

Model Checking CTL is Almost Always Inherently Sequential

Olaf Beyersdorff*, Arne Meier*, Michael Thomas*, Heribert Vollmer*, Martin Mundhenk[†] and Thomas Schneider[‡]

*Theoretical Computer Science, University of Hannover, Germany

{beyersdorff, meier, thomas, vollmer}@thi.uni-hannover.de

[†]Computer Science, University of Jena, Germany

martin.mundhenk@uni-jena.de

[‡]Computer Science, University of Manchester, UK

schneider@cs.man.ac.uk

Abstract—The model checking problem for CTL is known to be P-complete (Clarke, Emerson, and Sistla (1986), see Schnoebelen (2002)). We consider fragments of CTL obtained by restricting the use of temporal modalities or the use of negations—restrictions already studied for LTL by Sistla and Clarke (1985) and Markey (2004). For all these fragments, except for the trivial case without any temporal operator, we systematically prove model checking to be either inherently sequential (P-complete) or very efficiently parallelizable (LOGCFL-complete). For most fragments, however, model checking for CTL is already P-complete. Hence our results indicate that in most applications, approaching CTL model checking by parallelism will not result in the desired speed up.

We also completely determine the complexity of the model checking problem for all fragments of the extensions ECTL, CTL⁺, and ECTL⁺.

I. INTRODUCTION

Temporal logic was introduced by Pnueli [5] as a formalism to specify and verify properties of concurrent programs. Computation Tree Logic (CTL), the logic of branching time, goes back to Emerson and Clarke [6] and contains temporal operators for expressing that an event occurs at some time in the future (F), always in the future (G), in the next point of time (X), always in the future until another event holds (U), or as long as it is not released by the occurrence of another event (R), as well as path quantifiers (E, A) for speaking about computation paths. The full language obtained by these operators and quantifiers is called CTL^{*} [7]. In CTL, the interaction between the temporal operators and path quantifiers is restricted. The temporal operators in CTL are obtained by path quantifiers followed directly by any temporal operator, e.g., AF and AU are CTL-operators. Because they start with the universal path quantifier, they are called *universal CTL-operators*. Accordingly, EX and EG are examples for *existential CTL-operators*.

Since properties are largely verified automatically, the computational complexity of reasoning tasks is of great interest. Model checking (MC)—the problem of verifying whether a given formula holds in a state of a given model—is one of the most important reasoning tasks [2]. It is intractable for CTL^{*} (PSPACE-complete [8], [2]), but tractable for CTL (complete for polynomial time [1], [2]).

Although model checking for CTL is tractable, its P-hardness means that it is presumably not efficiently parallelizable. We therefore search for fragments of CTL with a model checking problem of lower complexity. We will consider all subsets of CTL-operators, and examine the complexity of the model checking problems for all resulting fragments of CTL. Further, we consider three additional restrictions affecting the use of negation and study the extensions ECTL, CTL⁺, and their combination ECTL⁺.

The complexity of model checking for fragments of temporal logics has been examined in the literature: Markey [4] considered satisfiability and model checking for fragments of Linear Temporal Logic (LTL). Under systematic restrictions to the temporal operators, the use of negation, and the interaction of future and past operators, Markey classified the two decision problems into NP-complete, coNP-complete, and PSPACE-complete. Further, [9] examined model checking for all fragments of LTL obtained by restricting the set of temporal and propositional operators. The resulting classification separated cases where model checking is tractable from those where it is intractable. For model checking paths in LTL an AC¹ algorithm is presented in [10].

Concerning CTL and its extension ECTL, our results in this paper show that most restricted versions of the model checking problem exhibit the same hardness as the general problem. More precisely, we show that apart from the trivial case where CTL-operators are completely absent, the complexity of CTL model checking is a dichotomy: it is either P-complete or LOGCFL-complete. Unfortunately, the latter case only occurs for a few rather weak fragments and hence there is not much hope that in practice, model checking can be sped up by using parallelism—it is inherently sequential.

Put as a simple rule, model checking for CTL is P-complete for every fragment that allows to express a universal *and* existential CTL-operator. Only for fragments involving the operators EX and EF (or alternatively AX and AG) model checking is LOGCFL-complete. This is visualized in Fig. 2 in Sect. V. Recall that LOGCFL is defined as the class of problems logspace-reducible to context-free languages,

and $NL \subseteq LOGCFL \subseteq NC^2 \subseteq P$. Hence, in contrast to inherently sequential P-hard tasks, problems in LOGCFL have very efficient parallel algorithms.

For the extensions CTL^+ and $ECTL^+$, the situation is more complex. In general, model checking CTL^+ and $ECTL^+$ is Δ_2^P -complete [11]. We show that for $T \subseteq \{A, E, X\}$, both model checking problems restricted to only allowed operators of T remain tractable, while for $T \not\subseteq \{A, E, X\}$, they become Δ_2^P -complete. Yet, for negation restricted fragments with only existential or only universal path quantifiers, we observe a complexity decrease to NP- resp. coNP-completeness.

This paper is organized as follows: Section II introduces CTL, its model checking problems, and the non-basics of complexity theory we use. Section III contains our main results, separated into upper and lower bounds. We also provide a refined analysis of the reductions between different model checking problems with restricted use of negation. The results are then generalized to extensions of CTL in Section IV. Finally, Section V concludes with a graphical overview of the results. For brevity, some proofs are omitted and will be included in the full version of this paper.

II. PRELIMINARIES

A. Temporal Logic

We inductively define CTL^* -formulae as follows. Let Φ be a finite set of atomic propositions. The symbols used are the atomic propositions in Φ , the constant symbol \top , the Boolean connectives \neg , \wedge , and \vee , and the temporal operator symbols A , E , X , F , G , U , and R .

