





Semi-competitive Differential Game Logic

Julia Butte^(✉)  and André Platzer 

Karlsruhe Institute of Technology, Karlsruhe, Germany
{julia.butte, platzer}@kit.edu

Abstract. This paper introduces *semi-competitive differential game logic* dGL_{sc} , which enables verification of safety-critical applications that involve interactions between two agents. In dGL_{sc} , these interactions are specified as games on hybrid systems with two players that *may* collaborate with each other when helpful and *may* compete when necessary. The players in the hybrid games of dGL_{sc} have individual goals that may overlap, leading to nonzero-sum games. This makes dGL_{sc} especially well-suited for verifying situations where players, e.g., share safety objectives but otherwise pursue different goals, so that zero-sum assumptions lead to overly conservative results. Additionally, dGL_{sc} solves the subtlety that even though each player may benefit from knowledge of the other player's goals, e.g., concerning shared safety objectives, unsafe situations might still occur if every player were to mutually assume the other player would act to avoid unsafety. The syntax and semantics, as well as a sound and relatively complete proof calculus are presented for dGL_{sc} . The relationship between dGL_{sc} and zero-sum differential game logic dGL is discussed and the purpose of dGL_{sc} illustrated in a canonical example.

Keywords: Differential game logic · Hybrid systems · Hybrid games · Nonzero-sum games · Cooperative games · Competitive games

1 Introduction

The safety of cyber-physical systems (CPS) is of significant interest to avoid damage to persons and goods due to faulty programs [9, 13, 24]. Cyber-physical systems include trains, planes, robots and autonomous cars [8, 13]. Particularly challenging are situations involving two CPSs, due to their possible interactions which frequently occur in aerial collision avoidance or steering of autonomous cars. To verify, for example, two planes on a collision course, different approaches can be used: The situation could be regarded as one hybrid system, and its safety verified using differential dynamic logic (dL) [12, 20, 21, 23–25]. This corresponds to all planes being centrally controlled, which is infeasible for larger numbers of planes. Or the situation can be modeled with games, regarding each plane as a player which represents the fact that planes are normally steered independently. Collision avoidance could then be verified assuming adversarial players using differential game logic (dGL) [9, 22]. This assumption results in a zero-sum game, i.e. exactly one player wins. But if one pilot is trying to avoid a collision, the

intruder pilot does *not* in fact have the opposite goal of causing a collision! Flying under a zero-sum assumption gives safe but unnecessarily conservative results.

Instead, many games are non-zero-sum, where players have individual goals which often overlap. The planes, for example, both do not want to crash, but they still want to fly into different directions. Such partially shared goals open up the possibility of cooperating, which leads to new winning strategies. Alas, the possibility and will to cooperate is no guarantee that everything goes smoothly. Misunderstandings might still happen, leading for example, to unintentional crashes.

To verify the safety of these two-player non-zero sum games, this paper introduces the *semi-competitive differential game logic* (\mathbf{dGL}_{sc}) which offers a natural way of expressing and reasoning about non-zero sum hybrid games. \mathbf{dGL}_{sc} supports players that behave *semi-competitively*, i.e., both players have individual goals and are open for cooperation but will compete if necessary. They both try to reach their goal while being aware of the other player's goal and can collaborate to jointly meet their respective goals. This strategy provides a solution to the challenging question of when players should cooperate. The possibility of cooperation adds more winning strategies thus enabling safety proofs for situations that are unsafe under conservative zero-sum game assumptions.

To underline the semantics' suitability, the paper derives an alternative representation of semi-competitive games in \mathbf{dGL} . Theoretically, this is possible, since all mixed (co)inductive concepts are definable in \mathbf{dGL} [22]. But as \mathbf{dGL} is fundamentally designed for zero-sum, this requires manual coding, causes redundancies, and duplicates verification effort.

The structure of this paper is as follows: First, Sect. 2 relates and compares relevant logics to \mathbf{dGL}_{sc} . After that, the logic \mathbf{dGL} whose correspondence to \mathbf{dGL}_{sc} will be proved later on, is reviewed in Sect. 3. Then, Sect. 4 introduces the notion of semi-competitiveness and defines the syntax and semantics of \mathbf{dGL}_{sc} . Section 5 establishes important properties of \mathbf{dGL}_{sc} and its relation to \mathbf{dGL} . Afterward, a proof calculus for \mathbf{dGL}_{sc} is introduced in Sect. 6 and its soundness and relative completeness proven.

2 Related Work

Non-zero sum games is a wide field of study that has been addressed by various communities. Game theory provides fundamental definitions of non-zero sum games and studies of their equilibria [15–17] which lay the basis for our work. In terms of game theory, the games played in \mathbf{dGL}_{sc} are non-zero sum sequential games with perfect information and binary payoffs. Whether a player can reach their goal can be considered by backward induction in \mathbf{dGL}_{sc} , similar to the determination of a subgame perfect equilibrium [5]. But \mathbf{dGL}_{sc} has the advantage that the backward induction is done without game trees, which might have uncountable, infinite breadth and unbounded depth due to continuous dynamics and consequently would be hard to handle.

In the field of synthesis, multi-player non-zero sum games played on a graph have been investigated by Fisman et al. who developed rational synthesis [10].

Although the assumptions about the players' behavior are very similar to \mathbf{dGL}_{sc} , their work pursues a different aim: Rational synthesis can be used to construct correct systems, while \mathbf{dGL}_{sc} instead can be used to verify already existing systems.

In purely discrete settings, non-zero sum games have already been addressed by Chatterjee et al. who introduced strategy logic [7]. This is a logic that allows modelling two-player non-zero sum games played on a graph. As explicit quantifications over strategies are possible, the logic is more powerful than e.g. alternating-time temporal logic (ATL) [3] or propositional game logic [18] but also more complex. Therefore, it is better suited for theoretical investigations than for practical use, which \mathbf{dGL}_{sc} supports. Additionally, continuous dynamics cannot be expressed in strategy logic.

For a hybrid setting, there exist rectangular hybrid games [11] or STORMED games [26] which both model players with hybrid automata. While STORMED games are able to express more varieties of continuous dynamics than rectangular hybrid games, it is possible for both of them to perform model checking and controller synthesis. But unlike \mathbf{dGL}_{sc} , both of these hybrid games assume the players to act adversarially. Another work on hybrid games which also assumes adversarial players, by Mitchell et al. [14], computes backwards reachable sets for hybrid games which resemble winning regions in \mathbf{dGL}_{sc} . Unlike \mathbf{dGL}_{sc} , they do not use a comparatively simple state-based semantics, but instead compute these sets by numerically solving Hamilton-Jacobi-Isaac partial differential equations.

