




Synchronized CTL over One-Counter Automata

Shaul Almagor   

Department of Computer Science, Technion, Israel

Daniel Assa 

Reichman University, Herzliya, Israel

Udi Boker   

Reichman University, Herzliya, Israel

Abstract

We consider the model-checking problem of *Synchronized Computation-Tree Logic* (CTL+Sync) over One-Counter Automata (OCAs). CTL+Sync augments CTL with temporal operators that require several paths to satisfy properties in a synchronous manner, e.g., the property “all paths should eventually see p at the same time”. The model-checking problem for CTL+Sync over finite-state Kripke structures was shown to be in $P^{NP^{NP}}$. OCAs are labelled transition systems equipped with a non-negative counter that can be zero-tested. Thus, they induce infinite-state systems whose computation trees are not regular. The model-checking problem for CTL over OCAs was shown to be PSPACE-complete.

We show that the model-checking problem for CTL+Sync over OCAs is decidable. However, the upper bound we give is non-elementary. We therefore proceed to study the problem for a central fragment of CTL+Sync, extending CTL with operators that require *all* paths to satisfy properties in a synchronous manner, and show that it is in EXP^{NEXP} (and in particular in EXPSPACE), by exhibiting a certain “segmented periodicity” in the computation trees of OCAs.

2012 ACM Subject Classification Theory of computation → Formal languages and automata theory

Keywords and phrases CTL, Synchronization, One Counter Automata, Model Checking

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2023.19

Related Version *Full Version*: <https://arxiv.org/pdf/2308.03308.pdf>

Funding *Shaul Almagor*: supported by the Israel Science Foundation grant 989/22.

Udi Boker: supported by the Israel Science Foundation grant 2410/22.

1 Introduction

Branching-time model checking is a central avenue in formal verification, as it enables reasoning about multiple computations of the system with both an existential and universal quantification. As systems become richer, the classical paradigm of e.g., CTL model checking over finite-state systems becomes insufficient. To this end, researchers have proposed extensions both of the logics [2, 6, 5, 3, 4, 1] and of the systems [7, 21, 9, 13].

In the systems’ frontier, of particular interest are infinite-state models. Typically, such models can quickly lead to undecidability (e.g., two-counter machines [18]). However, some models can retain decidability while still having rich modelling power. One such model that has received a lot of attention in recent years is One Counter Automata (OCAs) [20, 15] – finite state machines equipped with a non-negative counter that can be zero-tested. Model checking CTL over OCAs was studied in [13], where it was shown to be decidable in PSPACE. The main tool used there is the fact that despite the infinite configuration space, the computations of an OCA do admit some periodic behavior, which can be exploited to exhibit a small-model property for the satisfaction of Until formulas.



© Shaul Almagor, Daniel Assa, and Udi Boker;

licensed under Creative Commons License CC-BY 4.0

43rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2023).

Editors: Patricia Bouyer and Srikanth Srinivasan; Article No. 19; pp. 19:1–19:22



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In the logics' frontier, a useful extension of CTL is that of CTL with Synchronization operators (CTL+Sync), introduced in [4]. CTL+Sync extends CTL with operators that express synchronization properties of computation trees. Specifically, two new operators are introduced: *Until All* and *Until Exists*. The former, denoted by $\psi_1 UA \psi_2$, holds in state s if there is a uniform bound $k \in \mathbb{N}$ such that ψ_2 holds in all paths from s after exactly k steps, and ψ_1 holds in all paths up to step k . Thus, intuitively, it requires all the computations to synchronize the satisfaction of the Until operator. The latter, somewhat less natural operator, denoted by $\psi_1 UE \psi_2$, requires that there exists a uniform bound k such that in every level $j < k$ of the computation tree, some path satisfies ψ_1 and can be continued to satisfy ψ_2 at level k .

In comparison, the standard CTL operators $A\psi_1 U \psi_2$ and $E\psi_1 U \psi_2$ require that all paths/some path satisfy the Until formula, but there is no requirement that the bounds coincide. We illustrate the differences between the semantics in Figure 1. As discussed in [4], CTL+Sync can describe non ω -regular properties of trees, and hence goes beyond MSO, while retaining a decidable model-checking problem over finite Kripke structures.




Figure 1 The computation tree in (b) satisfies $A \text{true } U \bullet$, since every path eventually reaches \bullet . However, it does not satisfy $\text{true } UA \bullet$, since this eventuality does not happen synchronously. In contrast, the tree in (a) does satisfy $\text{true } UA \bullet$. The formula $E \circ U \textcircled{/}$ is satisfied in (a) (along the left branch), but not in (b), which only satisfies the weaker $\circ UE \textcircled{/}$. The boxed nodes on each level show that \circ holds along a path to $\textcircled{/}$ at level 6.

In this work, we show the decidability of CTL+Sync model checking over OCAs: Given an OCA \mathcal{A} and a CTL+Sync formula φ , the problem of whether \mathcal{A} satisfies φ . We thus combine the expressiveness of CTL+Sync with the rich modeling power of OCAs.

On the technical side, the approach taken in [4] to solve model-checking of CTL+Sync over Kripke structures does not seem to be very useful in our case. The solution there is based on the observation that every two levels of the computation tree that share the same set of Kripke-states must also share the same satisfaction value to every CTL+Sync formula. Hence, in that case, the algorithm can follow the computation tree of the powerset of the given Kripke structure, and terminate when encountering a level that has the same set of states as some previous level. In contrast, for OCAs, the unbounded counter prevents the ability to consider subsets of configurations.

On the other hand, the approach taken in [13] to solve model-checking of OCAs with respect to CTL is indeed useful in our case. Specifically, the algorithm in [13] is based on an analysis of the periodic behavior of the set of counter values that satisfy a given CTL formula in a given state of the OCA. We extend this approach, taking into account the additional complexity that stems from the synchronization requirements; see Section 5.

We start with establishing the decidability of CTL+Sync model checking using Presburger Arithmetic (see Section 4). This, however, yields a procedure with non-elementary runtime.

We then proceed to our main technical contribution (Section 5), providing an algorithm for model checking the central fragment of CTL+Sync that extends CTL with the UA operator, which requires *all* paths to satisfy properties in a synchronous manner. Its running time is in EXP^{NEXP} (and in particular in EXPSPACE), and for a fixed OCA and formulas of a fixed nesting depth, it is in PNP (and in particular in PSPACE).

Since CTL+Sync makes assertions on the behavior of different paths at the same time step (namely the same level of the computation tree), we need to reason about which configurations occur at each level of the tree. More precisely, in order to establish decidability we wish to exhibit a small-model property of the form *if the computation tree from some configuration (s, v) , for a state s and counter value v , satisfies the formula φ , then the computation tree from some configuration (s, v') for a small v' satisfies φ as well*. Unfortunately, the computation trees of an OCA from two configurations (s, v) and (s, v') cannot be easily obtained from one another using simple pumping arguments, due to the zero tests. This is in contrast to the case where one does not care about the length of a path, as in [13]. To overcome this, we show that computation trees of an OCA \mathcal{A} can be split into several segments, polynomially many in the size of \mathcal{A} , and that within each segment we can find a bounded prefix that is common to all trees after a certain counter threshold, and such that the remainder of the segment is periodic. Using this, we establish the small model property above. The toolbox used for proving this, apart from careful cycle analysis, includes 2TVASS – a variant of 2VASS studied in [17] that allows for one counter of the 2VASS to be zero-tested.

We believe that this structural result (Lemma 18) is of independent interest for reasoning about multiple traces in an OCA computation tree, when the length of paths plays a role.

Due to lack of space, Proposition proofs are omitted; they appear in the arXiv version.

2 Preliminaries

Let $\mathbb{N} = \{0, 1, \dots\}$ be the natural numbers. For a finite set $A \subseteq \mathbb{N}$ we denote by $\text{lcm}(A)$ the least common multiple of the elements in A .

For a finite sequence $\xi = t_0 t_1 \dots t_{h-1}$ and integers x, y , such that $0 \leq x \leq y \leq h-1$, we write $\xi[x..y]$ for the infix of ξ between positions x and y , namely for $t_x t_{x+1} \dots t_y$. We also use the parentheses “(” and “)” for a non-inclusive range, e.g., $[x..y) = [x..y-1]$, and abbreviate $\xi[x..x]$ by $\xi[x]$. We denote the length of ξ by $|\xi| = h$.

One Counter Automata

A *One Counter Automaton* (OCA) is a triple $\mathcal{A} = \langle S, \Delta, L \rangle$, where S is a finite set of states, $\Delta \subseteq (S \times \{=0, >0\} \times \{0, +1, -1\} \times S)$ is a transition relation, and $L : S \rightarrow AP$, for some finite set AP of *atomic propositions*, is the state labeling.

A pair $(s, v) \in S \times \mathbb{N}$ is a *configuration* of \mathcal{A} . We write $(s, v) \rightarrow_t (s', v')$ for a transition $t \in \Delta$ if one of the following holds:

- $t = (s, =0, e, s')$, with $e \in \{0, +1\}$, $v = 0$ and $v' = e$, or
- $t = (s, >0, e, s')$, with $e \in \{0, +1, -1\}$, $v > 0$ and $v' = v + e$.

We write $(s, v) \rightarrow (s', v')$ if $(s, v) \rightarrow_t (s', v')$ for some $t \in \Delta$.

We require that Δ is total, in the sense that for every configuration (s, v) we have $(s, v) \rightarrow (s', v')$ for some configuration (s', v') . Note that this is a syntactic requirement – every state should have outgoing transitions on both $=0$ and >0 . This corresponds to the standard requirement of Kripke structures that there are no deadlocks. We denote by $|\mathcal{A}|$ the number of states of \mathcal{A} . Note that the description size of \mathcal{A} is therefore polynomial in $|\mathcal{A}|$.

19:4 Synchronized CTL over One-Counter Automata

A *path* of \mathcal{A} is a sequence of transitions $\tau = t_1, \dots, t_k$ such that there exist states s_0, \dots, s_k where $t_i = (s_{i-1}, \bowtie_i, e_i, s_i)$ with $\bowtie_i \in \{=0, >0\}$ for all $1 \leq i \leq k$. We say that τ is *valid* from starting counter v_0 (or from configuration (s_0, v_0)), if there are counters $v_0, \dots, v_k \in \mathbb{N}$ such that for all $1 \leq i < k$ we have $(s_{i-1}, v_{i-1}) \rightarrow_{t_i} (s_i, v_i)$. We abuse notation and refer to the sequence of configurations also as a *path*, starting in (s_0, v_0) and ending in (s_k, v_k) . The *length* of the path τ is k , and we define its *effect* $(\pi) = \sum_{i=1}^k e_i$.

We also allow infinite paths, in which case there are no end configurations and the length is ∞ . In this case we explicitly mention that the path is infinite. We say that τ is *balanced/positive/negative* if $\text{effect}(\tau)$ is zero/positive/negative, respectively. It is a *cycle* if $s_0 = s_k$, and it *has a cycle* β , if β is a cycle and is a contiguous infix of τ .

CTL+Sync

A CTL+Sync formula φ is given by the following syntax, where q stands for an atomic proposition from a finite set AP of atomic propositions.

$$\varphi ::= \underbrace{\text{true} \mid q \mid \varphi \wedge \varphi \mid \neg\varphi \mid EX\varphi \mid E\varphi U\varphi \mid A\varphi U\varphi}_{\text{Standard CTL}} \mid \underbrace{\varphi UA\varphi \mid \varphi UE\varphi}_{\text{Sync operators}}$$

We proceed to the semantics. Consider an OCA $\mathcal{A} = \langle S, \Delta, L \rangle$, a configuration (s, v) , and a CTL+Sync formula φ . Then $\mathcal{A}^{(s,v)}$ satisfies φ , denoted by $\mathcal{A}^{(s,v)} \models \varphi$, as defined below.

