Linear Temporal Logic

Moonzoo Kim
CS Dept. KAIST

Review: Model checking

- ## Model checking
    - In a model-based approach, the system is represented by a model $\mathcal{M}$. The specification is again represented by a formula ϕ.
        - The verification consists of computing whether $\mathcal{M}$ satisfies ϕ $\mathcal{M} \vDash \phi$
            - Caution: $\mathcal{M} \vDash \phi$ represents satisfaction, not semantic entailment

- ## In model checking,
    - The model $\mathcal{M}$ is a transition systems and
    - the property ϕ is a formula in temporal logic
        - ex. $\square$ p, $\square$ q, $\Diamond$ q, $\square \Diamond$ q

Motivation for Temporal Logic

- So far, we have analyzed sequential programs only
  - `assert` is a convenient way of specify requirement properties
  - Safety properties are enough for sequential programs
    - "Bad thing never happens"
    - Ex. Mutual exclusion
- For concurrent programs, we need more than `assert` to specify important requirement properties conveniently
  - Liveness properties
    - "Good thing eventually happens"
    - Ex. Deadlock freedom
    - Ex. Starvation freedom
- Temporal logic is an adequate logic for describing requirement properties for concurrent system

Motivating Example (1/2)

- **Mutual exclusion protocol**
  - Alice and Bob are neighbors, and they share a yard.
  - Alice owns a cat and Bob owns a dog.
  - Alice and Bob should coordinate that both pets are never in the yard at the same time.

- We would like to design a mutual exclusion protocol to satisfy
  1. Mutual exclusion
     - pets are excluded from being in the yard at the same time
  2. Deadlock-freedom
     - Both pets want to enter the yard, then eventually at least one of them succeeds
  3. Starvation-freedom/lock-out freedom
   - If a pet wants to enter the yard, it will eventually succeed

Motivating Example (2/2)

- **One protocol design: Alice and Bob set up a flag pole at each house**
  - **Protocol @ Alice**
    1. Alice raises her flag
    2. When Bob's flag is lowered, she unleashes her cat
    3. When her cat comes back, she lowers her flag
  - **Protocol @ Bob**
    1. He raises his flag
    2. While Alice's flag is raised
       1. Bob lowers his flag
       2. Bob waits until Alice's flag is lowered
       3. Bob raises his flag
    3. As soon as his flag is raised and hers is down, he unleashes his dog
    4. When his dog comes back, he lowers his flag

Linear time temporal logic (LTL)

- **LTL models time as a sequence of states, extending infinitely into the future**
  - sometimes a sequence of states is called a computation path or an execution path, or simply a path

$$\boxed{\text{F p} \rightarrow \text{G r} \lor \neg \text{q U p}}$$

- **Def 3.1 LTL has the following syntax**
  - $\phi ::= \text{T} \mid \bot \mid \text{p} \mid \neg \phi \mid \phi \land \phi \mid \phi \lor \phi \mid \phi \rightarrow \phi$
    $\mid \text{X} \phi \mid \text{F} \phi \mid \text{G} \phi \mid \phi \text{U} \phi \mid \phi \text{W} \phi \mid \phi \text{R} \phi$

    where p is any propositional atom from some set Atoms
  - Operator precedence
    - the unary connectives bind most tightly. Next in the order come U, R, W, $\land$, $\lor$, and $\rightarrow$
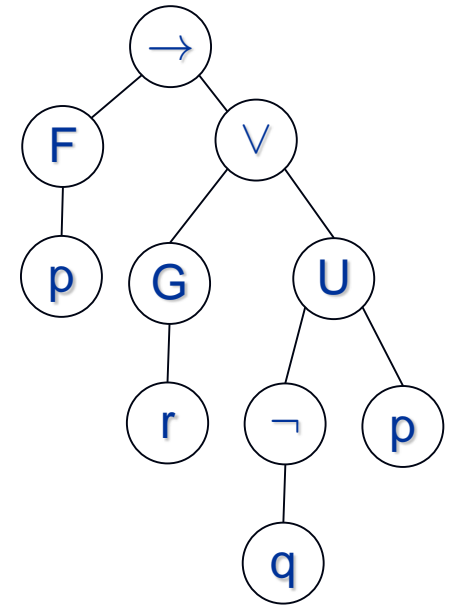
Semantics of LTL (1/3)

