André Platzer

Lecture Notes on Foundations of Cyber-Physical Systems

15-424/624/824 Foundations of Cyber-Physical Systems

Chapter 15
Winning Strategies & Regions

Synposis This chapter identifies a simple denotational semantics for hybrid games based on their winning regions, i.e. the set of states from which there is a winning strategy that wins the game for all strategies that the opponent might choose. Such a denotational semantics continues the successful trend in this book of understanding all operators in a compositional way. That is, the meaning of a compound hybrid game is a simple function of the meaning of its pieces. For repetitions in hybrid games, such a semantics will turn out to be surprisingly subtle, which will uncover a surprisingly rich complexity in hybrid games that is characteristically different from that of hybrid systems. This is the first indication that hybrid games come with their own unique sets of challenges beyond what hybrid systems already have in store for us.

15.1 Introduction

This chapter continues the study of hybrid games and their specification and verification logic, differential game logic [2], that Chap. 14 started. Chapter 14 saw the introduction of differential game logic with a primary focus on identifying and highlighting the new dynamical aspect of adversarial dynamics for modeling purposes. 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 game trees. While it is possible to turn such a tree-type semantics into an operational semantics for hybrid games [2], the resulting development is technically rather involved. Even if such an operational semantics is informative and touches on interesting concepts from descriptive set theory, it is quite unnecessarily complicated.

This chapter will, thus, be devoted to developing a much simpler yet rigorous semantics, a denotational semantics for hybrid games. Chapter 14 already highlighted subtleties how never-ending game play ruins determinacy (i.e. that one player always has a winning strategy), simply because there never is a state in which the winner would even be declared. Especially the aspect of repetition and its interplay with differential equations will need careful attention now. The denotational semantics will make this subtle aspect crystal-clear.

This chapter is based on previous work [2], where more information can be found on logic and hybrid games. The most important learning goals of this chapter 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 chapter 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 that each have a compositional semantics. Modularity makes it possible to generalize our rigorous reasoning principles for CPS to hybrid games while simultaneously taming their complexity. This chapter introduces the semantics of *differential game logic* dGL [2], which adds adversarial dynamics to the differential dynamic logic that has been used as the specification and verification language for CPS in the other parts of this textbook. Because of the rôle that alternation plays in hybrid games, this chapter also provides a perspective on advanced models of computation with alternating choices. Finally, the chapter will encourage us to reflect on the relationship of denotational and operational semantics.

CPS Skills: This chapter 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 chapters.

15.2 Semantics of Differential Game Logic

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 P that are used in the postconditions of dGL modal formulas $\langle\alpha\rangle P$ and $[\alpha]P$ define the winning

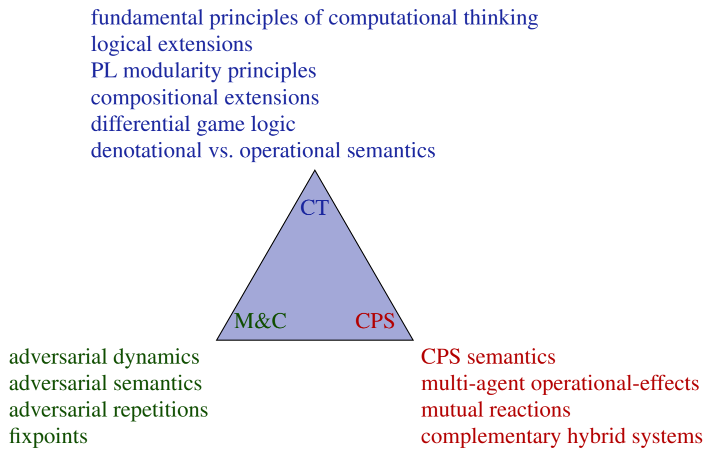

conditions for the hybrid game α. When playing the hybrid game α, we, thus, need to know the set of states in which the winning condition P is satisfied, because that is the region that the respective player wants to reach. That set of states in which P is true is denoted $[\![P]\!]$, which defines the semantics of dGL formula P. Recall that $\omega \in [\![P]\!]$ indicates that state ω is among the set of states in which P is true. The state ω in a hybrid game is still just a mapping that assigns real numbers to all variables, just as in hybrid programs, because that is what is needed to make sense of terms such as $x \cdot y + 2$ and formulas such as $x^2 \geq x \cdot y + 2$ in the hybrid game. A *state* ω is a mapping from variables to $\mathbb{R}$. The *set of states* is denoted $\mathscr{S}$.

15.2.1 Limits of Reachability Relations

The semantics of hybrid games is more subtle. The semantics of a hybrid program α is simply a reachability relation $[\![\alpha]\!] \subseteq \mathscr{S} \times \mathscr{S}$ where $(\omega, \nu) \in [\![\alpha]\!]$ indicates that final state ν is reachable from initial state ω by running HP α. That made it possible to define the semantics of the dL formula $\langle\alpha\rangle P$ via:

$$[\![\langle\alpha\rangle P]\!] = \{\omega \in \mathscr{S} : \nu \in [\![P]\!] \text{ for some } \nu \text{ with } (\omega, \nu) \in [\![\alpha]\!]\} \quad \text{for HP } \alpha \tag{15.1}$$

This approach does not suffice for hybrid games. First of all, the reachability relation $(\omega, \nu) \in [\![\alpha]\!]$ is only defined when α is a hybrid program, not when it is a hybrid game. And it is not even clear whether a reachability relation is all that it takes to understand the semantics of a hybrid game, because mere reachability information about states hardly retains enough information to understand the interactive aspects of gameplay that some choices are better than others for the respective players. But the deeper reason is that the shape (15.1) is too restrictive. Criteria of this shape would require Angel to single out a single state ω that satisfies the winning condition

$\nu \in [\![P]\!]$ and then get to that state ν by playing hybrid game α from ω. Yet, all that Demon then has to do to spoil this 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 state ν. More generally, winning into a single state is really difficult.

15.2.2 Set-valued Semantics of Differential Game Logic Formulas

Winning by leading the play into one of several states that satisfy the winning condition is more feasible. If we know the whole set of states $[\![P]\!]$ where postcondition P is true as the winning condition, then the hybrid game α uniquely determines the set of states from which Angel has a winning strategy in the game α to reach a state in $[\![P]\!]$. This *winning region* in hybrid game α for Angel's winning condition $[\![P]\!]$ will be denoted $\varsigma_\alpha([\![P]\!])$. More generally, for any set of states $X \subseteq \mathscr{S}$ will $\varsigma_\alpha(X)$ denote the set of states from which Angel has a winning strategy in the hybrid game α to reach a state in Angel's winning condition X. Accordingly, $\delta_\alpha(X)$ will denote the set of states from which Demon has a winning strategy in the hybrid game α to reach a state in Demon's winning condition X. Both sets will be defined in Sect. 15.2.3.

For a subset $X \subseteq \mathscr{S}$ the complement $\mathscr{S} \setminus X$ is denoted $X^\complement$. The notation ω_x^d from (2.8) on p. 46 still denotes the state that agrees with state ω except for the interpretation of variable x, which is changed to $d \in \mathbb{R}$. The value of term e in state ω is denoted by $\omega[\![e]\!]$ as in Definition 2.4. The denotational semantics of dGL formulas will be defined in Definition 15.1 by simultaneous induction along with the denotational semantics, $\varsigma_\alpha(\cdot)$ and $\delta_\alpha(\cdot)$, of hybrid games, defined in Definition 15.2, because dGL formulas are defined by simultaneous induction with hybrid games. The *(denotational) semantics of a hybrid game* α defines for each set of Angel's winning states $X \subseteq \mathscr{S}$ 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 15.1 (dGL semantics). The *semantics of a* dGL *formula* P is the subset $[\![P]\!] \subseteq \mathscr{S}$ of states in which P is true. It is defined inductively as follows:

1. $[\![e \geq \tilde{e}]\!] = \{\omega \in \mathscr{S} : \omega[\![e]\!] \geq \omega[\![\tilde{e}]\!]\}$
   That is, the set of states in which $e \geq \tilde{e}$ is true is the set in which the value of e is greater than or equal to the value of $\tilde{e}$.
2. $[\![\neg P]\!] = ([\![P]\!])^\complement$
   That is, the set of states in which $\neg P$ is true is the complement of the set of states in which P is true.
3. $[\![P \wedge Q]\!] = [\![P]\!] \cap [\![Q]\!]$
   That is, the set of states in which $P \wedge Q$ is true is the intersection of the set of states in which P is true with the set of states in which Q is true.

