommutativity coming from $2 \cdot \omega = 4 \cdot \omega$ and $\omega \cdot 2 < \omega \cdot 4$.

^aFor a moment read " $\omega = \infty$ " as infinity, but you will realize in an instant that this view does not go far enough, because there will be reason to distinguish different infinities.




Figure 5: Illustration of infinitely many ordinals up to ω^ω

6 Summary

This lecture saw the introduction of a proper formal semantics for differential game logic and hybrid games. This resulted in a simple denotational semantics, where the meaning of all formulas and hybrid games is a simple function of the meaning of its pieces. The only possible outlier was the semantics of repetition, which turned out to be rather subtle and ultimately defined by higher-ordinal iterations of winning region constructions. This led to an insightful appreciation for the complexities, challenges, and flexibilities of hybrid games.

The next lecture will revisit the semantics of repetition to find a simpler implicit characterization and leverage the semantic basis for the next leg in the logical trinity: axiomatics.

Exercises

Exercise 1. The formula (3) was shown to need $\omega + 1$ iterations of the winning region construction to terminate with the following answer justifying the validity of (3).

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

What happens if the winning region construction is used again to compute $\varsigma_{\alpha}^{\omega+2}([0, 1))$? How often does the winning region construction need to be iterated to justify validity of

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

Exercise 2. How often does the winning region construction need to be iterated to justify validity of

$$\langle (x := x - 1; y' = 1^d \cup y := y - 1; z' = 1^d \cup z := z - 1)^* \rangle (x < 0 \wedge y < 0 \wedge z < 0)$$

Exercise 3 (Clockwork ω).* How often does the winning region construction need to be iterated to justify validity of

$$\langle (?y < 0; x := x - 1; y' = 1^d \cup ?z < 0; y := y - 1; z' = 1^d \cup z := z - 1)^* \rangle (x < 0 \wedge y < 0 \wedge z < 0)$$

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.
- [Pla14] André Platzer. Differential game logic. *CoRR*, abs/1408.1980, 2014. arXiv:1408.1980.

- [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](https://doi.org/10.1007/978-3-642-31365-3_34).