Decidability Proof of LTL

The goal of this note is to explain why LTL is decidable. Given an LTL formula ψ we explain how to build a finite transition system S with a "partial" labelling function L (this is explained below) such that ψ has a model iff S, L is a model of ψ. In a sense S, L can be seen as a kind of minimal model (if there is one) of ψ.

This can be used to decide if a formula ϕ is valid in the usual way: we try to find a model for $\neg\phi$. If there is one, we know that ϕ is not valid. If our systematic attempt to find such a model fails, then we know that ϕ is valid.

To simplify the presentation we limit ourselves to the modalities F, G, X (no Until modality). We take also the following syntax for the formulae

$$\psi \ ::= \ \psi \wedge \psi \mid \psi \vee \psi \mid \mu \qquad \mu \ ::= \ p \mid \neg\, p \mid F\ \psi \mid G\ \psi \mid X\ \mu$$

It is clear that any formula can be put on this form, using de Morgan laws and the equivalences

$$X\ (\psi_1 \wedge \psi_2) \leftrightarrow X\ \psi_1 \wedge X\ \psi_2 \qquad X\ (\psi_1 \vee \psi_2) \leftrightarrow X\ \psi_1 \vee X\ \psi_2$$

A *state* will be a finite set of formulae Γ satisying the following properties

1. If $\psi_1 \wedge \psi_2 \in \Gamma$ then $\psi_1 \in \Gamma$ and $\psi_2 \in \Gamma$
2. If $\psi_1 \vee \psi_2 \in \Gamma$ then $\psi_1 \in \Gamma$ or $\psi_2 \in \Gamma$
3. We cannot have both $p \in \Gamma$ *and* $\neg p \in \Gamma$
4. If $G\ \psi \in \Gamma$ then $\psi \in \Gamma$ and $XG\ \psi \in \Gamma$
5. If $F\ \psi \in \Gamma$ then $\psi \in \Gamma$ or $XF\ \psi \in \Gamma$

The last two clauses reflect the equivalences

$$G\ \psi \leftrightarrow \psi \wedge XG\ \psi \qquad F\ \psi \leftrightarrow \psi \vee XF\ \psi$$

The main remark is that give a (finite) set of formulae Γ we can always find a finite number of states $\Gamma_1, \ldots, \Gamma_n$ such that $\wedge\Gamma$ is equivalent to $\wedge\Gamma_1 \vee \ldots \vee \wedge\Gamma_n$. (We can have $n = 0$ in which case Γ is incompatible.) There is furthermore a natural closure algorithm $C(\Gamma)$ that produces $\Gamma_1, \ldots, \Gamma_n$ from Γ, which can be specified by

1. $C(\Gamma) = \Gamma$ if Γ is a state
2. If $\psi_1 \wedge \psi_2 \in \Gamma$ then $C(\Gamma) = C(\Gamma, \psi_1, \psi_2)$
3. If $\psi_1 \vee \psi_2 \in \Gamma$ then $C(\Gamma) = C(\Gamma, \psi_1) \cup C(\Gamma, \psi_2)$
4. If $p, \neg p \in \Gamma$ then $C(\Gamma) = \emptyset$
5. If $G\ \psi \in \Gamma$ then $C(\Gamma) = C(\Gamma, \psi, XG\ \psi)$
6. If $F\ \psi \in \Gamma$ then $C(\Gamma) = C(\Gamma, \psi) \cup C(\Gamma, XF\ \psi)$

Some examples

If Γ is $\neg q \vee p, \neg p \vee r, q$ then $C(\Gamma)$ has only one element Γ, p, q, r.

If Γ is $p \vee q, \neg p \vee r$ then $C(\Gamma)$ has three elements Γ, p, r and $\Gamma, q, \neg p$ and Γ, q, r.

In the *propositional* case, we get a quite good algorithm for computing the conjunctive normal form in this way:

$$(\neg q \vee p) \wedge (\neg p \vee r) \wedge q \quad \leftrightarrow \quad p \wedge q \wedge r$$

$$(p \vee q) \wedge (\neg p \vee r) \quad \leftrightarrow \quad (p \wedge r) \vee (\neg p \wedge q) \vee (q \wedge r)$$

In this case, we can think of each state of $C(\Gamma)$ as a *partial* valuation which ensures the truth of all formulae in Γ. For instance, if Γ is $p \vee q, \neg p \vee r$ it is enough to take $p = r = 1$ to make all formulae in Γ to be true (we don't need to specify the value of q) or to take $p = 0, q = 1$ or to take $q = r = 1$.

