

Lecture Notes on Differential & Temporal Logics

André Platzer

Carnegie Mellon University
Lecture 16

1 Introduction

This course is devoted to the study of the [Foundations of Cyber-Physical Systems](#) [Pla12c, Pla12b]. [Lecture 3 on Choice & Control](#) explained hybrid programs, a program notation for hybrid systems [Pla08, Pla10, Pla12c, Pla12a]. [Lecture 4 on Safety & Contracts](#) defined differential dynamic logic [Pla08, Pla10, Pla12c, Pla12a] as a specification and verification logic for hybrid programs. [Lecture 5 on Dynamical Systems & Dynamic Axioms](#) and subsequent lectures studied proof principles for differential dynamic logic with which we can prove correctness properties of hybrid systems. In your labs, you have demonstrated aptly how you can model, specify, and verify quite sophisticated and challenging robots.

Yet, there was one rather puzzling phenomenon that we noticed in [Lecture 4](#) only then did not have a chance to consider any further. For a hybrid program α and differential dynamic logic formula ϕ , the modal formula

$$[\alpha]\phi$$

expresses that all *final* states reached by all runs of α satisfy the logical formula ϕ . The modal formula $[\alpha]\phi$ is, consequently, false exactly in those states from which α can reach a final state that violates the safety condition ϕ . Yet, what about states from which the final state reached by running α is safe but some intermediate state along the execution of α was not safe?

Shouldn't systems that violate safety condition ϕ at an intermediate state be considered unsafe as well?

The short answer is: that depends.

Does it even make a difference whether we study intermediate states as well or only worry about final states?

The short answer is again: that depends.

What exactly it depends on and how to systematically approach the general case of safety *throughout* the system execution is what today's lecture studies. The key to the answer will be understanding the temporal behavior of hybrid programs. The hybrid trace semantics of hybrid programs will also give us a deeper understanding of the hybrid aspect of time in hybrid systems.

This lecture is based on [Pla10, Chapter 4], which is a significant extension of [Pla07].

2 Temporalizing Hybrid Systems

In order to be able to distinguish whether a CPS is safe at the end of its run or safe always throughout its run, differential dynamic logic $d\mathcal{L}$ will be extended with additional temporal modalities. The resulting logic extends $d\mathcal{L}$ and is called *differential temporal dynamic logic* (dTL) [Pla10, Chapter 4]. The modal formula

$$[\alpha]\phi$$

of $d\mathcal{L}$ [Pla08, Pla12c] expresses that all *final* states reached by all runs of α satisfy the logical formula ϕ . The same $d\mathcal{L}$ formula $[\alpha]\phi$ is allowed in the logic dTL and has the same semantics [Pla10, Chapter 4]. The new temporal modal dTL formula

$$[\alpha]\Box\phi$$

instead, expresses that all states reached *all along* all traces of α satisfy ϕ . Those two modalities can be used to distinguish systems that are always throughout from those that are only safe in final states. For example, if the dTL formula

$$[\alpha]\phi \wedge \neg[\alpha]\Box\phi$$

is true in an initial state ν , then the system α will be safe (in the sense of satisfying ϕ) in all final states reached after running α from ν , but is not safe always throughout all traces of all runs of α from ν . Can that happen?

You should try to answer this question before it is discussed in a later part of these lecture notes.

3 Syntax of Differential Temporal Dynamic Logic

The *differential temporal dynamic logic* dTL extends differential dynamic logic [Pla08, Pla10, Pla12c] with temporal modalities for verifying temporal specifications of hybrid systems. Hence, dTL has two kinds of modalities:

Modal operators. Modalities of dynamic logic express statements about all possible behaviour ($[\alpha]\pi$) of a system α , or about the existence of a trace ($\langle\alpha\rangle\pi$), satisfying condition π . Unlike in standard dynamic logic, α is a model of a hybrid system.

The logic dTL uses hybrid programs to describe α as in previous lectures. Yet, unlike in standard dynamic logic [HKT00] or dL, π is a *trace formula* in dTL, and π can refer to all states that occur *during* a trace using temporal operators.

