Emergency kit for LTL

François Laroussinie

16 octobre 2024

draft !

1 La logique LTL

Let AP be a finite set of atomic propositions. LTL formulas are defined as follows :

$$\mathsf{LTL} \ni \varphi, \psi ::= p \mid \varphi \vee \psi \mid \neg\varphi \mid \mathbf{X}\,\varphi \mid \varphi\,\mathbf{U}\,\psi$$

with $p \in \mathsf{AP}$.

An LTL formula is interpreted over an infinite word $w = w_0 w_1 \ldots \in (2^{\mathsf{AP}})^\omega$. We use $w_{\geq i}$ with $i \geq 0$ to denote the infinite word $w_i w_{i+1} \ldots$.

— $w \models p$ iff $p \in w_0$
— $w \models \varphi \vee \psi$ iff $w \models \varphi$ or $w \models \psi$
— $w \models \neg\varphi$ iff $w \not\models \varphi$
— $w \models \mathbf{X}\,\varphi$ iff $w_{\geq 1} \models \varphi$
— $w \models \varphi\,\mathbf{U}\,\psi$ iff $\exists i \geq 0 : w_{\geq 0} \models \psi$ and $\forall 0 \leq k < i$, we have : $w_{\geq k} \models \varphi$

We also use standard abbreviations : $\top$, $\bot$, $\wedge$, $\Leftrightarrow$, $\Rightarrow$. And : $\mathbf{F}\,\_ \stackrel{\text{def}}{=} \top\,\mathbf{U}\,\_$, $\mathbf{G}\,\_ \stackrel{\text{def}}{=} \neg\mathbf{F}\,\neg\_$.

Definition 1. *A Kripke structure is a 4-tuple $\mathcal{S} = \langle S, s_0, \rightarrow, \ell\rangle$ where :*
— *S is a finite set of states, $s_0 \in S$ is the initial state,*
— *$\rightarrow \subseteq S \times S$ is the transition relation (we assume that $\forall s \in S, \exists s' \in S.\ s \rightarrow s'$),*
— *$\ell : S \rightarrow 2^{\mathsf{AP}}$ is a labelling function of the states with atomic propositions.*

An execution of $\mathcal{S}$ is an infinite sequence $\rho \in S^\omega$ s.t. $\forall i \geq 0$ we have $\rho(i) \rightarrow \rho(i+1)$. It describes an infinite word over 2^{AP} : $\ell(\rho(0))\ell(\rho(1))\ell(\rho(2))\ldots$ We use $W(\mathcal{S})$ to denote all these words associated with executions issued from the initial state s_0 of $\mathcal{S}$.

Verification problems We will consider the following problems :

Satisfiability : $\models \varphi$
input : $\varphi \in \mathsf{LTL}$
output : Yes iff there exists a word $w \in (2^{\mathsf{AP}})^\omega$ such that $w \models \varphi$.

Model-checking $\exists$: $\mathcal{S} \models_\exists \varphi$
input : $\mathcal{S}$ a Kripke structure, $\varphi \in \mathsf{LTL}$
output : Yes iff there exists $w \in W(\mathcal{S})$ such that $w \models \varphi$.

Model-checking $\forall$ **:** $\mathcal{S} \models_\forall \varphi$ **input :** $\mathcal{S}$ a Kripke structure, $\varphi \in \mathsf{LTL}$ **output :** Yes iff for every $w \in W(\mathcal{S})$, we have $w \models \varphi$.

2 Expressivity

To be done.

3 Generalized Büchi Automata for LTL

GBA. A generalized Büchi automaton (GBA) is a 5-tuple $\mathcal{A} = (Q, Q_0, \delta, \Sigma, \mathcal{F})$ where :
- Q is a finite set of states, $Q_0 \subseteq Q$ is the set of initial state,
- Σ is the alphabet,
- $\delta : Q \times \Sigma \mapsto 2^Q$ is the transition function, and
- $\mathcal{F} = \{F_1, \ldots, F_k\}$ is a generalized Büchi condition with $F_i \subseteq Q$ for any i.

