A Complete Axiomatization of Differential Game Logic for Hybrid Games

André Platzer

January 2013
CMU-CS-13-100

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

This report is an updated version superseding the earlier report CMU-CS-12-105 [Pla12b]. It is based on that earlier report yet also contains new interesting results (especially completeness).

This material is based upon work supported by the National Science Foundation under NSF CAREER Award CNS-1054246. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution or government. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of any sponsoring institution or government.

Keywords: game logic; hybrid dynamical systems; hybrid games; axiomatization

Abstract

We introduce *differential game logic* ($\mathsf{dG}\mathcal{L}$) for specifying and verifying properties of *hybrid games*, i.e. games on hybrid systems combining discrete and continuous dynamics. Unlike hybrid systems, hybrid games allow choices in the system dynamics to be resolved adversarially by different players with different objectives. The logic $\mathsf{dG}\mathcal{L}$ can be used to study the existence of winning strategies for such hybrid games. We present a simple sound and complete axiomatization of $\mathsf{dG}\mathcal{L}$ relative to the fixpoint logic of differential equations. We prove hybrid games to be determined and their winning regions to require higher closure ordinals and we identify separating axioms, i.e. axioms that distinguish hybrid games from hybrid systems.

1 Introduction

Hybrid systems [Hen96] are dynamical systems combining discrete dynamics and continuous dynamics, which are important, e.g., for modeling how computers control physical systems. Hybrid systems combine difference equations and differential equations with conditional switching, nondeterminism, and repetition. Hybrid systems are not semidecidable [Hen96], but nevertheless studied by many successful verification approaches. They have a complete axiomatization relative to differential equations in *differential dynamic logic* ($\mathsf{d}\mathcal{L}$) [Pla08, Pla12a], which extends Pratt's dynamic logic of conventional discrete programs [Pra76] to hybrid systems by adding differential equations and a reachability relation semantics on the real Euclidean space.

We consider *hybrid games* [TPS98, TLS00, BBC10, VPVD11], i.e. games of two players on a hybrid system, which have found a number of interesting applications [TPS98, TLS00, BBC10, PHP01, VPVD11, QP12]. Hybrid games extend hybrid systems by adding an adversarial resolution of the choices in the system dynamics. We obtain hybrid games from hybrid systems simply by adding the dual operator ^{d} for passing control between the players. Hybrid games without ^{d} are single player hybrid games, i.e. hybrid systems, because control never passes to the other player. Hybrid games using ^{d} give both players control over their respective choices (as indicated by ^{d}). They can play in reaction to the outcome that the previous choices by the players have had on the state of the system.

One of the most fundamental questions about a hybrid game is whether a player has a *winning strategy*[1], i.e. a way to resolve its choices that will lead to a state in which that player wins, no matter how the other player resolves his choices. We introduce *differential game logic* ($\mathsf{dG}\mathcal{L}$) [Pla12b] for studying the existence of winning strategies for hybrid games. It generalizes hybrid systems to hybrid games by adding the dual operator ^{d} and a winning strategy semantics on the real Euclidean space.

Games and logic have been shown to interact fruitfully in many ways [GS53, Ehr61, Par83, Par85, Aum95, HS97, Sti01, AHK02, PP03, CHP07, AG11, Vää11]. We focus on using logic to specify and verify properties of hybrid games. Our approach is inspired by Parikh's game logic [Par83, Par85, PP03]. Game logic generalizes (propositional discrete) dynamic logic to discrete games played on a finite state space. Game logic is elegant but very challenging. Its expressiveness has only begun to be understood after two decades [Ber03, BGL07].

Our logic $\mathsf{dG}\mathcal{L}$ generalizes differential dynamic logic ($\mathsf{d}\mathcal{L}$) [Pla08, Pla12a] from hybrid systems to hybrid games and, simultaneously, generalizes game logic [Par83, Par85, PP03] from games on discrete systems to hybrid systems with their differential equations, uncountable state spaces, uncountably many possible moves, and interacting discrete and continuous dynamics.

Hybrid games generalize discrete games [vNM55, Nas51]. The games we consider are reasonably tame sequential, non-cooperative, zero-sum two-player games of perfect information with payoffs ± 1, except that they are played on hybrid systems, which makes backwards induction for winning regions on hybrid games more difficult, because it requires higher closure ordinals. Hybrid games provide a complementary perspective on differential games. Differential games

[1] A closely related question is about ways to exhibit the winning strategy, for which existence is a prerequisite and a constructive proof an answer. If we know from which states a winning strategy exists, local search is enough.

formalize various notions of adversarial control on variables for a single differential equation [Isa67, Fri71, Pet93], including solutions based on a non-anticipatory measurable input to an integral interpretation of the differential equations [Fri71], joint limits of lower and upper limits of δ-anticipatory or δ-delayed strategies for $\delta \to 0$ [Pet93], and Pareto-optimal, Nash, or Stackelberg equilibria, whose computation requires solving PDEs that quickly become ill-posed (e.g., for feedback Nash equilibria unless for dimension one or linear-quadratic games); see Bressan [Bre10] for an overview. Hybrid games, instead, distinguish discrete versus continuous parts of the dynamics, which simplifies several concepts and, simultaneously, has been argued to make other aspects more realistic [TPS98, TLS00, BBC10, VPVD11, PHP01, QP12]. The situation is similar to hybrid systems, which provide a complementary perspective on dynamical systems [Hen96, Pla12a].

Our primary contributions are that we identify the logical essence of hybrid games and their game combinators by introducing differential game logic for hybrid games with a simple modal semantics and a simple proof calculus, which we prove to be a sound and complete axiomatization relative to the fixpoint logic of differential equations. Completeness for game logics is a subtle problem. Completeness of propositional game logic has been an open problem for 30 years [Par83]. We do not address this case, but focus on hybrid games and prove a generalization of Parikh's calculus to be relatively complete for hybrid games. Our completeness proof is constructive and identifies a fixpoint-style proof technique, which can be considered a modal analogue of characterizations in the Calculus of Constructions [CH88]. This technique is practical for hybrid games, and even more efficient for hybrid systems than previous complete proof techniques. These results suggest hybrid versions of influential views of understanding program invariants as fixpoints [CC77, Cla79]. In particular, Harel's convergence rule [HMP77], which poses significant practical challenges for hybrid systems verification, turns out to be unnecessary for hybrid games, hybrid systems, and programs.

We identify separating axioms capturing the logical difference of hybrid systems versus hybrid games. We prove hybrid games to be consistent and determined, i.e. in every state, exactly one player has a winning strategy, which is the basis for assigning classical truth to logical formulas that refer to winning strategies of hybrid games. We show that winning regions of hybrid games need higher closure ordinals.

We remark that a mere fragment of $\mathsf{dG}\mathcal{L}$ can be used to verify tricky case studies in robotic factory automation [QP12] that are out of scope for other approaches. But we focus on the theoretical development of the logic here, not its use.

2 Differential Game Logic

The hybrid games we consider have no draws. If a player is deadlocked, he loses right away. If the game completes without deadlock, the player who reaches one of his winning states wins. Thus, exactly one player wins each (completed) game play with complementary winning states. The games are zero-sum games, i.e. if one player wins, the other one loses, with player payoffs ± 1. Classically, the two players are called *Angel* and *Demon*. Our games are non-cooperative and sequential games. In non-cooperative games, players do not negotiate binding contracts, but can choose to act arbitrarily according to the rules represented in the game. Sequential (or dynamic) games are games that proceed in a series of steps, where, at each step, exactly one of the players can choose an action based on the outcome of the game so far. Concurrent games, where both players choose actions simultaneously, as well as equivalent games of imperfect information, are interesting but beyond the scope of this paper yet related [vNM55, AHK02, BP09].

2.1 Syntax

Differential game logic (dG$\mathcal{L}$) is for studying properties of many different hybrid games. The idea is to describe the game form, i.e. rules, dynamics, and choices of the particular hybrid game of interest, using a program notation and then study its properties by proving the validity of logical formulas that refer to the existence of winning strategies for objectives of those hybrid games. Even though hybrid game (forms) only describe the game form, we still simply refer to them as hybrid games. The objective is defined as part of the logical formula.

Definition 1 (Hybrid games)**.** The *hybrid games of differential game logic* dG$\mathcal{L}$ are defined by the following grammar (α, β are hybrid games, x a vector of variables, θ a vector of (polynomial) terms of the same dimension, H a formula of first-order real arithmetic, and ϕ is a dG$\mathcal{L}$ formula):

$$\alpha, \beta \ ::= \ x := \theta \mid x' = \theta \,\&\, H \mid ?\phi \mid \alpha \cup \beta \mid \alpha ; \beta \mid \alpha^* \mid \alpha^d$$

Definition 2 (dG$\mathcal{L}$ formulas)**.** The *formulas of differential game logic* dG$\mathcal{L}$ are defined by the following grammar (ϕ, ψ are dG$\mathcal{L}$ formulas, p is a predicate symbol, θ_i are (polynomial) terms, x a variable, and α is a hybrid game):

$$\phi, \psi \ ::= \ p(\theta_1, \dots, \theta_k) \mid \theta_1 \geq \theta_2 \mid \neg\phi \mid \phi \wedge \psi \mid \exists x\, \phi \mid \langle\alpha\rangle\phi \mid [\alpha]\phi$$

Operators $>, =, \leq, <, \vee, \rightarrow, \leftrightarrow, \forall x$ can be defined as usual, e.g., $\forall x\, \phi \equiv \neg\exists x\, \neg\phi$. Formula $\langle\alpha\rangle\phi$ expresses that Angel has a winning strategy to achieve ϕ in hybrid game α, i.e. Angel has a strategy to reach any of the states satisfying dG$\mathcal{L}$ formula ϕ when playing hybrid game α, no matter what strategy Demon chooses. The formula $[\alpha]\phi$ expresses that Demon has a winning strategy to achieve ϕ in hybrid game α, i.e. a strategy to reach any of the states satisfying ϕ, no matter what strategy Angel chooses. Note that the same game is played in $[\alpha]\phi$ as in $\langle\alpha\rangle\phi$ with the same choices resolved by the same players. The difference between both dG$\mathcal{L}$ formulas is the player whose winning strategy they refer to. Both use the set of states where dG$\mathcal{L}$ formula ϕ is true as the winning states for that player.

The atomic games of dG$\mathcal{L}$ are assignments, continuous evolutions, and tests. In the *deterministic assignment game* $x := \theta$, the value of variable x changes instantly and deterministically to that of θ by a discrete jump without any choices to resolve. In the *continuous evolution game* $x' = \theta \,\&\, H$, the system follows the differential equation $x' = \theta$ where the duration is Angel's choice, but Angel is not allowed to choose a duration that would make the state leave the region where formula H holds. In particular, Angel is deadlocked and loses if H does not hold in the current state, because she cannot even evolve for duration 0 then. The *test game* or *challenge* $?\phi$ has no effect on the state, except that Angel loses the game if dG$\mathcal{L}$ formula ϕ does not hold in the current state.

The compound games of dG$\mathcal{L}$ are sequential composition, choice, repetition, and duals. The *sequential game* $\alpha;\beta$ is the hybrid game that first plays hybrid game α and, when hybrid game α

terminates without a player having won already, continues by playing game β. In the *choice game* $\alpha \cup \beta$, Angel chooses whether to play hybrid game α or play hybrid game β. The *repeated game* α^* plays hybrid game α repeatedly and Angel chooses, after each play of α that terminates without a player having won already, whether to play the game again or not, albeit she cannot choose to play indefinitely but has to stop repeating ultimately. Most importantly, the *dual game* α^d is the same as playing the hybrid game α with the roles of the players swapped. That is Demon decides all choices in α^d that Angel has in α, and Angel decides all choices in α^d that Demon has in α. Players who are supposed to move but deadlock lose. Thus, while the test game $?\phi$ causes Angel to lose if formula ϕ does not hold, the *dual test game* (or *dual challenge game*) $(?\phi)^d$ causes Demon to lose if ϕ does not hold. The dual operator ^{d} is the only syntactic difference of dG$\mathcal{L}$ for hybrid games compared to d$\mathcal{L}$ for hybrid systems [Pla08, Pla12a], but a fundamental one, because it is the only operator where control passes from Angel to Demon or back. Without ^{d} all choices are resolved uniformly by one player.

The logic dG$\mathcal{L}$ only provides logically essential operators. Many other game interactions for games of perfect information can be defined from the elementary operators that dG$\mathcal{L}$ provides. *Demonic choice* between hybrid game α and β is $\alpha \cap \beta$, defined by $(\alpha^d \cup \beta^d)^d$, in which either the hybrid game α or the hybrid game β is played, by Demon's choice. *Demonic repetition* of hybrid game α is $\alpha^\times$, defined by $((\alpha^d)^*)^d$, in which α is repeated as often as Demon chooses to. In $\alpha^\times$, Demon chooses after each play of α whether to repeat the game, but cannot play indefinitely so he has to stop repeating ultimately. The *dual differential equation* $(x' = \theta \,\&\, H)^d$ follows the same dynamics as $x' = \theta \,\&\, H$ except that Demon chooses the duration, so he cannot choose a duration during which H stops to hold at any time. Hence he loses when H does not hold in the current state. Dual assignment $(x := \theta)^d$ is equivalent to $x := \theta$, because it involves no choices.

Observe that every (completed) play of a game is won or lost by exactly one player. Even a play of repeated game α^* has only one winner, because the game stops as soon as one player has won. This is different than the repetition of whole game plays (including winning/losing), where the purpose is for the players to repeat the same game over and over again, win and lose multiple times, and study who wins how often in the long run with mixed strategies. In our scenario, the overall game is played once (even if some part of it constitutes in repeating action choices) and stops as soon as either Angel or Demon have won. In applications, the system is already in trouble even if it loses the game only once, because that may entail that a safety-critical property has already been violated.

2.2 Semantics

A *state* s is a mapping from variables to $\mathbb{R}$. An *interpretation* I assigns a relation $I(p) \subseteq \mathbb{R}^k$ to each predicate symbol p of arity k. We let the interpretation also determine the set of states $\mathcal{S}$, which is isomorphic to a Euclidean space $\mathbb{R}^n$ when n is the number of relevant variables. For a subset $X \subseteq \mathcal{S}$ we abbreviate $\mathcal{S} \setminus X$ by $X^\complement$. We use s_x^d to denote the state that agrees with state s except for the interpretation of variable x, which is changed to $d \in \mathbb{R}$. The value of term θ in state s is denoted by $[\![\theta]\!]_s$. The denotational semantics of dG$\mathcal{L}$ formulas is defined by simultaneous induction with the denotational semantics, $\varsigma_\alpha(\cdot)$ and $\delta_\alpha(\cdot)$, of hybrid games, defined in Def. 4.

Definition 3 ($\mathsf{dG}\mathcal{L}$ semantics)**.** The *semantics of a* $\mathsf{dG}\mathcal{L}$ *formula* ϕ for each interpretation I is the subset $[\![\phi]\!]^I \subseteq \mathcal{S}$ of states in which ϕ is true. It is defined inductively as follows

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

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

Definition 4 (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[2]

1. $\varsigma_{x:=\theta}(X) = \{s \in \mathcal{S} : s_x^{[\![\theta]\!]_s} \in X\}$
2. $\varsigma_{x'=\theta\,\&\,H}(X) = \{\varphi(0) \in \mathcal{S} : \varphi(r) \in X$ for some $r \in \mathbb{R}_{\geq 0}$ and (differentiable) $\varphi : [0,r] \to \mathcal{S}$ such that $\varphi(\zeta) \in [\![H]\!]^I$ and $\frac{\mathrm{d}\,\varphi(t)(x)}{\mathrm{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ for all $0 \leq \zeta \leq r\}$
3. $\varsigma_{?\phi}(X) = [\![\phi]\!]^I \cap X$
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. $\delta_{x:=\theta}(X) = \{s \in \mathcal{S} : s_x^{[\![\theta]\!]_s} \in X\}$

[2] The semantics of a hybrid game is not just a reachability relation as for hybrid systems [Pla12a], because the dynamic interactions and nested choices of the players have to be taken into account.

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

3. $\delta_{?\phi}(X) = ([\![\phi]\!]^I)^{\complement} \cup X$

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. $\delta_{\alpha^d}(X) = (\delta_\alpha(X^{\complement}))^{\complement}$

We write $\varsigma_\alpha(X)$ instead of $\varsigma^I_\alpha(X)$ and $\delta_\alpha(X)$ instead of $\delta^I_\alpha(X)$, because the interpretation I that gives a semantics to predicate symbols in tests and evolution domains is clear from the context. Strategies do not occur explicitly in the dG$\mathcal{L}$ semantics, because it is based on the existence of winning strategies, not the strategies. The semantics is *compositional*, i.e. the semantics of a compound dG$\mathcal{L}$ formula is a simple function of the semantics of its pieces, and the semantics of a compound hybrid game is a function of the semantics of its pieces. This makes it possible to identify a compositional proof calculus. Furthermore, existence of a strategy in hybrid game α to achieve X is independent of any game and dG$\mathcal{L}$ formula surrounding α, but just depends on the remaining game α itself and the goal X. By a simple inductive argument, this shows that one can focus on memoryless strategies, because the existence of strategies does not depend on the context, hence, by working bottom up, the strategy itself cannot depend on past states and choices, only the current state, remaining game, and goal. This follows from a generalization of a classical result [Zer13], but is directly apparent in our logical setting.

The semantics is monotone, i.e. larger sets of winning states induce larger winning regions.

Lemma 1 (Monotonicity)**.** *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.

Hybrid games branch finitely when the players decide which game to play in $\alpha \cup \beta$ and $\alpha \cap \beta$, respectively. The games α^* and $\alpha^\times$ also branch finitely, because, after each repetition of α, the respective player (Angel for α^* and Demon for $\alpha^\times$) may decide whether to repeat again or stop. Repeated games may still lead to infinitely many branches, because a repeated game can be repeated any number of times. The game branches uncountably infinitely, however, when the players decide how long to evolve along differential equations in $x' = \theta \,\&\, H$ and $(x' = \theta \,\&\, H)^d$, because uncountably many nonnegative real number could be chosen as a duration (unless the system leaves H immediately). These choices can be made explicit by relating the simple denotational modal semantics of dGL to an equivalent operational game semantics (Appendix C).

Example 1. The following simple dGL formula

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

is true in all states from which there is a winning strategy for Angel to reach [0,1). It is Angel's choice whether to repeat (*) and, if she does, it is her choice ($\cup$) whether to increase x and then give Demon control over the duration of the differential equation $x' = x^2$ (left game) or whether to decrease x (right game). Unlike the following variation, formula (1) is valid:

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

3 Meta-Properties

3.1 Determinacy

Every particular game play in a hybrid game is won by exactly one player, because there are no draws. That alone does not imply 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.

In order to understand the importance of determinacy for classical logics, we consider the semantics of repetition, which is defined as a *least* fixpoint. It is crucial that this defines a well-founded repetition. Otherwise, the *filibuster formula* would not have a well-defined truth-value:

$$\langle (x := 0 \cap x := 1)^* \rangle x = 0 \tag{2}$$

It is Angel's choice whether to repeat (*), but it is Demon's choice ($\cap$) whether to do $x := 0$ or $x := 1$. The game in this formula never deadlocks (stalemates), because every player always has a remaining move (here even two). But, without the least fixpoint, the game would have perpetual checks, because no strategy helps either player win the game; see Fig. 1. Demon can move $x := 1$ and would win, but Angel observes this and decides to repeat, so Demon can again move $x := 1$. Thus (unless Angel is lucky starting from an initial state where she has won already) every strategy that one player has to reach $x = 0$ or $x = 1$ could be spoiled by the other player so the game would not be determined, i.e. no player has a winning strategy. Every player can let his opponent win, but would not have a strategy to win himself. Because of the least fixpoint $\varsigma_{\alpha^*}(X)$ in the semantics,

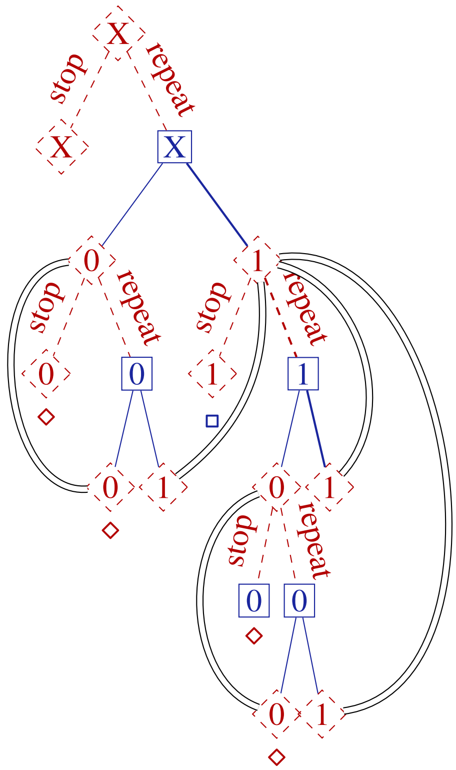

Figure 1: The filibuster game formula $\langle(x := 0 \cap x := 1)^*\rangle x = 0$ is false (unless $x = 0$ initially), but would be non-determined without least fixpoints (strategies follow thick actions). Angel's action choices are illustrated by dashed edges from dashed diamonds, Demon's action choices by solid edges from solid squares, and double lines indicate identical states with the same continuous state and a subgame of the same structure of subsequent choices. States where Angel wins are marked $\diamond$ and states where Demon wins by $\square$.

however, repetitions have to stop eventually (after an arbitrary unbounded number of rounds). Hence, in the example in Fig. 1, Demon still wins and formula (2) is *false*, unless $x = 0$ holds initially. In other words, the formula in (2) is equivalent to $x = 0$. Likewise, the dual filibuster game formula $x = 0 \rightarrow \langle(x := 0 \cup x := 1)^\times\rangle x = 0$ is (determined and) valid, because Demon has to stop repeating $^\times$ eventually so that Angel will win if she just patiently plays $x := 0$ each time. Similarly, it is important that Angel can only choose real durations $r \in \mathbb{R}_{\geq 0}$ for a continuous evolution game $x' = \theta H \,\&$, not infinity, so she will ultimately stop.