A and E are called a *path quantifiers*, temporal operators aside from A and E are *pure temporal operators*. The atomic propositions and the constants \top and \perp are *atomic formulae*. There are two kinds of formulae, *state formulae* and *path formulae*. Each atomic formula is a state formula, and each state formula is a path formula. If φ, ψ are state formulae and χ, π are path formulae, then $\neg\varphi$, $(\varphi \wedge \psi)$, $(\varphi \vee \psi)$, $A\chi$, $E\chi$ are state formulae, and $\neg\chi$, $(\chi \wedge \pi)$, $(\chi \vee \pi)$, $X\chi$, $F\chi$, $G\chi$, $[\chi U \pi]$, and $[\chi R \pi]$ are path formulae. The set of CTL^* -formulae (or *formulae*) consists of all state formulae.

A *Kripke structure* is a triple $K = (W, R, \eta)$, where W is a finite set of states, $R \subseteq W \times W$ a total relation (i. e., for each $w \in W$, there exists a w' such that $(w, w') \in R$), and $\eta: W \rightarrow \mathfrak{P}(\Phi)$ is a labelling function. A *path* x is an infinite sequence $x = (x_1, x_2, \dots) \in W^\omega$ such that $(x_i, x_{i+1}) \in R$, for all $i \geq 1$. For a path $x = (x_1, x_2, \dots)$ we denote by x^i the path (x_i, x_{i+1}, \dots) .

Let $K = (W, R, \eta)$ be a Kripke structure, $w \in W$ be a state, and $x = (x_1, x_2, \dots) \in W^\omega$ be a path. Further, let φ, ψ be state formulae and χ, π be path formulae. The truth of a CTL^* -formula w. r. t. K is inductively defined as:

$$\begin{aligned} K, w \models \top & \quad \text{always,} \\ K, w \models \perp & \quad \text{never,} \\ K, w \models p & \quad \text{iff } p \in \Phi \text{ and } p \in \eta(w), \end{aligned}$$

$$\begin{aligned} K, w \models \neg\varphi & \quad \text{iff } K, w \not\models \varphi, \\ K, w \models (\varphi \wedge \psi) & \quad \text{iff } K, w \models \varphi \text{ and } K, w \models \psi, \\ K, w \models (\varphi \vee \psi) & \quad \text{iff } K, w \models \varphi \text{ or } K, w \models \psi, \\ K, w \models A\chi & \quad \text{iff } K, x \models \chi \text{ for all paths} \\ & \quad x = (x_1, x_2, \dots) \text{ with } x_1 = w, \\ K, x \models \varphi & \quad \text{iff } K, x_1 \models \varphi, \\ K, x \models \neg\chi & \quad \text{iff } K, x \not\models \chi, \\ K, x \models (\chi \wedge \pi) & \quad \text{iff } K, x \models \chi \text{ and } K, x \models \pi, \\ K, x \models (\chi \vee \pi) & \quad \text{iff } K, x \models \chi \text{ or } K, x \models \pi, \\ K, x \models X\chi & \quad \text{iff } K, x^2 \models \chi \\ K, x \models [\chi U \pi] & \quad \text{iff there exists } k \in \mathbb{N} \text{ such that} \\ & \quad K, x^i \models \chi \text{ for } 1 \leq i < k \text{ and} \\ & \quad K, x^k \models \pi. \end{aligned}$$

The semantics of the remaining temporal operators is defined via the equivalences: $E\chi \equiv \neg A \neg \chi$, $F\chi \equiv [\top U \chi]$, $G\chi \equiv \neg F \neg \chi$, and $[\chi R \pi] \equiv \neg[\neg\chi U \neg\pi]$. A state formula φ is *satisfied by a Kripke structure* K if there exists $w \in W$ such that $K, w \models \varphi$. We will also denote this by $K \models \varphi$.

A *CTL-formula* is a CTL^* -formula in which each path quantifier is followed by exactly one pure temporal operator and each pure temporal operator is preceded by exactly one path quantifier. The set of CTL-formulae forms a strict subset of the set of all CTL^* -formulae. For example, $AGEFp$ is a CTL-formula, but $A(GFp \wedge Fq)$ is not. Pairs of path quantifiers and pure temporal operators are called *CTL-operators*. The operators AX , AF , AG , AU , and AR are *universal* CTL-operators, and EX , EF , EG , EU , and ER are *existential* CTL-operators. Note that $A[\psi U \chi] \equiv AF\chi \wedge \neg E[\neg\chi U (\neg\psi \wedge \neg\chi)]$, and thus $E[\psi R \chi] \equiv EG\chi \vee E[\chi U (\psi \wedge \chi)]$. Hence $\{AX, AF, AR\}$ is a minimal set of operators for CTL (in presence of all Boolean connectives), whereas $\{AX, AG, AU\}$ is not complete for CTL [12].

By $CTL(T)$ we denote the set of CTL-formulae using the connectives $\{\wedge, \vee, \neg\}$ and the CTL-operators in T only. Figure 1 shows the structure of sets of CTL-operators with respect to their expressive power. Moreover, we define the following fragments of $CTL(T)$:

- $CTL_{\text{pos}}(T)$ (positive): CTL-operators may not occur in the scope of a negation,
- $CTL_{\text{a.n.}}(T)$ (atomic negation): negation signs appear only directly in front of atomic propositions,
- $CTL_{\text{mon}}(T)$ (monotone): no negation signs allowed.

This restricted use of negation was introduced and studied in the context of linear temporal logic, LTL, by Sistla and Clarke [3] and Markey [4]. Their original notation was $\tilde{L}(T)$ for $CTL_{\text{a.n.}}(T)$ and $L^+(T)$ for $CTL_{\text{pos}}(T)$.

B. Model Checking

Now we define the *model checking problems* for the above mentioned fragments of CTL. Let \mathcal{L} be CTL, CTL_{mon} , $CTL_{\text{a.n.}}$, or CTL_{pos} .

