

Differential Game Logic

ANDRÉ PLATZER, Carnegie Mellon University

Differential game logic (dGL) is a logic for specifying and verifying properties of *hybrid games*, i.e. games that combine discrete, continuous, and adversarial dynamics. Unlike hybrid systems, hybrid games allow choices in the system dynamics to be resolved adversarially by different players with different objectives. The logic dGL can be used to study the existence of winning strategies for such hybrid games, i.e. ways of resolving the player's choices in *some* way so that he wins by achieving his objective for *all* choices of the opponent. Hybrid games are determined, i.e. from each state, one player has a winning strategy, yet computing their winning regions may take transfinitely many steps. The logic dGL, nevertheless, has a sound and complete axiomatization relative to any expressive logic. Separating axioms are identified that distinguish hybrid games from hybrid systems. Finally, dGL is proved to be strictly more expressive than the corresponding logic of hybrid systems by characterizing the expressiveness of both.

CCS Concepts: •**Theory of computation** → **Timed and hybrid models; Proof theory; Modal and temporal logics; Programming logic; Solution concepts in game theory;**

Additional Key Words and Phrases: Game logic, hybrid games, axiomatization, expressiveness

ACM Reference Format:

André Platzer. 2015. Differential game logic. *ACM Trans. Comput. Logic* 17, 1, Article 1 (November 2015), 51 pages.

DOI: <http://dx.doi.org/10.1145/2817824>

1. INTRODUCTION

Hybrid systems [Nerode and Kohn 1992; Alur et al. 1995; Branicky et al. 1998; Davoren and Nerode 2000] are dynamical systems combining discrete dynamics and continuous dynamics. They are widely important, e.g., for modeling how computers control physical systems such as cars [Deshpande et al. 1996], aircraft [Umeno and Lynch 2007] and other cyber-physical systems. Hybrid systems combine difference equations (or discrete assignments) and differential equations with conditional switching, nondeterministic choices, and repetition [Platzer 2010b]. Hybrid systems are not semidecidable [Henzinger et al. 1995], but nevertheless studied by many successful verification approaches [Platzer 2012c; Doyen et al. 2016]. They have a complete axiomatization relative to differential equations in *differential dynamic logic* (dL) [Platzer 2008; Platzer 2012a], which extends Pratt's dynamic logic of conventional discrete programs [Pratt 1976] to hybrid systems by adding differential equations and a reachability relation semantics on the real Euclidean space.

Hybrid games [Nerode et al. 1996; Tomlin et al. 1998; Henzinger et al. 1999; Tomlin et al. 2000; Dharmatti and Ramaswamy 2006; Bouyer et al. 2010; Vladimerou et al. 2011] are games of two players on a hybrid system. Hybrid games add an adversarial dynamics to hybrid systems, i.e. an adversarial way of resolving the choices in the system dynamics. Both players can make their respective choices arbitrarily. They are not assumed to cooperate towards a common goal but may compete. The prototypical

This material is based upon work supported by the National Science Foundation under NSF CAREER Award CNS-1054246. A preliminary version has appeared as a report [Platzer 2012b; Platzer 2013].

Author's address: A. Platzer, Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2015 Copyright held by the owner/author(s). 1529-3785/2015/11-ART1 \$15.00

DOI: <http://dx.doi.org/10.1145/2817824>

example of a hybrid game is RoboCup, where two (teams of) robots move continuously on a soccer field subject to the discrete decisions of their respective control programs, and they resolve their choices adversarially in active competition for scoring goals. Worst-case verification of many other situations leads to hybrid games. Two robots may already end up in a hybrid game if they simply do not know each other's objectives, because worst-case analysis assumes they might interfere, which makes them compete accidentally rather than on purpose. The former situation is *true competition*, the latter *analytic competition*, because possible competition was assumed for the sake of a worst-case analysis. Aircraft separation provides further natural scenarios for both true [Isaacs 1967] and analytic [Tomlin et al. 1998] competition. Hybrid games are also fundamental for security questions about hybrid systems, which intrinsically involve adversarial situations with more than one player. Many different variations of hybrid games are interesting for applications [Tomlin et al. 1998; Tomlin et al. 2000; Henzinger et al. 1999; Prandini et al. 2001; Bouyer et al. 2010; Vladimerou et al. 2011; Quesel and Platzer 2012], including games between controller and plant for control synthesis, hybrid pursuit-evader games, or hybrid games for verification of robot controllers against an uncertain environment or an external disturbance.

This article does not focus on one such fixed pattern of game interaction in hybrid systems, but considers a more general framework for hybrid game interactions of two players. While the results of this article show promise in practice, the focus of this article is on fundamental logical considerations for hybrid games. It develops a compositional programming language for hybrid games and a logic for hybrid games along with its fundamental compositional proof principles. The article analyzes hybrid games and contrasts them with hybrid systems in terms of their analytic complexity, axiomatizations, and expressiveness.

Approach. This article studies a compositional model of hybrid games obtained as a programming language from a compositional model of hybrid systems [Platzer 2008] by simply adding the duality operator d for passing control between the players. The dual game α^d is the same as the hybrid game α with the roles of the players swapped, much like what happens when turning a chessboard around by 180° so that players black and white swap sides. Hybrid games without d are single-player, like hybrid systems are, because d is the only operator where control passes to the other player. Hybrid games with d give both players control over their respective choices (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. The fact that d is an operator on hybrid games makes them fully symmetric. That is, they allow arbitrary combinations of all operators at arbitrary nesting depths to define the game, not just a single fixed pattern like, e.g., the separation into a single loop of a continuous plant player and a discrete controller player that has been predominant in other approaches.

Hybrid games are game-theoretically reasonably tame sequential, non-cooperative, zero-sum, two-player games of perfect information with payoffs ± 1 .¹ What makes them challenging is that they are played on hybrid systems, which causes reachability computations and the canonical game solution technique of backwards induction for winning regions to take infinitely many iterations ($\geq \omega_1^{\text{CK}}$) to terminate.

One of the most fundamental questions about a hybrid game is whether the player of interest has a *winning strategy*, i.e. a way of resolving his choices that will lead to a state in which that player wins, no matter how the opponent player resolves his respec-

¹Draws, coalitions, rewards, and payoffs different from ± 1 etc. are expressible in the logic developed in this article, which gains simplicity and elegance by focusing on the most fundamental case of hybrid games.

tive choices.² If the player has such a winning strategy, he can achieve his objectives no matter what the opponent does, otherwise he needs his opponent to cooperate.

This article introduces a logic and proof calculus for hybrid games and thereby decouples the questions of truth (existence of winning strategies) and proof (winning strategy certificates) and proof search (automatic construction of winning strategies). It studies provability (existence of proofs) and the proof theory of hybrid games and identifies what the right proof rules for hybrid games are (soundness & completeness).

This article presents *differential game logic* (dGL) and its axiomatization for studying the existence of winning strategies for hybrid games. It generalizes hybrid systems to hybrid games by adding the duality operator ^d and a winning strategy semantics on the real Euclidean space. Hybrid games simultaneously generalize hybrid systems [Nerode and Kohn 1992; Alur et al. 1995] and discrete games [Parikh 1983; Parikh 1985; Pauly and Parikh 2003]. Similarly, dGL simultaneously generalizes logics of hybrid systems and logics of discrete games. The logic dGL generalizes differential dynamic logic (dL) [Platzer 2008; Platzer 2012a] from hybrid systems to hybrid games with their adversarial dynamics and, simultaneously, generalizes Parikh’s propositional game logic [Parikh 1983; Parikh 1985; Pauly and Parikh 2003] from games on finite-state discrete systems to games on hybrid systems with their differential equations, uncountable state spaces, uncountably many possible moves, and interacting discrete and continuous dynamics.

Contributions. Every particular play of a hybrid game has exactly one winner (Section 2). From each state exactly one player has a winning strategy no matter how the opponent reacts (determinacy, Section 3). The dGL proof calculus can be used to find out which of the two players it is that has a winning strategy from which state (Section 4). The logic dGL for hybrid games is proved to be fundamentally more expressive than the logic dL for hybrid systems by characterizing the expressiveness of both (Section 5).

The primary contributions of this article are as follows. The logic dGL identifies the logical essence of hybrid games and their game combinators.³ It identifies a simple, algebraic, compositional model of hybrid games as a programming language for hybrid games by adding the control switching operator ^d to a programming language for hybrid systems [Platzer 2012a] and reinterpreting its compositional operators as operators on hybrid games. This article introduces differential game logic for hybrid games with a simple modal semantics and a simple compositional proof calculus, which is proved to be a sound and complete axiomatization relative to any expressive logic. Completeness for game logics is a subtle problem. Completeness of propositional discrete game logic has been an open problem for 30 years [Parikh 1983]. This article focuses on more general hybrid games and proves a generalization of Parikh’s calculus to be relatively complete for hybrid games. The 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 [Coquand and Huet 1988]. This technique is practical for hybrid games, and also easier for hybrid systems than previous proof techniques. These results suggest hybrid game versions of influential views

²A closely related question is about ways to exhibit that winning strategy, for which existence is a prerequisite and a constructive proof is a representation of that winning strategy. A proof is a certificate witnessing the existence of a winning strategy. As soon as one knows from which states a winning strategy exists, local search in the action space would be enough. But search may still be challenging in dense action spaces.

³Hybrid games only lead to a minor syntactic change compared to hybrid systems (the addition of ^d), yet one that entails pervasive semantical reconsiderations, because the semantic basis for assigning meaning to operators changes in the presence of adversarial interactions. This change leads to more expressiveness. It is a sign of logical robustness that this results in a surprisingly small change in the axiomatization. Overall, the changes induced by dualities are in some ways radical, yet, in other ways surprisingly smooth.

of understanding program invariants as fixpoints [Cousot and Cousot 1977; Clarke 1979]. Harel’s convergence rule [Harel et al. 1977], which poses practical challenges for hybrid systems verification, now turns out to be unnecessary for hybrid games, hybrid systems, and programs. All separating axioms are identified that capture the logical difference of hybrid systems versus hybrid games. Hybrid games are proved to be 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. Winning regions of hybrid games are characterized by fixpoints of a monotone operator, which can be obtained by iteration but that iteration may only stop after $\geq \omega_1^{\text{CK}}$ many steps. Hybrid games are shown to be a fundamental extension of hybrid systems by proving that the logic dGL for hybrid games is fundamentally more expressive than the corresponding logic dL for hybrid systems, which is related to long-standing unsolved questions in the propositional case [Parikh 1985; Berwanger et al. 2007]. This separation also characterizes the expressiveness of dL and of dGL .

Structure of this Article. The syntax and denotational semantics of dGL are introduced in Section 2. Section 3 establishes meta-properties, including determinacy (Section 3.1), equivalences of hybrid games and reductions eliminating evolution domains (Section 3.2), and an analysis proving that winning regions of hybrid games can be characterized by iteration of a monotone image computation operator until a fixpoint, which may only stop after $\geq \omega_1^{\text{CK}}$ many steps (Section 3.3) unlike hybrid systems whose corresponding closure ordinal is ω . Section 4 presents an axiomatization of dGL as a Hilbert-type proof calculus (Section 4.1) that is proved sound (Section 4.2) and complete (Section 4.3) relative to any differentially expressive logic, which are exemplified subsequently (Section 4.4). The axiomatic separation between hybrid systems and hybrid games is identified in Section 4.5 and proved in Appendix B. Hybrid games are proved more expressive than hybrid systems in Section 5, where the expressiveness of the hybrid games logic dGL and the hybrid systems logic dL are characterized. Related work is discussed in Section 6, concluding in Section 7.

Example proofs in the dGL calculus are shown in Appendix A. Appendix B proves the axiomatic separation from Section 4.5. For reference and to support the interactive intuition of game play, an operational semantics of hybrid games is shown in Electronic Appendix A. Alternative semantics for repetitions of hybrid games are contrasted in Electronic Appendix B. Finally, Electronic Appendix C provides concrete hybrid games to support the computational intuition for the higher closure ordinals proved generically in Section 3.3.

2. DIFFERENTIAL GAME LOGIC

A robot is a canonical example of a hybrid system. Suppose a robot, WALL-E, is running around on a planet collecting trash. His dynamics is that of a hybrid system, because his continuous dynamics comes from the differential equations describing his continuous physical motion in space, while his discrete dynamics comes from his computer-based control decisions about when to move in which direction and when to stop moving in order to gather trash. As soon as WALL-E meets another robot, EVE, however, her presence changes everything for him. If WALL-E neither knows how EVE is programmed nor exactly what her goal is, then the only safe thing he can assume about her is that she might do anything within her physical capabilities. It takes the study of a hybrid game to find out whether or not WALL-E can use his choices in some way to reach his goal, say, collecting trash and avoiding collisions with EVE, regardless of how EVE chooses her actions.

The hybrid games considered here have no draws.⁴ For any particular play of the WALL·E and EVE game, for example, either WALL·E achieves his objective or he does not. There is no in between. When a hybrid game expects a player to move, but the rules of the game do not permit any of his moves from the current state, then that player loses right away (he *deadlocks*). 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 for complementary winning states, because the games are uncountably-infinite state but of arbitrary unbounded but finite duration. So reaching a winning state in the limit after infinitely many steps is not enough to win the game, but the players have to win in any arbitrary unbounded finite amount of time and after an arbitrary unbounded finite number of steps. The games are *zero-sum* games, i.e. if one player wins, the other one loses, with player payoffs ± 1 . Losses or victories of different payoff are not considered explicitly, because they are representable with extra variables that track payoffs. The two players are classically called *Angel* and *Demon*. By considering aggregate players, these results generalize in the usual way to the case where Angel and Demon represent coalitions of agents that work together to achieve a common goal such as the aggregate Angel player for one team and the aggregate Demon player for the other team in the case of RoboCup.

Hybrid games are non-cooperative and sequential games. In non-cooperative games, the players can choose to act arbitrarily according to the rules represented in the game.⁵ 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 their equivalent⁶ games of imperfect information, are interesting but not considered explicitly here [Alur et al. 2002; Berwanger and Pinchinat 2009]. Imperfect information games lead to Henkin quantifiers, not first-order quantifiers.

2.1. Syntax

Differential game logic (dGL) is a logic for studying properties of 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 to 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* with its dynamics and rules and choices, not the objective, they are still called hybrid games. Hybrid game forms represent the rules of the game, so during the game play, players can never win but only lose (prematurely) by violating these rules. The actual objective for a hybrid game is defined in the modal logical formula that refers to that hybrid game form and is evaluated at the end of the game play. Hybrid games and differential game logic formulas are defined by simultaneous induction focusing on polynomial terms for simplicity.

⁴ For applications with *draws*, it is easy to follow Zermelo [1913] and compare two games, one for each player, that attribute draws pessimistically as losses. Draws in the original game result from those states from which both players would lose their games when considering draws pessimistically as their respective losses. Draws come from states where no player can make sure to win-and-not-draw but only to win-or-draw. Hybrid games with draws are represented using two modalities of the logic dGL developed in this article.

⁵ Applications with *cooperative games*, where players form coalitions or negotiate binding contracts, are representable in the rules of the game to track coalitions and limit player's choices according to the contracts.

⁶ *Concurrent games* in which the players choose actions simultaneously can be converted into sequential games of imperfect information in which the players choose sequentially yet without information about the opponent's choice.

Definition 2.1 (Hybrid games). The *hybrid games of differential game logic* dGL are defined by the following grammar (α, β are hybrid games, x a variable, θ a (polynomial) term, ψ is a dGL formula):

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

Definition 2.2 (dGL formulas). The *formulas of differential game logic* dGL are defined by the following grammar (ϕ, ψ are dGL formulas, p is a predicate symbol of arity k , θ_i are (polynomial) terms, x a variable, and α is a hybrid game):

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

Other operators $>, =, \leq, <, \vee, \rightarrow, \leftrightarrow, \forall x$ or \top and \perp for true and false can be defined as usual, e.g., $\forall x \phi \equiv \neg \exists x \neg \phi$, including all program operators [Platzer 2010b]. The modal formula $\langle \alpha \rangle \phi$ expresses that Angel has a winning strategy⁷ to achieve ϕ in hybrid game α , i.e. Angel has a strategy to reach a state satisfying dGL formula ϕ when playing hybrid game α , no matter what strategy Demon chooses. That is, $\langle \alpha \rangle \phi$ expresses that Angel can guarantee to reach into the set of states satisfying ϕ but she cannot usually predict which of the states satisfying ϕ she will reach, because that depends on Demon's choices. The modal formula $[\alpha] \phi$ expresses that Demon has a winning strategy to achieve ϕ in hybrid game α , i.e. a strategy to reach a state 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 dGL formulas is the player whose winning strategy they refer to. Both use the set of states where dGL formula ϕ is true as the winning states for that player. The winning condition is defined by the modal formula, α only defines the hybrid game form, not when the game is won, which is what ϕ does. Hybrid game α defines the rules of the game, including conditions on state variables that, if violated, cause the present player to lose for violation of the rules of the game. The dGL formulas $\langle \alpha \rangle \phi$ and $[\alpha] \neg \phi$ consider complementary winning conditions for Angel and Demon.

The *atomic games* of dGL are assignments, continuous evolutions, and tests. The *discrete assignment game* $x := \theta$ instantly changes the value of variable x to that of θ by a discrete jump without any choices to resolve. In the *continuous evolution game* $x' = \theta \ \& \ \psi$, 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, at any time, take the state outside the region where formula ψ holds. In particular, Angel is deadlocked and loses immediately if ψ does not hold in the current state, because she cannot even evolve for duration 0 then without going outside ψ .⁸ The *test game* or *challenge* $?\psi$ has no effect on the state, except that Angel loses the game immediately if dGL formula ψ does not hold in the current state. If Angel passes the challenge $?\psi$, the game continues from the same state, otherwise she loses immediately.

The *compound games* of dGL are sequential, 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 lost already (so no challenge in α failed),

⁷A strategy for a player can be thought of as a function that selects one option whenever that player has a choice during the game play. A winning strategy for a player is a way of resolving choices that will lead to a state in which that player wins, no matter how the opponent player resolves his respective choices. The semantics of dGL is a denotational semantics based on winning regions. A formal definition for strategies and winning strategies is, thus, unnecessary, yet shown in Appendix A for reference.

⁸The most common case for ψ is a formula of first-order real arithmetic. In Section 3.2, evolution domain constraints ψ turn out to be unnecessary, because they can be defined using hybrid games. In the ordinary differential equation $x' = \theta$, the term x' denotes the time-derivative of x and θ is a polynomial term that is allowed to mention x and other variables. Systems of differential equations are considered vectorially. More general forms of differential equations are possible [Platzer 2010a], but will not be considered explicitly.

continues by playing game β . When playing the *choice game* $\alpha \cup \beta$, Angel chooses whether to play hybrid game α or play hybrid game β . Like all the other choices, this choice is dynamic, i.e. every time $\alpha \cup \beta$ is played, Angel gets to choose again whether she wants to play α or β this time. The *repeated game* α^* plays hybrid game α repeatedly and Angel chooses, after each play of α that terminates without a player having lost already, whether to play the game again or not, albeit she cannot choose to play indefinitely but has to stop repeating ultimately. Angel is also allowed to stop α^* right away after zero iterations of α .

Most importantly, the *dual game* α^d is the same as playing the hybrid game α with the roles of the players swapped during game play. 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, hence Demon loses in the dual game α^d when he deadlocks because those correspond situations where Angel is supposed to move but deadlocks in α and vice versa. Thus, while the test game $?\psi$ causes Angel to lose if formula ψ does not hold, the *dual test game* (or *dual challenge*) $(?\psi)^d$ causes Demon to lose if ψ does not hold. It is exactly the same formula ψ whose truth-value decides about the fate of the game in both cases since d does not affect the meaning of formulas within tests, but the player who loses changes. In $?\psi$, it is Angel who loses when ψ does not hold because she was supposed to move. In $(?\psi)^d$, it is Demon who loses when ψ does not hold as he was supposed to move. For example, if α describes the game of chess, then α^d is chess where the players switch to control the other side. Recall that hybrid games are game forms so d only affects the actions (and premature losses if no action is possible), not the objective of the game, which is what the modal formulas define. So, Angel has the same goal ϕ in $\langle \alpha^d \rangle \phi$ and $\langle \alpha \rangle \phi$ but her actions in α^d switched to what used to be Demon's actions in α and vice versa, which is why it will turn out that $\langle \alpha^d \rangle \phi$ and $[\alpha] \phi$ are equivalent (Section 4). The dual operator d is the only syntactic difference of dGL for hybrid games compared to dL for hybrid systems [Platzer 2008; Platzer 2012a], 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 Angel without interaction. The presence of d requires a thorough semantic generalization throughout the logic, though.

The logic dGL only provides logically essential operators. Many other game interactions for games can be defined from the elementary operators that dGL 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 α^\times , defined by $((\alpha^d)^*)^d$, in which α is repeated as often as Demon chooses to. In α^\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* $(x' = \theta \& \psi)^d$ follows the same dynamics as $x' = \theta \& \psi$ except that Demon chooses the duration, so he cannot choose a duration during which ψ stops to hold at any time. Hence he loses when ψ does not hold in the current state. Dual assignment $(x := \theta)^d$ is equivalent to $x := \theta$, because it is deterministic so involves no choices. Other program operators are also definable [Platzer 2010b], e.g., non-deterministic assignment $x := *$ defined by $x' = 1; x' = -1$. Unary operators (including $^d, \forall x, [\alpha], \langle \alpha \rangle$) bind stronger than binary operators and $;$ binds stronger than \cup and \cap , so $\alpha; \beta \cup \gamma \equiv (\alpha; \beta) \cup \gamma$.

Note that, quite unlike in the case of α^* and unlike in differential games [Platzer 2015a], it is irrelevant whether Angel decides the duration for $x' = \theta \& \psi$ before or after that continuous evolution, because initial-value problems for $x' = \theta$ have unique solutions by Picard-Lindelöf as term θ is smooth.

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, e.g., because his opponent failed a test. 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 to completion, win and lose multiple times, and study who wins how often in the long run with mixed strategies. A hybrid game is played once (even if some part of it constitutes in repeating action choices) and it 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.

Example 2.3 (WALL·E and EVE). Consider a game of the robots WALL·E and EVE moving on a (one-dimensional) planet.

$$\begin{aligned} (w - e)^2 \leq 1 \wedge v = f \rightarrow & \langle ((u := 1 \cap u := -1); \\ & (g := 1 \cup g := -1); \\ & t := 0; (w' = v, v' = u, e' = f, f' = g, t' = 1 \& t \leq 1)^d \rangle^\times \\ & \rangle (w - e)^2 \leq 1 \end{aligned} \quad (1)$$

Robot WALL·E is at position w with velocity v and acceleration u and plays the part of Demon. Robot EVE is at e with velocity f and acceleration g and plays the part of Angel. The antecedent of (1) before the implication assumes that WALL·E and EVE start close to one another (distance at most 1) and with identical velocities. The objective of EVE, who plays Angel's part in (1), is to be close to WALL·E (i.e. $(w - e)^2 \leq 1$) as specified after the $\langle \cdot \rangle$ modality in the succedent. The hybrid game proceeds as follows. Demon WALL·E controls how often the hybrid game repeats by operator \times . In each iteration, Demon WALL·E first chooses (\cap) to accelerate ($u := 1$) or brake ($u := -1$), then Angel EVE chooses (\cup) whether to accelerate ($g := 1$) or brake ($g := -1$). Every time that the \times loop repeats, the players get to make that choice again. They are not bound by what they chose in the previous iterations. Yet, depending on the previous choices, the state will have evolved differently, which influences indirectly what moves a player needs to choose to win. After this sequence of choices of u and g by Demon and Angel, respectively, a clock variable t is reset to $t := 0$. Then the game follows a differential equation system such that the time-derivative of WALL·E's position w is his velocity v and the time-derivative of v is acceleration u , the time-derivative of EVE's position e is her velocity f and the time-derivative of f is acceleration g . The time-derivative of clock variable t is 1, yet the differential equation is restricted to the evolution domain $t \leq 1$. Angel controls the duration of differential equations. Yet, this differential equation is within a dual game by operator d , so Demon controls the duration of the continuous evolution. Here, both WALL·E and EVE evolve continuously but Demon WALL·E decides how long. He cannot chose durations > 1 , because that would make him violate the evolution domain constraint $t \leq 1$ and lose, so the players can change their control after at most one time unit, but Demon decides when exactly. Similar games can be studied for robot motion in higher dimensions using dGL.

The two players in Example 2.3 use various dualities to model a situation where Demon chooses discretely (\cap), then Angel chooses discretely (\cup), then Demon chooses the duration for the joint continuous evolution ($((w' = \dots \& t \leq 1)^d)$ and finally Demon decides about repetition (\times), which gives 6 dual operators with nesting depth 4 (or 3 when discounting dual assignments). Deeper nesting levels of hybrid game operators, which give rise to longer chains of d operators, can be used to describe hybrid games with more levels of interaction (any arbitrary number of nested $\cup, ^d, *$).

Example 2.4 (WALL·E and EVE and the world). The game in (1) accurately reflects the situation when WALL·E is in control of time since the (only) differential equation occurs within an odd number of d operators. But to design the model (1), EVE may have used a common modeling device to conservatively attribute the control of the differential equation to WALL·E, even if time is really under control of a third player, the external environment. EVE's reason for this model would be that she is not in control of time, so there is no reason to believe why time would help her. EVE, thus, conservatively ceases control of time to Demon, which corresponds to assuming that the third player of the external environment is allowed to collaborate with WALL·E to form an aggregate Demon player consisting of WALL·E and the environment.

When WALL·E wants to analyze his winning strategies with a $[\cdot]$ variation of (1), he could use the same modeling device to flip the differential equation over to Angel's control by removing the d to conservatively associate the environment to the opponent:

$$\begin{aligned} (w - e)^2 \leq 1 \wedge v = f \rightarrow & [((u := 1 \cap u := -1); \\ & (g := 1 \cup g := -1); \\ & t := 0; (w' = v, v' = u, e' = f, f' = g, t' = 1 \ \& \ t \leq 1))^\times \\ &] (w - e)^2 > 1 \end{aligned} \quad (2)$$

If, instead, WALL·E keeps the same hybrid game as in (2), just with $[\cdot]$, the game considers the situation when WALL·E has control over time, which would make some part of the game trivial, because, once he reached $(w - e)^2 > 1$, the Demon WALL·E can then just always evolve for 0 time units. Observe how a three-player game of WALL·E, EVE, and environment can be analyzed by combining the dGL formulas (1) and (2) propositionally, which then analyze the same game from different perspectives of possible collaborations. The dGL formula expressing that neither (1) nor (2) is true, for example, is true in exactly the states where WALL·E and EVE draw, because the disturbance of the external environment can choose the winner by helping either WALL·E or EVE.