For dG$\mathcal{L}$ to be a classical two-valued modal logic, hybrid games have no draws during any game play. But, because modalities refer to the existence of winning strategies, they only receive classical truth values if we additionally ensure that, from each state, one of the players has a winning strategy for complementary winning conditions of a hybrid game α. That is, we need to make sure that hybrid games are determined. Note that logic makes this proof very simple, without the need to use the deep Borel determinacy theorem for winning conditions that are Borel in the product topology induced on game trees by the discrete topology of actions [Mar75].

Theorem 2 (Consistency & determinacy)**.** *Hybrid games are consistent and determined, i.e.* $\vDash \neg\langle\alpha\rangle\neg\phi \leftrightarrow [\alpha]\phi$*.*

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

1. $\varsigma_{x:=\theta}(X^\complement)^\complement = \{s \in \mathcal{S} : s_x^{[\![\theta]\!]_s} \notin X\}^\complement = \varsigma_{x:=\theta}(X) = \delta_{x:=\theta}(X)$
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. $\varsigma_{?\phi}(X^\complement)^\complement = ([\![\phi]\!]^I \cap X^\complement)^\complement = ([\![\phi]\!]^I)^\complement \cup (X^\complement)^\complement = \delta_{?\phi}(X)$
4. $\varsigma_{\alpha\cup\beta}(X^\complement)^\complement = (\varsigma_\alpha(X^\complement) \cup \varsigma_\beta(X^\complement))^\complement = \varsigma_\alpha(X^\complement)^\complement \cap \varsigma_\beta(X^\complement)^\complement = \delta_\alpha(X) \cap \delta_\beta(X) = \delta_{\alpha\cup\beta}(X)$
5. $\varsigma_{\alpha;\beta}(X^\complement)^\complement = \varsigma_\alpha(\varsigma_\beta(X^\complement))^\complement = \varsigma_\alpha(\delta_\beta(X)^\complement)^\complement = \delta_\alpha(\delta_\beta(X)) = \delta_{\alpha;\beta}(X)$
6. $\varsigma_{\alpha^*}(X^\complement)^\complement = \left(\bigcap\{Z \subseteq \mathcal{S} : X^\complement \cup \varsigma_\alpha(Z) \subseteq Z\}\right)^\complement$
   $= \left(\bigcap\{Z \subseteq \mathcal{S} : (X \cap \varsigma_\alpha(Z)^\complement)^\complement \subseteq Z\}\right)^\complement$
   $= \left(\bigcap\{Z \subseteq \mathcal{S} : (X \cap \delta_\alpha(Z^\complement))^\complement \subseteq Z\}\right)^\complement$
   $= \bigcup\{Z \subseteq \mathcal{S} : Z \subseteq X \cap \delta_\alpha(Z)\} = \delta_{\alpha^*}(X)$. [3]
7. $\varsigma_{\alpha^d}(X^\complement)^\complement = (\varsigma_\alpha((X^\complement)^\complement)^\complement)^\complement = \delta_\alpha(X^\complement)^\complement = \delta_{\alpha^d}(X)$ ☐

One direction of Theorem 2 implies $\vDash \neg\langle\alpha\rangle\neg\phi \to [\alpha]\phi$, i.e. $\vDash \langle\alpha\rangle\neg\phi \vee [\alpha]\phi$, which means that, from any initial state, either Angel has a winning strategy to achieve $\neg\phi$ or Demon has a winning strategy to achieve ϕ. That is, hybrid games are determined, because there are no states from which none of the players has a winning strategy (for the same hybrid game α and complementary winning conditions $\neg\phi$ and ϕ, respectively). At least one player, thus, has a winning strategy for complementary winning conditions. The other direction of Theorem 2 implies $\vDash [\alpha]\phi \to \neg\langle\alpha\rangle\neg\phi$, i.e. $\vDash \neg([\alpha]\phi \wedge \langle\alpha\rangle\neg\phi)$, which means that there is no state from which Angel has a winning strategy to achieve $\neg\phi$ and, simultaneously, Demon has a winning strategy to achieve ϕ. That is, hybrid games are *consistent*, because at most one player has a winning strategy for complementary winning conditions. Along with modal congruence rules which hold for d$\mathcal{GL}$, Theorem 2 makes d$\mathcal{GL}$ a classical modal logic [Che80].

Hybrid games α and β are *equivalent* if $\varsigma_\alpha(X) = \varsigma_\beta(X)$ for all X and all I. By Theorem 2, α and β are equivalent iff $\delta_\alpha(X) = \delta_\beta(X)$ for all X and all I. Using the equivalences

$$(\alpha \cup \beta)^d \equiv \alpha^d \cap \beta^d,\ \ (\alpha;\beta)^d \equiv \alpha^d;\beta^d,\ \ (\alpha^*)^d \equiv (\alpha^d)^\times,\ \ \alpha^{dd} \equiv \alpha$$

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

on hybrid games, every hybrid game α can be transformed into an equivalent hybrid game in which ^{d} only occurs right after atomic games or as part of the definition of the derived operators $\cap$ and ${}^\times$. Other equivalences include $x' = \theta^* \equiv x' = \theta$.

3.2 Strategic Closure Ordinals

In order to examine whether we could directly implement the dG$\mathcal{L}$ semantics to compute winning regions for formulas by a reachability computation, we investigate how many iterations the fixpoint for the semantics $\varsigma_{\alpha^*}(X)$ of repetition needs. The interaction of the repetition operator and the dual operator in dG$\mathcal{L}$ makes things more difficult compared to hybrid systems without ^{d}. The combination of * and ^{d} is also more challenging than bounded hybrid games without *, which can be unfolded equivalently into logic for hybrid systems by a simple construction using our proof calculus in Section 4.

3.2.1 Scott-Continuity

Repetitions in classical hybrid systems only repeat any finite number of times [Pla12a]. If the semantics of dG$\mathcal{L}$ were Scott-continuous, we would know that this was the case for dG$\mathcal{L}$ as well. Dual-free α are indeed Scott-continuous.

Lemma 3 (Scott-continuity of ^{d}-free dG$\mathcal{L}$)**.** *For ^{d}-free α, the semantics is Scott-continuous, i.e.* $\varsigma_\alpha(\bigcup_{n\in J} X_n) = \bigcup_{n\in J} \varsigma_\alpha(X_n)$ *for all families* $\{X_n\}_{n\in J}$ *with any index set* J.

Proof. By monotonicity, $\bigcup_{n\in J} \varsigma_\alpha(X_n) \subseteq \varsigma_\alpha(\bigcup_{n\in J} X_n)$. We show the converse inclusion by a simple induction on the structure of α: $\varsigma_\alpha(\bigcup_{n\in J} X_n) \subseteq \bigcup_{n\in J} \varsigma_\alpha(X_n)$. IH is short for induction hypothesis.

1. $\varsigma_{x:=\theta}(\bigcup_{n\in J} X_n) = \{s \in \mathcal{S} : s_x^{[\![\theta]\!]_s} \in \bigcup_{n\in J} X_n\} \subseteq \bigcup_{n\in J}\{s \in \mathcal{S} : s_x^{[\![\theta]\!]_s} \in X_n\} = \bigcup_{n\in J} \varsigma_{x:=\theta}(X_n)$, since $s_x^{[\![\theta]\!]_s} \in \bigcup_{n\in J} X_n$ implies $s_x^{[\![\theta]\!]_s} \in X_n$ for some n.

2. $\varsigma_{x'=\theta\,\&\,H}(\bigcup_{n\in J} X_n) = \{\varphi(0) \in \mathcal{S} : \frac{\mathrm{d}\,\varphi(t)(x)}{\mathrm{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ and $\varphi(\zeta) \in [\![H]\!]^I$ for all $\zeta \le r$ for some (differentiable) $\varphi : [0,r] \to \mathcal{S}$ such that $\varphi(r) \in \bigcup_{n\in J} X_n\} \subseteq \bigcup_{n\in J} \varsigma_{x'=\theta\,\&\,H}(X_n) = \{\varphi(0) \in \mathcal{S} : \dots \varphi(r) \in X_n\}$, because $\varphi(r) \in \bigcup_{n\in J} X_n$ implies $\varphi(r) \in X_n$ for some n.

3. $\varsigma_{?\phi}(\bigcup_{n\in J} X_n) = [\![\phi]\!]^I \cap \bigcup_{n\in J} X_n = \bigcup_{n\in J}([\![\phi]\!]^I \cap X_n) = \bigcup_{n\in J} \varsigma_{?\phi}(X_n)$

4. $\varsigma_{\alpha\cup\beta}(\bigcup_{n\in J} X_n) = \varsigma_\alpha(\bigcup_{n\in J} X_n) \cup \varsigma_\beta(\bigcup_{n\in J} X_n) \overset{\text{IH}}{=} (\bigcup_{n\in J} \varsigma_\alpha(X_n)) \cup (\bigcup_{n\in J} \varsigma_\beta(X_n)) = \bigcup_{n\in J}(\varsigma_\alpha(X_n) \cup \varsigma_\beta(X_n)) = \bigcup_{n\in J} \varsigma_{\alpha\cup\beta}(X_n)$

5. $\varsigma_{\alpha;\beta}(\bigcup_{n\in J} X_n) = \varsigma_\alpha(\varsigma_\beta(\bigcup_{n\in J} X_n)) \overset{\text{IH}}{=} \varsigma_\alpha(\bigcup_{n\in J} \varsigma_\beta(X_n)) \overset{\text{IH}}{=} \bigcup_{n\in J} \varsigma_\alpha(\varsigma_\beta(X_n)) = \bigcup_{n\in J} \varsigma_{\alpha;\beta}(X_n)$

   $\varsigma_{\alpha^*}(\bigcup_{n\in J} X_n) = (\bigcup_{n\in J} X_n) \cup \varsigma_\alpha(\varsigma_{\alpha^*}(\bigcup_{n\in J} X_n))$ is the least fixpoint. We will show that $\bigcup_{n\in J} \varsigma_{\alpha^*}(X_n)$ also is a fixpoint, implying $\varsigma_{\alpha^*}(\bigcup_{n\in J} X_n) \subseteq \bigcup_{n\in J} \varsigma_{\alpha^*}(X_n)$. Indeed,

$(\bigcup_{n\in J} X_n) \cup \varsigma_\alpha(\bigcup_{n\in J} \varsigma_{\alpha^*}(X_n)) \overset{\text{IH}}{=} (\bigcup_{n\in J} X_n) \cup \bigcup_{n\in J} \varsigma_\alpha(\varsigma_{\alpha^*}(X_n)) =$
$\bigcup_{n\in J}(X_n \cup \varsigma_\alpha(\varsigma_{\alpha^*}(X_n)) \overset{\mu}{=} \bigcup_{n\in J} \varsigma_{\alpha^*}(X_n)$. The equation marked μ uses that $\varsigma_{\alpha^*}(X_n)$ is a fixpoint. □

Games with both ^{d} and *, however, do not generally have a Scott-continuous semantics nor an ω-chain continuous semantics, i.e. they are not even continuous for a monotonically increasing chain $X_0 \subseteq X_1 \subseteq X_2 \subseteq \dots$ with index set ω:

$$\mathbb{R} = \varsigma_{y:=y+1^\times}(\bigcup_{n=1}^{\infty}(-\infty, n]) \not\subseteq \bigcup_{n=1}^{\infty} \varsigma_{y:=y+1^\times}((-\infty, n]) = \emptyset$$
$$\text{since } \vDash \langle y := y + 1^\times\rangle \exists n : \mathbb{N}\, y \leq n$$
$$\text{but } \nvDash \exists n : \mathbb{N}\, \langle y := y + 1^\times\rangle y \leq n$$

This example shows that, even though Angel wins this game, there is no upper bound $< \omega$ on the number of iterations it takes her to win, because Demon could repeat $y := y + 1^\times$ arbitrarily often. This phenomenon is directly related to a failure of the Barcan axiom (Section 4.4). The quantifier $\exists n : \mathbb{N}$ over natural numbers is not essential here [Pla08].

A continuous variation of this argument shows that the semantics is not ω_1-based, where ω_1 is the first uncountable ordinal. A function τ on sets is κ-based, for an ordinal κ, if for all X, $x \in \tau(X)$ implies $x \in \tau(Y)$ for some $Y \subseteq X$ of cardinality $< \kappa$. The semantics $\varsigma_\alpha(\cdot)$ is not ω_1-based, because of Lemma 1 and removing just one state from the winning condition may lose states in the winning region:

$$[0, \infty) = \varsigma_{x'=1^d}([0, \infty))$$
$$\text{but } 0 \notin \varsigma_{x'=1^d}([0, \infty) \setminus \{a\}) = (a, \infty) \text{ for all } a > 0$$

3.2.2 Fixpoints

While ω may not be the number of iterations for the winning region $\varsigma_{\alpha^*}(X)$, Knaster-Tarski's seminal fixpoint theorem entails that there is some ordinal $\bar{\lambda}$ at which the iterations for the semantics of α^* stop. We use the following minor variation (starting with x at the bottom) of Kozen's formulation of the Knaster-Tarski theorem [HKT00, Theorem 1.12].

Let $\tau : L \to L$ a monotone operator on a partial order L, then $\tau^\lambda(x) \overset{\text{def}}{=} x \cup \bigcup_{\kappa<\lambda} \tau(\tau^\kappa(x))$ for all ordinals λ is equivalent to:

$$\tau^0(x) \overset{\text{def}}{=} x$$
$$\tau^{\kappa+1}(x) \overset{\text{def}}{=} x \cup \tau(\tau^\kappa(x))$$
$$\tau^\lambda(x) \overset{\text{def}}{=} \bigcup_{\kappa<\lambda} \tau^\kappa(x) \qquad \lambda \neq 0 \text{ a limit ordinal}$$

Yet, $\bigcup$ and $\tau^\lambda(x)$ are only guaranteed to exist if L is a complete partial order. If there is a $\bar\lambda$ such that $\tau^{\bar\lambda}(x) = \tau^{\bar\lambda+1}(x)$, then $\tau^{\bar\lambda}(x)$ is the least fixpoint above x and for all κ:

$$\tau^\dagger(x) \stackrel{\text{def}}{=} \bigcap\{z \in L : x \subseteq z, \tau(z) \subseteq z\} = \tau^{\bar\lambda}(x) = \tau^{\bar\lambda+\kappa}(x)$$

The least ordinal $\bar\lambda$ with that property is called *closure ordinal*. If τ is Scott-continuous on a complete partial order, then $\tau^\dagger(x) = \tau^\omega(x)$ by Kleene's fixpoint theorem, implying $\bar\lambda \leq \omega$. But $\varsigma_\alpha(\cdot)$ is not generally Scott-continuous, so $\bar\lambda$ might potentially be greater. If τ is countably-continuous on a complete partial order, then $\bar\lambda \leq \omega_1$. But this is not the case for $\varsigma_\alpha(\cdot)$ either, by the argument in Section 3.2.1.

Theorem 4 (Knaster-Tarski [HKT00]). *For every complete lattice L, there is an ordinal $\bar\lambda$ of cardinality $\leq |L|$ such that, for each monotone $\tau : L \to L$, i.e. $\tau(x) \subseteq \tau(y)$ for all $x \subseteq y$, the fixpoints of τ in L are a complete lattice and for all x and κ:*

$$\tau^\dagger(x) \stackrel{\textit{def}}{=} \bigcap\{z \in L : x \subseteq z, \tau(z) \subseteq z\} = \tau^{\bar\lambda}(x) = \tau^{\bar\lambda+\kappa}(x)$$

We first collect useful properties of $\tau^\kappa(\cdot)$. Since we use the extensive / inflationary definition of $\tau^\kappa(x)$, $\tau^\kappa(x)$ is not just monotone in x but also monotone and homomorphic in κ:

Lemma 5. *τ is inductive, i.e. $\tau^\kappa(x) \subseteq \tau^\lambda(x)$ for all $\kappa \leq \lambda$ and homomorphic in κ, i.e. $\tau^{\kappa+\lambda}(x) = \tau^\lambda(\tau^\kappa(x))$ for all κ, λ.*

Proof. Inductiveness, i.e. $\tau^\kappa(x) \subseteq \tau^\lambda(x)$ for $\kappa \leq \lambda$, which is monotonicity in κ, holds by definition [HKT00, Lemma 1.11]. Homomorphy in κ, i.e. $\tau^{\kappa+\lambda}(x) = \tau^\lambda(\tau^\kappa(x))$ can be proved by induction on λ, which is either 0, a successor ordinal (second line) or a limit ordinal $\neq 0$ (third line):

$$\begin{aligned}
\tau^{\kappa+0}(x) &= \tau^\kappa(x) = \tau^0(\tau^\kappa(x)) \\
\tau^{\kappa+(\lambda+1)}(x) &= x \cup \tau(\tau^{\kappa+\lambda}(x)) = x \cup \tau(\tau^\lambda(\tau^\kappa(x))) \\
&= \tau^\kappa(x) \cup \tau(\tau^\lambda(\tau^\kappa(x))) = \tau^{\lambda+1}(\tau^\kappa(x)) \\
\tau^{\kappa+\lambda}(x) &= \bigcup_{\iota<\kappa+\lambda} \tau^\iota(x) = \bigcup_{\iota<\kappa} \tau^\iota(x) \cup \bigcup_{\iota<\lambda} \tau^{\kappa+\iota}(x) \\
&= \bigcup_{\iota<\lambda} \tau^{\kappa+\iota}(x) = \bigcup_{\iota<\lambda} \tau^\iota(\tau^\kappa(x)) = \tau^\lambda(\tau^\kappa(x)) \Box
\end{aligned}$$

3.2.3 Higher Closure Ordinals

By Theorem 4, there is an ordinal $\bar\lambda$ of cardinality $\leq |\mathbb{R}|$ such that $\varsigma_{\alpha^*}(X) = \varsigma_\alpha^{\bar\lambda}(X)$ for all α and all X, because the powerset lattice is complete and $\varsigma_\alpha(\cdot)$ monotone by Lemma 1. This iterative definition $\varsigma_\alpha^{\bar\lambda}(X)$ corresponds to backward induction in classical game theory [vNM55, Aum95], yet it terminates at ordinal $\bar\lambda$ which may not be finite. How soon will this fixpoint iteration for winning regions stop? Unfortunately, hybrid games may have higher closure ordinals, because ω many repetitions of the operator (and even ω^n many) may not be enough to compute winning regions.

Theorem 6 (Closure ordinals)**.** *The semantics of* dG$\mathcal{L}$ *has a closure ordinal* $\geq \omega^\omega$*, i.e. for all* $\lambda < \omega^\omega$*, there are* α *and* X *such that* $\varsigma_{\alpha^*}(X) \neq \varsigma_\alpha^\lambda(X)$*.*

Proof. We show an easier proof that the closure ordinal is $\geq \omega \cdot 2$ and show the full proof for $\geq \omega^\omega$ in Appendix E. The specific dG$\mathcal{L}$ formulas considered for these increasing lower bounds also show that the closure ordinal is not a simple function of the syntactic structure, because minor syntactic variations lead to vastly different closure ordinals. To see that the closure ordinal is $> \omega$ even with just one variable, a single loop and dual, we consider the semantics of the following dG$\mathcal{L}$ formula, i.e. the set of states in which it is true:

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

We show that the winning regions for this dG$\mathcal{L}$ formula stabilize after $\omega \cdot 2$ iterations, because ω many iterations are necessary to show that *any* positive real can be reduced to $[0, 1)$ by choosing β sufficiently often, whereas another ω many iterations are needed to show that choice α, which makes progress ≥ 1 but possibly more under Demon's control, can turn x into a positive real. It is easy to see that $\varsigma^\omega_{\alpha\cup\beta}([0,1)) = \bigcup_{n\in\mathbb{N}} \varsigma^n_{\alpha\cup\beta}([0,1)) = [0,\infty)$, because $\varsigma^n_{\alpha\cup\beta}([0,1)) = [0,n)$ holds for all $n \in \mathbb{N}$ by a simple inductive argument:

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

But the iteration for the winning region does not stop at ω, because $\varsigma^{\omega+n}_{\alpha\cup\beta}([0,1)) = [-n,\infty)$ holds for all $n \in \mathbb{N}$ by another simple inductive argument:

$$\begin{aligned}
\varsigma^{\omega+n+1}_{\alpha\cup\beta}([0,1)) &= [0,1) \cup \varsigma_{\alpha\cup\beta}(\varsigma^{\omega+n}_{\alpha\cup\beta}([0,1)))\\
&= [0,1) \cup \varsigma_{\alpha\cup\beta}([-n,\infty))\\
&= [0,1) \cup \varsigma_\alpha([-n,\infty)) \cup \varsigma_\beta([-n,\infty))\\
&= [-n-1,\infty) \cup [-n,\infty)
\end{aligned}$$

Thus, $\varsigma^{\omega\cdot 2}_{\alpha\cup\beta}([0,1)) = \varsigma^{\omega+\omega}_{\alpha\cup\beta}([0,1)) = \bigcup_{n\in\mathbb{N}} \varsigma^{\omega+n}_{\alpha\cup\beta}([0,1)) = \mathbb{R} = \varsigma_{\alpha\cup\beta}(\mathbb{R})$. In this case, the closure ordinal is $\omega \cdot 2 > \omega$, since $\varsigma_{(\alpha\cup\beta)^*}([0,1)) = \mathbb{R} \neq \varsigma^{\omega+n}_{\alpha\cup\beta}([0,1))$ for all $n \in \mathbb{N}$. □

Consequently, the dG$\mathcal{L}$ semantics is more general than defining $\varsigma_{\alpha^*}(X)$ to be truncated to ω-repetition $\varsigma^\omega_\alpha(X) = \bigcup_{n\in\mathbb{N}} \varsigma^n_\alpha(X)$, which misses out on the existence of perfectly natural winning strategies. The semantics of dG$\mathcal{L}$ is also different than advance notice semantics; see Appendix D.