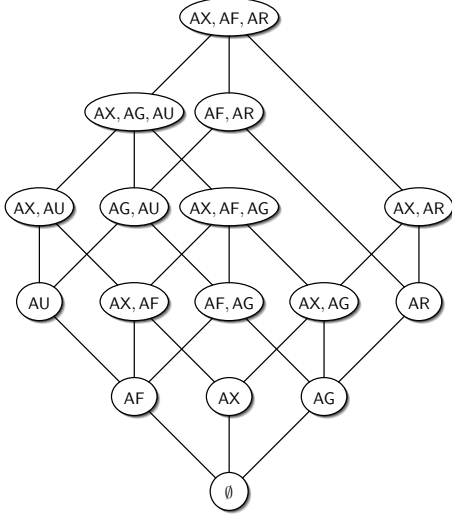


Figure 1. The expressive power of $\text{CTL}(T)$.

Problem: $\mathcal{L}\text{-MC}(T)$

Input: A Kripke structure $K = (W, R, \eta)$,
a state $w \in W$, and an $\mathcal{L}(T)$ -formula φ .

Question: Does $K, w \models \varphi$ hold?

C. Complexity Theory

We assume familiarity with standard notions of complexity theory (cf. [13]). Next we will introduce the notions from circuit complexity that we use for our results. All reductions in this paper are \leq_{cd} -reductions defined as follows: A language A is *constant-depth reducible* to B , $A \leq_{\text{cd}} B$, if there is a logtime-uniform AC^0 -circuit family with oracle gates for B that decides membership in A . That is, there is a circuit family $\mathcal{C} = (C_1, C_2, C_3, \dots)$ such that

- for every n , C_n computes the characteristic function of A for inputs of length n ,
- there is a polynomial p and a constant d such that for all input lengths n , the size of C_n is bounded by $p(n)$ and the depth of C_n is bounded by d ,
- each circuit C_n consists of unbounded fan-in AND and OR gates, negation gates, and gates that compute the characteristic function of B (the *oracle gates*),
- there is a linear-time Turing machine M that can check the structure of the circuit family, i.e., given a tuple $\langle n, g, t, h \rangle$ where n, g, h are binary numbers and $t \in \{\text{AND}, \text{OR}, \text{NOT}, \text{ORACLE}\}$, M accepts if C_n contains a gate g of type t with predecessor h .

Circuit families \mathcal{C} with this last property are called *logtime-uniform* (the name stems from the fact that the time needed by M is linear in the length of its input tuple, hence logarithmic in n). For background information we refer to [14], [15].

We easily obtain the following relations between model checking for fragments of CTL with restricted negation:

Lemma II.1. *For every set T of CTL-operators, we have $\text{CTL}_{\text{mon}}\text{-MC}(T) \leq_{\text{cd}} \text{CTL}_{\text{a.n.}}\text{-MC}(T) \leq_{\text{cd}} \text{CTL}_{\text{pos}}\text{-MC}(T)$. Further, for model checking, atomic negation can be eluded, i. e., $\text{CTL}_{\text{a.n.}}\text{-MC}(T) \leq_{\text{cd}} \text{CTL}_{\text{mon}}\text{-MC}(T)$.*

In Sect. III-C we complete the picture by showing that also $\text{CTL}_{\text{pos}}\text{-MC}(T) \leq_{\text{cd}} \text{CTL}_{\text{mon}}\text{-MC}(T)$.

The class P consists of all languages that have a polynomial-time decision algorithm. A problem is P-complete if it is in P and every other problem in P reduces to it. P-complete problems are sometimes referred to as *inherently sequential*, because P-complete problems most likely (formally: if $\text{P} \neq \text{NC}$) do not possess NC-algorithms, that is, algorithms running in polylogarithmic time on a parallel computer with a polynomial number of processors. Formally, NC contains all problems solvable by polynomial-size polylogarithmic-depth logtime-uniform families of circuits with bounded fan-in AND, OR, NOT gates.

There is an NC-algorithm for parsing context-free languages, that is, $\text{CFL} \subseteq \text{NC}$. Therefore, complexity theorists have studied the class LOGCFL of all problems reducible to context-free languages (the name “LOGCFL” refers to the original definition of the class in terms of logspace-reductions, however it is known that the class does not change if instead, as everywhere else in this paper, \leq_{cd} -reductions are used). Hence, $\text{LOGCFL} \subseteq \text{NC}$ (even $\text{LOGCFL} \subseteq \text{NC}^2$, the second level of the NC-hierarchy, where the depth of the occurring circuits is restricted to $O(\log^2 n)$). The class LOGCFL has a number of different maybe even somewhat surprising characterizations, e. g., languages in LOGCFL are those that can be decided by nondeterministic Turing machines operating in polynomial time that have a worktape of logarithmic size and additionally a stack whose size is not bounded.

More important for this paper is the characterization of LOGCFL as those problems computable by SAC^1 circuit families, that is families of circuits that

- have polynomial size and logarithmic depth,
- consist of unbounded fan-in OR gates and bounded fan-in AND gates and negation gates, but the latter are only allowed at the input-level,
- are logtime-uniform (as defined above).

Since the class LOGCFL is known to be closed under complementation, the second condition can equivalently be replaced to allow unbounded fan-in AND gates and restrict the fan-in of OR gates to be bounded.

To summarize:

$$\text{NC}^1 \subseteq \text{L} \subseteq \text{NL} \subseteq \text{LOGCFL} = \text{SAC}^1 \subseteq \text{NC}^2;$$

and problems in these classes possess very efficient parallel algorithms: they can be solved in time $O(\log^2 n)$ on a parallel machine with a tractable number of processors. For more background on these and related complexity classes, we refer the reader to [15].

III. MODEL CHECKING CTL AND CTL_{pos}

This section contains our main results on the complexity of model checking for CTL and CTL_{pos}. We defer the analysis of the fragments CTL_{a.n.} and CTL_{mon} to Sect. III-C where we will see that their model-checking problems are computationally equivalent to model checking for CTL_{pos}.

While model checking for CTL in general is known to be polynomial time solvable and in fact P-complete [1], [2], we improve the lower bound by showing that only one temporal operator is sufficient to obtain hardness for P.