Temporal operators. For dTL, the temporal trace formula $\Box\phi$ expresses that the formula ϕ holds all along a trace selected by $[\alpha]$ or $\langle\alpha\rangle$. For instance, the state formula $\langle\alpha\rangle\Box\phi$ says that the state formula ϕ holds at every state along at least one trace of α . Dually, the trace formula $\Diamond\phi$ expresses that ϕ holds at some point during such a trace. It can occur in a state formula $\langle\alpha\rangle\Diamond\phi$ to express that there is such a state in some trace of α , or as $[\alpha]\Diamond\phi$ to say that along each trace there is a state satisfying ϕ . The primary focus of attention in today's lecture is on homogeneous combinations of path and trace quantifiers like $[\alpha]\Box\phi$ or $\langle\alpha\rangle\Diamond\phi$.

The formulas of dTL are defined similarly to differential dynamic logic. However, the modalities $[\alpha]$ and $\langle\alpha\rangle$ accept trace formulas that refer to the temporal behavior of *all* states along a trace. Inspired by CTL and CTL* [EC82, EH86], dTL distinguishes between state formulas, which are true or false in states, and trace formulas, which are true or false for system traces. The sets Fml of state formulas and Fml_T of trace formulas with variables in Σ are simultaneously inductively defined in Def. 1.

Definition 1 (dTL formula). The (state) formulas of differential temporal dynamic logic (dTL) are defined by the grammar (where ϕ, ψ are dTL state formulas, π is a dTL trace formula, θ_1, θ_2 (polynomial) terms, x a variable, α a HP):

$$\phi, \psi ::= \theta_1 = \theta_2 \mid \theta_1 \geq \theta_2 \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi \mid \phi \rightarrow \psi \mid \forall x \phi \mid \exists x \phi \mid [\alpha]\pi \mid \langle\alpha\rangle\pi$$

The *trace formulas* of dTL are defined by the grammar (where ϕ is a dTL state formula):

$$\pi ::= \phi \mid \Box\phi \mid \Diamond\phi$$

Operators $>, \leq, <, \leftrightarrow$ can be defined as usual, e.g., $\phi \leftrightarrow \psi \equiv (\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$.

Formulas without \Box and \Diamond are *nontemporal formulas* and have the same semantics as the corresponding dL formulas. Unlike in CTL, dTL state formulas are true on a trace if they hold for the *last* state of a trace, not for the first. Thus, dTL formula $[\alpha]\phi$ expresses that ϕ is true at the end of each trace of α , which is the same as the dL formula $[\alpha]\phi$. In contrast, $[\alpha]\Box\phi$ expresses that ϕ is true *all along* all states of every trace of α . This combination gives a smooth embedding of nontemporal dL into dTL and makes it possible to define a compositional calculus. Like CTL, dTL allows nesting with a branching time semantics [EC82], e.g., $[\alpha]\Box(x \geq 2 \rightarrow \langle\beta\rangle\Diamond x \leq 0)$.

4 Trace Semantics of Hybrid Programs

In differential dynamic dL [Pla08, Pla12c] from Lecture 4, modalities only refer to the final states of system runs and the semantics is a reachability relation on states: State ω

is reachable from state ν using system α if there is a run of α which terminates in ω when started in ν . For dTL, however, formulas can refer to intermediate states of runs as well. To capture this, we change the semantics of a hybrid system α to be the set of its possible *traces*, i.e., successions of states that occur during the evolution of α . The relation between the initial and the final state alone is not sufficient.

States define the values of system variables during a hybrid evolution. A *state* is a map $\nu : \Sigma \rightarrow \mathbb{R}$. In addition, we distinguish a separate state Λ to denote the failure of a system run when it is *aborted* due to a test $? \chi$ that yields *false*. In particular, Λ can only occur at the end of an aborted system run and marks that no further extension of that trace is possible because of a failed test. The set of all states is denoted by \mathcal{S} .

Hybrid systems evolve along piecewise continuous traces in multi-dimensional space as time passes. Continuous phases are governed by differential equations, whereas discontinuities are caused by discrete jumps in state space. Unlike in discrete cases [Pra79, BS01], traces are not just sequences of states, since hybrid systems pass through uncountably many states even in bounded time. Beyond that, continuous changes are more involved than in pure real time [ACD90, HNSY92], because all variables can evolve along differential equations with different slopes. Generalizing the real-time traces of [HNSY92], the following definition captures hybrid behaviour by splitting the uncountable succession of states into periods σ_i that are regulated by the same control law. For discrete jumps, some of those periods are point flows of duration 0.