Other logics dealing with multi-player settings are e.g., coalition logic [19], ATL [3] or stochastic game logic [4]. These logics are able to verify that a coalition of arbitrarily many players is able to reach a certain goal. The game that is played is defined in the semantics. That has the disadvantage that validity demands a formula to be true for all games, which is only the case for very general statements. Our logic includes the game's definition in the formula, so that safety guarantees can be made for a specific game. Furthermore, the referenced works [3, 4, 19] only consider reachability of outcomes, whereas our work includes goals for all players.

Agotnes et al. [2] include a preference relation in their logic but keep reachability of outcomes and their preferability separate which leaves information unused. Moreover, games are not composable whereas our logic allows modularizing games so that proofs can be reused. Additionally, the players in our work take each other's goals into account to improve their strategies.

Closest to our work are the propositional game logic developed by Parikh [18] which addresses zero-sum two-player games, and differential game logic \mathbf{dGL} by Platzer [22] which extends game logic to include continuous dynamics. Relations to the μ -calculus are discussed in prior work [1]. Our work takes these logics a step further by allowing non-zero sum games instead.

3 Preliminaries

For better understanding of the content of the paper, the logic \mathbf{dGL} [22] will be recalled briefly in this section. In \mathbf{dGL} , there are two players called *Angel* and

Demon. These players play a *hybrid game* which is specified as part of a logical formula. In the following, syntax, semantics and a proof calculus for **dGL** will briefly be explained, based on the book [24].

The syntax of **dGL** is based on first-order logic. Additionally, there are two modalities $\langle \alpha \rangle P$ and $[\alpha]P$. The first one states that Angel can win the game α by achieving her goal P . The second modality means that Demon can win game α by reaching his goal P . Formally, the syntax is defined as follows:

Table 1. Hybrid games

Game	Name	Meaning
$x := e$	Assignment game	assigns e to x
$x' = f(x) \& Q$	Continuous game	Angel evolves ordinary differential equation (ODE) to change value of x while evolution domain constraint Q has to hold
$?Q$	Test game	tests if Angel fulfills Q , if not, she loses and Demon wins automatically
$\alpha \cup \beta$	Choice game	Angel chooses to play either α or β
$\alpha; \beta$	Sequential game	α and β are played sequentially
α^d	Dual game	controls in α are swapped between Angel and Demon
α^*	Repetition game	α is played repeatedly until Angel wants to stop after finite rounds

Definition 1 (dGL Syntax). *Formulas of dGL are defined by the grammar*

$$P, Q ::= e \geq \tilde{e} \mid \neg P \mid P \wedge Q \mid \forall x P \mid \exists x P \mid \langle \alpha \rangle P \mid [\alpha] P$$

where x is a variable, e, \tilde{e} terms, P, Q formulas and α a hybrid game (Table 1).

The formulas are interpreted over states. Each state is a function $\omega : \mathcal{V} \rightarrow \mathbb{R}$ which assigns a real value to each variable. The variable values in state ω_x^r are the same as in ω , except for x whose value is r . The semantics is defined as a function $\llbracket \cdot \rrbracket : Fml \rightarrow \mathcal{P}(\mathcal{S})$ which returns the states where a formula is true. The semantics for the first-order formulas is as usual. The semantics for the two modalities is defined using two functions that retrace from which states a player must have started to reach their goal at the end of the game.

Definition 2. (Semantics) *The dGL semantics is:*

- $\llbracket \langle \alpha \rangle P \rrbracket = \varsigma_\alpha(\llbracket P \rrbracket)$
- $\llbracket [\alpha] Q \rrbracket = \delta_\alpha(\llbracket Q \rrbracket)$

Angel's function for her winning region $\varsigma_\alpha(\cdot)$ is defined as follows:

- $\varsigma_{x:=e}(X) = \{\omega \in \mathcal{S} \mid \omega_x^{\omega[e]} \in X\}$

- $\varsigma_{x'=f(x) \& Q}(X) = \{\varphi(0) \in \mathcal{S} \mid \varphi(r) \in X \text{ for some } r \geq 0 \text{ and (differentiable)}$
 $\varphi : [0, r] \rightarrow \mathcal{S} \text{ such that } \varphi(s) \in \llbracket Q \rrbracket \text{ and } \frac{d\varphi(t)(x)}{dt}(s) = \varphi(s) \llbracket f(x) \rrbracket \text{ for all}$
 $0 \leq s \leq r\} \stackrel{\text{def}}{=} \{\varphi(0) \in \mathcal{S} \mid \varphi(r) \in X \text{ for some } r \geq 0 \text{ with } \varphi \models x' = f(x) \wedge Q\}$
- $\varsigma_{?Q}(X) = \llbracket Q \rrbracket \cap X$
- $\varsigma_{\alpha \cup \beta}(X) = \varsigma_{\alpha}(X) \cup \varsigma_{\beta}(X)$
- $\varsigma_{\alpha; \beta}(X) = \varsigma_{\alpha}(\varsigma_{\beta}(X))$
- $\varsigma_{\alpha^d}(X) = \varsigma_{\alpha}(X^{\mathbf{C}})^{\mathbf{C}}$
- $\varsigma_{\alpha^*}(X) = \bigcap \{Z \subseteq \mathcal{S} \mid X \cup \varsigma_{\alpha}(Z) \subseteq Z\}$

Demon's winning region is defined by $\delta_{\alpha}(X) = \varsigma_{\alpha}(X^{\mathbf{C}})^{\mathbf{C}}$ ([22, Th. 3.1]), i.e., Demon wins whenever Angel fails to reach the opposite of his goal, because he wins whenever Angel loses.

The proof calculus for $\text{d}\mathcal{GL}$ shown in Table 2, consists of all proof rules for first-order logic and one axiom for each game in Angel's modality. Additionally, there is also a determinacy axiom that links Angel's modality to Demon's. All axioms for Demon's modality can be derived using this axiom. Furthermore, the calculus contains a monotonicity rule and a fixpoint rule that handles repetition games.

Table 2. $\text{d}\mathcal{GL}$ proof calculus

$[\cdot]$ $[\alpha]P \leftrightarrow \neg \langle \alpha \rangle \neg P$	$\langle \cup \rangle \langle \alpha \cup \beta \rangle P \leftrightarrow \langle \alpha \rangle P \vee \langle \beta \rangle P$
$\langle := \rangle \langle x := e \rangle p(x) \leftrightarrow p(e)$	$\langle ; \rangle \langle \alpha; \beta \rangle P \leftrightarrow \langle \alpha \rangle \langle \beta \rangle P$
$\langle ' \rangle \langle x' = f(x) \rangle P \leftrightarrow \exists t \geq 0 \langle x := y(t) \rangle P \quad (y' = f(y))$	$\langle ^d \rangle \langle \alpha^d \rangle P \leftrightarrow \neg \langle \alpha \rangle \neg P$
$\langle ? \rangle \langle ?Q \rangle P \leftrightarrow Q \wedge P$	$\langle ^* \rangle \langle \alpha^* \rangle P \leftrightarrow P \vee \langle \alpha \rangle \langle \alpha^* \rangle P$
$\text{M} \frac{P \rightarrow Q}{\langle \alpha \rangle P \rightarrow \langle \alpha \rangle Q}$	$\text{FP} \frac{P \vee \langle \alpha \rangle Q \rightarrow Q}{\langle \alpha^* \rangle P \rightarrow Q}$

4 Semi-competitive Differential Game Logic

The logic $\text{d}\mathcal{GL}_{sc}$ supports two players called *Angel* and *Demon*. They semi-competitively play a game defined in the $\text{d}\mathcal{GL}_{sc}$ formula.