Theorem III.1. *For each nonempty set T of CTL-operators, CTL-MC(T) is P-complete. If $T = \emptyset$, then CTL-MC(T) is NC¹-complete.*

If we consider only formulae from CTL_{pos}, where no CTL-operators are allowed inside the scope of a negation, the situation changes and the complexity of model checking exhibits a dichotomous behavior. As long as EG or AF are expressible the model checking problem remains P-complete. Otherwise, its complexity drops to LOGCFL.

Theorem III.2. *Let T be any set of CTL-operators. Then CTL_{pos}-MC(T) is*

- NC¹-complete if $T = \emptyset$,
- LOGCFL-complete if $\emptyset \subsetneq T \subseteq \{\text{EX}, \text{EF}\}$ or $\emptyset \subsetneq T \subseteq \{\text{AX}, \text{AG}\}$, and
- P-complete otherwise.

We split the proofs of Theorems III.1 and III.2 into the upper and lower bounds in the following two subsections.

A. Upper Bounds

In general, model checking for CTL is known to be solvable in P [1]. While this upper bound also applies to CTL_{pos}-MC(T) (for every T), we improve it for positive CTL-formulae using only EX and EF, or only AX and AG.

Proposition III.3. *Let T be a set of CTL-operators such that $T \subseteq \{\text{EX}, \text{EF}\}$ or $T \subseteq \{\text{AX}, \text{AG}\}$. Then CTL_{pos}-MC(T) is in LOGCFL.*

Proof: First consider the case $T \subseteq \{\text{EX}, \text{EF}\}$. We claim that Algorithm 1 on page 4 recursively decides whether the Kripke structure $K = (W, R, \eta)$ satisfies the CTL_{pos}(T)-formula φ in state $w_0 \in W$. There, S is a stack that stores pairs $(\varphi, w) \in \text{CTL}_{\text{pos}}(T) \times W$ and R^* denotes the transitive closure of R .

Algorithm 1 always terminates because each subformula of φ is pushed to the stack S at most once. For correctness, an induction on the structure of formulae shows that Algorithm 1 returns **false** if and only if for the most recently popped pair (ψ, w) from S , we have $K, w \not\models \psi$. Thence, in particular, Algorithm 1 returns **true** iff $K, w \models \varphi$.

Algorithm 1 can be implemented on a nondeterministic polynomial-time Turing machine that besides its (unbounded)

Algorithm 1 Determine whether $K, w_0 \models \varphi$.

Require: a Kripke structure $K = (W, R, \eta)$, $w_0 \in W$, $\varphi \in \text{CTL}_{\text{pos}}(T)$

- 1: push($S, (\varphi, w_0)$)
- 2: **while** S is not empty **do**
- 3: $(\varphi, w) \leftarrow \text{pop}(S)$
- 4: **if** φ is a propositional formula **then**
- 5: **if** φ evaluates to false in w under η **then**
- 6: **return false**
- 7: **end if**
- 8: **else if** $\varphi = \alpha \wedge \beta$ **then**
- 9: push($S, (\beta, w)$)
- 10: push($S, (\alpha, w)$)
- 11: **else if** $\varphi = \alpha \vee \beta$ **then**
- 12: nondet. push($S, (\alpha, w)$) or push($S, (\beta, w)$)
- 13: **else if** $\varphi = \text{EX}\alpha$ **then**
- 14: nondet. choose $w' \in \{w' \mid (w, w') \in R\}$
- 15: push($S, (\alpha, w')$)
- 16: **else if** $\varphi = \text{EF}\alpha$ **then**
- 17: nondet. choose $w' \in \{w' \mid (w, w') \in R^*\}$
- 18: push($S, (\alpha, w')$)
- 19: **end if**
- 20: **end while**
- 21: **return true**

stack uses only logarithmic memory for the local variables. Thus CTL_{pos}-MC(T) is in LOGCFL.

The case $T \subseteq \{\text{AX}, \text{AG}\}$ is analogous and follows from closure of LOGCFL under complementation. ■

Finally, for the trivial case where no CTL-operators are present, model checking CTL(\emptyset)-formulae is equivalent to the problem of evaluating a propositional formula. This problem is known to be solvable in NC¹ [16].

B. Lower Bounds

The P-hardness of model checking for CTL is folklore in the model checking community (cf. [2]), but we could not find a formal proof.¹ We improve this lower bound and concentrate on the smallest fragments of monotone CTL—w. r. t. CTL-operators—with P-hard model checking.

Proposition III.4. *Let T denote a set of CTL-operators. Then CTL_{mon}-MC(T) is P-hard if T contains an existential and a universal CTL-operator.*

Proof: First, assume that $T = \{\text{AX}, \text{EX}\}$. We give a generic reduction from alternating Turing machines working in logarithmic space. Let M be such a machine and let x be an input to M . We may assume w. l. o. g. that each transition of M leads from an existential to a universal configuration and vice versa. Further we may assume that each computation

¹In [2], an informal proof sketch is given.

of M ends after the same number $p(n)$ of steps, where p is a polynomial and n is the length of M 's input.

Let $c_1, \dots, c_{q(n)}$ be an enumeration of all possible configurations of M on input x , starting with the initial configuration c_1 and polynomial q . We construct a Kripke structure $K := (W, R, \eta)$ by defining the set $W := \{c_i^j \mid 1 \leq i \leq q(n), 0 \leq j \leq p(n)\}$ and the relation $R \subseteq W \times W$ as

$$R := \left\{ (c_i^j, c_k^{j+1}) \mid \begin{array}{l} M \text{ reaches configuration } c_k \text{ from} \\ c_i \text{ in one step, } 0 \leq j < p(n) \end{array} \right\} \\ \cup \{ (c_i^{p(n)}, c_i^{p(n)}) \mid 1 \leq i \leq q(n) \}.$$

The labelling function η is defined as $\eta(w) := \{t\}$ iff w is an accepting configuration, and $\eta(w) = \emptyset$ otherwise. Then it holds that

$$M \text{ accepts } x \iff K, c_1^0 \models \psi_1 \left(\psi_2 \left(\dots \psi_{p(n)}(t) \dots \right) \right),$$

where $\psi_i(x) := \text{AX}(x)$ if M 's configurations after the i th step are universal, and $\psi_i(x) := \text{EX}(x)$ otherwise. Notice that the constructed CTL-formula does not contain any propositional operator. Since $p(n)$ and $q(n)$ are polynomials, the size of K and φ is polynomial in the size of (M, x) . Moreover, K and φ can be constructed from M and x using AC^0 -circuits. Thus, $A \leq_{\text{cd}} \text{CTL}_{\text{mon}}\text{-MC}(\{\text{AX}, \text{EX}\})$ for all $A \in \text{ALOGSPACE} = \text{P}$.

