Formal Specification and Verification

Temporal logic (1)

10.12.2018

Viorica Sofronie-Stokkermans

e-mail: sofronie@uni-koblenz.de

Formal specification

- Specification for program/system
- Specification for properties of program/system

Verification tasks:

Check that the specification of the program/system has the required properties.

Temporal logic

Motivation

The purpose of temporal logic (TL) is:

- reasoning about time (in philosophy), and
- reasoning about the behaviour of systems evolving over time (in computer science).

How to define a TL?

To define a temporal logic (TL), we need to specify:

- the language for talking about time or temporal systems;
- our model of time.

Motivation

What model of time should we use?

What is the structure of time?

Motivation

What model of time should we use?

What is the structure of time?

A very liberal definition:

A flow of time is a pair $(T, <)$, where T is a non-empty set of time points, and $<$ is an irreflexive and transitive binary relation on T.

Depending on the intended application, we often require additional properties. One of the most fundamental decisions is whether or not time should be linear.

$(T, <)$ is linear if, for all $x, y \in T$ with $x \neq y$, we have $x < y$ or $y < x$.

Models of time

Important additional properties for linear flows of time:

Boundedness: We have four options by combining:

- *Bounded to the past:* there exists an $x \in T$ such that $x \leq y$ for all $y \in T$ (genesis).
- *Bounded to the future:* there exists a an $x \in T$ such that $y \leq x$ for all $y \in T$ (doomsday).

Discreteness: Existence of direct predecessors and successors:

- If $x \in T$ is not genesis, then there exists a $y \in T$ such that $y < x$ and $y < z < x$ holds for no $z \in T$.
- If $x \in T$ is not doomsday, then there exists a $y \in T$ such that $x < y$ and $x < z < y$ holds for no $z \in T$.

It can be seen that one does not follow from the other.

Models of time

Important additional properties for linear flows of time:

Density: For all $x, y \in T$ with $x < y$, there is a $z \in T$ such that $x < z < y$.

Dedekind completeness: Any non-empty subset $S \subseteq T$ that has an upper bound has a least upper bound:

Definitions:

Upper bound for S: $x \in T$ with $y \leq x$ for all $y \in S$;

Least upper bound for S: upper bound x for S such that there is no $x' \in T$ with $x' < x$ and x' upper bound for S.

Models of time

The following are among the most natural linear flows of time:

- **The natural numbers $\mathbb{N}$ with the usual order $<$.**

  Linear, discrete, bounded to the past, not bounded to the future.

  Note that other flows of time have these properties as well:

  $T := \mathbb{N} \times \{0\} \cup \mathbb{Z} \times \{1\}$, where:

  $(x, a) < (y, b)$ if (i) $a < b$ or (ii) $a = b$ and $x < y$.

  NOTE: above example not Dedekind complete.

Models of time

The following are among the most natural linear flows of time:

- **The rational numbers $\mathbb{Q}$.**

  A natural dense flow of time, though with gaps (e.g. π).

  The unique countable linear dense flow of time without endpoints (up to isomorphism).

- **The real numbers $\mathbb{R}$.**

  Up to isomorphism, the unique dense, Dedekind-complete flow of time without end points that is separable:

  There exists a countable subset $D \subseteq T$ such that, for all $x, y \in T$ with $x < y$, there is a $z \in D$ with $t < z < u$.

Models of time

The alternative to linear time is branching time.

Time can be:

- **Branching to the future** reflecting that there are many possible futures;
- **Branching to the past** reflecting that many different histories are considered possible (due to incomplete knowledge).

Branching to the future and linear to the past is the most popular option

for each $x \in T$, the set $\{y \in T \mid y < x\}$ is linearly ordered by $<$.

We can identify additional properties similar to the linear case. Usually, branching time is assumed to be discrete and has a genesis.

Models of time

Which flow of time should we use?

Models of time

Which flow of time should we use?

This depends on the application!

Models of time

Which flow of time should we use?

This depends on the application!

The main application of TL in computer science is the verification of finite-state reactive and concurrent systems.

A state is a snapshot of the system capturing the values of the variables at an instant of time.

- Finite-state systems.
  Finite-state systems can only take finitely many states.
  (Often, infinite-state systems can be abstracted into finite-state ones by grouping the states into a finite number of partitions.)