4. $[\![\exists x P]\!] = \{\omega \in \mathscr{S} : \omega_x^r \in [\![P]\!] \text{ for some } r \in \mathbb{R}\}$
   That is, the states in which $\exists x P$ is true are those which only differ in the real value of x from a state in which P is true.
5. $[\![\langle\alpha\rangle P]\!] = \varsigma_\alpha([\![P]\!])$
   That is, the set of states in which $\langle\alpha\rangle P$ is true is Angel's winning region to achieve $[\![P]\!]$ 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 P holds.
6. $[\![[\alpha]P]\!] = \delta_\alpha([\![P]\!])$
   That is, the set of states in which $[\alpha]P$ is true is Demon's winning region to achieve $[\![P]\!]$ 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 P holds.

A dGL formula P is *valid*, written $\vDash P$, iff it is true in all states, i.e. $[\![P]\!] = \mathscr{S}$.

The semantics $\varsigma_\alpha(X)$ and $\delta_\alpha(X)$ of Angel's and Demon's winning regions for winning condition X in hybrid game α will be defined next

15.2.3 Winning Region Semantics of Hybrid Games

Definition 15.1 uses 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 Chap. 14), the winning regions of hybrid games can be defined directly, giving a denotational semantics to hybrid games.[1] The winning regions for Angel are illustrated in Fig. 15.1, for Demon in Fig. 15.2.

Definition 15.2 (Semantics of hybrid games without repetition). The *semantics of a hybrid game* α is a function $\varsigma_\alpha(\cdot)$ that, for each set of Angel's winning states $X \subseteq \mathscr{S}$, 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. $\varsigma_{x:=e}(X) = \{\omega \in \mathscr{S} : \omega_x^{\omega[\![e]\!]} \in X\}$
   That is, an assignment $x:=e$ wins a game into X from any state whose modification $\omega_x^{\omega[\![e]\!]}$ that changes the value of x to that of $\omega[\![e]\!]$ is in X.
2. $\varsigma_{x'=f(x)\,\&\,Q}(X) = \{\varphi(0) \in \mathscr{S} : \varphi(r) \in X \text{ for some solution } \varphi : [0,r] \to \mathscr{S}$ of any duration $r \in \mathbb{R}$ satisfying $\varphi \models x' = f(x) \land Q\}$
   That is, Angel wins the differential equation $x' = f(x)\,\&\,Q$ into X from

[1] The semantics of a hybrid game is not merely a reachability relation between states as for hybrid systems [1], because the adversarial dynamic interactions and nested choices of the players have to be taken into account. For brevity, the informal explanations sometimes say "win the game" when really they mean "have a winning strategy to win the game". The semantics of differential equations could be augmented to ignore the initial value of the differential symbol x' as in Part II. This is not pursued for simplicity, because considering $x' := *; x' = f(x)\,\&\,Q$ has the same effect.

any state $\varphi(0)$ from which there is a solution φ of $x' = f(x)$ of any duration r that remains in Q all the time and leads to a final state $\varphi(r) \in X$.

3. $\varsigma_{?Q}(X) = [\![Q]\!] \cap X$
   That is, Angel wins into X for a challenge $?Q$ from the states which satisfy Q to pass the challenge and are already in X, because challenges $?Q$ 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 game 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. $\varsigma_{\alpha^{\mathrm{d}}}(X) = (\varsigma_\alpha(X^\complement))^\complement$
   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 $X^\complement$.

Since the players switch sides in a dual game α^{d}, Angel's winning region $\varsigma_{\alpha^{\mathrm{d}}}(X)$ from which she has a winning strategy to achieve X in the dual game α^{d} is the same as the complement $(\varsigma_\alpha(X^\complement))^\complement$ of the set $\varsigma_\alpha(X^\complement)$ where Angel would have a winning strategy in the game α to achieve the region $X^\complement$ where she loses the dual game α^{d}. The winning region $\varsigma_\alpha(X^\complement)$ corresponds to Angel pretending to play for Demon's controls in α^{d} by playing Angel's controls in α but for Demon's objective $X^\complement$ instead of Angel's objective X. The complement of this region then is the winning region $\varsigma_{\alpha^{\mathrm{d}}}(X)$ where Angel has a winning strategy in the dual game α^{d} to achieve X, because she would not have had a winning strategy to achieve $X^\complement$ when pretend-playing for Demon in game α.

After having defined the winning region $\varsigma_\alpha(X)$ from which Angel has a winning strategy to achieve X in the hybrid game α, the next question is how to define the winning region $\delta_\alpha(X)$ from which Demon has a winning strategy to achieve X in the hybrid game α. Together, these define the functions used in the semantics of dGL formulas (Definition 15.1). For discrete assignments $x := e$, the winning region $\varsigma_{x:=e}(X)$ for Angel is the same as the winning region $\delta_{x:=e}(X)$ for Demon in the same game with the same winning condition X, because there are no choices to resolve in a discrete assignment.

But for differential equations, the winning regions are very different, because Angel is in control of the duration of the differential equation, so Demon only has a chance if the differential equation starts in X (because Angel could follow it for duration 0 from the evolution domain) and stays in X all the time (because Angel could follow it for any other duration within the evolution domain). Likewise, since Angel gets to decide how to resolve a choice $\alpha \cup \beta$, Demon can only win if he wins both subgames.

Definition 15.3 (Semantics of hybrid games without repetition, 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. $\delta_{x:=e}(X) = \{\omega \in \mathscr{S} : \omega_x^{\omega[\![e]\!]} \in X\}$
   That is, an assignment $x:=e$ wins a game into X from any state whose modification $\omega_x^{\omega[\![e]\!]}$ that changes the value of x to that of $\omega[\![e]\!]$ is in X.
2. $\delta_{x'=f(x)\,\&\,Q}(X) = \{\varphi(0) \in \mathscr{S} : \varphi(r) \in X$ for all durations $r \in \mathbb{R}$ and all solutions $\varphi : [0,r] \to \mathscr{S}$ satisfying $\varphi \models x' = f(x) \wedge Q\}$
   That is, Demon wins the differential equation $x' = f(x)\,\&\,Q$ into X from any state $\varphi(0)$ from which all solutions φ of $x' = f(x)$ of any duration r that remain within Q all the time lead to states $\varphi(r) \in X$ in the end.
3. $\delta_{?Q}(X) = ([\![Q]\!])^\complement \cup X$
   That is, Demon wins into X for a challenge $?Q$ from the states which violate Q so that Angel fails her challenge $?Q$ or that are already in X, because challenges $?Q$ 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 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. $\delta_{\alpha^d}(X) = (\delta_\alpha(X^\complement))^\complement$
   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 $X^\complement$.

Strategies do not occur explicitly in the dGL semantics, because the semantics is based on the existence of winning strategies, not on the strategies themselves. Just as the semantics of 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. Likewise, the semantics of a compound hybrid game is a simple function of the semantics of its pieces. Also observe how the 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 on the goal X.

Even if we will only prove the following monotonicity property of winning regions in Chap. 16 after having defined a semantics of repetition, we already state it now, because it provides useful intuition. The semantics is monotone [2], i.e. larger sets of winning states have larger winning regions, because it is easier to win into larger sets of winning states (Fig. 15.3).