For $T = \{\text{AF}, \text{EG}\}$ we take new atomic propositions $d_0, \dots, d_{p(n)}$ and modify the above reduction by defining the formulas η and ψ_i as follows:

$$\eta(w) := \{d_j \mid w = c_i^j, 1 \leq i \leq q(n)\} \cup \{t \mid w \text{ is an accepting configuration}\} \\ \psi_i(x) := \begin{cases} \text{AF}(d_{i+1} \wedge x), & \text{if } M\text{'s configurations in} \\ & \text{step } i \text{ are universal,} \\ \text{EG}(D_{i+1} \vee x), & \text{otherwise,} \end{cases} \quad (1)$$

where $D_i = \bigvee_{i \neq j \in \{0, \dots, p(n)\}} d_j$.

For the combinations of T being one of $\{\text{AF}, \text{EF}\}$, $\{\text{AF}, \text{EX}\}$, $\{\text{AG}, \text{EG}\}$, $\{\text{AG}, \text{EX}\}$, $\{\text{AX}, \text{EF}\}$, and $\{\text{AX}, \text{EG}\}$, the P-hardness of $\text{CTL}_{\text{mon}}\text{-MC}(T)$ is obtained using analogous modifications to η and the ψ_i 's.

For the remaining combinations involving the until or the release operator, observe that w. r. t. the Kripke structure K as defined in (1), $\text{AF}(d_i \wedge x)$ and $\text{EG}(D_i \vee x)$ are equivalent to $\text{A}[d_{i-1} \text{U}x]$ and $\text{E}[d_{i-1} \text{U}x]$, and R and U are duals. ■

In the presence of arbitrary negation, universal operators are definable by existential operators and vice versa. Hence, from Proposition III.4 we obtain the following corollary.

Corollary III.5. *The model checking problem $\text{CTL}\text{-MC}(T)$ is P-hard for each nonempty set T of CTL-operators.*

Returning to monotone CTL, in most cases even one operator suffices to make model checking P-hard:

Proposition III.6. *Let T denote a set of CTL-operators. Then $\text{CTL}_{\text{mon}}\text{-MC}(T)$ is P-hard if T contains at least one of the operators EG, EU, ER, AF, AU, or AR.*

The proof of this proposition proceeds similarly as the proof of Proposition III.4, but is technically more involved. In essence, it shows that both AX and EX can be simulated by using only EG.

By Lemma II.1, $\text{CTL}_{\text{mon}}\text{-MC}(T) \leq_{\text{cd}} \text{CTL}_{\text{pos}}\text{-MC}(T)$ and hence the above results directly translate to model checking for CTL_{pos} : for any set T of temporal operators, $\text{CTL}_{\text{pos}}\text{-MC}(T)$ is P-hard if $T \not\subseteq \{\text{EX}, \text{EF}\}$ or if $T \not\subseteq \{\text{AX}, \text{AG}\}$. These results cannot be improved w. r. t. T , as for $T \subseteq \{\text{EX}, \text{EF}\}$ and $T \subseteq \{\text{AX}, \text{AG}\}$ we obtain a LOGCFL upper bound for model checking from Proposition III.3. In the following proposition we prove the matching LOGCFL lower bound.

Proposition III.7. *For every nonempty set T of CTL-operators, the model checking problem $\text{CTL}_{\text{mon}}\text{-MC}(T)$ is LOGCFL-hard.*

Proof: As explained in Sect. II-C, LOGCFL can be characterized as the set of languages recognizable by logtime-uniform SAC^1 circuits, i. e., circuits of logarithmic depth and polynomial size consisting of \vee -gates with unbounded fan-in and \wedge -gates with fan-in 2. For every single CTL-operator O , we will show that $\text{CTL}_{\text{mon}}\text{-MC}(T)$ is LOGCFL-hard for all $T \supseteq \{O\}$ by giving a generic \leq_{cd} -reduction f from the word problem for SAC^1 circuits to $\text{CTL}_{\text{mon}}\text{-MC}(T)$.

First, consider $\text{EX} \in T$. Let C be a logtime-uniform SAC^1 circuit of depth ℓ with n inputs and let $x = x_1 \dots x_n \in \{0, 1\}^n$. Assume w. l. o. g. that C is layered into alternating layers of \wedge -gates and \vee -gates and that the output gate of C is an \vee -gate. We number the layers bottom-up, that is, the layer containing (only) the output gate has level 0, whereas the input-gates and negations of the input-gates are situated in layer ℓ . Denote the graph of C by $G := (V, E)$, where $V := V_{\text{in}} \uplus V_{\wedge} \uplus V_{\vee}$ is partitioned into the sets corresponding to the (possibly negated) input-gates, the \wedge -gates, and the \vee -gates, respectively. G is acyclic and directed with paths leading from the input to the output gates. From (V, E) , we construct a Kripke structure that allows to distinguish the two predecessors of an \wedge -gate from each other. This will be required to model proof trees using $\text{CTL}_{\text{mon}}(\{\text{EX}\})$ -formulae.

