

MODEL CHECKING FLAT FREEZE LTL ON ONE-COUNTER AUTOMATA

ANTONIA LECHNER¹, RICHARD MAYR², JOËL OUAKNINE³, AMAURY POULY³,
AND JAMES WORRELL⁴

¹ Computer Science Department, Université Libre de Bruxelles, Belgium
e-mail address: antonia.lechner@ulb.ac.be

² School of Informatics, LFCS, University of Edinburgh, UK
URL: <http://homepages.inf.ed.ac.uk/rmayr/>

³ Max Planck Institute for Software Systems, Saarbrücken, Germany
e-mail address: joel@mpi-sws.org
e-mail address: pamaury@mpi-sws.org

⁴ Department of Computer Science, University of Oxford, UK
e-mail address: james.worrell@cs.ox.ac.uk

ABSTRACT. Freeze LTL is a temporal logic with registers that is suitable for specifying properties of data words. In this paper we study the model checking problem for Freeze LTL on one-counter automata. This problem is known to be undecidable in general and PSPACE-complete for the special case of deterministic one-counter automata. Several years ago, Demri and Sangnier investigated the model checking problem for the flat fragment of Freeze LTL on several classes of counter automata and posed the decidability of model checking flat Freeze LTL on one-counter automata as an open problem. In this paper we resolve this problem positively, utilising a known reduction to a reachability problem on one-counter automata with parameterised equality and disequality tests. Our main technical contribution is to show decidability of the latter problem by translation to Presburger arithmetic.

1. INTRODUCTION

Runs of infinite-state machines, such as counter automata, can naturally be seen as *data words*, that is, sequences in which each position is labelled by a letter from a finite alphabet and a datum from an infinite domain. Freeze LTL is an extension of Linear Temporal Logic with registers and a binding mechanism, which has been introduced to specify properties of data words [3, 4, 8, 12]. The registers allow to compare data at different positions along the same computation.

2012 ACM CCS: [Theory of computation]: Formal languages and automata theory—Automata over infinite objects.

Key words and phrases: One-counter automata, disequality tests, reachability, Freeze LTL, Presburger arithmetic.

An example of a freeze LTL formula is

$$F(v \wedge \downarrow_r \text{XF}(v \wedge \uparrow_r)). \quad (1.1)$$

Evaluated on a run of a one-counter automaton, this formula is true if and only if there are at least two different positions in the run which both have control state v and the same counter value. Intuitively the operator \downarrow_r binds the current counter value to register r , while the operator \uparrow_r tests whether the current counter value is equal to the content of register r .

This paper concerns the model checking problem for Freeze LTL on one-counter automata. It is known that this problem is undecidable in general and PSPACE-complete if one restricts to *deterministic* one-counter automata [5]. Rather than restricting the class of one-counter automata, one can seek to identify decidable syntactic fragments of Freeze LTL. This approach was pursued in [6], which studied the *flat* fragment of Freeze LTL. The flatness condition places restrictions on the occurrence of the binding construct \downarrow_r in relation to the until operator (see Section 2.2 for details). For example, in a flat formula in negation normal form the binding operator \downarrow_r can occur within the scope of F but not G . (Thus formula (1.1) is flat.) The flatness restriction for Freeze LTL has a similar flavour to the respective flatness restrictions for constraint LTL [2] and for Metric Temporal Logic [1].

Demri and Sangnier [6] considered the decidability of model checking flat Freeze LTL across a range of different counter-machine models. For one-counter automata they showed decidability of model checking for a certain fragment of flat Freeze LTL and they left open the problem of model checking flat Freeze LTL in general.

The approach taken in [6] was to reduce the model checking problem for fragments of Freeze LTL on counter automata to reachability problems in counter automata augmented with certain kinds of parameterised tests. Specifically they reduce the model checking problem for flat Freeze LTL on one-counter automata to the problem of deciding reachability of Buchi objectives on one-counter automata extended with parameterised equality and disequality tests. The latter problem considers one-counter automata whose transitions may be guarded by equality or disequality tests that compare the counter value to integer-valued parameters, and it asks whether there exist parameter values such that there is an infinite computation that visits an accepting location infinitely many times. The parameterised tests are used to handle register binding in freeze LTL. The main technical contribution of this paper is to show decidability of the latter reachability problem by reduction to the decision problem for Presburger arithmetic. We thereby show that the model checking problem for flat Freeze LTL on one-counter automata is decidable.

A related work is [9], which considers one-counter automata with parameterised updates and equality tests. It is shown in [9] that reachability in this model is inter-reducible with the satisfiability problem for quantifier-free Presburger arithmetic with divisibility, and therefore decidable. In contrast to [9], in the present paper the counter automata do not have parameterised updates but they do have parameterised disequality tests. The results in this paper do not appear to be straightforwardly reducible to those of [9] nor *vice versa*. Both reachability problems can be seen as special cases of a long-standing open problem identified by Ibarra *et al.* [10], which asks to decide reachability on a class of automata with a single integer-valued counter, sign tests, and parameterised updates.

2. PRELIMINARIES

2.1. One-Counter Automata with Equality and Disequality Tests. We consider automata with a single counter that ranges over the nonnegative integers, equipped with both equality and disequality tests on counter values. Formally, a *one-counter automaton* (1-CA) is a tuple $\mathcal{C} = (V, E, \lambda, \tau)$, where V is a finite set of *states*, $E \subseteq V \times V$ is a finite set of *edges* between states, $\lambda : E \rightarrow Op$ labels each edge with an element from $Op = \{\text{add}(a) : a \in \mathbb{Z}\} \cup \{\text{eq}(a) : a \in \mathbb{N}\}$, and $\tau : V \rightarrow 2^{\mathbb{N}}$ maps each state v to a finite set $\tau(v)$ of *invalid counter values* at state v . Intuitively the operation $\text{add}(a)$ adds a to the counter and $\text{eq}(a)$ tests the counter for equality with a . The association of invalid counter values with each state can be seen as a type of disequality test. This last feature is not present in classical presentations of 1-CA, but we include it here to facilitate our treatment of Freeze LTL.

For any edge $e = (v, v')$, define $\text{start}(e) = v$ and $\text{end}(e) = v'$; moreover write $\text{weight}(e) = a$ if $\lambda(e) = \text{add}(a)$ and $\text{weight}(e) = 0$ if $\lambda(e) = \text{eq}(a)$. A *path* γ is a finite word on the alphabet E : $\gamma = e_1 \dots e_n$ such that $\text{end}(e_i) = \text{start}(e_{i+1})$ for all $1 \leq i < n$. The *length* of γ , denoted $|\gamma|$, is n . The *state sequence* of γ is $\text{start}(e_1), \text{end}(e_1), \text{end}(e_2), \dots, \text{end}(e_n)$. The *start* of γ , denoted $\text{start}(\gamma)$, is $\text{start}(e_1)$. The *end* of γ , denoted $\text{end}(\gamma)$, is $\text{end}(e_n)$. A path is *simple* if it contains no repeated states. The *weight* of γ , denoted by $\text{weight}(\gamma)$, is $\sum_{i=1}^n \text{weight}(e_i)$. A *subpath* γ' of γ is any factor of γ : $\gamma' = e_i e_{i+1} \dots e_j$. If γ and γ' are two paths such that $\text{end}(\gamma) = \text{start}(\gamma')$, $\gamma\gamma'$ is the concatenation of both paths.

A *cycle* ω is a path such that $\text{start}(\omega) = \text{end}(\omega)$. A cycle is *simple* if it has no repeated states except for the starting point, which appears twice. A cycle is *positive* if it has positive weight, *negative* if it has negative weight and *zero-weight* if it has weight zero. We denote by $\omega^k = \underbrace{\omega\omega \dots \omega}_{k \text{ times}}$ the sequence of k iterations of the cycle ω .

A *configuration* of a 1-CA $\mathcal{C} = (V, E, \lambda, \tau)$ is a pair (v, c) with $v \in V$ and $c \in \mathbb{Z}$. Intuitively, (v, c) corresponds to the situation where the 1-CA is in state v with counter value c . Configurations (v, c) with $c \geq 0$ and $c \notin \tau(v)$ are called *valid*, otherwise they are said to be *invalid*. The edge relation E induces an unlabelled transition relation between configurations: for any two configurations (v, c) and (v', c') , there is a transition $(v, c) \longrightarrow (v', c')$ if and only if there is an edge $e \in E$ such that $\text{start}(e) = v$, $\text{end}(e) = v'$, and $\text{weight}(e) = c' - c$. We will sometimes write $(v, c) \xrightarrow{e} (v', c')$ for such a transition. The transition is *valid* if both (v, c) and (v', c') are valid configurations and $c = a$ if $\lambda(e) = \text{eq}(a)$. Otherwise such a transition is *invalid*.

