On the Complexity of Temporal-Logic Path Checking*

Daniel Bundala and Joël Ouaknine

Department of Computer Science, University of Oxford
Wolfson Building, Parks Road, Oxford, OX1 3QD, UK

Abstract. Given a formula in a temporal logic such as LTL or MTL, a fundamental problem is the complexity of evaluating the formula on a given finite word. For LTL, the complexity of this task was recently shown to be in NC [9]. In this paper, we present an NC algorithm for MTL, a quantitative (or metric) extension of LTL, and give an AC^1 algorithm for UTL, the unary fragment of LTL. At the time of writing, MTL is the most expressive logic with an NC path-checking algorithm, and UTL is the most expressive fragment of LTL with a more efficient path-checking algorithm than for full LTL (subject to standard complexity-theoretic assumptions). We then establish a connection between LTL path checking and planar circuits, which we exploit to show that any further progress in determining the precise complexity of LTL path checking would immediately entail more efficient evaluation algorithms than are known for a certain class of planar circuits. The connection further implies that the complexity of LTL path checking depends on the Boolean connectives allowed: adding Boolean exclusive or yields a temporal logic with P-complete path-checking problem.

1 Introduction

One of the most fundamental problems in the fields of testing and verification is the *path-checking problem*: determine whether a given observation[1] of a system satisfies a given specification drawn from a fixed ambient logic. The complexity of this problem plays a key role in the design and analysis of offline monitoring and runtime verification procedures [6,12]. The path-checking problem also appears in testing [1] and in Monte-Carlo-based probabilistic verification [14].

Although the problem is simply stated, determining its precise complexity can prove to be quite challenging. The case of LTL was investigated more than a decade ago [5,13], and at the time is was conjectured that the straightforward polynomial-time dynamic-programming algorithm is not optimal.[2] And indeed, using reductions to planar circuits and tree-contraction algorithms, it was recently proved [9] that LTL path checking allows an efficient parallel algorithm

* Full version of the paper is available as [3].

[1] In this paper, all observations (paths, traces, words, etc.) considered are finite.

[2] The best known lower bound for LTL path checking is NC^1, which crudely arises from the NC^1-hardness of mere Boolean formula evaluation.

J. Esparza et al. (Eds.): ICALP 2014, Part II, LNCS 8573, pp. 86–97, 2014.

and lies in NC—in fact, in $\mathsf{AC}^1[\mathsf{logDCFL}]$. (This seminal result was rewarded by the ICALP 2009 best-paper award.) More recently, this work was extended to a very restricted metric extension of LTL, in which only temporal operators of the form $\mathrm{U}_{\leq b}$ are allowed [10].

In this paper, we give an algorithm for full Metric Temporal Logic (MTL) with the same complexity—$\mathsf{AC}^1[\mathsf{logDCFL}]$—known algorithm for LTL.

We reprise the strategy, introduced in [9], to represent temporal operators using a special class of planar monotone circuits, together with a generic algorithm [4] as a subroutine to evaluate those circuits. Such circuits have a very special form, which led the authors of [9] to ask whether the complexity of the path-checking algorithm can be improved by devising specialised circuit-evaluation algorithms. In this paper, we present evidence to the contrary, by showing that the evaluation of circuits drawn from a class of planar circuits studied in [11] is reducible to LTL path checking; any further progress in determining the precise complexity of the latter would therefore immediately entail more efficient evaluation algorithms than are known for this class of planar circuits. It is worth pointing out that augmenting this class of planar circuits with NOT gates makes the evaluation problem P-complete [7]. It follows that the complexity of path checking is sensitive to non-monotone connectives, as allowing Boolean exclusive-or in formulae enables the evaluation of circuits from this augmented class, and is therefore itself P-complete.

An examination of the algorithmic constructions of [9] shows that the most intricate parts arise in handling the Until operator. In this paper, we show that the removal of binary operators from the logic, yielding Unary Temporal Logic (UTL), leads to a much simpler path-checking problem, enabling us to devise an AC^1 algorithm for UTL path checking.

At the time of writing, our results provide (i) the most expressive known extension of LTL with an NC path-checking algorithm (MTL), (ii) the simplest known extension of LTL with a strictly harder path-checking problem (LTL + Xor), and (iii) the most expressive known fragment of LTL with a strictly more efficient path-checking algorithm than for full LTL (UTL).[3]

2 Preliminaries

We denote Boolean true and false by $\top$ and $\bot$, respectively. The set $\{\bot, \top\}$ is denoted by $\mathbb{B}$. A vector $v \in \mathbb{B}^n$ is ***downward monotone*** if $v(i+1) = \top \implies v(i) = \top$. It is ***upward monotone*** if $v(i-1) = \top \implies v(i) = \top$. A vector is ***monotone*** if it is upward or downward monotone. The set of monotone vectors is denoted by $\mathbb{M}$.

Temporal Logics: Let AP be a set of atomic propositions, $p \in \mathrm{AP}$ and $I \subseteq \mathbb{R}_{\geq 0}$ be an interval with endpoints in $\mathbb{N} \cup \{\infty\}$. The formulae of ***Metric Temporal Logic*** (MTL) are defined recursively as follows.

$$\varphi = p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathrm{X}_I\, \varphi \mid \mathrm{Y}_I\, \varphi \mid \varphi\, \mathrm{U}_I\, \varphi \mid \varphi\, \mathrm{S}_I\, \varphi \mid \varphi\, \mathrm{R}_I\, \varphi \mid \varphi\, \mathrm{T}_I\, \varphi$$

[3] Subject to standard complexity-theoretic assumptions.