When the role of such environments is not limited to influencing only differential equation durations, the appropriate placement of d operators around additional choices (\cup) or repetitions (*) or multiple differential equations gives models more flexibility. Similar phenomena happen when further players are added to the game as long as their role can be described by logical combinations of dGL formulas investigating the game from a set of aggregated two-player perspectives. Surprisingly, there is virtually no limit to how far logic around such essentially two-player hybrid games extends to seemingly substantially more general situations by appropriate modeling (Section 5).

2.2. Semantics

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

A *state* 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 . The interpretation further determines the set of states S , which is isomorphic to a Euclidean space \mathbb{R}^n when n is the number of relevant variables. For a subset $X \subseteq S$ the complement $S \setminus X$ is denoted X^c . Let s_x^r denote the state that agrees with state s except for the interpretation of variable x , which is changed to $r \in \mathbb{R}$. The value of term θ in state s is denoted by $\llbracket \theta \rrbracket_s$. The denotational semantics of dGL formulas will be defined in Def. 2.5 by simultaneous induction along

with the denotational semantics, $\varsigma_\alpha(\cdot)$ and $\delta_\alpha(\cdot)$, of hybrid games, defined in Def. 2.6, because dGL formulas are defined by simultaneous induction with hybrid games.

Definition 2.5 (dGL semantics). The semantics of a dGL formula ϕ for each interpretation I with a corresponding set of states \mathcal{S} is the subset $\llbracket \phi \rrbracket^I \subseteq \mathcal{S}$ of states in which ϕ is true. It is defined inductively as follows

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

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

Definition 2.6 (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 in α (whatever strategy Demon chooses). It is defined inductively as follows⁹

- (1) $\varsigma_{x=\theta}(X) = \{s \in \mathcal{S} : s_x^{\llbracket \theta \rrbracket_s} \in X\}$
- (2) $\varsigma_{x'=\theta \ \& \ \psi}(X) = \{\varphi(0) \in \mathcal{S} : \varphi(r) \in X \text{ for some } r \in \mathbb{R}_{\geq 0} \text{ and (differentiable) } \varphi : [0, r] \rightarrow \mathcal{S} \text{ such that } \varphi(\zeta) \in \llbracket \psi \rrbracket^I \text{ and } \frac{d\varphi(t)(x)}{dt}(\zeta) = \llbracket \theta \rrbracket_{\varphi(\zeta)} \text{ for all } 0 \leq \zeta \leq r\}$
- (3) $\varsigma_{? \psi}(X) = \llbracket \psi \rrbracket^I \cap X$
- (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) $\varsigma_{\alpha^d}(X) = (\varsigma_\alpha(X^c))^c$

The *winning region* of Demon, i.e. the set of states $\delta_\alpha(X)$ from which Demon has a winning strategy to achieve X in α (whatever strategy Angel chooses) is defined inductively as follows

- (1) $\delta_{x=\theta}(X) = \{s \in \mathcal{S} : s_x^{\llbracket \theta \rrbracket_s} \in X\}$
- (2) $\delta_{x'=\theta \ \& \ \psi}(X) = \{\varphi(0) \in \mathcal{S} : \varphi(r) \in X \text{ for all } r \in \mathbb{R}_{\geq 0} \text{ and (differentiable) } \varphi : [0, r] \rightarrow \mathcal{S} \text{ such that } \varphi(\zeta) \in \llbracket \psi \rrbracket^I \text{ and } \frac{d\varphi(t)(x)}{dt}(\zeta) = \llbracket \theta \rrbracket_{\varphi(\zeta)} \text{ for all } 0 \leq \zeta \leq r\}$
- (3) $\delta_{? \psi}(X) = (\llbracket \psi \rrbracket^I)^c \cup X$
- (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) $\delta_{\alpha^d}(X) = (\delta_\alpha(X^c))^c$

⁹ The semantics of a hybrid game is not merely a reachability relation between states as for hybrid systems [Platzer 2012a], because the adversarial dynamic interactions and nested choices of the players have to be taken into account.

The notation uses $\varsigma_\alpha(X)$ and $\delta_\alpha(X)$ instead of $\varsigma_\alpha^I(X)$ and $\delta_\alpha^I(X)$, because the interpretation I that gives a semantics to predicate symbols in tests and evolution domains is clear from the context. Strategies do not occur explicitly in the dGL semantics, because it is based on winning regions, i.e. the existence of winning strategies, not on the strategies themselves. The winning regions for Angel are illustrated in Fig. 1. The winning region $\varsigma_{\alpha^*}(X)$ is the smallest set Z around X that already contains all states from which one more $\varsigma_\alpha(\cdot)$ could win into Z , so $\varsigma_\alpha(\varsigma_{\alpha^*}(X)) \setminus \varsigma_{\alpha^*}(X) = \emptyset$. It includes all κ -fold iterations $\varsigma_\alpha^\kappa(X)$ of the winning region construction $\varsigma_\alpha(\cdot)$ (Section 3.3) for all κ .

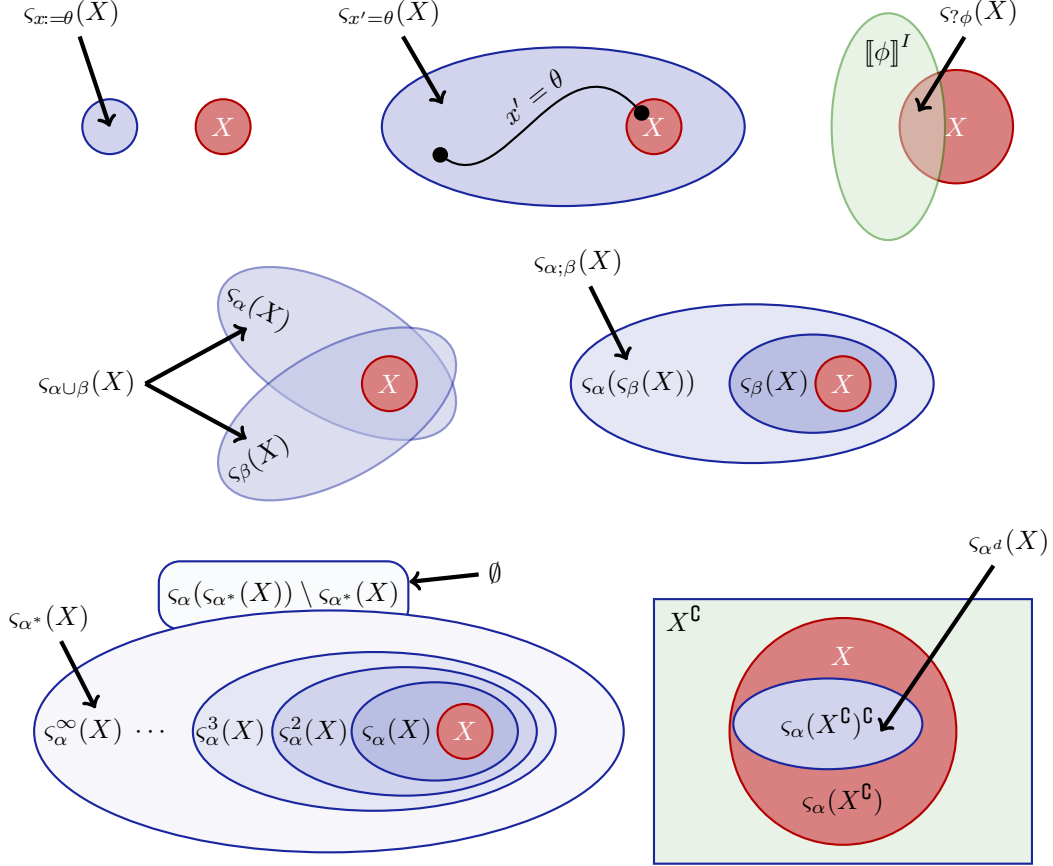


Fig. 1. Illustration of denotational semantics of hybrid games as winning regions

The semantics is *compositional*, i.e. the semantics of a compound dGL formula is a simple function of the semantics of its pieces, and the semantics of a compound hybrid game is a function of the semantics of its pieces. This guarantees referential transparency and enables a compositional proof calculus. Furthermore, existence of a strategy in hybrid game α to achieve X is independent of games and dGL formulas 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 [Zermelo 1913], but is directly apparent in a logical setting. Furthermore, the semantics is monotone, i.e. larger sets of winning states induce larger winning regions.

LEMMA 2.7 (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 $X^c \supseteq Y^c$, hence $\varsigma_\alpha(X^c) \supseteq \varsigma_\alpha(Y^c)$, so $\varsigma_{\alpha^d}(X) = (\varsigma_\alpha(X^c))^c \subseteq (\varsigma_\alpha(Y^c))^c = \varsigma_{\alpha^d}(Y)$. \square

Monotonicity implies that the least fixpoint in $\varsigma_{\alpha^*}(X)$ and the greatest fixpoint in $\delta_{\alpha^*}(X)$ are well-defined [Harel et al. 2000, 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. In particular, Angel cannot play a Zeno strategy with infinitely many steps in finite time. 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^c 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 α^\times also branch finitely, because, after each repetition of α , the respective player (Angel for α^* and Demon for α^\times) may decide whether to repeat again or stop. Repeated games still lead to infinitely many branches, because a repeated game can be repeated any arbitrary number of times. The game branches uncountably infinitely, however, when the players decide how long to evolve along differential equations in $x' = \theta \& \psi$ and $(x' = \theta \& \psi)^d$, because uncountably many nonnegative real number could be chosen as a duration (unless the system leaves ψ immediately). These choices can be made explicit by relating the simple denotational modal semantics of dGL to an equivalent operational game semantics that is technically much more involved but directly exposes the interactive intuition of game play. For reference, this approach has been made explicit in Appendix A.

Example 2.8. The following simple dGL formula

$$\langle (x := x + 1; (x' = x^2)^d \cup x := x - 1)^* \rangle (0 \leq x < 1) \quad (3)$$

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, every time she does, it is her choice (\cup) whether to increase x by 1 and then (after ;) give Demon control over the duration of the differential equation $x' = x^2$ (left game) or whether to instead decrease x by 1 (right game). Formula (3) is valid, because Angel has the winning strategy of choosing the left action $x := x + 1; (x' = x^2)^d$ until $x \geq 0$ followed by the right action $x := x - 1$ until $0 \leq x < 1$. By choosing the left action, $x \geq 0$ will ultimately happen, because $x := x + 1$ increases x by 1 and $x' = x^2$ can only make x bigger, because the derivative x^2 is nonnegative. Once $x \geq 0$, choosing the right action suitably often will reach the postcondition $0 \leq x < 1$ to allow Angel to win the game. Note that Angel also wins immediately with the left action from $x = -1$, since the differential equation is stuck at $x = 0$. The following minor variation, however, is not valid:

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

because Demon can spoil Angel's efforts by choosing $x := x - 2$ in his choice (\cap) to make x negative whenever $1 \leq x < 2$, and then increasing x to 1.5 again via $(x' = x^2)^d$ when Angel takes the left choice. Angel will never reach $0 \leq x < 1$ that way unless this was true initially already. This phenomenon is examined in Section 3.1 in more detail.

Example 2.9 (WALL-E and EVE). The dGL formula (1) from Example 2.3 is valid, because Angel EVE indeed has a winning strategy to get close to WALL-E by mimicking Demon's choices. Recall that Demon WALL-E controls the repetition \times , so the fact that the hybrid game starts EVE off close to WALL-E is not sufficient for EVE to win the game. The hybrid game in (1) would be trivial if Angel were to control the repetition (because she would then win just by choosing not to repeat) or were to control the differential equation (because she would then win by always evolving for duration 0). The analysis of (1) is more difficult if the first two lines in the hybrid game are swapped so that Angel EVE chooses g before Demon WALL-E chooses u .

3. META-PROPERTIES

This section analyzes meta-properties and semantical properties of the hybrid games of dGL, including determinacy of hybrid games, hybrid game equivalences such as reduction of evolution domains, and closure ordinals of hybrid games.

3.1. Determinacy

Every particular game play in a hybrid game is won by exactly one player, because hybrid games are zero-sum and there are no draws. That alone does not imply determinacy, i.e. that, from all initial situations, 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, consider the semantics of repetition, defined as a *least* fixpoint, which is crucial because that gives a well-founded repetition. Otherwise, the *filibuster formula* would not have a well-defined truth-value:

$$\langle (x := 0 \cap x := 1)^* \rangle x = 0 \quad (4)$$

It is Angel's choice whether to repeat (*), but every time Angel repeats, it is Demon's choice (\cap) whether to play $x := 0$ or $x := 1$. The game in this formula never deadlocks, 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 but just prevents the other player from winning; see Fig. 2.

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 $\zeta_{\alpha^*}(\cdot)$ in the semantics, however, repetitions are well-founded and, thus, have to stop eventually (after an arbitrary unbounded number of rounds). Hence, in the example in Fig. 2, Demon still wins and formula (4) is \perp , unless $x = 0$ holds initially (the unknown initial value is marked X in Fig. 2). In other words, the formula in (4) is equivalent to $x = 0$. The same phenomenon happens in Example 2.8. Likewise, the dual filibuster game formula

$$x = 0 \rightarrow \langle (x := 0 \cup x := 1)^\times \rangle x = 0 \quad (5)$$

is (determined and) valid, because Demon has to stop repeating \times eventually so that Angel wins if she patiently plays $x := 0$ each time. Similarly, the game in the following




Fig. 2. The filibuster game formula $\langle\langle x := 0 \cap x := 1 \rangle^*\rangle x = 0$ 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 bisimilar 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 .

hybrid filibuster formula would not be determined without the least fixpoint semantics

$$\langle\langle x := 0; x' = 1^d \rangle^*\rangle x = 0$$

because Demon could always evolve continuously to some state where $x > 0$ and Angel would never want to stop. Since Angel will have to stop eventually, she loses and the formula is \perp unless $x = 0$ holds initially.

It is also important that Angel can only choose real durations $r \in \mathbb{R}_{\geq 0}$ for a continuous evolution game $x' = \theta \& \psi$, not infinity ∞ , so she ultimately stops. Otherwise

$$\langle\langle x' = 1^d; x := 0 \rangle^*\rangle x = 0 \quad (6)$$

would not be determined, because Angel wants to repeat (unless $x = 0$ initially) and $x := 0$ will make her win once she stops after a nonzero number of repetitions. Yet, if Demon could choose ∞ as the duration for the continuous evolution game $x' = 1^d$, Angel would never get to play the subsequent $x := 0$ to win. Since durations need to be real numbers, though, each continuous evolution ultimately has to stop, so the formula in (6) is valid even if Demon can postpone Angel’s victory arbitrarily long.

In order to make sure that dGL is a classical two-valued modal logic, hybrid games have no draws. But, because modalities refer to the existence of winning strategies, they only receive classical truth values if, from each state, one of the players has a winning strategy for complementary winning conditions of a hybrid game α . The log-

ical setup of dGL makes this determinacy proof very simple, without the need to use, e.g., the deep [Friedman 1971b] Borel determinacy theorem for winning conditions that are Borel in the product topology induced on game trees by the discrete topology of actions [Martin 1975], which does not even fit to the structure of arbitrarily nested inductive and coinductive fixpoints of the winning region semantics of dGL.

THEOREM 3.1 (CONSISTENCY & DETERMINACY). *Hybrid games are consistent and determined, i.e. $\models \neg\langle\alpha\rangle\neg\phi \leftrightarrow [\alpha]\phi$.*

PROOF. The proof shows by a straightforward induction on the structure of α that $\varsigma_\alpha(X^{\mathbb{G}})^{\mathbb{G}} = \delta_\alpha(X)$ 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 $X \stackrel{\text{def}}{=} \llbracket\phi\rrbracket^I$.

- (1) $\varsigma_{x=\theta}(X^{\mathbb{G}})^{\mathbb{G}} = \{s \in \mathcal{S} : s_x^{\llbracket\theta\rrbracket^s} \notin X\}^{\mathbb{G}} = \varsigma_{x=\theta}(X) = \delta_{x=\theta}(X)$
- (2) $\varsigma_{x'=\theta \& \psi}(X^{\mathbb{G}})^{\mathbb{G}} = \{\varphi(0) \in \mathcal{S} : \varphi(r) \notin X \text{ for some } 0 \leq r \in \mathbb{R} \text{ and some (differentiable) } \varphi : [0, r] \rightarrow \mathcal{S} \text{ such that } \frac{d\varphi(t)(x)}{dt}(\zeta) = \llbracket\theta\rrbracket_{\varphi(\zeta)} \text{ and } \varphi(\zeta) \in \llbracket\psi\rrbracket^I \text{ for all } 0 \leq \zeta \leq r\}^{\mathbb{G}} = \delta_{x'=\theta \& \psi}(X)$, because the set of states from which there is no winning strategy for Angel to reach a state in $X^{\mathbb{G}}$ prior to leaving $\llbracket\psi\rrbracket^I$ along $x' = \theta \& \psi$ is exactly the set of states from which $x' = \theta \& \psi$ always stays in X (until leaving $\llbracket\psi\rrbracket^I$ in case that ever happens).
- (3) $\varsigma_{? \psi}(X^{\mathbb{G}})^{\mathbb{G}} = (\llbracket\psi\rrbracket^I \cap X^{\mathbb{G}})^{\mathbb{G}} = (\llbracket\psi\rrbracket^I)^{\mathbb{G}} \cup (X^{\mathbb{G}})^{\mathbb{G}} = \delta_{? \psi}(X)$
- (4) $\varsigma_{\alpha \cup \beta}(X^{\mathbb{G}})^{\mathbb{G}} = (\varsigma_\alpha(X^{\mathbb{G}}) \cup \varsigma_\beta(X^{\mathbb{G}}))^{\mathbb{G}} = \varsigma_\alpha(X^{\mathbb{G}})^{\mathbb{G}} \cap \varsigma_\beta(X^{\mathbb{G}})^{\mathbb{G}} = \delta_\alpha(X) \cap \delta_\beta(X) = \delta_{\alpha \cup \beta}(X)$
- (5) $\varsigma_{\alpha; \beta}(X^{\mathbb{G}})^{\mathbb{G}} = \varsigma_\alpha(\varsigma_\beta(X^{\mathbb{G}}))^{\mathbb{G}} = \varsigma_\alpha(\delta_\beta(X))^{\mathbb{G}} = \delta_\alpha(\delta_\beta(X)) = \delta_{\alpha; \beta}(X)$
- (6) $\varsigma_{\alpha^*}(X^{\mathbb{G}})^{\mathbb{G}} = \left(\bigcap\{Z \subseteq \mathcal{S} : X^{\mathbb{G}} \cup \varsigma_\alpha(Z) \subseteq Z\}\right)^{\mathbb{G}} = \left(\bigcap\{Z \subseteq \mathcal{S} : (X \cap \varsigma_\alpha(Z))^{\mathbb{G}} \subseteq Z\}\right)^{\mathbb{G}}$
 $= \left(\bigcap\{Z \subseteq \mathcal{S} : (X \cap \delta_\alpha(Z))^{\mathbb{G}} \subseteq Z\}\right)^{\mathbb{G}} = \bigcup\{Z \subseteq \mathcal{S} : Z \subseteq X \cap \delta_\alpha(Z)\} = \delta_{\alpha^*}(X)$.¹⁰
- (7) $\varsigma_{\alpha^d}(X^{\mathbb{G}})^{\mathbb{G}} = (\varsigma_\alpha((X^{\mathbb{G}})^{\mathbb{G}}))^{\mathbb{G}} = \delta_\alpha(X^{\mathbb{G}})^{\mathbb{G}} = \delta_{\alpha^d}(X) \quad \square$

One direction of Theorem 3.1 implies $\models \neg\langle\alpha\rangle\neg\phi \rightarrow [\alpha]\phi$, i.e. $\models \langle\alpha\rangle\neg\phi \vee [\alpha]\phi$, whose validity means that, from all initial states, 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 3.1 implies $\models [\alpha]\phi \rightarrow \neg\langle\alpha\rangle\neg\phi$, i.e. $\models \neg([\alpha]\phi \wedge \langle\alpha\rangle\neg\phi)$, whose validity means that there is no state from which Demon has a winning strategy to achieve ϕ and, simultaneously, Angel has a winning strategy to achieve $\neg\phi$. It cannot be that both players have a winning strategy for complementary conditions from the same state. 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 dGL, Theorem 3.1 makes dGL a classical (multi)modal logic [Chellas 1980], yet with modalities indexed by hybrid games.

Instead of giving a semantics to $[\cdot]$ in terms of the existence of a winning strategy for Demon, Theorem 3.1 could have been used as a definition of $[\cdot]$. That would have been easier, but would have obscured determinacy and the role of $[\cdot]$ as the winning strategy operator for Demon.

¹⁰The penultimate equation follows from the μ -calculus equivalence $\nu Z. \Upsilon(Z) \equiv \neg\mu Z. \neg\Upsilon(\neg Z)$ and the fact that least pre-fixpoints are fixpoints and that greatest post-fixpoints are fixpoints for monotone functions.

3.2. Hybrid Game Equivalences

As usual in programming languages, the same hybrid game can have multiple different syntactical representations. Some equivalence transformations on hybrid games can be useful to transform hybrid games into a conceptually simpler form.

Definition 3.2 (Hybrid game equivalence). Hybrid games α and β are *equivalent*, denoted $\alpha \equiv \beta$, if $\varsigma_\alpha(X) = \varsigma_\beta(X)$ for all X and all I .

By Theorem 3.1, α and β are equivalent iff $\delta_\alpha(X) = \delta_\beta(X)$ for all X and all I .

Remark 3.3. The equivalences

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

on hybrid games can transform every hybrid game α 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 $(x' = \theta)^* \equiv x' = \theta$ and $(x' = \theta \& \psi)^* \equiv ?\top \cup x' = \theta \& \psi$.

Quite unlike in hybrid systems and (poor test) differential dynamic logic [Platzer 2008; Platzer 2012a], every hybrid game containing a differential equation $x' = \theta \& \psi$ with evolution domain constraints ψ can be replaced equivalently by a hybrid game without evolution domain constraints (even using poor tests, i.e. each test $?\psi$ uses only first-order formulas ψ). Evolution domains are definable in hybrid games and can, thus, be removed equivalently from all hybrid games.

LEMMA 3.4 (DOMAIN REDUCTION). *Evolution domains of differential equations are definable as hybrid games: For every hybrid game there is an equivalent hybrid game that has no evolution domain constraints, i.e. all continuous evolutions are of the form $x' = \theta$.*

PROOF. For notational convenience, 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 $x' = \theta(x) \& \psi(x)$ is equivalent to the hybrid game:

$$t_0 := x_0; x' = \theta(x); (z := x; z' = -\theta(z))^d; ?(z_0 \geq t_0 \rightarrow \psi(z)) \quad (7)$$




Fig. 3. “There and back again game”: 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 vector x

See Fig. 3 for an illustration. Suppose the current player is Angel. The idea behind game equivalence (7) is that the fresh variable t_0 remembers the initial time x_0 , and Angel then evolves forward along $x' = \theta(x)$ for any arbitrary amount of time (Angel’s choice). Afterwards, the opponent Demon copies the state x into a fresh variable (vector) z that he can evolve backwards along $(z' = -\theta(z))^d$ for any arbitrary amount of time (Demon’s choice). The original player Angel must then pass the challenge

?($z_0 \geq t_0 \rightarrow \psi(z)$), i.e. Angel loses immediately if Demon was able to evolve backwards and leave region $\psi(z)$ 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). \square

Lemma 3.4 eliminates all evolution domain constraints equivalently in hybrid games from now on. While evolution domain constraints are fundamental parts of hybrid systems [Alur et al. 1992; Henzinger et al. 1995; Branicky et al. 1998; Platzer 2008], they turn out to be mere syntactic sugar for hybrid games. In that sense, hybrid games are more fundamental than hybrid systems, because they feature elementary operators and do not need built-in support for evolution domain constraints.

3.3. Strategic Closure Ordinals

In order to examine whether the dGL semantics could be implemented directly to compute winning regions for dGL formulas by a reachability computation or backwards induction for games, this section investigates how many iterations the fixpoint for the semantics $\varsigma_{\alpha^*}(X)$ of repetition needs. The number of required iterations marks a significant difference in analytic complexity of hybrid games compared to hybrid systems.

The semantics, $\varsigma_{\alpha^*}(X)$, of α^* is a least fixpoint and Knaster-Tarski's seminal fixpoint theorem entails that every least fixpoint of a monotone function on a complete lattice corresponds to some sufficiently large iteration. That is, there is some ordinal $\bar{\lambda}$ at which the $\bar{\lambda}$ th iteration, $\varsigma_{\alpha}^{\bar{\lambda}}(X)$, of $\varsigma_{\alpha}(\cdot)$ coincides with $\varsigma_{\alpha^*}(X)$, i.e. $\varsigma_{\alpha^*}(X) = \varsigma_{\alpha}^{\bar{\lambda}}(X)$; see Fig. 4. How big is $\bar{\lambda}$, i.e. how often does $\varsigma_{\alpha}(\cdot)$ need to iterate to obtain $\varsigma_{\alpha^*}(X)$?




Fig. 4. Least fixpoint $\varsigma_{\alpha^*}(X)$ corresponds to some higher iterate $\varsigma_{\alpha}^{\bar{\lambda}}(X)$ of $\varsigma_{\alpha}(\cdot)$ from winning condition X .

Recall that ordinals extend natural numbers and support (non-commutative) addition, multiplication, and exponentiation, ordered as:

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

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

3.3.1. Iterations and Fixpoints. For each hybrid game α , the semantics $\varsigma_\alpha(\cdot)$ is a monotone operator on the complete powerset lattice (Lemma 2.7). The κ th iterate, $\varsigma_\alpha^\kappa(\cdot)$, of $\varsigma_\alpha(\cdot)$ is defined by a minor variation of Kozen's formulation of Knaster-Tarski [Harel et al. 2000, Theorem 1.12], obtained by considering the sublattice with x at the bottom.

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

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

Yet, \bigcup and, thus, $\tau^\lambda(x)$ are only guaranteed to exist if L is a complete partial order.

THEOREM 3.5 (KNASTER-TARSKI [HAREL ET AL. 2000, THEOREM 1.12]). *For every complete lattice L , there is an ordinal $\bar{\lambda}$ of at most the cardinality of L such that, for each monotone $\tau : L \rightarrow L$, 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 $x \in L$ and all ordinals 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 the property in Theorem 3.5 is called *closure ordinal* of τ .