Lemma 15.1 (Monotonicity). *The* dGL *semantics is* monotone*, that is, both $\varsigma_\alpha(X) \subseteq \varsigma_\alpha(Y)$ and $\delta_\alpha(X) \subseteq \delta_\alpha(Y)$ for all $X \subseteq Y$.*

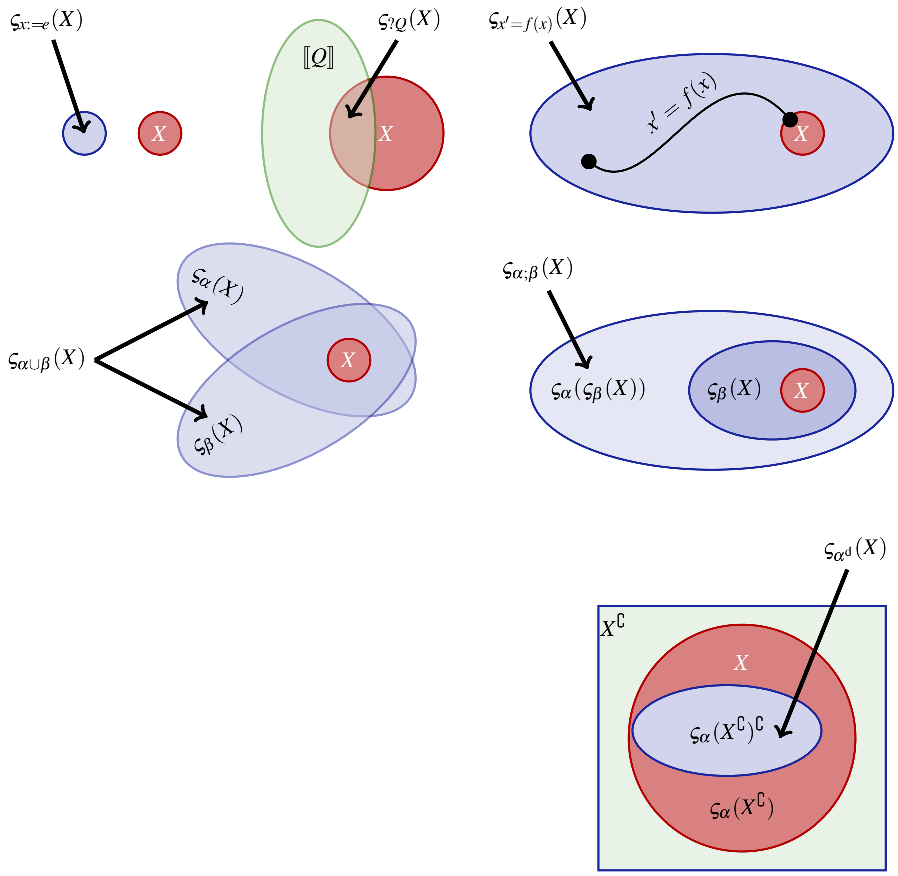

Fig. 15.1 Illustration of denotational semantics of hybrid games as Angel's winning regions

Note the big qualitative difference in the denotational semantics style of defining the winning region $\varsigma_\alpha(X)$ in Definition 15.2 compared to the operational semantics captured (informally) in Sect. 14.4. The denotational semantics directly associates with a hybrid game α and a winning condition X the set of states from which player Angel has a winning strategy in game α to achieve X. This results in a simple inductive definition of $\varsigma_\alpha(X)$ based on the structure of α. The game trees from the operational semantics in Sect. 14.4 give a more direct operational intuition of how a game can be played by following along the edges of its game graph. But rigorously defining the structure of these (infinite) graphs and what it means to have a winning strategy in it is technically more involved and complicates subsequent analysis of it, compared to the more convenient denotational semantics. Lemma 15.1 will be very easy to prove in Chap. 16. There are other circumstances where an operational semantics is more useful, so it is good to be familiar with both styles to choose the best fit for any question at hand.

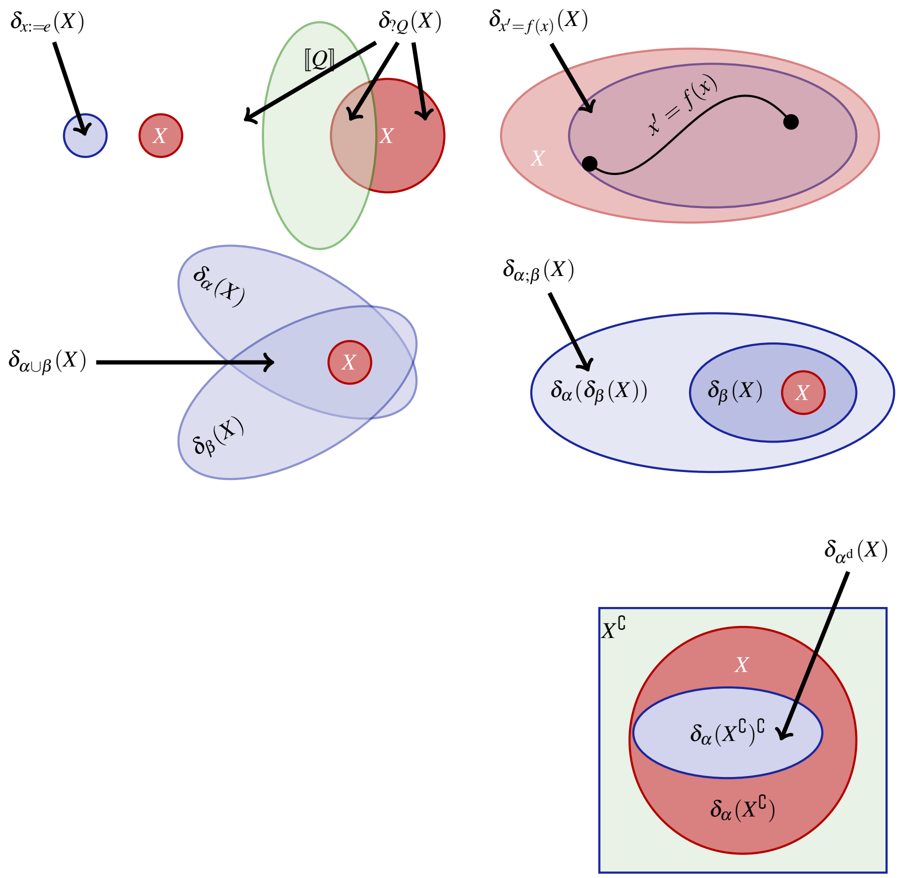

Fig. 15.2 Illustration of denotational semantics of hybrid games as Demon's winning regions

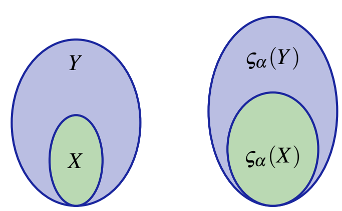

Fig. 15.3 Monotonicity: it is easier to win into larger sets of winning states $Y \supseteq X$

15.3 Semantics of Repetition in Hybrid Games

Before going any further we need to define a semantics for repetition, which will turn out to be surprisingly difficult. The final answer in Sect. 15.3.4 is not quite so complicated, but it takes considerable deliberation to get there. Since the insights along the way are of more general interest and illuminate interesting complexities with hybrid games nicely, we do not mind taking a careful route toward understanding the role of repetition in hybrid games.

15.3.1 Repetitions with Advance Notice

Definition 15.2 is still missing a definition for the semantics of repetition in hybrid games. With $\alpha^{n+1} \equiv \alpha^n;\alpha$ and $\alpha^0 \equiv ?true$, the semantics of repetition in hybrid systems was:

$$[\![\alpha^*]\!] = \bigcup_{n\in\mathbb{N}} [\![\alpha^n]\!]$$

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

$$\varsigma_{\alpha^*}(X) \stackrel{?}{=} \bigcup_{n<\omega} \varsigma_{\alpha^n}(X) \tag{15.2}$$

where ω is the first infinite ordinal (if you have never seen ordinals before, just read $n < \omega$ as n is in the 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 (15.2) 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* [3] 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 $\bigcup_{n<\omega} \varsigma_{\alpha^n}(X)$ and defines $\delta_{\alpha^*}(X)$ as $\bigcap_{n<\omega} \delta_{\alpha^n}(X)$. When playing α^*, Angel, thus, announces to Demon the number of repetitions $n < \omega$ when the game α^* starts and Demon announces the number of repetitions when the game $\alpha^\times$ starts. 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 we had meant with the operational game trees in Chap. 14.