A word $w = w_0 w_1 \ldots \in \Sigma^\omega$ is accepted by $\mathcal{A}$ iff there exists $\rho \in Q^\omega$ such that (1) $\rho(0) \in Q_0$, (2) $\rho(i+1) \in \delta(\rho(i), w_i)$ and (3) for every $1 \leq j \leq k$, $\mathsf{Inf}(\rho) \cap F_j \neq \varnothing$ where $\mathsf{Inf}(\rho)$ denotes the states that appear infinitely many times along ρ.

We use $\mathcal{L}(\mathcal{A})$ to denote the language of $\mathcal{A}$.

Automata construction. Given an LTL formula φ, we build a GBA $\mathcal{A}_\varphi$ whose language is precisely $\mathsf{mod}(\varphi)$ (*i.e.* the set of models of φ, that is $\{w \in (2^{\mathsf{AP}})^\omega \mid w \models \varphi\}$). Let S_φ be the set of all φ-subformulas and their negations (with $\neg\neg\psi = \psi$).

$\mathcal{A}_\varphi = (Q_\varphi, Q_0^\varphi, \delta_\varphi, (2^{\mathsf{AP}}), \mathcal{F}_\varphi)$.

1. $Q_\varphi \subseteq 2^{S_\varphi}$. The states of Q_φ are *maximal* and *consistent* subsets of S_φ :
   - A state $q \in Q_\varphi$ is consistent w.r.t. Boolean connectives when :
     - if $\varphi_1 \wedge \varphi_2 \in S_\varphi$, we have $(\varphi_1 \wedge \varphi_2 \in q) \Leftrightarrow (\varphi_1 \in q \text{ and } \varphi_2 \in q)$;
     - if $\varphi_1 \vee \varphi_2 \in S_\varphi$, we have $(\varphi_1 \vee \varphi_2 \in q) \Leftrightarrow (\varphi_1 \in q \text{ or } \varphi_2 \in q)$;
     - if $\psi \in q$ then $\neg\psi \notin q$.
   - A state $q \in Q_\varphi$ is maximal iff for any $\psi \in S_\varphi$, we have either $\psi \in q$ or $\neg\psi \in q$.
   - A state $q \in Q_\varphi$ is consistent w.r.t. temporal modalities when :
     - if $\varphi_1 \mathbf{U} \varphi_2 \in q$, then either $\varphi_2 \in q$ or $\psi_1 \in q$;
     - if $\varphi_1 \mathbf{U} \varphi_2 \in S_\varphi$ and $\varphi_2 \in q$, then $\varphi_1 \mathbf{U} \varphi_2 \in q$;
2. $Q_0^\varphi = \{q \in Q \mid \varphi \in q\}$;
3. let $q, q' \in Q$ and $\sigma \subseteq \mathsf{AP}$. We have :

$$q' \in \delta(q, \sigma) \quad \Leftrightarrow \quad \begin{cases} \bullet\ \sigma = q \cap \mathsf{AP} \\ \bullet\ \forall \mathbf{X}\,\psi \in S_\varphi,\ (\mathbf{X}\,\psi \in q \ \Leftrightarrow\ \psi \in q') \\ \bullet\ \forall \varphi_1 \mathbf{U} \varphi_2 \in S_\varphi,\ (\varphi_1 \mathbf{U} \varphi_2 \in q \ \Leftrightarrow\ (\varphi_2 \in q \vee \\ \quad (\varphi_1 \in q \ \wedge\ \varphi_1 \mathbf{U} \varphi_2 \in q')) \end{cases}$$

4. The acceptance condition is $\mathcal{F} = \{F_{\varphi_1 \mathbf{U} \varphi_2} \mid \varphi_1 \mathbf{U} \varphi_2 \in S_\varphi\}$ with :

$$F_{\varphi_1 \mathbf{U} \varphi_2} \ = \ \{q \in Q_\varphi \mid \varphi_1 \mathbf{U} \varphi_2 \notin q\ ;\ \text{or}\ \ \varphi_2 \in q\}$$