A *computation* π is a (finite or infinite) sequence of transitions:

$$\pi = (v_1, c_1) \longrightarrow (v_2, c_2) \longrightarrow (v_3, c_3) \longrightarrow \dots$$

We write $|\pi|$ for the length of π . If $(v_1, c_1) \xrightarrow{e_1} (v_2, c_2) \xrightarrow{e_2} \dots \xrightarrow{e_{n-1}} (v_n, c_n)$ is a finite computation, we will also write it as $(v_1, c_1) \xrightarrow{\gamma}^* (v_n, c_n)$, where $\gamma = e_1 e_2 \dots e_{n-1}$, or simply $(v_1, c_1) \longrightarrow^* (v_n, c_n)$. A computation π is *valid* if all transitions in the sequence are valid, otherwise it is *invalid*. If π is invalid, an *obstruction* is a configuration (v_i, c_i) such that either (v_i, c_i) is invalid or (v_i, c_i) is not the final configuration in π and $(v_i, c_i) \longrightarrow (v_{i+1}, c_{i+1})$ is an invalid transition.

Given a path γ and a counter value $c \in \mathbb{Z}$, the *path computation* $\gamma(c)$ is the (finite) computation starting at $(\text{start}(\gamma), c)$ and following the sequence of transitions that correspond to the edges in γ .

A *one-counter automaton with parameterised tests* is a tuple (V, E, X, λ, τ) , where V , E and λ are defined as before for 1-CA, X is a set of nonnegative integer parameters, $Op = \{\text{add}(a) : a \in \mathbb{Z}\} \cup \{\text{eq}(a), \text{eq}(x) : a \in \mathbb{N}, x \in X\}$ includes parameterised equality tests (but not parameterised updates), and $\tau : V \rightarrow 2^{\mathbb{N} \cup X}$ includes parameterised disequality tests. Note that $\tau(v)$ is still required to be finite for each $v \in V$.

For a given 1-CA $\mathcal{C} = (V, E, \lambda, \tau)$, an initial configuration (v, c) and a target configuration (v', c') , the *reachability* problem asks if there is a valid computation from (v, c) to (v', c') . When \mathcal{C} has sets $F_1, \dots, F_n \subseteq V$ of final states and an initial configuration (v, c) , the *generalised repeated control-state reachability* problem asks if there is a valid infinite computation from (v, c) which visits at least one state in each F_i infinitely often.

For a 1-CA $\mathcal{C} = (V, E, X, \lambda, \tau)$ with parameterised tests, initial configuration (v, c) , and target configuration (v', c') , the *reachability* problem asks if there exist values for the parameters such that there is a computation from (v, c) to (v', c') . Similarly, in the case where \mathcal{C} has sets $F_1, \dots, F_n \subseteq V$ of final states and an initial configuration (v, c) , the *generalised repeated control-state reachability* problem asks if there exist values for the parameters such that substituting these values satisfies the generalised repeated control-state reachability condition above.

Note that in our model of 1-CA, equality tests are defined on transitions (via the function λ) while disequality tests are defined on states (via the function τ). While this asymmetry may seem unnatural, it is technically convenient for the subsequent proofs. By contrast in the model of 1-CA in [6] both equality and disequality tests are defined on transitions. This model also allows multiple edges between states, which is excluded in our formalism. Nevertheless, from the point of view of reachability and repeated reachability the model of 1-CA that we use and that of [6] are easily seen to be equivalent (so that an algorithm for one type of 1-CA will work for the other type with only a polynomial overhead). For example, compare the 1-CA in Figure 1 (in figures, we write $+a$ for $\text{add}(a)$ and $= a?$ for $\text{eq}(a)$) and the 1-CA in Figure 2, which has disequality tests defined on transitions rather than states, as well as multiple edges between u_5 and u_6 . Then the computation

$$(v_1, 0) \longrightarrow (v_2, 10) \longrightarrow (v_2, 8) \longrightarrow (v_3, 5) \longrightarrow (v_5, 1) \longrightarrow (v_6, 1)$$

for the 1-CA in Figure 1 corresponds to the computation

$$\begin{aligned} (u_1, 0) &\longrightarrow (u_2, 10) \longrightarrow (u_3, 10) \longrightarrow (u_4, 10) \longrightarrow (u_2, 8) \\ &\longrightarrow (u_3, 8) \longrightarrow (u_4, 8) \longrightarrow (u_5, 5) \longrightarrow (u_6, 1) \longrightarrow (u_7, 1) \end{aligned}$$

for the 1-CA in Figure 2.

2.2. Model Checking Freeze LTL on One-Counter Automata. *Freeze LTL* [5] is an extension of Linear Temporal Logic that can be used to specify properties of data words. A data word is a (finite or infinite) sequence of symbols, each of which consists of a letter from a finite alphabet and another letter, often referred to as a *datum*, from an infinite alphabet. Freeze LTL is one of a variety of formalisms that arise by augmenting a temporal or modal logic with variable binding. Given a finite alphabet Σ and set of *registers* R , the formulas of Freeze LTL are given by the following grammar

$$\varphi ::= a \mid \uparrow_r \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U} \varphi \mid \downarrow_r \varphi,$$

where $a \in \Sigma$ and $r \in R$. In addition to the standard LTL connectives, Freeze LTL contains an atomic freeze formula \uparrow_r and a freeze operator \downarrow_r . We write LTL^\downarrow for the set of formulas




Figure 1: A simple 1-CA including a disequality test on the state v_2 and equality tests on the transitions (v_4, v_6) and (v_5, v_6) .




Figure 2: An automaton with disequality tests on transitions rather than states and with multiple edges which is equivalent to the one in Figure 1.

of Freeze LTL. A *sentence* is a formula in which each occurrence of a subformula \uparrow_r is in the scope of an operator \downarrow_r (for the same register r).

In general, formulas of LTL^\downarrow are interpreted over data words which have Σ as their finite alphabet and an arbitrary infinite alphabet. In this paper we are interested in a particular kind of data word—namely those arising from valid computations of 1-CA. We directly define the semantics of LTL^\downarrow over such computations, assuming that the alphabet Σ is the set of states of the 1-CA and that the infinite alphabet for data words is \mathbb{N} . In this context \downarrow_r can be seen as a binding construct that stores in register r the counter value at the current position in a computation, while \uparrow_r tests whether the counter value at the current position is equal to the content of register r . Formally, define a *register valuation* to be a partial function $f : R \rightarrow \mathbb{N}$ and consider a valid infinite computation

$$\pi = (v_1, c_1) \longrightarrow (v_2, c_2) \longrightarrow (v_3, c_3) \longrightarrow \dots$$

of a 1-CA \mathcal{C} . We define a satisfaction relation $\pi, i \models_f \varphi$ specifying when an LTL^\downarrow formula φ is satisfied at position i in π under valuation f :

$$\begin{aligned}
\pi, i \models_f a &\stackrel{\text{def}}{\iff} v_i = a \\
\pi, i \models_f \uparrow_r &\stackrel{\text{def}}{\iff} c_i = f(r) \\
\pi, i \models_f \neg\varphi &\stackrel{\text{def}}{\iff} \pi, i \not\models_f \varphi \\
\pi, i \models_f \varphi_1 \vee \varphi_2 &\stackrel{\text{def}}{\iff} \pi, i \models_f \varphi_1 \text{ or } \pi, i \models_f \varphi_2 \\
\pi, i \models_f \times\varphi &\stackrel{\text{def}}{\iff} \pi, i + 1 \models_f \varphi \\
\pi, i \models_f \varphi_1 \text{ U } \varphi_2 &\stackrel{\text{def}}{\iff} \pi, j \models_f \varphi_2 \text{ for some } j \geq i \text{ and } \pi, k \models_f \varphi_1 \text{ for all } i \leq k < j \\
\pi, i \models_f \downarrow_r \varphi &\stackrel{\text{def}}{\iff} \pi, i \models_{f[r \mapsto c_i]} \varphi
\end{aligned}$$

where $f[r \mapsto c]$ is the function that maps r to c and is otherwise equal to f . Note that the clauses for the Boolean and LTL connectives are defined in the same way as for standard LTL.

An occurrence of a subformula in an LTL^\downarrow formula is said to be *positive* if it lies within the scope of an even number of negations, otherwise it is *negative*. The *flat fragment* of LTL^\downarrow is the set of LTL^\downarrow formulas such that in every positive occurrence of a subformula $\varphi_1 \text{ U } \varphi_2$ the binding operator \downarrow_r does not appear in φ_1 , and in every negative occurrence of such a subformula, \downarrow_r does not appear in φ_2 for any register r .

The negation of many natural LTL^\downarrow specifications can be expressed by flat formulas. For example, consider the response property $\mathbf{G}(\downarrow_r(\text{req} \rightarrow \mathbf{F}(\text{serve} \wedge \uparrow_r)))$, expressing that every request is followed by a serve with the same associated ticket. (Here \mathbf{F} and \mathbf{G} are the “future” and “globally” modalities, defined by $\mathbf{F}\varphi := \text{true U } \varphi$ and $\mathbf{G}\varphi := \neg \mathbf{F} \neg \varphi$.) The negation of this formula is equivalent to $\mathbf{F}(\downarrow_r(\text{req} \wedge \mathbf{G}(\neg \text{serve} \vee \neg \uparrow_r)))$. The latter is easily seen to be flat after rewriting to the core LTL^\downarrow language with only the U temporal operator.

