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, Pla14], whose syntax was introduced in Lecture 20 on Hybrid Systems & Games and whose semantics was developed in Lecture 21 on Winning Strategies & Regions. Today's lecture furthers the development of differential game logic to the third leg of the logical trinity: its axiomatics. This lecture will focus on the development of rigorous reasoning techniques for hybrid games as models for CPS with adversarial dynamics. Without such analysis and reasoning techniques, a logic that only comes with syntax and semantics can be used as a specification language with a precise meaning, but would not be very helpful for actually analyzing and verifying hybrid games. It is the logical trinity of syntax, semantics, and axiomatics that gives logics the power of serving as well-founded specification and verification languages with a (preferably concise) syntax, an unambiguous semantics, and actionable analytic reasoning principles. Thus, today's lecture is the hybrid games analogue of Lecture 5 on Dynamical Systems and Dynamic Axioms. Indeed, after the logical sophistication we reached throughout the semester, this lecture will settle for a Hilbert-type calculus as in Lecture 5 on Dynamical Systems and Dynamic Axioms as opposed to the more refined and more automatable sequent calculus from Lecture 6 on Truth and Proof and subsequent lectures.

Before submerging completely into the development of rigorous reasoning techniques for hybrid games as models for CPS with adversarial dynamics, however, it will be wise to take a short detour by investigating a semantical simplification of the meaning of repetition by an implicit characterization of its winning region rather than the explicit construction by iteration from Lecture 21.

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 advance our understanding of the core principles behind CPS with hybrid games by understanding analytically and semantically how discrete, continuous, and the adversarial dynamics resulting, e.g., from multiple agents are integrated and interact in CPS. This lecture also uncovers nuances in the semantics of adversarial repetitions that makes them conceptually better behaved than the highly transfinite iterated winning region construction from Lecture 21. A byproduct of this development shows fixpoints in actions, which play a prominent role in the understanding of other classes of models as well and provides one important aspect for the subsequent development of reasoning techniques.

Computational Thinking: This lecture is devoted to the development of rigorous reasoning techniques for CPS models involving adversarial dynamics, which is critical to getting CPS with such interactions right. Hybrid games provide even more subtle interactions than hybrid systems did, which make it even more challenging to say for sure whether and why a design is correct without sufficient rigor in their analysis. After Lecture 21 captured the semantics of differential game logic and hybrid games compositionally, this lecture exploits the compositional meaning to develop compositional reasoning principles for hybrid games. This lecture systematically develops one reasoning principle for each of the operators of hybrid programs, resulting in a compositional verification approach. A compositional semantics is de facto a necessary but not a sufficient condition for the existence of compositional reasoning principles. Despite the widely generalized semantics of hybrid games compared to hybrid systems, this lecture will strive to generalize reasoning techniques for hybrid systems to hybrid games as smoothly as possible. This leads to a modular way of integrating adversariality into the realm of hybrid systems models also in terms of their analysis while simultaneously taming their complexity. This lecture provides an *axiomatization* of differential game logic dGℒ [Pla13, Pla14] to lift dGℒ from a specification language to a verification language for CPS with adversarial dynamics.

CPS Skills: We will develop a deep understanding of the semantics of CPS models with adversariality by carefully relating their semantics to their reasoning principles and aligning them in perfect unison. This understanding will also enable us to develop a better intuition for the operational effects involved in CPS. This lecture also shows insightful and influential nuances on the semantics of repetitions in CPS models with adversarial dynamics.

In our quest to develop rigorous reasoning principles for hybrid games, we will strive to identify compositional reasoning principles that align in perfect unison with the compositional semantics of hybrid games developed in Lecture 21 on Winning Strategies & Regions. This enterprise will be enlightening and, for the most part, quite successful. And, in fact, the reader is encouraged to start right away with the development of a proof calculus for differential game logic and later compare it with the one that these lecture notes develop. The part, where this will turn out to be rather difficult is repe-

tition, which is why the lecture notes take a scenic detour through characterizing their semantics.