4 Axiomatization

Simple dG$\mathcal{L}$ formulas can be checked by a tableau procedure that expands all choices and detects loops for termination as in our game tree examples (Fig. 1). This principle does not extend to more general hybrid games with inherently infinite state spaces [Hen96] and which need higher ordinals of iteration for computing winning regions by Theorem 6. Reachability computations for higher ordinals may not terminate, except when widening to find approximations and risk incompleteness [CC77]. Widening would be interesting, but we focus on axiomatizations in logic to identify the logical essence.

4.1 Proof Calculus

Fig. 2 presents a proof calculus for proving validity of dGL formulas as a more general symbolic proof technique.

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

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

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

$$\langle?\rangle\quad \langle ?\psi\rangle\phi \leftrightarrow (\psi \land \phi)$$

$$\langle\cup\rangle\quad \langle\alpha \cup \beta\rangle\phi \leftrightarrow \langle\alpha\rangle\phi \lor \langle\beta\rangle\phi$$

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

$$\langle^*\rangle\quad \phi \lor \langle\alpha\rangle\langle\alpha^*\rangle\phi \rightarrow \langle\alpha^*\rangle\phi$$

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

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

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

Figure 2: Differential game logic axiomatization

The proof calculus of dGL shares axioms with dL [Pla12a] and game logic [PP03]. It is based on the first-order Hilbert calculus (uniform substitution, modus ponens, and Bernays' $\forall$-generalization) with all instances of valid formulas of first-order logic as axioms, including first-order real arithmetic [Tar51].

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 game α. Axiom $\langle{:=}\rangle$ is Hoare's assignment rule. Formula $\phi(\theta)$ is obtained from $\phi(x)$ by *substituting* θ for 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 $?$ *binds* x. In 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 y is for Angel to decide. It goes without saying that variables like t are fresh in Fig. 2. Axioms $\langle ?\rangle$, $\langle \cup\rangle$, and $\langle ;\rangle$ are as in dynamic logic [Pra76] and d$\mathcal{L}$ [Pla12a] except that their meaning is different, because they refer to hybrid games instead of systems. The challenge axiom $\langle ?\rangle$ expresses that Angel has a winning strategy to achieve ϕ in the test game $?\psi$ exactly from those positions that are already in ϕ (because $?\psi$ does not change the state) and satisfy ψ for otherwise she would fail the test. Axiom $\langle \cup\rangle$ 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. 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 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 Angel has a winning strategy to achieve ϕ in α^*. The converse of $\langle {}^*\rangle$ can be derived[4] and is also denoted by $\langle {}^*\rangle$. 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. By combining axioms $\langle {}^d\rangle$ and $[\cdot]$ we obtain $\langle\alpha^d\rangle\phi \leftrightarrow [\alpha]\phi$, 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.

Rule M is the generalization rule of monotonic modal logic **C** [Che80]. It expresses that, if the implication $\phi \rightarrow \psi$ is valid, then, whenever Angel has a winning strategy in any hybrid game α to achieve ϕ, she also has a winning strategy to achieve ψ. Fixpoint rule FP characterizes $\langle\alpha^*\rangle\phi$ as a least pre-fixpoint. It says that, if ψ is any other 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 where $\langle\alpha^*\rangle\phi$ does, i.e. in all states from which Angel has a winning strategy in game α^* to achieve ϕ.

As usual, all substitutions in Fig. 2 are required to be *admissible* to avoid capture of variables, i.e. they require all variables x that are being replaced or that occur in their replacements not to occur in the scope of a quantifier or modality binding x. Recall that the uniform substitution rule from first-order logic substitutes *all* occurrences of predicate $p(\cdot)$ by a dG$\mathcal{L}$ formula $\psi(\cdot)$, i.e. it replaces all occurrences of $p(\theta)$ for any vectorial term θ by $\psi(\theta)$ simultaneously:

$$\text{(US)}\quad \frac{\phi}{\phi_{p(\cdot)}^{\psi(\cdot)}}$$

In particular, the uniform substitution rule requires all relevant substitutions of $\psi(\theta)$ for $p(\theta)$ to be admissible and requires that no $p(\theta)$ occurs in the scope of a quantifier or modality binding a variable of $\psi(\theta)$ other than those in θ; see [Chu56, §35,40]. If admissible, the formula $\psi(\theta)$ can use variables other than those in θ, hence, the case where p is a predicate symbol without arguments enables US to generate all instances from the dG$\mathcal{L}$ axioms, so that the axioms in Fig. 2 do not need to be considered as axiom schemes [Chu56, §35,40].

[4] $\phi \lor \langle\alpha\rangle\langle\alpha^*\rangle\phi \rightarrow \langle\alpha^*\rangle\phi$ derives by $\langle {}^*\rangle$. Thus, $\langle\alpha\rangle(\phi \lor \langle\alpha\rangle\langle\alpha^*\rangle\phi) \rightarrow \langle\alpha\rangle\langle\alpha^*\rangle\phi$ by M. Hence, $\phi \lor \langle\alpha\rangle(\phi \lor \langle\alpha\rangle\langle\alpha^*\rangle\phi) \rightarrow \phi \lor \langle\alpha\rangle\langle\alpha^*\rangle\phi$ by propositional congruence. Consequently, $\langle\alpha^*\rangle\phi \rightarrow \phi \lor \langle\alpha\rangle\langle\alpha^*\rangle\phi$ by FP.

The primary difference of the axiomatization of dGL compared to differential dynamic logic for hybrid systems [Pla12a] is the addition of axiom $\langle^d\rangle$ for dual games, the absence of axiom K, absence of Gödel's necessitation rule (dGL only has the monotonic modal rule M), absence of the Barcan formula (the converse Barcan formula is still derivable), and absence of the hybrid version of Harel's convergence rule [HMP77]. Given the big semantical difference of a hybrid system run versus a hybrid game, it is striking to see this concise difference in axioms. This is an indication that these are the appropriate logical characterizations. Due to the absence of K, we will see (in Section 4.4) why the induction axiom and the convergence axiom are absent in dGL, while corresponding proof rules are still valid. The induction rule (ind) is derivable from FP.

Lemma 7 (Invariance)**.** *Rule FP and the induction rule (ind) of dynamic logic are interderivable in the* dGL *calculus:*

$$\text{(ind)}\quad \frac{\psi \to [\alpha]\psi}{\psi \to [\alpha^*]\psi}$$

Proof. Rule ind derives from FP: We first derive the following minor variant

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

From $\psi \to [\alpha]\psi$ and $\psi \to \phi$ propositionally derive $\psi \to \phi \land [\alpha]\psi$, from which contraposition and propositional logic yield $\neg\phi \lor \neg[\alpha]\psi \to \neg\psi$. With $[\cdot]$, this gives $\neg\phi \lor \langle\alpha\rangle\neg\psi \to \neg\psi$. Now FP derives $\langle\alpha^*\rangle\neg\phi \to \neg\psi$, which, by $[\cdot]$, is $\neg[\alpha^*]\phi \to \neg\psi$, which gives $\psi \to [\alpha^*]\phi$ by contraposition. The classical $[]$-induction rule ind follows by $\phi \overset{\text{def}}{\equiv} \psi$. From ind, the variant ind_R is derivable again by M on $\psi \to \phi$.

Rule FP derives from ind: From $\phi \lor \langle\alpha\rangle\psi \to \psi$, propositionally derive $\phi \to \psi$ and $\langle\alpha\rangle\psi \to \psi$. By M, the former gives $\langle\alpha^*\rangle\phi \to \langle\alpha^*\rangle\psi$. By contraposition, the latter derives $\neg\psi \to \neg\langle\alpha\rangle\psi$, which gives $\neg\psi \to [\alpha]\neg\psi$ by $[\cdot]$. Now ind derives $\neg\psi \to [\alpha^*]\neg\psi$. By contraposition $\neg[\alpha^*]\neg\psi \to \psi$, which, by $[\cdot]$, is $\langle\alpha^*\rangle\psi \to \psi$. Thus, $\langle\alpha^*\rangle\phi \to \psi$ by the formula derived above. ☐

Example 2. The dual filibuster game formula (Section 3.1) proves easily by going back and forth between players:

$$\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{\cfrac{*}{x = 0 \to 0 = 0 \lor 1 = 0}\ \mathbb{R}}{x = 0 \to \langle x := 0\rangle x = 0 \lor \langle x := 1\rangle x = 0}\ \langle:=\rangle}{x = 0 \to \langle x := 0 \cup x := 1\rangle x = 0}\ \langle\cup\rangle}{x = 0 \to \neg\langle x := 0 \cap x := 1\rangle \neg x = 0}\ \langle^d\rangle}{x = 0 \to [x := 0 \cap x := 1]x = 0}\ [\cdot]}{x = 0 \to [(x := 0 \cap x := 1)^*]x = 0}\ \text{ind}}{x = 0 \to \langle(x := 0 \cup x := 1)^\times\rangle x = 0}\ \langle^d\rangle$$

More challenging hybrid games are provable in dGL; see [QP12] for a stress-test of a highly interactive 11-dimensional nonlinear hybrid game in robotic factory automation.

4.2 Soundness

Crucially, we prove soundness of the dG$\mathcal{L}$ calculus, i.e. all derivable formulas are valid. The proof uses the fact that the following congruence rule derives from two uses of rule M:

$$\text{(RE)} \quad \frac{\phi \leftrightarrow \psi}{\langle\alpha\rangle\phi \leftrightarrow \langle\alpha\rangle\psi}$$

Theorem 8 (Soundness)**.** *The* dG$\mathcal{L}$ *proof rules in Fig. 2 are sound, i.e. all provable formulas are valid.*

Proof. In order to prove soundness of an implication axiom $\phi \to \psi$, we fix any interpretation I with any set of states $\mathcal{S}$, and need to show $[\![\phi]\!]^I \subseteq [\![\psi]\!]^I$. To prove soundness of an equivalence axiom $\phi \leftrightarrow \psi$, we need to show $[\![\phi]\!]^I = [\![\psi]\!]^I$. To prove soundness of a proof rule

$$\frac{\phi}{\psi}$$

we assume that ϕ is valid, i.e. $[\![\phi]\!]^I = \mathcal{S}$ in all interpretations I with any set of states $\mathcal{S}$, and prove that ψ is valid, i.e. $[\![\psi]\!]^I = \mathcal{S}$ in all I with any $\mathcal{S}$. For most proof rules we prove the stronger condition of *local soundness*, i.e. for any interpretation I with any set of states $\mathcal{S}$: $[\![\phi]\!]^I = \mathcal{S}$ implies $[\![\psi]\!]^I = \mathcal{S}$. We use the μ-calculus notation in this proof where $\mu Z.\Upsilon(Z)$ denotes the least fixpoint of $\Upsilon(Z)$ and $\nu Z.\Upsilon(Z)$ denotes the greatest fixpoint. Soundness of modus ponens (MP) and $\forall$-generalization is standard and not shown.

- $[\cdot]$ $[\![[\alpha]\phi]\!]^I = [\![\neg\langle\alpha\rangle\neg\phi]\!]^I$ is a corollary to determinacy (Theorem 2).

- $\langle:=\rangle$ $[\![\langle x := \theta\rangle\phi(x)]\!]^I = \varsigma_{x:=\theta}([\![\phi(x)]\!]^I) = \{s \in \mathcal{S} : s_x^{[\![\theta]\!]_s} \in [\![\phi(x)]\!]^I\} = \{s \in \mathcal{S} : s \in [\![\phi(\theta)]\!]^I\} = [\![\phi(\theta)]\!]^I$, where the penultimate equation holds by the substitution lemma. The classical substitution lemma is sufficient for first-order logic $\phi(\theta)$. Otherwise the proof of the substitution lemma for d$\mathcal{L}$ [Pla10b, Lemma 2.2] generalizes to dG$\mathcal{L}$.

- $\langle'\rangle$ $[\![\langle x' = \theta\rangle\phi]\!]^I = \varsigma_{x'=\theta}([\![\phi]\!]^I) = \{\varphi(0) \in \mathcal{S} :$ for some $\varphi : [0, r] \to \mathcal{S}$ such that $\varphi(r) \in [\![\phi]\!]^I$ and $\frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ for all $\zeta \leq r\}$. Further, $[\![\exists t{\geq}0\,\langle x := y(t)\rangle\phi]\!]^I = \{s \in \mathcal{S} : s_t^r \in [\![\langle x := y(t)\rangle\phi]\!]^I$ for some $r \geq 0\} = \{s \in \mathcal{S} : s_t^r \in \{u \in \mathcal{S} : u_x^{[\![y(t)]\!]_u} \in [\![\phi]\!]^I\}$ some $r \geq 0\} = \{s \in \mathcal{S} : (s_t^r)_x^{[\![y(t)]\!]_{s_t^r}} \in [\![\phi]\!]^I$ for some $r \geq 0\}$. The inclusion "$\supseteq$" between both parts, because the function $\varphi(\zeta) := (s_t^\zeta)_x^{[\![y(t)]\!]_{s_t^\zeta}}$ solves the differential equation $x' = \theta$ by assumption. The inclusion "$\subseteq$" follows, because the solution of the smooth differential equation $x' = \theta$ is unique [Pla10b, Lemma 2.1].

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

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

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

⟨*⟩ Since $[\![\langle\alpha^*\rangle\phi]\!]^I = \varsigma_{\alpha^*}([\![\phi]\!]^I) = \mu Z.([\![\phi]\!]^I \cup \varsigma_\alpha(Z))$ is a fixpoint, we know $[\![\langle\alpha^*\rangle\phi]\!]^I = [\![\phi]\!]^I \cup \varsigma_\alpha([\![\langle\alpha^*\rangle\phi]\!]^I)$. Thus, $[\![\phi \lor \langle\alpha\rangle\langle\alpha^*\rangle\phi]\!]^I = [\![\phi]\!]^I \cup [\![\langle\alpha\rangle\langle\alpha^*\rangle\phi]\!]^I = [\![\phi]\!]^I \cup \varsigma_\alpha([\![\langle\alpha^*\rangle\phi]\!]^I) = [\![\langle\alpha^*\rangle\phi]\!]^I$. Consequently, $[\![\phi \lor \langle\alpha\rangle\langle\alpha^*\rangle\phi]\!]^I \subseteq [\![\langle\alpha^*\rangle\phi]\!]^I$.

⟨d⟩ $[\![\langle\alpha^d\rangle\phi]\!]^I = \varsigma_{\alpha^d}([\![\phi]\!]^I) = \varsigma_\alpha(([\![\phi]\!]^I)^\complement)^\complement = \varsigma_\alpha([\![\neg\phi]\!]^I)^\complement = ([\![\langle\alpha\rangle\neg\phi]\!]^I)^\complement = [\![\neg\langle\alpha\rangle\neg\phi]\!]^I$ by Def. 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 1).

FP Assume the premise $\phi \lor \langle\alpha\rangle\psi \to \psi$ is valid in I, i.e. $[\![\phi \lor \langle\alpha\rangle\psi]\!]^I \subseteq [\![\psi]\!]^I$. Thus, $[\![\phi]\!]^I \cup \varsigma_\alpha([\![\psi]\!]^I) = [\![\phi]\!]^I \cup [\![\langle\alpha\rangle\psi]\!]^I = [\![\phi \lor \langle\alpha\rangle\psi]\!]^I \subseteq [\![\psi]\!]^I$. That is, ψ is a pre-fixpoint of $Z = [\![\phi]\!]^I \cup \varsigma_\alpha(Z)$. Now using Lemma 1, $[\![\langle\alpha^*\rangle\phi]\!]^I = \varsigma_{\alpha^*}([\![\phi]\!]^I) = \mu Z.([\![\phi]\!]^I \cup \varsigma_\alpha(Z))$ is the least fixpoint and the least pre-fixpoint. Thus, $[\![\langle\alpha^*\rangle\phi]\!]^I \subseteq [\![\psi]\!]^I$, which implies that $\langle\alpha^*\rangle\phi \to \psi$ is valid in I.

US Standard soundness proofs for US [Chu56] generalize to dG$\mathcal{L}$. We show a proof in our notation, because it is based on an elegant use of the soundness of RE. Assume the premise ϕ is valid, i.e. $[\![\phi]\!]^I = \mathcal{S}$ in all interpretations I with any set of states $\mathcal{S}$. We can assume that the uniform substitution is admissible, otherwise rule US is not applicable and there is nothing to show. We prove that $\phi_{p(\cdot)}^{\psi(\cdot)}$ is valid, i.e. $[\![\phi_{p(\cdot)}^{\psi(\cdot)}]\!]^I = \mathcal{S}$ for all I with $\mathcal{S}$. Consider any particular interpretation J with set of states $\mathcal{S}$. Without loss of generality, we can assume p not to occur in $\psi(\cdot)$ (otherwise first replace all occurrences of p in $\psi(\cdot)$ by q and then use rule US again to replace those q by p). Thus, by uniform substitution, p does not occur in $\phi_{p(\cdot)}^{\psi(\cdot)}$ and the value of $J(p)$ is immaterial for the semantics of $\phi_{p(\cdot)}^{\psi(\cdot)}$. We can, therefore, pass to an interpretation I that modifies J by changing the semantics of p such that $[\![p(x)]\!]^I = [\![\psi(x)]\!]^J$ for all values of x. In particular, $[\![p(x)]\!]^I = [\![\psi(x)]\!]^I$ for all values of x, since p does not occur in $\psi(x)$. Thus, $I \models \forall x\,(p(x) \leftrightarrow \psi(x))$. Since M is locally sound, so is the congruence rule RE, which derives from M. The principle of substitution of equivalents [HC96, Chapter 13] (from $A \leftrightarrow B$ derive $\Upsilon(A) \leftrightarrow \Upsilon(B)$, where $\Upsilon(B)$ is the formula $\Upsilon(A)$ with some occurrences of A replaced by B), thus, generalizes to dG$\mathcal{L}$ and is locally sound. Hence, for any particular occurrence of $p(u)$ in ϕ, we have $I \models p(u) \leftrightarrow \psi(u)$, which implies $I \models \phi \leftrightarrow \phi_{p(u)}^{\psi(u)}$ for the ordinary replacement of $p(u)$ by $\psi(u)$. This process can be repeated for all occurrences of $p(u)$, leading to $I \models \phi \leftrightarrow \phi_{p(\cdot)}^{\psi(\cdot)}$. Thus, $\mathcal{S} = [\![\phi]\!]^I = [\![\phi_{p(\cdot)}^{\psi(\cdot)}]\!]^I$. Hence, $[\![\phi_{p(\cdot)}^{\psi(\cdot)}]\!]^J = \mathcal{S}$, because p no longer occurs after uniform substitution $\phi_{p(\cdot)}^{\psi(\cdot)}$, since all occurrences of p with any arguments will have been replaced at some point (since admissible). This implies that $\phi_{p(\cdot)}^{\psi(\cdot)}$ is valid since interpretation J with set of states $\mathcal{S}$ was arbitrary. □

The proof rules in Fig. 2 do not handle differential equations $x' = \theta \,\&\, H$ with evolution domain constraints H (other than $true$). Quite unlike in hybrid systems and (poor test) differential dynamic

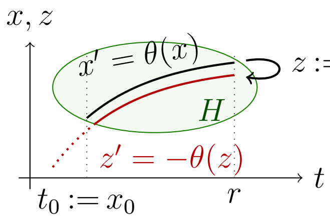

Figure 3: Angel evolves x forwards in time along $x' = \theta(x)$, Demon checks evolution domain backwards in time along $z' = -\theta(z)$ on a copy z of the state

logic [Pla08, Pla12a], however, every hybrid game containing a differential equation with evolution domain constraints can be replaced equivalently by a hybrid game without evolution domain constrains (even using poor tests, i.e. each test $?\phi$ uses only first-order formulas ϕ).

Lemma 9. *Evolution domains of differential equations are definable as hybrid games: For every hybrid game α, there is a hybrid game β that is equivalent (i.e. $\varsigma_\alpha(X) = \varsigma_\beta(X)$ for all X and all I) but has no evolution domain constraints.*

Proof. When, for notational convenience, we assume the (vectorial) differential equation $x' = \theta(x)$ to contain a clock $x_0' = 1$ and that t_0 and z are fresh variables, then $x' = \theta(x) \,\&\, H(x)$ is equivalent to the hybrid game:

$$t_0 := x_0; x' = \theta(x); (z := x; z' = -\theta(z))^d; ?(z_0 \geq t_0 \rightarrow H(z)) \tag{3}$$

See Fig. 3 for an illustration. Suppose the current player is Angel. The idea behind game equivalence (3) is that the fresh variable t_0 remembers the initial time x_0, and Angel then evolves along $x' = \theta(x)$ for any amount of time (Angel's choice). Afterwards, the opponent Demon copies the state x into a fresh variable (vector) z that it can evolve backwards along $(z' = -\theta(z))^d$ for any amount of time (Demon's choice). The original player Angel must then pass the challenge $?(z_0 \geq t_0 \rightarrow H(z))$, i.e. Angel loses immediately if Demon was able to evolve backwards and leave region $H(z)$ while satisfying $z_0 \geq t_0$, which checks that Demon did not evolve backward for longer than Angel evolved forward. Otherwise, when Angel passes the test, the extra variables t_0, z become irrelevant (they are fresh) and the game continues from the current state x that Angel chose in the first place (by selecting a duration for the evolution that Demon could not invalidate). □

Lemma 9 can eliminate all evolution domain constraints equivalently in hybrid games from now on.

4.3 Completeness

Completeness of dG$\mathcal{L}$ is a challenging question and related to a famous open problem about completeness of propositional game logic [Par83]. Based on Gödel's second incompleteness theorem, d$\mathcal{L}$ is incomplete [Pla08] and so is dG$\mathcal{L}$. So the right question to ask is that of relative completeness [Coo78, HMP77].