The (trace) semantics of hybrid programs is compositional, that is, the semantics of a complex program is defined as a simple function of the trace semantics of its parts. What a hybrid trace captures is the full temporal evolution of a hybrid system. Hybrid systems can behave in different ways, so their trace semantics will be a set of hybrid traces, each of which describes one particular temporal evolution over time. Time, however, is hybridized to a pair (i, ζ) of a discrete time index $i \in \mathbb{N}$ and a real time point $\zeta \in \mathbb{R}$. A single time component $\zeta \in \mathbb{R}$ itself would be an inadequate model of time for hybrid systems, because hybrid systems can make progress by a discrete transition without continuous time passing. That happens whenever discrete controls take action. Continuous time only passes during continuous evolutions along differential equations. Discrete actions only make discrete time index i pass.

Definition 2 (Hybrid trace). A *trace* is a (nonempty) finite or infinite sequence $\sigma = (\sigma_0, \sigma_1, \sigma_2, \dots)$ of functions $\sigma_i : [0, r_i] \rightarrow \mathcal{S}$ with their respective durations $r_i \in \mathbb{R}$ (for $i \in \mathbb{N}$). A *position* of σ is a pair (i, ζ) with $i \in \mathbb{N}$ and ζ in the interval $[0, r_i]$; the state of σ at (i, ζ) is $\sigma_i(\zeta)$. Positions of σ are ordered lexicographically by $(i, \zeta) \prec (j, \xi)$ iff either $i < j$, or $i = j$ and $\zeta < \xi$. Further, for a state $\nu \in \mathcal{S}$, $\hat{\nu} : 0 \mapsto \nu$ is the *point flow* at ν with duration 0. A trace *terminates* if it is a finite sequence $(\sigma_0, \sigma_1, \dots, \sigma_n)$ and $\sigma_n(r_n) \neq \Lambda$. In that case, the last state $\sigma_n(r_n)$ is denoted by *last* σ . The first state $\sigma_0(0)$ is denoted by *first* σ .

Unlike in [ACD90, HNSY92], the definition of traces also admits finite traces of bounded duration, which is necessary for compositionality of traces in $\alpha; \beta$. The semantics of

hybrid programs α as the set $\tau(\alpha)$ of its possible traces depends on valuations $\llbracket \cdot \rrbracket_\nu$ of formulas and terms at intermediate states ν . The valuation of terms and interpretations of function and predicate symbols are as for real arithmetic (Lecture 4). The valuation of formulas will be defined in Def. 6. Again, we use ν_x^d to denote the *modification* that agrees with state ν on all variables except for the symbol x , which is changed to $d \in \mathbb{R}$.

Definition 3 (Trace semantics of hybrid programs). The *trace semantics*, $\tau(\alpha)$, of a hybrid program α , is the set of all its possible hybrid traces and is defined inductively as follows:

1. $\tau(x := \theta) = \{(\hat{\nu}, \hat{\omega}) : \omega = \nu \text{ except that } \llbracket x \rrbracket_\omega = \llbracket \theta \rrbracket_\nu \text{ for } \nu \in \mathcal{S}\}$
2. $\tau(x' = \theta \ \& \ H) = \{(\varphi) : \varphi(t) \models x' = \theta \text{ and } \varphi(t) \models H \text{ for all } 0 \leq t \leq r \text{ for a solution } \varphi : [0, r] \rightarrow \mathcal{S} \text{ of any duration } r\}$; i.e., with $\varphi(t)(x') \stackrel{\text{def}}{=} \frac{d\varphi(\zeta)(x)}{d\zeta}(t)$, φ solves the differential equation and satisfies H at all times, see Lecture 2.
3. $\tau(?X) = \{(\hat{\nu}) : \llbracket X \rrbracket_\nu = \text{true}\} \cup \{(\hat{\nu}, \hat{\Lambda}) : \llbracket X \rrbracket_\nu = \text{false}\}$
4. $\tau(\alpha \cup \beta) = \tau(\alpha) \cup \tau(\beta)$
5. $\tau(\alpha; \beta) = \{\sigma \circ \varsigma : \sigma \in \tau(\alpha), \varsigma \in \tau(\beta) \text{ when } \sigma \circ \varsigma \text{ is defined}\}$;
the composition of $\sigma = (\sigma_0, \sigma_1, \sigma_2, \dots)$ and $\varsigma = (\varsigma_0, \varsigma_1, \varsigma_2, \dots)$ is

$$\sigma \circ \varsigma := \begin{cases} (\sigma_0, \dots, \sigma_n, \varsigma_0, \varsigma_1, \dots) & \text{if } \sigma \text{ terminates at } \sigma_n \text{ and last } \sigma = \text{first } \varsigma \\ \sigma & \text{if } \sigma \text{ does not terminate} \\ \text{not defined} & \text{otherwise} \end{cases}$$
6. $\tau(\alpha^*) = \bigcup_{n \in \mathbb{N}} \tau(\alpha^n)$, where $\alpha^{n+1} := (\alpha^n; \alpha)$ for $n \geq 1$, as well as $\alpha^1 := \alpha$ and $\alpha^0 := (?true)$.

