

Lecture Notes on Linear Temporal Logic

Ruben Martins

Carnegie Mellon University

Lecture 21

Thursday, April 4, 2024

1 Introduction

When using dynamic logic, we prove properties of (all or some) final states of the program under complete ignorance of what happens during the execution of the program. However, for instance, there is a potential difference between whether a data structure invariant holds literally always at all times during all runs of all its operations (which is essentially a prerequisite for uncorrupted concurrent usages) compared to whether the data structure invariant merely holds at the end of each of the operations if it was true before. There are ways of augmenting dynamic logics to temporal dynamic logics that provide explicit ways of proving formulas that are true, e.g., always throughout an execution.

While this works well and continues the deductive verification principles we saw so far, we will, instead, leverage the motivation of a temporal understanding of programs as a segway into studying temporal logics and their use in model checking.

Learning Goals

- Learn about the linear temporal logic and the semantics of the temporal operators $\Box P$, $\Diamond P$, $\mathbf{X}P$ and PUQ .
- Learn about Kripke structures, a widely-used finite-state structure for modeling computation in model checking.
- Learn about real-world tools for modeling and checking validity of temporal logic formulas.

2 Traces of Programs

If we want to study whether a formula such as a data structure invariant is true all the time always throughout the execution of a program, we will at least have to retain all the states that the program visits in its semantics. One possible approach for that is to make the semantics remember the set of all traces that a program can exhibit where a trace is a sequence of states that the program visited along the way. Another possible approach is to emphasize the structure that the transitions within the program took and then generate possible traces from that.

While the technical nuances of the following definitions are somewhat subtle, the main point is quite intuitive. With a little bit of care, we can retain all intermediate states during the execution of a program as an entire trace, instead of just retaining the fact that a particular final state was reachable from a particular initial state by running a program to completion.

A trace is either a finite sequence of states of some length $n \in \mathbb{N}$,

$$(\sigma_0, \sigma_1, \sigma_2, \dots, \sigma_n) \quad (1)$$

or it is an infinite sequence of states, one for each natural number, of length ∞ :

$$(\sigma_0, \sigma_1, \sigma_2, \sigma_3, \dots)$$

A trace terminates iff it is finite and its last state, σ_n in Eq. 1, is not the special failure state Λ , which indicates abortion due to a failed test. The special error state Λ does not provide values for any variables (so no formulas or terms can be evaluated in it), but it is used to mark the end of an aborted execution trace. No program ever continues from the error state Λ . For stylistic reasons, a trace of length n has $n + 1$ states.

Definition 1 (Trace semantics of programs). The *trace semantics*, $\tau(\alpha)$, of a program α , is the set of all its possible traces and is defined inductively as follows:

1. $\tau(x := e) = \{(\omega, \nu) : \nu = \omega \text{ except that } \nu(x) = \omega \llbracket e \rrbracket \text{ for } \omega \in S\}$
2. $\tau(?Q) = \{(\omega) : \omega \models Q\} \cup \{(\omega, \Lambda) : \omega \not\models Q\}$
3. $\tau(\text{if}(Q) \alpha \text{ else } \beta) = \{\sigma \in \tau(\alpha) : \sigma_0 \models Q\} \cup \{\sigma \in \tau(\beta) : \sigma_0 \not\models Q\}$
4. $\tau(\alpha; \beta) = \{\sigma \circ \varsigma : \sigma \in \tau(\alpha), \varsigma \in \tau(\beta)\};$
the composition of $\sigma = (\sigma_0, \sigma_1, \sigma_2, \dots)$ and $\varsigma = (\varsigma_0, \varsigma_1, \varsigma_2, \dots)$ is

$$\sigma \circ \varsigma := \begin{cases} (\sigma_0, \dots, \sigma_n, \varsigma_1, \varsigma_2, \dots) & \text{if } \sigma \text{ terminates in } \sigma_n \text{ and } \sigma_n = \varsigma_0 \\ \sigma & \text{if } \sigma \text{ does not terminate} \\ \text{not defined} & \text{otherwise} \end{cases}$$