One natural choice for an oracle logic for the completeness study is $L_{\mu\text{D}}$, the modal μ-calculus of differential equations:

$$\phi ::= X(\theta) \mid p(\theta) \mid \theta_1 \geq \theta_2 \mid \neg\phi \mid \phi \wedge \psi \mid \langle x' = \theta \rangle \phi \mid \mu X.\phi$$

where $\mu X.\phi$ requires all occurrences of X in ϕ to be positive. The semantics is the usual, e.g., $\mu X.\phi$ binds set variable X and real variable (vector) x and is interpreted as the least fixpoint X of ϕ, i.e. the smallest denotation of X such that $X(x) \leftrightarrow \phi$ holds for all x [Koz83, Lub89]. A more careful inspection of our proofs reveals that the two-variable fragment of $L_{\mu\text{D}}$ is enough. The *fixpoint logic* $L_{\mu D}$ *of differential equations* exposes the most natural interactivity on top of differential equations.

Lemma 10 (d$\mathcal{GL}$ expressibility)**.** *Logic* d$\mathcal{GL}$ *is* expressible *in $L_{\mu D}$: for each* d$\mathcal{GL}$ *formula ϕ there is a $L_{\mu D}$ formula $\phi^\flat$ that is equivalent, i.e. $\vDash \phi \leftrightarrow \phi^\flat$.*

Proof. By soundness of axiom $[\cdot]$, we do not need to consider the case $[\alpha]$, because $[\alpha]\phi \equiv \langle \alpha^d \rangle \phi$. Of course, $(p(\theta))^\flat = p(\theta)$ etc. The inductive cases are:

$$\begin{aligned}
(\neg\phi)^\flat &\equiv \neg(\phi^\flat)\\
(\phi \wedge \psi)^\flat &\equiv \phi^\flat \wedge \psi^\flat\\
(\exists x\, \phi)^\flat &\equiv \exists x\, (\phi^\flat)\\
(\langle x := \theta \rangle \phi)^\flat &\equiv \forall y\, (y = \theta \rightarrow (\phi_x^y)^\flat)\\
(\langle x' = \theta \rangle \phi)^\flat &\equiv \langle x' = \theta \rangle \phi^\flat\\
(\langle ?\psi \rangle \phi)^\flat &\equiv (\psi \wedge \phi)^\flat\\
(\langle \alpha \cup \beta \rangle \phi)^\flat &\equiv (\langle \alpha \rangle \phi \vee \langle \beta \rangle \phi)^\flat\\
(\langle \alpha ; \beta \rangle \phi)^\flat &\equiv (\langle \alpha \rangle \langle \beta \rangle \phi)^\flat\\
(\langle \alpha^* \rangle \phi)^\flat &\equiv \mu X.(\phi \vee \langle \alpha \rangle X(x))^\flat\\
(\langle \alpha^d \rangle \phi)^\flat &\equiv (\neg \langle \alpha \rangle \neg \phi)^\flat
\end{aligned}$$

It is easy to check that $\phi^\flat$ is equivalent to ϕ. Note that $(\phi \vee \psi)^\flat \equiv \phi^\flat \vee \psi^\flat$ is a consequence of the above definitions and the abbreviation $\phi \vee \psi \equiv \neg(\neg\phi \wedge \neg\psi)$. The substitution in the definition of $(\langle x := \theta \rangle \phi)^\flat$ is necessary, because a substitution of θ for x may not be admissible, but a variable renaming of fresh variable y for x in ϕ with the result ϕ_x^y is always admissible. Note that quantifiers are expressible in $L_{\mu\text{D}}$ via $\exists x\, \phi \equiv \langle x' = 1 \rangle \phi \vee \langle x' = -1 \rangle \phi$. Recall that $x' = \theta \,\&\, H$ is expressible by Lemma 9. The case $(\langle \alpha^* \rangle \phi)^\flat$ is defined as the least fixpoint of the reduction of $\phi \vee \langle \alpha \rangle X(x)$, where x are the variables of α using classical short notation [Lub89]. In particular, $(\langle \alpha^* \rangle \phi)^\flat$ satisfies $\phi \vee \langle \alpha \rangle (\langle \alpha^* \rangle \phi)^\flat \leftrightarrow (\langle \alpha^* \rangle \phi)^\flat$ and $(\langle \alpha^* \rangle \phi)^\flat$ is the formula with the smallest such interpretation, which is all that our subsequent proofs depend on. ☐

Theorem 11 (Relative completeness)**.** *The* d$\mathcal{GL}$ *calculus is a* sound and complete axiomatization *of hybrid games relative to $L_{\mu D}$, i.e. every valid* d$\mathcal{GL}$ *formula can be derived in the* d$\mathcal{GL}$ *calculus from $L_{\mu D}$ tautologies.*

Proof. We write $\vdash_{\mathcal{D}} \phi$ to indicate that dGL formula ϕ can be derived in the dGL proof calculus from valid $L_{\mu\mathrm{D}}$ formulas. It takes a moment's thought to conclude that soundness transfers to this case from Theorem 8, so we only need to show completeness. We have to prove that every valid dGL formula ϕ can be derived from $L_{\mu\mathrm{D}}$ axioms within the dGL calculus: from $\vDash \phi$ we have to prove $\vdash_{\mathcal{D}} \phi$. The proof proceeds as follows: By propositional recombination, we inductively identify fragments of ϕ that correspond to $\phi_1 \to \langle\alpha\rangle\phi_2$ or $\phi_1 \to [\alpha]\phi_2$ logically. Then, we express subformulas ϕ_i equivalently in $L_{\mu\mathrm{D}}$ by Lemma 10 as needed, and derive these first-order Angel or Demon properties. Finally, we prove that the original dGL formula can be re-derived from the subproofs in the dGL calculus.

By appropriate propositional derivations, we can assume ϕ to be given in conjunctive normal form. We assume that negations are pushed inside over modalities using the dualities $\neg[\alpha]\phi \equiv \langle\alpha\rangle\neg\phi$ and $\neg\langle\alpha\rangle\phi \equiv [\alpha]\neg\phi$ that are provable by axiom $[\cdot]$, and that negations are pushed inside over quantifiers using provable equivalences $\neg\forall x\,\phi \equiv \exists x\,\neg\phi$ and $\neg\exists x\,\phi \equiv \forall x\,\neg\phi$. The remainder of the proof follows an induction on a well-founded partial order $\prec$ induced on dGL formulas by the lexicographic ordering of the overall structural complexity of the hybrid games in the formula and the structural complexity of the formula itself and with $L_{\mu\mathrm{D}}$ at the bottom. $L_{\mu\mathrm{D}}$ is considered first-order, thus of lowest complexity, by relativity. Well-foundedness of $\prec$ is easy to see (formally from projections into concatenations of finite trees), because the overall structural complexity of hybrid games in any particular formula can only decrease finitely often at the expense of increasing the formula complexity, which can, in turn, only decrease finitely often to result in a $L_{\mu\mathrm{D}}$ formula. The only important property for us is that, if the structure of the hybrid games in ψ is simpler than those in ϕ (somewhere simpler and nowhere worse), then $\psi \prec \phi$ even if the logical formula structure of ψ is larger than that of ϕ, e.g., when ψ has more propositional connectives, quantifiers or modalities (but of smaller overall complexity hybrid games). In the following, *IH* is short for induction hypothesis.

0. If ϕ has no hybrid games, then ϕ is a first-order formula; hence provable by assumption (even decidable [Tar51] if in first-order real arithmetic, i.e. no uninterpreted predicate symbols occur).

1. ϕ is of the form $\neg\phi_1$; then ϕ_1 is first-order, as we assumed negations to be pushed inside, so case 0 applies.

2. ϕ is of the form $\phi_1 \wedge \phi_2$, then $\vDash \phi_1$ and $\vDash \phi_2$, so individually deduce simpler proofs for $\vdash_{\mathcal{D}} \phi_1$ and $\vdash_{\mathcal{D}} \phi_2$ by IH, which combine propositionally to a proof for $\vdash_{\mathcal{D}} \phi_1 \wedge \phi_2$.

3. The case where ϕ is of the form $\forall x\,\phi_2$, $\exists x\,\phi_2$, $[\alpha]\phi_2$ or $\langle\alpha\rangle\phi_2$ is included in case 4 with $\phi_1 \equiv \mathit{false}$.

4. ϕ is a disjunction and—without loss of generality—has one of the following forms (otherwise use provable associativity and commutativity to reorder disjunction):

$$\begin{aligned}&\phi_1 \vee [\alpha]\phi_2\\&\phi_1 \vee \langle\alpha\rangle\phi_2\\&\phi_1 \vee \exists x\,\phi_2\\&\phi_1 \vee \forall x\,\phi_2.\end{aligned}$$

Let $\phi_1 \vee \langle\!\![\alpha]\!\!\rangle \phi_2$ be a unified notation for those cases. Then, $\phi_2 \prec \phi$, since ϕ_2 has less modalities or quantifiers. Likewise, $\phi_1 \prec \phi$ because $\langle\!\![\alpha]\!\!\rangle \phi_2$ contributes one modality or quantifier to ϕ that is not part of ϕ_1. By Lemma 10 there are $L_{\mu\mathrm{D}}$ formulas $\phi_1^\flat, \phi_2^\flat$ with $\vDash \phi_i \leftrightarrow \phi_i^\flat$ for $i = 1, 2$, By congruence, the validity $\vDash \phi$ yields $\vDash \phi_1^\flat \vee \langle\!\![\alpha]\!\!\rangle \phi_2^\flat$, which implies $\vDash \neg\phi_1^\flat \rightarrow \langle\!\![\alpha]\!\!\rangle \phi_2^\flat$. By induction we now derive

$$\vdash_{\mathcal{D}} \neg\phi_1^\flat \rightarrow \langle\!\![\alpha]\!\!\rangle \wedge \phi_2^\flat. \tag{4}$$

Abbreviate the $L_{\mu\mathrm{D}}$ formula $\neg\phi_1^\flat$ by F and the $L_{\mu\mathrm{D}}$ formula $\phi_2^\flat$ by G, so that we need to prove $\vdash_{\mathcal{D}} F \rightarrow \langle\!\![\alpha]\!\!\rangle G$. Observe that all subsequent proofs except for $\langle\!\![x' = \theta]\!\!\rangle$ and $\exists x$ also work without encoding when simply using ϕ_1 as F and ϕ_2 as G.

(a) If $\langle\!\![\alpha]\!\!\rangle$ is the operator $\forall x$ then $\vDash F \rightarrow \forall x\, G$, where we can assume x not to occur in F by renaming. Hence, $\vDash F \rightarrow G$. Since $G \prec \forall x\, G$, because it has less quantifiers, also $F \rightarrow G \prec F \rightarrow \forall x\, G$, hence $\vdash_{\mathcal{D}} F \rightarrow G$ is derivable by IH. Then, $\vdash_{\mathcal{D}} F \rightarrow \forall x\, G$ derives by $\forall$-generalization of first-order logic, since x does not occur in F. It is even decidable if in first-order real arithmetic [Tar51].

In the sequel, we conclude $(F \rightarrow \psi) \prec (F \rightarrow \phi)$ from $\psi \prec \phi$ without further notice.

(b) If $\langle\!\![\alpha]\!\!\rangle$ is the operator $\exists x$ then $\vDash F \rightarrow \exists x\, G$, which is first-order (i.e. in $L_{\mu\mathrm{D}}$) and, thus, provable by IH, because F, G are $L_{\mu\mathrm{D}}$ formulas. It is even decidable if in first-order real arithmetic [Tar51].

(c) $\vDash F \rightarrow \langle x' = \theta \rangle G$ is an $L_{\mu\mathrm{D}}$ formula and hence is provable by assumption, because F, G are $L_{\mu\mathrm{D}}$ formulas. Similarly for $\vDash F \rightarrow [x' = \theta]G$.

(d) $\vDash F \rightarrow \langle x' = \theta \,\&\, H \rangle G$, then this formula is, by Lemma 9, equivalent to a formula without evolution domain restrictions. Using equation (3) from the proof of Lemma 9 as a definitory abbreviation concludes this case by induction hypothesis. Similarly for $\vDash F \rightarrow [x' = \theta \,\&\, H]G$.

(e) The cases where α is of the form $x := \theta$, $?\psi$, $\beta \cup \gamma$, or $\beta; \gamma$ are consequences of the soundness of the equivalence axioms $\langle:=\rangle$,$\langle?\rangle$,$\langle\cup\rangle$,$\langle;\rangle$ plus the duals obtained via duality axiom $[\cdot]$. Whenever their respective left-hand side is valid, their right-hand side is valid and of smaller complexity (the games get simpler), and hence derivable by IH. Thus, $F \rightarrow \langle\alpha\rangle G$ derives by applying the respective axiom. We explicitly show the cases that require some extra thought.

(f) $\vDash F \rightarrow \langle x := \theta \rangle G$ implies $\vDash F \wedge y = \theta \rightarrow G_x^y$ for a fresh variable y, where G_x^y is the result of substituting y for x. Since $F \wedge y = \theta \rightarrow G_x^y \prec \langle x := \theta \rangle G$, because there are less hybrid games, $\vdash_{\mathcal{D}} F \wedge y = \theta \rightarrow G_x^y$ is derivable by IH. Hence, $\langle:=\rangle$ derives $\vdash_{\mathcal{D}} F \wedge y = \theta \rightarrow \langle x := y \rangle G$. Propositional logic derives $\vdash_{\mathcal{D}} F \rightarrow (y = \theta \rightarrow \langle x := y \rangle G)$, from which $\vdash_{\mathcal{D}} F \rightarrow \forall y\, (y = \theta \rightarrow \langle x := y \rangle G)$ derives by $\forall$-generalization of first-order logic. Since y was fresh it does not appear in θ and G, so substitution validities of first-order logic derive $\vdash_{\mathcal{D}} F \rightarrow \langle x := \theta \rangle G$. Note that direct proofs by $\langle:=\rangle$ are possible when the resulting substitution is admissible, but the substitution in G_x^y is always admissible, because it is a variable renaming replacing x by y.

(g) $\vDash F \to \langle\beta\cup\gamma\rangle G$ implies $\vDash F \to \langle\beta\rangle G \vee \langle\gamma\rangle G$. Since $\langle\beta\rangle G \vee \langle\gamma\rangle G \prec \langle\beta\cup\gamma\rangle G$, because, even if the propositional and modal structure increased, the structural complexity of hybrid games β and γ is smaller than that of $\beta\cup\gamma$ (formula G did not change), $\vdash_{\mathcal{D}} F \to \langle\beta\rangle G \vee \langle\gamma\rangle G$ is derivable by IH. Hence, $\langle\cup\rangle$ derives $\vdash_{\mathcal{D}} F \to \langle\beta\cup\gamma\rangle G$.

(h) $\vDash F \to \langle\beta;\gamma\rangle G$, which implies $\vDash F \to \langle\beta\rangle\langle\gamma\rangle G$. Since $\langle\beta\rangle\langle\gamma\rangle G \prec \langle\beta;\gamma\rangle G$, because, even if the number of modalities increased, the overall structural complexity of the hybrid games decreased because there are less sequential compositions, $\vdash_{\mathcal{D}} F \to \langle\beta\rangle\langle\gamma\rangle G$ is derivable by IH. Hence, $\vdash_{\mathcal{D}} F \to \langle\beta;\gamma\rangle G$ derives by $\langle;\rangle$.

(i) $\vDash F \to \langle\beta^d\rangle G$ implies $\vDash F \to \neg\langle\beta\rangle\neg G$, which implies $\vDash F \to [\beta]G$. Since $[\beta]G \prec \langle\beta^d\rangle G$, because β^d is more complex than β, $\vdash_{\mathcal{D}} F \to [\beta]G$ can be derived by IH. Axiom $[\cdot]$, thus, derives $\vdash_{\mathcal{D}} F \to \neg\langle\beta\rangle\neg G$, from which axiom $\langle^d\rangle$ derives $\vdash_{\mathcal{D}} F \to \langle\beta^d\rangle G$.

(j) $\vDash F \to [\beta^d]G$ implies $\vDash F \to \neg\langle\beta^d\rangle\neg G$, hence $\vDash F \to \langle\beta\rangle G$. Since $\langle\beta\rangle G \prec [\beta^d]G$, because β^d is more complex than β, $\vdash_{\mathcal{D}} F \to \langle\beta\rangle G$ can be derived by IH. Consequently, $\vdash_{\mathcal{D}} F \to \neg\neg\langle\beta\rangle\neg\neg G$ can be derived using M on $\vdash G \to \neg\neg G$. Hence, $\langle^d\rangle$ derives $\vdash_{\mathcal{D}} F \to \neg\langle\beta^d\rangle\neg G$, from which axiom $[\cdot]$ derives $\vdash_{\mathcal{D}} F \to [\beta^d]G$.

(k) $\vDash F \to [\beta^*]G$ can be derived by induction as follows. Formula $[\beta^*]G$, which expresses that Demon has a winning strategy in game β^* to satisfy G, is an inductive invariant of β^*, because $[\beta^*]G \to [\beta][\beta^*]G$ is valid, even provable by the variation $[\beta^*]G \to G \wedge [\beta][\beta^*]G$ of $\langle^*\rangle$ that can be obtained from axioms $\langle^*\rangle$ and $[\cdot]$. Thus, its equivalent $L_{\mu\mathrm{D}}$ encoding according to Lemma 10 is also an inductive invariant:

$$\varphi \equiv ([\beta^*]G)^\flat.$$

$F \to \varphi$ and $\varphi \to G$ are valid (Angel controls *) and $(F \to \varphi) \prec \phi$ and $(\varphi \to G) \prec \phi$ by encoding, hence derivable by IH. By M, $\langle^d\rangle$ and $[\cdot]$, the latter derivation $\vdash_{\mathcal{D}} \varphi \to G$ extends to $\vdash_{\mathcal{D}} [\beta^*]\varphi \to [\beta^*]G$. As above, $\varphi \to [\beta]\varphi$ is valid, and thus derivable by IH, since β has less loops. Thus, ind, which derives from FP by Lemma 7, derives $\vdash_{\mathcal{D}} \varphi \to [\beta^*]\varphi$. The above derivations combine propositionally (cut with $[\beta^*]\varphi$ and φ) to $\vdash_{\mathcal{D}} F \to [\beta^*]G$.

(l) $\vDash F \to \langle\beta^*\rangle G$. Let x the vector of free variables of $\langle\beta^*\rangle G$. Since $\langle\beta^*\rangle G$ is the least pre-fixpoint, for any dG$\mathcal{L}$ formula ψ with free variables in x:

$$\vDash \forall x\,(G \vee \langle\beta\rangle\psi \to \psi) \to (\langle\beta^*\rangle G \to \psi)$$

by a variation of the soundness argument for FP, which is also derivable by the (semantic) deduction theorem from FP. In particular, this holds for a fresh predicate symbol p with arguments x:

$$\vDash \forall x\,(G \vee \langle\beta\rangle p(x) \to p(x)) \to (\langle\beta^*\rangle G \to p(x))$$

Using $\vDash F \to \langle\beta^*\rangle G$, this implies

$$\vDash \forall x\,(G \vee \langle\beta\rangle p(x) \to p(x)) \to (F \to p(x))$$

As $\forall x\,(G \lor \langle\beta\rangle p(x) \to p(x)) \to (F \to p(x)) \prec \phi$, because, even if the formula complexity increased, the structural complexity of the hybrid games decreased, because ϕ has one more loop, so this fact is derivable by IH:

$$\vdash_{\mathcal{D}} \forall x\,(G \lor \langle\beta\rangle p(x) \to p(x)) \to (F \to p(x))$$

By uniformly substituting $\langle\beta^*\rangle G$ with free variables x for $p(x)$, US derives using $p \notin F, G$:

$$\vdash_{\mathcal{D}} \forall x\,(G \lor \langle\beta\rangle\langle\beta^*\rangle G \to \langle\beta^*\rangle G) \to (F \to \langle\beta^*\rangle G) \tag{5}$$

Yet, $\langle{}^*\rangle$ derives $\vdash G \lor \langle\beta\rangle\langle\beta^*\rangle G \to \langle\beta^*\rangle G$, from which $\vdash \forall x\,(G \lor \langle\beta\rangle\langle\beta^*\rangle G \to \langle\beta^*\rangle G)$ derives by $\forall$-generalization. Now modus ponens with (5) derives $\vdash_{\mathcal{D}} F \to \langle\beta^*\rangle G$.

This concludes the derivation of (4). Further $\vDash \phi_1 \leftrightarrow \phi_1^\flat$ implies $\vDash \neg\phi_1 \to \neg\phi_1^\flat$, which is derivable by IH, because $\phi_1 \prec \phi$. We combine $\vdash_{\mathcal{D}} \neg\phi_1 \to \neg\phi_1^\flat$ with (4) (cut with $\neg\phi_1^\flat$) to derive

$$\vdash_{\mathcal{D}} \neg\phi_1 \to \langle\!\![\alpha]\!\!\rangle \phi_2^\flat. \tag{6}$$

Likewise $\vDash \phi_2 \leftrightarrow \phi_2^\flat$ implies $\vDash \phi_2^\flat \to \phi_2$, which is derivable by IH, as $\phi_2 \prec \phi$. From $\vdash_{\mathcal{D}} \phi_2^\flat \to \phi_2$ we derive $\vdash_{\mathcal{D}} \langle\!\![\alpha]\!\!\rangle\phi_2^\flat \to \langle\!\![\alpha]\!\!\rangle\phi_2$ by M if $\langle\!\![\alpha]\!\!\rangle$ is $\langle\alpha\rangle$, by M and $\langle{}^d\rangle$ if $\langle\!\![\alpha]\!\!\rangle$ is $[\alpha]$, by $\forall$-generalization if $\langle\!\![\alpha]\!\!\rangle$ is $\forall x$, and by $\forall$-generalization and duality if $\langle\!\![\alpha]\!\!\rangle$ is $\exists x$. Finally we combine the latter derivation propositionally with (6) by a cut with $\langle\!\![\alpha]\!\!\rangle\phi_2^\flat$ to derive $\vdash_{\mathcal{D}} \neg\phi_1 \to \langle\!\![\alpha]\!\!\rangle\phi_2$, from which $\vdash_{\mathcal{D}} \phi_1 \lor \langle\!\![\alpha]\!\!\rangle\phi_2$ derives propositionally.