Time passes differently during discrete and continuous change. During continuous evolution, the discrete step index i of positions (i, ζ) remains constant, whereas the continuous duration ζ remains 0 during discrete point flows. This permits multiple discrete state changes to happen at the same (super-dense) continuous time, unlike in other approaches [ACD90].

Example 4. For comparing the transition semantics of hybrid programs for dL from Lecture 3 and the trace semantics of hybrid programs for dTL from Def. 3, consider the following simple hybrid program α :

$$a := -2a; a := a^2.$$

The transition semantics is just the relation between initial and final states:

$$\rho(\alpha) \equiv \{(\nu, \omega) : \omega \text{ is like } \nu \text{ except that } \omega(a) = 4\nu(a)^2\}.$$

In particular, the $d\mathcal{L}$ formula $[\alpha]a \geq 0$ is valid, because all final states have a square as the value of a . In contrast, the trace semantics of α retains all intermediate states:

$$\tau(\alpha) \equiv \{(\hat{\nu}, \hat{s}, \hat{\omega}) : s \text{ is like } \nu \text{ except } s(a) = -2\nu(a) \\ \text{and } \omega \text{ is like } s \text{ except } \omega(a) = s(a)^2 = 4\nu(a)^2\}.$$

During these traces, $a \geq 0$ does not hold at all states. If the trace starts with a positive value ($\nu \models a > 0$), then it will become negative at the point flow s (where $s \models a < 0$), yet recover to a positive value ($\omega \models a > 0$) at the end.

Example 5. The previous example only had discrete jumps, and, thus, the traces only involved point flows. Now consider the hybrid program β from the train context:

$$a := -b; z' = v, v' = a; ?v \geq 0; a := A; z' = v, v' = a.$$

The transition semantics of this program only considers successful runs to completion. In particular, if $A > 0$, the velocity v will always be nonnegative at the end (otherwise the test $?v \geq 0$ in the middle fails and the program aborts), because the last differential equation will accelerate and increase the velocity again. Thus, the position z at the end of the program run will never be smaller than at the beginning.

If, instead, we consider the trace semantics of β , all intermediate states are in the set of traces:

$$\tau(\beta) \equiv \{(\hat{\mu}_0, \hat{\mu}_1, \varphi_1, \hat{\mu}_2, \hat{\mu}_3, \varphi_2) : \mu_1 = \mu_0[a \mapsto -\mu_0(b)] \text{ and} \\ \varphi_1 \text{ is a state flow of some duration } r_1 \geq 0 \text{ with } \varphi_1 \models z' = v \wedge v' = a \\ \text{starting in } \varphi_1(0) = \mu_1 \text{ and ending in a state with } \varphi_1(r_1)(v) \geq 0 \\ \text{and } \mu_2 = \varphi_1(r_1), \mu_3 = \varphi_1(r_1)[a \mapsto \varphi_1(r_1)(A)] \text{ and} \\ \varphi_2 \text{ is a state flow of some duration } r_2 \geq 0 \text{ with } \varphi_2 \models z' = v \wedge v' = a \\ \text{starting in } \varphi_2(0) = \mu_3 \text{ and ending in state } \varphi_2(r_2)\} \\ \cup \{(\hat{\mu}_0, \hat{\mu}_1, \varphi_1, \hat{\mu}_2, \hat{\Lambda}) : \mu_1 = \mu_0[a \mapsto -\mu_0(b)] \text{ and} \\ \varphi_1 \text{ is a state flow of some duration } r \geq 0 \text{ with } \varphi_1 \models z' = v \wedge v' = a \\ \text{starting in } \varphi_1(0) = \mu_1 \text{ and ending in a state with } \varphi_1(r)(v) < 0 \\ \text{further } \mu_2 = \varphi_1(r)\}.$$