Boolean Operators:

- $\mathcal{A}^{(s,v)} \models \text{true}$; $\mathcal{A}^{(s,v)} \not\models \text{false}$
- $\mathcal{A}^{(s,v)} \models \neg\varphi$ if $\mathcal{A}^{(s,v)} \not\models \varphi$.
- $\mathcal{A}^{(s,v)} \models q$ if $q \in L(s)$.
- $\mathcal{A}^{(s,v)} \models \varphi \wedge \psi$ if $\mathcal{A}^{(s,v)} \models \varphi$ and $\mathcal{A}^{(s,v)} \models \psi$.

CTL temporal operators:

- $\mathcal{A}^{(s,v)} \models EX\varphi$ if $(s, v) \rightarrow (s', v')$ for some configuration (s', v') such that $\mathcal{A}^{(s',v')} \models \varphi$.
- $\mathcal{A}^{(s,v)} \models E\varphi U\psi$ if there exists a valid path τ from (s, v) and $k \geq 0$, such that $\mathcal{A}^{\tau[k]} \models \psi$ and for every $j \in [0..k-1]$, we have $\mathcal{A}^{\tau[j]} \models \varphi$.
- $\mathcal{A}^{(s,v)} \models A\varphi U\psi$ if for every valid path τ from (s, v) there exists $k \geq 0$, such that $\mathcal{A}^{\tau[k]} \models \psi$ and for every $j \in [0..k-1]$, we have $\mathcal{A}^{\tau[j]} \models \varphi$.

Synchronization operators:

- $\mathcal{A}^{(s,v)} \models \varphi UA\psi$ if there exists $k \geq 0$, such that for every valid path τ from (s, v) of length k and for every $j \in [0..k-1]$ we have $\mathcal{A}^{\tau[j]} \models \varphi$ and $\mathcal{A}^{\tau[k]} \models \psi$.
- $\mathcal{A}^{(s,v)} \models \varphi UE\psi$ if there exists $k \geq 0$, such that for every $j \in [0..k-1]$ there exists a valid path τ from (s, v) of length k such that $\mathcal{A}^{\tau[j]} \models \varphi$ and $\mathcal{A}^{\tau[k]} \models \psi$.

► **Remark 1 (Additional operators).** Standard additional Boolean and CTL operators, e.g., \vee, EF, EG , etc. can be expressed by means of the given syntax. Similar shorthands can be defined for the synchronization operators, e.g., FE and GE , etc. We remark that one can also consider operators such as $XE\psi$ with the semantics “in the next step there exists a path satisfying ψ ”. However, the semantics of this coincides with the CTL operator EX .

Presburger Arithmetic

Presburger Arithmetic (PA) [19] is the first-order theory $\text{Th}(\mathbb{N}, 0, 1, +, <, =)$ of \mathbb{N} with addition and order. We briefly survey the results we need about PA, and refer the reader to [14] for a detailed survey.

For our purposes, a PA formula $\varphi(x_1, \dots, x_d)$, where x_1, \dots, x_d are free variables, is evaluated over \mathbb{N}^d , and *defines* the set $\{(a_1, \dots, a_d) \in \mathbb{N}^d \mid (a_1, \dots, a_d) \models \varphi(x_1, \dots, x_d)\}$. It is known that PA is decidable in 2-NEXP [8, 11].

A *linear set* is a set of the form $\text{Lin}(B, P) = \{\mathbf{b} + \lambda_1 \mathbf{p}_1 + \dots + \lambda_k \mathbf{p}_k \mid \mathbf{b} \in B, \lambda_1, \dots, \lambda_k \in \mathbb{N}\}$ where $B \subseteq \mathbb{N}^d$ is a finite *basis* and $P = \{\mathbf{p}_1, \dots, \mathbf{p}_k\} \subseteq \mathbb{N}^d$ are the *periods*. A *semilinear set* is then a finite union of linear sets. A fundamental result about PA [12] is that the sets definable in PA are exactly the semilinear sets, and moreover, one can effectively obtain from a PA formula φ a description of the semilinear set satisfying a formula φ , and vice versa.

► **Observation 2.** *In dimension 1, semilinear sets are finite unions of arithmetic progressions. By taking the lcm of the periods of the progressions and modifying the basis accordingly, we can assume a uniform period. That is, a semilinear set $S \subseteq \mathbb{N}$ is $\text{Lin}(B, \{p\})$ for effectively computable $B \subseteq \mathbb{N}$ and $p \in \mathbb{N}$.*

3 Periodicity and Flatness over OCAs

Recall that the configuration space of an OCA is $S \times \mathbb{N}$. The underlying approach we take to solve CTL+Sync model checking is to show that satisfaction of CTL+Sync formulas over these configurations exhibits some periodicity. Moreover, the run tree of the OCA can be captured, to an extent, using a small number of cycles (a property called *flatness*). These properties will be relied upon in proving our main results.

3.1 Periodicity

In this subsection we formalize our notions of ultimate periodicity, show how they suffice for model-checking, and cite important results about periodicity in CTL.

Consider a CTL+Sync formula φ . We say that φ is $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -*periodic* with respect to an OCA \mathcal{A} (or just *periodic*, if we do not care about the constants) if for every state $s \in S$ and counters $v, v' > \mathfrak{t}(\varphi)$, if $v \equiv v' \pmod{\mathfrak{p}(\varphi)}$ then $(s, v) \models \varphi \iff (s, v') \models \varphi$. We think of $\mathfrak{t}(\varphi)$ as its threshold and of $\mathfrak{p}(\varphi)$ as its period. We say that φ is *totally* $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -*periodic* with respect to \mathcal{A} if every subformula of φ (including φ itself) is $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -periodic with respect to \mathcal{A} . We usually omit \mathcal{A} , as it is clear from context.

Total-periodicity is tantamount to periodicity for each subformula, in the following sense.

► **Proposition 3.** *A CTL+Sync formula φ is totally $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -periodic if and only if every subformula ψ of φ is $(\mathfrak{t}'(\psi), \mathfrak{p}'(\psi))$ -periodic for some constants $\mathfrak{t}'(\psi), \mathfrak{p}'(\psi)$.*

For a totally periodic formula, model checking over OCA can be reduced to model checking over finite Kripke structures, as follows. Intuitively, we simply “unfold” the OCA and identify states with high counter values according to their modulo of $\mathfrak{p}(\varphi)$.

► **Proposition 4.** *Consider an OCA \mathcal{A} and a totally $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -periodic CTL+Sync formula φ , then we can effectively construct a Kripke structure \mathcal{K} of size $|\mathcal{A}| \cdot (\mathfrak{t}(\varphi) + \mathfrak{p}(\varphi))$ such that $\mathcal{A} \models \varphi$ if and only if $\mathcal{K} \models \varphi$.*

In [13, Theorem 1] it is proved that every CTL formula φ over OCA is periodic. Our goal is to give a similar result for CTL+Sync, which in particular contains CTL. In order to avoid replicating the proof in [13] for CTL, we observe that the proof therein is by structural induction over φ , and that moreover – the inductive assumption requires only periodicity of the subformulas of φ . We can thus restate [13, Theorem 1] with the explicit inductive assumption, so that we can directly plug our results about CTL+Sync into it.

Denote by $k = |S|$ the number of states of \mathcal{A} , and let $K = \text{lcm}(\{1, \dots, k\})$.

► **Theorem 5** (Theorem 1 in [13], restated). *Consider a CTL+Sync formula φ whose outermost operator is a CTL operator, and whose subformulas are periodic, then φ is periodic, and we have the following.*

- If $\varphi = \mathbf{true}$, $\varphi = \mathbf{false}$, or $\varphi = p$ for $p \in AP$, then φ is $(0, 1)$ -periodic.
- If $\varphi = \neg\psi$ then $\mathfrak{t}(\varphi) = \mathfrak{t}(\psi)$ and $\mathfrak{p}(\varphi) = \mathfrak{p}(\psi)$.
- If $\varphi = \psi_1 \wedge \psi_2$ then $\mathfrak{t}(\varphi) = \max\{\mathfrak{t}(\psi_1), \mathfrak{t}(\psi_2)\}$ and $\mathfrak{p}(\varphi) = \text{lcm}(\mathfrak{p}(\psi_1), \mathfrak{p}(\psi_2))$.
- If $\varphi = EX\psi$ then $\mathfrak{t}(\varphi) = \mathfrak{t}(\psi) + \mathfrak{p}(\psi)$ and $\mathfrak{p}(\varphi) = K \cdot \mathfrak{p}(\psi)$
- If $\varphi = E\psi_1 U \psi_2$ or $\varphi = A\psi_1 U \psi_2$ then¹ $\mathfrak{t}(\varphi) = \max\{\mathfrak{t}(\psi_1), \mathfrak{t}(\psi_2)\} + 2 \cdot k^2 \cdot \text{lcm}(K \cdot \mathfrak{p}(\psi_1), \mathfrak{p}(\psi_2))$ and $\mathfrak{p}(\varphi) = \text{lcm}(K \cdot \mathfrak{p}(\psi_1), \mathfrak{p}(\psi_2))$.

3.2 Linear Path Schemes

The runs of an OCA can take intricate shapes. Fortunately, however, we can use the results of [17] about *flatness* of a variant of 2-VASS with some zero tests, referred to as 2-TVASS, to obtain a simple form that characterizes reachability, namely linear path schemes.

A *linear path scheme* (LPS) is an expression of the form $\pi = \alpha_0 \beta_1^* \alpha_1 \cdots \beta_k^* \alpha_k$ where each $\alpha_i \in \Delta^*$ is a path in \mathcal{A} and each $\beta_i \in \Delta^*$ is a cycle in \mathcal{A} . The *flat length* of π is $|\pi| = |\alpha_0 \beta_1 \alpha_1 \cdots \beta_k \alpha_k|$, the *size* of π is k .

A concrete path τ in \mathcal{A} is π -shaped if there exist e_1, \dots, e_k such that $\tau = \alpha_0 \beta_1^{e_1} \alpha_1 \cdots \beta_k^{e_k} \alpha_k$.

Our first step is to use a result of [17] on 2-TVASS to show that paths of a fixed length in \mathcal{A} admit a short LPS. The idea is to transform the OCA \mathcal{A} to a 2-TVASS \mathcal{A}' by introducing a length-counting component. That is, in every transition of \mathcal{A}' as a 2-TVASS, the second component increments by 1.

► **Lemma 6.** *Let (s, v) and (s', v') be configurations of \mathcal{A} . If there exists a path τ of length ℓ from (s, v) to (s', v') , then there is also a π -shaped path τ' of length ℓ from (s, v) to (s', v') , where π is some linear path scheme of flat length $\text{poly}(|S|)$ and size $O(|S|^3)$.*

By Lemma 6, we can focus our attention to π -shaped paths where π is “short”. Henceforth, we call a path τ *basic* if it is π -shaped for some LPS π as per Lemma 6.

Using standard acceleration techniques (see e.g., [16, 10, 17]), we also get from Lemma 6 that the reachability relation of an OCA (including path length) is effectively semilinear. More precisely, we have the following.

► **Corollary 7.** *We can effectively compute, for every $s, s' \in S$, a PA formula $\text{Path}_{s,s'}(x, y, x', y')$ such that $(v, \ell, v', \ell') \models \text{Path}_{s,s'}(x, y, x', y')$ if and only if a path of length ℓ ending² at (s, v) can be continued to a path of length ℓ' ending at (s', v') .*

4 Model Checking CTL+Sync via Presburger Arithmetic

In this section we show that model checking CTL+Sync over OCAs is decidable, by reducing it to the satisfiability problem of a PA formula.

We start with a simple observation.

► **Lemma 8.** *Consider a totally $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -periodic CTL+Sync formula φ , then for every state $s \in S$ we can compute a PA formula $P_{\varphi,s}(x)$ such $v \models P_{\varphi,s}(x)$ if and only if $(s, v) \models \varphi$.*

¹ Note that in [13] the case of $A\psi_1 U \psi_2$ is not stated, but rather the dual Release operator $E\psi_1 R \psi_2$, which follows the same proof.