- Def 3.4 A transition system (called model) $\mathcal{M}$ = (S, $\rightarrow$, L)
  - a set of states S
  - a transition relation $\rightarrow$ (a binary relation on S)
    - such that every s $\in$ S has some s' $\in$ S with s $\rightarrow$ s'
  - a labeling function L: S $\rightarrow$ $\mathcal{P}$ (Atoms)
- Example
  - S={s_0,s_1,s_2}
  - $\rightarrow$={(s_0,s_1),(s_1,s_0),(s_1,s_2),(s_0,s_2),(s_2,s_2)}
  - L={(s_0,{p,q}),(s_1,{q,r}), (s_2,{r})}
- Def. 3.5 A path in a model $\mathcal{M}$ = (S, $\rightarrow$, L) is an infinite sequence of states s_{i_1}, s_{i_2}, s_{i_3},… in S s.t. for each j$\geq$ 1, $s_{i_j} \rightarrow s_{i_{j+1}}$. We write the path as $s_{i_1} \rightarrow s_{i_2} \rightarrow$ …
  - From now on if there is no confusion, we drop the subscript index i for the sake of simple description
- We write π^i for the suffix of a path starting at s_i.
  - ex. π^3 is $s_3 \rightarrow s_4 \rightarrow$ …

Semantics of LTL (2/3)

- Def 3.6 Let $\mathcal{M} = (S, \rightarrow, L)$ be a model and $\pi = s_1 \rightarrow \ldots$ be a path in $\mathcal{M}$. Whether π satisfies an LTL formula is defined by the satisfaction relation $\vDash$ as follows:
    - Basics: $\pi \vDash \top$, $\pi \nvDash \bot$, $\pi \vDash p$ iff $p \in L(s_1)$, $\pi \vDash \neg\phi$ iff $\pi \nvDash \phi$
    - Boolean operators: $\pi \vDash p \wedge q$ iff $\pi \vDash p$ and $\pi \vDash q$
        - similar for other boolean binary operators
    - $\pi \vDash X \phi$ iff $\pi^2 \vDash \phi$ (next $\bigcirc$)
    - $\pi \vDash G \phi$ iff for all $i \geq 1$, $\pi^i \vDash \phi$ (always $\square$)
    - $\pi \vDash F \phi$ iff there is some $i \geq 1$, $\pi^i \vDash \phi$ (eventually $\Diamond$)
    - $\pi \vDash \phi U \psi$ iff there is some $i \geq 1$ s.t. $\pi^i \vDash \psi$ and for all $j=1,\ldots,i-1$ we have $\pi^j \vDash \phi$ (strong until)
    - $\pi \vDash \phi W \psi$ iff either (weak until)
        - either there is some $i \geq 1$ s.t. $\pi^i \vDash \psi$ and for all $j=1,\ldots,i-1$ we have $\pi^j \vDash \phi$
        - or for all $k \geq 1$ we have $\pi^k \vDash \phi$
    - $\pi \vDash \phi R \psi$ iff either (release)
        - either there is some $i \geq 1$ s.t. $\pi^i \vDash \phi$ and for all $j=1,\ldots,i$ we have $\pi^j \vDash \psi$
        - or for all $k \geq 1$ we have $\pi^k \vDash \psi$

examples

[]p is satified at all locations in σ

<>p is satisfied at all locations in σ

[]<>p is satisfied at all locations in σ

<>q is satisfied at all locations except s_{n-1} and s_n

Xq is satisfied at s_{i+1} and at s_{i+3}

pUq (strong until) is satisfied at all locations except s_{n-1} and s_n

<>(pUq) (strong until) is satisfied at all locations except s_{n-1} and s_n

<>(pUq) (weak until) is satisfied at all locations

[]<>(pUq) (weak until) is satisfied at all locations

in model checking we are typically only
interested in whether a temporal logic formula
is satisfied for all runs of the system, starting
in the initial system state (that is: at s_0)

slide quoted from Caltech 101b.2 "Logic Model Checking" by Dr.G.Holzmann

visualizing LTL formulae

slide quoted from Caltech 101b.2 "Logic Model Checking" by Dr.G.Holzmann

slide quoted from Caltech 101b.2 "Logic Model Checking" by Dr.G.Holzmann

another example

LTL: (<>b1) -> (<>b2)

time ⟶

1. b1 never becomes true

formula satisfied

!b1

2. b1 and b2 both become true

formula satisfied

b2

b1

!b2

!b1

3. b1 becomes true but not b2

formula *not* satisfied
the property is *violated*

b1

!b2

!b1

Semantics of LTL (3/3)