The first set is the set of traces where the test $?v \geq 0$ in the middle succeeds and the system continues. The second set (after the union) is the set of traces that are aborted with $\hat{\Lambda}$ during their execution, because the middle test fails. Note that the traces in the first set have two continuous flows φ_1, φ_2 and four point flows $\hat{\mu}_0, \hat{\mu}_1, \hat{\mu}_2, \hat{\mu}_3$ in each trace. The traces in the second set have only one continuous flow φ_1 and three point flows $\hat{\mu}_0, \hat{\mu}_1, \hat{\mu}_2$, because the subsequent aborting point flow $\hat{\Lambda}$ does not terminate and aborts all further execution. In the trace semantics, $v < 0$ is possible in the middle of some traces, which is a fact that the transition semantics does not notice. Combining traces for $\alpha \cup \beta$, that is, for

$$(a := -2a; a := a^2) \cup (a := -b; z' = v, v' = a; ?v \geq 0; a := A; z' = v, v' = a)$$

is just the union $\tau(\alpha) \cup \tau(\beta)$ of the traces $\tau(\alpha)$ and $\tau(\beta)$ from Examples 4 and 5. Note that $a \leq 0$ will hold at least once during every trace of $\alpha \cup \beta$, either in the beginning, or after setting $a := -2a$ or $a := -b$, respectively, when we assume $b > 0$.

5 Semantics of State and Trace Formulas

In the semantics of dTL formulas, the dynamic modalities determine the set of traces according to the trace semantics of hybrid programs, and, independently, the temporal modalities determine at which points in time the respective postcondition needs to hold. The semantics of formulas is compositional and denotational, that is, the semantics of a complex formula is defined as a simple function of the semantics of its subformulas.

Definition 6 (dTL semantics). The *satisfaction relation* $\nu \models \phi$ for a dTL (state) formula ϕ in state ν is defined inductively:

- $\nu \models (\theta_1 = \theta_2)$ iff $\llbracket \theta_1 \rrbracket_\nu = \llbracket \theta_2 \rrbracket_\nu$.
- $\nu \models (\theta_1 \geq \theta_2)$ iff $\llbracket \theta_1 \rrbracket_\nu \geq \llbracket \theta_2 \rrbracket_\nu$.
- $\nu \models \neg\phi$ iff $\nu \not\models \phi$, i.e. if it is not the case that $\nu \models \phi$.
- $\nu \models \phi \wedge \psi$ iff $\nu \models \phi$ and $\nu \models \psi$.
- $\nu \models \phi \vee \psi$ iff $\nu \models \phi$ or $\nu \models \psi$.
- $\nu \models \phi \rightarrow \psi$ iff $\nu \not\models \phi$ or $\nu \models \psi$.
- $\nu \models \phi \leftrightarrow \psi$ iff $(\nu \models \phi \text{ and } \nu \models \psi)$ or $(\nu \not\models \phi \text{ and } \nu \not\models \psi)$.
- $\nu \models \forall x \phi$ iff $\nu_x^d \models \phi$ for all $d \in \mathbb{R}$.
- $\nu \models \exists x \phi$ iff $\nu_x^d \models \phi$ for some $d \in \mathbb{R}$.
- $\nu \models [\alpha]\pi$ iff for each trace $\sigma \in \tau(\alpha)$ that starts in first $\sigma = \nu$, if $\llbracket \pi \rrbracket_\sigma$ is defined, then $\llbracket \pi \rrbracket_\sigma = \text{true}$.
- $\nu \models \langle \alpha \rangle \pi$ iff there is a trace $\sigma \in \tau(\alpha)$ starting in first $\sigma = \nu$ such that $\llbracket \pi \rrbracket_\sigma$ is defined and $\llbracket \pi \rrbracket_\sigma = \text{true}$.

For trace formulas, the *valuation* $\llbracket \cdot \rrbracket_\sigma$ with respect to trace σ is defined inductively as:

1. If ϕ is a state formula, then $\llbracket \phi \rrbracket_\sigma = \llbracket \phi \rrbracket_{\text{last } \sigma}$ if σ terminates, whereas $\llbracket \phi \rrbracket_\sigma$ is *not defined* if σ does not terminate.
2. $\llbracket \Box \phi \rrbracket_\sigma = \text{true}$ iff $\sigma_i(\zeta) \models \phi$ holds for all positions (i, ζ) of σ with $\sigma_i(\zeta) \neq \Lambda$.
3. $\llbracket \Diamond \phi \rrbracket_\sigma = \text{true}$ iff $\sigma_i(\zeta) \models \phi$ holds for some position (i, ζ) of σ with $\sigma_i(\zeta) \neq \Lambda$.