All logics and results presented in this paper apply to temporal logics with past temporal operators. Note that negation is applied only to atomic propositions. Other operators are expressible using the following semantic equalities: $\mathrm{F}_I\varphi = \top\,\mathrm{U}_I\,\varphi$, $\mathrm{G}_I\varphi = \neg\,\mathrm{F}_I\,\neg\varphi$, $\varphi\,\mathrm{R}_I\,\psi = \neg(\neg\varphi\,\mathrm{U}_I\,\neg\psi)$ and $\varphi\,\mathrm{T}_I\,\psi = \neg(\neg\varphi\,\mathrm{S}_I\,\neg\psi)$. ***Linear Temporal Logic*** (LTL) is the subset of MTL in which I is always $[0,\infty)$ (and is omitted). The fragment UTL of LTL consists of all Boolean connectives and unary $(\mathrm{X},\mathrm{F},\mathrm{G})$ temporal operators and their past duals.

A ***trace*** π over AP of length n is a function $\pi : \{1,\ldots,n\}\times\mathrm{AP}\to\mathbb{B}$ assigning a truth value to every $p\in\mathrm{AP}$ at every index. We identify $p\in\mathrm{AP}$ with a vector in $\mathbb{B}^n$ and use $p(i)=\top$ if $\pi(i,p)=\top$. The proposition that is true only in the interval $[i,j]$ and false otherwise is denoted by $\chi_{i,j}$, i.e., $\chi_{i,j}(k)=\top$ if $i\le k\le j$ and $\chi_{i,j}(k)=\bot$ otherwise. To evaluate MTL formulae on π, we further associate with π a sequence of strictly-increasing ***timestamps*** $t_1<\ldots<t_n$.

Given an MTL formula φ and index $1\le i\le n$, the satisfaction relation $\pi,i\models\varphi$ is defined recursively as follows.

$$
\begin{array}{ll}
\pi,i\models p & \text{if } p(i)=\top\\
\pi,i\models\varphi_1\wedge\varphi_2 & \text{if } \pi,i\models\varphi_1 \text{ and } \pi,i\models\varphi_2\\
\pi,i\models\varphi_1\vee\varphi_2 & \text{if } \pi,i\models\varphi_1 \text{ or } \pi,i\models\varphi_2\\
\pi,i\models X_I\varphi & \text{if } i+1<n\wedge t_{i+1}-t_i\in I\wedge\pi,i+1\models\varphi\\
\pi,i\models Y_I\varphi & \text{if } i>1 \text{ and } t_i-t_{i-1}\in I \text{ and } \pi,i-1\models\varphi\\
\pi,i\models\varphi_1 U_I\varphi_2 & \text{if } \exists j\,.\,(i\le j\le n)\wedge\left(\begin{array}{l}\pi,j\models\varphi_2\\ t_j-t_i\in I\\ \forall k\,.\,i\le k<j\implies\pi,k\models\varphi_1\end{array}\right)\\
\pi,i\models\varphi_1 S_I\varphi_2 & \text{if } \exists j\,.\,(i\ge j\ge 1)\wedge\left(\begin{array}{l}\pi,j\models\varphi_2\\ t_i-t_j\in I\\ \forall k\,.\,i\ge k>j\implies\pi,k\models\varphi_1\end{array}\right)
\end{array}
$$

This paper studies the complexity of evaluating a given formula on a given trace.

Definition 1. *The* **path-checking problem** *for logic $\mathcal{L}$ is to determine, given a trace π and a formula φ of $\mathcal{L}$, whether $\pi,1\models\varphi$.*

Let φ be an MTL formula. Working from the smallest subformulae and using the above definitions to tabulate the values $\pi,i\models\psi$ for every i and subformula ψ yields a polynomial dynamic-programming algorithm evaluating φ on π.

Theorem 1 ([13]). *The path-checking problem for* MTL *is in* P.

Given a trace π and formula φ, we represent the value of φ on π as the vector $v\in\mathbb{B}^n$ such that $v(i)=\top$ if and only if $\pi,i\models\varphi$. We further represent LTL temporal operators as functions over vectors written in infix notation. For example, $\mathrm{U}:\mathbb{B}^n\times\mathbb{B}^n\to\mathbb{B}^n$ is a function such that $(p\,\mathrm{U}\,q)(i)=\top$ if and only if there is $i\le j\le n$ such that $q(j)=\top$ and $p(k)=\top$ for all $i\le k<j$.

A ***formula context*** $\varphi(X)$ is a formula with one occurrence of a proposition replaced by a variable X. If $\psi(X)$ is another formula context then $(\varphi\circ\psi)(X)$ is the context obtained by substituting $\psi(X)$ for X in $\varphi(X)$. If $q\in\mathrm{AP}$ is a proposition

then $\varphi(q)$ is obtained by substituting q for X. For example, $((p\mathrm{U}X)\circ(X\mathrm{S}q))(r) = (p\mathrm{U}(X\mathrm{S}q))(r) = p\mathrm{U}(r\mathrm{S}q)$. Composing formula contexts increases the size linearly as a formula context contain only one occurrence of X.

Circuits: A ***Boolean circuit*** (C, δ) consists of a set of ***gates*** C and a ***predecessor*** function $\delta : C \to \mathcal{P}(C)$. The type of a gate is either OR, AND, NOT, ID, ONE or ZERO. If c is of type τ and $\delta(c) = \{c_1, \ldots, c_n\}$ then we write $c = (\tau, c_1, \ldots, c_n)$. If $d \in \delta(c)$ then we say c ***depends*** on d or that there is a ***wire*** from d to c. The ONE and ZERO gates provide constants inputs. A gate is an ***input*** gate if it does not have a predecessor. A gate is an ***output*** gate if it is not a predecessor of any other gate. A circuit is ***monotone*** if it has no NOT gates. It is ***planar*** if the underlying DAG is planar. In this paper, all edges (wires) are straight-line segment and so a ***planar embedding*** is induced by a function $\gamma : C \to \mathbb{R}^2$ assigning a point in the plane to every gate.