2 Characterizing Winning Repetitions Implicitly

Lecture 21 on Winning Strategies & Regions culminated in a semantics of repetition defined as the union of all winning regions for all ordinals by an explicit (albeit wildly infinite) construction:

$$\varsigma_{\alpha^*}(X) = \varsigma^\infty(\alpha)X = \bigcup_{\kappa \text{ ordinal}} \varsigma_\alpha^\kappa(X) \quad \text{where}$$
$$\varsigma_\alpha^0(X) \overset{\text{def}}{=} X$$
$$\varsigma_\alpha^{\kappa+1}(X) \overset{\text{def}}{=} X \cup \varsigma_\alpha(\varsigma_\alpha^\kappa(X))$$
$$\varsigma_\alpha^\lambda(X) \overset{\text{def}}{=} \bigcup_{\kappa<\lambda} \varsigma_\alpha^\kappa(X) \qquad \lambda \neq 0 \text{ a limit ordinal}$$

Is there a more immediate way of characterizing the winning region $\varsigma_{\alpha^*}(X)$ of repetition implicitly rather than by explicit construction? This thought will lead to a beautiful illustration of Bertrand Russell's enlightening bonmot:

> The advantages of implicit definition over construction are roughly those of theft over honest toil. — Bertrand Russell (slightly paraphrased)

The above iterated winning region construction describes the semantics of repetition by iterating from below, i.e. starting from $\varsigma_\alpha^0(X) = X$ and adding states. Maybe the semantics of repetition could be characterized more indirectly but more concisely from above? With an implicit characterization.

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

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

This holds for any set $Z \subseteq \varsigma_{\alpha^*}(X)$. In particular, the set $Z \overset{\text{def}}{=} \varsigma_{\alpha^*}(X)$ itself satisfies

$$\varsigma_\alpha(\varsigma_{\alpha^*}(X)) \subseteq \varsigma_{\alpha^*}(X) \tag{1}$$

by Note 1. After all, repeating α once more from the winning region $\varsigma_{\alpha^*}(X)$ of repetition of α cannot give us any states that did not already have a winning strategy in α^*, because α^* could have been repeated one more time. Consequently, if a set $Z \subseteq \mathcal{S}$ claims to be the winning region $\varsigma_{\alpha^*}(X)$ of repetition, it at least has to satisfy

$$\varsigma_\alpha(Z) \subseteq Z \tag{2}$$

because, by (1), the true winning region $\varsigma_{\alpha^*}(X)$ does satisfy (2). Thus, strategyzing along α from Z does not give anything that Z would not already know about.

Is there anything else that such a set Z needs to satisfy to be the winning region $\varsigma_{\alpha^*}(X)$ of repetition? 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 exists, even though it may be rather boring. The empty set $Z \overset{\text{def}}{=} \emptyset$ looks like it would satisfy (2) because it is rather hard to win a game that requires Angel to enter the empty set of states to win.

On second thought, $\varsigma_\alpha(\emptyset) \subseteq \emptyset$ does not actually always hold for all hybrid games α. It is violated for states from which Angel can make sure Demon violates the rules of the game α by losing a challenge or failing to comply with evolution domain constraints. For example, when H is a nontrivial test like $x > 0$:

$$\varsigma_{?H^d}(\emptyset) = \varsigma_{?H}(\emptyset^\complement))^\complement = (\llbracket H \rrbracket^I \cap \mathcal{S})^\complement = (\llbracket H \rrbracket^I)^\complement = \llbracket \neg H \rrbracket^I \not\subseteq \emptyset$$