Correctness. We have the following theorem :

Theorem 2. *For any $\varphi \in \mathsf{LTL}$, we have $\mathcal{L}(\mathcal{A}_\varphi) = \mathsf{mod}(\varphi)$.*

Proof of $\mathcal{L}(\mathcal{A}_\varphi) \subseteq \mathsf{mod}(\varphi)$: For this, we prove the following Lemma :

Lemma 3. *For any* accepting *run $\rho = q_0q_1q_2\ldots$ over the word $w \in (2^{\mathsf{AP}})^\omega$, we have :*

$$\forall i \geq 0, \forall \psi \in S_\varphi, \textit{we have : } \Big(w_{\geq i} \models \psi \;\Leftrightarrow\; \psi \in \rho(i)\Big)$$

Proof. The proof is done by structural induction over ψ.

— $\psi = p \in \mathsf{A}$: Assume $w_{\geq i} \models p$, then we have $p \in w_i$ (by the semantics of LTL). As $\rho(i+1) \in \delta(\rho(i), w_i)$, we have $w_i = \rho(i) \cap \mathsf{AP}$ (by def. of the transitions), and thus $p \in \rho(i)$. Same argument for the other way.

— $\psi = \neg\psi_1$: if $w_{\geq i} \models \neg\psi_1$, then $w_{\geq i} \not\models \psi_1$ and by i.h. we get $\psi_1 \notin \rho(i)$, and then by def. of Q_φ, we get $\neg\psi_1 \in \rho(i)$. Same argument for the other way.

— $\psi = \psi_1 \wedge \psi_2$: If $w_{\geq i} \models \psi_1 \wedge \psi_2$, then (def. of the semantics) we have $w_{\geq i} \models \psi_1$ and $w_{\geq i} \models \psi_2$, and then by i.h. we have $\psi_1, \psi_2 \in \rho(i)$ and then by the def. of Q_φ, we get $\psi_1 \wedge \psi_2 \in \rho(i)$. Same argument for the other way.

— $\psi = \mathbf{X}\,\psi_1$. Assume $w_{\geq i} \models \mathbf{X}\,\psi_1$, then we have $w_{\geq +1i} \models \psi_1$ and by i.h. we get $\psi_1 \in \rho(i+1)$, and by the def. of the transitions we have $\mathbf{X}\,\psi_1 \in \rho(i)$. Same argument for the other way.

— $\psi = \psi_1 \,\mathbf{U}\, \psi_2$. Assume $w_{\geq i} \models \psi_1 \,\mathbf{U}\, \psi_2$. Then there exists $j \geq i$ such that $w_{\geq j} \models \psi_2$ and for any $i \leq k < j$ we have $w_{\geq k} \models \psi_1$. The induction hypothesis allows us to deduce that $\psi_2 \in \rho(j)$ and $\psi_1 \in \rho(k)$ for any $i \leq k < j$. By def. of Q_φ, we can deduce $\psi_1 \,\mathbf{U}\, \psi_2 \in \rho(j)$, and thus $\psi_1 \,\mathbf{U}\, \psi_2 \in \rho(j-1)$, and thus $\psi_1 \,\mathbf{U}\, \psi_2 \in \rho(j-2,),\ldots, \psi_1 \,\mathbf{U}\, \psi_2 \in \rho(i)$!
If $\psi_1 \,\mathbf{U}\, \psi_2 \in \rho(i)$, by def. of Q_φ, we know that either $\psi_2 \in \rho(i)$ (and then by i.h. we get $w_{\geq i} \models \psi_2$ and $w_{\geq i} \models \psi_1 \,\mathbf{U}\, \psi_2$), or $\psi_1 \in \rho(i)$ and $\psi_1 \,\mathbf{U}\, \psi_2 \in \rho(i+1)$. And then either $\psi_2 \in \rho(i+1)$ or $\psi_1 \in \rho(i+1)$ and $\psi_1 \,\mathbf{U}\, \psi_2 \in \rho(i+2)$ etc. As ρ is an accepting execution, it has to satisfy the Büchi condition $F_{\psi_1 \mathbf{U} \psi_2}$ and this ensures that for some position $j \geq i$ we will have $\psi_2 \in \rho(j)$ and the i.h. will allow us to conclude.