If we gave repetition an advance notice semantics, then that would be a big disadvantage for the player controlling repetitions. The following formula, for example, is valid in dGL, but would not be valid in the advance notice semantics (Fig. 15.4):

$$x = 1 \land a = 1 \rightarrow \langle((x:=a;a:=0) \cap x:=0)^*\rangle\, x \neq 1 \tag{15.3}$$

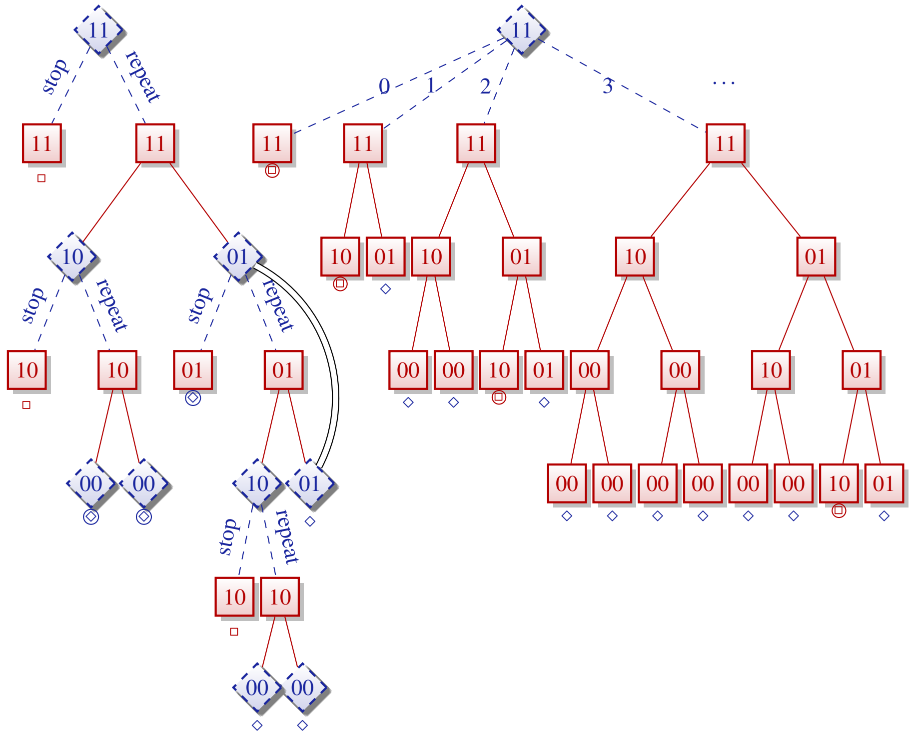

Fig. 15.4 Game trees for $x = 1 \land a = 1 \rightarrow \langle \alpha^* \rangle x \neq 1$ with game $\alpha \equiv (x := a; a := 0) \cap x := 0$ (notation: x, a). **(left)** valid in dGL by strategy "repeat once and repeat once more if $x = 1$, then stop" **(right)** false in advance notice semantics 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

The game starts with x and a both being 1 and asks whether Angel has a winning strategy to reach $x \neq 1$ with a repetition that she controls but Demon gets to choose whether 0 or a is put into x. The catch is that whenever the value of x is copied over to x in Demon's left choice, then a is zeroed out, so this only helps him once.

If, in the advance notice semantics, Angel announces when the repetition starts that she has chosen n repetitions of the game, then Demon wins with flying colors by choosing the right choice $x := 0$ option $n-1$ times followed by the left choice of $x := a; a := 0$ in the last repetition. This strategy would not work in the dGL semantics, where Angel is free to decide whether to repeat α^* *after* each repetition based on the resulting state of the game. Inspecting the state makes a big difference to decide whether to stop. If x has the value 0, then Angel decides to stop, otherwise she repeats. If Demon played the right choice $x := 0$, Angel stops. If he played the left choice $x := a; a := 0$, then Angel decides to repeat but will stop after the next iteration, regardless which option Demon chose. The winning strategy for (15.3) indicated with ⟐ in Fig. 15.4(left) shows that this dGL formula is valid.

Of course, there are also formulas that would be valid in the advance notice semantics but are not valid in dGL, for example, the dual of formula (15.3):

$$x=1 \land a=1 \to [((x:=a;a:=0) \cap x:=0)^*]\,x=1$$

Just like an advance notice semantics would make it easy for Demon to win α^* games with repetitions under Angel's control, it would also make it easy for Angel to win $\alpha^\times$ games with repetitions under Demon's control.

The advance notice semantics misses out on the existence of perfectly reasonable winning strategies, because it is just not interactive enough for proper hybrid game play. The dGL semantics is more general and gives the player in charge of repetition more control to inspect the state before having to decide on whether to repeat again. If you really need parts of a game where the number of repetitions are announced to the other player ahead of time, then it is easy to model them (Exercise 15.1).

Despite being built in direct analogy to the semantics of repetition in hybrid systems, the advance notice semantics is inappropriate for hybrid games, because it is very difficult for a CPS to predict ahead of time exactly how many iterations of a control cycle it will take to get to the goal.

For hybrid systems, it does not matter whether the number of iterations for a repetition is chosen ahead of time or afterwards, because there are no surprises during its evolution. All choices are consistently resolved by nondeterminism. This corresponds to all choices being resolved by Angel, which means she can always choose every choice in the best possible way. But for games, Demon can have a number of surprises in store for Angel, so that she will have to wait and see to decide how often to repeat.

15.3.2 Repetitions as Infinite Iterations

The trouble with the semantics in Sect. 15.3.1 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 how long the game will be played before he even has to do his first move.

Let's try to undo this. Instead of considering a choice over all n-fold repetitions α^n that reveals the chosen number n, we could consider a semantics that iterates n-times the winning region of α for any arbitrary finite number n (see Fig. 15.5):

$$\varsigma_{\alpha^*}(X) \overset{?}{=} \bigcup_{n<\omega} \varsigma_\alpha^n(X) \tag{15.4}$$

This semantics is called *ω-semantics* and also denoted $\varsigma_\alpha^\omega(X)$. All we need to do then is to define the iteration of the winning region construction. For any winning

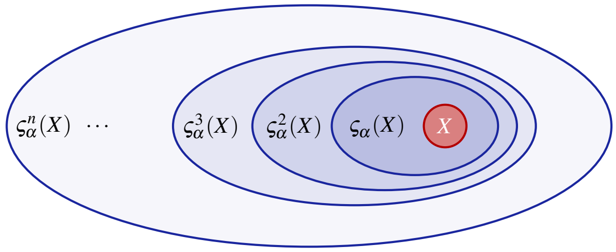

Fig. 15.5 Iteration $\varsigma_\alpha^n(X)$ of $\varsigma_\alpha(\cdot)$ from winning condition X.

condition $X \subseteq \mathscr{S}$, the n-times *iterated winning region* $\varsigma_\alpha^n(X)$ of α is defined by induction on n as:

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

The only states from which a repetition can win without actually repeating are the ones that start at the goal X already ($\varsigma_\alpha^0(X) = X$). The states from which a repetition can win into the set 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 $\varsigma_\alpha^\kappa(X)$. That is, the construction successively applies $\varsigma_\alpha(\cdot)$ while retaining the winning condition X:

$$\begin{aligned}
\varsigma_\alpha^0(X) &= X\\
\varsigma_\alpha^1(X) &= X \cup \varsigma_\alpha(X)\\
\varsigma_\alpha^2(X) &= X \cup \varsigma_\alpha(X \cup \varsigma_\alpha(X))\\
\varsigma_\alpha^3(X) &= X \cup \varsigma_\alpha(X \cup \varsigma_\alpha(X \cup \varsigma_\alpha(X)))\\
\varsigma_\alpha^4(X) &= X \cup \varsigma_\alpha(X \cup \varsigma_\alpha(X \cup \varsigma_\alpha(X \cup \varsigma_\alpha(X))))\\
&\vdots
\end{aligned}$$

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?

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.