5. $\tau(\text{while}(Q) \alpha) = \{\sigma^{(0)} \circ \sigma^{(1)} \circ \dots \circ \sigma^{(n)} : \text{for some } n \geq 0 \text{ such that for all } 0 \leq i < n:$
 - ① the loop condition is true $\sigma_0^{(i)} \models Q$ and ② $\sigma^{(i)} \in \llbracket \alpha \rrbracket$ and ③ $\sigma^{(n)}$ either does not

terminate or it terminates in $\sigma_m^{(n)}$ and $\sigma_m^{(n)} \not\models Q$ in the end}
 $\cup \{ \sigma^{(0)} \circ \sigma^{(1)} \circ \sigma^{(2)} \circ \dots : \text{for all } i \in \mathbb{N}: \textcircled{1} \sigma_0^{(i)} \models Q \text{ and } \textcircled{2} \sigma^{(i)} \in \llbracket \alpha \rrbracket \}$
 $\cup \{ (\omega) : \omega \not\models Q \}$

That is, the loop either runs a nonzero finite number of times with the last iteration either terminating or running forever, or the loop itself repeats infinitely often and never stops, or the loop does not even run a single time.

6. $\tau(\alpha^*) = \bigcup_{n \in \mathbb{N}} \tau(\alpha^n)$ where $\alpha^{n+1} \stackrel{\text{def}}{=} (\alpha^n; \alpha)$ for $n \geq 1$, and $\alpha^1 \stackrel{\text{def}}{=} \alpha$ and $\alpha^0 \stackrel{\text{def}}{=} (?true)$.

All cases in this definition are under the assumption that the respective compositions are defined. For example

$$\tau(\alpha; \beta) = \{ \sigma \circ \varsigma : \sigma \in \tau(\alpha), \varsigma \in \tau(\beta) \text{ when } \sigma \circ \varsigma \text{ is defined} \}$$

3 Linear Temporal Logic

Now that we have a set of traces such as the ones $\tau(\alpha)$ generated by a program α , we have more temporal information about the sequence of states that happened during the run of the program. That enables us to talk more about the way how truth-values change over time along such a trace.

Definition 2 (LTL). The formulas of linear temporal logic (LTL) with formulas over individual states F are defined by the following grammar:

$$P, Q ::= F \mid \neg P \mid P \wedge Q \mid \mathbf{X}P \mid \Box P \mid \Diamond P \mid P \mathbf{U} Q$$

The formula $\Box P$ means that P is always true in the future. The formula $\Diamond P$ means that P is sometimes true in the future, meaning at least at one point. The formula $\mathbf{X}P$ means that P is true in the next state. And the formula $P \mathbf{U} Q$ means that P is true until Q is true (which also will be true at some point).

The suffix of a trace σ starting at step $k \in \mathbb{N}$ is denoted σ^k and only defined if the trace has at least length k . That is

$$(\sigma_0, \sigma_1, \sigma_2, \dots, \sigma_{k-1}, \sigma_k, \sigma_{k+1}, \sigma_{k+2}, \dots)^k = (\sigma_k, \sigma_{k+1}, \sigma_{k+2}, \dots)$$

In particular σ^0 is the same as σ . Also $(\sigma_0) \circ \sigma^1 = \sigma$ if the trace has at least length 1 so that σ^1 is defined.

Definition 3. The truth of LTL formulas in a trace σ is defined inductively as follows:

1. $\sigma \models F$ iff $\sigma_0 \models F$ for a state formula F provided that $\sigma_0 \neq \Lambda$
2. $\sigma \models \neg P$ iff $\sigma \not\models P$, i.e. it is not the case that $\sigma \models P$
3. $\sigma \models P \wedge Q$ iff $\sigma \models P$ and $\sigma \models Q$
4. $\sigma \models \mathbf{X}P$ iff $\sigma^1 \models P$




Figure 1: Examples of traces satisfying LTL formulas

5. $\sigma \models \Box P$ iff $\sigma^i \models P$ for all $i \geq 0$
6. $\sigma \models \Diamond P$ iff $\sigma^i \models P$ for some $i \geq 0$
7. $\sigma \models P \mathbf{U} Q$ iff there is an $i \geq 0$ such that $\sigma^i \models Q$ and $\sigma^j \models P$ for all $0 \leq j < i$

In all cases, the truth-value of a formula is, of course, only defined if the respective suffixes of the traces are defined.