After the game has ended, *both* want to have achieved their goal. Each player has a separate goal which is independent of the other player's goal. The game is *non-zero sum*, both players can win or lose at the same time.

4.1 Semi-competitiveness

To describe the player's behavior in the logic, the notion of “semi-competitiveness” is introduced in this paper. *Semi-competitiveness means that players cooperate where it is helpful for them, and compete where it is necessary.* More precisely, a player will only choose options that help them reach their goal. They will never lose on purpose. If there are multiple options that all make the player

win, the player will *always* choose an option that also helps the other player win. If no option makes the player win, they will not altruistically help the other player instead. In terms of game theory, the players play a cooperative game, if collaboration is possible, and a non-cooperative game otherwise.

This gameplay is chosen to ensure the greatest cooperation among players, especially in games with multiple steps. Here, uncooperative behavior in a previous step might have negative consequences later. For example, if Angel and Demon pick candy for each other and Angel gives Demon a lemon candy, although he wanted a strawberry candy, he might in turn give Angel a candy she does not want because he already lost the game and will not help Angel anymore. If, on the other hand, Angel indeed gives him a strawberry candy, Demon will give her the candy she wants because now all options make him win, so he will help Angel.

Additionally, this gameplay ensures monotonicity of the logic, i.e., a greater goal means more possibilities to win. Monotonicity ensures that the logic behaves intuitively. Furthermore, it guarantees the existence of fixpoints which are needed later on to define parts of the semantics. Helping other players is a necessary assumption to obtain monotonicity: If the players' behavior does not also help the other player win and players choose randomly between all options they are indifferent to, then monotonicity is lost.

4.2 Syntax

The syntax of $\text{d}\mathcal{GL}_{sc}$ consists of first-order logic and two modalities. The modality $\langle \alpha \rangle(P, Q)$ expresses that Angel has a winning strategy to win the game α by achieving her goal P while knowing that Demon's goal is Q . $[\alpha](P, Q)$ means that Demon has a winning strategy to achieve Q by playing α while knowing that Angel's goal is P .

Definition 3. ($\text{d}\mathcal{GL}_{sc}$ Syntax). *The following grammar defines the $\text{d}\mathcal{GL}_{sc}$ formulas:*

- $\alpha, \beta ::= x := e \mid x' = f(x) \& Q \mid ?Q \mid \alpha \cup \beta \mid \alpha; \beta \mid \alpha^d \mid \alpha^*$
- $P, Q ::= e \geq \tilde{e} \mid \neg P \mid P \wedge Q \mid \forall x P \mid \exists x P \mid \langle \alpha \rangle(P, Q) \mid [\alpha](P, Q)$

where α, β are hybrid games, P, Q are $\text{d}\mathcal{GL}_{sc}$ formulas, x is a variable, f is a function and e, \tilde{e} are terms.

The hybrid games in Definition 3 have the same effects as the ones for $\text{d}\mathcal{GL}$ described in Table 1. So in the game $((x := x + 1 \cup x := x - 1); \{x' = -1\}^d)^*$, for example, the outermost game is a repetition game. Here, Angel chooses after each round, if she wants to play again. She might also choose to play 0 rounds. If she does play a round, Angel first chooses whether to increase or to decrease x by one. After that follows a continuous game. As it is marked with d , this game is under Demon's control, so he chooses how long to evolve the ODE to decrease x . That means, time passes and x changes according to the ODE until Demon stops time. In this example, the new value of x would be $x - t$ after time t . Then, Angel chooses again whether she wants to play another round or not.

The main difference now to \mathbf{dGL} is not in the effects of games but rather in the behavior of the players: If we have $\langle (x := 1 \cup x := 0)^d \rangle x = 1$ in \mathbf{dGL} , Demon will choose an option that hinders Angel at achieving her goal due to his adversarial behavior, i.e., $x := 0$ here. In \mathbf{dGL}_{sc} the formula could be $\langle (x := 1 \cup x := 0)^d \rangle (x = 1, \top)$, including the information that Demon's goal is *true*. As Demon can win in any case, Demon helps Angel by choosing $x := 1$ due to his semi-competitive behavior.

Example 1. As the verification of an aerial collision avoidance system is a fully-fledged case study, the two planes in this example are still standing on the ground. Angel and Demon are replenishing the stocks for their respective plane, but Angel is missing 3l of orange juice while Demon needs 5l of tomato juice (more would not fit in the tank, less and some passengers will stay thirsty). Fortunately, Angel has plenty of tomato juice and Demon's stock is full of orange juice (enough, so that they do not have to worry about running out). In \mathbf{dGL}_{sc} , filling up each other's stock can be modeled using two continuous games:

$$o = 0 \wedge t = 0 \rightarrow \langle \{t' = 1\}; \{o' = 1\}^d \rangle (o = 3, t = 5)$$

First, Angel, at a constant rate, fills tomato juice into Demon's tank. Then, Demon fills orange juice into her tank. Note how the goals of both players are visible at the end of the formula. Although these goals are not directly overlapping, the players' semi-competitive behavior ensures their cooperation which leads to them both winning the game. This would not have been possible if they played adversarially because they have no control over the variable that matters for their goal. Even with indifference towards Demon's goal, Angel could not have won the game. If she would not care about the amount of tomato juice she gave to Demon, Demon would probably have refused to cooperate with her, making her lose as well.

4.3 Semantics

The semantics of the formulas is defined in the following section. This is done via a function $\llbracket \cdot \rrbracket : Fml \rightarrow \mathcal{P}(\mathcal{S})$ that maps a formula to the set of states where the formula is true. \mathcal{S} is the set of all states. A state is a mapping $\omega : \mathcal{V} \rightarrow \mathbb{R}$ that maps all variables to a real number. The state ω_x^r denotes a state that coincides with the state ω in all variable values, except for x whose value is replaced by r .

The semantics of first-order operators is as expected. To define the semantics of Angel's and Demon's modality, the functions $\varsigma_\alpha(\cdot, \cdot)$ and $\delta_\alpha(\cdot, \cdot)$ resp., are used which describe the winning region of a game. These functions take two inputs each: The first one is Angel's goal and the second is Demon's goal. Since Angel and Demon know each other's goal, this knowledge is needed for determining the winning regions. Unlike in \mathbf{dGL} , Demon's semantics cannot be defined as a function of Angel's winning region $\varsigma_\alpha(\cdot, \cdot)$ anymore, because their winning regions are no longer complementary. Instead, their winning regions partially overlap and there are also states that belong to neither winning region.