- Def 3.8 Suppose $\mathcal{M}$ = (S, $\rightarrow$, L) is a model, s $\in$ S, and ϕ an LTL formula. We write $\mathcal{M},s \vDash \phi$ if for every execution path π of $\mathcal{M}$ starting at s, we have $\pi \vDash \phi$
  - If $\mathcal{M}$ is clear from the context, we write $s \vDash \phi$

- Example

  - $s_0 \vDash p \wedge q$ since $\pi \vDash p \wedge q$ for every path π beginning in s_0
  - $s_0 \vDash \neg r$, $s_0 \vDash \top$
  - $s_0 \vDash X\ r$, $s_0 \nvDash X\ (q \wedge r)$
  - $s_0 \vDash G\ \neg(p \wedge r)$, $s_2 \vDash G\ r$
  - For any s of $\mathcal{M}$, $s \vDash F(\neg q \wedge r) \rightarrow F\ G\ r$
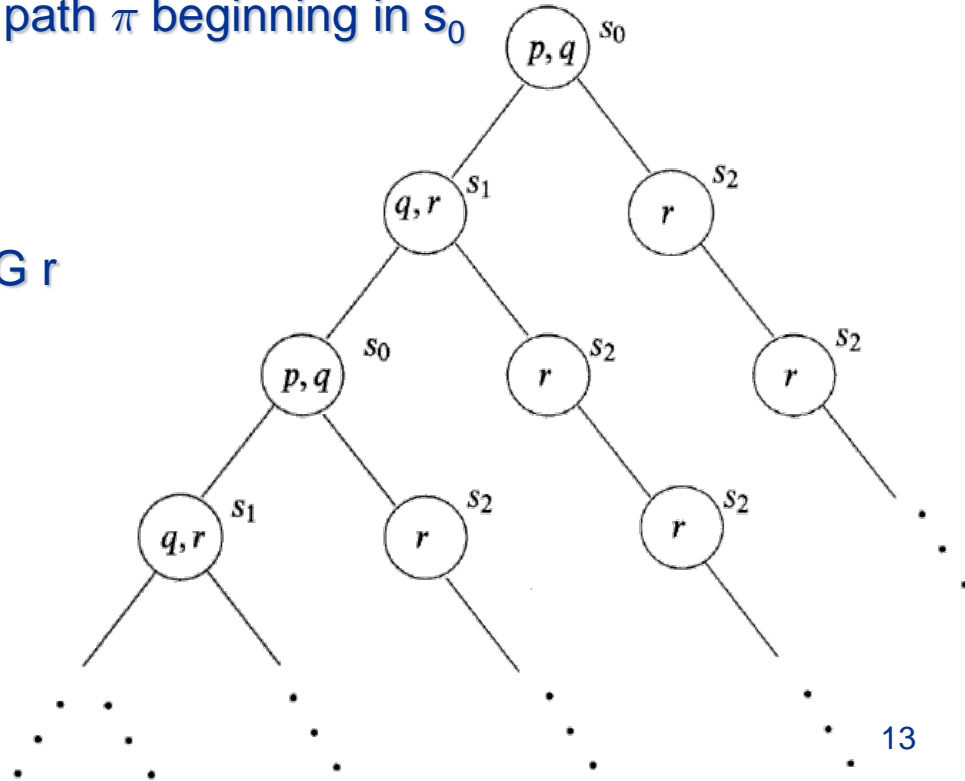   - Note that s_2 satisfies $\neg q \wedge r$
  - $s_0 \nvDash G\ F\ p$
    - $s_0 \rightarrow s_1 \rightarrow s_0 \rightarrow s_1 \dots \vDash G\ F\ p$
    - $s_0 \rightarrow s_2 \rightarrow s_2 \rightarrow s_2 \dots \nvDash G\ F\ p$
  - $s_0 \vDash G\ F\ p \rightarrow G\ F\ r$
  - $s_0 \nvDash G\ F\ r \rightarrow G\ F\ p$