² It is more natural to assume $\ell = 0$ and simply consider paths starting at (s, v) . However, our formulation makes things easier later on.

Next, we show that from a PA formula as above we can obtain a threshold and a period.

► **Lemma 9.** *Consider a CTL+Sync formula φ and PA formulas $P_{\varphi,s}(x)$, for every state s , such that $v \models P_{\varphi,s}(x)$ iff $(s, v) \models \varphi$. Then φ is $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -periodic for computable constants $\mathfrak{t}(\varphi)$ and $\mathfrak{p}(\varphi)$.*

Combining Theorem 5 and Lemma 8 we obtain that every CTL formula (without Sync) can be translated to PA formulas. We now turn to include the Sync operators.

Consider a CTL+Sync formula φ . We construct, by induction on the structure of φ , PA formulas $P_{s,\varphi}(v)$, for every state $s \in S$, such that $P_{s,\varphi}(v)$ holds if and only if $(s, v) \models \varphi$. For the Sync operators, this utilizes the PA formulas of Corollary 7.

- If $\varphi = p$ for an atomic proposition p , then $P_{s,\varphi}(v) = \text{True}$ if s is labeled with p and False otherwise.
- If $\varphi = \neg\psi$, then $P_{s,\varphi}(v) = \neg P_{s,\psi}(v)$.
- If $\varphi = \psi_1 \wedge \psi_2$, then $P_{s,\varphi}(v) = P_{s,\psi_1}(v) \wedge P_{s,\psi_2}(v)$.
- If $\varphi = EX\psi$, $\varphi = A\psi_1 U\psi_2$ or $\varphi = E\psi_1 U\psi_2$ then by the induction hypothesis, ψ , ψ_1 and ψ_2 have corresponding PA formulas, and by Lemma 9 we can compute $\mathfrak{t}(\psi)$ and $\mathfrak{p}(\psi)$ (and similarly for ψ_1 and ψ_2) such that ψ is $(\mathfrak{t}(\psi), \mathfrak{p}(\psi))$ -periodic. Then, by Theorem 5 we can compute $\mathfrak{t}(\varphi), \mathfrak{p}(\varphi)$ such that φ is $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -periodic. Thus, we can apply Lemma 8 to obtain PA formulas for φ .
- If $\varphi = \psi_1 UE\psi_2$, then $P_{s,\varphi}(v) = \exists \ell. \forall \ell' < \ell. \left(\bigvee_{s' \in S} \left(\exists v'. (\text{Path}_{s,s'}(v, 0, v', \ell') \wedge P_{s',\psi_1}(v')) \wedge \bigvee_{s'' \in S} \exists v''. \text{Path}_{s',s''}(v', \ell', v'', \ell) \wedge P_{s'',\psi_2}(v'') \right) \right)$
- If $\varphi = \psi_1 UA\psi_2$, then $P_{s,\varphi}(v) = \exists \ell. \left(\left(\exists v'. \bigvee_{s' \in S} (\text{Path}_{s,s'}(v, 0, v', \ell) \wedge P_{s',\psi_2}(v')) \right) \wedge \left(\bigwedge_{s' \in S} \forall v'. (\text{Path}_{s,s'}(v, 0, v', \ell) \rightarrow P_{s',\psi_2}(v')) \right) \right) \wedge \forall \ell' < \ell. \left(\bigwedge_{s' \in S} \forall v'. (\text{Path}_{s,s'}(v, 0, v', \ell') \rightarrow P_{s',\psi_1}(v')) \right)$

The semantics of the obtained PA formulas match the semantics of the respective Sync operators. By Lemma 9 we now have that for every CTL+Sync formula φ we can compute $\mathfrak{t}(\varphi), \mathfrak{p}(\varphi)$ such that φ is $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -periodic. By Proposition 3 we can further assume that φ is totally $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -periodic. Finally, by Proposition 4 we can decide whether $\mathcal{A} \models \varphi$.

► **Remark 10 (Complexity).** Observe that the complexity of our decision procedure via PA formulas is non-elementary. Indeed, when using a CTL subformula, we translate it, using Lemma 8, to a PA formula that may be exponential in the size of the formula and of the OCA. Thus, we might incur an exponential blowup in every step of the recursive construction, leading to a tower of exponents.

5 Model Checking the CTL+UA Fragment

In this section we consider the fragment of CTL+Sync induced by augmenting CTL with only the Sync operator $\psi_1 UA\psi_2$. For this fragment, we are able to obtain a much better upper bound for model checking, via careful analysis of the run tree of an OCA.

Throughout this section, we fix an OCA $\mathcal{A} = \langle S, \Delta, L \rangle$ with $n = |S| \geq 3$ states, and a CTL+UA formula φ . Consider a configuration (s, v) of \mathcal{A} and a CTL+UA formula $\varphi = \psi_1 UA\psi_2$ with UA being the outermost operator. The satisfaction of φ from (s, v) is determined by the computation tree of \mathcal{A} from (s, v) . Specifically, we have that $\mathcal{A}^{(s,v)} \models \varphi$ if there exists some bound $k \in \mathbb{N}$ such that ψ_2 holds in all configurations of level k of the computation tree, and ψ_1 holds in all configurations of levels ℓ for $0 \leq \ell < k$.

Therefore, in order to reason about the satisfaction of φ , it is enough to know which configurations appear in each level of the computation tree. This is in contrast with UE, where we would also need to consider the paths themselves. Fortunately, it means we can use the LPS of Lemma 6 to simplify the proofs.

► **Theorem 11.** *Given an OCA \mathcal{A} with $n = |\mathcal{A}|$ and a CTL+UA formula φ , we can compute a counter threshold cT and a counter period P , both single exponential in n and in the nesting depth of φ , such that φ is (cT, P) -periodic with respect to \mathcal{A} .*

Before we delve into the proof of Theorem 11, we show how it implies our main result.

► **Theorem 12.** *The model-checking problem for CTL+UA is decidable in EXP^{NEXP} .*

Proof. Consider a CTL+UA formula φ and an OCA \mathcal{A} . By Theorem 11 we can compute cT, P single exponential in $|\varphi|$ and $|\mathcal{A}|$, such that φ is (cT, P) -periodic. We then apply Proposition 4 to reduce the problem to model checking φ against a Kripke structure \mathcal{K} of size $|\mathcal{A}| \cdot (cT + P)$.

Finally, the proof of [4, Theorem 1] shows that model checking the CTL+UA fragment can be done in P^{NP} in the size of the Kripke structure and the formula, yielding an EXP^{NEXP} bound in our setting. ◀

► **Remark 13 (Lower Bounds).** The PSPACE-hardness of model checking CTL over OCA from [13] immediately implies PSPACE-hardness in our setting as well. However, tightening the gap between the lower and upper bounds remains an open problem.

The remainder of the paper is devoted to proving Theorem 11.

Cycle Manipulation and Slope Manipulation

A fundamental part of our proof involves delicately pumping and removing cycles to achieve specific counter values and/or lengths of paths. We do this with the following technical tools.

For a path τ , we define the *slope* of τ as $\frac{\text{effect}(\tau)}{|\tau|}$. Recall that a basic path is of the form $\tau = \alpha_0 \beta_1^{e_1} \alpha_1 \cdots \beta_k^{e_k} \alpha_k$ adhering to some LPS, where $k = O(|\mathcal{A}|)$ and each α_i and β_i is of length $\text{poly}(|\mathcal{A}|)$. We denote by b the maximum flat length of any LPS for a basic path. In particular, b bounds the flat length of the LPS, the size of it, and the length of any cycle or path in it.

We call a cycle *basic* if it is of length at most b . A slope of a path is *basic* if it may be the slope of a basic cycle, namely if it equals $\frac{x}{y}$, where $x \in [-b..b], y \in [1..b]$ and $|x| \leq y$. We denote the basic slopes by \mathbf{s}_i , starting with \mathbf{s}_1 for the smallest. For example for $b = 3$, the basic slopes are $(\mathbf{s}_1 = -1, \mathbf{s}_2 = -\frac{2}{3}, \mathbf{s}_3 = -\frac{1}{2}, \mathbf{s}_4 = -\frac{1}{3}, \mathbf{s}_5 = 0, \mathbf{s}_6 = \frac{1}{3}, \mathbf{s}_7 = \frac{1}{2}, \mathbf{s}_8 = \frac{2}{3}, \mathbf{s}_9 = 1)$. Observe that for every i , we have $|\mathbf{s}_i| \geq \frac{1}{b}$, and for every $j > i$, we have $\mathbf{s}_j - \mathbf{s}_i \geq \frac{1}{b^2}$ and when they are both negative, also $\frac{\mathbf{s}_j}{\mathbf{s}_i} \leq 1 - \frac{1}{b^2}$.

► **Proposition 14.** *Consider a basic path τ and basic cycles c_1, c_2, c_3 in τ with effects $e_1, e_2, e_3 \in [-b..b]$, respectively, and lengths $\ell_1, \ell_2, \ell_3 \in [1..b]$, respectively, such that $\frac{e_1}{\ell_1} \leq \frac{e_2}{\ell_2} \leq \frac{e_3}{\ell_3}$. Then there are numbers $k_1, k_3 \in [0..b^2]$, such that the combination of k_1 repetitions of c_1 and k_3 repetitions of c_3 yield an effect and length whose ratio is $\frac{e_2}{\ell_2}$.*

► **Proposition 15.** *Consider a path π with cycles c_1, c_2 with effects $e_1, e_2 \in [-b..b]$ and lengths $\ell_1, \ell_2 \in [1..b]$, respectively, such that $\frac{e_1}{\ell_1} < \frac{e_2}{\ell_2}$, and a length x that is divisible by $\text{lcm}[1..2b^2]$. Then there are numbers $k_1, k_2 \in [0..b \cdot x]$, such that the addition or removal of k_1 repetitions of c_1 and the addition or removal of k_2 repetitions of c_2 yield a path of the same effect as π and of a length shorter or longer, as desired, by x (provided enough cycle repetitions exist).*

In order to prove Theorem 11, we show that every computation tree of \mathcal{A} , starting from a big enough counter value $v > \text{cT}$, has a “segmented periodic” structure with respect to φ . That is, we can divide its levels into $\text{poly}(n)$ many *segments*, such that only the first $\text{sT} \in \text{Exp}(n, |\varphi|)$ levels in each segment are in the “core”, while every other level ℓ is a sort-of repetition of the level $\ell - P$, for a *period* P . We further show that there is a similarity between the cores of computation trees starting with counter values v and $v + P$. We depict the segmentation in Figure 2, and formalize it as follows.

Consider a formula $\varphi = \psi_1 U A \psi_2$, where ψ_1 and ψ_2 are $(\text{cT}(\psi_1), P(\psi_1))$ - and $(\text{cT}(\psi_2), P(\psi_2))$ -periodic, respectively. We define several constants to use throughout the proof.

Constants depending only on the number n of states in the OCA

- $b \in \text{Poly}(n)$: the bound on the length of a linear path scheme on \mathcal{A} .
- $B = \text{lcm}[1..2b^3]$.

Constants depending on n and the CTL+UA formula φ

- $P_{\text{prev}}(\varphi) = \text{lcm}(P(\psi_1), P(\psi_2))$ – the unified period of the subformulas.
- $\text{cT}_{\text{prev}}(\varphi) = \max(\text{cT}(\psi_1), \text{cT}(\psi_2))$ – the unified counter threshold of the subformulas.
- $P(\varphi) = B \cdot P_{\text{prev}}(\varphi)$ – the “period” of φ .
- $\text{sT}(\varphi) = b^9 \cdot P(\varphi)$ – the “segment threshold” of φ .
- $\text{cT}(\varphi) = b^{11} \cdot P(\varphi)$ – the “counter threshold” of φ .