The main subject of this paper is the decidability of the following model checking problem: given a 1-CA \mathcal{C} , a valid configuration (v, c) of \mathcal{C} , and a flat sentence $\varphi \in \text{LTL}^\downarrow$, does there exist a valid infinite computation π of \mathcal{C} , starting at (v, c) , such that $\pi, 1 \models_\emptyset \varphi$? Note that, following [6], we have given an existential formulation of the model checking problem. The model checking problem, as formulated above, is equivalent to asking whether $\neg\varphi$ holds along all valid infinite computations starting at (v, c) .

The model checking problem for flat LTL^\downarrow on 1-CA was reduced to a repeated reachability problem for 1-CA with parameterised tests in [6, Theorem 15]. The idea of the reduction is, given a 1-CA \mathcal{C} and a flat LTL^\downarrow sentence φ in negation normal form, to construct a 1-CA with parameterised tests which is the product of \mathcal{C} and φ . This product automaton includes a parameter x_r for each register r that is mentioned in a subformula of φ of type $\downarrow_r \varphi'$. This is where the restriction to the flat fragment of LTL^\downarrow is crucial, since it allows us to assume that the value stored in a register is never overwritten along any computation of \mathcal{C} , so that it can be represented by precisely one parameter. An occurrence of the binding operator \downarrow_r in φ is represented in the product automaton by an equality test $\text{eq}(x_r)$. A positive occurrence of a formula of the type \uparrow_r is likewise represented by an equality test $\text{eq}(x_r)$, while a negative occurrence of such a subformula is represented by a disequality test $\tau(v_r) = \{x_r\}$.

Rather than recapitulating the constructions and reasoning underlying [6, Theorem 15], we give below an extended example that demonstrates the main ideas behind that result and helps motivate the subsequent development in this paper.






 Figure 3: *
 \mathcal{C}_1

 Figure 4: *
 \mathcal{C}_2

 Figure 5: *
 \mathcal{C}_3

 Figure 6: *
 \mathcal{C}_4

Figure 7: The automata considered in Example 2.1

Example 2.1. Consider the flat LTL[↓] formula

$$\varphi \equiv \text{true} \text{ U } (\downarrow_r \text{ X } (\uparrow_r \wedge \text{ X } \uparrow_r))$$

and the four counter automata represented in Figure 7. We will assume that for every 1-CA in this example the initial counter value is c . We describe in detail a 1-CA with parameterised tests, denoted $\mathcal{C}_1^{(\varphi)}$, that arises as the product of \mathcal{C}_1 and φ . We moreover explain how this definition changes if we replace \mathcal{C}_1 with one of the remaining three 1-CA.

To construct $\mathcal{C}_1^{(\varphi)}$, the first step is to introduce a concrete representation of the syntax tree of φ , which we denote by T_φ (see Figure 8). This representation is convenient for distinguishing different occurrences of the same subformula within φ . Each node of T_φ is labelled with its *address* (a word from $\{0, 1\}^*$) together with an operator or atomic formula, corresponding to an occurrence of a subformula in φ . The address of the root node is ϵ . If a node has address w then its leftmost child (if it has a child) has address $w0$ and its second child (if it has two children) has address $w1$. The operator label of every node is assigned in the obvious way, taking the outermost connective of the corresponding subformula.


Next we introduce the notion of *atoms* of φ which are certain sets of (occurrences of) subformulas of φ . Formally an atom A is any subset of $\{0, 1\}^*$ that satisfies the following conditions:¹

- (1) If $w \in A$ and the node in T_φ with address w has label \wedge , then $w0, w1 \in A$.
- (2) If $w \in A$ and the node in T_φ with address w has label U , then $w0 \in A$ or $w1 \in A$.
- (3) If $w \in A$ and the node in T_φ with address w has label \downarrow_r , then $w0 \in A$.

The above conditions correspond to the intuition that an atom represents a set of subformulas of φ that hold at a certain position in a data word. We moreover define a transition relation between atoms by specifying that for atoms A, A' we have $A \longrightarrow A'$ if the following conditions hold:

- (1) If $w \in A$ and the node in T_φ with address w has label X then $w0 \in A'$.
- (2) If $w \in A$ and the node in T_φ with address w has label U then either $w1 \in A'$ or $w0 \in A'$ and $w \in A'$.

¹ In general there is also a condition for negation, but this is not relevant for the current simple example.

Figure 8: The formula tree T_φ


- (3) No atom A'' strictly included in A' satisfies the preceding two conditions (with A'' in place of A').

For the formula φ at hand, $\{\epsilon, 0\}$ is an atom and the set of atoms reachable from this one under the transition relation is $\{\{\epsilon, 0\}, \{\epsilon, 1, 10\}, \{100, 1000, 1001\}, \{10010\}, \emptyset\}$, with the transition relation being given by:

$$\begin{aligned} \{\epsilon, 0\} &\longrightarrow \{\epsilon, 0\} \\ \{\epsilon, 0\} &\longrightarrow \{\epsilon, 1, 10\} \\ \{\epsilon, 1, 10\} &\longrightarrow \{100, 1000, 1001\} \\ \{100, 1000, 1001\} &\longrightarrow \{10010\} \\ \{10010\} &\longrightarrow \emptyset \end{aligned}$$

Roughly speaking, the automaton $\mathcal{C}_1^{(\varphi)}$ arises as the product of \mathcal{C}_1 and the transition relation on atoms of φ . This automaton is shown in Figure 9. It has an initial state v_0 and auxiliary states $v_1^{aux}, \dots, v_6^{aux}$. All other states are of the form $\langle v, A \rangle$ where v is a state of \mathcal{C}_1 and A is an atom of φ . From the initial state v_0 there is a nondeterministic choice between the two atoms having ϵ as an element. The choice of $\{\epsilon, 1, 10\}$ means that $\downarrow_r X(\uparrow_r \wedge X \uparrow_r)$ (the right side of φ) holds in the initial state v_1 of \mathcal{C}_1 , and the choice of $\{\epsilon, 0\}$ means that this subformula will hold at some point in the future, i.e., in v_2 or v_3 . For any two consecutive edges

$$\langle v_i, A \rangle \longrightarrow v_k^{aux} \longrightarrow \langle v_j, A' \rangle$$


 Figure 9: The product automaton $\mathcal{C}_1^{(\varphi)}$ of \mathcal{C}_1 and φ .

in the product automaton $\mathcal{C}_1^{(\varphi)}$, the label on the second edge is equal to the label on the edge (v_i, v_j) in automaton \mathcal{C}_1 , and the label on the first edge is a test on the counter value. There is one set of final states, $F = \{\langle v_3, \emptyset \rangle\}$.²

Let us first follow the path from $\langle v_1, \{\epsilon, 1, 10\} \rangle$. The transition to the auxiliary state v_1^{aux} tests whether the current counter value, that is c , is equal to the parameter x . This equality test corresponds to node 1 in T_φ . Next, $\mathcal{C}_1^{(\varphi)}$ simulates a transition from v_1 to v_2 in \mathcal{C}_1 . The transitions to auxiliary states v_2^{aux} and v_6^{aux} include further tests for equality with x , which correspond to nodes 1000 and 10010 in T_φ , respectively. Clearly if we set $x = c$, there is a valid computation starting from (v_0, c) over this path which ends up visiting the final state $\langle v_3, \emptyset \rangle$ infinitely many times.

The paths starting at $\langle v_1, \{\epsilon, 0\} \rangle$ all correspond to cases where $\downarrow_r \mathbf{X}(\uparrow_r \wedge \mathbf{X} \uparrow_r)$ does not hold in v_1 , but only at a later time in v_2 or v_3 . It is easy to see that there is no value for x that allows reachability of the final state along any of these paths.

Recall the 1-CA \mathcal{C}_2 , \mathcal{C}_3 and \mathcal{C}_4 from Figure 7. In the 1-CA $\mathcal{C}_2^{(\varphi)}$, which is constructed in the same way as $\mathcal{C}_1^{(\varphi)}$ but with transition labels from \mathcal{C}_2 rather than \mathcal{C}_1 , there is a valid computation starting in (v_0, c) over the path that goes through v_3^{aux} and v_5^{aux} and ends up visiting $\langle v_3, \emptyset \rangle$ infinitely many times. This means that the right side of φ becomes true in \mathcal{C}_2 after one step. Similarly, in $\mathcal{C}_3^{(\varphi)}$, there is a valid computation over any path that goes through v_4^{aux} and v_5^{aux} to $\langle v_3, \emptyset \rangle$, which means that the right side of φ becomes true in \mathcal{C}_3 after two or more steps. Finally, in $\mathcal{C}_4^{(\varphi)}$, the state $\langle v_3, \emptyset \rangle$ can never be reached, since there is no computation from (v_1, c) in \mathcal{C}_4 that satisfies φ .

