About the expressive power of CTL combinators

F. Laroussinie, *
LIFIA-IMAG, Grenoble, France
and
Dept. of Comp. Sci., Aalborg University, Denmark

Key words: temporal logic, CTL, program specification, specification languages.

Abstract

We present a new and quite surprising result about the expressive power of the $\exists\_\mathsf{U}$ and $\forall\_\mathsf{U}$ combinators in CTL.

1 The CTL logic

CTL, the Computation Tree Logic proposed in [2] has been widely considered in literature for the specification of reactive systems [6, 5]. CTL is paradigmatic in the field of branching-time temporal logic because it admits efficient model checking algorithms (see [3]) while remaining very expressive.

CTL formulas are built using four combinators: $\exists\mathsf{X}$, $\forall\mathsf{X}$, $\exists\_\mathsf{U}$ and $\forall\_\mathsf{U}$, plus atomic propositions $a, b, \ldots$ and boolean combinators:

$$(CTL \ni) \;\; f, g \;::=\; \exists\mathsf{X}f \mid \forall\mathsf{X}f \mid \exists[f \,\mathsf{U}\, g] \mid \forall[f \,\mathsf{U}\, g] \mid \neg f \mid f \wedge g \mid a \mid b \mid \ldots$$

where we use the standard abbreviations: $\bot$, $\top$, $f \vee g$, $\ldots$ for resp. $a \wedge \neg a$, $\ldots$

CTL formulas are interpreted over Kripke structures, i.e. directed graphs where every vertice (the *states*) carries a boolean valuation for the atomic propositions. See e.g. [5] for formal definitions. Informally, $\exists\mathsf{X}f$ means "there exists a next state satisfying f" so that a state q satisfies $\exists\mathsf{X}f$ (written $q \models \exists\mathsf{X}f$) iff for some q', a successor state of q (written $q \to q'$), it holds that $q' \models f$. $\forall\mathsf{X}$ means "for all next states ...", so that $\forall\mathsf{X}f \equiv \neg\exists\mathsf{X}\neg f$ (we do not allow states with no successors). $\exists[f \,\mathsf{U}\, g]$ means "from the current state, there is a run π satisfying f until g holds", that is a run $q_0 \to q_1 \to \cdots$ s.t. $q_k \models g$ for some k and $q_i \models f$ for all $0 \leq i < k$. $\forall[f \,\mathsf{U}\, g]$ means "from the current state, all runs satisfy f until g".

Additional useful operators are $\exists\mathsf{F}$ and $\forall\mathsf{F}$, which are weak versions of $\exists\_\mathsf{U}$ (resp. $\forall\_\mathsf{U}$): $\exists\mathsf{F}f$ means "from the current state, there is a (complete) run along which f will hold at some point", while $\forall\mathsf{F}f$ means "along all runs f will eventually hold". Clearly $\exists\mathsf{F}f$ can be defined as $\exists\top\mathsf{U}f$ and $\forall\mathsf{F}f$ as $\forall\top\mathsf{U}f$.

2 A folk result on CTL

In this setting it is well-known that the $\forall\_\mathsf{U}$ operator can be expressed in terms of $\exists\_\mathsf{U}$ and $\forall\mathsf{F}$:

$$\forall[f \,\mathsf{U}\, g] \equiv \forall\mathsf{F}g \wedge \neg\exists[(\neg g) \,\mathsf{U}\, (\neg f \wedge \neg g)] \qquad (1)$$

This fact may be used to simplify proofs by induction over the structure of CTL formulas (see e.g. [1, 8]), or to ease model-checking (see e.g. [3, 4]).

In such cases $\forall\mathsf{F}$ is much simpler than $\forall\_\mathsf{U}$. $\exists\_\mathsf{U}$ also is often simpler to deal with than $\forall\_\mathsf{U}$. Indeed, looking at the quantifier alternation in the semantic clause for $\forall\_\mathsf{U}$, we see that it has the general form "$\forall\, \pi\, \exists\, k \forall\, i \ldots$" and then is (syntactically) in Π_3, whereas (the clause for) $\exists\_\mathsf{U}$ is in Σ_2 (witness "$\exists\, \pi\, \exists\, k\, \forall i$") and $\forall\mathsf{F}$ is in Π_2 (witness "$\forall\, \pi\, \exists k$"). Perhaps this explains why nobody (to our knowledge) has ever tried to express $\exists\_\mathsf{U}$ in terms of $\forall\_\mathsf{U}$ and $\exists\mathsf{F}$ even though many people think this is possible through a definition as simple as (1). We thought so for some time, but we were wrong.

In this note, we show that it is impossible to express $\exists\_\mathsf{U}$ with $\forall\_\mathsf{U}$ and $\exists\mathsf{F}$. This result is quite surprising, especially when one knows about (1) and considers the position of the combinators in the alternation hierarchy. We believe the impossibility proof is interesting. Indeed every time we found (1) in the literature, no mention

*LIFIA-IMAG, 46 Av. Félix Viallet, F-38000 Grenoble, FRANCE. Email:fl@lifia.imag.fr

was made of the dual question, not even through an open question. By contrast, the question pops up every time we teach CTL in the classroom.