A circuit is ***layered*** if it can be partitioned into ***layers*** $C_0, \ldots, C_n$ such that each wire goes from C_i to C_{i+1} for some i. Thus, C_0 contains only input gates. A layered circuit is ***stratified*** if all input gates appear in C_0. A circuit is ***upward planar*** if there is a planar embedding such that every edge monotonically increases in the upward direction—the direction of the evaluation of C. A circuit is ***upward layered (stratified)*** if it is both upward planar and layered (stratified). Each layer C_i of an upward-layered circuit consists of gates $\alpha_{i,j}$ in the left-to-right ordering. Each $\alpha_{i,j}$ depends on a contiguous block $\alpha_{i-1,l}, \ldots, \alpha_{i-1,r}$ layer below and the wires do not cross: if $\alpha_{i,j}$ depends on $\alpha_{i-1,q}$ and $\alpha_{i,k}$ depends on $\alpha_{i-1,r}$ then $j \leq k \iff q \leq r$. Fig. 3 shows upward stratified monotone circuits.

Given a circuit with one output gate, the ***circuit value problem***, abbreviated as ***CVP***, is the problem of determining the value of the output gate.

Complexity Classes: The class $\mathsf{logDCFL}$ consists of problems that are logspace many-one reducible to deterministic context-free languages. Equivalently, it is the class of problems decidable by a deterministic logspace Turing machine equipped with a stack and terminating in polynomial time. The circuit class AC^i for $i \in \mathbb{N}$ consists of problems decidable by polynomial-size unbounded fan-in circuits of depth $\log^i$. All circuits in this paper are ***uniform***—can be generated by a deterministic logspace Turing machine. Given a problem S and a complexity class C, we write $S \in \mathsf{AC}^1[C]$ if there is a family of AC^1 circuits with additional unbounded fan-in C-oracle gates that decide S. It is known that

$$\mathrm{L} \subseteq \mathsf{logDCFL} \subseteq \mathsf{AC}^1 \subseteq \mathsf{AC}^1[\mathsf{logDCFL}] \subseteq \mathsf{AC}^2 \subseteq \cdots \subseteq \mathsf{AC}^i \subseteq \mathsf{AC}^{i+1} \subseteq \cdots \subseteq \mathsf{P}$$

and that CVP for upward-stratified circuits is P-complete [7], CVP for monotone upward-stratified circuits is in $\mathsf{logDCFL}$ [4] and that CVP for monotone upward-layered circuits is in $\mathsf{AC}^1[\mathsf{logDCFL}]$ [11].

Tree Contraction: Let $T = (V, E)$ be a binary tree, the tree contraction algorithm [8] reduces T to a single node using a sequence of tree contraction steps. Let $l \in T$ be a leaf, p be its parent and s its sibling[4]. A tree contraction step

[4] If l does not have a sibling then we take s to be a fresh node.

collapses the triple (l, p, s) into a single node. Formally, a new tree $T' = (V', E')$ is obtained from T as follows: $V' = V \setminus \{l, p\}$

$$E' = \begin{cases} E \setminus \{(p,l),(p,s)\} & \text{if } p \text{ is the root of } T \\ (E \setminus \{(p,l),(p,s),(q,p)\}) \cup \{q,s\} & \text{otherwise } (q \text{ is the parent of } p) \end{cases}$$

Note that a contraction step is local and hence multiple non-interfering contractions can be performed in parallel. A tree contraction algorithm using only $\lceil \log n \rceil$ parallel steps exists [8]. Further, this algorithm can be implemented in AC^1.

Let φ be an LTL formula and π a trace. A tree contraction algorithm evaluating φ on π was given in [9]. The tree T used in [9] is the parse tree of φ. The leaves of T correspond to the atomic propositions and the internal nodes to Boolean or temporal operators. Each contraction step (l, p, s) partially evaluates the operator associated with p.

For example, suppose that the formula rooted at p is $\psi \,\mathrm{U}\, q$ where q is a proposition. Even if the value of ψ is unknown, we can still make some inferences. E.g., if $q(i) = \top$ then $(\psi \,\mathrm{U}\, q)(i) = \top$. If the last value $q(|\pi|) = \bot$ then $(\psi \,\mathrm{U}\, q)(|\pi|) = \bot$ and so on. The contraction step removes the nodes for ψ and U and then labels the node s by the partial evaluation of the function $(X \,\mathrm{U}\, q) \circ \psi$. It was shown in [9] how to represent, manipulate and evaluate these functions efficiently. When a subformula ψ is fully collapsed into a single node then the associated function is fully evaluated and the node is labelled by the constant $(\psi(1), \ldots, \psi(|\pi|)) \in \mathbb{B}^{|\pi|}$. The contraction algorithm eventually reduces the tree into a single node, which is labelled by $(\varphi(1), \ldots, \varphi(|\pi|)) \in \mathbb{B}^{|\pi|}$.

In general, a tree-contraction algorithm can evaluate a function f on a tree; each contraction step partially evaluating f on a subtree. In this paper, the evaluation is done as follows. Let $\mathcal{C}$ be the set of constants and $\mathcal{F}$ be a collection, closed under composition, of admissible functions $f : \mathcal{C} \to \mathcal{C}$.

- A constant $c_v \in \mathcal{C}$ is attached to every leaf v of T. The values of c_v for the initial leaves are given as a part of the input.
- A function $f_v \in \mathcal{F}$ is attached to every node v of T. Initially, f_v is the identity function.
- A tree contraction of (l, p, s) first builds $f' \in \mathcal{F}$ (depending on c_l and p) implementing the partial evaluation on p. Let $f'' = f_p \circ f'$. If s is a leaf then c_s is replaced by $f''(c_s)$. Otherwise, f_s is replaced by $f'' \circ f_s$.

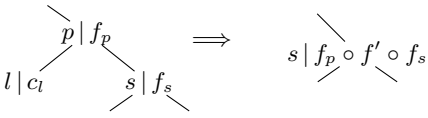

Fig. 1. An example of a tree contraction step

The output of the algorithm is the constant attached to the single remaining node. If each contraction step and admissible functions are in the complexity class C then, by [8], the contraction algorithm calculating c_{root} is in $\mathsf{AC}^1[C]$.