This completes the proof of completeness (Theorem 11). ☐

We highlight that the proof of Theorem 11 is constructive and nearly coding-free (except for $x' = \theta$, $\exists$ and $[\beta^*]$). Using US, the case for $\langle\beta^*\rangle G$ in the proof of Theorem 11 reveals an explicit $^\flat$-free reduction to a dG$\mathcal{L}$ formula with less loops, which can be considered a modal analogue of characterizations in the Calculus of Constructions [CH88]. These two observations easily reprove a classical result of Meyer and Halpern [MH82] about the semidecidability of termination assertions (logical formulas $F \to \langle\alpha\rangle G$ of uninterpreted dynamic logic with first-order F, G and regular programs α without differential equations). In fact, this proves a slightly stronger result about dynamic logic without any $[\alpha]\cdot$ with loops [Sch84], yet still without $\exists$. Theorem 11 shows that this result continues to hold for uninterpreted first-order game logic in the fragment where * only occurs with even ^{d}-polarity in $\langle\alpha\rangle$ and only of odd ^{d}-polarity in $[\alpha]$ (the conditions on tests in α are accordingly).

The completeness proof indicates a coding-free way of proving Angel properties $\langle\beta^*\rangle G$ that works efficiently in practice. Illustrative examples are shown in Appendix A. In particular, dG$\mathcal{L}$ does not need Harel's convergence rule [HMP77] for completeness and, thus, neither does logic for hybrid systems, even though it was previously based on it [Pla12a]. These results correspond to a hybrid reading of influential views of understanding program invariants as fixpoints [CC77, Cla79].

The coding-free constructive nature of Theorem 11 characterizes exactly which part of hybrid games proving is difficult: finding computationally succinct weaker invariants for $[\alpha^*]G$ and finding succinct differential (in)variants [Pla10a] for $[x' = \theta]$ and $\langle x' = \theta\rangle$ of which a solution is a special case [Pla12c]. The case $\exists x\, G$ is interesting in that a closer inspection of Theorem 11 reveals that its complexity depends on whether it supports Herbrand disjunctions, which is the case for uninterpreted first-order logic and first-order real arithmetic [Tar51], but not for $G \equiv [\alpha^*]\psi$, which already gives $\exists x\, G$ the full Π^1_1-complete complexity even for classical dynamic logic [HKT00].

An interesting question is whether dG$\mathcal{L}$ is complete relative to smaller logics, which Theorem 11 reduces to a study of expressing (two-variable) $L_{\mu\mathrm{D}}$. This gives hybrid versions of Parikh's completeness results for fragments of game logic [Par83].

Corollary 12 (Relative completeness of *-free dG$\mathcal{L}$)**.** *The* dG$\mathcal{L}$ *calculus is a* sound and complete axiomatization *of* **-free hybrid games relative to* d$\mathcal{L}$*.*

Proof. Lemma 10 reduces to d$\mathcal{L}$, even the first-order logic of differential equations [Pla12a], for *-free hybrid games. □

Corollary 13 (Relative completeness of ^{d}-free dG$\mathcal{L}$)**.** *The* dG$\mathcal{L}$ *calculus is a* sound and complete axiomatization *of* ^{d}*-free hybrid games relative to* d$\mathcal{L}$*.*

Proof. ^{d}-free loops are Scott-continuous by Lemma 3, so have closure ordinal ω and are, thus, equivalent to their d$\mathcal{L}$ form, and even expressible in the first-order logic of differential equations by [Pla12a, Theorem 9]. □

By Corollary 13, d$\mathcal{L}$ is relatively complete without the convergence rule that had been used before [Pla08]. In combination with the first and second relative completeness theorems of d$\mathcal{L}$ [Pla12a], it follows that the dG$\mathcal{L}$ calculus is a sound and complete axiomatization of *-free hybrid games and of ^{d}-free hybrid games relative to the first-order logic of differential equations. When adding the numerical Euler integration axiom [Pla12a], both are sound and complete axiomatizations of those classes of hybrid games relative to discrete dynamic logic [Pla12a]. Similar completeness results for dG$\mathcal{L}$ relative to d$\mathcal{L}$, and, thus, relative to first-order logic of differential equations, follow from Theorem 11, e.g., for the case of hybrid games with winning regions that are finite rank Borel sets.

As a corollary to Theorem 11 and an equi-expressibility result [Pla12a], dG$\mathcal{L}$ is complete with the Euler axiom relative to (first-order) discrete μ-calculus interpreted over the reals. This aligns the discrete and the continuous side of hybrid games in a constructive provably equivalent way similar to corresponding results about hybrid systems [Pla12a]. Yet, the interactivity of two-variable fixpoints still seems to stay.

4.4 Separating Axioms

In order to illustrate how and why dG$\mathcal{L}$ differs from differential dynamic logic d$\mathcal{L}$ [Pla08, Pla12a], i.e. how reasoning about hybrid games differs from reasoning about hybrid systems, we identify separating axioms, that is, axioms of d$\mathcal{L}$ that do not hold in dG$\mathcal{L}$. We investigate the difference in terms of important classes of modal logics; recall [HC96] or Appendix B.

Theorem 14. dG$\mathcal{L}$ *is a subregular, sub-Barcan, monotonic modal logic without the induction axiom of dynamic logic.*

The proof of Theorem 14 is in Appendix B, where, for each separating axiom, we give simple counterexamples illustrating what makes hybrid games different than hybrid systems.

Note that Harel's convergence rule is not a separating axiom, because it is sound for $\mathsf{dG}\mathcal{L}$, just unnecessary. Furthermore, in light of Theorem 6, it is questionable whether the convergence rule would even be relatively complete for hybrid games.

5 Related Work

Discrete games and the interaction of games and logic for various purposes have been studied with much success[vNM55, Par85, Aum95, HS97, Sti01, AHK02, PP03, Ber03, CHP07, BGL07, BP09, AG11, Vää11]. Propositional game logic [PP03] subsumes ΔPDL and CTL*. After more than two decades, it has been shown that the alternation hierarchy in propositional game logic is strict and encodes parity games that span the full alternation hierarchy of the (propositional) modal μ-calculus [Ber03] and that, being in the two variable fragment, it is less expressive than the (propositional) modal μ-calculus [BGL07]. Another influential propositional modal logic, ATL* has been used for model checking[AHK02] and is related to propositional game logic [BP09]. Applications and relations of game logic, ATL* [AHK02], and strategy logic [CHP07] have been discussed in the literature [AHK02, PP03, CHP07, BP09]. These logics for the propositional case are interesting, but it is not clear how their decision procedures should be generalized to the highly undecidable domain of hybrid games with differential equations, uncountable choices, and higher closure ordinals. The logic $\mathsf{dG}\mathcal{L}$ shows how such hybrid games can be proved and enjoys completeness.

Differential games have been studied with many different notions of solutions [Isa67, Fri71, Pet93, Bre10]. They are of interest when actions are in continuous time. We look at the complementary model of hybrid games where the underlying system is that of a hybrid system with interacting discrete and continuous dynamics, but the game actions are chosen at discrete instants of time, even if they take effect in continuous time.

Reachability aspects of games for hybrid systems have been studied before. A game view on hybrid systems verification has been proposed following a Hamilton-Jacobi-Bellman PDE formulation [TMBO03, MBT05], with subsequent extensions by Gao et al. [GLQ07]. Their primary focus is on adversarial choices in the continuous dynamics, which is very interesting, but not what we consider here. Axioms of a proof calculus are easier to get sound than numerical approximations of PDEs, which is an interesting but extremely challenging problem [Pla12a]. WCTL properties of STORMED hybrid games, which are restricted to evolve linearly in one "direction" all the time, have been shown to be decidable using bisimulation quotients [VPVD11]. STORMED hybrid games generalize o-minimal hybrid games which have been shown to be decidable earlier [BBC10]. The case of rectangular hybrid games is known to be decidable [HHM99]. Not all applications fall into decidable classes [QP12], so that a study of more general hybrid games is called for.

We take a complementary view and study logics and proofs for hybrid games instead of searching for decidable fragments using bisimulation quotients [HHM99, BBC10, VPVD11], which do not generally exist. We provide a proof-based verification technique for more general hybrid games with nonlinear dynamics. Our notion of hybrid games has more flexible nested hybrid game choices for the agents than the fixed controller-plant interaction considered in related work. We consider more general logical formulas with nested modal game operators. We do not consider concurrent games [BBC10], though, only sequential games.

There is more than one way how logic can be used to understand games of hybrid systems. Games can be added as separate constructs on top of unmodified differential dynamic logic [QP12], which focuses on the special case of advance notice semantics. We follow a different principle here. Instead of leaving differential dynamic logic untouched and adding several separate game constructs on top of full hybrid systems reachability operators as in [QP12], we modify the logic to be a game logic by adding a single operator ^{d} into the system dynamics. Our logic dG$\mathcal{L}$ results in a much simplified but nevertheless more general logic with a simpler more general semantics (and not restricted to advance notice) and simpler and more general calculus. We consider a Hilbert calculus and focus on fundamental logical properties instead of automation. For practical aspects like sequent calculus automation and a very challenging robotic factory automation case study that translates to dG$\mathcal{L}$, we refer to [QP12]. What is more difficult in dG$\mathcal{L}$ in comparison to that fragment [QP12], however, is the need to carefully identify which axioms are no longer sound for games, which is what we have pursued in Section 4.4.

The logic dG$\mathcal{L}$ we present here has some similarity with stochastic differential dynamic logic (Sd$\mathcal{L}$) [Pla11], because both may be used to verify properties of the hybrid system dynamics with partially uncertain behavior. Both approaches do, however, address uncertainty in fundamentally different ways. Sd$\mathcal{L}$ takes a probabilistic perspective on uncertainty in the system dynamics. The dG$\mathcal{L}$ approach put forth in this paper, instead, takes an adversarial perspective on uncertainty. Both views on how to handle uncertain behavior are useful but serve different purposes, depending on the nature of the system analysis question at hand. A probabilistic understanding of uncertainty can be superior whenever good information is available about the distribution of choices made by the environment and other agents. Whenever that is not possible, adversarial views may be more appropriate, since they do not lead to the inadequate biases that arbitrary probabilistic assumptions would impose.

6 Conclusions and Future Work

We have introduced differential game logic (dG$\mathcal{L}$) for hybrid games, which unifies logic of hybrid systems with Parikh's game logic. Despite the challenges of hybrid games like higher closure ordinals of winning regions, dG$\mathcal{L}$ has a simple modal semantics and a simple proof calculus, which we prove to be a sound and complete axiomatization of hybrid games relative to the fixpoint logic of differential equations. Combining dG$\mathcal{L}$ with axioms for differential equations [Pla10a, Pla12a] provides a way of handling hybrid games with nonlinear differential equations.

Our completeness proof is constructive and nearly coding-free, thereby exactly characterizing the difficult parts of hybrid games proving. The proof identifies an efficient fixpoint-style proof technique, which can be considered a modal analogue of characterizations in the Calculus of Constructions [CH88], and relates to hybrid versions of influential views of understanding program invariants as fixpoints [CC77, Cla79].

The relative completeness results show that dG$\mathcal{L}$ has all axioms for dealing with hybrid games. The study of (fragments of) dG$\mathcal{L}$ which are complete for smaller logics is interesting future work and, by completeness, reduces to questions of expressiveness that gives rise to interesting problems in descriptive set theory. The relation of the expressiveness of d$\mathcal{L}$ and dG$\mathcal{L}$ is an open question related to open questions for the propositional case [Par85, BGL07], yet partially characterized by smaller logics according to our expressibility results.

We observe that there is a striking similarity of the dG$\mathcal{L}$ calculus with the calculus for stochastic differential dynamic logic Sd$\mathcal{L}$ [Pla11], despite their fundamentally different semantical presuppositions (adversarial nondeterminism versus stochasticity), which indicates the existence of a deeper logical connection relating stochastic and adversarial uncertainty.

Acknowledgment

I thank Stephen Brookes, Frank Pfenning, and James Cummings for helpful discussions.

References

[AG11] Krzysztof R. Apt and Erich Grädel. *Lectures in Game Theory for Computer Scientists*. Cambridge Univ. Press, 2011.

[AHK02] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.

[Aum95] Robert Aumann. Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 8(1):263–300, 1995.

[Bar46] Ruth C. Barcan. A functional calculus of first order based on strict implication. *J. Symb. Log.*, 11(1):1–16, 1946.

[BBC10] Patricia Bouyer, Thomas Brihaye, and Fabrice Chevalier. O-minimal hybrid reachability games. *Logical Methods in Computer Science*, 6(1), 2010.

[Ber03] Dietmar Berwanger. Game logic is strong enough for parity games. *Studia Logica*, 75(2):205–219, 2003.

[BGL07] Dietmar Berwanger, Erich Grädel, and Giacomo Lenzi. The variable hierarchy of the μ-calculus is strict. *Theory Comput. Syst.*, 40(4):437–466, 2007.

[BP09] Dietmar Berwanger and Sophie Pinchinat. Game quantification patterns. In Ramaswamy Ramanujam and Sundar Sarukkai, editors, *ICLA*, volume 5378 of *LNCS*, pages 116–130. Springer, 2009.

[Bre10] Alberto Bressan. Noncooperative differential games. a tutorial, 2010.

[CC77] Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, pages 238–252, 1977.

[CH88] Thierry Coquand and Gérard P. Huet. The calculus of constructions. *Inf. Comput.*, 76(2/3):95–120, 1988.

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

[CHP07] Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Strategy logic. In Luís Caires and Vasco Thudichum Vasconcelos, editors, *CONCUR*, volume 4703 of *LNCS*, pages 59–73. Springer, 2007.

[Chu56] Alonzo Church. *Introduction to Mathematical Logic, Volume I*. Princeton University Press, Princeton, NJ, 1956.

[Cla79] Edmund M. Clarke. Program invariants as fixedpoints. *Computing*, 21(4):273–294, 1979.

[Coo78] Stephen A. Cook. Soundness and completeness of an axiom system for program verification. *SIAM J. Comput.*, 7(1):70–90, 1978.

[Ehr61] Andrzej Ehrenfeucht. An application of games to the completeness problem for formalized theories. *Fundam. Math.*, 49:129–141, 1961.

[Fri71] Avner Friedman. *Differential Games*. John Wiley, 1971.

[GLQ07] Y. Gao, J. Lygeros, and M. Quincampoix. On the reachability problem for uncertain hybrid systems. *IEEE T. Automat. Contr.*, 52(9):1572–1586, September 2007.

[GS53] David Gale and F. M. Stewart. Infinite games with perfect information. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games II*, volume 28 of *Ann. of Math. Stud.*, pages 246–266. Princeton Univ. Press, 1953.

[HC96] G. E. Hughes and M. J. Cresswell. *A New Introduction to Modal Logic*. Routledge, 1996.

[Hen96] Thomas A. Henzinger. The theory of hybrid automata. In *LICS*, pages 278–292, Los Alamitos, 1996. IEEE Computer Society.

[HHM99] Thomas A. Henzinger, Benjamin Horowitz, and Rupak Majumdar. Rectangular hybrid games. In Jos C. M. Baeten and Sjouke Mauw, editors, *CONCUR*, volume 1664 of *LNCS*, pages 320–335. Springer, 1999.

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

[HMP77] David Harel, Albert R. Meyer, and Vaughan R. Pratt. Computability and completeness in logics of programs (preliminary report). In *STOC*, pages 261–268. ACM, 1977.

[HS97] Jaakko Hintikka and Gabriel Sandu. Game-theoretical semantics. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Language*, pages 361–410. Elsevier, 1997.

[Isa67] Rufus Philip Isaacs. *Differential Games*. John Wiley, 1967.

[Koz83] Dexter Kozen. Results on the propositional μ-calculus. *Theor. Comput. Sci.*, 27:333–354, 1983.

[Lub89] Robert S. Lubarsky. μ-definable sets of integers. In *LICS*, pages 343–352, 1989.

[Mar75] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.

[MBT05] Ian Mitchell, Alexandre M. Bayen, and Claire Tomlin. A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games. *IEEE T. Automat. Contr.*, 50(7):947–957, 2005.

[MH82] Albert R. Meyer and Joseph Y. Halpern. Axiomatic definitions of programming languages: A theoretical assessment. *J. ACM*, 29(2):555–576, 1982.

[Nas51] John Nash. Non-cooperative games. *Ann. Math.*, 54(2):286–295, 1951.

[Par83] Rohit Parikh. Propositional game logic. In *FOCS*, pages 195–200. IEEE, 1983.

[Par85] Rohit Parikh. The logic of games and its applications. *Annals of Discrete Mathematics*, 24:111–140, 1985. M. Karpinski and J. van Leeuwen, eds., Topics in the Theory of Computation.

[Pet93] Leon A. Petrosjan. *Differential Games of Pursuit*. World Scientific, 1993.

[PHP01] Maria Prandini, João P. Hespanha, and George J. Pappas. Greedy control for hybrid pursuit games. In *ECC*, pages 2621–2626, 2001.

[Pla08] André Platzer. Differential dynamic logic for hybrid systems. *J. Autom. Reas.*, 41(2):143–189, 2008.

[Pla10a] André Platzer. Differential-algebraic dynamic logic for differential-algebraic programs. *J. Log. Comput.*, 20(1):309–352, 2010.

[Pla10b] André Platzer. *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer, Heidelberg, 2010.

[Pla11] André Platzer. Stochastic differential dynamic logic for stochastic hybrid programs. In Nikolaj Bjørner and Viorica Sofronie-Stokkermans, editors, *CADE*, volume 6803 of *LNCS*, pages 431–445. Springer, 2011.

[Pla12a] André Platzer. The complete proof theory of hybrid systems. In *LICS*, pages 541–550. IEEE, 2012.

[Pla12b] André Platzer. Differential game logic for hybrid games. Technical Report CMU-CS-12-105, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, March 2012.

[Pla12c] André Platzer. The structure of differential invariants and differential cut elimination. *Logical Methods in Computer Science*, 8(4):1–38, 2012.

[PP03] Marc Pauly and Rohit Parikh. Game logic - an overview. *Studia Logica*, 75(2):165–182, 2003.

[Pra76] Vaughan R. Pratt. Semantical considerations on Floyd-Hoare logic. In *FOCS*, pages 109–121. IEEE, 1976.

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

[Sch84] Peter H. Schmitt. Diamond formulas: A fragment of dynamic logic with recursively enumerable validity problem. *Information and Control*, 61(2):147–158, 1984.

[Sti01] Colin Stirling. *Modal and Temporal Properties of Processes*. Springer, 2001.

[Tar51] Alfred Tarski. *A Decision Method for Elementary Algebra and Geometry*. University of California Press, Berkeley, 2nd edition, 1951.

[TLS00] Claire J. Tomlin, John Lygeros, and Shankar Sastry. A game theoretic approach to controller design for hybrid systems. *Proc. IEEE*, 88(7):949–970, 2000.

[TMBO03] Claire Tomlin, Ian Mitchell, Alexandre M. Bayen, and Meeko Oishi. Computational techniques for the verification of hybrid systems. *Proc. IEEE*, 91(7):986–1001, 2003.

[TPS98] Claire Tomlin, George J. Pappas, and Shankar Sastry. Conflict resolution for air traffic management: a study in multi-agent hybrid systems. *IEEE T. Automat. Contr.*, 43(4):509–521, 1998.

[Vää11] Jouko Väänänen. *Models and Games*. Cambridge Univ. Press, 2011.

[vNM55] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton Univ. Press, 3rd edition, 1955.

[VPVD11] Vladimeros Vladimerou, Pavithra Prabhakar, Mahesh Viswanathan, and Geir E. Dullerud. Specifications for decidable hybrid games. *Theor. Comput. Sci.*, 412(48):6770–6785, 2011.

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

[Zer13] Ernst Friedrich Fernando Zermelo. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. In *Fifth International Congress of Mathematicians*, volume II, pages 501–504, 1913.

A Example dGL Proofs from Completeness Result

The completeness proof suggests the use of $\langle{}^*\rangle$ and US to prove $\langle\alpha^*\rangle$ properties. We show how easy this is in practice. Simple constructions and arithmetic close each of the following examples. Observe how logic programming style saturation with widening quickly proves the resulting arithmetic here.