The operator $\tau^\kappa(\cdot)$ enjoys useful properties. By its extensive / inflationary definition, $\tau^\kappa(x)$ is not just monotone in x but also monotone and homomorphic in κ . Since $\tau^0(x) = x$, this works for all ordinals.

LEMMA 3.6 (INDUCTIVE HOMOMORPHISM). *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 [Harel et al. 2000, 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):

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

By Theorem 3.5, there is an ordinal $\bar{\lambda}$ of cardinality at most that of \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 2.7. This iterative construction $\tau^{\bar{\lambda}}(X)$ corresponds to backward induction in classical game theory [von Neumann and Morgenstern 1955; Aumann 1995], yet it terminates at ordinal $\bar{\lambda}$ which is not necessarily finite.

3.3.2. Scott-Continuity. The semantics of repetitions in hybrid systems repeats some finite number of times [Platzer 2012a]. If the semantics of dGL were Scott-continuous, this would be the case for dGL as well, because the closure ordinal of Scott-continuous operators on a complete partial order is $\leq \omega$ by Kleene's fixpoint theorem. Dual-free α are indeed Scott-continuous, in particular, the closure ordinal for hybrid systems is ω .

LEMMA 3.7 (SCOTT-CONTINUITY OF d -FREE dGL). *The dGL semantics of d -free α 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 index set J .*

PROOF. By Lemma 2.7, $\bigcup_{n \in J} \varsigma_\alpha(X_n) \subseteq \varsigma_\alpha(\bigcup_{n \in J} X_n)$. The converse inclusion can be shown 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) $\varsigma_{x:=\theta}(\bigcup_{n \in J} X_n) = \{s \in \mathcal{S} : s_x^{\llbracket \theta \rrbracket_s} \in \bigcup_{n \in J} X_n\} \subseteq \bigcup_{n \in J} \{s \in \mathcal{S} : s_x^{\llbracket \theta \rrbracket_s} \in X_n\} = \bigcup_{n \in J} \varsigma_{x:=\theta}(X_n)$, since $s_x^{\llbracket \theta \rrbracket_s} \in \bigcup_{n \in J} X_n$ implies $s_x^{\llbracket \theta \rrbracket_s} \in X_n$ for some n .
- (2) $\varsigma_{x'=\theta \& \psi}(\bigcup_{n \in J} X_n) = \{\varphi(0) \in \mathcal{S} : \frac{d\varphi(t)(x)}{dt}(\zeta) = \llbracket \theta \rrbracket_{\varphi(\zeta)} \text{ and } \varphi(\zeta) \in \llbracket \psi \rrbracket^I \text{ for all } \zeta \leq r \text{ for some (differentiable) } \varphi : [0, r] \rightarrow \mathcal{S} \text{ such that } \varphi(r) \in \bigcup_{n \in J} X_n\} \subseteq \bigcup_{n \in J} \varsigma_{x'=\theta \& \psi}(X_n) = \{\varphi(0) \in \mathcal{S} : \dots \varphi(r) \in X_n\}$, because $\varphi(r) \in \bigcup_{n \in J} X_n$ implies $\varphi(r) \in X_n$ for some n .
- (3) $\varsigma_{\psi}(\bigcup_{n \in J} X_n) = \llbracket \psi \rrbracket^I \cap \bigcup_{n \in J} X_n = \bigcup_{n \in J} (\llbracket \psi \rrbracket^I \cap X_n) = \bigcup_{n \in J} \varsigma_{\psi}(X_n)$
- (4) $\varsigma_{\alpha \cup \beta}(\bigcup_{n \in J} X_n) = \varsigma_\alpha(\bigcup_{n \in J} X_n) \cup \varsigma_\beta(\bigcup_{n \in J} X_n) \stackrel{\text{IH}}{=} (\bigcup_{n \in J} \varsigma_\alpha(X_n)) \cup (\bigcup_{n \in J} \varsigma_\beta(X_n)) = \bigcup_{n \in J} (\varsigma_\alpha(X_n) \cup \varsigma_\beta(X_n)) = \bigcup_{n \in J} \varsigma_{\alpha \cup \beta}(X_n)$
- (5) $\varsigma_{\alpha; \beta}(\bigcup_{n \in J} X_n) = \varsigma_\alpha(\varsigma_\beta(\bigcup_{n \in J} X_n)) \stackrel{\text{IH}}{=} \varsigma_\alpha(\bigcup_{n \in J} \varsigma_\beta(X_n)) \stackrel{\text{IH}}{=} \bigcup_{n \in J} \varsigma_\alpha(\varsigma_\beta(X_n)) = \bigcup_{n \in J} \varsigma_{\alpha; \beta}(X_n)$
- (6) $\varsigma_{\alpha^*}(\bigcup_{n \in J} X_n) = (\bigcup_{n \in J} X_n) \cup \varsigma_\alpha(\varsigma_{\alpha^*}(\bigcup_{n \in J} X_n))$ is the least fixpoint. Prove that $\bigcup_{n \in J} \varsigma_{\alpha^*}(X_n)$ is a fixpoint, which implies $\varsigma_{\alpha^*}(\bigcup_{n \in J} X_n) \subseteq \bigcup_{n \in J} \varsigma_{\alpha^*}(X_n)$. Indeed, $(\bigcup_{n \in J} X_n) \cup \varsigma_\alpha(\bigcup_{n \in J} \varsigma_{\alpha^*}(X_n)) \stackrel{\text{IH}}{=} (\bigcup_{n \in J} X_n) \cup \bigcup_{n \in J} \varsigma_\alpha(\varsigma_{\alpha^*}(X_n)) = \bigcup_{n \in J} (X_n \cup \varsigma_\alpha(\varsigma_{\alpha^*}(X_n))) = \bigcup_{n \in J} \varsigma_{\alpha^*}(X_n)$. The last equation uses that $\varsigma_{\alpha^*}(X_n)$ is a fixpoint. \square

But $\varsigma_\alpha(\cdot)$ is not generally Scott-continuous, so $\bar{\lambda}$ might potentially be greater than ω for hybrid games. Games with both d and $*$ 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 ω as index set:

$$\mathbb{R} = \varsigma_{y:=y+1^\times} \left(\bigcup_{n < \omega} (-\infty, n] \right) \not\subseteq \bigcup_{n < \omega} \varsigma_{y:=y+1^\times} ((-\infty, n]) = \emptyset$$

$$\text{hence } \models \langle y := y + 1^\times \rangle \exists n : \mathbb{N} \ y \leq n \text{ but } \not\models \exists n : \mathbb{N} \ \langle y := y + 1^\times \rangle y \leq n$$

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 4.5). The quantifier $\exists n : \mathbb{N}$ over natural numbers is not essential here [Platzer 2008] but mere convenience to make both lines above match directly.

If τ is countably-continuous, i.e. continuous for families with countable index sets, on a complete partial order, then its closure ordinal is $\bar{\lambda} \leq \omega_1$. But this is not the case for $\varsigma_\alpha(\cdot)$ either, by the above counterexample with countable index set ω .

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 $< \kappa$. If τ is κ -based, then its closure ordinal is $\leq \kappa$ [Aczel 1977, Proposition 1.3.4]. The semantics $\varsigma_\alpha(\cdot)$ is not even ω_1 -based, however, because of Lemma 2.7 and removing just one state from the winning condition may lose states in the winning region:

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

Consequently, classical bounds on closure ordinals all fail to apply due to the combination of $^d, ^*$, and differential equations that makes hybrid games challenging.

3.3.3. Transfinite Closure Ordinals. When will the iteration for the fixpoints in the winning region definitions stop? Hybrid games may have higher closure ordinals, because ω many repetitions of the operator (and even $< \omega_1^{\text{CK}}$ many) may not be enough to compute winning regions. In other words, $\varsigma_{\alpha^*}(X)$ will coincide with iterations $\varsigma_{\alpha}^{\kappa}(X)$ as illustrated in Fig. 4, but this may need many more than ω iterations to terminate.

THEOREM 3.8 (CLOSURE ORDINALS). *The semantics of dGL has a closure ordinal $\geq \omega_1^{\text{CK}}$, i.e. for all $\lambda < \omega_1^{\text{CK}}$, there are α and X such that $\varsigma_{\alpha^*}(X) \neq \varsigma_{\alpha}^{\lambda}(X)$.*

PROOF. For concreteness, the proof first shows the weaker bound $\geq \omega \cdot 2$. Minor syntactic variations lead to vastly different closure ordinals (Appendix C), so the closure ordinal is not a simple function of the syntactic structure.

To see that the closure ordinal is $> \omega$ even with just one variable, a single loop and dual, consider the set of states in which the following dGL formula is true:

$$\langle \underbrace{(x := x + 1; x' = 1^d)}_{\alpha} \cup \underbrace{(x := x - 1)}_{\beta} \rangle^* (0 \leq x < 1) \quad (8)$$

The winning regions for this dGL formula stabilize after $\omega \cdot 2$ iterations, because ω many iterations are necessary to show that *all* positive reals 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 some positive real. It is easy to see that $\varsigma_{\alpha \cup \beta}^{\omega}([0, 1)) = \bigcup_{n < \omega} \varsigma_{\alpha \cup \beta}^n([0, 1)) = [0, \infty)$, because $\varsigma_{\alpha \cup \beta}^n([0, 1)) = [0, n + 1)$ holds for all $n \in \mathbb{N}$ by a simple inductive argument:

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

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

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

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

To show that the closure ordinal is $\geq \omega_1^{\text{CK}}$, fix an ordinal $\lambda < \omega_1^{\text{CK}}$, i.e. a recursive ordinal. Let $\prec \subseteq M \times M$ be a corresponding recursive well-order of order type λ on a corresponding set $M \subseteq \mathbb{R}$.¹¹ That is, let f_{\prec} a recursive function such that the relation $x \prec y$ given by $f_{\prec}(x, y) = 0$ defines a well-order on the set $M \stackrel{\text{def}}{=} \{x \in \mathbb{R} : f_{\prec}(x, y) = 0 \text{ or } f_{\prec}(y, x) = 0 \text{ for some } y \in \mathbb{R}\}$. Without loss of generality, assume that $0 \in M$ is the least element of M with respect to \prec . Since \prec is recursive, denote by $?f_{\prec}(x, y) = 0$

¹¹A *well-order* is a linear order \prec on M in which every non-empty subset has a least element. Two sets M, N have equal *order type* iff they have an order-isomorphism $\varphi : M \rightarrow N$, i.e. a monotone bijection with monotone inverse. More background can be found in the literature [Rogers 1987].

the program that does not change the value of variables x, y and that implements the recursive function that terminates if $x \in M$ and either $x \prec y$ or $y \notin M$ and that otherwise fails (like $?(0 = 1)$ would). Consider the dGL formula

$$\langle \underbrace{(y := x; (x' = 1; x' = -1; ?f_{\prec}(x, y) = 0)^d)^*}_{\alpha} \rangle x = 0 \quad (9)$$

By definition of $?f_{\prec}(x, y) = 0$, formula (9) is valid, because x is in M after each successful run of $?f_{\prec}(x, y) = 0$, and \prec is a well-order on M with least element 0. By construction, $\varsigma_{\alpha}(X) = \{a \in \mathbb{R} : b \in X \text{ for all } b \text{ with } f_{\prec}(b, a) = 0\}$ for $X \subseteq \mathbb{R}$. Since \prec has order type λ , $\varsigma_{\alpha}^{\kappa}(\{0\}) \neq \varsigma_{\alpha}^{\lambda}(\{0\}) = M$ for all $\kappa < \lambda$, otherwise the $\varsigma_{\alpha}^{\kappa}(\{0\})$ would induce a monotone injection (even order-isomorphism) from M to $\kappa < \lambda$, which is a contradiction. Indeed, $\varphi : M \rightarrow \kappa; x \mapsto \inf\{\iota : x \in \varsigma_{\alpha}^{\iota}(\{0\})\}$ would otherwise be a monotone injection as $x \prec y$ in M implies $\varphi(x) < \varphi(y)$, because $\varphi(x) \geq \varphi(y)$ implies $y \in \varsigma_{\alpha}(X)$ for a set $X = \varsigma_{\alpha}^{\varphi(y)-1}(\{0\})$ that does not contain x , contradicting $x \prec y$. Note that $\varphi(y)$ is a successor ordinal and hence $\varphi(y) - 1$ defined, since φ maps into successor ordinals and 0 by the definition of φ . Consequently, $\varsigma_{\alpha}^{\lambda}(\{0\}) = M \neq \varsigma_{\alpha}^{\lambda+1}(\{0\}) = \varsigma_{\alpha}(M) = \mathbb{R} = \varsigma_{\alpha^*}(\{0\})$, where $M \neq \mathbb{R}$ because λ is recursive hence countable and \prec a linear order on M . Thus, the closure ordinal for formula (9) is $\lambda + 1 > \lambda$. Hence, for all recursive ordinals λ , there is a hybrid game with a bigger closure ordinal. So, the closure ordinal is $\geq \omega_1^{\text{CK}}$. \square

By Theorem 3.8, the closure ordinal for dGL is between ω_1^{CK} and ordinals of the cardinality of the reals (Theorem 3.5). The same proof works for other well-orderings that are definable in hybrid games, not just those that are definable by classical recursive functions. The proof does not permit arbitrary well-orderings of the real numbers, however, because those may not be definable by hybrid games. Hence, the closure ordinal for dGL is at least ω_1^{HG} , defined as the first ordinal λ that does not have a well-ordering of order type λ that is definable by hybrid games. This ordinal satisfies $\omega_1^{\text{CK}} \leq \omega_1^{\text{HG}}$ and is at most of the cardinality of the reals. A more precise grasp on ω_1^{HG} is in Section 5.

The fact that hybrid games require highly transfinite closure ordinals has a number of consequences. It makes reachability computations and backwards induction difficult, because they only terminate after significantly more than ω -infinitely many steps. It requires higher bounds on the number of repetitions played in hybrid games. It causes classical arguments for relative completeness to fail (Section 4.3). And it causes acute semantical differences that are only visible in hybrid games, not in hybrid systems. For example, the dGL semantics is more general than defining $\varsigma_{\alpha^*}(X)$ to be truncated to ω -repetition $\varsigma_{\alpha}^{\omega}(X) = \bigcup_{n < \omega} \varsigma_{\alpha}^n(X)$, which misses out on the existence of perfectly natural winning strategies. The semantics of dGL is also different than advance notice semantics. For reference, both comparisons are elaborated in Appendix B.

4. AXIOMATIZATION

Section 2 has defined dGL so that every game play has exactly one winner. Section 3 has shown that hybrid games are determined, i.e. from every state, exactly one of the players has a winning strategy for complementary winning conditions. But how can one find out which of the players that is? In principle, one could follow the iterated winning region construction according to the semantics until reaching a fixpoint (Section 3.3), which corresponds to the reachability computation underlying model checking as well as to the backwards induction technique in games. But iterated winning region constructions would not generally terminate in finite time, because the closure ordinal is highly transfinite by Theorem 3.8.

Every dGL sentence without free variables or predicate symbols is either true or false, because dGL is a classical logic. But the semantics of dGL formulas and hybrid

games is ineffective, because computing the semantics, like classical model checking or game solving would, requires transfinite computations for the winning regions. This calls for other ways of proving the validity of dGL formulas.

Simple dGL formulas can be checked by a tableau procedure that expands all choices and detects loops for termination as in the game tree examples (Fig. 2 and Appendix). This principle, however, does not extend to more general hybrid games with differential equations, inherently infinite state spaces, and which need higher ordinals of iteration for computing winning regions by Theorem 3.8.

4.1. Proof Calculus

A Hilbert-type proof calculus for proving validity of dGL formulas is presented in Fig. 5.

$$\begin{array}{l}
[\cdot] \quad [\alpha]\phi \leftrightarrow \neg\langle\alpha\rangle\neg\phi \\
\langle:=\rangle \quad \langle x := \theta \rangle\phi(x) \leftrightarrow \phi(\theta) \\
\langle'\rangle \quad \langle x' = \theta \rangle\phi \leftrightarrow \exists t \geq 0 \langle x := y(t) \rangle\phi \quad (y'(t) = \theta) \\
\langle?\rangle \quad \langle ?\psi \rangle\phi \leftrightarrow (\psi \wedge \phi) \\
\langle\cup\rangle \quad \langle \alpha \cup \beta \rangle\phi \leftrightarrow (\langle\alpha\rangle\phi \vee \langle\beta\rangle\phi) \\
\langle;\rangle \quad \langle \alpha; \beta \rangle\phi \leftrightarrow \langle\alpha\rangle\langle\beta\rangle\phi \\
\langle*\rangle \quad (\phi \vee \langle\alpha\rangle\langle\alpha^*\rangle\phi) \rightarrow \langle\alpha^*\rangle\phi \\
\langle^d\rangle \quad \langle\alpha^d\rangle\phi \leftrightarrow \neg\langle\alpha\rangle\neg\phi \\
\text{M} \quad \frac{\phi \rightarrow \psi}{\langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\psi} \\
\text{FP} \quad \frac{(\phi \vee \langle\alpha\rangle\psi) \rightarrow \psi}{\langle\alpha^*\rangle\phi \rightarrow \psi}
\end{array}$$

Fig. 5. Differential game logic axiomatization

The logic dGL simultaneously generalizes logics of hybrid systems and logics of discrete games and so does its proof calculus. The proof calculus of dGL shares axioms with differential dynamic logic [Platzer 2012a] and discrete game logic [Pauly and Parikh 2003]. It is based on the first-order Hilbert calculus (modus ponens, uniform substitution, and Bernays' \forall -generalization) with all instances of valid formulas of first-order logic as axioms, including (decidable) first-order real arithmetic [Tarski 1951]. Write $\vdash \phi$ iff dGL formula ϕ can be *proved* with the dGL proof rules from dGL axioms (Fig. 5). That is, a dGL formula is inductively defined to be *provable* in the dGL calculus if it is an instance of a dGL axiom or if it is the conclusion (below the rule bar) of an instance of one of the dGL proof rules M, FP, modus ponens, uniform substitution, or \forall -generalization, whose premises (above the rule bar) are all provable.

The *determinacy axiom* $[\cdot]$ describes the duality of winning strategies for complementary winning conditions of Angel and Demon, i.e. that Demon has a winning strategy to achieve ϕ in hybrid game α if and only if Angel does not have a counter strategy, i.e. winning strategy to achieve $\neg\phi$ in the same game α . Axiom $\langle:=\rangle$ is for *assignments by substitution*. Formula $\phi(\theta)$ is obtained from $\phi(x)$ by *substituting* θ for x at all occurrences of x , provided x does not occur in the scope of a quantifier or modality binding

x or a variable of θ . A modality containing $x :=$ or x' outside the scope of tests $?\psi$ or evolution domain constraints *binds* x , because it may change the value of x . In the *differential equation axiom* $\langle' \rangle$, $y(\cdot)$ is the unique [Walter 1998, Theorem 10.VI] solution of the symbolic initial value problem $y'(t) = \theta, y(0) = x$. The duration t how long to follow solution y is for Angel to decide, hence existentially quantified. It goes without saying that variables like t are fresh in Fig. 5.

Axioms $\langle ? \rangle$, $\langle \cup \rangle$, and $\langle ; \rangle$ are as in dynamic logic [Pratt 1976] and differential dynamic logic [Platzer 2012a] except that their meaning is quite different, because they refer to winning strategies of hybrid games instead of reachability relations of systems. The *challenge axiom* $\langle ? \rangle$ expresses that Angel has a winning strategy to achieve ϕ in the test game $? \psi$ exactly from those positions that are already in ϕ (because $? \psi$ does not change the state) and that satisfy ψ for otherwise she would fail the test and lose the game immediately. The *axiom of choice* $\langle \cup \rangle$ expresses that Angel has a winning strategy in a game of choice $\alpha \cup \beta$ to achieve ϕ iff she has a winning strategy in either hybrid game α or in β , because she can choose which one to play. The *sequential game axiom* $\langle ; \rangle$ expresses that Angel has a winning strategy in a sequential game $\alpha; \beta$ to achieve ϕ iff she has a winning strategy in game α to achieve $\langle \beta \rangle \phi$, i.e. to get to a position from which she has a winning strategy in game β to achieve ϕ . The *iteration axiom* $\langle * \rangle$ characterizes $\langle \alpha^* \rangle \phi$ as a pre-fixpoint. It expresses that, if the game is already in a state satisfying ϕ or if Angel has a winning strategy for game α to achieve $\langle \alpha^* \rangle \phi$, i.e. to get to a position from which she has a winning strategy for game α^* to achieve ϕ , then, either way, Angel has a winning strategy to achieve ϕ in game α^* . The converse of $\langle * \rangle$ can be derived¹² and is also denoted by $\langle * \rangle$. The *dual axiom* $\langle ^d \rangle$ characterizes dual games. It says that Angel has a winning strategy to achieve ϕ in dual game α^d iff Angel does not have a winning strategy to achieve $\neg \phi$ in game α . Combining dual game axiom $\langle ^d \rangle$ with the determinacy axiom $\langle \cdot \rangle$ yields $\langle \alpha^d \rangle \phi \leftrightarrow \langle \alpha \rangle \phi$, i.e. that Angel has a winning strategy to achieve ϕ in α^d iff Demon has a winning strategy to achieve ϕ in α . Similar reasoning derives $\langle \alpha^d \rangle \phi \leftrightarrow \langle \alpha \rangle \phi$.

Monotonicity rule M is the generalization rule of monotone modal logic C [Chellas 1980]. It expresses that, if the implication $\phi \rightarrow \psi$ is valid, then, from wherever Angel has a winning strategy in a hybrid game α to achieve ϕ , she also has a winning strategy to achieve ψ , because ψ holds wherever ϕ does. So rule M expresses that easier objectives are easier to win. *Fixpoint rule* FP characterizes $\langle \alpha^* \rangle \phi$ as a *least* pre-fixpoint. It says that, if ψ is another formula that is a pre-fixpoint, i.e. that holds in all states that satisfy ϕ or from which Angel has a winning strategy in game α to achieve that condition ψ , then ψ also holds wherever $\langle \alpha^* \rangle \phi$ 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. 5 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 to not occur in the scope of a quantifier or modality binding x . The *uniform substitution* rule US [Church 1956, §35,40] from first-order logic substitutes *all* occurrences of predicate $p(\cdot)$ by a dGL formula $\psi(\cdot)$, i.e. it replaces all occurrences of $p(\theta)$, for a vectorial term θ , by the corresponding $\psi(\theta)$ simultaneously:

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

In particular, rule US 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

¹² $\phi \vee \langle \alpha \rangle \langle \alpha^* \rangle \phi \rightarrow \langle \alpha^* \rangle \phi$ derives by $\langle * \rangle$. Thus, $\langle \alpha \rangle (\phi \vee \langle \alpha \rangle \langle \alpha^* \rangle \phi) \rightarrow \langle \alpha \rangle \langle \alpha^* \rangle \phi$ by M. Hence, $\phi \vee \langle \alpha \rangle (\phi \vee \langle \alpha \rangle \langle \alpha^* \rangle \phi) \rightarrow \phi \vee \langle \alpha \rangle \langle \alpha^* \rangle \phi$ by propositional congruence. Consequently, $\langle \alpha^* \rangle \phi \rightarrow \phi \vee \langle \alpha \rangle \langle \alpha^* \rangle \phi$ by FP.

variable of $\psi(\theta)$ other than the occurrences in θ . 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 rule US to generate all formula instances from the dGL axioms. Rule US turns axioms into axiom schemes [Church 1956, §35,40], which is a powerful principle that extends to modalities with program constants but is beyond the scope of this article and is pursued in followup work [Platzer 2015b].

Despite their fundamentally different semantics (reachability relations on states of hybrid system runs versus existence of winning strategies into sets of states of interactive hybrid game play) and different dynamical effects (mixed discrete, continuous, and adversarial dynamics), the axiomatization of dGL ends up surprisingly close to that of the logic dL for hybrid systems [Platzer 2012a]. The primary difference of the axiomatization of dGL compared to that of dL 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 monotone modal rule M), absence of the Barcan formula (the converse Barcan formula is still derivable¹³), absence of vacuity V, and absence of the hybrid version of Harel's convergence rule [Harel et al. 1977]. Due to the absence of K, the induction axiom and the convergence axiom are absent in dGL, while corresponding proof rules are still valid; see Section 4.5 for details. The induction rule (ind) is derivable from FP.

A proof of a classical result about the interderivability of FP with the induction rule ind is included for the sake of completeness.

LEMMA 4.1 (INVARIANCE). *Rule FP and the induction rule (ind) of dynamic logic are interderivable in the dGL calculus:*

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

PROOF. Rule ind derives from FP: First derive the following minor variant

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

From $\psi \rightarrow [\alpha]\psi$ and $\psi \rightarrow \phi$ propositionally derive $\psi \rightarrow \phi \wedge [\alpha]\psi$, from which contraposition and propositional logic yield $\neg\phi \vee \neg[\alpha]\psi \rightarrow \neg\psi$. With $[\cdot]$, this gives $\neg\phi \vee \langle\alpha\rangle\neg\psi \rightarrow \neg\psi$. Now FP derives $\langle\alpha^*\rangle\neg\phi \rightarrow \neg\psi$, which, by $[\cdot]$, is $\neg[\alpha^*]\phi \rightarrow \neg\psi$, which gives $\psi \rightarrow [\alpha^*]\phi$

by contraposition. The classical \square -induction rule ind follows by $\phi \stackrel{\text{def}}{=} \psi$. From ind, the variant ind_R is derivable again by M on $\psi \rightarrow \phi$.

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

Hence, the dGL calculus could have been equipped with rule ind instead of FP.

¹³ 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 $\exists x \langle\alpha\rangle\phi \rightarrow \langle\alpha\rangle\exists x \phi$, since converse Barcan assumes that x is not free in the succedent.