Eventually, these periodicity constants are plugged into the inductive cases of Theorem 5, as shown in the proof of Theorem 11. Then, all the constants are single-exponential in n and the nesting depth of φ . Notice that the following relationship holds between the constants.

► **Proposition 16.** $P(\varphi) > \text{cT}_{\text{prev}}(\varphi)$ for $b \geq 2$.

When clear from the context, we omit the parameter φ from P , sT , cT , P_{prev} , cT_{prev} .

We provide below an intuitive explanation for the choice of constants above.


Intuition for the period P

The period has two different roles: *level-periodicity* within each segment of a computation tree, and *counter-value periodicity* between two computation trees starting with different counter values.

Level periodicity within a segment: For lengthening or shortening a basic path by P , we add and/or remove some copies of its cycles. Adding or removing P_{prev} copies of the same cycle guarantees that the end counter values of the original and new paths are equivalent modulo P_{prev} . Since the cycles in a basic path are of length in $[0..b]$, setting P to be divisible by $P_{\text{prev}} \cdot \text{lcm}[1..b]$, allows to add or remove $\frac{P}{|c|}$ copies of a cycle c , where $\frac{P}{|c|}$ is divisible by P_{prev} , as desired. Yet, we might need to add copies of one cycle and remove copies of another, thus, as per Proposition 15, we need P to be divisible by $P_{\text{prev}} \cdot \text{lcm}[1..2b^2]$.

Counter periodicity between computation trees: We change a path τ that starts with a counter value v to a path that starts with a counter value $v + P$, or vice versa, by lengthening or shortening it by $\frac{P}{s}$, respectively, where s is a positive basic slope. In some cases, we need to also make sure that the longer or shorter path has a drop bigger or smaller, respectively, than τ by exactly P .

As $\frac{P}{s}$ is bounded by $b \cdot P$, if there are at least $b \cdot P$ repetitions of a cycle c in τ whose slope is $-s$, we can just add or remove $\frac{b \cdot P}{|c|}$ copies of c , so we need P to be divisible by $P_{\text{prev}} \cdot \text{lcm}[1..b]$, for guaranteeing that the counter values at the end of the original and new paths are equivalent. Yet, in some cases we need to combine two cycles, as per Proposition 14. As the combination of the two cycles might be of length up to $2b^3$, we need P to be divisible by $P_{\text{prev}} \cdot \text{lcm}[1..2b^3]$.



■ **Figure 2** A computation tree from (s, v) above and from $(s, v + P)$ below, demonstrating the *segmented periodicity*. Each segment i starts at ξ_i and ends at $\xi_{i+1} - 1$ (except for the last, which never ends). The *core* of the computation tree is the union of the first sT positions from each segment. Following these sT positions, there is within each segment periodicity of length P , meaning that for every path of length ℓ , there exists an equivalent path of length $\ell + P$, as shown for the three points with the same state s and equivalent counter values $v_1 \equiv v_2 \equiv v_3$. Between the trees, there is equivalence between the core positions, mapped via the *shift* function, as shown for the two points with the same state s' and the equivalent counter values $v' \equiv v''$.

Intuition for the counter threshold cT and the segment threshold sT

In order to apply Propositions 14 and 15, we need to have in the handled path many repetitions of two cycles of different slopes. We thus choose cT and sT to be large enough so that paths in which only one (negative) cycle slope is repeated many times must hit zero within a special region called the “core” of the tree, as defined below.

The core of a computation tree

For every counter value $v > cT$, the “core” with respect to a fixed formula φ , denoted by $\text{core}(v) \subseteq \mathbb{N}$, of a computation tree of \mathcal{A} that starts with a counter value v consists of $m + 1 < b^2$ segments, with each segment corresponding to a negative basic slope and having sT consequent numbers. For every $i \in [0..m]$, the start of Segment i depends on the initial counter value v of the computation tree, it is denoted by $\xi_i(v)$, and it is defined as follows:

- $\xi_0(v) = 0$,
- For $i \in [1..m]$, we set $\xi_i(v) = \frac{-1}{s_i}(v - cT_{prev}) - b^8 \cdot P$.
- For convenience, we also define $\xi_{m+1}(v) = \infty$.

Then, we define $\text{core}(v) = \bigcup_{i=0}^m [\xi_i(v)..(\xi_i(v) + sT - 1)]$.

Observe that the core of every tree is an ordered list of $((m + 1) \cdot sT)$ numbers (levels), while just the starting level of every segment depends on the initial counter value v . We can thus define a bijection $\text{shift} : \text{core}(v) \rightarrow \text{core}(v + P)$ that maps the i -th number in $\text{core}(v)$ to the i -th number in $\text{core}(v + P)$ (see also Figure 2).

Recall that we define sT so that, intuitively, if a path is long enough to reach $\xi_i(v) + sT$ without reaching counter value cT_{prev} , then the path must have many cycles with a slope larger than s_i , and if the path manages to reach cT_{prev} before $\xi_{i+1}(v)$ (namely the end of Segment i), then it must have many cycles with a slope at least as small as s_i . This is formalized as follows.

► **Lemma 17.** *Let τ be a basic path, or a prefix of it, of length ℓ starting from counter value $v > cT$, and let $i \in [0..m]$.*

1. *If $\ell \geq \xi_i(v) + \frac{sT}{2}$ and the counter values of τ stay above cT_{prev} , then τ has a cycle with slope s_j for $j > i$ that repeats at least $b^4 \cdot P$ times.*
2. *If $\ell < \xi_{i+1}(v)$ and the counter values of τ reach cT_{prev} , then τ has a cycle with slope s_j for $j \leq i$ that repeats at least $b^4 \cdot P$ times.*

As a sanity check, the lemma states that if a path τ reaches cT_{prev} for the first time at length $\ell \in [\xi_1(v) + \frac{sT}{2} .. \xi_2(v))$, then it has many cycles with slope $s_1 = -1$ (enough to decrease down to cT_{prev} before $\xi_2(v)$), as well as many cycle with slope at least $s_2 = \frac{-(b-1)}{b}$ (enough to keep above cT_{prev} through $\xi_1(v) + \frac{sT}{2}$). Observe also that a path cannot reach cT_{prev} at Segment 0, namely before $\xi_1(v)$.

The segment and shift periodicity

Consider a threshold T and period P . We say that counter values u, v are (T, P) -equivalent, denoted by $u \equiv_{T, P} v$ if either $u, v \geq T$ and P divides $|u - v|$, or $u, v < T$ and $u = v$. That is, either both u, v are greater than T , in which case they are equivalent modulo P , or they are both smaller than T and are equal.

The segment periodicity within a computation tree is then stated as Claim 1 in Lemma 18 below, while the similarity between computation trees starting from counters v and $v + P$ as Claim 2. (By $(s, v) \xrightarrow{\ell} (s', v')$ we mean that the computation tree starting with state s and counter value v has a path of length ℓ ending in state s' and counter value v' .)

► **Lemma 18.** *Consider states s and e , a counter value $v > cT$, an arbitrary counter value u , and an arbitrary path length ℓ .*

1. *If $\ell \notin \text{core}(v)$ then:*
 - a. $(s, v) \xrightarrow{\ell} (e, u) \implies (s, v) \xrightarrow{\ell-P} (e, u')$, and
 - b. $(s, v) \xrightarrow{\ell-P} (e, u) \implies (s, v) \xrightarrow{\ell} (e, \tilde{u})$,
for some counter values u' and \tilde{u} , such that $u \equiv_{cT_{prev}, P_{prev}} u' \equiv_{cT_{prev}, P_{prev}} \tilde{u}$.
2. *If $\ell \in \text{core}(v)$ then:*
 - a. $(s, v + P) \xrightarrow{\text{shift}(\ell)} (e, u) \implies (s, v) \xrightarrow{\ell} (e, u')$, and
 - b. $(s, v) \xrightarrow{\ell} (e, u) \implies (s, v + P) \xrightarrow{\text{shift}(\ell)} (e, \tilde{u})$
for some counter values u' and \tilde{u} , such that $u \equiv_{cT_{prev}, P_{prev}} u' \equiv_{cT_{prev}, P_{prev}} \tilde{u}$.

Proof. Throughout the proof, we will abbreviate $u \equiv_{cT_{prev}, P_{prev}} u'$ by $u \equiv u'$. We split the proof into four parts, each devoted to one of the four stated implications. In each of them, we assume the existence of a path τ that witnesses the left side of the implication, say $(s, v) \xrightarrow{\ell} (e, u)$, and show that there exists a path τ' that witnesses the right side of the implication, say $(s, v) \xrightarrow{\ell-P} (e, u')$, where $u \equiv u'$. By Lemma 6, we assume that τ is a basic path. We present some of the cases; the remaining parts are in the appendix.

Proof of Lemma 18.1a

Let τ be a basic path of length $\ell \notin \text{core}(v)$ such that $(s, v) \xrightarrow{\ell} (e, u)$ via τ . We construct from τ a path τ' for $(s, v) \xrightarrow{\ell-P} (e, u')$, such that $u \equiv u'$. The proof is divided to two cases.

Case 1a.1: The counter values in τ stay above cT_{prev}

If there is no position in τ with counter value cT_{prev} , then in particular τ has no zero-transitions. Since $\ell \notin \text{core}(v)$, then in particular $\ell \geq sT > 3b^5P$. Thus, there are at least $3b^4P$ cycle repetitions in τ .

If there is a non-positive cycle c that is repeated at least P times, we can obtain τ' by removing $\frac{P}{|c|}$ copies of c , as the counter values along τ' are at least as high as the corresponding ones in τ . Observe that τ' is of length $\ell - P$ from (s, v) to (e, u') with $u' = u - \text{effect}(c) \frac{P}{|c|}$. Since we have $u' \geq u \geq cT_{prev}$, then $u \equiv u'$.

Otherwise, each non-positive cycle in τ is taken at most P times. Thus, the positive cycles are repeated at least $3b^4P - bP \geq 3b^3P$ times. In particular, there exists a positive cycle c that repeats at least $3b^2P$ times. By removing $\frac{P}{|c|}$ occurrences of it, we obtain a path τ' of length $\ell - P$. Notice first that this path is valid. Indeed, up until the cycle c is taken, the path τ' coincides with τ , so the counter remains above cT_{prev} . Since c is a positive cycle, after completing its iterations, the counter value becomes at least $3b^2P - P + cT_{prev}$. Then, even if all remaining transitions in the negative cycles have effect -1 , the counter value is reduced by at most b^2P (as there are at most $(b-1)P$ remaining cycles, each of effect at least $-b$, and the simple paths in τ can reduce by another b at most). Thus, the value of the counter remains at least $3b^2P - P + cT_{prev} - b^2P > cT_{prev}$. Finally, let (e, u') be the configuration reached at the end of τ' , then $u - u' = \text{effect}(c) \frac{P}{|c|}$, so $u \equiv u'$.

Case 1a.2: τ reaches counter value cT_{prev}

Let $0 \leq z_f \leq z_u \leq \ell$ be the first and ultimate positions in τ where the counter value is exactly cT_{prev} . We split τ into three parts: $\tau_1 = \tau[0 .. z_f)$, $\tau_2 = \tau[z_f .. z_u)$, $\tau_3 = \tau[z_u .. \ell]$ (it could be that $z_f = z_u$, in which case the middle part is empty). Since τ is of length $\ell \geq sT \geq b^9 \cdot P$, then at least one of the parts above is of length at least $b^8 \cdot P$ (recall $b \geq 3$). We split according to which part that is. For simplicity, we start with the cases that τ_2 or τ_3 are long, and only then handle the case of a long τ_1 .

1. *The middle part $\tau_2 = \tau[z_f .. z_u]$ is of length at least b^8P .*

As τ_2 is of length at least $b^8 \cdot P$, some cycle c in it must repeat at least $b^6 \cdot P$ times.

If c is balanced, we can obtain τ' by removing $\frac{P}{|c|}$ of its repetitions.