Models of time

Which flow of time should we use?

This depends on the application!

The main application of TL in computer science is the verification of finite-state reactive and concurrent systems.

A state is a snapshot of the system capturing the values of the variables at an instant of time.

- Reactive Systems.
  A reactive system interacts with the environment frequently and usually does not terminate. Its correctness is defined via these interactions. This is in contrast to a classical algorithm that takes an input initially and then eventually terminates producing a result.

Models of time

Which flow of time should we use?

This depends on the application!

The main application of TL in computer science is the verification of finite-state reactive and concurrent systems.

A state is a snapshot of the system capturing the values of the variables at an instant of time.

- Concurrent Systems.
  Systems consisting of multiple, interacting processes. One process does not know about the internal state of the others. May be viewed as a collection of reactive systems.

Models of time

Which flow of time should we use?

This depends on the application!

The main application of TL in computer science is the verification of finite-state reactive and concurrent systems.

Task: Verificaton.

Given the (formal) description of a system and of its intended behaviour, check whether the system indeed complies with this behaviour.

Transition systems

We use an abstract model of reactive and concurrent systems.

Definition (Transition system, simplified version)

Let Π be a finite set of propositional variables.

A transition system is a tuple $(S, \rightarrow, S_i, L)$ with

- S a non-empty set of states;
- $\rightarrow \subseteq S \times S$ is a transition relation that is total, i.e. for each state $s \in S$, there is a state $s' \in S$ such that $s \rightarrow s'$;
- $S_i \subseteq S$ is a set of initial states;
- $L : S \rightarrow \{0, 1\}^{AP}$ is a valuation function which we will also regard as a function $L : AP \times S \rightarrow \{0, 1\}$

Example

Consider the following simple mutual-exclusion protocol:

```
    task body ProcA is
    begin
    loop
(0) Non_Critical_Section_A;
(1) loop [exit when Turn = 0] end loop;
(2) Critical_Section_A;
(3) Turn := 1;
    end loop;
    end ProcA;

    task body ProcB is
    begin
    loop
(0) Non_Critical_Section_B;
(1) loop [exit when Turn = 1] end loop;
(2) Critical_Section_B;
(3) Turn := 0;
    end loop;
    end ProcA;
```

Assume that the processes run asynchronously, i.e., either Process A or B makes a step, but not both. The order of executions is undetermined.

Example

$$\Pi = \{(T = i) \mid i \in \{0,1\}\} \cup \{(X = i) \mid X \in \{A,B\}, i \in \{0,1,2,3\}\}$$

$(T = i)$ means that Turn is set to i, and

$(X = i)$ means the process X is currently in Line i.

Example

We define the following transition system $(S, \rightarrow, S_i, L)$:

- $S = \{0,1\} \times \{0,1,2,3\} \times \{0,1,2,3\}$

  $(t,i,j) \in S$: state in which Turn $= t$, A is at line i, B is at line j

- $S_i = \{(0,0,0),(1,0,0)\}$

- $\rightarrow = R_A \cup R_B$, where

$$
\begin{aligned}
R_A = \{((t,i,j),(t',i',j)) \mid \quad & (i \in \{0,2,3\} \wedge t' = t) \rightarrow i' = i+1 \ (mod 4), \\
& t = 0, i = 1 \rightarrow i' = 2 \\
& t = 1, i = 1 \rightarrow i' = 1 \\
& i = 3 \rightarrow t' = 1\}
\end{aligned}
$$

  and R_B is defined similarly

- $L((T = t'),(t,i,j)) = 1$ iff $t' = t$
  $L((A = i'),(t,i,j)) = 1$ iff $i' = i$
  $L((B = j'),(t,i,j)) = 1$ iff $j' = j$

Computations

Let $TS = (S, \rightarrow, S_i, L)$ be a transition system.

A computation (or execution) of TS is an infinite sequence $s_0 s_1 \ldots$ of states such that $s_0 \in S_i$ and $s_i \rightarrow s_{i+1}$ for all $i \geq 0$.

Example: computation (execution) of the transition system from the previous example:

$(0,0,0),(0,1,0),(0,1,1),(0,2,1),(0,3,1),(1,0,1),(1,0,2),\ldots$

This corresponds to an (asynchronous) execution of the concurrent system with Processes A and B.