As usual, a (state) formula is *valid* if it is true in all states. If $\nu \models \phi$, then we say that dTL state formula ϕ is true at ν or that ν is a model of ϕ . A (state) formula ϕ is *valid*, written $\models \phi$, iff $\nu \models \phi$ for all states ν . A formula ϕ is a *consequence* of a set of formulas Γ , written $\Gamma \models \phi$, iff, for each ν : $(\nu \models \psi \text{ for all } \psi \in \Gamma)$ implies that $\nu \models \phi$. Likewise, for trace formula π and trace σ we write $\sigma \models \pi$ iff $\llbracket \pi \rrbracket_\sigma = \text{true}$ and $\sigma \not\models \pi$ iff $\llbracket \pi \rrbracket_\sigma = \text{false}$. In particular, we only write $\sigma \models \pi$ or $\sigma \not\models \pi$ if $\llbracket \pi \rrbracket_\sigma$ is defined, which it is not the case if π is a state formula and σ does not terminate. The points where a dTL property ϕ has to hold for the various combinations of temporal and dynamic modalities are illustrated in Fig. 1.

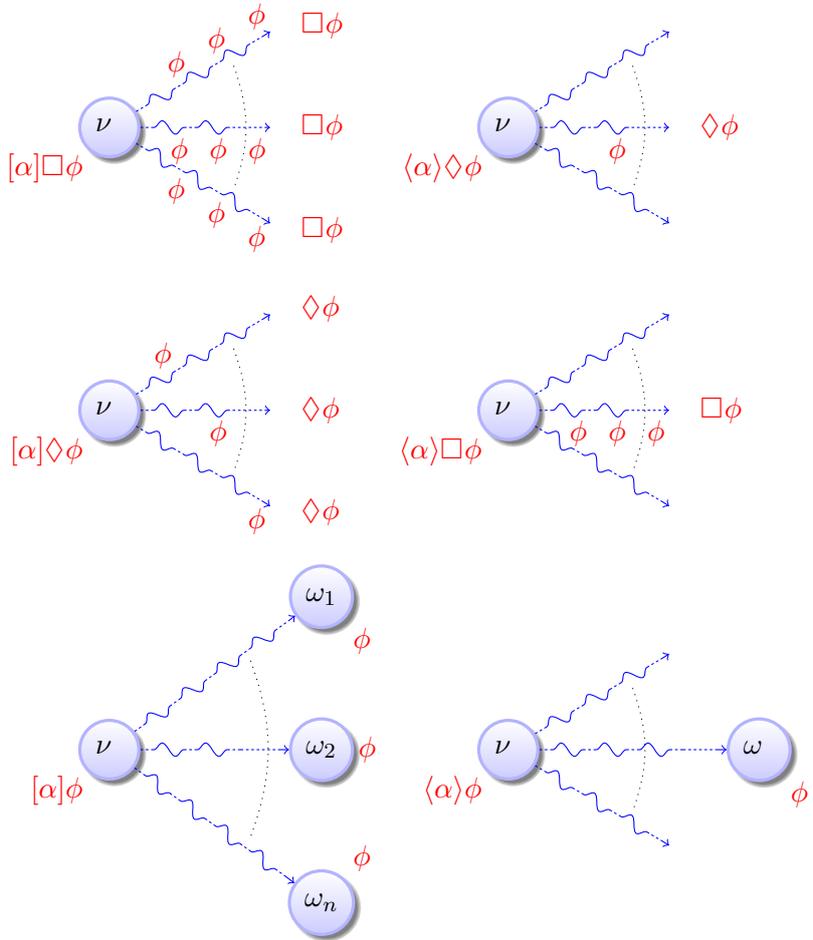


Figure 1: Trace semantics of dTL formulas

6 Conservative Temporal Extension

The following result shows that the extension by temporal operators that dTL provides does not change the meaning of nontemporal $d\mathcal{L}$ formulas. The trace semantics given in Def. 6 is equivalent to the final state reachability relation semantics given in Lecture 4 for the sublogic $d\mathcal{L}$ of dTL.