Yet, then the set of states that make Demon violate the rules satisfies (2) instead of $\emptyset$:

$$\varsigma_{?H^d}(\llbracket \neg H \rrbracket^I) = \varsigma_{?H}((\llbracket \neg H \rrbracket^I)^\complement))^\complement = \varsigma_{?H}(\llbracket H \rrbracket^I))^\complement = (\llbracket H \rrbracket^I \cap \llbracket H \rrbracket^I)^\complement = \llbracket \neg H \rrbracket^I \subseteq \llbracket \neg H \rrbracket^I$$

But even if the empty set $\emptyset$ satisfies (2), it may be a bit small. Likewise, even if $\llbracket \neg H \rrbracket^I$ satisfies (2) for $\alpha \equiv ?H^d$, the set $\llbracket \neg H \rrbracket^I$ may still be a bit small. Angel is still in charge of repetition and can decide how often to repeat and whether to repeat at all. 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 when already starting in X by simply suggesting Angel should repeat zero times. Angel would certainly love to do that, because it does not sound like a lot of work to repeat something zero times. Consequently, if a set $Z \subseteq \mathcal{S}$ claims to be the winning region $\varsigma_{\alpha^*}(X)$, then it has to satisfy (2) and

$$X \subseteq Z \tag{3}$$

Both conditions (2) and (3) together can be summarized in a single condition as follows:

Note 2 (Pre-fixpoint). *Every candidate Z for the winning region $\varsigma_{\alpha^*}(X)$ satisfies:*

$$X \cup \varsigma_\alpha(Z) \subseteq Z \tag{4}$$

A set Z satisfying condition (4) is called a pre-fixpoint.

Again: what is this set Z that satisfies (4)? Is there only one choice? Or multiple? If there are multiple choices, which Z is the right one for the semantics of repetition? 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 $X = \emptyset$ (and even then $\emptyset$ actually only works if Demon cannot be tricked into violating the rules of the game). The set X itself is too small as well unless the game has no incentive to start repeating, because $\varsigma_\alpha(X) \subseteq X$. But the full state space $Z \overset{\text{def}}{=} \mathcal{S}$ always satisfies (4) trivially so (4) has a solution. Now, the whole space is a little too 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

$$\langle (x := 1; x' = 1^d \cup x := x - 1)^* \rangle (0 \leq x < 1) \tag{5}$$

the full state space is hardly the right winning region for any arbitrary α^*. It definitely depends on the hybrid game α and the winning condition ϕ whether Angel has a winning strategy for $\langle\alpha\rangle\phi$ or not. 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 (4) hardly qualifies.

So which solution Z of (4) do we define to be $\varsigma_{\alpha^*}(X)$ now?

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

Among the many Z that solve (4), the largest one is not informative, because the largest Z simply degrades to $\mathcal{S}$. So smaller solutions Z are preferable. Which one? How do multiple solutions even relate to each other? Suppose Y, Z are both solutions of (4). That is

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

Then, by the monotonicity lemma from Lecture 21 (repeated in Lemma 3 below):

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

Hence, by (8), the intersection $Y \cap Z$ of solutions Y and Z of (4) also is a solution of (4):

Lemma 1 (Intersection closure)**.** *For any two solutions Y, Z of the prefix condition (4), a smaller solution of (4) can be obtained by intersection $Y \cap Z$.*

Whenever there are two solutions Z_1, Z_2 of (4), their intersection $Y_1 \cap Z_2$ solves (4) as well. When there's yet another solution Z_3 of (4), their intersection $Y_1 \cap Y_2 \cap Y_3$ also solves (4). Similarly for *any* larger family of solutions whose intersection will solve (4). If we keep on intersecting solutions, we will arrive at smaller and smaller solutions until, some fine day, there's not going to be a smaller one. This yields the smallest solution Z of (4) which can be characterized directly.

Note 4 (Semantics of repetitions)**.** *Among the many Z that solve (4), $\varsigma_{\alpha^*}(X)$ is defined to be the smallest Z that solves (4):*

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

The characterization in terms of iterated winning regions from Lecture 21 leads to the same set $\varsigma_{\alpha^*}(X)$, but the (least pre-fixpoint or) fixpoint characterization (9) is easier.