For $i \in \{1, 2\}$, let $V_{\text{in}}^i := \{v^i \mid v \in V_{\text{in}}\}$, $V_{\vee}^i := \{v^i \mid v \in V_{\vee}\}$ and define $V_{\text{in}, \vee}^i := V_{\text{in}}^i \cup V_{\vee}^i$. Further define

$$E' := \{(v, u^i) \in V_{\wedge} \times V_{\text{in}, \vee}^i \mid (u, v) \in E \text{ and } u \text{ is the } i\text{th} \\ \text{predecessor of } v\} \cup \{(v, v) \mid v \in V_{\text{in}}^1 \cup V_{\text{in}}^2\} \cup \\ \bigcup_{i \in \{1, 2\}} \{(v^i, u) \in V_{\text{in}, \vee}^i \times V_{\wedge} \mid (u, v) \in E\},$$

where the ordering of the predecessors is implicitly given in the encoding of C . We now define a Kripke structure $K := (V', E', \eta)$ with states $V' := V_{\text{in}, \vee}^1 \cup V_{\text{in}, \vee}^2 \cup V_{\wedge}$,

relation E' , and labelling function $\eta : V' \rightarrow \mathfrak{P}(\{1, 2, t\})$,

$$\eta(v) := \begin{cases} \{i, t\}, & \text{if } v = v_{\text{in}_j} \in V_{\text{in}}^i \text{ and } x_j = 1, \\ \{i, t\}, & \text{if } v = \bar{v}_{\text{in}_j} \in V_{\text{in}}^i \text{ and } x_j = 0, \\ \{i\}, & \text{if } v = v_{\text{in}_j} \in V_{\text{in}}^i \text{ and } x_j = 0, \\ \{i\}, & \text{if } v = \bar{v}_{\text{in}_j} \in V_{\text{in}}^i \text{ and } x_j = 1, \\ \{i\}, & \text{if } v \in V_{\vee}^i, \\ \emptyset, & \text{otherwise,} \end{cases}$$

where $i = 1, 2$, $j = 1, \dots, n$ and $v_{\text{in}_1}, \dots, v_{\text{in}_n}, \bar{v}_{\text{in}_1}, \dots, \bar{v}_{\text{in}_n}$ enumerate the input gates and their negations. The formula φ that is to be evaluated on K will consist of atomic propositions 1, 2 and t , Boolean connectives \wedge and \vee , and the CTL-operator EX. To construct φ we recursively define formulae $(\varphi_i)_{0 \leq i \leq \ell}$ by

$$\varphi_i := \begin{cases} t, & \text{if } i = \ell, \\ \text{EX}\varphi_{i+1}, & \text{if } i \text{ is even } (\vee\text{-layers}), \\ \bigwedge_{i=1,2} \text{EX}(i \wedge \varphi_{i+1}), & \text{if } i \text{ is odd } (\wedge\text{-layers}). \end{cases}$$

We define the reduction function f as the mapping $(C, x) \mapsto (K, v_0, \varphi)$, where v_0 is the node corresponding to the output gate of C and $\varphi := \varphi_0$. We stress that the size of φ is polynomial, for the depth of C is logarithmic only. Clearly, each minimal accepting subtree (cf. [17] or [15, Definition 4.15]) of C on input x translates into a sub-structure K' of K such that $K', v_0 \models \varphi$ where

- 1) K' includes v_0 ,
- 2) K' includes one successor for every node corresponding to an \vee -gate, and
- 3) K' includes the two successors of every node corresponding to an \wedge -gate.

As $C(x) = 1$ iff there exists a minimal accepting subtree of C on x , the LOGCFL-hardness of $\text{CTL}_{\text{mon}}\text{-MC}(T)$ for $\text{EX} \in T$ follows.

Second, consider $\text{EF} \in T$. We have to extend our Kripke structure to contain information about the depth of the corresponding gate. We may assume w.l.o.g. that C is encoded such that each gate contains an additional counter holding the distance to the output gate (which is equal to the number of the layer it is contained in, cf. [15]). We extend η to include this distance i , $1 \leq i \leq \ell$ into ‘‘depth-propositions’’ d_i as in the proof of Proposition III.4. Denote this modified Kripke structure by K' . Further, we define $(\varphi'_i)_{0 \leq i \leq \ell}$ as

$$\varphi'_i := \begin{cases} d_\ell \wedge t, & \text{if } i = \ell, \\ \text{EF}(d_{i+1} \wedge \varphi'_{i+1}), & \text{if } i \text{ is even,} \\ \bigwedge_{i=1,2} \text{EF}(d_{i+1} \wedge i \wedge \varphi'_{i+1}), & \text{if } i \text{ is odd.} \end{cases}$$

Redefining the reduction f as $(C, x) \mapsto (K', v_0, \varphi'_0)$ yields the LOGCFL-hardness of $\text{CTL}_{\text{mon}}\text{-MC}(T)$ for $\text{EF} \in T$.

Third, let $\text{AX} \in T$. LOGCFL-hardness follows with similar arguments as for $\text{EX} \in T$, using that LOGCFL is closed under complement. An analogous argument works for the

case $\text{AG} \in T$. The remaining fragments are even P-complete by Prop. III.6. \blacksquare

Using Lemma II.1 we obtain LOGCFL-hardness of $\text{CTL}_{\text{pos}}\text{-MC}(T)$ for all nonempty sets T of CTL-operators.

In the absence of CTL-operators, the lower bound for the model checking problem again follows from the lower bound for evaluating monotone propositional formulae. This problem is known to be hard for NC^1 [16], [18].

C. The Power of Negation

We will now show that model checking for the fragments $\text{CTL}_{\text{a.n.}}$ and CTL_{pos} is computationally equivalent to model checking for CTL_{mon} , for any set T of CTL-operators. Since we consider cd-reductions, this is not immediate.