A tree contraction algorithm for LTL path checking [9] runs in $\mathsf{AC}^1[\mathsf{logDCFL}]$. Constants $\mathcal{C} = \mathbb{B}^n$ denote the truth values of propositions and subformulae. Functions $\mathcal{F}$ are represented by upward stratified circuits with n input and n output gates (***transducer circuits***), which are closed under composition [9] and their evaluation and composition is in logDCFL [2]. For a fixed $s \in \mathbb{B}^n$, [9] gives transducer circuits for $s \wedge x, s \vee x, s\,\mathrm{U}\,x$, and $s\,\mathrm{R}\,x$ as the functions of $x \in \mathbb{B}^n$. In Section 4, we give transducer circuits for MTL temporal operators.

3 Reduction from Upward Layered CVP to LTL Path Checking

Given an upward layered monotone circuit C with n gates and m wires we show how to build an LTL formula φ over at most $2n$ propositions and a trace π of length $|\pi| \leq m$ such that C evaluates to $\top$ if and only if $\pi \models \varphi$.

Denote the layers of C by $C_0, \ldots, C_k$ and the size of each C_i by n_i. Let $\alpha_{i,j}$ be the gates in C_i in the left-to-right order in the upward planar embedding of C. For each layer, we partition the trace into blocks—each of which stores the outputs of a gate in the layer. Fig. 2 shows a valid partitioning. In the figure, gate a occupies block $[1,1]$, gate e occupies $[3,5]$, gate g occupies $[1,7]$, etc.

In general, a valid partitioning consists of a trace π and intervals $v(i,j)$ associated with each gate $\alpha_{i,j}$ such that $v(i,j)$ overlaps precisely with the blocks of the gates the gate $\alpha_{i,j}$ depends on. Formally,

- intervals $v(i,1), v(i,2), \ldots, v(i,n_i)$ are disjoint and partition $[1, |\pi|]$ for every i,
- if $\alpha_{i+1,j}$ depends on $\alpha_{i,p}, \alpha_{i,p+1}, \ldots, \alpha_{i,q}$ then $v(i+1,j) \subseteq \cup_{r=p,\ldots,q} v(i,r)$ and $v(i+1,j)$ overlaps with each $v(i,r)$ for $p \leq r \leq q$,

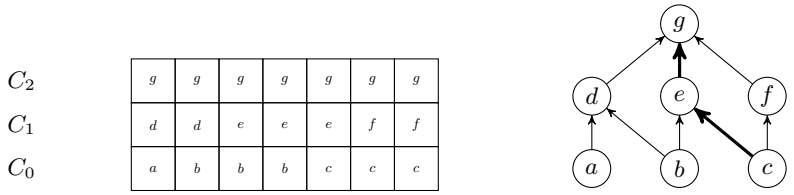

Fig. 2. An upward layered circuit (on the right) with its partition (on the left). The path π for the gate labelled e is highlighted.

Suppose we are given a valid partitioning. Then for $i > 0$ and every $1 \leq j \leq n_i$ we build a formula context $\varphi_{i,j}$ mimicking the evaluation of the gate $\alpha_{i,j}$.

For example, suppose that the gate e in Fig. 2 is an OR gate and the values of the block in the first layer is $r = (a, b, b, b, c, c, c) \in \mathbb{B}^7$ for some $a, b, c \in \mathbb{B}$. Recall that $(\varphi\,\mathrm{U}\,\psi)(i) = \psi(i) \vee (\varphi(i) \wedge (\varphi\,\mathrm{U}\,\psi)(i+1))$. Hence, if $\varphi(i) = \bot$ then $(\varphi\,\mathrm{U}\,\psi)(i) = \psi(i)$ and if $\varphi(i) = \top$ then $(\varphi\,\mathrm{U}\,\psi)(i) = \psi(i) \vee (\varphi\,\mathrm{U}\,\psi)(i+1)$. Further recall that $\chi_{i,j}$ is a proposition that is true on $[i,j]$ and false otherwise. Hence, $(\chi_{3,4}\,\mathrm{U}\,r)(1) = a, (\chi_{3,4}\,\mathrm{U}\,r)(2) = b$ and $(\chi_{3,4}\,\mathrm{U}\,r)(5,6,7) = c$. Also, $(\chi_{3,4}\,\mathrm{U}\,r)(4) = r(4) \vee (\chi_{3,4}\,\mathrm{U}\,r)(5) = b \vee c$. Finally, $(\chi_{3,4}\,\mathrm{U}\,r)(3) = r(3) \vee (\chi_{3,4}\,\mathrm{U}$

$r)(4) = b \vee (b \vee c) = b \vee c$. So $\chi_{3,4}\,\mathrm{U}\,r = (a, b, b \vee c, b \vee c, c, c, c)$. Performing a similar calculation backwards, we get $\chi_{4,5}\,\mathrm{S}\,(\chi_{3,4}\,\mathrm{U}\,r) = (a, b, b \vee c, b \vee c, b \vee c, c, c)$ which gives the value of block e in Fig. 2 and leaves other blocks unchanged.

Denote the type of $\alpha_{i,j}$ by τ and the left and the right endpoint of $v(i,j)$ by l and r, respectively. Then $\varphi_{i,j}$ is constructed as follows:

- If $\tau = \text{ONE}$ then $\varphi_{i,j}(X) = \chi_{l,r} \vee X$.
- If $\tau = \text{ZERO}$ then $\varphi_{i,j}(X) = (\neg\chi_{l,r}) \wedge X$.
- If $\tau = \text{ID}$ then $\varphi_{i,j}(X) = X$.
- If $\tau = \text{OR}$ then $\varphi_{i,j}(X) = \chi_{l+1,r}\,\mathrm{S}\,(\chi_{l,r-1}\,\mathrm{U}\,X)$.
- If $\tau = \text{AND}$ then $\varphi_{i,j}(X) = \chi_{l+1,r}\,\mathrm{T}\,(\chi_{l,r-1}\,\mathrm{R}\,X)$.