Note that our formalization allows computations that are unfair, e.g., in which Process B is never executed. Such issues are not adressed on the level of transition systems.

Example

Interesting properties that can be verified in this Example include the following:

- **Mutual exclusion:** can A and B be at Line (2) at the same time?

- **Guaranteed accessibility:** if process $X \in \{A, B\}$ is at Line (2), is it guaranteed that it will eventually reach Line (3)?

  (holds, but only in computations that execute both Process A and Process B infinitely often)

Later, we will express such properties as temporal logic formulas.

Computation trees

Transition systems can be non-deterministic, i.e., for an $s \in S$, the set $\{s' \mid s \to s'\}$ can have arbitrary cardinality > 0.

Thus, in general there is more than a single computation.

Instead of considering single computations in isolation, we can arrange all of them in a computation tree.

Informally, for $s \in S_i$, the (infinite) computation tree $T(TS, s)$ of TS at $s \in S$ is inductively constructed as follows:

- use s as the root node;
- for each leaf s' of the tree, add successors $\{t \in S \mid s' \to t\}$.

Computation trees

The computation tree of the transition system from the previous example starting at state (0, 0, 0) is:

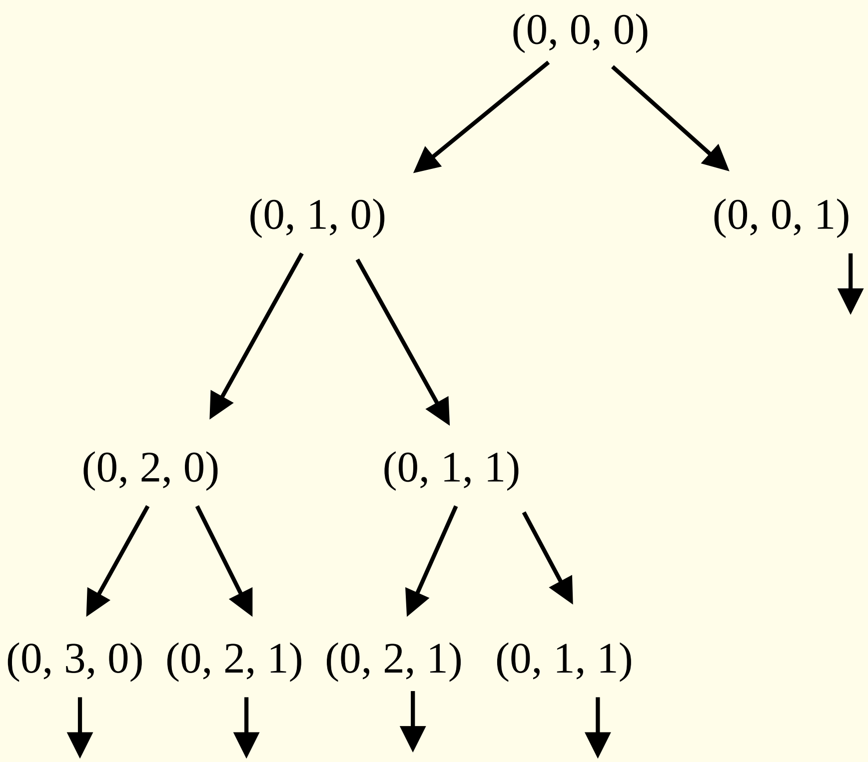

Linear Time Logic

Syntax

Π set of propositional variables.

The set of LTL (linear time logic) formulae is the smallest set such that:

- $\bot, \top$ and each propositional variable $P \in \Pi$ are formulae;
- if F, G are formulae, then so are $F \wedge G, F \vee G, \neg F$;
- if F, G are formulae, then so are $\bigcirc F$ and $F \mathcal{U} G$

Linear Time Logic

Syntax

Π set of propositional variables.

The set of LTL (linear time logic) formulae is the smallest set such that:

- $\bot, \top$ and each propositional variable $P \in \Pi$ are formulae;
- if F, G are formulae, then so are $F \wedge G, F \vee G, \neg F$;
- if F, G are formulae, then so are $\bigcirc F$ and $F \mathcal{U} G$

Remark: Instead of $\bigcirc F$ in some books also XF is used.