If φ featured any negated atomic formulas of the form $\downarrow_{r'} \varphi'$, auxiliary states with disequality tests would be needed in the product automaton. For a general description of how to construct the product automaton and a proof, see [6].

Note that the definition of 1-CA with parameterised tests in [6] includes parameterised equality and disequality tests (as in the present paper) together with parameterised inequality tests, i.e., testing whether the counter value is less than or greater than the value of a

²In general, there is a set of final states for each U operator in the formula.

parameter. However, it is clear from the details of the reduction that only equality and disequality tests are needed, and thus we do not consider inequality tests in this paper.

2.3. Presburger Arithmetic. *Presburger arithmetic* is the first-order logic over the structure $\langle \mathbb{Z}, +, <, 0, 1 \rangle$, where $+$ and $<$ are the standard addition and ordering on integers. Presburger arithmetic is known to be decidable [13]. Using shorthand notation, we can assume that the atomic formulas of Presburger arithmetic are equalities or inequalities between linear polynomials with integer coefficients.

3. NORMAL FORM FOR PATHS

In this section, we show that any valid finite computation of a 1-CA $\mathcal{C} = (V, E, \lambda, \tau)$ can be rewritten to a normal form whose shape only depends on the automaton and such that the initial and final configuration of the computation are preserved. Informally, any such computation can be described as a sequence of “take this transition” and “take this cycle k times”. We show that the maximum length of a description of this kind is independent of the original computation. Such a description is similar in spirit to the semilinear path schemes described in [11].

We give a brief overview of the technical development below. The first step (Lemma 3.1) is to bound the number of equality tests along a minimum-length computation between two configurations of a 1-CA. Thereafter we focus on computations that are free of equality tests. To obtain a succinct representation of such computations we define a rewriting system that reorders computations by gathering together in the same place executions of the same simple cycle; we moreover introduce a compressed representation of iterated simple cycles, leading to the notion of *folded paths*. Lemmas 3.2 and 3.3 show that the rewriting rules are sound (i.e., preserve validity of computations) and terminating. We then concentrate on bounding the length of folded paths which cannot be further rewritten. To this end we identify a set of “critical” configurations that block application of the rewriting rules, and we bound the number of such configurations (Lemma 3.5). This leads to an upper bound on the length of a folded path that cannot be rewritten (Lemmas 3.6 and 3.7). Finally the whole analysis, including equality tests, is summarised in Theorem 3.8 which gives the required upper bound on the length of folded paths.

In the rest of this section we consider a fixed 1-CA $\mathcal{C} = (V, E, \lambda, \tau)$. First we show that, without loss of generality, any computation in \mathcal{C} can be broken down into a small number of segments that do not contain any transitions with equality tests. The idea is that any segment between two identical equality tests can be omitted.

Lemma 3.1 (Equality-test isolation). *Let π be a valid finite computation from (v, c) to (v', c') . Then there exists a path γ such that $\gamma(c)$ is a valid computation from (v, c) to (v', c') and γ is of the form $\gamma = \gamma_0 e_1 \gamma_1 e_2 \cdots e_n \gamma_n$, where e_i is an edge with an equality test, γ_i is a path without equality tests and $n \leq |E|$.*

Proof. Let π' be the shortest valid computation from (v, c) to (v', c') . We can decompose it as

$$\pi' = (v, c) \xrightarrow{\gamma_0} (v_1, c_1) \xrightarrow{e_1} (v'_1, c_1) \xrightarrow{\gamma_1} (v_2, c_2) \cdots (v'_n, c_n) \xrightarrow{\gamma_n} (v', c')$$

where for every i , γ_i is a path without any equality tests and $e_i = (v_i, \text{eq}(c_i), v'_i) \in E$ is an equality test. Then clearly $\pi' = \gamma(c)$ where

$$\gamma = \gamma_0 e_1 \gamma_1 e_2 \cdots e_n \gamma_n.$$

Assume for a contradiction that $n > |E|$. Then by the pigeonhole principle, there exists $i < j$ such that $e_i = e_j$. But since π' is a valid computation, the two transitions $(v_i, c_i) \xrightarrow{\text{eq}(c_i)} (v'_i, c_i)$ and $(v_j, c_j) \xrightarrow{\text{eq}(c_i)} (v'_j, c_j)$ are the same and $(v_i, c_i) = (v_j, c_j)$. Thus we can delete part of the computation and define

$$\gamma' = \gamma_0 e_1 \gamma_1 \cdots e_i \gamma_j e_{j+1} \cdots e_n \gamma_n.$$

Then $\gamma'(c)$ is a valid computation from (v, c) to (v', c') and is shorter than π' , which is a contradiction. \square

We need to introduce some terminology to formalise our notion of normal form. Write SC for the set of all equality-free simple cycles in \mathcal{C} . We moreover denote by $\text{SC}^+ \subseteq \text{SC}$ the set of equality-free simple cycles that have positive weight and likewise by $\text{SC}^- \subseteq \text{SC}$ the set of cycles with negative weight.

The *cycle alphabet*, denoted C , consists of symbols of the form $\underline{\omega}^k$ where $\omega \in \text{SC}$ and $k \in \mathbb{N}$. Note that this alphabet is infinite. Also note that $\underline{\omega}^k$ is a single symbol, underlined to distinguish it from the cycle ω^k , which consists of $|\omega|k$ symbols from E . For convenience, we use $\underline{\omega}$ as shorthand for $\underline{\omega}^1$. We naturally define the start and end of symbol $\underline{\omega}^k$ by the start of ω : $\text{start}(\underline{\omega}^k) = \text{end}(\underline{\omega}^k) = \text{start}(\omega)$.

A *folded path* χ is a word over the alphabet $E \cup C$: $\chi = s_1 \cdots s_n$ such that $\text{end}(s_i) = \text{start}(s_{i+1})$ for every $i < n$. We also define the natural unfolding folded paths via a monoid homomorphism $\text{unfold} : (E \cup C)^* \rightarrow E^*$ such that $\text{unfold}(e) = e$ for $e \in E$ and $\text{unfold}(\underline{\omega}^k) = \omega^k$ for $\underline{\omega}^k \in C$. The weight of a folded path is the weight of its unfolding.

For the rest of this section we fix an initial counter value $c \in \mathbb{N}$ and we only consider computations starting at c that do not feature equality tests. We refer to a folded path χ as being *valid* if $\text{unfold}(\chi)(c)$ is a valid computation.

Define the following nondeterministic rewriting system on folded paths. Each rule of the system has a name, a pattern to match against, a condition that must be satisfied for the rule to apply, and the result of the rule. We denote by $\chi \rightsquigarrow \chi'$ the fact that χ rewrites to χ' .

Rule	Pattern	Result	Condition
fold	$\psi \omega \phi$	$\psi \underline{\omega} \phi$	ω is a simple cycle of nonzero weight.
simplify	$\psi \rho \phi$	$\psi \phi$	Nonempty ρ , $\text{weight}(\text{unfold}(\rho)) = 0$ and $\text{end}(\psi) = \text{start}(\phi)$.
gather ⁺	$\psi \underline{\omega}^k \rho \underline{\omega}^\ell \phi$	$\psi \underline{\omega}^{k+1} \rho \underline{\omega}^{\ell-1} \phi$	Result is valid, ω is a positive simple cycle and $\ell > 0$.
gather ⁻	$\psi \underline{\omega}^k \rho \underline{\omega}^\ell \phi$	$\psi \underline{\omega}^{k-1} \rho \underline{\omega}^{\ell+1} \phi$	Result is valid, ω is a negative simple cycle and $k > 0$.

Lemma 3.2 (Soundness). *If χ is a valid folded path that rewrites to χ' , then χ' is also valid. Furthermore, χ and χ' start and end at the same state and $\text{weight}(\text{unfold}(\chi)) = \text{weight}(\text{unfold}(\chi'))$.*

Proof. This is easily checked for each rule:

- **fold:** Clearly $\text{unfold}(\chi) = \text{unfold}(\chi')$.
- **simplify:** First note that the result is well-formed because of the condition on start and end. The unfolding of the first part (ψ) of the path is unchanged, so it remains valid and with the same starting state. Since the second part of the path (ρ) has weight 0, the counter value is the same at the beginning and end of ρ , so the unfolding of the third part

(ϕ) stays the same, and thus valid with the same end state. The weight of the unfolded path remains unchanged as the removed part ρ has weight 0.

