A Tour of Temporal Logic

Sean Watters

MSP Group, Strathclyde

10/11/2021

Model Checking

The Goal

Validate specifications of dynamical systems.

Our Approach

1. Model the behaviour of the system.
2. Express the specification as a formula in some logic.
3. Check the validity of the formula in the model.

Constraints

- Model checking should be decidable (at minimum!). Ideally, it should even be *fast*.
- The logic should be expressive — we want to reason about ongoing behaviour (which may never terminate).

Example Properties

Safety

Something bad never happens.

Liveness

Something good keeps happening. Or, more precisely:

At some point in the future, a good thing will happen. At some point after that, another good thing will happen, and so on.

Wednesday Evening

At some point before I next fall asleep, I will have a beer. After I have a beer, I will be happy until I fall asleep. And by the way, I will actually fall asleep eventually.

Normal Modal Logic: Kripke Models

Definition

A Kripke model, $\mathcal{M} = \langle\ S,\ R,\ V\ \rangle$, where:

- S, a set of states.
- $R \subseteq S \times S$, an edge relation.
- $V : S \rightarrow \mathcal{P}(At)$, a valuation function for a countable set of propositional atoms, At.

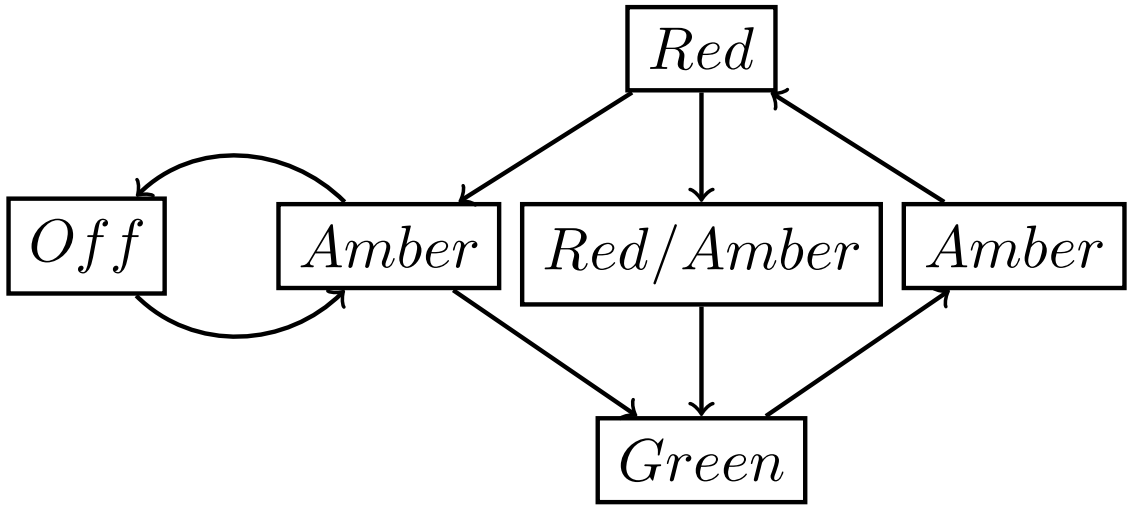

Normal Modal Logic: Syntax and Semantics

Syntax

$$\varphi \;:=\; \top \mid \bot \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \Box\varphi \mid \Diamond\varphi$$

Semantics of $\Box$ and $\Diamond$

Recall: $\mathcal{M} = \langle\, S, \quad R \subseteq S \times S, \quad V : S \to \mathcal{P}(At) \,\rangle$

$$[\![\Box\varphi]\!] \;=\; \{s \in S \mid \forall s' \in S.\ (s, s') \in R \to s' \in [\![\varphi]\!]\}$$

$$[\![\Diamond\varphi]\!] \;=\; \{s \in S \mid \exists s' \in S.\ (s, s') \in R \wedge s' \in [\![\varphi]\!]\}$$

Uh-oh!

This only allows us to look finitely many steps into the future, and each step adds to the length of the formula.

(Non-)Example Properties

Safety

Something bad never happens.

Something bad, φ. Then:

- φ is false at the current state: $\neg\varphi$
- φ is false at every successor: $\Box\neg\varphi$
- φ is false at every successor of every successor: $\Box\Box\neg\varphi$
- etc...