From Lemma II.1 it follows that the hardness results for $\text{CTL}_{\text{mon}}\text{-MC}(T)$ also hold for $\text{CTL}_{\text{a.n.}}\text{-MC}(T)$ and $\text{CTL}_{\text{pos}}\text{-MC}(T)$. Moreover, the algorithms for $\text{CTL}_{\text{pos}}\text{-MC}(T)$ also work for $\text{CTL}_{\text{mon}}\text{-MC}(T)$ and $\text{CTL}_{\text{a.n.}}\text{-MC}(T)$ without using more computation resources. Both observations together yield the same completeness results for all CTL-fragments with restricted negations.

Theorem III.8. *Let T be any set of CTL-operators. Then $\text{CTL}_{\text{mon}}\text{-MC}(T)$, $\text{CTL}_{\text{a.n.}}\text{-MC}(T)$, and $\text{CTL}_{\text{pos}}\text{-MC}(T)$ are*

- NC^1 -complete if T is empty,
- LOGCFL-complete if $\emptyset \subsetneq T \subseteq \{\text{EX}, \text{EF}\}$ or $\emptyset \subsetneq T \subseteq \{\text{AX}, \text{AG}\}$,
- P-complete otherwise.

Moreover, the problems $\text{CTL}_{\text{mon}}\text{-MC}(T)$, $\text{CTL}_{\text{a.n.}}\text{-MC}(T)$, and $\text{CTL}_{\text{pos}}\text{-MC}(T)$ are equivalent w.r.t. cd-reducibility.

This equivalence extends Lemma II.1. We remark that this equivalence is not straightforward. Simply applying de Morgan’s laws to transform one problem into another requires counting the number of negations on top of \wedge - and \vee -connectives. This counting cannot be achieved by an AC^0 -circuit and does not lead to the aspired reduction. Here we obtain equivalence of the problems as a consequence of our generic hardness proofs in Sect. III-B.

IV. MODEL CHECKING EXTENSIONS OF CTL

What CTL lacks in practice is the ability to express fairness properties. To address this shortcoming, Emerson and Halpern introduced ECTL in [7].

ECTL extends CTL with the $\tilde{\text{F}}$ -operator, which states that for every moment in the future, the enclosed formula will eventually be satisfied again: for a Kripke structure K , a path $x = (x_1, x_2, \dots)$, and a path formula χ

$$K, x \models \tilde{\text{F}}\chi \quad \text{iff} \quad K, x^i \models \text{F}\chi \quad \text{for all } i \in \mathbb{N}.$$

The dual operator $\tilde{\text{G}}$ is defined analogously. As for CTL, model checking for ECTL is known to be tractable. Moreover, our next result shows that even for all fragments, model checking for ECTL is not harder than for CTL.

Theorem IV.1. Let T be a set of temporal operators. Then $\text{ECTL-MC}(T) \equiv_{\text{cd}} \text{CTL-MC}(T')$ and $\text{ECTL}_{\text{pos}}\text{-MC}(T) \equiv_{\text{cd}} \text{CTL}_{\text{pos}}\text{-MC}(T')$, where T' is obtained from T by substituting \bar{F}, \bar{G} with F, G .

Another extension of CTL is CTL^+ where Boolean combinations of pure temporal operators are allowed in the scope of path quantifiers. In contrast to CTL, model checking for CTL^+ is not tractable, but Δ_2^{P} -complete [11]. Below we classify the complexity of model checking for both full and positive fragments of CTL^+ .

Theorem IV.2. Let T be a set of temporal operators. $\text{CTL}^+\text{-MC}(T)$ is

- NC^1 -complete if $T \subseteq \{A, E\}$ or $T = \{X\}$,
- P-complete if $\{X\} \subsetneq T \subseteq \{A, E, X\}$, and
- Δ_2^{P} -complete otherwise.

For the positive fragments of CTL^+ we obtain:

Theorem IV.3. Let T be a set of temporal operators. $\text{CTL}_{\text{pos}}^+\text{-MC}(T)$ is

- NC^1 -complete if $T \subseteq \{A, E\}$ or $T = \{X\}$,
- LOGCFL-complete if $T = \{A, X\}$ or $T = \{E, X\}$,
- P-complete if $T = \{A, E, X\}$,
- NP-complete if $E \in T$, $A \notin T$, and T contains a pure temporal operator aside from X ,
- coNP-complete if $A \in T$, $E \notin T$, and T contains a pure temporal operator aside from X , and
- Δ_2^{P} -complete otherwise.

Finally, ECTL^+ is the combination of ECTL and CTL^+ . For its model checking problem we obtain:

Corollary IV.4. Let T be a set of temporal operators. Then $\text{ECTL}^+\text{-MC}(T) \equiv_{\text{cd}} \text{CTL}^+\text{-MC}(T')$ and $\text{ECTL}_{\text{pos}}^+\text{-MC}(T) \equiv_{\text{cd}} \text{CTL}_{\text{pos}}^+\text{-MC}(T')$, where T' is obtained from T by substituting \bar{F}, \bar{G} with F, G .

V. CONCLUSION

We have shown (Theorem III.2) that model checking for $\text{CTL}_{\text{pos}}(T)$ is already P-complete for most fragments of CTL. Only for some weak fragments, model checking becomes easier: if $T \subseteq \{EX, EF\}$ or $T \subseteq \{AX, AG\}$, then $\text{CTL}_{\text{pos}}\text{-MC}(T)$ is LOGCFL-complete. In the case that no CTL-operators are used, NC^1 -completeness of evaluating propositional formulae applies. As a direct consequence (Theorem III.1), model checking for $\text{CTL}(T)$ is P-complete for every nonempty T . This shows that for the majority of interesting fragments, model checking $\text{CTL}(T)$ is inherently sequential and cannot be sped up by using parallelism.