Example 4.2. The dual filibuster game formula (5) from Section 3.1 proves easily by going back and forth between players:

$$\begin{array}{l}
\mathbb{R} \frac{*}{x = 0 \rightarrow 0 = 0 \vee 1 = 0} \\
\langle := \rangle \frac{x = 0 \rightarrow \langle x := 0 \rangle x = 0 \vee \langle x := 1 \rangle x = 0}{x = 0 \rightarrow \langle x := 0 \cup x := 1 \rangle x = 0} \\
\langle \cup \rangle \frac{x = 0 \rightarrow \langle x := 0 \cup x := 1 \rangle x = 0}{x = 0 \rightarrow \neg \langle (x := 0 \cup x := 1)^d \rangle \neg x = 0} \\
\langle^d \rangle \frac{x = 0 \rightarrow \neg \langle (x := 0 \cup x := 1)^d \rangle \neg x = 0}{x = 0 \rightarrow \neg \langle x := 0 \cap x := 1 \rangle \neg x = 0} \\
[\cdot] \frac{x = 0 \rightarrow [x := 0 \cap x := 1] x = 0}{x = 0 \rightarrow [(x := 0 \cap x := 1)^*] x = 0} \\
\text{ind} \frac{x = 0 \rightarrow [(x := 0 \cap x := 1)^*] x = 0}{x = 0 \rightarrow \neg \langle (x := 0 \cap x := 1)^* \rangle \neg x = 0} \\
[\cdot] \frac{x = 0 \rightarrow \neg \langle (x := 0 \cap x := 1)^* \rangle \neg x = 0}{x = 0 \rightarrow \langle (x := 0 \cap x := 1)^{*d} \rangle x = 0} \\
\langle^d \rangle \frac{x = 0 \rightarrow \langle (x := 0 \cap x := 1)^{*d} \rangle x = 0}{x = 0 \rightarrow \langle (x := 0 \cup x := 1)^\times \rangle x = 0}
\end{array}$$

The unmarked proof steps expand the definitions for \cap and $^\times$. By pushing dualities through with Remark 3.3, for example, the goal formula (5) at the bottom is equivalent to $x = 0 \rightarrow \langle (x := 0 \cap x := 1)^{*d} \rangle x = 0$, since assignments are unaffected by d .

A proof of a $\langle \alpha^* \rangle$ property will be considered later, because the proof technique for those properties comes from the completeness proof.

4.2. Soundness

Soundness studies whether all provable formulas are valid, which is crucial for ensuring that dGL proofs always produce correct verification results about hybrid games. The soundness proof uses that the following modal congruence rule derives from two uses of the monotonicity rule M:

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

THEOREM 4.3 (SOUNDNESS). *The dGL proof calculus in Fig. 5 is sound, i.e. all provable formulas are valid.*

PROOF. The dGL proof calculus is sound if all instances of axioms and proof rules are sound. Proving soundness of an implication axiom $\phi \rightarrow \psi$ considers an interpretation I with a set of states S and requires showing $\llbracket \phi \rrbracket^I \subseteq \llbracket \psi \rrbracket^I$. Proving soundness of an equivalence axiom $\phi \leftrightarrow \psi$ requires showing $\llbracket \phi \rrbracket^I = \llbracket \psi \rrbracket^I$. Proving soundness of a rule

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

assumes that premise ϕ is valid, i.e. $\llbracket \phi \rrbracket^I = S$ in all interpretations I with a set of states S , and requires showing that conclusion ψ is valid, i.e. $\llbracket \psi \rrbracket^I = S$ in all I with S . All proof rules of dGL except US satisfy the stronger condition of *local soundness*, i.e. for all interpretations I with a set of states S : $\llbracket \phi \rrbracket^I = S$ implies $\llbracket \psi \rrbracket^I = S$. For the proof, recall the μ -calculus notation 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 (from ϕ derive $\forall x \phi$) is standard and not shown. The other axioms and rules are proved to be sound subsequently.

$[\cdot] \llbracket [\alpha] \phi \rrbracket^I = \llbracket \neg \langle \alpha \rangle \neg \phi \rrbracket^I$ is a corollary to determinacy (Theorem 3.1).

$\langle := \rangle$ $\llbracket \langle x := \theta \rangle \phi(x) \rrbracket^I = \varsigma_{x:=\theta}(\llbracket \phi(x) \rrbracket^I) = \{s \in \mathcal{S} : s_x^{\llbracket \theta \rrbracket^I} \in \llbracket \phi(x) \rrbracket^I\} = \{s \in \mathcal{S} : s \in \llbracket \phi(\theta) \rrbracket^I\} = \llbracket \phi(\theta) \rrbracket^I$, 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} [Platzer 2010b, Lemma 2.2] generalizes to d \mathcal{GL} or follows from uniform substitution lemmas [Platzer 2015b].

$\langle \prime \rangle$ $\llbracket \langle x' = \theta \rangle \phi \rrbracket^I = \varsigma_{x'=\theta}(\llbracket \phi \rrbracket^I) = \{\varphi(0) \in \mathcal{S} : \text{for some } \varphi: [0, r] \rightarrow \mathcal{S} \text{ so that } \varphi(r) \in \llbracket \phi \rrbracket^I \text{ and } \frac{d\varphi(t)(x)}{dt}(\zeta) = \llbracket \theta \rrbracket_{\varphi(\zeta)} \text{ for all } \zeta \leq r\}$. Also, $\llbracket \langle \exists t \geq 0 \rangle \langle x := y(t) \rangle \phi \rrbracket^I = \{s \in \mathcal{S} : s_t^I \in \llbracket \langle x := y(t) \rangle \phi \rrbracket^I \text{ for some } r \geq 0\} = \{s \in \mathcal{S} : s_t^I \in \{u \in \mathcal{S} : u_x^{\llbracket y(t) \rrbracket^I} \in \llbracket \phi \rrbracket^I\} \text{ for } r \geq 0\} = \{s \in \mathcal{S} : (s_t^I)_x^{\llbracket y(t) \rrbracket^I} \in \llbracket \phi \rrbracket^I \text{ for some } r \geq 0\}$. The inclusion “ \supseteq ” between both parts

holds, because the function $\varphi(\zeta) := (s_t^I)_x^{\llbracket y(t) \rrbracket^I}$ solves the differential equation $x' = \theta$ by assumption. The inclusion “ \subseteq ” follows, because the solution of the (smooth) differential equation $x' = \theta$ is unique [Platzer 2010b, Lemma 2.1].

$\langle ? \rangle$ $\llbracket \langle ? \psi \rangle \phi \rrbracket^I = \varsigma_{? \psi}(\llbracket \phi \rrbracket^I) = \llbracket \psi \rrbracket^I \cap \llbracket \phi \rrbracket^I = \llbracket \psi \wedge \phi \rrbracket^I$

$\langle \cup \rangle$ $\llbracket \langle \alpha \cup \beta \rangle \phi \rrbracket^I = \varsigma_{\alpha \cup \beta}(\llbracket \phi \rrbracket^I) = \varsigma_{\alpha}(\llbracket \phi \rrbracket^I) \cup \varsigma_{\beta}(\llbracket \phi \rrbracket^I) = \llbracket \langle \alpha \rangle \phi \rrbracket^I \cup \llbracket \langle \beta \rangle \phi \rrbracket^I = \llbracket \langle \alpha \rangle \phi \vee \langle \beta \rangle \phi \rrbracket^I$

$\langle ; \rangle$ $\llbracket \langle \alpha ; \beta \rangle \phi \rrbracket^I = \varsigma_{\alpha; \beta}(\llbracket \phi \rrbracket^I) = \varsigma_{\alpha}(\varsigma_{\beta}(\llbracket \phi \rrbracket^I)) = \varsigma_{\alpha}(\llbracket \langle \beta \rangle \phi \rrbracket^I) = \llbracket \langle \alpha \rangle \langle \beta \rangle \phi \rrbracket^I$.

$\langle * \rangle$ Since $\llbracket \langle \alpha^* \rangle \phi \rrbracket^I = \varsigma_{\alpha^*}(\llbracket \phi \rrbracket^I) = \mu Z.(\llbracket \phi \rrbracket^I \cup \varsigma_{\alpha}(Z))$ is a fixpoint, have $\llbracket \langle \alpha^* \rangle \phi \rrbracket^I = \llbracket \phi \rrbracket^I \cup \varsigma_{\alpha}(\llbracket \langle \alpha^* \rangle \phi \rrbracket^I)$. Thus, $\llbracket \phi \vee \langle \alpha \rangle \langle \alpha^* \rangle \phi \rrbracket^I = \llbracket \phi \rrbracket^I \cup \llbracket \langle \alpha \rangle \langle \alpha^* \rangle \phi \rrbracket^I = \llbracket \phi \rrbracket^I \cup \varsigma_{\alpha}(\llbracket \langle \alpha^* \rangle \phi \rrbracket^I) = \llbracket \langle \alpha^* \rangle \phi \rrbracket^I$. Consequently, $\llbracket \phi \vee \langle \alpha \rangle \langle \alpha^* \rangle \phi \rrbracket^I \subseteq \llbracket \langle \alpha^* \rangle \phi \rrbracket^I$.

$\langle d \rangle$ $\llbracket \langle \alpha^d \rangle \phi \rrbracket^I = \varsigma_{\alpha^d}(\llbracket \phi \rrbracket^I) = \varsigma_{\alpha}(\llbracket \phi \rrbracket^I)^{\mathcal{C}} = \varsigma_{\alpha}(\llbracket \neg \phi \rrbracket^I)^{\mathcal{C}} = \llbracket \langle \alpha \rangle \neg \phi \rrbracket^I = \llbracket \neg \langle \alpha \rangle \neg \phi \rrbracket^I$ by Def. 2.6.

M Assume the premise $\phi \rightarrow \psi$ is valid in interpretation I , i.e. $\llbracket \phi \rrbracket^I \subseteq \llbracket \psi \rrbracket^I$. Then the conclusion $\langle \alpha \rangle \phi \rightarrow \langle \alpha \rangle \psi$ is valid in I , i.e. $\llbracket \langle \alpha \rangle \phi \rrbracket^I = \varsigma_{\alpha}(\llbracket \phi \rrbracket^I) \subseteq \varsigma_{\alpha}(\llbracket \psi \rrbracket^I) = \llbracket \langle \alpha \rangle \psi \rrbracket^I$ by monotonicity (Lemma 2.7).

FP Assume the premise $\phi \vee \langle \alpha \rangle \psi \rightarrow \psi$ is valid in I , i.e. $\llbracket \phi \vee \langle \alpha \rangle \psi \rrbracket^I \subseteq \llbracket \psi \rrbracket^I$. That is, $\llbracket \phi \rrbracket^I \cup \varsigma_{\alpha}(\llbracket \psi \rrbracket^I) = \llbracket \phi \rrbracket^I \cup \llbracket \langle \alpha \rangle \psi \rrbracket^I = \llbracket \phi \vee \langle \alpha \rangle \psi \rrbracket^I \subseteq \llbracket \psi \rrbracket^I$. Thus, ψ is a pre-fixpoint of $Z = \llbracket \phi \rrbracket^I \cup \varsigma_{\alpha}(Z)$. Now using Lemma 2.7, $\llbracket \langle \alpha^* \rangle \phi \rrbracket^I = \varsigma_{\alpha^*}(\llbracket \phi \rrbracket^I) = \mu Z.(\llbracket \phi \rrbracket^I \cup \varsigma_{\alpha}(Z))$ is the least fixpoint and the least pre-fixpoint. Thus, $\llbracket \langle \alpha^* \rangle \phi \rrbracket^I \subseteq \llbracket \psi \rrbracket^I$, which implies that $\langle \alpha^* \rangle \phi \rightarrow \psi$ is valid in I .

US Standard soundness proofs for US [Church 1956] generalize to d \mathcal{GL} . A new proof based on an elegant use of the soundness of RE is shown here. Assume the premise ϕ is valid, i.e. $\llbracket \phi \rrbracket^I = \mathcal{S}$ in all interpretations I with a set of states \mathcal{S} . Assume that the uniform substitution is admissible, otherwise rule US is not applicable and there is nothing to show. It needs to be shown that $\phi_{p(\cdot)}^{\psi(\cdot)}$ is valid, i.e. $\llbracket \phi_{p(\cdot)}^{\psi(\cdot)} \rrbracket^I = \mathcal{S}$ for all I with \mathcal{S} . Consider any particular interpretation J with set of states \mathcal{S} . Without loss of generality, 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 $\phi_{p(\cdot)}^{\psi(\cdot)}$. Therefore, pass to an interpretation I that modifies J by changing the semantics of p such that $\llbracket p(x) \rrbracket^I = \llbracket \psi(x) \rrbracket^J$ for all values of x . In particular, $\llbracket p(x) \rrbracket^I = \llbracket \psi(x) \rrbracket^I$ 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 [Hughes and Cresswell 1996, 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, general-

izes to dGL and is locally sound. Hence, for all particular occurrences of $p(u)$ in ϕ , 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, $\mathcal{S} = \llbracket \phi \rrbracket^I = \llbracket \phi_{p(\cdot)}^{\psi(\cdot)} \rrbracket^I$. Hence, $\llbracket \phi_{p(\cdot)}^{\psi(\cdot)} \rrbracket^J = \mathcal{S}$, 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.

This concludes the soundness proofs for all axioms and proof rules of the dGL proof calculus, which is, thus, sound. \square

The proof calculus in Fig. 5 does not handle differential equations $x' = \theta \ \& \ \psi$ with evolution domain constraints ψ (other than \top). Yet, Lemma 3.4 from Section 3.2 eliminates all evolution domain constraints equivalently from hybrid games, so that evolution domains no longer occur after this equivalence transformation.

4.3. Completeness

The converse of soundness is completeness, which is the question whether all valid formulas are provable. Completeness of dGL is a challenging question related to a famous open problem about completeness of propositional game logic [Parikh 1983]. Based on Gödel's second incompleteness theorem [Gödel 1931], dL is incomplete [Platzer 2008, Theorem 2] and so is dGL. Hence, the right question to ask is that of *relative completeness* [Cook 1978; Harel et al. 1977], i.e. completeness relative to an oracle logic L . Relative completeness studies the question whether a proof calculus has all proof rules that are required for proving all valid formulas in the logic from tautologies in L . Using a notion similar to Leivant's [2009], the question of relative completeness can be separated from that of expressivity. Relative completeness can be shown *schematically* for dGL, i.e. the dGL calculus is complete relative to *any* expressive logic. This is to be contrasted with dL, whose relative completeness proof was dependent on the particular base logic and the specifics of its encoding [Platzer 2008]. In particular, the dGL completeness result is coding-free [Moschovakis 1974], i.e. independent of the particular encoding. It only depends on the ability to express formulas.

Definition 4.4 (Expressive). A logic L is *expressive* (for dGL) if, for each dGL formula ϕ there is a formula ϕ^b of L that is equivalent, i.e. $\models \phi \leftrightarrow \phi^b$. Logic L is *constructively expressive* if, in addition, the mapping $\phi \mapsto \phi^b$ is effective. The logic L is *differentially expressive* for a given proof calculus if L is expressive and all equivalences of the form $\langle x' = \theta \rangle G \leftrightarrow \langle \langle x' = \theta \rangle G \rangle^b$ are provable in that calculus. The logic L is assumed to be closed under the connectives of first-order logic.¹⁴

Differential expressiveness ensures that the expressive logic L is equipped with proof rules for concluding properties of differential equations from their equivalent expressions in L . The differential equation axiom $\langle ' \rangle$ is available for that purpose but limited to expressible solutions. More general ways of concluding properties of differential equations for differential expressiveness include differential invariants and differential cuts [Platzer 2010a], differential ghosts [Platzer 2012d], and the Euler axiom [Platzer 2012a]. Concrete examples of differentially expressive logics are developed in Section 4.4 after proving completeness schematically relative to any arbitrary differentially expressive logic.

¹⁴ Alternatively, the equivalence $\exists x \phi \equiv \langle x' = 1 \rangle \phi \vee \langle x' = -1 \rangle \phi$ can be used to consider quantifiers as abbreviations in differentially expressive logics L .

The classical approach for completeness proofs [Cook 1978; Harel et al. 1977] proceeds in stages of first-order safety assertions, first-order termination assertions, and then the repeated use of those to prove the general case. That approach does not work for dGL, because hybrid games are so highly symmetric that they may contain operators whose proof depends on proofs about all other operators. A proof of $F \rightarrow \langle \alpha \rangle G$, for example, may require proofs of formulas of the form $A \rightarrow [\beta] B$, e.g., when α is β^d . Such an attempt of proving completeness for $\langle \alpha \rangle$ formulas would need to assume completeness for $[\beta]$ formulas and vice versa, which is a cyclic assumption. Even more involved cyclic arguments result from trying to prove completeness of $\langle \alpha^* \rangle$ and $[\alpha^*]$ formulas that way. Furthermore, the previous arguments for completeness of $\langle \alpha^* \rangle$ formulas [Cook 1978; Harel et al. 1977; Platzer 2008] depend on proofs about repetition counts. Those do not work in a hybrid game setting, either, because guaranteed repetition bounds for winning repetition games can be recursively transfinite (Theorem 3.8). Also compare how the semantical discrepancies discussed in Appendix B relate to repetition bounds.

Instead, completeness for all dGL formulas of all kinds can be proved simultaneously, yet with a more involved well-founded partial order on formulas that ensures that the inductive argument in the completeness proof stays well-founded. This generality has beneficial side-effects, though, because the resulting proof architecture enables a result with minimal coding that makes it possible to exactly identify all complex cases.

THEOREM 4.5 (RELATIVE COMPLETENESS). *The dGL calculus is a sound and complete axiomatization of hybrid games relative to any differentially expressive logic L , i.e. every valid dGL formula is provable in the dGL calculus from L tautologies.*

PROOF. Write $\vdash_L \phi$ to indicate that dGL formula ϕ can be derived in the dGL proof calculus from valid L formulas. It takes a moment's thought to conclude that soundness transfers to this case from Theorem 4.3, so it remains to prove completeness. For every valid dGL formula ϕ it has to be proved that ϕ can be derived from valid L tautologies within the dGL calculus: from $\models \phi$ prove $\vdash_L \phi$. The proof proceeds as follows: By propositional recombination, inductively identify fragments of ϕ that correspond to $\phi_1 \rightarrow \langle \alpha \rangle \phi_2$ or $\phi_1 \rightarrow [\alpha] \phi_2$ logically. Find structurally simpler formulas from which these Angel or Demon properties can be derived taking care that the resulting formulas are simpler than the original one in a well-founded order. Finally, prove that the original dGL formula can be re-derived from the subproofs in the dGL calculus.

By appropriate propositional derivations, assume ϕ to be given in conjunctive normal form. 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 first-order 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, with the logic L placed at the bottom of the partial order \prec . The base logic L is considered of lowest complexity by relativity, because $\models F$ immediately implies $\vdash_L F$ for all formulas F of L . Well-foundedness of \prec follows (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 formula in L . The only important property 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. The proof follows the syntactic structure of dGL formulas.

- (0) If ϕ has no hybrid games, then ϕ is a first-order formula; hence provable by assumption (even decidable [Tarski 1951] 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 and quantifier-free, as negations are assumed 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 $\vdash_L \phi_1$ and $\vdash_L \phi_2$ by IH, which combine propositionally to a proof for $\vdash_L \phi_1 \wedge \phi_2$.
- (3) The case where ϕ is of the form $\exists x \phi_2$, $\forall x \phi_2$, $\langle \alpha \rangle \phi_2$ or $[\alpha] \phi_2$ is included in case 4 with $\phi_1 \equiv \perp$.
- (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):

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

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 . When abbreviating the simpler formulas $\neg\phi_1$ by F and ϕ_2 by G , the validity $\vDash \phi$ yields $\vDash \neg F \vee \langle \alpha \rangle G$, so $\vDash F \rightarrow \langle \alpha \rangle G$, from which the remainder of the proof inductively derives

$$\vdash_L F \rightarrow \langle \alpha \rangle G. \quad (10)$$

The proof of (10) is by induction on the syntactic structure of $\langle \alpha \rangle$.

- (a) If $\langle \alpha \rangle$ is the operator $\forall x$ then $\vDash F \rightarrow \forall x G$, where x can be assumed 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 $\vdash_L F \rightarrow G$ is derivable by IH. Then, $\vdash_L F \rightarrow \forall x G$ 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 [Tarski 1951]. The remainder of the proof concludes $(F \rightarrow \psi) \prec (F \rightarrow \phi)$ from $\psi \prec \phi$ without further notice.
- (b) If $\langle \alpha \rangle$ is the operator $\exists x$ then $\vDash F \rightarrow \exists x G$. If F and G are L formulas, then, since L is closed under first-order connectives, so is the valid formula $F \rightarrow \exists x G$, which is, then, provable by IH and even decidable if in first-order real arithmetic [Tarski 1951].

Otherwise, F, G correspond to L formulas by expressiveness of L (Def. 4.4), which implies the existence of an L formula G^b such that $\vDash G^b \leftrightarrow G$. Since L is closed under first-order connectives, the valid formula $F \rightarrow \exists x (G^b)$ is provable by IH, because $(F \rightarrow \exists x (G^b)) \prec (F \rightarrow \exists x G)$ since $G^b \in L$ while $G \notin L$. Now, $\vDash G^b \leftrightarrow G$ implies $\vDash G^b \rightarrow G$, which is derivable by IH, because $(G^b \rightarrow G) \prec \phi$ since G^b is in L . From $\vdash_L G^b \rightarrow G$, the derivable dual of \forall -generalization derives $\vdash_L \exists x (G^b) \rightarrow \exists x G$, which combines with $\vdash_L F \rightarrow \exists x (G^b)$ by modus ponens to $\vdash_L F \rightarrow \exists x G$.¹⁵

- (c) $\vDash F \rightarrow \langle x' = \theta \rangle G$ implies $\vDash F \rightarrow (\langle x' = \theta \rangle G)^b$, which is derivable by IH, because $(F \rightarrow (\langle x' = \theta \rangle G)^b) \prec \phi$ since $(\langle x' = \theta \rangle G)^b$ is in L . Since L is differentially ex-

¹⁵ Expressiveness could also render F and G as F^b, G^b into L in this and other cases of this proof and finally come back to G using rule M instead of \exists -generalization. But the other cases have direct proofs.

pressive, $\vdash_L \langle x' = \theta \rangle G \leftrightarrow (\langle x' = \theta \rangle G)^b$ is provable, so $\vdash_L F \rightarrow \langle x' = \theta \rangle G$ derives from $\vdash_L F \rightarrow (\langle x' = \theta \rangle G)^b$ by modus ponens.

- (d) $\models F \rightarrow [x' = \theta]G$ implies $\models F \rightarrow \neg \langle x' = \theta \rangle \neg G$. Thus, $\models F \rightarrow \neg (\langle x' = \theta \rangle \neg G)^b$, which is derivable by IH, because $(F \rightarrow \neg (\langle x' = \theta \rangle \neg G)^b) \prec \phi$ since $(\langle x' = \theta \rangle \neg G)^b$ is in L . Since L is differentially expressive, $\vdash_L \langle x' = \theta \rangle \neg G \leftrightarrow (\langle x' = \theta \rangle \neg G)^b$ is provable, so $\vdash_L F \rightarrow \neg \langle x' = \theta \rangle \neg G$ derives from $\vdash_L F \rightarrow \neg (\langle x' = \theta \rangle \neg G)^b$ by propositional congruence. Axiom $[\cdot]$, thus, derives $\vdash_L F \rightarrow [x' = \theta]G$.
- (e) $\models F \rightarrow \langle x' = \theta \& \psi \rangle G$, then this formula is, by Lemma 3.4, equivalent to a formula without evolution domain restrictions. Using equation (7) from the proof of Lemma 3.4 as a definitorial abbreviation concludes this case by IH. Similarly for $\models F \rightarrow [x' = \theta \& \psi]G$.
- (f) 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 the 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. This proof focuses on the $\langle \rangle$ cases, because \square cases derive by axiom $[\cdot]$ from the $\langle \rangle$ equivalences.
- (g) $\models F \rightarrow \langle x := \theta \rangle G$ implies $\models F \wedge y = \theta \rightarrow G_x^y$ 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, $\vdash_L F \wedge y = \theta \rightarrow G_x^y$ is derivable by IH. Hence, $\langle := \rangle$ derives $\vdash_L F \wedge y = \theta \rightarrow \langle x := y \rangle G$. Propositional logic derives $\vdash_L F \rightarrow (y = \theta \rightarrow \langle x := y \rangle G)$, from which $\vdash_L F \rightarrow \forall y (y = \theta \rightarrow \langle x := y \rangle G)$ derives by \forall -generalization of first-order logic as y is not in F . Since y was fresh it does not appear in θ and G , so substitution validities of first-order logic derive $\vdash_L F \rightarrow \langle x := \theta \rangle G$. Note that direct proofs of $F \rightarrow \langle x := \theta \rangle G$ by $\langle := \rangle$ are possible when the resulting substitution is admissible, but the substitution forming G_x^y is always admissible, because it is a variable renaming replacing x by the fresh y .
- (h) $\models F \rightarrow \langle ? \psi \rangle G$ implies $\models F \rightarrow \psi \wedge G$. Since $(\psi \wedge G) \prec \langle ? \psi \rangle G$, because it has less modalities, $\vdash_L F \rightarrow \psi \wedge G$ is derivable by IH. Hence, $\langle ? \rangle$ derives $\vdash_L F \rightarrow \langle ? \psi \rangle G$ by propositional congruence, which is used without further notice subsequently.
- (i) $\models F \rightarrow \langle \beta \cup \gamma \rangle G$ implies $\models F \rightarrow \langle \beta \rangle G \vee \langle \gamma \rangle G$. 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 both hybrid games β and γ is smaller than that of $\beta \cup \gamma$ (formula G did not change), $\vdash_L F \rightarrow \langle \beta \rangle G \vee \langle \gamma \rangle G$ is derivable by IH. Hence, $\langle \cup \rangle$ derives $\vdash_L F \rightarrow \langle \beta \cup \gamma \rangle G$.
- (j) $\models F \rightarrow \langle \beta; \gamma \rangle G$, which implies $\models F \rightarrow \langle \beta \rangle \langle \gamma \rangle G$. 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, $\vdash_L F \rightarrow \langle \beta \rangle \langle \gamma \rangle G$ is derivable by IH. Hence, $\vdash_L F \rightarrow \langle \beta; \gamma \rangle G$ derives by $\langle ; \rangle$.
- (k) $\models F \rightarrow [\beta^*]G$ 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 $[\beta^*]G \rightarrow [\beta][\beta^*]G$ is valid, even provable by the variation $[\beta^*]G \rightarrow G \wedge [\beta][\beta^*]G$ of $\langle * \rangle$ that can be obtained from axioms $\langle * \rangle$ and $[\cdot]$. Thus, its equivalent L encoding from Def. 4.4 is also an inductive invariant:

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

Then $F \rightarrow \varphi$ and $\varphi \rightarrow G$ are valid (Angel controls $*$), so derivable by IH, since $(F \rightarrow \varphi) \prec \phi$ and $(\varphi \rightarrow G) \prec \phi$ hold, because φ is in L . By \mathbf{M} , $\langle^d \rangle$ and $[\cdot]$, the latter derivation $\vdash_L \varphi \rightarrow G$ extends to $\vdash_L [\beta^*]\varphi \rightarrow [\beta^*]G$. As above, $\varphi \rightarrow [\beta]\varphi$ is valid, and thus derivable by IH, since β has less loops than β^* . Thus, ind,

which derives from FP by Lemma 4.1, derives $\vdash_L \varphi \rightarrow [\beta^*]\varphi$. The above derivations $\vdash_L F \rightarrow \varphi$, $\vdash_L \varphi \rightarrow [\beta^*]\varphi$, and $\vdash_L [\beta^*]\varphi \rightarrow [\beta^*]G$ combine by modus ponens to $\vdash_L F \rightarrow [\beta^*]G$.