Still, the semantics of the assignment, test, and sequential game is similar to \mathbf{dGL} as the effects of the game are the same, and the players' decisions cannot differ from those they would make in \mathbf{dGL} because these games do not involve choices. A little tweak for the semantics of the sequential game is needed nonetheless, as the winning regions in \mathbf{dGL}_{sc} take two arguments instead of one. To win the sequential game $\alpha; \beta$, the player has to reach their winning region for β after α , so they can reach their goal after β . The same applies for the other player except their respective winning region operator is needed, instead, so the semantics of the sequential game is:

$$\varsigma_{\alpha;\beta}(X, Y) = \varsigma_{\alpha}(\varsigma_{\beta}(X, Y), \delta_{\beta}(X, Y)) \text{ and } \delta_{\alpha;\beta}(X, Y) = \delta_{\alpha}(\varsigma_{\beta}(X, Y), \delta_{\beta}(X, Y))$$

For Angel, the semantics of the continuous game and the choice game also correspond to the semantics in \mathbf{dGL} . The winning region only returns states in which Angel *has* an option so that she can win but which one she will *actually* choose is invisible. Locally, only her goal matters to her because she has control which is also why she cannot expect help from Demon in this step. Globally, semi-competitiveness comes into play through nested games, e.g., the games Angel chooses from in the choice game. Especially important here is the dual game because in this game Demon's goals matter for Angel which then inductively matter for Angel through the whole game.

Demon's semantics of the continuous game and choice game does differ from the one in \mathbf{dGL} though. His winning regions are extended by states where Angel, who is in control, will help him. As they act semi-competitively, Angel will help him in states where both of their goals can be reached simultaneously. Therefore, Demon's semantics for the continuous game $\delta_{x'=f(x) \& Q}(X, Y)$ is:

$$\begin{aligned} & \{ \varphi(0) \in \mathcal{S} \mid \varphi(r) \in Y \text{ for all } r \text{ with } \varphi \models x' = f(x) \wedge Q \} \\ & \cup \{ \varphi(0) \in \mathcal{S} \mid \varphi(r) \in X \cap Y \text{ for some } r \text{ with } \varphi \models x' = f(x) \wedge Q \} \end{aligned}$$

The definition of the semantics uses a function from time to states φ . Over the course of this function, x changes such that $x' = f(x)$ and the evolution domain constraint Q holds in every state in φ . If Demon stays in his goal Y for all states in φ , he can definitely win the game (left-hand side of the union). If Angel can reach a state by evolving the ODE (i.e. there is a state in φ) where both of their goals are fulfilled, she will stop there because she behaves semi-competitively, making both of them win (right-hand side of the union). The definition for the semantics of the choice game follows a corresponding idea:

$$\begin{aligned} \delta_{\alpha \cup \beta}(X, Y) &= (\delta_{\alpha}(X, Y) \cap \delta_{\beta}(X, Y)) \cup (\delta_{\alpha}(X, Y) \cap \varsigma_{\alpha}(X, Y)) \\ &\quad \cup (\delta_{\beta}(X, Y) \cap \varsigma_{\beta}(X, Y)) \end{aligned}$$

If Demon can win both α and β , i.e. he is in the winning region of α and β , he can definitely win the game. If he and Angel can both win α , i.e. the current state is part of both Angel's and Demon's winning region for α , Demon can also win the game because Angel will choose α due to her semi-competitive behavior. A similar argument applies, if both of them can win β .

In the dual game, the controls between Angel and Demon are swapped. This is simulated in the semantics by letting Angel keep control but making her play for Demon, i.e. trying to achieve his goal. In exchange, Demon plays with Angel's goal. As Angel now wins, if Demon's goal is achieved, her winning region actually denotes the states where Demon wins, as well as vice versa. Consequently, Angel and Demon swap goals and winning regions in the definitions of the dual game:

$$\varsigma_{\alpha^d}(X, Y) = \delta_{\alpha}(Y, X) \text{ and } \delta_{\alpha^d}(X, Y) = \varsigma_{\alpha}(Y, X)$$

The main idea to define the semantics of the repetition game is that allowing one more round of α should not change the winning region. Consequently, the winning region for Angel should be some kind of fixpoint of the form $X \cup \varsigma_{\alpha}(\varsigma_{\alpha}^*(X, Y), \delta_{\alpha}^*(X, Y)) = \varsigma_{\alpha}^*(X, Y)$ as Angel either is already in her goal, or she can reach it after one more round of α . The problem is that by the same idea, Demon's winning region should also be a fixpoint, so using a definition like this would result in complex nested fixpoints. Additionally, this fixpoint is not unique yet. To solve the first problem, a case distinction is made between competitive and cooperative behavior, resulting in two fixpoints. If they compete, the other's goal is the complement of the fixpoint because they have complementary goals, while in the cooperative case, the other's goal is the same fixpoint, as they share the same goal. Consequently, a candidate Z for Angel's first fixpoint should fulfill $X \cup \varsigma_{\alpha}(Z, Z^C) = Z$ and for the second it should fulfill $(X \cap Y) \cup (\varsigma_{\alpha}(Z, Z) \cap \delta_{\alpha}(Z, Z)) = Z$. The intersection of their winning regions rules out that one player wins and fulfills both goals (otherwise they would not cooperate), but the other one loses, e.g., by failing a test. To make the fixpoints unique, either the greatest or the least fixpoint can be chosen. \mathcal{S} as the greatest possible set fulfills the equation, but this does not contain much information, so the least fixpoint is used. It can be computed as the intersection of all pre-fixpoints (sets that fulfill the equation only with set inclusion). Taken together, the semantics for Angel is:

$$\begin{aligned} \varsigma_{\alpha}^*(X, Y) = & \bigcap \{Z \subseteq \mathcal{S} \mid X \cup \varsigma_{\alpha}(Z, Z^C) \subseteq Z\} \\ & \cup \bigcap \{Z \subseteq \mathcal{S} \mid (X \cap Y) \cup (\varsigma_{\alpha}(Z, Z) \cap \delta_{\alpha}(Z, Z)) \subseteq Z\} \end{aligned}$$

The definition of Demon's winning region follows the same idea. The second fixpoint is identical because here, both of them win. If Demon competes with Angel, he can only win for sure if he always stays in his goal, no matter how many rounds Angel plays. Consequently, Demon needs to be in his goal now, and he should be in his winning region, even if one more round of α is allowed: $Z = Y \cap \delta_{\alpha}(Z^C, Z)$. Because the least fixpoint in this case is the empty set, which does not contain much information, the greatest fixpoint is used. It can be computed by union of all pre-fixpoints. Thus, the semantics for Demon are:

$$\delta_\alpha(X, Y) = \bigcup \{Z \subseteq \mathcal{S} \mid Z \subseteq Y \cap \delta_\alpha(Z^C, Z)\} \\ \cup \bigcap \{Z \subseteq \mathcal{S} \mid (X \cap Y) \cup (\varsigma_\alpha(Z, Z) \cap \delta_\alpha(Z, Z))\}$$

The full definition of the semantics can be found in the long version [6].