Linear Time Logic

Semantics

- Transition systems $(S, \rightarrow, L)$
  (with the property that for every $s \in S$ there exists $s' \in S$ with $s \rightarrow s'$ i.e. no state of the system can "deadlock"[a])

  Transition systems are also simply called models in what follows.

---

[a] This is a technical convenience, and in fact it does not represent any real restriction on the systems we can model. If a system did deadlock, we could always add an extra state s_d representing deadlock, together with new transitions $s \rightarrow s_d$ for each s which was a deadlock in the old system, as well as $s_d \rightarrow s_d$.

Linear Time Logic

Semantics

- Transition systems $(S, \to, L)$
  (with the property that for every $s \in S$ there exists $s' \in S$ with $s \to s'$
  i.e. no state of the system can "deadlock"[a])

  Transition systems are also simply called models in what follows.

- Computation (execution, path) in a model $(S, \to, L)$
  infinite sequence of states $\pi = s_0, s_1, s_2, \ldots$ in S such that for each $i \geq 0$, $s_i \to s_{i+1}$.
  We write the path as $s_0 \to s_1 \to s_2 \to \ldots$.

---

[a] This is a technical convenience, and in fact it does not represent any real restriction on the systems we can model. If a system did deadlock, we could always add an extra state s_d representing deadlock, together with new transitions $s \to s_d$ for each s which was a deadlock in the old system, as well as $s_d \to s_d$.

Linear Time Logic

Consider the path $\pi = s_0 \to s_1 \to ....$

It represents a possible future of our system.

We write π^i for the suffix starting at s_i, e.g.,

$$\pi^3 = s_3 \to s_4 \to ....$$

Linear Time Logic

Semantics

Let $TS = (S, \rightarrow, L)$ be a model and $\pi = s_0 \rightarrow \ldots$ be a path in TS.

Whether π satisfies an LTL formula is defined by the satisfaction relation $\models$ as follows:

- $\pi \models \top$
- $\pi \not\models \bot$
- $\pi \models p$ iff $p \in L(s_0)$,   if $p \in \Pi$
- $\pi \models \neg F$ iff $\pi \not\models F$
- $\pi \models F \wedge G$ iff $\pi \models F$ and $\pi \models G$
- $\pi \models F \vee G$ iff $\pi \models F$ or $\pi \models G$
- $\pi \models \bigcirc F$ iff $\pi^1 \models F$
- $\pi \models F \mathcal{U} G$ iff $\exists m \geq 0$ $s.t.$ $\pi^m \models G$ and $\forall k \in \{0, \ldots, m-1\} : \pi^k \models F$

Linear Time Logic

Alternative way of defining the semantics:

An LTL structure M is an infinite sequence $S_0 S_1 \ldots$ with $S_i \subseteq \Pi$ for all $i \geq 0$. We define satisfaction of LTL formulas in M at time points $n \in \mathbb{N}$ as follows:

- $M, n \models p$ iff $p \in S_n, \qquad$ if $p \in \Pi$
- $M, n \models F \wedge G$ iff $M, n \models F$ and $M, n \models G$
- $M, n \models F \vee G$ iff $M, n \models F$ or $M, n \models G$
- $M, n \models \neg F$ iff $M, n \not\models F$
- $M, n \models \bigcirc F$ iff $M, n+1 \models F$
- $M, n \models F \mathcal{U} G$ iff $\exists m \geq n \; s.t. \; M, m \models G$ and $\forall k \in \{n, \ldots, m-1\} : M, k \models F$

Note that the time flow $(\mathbb{N}, <)$ is implicit.

Transition systems and LTL models

The connection between transition systems and LTL structures is as follows:

Every computation (evolution, path) of a transition system $s_0 \rightarrow s_1 \ldots$ gives rise to an LTL structure.

To see this, let $TS = (S, \rightarrow, L)$ be a transition system.

A computation $s_0, s_1, ...$ of TS induces an LTL structure $L(s_0)L(s_1)\ldots$

Such an LTL structure is called a trace of TS.