Example 3. The simple non-game dGL formula

$$x \geq 0 \rightarrow \langle (x := x-1)^* \rangle 0 \leq x < 1$$

is provable as shown in Fig. 4, where $\langle\alpha^*\rangle 0 \leq x < 1$ is short for $\langle (x := x-1)^* \rangle 0 \leq x < 1$.

$$\begin{array}{l r}
 & * \\ \hline
\mathbb{R} & \forall x\,(0 \leq x < 1 \lor p(x-1) \rightarrow p(x)) \rightarrow (x \geq 0 \rightarrow p(x)) \\ \hline
\langle{:=}\rangle & \forall x\,(0 \leq x < 1 \lor \langle x := x-1\rangle p(x) \rightarrow p(x)) \rightarrow (x \geq 0 \rightarrow p(x)) \\ \hline
\text{US} & \forall x\,(0 \leq x < 1 \lor \langle x := x-1\rangle\langle\alpha^*\rangle 0 \leq x < 1 \rightarrow \langle\alpha^*\rangle 0 \leq x < 1) \rightarrow (x \geq 0 \rightarrow \langle\alpha^*\rangle 0 \leq x < 1) \\ \hline
\langle{}^*\rangle,\forall & \forall x\,(0 \leq x < 1 \lor \langle x := x-1\rangle\langle\alpha^*\rangle 0 \leq x < 1 \rightarrow \langle\alpha^*\rangle 0 \leq x < 1) \\ \hline
\text{MP} & x \geq 0 \rightarrow \langle\alpha^*\rangle 0 \leq x < 1
\end{array}$$

Figure 4: dGL Angel proof for Example 3 using technique from completeness proof

Example 4. The dGL formula

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

which comes from (9) on p. 43 is provable as shown in Fig. 5, where $\beta \cap \gamma$ is short for $x := a; a := 0 \cap x := 0$ and $\langle (\beta \cap \gamma)^* \rangle x \neq 1$ short for $\langle (x := a; a := 0 \cap x := 0)^* \rangle x \neq 1$:

$$\begin{array}{l r}
 & * \\ \hline
\mathbb{R} & \forall x\,(x \neq 1 \lor p(a,0) \land p(0,a) \rightarrow p(x,a)) \rightarrow (\mathit{true} \rightarrow p(x,a)) \\ \hline
\langle;\rangle,\langle{:=}\rangle & \forall x\,(x \neq 1 \lor \langle\beta\rangle p(x,a) \land \langle\gamma\rangle p(x,a) \rightarrow p(x,a)) \rightarrow (\mathit{true} \rightarrow p(x,a)) \\ \hline
\langle\cup\rangle,\langle{}^d\rangle & \forall x\,(x \neq 1 \lor \langle\beta \cap \gamma\rangle p(x,a) \rightarrow p(x,a)) \rightarrow (\mathit{true} \rightarrow p(x,a)) \\ \hline
\text{US} & \forall x\,(x \neq 1 \lor \langle\beta \cap \gamma\rangle\langle(\beta \cap \gamma)^*\rangle x \neq 1 \rightarrow \langle(\beta \cap \gamma)^*\rangle x \neq 1) \rightarrow (\mathit{true} \rightarrow \langle(\beta \cap \gamma)^*\rangle x \neq 1) \\ \hline
\langle{}^*\rangle,\forall,\text{MP} & \mathit{true} \rightarrow \langle(\beta \cap \gamma)^*\rangle x \neq 1 \\ \hline
\mathbb{R} & x = 1 \land a = 1 \rightarrow \langle(\beta \cap \gamma)^*\rangle x \neq 1
\end{array}$$

Figure 5: dGL Angel proof for Example 4 using technique from completeness proof

Example 5. The dGL formula

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

which comes from (10) on p. 45 is provable as shown in Fig. 6, where the notation $\langle(\beta \cup \gamma)^*\rangle 0 \leq x < 1$ is short for $\langle (x := 1; x' = 1^d \cup x := x-1)^* \rangle 0 \leq x < 1$: The proof steps for β use in $\langle{}'\rangle$ that

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{*}{\forall x\,(0 \le x < 1 \vee \forall t{\ge}0\, p(0+t) \vee p(x-1) \to p(x)) \to (\mathit{true} \to p(x))}\,\mathbb{R}}{\forall x\,(0 \le x < 1 \vee \langle x := 1\rangle \neg \exists t{\ge}0\, \langle x := x+t\rangle \neg p(x) \vee p(x-1) \to p(x)) \to (\mathit{true} \to p(x))}\,\langle:=\rangle}{\forall x\,(0 \le x < 1 \vee \langle x := 1\rangle \neg \langle x' = 1\rangle \neg p(x) \vee p(x-1) \to p(x)) \to (\mathit{true} \to p(x))}\,\langle'\rangle}{\forall x\,(0 \le x < 1 \vee \langle\beta\rangle p(x) \vee \langle\gamma\rangle p(x) \to p(x)) \to (\mathit{true} \to p(x))}\,\langle;\rangle,\langle^{d}\rangle}{\forall x\,(0 \le x < 1 \vee \langle\beta \cup \gamma\rangle p(x) \to p(x)) \to (\mathit{true} \to p(x))}\,\langle\cup\rangle}{\forall x\,(0 \le x < 1 \vee \langle\beta \cup \gamma\rangle\langle(\beta \cup \gamma)^*\rangle 0 \le x < 1 \to \langle(\beta \cup \gamma)^*\rangle 0 \le x < 1) \to (\mathit{true} \to \langle(\beta \cup \gamma)^*\rangle 0 \le x < 1)}\,\text{US}}{\mathit{true} \to \langle(\beta \cup \gamma)^*\rangle 0 \le x < 1}\,\langle^*\rangle,\forall,\text{MP}$$

Figure 6: dG$\mathcal{L}$ Angel proof for Example 5 using technique from completeness proof

$t \mapsto x + t$ is the solution of the differential equation, so the subsequent use of $\langle:=\rangle$ substitutes 1 in to obtain $t \mapsto 0 + t$. Recall that the winning regions for formula (10) need $> \omega$ iterations to converge. It is still provable easily. A variation of this proof shows dG$\mathcal{L}$ formula (1) from p. 7, where the handling of the nonlinear differential equation is a bit more complicated.

B Proof of Separating Axioms

This section shows a proof of Theorem 14 with an emphasis on simple counterexamples for each separating axiom.

B.0.1 Subnormal Modal Logic

First, we show that, unlike d$\mathcal{L}$, dG$\mathcal{L}$ is not a normal modal logic [HC96]. Axiom K, the modal modus ponens from normal modal logic [HC96], dynamic logic [Pra76], and differential dynamic logic [Pla12a], i.e.

$$[\alpha](\phi \to \psi) \to ([\alpha]\phi \to [\alpha]\psi)$$

is not sound for dG$\mathcal{L}$ as witnessed using the choice $\alpha \equiv (x := 1 \cap x := 0); y := 0$ and $\phi \equiv x = 1$, $\psi \equiv y = 1$; see Fig. 7. The global version of K, i.e. the implicative version of Gödel's

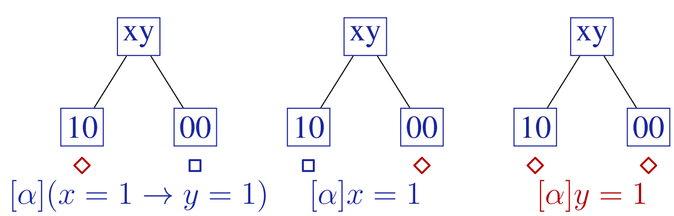

Figure 7: Game trees for counterexample to axiom K using $\alpha \equiv (x := 1 \cap x := 0); y := 0$.

generalization rule is still sound and derives with $\langle^{d}\rangle$ and $[\cdot]$ from M using $\alpha \equiv \beta^{d}$

$$\frac{\phi \to \psi}{[\beta]\phi \to [\beta]\psi}$$

The normal Gödel generalization rule G, i.e.

$$\frac{\phi}{[\alpha]\phi}$$

however, is not sound for dG$\mathcal{L}$ as witnessed by the choice $\alpha \equiv (?\mathit{false})^d, \phi \equiv \mathit{true}$.

B.0.2 Subregular Modal Logic

Regular modal logics are monotonic modal logics [Che80] that are weaker than normal modal logics. But the regular modal generalization rule [Che80], i.e.

$$\frac{\phi_1 \land \phi_2 \to \psi}{[\alpha]\phi_1 \land [\alpha]\phi_2 \to [\alpha]\psi}$$

is not sound for dG$\mathcal{L}$ either as witnessed by the choice $\alpha \equiv (x := 1 \cap x := 0); y := 0$, $\phi_1 \equiv x = 1, \phi_2 \equiv x = y, \psi \equiv x = 1 \land x = y$; see Fig. 8.

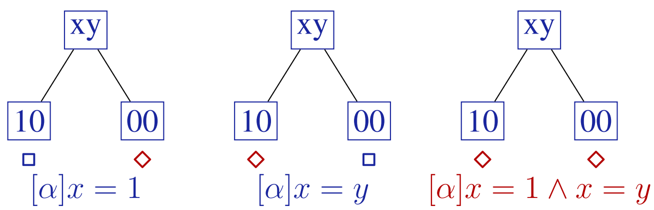

Figure 8: Game trees for counterexample to regular modal rule using $\alpha \equiv (x := 1 \cap x := 0); y := 0$.

B.0.3 Monotonic Modal Logic

The axiom that is closest to K but still sound for dG$\mathcal{L}$ is a monotonicity axiom. This axiom is sound for dG$\mathcal{L}$, yet already included in the monotonicity rule M:

Lemma 15 ([Che80, Theorem 8.13])**.** *In the presence of rule RE from p. 16, rule M is interderivable with axiom M:*

$$\langle\alpha\rangle\phi \lor \langle\alpha\rangle\psi \to \langle\alpha\rangle(\phi \lor \psi)$$

Proof. Axiom M derives from rule M: From $\phi \to \phi \lor \psi$, M derives $\langle\alpha\rangle\phi \to \langle\alpha\rangle(\phi \lor \psi)$. From $\psi \to \phi \lor \psi$, M derives $\langle\alpha\rangle\psi \to \langle\alpha\rangle(\phi \lor \psi)$, from which propositional logic yields $\langle\alpha\rangle\phi \lor \langle\alpha\rangle\psi \to \langle\alpha\rangle(\phi \lor \psi)$.

Conversely, rule M derives from axiom M and rule RE: From $\phi \to \psi$ propositional logic derives $\phi \lor \psi \leftrightarrow \psi$, from which RE derives $\langle\alpha\rangle(\phi \lor \psi) \leftrightarrow \langle\alpha\rangle\psi$. From axiom M, propositional logic, thus, derives $\langle\alpha\rangle\phi \to \langle\alpha\rangle\psi$. □

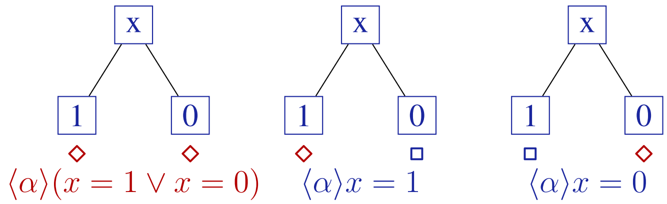

Figure 9: Game trees for counterexample to converse monotone axiom using $\alpha \equiv x := 1 \cap x := 0$.

The converse of axiom M is sound for d$\mathcal{L}$ but not for dG$\mathcal{L}$, however, as witnessed by $\alpha \equiv x := 1 \cap x := 0$, $\phi \equiv x = 1, \psi \equiv x = 0$; see Fig. 9:

$$\langle\alpha\rangle(\phi \vee \psi) \rightarrow \langle\alpha\rangle\phi \vee \langle\alpha\rangle\psi$$

The presence of the regular congruence rule RE and the fact that $[\alpha]\phi \leftrightarrow \neg\langle\alpha\rangle\neg\phi$ still make dG$\mathcal{L}$ a classical modal logic [Che80]. Rule M even makes dG$\mathcal{L}$ a monotone modal logic [Che80].

B.0.4 Sub-Barcan

The most important axioms about the interaction of quantifiers and modalities in first-order modal logic are the Barcan and converse Barcan axioms [Bar46], which, together, characterize constant domain in normal first-order modal logics [HC96]. The Barcan axiom B, which characterizes anti-monotonic domains in first-order modal logic [HC96], is sound for constant-domain first-order dynamic logic and for differential dynamic logic d$\mathcal{L}$ when x does not occur in α [Pla12a]:

$$\langle\alpha\rangle\exists x\,\phi \rightarrow \exists x\,\langle\alpha\rangle\phi \qquad (x \notin \alpha)$$

but the Barcan axiom is not sound for dG$\mathcal{L}$ as witnessed by the choice $\alpha \equiv y := y + 1^{\times}$ or $\alpha \equiv y' = 1^{d}$ and $\phi \equiv x \geq y$. The equivalent Barcan formula

$$\forall x\,[\alpha]\phi \rightarrow [\alpha]\forall x\,\phi \qquad (x \notin \alpha)$$

is not sound for dG$\mathcal{L}$ as witnessed by the choice $\alpha \equiv y := y + 1^{\times}$ or $\alpha \equiv y' = 1^{d}$ and $\phi \equiv y \geq x$. The converse Barcan formula of first-order modal logic, which characterizes monotonic domains [HC96], is sound for dG$\mathcal{L}$ and can be derived[5] when x does not occur in α:

$$(\overleftarrow{\mathrm{B}}) \quad \exists x\,\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\exists x\,\phi \quad \text{where } x \notin \alpha$$

B.0.5 No Induction Axiom

The induction axiom

$$[\alpha^*](\phi \rightarrow [\alpha]\phi) \rightarrow (\phi \rightarrow [\alpha^*]\phi) \tag{7}$$

holds for d$\mathcal{L}$, but, unlike induction rule ind, does not hold for dG$\mathcal{L}$ as witnessed by the choice $\alpha^* \equiv ((x := a; a := 0) \cap x := 0)^*$ and $\phi \equiv x = 1$; see Fig. 10. Note that the failure of the induction

[5]From $\phi \rightarrow \exists x\,\phi$, derive $\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\exists x\,\phi$ by M, from which first-order logic derives $\forall x\,(\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\exists x\,\phi)$ and then derives $\exists x\,\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\exists x\,\phi$, as x is not free in the succedent.

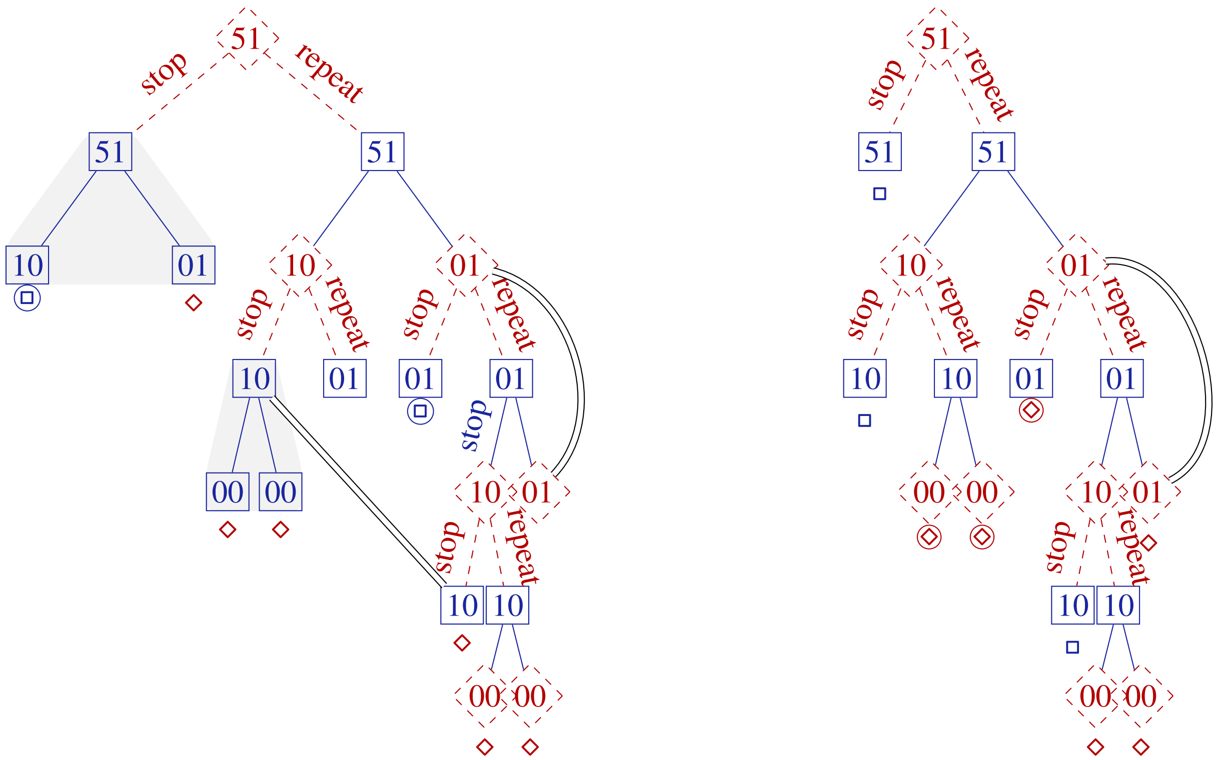

Figure 10: Game trees for counterexample to induction axiom (notation: x, a) with game $\alpha \equiv (x := a; a := 0) \cap x := 0$. **(top)** $[\alpha^*](x = 1 \rightarrow [\alpha]x = 1)$ is true by the strategy "if Angel chose stop, choose $x := a; a := 0$, otherwise always choose $x := 0$" **(bottom)** $[\alpha^*]x = 1$ is false by strategy "repeat once and repeat once more if $x = 1$, then stop." If a winning state can be reached by a winning strategy, we enclose the mark in a circle ◈ or ⊡, respectively.

axiom in the counterexample for (7) hinges on the fact that Angel is free to decide whether or not to repeat α after each round depending on the state. This would be different if we had chosen an advance notice semantics for α^*; see Appendix D. By a variation of the soundness argument for FP, it can be shown, however, that a variation of the induction axiom is still sound if we translate the induction rule ind into an axiom using the universal closure, denoted $\mathrm{Cl}_\forall$, with respect to all variables bound in α:

$$\mathrm{Cl}_\forall(\phi \to [\alpha]\phi) \to (\phi \to [\alpha^*]\phi)$$

This trick with the universal closure does not work for the dual of the induction axiom, which is called first arrival axiom. The first arrival axiom, $\langle\alpha^*\rangle\phi \to \phi \vee \langle\alpha^*\rangle(\neg\phi \wedge \langle\alpha\rangle\phi)$, which holds for d$\mathcal{L}$, expresses that, if ϕ holds after a repetition of α, then it either holds right away or α can be repeated so that ϕ does not hold yet but can hold after one more repetition [PP03]. This axiom does not hold, however, for dG$\mathcal{L}$ as witnessed by $\alpha^* \equiv ((x := x - y \cap x := 0); y := x)^*$ and $\phi \equiv x = 0$, since two iterations surely yield $x = 0$, but one iteration may or may not yield $x = 0$, depending on Demon's choice; see Fig. 11.

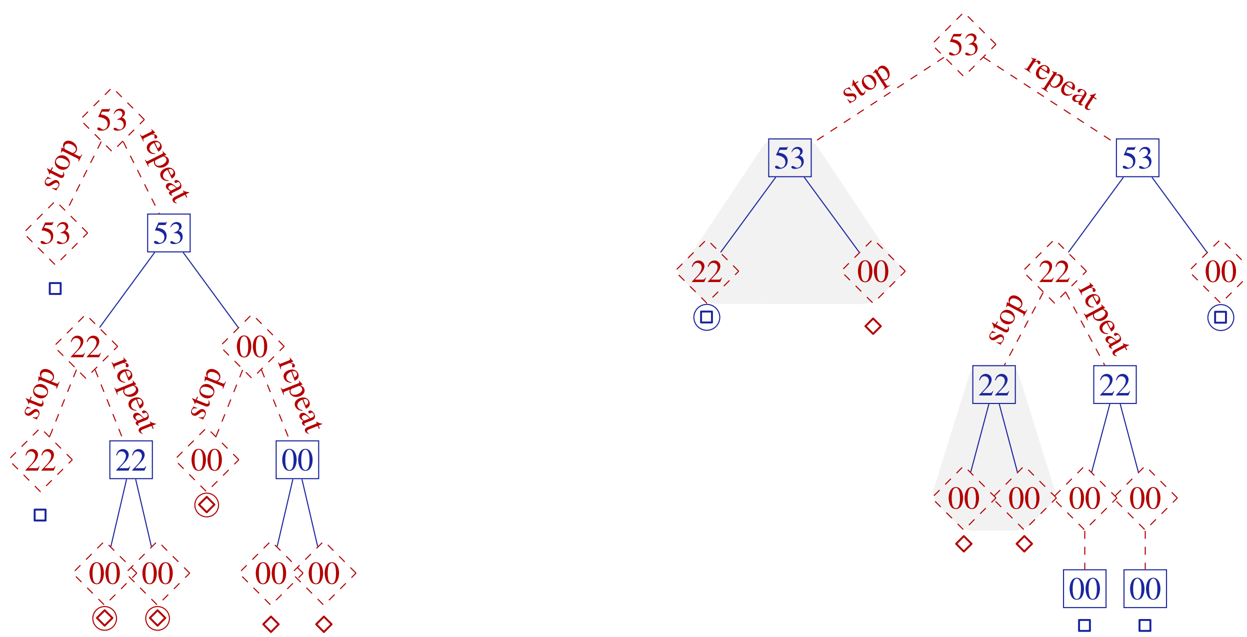

Figure 11: Game trees for counterexample to first arrival axiom (notation: x, y) with game $\alpha \equiv (x := x - y \cap x := 0); y := x$. **(top)** $\langle\alpha^*\rangle x = 0$ is true no matter which choices Demon makes **(bottom)** $\langle\alpha^*\rangle(x \neq 0 \wedge \langle\alpha\rangle x = 0)$ is false, because stop can be defeated by $x := x - y$ and repeat can be defeated by $x := 0$.

C Operational Game Semantics

In order to relate the intuition of interactive game play to the denotational semantics of hybrid games, we show an operational semantics for hybrid games that is more complicated than the modal semantics from Section 2.2 but makes strategies explicit and more directly reflects the intuition how hybrid games are played successively. The modal semantics is beneficial, because it is simpler. The results in this section are not needed in the rest of the paper and play an informative role. The operational semantics formalizes the intuition behind the game tree in Fig. 1 and relates to standard notions in game theory and descriptive set theory. We prove in Theorem 16 below that the operational game semantics is equivalent to the modal semantics from Section 2.2. The (denotational) modal semantics is much simpler but the operational semantics makes winning strategies explicit. As the set of actions A for a hybrid game, we choose:

$$\{\mathfrak{l}, \mathfrak{r}, \mathfrak{s}, \mathfrak{g}, \mathfrak{d}\} \cup \{(x := \theta) : x \text{ variable}, \ \theta \text{ term}\}$$
$$\cup \{(x' = \theta \,\&\, H @ r) : x \text{ variable}, \ \theta \text{ term}, H \text{ formula}, r \in \mathbb{R}_{\geq 0}\}$$
$$\cup \{?\phi : \phi \text{ formula}\}$$

For game $\alpha \cup \beta$, action $\mathfrak{l}$ decides to descend left into α, $\mathfrak{r}$ is the action of descending right into β. In game α^*, action $\mathfrak{s}$ decides to stop repeating, action $\mathfrak{g}$ decides to go back and repeat. Action $\mathfrak{d}$ starts and ends a dual game for α^d. The other actions represent the actions for atomic games: assignment actions, continuous evolution actions (in which time r is the critical decision), and test actions.