Remark 1. As in \mathbf{dGL} , goals can still be converted to tests at the end, i.e. there is an equivalent formulation for each modality where the goals are replaced by *true* and the original goals appear as tests at the end of the program. As Angel's and Demon's winning regions are not complementary anymore, this conversion is much more subtle than in \mathbf{dGL} because Angel does not always win if Demon loses. Additionally, Demon's ability to win influences his choices and thus Angel's ability to win and vice versa. Therefore, tests for both goals must be incorporated, and Angel and Demon, resp., need to pass them in any order (either test Angel first, or test Demon first) to ensure that they do not only win because the other one loses. This also demonstrates that the opponent's goals matter to both players. For the full statement, reference the long version [6, L. 8].

5 \mathbf{dGL}_{sc} Properties and Relation to Zero-Sum \mathbf{dGL}

In this section, some important results for \mathbf{dGL}_{sc} are proven. The first result proves monotonicity which ensures that the logic behaves in an intuitive way: A larger goal can be reached from a larger winning region. The first lemma states that the logic is monotone if both goals are increased. In this case, both Angel's and Demon's winning region will increase. The proof for this lemma is a straightforward structural induction over α .

Lemma 1. (Monotonicity). *The logic \mathbf{dGL}_{sc} is monotone, i.e.*

$$X \subseteq A, Y \subseteq B \text{ implies } \varsigma_\alpha(X, Y) \subseteq \varsigma_\alpha(A, B) \text{ and } \delta_\alpha(X, Y) \subseteq \delta_\alpha(A, B)$$

See proof in the long version [6].

The second lemma states that if Angel's and Demon's goals are disjoint, Angel's winning region expands by increasing her goal while Demon's goal may change arbitrarily. Intuitively, all parts in the semantics covering cooperation are empty since the goals are disjoint. Then, proving that the remaining sets grow by increasing the goal is fairly easy. After that, the parts where Angel and Demon cooperate, can be added. Thus, the winning region can only increase. The proof for Demon's statement is conducted similarly.

The following lemma is especially helpful if Angel and Demon have opposing goals, before and after the increase of one goal. In this case, enlarging one goal means that the other shrinks. Therefore, the first monotonicity Lemma 1 is not applicable but Lemma 2 can be used.

Lemma 2. (Disjoint Monotonicity). *If the goals of Angel and Demon are initially disjoint, then \mathbf{dGL}_{sc} is monotone in each argument:*

$$X \subseteq A, X \cap Y = \emptyset \text{ implies } \varsigma_\alpha(X, Y) \subseteq \varsigma_\alpha(A, B) \text{ and} \\ Y \subseteq B, X \cap Y = \emptyset \text{ implies } \delta_\alpha(X, Y) \subseteq \delta_\alpha(A, B)$$

Sets A and B may overlap. For each case the other goal can be chosen arbitrarily.

See proof in the long version [6].

Another important property is that the logic behaves as intended. But how can this be checked? If the players have complementary goals, they have to compete and will behave adversarially. The same setting can also be expressed in \mathbf{dGL} (which is assumed to be defined suitably). Consequently, both formulas should also have the same truth values. This holds indeed, as Lemma 3 shows:

Lemma 3. (Complementary \mathbf{dGL} Equivalence). *If Angel and Demon have complementary goals, the logic \mathbf{dGL}_{sc} reduces to \mathbf{dGL} , i.e. it holds that*

$$\varsigma_\alpha(X, X^C) = \varsigma_\alpha X \quad \text{and} \quad \delta_\alpha(X^C, X) = \delta_\alpha X$$

See proof in the long version [6].

In the proof it is shown via structural induction that all parts in the semantics that differ between \mathbf{dGL}_{sc} and \mathbf{dGL} cancel out, if the goals oppose.

Remark 2. If there were a super-logic that contained both \mathbf{dGL}_{sc} and \mathbf{dGL} , the previous lemma could be rendered as the syntactic equivalences

$$\langle \alpha \rangle(P, \neg P) \leftrightarrow \langle \alpha \rangle P \quad \text{and} \quad [\alpha](\neg P, P) \leftrightarrow [\alpha] P$$

Notation 1. *In the following, $\langle \alpha \rangle(P, \neg P)$ will be abbreviated with $\langle \alpha \rangle P$ and $[\alpha](\neg Q, Q)$ with $[\alpha]Q$, as well as $\varsigma_\alpha(X, X^C)$ with $\varsigma_\alpha X$ and $\delta_\alpha(Y^C, Y)$ with $\delta_\alpha Y$. These two formulas express the same and have the same winning regions in \mathbf{dGL}_{sc} and \mathbf{dGL} respectively (Lemma 3). The abbreviations allow for staying in the logic \mathbf{dGL}_{sc} without having to write P, Q, X, Y twice.*

Consequently, determinacy holds even though it only does for complementary goals, which will be useful for the proof calculus later on.

Corollary 1 (Complementary Determinacy). *Determinacy holds if Angel and Demon have complementary goals, i.e. the following equivalence is valid:*

$$\neg \langle \alpha \rangle(P, \neg P) \leftrightarrow [\alpha](P, \neg P)$$

See proof in the long version [6].

To show that \mathbf{dGL}_{sc} 's definition is reasonable for the general case, first observe the following: If Angel and Demon compete in a later game, they will not cooperate in an earlier game because they will not reach both goals at the same time anyway. Similarly, if Angel and Demon cooperate in a later game, they can already do so for earlier games because they can win together. Inductively, this means that Angel and Demon can already decide at the start whether to cooperate or to compete! As competing players try to achieve mutually exclusive goals, each of them might as well play against an adversarial player because the other one might hinder them at achieving their goal in order to reach their own. Two cooperating players can also be modeled as an adversarial game: The coalition is regarded as one player and tries to reach both goals, playing against

an adversarial player with no control. Since one of the two cases always applies, every semi-competitive game can be split into two adversarial games. Using this alternative description, the definition of \mathbf{dGL}_{sc} can be checked against \mathbf{dGL} again. The equivalence holds indeed, as Lemma 4 shows. The $^{-d}$ indicates the removal of all dual operators to transfer all controls to the coalition player. Although the conversion is possible, there are *two* winning regions in \mathbf{dGL} instead of one in \mathbf{dGL}_{sc} which doubles the work for proving formulas. On top, the formulas are linked by disjunction which makes it impossible to get rid of one part if each formula only holds for certain cases. Additionally, \mathbf{dGL}_{sc} offers a much more intuitive way of expressing games where both players have an individual goal.