For example, $\mathbf{X}P$ only has a truth-value in trace σ if σ^1 is defined, which means that the trace σ has length 1 (recall that this means it has at least 1+1 states). So $\mathbf{X}P$ is neither true or false but simply meaningless in a trace such as σ_0 that does not actually have a next state. Likewise, $\mathbf{X}p$ is meaningful (and either true or false depending on whether p is true in σ_1) in a trace $\sigma = (\sigma_0, \sigma_1)$, but $\mathbf{X}\mathbf{X}p$ is not meaningful in the same trace because it's not long enough to have a successor of a successor.

Note that the meaning of the box and diamond modalities of LTL is quite analogous to the meaning that the box and diamond modalities already have in dynamic logic. The only difference is that dynamic logic modalities range over the runs of a concrete program while the modalities of LTL range over time (along a fixed trace of something). This is not a coincidence. Both are versions of modal logics, which differ in terms of what the box and diamond modalities range over but are otherwise built similarly.

For the cases $\mathbf{X}P, \Box P, \Diamond P, P \mathbf{U} Q$ It is, of course, very important to retain the entire suffix of the trace for the semantics (not just a single state) in case the subformulas P and Q themselves mention further temporal operators. For example, LTL formula

$$\Box \Diamond P$$

expresses that P is true infinitely often when referring to an infinite trace. On a finite trace, it merely means that P is true in the last (non-failure) state.

The LTL formula

$$\Diamond \Box P$$

expresses that P is eventually true all the time (so is true almost always, so except at finitely many exception states) when referring to an infinite trace. On a finite trace, it also merely means that P is true in the last (non-failure) state.

4 Exercises

We can use LTL semantics to reason about the validity of LTL formulas. For instance, consider the following LTL equivalences that characterize distributive properties of temporal operators:

1. $\Diamond(P \vee Q) \leftrightarrow \Diamond P \vee \Diamond Q$
2. $\Box(P \vee Q) \leftrightarrow \Box P \vee \Box Q$

We will show that (1) is valid using the semantics of the LTL operators while (2) is incorrect by providing a counterexample.

$$\Diamond(P \vee Q) \leftrightarrow \Diamond P \vee \Diamond Q$$

This formula is valid. We will need to prove both directions. In this lecture, we show how to prove “ \rightarrow ” and leave the other direction as an exercise.

1. Assume the left hand side: $\Diamond(P \vee Q)$
2. By the semantics of \Diamond : $\sigma \models \Diamond(P \vee Q)$ iff $\sigma^i \models (P \vee Q)$ for some $i \geq 0$
3. By the semantics of \vee : $\sigma^i \models P \vee \sigma^i \models Q$ for some $i \geq 0$
4. By the semantics of \Diamond : $\sigma \models \Diamond P \vee \sigma \models \Diamond Q$
5. By the semantics of \vee : $\sigma \models \Diamond P \vee \Diamond Q$

$$\Box(P \vee Q) \leftrightarrow \Box P \vee \Box Q$$

This formula is not correct. Consider an infinite trace $\sigma = (\sigma_0, \sigma_1, \dots)$ where $\sigma_0 \models (Q \wedge \neg P)$, and $\sigma_i \models (\neg Q \wedge P)$ for all $i \geq 1$.

- We have that $\sigma \models \Box(P \vee Q)$ because at every state in the trace, either P or Q is satisfied.
- However, $\sigma \not\models \Box P$ since P is not satisfied in σ_0 .
- Also, $\sigma \not\models \Box Q$ since Q is not satisfied in many states (e.g., σ_1).
- Since $\sigma \not\models \Box P$ and $\sigma \not\models \Box Q$ then $\sigma \not\models \Box P \vee \Box Q$

5 LTL Formulas on Program Traces

The following very clever program solves the issue of subtracting from negative numbers by first turning them into positive numbers and then adding, while ultimately flipping the sign again.

$$\begin{aligned} x &\leftarrow -x; \\ x &\leftarrow x + 7; \\ x &\leftarrow -x \end{aligned}$$