Practical patterns of specification

- For any state, if a request occurs, then it will eventually be acknowledge
  - G(requested $\rightarrow$ F acknowledged)
- A certain process is enabled infinitely often on every computation path
  - G F enabled
- Whatever happens, a certain process will eventually be permanently deadlocked
  - F G deadlock
- If the process is enabled infinitely often, then it runs infinitely often
  - G F enabled $\rightarrow$ G F running
- An upwards traveling lift at the second floor does not change its direction when it has passengers wishing to go to the fifth floor
  - G (fllor2 $\wedge$ directionup $\wedge$ ButtonPressed5 $\rightarrow$ (directionup U floor5)

- It is impossible to get to a state where a system has started but is not ready
  - ϕ = G ¬(started $\wedge$ ¬ready)
  - What is the meaning of (intuitive) negation of ϕ ?
    - For every path, it is possible to get to such a state (started$\wedge$¬ready).
    - There exists a such path that gets to such a state.
      - we cannot express this meaning directly
- LTL has limited expressive power
  - For example, LTL cannot express statements which assert the existence of a path
    - From any state s, there exists a path π starting from s to get to a restart state
    - The lift can remain idle on the third floor with its doors closed
  - Computation Tree Logic (CTL) has operators for quantifying over paths and can express these properties

Summary of practical patterns

G p	always p	invariance
F p	eventually p	guarantee
p → (F q)	p implies eventually q	response
p → (q U r)	p implies q until r	precedence
G F p	always, eventually p	recurrence (progress)
F G p	eventually, always p	stability (non-progress)
F p → F q	eventually p implies eventually q	correlation

Equivalences between LTL formulas

- Def 3.9 $\phi \equiv \psi$ if for all models $\mathcal{M}$ and all paths π in $\mathcal{M}$: $\pi \vDash \phi$ iff $\pi \vDash \psi$

- $\neg G\ \phi \equiv F\ \neg\phi$, $\neg F\ \phi \equiv G\ \neg\phi$, $\neg X\ \phi \equiv X\ \neg\phi$

- $\neg\ (\phi\ U\ \psi) \equiv \neg\phi\ R\ \neg\psi$, $\neg(\phi\ R\ \psi) \equiv \neg\phi\ U\ \neg\psi$

- $F\ (\phi \vee \psi) \equiv F\ \phi \vee F\ \psi$

- $G\ (\phi \wedge \psi) \equiv G\ \phi \wedge G\ \psi$

- $F\ \phi \equiv T\ U\ \phi$, $G\ \phi \equiv \bot\ R\ \phi$

- $\phi\ U\ \psi \equiv \phi\ W\ \psi \wedge F\ \psi$

- $\phi\ W\ \psi \equiv \phi\ U\ \psi \vee G\ \phi$

- $\phi\ W\ \psi \equiv \psi\ R\ (\phi \vee \psi)$

- $\phi\ R\ \psi \equiv \psi\ W\ (\phi \wedge \psi)$

Adequate sets of connectives for LTL (1/2)

- X is completely orthogonal to the other connectives
    - X does not help in defining any of the other connectives.
    - The other way is neither possible
- Each of the sets {U,X}, {R,x}, {W,X} is adequate
    - {U,X}
        - ϕ R $\psi \equiv \neg (\neg \phi$ U $\neg \psi)$
        - ϕ W $\psi \equiv \psi$ R $(\phi \vee \psi) \equiv \neg (\neg\psi$ U $\neg(\phi \vee \psi))$
    - {R,X}
        - ϕ U $\psi \equiv \neg (\neg\phi$ R $\neg\psi)$
        - ϕ W $\psi \equiv \psi$ R $(\phi \vee \psi)$
    - {W,X}
        - ϕ U $\psi \equiv \neg (\neg \phi$ R $\neg \psi)$
        - ϕ R $\psi \equiv \psi$ W $(\phi \wedge \psi)$

Adequate sets of connectives for LTL (2/2)