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 (4):

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

Also compare this with where we came from when we argued for (1). 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 such set. That is, 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 $\varsigma_{\alpha^*}(X)$ actually satisfies not just the inclusion (4) but even the equation

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

Note 5 (Semantics of repetitions, fixpoint formulation). *That is,* $\varsigma_{\alpha^*}(X)$ *is 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)*, i.e. it satisfies*

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

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, Pla14] 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, Pla14] way beyond the first infinite ordinal ω.

3 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), or in terms of transfinitely iterated winning region constructions, the winning regions of hybrid games can be defined directly, giving a denotational semantics to hybrid games.

The only difference of the following semantics 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

1. $\varsigma_{x:=\theta}(X) = \{\nu \in \mathcal{S} : \nu_x^{[\![\theta]\!]_\nu} \in X\}$
2. $\varsigma_{x'=\theta\,\&\,H}(X) = \{\varphi(0) \in \mathcal{S} : \varphi(r) \in X \text{ for some } r \in \mathbb{R}_{\geq 0} \text{ and (differentiable) } \varphi : [0,r] \to \mathcal{S} \text{ such that } \varphi(\zeta) \in [\![H]\!]^I \text{ and } \frac{\mathsf{d}\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)} \text{ for all } 0 \leq \zeta \leq r\}$
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. $\delta_{x'=\theta\,\&\,H}(X) = \{\varphi(0) \in \mathcal{S} : \varphi(r) \in X \text{ for all } r \in \mathbb{R}_{\geq 0} \text{ and (differentiable) } \varphi : [0,r] \to \mathcal{S} \text{ such that } \varphi(\zeta) \in [\![H]\!]^I \text{ and } \frac{\mathsf{d}\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)} \text{ for all } 0 \leq \zeta \leq r\}$
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. $\delta_{\alpha^d}(X) = (\delta_\alpha(X^\complement))^\complement$

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.

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. Fur-

thermore, 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 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.

4 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, Pla14]). *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 $X \overset{\text{def}}{=} [\![\phi]\!]^I$.

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 $\frac{\mathrm{d}\varphi(t)(x)}{\mathrm{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ 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. $\varsigma_{\alpha^*}(X^\complement)^\complement = \Big(\bigcap\{Z \subseteq \mathcal{S} : X^\complement \cup \varsigma_\alpha(Z) \subseteq Z\}\Big)^\complement = \Big(\bigcap\{Z \subseteq \mathcal{S} : (X \cap \varsigma_\alpha(Z)^\complement)^\complement \subseteq Z\}\Big)^\complement$
   $= \Big(\bigcap\{Z \subseteq \mathcal{S} : (X \cap \delta_\alpha(Z^\complement))^\complement \subseteq Z\}\Big)^\complement = \bigcup\{Z \subseteq \mathcal{S} : Z \subseteq X \cap \delta_\alpha(Z)\} = \delta_{\alpha^*}(X)$. [1]
7. $\varsigma_{\alpha^d}(X^\complement)^\complement = (\varsigma_\alpha((X^\complement)^\complement)^\complement)^\complement = \delta_\alpha(X^\complement)^\complement = \delta_{\alpha^d}(X)$ ∎

[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.

5 Hybrid Game Axioms

An axiomatization for differential game logic has been found in [Pla13, Pla14], where we refer to for more details.

Note 9 (Differential game logic axiomatization [Pla13, Pla14]).