3 A surprising remark

(1) entails the well known result:

Folk theorem 1 *CTL is no more expressive than the fragment L_E where only $\exists\_\mathsf{U}$, $\forall\mathsf{F}$, $\exists\mathsf{X}$ are allowed.*

By contrast

Theorem 2 *CTL is strictly more expressive than the fragment L_A where only $\exists\mathsf{F}$, $\forall\_\mathsf{U}$, $\exists\mathsf{X}$ are allowed. Specifically $\exists[a \,\mathsf{U}\, b]$ has no equivalent in L_A.*

For a proof, consider the Kripke structure M given in Figure 1.

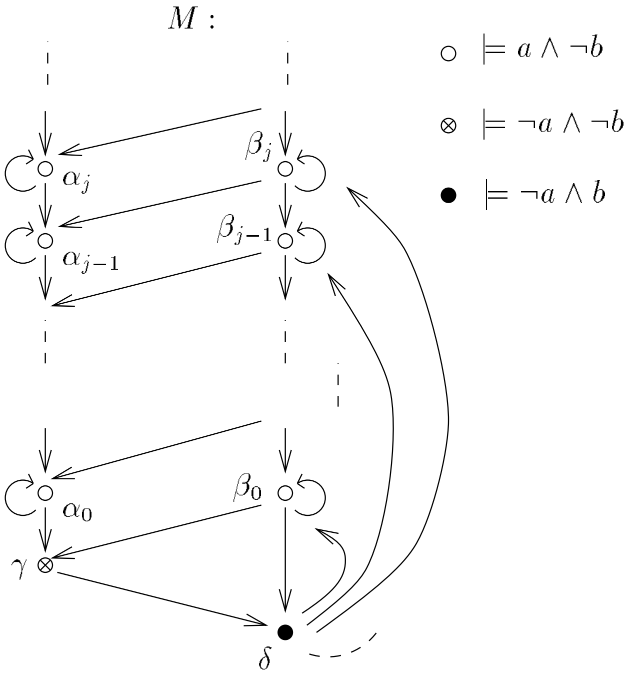

Figure 1: The valuation of the states are indicated by their color.

We prove that considering larger and larger j's, the states α_j and β_j agree on larger and larger L_A formulas. Formally, writing $|f|$ for the height of nested temporal operators in f, we have:

Lemma 3 *For all k and all $j \geq k$, α_j and β_j agree on any $f \in L_A$ s.t. $|f| \leq k$.*

Proof By induction on k. The base case where $k = 0$ is clear: α_j and β_j agree on atomic propositions.

Now assume $k = k' + 1$: by ind. hyp. α_j and β_j agree on all f s.t. $|f| < k$. Now consider some f with $|f| = k$. We prove that α_j and β_j agree on f by induction over the structure of f:

- The cases where f has the form $\neg g$ or $g_1 \wedge g_2$ are obvious.
- f is some $\exists\mathsf{X}g$ (and then $|g| \leq k'$): Assume $\beta_j \models \exists\mathsf{X}g$. There are three distinct cases depending on which successor of β_j satisfies g: it may be α_{j-1}, β_j or β_{j-1}. Then α_{j-1}, or by ind. hyp. α_j or α_{j-1} satisfies g so that $\alpha_j \models \exists\mathsf{X}g$. The reverse direction (assuming $\alpha_j \models \exists\mathsf{X}g$) is dealt with in the same way.
- f is some $\exists\mathsf{F}g$: All reachable states from α_j are also reachable from β_j and conversely (because β_j is reachable from δ). Therefore α_j and β_j agree on f.
- f is some $\forall[g \,\mathsf{U}\, g']$: If $\alpha_j \models \forall[g \,\mathsf{U}\, g']$, then $\alpha_j \models g'$ because there is a run looping on α_j. By ind. hyp., $\beta_j \models g'$ and then $\beta_j \models \forall[g \,\mathsf{U}\, g']$. The reverse direction is proved in the same way.

$\square$

Now pick any L_A formula f and write k for $|f|$. The previous lemma states that if $j \geq k$, then $\alpha_j \models f$ iff $\beta_j \models f$. But $\alpha_j \not\models \exists[a \,\mathsf{U}\, b]$ and $\beta_j \models \exists[a \,\mathsf{U}\, b]$. Therefore f is not equivalent to $\exists[a \,\mathsf{U}\, b]$ and this completes the proof of Theorem 2.

4 Variants and extensions

Clearly the above results apply to the logic $CTL\backslash\mathsf{X}$, i.e. CTL without a "NextTime" operator, a logic suited to reasoning modulo stuttering [7].