Would the following dGL formula be valid with the semantics from (15.4)?

$$\langle(\underbrace{\underbrace{x:=1;x'=1^{\mathrm{d}}}_{\beta}\cup\underbrace{x:=x-1}_{\gamma}}_{\alpha})^*\rangle\,(0\le x<1) \tag{15.5}$$

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

As usual, $[a,b)$ denotes the interval from a inclusive to b exclusive. Using the abbreviations indicated in (15.12) such as $\alpha \equiv \beta\cup\gamma$, it is easy to see that $\varsigma^n_\alpha([0,1)) = [0,n+1)$ for all $n\in\mathbb{N}$ by a simple inductive proof:

$$\varsigma^0_{\beta\cup\gamma}([0,1)) = [0,1)$$
$$\varsigma^{n+1}_{\beta\cup\gamma}([0,1)) = [0,1)\cup\varsigma_{\beta\cup\gamma}(\varsigma^n_{\beta\cup\gamma}([0,1))) \overset{\text{IH}}{=} [0,1)\cup\varsigma_{\beta\cup\gamma}([0,n+1))$$
$$= [0,1)\cup\varsigma_\beta([0,n+1))\cup\varsigma_\gamma([0,n)) = [0,1)\cup\emptyset\cup[1,n+2) = [0,n+1+1)$$

Consequently, the ω-semantics from (15.4) consists of all nonnegative reals:

$$\bigcup_{n<\omega}\varsigma^n_\alpha([0,1)) = \bigcup_{n<\omega}[0,n+1) = [0,\infty) \tag{15.6}$$

Hence, the $\omega-$semantics from (15.4) would indicate that the hybrid game (15.12) can be won only from initial states in $[0,\infty)$, that is, for initial states that satisfy $0\le x$.

Unfortunately, this is complete nonsense. True, the hybrid game in dGL formula (15.12) can be won from all initial states that satisfy $0\le x$. But it can also be won from many other initial states! There are cases, where the ω-semantics is minuscule compared to the true winning region and arbitrarily far away from the truth [2].

For the formula (15.12), the ω-semantics misses out on Angel's perfectly reasonable winning strategy "first choose $x:=1;x'=1^{\mathrm{d}}$ and then always choose $x:=x-1$ until stopping at $0\le x<1$". This winning strategy wins from every initial state in $\mathbb{R}$, which is a much bigger set than the set of nonnegative reals from (15.6).

This winning strategy justifies that the dGL formula (15.12) is valid. Yet, is there a direct way to see that (15.6) is not the final answer for (15.12) without having to put the winning region computations aside and constructing a separate ingenious winning strategy, which would undermine the whole point of using winning regions for the semantics?

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

The crucial observation comes from a closer inspection of what exactly we did to arrive at (15.6). The fact (15.6) shows that the hybrid game in (15.12) can be won from all nonnegative initial values with at most ω (that is "first countably infinitely many") steps. The induction step proving $\varsigma^n_\alpha([0,1)) = [0,n)$ for all $n\in\mathbb{N}$ 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 71 (+1 argument) *Whenever a set Z is in the winning region $\varsigma_{\alpha^*}(X)$ of repetition, then $\varsigma_\alpha(Z)$ is in the winning region $\varsigma_{\alpha^*}(X)$ as well, because it is just one round away from Z and α^* could simply repeat once more. That is:*

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

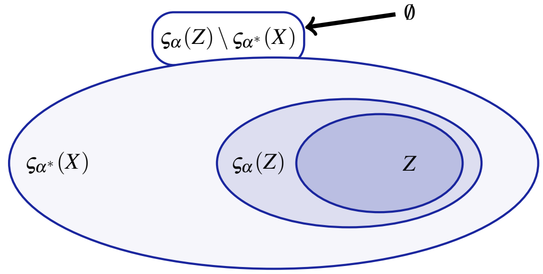

Fig. 15.6 Winning regions $\varsigma_\alpha(Z)$ of sets $Z \subseteq \varsigma_{\alpha^*}(X)$ are already included in $\varsigma_{\alpha^*}(X)$ since $\varsigma_\alpha(Z)$ is just one more round away from Z

Applying Note 71, which is illustrated in Fig. 15.6, to the situation at hand works as follows. The fact (15.6) explains that at least $[0,\infty) \subseteq \varsigma_{(\beta\cup\gamma)^*}([0,1))$ is in the winning region of repetition. By Note 71, the winning region $\varsigma_{(\beta\cup\gamma)^*}([0,1))$ also contains the one-step winning region $\varsigma_{\beta\cup\gamma}([0,\infty)) \subseteq \varsigma_{(\beta\cup\gamma)^*}([0,1))$ of $[0,\infty)$. 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 (since there is only one variable in this hybrid game) and it is kind of hard to add anything to that. 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 (15.12) 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. It took us more than infinitely many steps to find the winning region of the hybrid game described in (15.12). After infinitely many iterations to arrive at $\varsigma_\alpha^\omega([0,1)) = \bigcup_{n<\omega} \varsigma_\alpha^n([0,1)) = [0,\infty)$, 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. 15.7 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 than just a little more than first countable infinitely many [2].

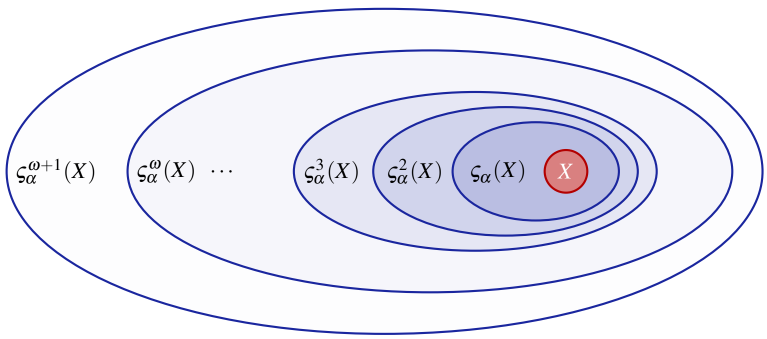

Fig. 15.7 Iteration $\varsigma_\alpha^{\omega+1}(X)$ of $\varsigma_\alpha(\cdot)$ from winning condition $X=[0,1)$ stops when applying $\varsigma_\alpha(\cdot)$ to the ωth infinite iteration $\varsigma_\alpha^{\omega}(X)$.

The existence of the above winning strategy is found at level $\varsigma_\alpha^{\omega+1}([0,1)) = \mathbb{R}$. 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. Angel cannot even 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 for (15.12).

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

15.3.3 Inflationary Semantics of Repetition

Despite the quite discouraging fact that infinitely many iterations of the winning region construction $\varsigma_\alpha(\cdot)$ do not suffice to accurately describe the winning region of repetition α^*, there still is a way of rescuing the situation if we simply keep iterating. We just need to repeat the construction more than infinitely often, leading us into the wonderful world of ordinals. Even if we will ultimately discard this higher iteration semantics with ordinals in favor of a simpler semantics of repetition in Sect. 15.3.4, we still learn interesting subtle nuances about hybrid games by pursuing more iteration at first.

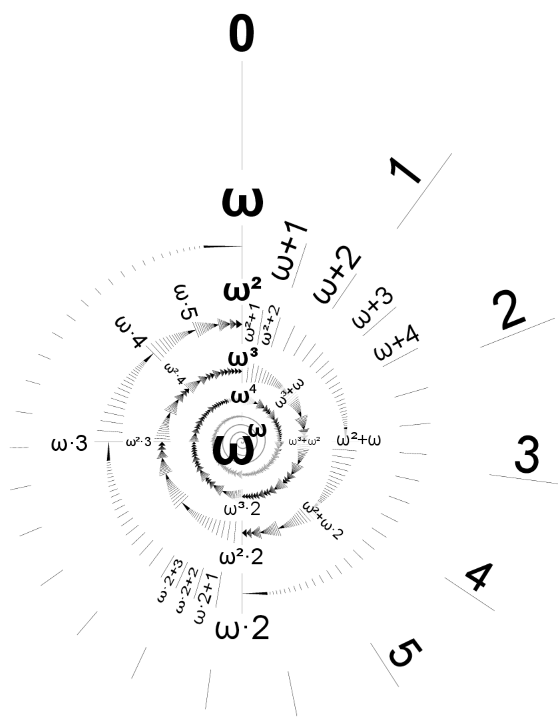

Fig. 15.8 Illustration of infinitely many ordinals up to ω^ω, including $0 < 1 < 2 < \cdots < \omega < \omega+1 < \cdots < \omega\cdot 2 < \omega\cdot 2+1 < \cdots < \omega^2 < \omega^2+1 < \cdots < \omega^2+\omega < \omega^2+\omega+1 < \cdots$

The key to understanding *ordinals* is that each ordinal κ always has a *successor ordinal* $\kappa+1$ but every set of ordinals also has a least upper bound λ, called *limit ordinal* if it is not already a successor ordinal. For example, ω is the first infinite ordinal, and the smallest ordinal that is bigger than all natural numbers. But ω also has a successor ordinal $\omega+1$, which, in turn, has a successor ordinal $\omega+2$, and all those have a least upper bound $\omega\cdot 2$ (Fig. 15.8). All ordinals starting at ω are called *transfinite ordinals*, because there are infinitely many smaller ordinals.