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

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

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

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

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

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

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

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

$$\text{(M)} \quad \frac{\phi \rightarrow \psi}{\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\psi}$$

$$\text{(FP)} \quad \frac{\phi \vee \langle\alpha\rangle\psi \rightarrow \psi}{\langle\alpha^*\rangle\phi \rightarrow \psi}$$

$$\text{(ind)} \quad \frac{\phi \rightarrow [\alpha]\phi}{\phi \rightarrow [\alpha^*]\phi}$$

The determinacy axiom $[\cdot]$ describes the duality of winning strategies for complementary winning conditions of Angel and Demon, i.e. that Demon has a winning strategy to achieve ϕ in hybrid game α if and only if Angel does not have a counter strategy, i.e. winning strategy to achieve $\neg\phi$ in the same game α. The determinacy axiom $[\cdot]$ internalizes Theorem 4. Axiom $\langle:=\rangle$ is Hoare's assignment rule. Formula $\phi(\theta)$ is obtained from $\phi(x)$ by *substituting* θ for x at all occurrences of x, provided x does not occur in the scope of a quantifier or modality binding x or a variable of θ. A modality containing $x :=$ or x' outside the scope of tests $?H$ or evolution domain constraints *binds* x, because it may change the value of x. In the differential equation axiom $\langle'\rangle$, $y(\cdot)$ is the unique [Wal98, Theorem 10.VI] solution of the symbolic initial value problem $y'(t) = \theta, y(0) = x$. The duration t how long to follow solution y is for Angel to decide, hence existentially quantified. It goes without saying that variables like t are fresh in Fig. 9.

Axioms $\langle?\rangle$, $\langle\cup\rangle$, and $\langle;\rangle$ are as in differential dynamic logic [Pla12] except that their meaning is quite different, because they refer to winning strategies of hybrid games instead of reachability relations of systems. The challenge axiom $\langle?\rangle$ expresses that Angel

has a winning strategy to achieve ϕ in the test game $?H$ exactly from those positions that are already in ϕ (because $?H$ does not change the state) and that satisfy H for otherwise she would fail the test and lose the game immediately. The axiom of choice $\langle\cup\rangle$ expresses that Angel has a winning strategy in a game of choice $\alpha \cup \beta$ to achieve ϕ iff she has a winning strategy in either hybrid game α or in β, because she can choose which one to play. The sequential game axiom $\langle;\rangle$ expresses that Angel has a winning strategy in a sequential game $\alpha;\beta$ to achieve ϕ iff she has a winning strategy in game α to achieve $\langle\beta\rangle\phi$, i.e. to get to a position from which she has a winning strategy in game β to achieve ϕ. The iteration axiom $\langle^*\rangle$ characterizes $\langle\alpha^*\rangle\phi$ as a pre-fixpoint. It expresses that, if the game is already in a state satisfying ϕ or if Angel has a winning strategy for game α to achieve $\langle\alpha^*\rangle\phi$, i.e. to get to a position from which she has a winning strategy for game α^* to achieve ϕ, then, either way, Angel has a winning strategy to achieve ϕ in game α^*. The converse of $\langle^*\rangle$ can be derived[2] and is also denoted by $\langle^*\rangle$. The dual axiom $\langle^d\rangle$ characterizes dual games. It says that Angel has a winning strategy to achieve ϕ in dual game α^d iff Angel does not have a winning strategy to achieve $\neg\phi$ in game α. Combining dual game axiom $\langle^d\rangle$ with the determinacy axiom $[\cdot]$ yields $\langle\alpha^d\rangle\phi \leftrightarrow [\alpha]\phi$, i.e. that Angel has a winning strategy to achieve ϕ in α^d iff Demon has a winning strategy to achieve ϕ in α. Similar reasoning derives $[\alpha^d]\phi \leftrightarrow \langle\alpha\rangle\phi$.