Corollary 4
- *$CTL\backslash\mathsf{X}$ is no more expressive than the fragment where only $\exists\_\mathsf{U}$ and $\forall\mathsf{F}$ are allowed.*
- *$CTL\backslash\mathsf{X}$ is strictly more expressive than the fragment where only $\forall\_\mathsf{U}$ and $\exists\mathsf{F}$ are allowed.*

Another interesting logic considered in the literature is the $ECTL$ logic from [6]. $ECTL$ is "Extended CTL", or CTL with fairness. It contains the $\exists \overset{\infty}{\mathsf{F}}$ operator in addition to CTL operators. Informally $\exists \overset{\infty}{\mathsf{F}} f$ means "there is a run (starting from the current state) along which f holds infinitely often". ($ECTL$ also has a $\forall \overset{\infty}{\mathsf{F}}$ combinator but $\forall \overset{\infty}{\mathsf{F}} f$ can be written as $\neg\exists\mathsf{F}\neg\forall\mathsf{F}f$.)

Corollary 5
- *$ECTL$ is no more expressive than the fragment where only $\exists\_\mathsf{U}$, $\forall\mathsf{F}$, $\exists \overset{\infty}{\mathsf{F}}$, $\exists\mathsf{X}$ are allowed.*
- *$ECTL$ is strictly more expressive than the fragment where only $\forall\_\mathsf{U}$, $\exists\mathsf{F}$, $\exists \overset{\infty}{\mathsf{F}}$ and $\exists\mathsf{X}$ are allowed.*

Proof It is easy to extend lemma 3 with $\exists \overset{\infty}{\mathsf{F}}$: for any run π starting from α_j (resp. β_j) there is a run π' starting from β_j (resp. α_j) and having π as a suffix. Clearly any property holding infinitely often along π also holds infinitely often along π'. $\square$

A well-known variant of the "Until" combinator is the "Unless" (see e.g. [5]), or "weak Until", written W, and differing from U in that a run satisfies $f\mathsf{W}g$ is f holds as long as g does not hold, but g is not required to hold eventually. Then CTL can be defined with W in place of U because one can express $\exists\_\mathsf{U}$ and $\forall\_\mathsf{U}$ with $\exists\_\mathsf{W}$ and $\forall\_\mathsf{W}$ (and conversely):

$$\begin{aligned}
\exists[f \mathsf{U} g] &\equiv \neg\forall[(\neg g)\mathsf{W}(\neg f \wedge \neg g)]\\
\exists[f\mathsf{W}g] &\equiv \neg\forall[(\neg g) \mathsf{U} (\neg f \wedge \neg g)]\\
\forall[f \mathsf{U} g] &\equiv \neg\exists[(\neg g)\mathsf{W}(\neg f \wedge \neg g)]\\
\forall[f\mathsf{W}g] &\equiv \neg\exists[(\neg g) \mathsf{U} (\neg f \wedge \neg g)]
\end{aligned}$$

Then a consequence of our previous results is that CTL is strictly more expressive than the logic where only $\exists\_\mathsf{W}$, $\forall\_\mathsf{W}\bot$ and $\exists\mathsf{X}$ are allowed, while it is no more expressive than the logic where only $\forall\_\mathsf{W}$, $\exists\_\mathsf{W}\bot$ and $\exists\mathsf{X}$ are allowed.

Acknowledgements

We would like to thank Philippe Schnoebelen and Sophie Pinchinat for their many helpful suggestions and comments on this work.

References

[1] A. Bouajjani, R. Echahed, and J. Sifakis. On model checking for real-time properties with durations. In *Proc. 8th IEEE Symp. Logic in Computer Science (LICS'93), Montreal, Canada, June 1993*, 1993.

[2] E. M. Clarke and E. A. Emerson. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Proc. Logics of Programs Workshop, Yorktown Heights, LNCS 131*, pages 52–71. Springer-Verlag, May 1981.

[3] E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, April 1986.

[4] E. M. Clarke and O. Grümberg. Research on automatic verification of finite-state concurrent systems. *Ann. Rev. Comput. Sci.*, 2:269–290, 1987.

[5] E. A. Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, vol. B*, chapter 16, pages 995–1072. Elsevier Science Publishers, 1990.

[6] E. A. Emerson and J. Y. Halpern. "Sometimes" and "Not Never" revisited: On branching versus linear time temporal logic. *Journal of the ACM*, 33(1):151–178, 1986.

[7] L. Lamport. What good is temporal logic ? In R. E. A. Mason, editor, *Information Processing'83. Proc. IFIP 9th World Computer Congress, Sep. 1983, Paris, France*, pages 657–668. North-Holland, 1983.

[8] F. Laroussinie and Ph. Schnoebelen. A hierarchy of temporal logics with past. In *Proc. 11th Ann. Symp. Theoretical Aspects of Computer Science (STACS'94), Caen, France, Feb. 1994*, volume 775 of *Lecture Notes in Computer Science*, pages 47–58. Springer-Verlag, 1994.