If c is positive, starting at position x with counter value v_x , then the counter value at position y where c 's repetitions end is at least $v_x + b^6 \cdot P$. As τ_2 eventually gets down to $cT_{prev} < P$, there must be a negative cycle c_- that repeats at least $b \cdot P$ times between position y and the first position after y that has the counter value $v_x + \frac{b^6 \cdot P}{2}$. Hence, we can obtain τ' by removing repetitions of c and c_- , as per Proposition 15, ensuring that the only affected values are above v_x .

If c is negative, starting at position x with counter value v_x , then $v_x \geq b^6 \cdot P$. As τ_2 starts with counter value cT_{prev} , and $cT_{prev} < P$ (Proposition 16), there must be a positive cycle c_+ that repeats at least $b \cdot P$ times between the last position with counter value $\frac{b^6 \cdot P}{2}$ and x . Hence, we can obtain τ' by removing repetitions of c and c_+ , as per Proposition 15, ensuring that the only affected values are above $b^5 \cdot P$.

2. *The last part $\tau_3 = \tau[z_u .. \ell]$ is of length at least b^8P .*

As in the previous case, a cycle c must repeat in τ_3 at least $b^6 \cdot P$ times. If c is balanced, we can remove $\frac{P}{|c|}$ of its repetitions, getting the desired path τ' .

Otherwise, it must be that τ_3 stays above cT_{prev} , and reaches a value at least $b^6 \cdot P$. Indeed, if c is positive then its repetitions end at some position x with a counter value at least that high, and if it is negative it starts at some position x with a counter value at least that high.

If the counter value also drops to $\frac{b^6 \cdot P}{2}$ after position x , then we can remove positive and negative cycles exactly as in the previous case. Otherwise, we can just remove $\frac{P}{|c|}$ repetitions of c , guaranteed that the counter value at the end of τ_3 is above cT_{prev} .

3. *Only the first part $\tau_1 = \tau[0 .. z_f]$ is of length at least b^8P .*

If any of the other parts is long, we shorten them. Otherwise, their combined length is less than $2b^8 \cdot P < \frac{sT}{2}$, implying that the first part τ_1 is longer than $\xi_i(v) + \frac{sT}{2}$.

Hence, by Lemma 17, there are “fast” and “slow” cycles c_f and c_s , respectively, of slopes $s_f < s_s$, such that each of them repeats at least $b^4 \cdot P$ times in τ_1 .

Thus, by Proposition 15, we can add and/or remove some repetitions of c_f and c_s , such that τ_1 is shorten by exactly P . Yet, we should ensure that the resulting path τ' is valid, in the sense that its corresponding first part τ'_1 cannot get the counter value to 0. We show it by cases:

- If c_f or c_s are balanced cycles, then we can remove the balanced cycle only, without changing the remaining counter values.
- If there is a positive cycle c_+ that repeats at least $2b^2 \cdot P$ times, then the counter value climbs by at least $2b^2 \cdot P$ from its value v_x at position x where c_+ starts and the position y where its repetitions end. As the counter gets down to cT_{prev} at the end of τ_1 , and $cT_{prev} < P$ (Proposition 16), there must be a negative cycle c_- that repeats at least $b \cdot P$ times between position y and the first position after y that has the counter value $v_x + b^2 \cdot P$. Hence, we can remove repetitions of c_+ and c_- , as per Proposition 15, ensuring that the only affected values are above v_x .
- Otherwise, both c_f and c_s are negative, implying that we add some repetitions of c_f and remove some repetitions of c_s . We further split into two subcases:
 - If c_s appears before c_f then there is no problem, as the only change of values will be their increase, and all the values were nonzero to begin with (as we are before z_f).
 - If c_f appears first, ending at some position x , while c_s starts at some later position y , then a-priori it might be that repeating c_f up to $b \cdot P$ more times, as per Proposition 15, will take the counter value to 0.

However, observe that since there are at most $b - 2$ positive cycles, and each of them can repeat at most $2b^2 \cdot P - 1$ times, the counter value v_x at position x , and along the way until position y , is at least $v_y - (b - 2)2b^2 \cdot P$, where v_y is the counter value at position y . As c_s repeats at least $b^4 \cdot P$ times, we have $v_y \geq b^4 \cdot P$. Thus $v_x \geq b^4 \cdot P - 2(b - 2)b^2 \cdot P > b^2P$. Hence, repeating c_f up to $b \cdot P$ more times at position x cannot take the counter value to 0, until position y , as required.

Proof of Lemma 18.2a

The case of Segment ξ_0

For a path of length ℓ , we have in Segment 0 that $\text{shift}(\ell) = \ell$, and indeed a path from v is valid from $v + P$ and vice versa, as they do not hit cT_{prev} : Their maximal drop is sT , while $v \geq cT > P + sT > cT_{prev} + sT$.

We turn to the i th segment, for $i \geq 1$. Consider a basic path τ for $(s, v + P) \xrightarrow{\text{shift}(\ell)} (e, u)$. Recall that $\text{shift}(\ell) = \ell + \frac{P}{-s_i} \in [\S_i(v + P) .. \S_i(v + P) + \mathbf{sT}]$. We construct from τ a path τ' for $(s, v) \xrightarrow{\ell} (e, u')$, such that $u \equiv u'$, along the following cases.

Case 2a.1: The counter values in τ stay above cT_{prev}

As there is no position in τ with counter value cT_{prev} , then in particular τ has no zero-transitions. We further split into two subcases:

1. If τ does not have $b \cdot P$ repetitions of a “relatively fast” cycle with slope \mathbf{s}_j for $j \leq i$, then the drop of τ , and of every prefix of it, is at most $X + Y$, where X stands for the drop outside “slow” cycles of slope \mathbf{s}_h for $h > i$, and Y for the rest of the drop. We have $X < b^3 \cdot P$ and $Y < (\ell + \frac{P}{-s_i})(-\mathbf{s}_{i+1})$.

We claim that we can obtain τ' by removing $\frac{P}{-s_i \cdot |c|}$ repetitions of any cycle c , which repeats enough in τ , having that the drop D of τ' is less than $v - cT_{prev}$.

Indeed, the maximal such drop D might be the result of removing only cycles of slope $(+1)$, whose total effect is $\frac{P}{-s_i}$, having $D \leq \frac{P}{-s_i} + X + Y = \frac{P}{-s_i} + b^3 \cdot P + (\ell + \frac{P}{-s_i})(-\mathbf{s}_{i+1}) \leq b \cdot P + b^3 \cdot P + (\ell + \frac{P}{-s_i})(-\mathbf{s}_{i+1})$. Since $\ell < \S_{i+1}(v + P) = \frac{1}{-s_{i+1}}(v + P - cT_{prev}) - b^8 \cdot P$, we have $D \leq (b^3 + b) \cdot P + (\frac{1}{-s_{i+1}}(v + P - cT_{prev}) - b^8 \cdot P) + \frac{P}{-s_i}(-\mathbf{s}_{i+1}) = (b^3 + b) \cdot P + v + P - cT_{prev} - (-\mathbf{s}_{i+1}) \cdot b^8 \cdot P + \frac{(-\mathbf{s}_{i+1})}{-s_i} \cdot P = (b^3 + b + 1 + \frac{(-\mathbf{s}_{i+1})}{-s_i}) \cdot P + v - cT_{prev} - (-\mathbf{s}_{i+1}) \cdot b^8 \cdot P < b^4 \cdot P - b^7 \cdot P + v - cT_{prev}$. It is thus left to show that $b^4 \cdot P < b^7 \cdot P$, which obviously holds.

2. Otherwise, namely when τ does have $b \cdot P$ repetitions of a “relatively fast” cycle with slope \mathbf{s}_j for $j \leq i$, let c be the first such cycle in τ . We can obtain τ' by removing $\frac{P}{-s_i \cdot |c|}$ repetitions of c : The counter value in τ' , which starts with counter value v , at the position after the repetitions of c will be at least as high as the counter value in τ , which starts with counter value $v + P$, after the repetitions of c . Notice that the counter value cannot hit cT_{prev} before arriving to the repetitions of c by the argument of the previous subcase.

Case 2a.2: τ reaches counter value cT_{prev}

Again let $\tau_1 = \tau[0 .. z_f], \tau_2 = \tau[z_f .. z_u], \tau_3 = \tau[z_u .. \text{shift}(\ell)]$ as in 1a.

In order to handle possible zero transitions, we shorten τ_1 , such that the resulting first part τ'_1 of τ' , which starts with counter value v , also ends with counter value exactly cT_{prev} . Since τ_1 reaches cT_{prev} and is shorter than $\S_{i+1}(v + P)$, it has by Lemma 17.2 at least $b^4 \cdot P$ repetitions of a “fast” cycle of slope $\mathbf{s}_f \leq \mathbf{s}_i$. Let c_f be the first such cycle. We split to cases.

1. If $\mathbf{s}_f = \mathbf{s}_i$ or τ_2 or τ_3 are of length at least $b^5 \cdot P$, we can remove $\frac{P}{-s_f \cdot |c_f|}$ repetitions of c_f in τ_1 . Note that the resulting first part τ'_1 of τ' indeed ends with counter value cT_{prev} . However, while when $\mathbf{s}_f = \mathbf{s}_i$ the resulting length of τ' will be ℓ , as required, in the case that $\mathbf{s}_f < \mathbf{s}_i$, we have that τ' will be longer than ℓ . Nevertheless, in this case, as τ_2 or τ_3 are of length at least $b^5 \cdot P$, we can further shorten τ_2 or τ_3 without changing their effect, as per Proposition 15, analogously to 1 or 2, respectively, in the proof of Lemma 18.1a.2.
2. Otherwise, we are in the case that τ_1 has a “really fast” cycle of slope $\mathbf{s}_f < \mathbf{s}_i$ that repeats at least $b^4 \cdot P$ times, and both τ_2 or τ_3 are of length less than $b^5 \cdot P$. We claim that in this case τ_1 must also have $b^4 \cdot P$ repetitions of a “relatively slow” cycle c_s of slope $\mathbf{s}_s \geq \mathbf{s}_i$. Indeed, assume toward contradiction that τ_1 has less than $b^4 \cdot P$ repetitions of a cycle with slope \mathbf{s}_s for $s \geq i$. Then the longest such path has less than b transitions of $(+1)$ out of cycles, $b^6 \cdot P$ such transitions in cycles, and the rest of it consists of “fast” cycles with slope indexed lower than i .

Thus its length is at most $X + L$, where $X = b^6 \cdot P$ is the $(+1)$ transitions, and L is the longest length to drop from counter value $v + P + X$ to cT_{prev} with “fast” cycles. Thus, $L \leq \frac{1}{-s_{i-1}}(v + P + X - cT_{prev})$.

Now, we have that the length of τ is at least $\xi_i(v + P) = \frac{1}{-s_i}(v - cT_{prev}) - b^8 \cdot P$. Thus, parts τ_2 and τ_3 of τ are of length at least $Z = \frac{1}{-s_i}(v - cT_{prev}) - b^8 \cdot P - \frac{1}{-s_{i-1}}(v + P + X - cT_{prev}) - X = (\frac{1}{-s_i} - \frac{1}{-s_{i+1}})(v - cT_{prev}) - b^8 \cdot P - (1 + \frac{1}{-s_{i+1}})b^6 \cdot P$.

Since $(\frac{1}{-s_i} - \frac{1}{-s_{i+1}}) \geq \frac{1}{b^2}$, $(1 + \frac{1}{-s_{i+1}}) \leq b+1$, and $v - cT_{prev} \geq cT - cT_{prev} > cT - P > 3b^{10}P$, we have $Z \geq \frac{1}{b^2}(3b^{10} \cdot P) - b^8 \cdot P - (b+1)b^6 \cdot P = 2b^8 \cdot P - (b+1)b^6 \cdot P$. Therefore, at least one of τ_2 and τ_3 is of length at least $b^7 \cdot P$, leading to a contradiction.