Monotonicity rule M is the generalization rule of monotonic modal logic **C** [Che80] and internalizes Lemma 3. It expresses that, if the implication $\phi \to \psi$ is valid, then, from wherever Angel has a winning strategy in a hybrid game α to achieve ϕ, she also has a winning strategy to achieve ψ, because ψ holds wherever ϕ does. So rule M expresses that easier objectives are easier to win. Fixpoint rule FP characterizes $\langle\alpha^*\rangle\phi$ as a *least* pre-fixpoint. It says that, if ψ is another formula that is a pre-fixpoint, i.e. that holds in all states that satisfy ϕ or from which Angel has a winning strategy in game α to achieve that condition ψ, then ψ also holds whereever $\langle\alpha^*\rangle\phi$ does, i.e. in all states from which Angel has a winning strategy in game α^* to achieve ϕ.

The proof rules FP and the induction rule ind are equivalent in the sense that one can be derived from the other in the dGℒ calculus [Pla13, Pla14].

Example 5. The dual filibuster game formula from Lecture 20 proves easily in the dGℒ calculus by going back and forth between players [Pla13] using the abbreviations $\cap, ^\times$:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{*}{x=0 \vdash 0=0 \vee 1=0}\,\mathbb{R}}{x=0 \vdash \langle x:=0\rangle x=0 \vee \langle x:=1\rangle x=0}\,\langle:=\rangle}{x=0 \vdash \langle x:=0 \cup x:=1\rangle x=0}\,\langle\cup\rangle}{x=0 \vdash \neg\langle x:=0 \cap x:=1\rangle\neg x=0}\,\langle^d\rangle}{x=0 \vdash [x:=0 \cap x:=1]x=0}\,[\cdot]}{x=0 \vdash [(x:=0 \cap x:=1)^*]x=0}\,\text{ind}}{x=0 \vdash \langle(x:=0 \cup x:=1)^\times\rangle x=0}\,\langle^d\rangle$$

[2] $\phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi \to \langle\alpha^*\rangle\phi$ derives by $\langle^*\rangle$. Thus, $\langle\alpha\rangle(\phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi) \to \langle\alpha\rangle\langle\alpha^*\rangle\phi$ by M. Hence, $\phi \vee \langle\alpha\rangle(\phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi) \to \phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi$ by propositional congruence. Consequently, $\langle\alpha^*\rangle\phi \to \phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi$ by FP.

Theorem 6 (Soundness [Pla13, Pla14]). *The* dGL *proof calculus in Fig. 9 is sound, i.e. all provable formulas are valid.*

Proof. The full proof can be found in [Pla13, Pla14]. We just consider a few cases to exemplify the fundamentally more general semantics of hybrid games arguments compared to hybrid systems arguments. To prove soundness of an equivalence axiom $\phi \leftrightarrow \psi$, show $[\![\phi]\!]^I = [\![\psi]\!]^I$ for all interpretations I with any set of states $\mathcal{S}$.

$\langle\cup\rangle$ $[\![\langle\alpha \cup \beta\rangle\phi]\!]^I = \varsigma_{\alpha\cup\beta}([\![\phi]\!]^I) = \varsigma_\alpha([\![\phi]\!]^I) \cup \varsigma_\beta([\![\phi]\!]^I) = [\![\langle\alpha\rangle\phi]\!]^I \cup [\![\langle\beta\rangle\phi]\!]^I = [\![\langle\alpha\rangle\phi \vee \langle\beta\rangle\phi]\!]^I$

$\langle;\rangle$ $[\![\langle\alpha;\beta\rangle\phi]\!]^I = \varsigma_{\alpha;\beta}([\![\phi]\!]^I) = \varsigma_\alpha(\varsigma_\beta([\![\phi]\!]^I)) = \varsigma_\alpha([\![\langle\beta\rangle\phi]\!]^I) = [\![\langle\alpha\rangle\langle\beta\rangle\phi]\!]^I$.

$\langle?\rangle$ $[\![\langle?H\rangle\phi]\!]^I = \varsigma_{?H}([\![\phi]\!]^I) = [\![H]\!]^I \cap [\![\phi]\!]^I = [\![H \wedge \phi]\!]^I$

$[\cdot]$ is sound by Theorem 4.