Proposition 7 (Conservative temporal extension [Pla10, Proposition 4.1]). *The logic dTL is a conservative extension of nontemporal $d\mathcal{L}$, i.e., the set of valid $d\mathcal{L}$ formulas is the same with respect to transition reachability semantics of $d\mathcal{L}$ (Lecture 4) as with respect to the trace semantics of dTL (Def. 6).*

The proof is by induction using that the reachability relation fits to the trace semantics. That is, the reachability relation semantics of hybrid programs agrees with the first and last states of the traces in the trace semantics.

Lemma 8 (Trace relation [Pla10, Lemma 4.1]). *For hybrid programs α :*

$$\rho(\alpha) = \{(\text{first } \sigma, \text{last } \sigma) : \sigma \in \tau(\alpha) \text{ terminates}\}.$$

In particular, the trace semantics from today's lecture fits seamlessly to the original reachability semantics that was the basis for the previous lectures. The trace semantics exactly satisfies the objective of characterizing the same reachability relation between initial and final states, while, in addition, keeping a trace of all intermediate states around. For nontemporal dTL formulas and for $d\mathcal{L}$ formulas, this full trace with intermediate states is not needed, because the reachability relation between initial and final states is sufficient to define the meaning. For temporal dTL formulas, instead, the trace is crucial to give a meaning to \Box and \Diamond .

7 Summary

This lecture introduced a temporal extension of the logic $d\mathcal{L}$ and a trace semantics of hybrid programs. This extends the syntax and semantics to the presence of temporal modalities. The next lecture investigates how to prove temporal properties of hybrid systems.

Exercises

Exercise 1. Can you give a formula of the following form that is valid?

$$[\alpha]\Box\phi \wedge \neg[\alpha]\phi$$

Exercise 2. In which case does the temporal $[\alpha]\Box\phi$ differ from the nontemporal $[\alpha]\phi$.

References

- [ACD90] Rajeev Alur, Costas Courcoubetis, and David L. Dill. Model-checking for real-time systems. In *LICS*, pages 414–425. IEEE Computer Society, 1990.
- [BS01] Bernhard Beckert and Steffen Schlager. A sequent calculus for first-order dynamic logic with trace modalities. In Rajeev Goré, Alexander Leitsch, and Tobias Nipkow, editors, *IJCAR*, volume 2083 of *LNCS*, pages 626–641. Springer, 2001.
- [DBL12] *Proceedings of the 27th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25–28, 2012*. IEEE, 2012.
- [EC82] E. Allen Emerson and Edmund M. Clarke. Using branching time temporal logic to synthesize synchronization skeletons. *Sci. Comput. Program.*, 2(3):241–266, 1982.
- [EH86] E. Allen Emerson and Joseph Y. Halpern. “Sometimes” and “Not Never” revisited: on branching versus linear time temporal logic. *J. ACM*, 33(1):151–178, 1986.
- [HKT00] David Harel, Dexter Kozen, and Jerzy Tiuryn. *Dynamic logic*. MIT Press, 2000.
- [HNSY92] Thomas A. Henzinger, Xavier Nicollin, Joseph Sifakis, and Sergio Yovine. Symbolic model checking for real-time systems. In *LICS*, pages 394–406. IEEE Computer Society, 1992.
- [Pla07] André Platzer. A temporal dynamic logic for verifying hybrid system invariants. In Sergei N. Artëmov and Anil Nerode, editors, *LFCS*, volume 4514 of *LNCS*, pages 457–471. Springer, 2007. doi:10.1007/978-3-540-72734-7_32.
- [Pla08] André Platzer. Differential dynamic logic for hybrid systems. *J. Autom. Reas.*, 41(2):143–189, 2008. doi:10.1007/s10817-008-9103-8.
- [Pla10] André Platzer. *Logical Analysis of Hybrid Systems: Proving Theorems for Complex Dynamics*. Springer, Heidelberg, 2010. doi:10.1007/978-3-642-14509-4.
- [Pla12a] André Platzer. The complete proof theory of hybrid systems. In *LICS* [DBL12], pages 541–550. doi:10.1109/LICS.2012.64.
- [Pla12b] André Platzer. Dynamic logics of dynamical systems. *CoRR*, abs/1205.4788, 2012. arXiv:1205.4788.
- [Pla12c] André Platzer. Logics of dynamical systems. In *LICS* [DBL12], pages 13–24. doi:10.1109/LICS.2012.13.

[Pra79] Vaughan R. Pratt. Process logic. In *POPL*, pages 93–100, 1979.