- (l) $\models F \rightarrow \langle \beta^* \rangle G$. Let x the vector of free variables of $\langle \beta^* \rangle G$. Since $\langle \beta^* \rangle G$ is the least pre-fixpoint, for all dGL formulas ψ with free variables in x :

$$\models \forall x (G \vee \langle \beta \rangle \psi \rightarrow \psi) \rightarrow (\langle \beta^* \rangle G \rightarrow \psi)$$

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 :

$$\models \forall x (G \vee \langle \beta \rangle p(x) \rightarrow p(x)) \rightarrow (\langle \beta^* \rangle G \rightarrow p(x))$$

Using $\models F \rightarrow \langle \beta^* \rangle G$, this implies

$$\models \forall x (G \vee \langle \beta \rangle p(x) \rightarrow p(x)) \rightarrow (F \rightarrow p(x))$$

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

$$\vdash_L \forall x (G \vee \langle \beta \rangle p(x) \rightarrow p(x)) \rightarrow (F \rightarrow p(x))$$

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

$$\vdash_L \forall x (G \vee \langle \beta \rangle \langle \beta^* \rangle G \rightarrow \langle \beta^* \rangle G) \rightarrow (F \rightarrow \langle \beta^* \rangle G) \quad (11)$$

Yet, $\langle * \rangle$ derives $\vdash G \vee \langle \beta \rangle \langle \beta^* \rangle G \rightarrow \langle \beta^* \rangle G$, from which $\vdash \forall x (G \vee \langle \beta \rangle \langle \beta^* \rangle G \rightarrow \langle \beta^* \rangle G)$ derives by \forall -generalization. Now modus ponens derives $\vdash_L F \rightarrow \langle \beta^* \rangle G$ by (11).

- (m) $\models F \rightarrow \langle \beta^d \rangle G$ implies $\models F \rightarrow \neg \langle \beta \rangle \neg G$, which implies $\models F \rightarrow [\beta]G$. Since $[\beta]G \prec \langle \beta^d \rangle G$, because β^d is more complex than β even if the modality changed, $\vdash_L F \rightarrow [\beta]G$ can be derived by IH. Axiom $[\cdot]$, thus, derives $\vdash_L F \rightarrow \neg \langle \beta \rangle \neg G$, from which axiom $\langle^d \rangle$ derives $\vdash_L F \rightarrow \langle \beta^d \rangle G$.
- (n) $\models F \rightarrow [\beta^d]G$ implies $\models F \rightarrow \neg \langle \beta^d \rangle \neg G$, hence $\models F \rightarrow \langle \beta \rangle G$. Since $\langle \beta \rangle G \prec [\beta^d]G$, because β^d is more complex than β even if the modality changed, $\vdash_L F \rightarrow \langle \beta \rangle G$ can be derived by IH. Consequently, $\vdash_L F \rightarrow \neg \neg \langle \beta \rangle \neg \neg G$ can be derived using M on $\vdash G \rightarrow \neg \neg G$. Hence, $\langle^d \rangle$ derives $\vdash_L F \rightarrow \neg \langle \beta^d \rangle \neg G$, from which axiom $[\cdot]$ derives $\vdash_L F \rightarrow [\beta^d]G$.

This concludes the derivation of (10), because all operators $\langle \alpha \rangle$ for the form (10) have been considered. From (10), which is $\vdash_L \neg \phi_1 \rightarrow \langle \alpha \rangle \phi_2$ after resolving abbreviations, $\vdash_L \phi_1 \vee \langle \alpha \rangle \phi_2$ derives propositionally.

This completes the proof of completeness (Theorem 4.5), because all syntactical forms of dGL formulas have been covered. \square

The proof of Theorem 4.5 is constructive, so Theorem 4.5 is constructive if L is constructively expressive. The proof is Moschovakis coding-free [Moschovakis 1974]. It even works entirely without coding, except for $x' = \theta$, \exists and $[\beta^*]$. Using US, the case for $\langle \beta^* \rangle G$ in the proof of Theorem 4.5 reveals an explicit β -free reduction to a dGL formula with less loops, which can be considered a modal analogue of characterizations in the Calculus of Constructions [Coquand and Huet 1988]. Using Theorem 4.5, these observations easily reprove a classical result of Meyer and Halpern [Meyer and Halpern 1982] about the semidecidability of termination assertions (logical formulas $F \rightarrow \langle \alpha \rangle G$ of uninterpreted dynamic logic with first-order F, G and regular programs α without differential equations). In fact, this proves a stronger result about semidecidability of dynamic logic without any $[\alpha]$: with loops [Schmitt 1984]. Theorem 4.5 shows that this result continues to hold for uninterpreted game logic in the fragment where $*$ only oc-

curs 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 constructive nature of Theorem 4.5 characterizes exactly which part of hybrid games proving is difficult: finding computationally succinct weaker invariants for $[\beta^*]G$ and finding succinct differential (in)variants [Platzer 2010a] for $[x' = \theta]$ and $\langle x' = \theta \rangle$, of which a solution is a special case [Platzer 2012d]. The case $\exists x G$ is interesting in that a closer inspection of Theorem 4.5 reveals that its complexity depends on whether that quantifier supports Herbrand disjunctions. That is the case for uninterpreted first-order logic and first-order real arithmetic [Tarski 1951], but not for $G \equiv [\beta^*]\psi$, which already gives $\exists x G$ the full Π_1^1 -complete complexity even for classical dynamic logic [Harel et al. 2000, Theorems 13.1,13.2]. Herbrand disjunctions for $\exists x G$ justify how Theorem 4.5 implies the result of Schmitt [Schmitt 1984].

The proof of Theorem 4.5 uses *minimal* coding. The case $[\beta^*]$ needs encoding, because $F \rightarrow [\beta^*]G$ validity is already Π_2^0 -complete for classical dynamic logic [Harel et al. 2000, Theorem 13.5]. The case \exists needs encoding in the presence of $[\beta^*]$, because $\exists x [\beta^*]G$ validity is Π_1^1 -complete for classical dynamic logic [Harel et al. 2000, Theorems 13.1]. The case $x' = \theta$ leads to classical Δ_1^1 -hardness over \mathbb{N} [Platzer 2008, Lemma 4].

The completeness proof indicates a coding-free way of proving Angel properties $\langle \beta^* \rangle G$ that is similar to characterizations in the Calculus of Constructions and works in practice (Appendix A). In particular, dGL does not need Harel's convergence rule [Harel et al. 1977] for completeness and, thus, neither does logic for hybrid systems, even though it was previously based on it [Platzer 2012a]. These results correspond to a hybrid game reading of influential views of understanding program invariants as fixpoints [Cousot and Cousot 1977; Clarke 1979].

4.4. Expressibility

The dGL calculus is complete relative to any differentially expressive logic L (Theorem 4.5). One natural choice for an oracle logic is $L_{\mu D}$, the modal μ -calculus of differential equations (*fixpoint logic 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 the least fixpoint $\mu X. \phi$ 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 [Kozen 1983; Lubarsky 1989]. A more careful inspection of the proofs in this article reveals that the two-variable fragment of $L_{\mu D}$ is enough, which gives a stronger statement as long as the variable hierarchy for $L_{\mu D}$ does not collapse [Berwanger et al. 2007]. The logic $L_{\mu D}$ is considered in this context, because it exposes the most natural interactivity on top of differential equations and makes the constructions most apparent and minimally coding themselves.

LEMMA 4.6 (CONTINUOUS EXPRESSIBILITY). *$L_{\mu D}$ is constructively differentially expressive for dGL.*

PROOF. Of course, $(p(\theta))^b = p(\theta)$ etc. The inductive cases are shown in Fig. 6. It is easy to check that ϕ^b is equivalent to ϕ , e.g. based on the soundness of the dGL axioms. Note that $(\phi \vee \psi)^b \equiv \phi^b \vee \psi^b$ is a consequence of the above definitions and the abbreviation $\phi \vee \psi \equiv \neg(\neg\phi \wedge \neg\psi)$. The quantifier in the definition of $(\langle x := \theta \rangle \phi)^b$ is not necessary if the substitution of θ for x is admissible. The variable renaming of fresh variable y for x in ϕ with the result ϕ_x^y is always admissible. Quantifiers are expressible in $L_{\mu D}$ using Footnote 14. Also $x' = \theta \& \psi$ is expressible by Lemma 3.4. The case $(\langle \alpha^* \rangle \phi)^b$ is defined as the least fixpoint of the reduction of $\phi \vee \langle \alpha \rangle X(x)$, where x are

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

Fig. 6. Inductive cases for constructive expressivity of $L_{\mu D}$.

the variables of α using classical short notation [Lubarsky 1989]. In particular, $(\langle \alpha^* \rangle \phi)^b$ satisfies $\phi \vee \langle \alpha \rangle (\langle \alpha^* \rangle \phi)^b \leftrightarrow (\langle \alpha^* \rangle \phi)^b$ and $(\langle \alpha^* \rangle \phi)^b$ is the formula with the smallest such interpretation, which is all that these proofs depend on. Finally, $L_{\mu D}$ is differentially expressive, because it includes all formulas of the form $\langle x' = \theta \rangle \phi$. \square

A discrete analog of Lemma 4.6 follows from a (constructive) equi-expressibility result [Platzer 2012a, Theorem 9] using the *Euler axiom*, which relates properties of differential equations to properties of their Euler discretizations [Platzer 2012a].

COROLLARY 4.7 (DISCRETE EXPRESSIBILITY). *The (first-order) discrete μ -calculus over \mathbb{R} is constructively differentially expressive for dGL (with the Euler axiom).*

This aligns the discrete and the continuous side of hybrid games in a constructive provably equivalent way similar to corresponding results about hybrid systems [Platzer 2012a]. Yet, the interactivity of two-variable fixpoints stays, which turns out to be necessary (Section 5).

COROLLARY 4.8 (RELATIVE COMPLETENESS). *The dGL calculus is a sound and complete axiomatization of dGL relative to $L_{\mu D}$. With the Euler axiom, the dGL calculus is a sound and complete axiomatization of dGL relative to the discrete μ -calculus over \mathbb{R} .*

PROOF. Follows from Theorem 4.5, Lemma 4.6, and Corollary 4.7. \square

An interesting question is whether fragments of dGL are complete relative to smaller logics, which Theorem 4.5 and Lemma 4.6 reduce solely to a study of expressing (two-variable) $L_{\mu D}$. This yields the following hybrid versions of Parikh's completeness results for fragments of game logic [Parikh 1983].

COROLLARY 4.9 (RELATIVE COMPLETENESS OF *-FREE dGL). *The dGL calculus is a sound and complete axiomatization of *-free hybrid games relative to dL.*

PROOF. Lemma 4.6 reduces to dL, even the first-order logic of differential equations [Platzer 2012a], for *-free hybrid games. \square

COROLLARY 4.10 (RELATIVE COMPLETENESS OF ^d-FREE dGL). *The dGL calculus is a sound and complete axiomatization of ^d-free hybrid games relative to dL.*

PROOF. All d -free loops are Scott-continuous by Lemma 3.7, 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 [Platzer 2012a, Theorem 9]. \square

By Corollary 4.10, $d\mathcal{L}$ is relatively complete without the convergence rule that had been used before [Platzer 2008]. In combination with the first and second relative completeness theorems of $d\mathcal{L}$ [Platzer 2012a], it follows that the $d\mathcal{GL}$ 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 Euler axiom [Platzer 2012a], both are sound and complete axiomatizations of those classes of hybrid games relative to discrete dynamic logic [Platzer 2012a]. Similar completeness results for $d\mathcal{GL}$ relative to $d\mathcal{L}$, and, thus, relative to the first-order logic of differential equations, follow from Theorem 4.5 with some more thought, e.g., for the case of hybrid games with winning regions that are finite rank Borel sets.

4.5. Separating Axioms

In order to illustrate how and why reasoning about hybrid games differs from reasoning about hybrid systems, this section identifies separating axioms, i.e. axioms of $d\mathcal{L}$ [Platzer 2008; Platzer 2012a] that do not hold in $d\mathcal{GL}$. The following result identifies the *axiomatic separation*, i.e. all axioms differing in the respective complete axiomatizations of hybrid systems and hybrid games. It investigates the difference in terms of important classes of modal logics; recall [Hughes and Cresswell 1996] or Appendix B.

THEOREM 4.11 (AXIOMATIC SEPARATION). *The axiomatic separation of hybrid games compared to hybrid systems is exactly the Kripke axiom K , the loop induction axiom I , Harel's loop convergence axiom C , the Barcan axiom B , the vacuous axiom V , and the normal Gödel generalization rule G . Hence, $d\mathcal{GL}$ is a subregular, sub-Barcan, monotone modal logic without the loop induction loop convergence axioms and vacuity.*

The proof of Theorem 4.11 is in Appendix B, where a simple counterexample for each separating axiom illustrates what makes hybrid games different than hybrid systems. The difference in axioms is summarized in Fig. 7, where Cl_\forall is the universal closure with respect to all variables bound in hybrid game α . Besides the axiomatic separation, Fig. 7 shows additional related axioms or proof rules for illustration purposes.

While explicit counterexamples proving the separation in Theorem 4.11 are in Appendix B, the sequel explains the intuition for the difference causing unsoundness in hybrid games of the axioms identified in Theorem 4.11. Kripke's K is unsound (for hybrid games): even if Demon can play RoboCup so that his robots score a goal every time they pass the ball (just never try) and Demon can also play RoboCup so that his robots always pass the ball (somewhere), that does not mean Demon would have a strategy to always score goals. The converse monotonicity axiom \overline{M} is unsound: just because Angel WALL·E has a strategy to be close to EVE or far away does not mean WALL·E would either have a strategy to always end up close to E or a strategy that is always far away. The induction axiom I is unsound: just because Demon has a strategy for his RoboCup robots (e.g. power down) that, no matter how often α^* repeats, Demon still has a strategy such that his robots do not run out of battery for just one more control cycle, that does not mean he has a strategy to keep his robots' batteries nonempty all the time. Harel's convergence rule C is unsound: even if Demon may have a strategy (e.g. waiting) such that after any number of rounds of the game, he has a strategy to move his robots closer to the goal for one control cycle, he still may not have a strategy to ultimately reach the goal, because that requires many rounds of guaranteed progress not just one. The Barcan axiom B is unsound: just because the winner of a RoboCup tournament can be chosen for x after the robot game α does not mean it

\mathbf{K} $[\alpha](\phi \rightarrow \psi) \rightarrow ([\alpha]\phi \rightarrow [\alpha]\psi)$	$\mathbf{M}_{[\cdot]}$ $\frac{\phi \rightarrow \psi}{[\alpha]\phi \rightarrow [\alpha]\psi}$
\mathbf{M} $\langle \alpha \rangle(\phi \vee \psi) \rightarrow \langle \alpha \rangle\phi \vee \langle \alpha \rangle\psi$	\mathbf{M} $\langle \alpha \rangle\phi \vee \langle \alpha \rangle\psi \rightarrow \langle \alpha \rangle(\phi \vee \psi)$
\mathbf{X} $[\alpha^*](\phi \rightarrow [\alpha]\phi) \rightarrow (\phi \rightarrow [\alpha^*]\phi)$	$\forall \mathbf{I}$ $\text{Cl}_\forall (\phi \rightarrow [\alpha]\phi) \rightarrow (\phi \rightarrow [\alpha^*]\phi)$
\mathcal{C} $[\alpha^*]\forall v > 0 (\varphi(v) \rightarrow \langle \alpha \rangle\varphi(v-1)) \rightarrow \forall v (\varphi(v) \rightarrow \langle \alpha^* \rangle \exists v \leq 0 \varphi(v)) \quad (v \notin \alpha)$	
\mathbf{B} $\langle \alpha \rangle \exists x \phi \rightarrow \exists x \langle \alpha \rangle \phi \quad (x \notin \alpha)$	$\overleftarrow{\mathbf{B}}$ $\exists x \langle \alpha \rangle \phi \rightarrow \langle \alpha \rangle \exists x \phi \quad (x \notin \alpha)$
\mathbf{V} $\phi \rightarrow [\alpha]\phi \quad (\mathbf{FV}(\phi) \cap \mathbf{BV}(\alpha) = \emptyset)$	\mathbf{VK} $\phi \rightarrow ([\alpha]\top \rightarrow [\alpha]\phi) \quad (\mathbf{FV}(\phi) \cap \mathbf{BV}(\alpha) = \emptyset)$
\mathcal{G} $\frac{\phi}{[\alpha]\phi}$	$\mathbf{M}_{[\cdot]}$ $\frac{\phi \rightarrow \psi}{[\alpha]\phi \rightarrow [\alpha]\psi}$
\mathbf{R} $\frac{\phi_1 \wedge \phi_2 \rightarrow \psi}{[\alpha]\phi_1 \wedge [\alpha]\phi_2 \rightarrow [\alpha]\psi}$	$\mathbf{M}_{[\cdot]}$ $\frac{\phi_1 \wedge \phi_2 \rightarrow \psi}{[\alpha](\phi_1 \wedge \phi_2) \rightarrow [\alpha]\psi}$
\mathbf{FA} $\langle \alpha^* \rangle \phi \rightarrow \phi \vee \langle \alpha^* \rangle (\neg \phi \wedge \langle \alpha \rangle \phi)$	

Fig. 7. Separating axioms: The axioms and rules on the left are sound for hybrid systems but not for hybrid games. The related axioms and rules on the right are sound for hybrid games.

would be possible to predict this winner x before the game α . By contrast, the converse Barcan axiom $\overleftarrow{\mathbf{B}}$ is sound for hybrid games since, if known before the game α , selecting the winner for x can still be postponed until after the game x . The vacuous axiom \mathbf{V} , in which no free variable of ϕ is bound by α , is unsound: even if ϕ does not change its truth-value during α does not mean it would be possible for Demon to reach any final state at all without being tricked into violating the rules of the game along the way. With an additional assumption ($[\alpha]\top$) that Demon has a winning strategy to reach any final state at all (in which \top , i.e. true, holds which imposes no condition), the possible vacuous axiom \mathbf{VK} is sound. Gödel's rule \mathcal{G} is unsound: even if ϕ holds in all states, Demon may still fail to win $[\alpha]\phi$ if he loses prematurely since Angel tricks Demon into violating the rules during the hybrid game α . Regularity rule \mathbf{R} is unsound: just because Demon's RoboCup robots have a strategy to focus the robots on strong defense and another strategy to, instead, focus them on strong offense that does not mean he would have a strategy to win RoboCup even if simultaneously strong defense and strong offense together might imply victory, because offensive and defensive strategies are in conflict. First arrival \mathbf{FA} is unsound: just because Angel's robot has a strategy to ultimately capture Demon's faster robot with less battery does not mean she would either start with capture or would have a strategy to repeat her control cycle so that she exactly captures Demon's robot during the next control cycle, as Demon might save up his energy and speed up just when Angel predicted to catch him. Having a better battery, Angel will still ultimately win even if Demon sped ahead, but not in the round she thought to be able to predict.

Unlike Harel's convergence axiom, Harel's convergence rule [Harel et al. 1977] is not a separating axiom, because it is sound for dGL , just unnecessary. In light of Theorem 3.8, it is questionable whether the convergence rule would be relatively complete for hybrid games, because it is based on the existence of bounds on the repetition count. The hybrid version of Harel's convergence rule [Platzer 2008] reads as follows

(it assumes that v does not occur in α):

$$\frac{\varphi(v+1) \wedge v+1 > 0 \rightarrow \langle \alpha \rangle \varphi(v)}{\exists v \varphi(v) \rightarrow \langle \alpha^* \rangle \exists v \leq 0 \varphi(v)}$$

If the convergence rule could prove, e.g., dGL formula (8) from Theorem 3.8, then $\varphi(\cdot)$ would yield a bound on the number of repetitions, which, by the proof of Theorem 3.8 does not exist below closure ordinal $\omega \cdot 2$. The premise of the convergence rule makes the bound induced by $\varphi(v)$ progress by 1 in each iteration. The postcondition in the conclusion makes it terminate for $v \leq 0$. And the conclusion's antecedent requires a real number for the initial bound. Thus, the convergence rule only permits bounds below ω , not the required transfinite ordinal $\omega \cdot 2$.

These thoughts further suggest a transfinite version of the convergence rule with an extra inductive premise for limit ordinals. That would be interesting, but is technically more involved than the simple dGL axiomatization, because it would require multi-sorted quantifiers and proof rules for ordinal arithmetic.

5. EXPRESSIVENESS

Differential game logic dGL is a logic for hybrid games whose axiomatic separation to differential dynamic logic dL for hybrid systems [Platzer 2008; Platzer 2012a] has been characterized in Section 4.5. How does dGL compare in expressiveness to differential dynamic logic dL, which is the corresponding logic for hybrid systems? Hybrid systems are expected to be single-player hybrid games where one of the players never gets to decide. And, dL is expected to be a sublogic of dGL. But what about the converse? How the expressiveness of dGL relates to that of dL is related to classical long-standing open questions for the propositional case [Parikh 1985; Berwanger et al. 2007]. Note that even known classical results about expressiveness for the propositional case do not transfer to dGL, because they hinge on finite state [Parikh 1985].

The notation $L_1 \leq L_2$ signifies that logic L_2 is expressive for logic L_1 (Def. 4.4). Likewise, $L_1 \equiv L_2$ signifies equivalent expressiveness, i.e. $L_1 \leq L_2$ and $L_2 \leq L_1$. Further, $L_1 < L_2$ means that L_1 is strictly less expressive than L_2 , i.e. $L_1 \leq L_2$ but not $L_2 \leq L_1$.

LEMMA 5.1 (SINGLE-PLAYER HYBRID GAMES). $dL \leq dGL$ by syntactic embedding.

PROOF. Hybrid systems form single-player hybrid games, i.e. d -free hybrid games. The identity function is a syntactic embedding of dL into dGL, which preserves the semantics as follows. With Lemma 3.7, Kleene's fixpoint theorem implies that ω is the closure ordinal for d -free hybrid games α . Hence, for d -free α , a simple induction shows

$$\varsigma_{\alpha^*}(X) = \varsigma_{\alpha}^{\omega}(X) = \bigcup_{n < \omega} \varsigma_{\alpha}^n(X) = \bigcup_{n < \omega} \varsigma_{\alpha^n}(X) \quad (12)$$

where α^n is the n -fold sequential composition of α given by $\alpha^0 \equiv \top$ and $\alpha^{n+1} \equiv \alpha; \alpha^n$. The semantics of d -free dGL agrees with that defined for dL originally [Platzer 2008; Platzer 2012a] by a simple comparison using (12) for the crucial case α^* . \square

What about the converse? Is the logic dGL truly new or could it have been expressed in dL? Unlike dL, dGL is meant for hybrid games and makes it more convenient to refer directly to questions about hybrid games.¹⁶ Does dGL provide features strictly necessary for hybrid games that dL is missing? Finitely bounded hybrid games are

¹⁶ Even if a logic is not strictly more expressive but “only” more convenient, it is still often strongly preferable. Program logics and their cousins, for example, are used widely, even though first-order integer arithmetic would theoretically suffice [Harel et al. 1977; Harel and Kozen 1984; Harel et al. 2000].

expressible in $d\mathcal{L}$ by Corollary 4.9. What about other hybrid games? Both possible outcomes are interesting. If $d\mathcal{L} \equiv d\mathcal{GL}$, then Theorem 4.5 implies that $d\mathcal{GL}$ is complete relative to $d\mathcal{L}$ and relative to the smaller logics that $d\mathcal{L}$ is complete for [Platzer 2012a]. If $d\mathcal{L} < d\mathcal{GL}$, instead, then $d\mathcal{GL}$ is a provably more expressive logic with features that are strictly necessary for hybrid games. The answer takes some preparations but it also characterizes the general expressiveness of $d\mathcal{L}$ and $d\mathcal{GL}$ as a byproduct.

Let $\langle y_0, \dots, y_n \rangle$ denote a \mathbb{R} -Gödel encoding, i.e. a bijective function pairing $(n+1)$ -tuples of real numbers y_0, y_1, \dots, y_n into a single real, $\langle y_0, \dots, y_n \rangle$, that, along with its inverse, is definable in FOD [Platzer 2008, Lemma 4]. FOD is the *first-order logic of differential equations* [Platzer 2008], i.e. the first-order fragment of $d\mathcal{L}$ and $d\mathcal{GL}$ where all hybrid games α are of the form $x' = \theta$. By Lemma 5.1, FOD is a sublogic of $d\mathcal{GL}$ and, thus, \mathbb{R} -Gödel encodings are definable in $d\mathcal{GL}$. (*Rich-test*) *regular dynamic logic* (DL) [Harel et al. 2000; Platzer 2012a] over \mathbb{R} is the fragment of $d\mathcal{L}$ (and by Lemma 5.1 of $d\mathcal{GL}$) without d and without differential equations. Both FOD [Platzer 2008, Lemma 4] and DL [Platzer 2012a, Theorem 9] can define \mathbb{R} -Gödel encodings. *Acceptable structures* are structures in which elementary \mathbb{R} -Gödel encodings are definable [Moschovakis 1974].

The open *recursive game quantifier* \mathcal{D} of length ω applied to formula $\varphi(x, y)$ is

$$\mathcal{D}y \varphi(x, y) \stackrel{\text{def}}{=} \forall y_0 \exists y_1 \forall y_2 \exists y_3 \dots \bigvee_{n < \omega} \varphi(x, \langle y_0, \dots, y_n \rangle) \quad (13)$$

The semantics of this infinitary formula with its ω many quantifiers and its infinitary disjunction of length ω is defined by a Gale-Stewart [1953] game in which two players alternate in choosing values for the ω many variables y_{2i} (for player \forall) and y_{2i+1} (for player \exists). Player \exists wins if $\varphi(x, \langle y_0, \dots, y_n \rangle)$ holds for some $n < \omega$, i.e. $\varphi(x, y)$ holds when y has the value $\langle y_0, \dots, y_n \rangle$ that is a \mathbb{R} -Gödel encoding of the tuple (y_0, \dots, y_n) . If player \exists has a winning strategy, the infinitary disjunction $\bigvee_{n < \omega} \varphi(x, \langle y_0, \dots, y_n \rangle)$ is satisfied; see [Moschovakis 1974; Väänänen 2011] for details.