M Assume the premise $\phi \to \psi$ is valid in interpretation I, i.e. $[\![\phi]\!]^I \subseteq [\![\psi]\!]^I$. Then the conclusion $\langle\alpha\rangle\phi \to \langle\alpha\rangle\psi$ is valid in I, i.e. $[\![\langle\alpha\rangle\phi]\!]^I = \varsigma_\alpha([\![\phi]\!]^I) \subseteq \varsigma_\alpha([\![\psi]\!]^I) = [\![\langle\alpha\rangle\psi]\!]^I$ by monotonicity (Lemma 3).

FP Assume the premise $\phi \vee \langle\alpha\rangle\psi \to \psi$ is valid in I, i.e. $[\![\phi \vee \langle\alpha\rangle\psi]\!]^I \subseteq [\![\psi]\!]^I$. That is, $[\![\phi]\!]^I \cup \varsigma_\alpha([\![\psi]\!]^I) = [\![\phi]\!]^I \cup [\![\langle\alpha\rangle\psi]\!]^I = [\![\phi \vee \langle\alpha\rangle\psi]\!]^I \subseteq [\![\psi]\!]^I$. Thus, ψ is a pre-fixpoint of $Z = [\![\phi]\!]^I \cup \varsigma_\alpha(Z)$. Now using Lemma 3, $[\![\langle\alpha^*\rangle\phi]\!]^I = \varsigma_{\alpha^*}([\![\phi]\!]^I)$ is the least fixpoint and the least pre-fixpoint [Koz06, Appendix A] of $Z = [\![\phi]\!]^I \cup \varsigma_\alpha(Z)$. Since $[\![\langle\alpha^*\rangle\phi]\!]^I$ is the least pre-fixpoint and $[\![\psi]\!]^I$ is a pre-fixpoint, $[\![\langle\alpha^*\rangle\phi]\!]^I \subseteq [\![\psi]\!]^I$ holds, which implies that $\langle\alpha^*\rangle\phi \to \psi$ is valid in I. □

So that is quite wonderful. This gives us a sound proof approach for CPSs that are as challenging as hybrid games. Now what exactly did we prove the axioms sound for? What does sound mean and entail exactly?

Note 11 (The miracle of soundness)**.** *Soundness of the* dGℒ *proof calculus means that all* dGℒ *formulas that are provable using the* dGℒ *calculus are valid. A* condition sine qua non *for logic, i.e. a condition without which logic could not be. It would not make sense to prove a formula if that would not even entail its validity.*

For a proof calculus to be sound, every formula that it proves with any proof has to be valid. Fortunately, proofs are composed by proof rules from axioms. So all we need to do to ensure that a proof calculus is sound is to prove the few axioms to be sound and then everything we ever derive from them by sound proof rules is correct as well, no matter how big and complicated. A proof is a long combination of many simple arguments, each of which just involve one of the axioms or proof rules. Once each of those finitely many axioms and proof rules are proved to be sound, all those infinitely many proofs that can be conducted in the dGℒ *proof calculus become sound as well. That is compositionality in its finest form for the soundness argument.* It is soundness that ultimately links semantics and axiomatics into perfect unison[a] so that axiomatic proof analysis coincides with semantic truth*, an important aspect of the logical trinity.*

One minor subtlety is that a proof could use many possible instances of the same finite axiom list, so that the soundness proof for the axioms has to work for any instance. This aspect is often left implicit in soundness arguments, although a rigorous treatment can be given by distinguishing axioms from axiom schemes [Pla13, Pla14].

[a] In search of perfection, completeness is another important aspect in achieving perfect unison, which, incidentally, holds for differential game logic as well [Pla13, Pla14]

6 Relating Differential Game Logic and Differential Dynamic Logic

Now that we have come to appreciate the value of soundness, couldn't we have known about that, for the most part, before Theorem 6? Most dGℒ axioms look rather familiar, except for $\langle\cdot\rangle$ versus $[\cdot]$ dualities, when we compare them to the dℒ axioms from Lecture 5 on Dynamical Systems and Dynamic Axioms. Does that not mean that these same axioms are already trivially sound? Why did we go through the (rather minor) trouble of proving Theorem 6?