Example 1

If Γ is $G\ p,\ F\ q,\ G\ (\neg p \vee \neg q)$ then $C(\Gamma)$ has only one element

$$\Gamma,\ p,\ \neg q,\ XG\ (\neg p \vee \neg q),\ XF\ q,\ XG\ p$$

Example 2

If Γ is $G\ (\neg p \vee X\ p),\ p,\ F\ (\neg p)$ then $C(\Gamma)$ has only one element

$$\Gamma,\ X\ p,\ XG\ (\neg p \vee X\ p),\ XF\ (\neg p)$$

Example 3

If Γ is $G\ (p \vee q),\ F\ (\neg p),\ F\ (\neg q)$ then $C(\Gamma)$ has for elements

$\Gamma_1 = \Gamma,\ p,\ \neg q,\ XG\ (p \vee q),\ XF\ (\neg p)$
$\Gamma_2 = \Gamma,\ p,\ XG\ (p \vee q),\ XF\ (\neg p),\ XF\ (\neg q)$
$\Gamma_3 = \Gamma,\ q,\ \neg\ p,\ XG\ (p \vee q),\ XF\ (\neg q)$
$\Gamma_4 = \Gamma,\ q,\ XG\ (p \vee q),\ XF\ (\neg p),\ XF\ (\neg q)$

Transition relation and minimal potential models

If Γ is a set of formulae, we write $X^{-1}(\Gamma)$ the set of formulae μ such that $X\ \mu \in \Gamma$.

The transition relation is now defined as $\Gamma \rightarrow \Gamma'$ iff Γ' is one of the state in $C(X^{-1}(\Gamma))$.

We can now define the minimal potential model of a set of formulae Γ. The initial states are the elements of $C(\Gamma)$, and the transition system is obtained by taking the states related to these initial states by the transitive closure of the relation $\Delta \rightarrow \Delta'$.

This is a finite transition system, which can be called the *minimal potential* model of Γ. To be a model of Γ we have to find a path

$$\sigma \ = \ \Gamma_1 \rightarrow \Gamma_2 \rightarrow \dots$$

in this transition system which satisfies: if $F\ \mu \in \Gamma_i$ then there exists $j \geqslant i$ such that $\mu \in \Gamma_j$. This is a fairness condition, and the existence of such a path can be checked in the following way. We say that Δ is good for μ iff $F\ \mu \in \Delta$ implies $\mu \in \Delta$. We list then the subformulae

$F\ \mu_1, \ldots,\ F\ \mu_k$ of Γ and the condition is that there is a path $\Delta_1 \rightarrow^* \Delta_2 \ldots \rightarrow^* \Delta_k \rightarrow^* \Delta_1$ where Δ_i is good for μ_i.

It is then possible to show that this method is *sound*: if we have such a path, then we have a model for Γ. For this, one consider the path

$$\sigma \ = \ \Gamma_1 \rightarrow \Gamma_2 \rightarrow \ldots$$

and one shows by induction on ψ that $\sigma^k \Vdash \psi$ if $\psi \in \Gamma_k$, where one takes $L(\Gamma_k)$ to be the set of atomic formulae p such that p is in Γ_k. What matters really is that we have $\sigma_k \Vdash p$ if p is in Γ_k and $\sigma_k \Vdash \neg p$ if $\neg p$ is in Γ_k. The value of q at σ_k actually does not matter if neither q nor $\neg q$ figures in Γ_k. The fact that $\sigma^k \Vdash \psi$ if $\psi \in \Gamma_k$ is clear if ψ is p or $\neg p$, and it holds by induction if ψ is a conjunction or a disjunction. It holds also by induction if ψ is of the form $X\ \mu$. If $\psi = G\ \psi_1$ we have by induction $\sigma_l \Vdash \psi_1$ for all $l \geqslant k$ and hence $\sigma_k \Vdash \psi$ if ψ is in Γ_k. Finally if $\psi = F\ \psi_1$ and ψ is in Γ_k then there exists $l \geqslant k$ such that we have both $F\ \psi_1$ and ψ_1 in Γ_l and then we have by induction $\sigma_l \Vdash \psi_1$ and hence $\sigma_k \Vdash \psi$ as desired.

One can show also that this method is *complete*: if there is a model $M, \pi = s_1 \rightarrow s_2 \rightarrow \ldots$ then it is possible to approximate this model by a path

$$\sigma \ = \ \Gamma_1 \rightarrow \Gamma_2 \rightarrow \ldots$$

such that M, π^k validates all formulae of Γ_k. Indeed, M, s_1 validates all formulae of Γ and hence it is possible to find Γ_1 in $C(\Gamma)$ such that M, s_1 validates all formulae in Γ_1. It then follows that M, s_2 validates all formulae in $X^{-1}(\Gamma_1)$ and hence it is possible to find Γ_2 in $C(X^{-1}(\Gamma_1))$ such that M, s_2 validates all formulae in Γ_2, and so on. Furthemore if $XF\ \mu$ is in Γ_k and s_{k+1} validates μ then we can choose Γ_{k+1} such that both $F\ \mu$ and μ are in Γ_{k+1}. if $XF\ \mu$ is in Γ_k and s_{k+1} does *not* validate μ then it validates $XF\ \mu$ and we have $XF\ \mu$ in Γ_{k+1}. Since M, π^k is a model of all formulae in Γ_k eventually we find $l \geqslant k$ such that M, s_l validates μ. Hence we can choose σ such that there are infinitely many good states for each μ, where μ is a subformula of one formula in Γ.

Some examples

It is actually possible to run this method by hand on some small examples.