When we apply the winning region construction $\varsigma_\alpha(\cdot)$ for each successor ordinal $\kappa+1$, but take the union of all previous winning regions at limit ordinals λ such as ω, the semantics of repetition can be defined using transfinite iteration (Fig. 15.9):

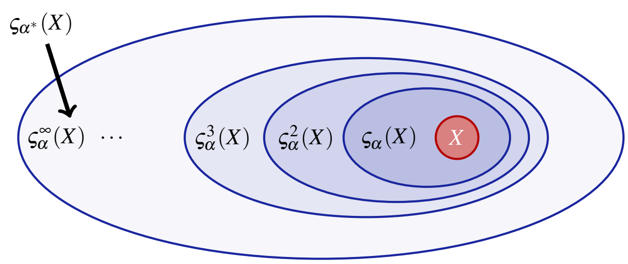

Fig. 15.9 Transfinite iteration $\varsigma_\alpha^\infty(X)$ of $\varsigma_\alpha(\cdot)$ from winning condition X results in winning region $\varsigma_{\alpha^*}(X)$ of repetition.

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

The semantics of repetition is the union of all winning regions for all ordinals:

$$
\varsigma_{\alpha^*}(X) = \varsigma_\alpha^\infty(X) \stackrel{\text{def}}{=} \bigcup_{\kappa \text{ ordinal}} \varsigma_\alpha^\kappa(X) \tag{15.7}
$$

Note 72 (Infinite iterations infinitely often) *Unfortunately, hybrid games require rather big infinite ordinals until this inflationary style of computing their winning regions stops [2]. That translates into an infinite amount of work and then some more, infinitely often, to compute the winning region. Hardly the sort of thing we would like to wait for until we finally know who wins a game.*

This semantics for repetition from (15.7) provides the correct answer if we do not mind the highly transfinite number of iterations it needs. Unfortunately, even the pretty infinite ordinal ω^ω is not enough for all hybrid games [2, Theorem 3.8].

Expedition 15.1 (Ordinal numbers)

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. 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[a] infinite ordinal: $0 < 1 < 2 < 3 < \cdots < \omega$ 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.

Ordinals are keen on ensuring that every ordinal has a successor and for every set of ordinals there is a bigger ordinal. So, ω must have a successor, which is the successor ordinal $\omega+1$, the successor of which is $\omega+2$ etc.:

$$0 < 1 < 2 < 3 < 4 < \cdots < \omega < \omega+1 < \omega+2 < \omega+3 < \omega+4 < \cdots$$

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 $\omega\cdot 2+1$:

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

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 still have a limit ordinal that's bigger than all of them, which is $\omega\cdot\omega = \omega^2$. That one again has a successor ω^2+1 and so on (Fig. 15.8):

$$0<1<2<\cdots\omega<\omega+1<\omega+2<\cdots\omega\cdot 2<\omega\cdot 2+1<\cdots\omega\cdot 3<\omega\cdot 3+1<\cdots$$
$$\omega^2<\omega^2+1<\cdots\omega^2+\omega<\omega^2+\omega+1<\cdots\omega^\omega<\cdots\omega^{\omega^\omega}<\cdots\omega_1^{\mathrm{CK}}<\cdots\omega_1\cdots$$

The first infinite ordinal is ω, the Church-Kleene ordinal ω_1^{CK} is 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.

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

The semantics we pursue in Sect. 15.3.4 is much easier in this respect, it just does not provide the same insights about the iterative complexities of hybrid games.

With this refined understanding of iteration, look back at dGL formula (15.12) 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 (15.12). 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 this, but suppose she were to brag to impress Demon by saying she could win (15.12) within $n \in \mathbb{N}$ repetitions, then it would be impossible for her to keep that promise. No matter how big a bound $n \in \mathbb{N}$ she were to choose, Demon could still always spoil it from any negative initial state by evolving his differential equation $x' = 1^{\mathrm{d}}$ for much longer than n time units so that it takes Angel more than n rounds to decrease the resulting value down to the interval $[0,1)$ again.

This illustrates the dual of the discussion on the advance notice semantics in Sect. 15.3.1, which showed that Demon could make Angel win faster than she announced just to make her lose in the final round. In (15.12), 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 (15.12). 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.

15.3.4 Characterizing Winning Repetitions Implicitly

Section 15.3.3 culminated in a semantics of repetition defined as the union of all winning regions for all ordinals by an explicit (albeit wildly infinite) construction (15.7). 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 iterated winning region construction (15.7) describes the semantics of repetition by iterating from below, i.e. starting from $\varsigma^0_\alpha(X) = X$ and adding states. Could the semantics of repetition be characterized more indirectly but more concisely from above? With an implicit characterization instead of an explicit construction.

The +1 argument (Note 71) implies $\varsigma_\alpha(Z) \subseteq \varsigma_{\alpha^*}(X)$ for any set $Z \subseteq \varsigma_{\alpha^*}(X)$. In particular, the set $Z \stackrel{\text{def}}{=} \varsigma_{\alpha^*}(X)$ itself satisfies:

$$\varsigma_\alpha(\varsigma_{\alpha^*}(X)) \subseteq \varsigma_{\alpha^*}(X) \tag{15.8}$$

Expedition 15.2 (Ordinal arithmetic)

Ordinals support addition, multiplication, and exponentiation, which can be defined by induction on its second argument quite similar to how they are defined for natural numbers. The only oddity is that these operations are noncommutative. The constructions distinguish the case of successor ordinals that are direct successors of a smaller compared to limit ordinals which are the least upper bounds over all smaller ordinals:

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

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 $\omega \cdot 2 < \omega \cdot 4$.

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 just been repeated one more time itself. Consequently, if a set $Z \subseteq \mathscr{S}$ claims to be the winning region $\varsigma_{\alpha^*}(X)$ of repetition, it at least has to satisfy

$$\varsigma_\alpha(Z) \subseteq Z \tag{15.9}$$

because, by (15.8), the true winning region $\varsigma_{\alpha^*}(X)$ does satisfy (15.9). Thus, strategizing 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 qualify for being the winning region $\varsigma_{\alpha^*}(X)$ of repetition? Is there only one choice for Z? 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 $Z \stackrel{\text{def}}{=} \emptyset$ looks like it would satisfy (15.9) because it is rather hard to win a game that requires Angel to enter the empty set of states $\emptyset$ 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. When Q is a nontrivial formula like $x > 0$ Demon fails $?Q^{d}$ sometimes:

$$\varsigma_{?Q^{d}}(\emptyset) = (\varsigma_{?Q}(\emptyset^{\complement}))^{\complement} = ([\![Q]\!] \cap \mathscr{S})^{\complement} = ([\![Q]\!])^{\complement} = [\![\neg Q]\!] \not\subseteq \emptyset$$

Yet, then the set of states $[\![\neg Q]\!]$ that make Demon violate the rules satisfies (15.9):

$$\varsigma_{?Q^{d}}([\![\neg Q]\!]) = (\varsigma_{?Q}([\![\neg Q]\!]^{\complement}))^{\complement} = (\varsigma_{?Q}([\![Q]\!]))^{\complement} = ([\![Q]\!] \cap [\![Q]\!])^{\complement} = [\![\neg Q]\!] \subseteq [\![\neg Q]\!]$$