LEMMA 5.2 (GAME QUANTIFIER). *Recursive game quantifier \mathcal{D} is definable in $d\mathcal{GL}$.*

PROOF. Let $\varphi(x, y)$ a $d\mathcal{GL}$ formula, which, to simplify notation, is assumed to check the sequence that y encodes only at odd indices n . Then $\mathcal{D}y \varphi(x, y)$ is definable in $d\mathcal{GL}$:

$$\langle y := \langle \rangle; (z' = 1^d; z' = -1^d; y := \langle y, z \rangle; z' = 1; z' = -1; y := \langle y, z \rangle)^* \rangle \varphi(x, y) \quad (14)$$

This $d\mathcal{GL}$ formula uses $\langle y_0, y_1, y_2, \dots, y_n \rangle$ reordered as $\langle \dots \langle \langle \langle \langle \rangle, y_0 \rangle, y_1 \rangle, y_2 \rangle, \dots, y_n \rangle \rangle$ by a recursive permutation starting from the empty tuple encoding $\langle \rangle$ for simplicity reasons. Angel and Demon alternate differential equations for z in (14) that get successively paired into y by the pairing assignment $y := \langle y, z \rangle$, which is definable [Platzer 2008, Lemma 4]. This alternation of differential equations corresponds to the alternation of quantifiers in \mathcal{D} . The number of actual alternations played in (14) can be exactly any arbitrary $n < \omega$, because the semantics of $\langle \alpha^* \rangle$ is a least fixpoint, so well-founded. In each round, Demon first changes z to an arbitrary value by evolving along $z' = 1^d$ and then $z' = -1^d$ for a suitable amount of time. Subsequently, Angel changes z to an arbitrary value, and both values of z are paired into y using the \mathbb{R} -Gödel encoding. \square

Note that (14) equivalently defines (13), even though the (14) is a finite $d\mathcal{GL}$ formula while (13) is an infinite formula in an infinitary logic augmented with the game quantifier [Väänänen 2011], so (14) is infinitely more concise. The closed recursive game quantifier $\neg \mathcal{D}y \neg \phi(x, y)$ is definable in $d\mathcal{GL}$ by duality as well, noting that open as well as closed Gale-Stewart games are determined [Gale and Stewart 1953]. Finally observe that (14) would not define (13) in the (weaker) advance notice semantics (Appendix B),

which corresponds to swapping the quantifier alternation with $\bigvee_{n < \omega}$ to the finitary:

$$\bigvee_{n < \omega} \forall y_0 \exists y_1 \forall y_2 \exists y_3 \dots \varphi(x, (y_0, \dots, y_n))$$

With this preparation, $d\mathcal{GL}$ can be proved to be strictly more expressive than $d\mathcal{L}$, which means that hybrid games are fundamentally more expressive than hybrid systems. In passing, the expressiveness of $d\mathcal{GL}$ and $d\mathcal{L}$ are characterized in terms of inductive and first-order definability, respectively, over acceptable reals.

THEOREM 5.3 (EXPRESSIVE POWER). $d\mathcal{L} < d\mathcal{GL}$.

PROOF. By Lemma 5.1, it only remains to refute $d\mathcal{GL} \leq d\mathcal{L}$. \mathbb{R} -Gödel encodings etc. are elementarily definable in FOD [Platzer 2008, Lemma 4], thus, also in DL over \mathbb{R} [Platzer 2012a, Theorem 9]. This makes \mathbb{R} an acceptable structure [Moschovakis 1974] when augmented with the corresponding definitions from FOD or DL over \mathbb{R} . Further, $d\mathcal{L} \equiv \text{FOD}$ [Platzer 2008] and $d\mathcal{L} \equiv \text{DL over } \mathbb{R}$ [Platzer 2012a, Theorem 9]. On acceptable structures, DL defines exactly all first-order definable relations [Harel and Kozen 1984, Theorems 3 and 4]. On acceptable structures, the open recursive game quantifier $\exists y \varphi(x, y)$ for first-order formulas $\varphi(x, y)$ exactly defines all (positive first-order) inductively definable relations [Moschovakis 1972; Moschovakis 1974, Theorem 5C.2].¹⁷ Game quantifier \exists is definable in $d\mathcal{GL}$ by Lemma 5.2, and so are all inductive relations. In acceptable structures, not all inductively definable relations are first-order definable [Moschovakis 1974, Theorem 5B.2]. Thus, $d\mathcal{GL}$ defines an inductive relation that DL cannot define over \mathbb{R} , so neither can $d\mathcal{L}$. Hence, $d\mathcal{L} \equiv \text{DL} < d\mathcal{GL}$ over \mathbb{R} . \square

Thus, hybrid games can characterize relations that hybrid systems cannot, because $d\mathcal{GL}$ defines all inductive relations over (augmented) \mathbb{R} , while $d\mathcal{L}$ defines exactly all first-order definable relations. The proof of Theorem 3.8 implies that ω_1^{HG} exceeds all order types of all inductive well-orders, because all inductive relations can be characterized in $d\mathcal{GL}$. All closure ordinals of inductive relations occur as order types of some inductive well-order, because the staging order of inductive definitions is well-founded [Moschovakis 1974, Theorems 3A.3, 3C.1]. So ω_1^{HG} equals the closure ordinal of the underlying structure.

The game quantifier and its characterization in the proof of Lemma 5.2 along with the differential equation characterization of Gödel encodings [Platzer 2008, Lemma 4] implies the existence of a smaller syntactic fragment of $d\mathcal{GL}$ that is differentially expressive, so that $d\mathcal{GL}$ is complete relative to this fragment of $d\mathcal{GL}$ by Theorem 4.5. By (13), alternating differential equations in a single loop are the dominant feature of this fragment. The only modification to the proof of Lemma 4.6 is the case of $(\langle \alpha^* \rangle \phi)^b$ which then uses (13) with a (definable) formula $\varphi(x, (y_0, \dots, y_n))$ that simply checks whether the decision sequence y_0, \dots, y_n gives a valid play of hybrid game α^* in which Angel wins. The fact that \exists assumes strict alternation of the players is easily overcome by choosing φ to be independent of y_i when the player for its quantifier does not get to choose at step i in α^* . The actions can be chosen, e.g., as discussed in Appendix A.

6. RELATED WORK

Games and logic have been shown to interact fruitfully in many ways [Gale and Stewart 1953; Ehrenfeucht 1961; Parikh 1983; Parikh 1985; Aumann 1995; Hintikka and Sandu 1997; Stirling 2001; Alur et al. 2002; Pauly and Parikh 2003; Apt and Grädel 2011; Väänänen 2011]. The present article focuses on using logic to specify and verify

¹⁷ The game quantifier in [Moschovakis 1972] starts with $\exists y_1$, which is a difference easily overcome.

properties of hybrid games, inspired by Parikh's propositional game logic for finite-state discrete games [Parikh 1983; Parikh 1985; Pauly and Parikh 2003].

Parikh's game logic generalizes (propositional discrete) dynamic logic to discrete games played on a finite state space, and subsumes Δ PDL and CTL* [Pauly and Parikh 2003]. After more than two decades, its expressiveness has only begun to be understood. 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 [Berwanger 2003] and that, being in the two variable fragment, it is less expressive than the (propositional) modal μ -calculus [Berwanger et al. 2007]. Another influential propositional modal logic, ATL* has been used for model checking finite-state systems [Alur et al. 2002] and is related to propositional game logic [Berwanger and Pinchinat 2009]. Applications and relations of game logic, ATL* [Alur et al. 2002], and strategy logics with explicit strategies [Chatterjee et al. 2010; Mogavero et al. 2012] are discussed in the literature [Alur et al. 2002; Pauly and Parikh 2003; Berwanger and Pinchinat 2009; Chatterjee et al. 2010; Mogavero et al. 2012; Bulling and Jamroga 2014]. Completeness of the ATL fragment of ATL* has been considered [Goranko and van Drimmlen 2006] as well as its expressiveness and complexity [Laroussinie et al. 2008]. But logical investigations of ATL-type logics are scarce according to a recent survey [Bulling and Jamroga 2014] with more detailed comparisons and progress on the satisfiability problem for ATL with imperfect information. These logics for the propositional case of finite-state discrete games 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 dGL shows how such hybrid games can be proved, enjoys compositionality, completeness, and comes with a rich theory.

Differential games have been studied with many different notions of solutions [Isaacs 1967; Friedman 1971a; Petrosjan 1993; Bressan 2010]. They are of interest when actions are solely in continuous time. The present article considers 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 their outcomes take effect in continuous time.

Hybrid games provide a complementary perspective on differential games, just like hybrid systems provide a complementary perspective on continuous dynamical systems. Differential games formalize various notions of adversarial control on variables for a single differential equation [Isaacs 1967; Friedman 1971a; Petrosjan 1993], including solutions based on a non-anticipatory measurable input to an integral interpretation of the differential equations [Friedman 1971a], joint limits for $\delta \rightarrow 0$ of lower and upper limits of δ -anticipatory or δ -delayed strategies [Petrosjan 1993], and Pareto-optimal, Nash, or Stackelberg equilibria, whose computation requires solving PDEs that quickly become ill-posed (already for feedback Nash equilibria except in very special cases); see Bressan [2010] for an overview. Hybrid games, instead, distinguish discrete versus continuous parts of the dynamics, which simplifies the concepts, because easier pieces are involved, and, simultaneously, have been argued to make other aspects like delays in decisions and the integration of computer-decision into continuous physics more realistic [Tomlin et al. 1998; Tomlin et al. 2000; Bouyer et al. 2010; Vladimerou et al. 2011; Prandini et al. 2001; Quesel and Platzer 2012]. The situation is similar to hybrid systems, which provide a complementary perspective on continuous dynamical systems [Nerode and Kohn 1992; Alur et al. 1995; Branicky et al. 1998] that can model more complicated systems as a combination of simpler concepts [Platzer 2012a] and can model computational effects more realistically.

Some 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 [Tomlin et al. 2003; Mitchell et al. 2005], with subsequent extensions by Gao et al. [2007]. Their primary focus is on adversarial choices in the continuous dynamics not on interactions with the discrete dynamics or on interactive game play. Similar observations apply to the viability theory approach to differential games, which gives powerful answers when the differential game and its winning conditions satisfy a number of conditions [Cardaliaguet et al. 2007].

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 [Vladimerou et al. 2011]. STORMED hybrid games generalize o-minimal hybrid games which have been shown to be decidable before [Bouyer et al. 2010]. Timed games [Chatterjee et al. 2011] as well as initialized rectangular hybrid games are known to be decidable [Henzinger et al. 1999], which is limited to the case where all evolution domains and jump constraints are bounded rectangles independent of the previous state of the system and when the controller can only either disable transitions or decide when to take transitions, not both [Henzinger et al. 1999]. Many applications do not fall into these decidable classes [Quesel and Platzer 2012], so that a study of more general hybrid games is called for. The results in this article have implications for such reachability analyses. They show, for example, that reachability computations and backwards induction for hybrid games require highly transfinite closure ordinals $\geq \omega_1^{\text{CK}}$. The completeness proof further exactly characterizes the challenging cases in hybrid games verification.

This article takes a complementary view and studies logics and proofs for hybrid games instead of searching for decidable fragments using bisimulation quotients [Henzinger et al. 1999; Bouyer et al. 2010; Vladimerou et al. 2011], which cannot generally exist. It provides a proof-based and compositional verification technique for more general hybrid games with nonlinear dynamics. This article’s notion of hybrid games is more flexible, because it allows arbitrary nested hybrid game choices rather than one fixed pattern of interaction such as the game of a discrete controller against a continuous plant considered in related work. This results in dGL’s ability to express more general logical formulas with the flexibility expected from a logic and programming language, including arbitrarily nested game operators and nested modalities, which leads to a rich logical theory.

There is more than one way how logic can help to understand games of hybrid systems. Concurrent work has shown that games can also be added as separate constructs on top of unmodified differential dynamic logic [Quesel and Platzer 2012], focusing on the special case of advance notice semantics (Appendix B). The present article follows an entirely different principle. Instead of leaving differential dynamic logic untouched and adding several separate game constructs on top of full hybrid systems reachability modalities [Quesel and Platzer 2012], the logic dGL becomes a proper game logic by adding a single operator d for adversariality into the system dynamics. The logic dGL results in a much simplified but nevertheless more general logic with a simpler and more general semantics and simpler and more general proof calculus. The present article studies a Hilbert calculus and focuses on fundamental logical properties and theory. See [Quesel and Platzer 2012] for practical aspects like a very challenging robotic factory automation case study that translates to dGL. Since dGL is a gentle extension with the single operator d , it is more elegant and significantly easier to implement. What is more difficult in dGL in comparison to that fragment [Quesel and Platzer 2012], however, is the need to carefully identify which axioms are no longer sound for games, which has been pursued in Section 4.5.

The logic dGL presented here has similarities with stochastic differential dynamic logic (SdL) [Platzer 2011], 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. $\text{Sd}\mathcal{L}$ takes a probabilistic perspective on uncertainty in the system dynamics. The dGL 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 quite 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. Adversarial dynamics is also called for in cases of true competition, like in RoboCup.

7. CONCLUSIONS AND FUTURE WORK

This article introduced differential game logic (dGL) for hybrid games that combine discrete, continuous, and adversarial dynamics. Just like hybrid games unify hybrid systems with discrete games, dGL unifies logic of hybrid systems with Parikh's propositional game logic of finite-state discrete games. Hybrid games are challenging, since computing their winning regions may require closure ordinals $\geq \omega_1^{\text{CK}}$. The logic dGL for hybrid games is fundamentally more expressive than the corresponding logic dL for hybrid systems, because it defines all inductive relations over the augmented structure of \mathbb{R} rather than exactly the first-order definables. Nevertheless, dGL has a simple modal semantics and a simple proof calculus, which is proved to be a sound and complete axiomatization of hybrid games relative to any (differentially) expressive logic.

The completeness proof is constructive with minimal coding, thereby exactly characterizing all 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 [Coquand and Huet 1988], and relates to hybrid game versions of influential views of understanding program invariants as fixpoints [Cousot and Cousot 1977; Clarke 1979]. Relative completeness shows that dGL has all axioms and proof rules for dealing with hybrid games and only the base games of differential equations themselves are difficult. The study of (fragments of) dGL which are complete for smaller logics is interesting future work. By the schematic completeness result, this reduces solely to questions of expressiveness, which give rise to interesting questions in descriptive set theory.

It is intriguing to observe the overwhelming impact of the innocent addition of a duality operator. Yet, it is also reassuring to find that logical robustness makes logical foundations continue to work despite the formidable extra challenges of hybrid games. To wit, this article contrasted hybrid games with hybrid systems in terms of their analytic complexity, axiomatizations, and expressiveness.

The dGL axiomatization is strikingly similar to the calculus for stochastic differential dynamic logic SdL [Platzer 2011], despite their fundamentally different semantical presuppositions (adversarial nondeterminism versus stochasticity), which indicates the existence of a deeper logical connection relating stochastic and adversarial uncertainty despite their different mathematical basis (fixpoints, closure ordinals, acceptable structures, inductive definability, and game theory versus stochastic processes, martingales, Markov times, and infinitesimal generators). Because of the axiomatic similarity, the rich theory of dGL may shed light on the logical theory of stochastic hybrid systems, which so far remained elusive.

The logic of hybrid games opens up many directions for future work, including the study of computationally bounded winning strategies, e.g., only strategies that are constructible with small closure ordinals, or with finite rank Borel winning regions, as well as a study of constructive dGL to retain the winning strategies as explicit proof terms. Yet, challenges abound, given the ability of dGL to define closed elementary

games won by a player for whom no hyperelementary quasiwinning strategies exist, which follows from Theorem 5.3 by [Moschovakis 1974, Chapter 7].

Draws, coalitions, rewards, and payoffs different from ± 1 can be expressed easily in dGL using extra variables, but it may be useful to include direct syntactical support. Theorem 5.3 shows that all inductively definable game concepts are expressible in dGL. Rather than including direct support for each, dGL focuses on the most fundamental aspects of hybrid games for reasons of simplicity and elegance. Concurrent games and their equivalent sequential imperfect information games are interesting but are challenging even in the discrete case, because imperfect information leads to Henkin quantifiers. By Theorem 5.3, the challenge is not to add concurrent games but rather to sustain dGL's compositional verification principles. The logic dGL presented here can be augmented with differential games as a new kind of atomic games [Platzer 2015a]. Thanks to its compositional semantics, this results in a modular construction, but is not pursued in this article, because it requires a separate body of mathematics. Combining dGL with axioms for differential equations [Platzer 2010a; Platzer 2012a] already provides a way of handling hybrid games with nonlinear differential equations, differential-algebraic inequalities and differential equations with input.

A. EXAMPLE dGL PROOFS

The completeness proof suggests the use of iteration axiom $\langle^* \rangle$ and US to prove $\langle \alpha^* \rangle$ properties. The following examples illustrate how this works in practice. Observe how logic programming saturation with widening quickly proves the resulting arithmetic.

Example A.1 (Non-game system). The simple non-game dGL formula

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

is provable, shown in Fig. 8, where $\langle \alpha^* \rangle_{0 \leq x < 1}$ is short for $\langle (x := x - 1)^* \rangle (0 \leq x < 1)$. The MP use in Fig. 8 is Hilbert-style, i.e. combines the two lines above by modus ponus.

$$\begin{array}{c}
 \mathbb{R} \frac{}{\forall x (0 \leq x < 1 \vee p(x-1) \rightarrow p(x)) \rightarrow (x \geq 0 \rightarrow p(x))} \\
 \langle := \rangle \frac{}{\forall x (0 \leq x < 1 \vee \langle x := x - 1 \rangle p(x) \rightarrow p(x)) \rightarrow (x \geq 0 \rightarrow p(x))} \\
 \text{US} \frac{}{\forall x (0 \leq x < 1 \vee \langle x := x - 1 \rangle \langle \alpha^* \rangle_{0 \leq x < 1} \rightarrow \langle \alpha^* \rangle_{0 \leq x < 1}) \rightarrow (x \geq 0 \rightarrow \langle \alpha^* \rangle_{0 \leq x < 1})} \\
 \langle^* \rangle, \forall \frac{}{\forall x (0 \leq x < 1 \vee \langle x := x - 1 \rangle \langle \alpha^* \rangle_{0 \leq x < 1} \rightarrow \langle \alpha^* \rangle_{0 \leq x < 1})} \\
 \text{MP} \frac{}{x \geq 0 \rightarrow \langle \alpha^* \rangle_{0 \leq x < 1}}
 \end{array}$$

Fig. 8. dGL Angel proof for Example A.1 using technique from completeness proof

Example A.2 (Choice game). The dGL formula

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

is provable as shown in Fig. 9, 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 $\langle (x := a; a := 0 \cap x := 0)^* \rangle x \neq 1$:

Example A.3 (Hybrid game). The dGL formula

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

is provable as shown in Fig. 10, where the notation $\langle (\beta \cup \gamma)^* \rangle_{0 \leq x < 1}$ is short for $\langle (x := 1; x' = 1^d \cup x := x - 1)^* \rangle (0 \leq x < 1)$: Here and in Fig. 9, the $\forall, \langle^* \rangle, \text{MP}$ steps conclude as in Fig. 8. The proof step $\langle^* \rangle$ uses that $t \mapsto x + t$ is the solution of the differential equation, so the subsequent use of $\langle := \rangle$ substitutes 1 in for x to obtain $t \mapsto 1 + t$. Recall

$$\begin{array}{c}
\mathbb{R} \frac{\quad}{\forall x (x \neq 1 \vee p(a, 0) \wedge p(0, a) \rightarrow p(x, a)) \rightarrow (\top \rightarrow p(x, a))} * \\
\langle ; \rangle, \langle := \rangle \frac{\quad}{\forall x (x \neq 1 \vee \langle \beta \rangle p(x, a) \wedge \langle \gamma \rangle p(x, a) \rightarrow p(x, a)) \rightarrow (\top \rightarrow p(x, a))} \\
\langle ^d \rangle, \langle \cup \rangle \frac{\quad}{\forall x (x \neq 1 \vee \langle \beta \cap \gamma \rangle p(x, a) \rightarrow p(x, a)) \rightarrow (\top \rightarrow p(x, a))} \\
\text{US} \frac{\quad}{\forall x (x \neq 1 \vee \langle \beta \cap \gamma \rangle \langle (\beta \cap \gamma)^* \rangle x \neq 1 \rightarrow \langle (\beta \cap \gamma)^* \rangle x \neq 1) \rightarrow (\top \rightarrow \langle (\beta \cap \gamma)^* \rangle x \neq 1)} \\
\langle * \rangle, \forall, \text{MP} \frac{\quad}{\top \rightarrow \langle (\beta \cap \gamma)^* \rangle x \neq 1} \\
\mathbb{R} \frac{\quad}{x = 1 \wedge a = 1 \rightarrow \langle (\beta \cap \gamma)^* \rangle x \neq 1}
\end{array}$$

Fig. 9. dGL Angel proof for Example A.2 using technique from completeness proof

$$\begin{array}{c}
\mathbb{R} \frac{\quad}{\forall x (0 \leq x < 1 \vee \forall t \geq 0 p(1+t) \vee p(x-1) \rightarrow p(x)) \rightarrow (\top \rightarrow p(x))} * \\
\langle := \rangle \frac{\quad}{\forall x (0 \leq x < 1 \vee \langle x := 1 \rangle \neg \exists t \geq 0 \langle x := x+t \rangle \neg p(x) \vee p(x-1) \rightarrow p(x)) \rightarrow (\top \rightarrow p(x))} \\
\langle ' \rangle \frac{\quad}{\forall x (0 \leq x < 1 \vee \langle x := 1 \rangle \neg \langle x' = 1 \rangle \neg p(x) \vee p(x-1) \rightarrow p(x)) \rightarrow (\top \rightarrow p(x))} \\
\langle ; \rangle, \langle ^d \rangle \frac{\quad}{\forall x (0 \leq x < 1 \vee \langle \beta \rangle p(x) \vee \langle \gamma \rangle p(x) \rightarrow p(x)) \rightarrow (\top \rightarrow p(x))} \\
\langle \cup \rangle \frac{\quad}{\forall x (0 \leq x < 1 \vee \langle \beta \cup \gamma \rangle p(x) \rightarrow p(x)) \rightarrow (\top \rightarrow p(x))} \\
\text{US} \frac{\quad}{\forall x (0 \leq x < 1 \vee \langle \beta \cup \gamma \rangle \langle (\beta \cup \gamma)^* \rangle 0 \leq x < 1 \rightarrow \langle (\beta \cup \gamma)^* \rangle 0 \leq x < 1) \rightarrow (\top \rightarrow \langle (\beta \cup \gamma)^* \rangle 0 \leq x < 1)} \\
\langle * \rangle, \forall, \text{MP} \frac{\quad}{\top \rightarrow \langle (\beta \cup \gamma)^* \rangle 0 \leq x < 1}
\end{array}$$

Fig. 10. dGL Angel proof for Example A.3 using technique from completeness proof

that the winning regions for the formula needs $>\omega$ iterations to converge (Electronic Appendix B.2). It is still provable easily. A variation of this proof shows dGL formula (3) from p. 12, where the handling of the nonlinear differential equation is a bit more complicated.

A variation of Example A.3 also proves dGL formula (8) from the proof of Theorem 3.8, whose closure ordinal is $\omega \cdot 2$.

B. PROOF OF SEPARATING AXIOMS

This section proves Theorem 4.11 with an emphasis on simple counterexamples for each separating axiom to identify the logical essence of the informal explanations shown in Section 4.5.

Subnormal Modal Logic. Unlike dL, dGL is not a normal modal logic [Hughes and Cresswell 1996]. Axiom K, the modal modus ponens from normal modal logic [Hughes and Cresswell 1996], dynamic logic [Pratt 1976], and differential dynamic logic [Platzer 2012a], i.e.

$$\mathbf{K} \quad [\alpha](\phi \rightarrow \psi) \rightarrow ([\alpha]\phi \rightarrow [\alpha]\psi)$$

is not sound for dGL 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. 11. The global version of K, i.e. the implicative version of Gödel's generalization rule is still sound and derives with $\langle ^d \rangle$ and $[\cdot]$ from M using $\alpha \equiv \beta^d$

$$\mathbf{M}_{[\cdot]} \quad \frac{\phi \rightarrow \psi}{[\beta]\phi \rightarrow [\beta]\psi}$$




Fig. 11. Game trees for counterexample to axiom K using $\alpha \equiv (x := 1 \cap x := 0); y := 0$.

The normal Gödel generalization rule G, i.e.

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

however, is not sound for dGL as witnessed by the choice $\alpha \equiv (? \perp)^d, \phi \equiv \top$.

The vacuous axiom V, which expresses that formulas do not change their truth-value along a program if their free variables are not bound, is sound for dynamic logics and differential dynamic logic dL when no free variable of ϕ is bound by α [Platzer 2012a]:

$$\mathcal{V} \phi \rightarrow [\alpha]\phi \quad (\text{FV}(\phi) \cap \text{BV}(\alpha) = \emptyset)$$

but the vacuous axiom is not sound for dGL as witnessed by the choice $\phi \equiv x = 0$ and $\alpha \equiv y := 0; (?y = 1)^d$. With an additional assumption $[\alpha]\top$ expressing that the game can be played to a final state at all, the possible vacuous axiom VK is sound for dGL:

$$(\text{VK}) \phi \rightarrow ([\alpha]\top \rightarrow [\alpha]\phi) \quad (\text{FV}(\phi) \cap \text{BV}(\alpha) = \emptyset)$$

If Demon can always finish the game ($[\alpha]\top$) then ϕ will continue to hold if it was true initially since α only changes bound variables and ϕ only depends on its free variables.

The closest counterpart to G that is sound for hybrid games is $M_{[\cdot]}$ and the closest counterpart to V that is sound is VK. Both require the extra assumption $[\alpha]\top$, which is only trivial for hybrid systems not for hybrid games.