Example 1

If Γ is $G\ p,\ F\ q,\ G\ (\neg p \vee \neg q)$ then $C(\Gamma)$ has only one element

$$\Gamma_1 = \Gamma,\ p,\ \neg q,\ XG\ (\neg p \vee \neg q),\ XF\ q,\ XG\ p$$

We get a transition system with only one transition $\Gamma_1 \rightarrow \Gamma_1$. Since Γ_1 is not good for q, this is not a model. Hence there is *no* model and the set $G\ p,\ F\ q,\ G\ (\neg p \vee \neg q)$ is *incompatible* which means that we have $G\ p \wedge F\ q \rightarrow F\ (p \vee q)$.

Example 2

If Γ is $G\ (\neg p \vee X\ p),\ p,\ F\ (\neg p)$ then $C(\Gamma)$ has only one element

$$\Gamma_1 = \Gamma,\ X\ p,\ XG\ (\neg p \vee X\ p),\ XF\ (\neg p)$$

We get a transition system with only one transition $\Gamma_1 \rightarrow \Gamma_1$. Since Γ_1 is not good for $\neg p$, this is not a model. Hence there is *no* model and the set $G\ (\neg p \vee X\ p),\ p,\ F\ (\neg p)$ is *incompatible* which means that we have $G\ (p \rightarrow X\ p) \wedge p \rightarrow G\ p$.

Example 3

If Γ is $G\ (p \vee q),\ F\ (\neg p),\ F\ (\neg q)$ then $C(\Gamma)$ has for elements

$\Gamma_1 = \Gamma,\ p,\ \neg q,\ XG\ (p \vee q),\ XF\ (\neg p)$
$\Gamma_2 = \Gamma,\ p,\ XG\ (p \vee q),\ XF\ (\neg p),\ XF\ (\neg q)$
$\Gamma_3 = \Gamma,\ q,\ \neg\ p,\ XG\ (p \vee q),\ XF\ (\neg q)$
$\Gamma_4 = \Gamma,\ q,\ XG\ (p \vee q),\ XF\ (\neg p),\ XF\ (\neg q)$

For building the minimal potential model, we need to consider the closures of $X^{-1}(\Gamma_i)$. Notice that $X^{-1}(\Gamma_2) = X^{-1}(\Gamma_4) = \Gamma$. We have $X^{-1}(\Gamma_1) = G\ (p \vee q),\ F\ (\neg p)$ which generates

$\Gamma_5 = G\ (p \vee q),\ F\ (\neg p),\ p,\ XG\ (p \vee q),\ XF\ (\neg\ p)$
$\Gamma_6 = G\ (p \vee q),\ F\ (\neg p),\ q,\ \neg\ p,\ XG\ (p \vee q)$
$\Gamma_7 = G\ (p \vee q),\ F\ (\neg p),\ q,\ XG\ (p \vee q),\ XF\ (\neg\ p)$

and $X^{-1}(\Gamma_3) = G\ (p \vee q),\ F\ (\neg q)$ which generates

$\Gamma_8 = G\ (p \vee q),\ F\ (\neg q),\ q,\ XG\ (p \vee q),\ XF\ (\neg\ q)$
$\Gamma_9 = G\ (p \vee q),\ F\ (\neg q),\ p,\ \neg\ q,\ XG\ (p \vee q)$
$\Gamma_{10} = G\ (p \vee q),\ F\ (\neg q),\ p,\ XG\ (p \vee q),\ XF\ (\neg\ q)$

We need then to add the states

$\Gamma_{11} = G\ (p \vee q),\ p,\ XG\ (p \vee q)$
$\Gamma_{12} = G\ (p \vee q),\ q,\ XG\ (p \vee q)$

We find then the model

$$\Gamma_1 \rightarrow \Gamma_6 \rightarrow \Gamma_{11} \rightarrow \Gamma_{11} \rightarrow \dots$$

which shows that Γ is not incompatible. Hence we conclude from this that the formula

$$G\ (p \vee q) \rightarrow G\ p \vee G\ q$$

is *not* valid (it has a counter-model).

Example 4

The reader can now test this method on the example $GF\ p,\ FG\ (\neg p)$ (we find one model) and $FG\ p,\ FG\ (\neg p)$ (no model).

Connection with first-order logic

There is a natural interpretation of LTL in the first-order logic over the language with one successor symbol, one relation symbol ($\leqslant$) and where each atomic formula p is interpreted as a unary predicate $p(x)$.

For instance $G\ (p \wedge q) \rightarrow G\ p \wedge G\ q$ becomes

$$(\forall x.(p(x) \wedge q(x))) \rightarrow \forall x.p(x) \wedge \forall x.q(x)$$

and $G\ (p \rightarrow X\ p) \wedge p \rightarrow G\ p$ becomes

$$\forall x.(p(x) \rightarrow p(s\ x)) \wedge p(z) \rightarrow \forall y.z \leqslant y \rightarrow p(y)$$

We have just given a *decision procedure* for this fragment of first-order logic: monadic (only unary predicates) theory of integers.

By considering a version of LTL with two next operations X_0, X_1 it would be possible similarly to give a decision procedure for the corresponding fragment of first-order logic: monadic theory of binary words.