Abbreviations

- The future diamond

  $\Diamond\phi := \top\mathcal{U}\phi$

  $\pi \models \Diamond\phi$ iff $\exists m \geq 0 : \pi^m \models \phi$

- The future box

  $\Box\phi := \neg\Diamond\neg\phi$

  $\pi \models \Box\phi$ iff $\forall m \geq 0 : \pi^m \models \phi$

Abbreviations

- The future diamond

  $\Diamond\phi := \top\mathcal{U}\phi$ Sometimes denoted also $F\phi$

  $\pi \models \Diamond\phi$ iff $\exists m \geq 0 : \pi^m \models \phi$ $M, n \models \Diamond\phi$ iff $\exists m \geq n : M, m \models \phi$

- The future box

  $\Box\phi := \neg\Diamond\neg\phi$ Sometimes also denoted $G\phi$

  $\pi \models \Box\phi$ iff $\forall m{\geq}0 : \pi^m \models \phi$ $M, n \models \Box\phi$ iff $\forall m{\geq}n : M, m \models \phi$

Abbreviations

- The infinitely often operator

  $\Diamond^{\infty}\phi := \Box\Diamond\phi$

  $\pi \models \Diamond^{\infty}\phi$ iff $\{m \geq 0 \mid \pi^m \models \phi\}$ is infinite

  $M, n \models \Diamond^{\infty}\phi$ iff $\{m \geq n \mid M, m \models \phi\}$ is infinite

- The almost everywhere operator

  $\Box^{\infty}\phi := \Diamond\Box\phi$

  $\pi \models \Box^{\infty}\phi$ iff $\{m \geq 0 \mid \pi^m \not\models \phi\}$ is finite.

  $M, n \models \Box^{\infty}\phi$ iff $\{m \geq n \mid M, m \not\models \phi\}$ is finite.

Abbreviations

- The release operator
  $\phi\mathcal{R}\psi := \neg(\neg\phi\mathcal{U}\neg\psi)$

  $\pi \models \phi\mathcal{R}\psi$ iff ($\exists m \geq 0 : \pi^m \models \phi$ and $\forall k \leq m$: $\pi^k \models \psi$) or
  ($\forall k \geq 0 : \pi^k \models \psi$)

  $M, n \models \phi\mathcal{R}\psi$ iff ($\exists m \geq n : M, m \models \phi$ and $\forall k \leq m : M, m \models \psi$) or
  ($\forall k \geq m : M, k \models \psi$)

  Read as
  "ψ always holds unless released by ϕ" i.e.,
  "ψ holds permanently up to and including the first point where ϕ holds (such an ϕ-point need not exist at all)".

Abbreviations

- The strict until operator:
  $F\mathcal{U}^{<}G := \bigcirc(F\mathcal{U}G)$

  $\pi \models F\mathcal{U}^{<}G$ iff $\exists m > 0 : \pi^m \models G \wedge \forall k \in \{1, 2, \ldots, m-1\}, \pi^k \models F$

  $M, n \models F\mathcal{U}^{<}G$ iff $\exists m > n : M, m \models G \wedge \forall k \in \{n+1, ..., m-1\}, M, k \models F$

The difference between standard and strict until is that strict until requires G to happen in the strict future and that F needs not hold true of the current point.

Equivalence

We say that two LTL formulas F and G are (globally) equivalent
(written $F \equiv G$)
if, for all LTL structures M and $i \geq 0$, we have $M, i \models F$ iff $M, i \models G$.

equivalently:

if for all transition systems T and all paths π in T we have:
$\pi \models F$ iff $\pi \models G$.

Note that:

$\bigcirc F \equiv \bot\, \mathcal{U}^{<} F$ and

$F \mathcal{U} G \equiv G \vee (F \wedge (F \mathcal{U}^{<} G))$

Thus, an equally expressive version of LTL is obtained by using $\mathcal{U}^{<}$ as the only temporal operator.

This cannot be done with the standard until

Equivalence

Some useful equivalences that will be useful later on (exercise: prove them):

$\neg \bigcirc F \equiv \bigcirc \neg F$	(self-duality of next)
$\Diamond\Diamond F \equiv \Diamond F$	(idempotency of diamond)
$\bigcirc \Diamond F \equiv \Diamond \bigcirc F$	(commutation of next with Diamond)
$\Diamond\Diamond^{\infty} F \equiv \Diamond^{\infty} F \equiv \Diamond^{\infty} \Diamond F$	(absorption of diamonds by "infinitely ofte
$F \mathcal{U} G \equiv \neg(\neg F \mathcal{R} \neg G)$	(until and release are duals)
$F \mathcal{U} G \equiv G \vee (F \wedge \bigcirc(F \mathcal{U} G))$	(unfolding of until)
$F \mathcal{R} G \equiv (F \wedge G) \vee (G \wedge \bigcirc(F \mathcal{R} G))$	(unfolding of release)

Temporal Properties

A temporal property is a set of LTL structures
(those on which the property is true).