While all the results above can be transferred to ECTL (Theorem IV.1), CTL^+ and ECTL^+ exhibit different properties. For both logics, the general model checking problem was shown to be complete for Δ_2^{P} in [11]. Here we proved that model checking fragments of $\text{CTL}^+(T)$ and $\text{ECTL}^+(T)$ for $T \subseteq \{A, E, X\}$ remains tractable, while the existential




Figure 2. Complexity of $\text{CTL}_{\text{pos}}\text{-MC}(T)$ for all sets T of CTL-operators (depicted as a “finite automaton” where states indicate completeness results and arrows indicate an increase of the set of CTL-operators).

and the universal fragments of $\text{CTL}_{\text{pos}}^+(T)$ and $\text{ECTL}_{\text{pos}}^+(T)$ containing temporal operators other than X are complete for NP and coNP, respectively.

Instead of restricting only the use of negation as done in this paper, one might go one step further and restrict the allowed Boolean connectives in an arbitrary way. One might, e. g., allow the exclusive-OR as the only propositional connective. This has been done for the case of linear temporal logic LTL in [9], where the complexity of $\text{LTL-MC}(T, B)$ for an arbitrary set T of temporal operators and B of propositional connectives was studied. We think that a corresponding study for CTL (or CTL^*) is an interesting topic for further research. E.g., restricting the Boolean connectives to only one of the functions AND or OR leads to many NL-complete fragments, but a full classification is still open. The complexity of the corresponding satisfiability problems $\text{CTL-SAT}(T, B)$ and $\text{CTL}^*\text{-SAT}(T, B)$ has been completely determined in [19].

Another interesting topic for future research is to determine the exact complexity of model checking problems where either the formula is fixed and only the Kripke structure varies (system complexity, cf. [20]), or where the Kripke structure is fixed and the formula varies (specification complexity). In contrast, our results in this paper concern the joint complexity which is measured in terms of both inputs.

ACKNOWLEDGMENT

This paper is supported in part by grants DFG VO 630/6-1, DAAD-ARC D/08/08881 and BC-ARC 1323.

REFERENCES

- [1] E. Clarke, E. A. Emerson, and A. Sistla, “Automatic verification of finite-state concurrent systems using temporal logic specifications,” *ACM Transactions on Programming Languages and Systems*, vol. 8, no. 2, pp. 244–263, 1986.
- [2] P. Schnoebelen, “The complexity of temporal logic model checking,” in *Advances in Modal Logic*, vol. 4, 2002, pp. 393–436.

- [3] A. Sistla and E. Clarke, “The complexity of propositional linear temporal logics,” *Journal of the ACM*, vol. 32, no. 3, pp. 733–749, 1985.
- [4] N. Markey, “Past is for free: on the complexity of verifying linear temporal properties with past,” *Acta Informatica*, vol. 40, no. 6-7, pp. 431–458, 2004.
- [5] A. Pnueli, “The temporal logic of programs,” in *Proc. 18th Symposium on Foundations of Computer Science*. IEEE Computer Society Press, 1977, pp. 46–57.
- [6] E. A. Emerson and E. M. Clarke, “Using branching time temporal logic to synthesize synchronization skeletons,” *Science of Computer Programming*, vol. 2, no. 3, pp. 241–266, 1982.
- [7] E. A. Emerson and J. Y. Halpern, ““Sometimes” and “not never” revisited: on branching versus linear time temporal logic,” *Journal of the ACM*, vol. 33, no. 1, pp. 151–178, 1986.
- [8] E. A. Emerson and C.-L. Lei, “Modalities for model checking: Branching time logic strikes back,” *Science of Computer Programming*, vol. 8, no. 3, pp. 275–306, 1987.
- [9] M. Bauland, M. Mundhenk, T. Schneider, H. Schnoor, I. Schnoor, and H. Vollmer, “The tractability of model checking for LTL: the good, the bad, and the ugly fragments,” in *Methods for Modalities 5 (M4M5) 2007*, ser. Electronic Notes in Theoretical Computer Science, vol. 231, 2009, pp. 277–292.
- [10] L. Kuhtz and B. Finkbeiner, “LTL Path Checking is Efficiently Parallelizable,” *International Colloquium on Automata, Languages and Programming*, 2009.
- [11] F. Laroussinie, N. Markey, and P. Schnoebelen, “Model Checking CTL+ and FCTL is Hard,” *Proc. Foundations of Software Science and Computation Structure (FOSSACS’2001), Genova, Italy*, vol. 2030, pp. 318–331, 2001. [Online]. Available: <http://citeseer.ist.psu.edu/laroussinie01model.html>
- [12] F. Laroussinie, “About the expressive power of CTL combinators,” *Information Processing Letters*, vol. 54, no. 6, pp. 343–345, 1995.
- [13] C. H. Papadimitriou, *Computational Complexity*. Reading, MA: Addison-Wesley, 1994.
- [14] K. Regan and H. Vollmer, “Gap-languages and log-time complexity classes,” *Theoretical Computer Science*, vol. 188, pp. 101–116, 1997.
- [15] H. Vollmer, *Introduction to Circuit Complexity – A Uniform Approach*, ser. Texts in Theoretical Computer Science. Berlin Heidelberg: Springer Verlag, 1999.
- [16] S. R. Buss, “The Boolean formula value problem is in ALOGTIME,” in *Proceedings 19th Symposium on Theory of Computing*. ACM Press, 1987, pp. 123–131.
- [17] W. L. Ruzzo, “Tree-size bounded alternation,” *Journal of Computer and System Sciences*, vol. 21, pp. 218–235, 1980.
- [18] H. Schnoor, “The complexity of the Boolean formula value problem,” Theoretical Computer Science, University of Hannover, Tech. Rep., 2005.
- [19] A. Meier, M. Mundhenk, M. Thomas, and H. Vollmer, “The complexity of satisfiability for fragments of CTL and CTL*,” in *Proceedings of the 2nd Workshop on Reachability Problems in Computational Models (RP 2008)*, ser. Electronic Notes in Theoretical Computer Science, vol. 223, 2008, pp. 201–213.
- [20] O. Kupferman, M. Y. Vardi, and P. Wolper, “An automata-theoretic approach to branching-time model checking,” *Journal of the ACM*, vol. 47, no. 2, pp. 312–360, 2000.