Subregular Modal Logic. Regular modal logics are monotone modal logics [Chellas 1980] that are weaker than normal modal logics. But the regular modal generalization rule [Chellas 1980], i.e.

$$\mathcal{R} \frac{\phi_1 \wedge \phi_2 \rightarrow \psi}{[\alpha]\phi_1 \wedge [\alpha]\phi_2 \rightarrow [\alpha]\psi}$$

is not sound for dGL either as witnessed by the choice $\alpha \equiv (x := 1 \cap x := 0); y := 0$, $\phi_1 \equiv x = 1, \phi_2 \equiv x = y, \psi \equiv x = 1 \wedge x = y$; see Fig. 12.




Fig. 12. Game trees for counterexample to regular modal rule using $\alpha \equiv (x := 1 \cap x := 0); y := 0$.

Monotone Modal Logic. The axiom that is closest to K but still sound for dGL is a monotonicity axiom. This axiom is sound for dGL, yet already included in the monotonicity rule M:

LEMMA B.1 ([CHELLAS 1980, THEOREM 8.13]). *In the presence of rule RE from p. 25, rule M is interderivable with axiom M:*

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

PROOF. Axiom M derives from rule M: From $\phi \rightarrow \phi \vee \psi$, M derives $\langle \alpha \rangle \phi \rightarrow \langle \alpha \rangle (\phi \vee \psi)$. From $\psi \rightarrow \phi \vee \psi$, M derives $\langle \alpha \rangle \psi \rightarrow \langle \alpha \rangle (\phi \vee \psi)$, from which propositional logic yields $\langle \alpha \rangle \phi \vee \langle \alpha \rangle \psi \rightarrow \langle \alpha \rangle (\phi \vee \psi)$.

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

The converse of axiom M is sound for dL but not for dGL, however, as witnessed by $\alpha \equiv x := 1 \cap x := 0$, $\phi \equiv x = 1$, $\psi \equiv x = 0$; see Fig. 13:

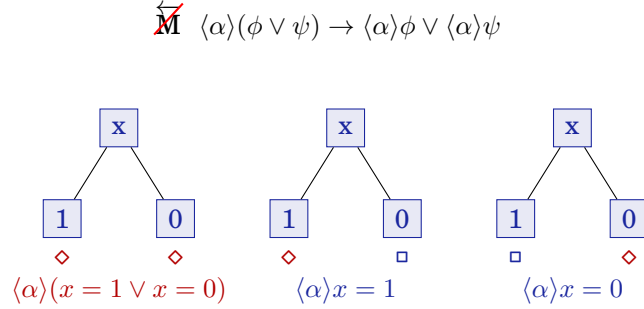


Fig. 13. Game trees for counterexample to converse monotone axiom using $\alpha \equiv x := 1 \cap x := 0$.

The presence of the regular congruence rule RE and the fact that $[\alpha]\phi \leftrightarrow \neg \langle \alpha \rangle \neg \phi$ by determinacy (Theorem 3.1) still make dGL a classical modal logic [Chellas 1980]. Rule M even makes dGL a monotone modal logic [Chellas 1980].

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 [Barcan 1946], which, together, characterize constant domain in normal first-order modal logics [Hughes and Cresswell 1996]. The Barcan axiom B, which characterizes anti-monotone domains in first-order modal logic [Hughes and Cresswell 1996], is sound for constant-domain first-order dynamic logic and for differential dynamic logic dL when x does not occur in α [Platzer 2012a]:

$$B \langle \alpha \rangle \exists x \phi \rightarrow \exists x \langle \alpha \rangle \phi \quad (x \notin \alpha)$$

but the Barcan axiom is not sound for dGL 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

$$B \forall x [\alpha]\phi \rightarrow [\alpha]\forall x \phi \quad (x \notin \alpha)$$

is not sound for dGL 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 monotone domains [Hughes and Cresswell 1996], is sound for dGL and can be derived when x does not occur in α (see Footnote 13 on p. 24):

$$\overline{B} \exists x \langle \alpha \rangle \phi \rightarrow \langle \alpha \rangle \exists x \phi \quad (x \notin \alpha)$$

No Induction Axiom. The induction axiom

$$\not\vdash [\alpha^*](\phi \rightarrow [\alpha]\phi) \rightarrow (\phi \rightarrow [\alpha^*]\phi) \quad (15)$$

holds for $d\mathcal{L}$, but, unlike the induction rule ind , does not hold for $d\mathcal{GL}$ as witnessed by $\alpha \equiv ((x := a; a := 0) \cap x := 0)$ and $\phi \equiv (x = 1)$; see Fig. 14. The failure of the induction




Fig. 14. Game trees for counterexample to induction axiom (notation: x, a) with game $\alpha \equiv (x := a; a := 0) \cap x := 0$. **(left)** $[\alpha^*](x = 1 \rightarrow [\alpha]x = 1)$ is true by the strategy “if Angel chose stop, choose $x := a; a := 0$, otherwise always choose $x := 0$ ” **(right)** $[\alpha^*]x = 1$ is false by the strategy “repeat once and repeat once more if $x = 1$, then stop.” If a winning state can be reached by a winning strategy, the mark is enclosed in a circle \odot or \ominus , respectively.

axiom in the counterexample for (15) hinges on the fact that Angel is free to decide whether to repeat α after each round depending on the state. This would be different for an advance notice semantics for α^* ; see Appendix B. By a variation of the soundness argument for FP or the semantic deduction theorem applied to the ind rule, it can be shown, however, that a variation of the induction axiom is still sound if the induction rule ind is translated into an axiom using the universal closure, denoted Cl_\forall , with respect to all variables bound in α :

$$(\forall I) \text{Cl}_\forall (\phi \rightarrow [\alpha]\phi) \rightarrow (\phi \rightarrow [\alpha^*]\phi)$$

Universal closures do not rescue the first arrival axiom, a dual of induction axiom I:

$$\not\vdash \langle \alpha^* \rangle \phi \rightarrow \phi \vee \langle \alpha^* \rangle (\neg \phi \wedge \langle \alpha \rangle \phi)$$

This axiom holds for $d\mathcal{L}$. It 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 [Pauly and Parikh 2003]. This axiom does not hold, however, for dGL 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. 15. Observe how the failure of the first arrival axiom in dGL relates to the impossibility of predicting precise enough repetition counts in hybrid games (recall corresponding discussions for Theorem 3.8, Section 4.3, and Appendix B).

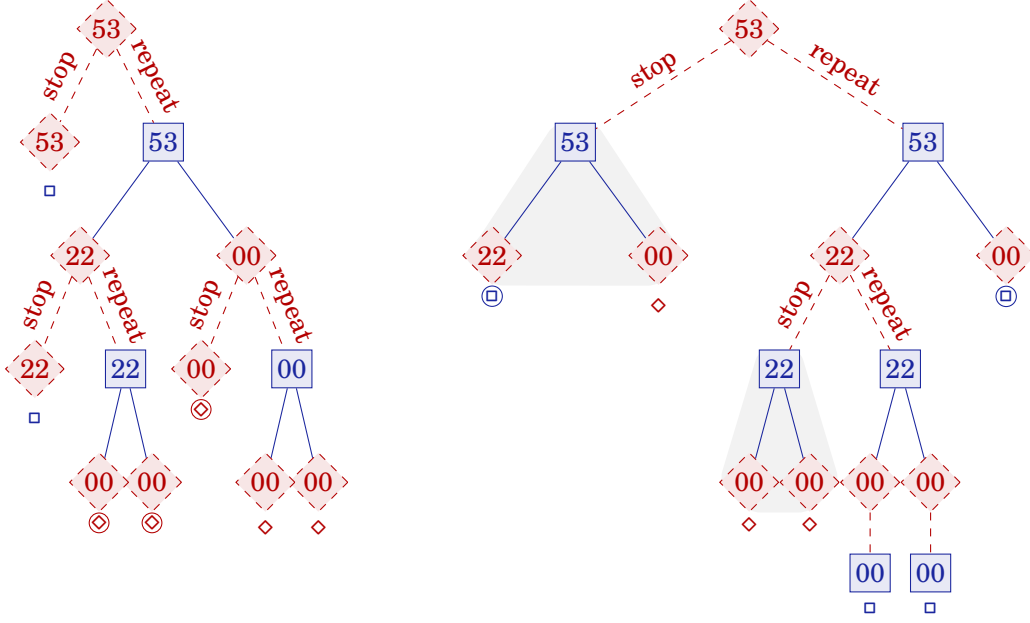


Fig. 15. Game trees for counterexample to first arrival axiom with game $\alpha \equiv (x := x - y \cap x := 0); y := x$ (notation: x, y). **(left)** $\langle \alpha^* \rangle x = 0$ is true no matter what Demon chooses **(right)** $\langle \alpha^* \rangle (x \neq 0 \wedge \langle \alpha \rangle x = 0)$ is false, because stop can be defeated by $x := x - y$ and repeat can be defeated by $x := 0$.

The hybrid systems axiom version [Platzer 2012a] of Harel's convergence rule [Harel et al. 1977], in which v does not occur in α (written $v \notin \alpha$),

$$\not\vdash [\alpha^*] \forall v > 0 (\varphi(v) \rightarrow \langle \alpha \rangle \varphi(v - 1)) \rightarrow \forall v (\varphi(v) \rightarrow \langle \alpha^* \rangle \exists v \leq 0 \varphi(v)) \quad (v \notin \alpha)$$

holds for dL, but not for dGL as witnessed by $\alpha \equiv (x := x - y; y := 0 \cap x := x - 1)$ and $\varphi(v) \equiv (x \leq v)$. In a state where $y = 1, x \geq 2$, $[\alpha^*] \forall v > 0 (\varphi(v) \rightarrow \langle \alpha \rangle \varphi(v - 1))$ is true by the strategy "always choose $x := x - 1$ " for Demon for $[\alpha^*]$ and arbitrary strategies for Angel for the nested $\langle \alpha \rangle$. Yet, $\langle \alpha^* \rangle \exists v \leq 0 \varphi(v)$ is false by the strategy "always choose $x := x - y; y := 0$ ", because x will no longer change after the first iteration then. The hybrid version of Harel's convergence rule is sound but unnecessary (Section 4.5).

This completes the proof of Theorem 4.11 by inspecting the complete axiomatization of dGL from Theorem 4.5 compared to the complete axiomatization of hybrid systems [Platzer 2012a]. Each axiom of hybrid systems has been considered and either continues to hold for hybrid games (Theorem 4.3) or has been refuted with a counterexample (K, I, C, B, V, G) or continues to hold but is unnecessary for completeness (Harels' convergence rule). A few additional axioms and rules that are not part of the hybrid systems axiomatization have been considered for illustration purposes, because they are closely related and highlight interesting aspects of the axiomatic similarity (axiom M, $\forall I$, \overleftarrow{B} , VK, $M_{[1]}$) or difference (\overleftarrow{M} , R, FA) between hybrid systems and hybrid games.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library.

ACKNOWLEDGMENTS

I thank Stephen Brookes, Frank Pfenning, James Cummings, and Erik Zawadzki for helpful discussions and I appreciate the helpful comments by the reviewers.

REFERENCES

- Peter Aczel. 1977. An Introduction to Inductive Definitions. In *Handbook of Mathematical Logic*, Jon Barwise (Ed.). Studies in Logic and the Foundations of Mathematics, Vol. 90. Elsevier, 739 – 782. DOI: [http://dx.doi.org/10.1016/S0049-237X\(08\)71120-0](http://dx.doi.org/10.1016/S0049-237X(08)71120-0)
- Rajeev Alur, Costas Courcoubetis, Nicolas Halbwachs, Thomas A. Henzinger, Pei-Hsin Ho, Xavier Nicollin, Alfredo Olivero, Joseph Sifakis, and Sergio Yovine. 1995. The Algorithmic Analysis of Hybrid Systems. *Theor. Comput. Sci.* 138, 1 (1995), 3–34.
- Rajeev Alur, Costas Courcoubetis, Thomas A. Henzinger, and Pei-Hsin Ho. 1992. Hybrid Automata: An Algorithmic Approach to the Specification and Verification of Hybrid Systems.. In *Hybrid Systems*. 209–229.
- Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *J. ACM* 49, 5 (2002), 672–713. DOI: <http://dx.doi.org/10.1145/585265.585270>
- Krzysztof R. Apt and Erich Grädel. 2011. *Lectures in Game Theory for Computer Scientists*. Cambridge Univ. Press.
- Robert Aumann. 1995. Backward Induction and Common Knowledge of Rationality. *Games and Economic Behavior* 8, 1 (1995), 263–300.
- Ruth C. Barcan. 1946. A Functional Calculus of First Order Based on Strict Implication. *J. Symb. Log.* 11, 1 (1946), 1–16.
- Dietmar Berwanger. 2003. Game Logic is Strong Enough for Parity Games. *Studia Logica* 75, 2 (2003), 205–219. DOI: <http://dx.doi.org/10.1023/A:1027358927272>
- Dietmar Berwanger, Erich Grädel, and Giacomo Lenzi. 2007. The Variable Hierarchy of the μ -calculus is strict. *Theory Comput. Syst.* 40, 4 (2007), 437–466.
- Dietmar Berwanger and Sophie Pinchinat. 2009. Game Quantification Patterns. In *ICLA (LNCS)*, Ramaswamy Ramanujam and Sundar Sarukkai (Eds.), Vol. 5378. Springer, 116–130.
- Patricia Bouyer, Thomas Brihaye, and Fabrice Chevalier. 2010. O-Minimal Hybrid Reachability Games. *Log. Meth. Comput. Sci.* 6, 1 (2010).
- Michael S. Branicky, Vivek S. Borkar, and Sanjoy K. Mitter. 1998. A Unified Framework for Hybrid Control: Model and Optimal Control Theory. *IEEE T. Automat. Contr.* 43, 1 (1998), 31–45.
- Alberto Bressan. 2010. Noncooperative Differential Games. A Tutorial. (2010).
- Nils Bulling and Wojciech Jamroga. 2014. Comparing variants of strategic ability: how uncertainty and memory influence general properties of games. *Autonomous Agents and Multi-Agent Systems* 28, 3 (2014), 474–518. DOI: <http://dx.doi.org/10.1007/s10458-013-9231-3>
- Pierre Cardaliaguet, Marc Quincampoix, and Patrick Saint-Pierre. 2007. Differential Games Through Viability Theory: Old and Recent Results. In *Advances in Dynamic Game Theory*, Steffen Jørgensen, Marc Quincampoix, and Thomas L. Vincent (Eds.). Ann. Int. Soc. Dyn. Game., Vol. 9. Birkhäuser, 3–35. DOI: http://dx.doi.org/10.1007/978-0-8176-4553-3_1
- Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. 2010. Strategy logic. *Inf. Comput.* 208, 6 (2010), 677–693. DOI: <http://dx.doi.org/10.1016/j.ic.2009.07.004>
- Krishnendu Chatterjee, Thomas A. Henzinger, and Vinayak S. Prabhu. 2011. Timed Parity Games: Complexity and Robustness. *Logical Methods in Computer Science* 7, 4 (2011). DOI: [http://dx.doi.org/10.2168/LMCS-7\(4:8\)2011](http://dx.doi.org/10.2168/LMCS-7(4:8)2011)
- Brian F. Chellas. 1980. *Modal Logic: An Introduction*. Cambridge Univ. Press.
- Alonzo Church. 1956. *Introduction to Mathematical Logic, Volume I*. Princeton University Press, Princeton, NJ.
- Edmund M. Clarke. 1979. Program Invariants as Fixedpoints. *Computing* 21, 4 (1979), 273–294.
- Stephen A. Cook. 1978. Soundness and Completeness of an Axiom System for Program Verification. *SIAM J. Comput.* 7, 1 (1978), 70–90.
- Thierry Coquand and Gérard P. Huet. 1988. The Calculus of Constructions. *Inf. Comput.* 76, 2/3 (1988), 95–120. DOI: [http://dx.doi.org/10.1016/0890-5401\(88\)90005-3](http://dx.doi.org/10.1016/0890-5401(88)90005-3)

- Patrick Cousot and Radhia Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *POPL*. 238–252.
- Jennifer M. Davoren and Anil Nerode. 2000. Logics for Hybrid Systems. *IEEE* 88, 7 (July 2000), 985–1010.
- Akash Deshpande, Aleks Göllü, and Pravin Varaiya. 1996. SHIFT: A Formalism and a Programming Language for Dynamic Networks of Hybrid Automata. In *Hybrid Systems*. 113–133.
- S. Dharmatti and M. Ramaswamy. 2006. Zero-Sum Differential Games Involving Hybrid Controls. *J. Optimiz. Theory App.* 128, 1 (2006), 75–102. DOI: <http://dx.doi.org/10.1007/s10957-005-7558-x>
- Laurent Doyen, Goran Frehse, George J. Pappas, and André Platzer. 2016. Verification of Hybrid Systems. In *Handbook of Model Checking*, Edmund M. Clarke, Thomas A. Henzinger, and Helmut Veith (Eds.). Springer, Chapter 28.
- Andrzej Ehrenfeucht. 1961. An application of games to the completeness problem for formalized theories. *Fundam. Math.* 49 (1961), 129–141.
- Avner Friedman. 1971a. *Differential Games*. John Wiley.
- Harvey Friedman. 1971b. Higher set theory and mathematical practice. *Ann. Math. Log.* 2, 3 (1971), 325–357. DOI: [http://dx.doi.org/10.1016/0003-4843\(71\)90018-0](http://dx.doi.org/10.1016/0003-4843(71)90018-0)
- David Gale and F. M. Stewart. 1953. Infinite games with perfect information. In *Contributions to the Theory of Games II (Ann. of Math. Stud.)*, H. W. Kuhn and A. W. Tucker (Eds.), Vol. 28. Princeton Univ. Press, 246–266.
- Y. Gao, J. Lygeros, and M. Quincampoix. 2007. On the Reachability Problem for Uncertain Hybrid Systems. *IEEE T. Automat. Contr.* 52, 9 (Sept. 2007), 1572–1586. DOI: <http://dx.doi.org/10.1109/TAC.2007.904449>
- Kurt Gödel. 1931. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Mon. hefte Math. Phys.* 38 (1931), 173–198.
- Valentin Goranko and Govert van Drimmelen. 2006. Complete axiomatization and decidability of Alternating-time temporal logic. *Theor. Comput. Sci.* 353, 1-3 (2006), 93–117. DOI: <http://dx.doi.org/10.1016/j.tcs.2005.07.043>
- David Harel and Dexter Kozen. 1984. A Programming Language for the Inductive Sets, and Applications. *Information and Control* 63, 1/2 (1984), 118–139. DOI: [http://dx.doi.org/10.1016/S0019-9958\(84\)80045-5](http://dx.doi.org/10.1016/S0019-9958(84)80045-5)
- David Harel, Dexter Kozen, and Jerzy Tiuryn. 2000. *Dynamic Logic*. MIT Press.
- David Harel, Albert R. Meyer, and Vaughan R. Pratt. 1977. Computability and Completeness in Logics of Programs (Preliminary Report). In *STOC*. ACM, 261–268.
- Thomas A. Henzinger, Benjamin Horowitz, and Rupak Majumdar. 1999. Rectangular Hybrid Games. In *CONCUR (LNCS)*, Jos C. M. Baeten and Sjouke Mauw (Eds.), Vol. 1664. Springer, 320–335. DOI: http://dx.doi.org/10.1007/3-540-48320-9_23
- Thomas A. Henzinger, Peter W. Kopke, Anuj Puri, and Pravin Varaiya. 1995. What’s decidable about hybrid automata?. In *STOC*, Frank Thomson Leighton and Allan Borodin (Eds.). ACM, 373–382. DOI: <http://dx.doi.org/10.1145/225058.225162>
- Jaakko Hintikka and Gabriel Sandu. 1997. Game-Theoretical Semantics. In *Handbook of Logic and Language*, Johan van Benthem and Alice ter Meulen (Eds.). Elsevier, 361–410.
- George E. Hughes and Max J. Cresswell. 1968. *A New Introduction to Modal Logic*. Routledge.
- Rufus Philip Isaacs. 1967. *Differential Games*. John Wiley.
- Alexander S. Kechris. 1994. *Classical Descriptive Set Theory*. Grad. Texts in Math., Vol. 156. Springer.
- Dexter Kozen. 1983. Results on the Propositional μ -Calculus. *Theor. Comput. Sci.* 27 (1983), 333–354.
- François Laroussinie, Nicolas Markey, and Ghassan Oreiby. 2008. On the Expressiveness and Complexity of ATL. *Logical Methods in Computer Science* 4, 2 (2008).
- Daniel Leivant. 2009. On the Completeness of Dynamic Logic. In *FOSSACS (LNCS)*, Luca de Alfaro (Ed.), Vol. 5504. Springer, 78–91.
- LICS 2012. *Proceedings of the 27th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25–28, 2012*. IEEE.
- Robert S. Lubarsky. 1989. μ -Definable Sets of Integers. In *LICS*. IEEE Computer Society, 343–352. DOI: <http://dx.doi.org/10.1109/LICS.1989.39189>
- Donald A. Martin. 1975. Borel Determinacy. *Annals of Mathematics* 102, 2 (1975), 363–371. DOI: <http://dx.doi.org/10.2307/1971035>
- Albert R. Meyer and Joseph Y. Halpern. 1982. Axiomatic Definitions of Programming Languages: A Theoretical Assessment. *J. ACM* 29, 2 (1982), 555–576. DOI: <http://dx.doi.org/10.1145/322307.322324>
- Ian Mitchell, Alexandre M. Bayen, and Claire Tomlin. 2005. A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games. *IEEE T. Automat. Contr.* 50, 7 (2005), 947–957.

- Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y. Vardi. 2012. What Makes ATL* Decidable? A Decidable Fragment of Strategy Logic. In *CONCUR (LNCS)*, Maciej Koutny and Irek Ulidowski (Eds.), Vol. 7454. Springer, 193–208. DOI: http://dx.doi.org/10.1007/978-3-642-32940-1_15
- Yiannis N. Moschovakis. 1972. The Game Quantifier. *Proc. Amer. Math. Soc.* 31, 1 (1972), 245–250.
- Yannis N. Moschovakis. 1974. *Elementary Induction on Abstract Structures*. Vol. 77. Elsevier.
- Anil Nerode and Wolf Kohn. 1992. Models for Hybrid Systems: Automata, Topologies, Controllability, Observability. In *Hybrid Systems*. 317–356.
- Anil Nerode, Jeffrey B. Rummel, and Alexander Yakhnis. 1996. Hybrid System Games: Extraction of Control Automata with Small Topologies. In *Hybrid Systems*. 248–293. DOI: <http://dx.doi.org/10.1007/BFb0031565>
- Rohit Parikh. 1983. Propositional Game Logic. In *FOCS*. IEEE, 195–200. DOI: <http://dx.doi.org/10.1109/SFCS.1983.47>
- Rohit Parikh. 1985. The Logic of Games and Its Applications. *Annals of Discrete Mathematics* 24 (1985), 111–140. M. Karpinski and J. van Leeuwen, eds., Topics in the Theory of Computation.
- Marc Pauly and Rohit Parikh. 2003. Game Logic - An Overview. *Studia Logica* 75, 2 (2003), 165–182. DOI: <http://dx.doi.org/10.1023/A:1027354826364>
- Leon A. Petrosjan. 1993. *Differential Games of Pursuit*. World Scientific.
- André Platzer. 2008. Differential Dynamic Logic for Hybrid Systems. *J. Autom. Reas.* 41, 2 (2008), 143–189. DOI: <http://dx.doi.org/10.1007/s10817-008-9103-8>
- André Platzer. 2010a. Differential-Algebraic Dynamic Logic for Differential-Algebraic Programs. *J. Log. Comput.* 20, 1 (2010), 309–352. DOI: <http://dx.doi.org/10.1093/logcom/exn070>
- André Platzer. 2010b. *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer, Heidelberg. DOI: <http://dx.doi.org/10.1007/978-3-642-14509-4>
- André Platzer. 2011. Stochastic Differential Dynamic Logic for Stochastic Hybrid Programs. In *CADE (LNCS)*, Nikolaj Bjørner and Viorica Sofronie-Stokkermans (Eds.), Vol. 6803. Springer, 431–445. DOI: http://dx.doi.org/10.1007/978-3-642-22438-6_34
- André Platzer. 2012a. The Complete Proof Theory of Hybrid Systems, See LICS [2012], 541–550. DOI: <http://dx.doi.org/10.1109/LICS.2012.64>
- André Platzer. 2012b. *Differential Game Logic for Hybrid Games*. Technical Report CMU-CS-12-105. School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.
- André Platzer. 2012c. Logics of Dynamical Systems, See LICS [2012], 13–24. DOI: <http://dx.doi.org/10.1109/LICS.2012.13>
- André Platzer. 2012d. The Structure of Differential Invariants and Differential Cut Elimination. *Log. Meth. Comput. Sci.* 8, 4 (2012), 1–38. DOI: [http://dx.doi.org/10.2168/LMCS-8\(4:16\)2012](http://dx.doi.org/10.2168/LMCS-8(4:16)2012)
- André Platzer. 2013. *A Complete Axiomatization of Differential Game Logic for Hybrid Games*. Technical Report CMU-CS-13-100R. School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.
- André Platzer. 2015a. Differential Hybrid Games. *CoRR* abs/ (2015).
- André Platzer. 2015b. A Uniform Substitution Calculus for Differential Dynamic Logic. In *CADE (LNCS)*, Amy Felty and Aart Middeldorp (Eds.), Vol. 9195. Springer, 467–481. DOI: http://dx.doi.org/10.1007/978-3-319-21401-6_32
- Maria Prandini, João P. Hespanha, and George J. Pappas. 2001. Greedy Control for Hybrid Pursuit Games. In *ECC*. 2621–2626.
- Vaughan R. Pratt. 1976. Semantical Considerations on Floyd-Hoare Logic. In *FOCS*. IEEE, 109–121.
- Jan-David Quesel and André Platzer. 2012. Playing Hybrid Games with KeYmaera. In *IJCAR (LNCS)*, Bernhard Gramlich, Dale Miller, and Uli Sattler (Eds.), Vol. 7364. Springer, 439–453. DOI: http://dx.doi.org/10.1007/978-3-642-31365-3_34
- Hartley Rogers. 1987. *Theory of Recursive Functions and Effective Computability*. MIT Press.
- Peter H. Schmitt. 1984. Diamond Formulas: A Fragment of Dynamic Logic with Recursively Enumerable Validity Problem. *Information and Control* 61, 2 (1984), 147–158. DOI: [http://dx.doi.org/10.1016/S0019-9958\(84\)80055-8](http://dx.doi.org/10.1016/S0019-9958(84)80055-8)
- Colin Stirling. 2001. *Modal and Temporal Properties of Processes*. Springer.
- Alfred Tarski. 1951. *A Decision Method for Elementary Algebra and Geometry* (2nd ed.). University of California Press, Berkeley.
- Claire Tomlin, Ian Mitchell, Alexandre M. Bayen, and Meeko Oishi. 2003. Computational techniques for the verification of hybrid systems. *Proc. IEEE* 91, 7 (2003), 986–1001.
- Claire Tomlin, George J. Pappas, and Shankar Sastry. 1998. Conflict resolution for air traffic management: a study in multi-agent hybrid systems. *IEEE T. Automat. Contr.* 43, 4 (1998), 509–521.