- **gather** $^\pm$: The condition ensures the result is valid. The start and end state clearly do not change, and neither does the weight, since $\text{unfold}(\chi')$ contains the same edges as $\text{unfold}(\chi)$, only in a different order. \square

Lemma 3.3 (Termination). *There are no infinite chains of rewriting.*

Proof. First we give an informal explanation. The first thing to notice is that the length of a folded path (over alphabet $E \cup C$) never increases after a rewriting operation. The second thing is that the length of a folded path over E (i.e., ignoring symbols from C) never increases either. Since rule **simplify** strictly decreases the length, it can only be applied finitely many times. Similarly, rule **fold** strictly decreases the length over E because it replaces a symbol from E by one from C . Rules **gather** $^\pm$ are more difficult to analyse because they only reorder the path by replacing symbols from C . But notice that a symbol $\underline{\omega}$, where ω is a positive cycle, can only move left, and similarly a negative cycle can only move right. Intuitively, this process must be finite because once a positive (negative) cycle reaches the leftmost (rightmost) position, it cannot move anymore.

Formally, we will define a valuation over folded paths and show that it decreases after each application of a rule. First, for any folded path χ and any given simple cycle ω , define the ω -projection $p_\omega(\chi)$ of χ to be the subword consisting only of symbols of the form $\underline{\omega}^k$:

$$p_\omega(e\chi) = p_\omega(\chi) \text{ if } e \in E \quad p_\omega(\underline{\omega}^k\chi) = \underline{\omega}^k p_\omega(\chi) \quad p_\omega(\underline{\theta}^k\chi) = p_\omega(\chi) \text{ if } \theta \neq \omega.$$

For any folded path χ , define:

$$(|\chi|) = (|\chi|, |\chi|_E, \sigma(\chi)), \quad \text{where} \quad \sigma(\chi) = \sum_{\omega \in \text{SC}} \sigma_\omega(p_\omega(\chi)),$$

$|\chi|$ is the word length of χ (over alphabet $E \cup C$), $|\chi|_E$ is the word length of χ only counting symbols in E , and $\sigma_\omega(p_\omega(\chi))$ is defined as follows:

$$\sigma_\omega \left(\underline{\omega}^{k_1} \underline{\omega}^{k_2} \dots \underline{\omega}^{k_n} \right) = \begin{cases} \sum_{i=1}^n i k_i & \text{if } \text{weight}(\omega) > 0 \\ 0 & \text{if } \text{weight}(\omega) = 0 \\ \sum_{i=1}^n (n+1-i) k_i & \text{if } \text{weight}(\omega) < 0. \end{cases}$$

We will now show that $(|\chi|)$ decreases in lexicographic order each time a rule is applied. In the case of rule **fold**, if $|\omega| \geq 2$ then clearly $|\chi|$ decreases because we replace several symbols with just one. If $|\omega| = 1$ then $|\chi|$ stays constant but $|\chi|_E$ decreases by one because we replace one symbol from E by one symbol from C . Similarly, rule **simplify** decreases $|\chi|$ because we remove a nonzero-length subpath. Since rules **gather** $^+$ and **gather** $^-$ are symmetric, we only consider **gather** $^+$. Note that the rule does not change $|\chi|$ or $|\chi|_E$ because it only replaces symbols from C with different symbols from C , so we are only concerned with $\sigma(\chi)$.

Assume the rule rewrites $\psi \underline{\omega}^k \rho \underline{\omega}^\ell \phi$ into $\psi \underline{\omega}^{k+1} \rho \underline{\omega}^{\ell-1} \phi$. First note that if $\theta \neq \omega$ is a simple cycle, then the θ -projection is the same before and after the rule because the rule does not replace any symbols of the form $\underline{\theta}^k$, so σ_θ does not change. The case of σ_ω is slightly more involved and we need to introduce some notations:

$$p_\omega(\psi) = \underline{\omega}^{u_1} \dots \underline{\omega}^{u_n}, \quad p_\omega(\rho) = \underline{\omega}^{u_{n+2}} \dots \underline{\omega}^{u_m}, \quad p_\omega(\phi) = \underline{\omega}^{u_{m+2}} \dots \underline{\omega}^{u_q}$$

and

$$u_{n+1} = k, \quad u_{m+1} = \ell, \quad u'_{n+1} = k + 1, \quad u'_{m+1} = \ell - 1, \quad u'_i = u_i \text{ if } i \neq n + 1, m + 1.$$

Then we can observe that:

$$\sigma_\omega \left(p_\omega \left(\psi \underline{\omega^k} \rho \underline{\omega^\ell} \phi \right) \right) = \sigma_\omega \left(\underline{\omega^{u_1}} \cdots \underline{\omega^{u_q}} \right) = \sum_{i=1}^q i u_i, \quad (3.1)$$

$$\sigma_\omega \left(p_\omega \left(\psi \underline{\omega^{k+1}} \rho \underline{\omega^{\ell-1}} \phi \right) \right) = \sigma_\omega \left(\underline{\omega^{u'_1}} \cdots \underline{\omega^{u'_q}} \right) = \sum_{i=1}^q i u'_i. \quad (3.2)$$

Thus:

$$\begin{aligned} (3.1) - (3.2) &= \sum_{i=1}^q i(u_i - u'_i) \\ &= (n+1)(u_{n+1} - u'_{n+1}) + (m+1)(u_{m+1} - u'_{m+1}) \\ &= -(n+1) + (m+1) \\ &> 0 \text{ because } m > n. \end{aligned}$$

Thus $\sigma_\omega(\chi)$ decreases after the rule is applied and thus $\sigma(\chi)$ also decreases. \square

Lemma 3.4 (Size of cycle-free subpaths). *If $\psi\rho\phi$ is a folded path such that $\rho \in E^*$ and no rewriting rule applies, then $|\rho| < |V|$.*

Proof. Assume the contrary: if ρ only consists of edges and has length $\geq |V|$, then some state is repeated in the state sequence of ρ . Thus ρ contains a cycle and thus a simple cycle. So rule **fold** applies if the cycle has nonzero weight, or rule **simplify** applies if it has weight zero. \square

The next lemma analyses situations in which the pattern of one of the rules **gather**⁺ and **gather**⁻ matches a factor of a folded word, but application of the rule leads to an invalid computation. The idea is to identify a set of so-called critical configurations which can potentially prevent application of one of these two rules and then to bound the number of such critical configurations. As we observe below, both rules are sound with respect to the requirement that counter values be nonnegative and can only cause a computation to become invalid through the presence of disequality tests.

Given a state v of \mathcal{C} , we define a set $B^+(v)$ of *critical values* for positive cycles and a set $B^-(v)$ of critical values for negative cycles. These sets represent valid configurations (v, c) from which some simple cycle cannot be executed due to a disequality test. Formally, for $S \subseteq \mathbb{Z}$ and $x \in \mathbb{Z}$, write $S - x$ to denote $\{y - x \mid y \in S\}$; then we define $B^+(v)$ to be the union of the sets $\tau(\text{end}(\gamma)) - \text{weight}(\gamma)$ for γ a non-empty prefix of some positive cycle starting at v . Likewise we define $B^-(v)$ to be the union of the sets $\tau(\text{end}(\gamma)) - \text{weight}(\gamma)$ for γ a non-empty prefix of some negative cycle starting at v .

Lemma 3.5 (Obstructions in irreducible paths with cycles). *Let ω be a positive (resp. negative) cycle and assume that rule **gather**⁺ (resp. **gather**⁻) does not apply to $\psi \underline{\omega^k} \rho \underline{\omega^\ell} \phi$ (which we assume is valid and $k, \ell > 0$) for this particular pattern. Then there exists a (potentially empty) prefix μ of ρ such that $\text{unfold}(\psi \underline{\omega^k} \mu)(c)$ has the form $(v, c) \xrightarrow{*} (v', c')$ where c' is critical for v' for positive (resp. negative) cycles, i.e. $c' \in B^+(v')$ (resp. $c' \in B^-(v')$).*

$B^-(v')$). Furthermore $B^+(v')$ and $B^-(v')$ only depend on the automaton and

$$|B^+(v')| \leq |\text{SC}^+| \sum_{u \in V} |\tau(u)| \quad \text{and} \quad |B^-(v')| \leq |\text{SC}^-| \sum_{u \in V} |\tau(u)|.$$

Proof. We first show the result for positive cycles. Let $\pi = \text{unfold}(\psi \omega^k \rho \omega^\ell \phi)(c)$ and $\pi' = \text{unfold}(\psi \omega^{k+1} \rho \omega^{\ell-1} \phi)(c)$. To make things slightly easier to understand, note that:

$$\begin{aligned} \pi &= [\text{unfold}(\psi) \omega^k \text{unfold}(\rho) \omega \omega^{\ell-1} \text{unfold}(\phi)](c) \\ \pi' &= [\text{unfold}(\psi) \omega^k \omega \text{unfold}(\rho) \omega^{\ell-1} \text{unfold}(\phi)](c). \end{aligned}$$

Since $\text{unfold}(\rho) \omega$ and $\omega \text{unfold}(\rho)$ have the same weight, it is clear that the first $(\text{unfold}(\psi) \omega^k)$ and last $(\omega^{\ell-1} \text{unfold}(\phi))$ parts of the computation are the same in π and π' , i.e., they have the same counter values. Consequently, if they are valid in π , the same parts are also valid in π' . Since by the hypothesis **gather**⁺ does not apply, π' is invalid. So there must be an obstruction (u, d) in the middle part $(\omega \text{unfold}(\rho))$ of π' . There are two possibilities.