Thus, a temporal property P can be defined using an LTL formula F:

$$P = \{M \mid M, 0 \models F\}.$$

When given a transition system TS representing a reactive system and an LTL formula F representing a temporal property,

TS satisfies F if $M, 0 \models F$ for all traces M of TS.

In this case, we write $TS \models F$.

Typical properties of reactive systems that need to be checked during verification are safety properties, liveness properties, and fairness properties.

Safety properties

Intuitively, a safety property asserts that "nothing bad happens"

general form: Condition $\rightarrow \Box F_{\text{Safe}}$

Examples of safety properties:

- **Mutual Exclusion.** For the example:

$$\Box(\neg((A = 2) \wedge (B = 2)))$$

- **Freedom from Deadlocks:** At any time, some process should be enabled:

$$\Box(enabled_1 \vee \cdots \vee enabled_k)$$

- **Partial Correctness:** If F is satisfied when the program starts, then G will be satisfied if the program reaches a distinguished state:

$$F \rightarrow \Box(\text{Dist} \rightarrow G)$$

where Dist $\in \Pi$ marks the distinguished state.

Liveness properties

Intuitively, a liveness property asserts that "something good will happen"

Examples of liveness properties:

- **Guaranteed Accessibility.** For the example:

$$\Box(A = 1 \rightarrow \Diamond(A = 2)) \wedge \Box(B = 1 \rightarrow \Diamond(B = 2))$$

- **Responsiveness:** If a request is issued, it will eventually be granted:

$$\Box(\mathrm{req} \rightarrow \Diamond \mathrm{grant})$$

- **Total Correctness:** If F is satisfied when the program starts, then the program terminates in a distinguished state where G is satisfied:

$$\phi \rightarrow \Diamond(\mathrm{Dist} \wedge G)$$

Note that, in contrast, partial correctness is a safety property.

Fairness properties

When modelling concurrent systems, it is usually important to make some fairness assumptions. Assume that there are k processes, that $\mathsf{enabled}_i \in \Pi$ is true in a state s if process $\#i$ is enabled in s for execution, and that $\mathsf{executed}_i$ is true in a state s if process $\#i$ has been executed to reach s.

Examples of fairness properties

- **Unconditional Fairness:** Every process is executed infinitely often:

  $$\bigwedge_{1 \leq i \leq k} \Diamond^{\infty} \mathsf{executed}_i$$

  Unconditional fairness is appropriate when processes can (and should!) be executed and any time. This is not always the case.

Fairness properties

When modelling concurrent systems, it is usually important to make some fairness assumptions. Assume that there are k processes, that $\mathsf{enabled}_i \in \Pi$ is true in a state s if process $\#i$ is enabled in s for execution, and that $\mathsf{executed}_i$ is true in a state s if process $\#i$ has been executed to reach s.

Examples of fairness properties

- **Strong Fairness:** Every process enabled infinitely often is executed infinitely often:

  $$\bigwedge_{1 \le i \le k} (\Diamond^{\infty}\mathsf{enabled}_i \rightarrow \Diamond^{\infty}\mathsf{executed}_i)$$

  Processes enabled only finitely often need not be guaranteed to be executed: they eventually and forever retract being enabled.

Fairness properties

When modelling concurrent systems, it is usually important to make some fairness assumptions. Assume that there are k processes, that $\text{enabled}_i \in \Pi$ is true in a state s if process $\#i$ is enabled in s for execution, and that executed_i is true in a state s if process $\#i$ has been executed to reach s.

Examples of fairness properties

- **Weak Fairness:** Every process enabled almost everywhere is executed infinitely often.

  $$\bigwedge_{1 \leq i \leq k} (\Box^{\infty}\text{enabled}_i \rightarrow \Diamond^{\infty}\text{executed}_i)$$

  This means that a process cannot be enabled constantly in an infinite interval without being executed in this interval.