"Liveness"

Something good eventually happens.

Dually: something good, φ. Then:

$$\varphi \vee \Diamond\varphi \vee \Diamond\Diamond\varphi \vee \cdots$$

Introducing: Temporal Logic!

Idea

We're really trying to reason about things happening over time, so let's look at some other modal logics which do exactly that — temporal logics!

Two Approaches

- Branching-time: Models capture all possible execution paths. Formulas quantify over paths. *e.g.* CTL, CTL*, μ-calculus.
- Linear-time: Model is a single possible execution path. Formulas relate only to that path. Formulas then commonly checked against all possible paths. *e.g.* LTL.

Linear-Time Temporal Logic (LTL): Models

Definition

For a set of atoms At, an LTL model is an ω-word on the alphabet $\mathcal{P}(At)$.

Remarks

- An LTL model can be viewed as a single, infinite path through a Kripke model.
- LTL can be adapted to be interpreted over finite paths (LTL_f), but we'll only look at the standard semantics here.

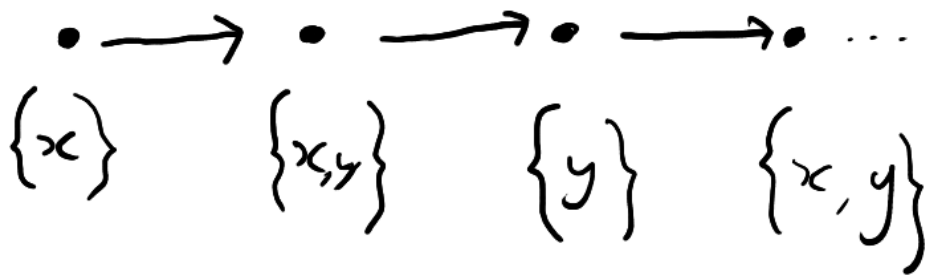

LTL: Syntax and Semantics

Syntax

$$\varphi := \top \mid \bot \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U} \varphi$$

Primitive Temporal Operators

Interpret formulas on an infinite word $w^i = a_i a_{i+1} \cdots$

- $\mathbf{X}\varphi$ — φ is true at the **next** state.

$$w^i \vDash \mathbf{X}\varphi \iff w^{i+1} \vDash \varphi$$

- $\varphi \mathbf{U} \psi$ — φ is true **until** (not necessarily including) ψ becomes true, and ψ must become true in finitely many steps.

$$w^i \vDash \varphi \mathbf{U} \psi \iff \exists n \geq i.\ (w^n \vDash \psi) \wedge (\forall m.\ i \leq m < n \rightarrow w^m \vDash \varphi)$$

LTL: Unary Temporal Operators

Definition

- $\mathbf{F}\varphi$ — φ becomes true at some point in the **future** (in finitely many steps).

$$\mathbf{F}\varphi = \top \mathbf{U} \varphi$$

Definition

- $\mathbf{G}\varphi$ — φ is true **globally** (ie, at all states, forever). It's just the dual of $\mathbf{F}$.

$$\mathbf{G}\varphi = \neg \mathbf{F} \neg \varphi$$

Remark

It might be tempting to define $\mathbf{G}\varphi = \varphi \mathbf{U} \bot$, but something goes wrong...

LTL: Strong and Weak Binary Operators

Strong Until, $\varphi \mathbf{U} \psi$

φ holds until ψ holds, and ψ must eventually hold.

Weak Release, $\varphi \mathbf{R} \psi \ = \ \neg(\neg\varphi \mathbf{U} \neg\psi)$

ψ holds until (and including) the point where φ first holds. If φ never becomes true, then ψ holds forever.

Weak Until, $\varphi \mathbf{W} \psi \ = \ \varphi \mathbf{U} (\psi \vee \mathbf{G}\varphi) \ = \ \psi \mathbf{R} (\varphi \vee \psi)$

φ holds until ψ holds. If ψ never becomes true, φ will hold forever.

$$\mathbf{G}\varphi = \varphi \mathbf{W} \bot$$

Strong Release, $\varphi \mathbf{M} \psi \ = \ \neg(\neg\varphi \mathbf{W} \neg\psi) \ = \ \psi \mathbf{U} (\varphi \wedge \psi)$

ψ holds until (and including) the point where φ first holds, and φ must eventually hold.