The first case is when the obstruction (u, d) is in the $\text{unfold}(\rho)$ part of π' . Note that $d = c^* + \text{weight}(\omega)$, where (u, c^*) is the corresponding configuration in the $\text{unfold}(\rho)$ part of π . Since ω is a positive cycle, $d > c^*$ cannot be negative (since (u, c^*) occurs in π , which is valid). Since we assumed that all computations are free of equality tests, the obstruction must be because of a disequality, i.e., it must be that $d = c^* + \text{weight}(\omega) \in \tau(u)$. Thus $c^* \in \tau(u) - \text{weight}(\omega)$ and c^* is critical for u . Then there exists a prefix μ of ρ such that $\text{unfold}(\psi \omega^k \mu)(c) = (v, c) \xrightarrow{*} (u, c^*)$ and this shows the result.

The second case is when (u, d) is in the ω part of the middle part $(\omega \text{unfold}(\rho))$ of π' . Again, it is impossible that the counter value d be negative. Indeed, remember that ω is a positive cycle and $k > 0$, thus

$$\begin{aligned} \pi' &= [\text{unfold}(\psi) \omega^{k+1} \text{unfold}(\rho) \omega^{\ell-1} \text{unfold}(\phi)](c) \\ &= [\text{unfold}(\psi) \omega^{k-1} \omega \text{unfold}(\rho) \omega^{\ell-1} \text{unfold}(\phi)](c) \\ &= (v, c) \xrightarrow{\text{unfold}(\psi) \omega^{k-1}}^* (v_1, c_1) \xrightarrow{\omega}^* (v_1, c_2) \xrightarrow{\omega}^* (v_1, c_3) \xrightarrow{\text{unfold}(\rho) \omega^{\ell-1} \text{unfold}(\phi)}^* (v'', c''). \end{aligned}$$

We already argued that $(v, c) \xrightarrow{*} (v_1, c_2)$ is valid, so in particular $(v_1, c_1) \xrightarrow{\omega}^* (v_1, c_2)$ is valid. Note that the obstruction is in the second iteration of ω : $(v_1, c_2) \xrightarrow{\omega}^* (v_1, c_3)$. Since ω is a positive cycle, $c_2 > c_1$. Note that initially the cycle ω was feasible (with the counter not going negative) starting with a lower counter value (c_1) so the counter cannot possibly become negative on the second iteration starting with a higher counter value (c_2) . Thus, again, the obstruction happens because of a disequality. That is, we can write $\omega = \gamma \gamma'$ such that:

$$\pi' = (v, c) \xrightarrow{\text{unfold}(\psi) \omega^k}^* (v_1, c_2) \xrightarrow{\gamma}^* (u, d) \xrightarrow{\gamma'}^* (v_1, c_3) \xrightarrow{\text{unfold}(\rho) \omega^{\ell-1} \text{unfold}(\phi)}^* (v'', c'')$$

and the obstruction happens because $d \in \tau(u)$. Note however that $d = c_2 + \text{weight}(\gamma)$ and thus $c_2 \in \tau(u) - \text{weight}(\gamma)$. In this case, c_2 is critical for v_1 . Choose μ to be the empty word, so that $\text{unfold}(\psi \omega^k \mu)(c) = (v, c) \xrightarrow{*} (v_1, c_2)$ to show the result.

Observe that the definition of critical values only depends on the automaton itself. Furthermore, the size of $B^+(v)$ can easily be bounded. Indeed, there are $|\text{SC}^+|$ positive simple cycles, for each such cycle its non-empty prefixes all end in different states, and a

prefix ending in a state u contributes $|\tau(u)|$ elements to $B^+(v)$. It follows that $|B^+(v)| \leq |\text{SC}^+| \sum_{u \in V} |\tau(u)|$.

The proof is exactly the same in the negative case except for one detail. This time we move negative cycles to the right so that the middle part of π' ($\text{unfold}(\rho)\omega$) can only get higher counter values than the middle part of π ($\omega \text{unfold}(\rho)$), as in the positive case. \square

The following lemma is a step towards bounding the length of a folded path to which no rewriting rule applies. We use this lemma to obtain such a bound in Lemma 3.7.

Lemma 3.6 (Length of irreducible paths). *Let χ be a folded path such that no rewriting rule applies on χ . Let $Y = \text{SC}^+$ or $Y = \text{SC}^-$. Then for every $\omega \in Y$, the number of symbols in χ of the form $\underline{\omega}$ (the exponent does not matter) is bounded by*

$$|V||Y| \left(1 + \sum_{v \in V} |\tau(v)|\right).$$

Proof. Without loss of generality, we show the result for $X = \text{SC}^+$. First note that if $\underline{\omega^k}$ appears in χ and no rule applies, then $k > 0$, otherwise we could apply **simplify** to remove $\underline{\omega^0}$. We can thus decompose the path as:

$$\chi = \phi_0 \underline{\omega^{k_1}} \phi_1 \underline{\omega^{k_2}} \phi_2 \cdots \phi_{n-1} \underline{\omega^{k_n}} \phi_n$$

where $k_i > 0$ and ϕ_i does not contain any $\underline{\omega}$ symbol. Since no rule applies, by Lemma 3.5, there exist prefixes $\mu_1, \mu_2, \dots, \mu_{n-1}$ of $\phi_1, \phi_2, \dots, \phi_{n-1}$ respectively, such that for each i :

$$(v, c) \xrightarrow{\phi_0 \underline{\omega^{k_1}} \phi_1 \cdots \phi_{i-1} \underline{\omega^{k_i}} \mu_i}^* (v_i, c_i) \quad \text{where } c_i \in B^+(v_i).$$

Assume for a contradiction that there is a repeated configuration among the (v_i, c_i) . Then there exists $i < j$ such that $v_i = v_j$ and $c_i = c_j$. Let $\phi_i = \mu_i \rho$ and $\phi_j = \mu_j \rho'$, and observe that:

$$(v, c) \xrightarrow{\phi_0 \underline{\omega^{k_1}} \phi_1 \cdots \phi_{i-1} \underline{\omega^{k_i}} \mu_i}^* (v_i, c_i) \xrightarrow{\rho \underline{\omega^{k_{i+1}}} \phi_{i+1} \cdots \phi_{j-1} \underline{\omega^{k_j}} \mu_j}^* (v_i, c_i) \xrightarrow{\rho' \underline{\omega^{k_{j+1}}} \phi_{j+1} \cdots \phi_{n-1} \underline{\omega^{k_n}} \phi_n}^* (v', c').$$

Thus the subpath $\rho \underline{\omega^{k_{i+1}}} \phi_{i+1} \cdots \phi_{j-1} \underline{\omega^{k_j}} \mu_j$ has weight 0 and rule **simplify** must apply:

$$\chi \rightsquigarrow \phi_0 \underline{\omega^{k_1}} \phi_1 \cdots \phi_{i-1} \underline{\omega^{k_i}} \mu_i \rho' \underline{\omega^{k_{j+1}}} \phi_{j+1} \cdots \phi_{n-1} \underline{\omega^{k_n}} \phi_n$$

which is a contradiction because we assumed that no rule can apply to χ .

Consequently, for any $i \neq j$, we have $(v_i, c_i) \neq (v_j, c_j)$. But remember that $c_i \in B^+(v_i)$, thus $(v_i, c_i) \in A$ where:

$$A = \bigcup_{v \in V} \{v\} \times B^+(v).$$

This shows that $n - 1 \leq |A|$. Indeed, by the pigeonhole principle, some pair (v_i, c_i) would be repeated if $n - 1 > |A|$. We can easily bound the size of A using the bound on $B^+(v)$ from Lemma 3.5:

$$|A| \leq \sum_{v \in V} |B^+(v)| \leq |V| |\text{SC}^+| \sum_{v \in V} |\tau(v)|.$$

Finally we have

$$n \leq |V| |\text{SC}^+| \sum_{v \in V} |\tau(v)| + 1 \leq |V| |\text{SC}^+| \left(1 + \sum_{v \in V} |\tau(v)|\right)$$

because $|V| \geq 1$ and $|\text{SC}^+| \geq 1$ unless there are no positive cycles, in which case $n = 0$ anyway. \square