- Claire J. Tomlin, John Lygeros, and Shankar Sastry. 2000. A Game Theoretic Approach to Controller Design for Hybrid Systems. *Proc. IEEE* 88, 7 (2000), 949–970.
- Shinya Umeno and Nancy A. Lynch. 2007. Safety Verification of an Aircraft Landing Protocol: A Refinement Approach. In *HSCC (LNCS)*, Alberto Bemporad, Antonio Bicchi, and Giorgio C. Buttazzo (Eds.), Vol. 4416. Springer, 557–572. DOI : http://dx.doi.org/10.1007/978-3-540-71493-4_43
- Jouko Väänänen. 2011. *Models and Games*. Cambridge Univ. Press.
- Vladimeros Vladimerou, Pavithra Prabhakar, Mahesh Viswanathan, and Geir E. Dullerud. 2011. Specifications for decidable hybrid games. *Theor. Comput. Sci.* 412, 48 (2011), 6770–6785. DOI : <http://dx.doi.org/10.1016/j.tcs.2011.08.036>
- John von Neumann and Oskar Morgenstern. 1955. *Theory of Games and Economic Behavior* (3rd ed.). Princeton Univ. Press.
- Wolfgang Walter. 1998. *Ordinary Differential Equations*. Springer.
- Ernst Friedrich Fernando Zermelo. 1913. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. In *Fifth International Congress of Mathematicians*, Vol. II. 501–504.

Received August 2014; revised February 2015; accepted July 2015

Online Appendix to: Differential Game Logic

ANDRÉ PLATZER, Carnegie Mellon University

For reference and to support the interactive intuition of game play, an operational semantics of hybrid games is shown in Appendix A. Alternative semantics for repetitions of hybrid games are contrasted in Appendix B. Finally, Appendix C provides concrete hybrid games to support the computational intuition for the higher closure ordinals proved generically in Section 3.3.

A. OPERATIONAL GAME SEMANTICS

In order to relate the intuition of interactive game play to the denotational semantics of hybrid games, this section shows an operational semantics for hybrid games that is more complicated than the modal semantics from Section 2.2 but makes strategies explicit and directly reflects the intuition how hybrid games are played interactively. 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. 2 and relates to standard notions in game theory and descriptive set theory. Theorem A.2 below proves 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 choose:

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

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.

The operational game semantics uses standard notions from descriptive set theory [Kechris 1994]. 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 $(s_1, \dots, s_n, t_1, \dots, t_m)$ if $s = (s_1, \dots, s_n)$ and $t = (t_1, \dots, t_m)$. For an $a \in A$, write $a \hat{t}$ for $(a) \hat{t}$ and write $t \hat{a}$ for $t \hat{(a)}$. For a set $S \subseteq A^{(\mathbb{N})}$, write $S \hat{t}$ for $\{s \hat{t} : s \in S\}$ and $t \hat{S}$ for $\{t \hat{s} : s \in S\}$. The state $\llbracket t \rrbracket_s$ 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) $\llbracket x := \theta \rrbracket_s = s_x^{\llbracket \theta \rrbracket_s}$
- (2) $\llbracket x' = \theta \ \& \ \psi @ r \rrbracket_s = \varphi(r)$ for the unique $\varphi : [0, r] \rightarrow \mathcal{S}$ differentiable, $\varphi(0) = s$, $\frac{d\varphi(t)(x)}{dt}(\zeta) = \llbracket \theta \rrbracket_{\varphi(\zeta)}$ and $\varphi(\zeta) \in \llbracket \psi \rrbracket^I$ for all $\zeta \leq r$. Note that $\llbracket x' = \theta \ \& \ \psi @ r \rrbracket_s$ is not defined if no such φ of duration r exists.
- (3) $\llbracket ?\psi \rrbracket_s = \begin{cases} s & \text{if } s \in \llbracket \psi \rrbracket^I \\ \text{not defined} & \text{otherwise} \end{cases}$

- (4) $\lfloor l \rfloor_s = \lfloor \tau \rfloor_s = \lfloor \mathfrak{s} \rfloor_s = \lfloor \mathfrak{g} \rfloor_s = \lfloor \mathfrak{d} \rfloor_s = \lfloor () \rfloor_s = s$
 (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 $a \in A$. Denote by $\text{leaf}(T)$ the set of all leaves of T , i.e. nodes $t \in T$ that have no successor in T .

Definition A.1 (Operational game semantics). The operational game semantics of hybrid game α is, for each state s of each interpretation I , a tree $\mathfrak{g}(\alpha)(s) \subseteq A^{(\mathbb{N})}$ defined as follows (see Fig. 16 for a schematic illustration):

- (1) $\mathfrak{g}(x := \theta)(s) = \{(x := \theta)\}$
 (2) $\mathfrak{g}(x' = \theta \ \& \ \psi)(s) = \{(x' = \theta \ \& \ \psi @ r) : r \in \mathbb{R}, r \geq 0, \varphi(0) = s \text{ for some (differentiable)}$
 $\varphi : [0, r] \rightarrow \mathcal{S} \text{ such that } \frac{d\varphi(t)(x)}{dt}(\zeta) = \llbracket \theta \rrbracket_{\varphi(\zeta)} \text{ and } \varphi(\zeta) \in \llbracket \psi \rrbracket^I \text{ for all } \zeta \leq r\}$
 (3) $\mathfrak{g}(\ ? \ \psi)(s) = \{(\ ? \ \psi)\}$
 (4) $\mathfrak{g}(\alpha \cup \beta)(s) = \mathfrak{r} \hat{\ } \mathfrak{g}(\alpha)(s) \cup \mathfrak{r} \hat{\ } \mathfrak{g}(\beta)(s)$
 (5) $\mathfrak{g}(\alpha; \beta)(s) = \mathfrak{g}(\alpha)(s) \cup \bigcup_{t \in \text{leaf}(\mathfrak{g}(\alpha)(s))} \mathfrak{g}(\beta)(\lfloor t \rfloor_s)$
 (6) $\mathfrak{g}(\alpha^*)(s) = \bigcup_{n < \omega} f^n(\{(s), (\mathfrak{g})\})$
 where f^n is the n -fold composition of the function
 $f(Z) \stackrel{\text{def}}{=} Z \cup \bigcup_{t \in \text{leaf}(Z)} t \hat{\ } \mathfrak{g} \hat{\ } \mathfrak{g}(\alpha)(\lfloor t \hat{\ } \mathfrak{g} \rfloor_s) \hat{\ } \{(s), (\mathfrak{g})\}$
 (7) $\mathfrak{g}(\alpha^d)(s) = \mathfrak{d} \hat{\ } \mathfrak{g}(\alpha)(s) \hat{\ } \mathfrak{d}$

Note the implicit closure under prefixes in the definition of $\mathfrak{g}(\alpha)(s)$ for readability. For example, $\mathfrak{g}(\alpha^d)(s) = \mathfrak{d} \hat{\ } \mathfrak{g}(\alpha)(s) \hat{\ } \mathfrak{d}$ means $\mathfrak{g}(\alpha^d)(s) = \{(), (\mathfrak{d})\} \cup \mathfrak{d} \hat{\ } \mathfrak{g}(\alpha)(s) \cup \mathfrak{d} \hat{\ } \mathfrak{g}(\alpha)(s) \hat{\ } \mathfrak{d}$.

Angel gets to choose which action to take at node $t \in \mathfrak{g}(\alpha)(s)$ if t has an even number of occurrences of \mathfrak{d} , otherwise Demon gets to choose. In the former case *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 $t \in \mathfrak{g}(\alpha)(s)$ is *deadlocked at t* if $t \notin \text{leaf}(\mathfrak{g}(\alpha)(s))$ 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 $t \in \mathfrak{g}(\alpha^*)(s)$ cannot choose \mathfrak{g} infinitely often for that loop because $n < \omega$.

The players use Markov strategies, i.e. their choices only depend on the current state of the system and they have no additional information about the strategy of the other player. A *strategy for Angel* from initial state s is a nonempty subtree $\sigma \subseteq \mathfrak{g}(\alpha)(s)$ that accepts all of Demon's actions at nodes t where Demon acts and selects a unique Angel action when Angel acts at t :

- (1) for all $t \in \sigma$ at which Demon acts, $t \hat{\ } a \in \sigma$ for all $a \in A$ such that $t \hat{\ } a \in \mathfrak{g}(\alpha)(s)$.
 (2) for all $t \in \sigma$ at which Angel acts, if $t \notin \text{leaf}(\mathfrak{g}(\alpha)(s))$, 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 $(a_1, \dots, a_n) \in A^{(\mathbb{N})}$ of

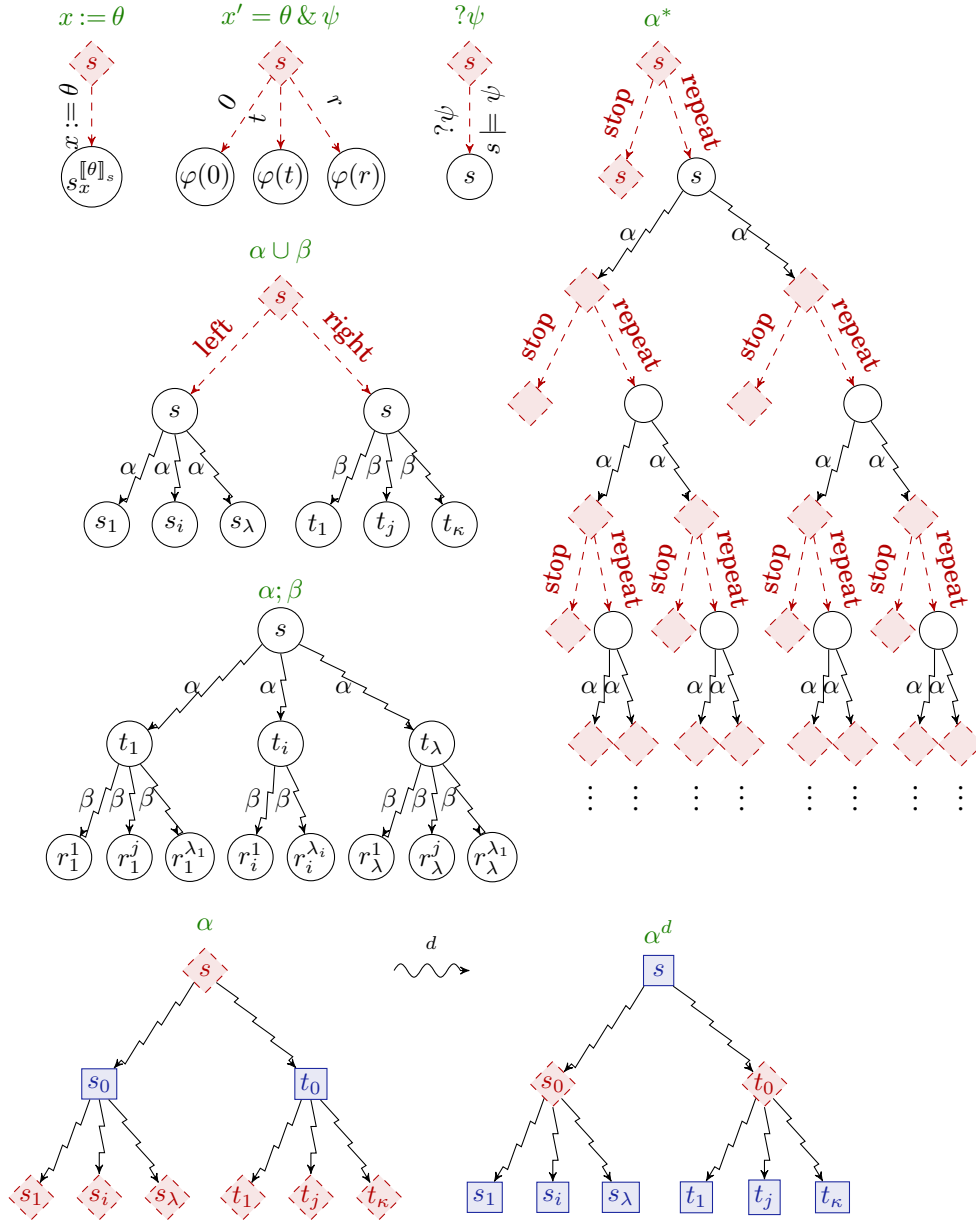


Fig. 16. Operational game semantics for hybrid games of dGL

maximal length such that

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

By definition of strategies, $\sigma \oplus \tau$ is unique. A *winning strategy for Angel* for winning condition $X \subseteq S$ from state s in interpretation I is a strategy $\sigma \subseteq g(\alpha)(s)$ for Angel

from s such that, for all strategies $\tau \subseteq g(\alpha)(s)$ for Demon from s : Demon deadlocks or $[\sigma \oplus \tau]_s \in X$. A *winning strategy for Demon* for (Demon's) winning condition $X \subseteq S$ from state s in interpretation I is a strategy $\tau \subseteq g(\alpha)(s)$ for Demon from s such that, for all strategies $\sigma \subseteq g(\alpha)(s)$ for Angel from s : Angel deadlocks or $[\sigma \oplus \tau]_s \in X$.

The denotational modal semantics (Section 2.2) is equivalent to the operational semantics:

THEOREM A.2 (EQUIVALENT SEMANTICS). *The modal semantics of dGL is equivalent to the operational game-tree semantics of dGL, i.e. for each hybrid game α , each initial state s in each interpretation I , and each winning condition $X \subseteq S$:*

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

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

PROOF. Proceed by simultaneous induction on the structure of α and prove equivalence. As part of the equivalence proof, 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. 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 $g(\alpha)(s)$. Theorem 3.1 implies $\delta_\alpha(X^\complement) = \varsigma_\alpha(X)^\complement$.

- (1) $s \in \varsigma_{x:=\theta}(X) \iff s_x^{[\theta]} \in X \iff [\sigma \oplus \tau]_s = [x := \theta]_s = s_x^{[\theta]} \in X$, using $\sigma \stackrel{\text{def}}{=} \{(x := \theta)\} = g(x := \theta)(s)$. The converse direction follows, because the strategy σ follows the only permitted strategy.
- (2) $s \in \varsigma_{x'=\theta \& \psi}(X) \iff s = \varphi(0), \varphi(r) \in X$ for some $r \in \mathbb{R}$ and some (differentiable) $\varphi : [0, r] \rightarrow S$ such that $\frac{d\varphi(t)(x)}{dt}(\zeta) = \llbracket \theta \rrbracket_{\varphi(\zeta)}$ and $\varphi(\zeta) \in \llbracket \psi \rrbracket^I$ for all $\zeta \leq r$
 $\iff [\sigma \oplus \tau]_s = [x' = \theta \& \psi @ r]_s = \varphi(r) \in X$, using $\sigma \stackrel{\text{def}}{=} \{(x' = \theta \& \psi @ r)\} \subseteq g(x' = \theta \& \psi)(s)$. The converse direction follows, since this σ has the only permitted form for a strategy and different values of r that lead to X are equally useful.
- (3) $s \in \varsigma_{?\psi}(X) = \llbracket \psi \rrbracket^I \cap X \iff [\sigma \oplus \tau]_s = [?\psi]_s = s \in X$, with $s \in \llbracket \psi \rrbracket^I$ using $\sigma \stackrel{\text{def}}{=} \{(\psi)\} = g(?\psi)(s)$. The converse direction uses that this σ is the only permitted strategy and it deadlocks exactly if $s \notin \llbracket \psi \rrbracket^I$.
- (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 $\sigma_\alpha \subseteq g(\alpha)(s)$ for Angel for X from s or there is a winning strategy $\sigma_\beta \subseteq g(\beta)(s)$ for Angel for X from s . This is equivalent to $\sigma \subseteq g(\alpha \cup \beta)(s)$ being a winning strategy for Angel for X from s , using $\sigma \stackrel{\text{def}}{=} \hat{\tau} \sigma_\alpha$ or $\sigma \stackrel{\text{def}}{=} \hat{\tau} \sigma_\beta$, respectively.
- (5) $s \in \varsigma_{\alpha; \beta}(X) = \varsigma_\alpha(\varsigma_\beta(X))$. By induction hypothesis, this is equivalent to the existence of a strategy $\sigma_\alpha \subseteq g(\alpha)(s)$ for Angel such that for all strategies $\tau \subseteq g(\alpha)(s)$ for Demon: $[\sigma_\alpha \oplus \tau]_s \in \varsigma_\beta(X)$. By induction hypothesis, $[\sigma_\alpha \oplus \tau]_s \in \varsigma_\beta(X)$ is equivalent to the existence of a winning strategy σ_τ for Angel (which depends on the state $[\sigma_\alpha \oplus \tau]_s$ that the previous α game led to) with winning condition X from $[\sigma_\alpha \oplus \tau]_s$. This is equivalent to $\sigma \subseteq g(\alpha; \beta)(s)$ being a winning strategy for Angel for X from s , using

$$\sigma \stackrel{\text{def}}{=} \sigma_\alpha \cup \bigcup (\sigma_\alpha \oplus \tau) \hat{\tau} \sigma_\tau \quad (16)$$

The union is over all leaves $\sigma_\alpha \oplus \tau \in \text{leaf}(g(\alpha)(s))$ 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 $[\sigma_\alpha \oplus \tau]_s$ that (when α terminates) was reached by the actions $\sigma_\alpha \oplus \tau$ according to the strategy τ that Demon was observed to have played during α . The converse direction uses that strategies do not depend on moves that have not been played yet and that all strategies can be factorized by prefixes of what has actually been played to be coerced into the form (16).

- (6) Both inclusions of the case α^* are proved separately. If W denotes the set of states from which Angel has a winning strategy in $g(\alpha^*)(s)$ to achieve X , then need to show that $\varsigma_{\alpha^*}(X) = W$. For $\varsigma_{\alpha^*}(X) \subseteq W$, it is enough to show that W is a prefixpoint, i.e. $X \cup \varsigma_\alpha(W) \subseteq W$, because $\varsigma_{\alpha^*}(X)$ is the least (pre-)fixpoint. Consider a $s \in X \cup \varsigma_\alpha(W)$. If $s \in X$ then $s \in W$ with the winning strategy $\sigma \stackrel{\text{def}}{=} \{(s)\}$ for Angel to achieve X in α^* from s . Otherwise, $s \in \varsigma_\alpha(W)$ implies, by induction hypothesis, that there is a winning strategy $\sigma_\alpha \subseteq g(\alpha)(s)$ for Angel in α to achieve W from s . By definition of W , Angel has a winning strategy in $g(\alpha^*)(s)$ to achieve X from all states reached after playing α from s according to σ_α , i.e. $[\sigma_\alpha \oplus \tau]_s \in W$ for all strategies τ of Demon. As in the case for $\alpha; \beta$, 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 , thus leads to a winning strategy of the form

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

for Angel to achieve X in α^* from s , where the union is over all leaves $\sigma_\alpha \oplus \tau \in \text{leaf}(g(\alpha)(s))$ in all strategies 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, recall $\varsigma_{\alpha^*}(X)^\complement = \delta_{\alpha^*}(X^\complement) = \bigcup \{Z \subseteq \mathcal{S} : Z \subseteq X^\complement \cap \delta_\alpha(Z)\}$ by Theorem 3.1. 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, obtain a winning strategy τ for Demon to achieve Z in α^* for any arbitrary 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 3.1. 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 $\sigma \subseteq g(\alpha)(s)$ for Angel winning X^\complement in α from s . Since $\varsigma_{\alpha^d}(X) = \delta_\alpha(X)$ by Theorem 3.1, this is equivalent to: there is a winning strategy $\tau \subseteq g(\alpha)(s)$ 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 $\sigma \subseteq g(\alpha^d)(s)$ for Angel winning X in α^d from s using $\sigma \stackrel{\text{def}}{=} \partial \hat{\tau} \hat{\partial}$. The converse direction uses that all strategies permitted for α^d begin and end with ∂ . \square

B. ALTERNATIVE SEMANTICS

To elaborate why the dGL semantics is both natural and general, this section briefly considers alternative choices for the semantics, focusing on the role of repetition in the context of hybrid games. It turns out that alternative semantics require prior bounds of repetitions of $<\omega$ (Appendix B.1) and ω (Appendix B.2), respectively.

followed by one choice of $x := a; a := 0$ in the last repetition. This strategy would not work in the dGL semantics, because Angel is free to decide whether to repeat α^* after each repetition based on the resulting state of the game.

Conversely, the dual formula would be valid in the advance notice semantics but is not valid in dGL:

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

The dGL semantics is more general, because it gives the player in charge of repetition more control as the state can be inspected before deciding on whether to repeat again. Advance notice semantics, instead, only allows the choice of a fixed number of repetitions. The advance notice games can be expressed easily in dGL by having the players choose a counter c before the loop that decreases to 0 during the repetition. The advance notice semantics can be expressed in dGL, e.g., for (17) as

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

B.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 $\varsigma_{\alpha}^{\omega}(X) = \bigcup_{n < \omega} \varsigma_{\alpha}^n(X)$ 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 α^* . Nevertheless, the semantics are different. The ω -strategic semantics would make the following valid dGL formula invalid:

$$\langle (x := 1; x' = 1^d \cup x := x - 1)^* \rangle (0 \leq x < 1) \quad (18)$$

By a simple variation of the argument in the proof of Theorem 3.8, $\varsigma_{\alpha}^{\omega}([0, 1)) = [0, \infty)$, 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 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 all positive values of x . There is no smaller upper bound on the number of iterations it takes Angel to win, in particular Angel cannot promise ω as a bound on the repetition count, which is what the ω -semantics would require her to do. But strategies do converge after $\omega + 1$ iterations. According to Theorem 3.8, similar shortcomings would apply for a semantics that cuts winning region iteration of at higher transfinite ordinals below ω_1^{HG} .

The dGL semantics is also more general, because, by Theorem 3.8, its closure ordinal is $\geq \omega_1^{\text{CK}}$, in contrast to the ω -semantics, which has closure ordinal ω by construction. The same observation shows a fundamental difference between the dGL semantics and the advance notice semantics, which has closure ordinal $\leq \omega$.

C. PROOF OF HIGHER CLOSURE ORDINALS

This section illustrates that closure ordinals are not a simple function of the syntactic structure, because minor syntactic variations lead to vastly different closure ordinals.

PROOF OF THEOREM 3.8. In this proof, proceed in stages of increasing difficulty. That the closure ordinal is $\geq \omega \cdot 2$ has already been shown on p.20. Now 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, follow a similar argument using more variables. Consider the family of formulas (for some $N \in \mathbb{N}$) of the form

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

The winning regions for this dGL formula stabilizes after $\omega \cdot N$ iterations, because ω many iterations are necessary to show that *all* 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. The vector space of variables (x_N, \dots, x_1) is used in that order. It is easy to see that $\zeta^\omega(\alpha)(-\infty, 0)^N = \bigcup_{n < \omega} \zeta^n(\alpha)(-\infty, 0)^N = (-\infty, 0)^{N-1} \times \mathbb{R}$, because $\zeta^{n+1}(\alpha)(-\infty, 0) = (-\infty, 0)^{N-1} \times (-\infty, n)$ holds for all $n \in \mathbb{N}$, n by a simple inductive argument:

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

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

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

Consequently, $\zeta_{\alpha^*}((-\infty, 0)^N) = \zeta^{\omega \cdot N}(\alpha)(-\infty, 0)^N \neq \zeta^{\omega \cdot (N-1) + n}(\alpha)(-\infty, 0)^N$, which then makes $\omega \cdot N$ the closure ordinal for α . Since hybrid games α of the above form can be considered 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 , follow an argument expanding on the previous case. Consider the family of formulas (for some $N \in \mathbb{N}$) of the form

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

The winning region for this ‘‘clockwork ω ’’ formula stabilizes after ω^N iterations, ω many iterations are necessary to show that *all* 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.

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

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

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

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

by induction on the lexicographical order of j and n . Let ①, ②, ③ denote the if conditions on the right, respectively. Note that, in the case ③, there are some subordinate cases which do not need to be tracked in this 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.

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

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

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

Instead, prove base case $j = 1, n = 0$, as the extra offset k has been overcome at ω :

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

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 do not need to be tracked, 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 $j + 1$ prove the case $n = 0$ as follows.

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

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

$$\begin{aligned} \zeta^{\omega^j(n+1)}(\alpha) b_N \dots b_j \dots b_1 &= \zeta^{\omega^j n + \omega^j}(\alpha) b_N \dots b_j \dots b_1 \stackrel{\text{Lemma 3.6}}{=} \zeta^{\omega^j}(\alpha) \zeta^{\omega^j n}(\alpha) b_N \dots b_j \dots b_1 \\ &\stackrel{IH}{=} \begin{cases} \zeta^{\omega^j}(\alpha) b_N \dots (b_{j+1} + n - 1) \infty^j & \text{if } \textcircled{1} \\ \zeta^{\omega^j}(\alpha) b_N \dots (b_{j+1} + n) \infty^j & \text{if } \textcircled{2} \\ \zeta^{\omega^j}(\alpha) b_N \dots (b_{k+1} + 1) 1^{k-j-1} n \infty^j \cup \vec{b} & \text{if } \textcircled{3} \end{cases} \\ &\stackrel{IH}{=} \begin{cases} b_N \dots (b_j + n) \infty^j & \text{if } \textcircled{1} \\ b_N \dots (b_j + n + 1) \infty^j & \text{if } \textcircled{2} \\ b_N \dots (b_{k+1} + 1) 1^{k-j-1} (n+1) \infty^j \cup \vec{b} & \text{if } \textcircled{3} \end{cases} \end{aligned}$$

Consequently, $\zeta_{\alpha^*}((-\infty, 0)^N) = \zeta^{\omega^N}(\alpha)(-\infty, 0)^N = \mathbb{R}^N \neq \zeta^{\omega^{N-1} \cdot n}(\alpha)(-\infty, 0)^N$ for all $n \in \mathbb{N}$, which makes ω^N the closure ordinal for α . Since hybrid games α of the above form can be considered 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 $\geq \omega^\omega$. \square