It can be shown that the formula context $\varphi_{i,j}$ updates the block $v(i,j)$ and leaves the other blocks unchanged. Hence, the formula context $\psi_i(X) = \varphi_{i,1} \circ \varphi_{i,2} \circ \cdots \circ \varphi_{i,n_i}$ evaluates the i-th layer C_i of C.

Formally, for each layer C_i let $r_i \in \mathbb{B}^n$ be a proposition such that $r_i(k) = \top$ if $k \in v(i,j)$ for some j and $\alpha_{i,j}$ evaluates to $\top$ and $r_i(k) = \bot$, otherwise. Then, the formula $\varphi = (\psi_k \circ \psi_{k-1} \circ \cdots \circ \psi_1)(r_0)$ computes the output of the circuit.

Lemma 1. *Let ψ_i, φ be as above. Then $\psi_i(r_{i-1}) = r_i$ and $\varphi(r_0)(1) = \top$ if and only if C evaluates to $\top$. Moreover, φ can be built in L.*

Finally, we show how to devise $v(i,j)$'s – the partitioning of the trace. Without loss of generality, connecting to a gate in the previous layer if necessary, we assume that all ONE and ZERO gates not in C_0 have at least one predecessor.

Given a gate $\alpha_{i,j}$ there is unique rightmost gate in the layer C_{i+1} that $\alpha_{i,j}$ is connected to by a wire. Now, start at $\alpha_{i,j}$ and take the rightmost wires until the sink is reached. Denote the traversed path by π_u. Similarly, there is unique rightmost gate in the layer C_{i-1} that $\alpha_{i,j}$ is connected to by a wire. Start at $\alpha_{i,j}$ and take the rightmost wires going down until a gate in C_0 is reached. Denote the traversed path by π_d. Let π be the concatenation of π_d and π_u. (See Fig. 2)

Let $k_{i,j}$ be the number of wires to the left of π. A wire from $\alpha_{i,j}$ to $\alpha_{i+1,k}$ is to the left of the wire from $\alpha_{i,a}$ to $\alpha_{i+1,b}$ if $j < a$ or $k < b$. We store the output of gate $\alpha_{i,j}$ in the block $v(i,j) := [k_{i,j-1} + 1, k_{i,j} + 1]$. We use $k_{i,0} = 0$.

Fig. 2 shows a circuit and the partitioning obtained by the above procedure. The rightmost wire going up and down from e are $e \to g$ and $c \to e$, respectively. Thus, $\pi_u = e \to g$ and $\pi_d = c \to e$. The path $\pi = c \to e \to g$ is highlighted in the figure. Four wires $a \to d, b \to d, b \to e, d \to g$ are to the left of π. We associate the block $[3, 5]$ with gate e. All blocks, grouped by layers, are shown in Fig. 2.

The following lemma summarises the important properties of $k_{i,j}$'s.

Lemma 2. *Let $k_{i,j}$'s and $v(i,j)$'s be as above. Then the following hold:*

- *$k_{i,j-1} < k_{i,j}$ for every i and j,*
- *$k_{i,n_i} = k_{j,n_j}$ for every i and j,*
- *$k_{i,n_i} \leq m$ for every i,*
- *for every i and $j = 1, \ldots, n_i$ the intervals $v(i,j)$'s partition $[1, k_{i,n_i}]$,*
- *if $\alpha_{i+1,j}$ depends on $\alpha_{i,p}, \alpha_{i,p+1}, \ldots, \alpha_{i,q}$ then $v(i+1,j) \subseteq \cup_{r=p,\ldots,q} v(i,r)$ and $v(i+1,j)$ overlaps with each $v(i,r)$ for $p \leq r \leq q$,*
- *each $k_{i,j}$ can be computed in L.*

This finishes the reduction from upward-layered CVP to LTL path checking. It was shown in [9] that the latter is in $\mathsf{AC}^1[\mathsf{logDCFL}]$. Therefore:

Theorem 2. *The CVP for upward-layered monotone circuits is in* $\mathsf{AC}^1[\mathsf{logDCFL}]$.

An alternative proof of Theorem 2 already appeared in [11]. Moreover, the relationship shows that any improvement in LTL path checking would entail an improvement in the evaluation of upward-layered monotone circuits.

The above reduction assumes the monotonicity of the input circuit. However, if the target logic LTL is extended to include binary exclusive or (xor) as a connective, then evaluating NOT gates becomes possible using $\varphi_{i,j}(X) = \chi_{l,r} \oplus X$ as a formula context for NOT gate $\alpha_{i,j}$. Noting that CVP is P-complete for general (non-monotone) upward stratified circuits [7], we have the following:

Theorem 3. $\mathsf{LTL} + \mathsf{Xor}$ *path checking is* P*-complete.*

Thus, the complexity of LTL path checking depends on the monotonicity of the Boolean connectives present in the formula.

4 MTL Path Checking is Efficiently Parallelisable

We now show how the tree-contraction method of [9] extends to full MTL; giving an $\mathsf{AC}^1[\mathsf{logDCFL}]$ path-checking algorithm for MTL. By [9], summarised in Section 2, it suffices to give upward stratified transducer circuits for U_I and its duals.

Let π be the input trace with (floating-point) timestamps $t_1, \ldots, t_n$. Fix an interval I and consider the U_I operator. We now describe a dynamic-programming approach that yields planar circuits calculating $(\psi_1 \,\mathrm{U}_I\, \psi_2)(i)$. For $i \neq j$ the values $(\psi_1 \,\mathrm{U}_I\, \psi_2)(i)$ and $(\psi_1 \,\mathrm{U}_I\, \psi_2)(j)$ depend on the values of subformulae in some future intervals. In general, these intervals overlap and so naive constructions of transducer circuits are not planar. See Fig. 3 for the kind of circuits we build.