Lemma 4 (General \mathbf{dGL} Equivalence). *Winning regions in \mathbf{dGL}_{sc} correspond to those in \mathbf{dGL} as follows:*

$$\varsigma_{\alpha}(X, Y) = \varsigma_{\alpha^{-d}}(X \cap Y) \cup \varsigma_{\alpha}(X) \quad \text{and} \quad \delta_{\alpha}(X, Y) = \varsigma_{\alpha^{-d}}(X \cap Y) \cup \delta_{\alpha}(Y)$$

where α is a hybrid game, X, Y are sets of states and α^{-d} is defined in Table 3.

See proof in the long version [6].

As a corollary for Lemma 4, it can be stated that players will cooperate, iff both can win the game.

Corollary 2 (Joint Cooperation). *It holds that*

$$\varsigma_{\alpha}(X, Y) \cap \delta_{\alpha}(X, Y) = \varsigma_{\alpha^{-d}}(X \cap Y)$$

See proof in the long version [6].

6 Proof Calculus

For making practical use of \mathbf{dGL}_{sc} , a proof calculus is provided in this section. The proof calculus includes two axioms for each hybrid game, one for Angel and one for Demon. They correspond to the game's semantic definitions and decompose formulas into smaller parts.

Table 3. Hybrid system α^{-d} is the *systematization* of hybrid game α

$(x := e)^{-d} \equiv x := e$	$(\alpha; \beta)^{-d} \equiv \alpha^{-d}; \beta^{-d}$
$(x' = f(x) \& Q)^{-d} \equiv x' = f(x) \& Q$	$(\alpha^d)^{-d} \equiv \alpha^{-d}$
$(?Q)^{-d} \equiv ?Q$	$(\alpha^*)^{-d} \equiv (\alpha^{-d})^*$
$(\alpha \cup \beta)^{-d} \equiv \alpha^{-d} \cup \beta^{-d}$	

$\langle := \rangle$	$\langle x := e \rangle(p(x), q(x)) \leftrightarrow p(e)$	
$\langle ' \rangle$	$\langle x' = f(x) \rangle(P, Q) \leftrightarrow \exists t \geq 0 \langle x := y(t) \rangle(P, Q)$	$(y' = f(y))$
$\langle ? \rangle$	$\langle ?R \rangle(P, Q) \leftrightarrow R \wedge P$	
$\langle \cup \rangle$	$\langle \alpha \cup \beta \rangle(P, Q) \leftrightarrow \langle \alpha \rangle(P, Q) \vee \langle \beta \rangle(P, Q)$	
$\langle ; \rangle$	$\langle \alpha; \beta \rangle(P, Q) \leftrightarrow \langle \alpha \rangle(\langle \beta \rangle(P, Q), [\beta](P, Q))$	
$\langle ^d \rangle$	$\langle \alpha^d \rangle(P, Q) \leftrightarrow [\alpha](Q, P)$	
$\langle * \rangle$	$\langle \alpha^* \rangle(P, Q) \leftrightarrow P \vee \langle \alpha; \alpha^* \rangle(P, Q)$	
FP	$\frac{P \vee \langle \alpha \rangle R_1 \rightarrow R_1 \quad (P \wedge Q) \vee (\langle \alpha \rangle(R_2, R_2) \wedge [\alpha](R_2, R_2)) \rightarrow R_2}{\langle \alpha^* \rangle(P, Q) \rightarrow R_1 \vee R_2}$	
$[:=]$	$[x := e](p(x), q(x)) \leftrightarrow q(e)$	
$[']$	$[x' = f(x)](P, Q) \leftrightarrow \forall t \geq 0 [x := x(t)](P, Q) \vee \exists t \geq 0 \langle x := x(t) \rangle(P \wedge Q, Q)$	$(x' = f(x))$
$[?]$	$[?R](P, Q) \leftrightarrow \neg R \vee Q$	
$[\cup]$	$[\alpha \cup \beta](P, Q) \leftrightarrow ([\alpha](P, Q) \wedge [\beta](P, Q)) \vee ([\alpha](P, Q) \wedge \langle \alpha \rangle(P, Q)) \vee ([\beta](P, Q) \wedge \langle \beta \rangle(P, Q))$	
$[:]$	$[\alpha; \beta](P, Q) \leftrightarrow [\alpha](\langle \beta \rangle(P, Q), [\beta](P, Q))$	
$[*]$	$[\alpha^*](P, Q) \leftrightarrow (P \wedge Q) \vee (\langle \alpha; \alpha^* \rangle(P, Q) \wedge [\alpha; \alpha^*](P, Q)) \vee (Q \wedge [\alpha; \alpha^*]Q)$	
M $\langle \rangle$	$\frac{P_1 \rightarrow P_2 \quad Q_1 \rightarrow Q_2}{\langle \alpha \rangle(P_1, Q_1) \rightarrow \langle \alpha \rangle(P_2, Q_2)}$	
M ₂ $\langle \rangle$	$\frac{P_1 \rightarrow P_2 \quad P_1 \wedge Q_1 \rightarrow \perp}{\langle \alpha \rangle(P_1, Q_1) \rightarrow \langle \alpha \rangle(P_2, Q_2)}$	
det	$\neg \langle \alpha \rangle P \leftrightarrow [\alpha] \neg P$	

Fig. 1. Proof calculus for dGL_{sc}

Additionally, there is the determinacy axiom det , which corresponds to Corollary 1 and links Angel and Demon's modality, if they have opposing goals. The fixpoint rule FP is, like the other rules, written with the premises at the top and the conclusion at the bottom. The rule characterizes $\langle \alpha^* \rangle(P, Q)$ as the union of two least fixpoints. If R_1 is a pre-fixpoint for the first fixpoint and R_2 is a pre-fixpoint for the second, then $R_1 \vee R_2$ holds whenever $\langle \alpha^* \rangle(P, Q)$ holds. As a subtle detail, R_1 and R_2 do not need to be the same because Angel can rarely achieve the same by competing and by cooperating with Demon. Unlike axiom $\langle * \rangle$, the rule does not feature more instances of repetition games in the premises. Demon's induction rule is not included because it can be derived from FP. The other two proof rules M $\langle \rangle$ and M₂ $\langle \rangle$ correspond to the two monotonicity lemmas. Unlike the axioms, the monotonicity rules for Demon can be derived by using $\langle ^d \rangle$. For the monotonicity rules, it is irrelevant that this changes the game. But for Demon's axioms, the dual game introduced by $\langle ^d \rangle$ prevents the use of any other axiom to derive them.