We use standard notions from descriptive set theory. The set of finite sequences of actions is denoted by $A^{(\mathbb{N})}$, the set of countably infinite sequences by $A^{\mathbb{N}}$. The empty sequence of actions is $()$. The concatenation, $s\hat{\ }t$, of sequences $s, t \in A^{(\mathbb{N})}$ is defined as $(s_1, \ldots, s_n, t_1, \ldots, t_m)$ if $s = (s_1, \ldots, s_n)$ and $t = (t_1, \ldots, t_m)$. For an $a \in A$, we write $a\hat{\ }t$ for $(a)\hat{\ }t$ and write $t\hat{\ }a$ for $t\hat{\ }(a)$. For a set $S \subseteq A^{(\mathbb{N})}$, we write $S\hat{\ }t$ for $\{s\hat{\ }t : s \in S\}$ and $t\hat{\ }S$ for $\{t\hat{\ }s : s \in S\}$. The state $\lfloor t \rfloor_s$ reached by *playing* a sequence of actions $t \in A^{(\mathbb{N})}$ from a state s in interpretation I is inductively defined by applying the actions sequentially, i.e. as follows:

1. $\lfloor x := \theta \rfloor_s = s_x^{[\![\theta]\!]_s}$
2. $\lfloor x' = \theta \,\&\, H @ r \rfloor_s = \varphi(r)$ for the unique $\varphi : [0, r] \to \mathcal{S}$ differentiable, $\varphi(0) = s$, $\frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ and $\varphi(\zeta) \in [\![H]\!]^I$ for all $\zeta \leq r$. Note that $\lfloor x' = \theta \,\&\, H @ r \rfloor_s$ is not defined if no such φ exists.
3. $\lfloor ?\phi \rfloor_s = \begin{cases} s & \text{if } s \in [\![\phi]\!]^I \\ \text{not defined} & \text{otherwise} \end{cases}$
4. $\lfloor \mathfrak{l} \rfloor_s = \lfloor \mathfrak{r} \rfloor_s = \lfloor \mathfrak{s} \rfloor_s = \lfloor \mathfrak{g} \rfloor_s = \lfloor \mathfrak{d} \rfloor_s = \lfloor () \rfloor_s = s$
5. $\lfloor a\hat{\ }t \rfloor_s = \lfloor t \rfloor_{(\lfloor a \rfloor_s)}$ for $a \in A$ and $t \in A^{(\mathbb{N})}$

A *tree* is a set $T \subseteq A^{(\mathbb{N})}$ that is closed under prefixes, that is, whenever $t \in T$ and s is a prefix of t (i.e. $t = s\hat{\ }r$ for some $r \in A^{(\mathbb{N})}$), then $s \in T$. A node $t \in T$ is a successor of node $s \in T$ iff $t = s\hat{\ }a$ for some $a \in A$. By $\mathrm{leaf}(T)$ we denote the set of all leaves of T, i.e. nodes $t \in T$ that have no successor in T.

Definition 5 (Operational game semantics)**.** The *operational game semantics* of hybrid game α is, for each state s of each interpretation I, a tree $\mathrm{g}(\alpha)(s) \subseteq A^{(\mathbb{N})}$ defined as follows (see Fig. 12 for a schematic illustration):

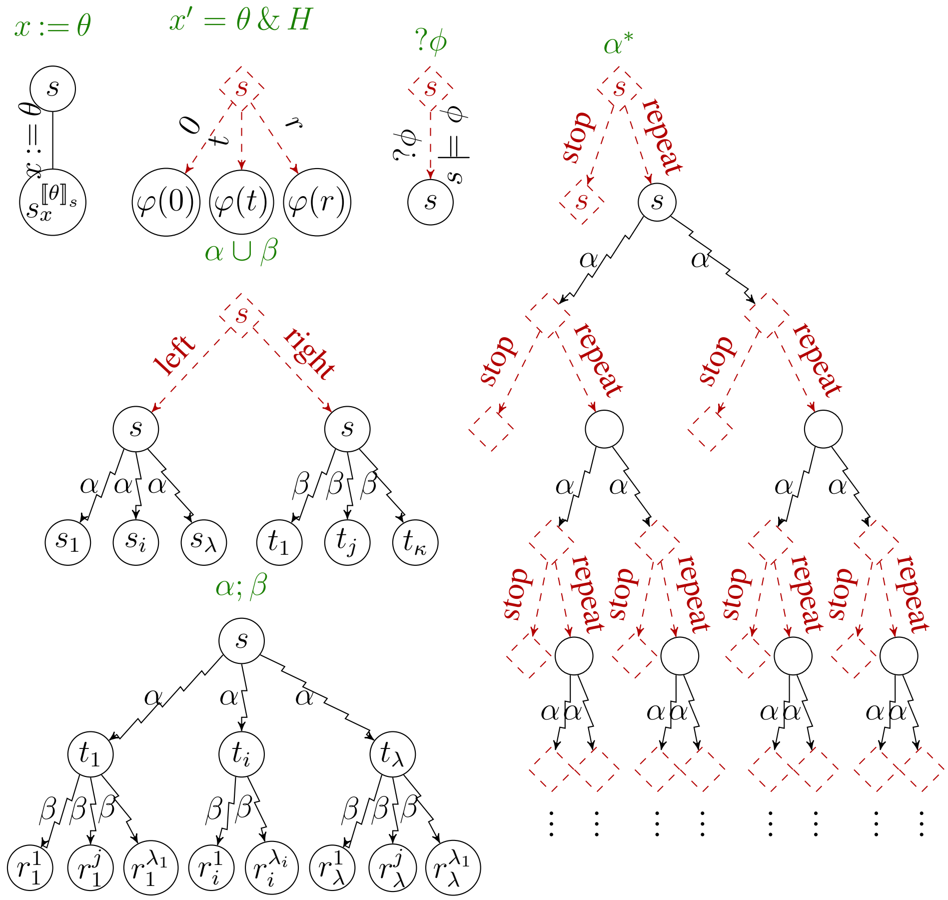

Figure 12: Operational game semantics for hybrid games of dG$\mathcal{L}$

1. $\mathsf{g}(x := \theta)(s) = \{(x := \theta)\}$

2. $\mathsf{g}(x' = \theta \,\&\, H)(s) = \{(x' = \theta \,\&\, H@r) : r \in \mathbb{R}, r \geq 0, \varphi(0) = s$ for some (differentiable) $\varphi : [0, r] \to \mathcal{S}$ such that $\frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ and $\varphi(\zeta) \in [\![H]\!]^I$ for all $\zeta \leq r\}$

3. $\mathsf{g}(?\phi)(s) = \{(?\phi)\}$

4. $\mathsf{g}(\alpha \cup \beta)(s) = \mathfrak{l}\hat{\ }\mathsf{g}(\alpha)(s) \cup \mathfrak{r}\hat{\ }\mathsf{g}(\beta)(s)$

5. $\mathsf{g}(\alpha; \beta)(s) = \mathsf{g}(\alpha)(s) \cup \bigcup_{t \in \mathrm{leaf}(\mathsf{g}(\alpha)(s))} \mathsf{g}(\beta)(\lfloor t \rfloor_s)$

6. $\mathsf{g}(\alpha^*)(s) = \bigcup_{n \in \mathbb{N}} f^n(\{(\mathfrak{s}), (\mathfrak{g})\})$
   where f^n is the n-fold composition of the function
   $f(Z) \stackrel{\text{def}}{=} Z \cup \bigcup_{t\hat{\ }\mathfrak{g} \in \mathrm{leaf}(Z)} t\hat{\ }\mathfrak{g}\hat{\ }\mathsf{g}(\alpha)(\lfloor t\hat{\ }\mathfrak{g} \rfloor_s)\hat{\ }\{(\mathfrak{s}), (\mathfrak{g})\}$

7. $\mathsf{g}(\alpha^d)(s) = \mathfrak{d}\hat{\ }\mathsf{g}(\alpha)(s)\hat{\ }\mathfrak{d}$

In the definition of $\mathsf{g}(\alpha)(s)$, note that we implicitly close under prefixes as necessary for readability reasons. For example, we write $\mathsf{g}(\alpha^d)(s) = \mathfrak{d}\hat{\ }\mathsf{g}(\alpha)(s)\hat{\ }\mathfrak{d}$ to mean $\mathsf{g}(\alpha^d)(s) = \{(), (\mathfrak{d})\} \cup \mathfrak{d}\hat{\ }\mathsf{g}(\alpha)(s) \cup \mathfrak{d}\hat{\ }\mathsf{g}(\alpha)(s)\hat{\ }\mathfrak{d}$.

Angel gets to choose which action to take at node $t \in \mathsf{g}(\alpha)(s)$ if t has an even number of occurrences of $\mathfrak{d}$, otherwise Demon gets to choose. In the former case we say *Angel acts at* t, in the latter *Demon acts at* t. Thus, at every t, exactly one of the players acts at t. If the player who acts at t is deadlocked, then that player loses immediately. A player who acts at $t \in \mathsf{g}(\alpha)(s)$ is *deadlocked* at t if $t \notin \mathrm{leaf}(\mathsf{g}(\alpha)(s))$ and no successor s is enabled, i.e. $\lfloor s \rfloor_s$ is not defined. This can happen if the last action in s has a condition that is not satisfied like $?x \geq 0$ or $x' = \theta \,\&\, x \geq 0$ at a state where $x < 0$. Note that the player who acts at $t \in \mathsf{g}(\alpha^*)(s)$ cannot choose $\mathfrak{g}$ infinitely often for that loop.

A *strategy for Angel* from initial state s is a nonempty subtree $\sigma \subseteq \mathsf{g}(\alpha)(s)$ such that

1. for all $t \in \sigma$ at which Demon acts, $t\hat{\ }a \in \sigma$ for all $a \in A$ such that $t\hat{\ }a \in \mathsf{g}(\alpha)(s)$.

2. for all $t \in \sigma$ at which Angel acts, if $t \notin \mathrm{leaf}(\mathsf{g}(\alpha)(s))$, then there is a unique $a \in A$ with $t\hat{\ }a \in \sigma$.

Strategies for Demon are defined accordingly, with "Angel" and "Demon" swapped. The action sequence $\sigma \oplus \tau$ played from state s in interpretation I when Angel plays strategy σ and Demon plays strategy τ from s is defined as the sequence $(a_1, \ldots, a_n) \in A^{(\mathbb{N})}$ of maximal length such that

$$a_{n+1} := \begin{cases} a & \text{if Angel acts at } (a_1, \ldots, a_n) \\ & \text{and } (a_1, \ldots, a_n)\hat{\ }a \in \sigma \\ a & \text{if Demon acts at } (a_1, \ldots, a_n) \\ & \text{and } (a_1, \ldots, a_n)\hat{\ }a \in \tau \\ \text{not defined} & \text{otherwise} \end{cases}$$

By definition of a strategy for Angel/Demon, the a is unique. A *winning strategy for Angel* for winning condition $X \subseteq \mathcal{S}$ from state s in interpretation I is a strategy $\sigma \subseteq \mathsf{g}(\alpha)(s)$ for Angel from s such that, for all strategies $\tau \subseteq \mathsf{g}(\alpha)(s)$ for Demon from s: Demon deadlocks or $\lfloor \sigma \oplus \tau \rfloor_s \in X$. A *winning strategy for Demon* for (Demon's) winning condition $X \subseteq \mathcal{S}$ from state s in interpretation I is a strategy $\tau \subseteq \mathsf{g}(\alpha)(s)$ for Demon from s such that, for all strategies $\sigma \subseteq \mathsf{g}(\alpha)(s)$ for Angel from s: Angel deadlocks or $\lfloor \sigma \oplus \tau \rfloor_s \in X$.

We show that the denotational modal semantics from Section 2.2 is equivalent to the operational semantics:

Theorem 16 (Equivalent semantics)**.** *The modal semantics of* dG$\mathcal{L}$ *is equivalent to the game tree operational semantics of* dG$\mathcal{L}$*, i.e. for each hybrid game α, each initial state s in each interpretation I, and each winning condition $X \subseteq \mathcal{S}$:*

$$s \in \varsigma_\alpha(X) \iff \textit{there is a winning strategy } \sigma \subseteq \mathsf{g}(\alpha)(s) \textit{ for Angel to achieve } X \textit{ from } s$$

$$s \in \delta_\alpha(X^\complement) \iff \textit{there is a winning strategy } \tau \subseteq \mathsf{g}(\alpha)(s) \textit{ for Demon to achieve } X^\complement \textit{ from } s$$

Proof. We proceed by simultaneous induction on the structure of α and prove equivalence. As part of the equivalence proof, we construct a winning strategy σ achieving X using that $s \in \varsigma_\alpha(X)$. The simultaneous induction steps for $\delta_\alpha(X^\complement)$ are simple dualities, except for the case of α^*. It is easy to see that Angel and Demon cannot both have a winning strategy from the same state s for complementary winning conditions X and $X^\complement$ in the same game $\mathsf{g}(\alpha)(s)$. By Theorem 2, we further know $\delta_\alpha(X^\complement) = \varsigma_\alpha(X)^\complement$.

1. $s \in \varsigma_{x:=\theta}(X) \iff s_x^{[\![\theta]\!]_s} \in X \iff \lfloor \sigma \oplus \tau \rfloor_s = \lfloor x := \theta \rfloor_s = s_x^{[\![\theta]\!]_s} \in X$, using $\sigma \stackrel{\text{def}}{=} \{(x := \theta)\} = \mathsf{g}(x := \theta)(s)$. The converse direction follows, because the strategy σ follows the only permitted strategy.

2. $s \in \varsigma_{x'=\theta \,\&\, H}(X) \iff s = \varphi(0), \varphi(r) \in X$ for some $r \in \mathbb{R}$ and some (differentiable) $\varphi : [0, r] \to \mathcal{S}$ such that $\frac{\mathsf{d}\,\varphi(t)(x)}{\mathsf{d}t}(\zeta) = [\![\theta]\!]_{\varphi(\zeta)}$ and $\varphi(\zeta) \in [\![H]\!]^I$ for all $\zeta \leq r \iff \lfloor \sigma \oplus \tau \rfloor_s = \lfloor x' = \theta \,\&\, H@r \rfloor_s = \varphi(r) \in X$, using $\sigma \stackrel{\text{def}}{=} \{(x' = \theta \,\&\, H@r)\} \subseteq \mathsf{g}(x' = \theta \,\&\, H)(s)$. The converse direction follows, because this σ has the only permitted form for a strategy where different values of r that lead to X are equivalently useful.

3. $s \in \varsigma_{?\phi}(X) = [\![\phi]\!]^I \cap X \iff \lfloor \sigma \oplus \tau \rfloor_s = \lfloor ?\phi \rfloor_s = s \in X$, with $s \in [\![\phi]\!]^I$ using $\sigma \stackrel{\text{def}}{=} \{(?\phi)\} = \mathsf{g}(?\phi)(s)$. The converse direction uses that this σ is the only permitted strategy and it deadlocks exactly if $s \notin [\![\phi]\!]^I$.

4. $s \in \varsigma_{\alpha\cup\beta}(X) = \varsigma_\alpha(X) \cup \varsigma_\beta(X) \iff s \in \varsigma_\alpha(X)$ or $s \in \varsigma_\beta(X)$. By induction hypothesis, this is equivalent to: there is a winning strategy $\sigma_\alpha \subseteq \mathsf{g}(\alpha)(s)$ for Angel for X from s or there is a winning strategy $\sigma_\beta \subseteq \mathsf{g}(\beta)(s)$ for Angel for X from s. This is equivalent to $\sigma \subseteq \mathsf{g}(\alpha \cup \beta)(s)$ being a winning strategy for Angel for X from s, using either $\sigma \stackrel{\text{def}}{=} \mathfrak{l}\hat{}\sigma_\alpha$ or $\sigma \stackrel{\text{def}}{=} \mathfrak{r}\hat{}\sigma_\beta$.

5. $s \in \varsigma_{\alpha;\beta}(X) = \varsigma_\alpha(\varsigma_\beta(X))$. By induction hypothesis, this is equivalent to the existence of a strategy $\sigma_\alpha \subseteq \mathsf{g}(\alpha)(s)$ for Angel such that for all strategies $\tau \subseteq \mathsf{g}(\alpha)(s)$ for Demon: $\lfloor \sigma_\alpha \oplus \tau \rfloor_s \in \varsigma_\beta(X)$. By induction hypothesis, $\lfloor \sigma_\alpha \oplus \tau \rfloor_s \in \varsigma_\beta(X)$ is equivalent to the existence of a winning strategy σ_τ for Angel (which depends on the state $\lfloor \sigma_\alpha \oplus \tau \rfloor_s$ that the previous α game led to) with winning condition X from $\lfloor \sigma_\alpha \oplus \tau \rfloor_s$. This is equivalent to $\sigma \subseteq \mathsf{g}(\alpha;\beta)(s)$ being a winning strategy for Angel for X from s, using

$$\sigma \stackrel{\text{def}}{=} \sigma_\alpha \cup \bigcup (\sigma_\alpha \oplus \tau)\hat{\ }\sigma_\tau \tag{8}$$

The union is over all leaves $\sigma_\alpha \oplus \tau \in \mathrm{leaf}(\mathsf{g}(\alpha)(s))$ for which the game is not won by a player yet. Note that σ is a winning strategy for X, because, for all plays for which the game is decided during α, the strategy σ_α already wins the game. For the others, σ_τ wins the game from the respective state $\lfloor \sigma_\alpha \oplus \tau \rfloor_s$ that was reached by the actions $\sigma_\alpha \oplus \tau$ according to the strategy τ that Demon was observed (when α terminates) to have played during α. The converse direction uses that strategies do not depend on moves that have not been played yet and that any strategy can be factorized by prefixes of what has actually been played to be coerced into the form (8).

6. We prove both inclusions of the case α^* separately. If W denotes the set of states from which Angel has a winning strategy in $\mathsf{g}(\alpha^*)(s)$ to achieve X, then we need to show that $\varsigma_{\alpha^*}(X) = W$. For $\varsigma_{\alpha^*}(X) \subseteq W$, it is enough to show that W is a pre-fixpoint, i.e. $X \cup \varsigma_\alpha(W) \subseteq W$, because $\varsigma_{\alpha^*}(X)$ is the least (pre-)fixpoint. Consider any $s \in X \cup \varsigma_\alpha(W) \subseteq W$. If $s \in X$ then $s \in W$ with the winning strategy $\sigma \stackrel{\text{def}}{=} \{(\mathfrak{s})\}$ for Angel to achieve X in α^* from s. Otherwise, $s \in \varsigma_\alpha(W) \subseteq W$ implies, by induction hypothesis, that there is a winning strategy $\sigma_\alpha \subseteq \mathsf{g}(\alpha)(s)$ for Angel in α to achieve W from s. By definition of W, Angel has a winning strategy in $\mathsf{g}(\alpha^*)(s)$ to achieve X from all states reached after playing α from s according to σ_α, i.e. $\lfloor \sigma_\alpha \oplus \tau \rfloor_s \in W$ for all strategies τ of Demon. Thus, by composing σ_α with the respective (state-dependent) winning strategies σ_τ for all possible resulting states (which are all in W) corresponding to the respective possible strategies τ that Demon could play during the first α, we obtain a winning strategy of the form

$$\sigma \stackrel{\text{def}}{=} \mathfrak{g}\hat{\ }\sigma_\alpha \cup \bigcup \mathfrak{g}\hat{\ }(\sigma_\alpha \oplus \tau)\hat{\ }\sigma_\tau$$

for Angel to achieve X in α^* from s, where the union is over all leaves $\sigma_\alpha \oplus \tau \in \mathrm{leaf}(\mathsf{g}(\alpha)(s))$ in any strategy τ of Demon for which the game is not won by a player yet during the first α.

The converse inclusion $\varsigma_{\alpha^*}(X) \supseteq W$ is equivalent to $\varsigma_{\alpha^*}(X)^\complement \subseteq W^\complement$. For this, we recall that $\varsigma_{\alpha^*}(X)^\complement = \delta_{\alpha^*}(X^\complement) = \bigcup \{Z \subseteq \mathcal{S} : Z \subseteq X^\complement \cap \delta_\alpha(Z)\}$ by Theorem 2. Thus, since $\varsigma_{\alpha^*}(X)^\complement$ is a greatest (post-)fixpoint, it is enough to show $Z \subseteq W^\complement$ for all Z with $Z \subseteq X^\complement \cap \delta_\alpha(Z)$. Since, $Z \subseteq \delta_\alpha(Z)$, Demon has a winning strategy in α to achieve Z from all $s \in Z$, by induction hypothesis. By composing the respective winning strategies for Demon, we obtain a winning strategy τ for Demon to achieve Z in α^* for *any* number of repetitions that Angel chooses (recall that Angel cannot choose to repeat α^* infinitely often to win). Since $Z \subseteq X^\complement$, Angel cannot have a winning strategy to achieve X in α^* from any $s \in Z$ by Theorem 2. Thus, $Z \subseteq W^\complement$.

7. $s \in \varsigma_{\alpha^d}(X) = \varsigma_\alpha(X^\complement)^\complement$. $\iff$ $s \notin \varsigma_\alpha(X^\complement)$. By induction hypothesis, this is equivalent to: there is no winning strategy $\sigma \subseteq \mathtt{g}(\alpha)(s)$ for Angel winning $X^\complement$ in α from s. Since $\varsigma_{\alpha^d}(X) = \delta_\alpha(X)$ by Theorem 2, this is equivalent to: there is a winning strategy $\tau \subseteq \mathtt{g}(\alpha)(s)$ for Demon winning X in α from s. Since the nodes where Angel acts swap with the nodes where Demon acts when moving from α to α^d, this is equivalent to: there is a winning strategy $\sigma \subseteq \mathtt{g}(\alpha^d)(s)$ for Angel winning X in α^d from s using $\sigma \stackrel{\text{def}}{=} \mathfrak{d}\hat{}\tau\hat{}\mathfrak{d}$. The converse direction uses that all strategies permitted for α^d begin and end with $\mathfrak{d}$. □

D Alternative Semantics

To argue why the dG$\mathcal{L}$ semantics is both natural and general, we briefly discuss alternative choices for the semantics, focusing on the role of repetition in the context of hybrid games.