□

Proof of $\mathsf{mod}(\varphi) \subseteq \mathcal{L}(\mathcal{A}_\varphi)$: Assume $w \in (2^{\mathsf{AP}})^\omega$ satisfies φ. Let $\rho \in (2^{S_\varphi})^\omega$ defined as follows : $\forall i \geq 0$, $\rho(i) = \{\psi \in S_\varphi \mid w_{\geq i} \models \psi\} \cup \{\neg\psi \mid w_{\geq i} \not\models \psi\}$. For any i, $\rho(i)$ is maximal an consistent. Moreover we have $\rho(i+1) \in \delta(\rho(i), w_i)$. And it is accepting : for every $\varphi_1 \,\mathbf{U}\, \varphi_2$ subformula belonging to some $\rho(i)$, there exists a position $j \geq i$ s.t. $\varphi_2 \in \rho(j)$ (because by def. $\rho(i)$ contains only formulas that are satisfied at position i). Thus the acceptance condition is satisfied.

4 Decision procedures for LTL

The verification problems for LTL reduce to decision problems over $\mathcal{A}_\varphi$:

— $\models \varphi$ is equivalent to decide whether $\mathcal{L}(\mathcal{A}_\varphi) \neq \varnothing$.

— $\mathcal{S} \models_\exists \varphi$ is equivalent to decide whether $W(\mathcal{S}) \cap \mathcal{L}(\mathcal{A}_\varphi) \neq \varnothing$.

— $\mathcal{S} \models_\forall \varphi$ is equivalent to decide whether $W(\mathcal{S}) \subseteq \mathcal{L}(\mathcal{A}_\varphi)$ or $W(\mathcal{S}) \cap \mathcal{L}(\mathcal{A}_{\neg\varphi}) = \varnothing$.

All these problems are PSPACE-complete.

PSPACE-hardness. Let $\mathcal{M} = (\Sigma, Q, q_0, \Delta, \{q_{acc}\})$ be a deterministic polynomially bounded Turing Machine and $w \in \Sigma^n$. We reduce the problem to decide whether $w \in \mathcal{L}(\mathcal{M})$ to some model-checking problem $\mathcal{S}_\mathcal{M} \models_\exists \Phi_{\mathcal{M},w_0}$. Let p be the polynomial function associated with $\mathcal{M}$: the computation of $\mathcal{M}$ over w uses at most $p(n)$ cells on the tape ($|w| = n$).

We assume that the Turing machine stays forever in the accepting state q_{acc} when it is reached. A configuration of $\mathcal{M}$ over w_0 is a word over $\Sigma' = \Sigma \cup \{\#\}$ of length n (the tape), a control state and a position for the tape head. We use the following set of $\mathsf{AP} = \Sigma \cup \{\#, \mathsf{begin}, \mathsf{end}, \} \cup Q$.

Consider the Kripke structure $\mathcal{S}_\mathcal{M}$ is described in Figure 1. The state 0 is followed by a state describing the contents of the first cell : it is labels by either a letter in $\Sigma \cup \{\#\}$ or a letter *and* a state (indicating what is the current control state and the current position of the head). The states that follow the state 1 describe the second cell, etc. Until the $p(n)$-th cell. A configuration of the machine can be described as a path in the structure between state 0 and state $p(n)$, and a computation can be encoded as a sequence of such configurations... The state 0 is labeled with begin and the state $p(n)$ is labeled by end.