Lemma 3.7 (Length of equality-free computations). *Let π be a valid finite computation (without equality tests) from (v, c) to (v', c') . Then there exists a folded path χ such that $\text{unfold}(\chi)(c)$ is a valid computation from (v, c) to (v', c') , the length of $\text{unfold}(\chi)(c)$ is at most that of π and the word length of χ is bounded by:*

$$|V| + |V|^2 |\text{SC}|^2 \left(1 + \sum_{v \in V} |\tau(v)|\right)$$

Proof. Let χ_0 be the path defined by π : it is a word over alphabet E and is thus a (trivial) folded path. By definition $\text{unfold}(\chi_0(c)) = \pi$ is a valid computation from (v, c) to (v', c') and the length of $\text{unfold}(\chi_0(c))$ is equal to that of π . Let χ be any rewriting of χ_0 such that no rule applies on χ : it exists because there are no infinite rewriting chains by Lemma 3.3. By Lemma 3.2, $\text{unfold}(\chi(c))$ is still a valid computation from (v, c) to (v', c') . Let ω be a simple cycle: note that it is either positive or negative, because rule **simplify** removes zero-weight cycles. Then by Lemma 3.6, the number of symbols of the form $\underline{\omega}$ appearing in χ is bounded by³:

$$|V| |\text{SC}| \left(1 + \sum_{v \in V} |\tau(v)|\right) \quad (3.3)$$

and thus the total number of symbols in χ of the form $\underline{\omega}$ for any ω is bounded by:

$$|V| |\text{SC}|^2 \left(1 + \sum_{v \in V} |\tau(v)|\right). \quad (3.4)$$

Furthermore, inbetween symbols of the form $\underline{\omega}$, there can be subpaths consisting of symbols in E only, so χ is of the form

$$\chi = \phi_0 \underline{\omega_1^{k_1}} \phi_1 \underline{\omega_2^{k_2}} \cdots \underline{\omega_n^{k_n}} \phi_n$$

where $\phi_i \in E^*$ and $\omega_i \in \text{SC}$ for all i . By the reasoning above, $n \leq (3.4)$. Furthermore, by Lemma 3.4, $\phi_i < |V|$ for all i . It follows that the total length of χ is bounded by

$$\begin{aligned} (n+1)(|V|-1) + n &\leq |V| + n|V| \\ &\leq |V| + |V|^2 |\text{SC}|^2 \left(1 + \sum_{v \in V} |\tau(v)|\right). \end{aligned}$$

Finally the length of $\text{unfold}(\chi(c))$ at most that of π because the rewriting system does not increase the length of the path and the length of $\text{unfold}(\chi_0(c))$ is equal to that of π . \square

The main result of this section shows that any valid computation has an equivalent valid computation given by a folded path whose length only depends on the automaton.

Theorem 3.8 (Length of computations). *Let π be a valid finite computation from (v, c) to (v', c') . Then there exists a folded path χ such that $\chi(c)$ is a valid computation from (v, c) to (v', c') , the length of $\text{unfold}(\chi(c))$ is at most that of π and the word length of χ is bounded by:*

$$|E| \left(1 + |V| + |V|^2 |\text{SC}|^2 \left(1 + \sum_{v \in V} |\tau(v)|\right)\right).$$

³Since obviously $\max(|\text{SC}^+|, |\text{SC}^-|) \leq |\text{SC}|$.

Proof. Apply Lemma 3.1 to isolate the equality tests (at most $|E|$ of them) and apply Lemma 3.7 to each equality-free subcomputation. We can improve the bound slightly by noticing that there can only be up to $|E|$ equality-free subcomputations (and not $|E| + 1$). Indeed, if there are $|E|$ different equality tests in the path, there are no further edges available for equality-free computations, and the word length is at most $|E|$. \square

4. REACHABILITY WITH PARAMETERISED TESTS

In this section we will show that both the reachability problem and the generalised repeated control-state reachability problem for 1-CA with parameterised tests are decidable, via a symbolic encoding of folded paths, making use of the normal form from the previous section. The result of this encoding is a formula of Presburger arithmetic.

Recall that $C = \{\omega^k : \omega \in \text{SC}, k \in \mathbb{N}\}$. Let $C' = \{\underline{\omega} : \omega \in \text{SC}\}$. We define a *path shape* to be a word over the alphabet $E \cup C'$: $\xi = t_1 \dots t_n$ such that $\text{end}(t_i) = \text{start}(t_{i+1})$, where $\text{start}(\underline{\omega}) = \text{end}(\underline{\omega}) = \text{start}(\omega)$. Given a path shape $\xi = \gamma_0 \underline{\omega}_1 \gamma_1 \dots \underline{\omega}_n \gamma_n$ with $\gamma_i \in E^*$, we write $\xi(k_1, \dots, k_n)$ for the folded path $\gamma_0 \underline{\omega}_1^{k_1} \gamma_1 \dots \underline{\omega}_n^{k_n} \gamma_n$. The advantage of working with path shapes rather than folded paths is that the former are words over a finite alphabet.

Lemma 4.1 (Encoding computations). *Given a 1-CA $\mathcal{C} = (V, E, X, \lambda, \tau)$ with parameterised tests and configurations (v, c) and (v', c') , and given a path shape $\xi = t_1 t_2 \dots t_n \in (E \cup C')^*$, there exists a Presburger arithmetic formula $\varphi_{\text{comp}}^{(\xi), (v, c), (v', c')}(\mathbf{k}, \mathbf{x})$, with free variables \mathbf{x} corresponding to the parameters X and \mathbf{k} corresponding to exponents to be substituted in ξ , which evaluates to true if and only if $\text{unfold}(\xi(\mathbf{k}))(c)$ is a valid computation from (v, c) to (v', c') .*

Proof. Assume first that ξ does not include any equality tests. We define a formula $\varphi_{\text{valid, noeq}}^{(t)}(\mathbf{k}, \mathbf{x}, y)$ which, given an equality-free symbol $t \in E \cup C'$ and an integer y , evaluates to true if and only if $\text{unfold}(t(\mathbf{k}))(y)$ is a valid computation. There are two cases:

- $t \in E$. Then $\varphi_{\text{valid, noeq}}^{(t)}(\mathbf{x}, y) \equiv y \geq 0 \wedge y + \text{weight}(t) \geq 0 \wedge y \notin \tau(\text{start}(t))$.
- $t \in C'$, i.e., $t(\mathbf{k}) = \underline{\omega}^k$ for some simple cycle $\omega = e_1 e_2 \dots e_\ell$ and $k \in \mathbf{k}$. Then

$$\varphi_{\text{valid, noeq}}^{(t)}(\mathbf{k}, \mathbf{x}, y) \equiv \forall k' (0 \leq k' < k) \Rightarrow \bigwedge_{i=1}^{\ell} \left(y + k' \text{weight}(\omega) + \sum_{j=1}^{i-1} \text{weight}(e_j) \geq 0 \wedge y + k' \text{weight}(\omega) + \sum_{j=1}^{i-1} \text{weight}(e_j) \notin \tau(\text{start}(e_i)) \right) \wedge y + k \text{weight}(\omega) \geq 0.$$

Note that for each edge $e \in E$, $\text{weight}(e)$ is a constant, given by the automaton, and $\text{weight}(\omega)$ is a shorthand for $\sum_{i=1}^{\ell} \text{weight}(e_i)$, which is also a constant. So the only type of multiplication in the formula is by a constant. A formula of the form $a \notin \tau(u)$ is a shorthand for $\bigwedge_{b \in \tau(u)} a \neq b$, which is clearly a Presburger arithmetic formula. Since \mathcal{C} has parameterised tests, in general some of these disequalities include variables from \mathbf{x} . We can now define a formula with the required property in the case where ξ does not include any

equality tests:

$$\begin{aligned} \varphi_{comp, noeq}^{(\xi), (v, c), (v', c')}(\mathbf{k}, \mathbf{x}) &\equiv \left(\bigwedge_{i=1}^{n-1} \text{end}(t_i) = \text{start}(t_{i+1}) \right) \wedge \text{start}(t_1) = v \wedge \text{end}(t_n) = v' \wedge \\ &\sum_{i=1}^n \text{weight}(t_i(\mathbf{k})) = c' - c \wedge \bigwedge_{i=1}^n \varphi_{valid, noeq}^{(t_i)}(\mathbf{k}, \mathbf{x}, c + \sum_{j=1}^{i-1} \text{weight}(t_j(\mathbf{k}))), \end{aligned}$$

where we use the shorthand $\text{weight}(s)$ for $s \in E \cup C$: if $s \in E$ then $\text{weight}(s)$ is a constant as above, and if $s \in C$ then it is of the form $\underline{\omega}^k$ and $\text{weight}(s) = k \sum_{e \in \omega} \text{weight}(e)$. Again, the only multiplications are by constants, so the resulting formula is a formula of Presburger arithmetic.