D.1 Advance Notice Semantics

One alternative semantics is the *advance notice semantics* for α^*, which requires the players to announce the number of times that game α will be repeated when the loop begins. The advance notice semantics defines $\varsigma_{\alpha^*}(X)$ as $\bigcup_{n\in\mathbb{N}} \varsigma_{\alpha^n}(X)$ where $\alpha^{n+1} \equiv \alpha^n; \alpha$ and $\alpha^0 \equiv ?true$ and defines $\delta_{\alpha^*}(X)$ as $\bigcap_{n\in\mathbb{N}} \delta_{\alpha^n}(X)$. When playing α^*, Angel, thus, announces to Demon how many repetitions n are going to be played when the game α^* begins and Demon announces how often to repeat $\alpha^\times$. This advance notice makes it easier for Demon to win loops α^* and easier for Angel to win loops $\alpha^\times$, because the opponent announces an important feature of their strategy immediately as opposed to revealing whether or not to repeat the game once more one iteration at a time as in Def. 4.

In hybrid systems, the advance notice semantics and the least fixpoint semantics are equivalent (Lemma 3), but the advance notice semantics and dG$\mathcal{L}$'s least fixpoint semantics are different for hybrid games. The following formula is valid in dG$\mathcal{L}$ (see Fig. 13), but would not be valid in the advance notice semantics:

$$x = 1 \land a = 1 \rightarrow \langle((x := a; a := 0) \cap x := 0)^*\rangle x \neq 1 \tag{9}$$

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

Conversely, the dual formula would be valid in the advance notice semantics but is not valid in dG$\mathcal{L}$:

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

The dG$\mathcal{L}$ semantics is more general, because advance notice games can be expressed easily in dG$\mathcal{L}$ by having the players choose a counter c before the loop that decreases to 0 during the repetition.

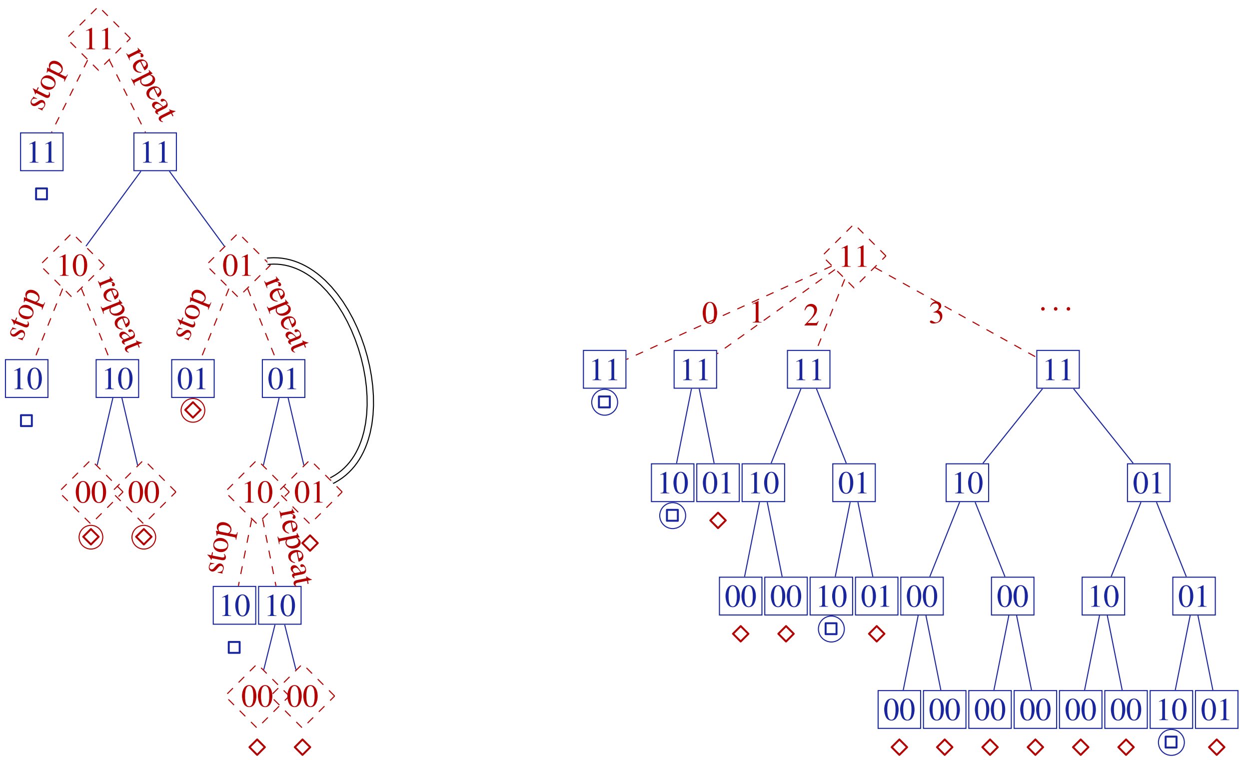

Figure 13: Game trees for $x = 1 \wedge a = 1 \rightarrow \langle \alpha^* \rangle x \neq 1$ with game $\alpha \equiv (x := a; a := 0) \cap x := 0$ (notation: x, a). **(top)** valid in d$\mathcal{GL}$ by strategy "repeat once and repeat once more if $x = 1$, then stop" **(bottom)** false in advance notice 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 advance notice semantics can be expressed in dGL, e.g., for (9) as

$$x = 1 \land a = 1 \to \langle c := 0; c := c + 1^*; \\ (((x := a; a := 0) \cap x := 0); c := c - 1)^*; ?c = 0\rangle x \neq 1$$

The dGL semantics cannot, however, be expressed conversely in an advance notice semantics, so the dGL semantics is strictly more general.

D.2 ω-Strategic Semantics

Another alternative choice for the semantics would have been to allow only arbitrary finite iterations of the strategy function for computing the winning region by using the *ω-strategic semantics*, which defines $\varsigma_{\alpha^*}(X)$ as $\varsigma_\alpha^\omega(X) = \bigcup_{n\in\mathbb{N}} \varsigma_\alpha^n(X)$ along with a corresponding definition for $\delta_{\alpha^*}(X)$. Like the dGL semantics, but quite unlike the advance notice semantics, the ω-strategic semantics does not require Angel to disclose how often she is going to repeat when playing α^*. Similarly, Demon does not have to announce how often to repeat when playing $\alpha^\times$. Nevertheless, the semantics are different. The ω-strategic semantics would make the following valid dGL formula invalid:

$$\langle (x := 1; x' = 1^d \cup x := x - 1)^* \rangle\, (0 \leq x < 1) \tag{10}$$

By a simple variation of the argument in the proof of Theorem 6, $\varsigma_\alpha^\omega([0,1)) = [0,\infty)$, because $\varsigma_\alpha^n([0,1)) = [0,n)$ for all $n \in \mathbb{N}$. Yet, this ω-level of iteration of the strategy function for winning regions misses out on the perfectly reasonable winning strategy "first choose $x := 1; x' = 1^d$ and then always choose $x := x - 1$ until stopping at $0 \leq x < 1$". The existence of this winning strategy is only found at the level $\varsigma_\alpha^{\omega+1}([0,1)) = \varsigma_\alpha([0,\infty)) = \mathbb{R}$. Even though any particular use of the winning strategy in any game play uses only some finite number of repetitions of the loop, the argument why it will always work requires $> \omega$ many iterations of $\varsigma_\alpha(\cdot)$, because Demon can change x to an arbitrarily big value, so that ω many iterations of $\varsigma_\alpha(\cdot)$ are needed to conclude that Angel has a winning strategy for any positive value of x. There is no upper bound $< \omega$ on the number of iterations it takes Angel to win. But it does converge after $\omega + 1$ iterations. According to Theorem 6, the same shortcomings of the ω-semantics apply at higher closure ordinals.

E Proof of Higher Closure Ordinals

Proof of Theorem 6. In this proof, we proceed in stages of increasing difficulty. We have already shown above that the closure ordinal is $\geq \omega \cdot 2$ in Appendix **??**. Now we prove the bounds $\geq \omega^2$ and finally $\geq \omega^\omega$. In order to see that the closure ordinal is at least ω^2 even for a single nesting layer of dual and loop, we follow a similar argument using more variables. Consider the family of formulas (for some $N \in \mathbb{N}$) of the form

$$\langle(\underbrace{x_N := x_N - 1; x'_{N-1} = 1^d \cup \ldots \cup x_2 := x_2 - 1; x'_1 = 1^d \cup x_1 := x_1 - 1}_{\alpha})^*\rangle \bigwedge_{i=1}^{N} x_i < 0$$

We show that the winning regions for this dG$\mathcal{L}$ formula stabilize after $\omega \cdot N$ iterations, because ω many iterations are necessary to show that *any* x_1 can be reduced to $(-\infty, 0)$ by choosing the last action sufficiently often, whereas another ω many iterations are needed to show that x_2 can then be reduced to $(-\infty, 0)$ by choosing the second-to-last action sufficiently often, increasing x_1 arbitrarily under Demon's control, which can still be won because this adversarial increase in x_1 can be compensated for by the first part of the winning strategy. We use the vector space of variables $(x_N, \ldots, x_1)$ in that order. It is easy to see that $\varsigma_\alpha^\omega((-\infty,0)^N) = \bigcup_{n\in\mathbb{N}} \varsigma_\alpha^n((-\infty,0)^N) = (-\infty,0)^{N-1} \times \mathbb{R}$, because $\varsigma_\alpha^{n+1}((-\infty,0)) = (-\infty,0)^{N-1} \times (-\infty,n)$ holds for all $n \in \mathbb{N}, n$ by a simple inductive argument:

$$\begin{aligned}
\varsigma_\alpha^1((-\infty,0)^N) &= (-\infty,0)^N \\
\varsigma_\alpha^{n+1}((-\infty,0)^N) &= (-\infty,0)^N \cup \varsigma_\alpha(\varsigma_\alpha^n((-\infty,0)^N)) = (-\infty,0)^N \cup \varsigma_\alpha((-\infty,0)^{N-1} \times (-\infty,n-1)) \\
&= (-\infty,0)^{N-1} \times (-\infty,n)
\end{aligned}$$

Inductively, $\varsigma_\alpha^{\omega\cdot(k+1)}((-\infty,0)^N) = \bigcup_{n\in\mathbb{N}} \varsigma_\alpha^{\omega\cdot k+n}((-\infty,0)^N) = (-\infty,0)^{N-k-1} \times \mathbb{R}^{k+1}$, because $\varsigma_\alpha^{\omega\cdot k+n+1}((-\infty,0)) = (-\infty,0)^{N-k-1} \times (-\infty,n) \times \mathbb{R}^k$ holds for all $n \in \mathbb{N}$ by a simple inductive argument:

$$\begin{aligned}
\varsigma_\alpha^{\omega\cdot k+n+1}((-\infty,0)^N) &= (-\infty,0)^N \cup \varsigma_\alpha(\varsigma_\alpha^{\omega\cdot k+n}((-\infty,0)^N)) = (-\infty,0)^N \cup \varsigma_\alpha((-\infty,0)^{N-k-1} \times (-\infty,n-1) \times \mathbb{R}^k) \\
&= (-\infty,0)^{N-k-1} \times (-\infty,n) \times \mathbb{R}^k
\end{aligned}$$

Consequently, $\varsigma_{\alpha^*}((-\infty,0)^N) = \varsigma_\alpha^{\omega\cdot N}((-\infty,0)^N) \neq \varsigma_\alpha^{\omega\cdot(N-1)+n}((-\infty,0)^N)$, which makes $\omega \cdot N$ the closure ordinal for α. Since we can consider hybrid games α of the above form with arbitrarily big $N \in \mathbb{N}$, the common closure ordinal has to be $\geq \omega \cdot N$ for all $N \in \mathbb{N}$, i.e. it has to be $\geq \omega^2$.

In order to see that the closure ordinal is at least ω^ω, we follow an argument expanding on the previous case. Consider the family of formulas (for some $N \in \mathbb{N}$) of the form

$$\langle(\underbrace{?x_{N-1} < 0; x'_{N-1} = 1^d; x_N := x_N - 1 \cup \ldots \cup ?x_1 < 0; x'_1 = 1^d; x_2 := x_2 - 1 \cup x_1 := x_1 - 1}_{\alpha})^*\rangle \bigwedge_{i=1}^{N} x_i < 0$$

We prove that the winning regions for this "clockwork ω" formula stabilize after ω^N iterations, ω many iterations are necessary to show that *any* x_1 can be reduced to $(-\infty, 0)$ by choosing the last action sufficiently often, whereas another ω many iterations are needed to show that x_2 can then be reduced to $(-\infty, 0)$ by choosing the second-to-last action sufficiently often in case x_1 has already been reduced to $(-\infty, 0)$. Every time the second-to-last action is chosen, however, Demon increases x_1 arbitrarily, which again takes ω many steps of the last action to understand how x_1 can again be reduced to $(-\infty, 0)$ before the second-to-last action can be chosen again to decrease x_2 further. This phenomenon that ω many actions on x_{i-1} are needed before x_i can be decreased by 1 holds for all i recursively. Note that in any particular game play, Demon can only increase x_i by some finite amount. But Angel does not have a finite bound on that increment, so she will first have to convince herself that she has a winning strategy that could tolerate any change in x_i, which takes ω many iterations of the previous argument.

We use the vector space of variables $(x_N, \dots, x_1)$ in that order. For $b_N, \dots, b_1 \in \mathbb{N} \cup \{\infty\}$, we use the short hand notation

$$b_N \dots b_2 b_1 \overset{\text{def}}{=} (-\infty, b_N) \times \cdots \times (-\infty, b_2) \times (-\infty, b_1)$$

and also write b_i^n for $(-\infty, b_i)^n$ in that context. Let $\vec{b} = (b_N, \dots, b_1)$. We prove that $\forall n \in \mathbb{N}\ \forall j \in \mathbb{N}, j > 0$

$$\begin{aligned}
\varsigma_\alpha^{\omega^j(n+1)}(b_N \dots b_j \dots b_1) &= b_N \dots (b_{j+1} + n)\infty^j && \text{if ① } b_N, \dots, b_j < \infty, j > 0\\
\varsigma_\alpha^{\omega^j(n+1)}(b_N \dots b_{j+1}\infty^j) &= b_N \dots (b_{j+1} + n + 1)\infty^j && \text{if ② } b_N, \dots, b_{j+1} < \infty, b_j = \infty = \dots b_1\\
\varsigma_\alpha^{\omega^j(n+1)}(b_N \dots b_{k+1}\infty^{k-j}\infty^j) &= b_N \dots (b_{k+1} + 1)1^{k-j-1}(n+1)\infty^j \cup \vec{b} && \text{if ③ } b_N, \dots, b_{k+1} < \infty, b_k = \infty, k > j
\end{aligned}$$

by induction on the lexicographical order of j and n. Note that, in the case ③, there are some subordinate cases which we do not need to track in our analysis, because they are strategic dead ends. IH is short for induction hypothesis.

The base case $j = 0, n = 0$ is vacuous for ① and can be checked easily for ②.

$$\begin{aligned}
\varsigma_\alpha^{\omega^0 1}(b_N \dots b_1\infty^0) &= \varsigma_\alpha^1(b_N \dots b_1) = b_N \dots (b_1 + 1) = b_N \dots (b_1 + 1)\infty^0\\
\varsigma_\alpha^{\omega^0(n+1)}(b_N \dots b_1\infty^0) &= \vec{b} \cup \varsigma_\alpha(\varsigma_\alpha^n(b_N \dots b_1)) = \vec{b} \cup \varsigma_\alpha(b_N \dots (b_1 + n)) = b_N \dots (b_1 + n + 1)
\end{aligned}$$

For ③, the case $j = 0$ holds only after an extra offset k, however:

$$\begin{aligned}
\varsigma_\alpha^1(b_N \dots b_{k+1}\infty^k) &= \vec{b} \cup b_N \dots (b_{k+1} + 1)0\infty^{k-1}\\
\varsigma_\alpha^{n+1}(b_N \dots b_{k+1}\infty^k) &= \varsigma_\alpha^n(b_N \dots b_{k+1}\infty^k) \cup b_N \dots (b_{k+1} + 1)1^n 0\infty^{k-n-1} \quad \text{for } n < k\\
\varsigma_\alpha^{k+n+1}(b_N \dots b_{k+1}\infty^k) &= \varsigma_\alpha^{k+n}(b_N \dots b_{k+1}\infty^k) \cup b_N \dots (b_{k+1} + 1)1^{k-1}(n+1)
\end{aligned}$$

So instead, we prove base case $j = 1, n = 0$, because the finite extra offset k has been overcome at ω:

$$\begin{aligned}
\varsigma_\alpha^{\omega^1 1}(b_N \dots b_1) &= \bigcup_{n\in\mathbb{N}} \varsigma_\alpha^{\omega^0(n+1)}(b_N \dots b_1\infty^0) = \bigcup_{n\in\mathbb{N}} b_N \dots (b_1 + n + 1) = b_N \dots b_2\infty && \text{if ①}\\
\varsigma_\alpha^{\omega^1 1}(b_N \dots b_2\infty) &= \bigcup_{n\in\mathbb{N}} \varsigma_\alpha^{\omega^0(n+1)}(b_N \dots b_2\infty^1) = b_N \dots (b_2 + 1)\infty && \text{if ②}\\
\varsigma_\alpha^{\omega^1 1}(b_N \dots b_{k+1}\infty^k) &= \bigcup_{n\in\mathbb{N}} \varsigma_\alpha^{\omega^0(n+1)}(b_N \dots b_{k+1}\infty^k) = \bigcup_{n\in\mathbb{N}} b_N \dots (b_{k+1} + 1)1^{k-1}(n+1) \cup \vec{b}\\
&= b_N \dots (b_{k+1} + 1)1^{k-1}\infty \cup \vec{b} && \text{if ③}
\end{aligned}$$

In case ③, there are some subordinate cases $\cup\vec{b}$ coming from mixed occurrences $b_N \dots (b_{k+1} + 1)^i 0 \infty^{k-i-1}$, but we do not need to track them in our analysis, because they are strategic dead ends. By construction of α, no counter can be changed without resetting all smaller variables to 0 first as indicated.

$j \curvearrowright j+1, n=0$: For the step from j to $j+1$ we prove the case $n=0$ as follows.

$$\varsigma_\alpha^{\omega^{j+1}\cdot(0+1)}(b_N \dots b_j \dots b_1) = \varsigma_\alpha^{\omega^j\cdot\omega}(b_N \dots b_j \dots b_1) = \bigcup_{n\in\mathbb{N}} \varsigma_\alpha^{\omega^j\cdot(n+1)}(b_N \dots b_j \dots b_1)$$

$$\overset{IH}{=} \begin{cases} \bigcup_{n\in\mathbb{N}} b_N \dots (b_{j+1}+n)\infty^j & \text{if ①} \\ \bigcup_{n\in\mathbb{N}} b_N \dots (b_{j+1}+n+1)\infty^j & \text{if ②} \\ \bigcup_{n\in\mathbb{N}} b_N \dots (b_{k+1}+1)1^{k-j-1}(n+1)\infty^j \cup \vec{b} & \text{if ③} \end{cases}$$

$$\overset{IH}{=} \begin{cases} b_N \dots b_{j+2}\infty^{j+1} & \text{if } b_N,\dots,b_j < \infty \\ b_N \dots b_{j+2}\infty^{j+1} & \text{if } b_N,\dots,b_{j+1} < \infty \\ b_N \dots (b_{j+2}+1)\infty^{j+1} & \text{if } b_N,\dots,b_{j+2} < \infty, b_{j+1} = \infty, k = j+1 \\ b_N \dots (b_{k+1}+1)1^{k-j-2}1\infty^{j+1} \cup \vec{b} & \text{if } b_N,\dots,b_{k+1} < \infty, b_k = \infty, k > j+1 \end{cases}$$

$n \curvearrowright n+1$: Within any level j, we prove the step from n to $n+1$ as follows. If $n=0$, then $\varsigma_\alpha^{\omega^j(n+1)}(b_N \dots b_j \dots b_1) = \varsigma_\alpha^{\omega^j}(b_N \dots b_j \dots b_1)$ already has the property by induction hypothesis. Otherwise $n > 0$, which allows us to conclude:

$$\varsigma_\alpha^{\omega^j(n+1)}(b_N \dots b_j \dots b_1) = \varsigma_\alpha^{\omega^j n + \omega^j}(b_N \dots b_j \dots b_1) \overset{Lemma\,5}{=} \varsigma_\alpha^{\omega^j}(\varsigma_\alpha^{\omega^j n}(b_N \dots b_j \dots b_1))$$

$$\overset{IH}{=} \begin{cases} \varsigma_\alpha^{\omega^j}(b_N \dots (b_{j+1}+n-1)\infty^j) & \text{if ①} \\ \varsigma_\alpha^{\omega^j}(b_N \dots (b_{j+1}+n)\infty^j) & \text{if ②} \\ \varsigma_\alpha^{\omega^j}(b_N \dots (b_{k+1}+1)1^{k-j-1}n\infty^j \cup \vec{b}) & \text{if ③} \end{cases}$$

$$\overset{IH}{=} \begin{cases} b_N \dots (b_j+n)\infty^j & \text{if ①} \\ b_N \dots (b_j+n+1)\infty^j & \text{if ②} \\ b_N \dots (b_{k+1}+1)1^{k-j-1}(n+1)\infty^j \cup \vec{b} & \text{if ③} \end{cases}$$

Consequently, $\varsigma_{\alpha^*}((-\infty,0)^N) = \varsigma_\alpha^{\omega^N}((-\infty,0)^N) = \mathbb{R}^N \neq \varsigma_\alpha^{\omega^{N-1}\cdot n}((-\infty,0)^N)$ for all $n \in \mathbb{N}$, which makes ω^N the closure ordinal for α. Since we can consider hybrid games α of the above form with arbitrarily big $N \in \mathbb{N}$, the common closure ordinal has to be $\geq \omega^N$ for all $N \in \mathbb{N}$, i.e. it has to be $\geq \omega^\omega$. □