But even in cases where the empty set $\emptyset$ satisfies (15.9), it may be too small. Likewise, even if the set of states where Demon violates the rules immediately satisfies (15.9), this set may still be too 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 also contain the winning condition X, because the winning condition X is particularly easy to reach when already starting in X because Angel could then simply decide to stop fooling around and just repeat zero times. Consequently, if a set $Z \subseteq \mathscr{S}$ claims to be the winning region $\varsigma_{\alpha^*}(X)$, then it has to satisfy (15.9) and also satisfy:

$$X \subseteq Z \tag{15.10}$$

Both conditions (15.9), (15.10) together can be summarized in one condition:

Note 73 (Pre-fixpoint) *Every candidate Z for the winning region $\varsigma_{\alpha^*}(X)$ satisfies the* pre-fixpoint *condition:*

$$X \cup \varsigma_\alpha(Z) \subseteq Z \tag{15.11}$$

Again: what is this set Z that satisfies (15.11)? 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 $Z \stackrel{\text{def}}{=} \mathscr{S}$ always satisfies (15.11) trivially so (15.11) definitely 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

$$\langle (x := 1; x' = 1^{d} \cup x := x - 1)^{*} \rangle \, (0 \leq x < 1) \tag{15.12}$$

the full state space is hardly the right winning region for any arbitrary hybrid game α^*. It definitely depends on the hybrid game α and the winning condition P whether Angel has a winning strategy for $\langle\alpha\rangle P$ or not. For example for Demon's favorite game where he always wins, the winning region $\varsigma_{\alpha^*}(X)$ had better be $\emptyset$, not $\mathscr{S}$. Thus, the largest solution Z of (15.11) hardly qualifies.

So which solution Z of (15.11) 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 (15.11), the largest one is not informative, because the largest Z simply degrades to the full state space $\mathscr{S}$. So smaller solutions Z are preferable. Which one? How do multiple solutions even relate to each other? Suppose Y,Z are both solutions of (15.11). That is

$$X \cup \varsigma_\alpha(Y) \subseteq Y \tag{15.13}$$

$$X \cup \varsigma_\alpha(Z) \subseteq Z \tag{15.14}$$

Then, by the monotonicity lemma Lemma 15.1:

$$X \cup \varsigma_\alpha(Y \cap Z) \overset{\text{mon}}{\subseteq} X \cup (\varsigma_\alpha(Y) \cap \varsigma_\alpha(Z)) \overset{(15.13),(15.14)}{\subseteq} Y \cap Z \tag{15.15}$$

Hence, by (15.15), the intersection $Y \cap Z$ of solutions Y and Z of (15.11) also is a solution of (15.11):

Lemma 15.2 (Intersection closure). *For any two solutions Y,Z of the prefix condition* (15.11)*, the intersection $Y \cap Z$ is a solution of* (15.11) *as well.*

Whenever there are two solutions Z_1, Z_2 of (15.11), their intersection $Z_1 \cap Z_2$ solves (15.11) as well. When there's yet another solution Z_3 of (15.11), their intersection $Z_1 \cap Z_2 \cap Z_3$ also solves (15.11). Similarly the intersection of any larger family of solutions solves (15.11). 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 (15.11) which is $\varsigma_{\alpha^*}(X)$.

Note 74 (Semantics of repetitions) *Among the many Z that solve* (15.11)*, $\varsigma_{\alpha^*}(X)$ is defined to be the smallest Z that solves prefix condition* (15.11)*:*

$$\varsigma_{\alpha^*}(X) = \bigcap \{Z \subseteq \mathscr{S} : X \cup \varsigma_\alpha(Z) \subseteq Z\} \tag{15.16}$$

In other words, the winning region $\varsigma_{\alpha^*}(X)$ is the smallest set Z that already contains the winning condition X and the set of states $\varsigma_\alpha(Z)$ from which Angel can win into Z with one more round of game α. Hence, adding to Z the set of states $\varsigma_\alpha(Z)$ where one more round would win does not change the set Z, as illustrated in Fig. 15.10.

The fact that $\varsigma_{\alpha^*}(X)$ is defined as the *smallest* of all these sets makes sure that Angel only wins games by a well-founded number of repetitions. That is, she only

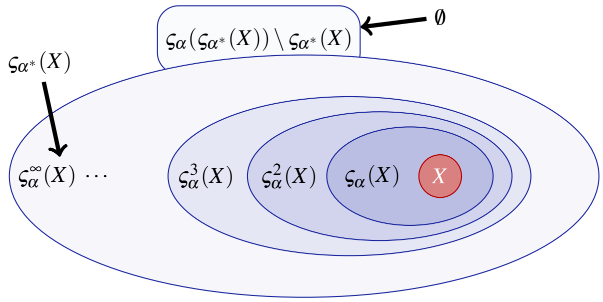

Fig. 15.10 Illustration of denotational semantics of winning region of hybrid game repetitions

wins a repetition if she ultimately stops repeating, not by postponing termination forever with a filibuster [2].

The characterization in terms of iterated winning regions from Sect. 15.3.3 leads to the same set $\varsigma_{\alpha^*}(X)$, but the (least pre-fixpoint or) least fixpoint characterization (15.16) is easier to describe and reason with. Understanding why the two styles of defining the semantics lead to the same result will take some thoughts.

The set on the right-hand side of (15.16) is an intersection of solutions, thus, a solution by Lemma 15.2 (or its counterpart for arbitrary families of solutions). Hence $\varsigma_{\alpha^*}(X)$ itself satisfies the prefix condition (15.11):

$$X \cup \varsigma_\alpha(\varsigma_{\alpha^*}(X)) \subseteq \varsigma_{\alpha^*}(X) \tag{15.17}$$

Also compare this with where we came from when we argued for (15.8). Could it be the case that the inclusion in (15.17) is strict, i.e. not equals? No this cannot happen, because $\varsigma_{\alpha^*}(X)$ is the smallest such set. That is, by (15.17), the set $Z \overset{\text{def}}{=} X \cup \varsigma_\alpha(\varsigma_{\alpha^*}(X))$ satisfies $Z \subseteq \varsigma_{\alpha^*}(X)$ and, thus, by Lemma 15.1:

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

Thus, the set $Z \overset{\text{def}}{=} X \cup \varsigma_\alpha(\varsigma_{\alpha^*}(X))$ satisfies the condition $X \cup \varsigma_\alpha(Z) \subseteq Z$ from (15.16). Since $\varsigma_{\alpha^*}(X)$ is the smallest such set by (15.16), it is smaller or equal Z:

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

Consequently, with (15.17), this implies that both inclusions hold, so $\varsigma_{\alpha^*}(X) = Z$. Thus, the least pre-fixpoint $\varsigma_{\alpha^*}(X)$ satisfies not just the pre-fixpoint inclusion (15.11) but it even satisfies the fixpoint equation:

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

Note 75 (Semantics of repetitions, fixpoint formulation) *The winning region* $\varsigma_{\alpha^*}(X)$ *of repetition is a* fixpoint *solving the equation*

$$X \cup \varsigma_\alpha(Z) = Z \tag{15.18}$$

It is the least fixpoint*, i.e. the smallest set* Z *solving the equation* (15.18)*. That is, it satisfies*

$$\varsigma_{\alpha^*}(X) = \bigcap\{Z \subseteq \mathscr{S} : X \cup \varsigma_\alpha(Z) = Z\} \tag{15.19}$$

Due to the seminal fixpoint theorem of Knaster-Tarski [4], the least fixpoint semantics $\varsigma_{\alpha^*}(X)$ from (15.16) and its equivalent reformulation (15.19), gives the same set of states as the inflationary semantics (15.7), because the semantics is monotone (Lemma 15.1). That is, after iterating the winning region construction in Sect. 15.3.3 for all (large enough) ordinals starting from X, the result will be the least fixpoint of (15.18).

Lemma 15.3 (Transfinite inflation leads to a least fixpoint).

$$\varsigma_{\alpha^*}(X) \stackrel{\text{def}}{=} \bigcap\{Z \subseteq \mathscr{S} : X \cup \varsigma_\alpha(Z) \subseteq Z\} \;=\; \varsigma_\alpha^\infty(X) \stackrel{\text{def}}{=} \bigcup_{\kappa \text{ ordinal}} \varsigma_\alpha^\kappa(X)$$

But the iterated winning region constructions go significantly transfinite [2], way beyond the first infinite ordinal ω.