Recall, that the tree contraction is applied only to a leaf, its parent and its sibling. Let $s \in \mathbb{B}^n$ be a vector. We need to construct only circuits for $s \,\mathrm{U}_I\, \varphi$ and $\varphi \,\mathrm{U}_I\, s$ for known s. First consider the case $s \,\mathrm{U}_I\, \varphi$. (see left part of Fig. 3)

For index $1 \leq i \leq n$ the formula $(s \,\mathrm{U}_I\, \varphi)(i)$ is true if there is $j \geq i$ such that $t_j \in t_i + I$ and $\varphi(j) = \top$ and $s(k) = \top$ for all $i \leq k < j$. So let $T_i = \{j \mid t_j \in t_i + I\}$ be the set of indices in $t_i + I$. If $T_i = \emptyset$ then $(s \,\mathrm{U}_I\, \varphi)(i) = \bot$.

Otherwise, let $\mathrm{first}(i) = \min T_i$ and $\mathrm{last}(i) = \max T_i$ be the first and the last index in the interval $t_i + I$, respectively. So $(s \,\mathrm{U}_I\, \varphi)(i)$ is true if there exists $\mathrm{first}(i) \leq j \leq \mathrm{last}(i)$ such that $\varphi(j) = \top$ and $s(k) = \top$ for all $i \leq k < j$.

Now, the value of s is known. So let $\mathrm{seg}(i) = \min\{j \mid j \geq i \wedge s(j) = \bot\}$ be the first index no smaller than i such that $s(j)$ evaluates to false, i.e., $s(j)$ is true from i to $\mathrm{seg}(i) - 1$. Thus, $(s \,\mathrm{U}_I\, \varphi)(i)$ is true if there exists $\mathrm{first}(i) \leq j \leq \mathrm{last}(i)$ such that $\varphi(j) = \top$ and $j \leq \mathrm{seg}(i)$. So take $L_i = \mathrm{first}(i)$ and $R_i = \min(\mathrm{last}(i), \mathrm{seg}(i))$. Then $(s \,\mathrm{U}_I\, \varphi)(i)$ is true if $\bigvee_{L_i \leq j \leq R_i} \varphi(j)$ is true.

To build the circuits, we formalise the intuition from the left half of Fig. 3. The circuit C consists of internal gates $d_{p,q}$ and output gates o_i for each $1 \leq i \leq n$.

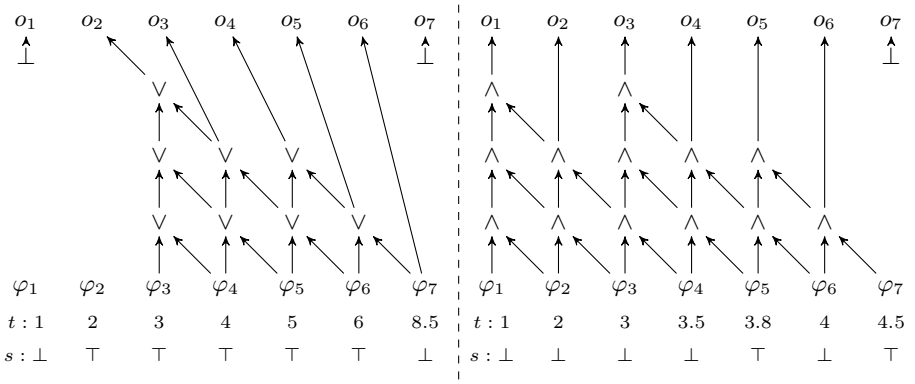

Fig. 3. Transducer circuits for $s\,\mathrm{U}_{[1,5]}\,\varphi$ and $\varphi\,\mathrm{U}_{[1,5]}\,s$. The first line below the circuits are timestamps, the second row are values of s. Note that different timestamps and s are used in the two examples. The inputs and the outputs of the circuits are denoted φ_i and o_i respectively.

Each internal gate $d_{p,q}$ calculates $\varphi_p \vee \cdots \vee \varphi_q$. Precisely, $d_{p,q}$ is present in the circuit if there is an i such that $L_i \le p \le q \le R_i$. If $p = q$ then $l(d_{p,q}) = (\mathrm{ID}, \varphi_p)$. Otherwise, $l(d_{p,q}) = (\mathrm{OR}, d_{p,q-1}, d_{p+1,q})$.

For the output gates, we define o_i so that $o_i = \bigvee_{L_i \le j \le R_i} \varphi(j) = (s\,\mathrm{U}_I\,\varphi)(i)$. Specifically, if $T_i = \emptyset$ then we set $l(o_i) = \bot$, otherwise, $l(o_i) = (\mathrm{ID}, d_{L_i,R_i})$.

An embedding $\gamma : C \to \mathbb{R}^2$ for the circuit C is $\gamma(o_i) = (i, n)$, $\gamma(\varphi_i) = (i, 0)$ and $\gamma(d_{p,q}) = (p, q - p + 1)$. Observe that $L_i \le L_{i+1}$ and $R_i \le R_{i+1}$. Hence, it cannot happen that $L_i < L_j \le R_j < R_i$ for some i and j. So the intervals may overlap but never is one properly contained in another. This ensures that the embedding is planar.

Finally, note that it is possible to compute L_i and R_i for every i in logarithmic space. Hence, the circuit construction can be carried out in logarithmic space.

Lemma 3. *Let p be any proposition. For each i, set the input φ_i of the circuit to $p(i)$. Then for each j, the value of o_j is true if and only if $(s\,\mathrm{U}_I\,p)(j)$ is true.*

We now give an analogous derivation and circuit construction for $\varphi\,\mathrm{U}_I\,s$. See the right side of Fig. 3 for an example of a resulting circuit.

For index $1 \le i \le n$ the formula $(\varphi\,\mathrm{U}_I\,s)(i)$ is true if there exists $j \ge i$ such that $t_j \in t_i + I$ and $s(j) = \top$ and $\varphi(k) = \top$ for all $i \le k < j$. Since s is known, we choose the first possible j. So let $\mathrm{limit}(i) = \min\{j \mid \mathrm{first}(i) \le j \le \mathrm{last}(i) \wedge s(j) = \top\}$ be the first j in the interval $t_i + I$ such that $s(j)$ is true.