LTL formulas are used to specify that :

— the machine starts with w on the tape, q_0 as initial state and the head at position 1 ;
— two successive configurations are consistent w.r.t. the transitions of the machine ;
— the machine ends in the accepting state.

For the first property, we can use the following formula :

$$\Phi_0 \stackrel{\text{def}}{=} \mathbf{X}\,(q_0 \wedge w_0 \wedge \Big(\bigwedge_{1\leq i\leq n-1} \mathbf{X}^{2i} w_i\Big) \wedge \mathbf{X}^{2n}\Big(\bigwedge_{0\leq i\leq p(n)-n} \mathbf{X}^{2i}\sharp\Big)$$

Reaching the accepting state is ensured with : $\Phi_f \stackrel{\text{def}}{=} \mathbf{F}\, q_{acc}$.

It remains to verify that two consecutive configurations are consistent. In general, a cell i may depend on the cells $i-1$, i and $i+1$ of the previous configuration. We enumerate all possibilities for three cells and the change they induce for the middle cell. Let $\bar{q}$ be the abbreviation for $\bigwedge_{q\in Q} \neg q$. We first consider the case of three cells without any control state, the middle cell cannot change :

$$\Phi_1 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1,\alpha_2,\alpha_3\in\Sigma} \mathbf{G}\,\Big((\alpha_1 \wedge \mathbf{X}^2\alpha_2 \wedge \mathbf{X}^4\alpha_3) \;\Rightarrow\; \mathbf{X}^{p(n)+3}\alpha_2\Big)$$

The borders of the tape require dedicated formulas :

$$\Phi_2 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1,\alpha_2\in\Sigma} \mathbf{G}\,\Big((\mathsf{begin} \wedge \mathbf{X}^2\alpha_1 \wedge \mathbf{X}^4\alpha_2) \;\Rightarrow\; \mathbf{X}^{p(n)+3}\alpha_1\Big)$$

And :

$$\Phi_3 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1,\alpha_2\in\Sigma} \mathbf{G}\,\Big((\alpha_1 \wedge \mathbf{X}^2\alpha_2 \wedge \mathbf{X}^4\mathsf{end}) \;\Rightarrow\; \mathbf{X}^{p(n)+3}\alpha_2\Big)$$