So, we are in the case that τ_1 has at least $b^4 \cdot P$ repetitions of a “really fast” cycle c_f of slope $s_f < s_i$ as well as $b^4 \cdot P$ repetitions of a “relatively slow” cycle c_s of slope $s_s \geq s_i$. By analyzing the different possible orders of c_s and c_f , we can cut and repeat the cycles far enough from 0 so as to construct valid paths. See the Appendix for details. ◀

5.1 Proof of Theorem 11

Proof. The proof is by induction over the structure of φ , where Theorem 5 already provides the periodicity for all CTL operators.

It remains to plug UA into the induction by showing (1) the (cT, P) -periodicity of a formula $\varphi = \psi_1 U A \psi_2$ with respect to an OCA \mathcal{A} , provided that its subformulas are (cT_{prev}, P_{prev}) -periodic; and (2) by showing that the period P and threshold cT are single-exponential in $n = |\mathcal{A}|$ and in the nesting depth of φ .

1. We show that for every state $s \in S$ and counters $v, v' > cT$, if $v \equiv v' \pmod{P}$ then $(s, v) \models \varphi \iff (s, v') \models \varphi$. Without loss of generality, write $v' = v + z \cdot P$, for some $z \in \mathbb{N}$.
 - If $(s, v) \models \varphi$ then by the definition of the AU operator and the completeness of \mathcal{A} we have i) there is a level ℓ , such that for every state e and counter value u , if $(s, v) \xrightarrow{\ell} (e, u)$ then $(e, u) \models \psi_2$, and ii) for every level $m < \ell$, state h and counter value x , if $(s, v) \xrightarrow{m} (h, x)$ then $(h, x) \models \psi_1$.

Observe first that if $\ell \notin \text{core}(v)$, then there also exists a level $\hat{\ell} < \ell$ witnessing $(s, v) \models \varphi$, such that $\hat{\ell} \in \text{core}(v)$. Indeed, we obtain $\hat{\ell}$, by choosing the largest level $\hat{\ell}$ in the core of v 's segment, such that $\ell \equiv \hat{\ell} \pmod{P}$. As $\hat{\ell} < \ell$, it directly follows that for every level $\hat{m} < \hat{\ell}$, state \hat{h} and counter value \hat{x} , if $(s, v) \xrightarrow{\hat{m}} (\hat{h}, \hat{x})$ then $(\hat{h}, \hat{x}) \models \psi_1$. Now, assume toward contradiction that there is a state \hat{e} and a counter value \hat{u} , such that $(s, v) \xrightarrow{\hat{\ell}} (\hat{e}, \hat{u})$ and $(\hat{e}, \hat{u}) \not\models \psi_2$. Then by (possibly several applications of) Lemma 18.1b, there is also a counter value $\hat{u} \equiv_{cT_{prev}, P_{prev}} \hat{u}$, such that $(s, v) \xrightarrow{\hat{\ell}} (\hat{e}, \hat{u})$. As ψ_2 is (cT_{prev}, P_{prev}) -periodic, we have $(\hat{e}, \hat{u}) \not\models \psi_2$, leading to a contradiction.

Next, we claim that the level $\ell' = \text{shift}^z(\hat{\ell})$, obtained by z (recall that $z = \frac{v'-v}{P}$) applications of the shift function on $\hat{\ell}$, witnesses $(s, v') \models \varphi$, namely that i) for every state e' and counter value u' , if $(s, v') \xrightarrow{\ell'} (e', u')$ then $(e', u') \models \psi_2$, and ii) for every level $m' < \ell'$, state h' and counter value x' , if $(s, v') \xrightarrow{m'} (h', x')$ then $(h', x') \models \psi_1$.

Indeed, i) were it the case that $(e', u') \not\models \psi_2$ then by (z applications of) Lemma 18.2a, there was also a counter value $u'' \equiv_{cT_{prev}, P_{prev}} u'$, such that $(s, v) \xrightarrow{\hat{\ell}} (e', u'')$ and therefore $(e', u'') \not\models \psi_2$, leading to a contradiction; and ii) were it the case that $(s, v') \xrightarrow{m'} (h', x')$ and $(h', x') \not\models \psi_1$ then a) by Lemma 18.1a, as in the argument above, there is also a level $\tilde{m} \leq m'$, such that \tilde{m} is in the core of m' 's segment and $(s, v') \xrightarrow{\tilde{m}} (h', \tilde{x})$ where $\tilde{x} \equiv_{cT_{prev}, P_{prev}} x$, and b) by Lemma 18.2a there is a level $\hat{m} < \hat{\ell}$, such that $(s, v) \xrightarrow{\hat{m}} (h', \hat{x})$ where $\hat{x} \equiv_{cT_{prev}, P_{prev}} x$ and therefore $(h', \hat{x}) \not\models \psi_1$, leading to a contradiction.

- If $(s, v') \models \varphi$ then we have $(s, v) \models \varphi$ by an argument analogous to the above, while using Lemma 18.2b instead of Lemma 18.2a.
- 2. The threshold cT and period P are calculated along the induction on the structure of the formula φ . They start with threshold 0 and period 1, and their increase in each step of the induction depends on the outermost operator.

Observe first that we can take as the worst case the same increase in every step, that of the UA case, since it guarantees the others. Namely, its required threshold, based on the threshold and period of the subformulas, is bigger than the threshold required in the other cases, and its required period is divisible by the periods required for the other cases. Next, notice that both the threshold and period in the UA case only depend on the periods of the subformulas (i.e., not on their thresholds), so it is enough to show that the period is singly exponential in n and the nesting depth of φ .

The period in the UA case is defined to be $P(\varphi) = B \cdot \text{lcm}(P(\psi_1), P(\psi_2))$, where $B = \text{lcm}[1..2b^3]$, and b is the bound on the length of a linear path scheme for \mathcal{A} . By [17], b is polynomial in n , and as $\text{lcm}[1..2b^3] < 4^{2b^3}$, we get that B is singly exponential in n .

Considering $\text{lcm}(P(\psi_1), P(\psi_2))$, while in general $\text{lcm}(x, y)$ of two numbers x and y might be equal to their multiplication, in our case, as both ψ_1 and ψ_2 are calculated along the induction via the same scheme above, they are both an exponent of B . Hence, $\text{lcm}(P(\psi_1), P(\psi_2)) = \max(P(\psi_1), P(\psi_2))$. Thus, we get that $P(\varphi) \leq B^x$, where x is bounded by the nesting depth of φ . ◀

References

- 1 Shaull Almagor, Udi Boker, and Orna Kupferman. Formalizing and reasoning about quality. *Journal of the ACM*, 63(3):24:1–24:56, 2016.
- 2 Roland Axelsson, Matthew Hague, Stephan Kreutzer, Martin Lange, and Markus Latte. Extended computation tree logic. In *Logic for Programming, Artificial Intelligence, and Reasoning: 17th International Conference, LPAR-17, Yogyakarta, Indonesia, October 10-15, 2010. Proceedings 17*, pages 67–81. Springer, 2010.
- 3 Udi Boker, Krishnendu Chatterjee, Thomas A. Henzinger, and Orana Kupferman. Temporal specifications with accumulative values. *ACM Trans. Comput. Log.*, 15(4):27:1–27:25, 2014.
- 4 Krishnendu Chatterjee and Laurent Doyen. Computation tree logic for synchronization properties. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPICs*, pages 98:1–98:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016.
- 5 Michael R Clarkson, Bernd Finkbeiner, Masoud Koleini, Kristopher K Micinski, Markus N Rabe, and César Sánchez. Temporal logics for hyperproperties. In *Principles of Security and Trust: Third International Conference, POST 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings 3*, pages 265–284. Springer, 2014.
- 6 Michael R Clarkson and Fred B Schneider. Hyperproperties. *Journal of Computer Security*, 18(6):1157–1210, 2010.
- 7 Byron Cook, Heidy Khlaaf, and Nir Piterman. Verifying increasingly expressive temporal logics for infinite-state systems. *Journal of the ACM (JACM)*, 64(2):1–39, 2017.
- 8 David C Cooper. Theorem proving in arithmetic without multiplication. *Machine intelligence*, 7(91-99):300, 1972.
- 9 Stéphane Demri, Alain Finkel, Valentin Goranko, and Govert van Drimmelen. Model-checking CTL* over flat Presburger counter systems. *Journal of Applied Non-Classical Logics*, 20(4):313–344, 2010.

- 10 Alain Finkel, Jérôme Leroux, and Grégoire Sutre. Reachability for two-counter machines with one test and one reset. In *FSTTCS 2018-38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, volume 122, pages 31–1. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2018.
- 11 Michael J Fischer and Michael O Rabin. Super-exponential complexity of Presburger arithmetic. In *Quantifier Elimination and Cylindrical Algebraic Decomposition*, pages 122–135. Springer, 1998.
- 12 Seymour Ginsburg and Edwin H Spanier. Bounded ALGOL-like languages. *Transactions of the American Mathematical Society*, 113(2):333–368, 1964.
- 13 Stefan Göller and Markus Lohrey. Branching-time model checking of one-counter processes and timed automata. *SIAM J. Comput.*, 42(3):884–923, 2013.
- 14 Christoph Haase. A survival guide to Presburger arithmetic. *ACM SIGLOG News*, 5(3):67–82, 2018.
- 15 Christoph Haase, Stephan Kreutzer, Joël Ouaknine, and James Worrell. Reachability in succinct and parametric one-counter automata. In *CONCUR 2009 – Concurrency Theory: 20th International Conference, CONCUR 2009, Bologna, Italy, September 1-4, 2009. Proceedings 20*, pages 369–383. Springer, 2009.
- 16 Jérôme Leroux and Grégoire Sutre. Flat counter automata almost everywhere! In *International Symposium on Automated Technology for Verification and Analysis*, pages 489–503. Springer, 2005.
- 17 Jérôme Leroux and Grégoire Sutre. Reachability in two-dimensional vector addition systems with states: One test is for free. In *31st International Conference on Concurrency Theory (CONCUR 2020)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020.
- 18 Marvin L. Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, Englewood Cliffs, N.J., 1967.
- 19 Mojzesz Presburger. Über die Vollständigkeit eines gewissen Systems der Arithmetik ganzer Zahlen, in welchem die Addition als einzige Operation hervortritt. *Comptes Rendus du I Congrès de Mathématiciens des Pays Slaves*, pages 92–101, 1929.
- 20 Leslie G Valiant and Michael S Paterson. Deterministic one-counter automata. *Journal of Computer and System Sciences*, 10(3):340–350, 1975.
- 21 Igor Walukiewicz. Model checking CTL properties of pushdown systems. In *International Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 127–138. Springer, 2000.

A Appendix – Omitted Proofs

A.1 Proofs for Sections 3 and 4

Proof of Lemma 6. In [17], the authors consider a model called 2-TVASS, comprising a 2-dimensional vector addition system where the first counter can also be tested for zero. In our setting, this corresponds to an OCA equipped with an additional counter that cannot be tested for zero. A configuration of a 2-TVASS is (s, x_1, x_2) describing the state and the values of the two counters.

We transform the OCA \mathcal{A} to a 2-TVASS \mathcal{A}' by introducing a length-counting component. That is, in every transition of \mathcal{A}' as a 2-TVASS, the second component increments by 1. We thus have the following: there is a path of length ℓ from (s, v) to (s', v') in \mathcal{A} if and only if there is a path from $(s, v, 0)$ to (s', v', ℓ) in \mathcal{A}' .