If there is no such index then $(\varphi\,\mathrm{U}_I\,s)(i) = \bot$. Otherwise, $(\varphi\,\mathrm{U}_I\,s)(i)$ is true if $\varphi(k) = \top$ for all $i \le k < \mathrm{limit}(i)$. That is, $(\varphi\,\mathrm{U}_I\,s)(i) = \bigwedge_{i \le j < \mathrm{limit}(i)} \varphi_j$.

Now, the circuit C (see right half of Fig. 3) consists of gates $c_{p,q}$ calculating $\varphi_p \wedge \cdots \wedge \varphi_q$ and output gates o_i for $i = 1 \ldots n$. The gate $c_{p,q}$ is present in C if there is i such that $i \le p \le q < \mathrm{limit}(i)$. If $p = q$ then $l(c_{p,q}) = (\mathrm{ID}, \varphi_p)$. Otherwise, $l(c_{p,q}) = (\mathrm{AND}, c_{p,q-1}, c_{p+1,q})$.

For output, we set o_i so that $o_i = \bigwedge_{i \le j < \mathrm{limit}(i)} \varphi_j = (\varphi \mathrm{U}_I s)(i)$. If $\mathrm{limit}(i) = \infty$ then $l(o_i) = \bot$, if $\mathrm{limit}(i) = i$ then $l(o_i) = \top$ and else $l(o_i) = (\mathrm{ID}, c_{i,\mathrm{limit}(i)-1})$.

The embedding $\gamma : C \to \mathbb{R}^2$ of the circuit C is the same as above, $\gamma(o_i) = (i, n)$, $\gamma(\varphi_i) = (i, 0)$ and $\gamma(c_{p,q}) = (p, q - p + 1)$. Since, $i < j$ implies $\mathrm{limit}(i) \le \mathrm{limit}(j)$, the embedding is planar.

This finishes the construction of circuits for U_I. Circuits for the dual operators of U_I are obtained either by dualising OR and AND gates (Release operator), by performing the construction backwards in time (Since operator) or both (Trigger operator). Therefore,

Theorem 4. MTL *path checking is in* $\mathsf{AC}^1[\mathsf{logDCFL}]$.

A considerably weaker result appeared in [10], where the authors gave circuits and an $\mathsf{AC}^1[\mathsf{logDCFL}]$ algorithm only for a fragment of MTL interpreted over traces with integral timestamps $t_i = i$ and intervals of the form $[0, a]$ for $a \in \mathbb{N}$.

5 UTL

The most complicated circuits in the LTL path-checking algorithm [9] correspond to $s \mathrm{U} \psi$ and $\psi \mathrm{U} s$ formulae. As in the case of MTL, the circuits are also not uniform but depend on s. In this section, we devise an AC^1 tree-contraction algorithm for UTL—the fragment of LTL obtained by omitting binary temporal operators. The algorithm works even if XOR is allowed and is based on the analysis of functions arising in the tree contraction algorithm applied to UTL formulae. First consider the future-only fragment of UTL.

Let $p \in \mathbb{B}^n$ be any proposition. If $p(i) = \bot$ for every i then $(\mathrm{F}p)(i) = \bot$ for every i. Otherwise, let i be the largest index such that $p(i) = \top$. Then, $(\mathrm{F}p)(j) = \top$ for all $j \le i$. By construction, $p(k) = \bot$ for all $k > i$. Hence, $(\mathrm{F}p)(k) = \bot$ for all $k > i$. Thus, Fp is downward monotone and depends only on the largest i with $p(i) = \top$. In particular, only $n + 1$ possible values exist for Fp.

Similarly, let t be the largest index such that $p(t) = \bot$. Then $p(j) = \top$ for all $j > t$. Hence $(\mathrm{G}p)(j) = \top$ for all $j > t$. Since $p(t) = \bot$ we have $(\mathrm{G}p)(k) = \bot$ for all $k \le t$. Thus, Gp is upward monotone and depends only on the largest t with $p(t) = \bot$. In particular, only $n + 1$ possible values exist for Gp.

So for any formula ψ the value of $\mathrm{F} \circ \psi$ or $\mathrm{G} \circ \psi$ is a monotone vector—of which there are only $2n$ many. Hence for any formula context $\varphi(X)$, the formula contexts $\varphi \circ (\mathrm{F}X)$ and $\varphi \circ (\mathrm{G}X)$ can be represented as $g \circ F$ or $g \circ G$ where $g : \mathbb{M} \to \mathbb{B}^n$ is a ***function with monotone domain***. Since $|\mathbb{M}| = O(n)$, enumerating all outputs of g explicitly requires only $|g| = O(n^2)$ space. Similar results hold for the past equivalents of G and F.

Now, Boolean operators are applied componentwise and obey the usual identities: $\bot \wedge p = \bot, \top \wedge p = p, \bot \vee p = p, \top \vee p = \top, \bot \oplus p = p$ and $\top \oplus p = \neg p$. Therefore, to represent partial evaluation of conjunction $(p \wedge X, x \wedge X)$, disjunction $(p \vee X, X \vee p)$ and xor $(p \oplus X, X \oplus p)$ it suffices to keep track whether each component is $\bot, \top$ or equal to the original or the negation of the value in X.

Furthermore, Next (Xp) and Yesterday (Yp) temporal operators shift p by 1 and -1, respectively. Let m be the size of the input formula. The last two paragraphs motivate the definition of filters: let $v \in \{\bot, \top, \mathrm{ID}, \mathrm{NOT}\}^n$ and $k \in [-m, m]$ satisfy $v(i) \in \mathbb{B}$ if $i + k \notin \{1, \ldots, n\}$. Then a ***filter with offset*** k ***and pattern*** v is the function $f_{v,k} : \mathbb{B}^n \to \mathbb{B}^n$ such that

$$f_{v,k}(p)(i) = \begin{cases} \bot & \text{if } v(i) = \bot \\ \top & \text{if } v(i) = \top \\ p(i+k) & \text{if } v(i) = \mathrm{ID} \\ \neg p(i+k) & \text{if } v(i) = \mathrm{NOT} \end{cases}$$

The identity function as well as the partial evaluation of conjunction, disjunction, and xor are expressible as filters with offset 0. Temporal operators Next and Yesterday are identity filters with offsets 1 and -1, respectively. Note that filters are closed under composition.