LTL: Example Properties

Safety

Something bad never happens: $\mathbf{G}(\neg\varphi)$

Liveness

Something good keeps happening: $\mathbf{GF}\varphi$

Wednesday Evening, just for fun

- I will drink beer some time before I sleep:
  $\neg\ asleep\ \mathbf{U}\ beer$
- I will fall asleep some time after drinking beer:
  $(\neg\ asleep\ \mathbf{U}\ beer) \wedge \mathbf{F}\ asleep$
- Between drinking beer and falling asleep, I will be happy:
  $\neg\ asleep\ \mathbf{U}(beer\ \wedge (happy\ \mathbf{U}\ asleep))$

Interpreting LTL on Kripke Models

Definition

We call a Kripke model *serial* if $\forall s \in S.\ \exists t \in S.\ (s, t) \in R$. That is, every state has at least one successor. Therefore, the model cannot contain terminal states, and all paths through the model must be ω-paths.

LTL on Serial Kripke Models

Consider a pair of a serial Kripke model and a state in that model, $(\mathcal{M}, s)$, and an LTL formula φ. Let:

$$(\mathcal{M}, s) \vDash \varphi \iff \forall p \in paths(\mathcal{M}, s).\ p \vDash \varphi$$

...where $paths$ is a function returning all ω-paths in $\mathcal{M}$ starting at s, mapped with the valuation function s.t. they satisfy the definition of an LTL model.

Computation Tree Logic (CTL): Models

And now, branching time.

- Interpreting LTL formulas on Kripke models in this manner does not fully exploit the structure of the model.
- Instead of quantifying the formula over all paths, let's now look at a *branching-time* temporal logic, which can also do existential quantification over paths.

Recall

A (serial) Kripke model: $\mathcal{M} = \langle\ S, \quad R \subseteq S \times S, \quad V : S \to \mathcal{P}(At)\ \rangle$

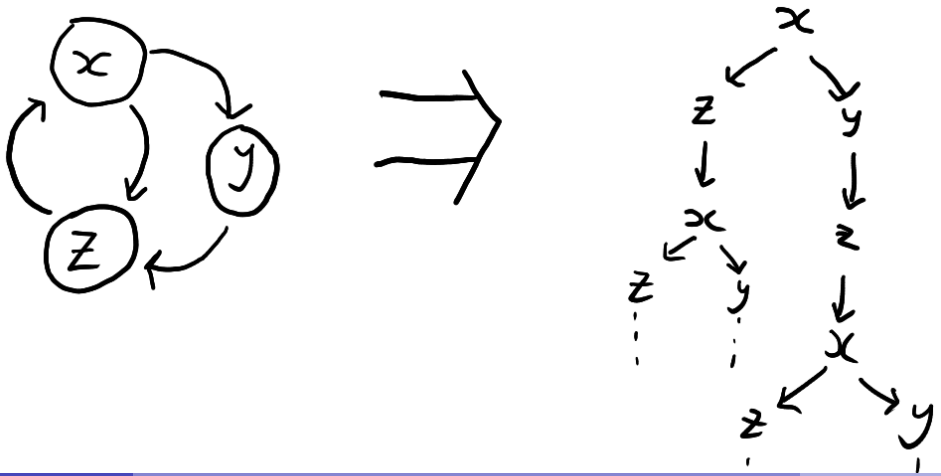

CTL: Syntax and Semantics

Syntax

$$\varphi := \top \mid \bot \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \mathbf{AX}\varphi \mid \mathbf{A}[\varphi \mathbf{U} \varphi] \mid \mathbf{E}[\varphi \mathbf{U} \varphi]$$

Primitive Quantified Temporal Operators

- $\mathbf{AX}\varphi$ — On all linear paths starting from the current state, $\mathbf{X}\varphi$ holds. Or: φ is true at all successor states.
- $\mathbf{A}[\varphi \mathbf{U} \psi]$ — On all linear paths starting from the current state, $\varphi \mathbf{U} \psi$ holds.
- $\mathbf{E}[\varphi \mathbf{U} \psi]$ — There exists a linear path starting from the current state where $\varphi \mathbf{U} \psi$ holds.