From [17, Corollary 16], using the fact that \mathcal{A}' has only weights in $\{-1, 0, 1\}$, we have that if there is a path from (s, x_1, x_2) to (s', y_1, y_2) in \mathcal{A}' , then there is also such a π -shaped path where π is of flat length $\text{poly}(|S|)$ and size $O(|S|^3)$, which concludes the proof. ◀

19:18 Synchronized CTL over One-Counter Automata

Proof of Lemma 8. Consider the set $V_{\text{init}} = \{v \mid v \leq \mathfrak{t}(\varphi) \text{ and } (s, v) \models \varphi\}$. We define

$$P_{\varphi, s}(x) := \bigvee_{u \in V_{\text{init}}} \exists m. x = u + m \cdot \mathfrak{p}(\varphi)$$

The correctness of this formula is immediate from the definition of V_{init} , $\mathfrak{t}(\varphi)$, and $\mathfrak{p}(\varphi)$. In order to compute $P_{\varphi, s}(x)$, we need to compute V_{init} . This is done using Proposition 4, by evaluating φ from every state of the Kripke structure. \blacktriangleleft

Proof of Lemma 9. Recall from Observation 2 that the set of satisfying assignments for a PA formula in dimension 1 is effectively semilinear and can be written as $\text{Lin}(B, \{r\}) = \{b + \lambda r \mid b \in B, \lambda \in \mathbb{N}\}$ for effectively computable $B \subseteq \mathbb{N}$ and $r \in \mathbb{N}$.

For every state $s \in S$, consider therefore the set $\text{Lin}(B_s, \{r_s\})$ corresponding to $P_{\varphi, s}$. Define $\mathfrak{t}(\varphi) = \max\{b \mid b \in B_s, s \in S\}$ and $\mathfrak{p}(\varphi) = \text{lcm}(\{r_s\}_{s \in S}) \cdot \mathfrak{t}(\varphi)$.

We claim that φ is $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -periodic. Indeed, consider $v, v' > \mathfrak{t}(\varphi)$ such that $v \equiv v' \pmod{\mathfrak{p}(\varphi)}$, and let $s \in S$. Without loss of generality assume $v > v'$, so we can write $v - v' = K \cdot \mathfrak{p}(\varphi)$ for some $K \in \mathbb{N}$.

Now, if $(s, v') \models \varphi$, then $v' = b'_s + \lambda' r_s$ for some $b'_s \in B_s$ and $\lambda' \in \mathbb{N}$. Then, $v = b'_s + \lambda' r_s + K \cdot \mathfrak{p}(\varphi) = b'_s + (\lambda' + K \frac{\mathfrak{p}(\varphi)}{r_s}) r_s$, and note that r_s divides $\mathfrak{p}(\varphi)$, so we have $v \in \text{Lin}(B_s, \{r_s\})$ and therefore $(s, v) \models \varphi$.

Conversely, if $(s, v) \models \varphi$ then $v = b_s + \lambda r_s$ for some $b_s \in B_s$ and $\lambda \in \mathbb{N}$. Then, $v' = b_s + \lambda r_s - K \cdot \mathfrak{p}(\varphi) = b_s + (\lambda - K \frac{\mathfrak{p}(\varphi)}{r_s}) r_s$. Since $v' > \mathfrak{t}(\varphi) \geq b_s$, then $\lambda - K \frac{\mathfrak{p}(\varphi)}{r_s} \geq 1$. Thus, $v' \in \text{Lin}(B_s, \{r_s\})$ so $(v', s) \models \varphi$. Hence, φ is $(\mathfrak{t}(\varphi), \mathfrak{p}(\varphi))$ -periodic. \blacktriangleleft

Proof of Lemma 17. For $i = 0$, the first claim trivially holds, as it claims for a cycle with a slope at least \mathfrak{s}_1 , which holds for any cycle, and since $\frac{\mathfrak{s}T}{2} > 2b^5 \cdot P$, it must repeat at least $b^4 \cdot P$ times. The second claim vacuously holds, as no path can reach $\mathfrak{cT}_{\text{prev}}$ at Segment 0.

Consider $i \in [1..m]$. We prove each of the claims separately.

We show that there are at least $b^5 \cdot P$ total cycle repetitions of the required slope, and since there are at most b different cycles, one of them should repeat at least $b^4 \cdot P$ times.

1. Consider a basic path τ of length $\ell \geq \mathfrak{s}_i(v) + \frac{\mathfrak{s}T}{2}$, such that the counter values of τ stay above $\mathfrak{cT}_{\text{prev}}$. Assume by way of contradiction that τ has less than $b^5 \cdot P$ cycles with \mathfrak{s}_j for $j > i$. We will show that ℓ must fall short of $\mathfrak{s}_i(v) + \frac{\mathfrak{s}T}{2}$.

Since there are at most $b^5 \cdot P - 1$ cycles with \mathfrak{s}_j with $j > i$, each of length at most b , then the total length of τ spent in such cycles and in the simple paths of τ (of total length at most b) is $X \leq b(b^5 \cdot P - 1) + b = b^6 \cdot P$. The effect accumulated by the transitions in X is at most X , if all the relevant transitions are $+1$.

The remaining part of the path, whose length is denoted by $Y = \ell - X$, is spent in cycles with \mathfrak{s}_j for $j \leq i$. Note that its effect is at most $(v + X - \mathfrak{cT}_{\text{prev}})$, and therefore its length $Y \leq \frac{-1}{\mathfrak{s}_i}(v + X - \mathfrak{cT}_{\text{prev}})$.

Observe that any actual path τ that satisfies the required constraints cannot be shorter than the above theoretical path in which there are X transitions of effect $(+1)$.

Therefore, we have $\ell \leq X + Y \leq b^6 \cdot P + \frac{-1}{\mathfrak{s}_i}(v + b^6 \cdot P - \mathfrak{cT}_{\text{prev}}) = \frac{-1}{\mathfrak{s}_i}(v - \mathfrak{cT}_{\text{prev}}) + (\frac{-1}{\mathfrak{s}_i} + 1)b^6 \cdot P \leq \frac{-1}{\mathfrak{s}_i}(v - \mathfrak{cT}_{\text{prev}}) + (b^7 + b^6) \cdot P$.

Recall, however, that $\ell \geq \mathfrak{s}_i(v) + \frac{\mathfrak{s}T}{2} = (\frac{-1}{\mathfrak{s}_i}(v - \mathfrak{cT}_{\text{prev}}) - b^8 \cdot P) + (\frac{1}{2}b^9 \cdot P)$, leading to a contradiction, as $\frac{b^9}{2} > b^6 + b^7 + b^8$, for $b > 2$.

2. Consider a basic path τ of length $\ell < \mathfrak{s}_{i+1}(v)$, such that the counter values of τ reach $\mathfrak{cT}_{\text{prev}}$. Without loss of generality, we can assume that ℓ is the first time when τ reaches $\mathfrak{cT}_{\text{prev}}$ (otherwise we will look at the prefix of τ that satisfies this). Assume by way of contradiction that τ has less than $b^5 \cdot P$ cycles with \mathfrak{s}_j for $j < i$, we show that ℓ must fall after $\mathfrak{s}_{i+1}(v)$.

Since there are at most $b^5 \cdot P - 1$ cycles with \mathbf{s}_j for $j \leq i$, each of length at most b , then the total length of τ spent in such cycles and in the simple paths of τ (of total length at most b) is $X \leq b(b^5 \cdot P - 1) + b = b^6 \cdot P$. The effect accumulated by the transitions in X is at least $-X$, if all the relevant transitions are -1 .

The remaining part of the path, whose length is denoted by $Y = \ell - X$, is spent in cycles with \mathbf{s}_j for $j > i$. Note that its effect is at least $(v - X - cT_{prev})$, and therefore its length $Y \geq \frac{-1}{\mathbf{s}_{i+1}}(v - X - cT_{prev})$.

Observe that any actual path τ that satisfies the required constraints cannot be shorter than the above theoretical path in which there are X transitions of effect (-1) .

Therefore, we have $\ell \geq X + Y \geq b^6 \cdot P + \frac{-1}{\mathbf{s}_{i+1}}(v - b^6 \cdot P - cT_{prev}) = \frac{-1}{\mathbf{s}_{i+1}}(v - cT_{prev}) + (\frac{-1}{\mathbf{s}_{i+1}} - 1)b^6 \cdot P \geq \frac{-1}{\mathbf{s}_{i+1}}(v - cT_{prev}) - b^7 \cdot P$.

Recall, however, that $\ell < \xi_{i+1}(v) = \frac{-1}{\mathbf{s}_{i+1}}(v - cT_{prev}) - b^8 \cdot P$, leading to a contradiction. \blacktriangleleft

A.2 Remaining Cases in the Proof of Lemma 18

We present remaining arguments for the proof.

A.2.1 Proof of Lemma 18.1b

Consider a length $\ell \notin \text{core}(v)$, and let τ be a basic path of length $\ell - P$ such that $(s, v) \xrightarrow{\ell-P} (e, u)$. We construct from τ a path τ' for $(s, v) \xrightarrow{\ell} (e, u')$, such that $u \equiv u'$.

This part of the proof, showing that a path can be properly lengthened by P , is partially analogous to part 1a of the proof, which shows that a path can be shortened by P . We split the proof to the same cases as in part 1a, and whenever the proofs are similar enough, we refer to the corresponding case in part 1a.

Case 1b.1: The counter values in τ stay above cT_{prev}

Since there is no position in τ with counter value cT_{prev} , then in particular τ has no zero-transitions. We split to subcases.

- If there is a non-negative cycle c in τ , we can obtain τ' by adding $\frac{P}{|c|}$ copies of c , having that the counter values in τ' are at least as high as the corresponding ones in τ . Observe that τ' is of length $\ell + P$ from (s, v) to (e, u') with $u' = u + \text{effect}(c) \frac{P}{|c|}$. Since we have $u' \geq u \geq cT_{prev}$, then $u \equiv u'$.
- If there are in τ two cycles of different slopes, such that each of them repeats at least $b \cdot P$ times, we can use Proposition 14 to properly lengthen τ into τ' .
- Otherwise, we are in the case that only a single cycle c with some negative slope \mathbf{s} repeats at least $b \cdot P$ times. Since $\ell - P > \xi_i(v) + \frac{\mathbf{s}T}{2}$, by Lemma 17.1, we have $\mathbf{s} > \mathbf{s}_i$. We can thus add $\frac{P}{|c|}$ copies of c , guaranteed that the counter values in τ' remain above cT_{prev} . Indeed, the maximal drop (i.e., the negation of the minimum cumulative counter effect in the path) D of τ is $X + Y$, where X stands for the repetitions of c , and Y for the rest. We have $X \leq -\mathbf{s} \cdot (\ell - P)$ (notice that \mathbf{s} is negative) and $Y \leq b + (b - 1)b^2P$, considering the worst case in which all transitions out of cycles, as well as in the cycles that are not c are of effect (-1) . Thus, $D \leq -\mathbf{s} \cdot (\ell - P) + b^3P$.
As $\ell - P < \ell < \xi_{i+1}(v) = \frac{(v - cT_{prev})}{-\mathbf{s}_{i+1}} - b^8 \cdot P$, we have $D \leq \frac{\mathbf{s}}{\mathbf{s}_{i+1}}(v - cT_{prev}) + b^3 \cdot P - (-\mathbf{s})b^8 \cdot P$. Since $\mathbf{s} > \mathbf{s}_i$, it follows that $\mathbf{s} \geq \mathbf{s}_{i+1}$, thus (as both \mathbf{s} and \mathbf{s}_{i+1} are negative), we have $D \leq (v - cT_{prev}) + b^3 \cdot P - (-\mathbf{s})b^8 \cdot P$.

19:20 Synchronized CTL over One-Counter Automata

Recall that we aim to show that $v - D$ can “survive” a drop of up to P additional copies of c (i.e., at most $b \cdot P$), while keeping the counter value above cT_{prev} . Thus, we need to show that $v - ((v - cT_{prev}) + b^3 \cdot P - (-s)b^8P) - b \cdot P > cT_{prev}$. It is left to show that $(-s)b^8 \cdot P > (b^3 + b) \cdot P$, which obviously holds, as $-s \geq \frac{1}{b}$.