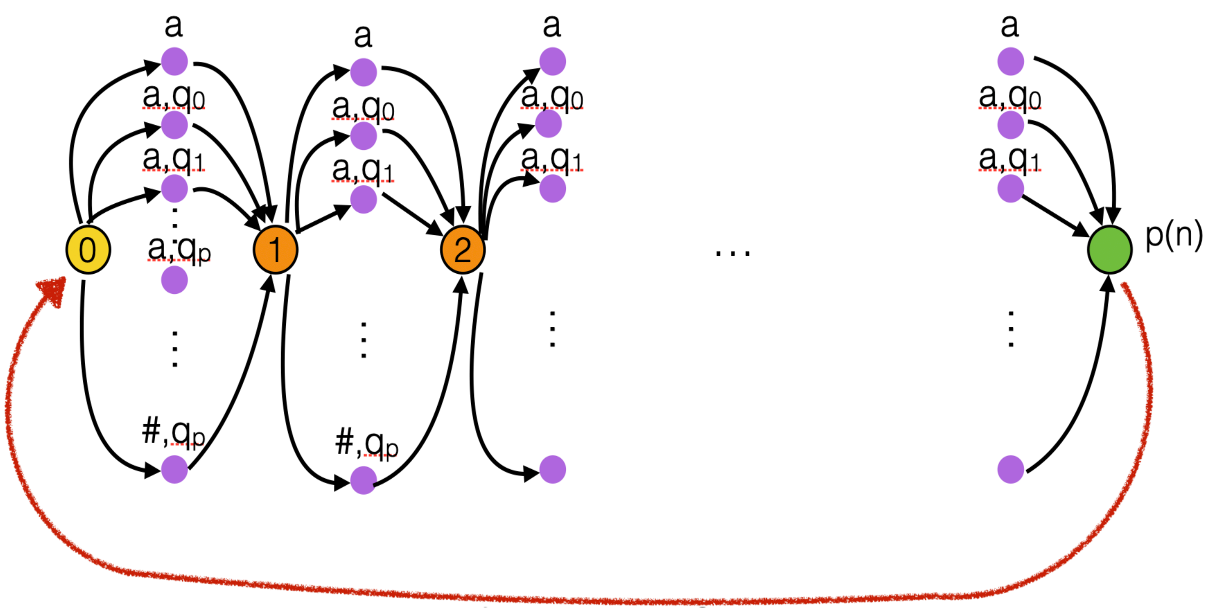

FIGURE 1 – Kripke structure for encoding the Turing machine.

Now for every transition $(q, \sigma, q', \sigma', R)$, we have the three following formulas :

$$\Phi_4 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1,\alpha_2 \in \Sigma'} \mathbf{G}\left((q \wedge \sigma \wedge \mathbf{X}^2\alpha_1 \wedge \mathbf{X}^4\alpha_2) \Rightarrow \mathbf{X}^{p(n)+3}(q' \wedge \alpha_1)\right)$$
$$\Phi_5 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1,\alpha_2 \in \Sigma'} \mathbf{G}\left((\alpha_1 \wedge \mathbf{X}^2(q \wedge \sigma) \wedge \mathbf{X}^4\alpha_2) \Rightarrow \mathbf{X}^{p(n)+3}(\sigma')\right)$$
$$\Phi_4 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1,\alpha_2 \in \Sigma'} \mathbf{G}\left((\alpha_1 \wedge \mathbf{X}^2\alpha_2 \wedge \mathbf{X}^4(q \wedge \sigma)) \Rightarrow \mathbf{X}^{p(n)+3}(\alpha_2)\right)$$
$$\Phi_5 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1 \in \Sigma'} \mathbf{G}\left((\mathsf{begin} \wedge \mathbf{X}^2(q \wedge \sigma) \wedge \mathbf{X}^4(\alpha_1)) \Rightarrow \mathbf{X}^{p(n)+3}(\sigma')\right)$$

And for every transition $(q, \sigma, q', \sigma', L)$, we have three formulas as follows :

$$\Phi_6 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1,\alpha_2 \in \Sigma'} \mathbf{G}\left((q \wedge \sigma \wedge \mathbf{X}^2\alpha_1 \wedge \mathbf{X}^4\alpha_2) \Rightarrow \mathbf{X}^{p(n)+3}(\alpha_1)\right)$$
$$\Phi_7 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1,\alpha_2 \in \Sigma'} \mathbf{G}\left((\alpha_1 \wedge \mathbf{X}^2(q \wedge \sigma) \wedge \mathbf{X}^4\alpha_2) \Rightarrow \mathbf{X}^{p(n)+3}(\sigma')\right)$$
$$\Phi_8 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1,\alpha_2 \in \Sigma'} \mathbf{G}\left((\alpha_1 \wedge \mathbf{X}^2\alpha_2 \wedge \mathbf{X}^4(q \wedge \sigma)) \Rightarrow \mathbf{X}^{p(n)+3}(q' \wedge \alpha_2)\right)$$
$$\Phi_9 \stackrel{\text{def}}{=} \bigwedge_{\alpha_1 \in \Sigma'} \mathbf{G}\left((\alpha_1 \wedge \mathbf{X}^2(q \wedge \sigma) \wedge \mathbf{X}^4(\mathsf{end})) \Rightarrow \mathbf{X}^{p(n)+3}(\sigma')\right)$$

It remains to show that $w \in \mathcal{L}(\mathcal{M})$ iff $\mathcal{S}_{\mathcal{M}} \models_{\exists} \bigwedge_{0 \leq i \leq 9} \Phi_i \wedge \Phi_f$.