- Thm 4.10 $\phi \cup \psi \equiv \neg(\neg\psi \cup (\neg\phi \wedge \neg\psi)) \wedge F \psi$
- Proof: take any path $s_0 \to s_1 \to s_2 \to \ldots$ in any model
  - Suppose $s_0 \vDash \phi \cup \psi$
    - Let **n** be the smallest number s.t. $s_n \vDash \psi$
      - We know that such n exists from $\phi \cup \psi$. Thus, $s_0 \vDash F \psi$
      - For each $k < n$, $s_k \vDash \phi$ since $\phi \cup \psi$
    - We need to show $s_0 \vDash \neg(\neg\psi \cup (\neg\phi \wedge \neg\psi))$
      - case 1: for all i, $s_i \nvDash \neg\phi \wedge \neg\psi$. Then, $s_0 \vDash \neg(\neg\psi \cup (\neg\phi \wedge \neg\psi))$
      - case 2: for some i, $s_i \vDash \neg\phi \wedge \neg\psi$. Then, we need to show
        - (*)for each $i > 0$, if $s_i \vDash \neg\phi \wedge \neg\psi$, then there is some $j < i$ with $s_j \nvDash \neg\psi$ (i.e. $s_j \vDash \psi$)
        - Take any $i > 0$ with $s_i \vDash \neg\phi \wedge \neg\psi$. We know that $i > n$ since $s_0 \vDash \phi \cup \psi$. So we can take j=n and have $s_j \vDash \psi$
  - Conversely, suppose $s_0 \vDash \neg(\neg\psi \cup (\neg\phi \wedge \neg\psi)) \wedge F \psi$
    - Since $s_o \vDash F \psi$, we have a minimal **n** as before s.t. $s_n \vDash \psi$
      - case 1: for all i, $s_i \nvDash \neg\phi \wedge \neg\psi$ (i.e. $s_i \vDash \phi \vee \psi$). Then $s_0 \vDash \phi \cup \psi$
      - case 2: for some i, $s_i \vDash \neg\phi \wedge \neg\psi$. We need to prove for any $i < n$, $s_i \vDash \phi$
        - Suppose $s_i \nvDash \phi$ (i.e., $s_i \vDash \neg\phi$). Since n is minimal, we know $s_i \vDash \neg\psi$. So by (*) there is some $j < i < n$ with $s_j \vDash \psi$, contradicting the minimality of n. Contradiction

18

Mutual exclusion example

- When concurrent processes share a resource, it may be necessary to ensure that they do not have access to the common resource at the same time
  - We need to build a protocol which allows only one process to enter critical section
- Requirement properties
  - Safety:
    - Only one process is in its critical section at anytime
  - Liveness:
    - Whenever any process requests to enter its critical section, it will eventually be permitted to do so
  - Non-blocking:
    - A process can always request to enter its critical section
  - No strict sequencing:
   - processes need not enter their critical section in strict sequence

- We model two processes
  - each of which is in
    - non-critical state (n) or
    - trying to enter its critical state (t) or
    - critical section (c)
  - No self edges
- each process executes like $n \rightarrow t \rightarrow c \rightarrow n \rightarrow \ldots$
  - but the two processes interleave with each other
    - only one of the two processes can make a transition at a time (asynchronous interleaving)

1st model for mutual exclusion

- Safety: $s_0 \vDash G \neg (c_1 \wedge c_2)$
- Liveness $s_0 \nvDash G(t_1 \to F c_1)$
  - see $s_0 \to s_1 \to s_3 \to s_7 \to s_1 \to s_3 \to s_7 \dots$
- Non-blocking
  - for every state satisfying n_i,
    there is a successor satisfying t_i
    - s_0 satisfies this property
  - We cannot express this property
    in LTL but in CTL
    - Note that LTL specifies that ϕ is satisfied for all paths
- No strict ordering
  - there is a path where c_1 and c_2 do not occur in strict order
  - Complement of this is
    - $G(c_1 \to c_1 \ W \ (\neg c_1 \wedge \underline{\neg c_1 \ W \ c_2}))$
    - anytime we get into a c_1 state, either that condition persists indefinitely, or it ends
      with a non-c_1 state and in that case there is no further c_1 state unless and until we
      obtain a $\underline{c_2}$ state

2nd model for mutual exclusion

- **All 4 properties are satisfied**
  - Safety
  - Liveness
  - Non-blocking
  - No strict sequencing