Storing v explicitly and k in unary requires $O(n + |\varphi|)$ bits per filter. By fully expanding the definition, we can evaluate and compose two filters in AC^0. Moreover, if $g : \mathbb{M} \to \mathbb{B}^n$ is a function with monotone domain then $(f_{v,k} \circ g) : \mathbb{M} \to \mathbb{B}^n$ is also a function with monotone domain and the composition in AC^0.

Lemma 4. *There are uniform* AC^0 *circuits calculating* $f_{v,k} \circ f_{v',k'}$ *and* $f_{v,k}(p)(i)$ *and* $f_{v,k} \circ g$ *and* $\mathrm{F} \circ g$ *and* $\mathrm{G} \circ g$*, where* f*'s are filters and* g *is a function with monotone domain.*

We represent the functions arising in the tree-contraction algorithm as follows. If the contracted subtree S does not contain F or G operators then it is representable by a filter. If it contains F or G then let T be the first such occurrence. Then the segment from the leaves to T is representable by a filter and the segment above T is representable by a function with monotone domain. Thus, the function h associated with S can be represented as:

$$h = \begin{cases} \text{filter} & \text{no temporal operator} \\ f \circ T \circ \text{filter} & T \text{ is the first temporal operator; } f : \mathbb{M} \to \mathbb{B}^n \end{cases}$$

Now, if the contracted node is a Boolean connective, X or Y then we calculate $f_{v,k} \circ h$ for an appropriate filter. If the contracted node is F or G then we calculate $\mathrm{F} \circ h$ or $\mathrm{G} \circ h$. In either case, the resulting function is representable using the above format. Moreover, by Lemma 4, the composition is in AC^0. Hence, the complexity of the tree contraction algorithm is $\mathsf{AC}^1[\mathsf{AC}^0] = \mathsf{AC}^1$.

Theorem 5. UTL *path checking is in* AC^1*.*

Same results apply to past temporal operators. Note that the construction works also when the negation is applied to arbitrary subformulae, and not only to propositions. Also note that $F_{[a,\infty)}p$ is downward monotone and the corresponding circuits are constructible in logarithmic space. Therefore, the above arguments apply to the more powerful logic $\mathsf{UTL}_{\geq}$ obtained by allowing $F_{[a,\infty)}$ and $G_{[b,\infty)}$ operators. To the best of our knowledge, $\mathsf{UTL}_{\geq}$ is the most expressive and powerful fragment of LTL with a sub-$\mathsf{AC}^1[\mathsf{logDCFL}]$ path-checking problem.

6 Conclusion

The results obtained in this paper shed further light on the complexity landscape of temporal-logic path-checking problems. Several open questions however remain, the main one being to determine the precise complexity of LTL path checking. In particular, there has been no progress on the trivial NC^1 lower bound over the past ten years. Furthermore, might it be possible to separate the complexity of LTL and MTL, or of these logics and their future-only fragment?

Acknowledgments. This research was financially supported by EPSRC.

References

1. Artho, C., Barringer, H., Goldberg, A., Havelund, K., Khurshid, S., Lowry, M., Pasareanu, C., Rosu, G., Sen, K., Visser, W., Washington, R.: Combining test case generation and runtime verification. Theoretical Computer Science 336(2-3) (2005)
2. Barrington, D.A.M., Lu, C.-J., Miltersen, P.B., Skyum, S.: On monotone planar circuits. In: Proceedings of Fourteenth Annual IEEE Conference on Computational Complexity (1999)
3. Bundala, D., Ouaknine, J.: On the complexity of path checking of temporal logics (full version). CoRR, abs/1312.7603
4. Chakraborty, T., Datta, S.: One-input-face MPCVP is hard for L, but in LogDCFL. In: Arun-Kumar, S., Garg, N. (eds.) FSTTCS 2006. LNCS, vol. 4337, pp. 57–68. Springer, Heidelberg (2006)
5. Demri, S., Schnoebelen, P.: The complexity of propositional linear temporal logics in simple cases. Information and Computation 174(1), 84–103 (2002)
6. Finkbeiner, B., Sipma, H.: Checking finite traces using alternating automata. Formal Methods in System Design 24(2), 101–127 (2004)
7. Goldschlager, L.M.: The monotone and planar circuit value problems are log space complete for P. SIGACT News 9 (July 1977)
8. Karp, R.M., Ramachandran, V.: Parallel algorithms for shared-memory machines. In: Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity (A) (1990)
9. Kuhtz, L., Finkbeiner, B.: LTL path checking is efficiently parallelizable. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikoletseas, S., Thomas, W. (eds.) ICALP 2009, Part II. LNCS, vol. 5556, pp. 235–246. Springer, Heidelberg (2009)
10. Kuhtz, L., Finkbeiner, B.: Efficient parallel path checking for linear-time temporal logic with past and bounds. Logical Methods in Computer Science 8(4) (2012)
11. Limaye, N., Mahajan, M., Sarma, J.M.N.: Evaluating monotone circuits on cylinders, planes and tori. In: Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 660–671. Springer, Heidelberg (2006)
12. Maler, O., Nickovic, D.: Monitoring temporal properties of continuous signals. In: Lakhnech, Y., Yovine, S. (eds.) FORMATS 2004 and FTRTFT 2004. LNCS, vol. 3253, pp. 152–166. Springer, Heidelberg (2004)
13. Markey, N., Schnoebelen, P.: Model checking a path (Preliminary report). In: Amadio, R.M., Lugiez, D. (eds.) CONCUR 2003. LNCS, vol. 2761, pp. 251–265. Springer, Heidelberg (2003)
14. Younes, H.L.S., Simmons, R.G.: Probabilistic verification of discrete event systems using acceptance sampling. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 223–235. Springer, Heidelberg (2002)