Semantics of AX, AU, EU

$\mathbf{AX}\varphi$ — On all paths, $\mathbf{X}\varphi$

$$[\![\mathbf{AX}\varphi]\!] = \{s \in S \mid \forall s' \in S.\ (s, s') \in R \to s' \in [\![\varphi]\!]\}$$

$\mathbf{A}[\varphi \mathbf{U} \psi]$ — On all paths, $\varphi \mathbf{U} \psi$

$$[\![\mathbf{A}[\varphi \mathbf{U} \psi]]\!] =$$
$$\{s \in S \mid \forall p \in paths(\mathcal{M}, s).\ \exists n \in \mathbb{N}.\ p^n \in [\![\psi]\!] \wedge (\forall m < n.\ p^m \in [\![\varphi]\!])\}$$

$\mathbf{E}[\varphi \mathbf{U} \psi]$ — There exists a path where $\varphi \mathbf{U} \psi$

$$[\![\mathbf{E}[\varphi \mathbf{U} \psi]]\!] =$$
$$\{s \in S \mid \exists p \in paths(\mathcal{M}, s).\ \exists n \in \mathbb{N}.\ p^n \in [\![\psi]\!] \wedge (\forall m < n.\ p^m \in [\![\varphi]\!])\}$$

The Quantified Temporal Operator Zoo

- $\mathbf{EX}\varphi = \neg\mathbf{AX}\neg\varphi$ — There exists a successor state where φ holds.
- $\mathbf{AF}\varphi = \mathbf{A}[\top\mathbf{U}\varphi]$ — Inevitability. φ eventually holds on all paths.
- $\mathbf{EF}\varphi = \mathbf{E}[\top\mathbf{U}\varphi]$ — Possibility. φ eventually holds on at least one path.
- $\mathbf{AG}\varphi = \neg\mathbf{EF}\neg\varphi$ — Invariance. φ is always true on every path.
- $\mathbf{EG}\varphi = \neg\mathbf{AF}\neg\varphi$ — There exists at least one path where φ is always true.

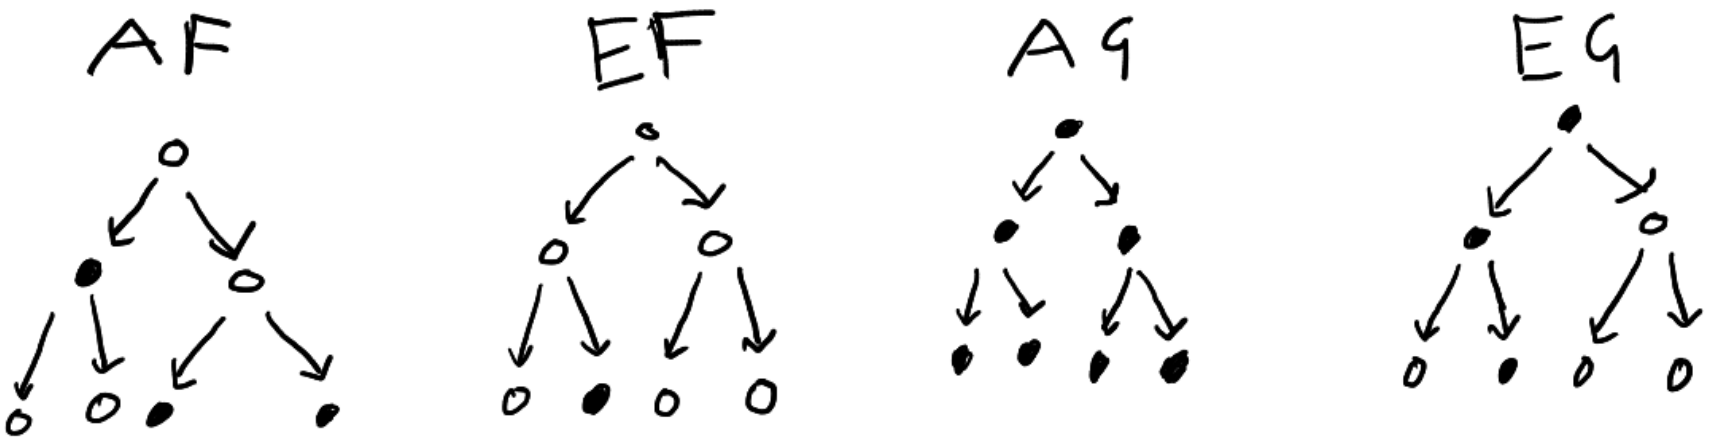

Expressive Power of LTL and CTL (1)

CTL $\subseteq$ LTL?