The situation for Demon's winning region for repetition is correspondingly. The difference is that Angel controls repetition α^*, so Demon only has a winning strategy to achieve X if he starts in X (because Angel could repeat 0 times) and has a winning strategy to stay in X all the time. Postponing termination forever will make Demon win if only he stays in X, because Angel is in charge of repetition and will ultimately have to stop repeating. Consequently, the winning region for Demon for Angel's repetition is the largest fixpoint.

Note 76 (Demon's winning region for repetition)

$$\delta_{\alpha^*}(X) = \bigcup\{Z \subseteq \mathscr{S} : X \cap \delta_\alpha(Z) = Z\} = \bigcup\{Z \subseteq \mathscr{S} : Z \subseteq X \cap \delta_\alpha(Z)\}$$

The winning region $\delta_{\alpha^*}(X)$ is the largest set Z that is contained in the winning condition X and in the set of states $\delta_\alpha(Z)$ where Demon has a winning strategy to remain in Z for one more round of game α. This set is the largest fixpoint, because Demon would not mind repeating indefinitely, since he knows that Angel will ultimately have to stop repeating at some point anyhow. He only needs to make sure not to have left the winning condition X, because he cannot know how often Angel will repeat.

15.4 Semantics of Hybrid Games

The semantics of hybrid games from Sect. 15.2.3 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 an operational game tree semantics (as in Chap. 14), or in terms of transfinitely iterated winning region constructions (Sect. 15.3.3), the winning regions of hybrid games can be defined directly (Sect. 15.3.4), giving a denotational semantics to hybrid games.

The only difference of the following semantics compared to the previous Definition 15.2 is the new case of repetition α^* illustrated in Fig. 15.10.

Definition 15.4 (Semantics of hybrid games). The *semantics of a hybrid game* α is a function $\varsigma_\alpha(\cdot)$ that, for each set of Angel's winning states $X \subseteq \mathscr{S}$, 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. $\varsigma_{x:=e}(X) = \{\omega \in \mathscr{S} : \omega_x^{\omega[\![e]\!]} \in X\}$
   That is, an assignment $x:=e$ wins a game into X from any state whose modification $\omega_x^{\omega[\![e]\!]}$ that changes the value of x to that of $\omega[\![e]\!]$ is in X.
2. $\varsigma_{x'=f(x)\,\&\,Q}(X) = \{\varphi(0) \in \mathscr{S} : \varphi(r) \in X$ for some solution $\varphi : [0,r] \to \mathscr{S}$ of any duration $r \in \mathbb{R}$ satisfying $\varphi \models x' = f(x) \land Q\}$
   That is, Angel wins the differential equation $x' = f(x)\,\&\,Q$ into X from any state $\varphi(0)$ from which there is a solution φ of $x' = f(x)$ of any duration r that remains in Q all the time and leads to a final state $\varphi(r) \in X$.
3. $\varsigma_{?Q}(X) = [\![Q]\!] \cap X$
   That is, Angel wins into X for a challenge $?Q$ from the states which satisfy Q to pass the challenge and are already in X, because challenges $?Q$ 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. $\varsigma_{\alpha^*}(X) = \bigcap\{Z \subseteq \mathscr{S} : X \cup \varsigma_\alpha(Z) \subseteq Z\}$
   That is, Angel wins a game of repetition α^* into X from the smallest set of states Z that includes both X and the set of states $\varsigma_\alpha(Z)$ from which Angel can achieve Z in one more round of game α.
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 $X^\complement$.

Definition 15.5 (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:

1. $\delta_{x:=e}(X) = \{\omega \in \mathscr{S} : \omega_x^{\omega[\![e]\!]} \in X\}$
   That is, an assignment $x := e$ wins a game into X from any state whose modification $\omega_x^{\omega[\![e]\!]}$ that changes the value of x to that of $\omega[\![e]\!]$ is in X.
2. $\delta_{x'=f(x)\,\&\,Q}(X) = \{\varphi(0) \in \mathscr{S} : \varphi(r) \in X$ for all durations $r \in \mathbb{R}$ and all solutions $\varphi : [0,r] \to \mathscr{S}$ satisfying $\varphi \models x' = f(x) \land Q\}$
   That is, Demon wins the differential equation $x' = f(x)\,\&\,Q$ into X from any state $\varphi(0)$ from which all solutions φ of $x' = f(x)$ of any duration r that remain within Q all the time lead to states $\varphi(r) \in X$ in the end.
3. $\delta_{?Q}(X) = ([\![Q]\!])^\complement \cup X$
   That is, Demon wins into X for a challenge $?Q$ from the states which violate Q so that Angel fails her challenge $?Q$ or that are already in X, because challenges $?Q$ 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 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. $\delta_{\alpha^*}(X) = \bigcup\{Z \subseteq \mathscr{S} : Z \subseteq X \cap \delta_\alpha(Z)\}$
   That is, Demon wins a game of repetition α^* into X from the biggest set of states Z that is included both in X and in the set of states $\delta_\alpha(Z)$ from which Demon can achieve Z in one more round of game α.
7. $\delta_{\alpha^d}(X) = (\delta_\alpha(X^\complement))^\complement$
   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 $X^\complement$.

The semantics of dGL is still *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.

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 Angel's victory forever, because Angel ultimately has to stop repeating. Thus, for the formula $[\alpha^*]P$, Demon already has a winning strategy if he only has a strategy that is not losing by preventing P indefinitely, because Angel eventually has to stop repeating anyhow and will then end up in a state not satisfying P, which makes her lose.

For Demon's repetition $[\alpha^\times]P$ the situation is dual, so Demon will ultimately have to stop repeating and get to state P in finite time. But Angel is happy to postpone Demon's victory forever, because Demon will eventually have to stop since he is in charge of Demon's repetition $\alpha^\times$.

15.5 Summary

This chapter 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 required higher-ordinal iterations of winning region constructions. This led to an insightful appreciation for the complexities, challenges, and flexibilities of hybrid games. But the final word on the semantics of repetition was a simpler implicit characterization. The next chapter will leverage the semantic basis for the next leg in the logical trinity: axiomatics. That will enable us to succinctly reason about hybrid games and whether our player of interest has a winning strategy.

The concepts that we touched upon in this chapter are of independent interest. Fixpoints play a huge role in many areas of science. Ordinals are also of more general interest. Differences between operational and denotational styles of giving a semantics are also more broadly impactful.

Exercises

15.1 (Modeling advance notice semantics). The advance notice semantics from Sect. 15.3.1 was discarded in favor of the more general semantics of repetition in Sect. 15.3.4 that allows the player controlling repetition to decided based on observing the state. Suppose you have a game where you want to allow Angel to repeat α any (finite) number of times but you require that she announces the number of repetitions of α ahead of time, just like in the advance notice semantics. Construct a hybrid game that requires Angel to disclose the intended number of repetitions of α to Demon ahead of time even in the semantics of Definition 17.1. Hint: you can use additional variables.

15.2. The formula (15.12) was shown to need $\omega+1$ iterations of the winning region construction to terminate with the following answer justifying the validity of (15.12).

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

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

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

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

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

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

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

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

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

15.6. Can you find dGL formulas for which the winning region construction takes even longer to terminate? How far can you push this?

References

1. Platzer, A. *The Complete Proof Theory of Hybrid Systems* in *LICS* (IEEE, 2012), 541–550. doi:10.1109/LICS.2012.64.
2. Platzer, A. Differential Game Logic. *ACM Trans. Comput. Log.* **17,** 1:1–1:51 (2015).
3. Quesel, J.-D. & Platzer, A. *Playing Hybrid Games with KeYmaera.* in *IJCAR* (eds Gramlich, B., Miller, D. & Sattler, U.) **7364** (Springer, 2012), 439–453. doi:10.1007/978-3-642-31365-3_34.
4. Tarski, A. A lattice-theoretical fixpoint theorem and its applications. *Pacific J. Math.* **5,** 285–309 (1955).