This program does correctly subtract 7 from a negative number as witnessed by a corresponding proof of the following dynamic logic formula:

$$x = x_0 \rightarrow [x \leftarrow -x; x \leftarrow x + 7; x \leftarrow -x] x = x_0 - 7$$

This formula means that whenever x_0 equals the initial value of variable x then *after* running the program, the resulting value of x will be the result of subtracting 7 from x_0 , which, since it didn't change, still is the initial value of x . We can write this as an LTL formula in terms of its traces. In particular, we can say that all traces σ of the program satisfy the following LTL:

$$\sigma \models x = x_0 \rightarrow \Box \Diamond (x = x_0 - 7)$$

Inspecting the semantics of this formula, it says that:

- If $\sigma \models x = x_0$, which is true if and only if $\sigma_0 \models x = x_0$,
- then $\sigma \models \Box \Diamond (x = x_0 - 7)$. In other words, $\sigma^i \models \Diamond (x = x_0 - 7)$ for all $i \geq 0$, so there always exists $j \geq i$ where $\sigma^j \models x = x_0 - 7$.

Because the traces of this program are finite, this imposes the same condition as the DL box modality, that the final state will satisfy $x = x_0 - 7$ when the initial state satisfies $x = x_0$.

The program also satisfies the property that if x is initially negative then x is finally negative:

$$x < 0 \rightarrow [x \leftarrow -x; x \leftarrow x + 7; x \leftarrow -x] x < 0$$

But it *does not* satisfy that x is negative always at all times while running the program, because the whole point is that the first assignment flips the sign of x into a positive number. In fact, all traces of this program are of the following form:

$$\tau(x \leftarrow -x; x \leftarrow x + 7; x \leftarrow -x) = \{(\omega, \omega_x^{-\omega(x)}, \omega_x^{-\omega(x)+7}, \omega_x^{-(\omega(x)+7)}) : \omega \text{ is any state}\}$$

Consequently, if $\sigma \in \tau(x \leftarrow -x; x \leftarrow x + 7; x \leftarrow -x)$ is a trace of this program starting in an initial state σ_0 with negative initial value of x so $\sigma_0(x) < 0$, then the LTL formula

$\Box(x < 0)$ is *not* true for it even if it is true initially and in the end. Indeed, all traces $\sigma \in \tau(x \leftarrow -x; x \leftarrow x + 7; x \leftarrow -x)$ of the program satisfy:

$$\sigma \not\models x < 0 \rightarrow \Box(x < 0)$$

That is, the following condition is false for σ :

if $x < 0$ is true (initially, because there's no temporal operator on the left hand side of the implication), then $x < 0$ is true always in the future (of σ).

But what is, indeed, true for all traces σ of the program is:

$$\sigma \models x < 0 \rightarrow \Box(x \neq 0)$$

That is, if x starts negative then it will always be nonzero at every point in time throughout the entire trace σ .

6 Modeling Computation Over Time

So far we have always modeled computation with programs, written in a language with semantics that allow us to reason about their behavior mathematically. However, in the context of model checking, it is often the case that complex systems are modeled in terms of transition systems called *Kripke structures*, given in Definition 4.

Definition 4 (Kripke structure). A *Kripke frame* (W, \leadsto) consists of a set W with a transition relation $\leadsto \subseteq W \times W$ where $s \leadsto t$ indicates that there is a direct transition from s to t in the Kripke frame (W, \leadsto) . The elements $s \in W$ are also called states. A *Kripke structure* $K = (W, \leadsto, v, I)$ is a Kripke frame (W, \leadsto) with a mapping $v : W \rightarrow 2^V$, where 2^V is the powerset of V assigning truth-values to all the propositional atoms in all states. Moreover, a Kripke structure has a set of initial states $I \subseteq W$.

A Kripke structure $K = (W, \leadsto, v, I)$ is called a *computation structure* if W is a finite set of states and every element $s \in W$ has at least one direct successor $t \in W$ with $s \leadsto t$. A (computation) *path* is an infinite sequence $s_0, s_1, s_2, s_3, \dots$ of states $s_i \in W$ such that $s_i \leadsto s_{i+1}$ for all i . We will always assume that the structures used in model checking are computation structures, unless otherwise noted.