No; there are CTL formulas not expressible in LTL. Example:

$$\mathbf{EX}p$$

- An LTL-expressible CTL formula must represent a property of the form "for all paths, φ holds".
- So a strictly positive CTL formula which only contains existential quantifiers cannot be LTL-expressible.

Note

This does not mean that all CTL formulas with existential quantifiers are not LTL-expressible. For example, $\neg\mathbf{EX}\neg p = \mathbf{AX}p$ by duality, and $\mathbf{AX}p$ corresponds to $\mathbf{X}p$ in LTL.

Expressive Power of LTL and CTL (2)

LTL $\subseteq$ CTL?

No; there are LTL formulas not expressible in CTL. Example:

$$\mathbf{FG}p$$

- CTL requires every operator be individually quantified.
- We want to say $\mathbf{A}(\mathbf{FG}p)$ — "on all paths, past a certain point p always holds".
- But our only options are $\mathbf{AF}(\mathbf{AG}p)$ (too strong), or $\mathbf{AF}(\mathbf{EG}p)$ (too weak).

Note

$\mathbf{AF}(\mathbf{EG}p)$ seems promising, but $\mathbf{AF}(\mathbf{EG}p) \wedge \mathbf{AG}(\neg p) \neq \bot$, because there can exist a path where p is never true so long as $\mathbf{EG}p$ is true at some point on that path.

Where is normal modal logic?

Note

LTL and CTL as defined here are restricted to serial models. For a cleaner comparison, let's also restrict normal modal logic to serial models.

This logic is called "D", and has $\Diamond\top$ as a theorem (or axiom).

$D \subset CTL$

The propositional fragments are trivially the same. Therefore we only need to translate $\Box$ and $\Diamond$ into CTL. Unsurprisingly:

- $[\![\Box\varphi]\!] = [\![\mathbf{AX}\varphi]\!]$
- $[\![\Diamond\varphi]\!] = [\![\mathbf{EX}\varphi]\!]$

$D \not\subset LTL$

$[\![\Diamond\varphi]\!] = [\![\mathbf{EX}\varphi]\!]$, and $\mathbf{EX}\varphi$ is not LTL-expressible.

Unifying LTL and CTL

Can't just drop the A's

Let's look again at $\mathbf{AF}(\mathbf{AG}p)$. It's a stronger statement than $\mathbf{A}(\mathbf{FG}p)$, as we saw. Any model that satisfies $\mathbf{AF}(\mathbf{AG}p)$ must also satisfy $\mathbf{A}(\mathbf{FG}p)$, but the same is not true in reverse.

If existential path quantification was the only difference between CTL and LTL, then we might have though that these formulas would be equivalent. But they're not! But that intuition wasn't entirely off-base...

Theorem

If a CTL formula φ is expressible in LTL, then it is expressible in LTL exactly by deleting all path quantifiers.

If we are to prove this theorem, we need a logic which subsumes *both* LTL and CTL, and thereby allows us to directly reason about the differences between $\mathbf{A}(\mathbf{FG}p)$ and $\mathbf{AF}(\mathbf{AG}p)$.

CTL*: Arbitrary Path Quantification

Syntax

The grammar is a little more complex, because we need to ensure that all temporal operators are quantified. CTL* formulas are generated by φ.

$$\varphi := \bot \mid \top \mid p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathbf{A}\psi \mid \mathbf{E}\psi$$
$$\psi := \varphi \mid \neg\psi \mid \psi \wedge \psi \mid \psi \vee \psi \mid \mathbf{X}\psi \mid \psi\mathbf{U}\psi$$

The LTL Fragment

If φ is an LTL formula, then $\mathbf{A}\varphi$ is the equivalent formula in CTL*.