Additionally, the proof calculus includes all rules for first-order logic. The proof calculus' rules are listed in Fig. 1, except for the first-order rules which are left out for reasons of space limitations and readability.

Example 2. With these proof rules on hand, it can now be proved that Angel is actually able to win the game seen in Example 1 and gets enough orange juice to supply all her passengers but not so much that it spills everywhere. In the

proof, the hybrid game is straightforwardly broken down using the continuous, dual and assignment axioms:

$$\begin{array}{c}
 \frac{\mathbb{R} \quad *}{\text{VR, wR, } \langle := \rangle, \exists \text{R}} \frac{o = 0, t = 0 \vdash y + 5 = 5 \wedge o + 3 = 3}{o = 0, t = 0 \vdash \forall s \geq 0 [o := o + s](t + 5 = 5, o = 3)} \\
 \frac{\quad \vee \exists s \geq 0 \langle o = o + s \rangle (t + 5 = 5 \wedge o = 3, o = 3)}{\text{[']}} \\
 \frac{\text{[']}}{\text{ER, } \langle := \rangle} \frac{o = 0, t = 0 \vdash [x' = 1](t + 5 = 5, o = 3)}{o = 0, t = 0 \vdash \exists s \geq 0 \langle t := t + s \rangle ([x' = 1](t = 5, o = 3),} \\
 \frac{\quad \langle x' = 1 \rangle (t = 5, o = 3))}{\langle ' \rangle} \\
 \frac{\langle ' \rangle}{\langle ; \rangle, \langle ^d \rangle, [^d]} \frac{o = 0; t = 0 \vdash \langle y' = 1 \rangle ([x' = 1](t = 5, o = 3),}{o = 0, t = 0 \vdash \langle \{t' = 1\}; \{o' = 1\}^d \rangle (o = 3, t = 5)}
 \end{array}$$

Notably, the point where the proof splits into two cases (cooperation or competition), only occurs in the penultimate step of the proof. In this case it is already conceivable which option is provable, and the other branch can be eliminated, but usually this is hard to see. Therefore, this saves considerable effort compared to using dGL where both cases had to be considered from the start.

To be of any use, soundness of the proof calculus is crucial. Soundness means that all formulas that can be proven in the proof calculus are actually valid. If a false formula could be proven, anything could be, rendering it useless. Therefore, the soundness of the proof calculus is proven in the following theorem:

Theorem 1 (Soundness). *The dGL_{sc} proof calculus is sound.*

See proof in the long version [6].

Another important property of a proof calculus is completeness. Completeness means that any valid formula can be proven in the proof calculus. But how can this property be shown for dGL_{sc} ? For dGL it is already known that it is relatively complete [22]. This means that any valid formula can be proved if all tautologies may be assumed in the antecedent, i.e. the notion is less strict than completeness. For dGL_{sc} only relative completeness can be proven, as this logic builds on top of dGL . Adapting the definitions from [22], relative completeness is defined as:

Definition 4 (Expressiveness [22]). *A logic L is expressive (for dGL_{sc}), if for every formula ϕ , there exists an equivalent formula ϕ^b of L , i.e. $\models \phi \leftrightarrow \phi^b$. The logic L is called differentially expressive, if it is expressive and all equivalences of the form $\langle x' = \theta \rangle (P, Q) \leftrightarrow (\langle x' = \theta \rangle (P, Q))^b$ and $[x' = \theta](P, Q) \leftrightarrow ([x' = \theta](P, Q))^b$ are provable in the proof calculus. It is assumed that the logic L is closed under first-order logic.*

Definition 5 (Relative completeness [22]). *A logic is complete relative to an expressive logic L , if every valid formula can be proved in the calculus from L tautologies.*

If now, the \mathbf{dGL}_{sc} proof calculus could be converted to the \mathbf{dGL} proof calculus, relative completeness could be proven almost for free. In the semantics this conversion is possible by Lemma 4. The proof calculus though, does not contain any axioms that syntactically rephrase this lemma. The inclusion of a corresponding axiom would also run counter to the intended compositionality principles of \mathbf{dGL}_{sc} . Fortunately, this is also unnecessary, as Theorem 2 shows a core insight: The two complementarization axioms (C.A. 1, 2) that split one modality with two individual player goals into two modalities with complementary goals, are *admissible*! Consequently, the set of provable formulas is unaffected by adding them. Hence, similar to Gentzen’s cut elimination theorem, the compositional axioms of \mathbf{dGL}_{sc} prove every \mathbf{dGL}_{sc} formula that the complementarization axioms would reduce to \mathbf{dGL} , thereby guaranteeing relative completeness for free.

Theorem 2 (Complementarization Elimination). *The complementarization axioms are admissible in the proof calculus:*

- Complementarization axiom 1: $\langle \alpha \rangle(P, Q) \leftrightarrow \langle \alpha^{-d} \rangle(P \wedge Q) \vee \langle \alpha \rangle P$
- Complementarization axiom 2: $[\alpha](P, Q) \leftrightarrow \langle \alpha^{-d} \rangle(P \wedge Q) \vee [\alpha]Q$

See proof in the long version [6].

Using Theorem 3, the relative completeness of \mathbf{dGL}_{sc} can easily be proven:

Theorem 3 (Completeness). *The \mathbf{dGL}_{sc} proof calculus is complete relative to any differentially expressive logic L .*

Proof (Sketch) With complementarization, every formula of \mathbf{dGL}_{sc} can be transformed to a formula where all modalities feature opposing goals. The complementarization axioms are admissible by Theorem 2, so the rewrite can be mimicked using only the rules available in the \mathbf{dGL}_{sc} proof calculus. For opposing goals, the rules in the proof calculus correspond to the rules of the proof calculus for \mathbf{dGL} which is complete relative to any differentially expressive logic L [22]. Consequently, the proof calculus for \mathbf{dGL}_{sc} is also relatively complete.

See full proof in the long version [6].

7 Conclusion

This paper introduces the logic \mathbf{dGL}_{sc} for reasoning about hybrid games with two players that have individual goals. \mathbf{dGL}_{sc} helps determine if a player is able to achieve their goal after playing a given game while taking into account the goal of the other player to *compete or collaborate if needed*. Additionally, the paper introduces the notion of “semi-competitiveness” to characterize the players’ behavior.

Syntax, semantics and a proof calculus for \mathbf{dGL}_{sc} are given in this paper. Monotonicity is shown to hold for \mathbf{dGL}_{sc} and the winning regions of \mathbf{dGL}_{sc} coincide with those of zero-sum \mathbf{dGL} for complementary goals. Furthermore, complementarization has been proven admissible, illustrating the semi-competitive behavior of the players and justifying that the semantics is defined suitably.