Case 1b.2: τ reaches counter value cT_{prev}

Let $0 \leq z_f \leq z_u \leq \ell - P$ be the first and ultimate positions in τ where the counter value is exactly cT_{prev} . We split τ into three parts: $\tau_1 = \tau[0 .. z_f)$, $\tau_2 = \tau[z_f .. z_u)$, $\tau_3 = \tau[z_u .. \ell - P]$ as in the analogous case in 1a.

1. *The middle part $\tau_2 = \tau[z_f .. z_u]$ or last part $\tau_3 = \tau[z_u .. \ell - P]$ are of length at least b^8P .* These cases are analogous to their counterparts 1 and 2 of 1a.2, by adding cycle repetitions instead of removing them.
2. *Only the first part $\tau_1 = \tau[0 .. z_f)$ of length at least b^8P*

If any of the other parts is long, we lengthen them. Otherwise, their combined length is less than $2b^8 \cdot P < \frac{sT}{2}$, implying that the first part τ_1 is longer than $\frac{sT}{2}(v) + \frac{sT}{2}$.

Hence, by Lemma 17, there are “fast” and “slow” cycles c_f and c_s , respectively, of slopes $s_f < s_s$, such that each of them repeats at least $b^4 \cdot P$ times in τ_1 .

Thus, by Proposition 15, we can add and/or remove some repetitions of c_f and c_s , such that τ_1 is longer by exactly P . Yet, we should ensure that τ' is valid, in the sense that its corresponding first part τ'_1 cannot get the counter value to 0. We show it by cases:

- If c_f or c_s are balanced, we can lengthen without changing the remaining counters.
- If there is a positive cycle c_+ that repeats at least $2b^2 \cdot P$ times, then the counter value climbs by at least $2b^2 \cdot P$ from its value v_x at position x where c_+ starts and the position y where its repetitions end. As the counter gets down to $cT_{prev} < P$ at the end of τ_1 , and $cT_{prev} < P$ (Proposition 16), there must be a negative cycle c_- that repeats at least $b \cdot P$ times between position y and the first position after y that has the counter value $v_x + b^2 \cdot P$. Hence, we can add repetitions of c_+ and c_- , as per Proposition 15, ensuring that the only affected values are above v_x .
- Otherwise, both c_f and c_s are negative, implying that we add some repetitions of c_s and remove some repetitions of c_f . We further split into two subcases:
 - If c_f appears before c_s then there is no problem, as the only change of values will be their increase.
 - If c_s appears first, ending at some position x , while c_f starts at some position y , then a-priori it might be that repeating c_s up to $b \cdot P$ more times, as per Proposition 15, will take the counter value to 0.

Yet, observe that since there are at most $b - 2$ positive cycles, and each of them can repeat at most $2b^2 \cdot P - 1$ times, the counter value v_x at position x , and along the way until position y , is at least $v_y - (b - 2)2b^2 \cdot P$, where v_y is the counter value at position y . As c_f repeats at least $b^4 \cdot P$ times, we have $v_y \geq b^4 \cdot P$. Thus $v_x \geq b^4 \cdot P - 2(b - 2)b^2 \cdot P > b^2P$. Hence, repeating c_s up to $b \cdot P$ more times at position x cannot take the counter value to 0, until position y , as required.

Continuation of the Proof of Lemma 18.2a

We detail the subcases of Case 2a.2.2 that were not detailed in the main text:

- If there is such a relatively slow cycle c_s that ends in a position x before a position y in which the first really fast cycle c_f starts, then: If $\mathbf{s}_s = s_i$, we can simply remove $\frac{P}{-\mathbf{s}_s \cdot |c_s|}$ repetitions of c_s , getting the desired counter and lengths changes.

Otherwise, namely when $\mathbf{s}_s > s_i$, we claim that the counter value between positions x and y is high enough, allowing to remove some repetitions of c_s and c_f , as per Proposition 14. As c_s might be positive, removal of its repetitions might decrease the counter; the worst case is when all its transitions are of effect $(+1)$, and we shorten the path the most, which is bounded by $b \cdot P$. To ensure that the counter remains above the value P after the removal and until position y , we thus need it to be at least $b \cdot P + P$ at all positions until y . Indeed, the path has only cycles of slope at least \mathbf{s}_{i+1} until position $y < \xi_i + \mathbf{sT}$. Therefore, its drop D is at most b , for the non-cyclic part, plus $(-\mathbf{s}_{i+1})(\xi_i(v + P) + \mathbf{sT})$ for the cyclic part. (Notice that \mathbf{s}_{i+1} is negative.) Therefore, $D \leq b + (-\mathbf{s}_{i+1})(\xi_i(v + P) + \mathbf{sT})$. Taking the value of $\xi_i(v + P)$, we have $D \leq b + (-\mathbf{s}_{i+1})(\frac{-1}{\mathbf{s}_i}(v + P - \mathbf{cT}_{prev}) - b^8 \cdot P + \mathbf{sT}) = b + \frac{\mathbf{s}_{i+1}}{\mathbf{s}_i}(v + P - \mathbf{cT}_{prev}) + (\mathbf{s}_{i+1})(b^8 \cdot P + \mathbf{sT})$. Since $\mathbf{s}_{i+1} \leq \frac{1}{b}$ and $\frac{\mathbf{s}_i}{\mathbf{s}_i} \leq 1 - \frac{1}{b^2}$, we have $D \leq b + (1 - \frac{1}{b^2})(v + P - \mathbf{cT}_{prev}) + \frac{1}{b}(b^8 \cdot P + \mathbf{sT}) \leq b + (1 - \frac{1}{b^2})(v + P - \mathbf{cT}_{prev}) + (b^7 \cdot P + b^8 \cdot P)$. Now, as the counter starts with value $v + P$, we need to show that $E = v + P - D - (b^2 + 1)P \geq 0$. Indeed, $E \geq \frac{1}{b^2}(v + P) - b^7 \cdot P - b^8 \cdot P - (b^2 + 1)P$. As $v \geq \mathbf{cT} = b^{11} \cdot P$, we have $E \geq (b^9 - b^7 - b^8 - b^2 - 1)P > 0$, as required.

- Otherwise, namely when there is no relatively slow cycle c_s that repeats at least $b^4 \cdot P$ times before the position y in which the first really fast cycle c_f starts:

Let z be the position after y , in which the path has the smallest counter value v_z before a relatively slow cycle that repeats $b^4 \cdot P$ times.

If $v_z > P$, we can remove some repetitions of the really fast and relatively slow cycles, as per Proposition 14; We only remove repetitions of the fast decreasing cycle, so as a result the counter will only grow, and if it is guaranteed to be above P when starting the path from counter value $v + P$, it will not reach 0 when starting from counter value v .

Otherwise, namely when $v_z \leq P$, we can remove $\frac{P}{-\mathbf{s}_f \cdot |c_f|}$ repetitions of the really fast cycle c_f , guaranteeing that τ'_1 remains above 0 until position y .

Observe, however, that the resulting path τ' will be longer than ℓ . Nevertheless, we claim that in this case, the portion of τ_1 from position z until its last position with counter value v_z can be further shorten without changing the path's effect. Indeed, there is a cycle that repeats at least $b^4 \cdot P$ times in this part, having an absolute effect of at least $b^3 \cdot P$, implying that this part of τ_1 climbs up to a value of at least $b^3 \cdot P$, and down again, allowing for the removal of positive and negative cycles, as per Proposition 15, analogously to 1 in the proof of Lemma 18.1a.2.

A.2.2 Proof of Lemma 18.2b

We remark that Segment 0 is identical in the shortening and lengthening argument, and this is a repetition of the main text. For a path of length ℓ , we have in Segment 0 that $\text{shift}(\ell) = \ell$, and indeed a path from v is valid from $v + P$ and vice versa, as they do not hit \mathbf{cT}_{prev} : Their maximal drop is \mathbf{sT} , while $v \geq \mathbf{cT} > P + \mathbf{sT} > \mathbf{cT}_{prev} + \mathbf{sT}$.

We turn to the i th segment, for $i \geq 1$. Consider a basic path τ for $(s, v) \xrightarrow{\ell} (e, u)$. Recall that $\text{shift}(\ell) = \ell + \frac{P}{-\mathbf{s}_i} \in [\xi_i(v + P) .. \xi_i(v + P) + \mathbf{sT}]$. We construct from τ a path τ' for $(s, v + P) \xrightarrow{\text{shift}(\ell)} (e, u')$, such that $u \equiv u'$, along the following cases.

Case 2b.1: The counter values in τ stay above cT_{prev}

Since there is no position in τ with counter value cT_{prev} , then in particular τ has no zero-transitions. Since $\ell > \xi_{i-1}(v) + \frac{sT}{2}$, by Lemma 17.1, there is a cycle c in τ of slope $\mathbf{s} > \mathbf{s}_{i-1}$, namely of slope $\mathbf{s} \geq \mathbf{s}_i$. We can thus get the required path τ' , by adding $\frac{P}{-\mathbf{s} \cdot |c|}$ repetitions of c . Indeed, the counter value of τ' at the position that c starts will be bigger by P than the counter value of τ at that position, while the counter of τ' after the repetitions of c will be at least as high as the counter value of τ at the corresponding position.

Case 2b.2: τ reaches counter value cT_{prev}

Let $0 \leq z_f \leq z_u \leq \ell$ be the first and ultimate positions in τ where the counter value is exactly cT_{prev} . We split τ into three parts: $\tau_1 = \tau[0 .. z_f)$, $\tau_2 = \tau[z_f .. z_u)$, $\tau_3 = \tau[z_u .. \ell]$ (it could be that $z_f = z_u$, in which case the middle part is empty).

In order to accommodate with possible zero transitions, we should lengthen τ_1 , such that the resulting first part τ'_1 of the new path τ' , which starts with counter value $v + P$, will also end with counter value exactly cT_{prev} . Since τ_1 reaches cT_{prev} and it is shorter than $\xi_{i+1}(v)$, it has by Lemma 17.2 at least $b^4 \cdot P$ repetitions of a “fast” cycle c_f of slope $\mathbf{s}_f \leq \mathbf{s}_i$.

1. If $\mathbf{s}_f = \mathbf{s}_i$ or τ_2 or τ_3 are of length at least $b^5 \cdot P$, we can add $\frac{P}{-\mathbf{s}_f \cdot |c_f|}$ repetitions of c_f in τ_1 . Note that the resulting first part τ'_1 of the new path τ' indeed ends with counter value cT_{prev} . However, while when $\mathbf{s}_f = \mathbf{s}_i$ the resulting length of τ' will be $\text{shift}(\ell)$, as required, in the case that $\mathbf{s}_f < \mathbf{s}_i$, we have that the resulting path τ' will be shorter than $\text{shift}(\ell)$. Nevertheless, in this case, as τ_2 or τ_3 are of length at least $b^5 \cdot P$, we can further lengthen τ_2 or τ_3 without changing their effect, as per Proposition 15, analogously to 1 or 2, respectively, in the proof of Lemma 18.1a.2.
2. Otherwise, we are in the case that τ_1 has a “really fast” cycle of slope $\mathbf{s}_f < \mathbf{s}_i$ that repeats at least $b^4 \cdot P$ times, and both τ_2 or τ_3 are of length less than $b^5 \cdot P$. We claim that in this case τ_1 must also have $b^4 \cdot P$ repetitions of a “relatively slow” cycle c_s of slope $\mathbf{s}_s \geq \mathbf{s}_i$. The argument for that is the same as in 2 of part 2a of the proof. As the case of $\mathbf{s}_s = \mathbf{s}_i$ is handled in the previous subcase, we may further assume that $\mathbf{s}_s > \mathbf{s}_i$. We may thus add repetitions of c_f and c_s , as per Proposition 14, lengthening τ_1 by exactly $\text{shift}(\ell) - \ell$, while also ensuring that it ends with counter value exactly cT_{prev} .