The CTL Fragment

If φ is a CTL formula, then φ is also a CTL* formula.

The Big Picture

A CTL* formula not in LTL or CTL

$$\mathbf{E}(\mathbf{FG}p)$$

- Existential path quantification not expressible in LTL.
- Quantification over arbitrary subformulas not allowed in CTL.

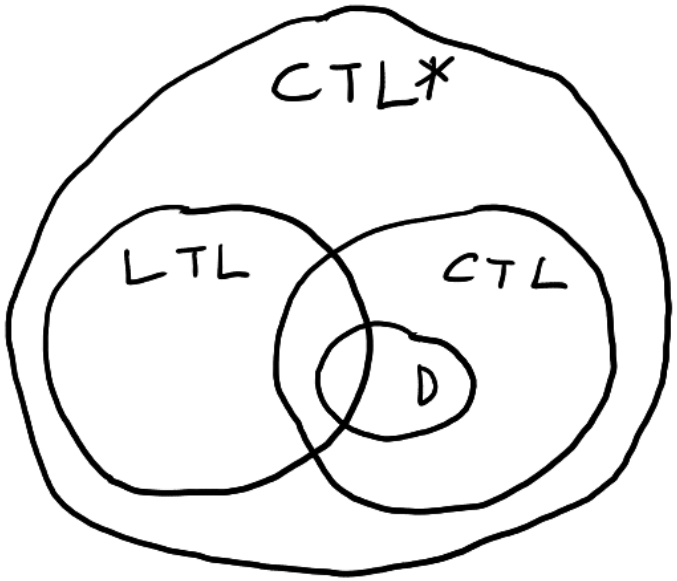

CTL* to LTL (1)

Definition

Let φ^d denote the result of replacing all occurrences of $\{\mathbf{A}, \mathbf{E}\}\psi$ with ψ in the CTL* formula φ.

Theorem

Let φ be a CTL formula. Then φ is LTL-expressible iff $[\![\varphi]\!] = [\![\boldsymbol{A}(\varphi^d)]\!]$.*

Lemma

Let $\mathcal{M}$ be a Kripke model and p an infinite path in $\mathcal{M}$. Then there exists a prefix xy of p such that xy^ω is an infinite path in $\mathcal{M}$, and p and xy^ω prove the same linear formulas.

CTL* to LTL (2)

Proof Sketch — full version in Clarke & Draghicescu (1989)

- Assume that φ is equivalent to $\mathbf{A}\psi$, where ψ is an LTL formula.
- Fix a model and state for which φ holds: $\mathcal{M}, s \vDash \varphi$.
- By the assumption, for all paths p starting at s in $\mathcal{M}$: $p \vDash \psi$.
- By Lemma: For all paths of the form xy^ω starting at s in $\mathcal{M}$: $xy^\omega \vDash \psi$.
- Let $\mathcal{M}(p)$ be the single-path Kripke model formed by the path p in model $\mathcal{M}$. Then for all paths xy^ω in $\mathcal{M}$: $\mathcal{M}(xy^\omega), s \vDash \varphi$.
- $\mathcal{M}(xy^\omega)$ is a single-path model, therefore path quantification is meaningless. So for all paths xy^ω in $\mathcal{M}$: $xy^\omega \vDash \varphi^d$.
- By Lemma: for all paths p in $\mathcal{M}$: $p \vDash \varphi^d$
- Therefore: $\mathcal{M}, s \vDash \mathbf{A}\varphi^d$. QED.

Thanks!

Final Thoughts

- Models need to be extracted from real-world systems/designs, so lots of research goes into doing this efficiently.
- That also means researching temporal logics for other types of model, like Petri nets.
- With appropriate changes to the semantics, we can also consider non-total models, continuous models, games with players, etc.
- Of course, you can do it all coalgebraically too.

Further Reading

- Vardi: "Branching vs Linear Time: Final Showdown"
- Emerson: "Temporal and Modal Logic"
- Maidle: "The Common Fragment of ACTL and LTL"
- Clarke & Draghicescu: "Expressibility Results for Linear-Time and Branching-Time Logics"