Example 5. An example of a Kripke structure that represents a vending machine is shown in Figure 2.

The set of states W represented in Figure 2 are $W = \{s_0, s_1, s_2, s_3\}$. The propositional atoms V that appear in those states are $V = \{\text{coin}, \text{select}, \text{coffee}, \text{tea}\}$. The initial state $I = \{s_0\}$. The mapping v is represented as follows:

$$\begin{aligned} s_0 &\rightarrow \{\text{coin}\} \\ s_1 &\rightarrow \{\text{select}\} \\ s_2 &\rightarrow \{\text{coffee}\} \\ s_3 &\rightarrow \{\text{tea}\} \end{aligned}$$




Figure 2: Computation structure describing the operation of a vending machine.

Note that we only show the propositional atoms that are assigned the truth value *true* but the remaining atoms would be assigned truth value *false*. Finally, the transition relation \leadsto is defined as: $\{s_0 \leadsto s_1, s_1 \leadsto s_2, s_1 \leadsto s_3, s_2 \leadsto s_0, s_3 \leadsto s_0\}$.

What are examples of properties that we might want to check against such a computation? For one, being able to conclude that the vending machine does not dispense coffee or tea until after a user produces a coin would be helpful. Likewise, we may want to check that once a coin has been presented, the machine will eventually dispense *some* beverage. These are easy to check by visual inspection of Figure 2. In general, this will not be the case, as the number of states and transitions comprising a system are too large to simply inspect.

7 LTL Model Checking in Practice

In this lecture, we will not cover algorithms for checking if a LTL formula is valid or if a LTL formula satisfies a given Kripke structure. However, we use existing tools to explore LTL model checking in practice. An old but still functional tool is NuSMV that is available at <https://nusmv.fbk.eu/>.

Let's start by modeling the Kripke structure of the vending machine in Figure 2. We start by defining the state variables and the propositional atoms:

```

VAR
  state: {s0, s1, s2, s3};
  input: {coin, select, coffee, tea};

```

Next, we specify which is the initial state and the propositional atoms that are true in that state:

```

ASSIGN
  init(state) := {s0};
  init(input) := {coin};

```


Then, we encode the transition relation \curvearrowright :

```
next(state) := case
  state = s0 & input = coin : s1;
  state = s1 & input = select : {s2,s3};
  state = s2 & input = coffee : s0;
  state = s3 & input = tea : s0;
  TRUE : state;
esac;
```

And finally, we specify which propositional atom is true in each state:

```
next(input) := case
  state = s0 : coin;
  state = s1 : select;
  state = s2 : coffee;
  state = s3 : tea;
  TRUE : input;
esac;
```

The entire example is available at <https://www.cs.cmu.edu/~15414/lectures/21-ltl/vending.smv>

We can now specify LTL formulas and check their validity. For instance, the formula $\Box \text{coffee}$ can be written as:

LTLSPEC $G (\text{input} = \text{coffee});$

For which NuSMV would return a counterexample to show that this formula is false:

```
-- specification G input = coffee is false
-- as demonstrated by the following execution sequence
```

Trace Description: LTL Counterexample

Trace Type: Counterexample

```
-- Loop starts here
```

```
-> State: 1.1 <-
```

```
  state = s0
```

```
  input = coin
```

```
-> State: 1.2 <-
```

```
  state = s1
```

```
-> State: 1.3 <-
```

```
  input = select
```

```
-> State: 1.4 <-
```

```
  state = s2
```

```
-> State: 1.5 <-
```

```
  input = coffee
```

```
-> State: 1.6 <-
```

```
  state = s0
```

```
-> State: 1.7 <-
```

`input = coin`

On the other hand, the formula $\Box((coin \wedge \mathbf{X}select) \rightarrow \Diamond(coffee \vee tea))$ can be written as:

LTLSPEC `G (input = coin & X input = select ->
F (input = coffee | input = tea))`

For which the prover will return:

```
-- specification  G ((input = coin &  X input = select) ->  
  F (input = coffee | input = tea))  is true
```

NuSMV can be used to model more complex systems and we will see other examples when we talk about CTL in the next lecture.