In the future, several extensions of $\text{d}\mathcal{GL}_{sc}$ would be interesting. First, players might behave suboptimally by not always cooperating when necessary. In this case, players would need to actively agree on cooperating. As zero-sum three player games can be reduced to non-zero sum two-player games [17], $\text{d}\mathcal{GL}_{sc}$ could also be extended to a three-player logic. This logic could then be used for safety proofs of more complex scenarios with more players like overtaking maneuvers.

Acknowledgements. This work was funded by an Alexander von Humboldt Professorship and by KiKIT, the Pilot Program for Core-Informatics at the KIT of the Helmholtz Association.

References

1. Abou El Wafa, N., Platzer, A.: Complete game logic with sabotage. In: Sobocinski, P., Lago, U.D., Esparza, J. (eds.) LICS, pp. 1:1–1:15. ACM, New York (2024). <https://doi.org/10.1145/3661814.3662121>
2. Ågotnes, T., van der Hoek, W., Wooldridge, M.J.: On the logic of coalitional games. In: Nakashima, H., Wellman, M.P., Weiss, G., Stone, P. (eds.) 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006), Hakodate, Japan, May 8–12, 2006, pp. 153–160. ACM (2006). <https://doi.org/10.1145/1160633.1160659>
3. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. J. ACM **49**(5), 672–713 (2002). <https://doi.org/10.1145/585265.585270>
4. Baier, C., Brázdil, T., Größer, M., Kucera, A.: Stochastic game logic. In: Fourth International Conference on the Quantitative Evaluation of Systems (QEST 2007), 17–19 September 2007, Edinburgh, Scotland, UK, pp. 227–236. IEEE Computer Society (2007). <https://doi.org/10.1109/QEST.2007.38>
5. Bielefeld, R.S.: Reexamination of the perfectness concept for equilibrium points in extensive games. In: Models of Strategic Rationality, pp. 1–31. Springer Netherlands, Dordrecht (1988). https://doi.org/10.1007/978-94-015-7774-8_1
6. Butte, J., Platzer, A.: Semi-competitive differential game logic. CoRR abs/2505.14688 (2025). <https://doi.org/10.48550/ARXIV.2505.14688>
7. Chatterjee, K., Henzinger, T.A., Piterman, N.: Strategy logic. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007, Lisbon, Portugal, September 3–8, 2007, Proceedings. LNCS, vol. 4703, pp. 59–73. Springer (2007). https://doi.org/10.1007/978-3-540-74407-8_5
8. Chattopadhyay, A., Lam, K.: Security of autonomous vehicle as a cyber-physical system. In: 7th International Symposium on Embedded Computing and System Design, ISED 2017, Durgapur, India, December 18–20, 2017, pp. 1–6. IEEE (2017). <https://doi.org/10.1109/ISED.2017.8303906>
9. Cleaveland, R., Mitsch, S., Platzer, A.: Formally verified next-generation airborne collision avoidance games in ACAS x. ACM Trans. Embed. Comput. Syst. **22**(1), 1–30 (2023). <https://doi.org/10.1145/3544970>
10. Fisman, D., Kupferman, O., Lustig, Y.: Rational synthesis. In: Esparza, J., Majumdar, R. (eds.) TACAS 2010, Paphos, Cyprus, March 20–28, 2010. Proceedings. LNCS, vol. 6015, pp. 190–204. Springer (2010). https://doi.org/10.1007/978-3-642-12002-2_16

11. Henzinger, T.A., Horowitz, B., Majumdar, R.: Rectangular hybrid games. In: Baeten, J.C.M., Mauw, S. (eds.) CONCUR 1999, Eindhoven, The Netherlands, August 24–27, 1999, Proceedings. LNCS, vol. 1664, pp. 320–335. Springer (1999). <https://doi.org/10.1007/3-540-48320-9-23>
12. Jeannin, J., et al.: A formally verified hybrid system for the next-generation airborne collision avoidance system. In: Baier, C., Tinelli, C. (eds.) TACAS. LNCS, vol. 9035, pp. 21–36. Springer (2015). <https://doi.org/10.1007/978-3-662-46681-0-2>
13. Kabra, A., Mitsch, S., Platzer, A.: Verified train controllers for the Federal Railroad Administration train kinematics model: balancing competing brake and track forces. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* **41**(11), 4409–4420 (2022). <https://doi.org/10.1109/TCAD.2022.3197690>
14. Mitchell, I., Bayen, A., Tomlin, C.: A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games. *IEEE Trans. Autom. Control* **50**(7), 947–957 (2005). <https://doi.org/10.1109/TAC.2005.851439>
15. Nash, J.F.: Equilibrium points in n-person games. *Proc. Natl. Acad. Sci.* **36**(1), 48–49 (1950). <https://doi.org/10.1073/pnas.36.1.48>
16. Nash, J.F.: Non-cooperative games. In: *The Foundations of Price Theory Vol 4*, pp. 329–340. Routledge (2024). <https://doi.org/10.2307/1969529>
17. von Neumann, J., Morgenstern, O.: *Theory of Games and Economic Behavior*. Princeton University Press, Princeton (2004). <https://doi.org/10.1515/9781400829460>
18. Parikh, R.: The logic of games and its applications. In: Karplinski, M., van Leeuwen, J. (eds.) *Topics in the Theory of Computation*, North-Holland Mathematics Studies, vol. 102, pp. 111–139. North-Holland (1985). [https://doi.org/10.1016/S0304-0208\(08\)73078-0](https://doi.org/10.1016/S0304-0208(08)73078-0)
19. Pauly, M.: A modal logic for coalitional power in games. *J. Log. Comput.* **12**(1), 149–166 (2002). <https://doi.org/10.1093/LOGCOM/12.1.149>
20. Platzer, A.: Differential dynamic logic for hybrid systems. *J. Autom. Reas.* **41**(2), 143–189 (2008). <https://doi.org/10.1007/s10817-008-9103-8>
21. Platzer, A.: *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer, Heidelberg (2010). <https://doi.org/10.1007/978-3-642-14509-4>
22. Platzer, A.: Differential game logic. *ACM Trans. Comput. Log.* **17**(1), 1:1–1:51 (2015). <https://doi.org/10.1145/2817824>
23. Platzer, A.: A complete uniform substitution calculus for differential dynamic logic. *J. Autom. Reason.* **59**(2), 219–265 (2016). <https://doi.org/10.1007/s10817-016-9385-1>
24. Platzer, A.: *Logical Foundations of Cyber-Physical Systems*. Springer, Cham (2018). <https://doi.org/10.1007/978-3-319-63588-0>
25. Platzer, A., Tan, Y.K.: Differential equation invariance axiomatization. *J. ACM* **67**(1), 6:1–6:66 (2020). <https://doi.org/10.1145/3380825>
26. Vladimerou, V., Prabhakar, P., Viswanathan, M., Dullerud, G.E.: Specifications for decidable hybrid games. *Theor. Comput. Sci.* **412**(48), 6770–6785 (2011). <https://doi.org/10.1016/J.TCS.2011.08.036>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