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

It is not quite so easy. After all, we could have given the same syntactical operator $\cup$ an entirely different meaning for hybrid games than before for hybrid systems. Maybe we could have been silly and flipped the meaning of ; and $\cup$ around The fact of the matter is, of course, that we did not. The operator $\cup$ still means choice, just for hybrid games rather than hybrid systems. So could we deduce the soundness of the dG$\mathcal{L}$ axioms in Fig. 9 from the soundness of the corresponding d$\mathcal{L}$ axioms from Lecture 5 on Dynamical Systems and Dynamic Axioms and focus on the new axioms, only?

Before we do anything of the kind, we first need to convince ourselves that the d$\mathcal{L}$ semantics really coincides with the more general dG$\mathcal{L}$ semantics in case there are no games involved. How could that be done? Maybe by proving validity of all formulas of the following form

$$\underbrace{\langle\alpha\rangle\phi}_{\text{in d}\mathcal{L}} \leftrightarrow \underbrace{\langle\alpha\rangle\phi}_{\text{in dG}\mathcal{L}} \tag{13}$$

for dual-free hybrid games α, i.e. those that do not mention ^{d} (not even indirectly hidden in the abbreviation $\cap,^\times$).

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

The problem with (13) is that it is not directly a formula in any logic, because the $\leftrightarrow$ operator could hardly be applied meaningfully to two formulas from different logics. Well, of course, every dL formula is a dGL formula, so the left-hand side of (13) could be embedded into dGL, but then (13) becomes well-defined but is only stating a mere triviality.

Instead, a proper approach would be to rephrase the well-intended (13) semantically:

$$\underbrace{\nu \models \langle\alpha\rangle\phi}_{\text{in dL}} \text{ iff } \underbrace{\nu \in [\![\langle\alpha\rangle\phi]\!]^I}_{\text{in dGL}} \tag{14}$$

which is equivalent to

$$\big(\underbrace{\omega \models \phi \text{ for some } \omega \text{ with } (\nu,\omega) \in \rho(\alpha)}_{\text{statement about reachability in dL}}\big) \text{ iff } \underbrace{\nu \in \varsigma_\alpha([\![\phi]\!]^I)}_{\text{winning in dGL}}$$

Equivalence (14) can be shown. In fact, an exercise in Lecture 3 on Choice & Control already developed an understanding of the dL semantics based on sets of states, preparing for(14).

The trouble is that, besides requiring a proof itself, the equivalence (14) will still not quite justify soundness of the dGL axioms in Fig. 9 that look innocuously like dL axioms. Equivalence (14) is for dual-free hybrid games α. But even if the top-level operator in axiom $\langle\cup\rangle$ is not ^{d}, that dual operator could still occur within α or β, which requires a game semantics to make sense of.

Consequently, we are better off proving soundness for the dGL axioms according to their actual semantics, like in Theorem 6, as opposed to trying half-witted ways out that only make soundness matters worse.

Exercises

Exercise 1. Explain how often you will have to repeat the winning region construction to show that the following dGL formula is valid:

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

Exercise 2. Can you find dGL 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.

Exercise 4. Prove the elided cases of Lemma 3.

Exercise 5. Find the appropriate soundness notion for the axioms of dGL and prove that the axioms are sound.

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

References

[Che80] Brian F. Chellas. *Modal Logic: An Introduction*. Cambridge Univ. Press, 1980.

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

[Koz06] Dexter Kozen. *Theory of Computation*. Springer, 2006.

[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.

[Wal98] Wolfgang Walter. *Ordinary Differential Equations*. Springer, 1998.