Finally, in the case where ξ includes equality tests, we split $\text{unfold}(\xi)$ at the t_i which are equality tests, and construct a formula $\varphi_{comp, noeq}$ as above for each equality-free part of ξ . $\varphi_{comp}^{(\xi), (v, c), (v', c')}(\mathbf{k}, \mathbf{x})$ is the conjunction of these formulas. \square

Remark 4.2 (Removing the universal quantification). For simplicity, we have used a universal quantifier in $\varphi_{valid, noeq}^{(t)}(\mathbf{k}, \mathbf{x}, y)$ to express that k iterations of a cycle yield a valid computation. In fact it is possible to rewrite $\varphi_{valid, noeq}^{(t)}(\mathbf{k}, \mathbf{x}, y)$ as a purely existential formula, with a polynomial blowup. Let $\omega = e_1 \cdots e_\ell$ be a cycle and suppose we want to check that $\omega^k(y)$ is a valid computation. Let $u = \text{start}(e_i)$ be a state on the cycle. First we need to express that the counter value at u is never negative along $\omega^k(y)$. Since the counter value at u is monotone during the k iterations of the cycle (it increases if ω is positive and decreases if ω is negative), we only need check that it is nonnegative at the first and last iteration:

$$y + \sum_{j=1}^{i-1} \text{weight}(e_j) \geq 0 \wedge y + (k-1) \text{weight}(\omega) + \sum_{j=1}^{i-1} \text{weight}(e_j) \geq 0.$$

Next, for each $b \in \tau(u)$, we need to check that the cycle avoids b in u . Without loss of generality, assume that ω is positive. Then the counter value at u increases after each iteration. We can now perform a case analysis on the three ways to satisfy a disequality test during the k iterations of ω :

- The value at the first iteration is already bigger than b :

$$y + \sum_{j=1}^{i-1} \text{weight}(e_j) > b.$$

- The value at the last iteration is less than b :

$$y + (k-1) \text{weight}(\omega) + \sum_{j=1}^{i-1} \text{weight}(e_j) < b.$$

- There is an iteration k' , with $0 \leq k' < k - 1$, at which the counter value is less than b , but where at the next iteration $k' + 1$ the counter value is bigger than b :

$$\begin{aligned} \exists k' (0 \leq k' < k - 1) \wedge y + k' \text{weight}(\omega) + \sum_{j=1}^{i-1} \text{weight}(e_j) < b \\ \wedge y + (k' + 1) \text{weight}(\omega) + \sum_{j=1}^{i-1} \text{weight}(e_j) > b. \end{aligned}$$

Finally, we can use a conjunction over all states in ω to get a formula which is equivalent to $\varphi_{\text{valid, noeq}}^{(t)}(\mathbf{k}, \mathbf{x}, y)$ but has no universal quantifiers.

Lemma 4.3 (Encoding reachability). *Let $\mathcal{C} = (V, E, X, \lambda, \tau)$ be a 1-CA with parameterised tests, and let (v, c) and (v', c') be given configurations of \mathcal{C} . Then there exists a Presburger arithmetic formula $\varphi_{\text{reach}}^{(v,c),(v',c')}(\mathbf{x})$ which evaluates to true if and only if there is a valid computation from (v, c) to (v', c') in \mathcal{C} , as well as a formula $\varphi_{\text{reach}_+}^{(v,c),(v',c')}$ which is true if and only if there is such a computation of length at least 1.*

Proof. Note that the bounds on the length of computations in 1-CA from the previous section do not depend on the values occurring in equality or disequality tests. That is, if there is a valid computation $(v, c) \xrightarrow{\pi}^* (v', c')$ for any given values of the parameters, then there is a folded path χ of word length at most $p(\mathcal{C})$ such that $(v, c) \xrightarrow{\text{unfold}(\chi(c))}^* (v', c')$ is a valid computation, where p is the polynomial function given in Theorem 3.8. Equivalently, there is a path shape ξ of word length at most $p(\mathcal{C})$ and there exist values \mathbf{k} such that $(v, c) \xrightarrow{\text{unfold}(\xi(\mathbf{k})(c))}^* (v', c')$ is a valid computation.

Since path shapes are words over a finite alphabet, we can express this property as a finite disjunction

$$\varphi_{\text{reach}}^{(v,c),(v',c')}(\mathbf{x}) \equiv \exists \mathbf{k} \bigvee_{|\xi| \leq p(\mathcal{C})} \varphi_{\text{comp}}^{(\xi),(v,c),(v',c')}(\mathbf{k}, \mathbf{x}).$$

For φ_{reach_+} , we simply change the disjunction to be over all ξ such that $1 \leq |\xi| \leq p(\mathcal{C})$. \square

Lemma 4.4 (Encoding repeated control-state reachability). *Let $\mathcal{C} = (V, E, X, \lambda, \tau)$ be a 1-CA with parameterised tests, let $F \subseteq V$ be a set of final states, and let (v, c) be the initial configuration of \mathcal{C} . Then there exists a Presburger arithmetic formula $\varphi_{\text{rep-reach}}^{(v,c),(F)}(\mathbf{x})$ which evaluates to true if and only if there is a valid infinite computation π which starts in (v, c) and visits at least one state in F infinitely often.*

Proof. Suppose there is an infinite computation which starts in (v, c) and visits a state $u \in F$ infinitely often. Equivalently, there is a counter value $d \in \mathbb{N}$ such that $(v, c) \xrightarrow{}^* (u, d)$ is a valid (finite) computation, and there is a cycle ω with $\text{start}(\omega) = u$ such that $\omega^k(d)$ is a valid computation for all $k \in \mathbb{N}$. There are two possible cases:

- $\text{weight}(\omega) = 0$, so $\omega^k(d)$ is valid for all k if and only if $\omega(d)$ is valid.
- $\text{weight}(\omega) > 0$, so it might be possible to start from (u, d) and follow the edges of ω a finite number of times before an obstruction occurs. However, if ω can be taken an arbitrary number of times, then the counter value will tend towards infinity, so we are free to choose

ω to be an equality-free simple cycle, and d to be high enough to guarantee that if ω can be taken once without obstructions, it can be taken infinitely many times.

The resulting formula is then

$$\varphi_{rep\text{-}reach}^{(v,c),(F)}(\mathbf{x}) \equiv \exists d \bigvee_{u \in F} \left(\varphi_{reach}^{(v,c),(u,d)}(\mathbf{x}) \wedge \left(\varphi_{reach_+}^{(u,d),(u,d)}(\mathbf{x}) \vee \left(d > M(\mathbf{x}) \wedge \exists d' \bigvee_{\omega \in SC^+} \varphi_{comp, noeq}^{(\omega),(u,d),(u,d')}(1, \mathbf{x}) \right) \right) \right)$$

where $M(\mathbf{x}) = \max(\bigcup_{v \in V} \tau(v)) - \sum\{\text{weight}(e) : e \in E, \text{weight}(e) < 0\}$. The sum over negative edge weights ensures that the counter always stays above $\max(\bigcup_{v \in V} \tau(v))$ along the computation $\omega(d)$, since each edge is taken at most once in ω . Since ω is a positive cycle, this implies that the counter always stays above all bad values along $\omega(d^k)$ for each $k \in \mathbb{N}$, so no obstructions can occur. \square

Theorem 4.5 (Decidability of reachability problems). *Both the reachability problem and the generalised repeated control-state reachability problem are decidable for 1-CA with parameterised tests.*

Proof. Given a 1-CA $\mathcal{C} = (V, E, X, \lambda, \tau)$ with parameterised tests and configurations (v, c) and (v', c') , to check if there exist values for the parameters X such that there is a valid computation from (v, c) to (v', c') , we use Lemma 4.3 to construct the formula $\exists \mathbf{x} \varphi_{reach}^{(v,c),(v',c')}(\mathbf{x})$.

To solve the generalised repeated control-state reachability problem for a 1-CA $\mathcal{C} = (V, E, X, \lambda, \tau)$ with sets of final states $F_1, \dots, F_n \subseteq V$ and initial configuration (v, c) , note that this problem can easily be reduced to the simpler case where $n = 1$, using a translation similar to the standard translation from generalised Büchi automata to Büchi automata. In the case where $n = 1$, we can use Lemma 4.4 to construct the formula $\exists \mathbf{x} \varphi_{rep\text{-}reach}^{(v,c),(F_1)}(\mathbf{x})$. \square

Corollary 4.6 (Decidability of model checking flat Freeze LTL). *The existential model checking problem for flat Freeze LTL on 1-CA is decidable.*

5. CONCLUSION

The main result of this paper is that the model checking problem for the flat fragment of Freeze LTL on one-counter automata is decidable. We have concentrated on showing decidability rather than achieving optimal complexity. For example, we have reduced the model checking problem to the decision problem for the class of sentences of Presburger arithmetic with quantifier prefix $\exists^* \forall^*$. We explained in Remark 4.2 that in fact the reduction can be refined to yield a (polynomially larger) purely existential sentence.

Another important determinant of the complexity of our procedure is the dependence of the symbolic encoding of computations (via path shapes) in Section 4 on the number of simple cycles in the underlying control graph of the one-counter automaton. The number of such cycles may be exponential in the number of states. It remains to be seen whether it is possible to give a more compact symbolic representation, e.g., in terms of the Parikh image of paths. As it stands, our procedure for model checking flat Freeze LTL formulas on classical one-counter automata works as follows. From the flat Freeze LTL formula and

the automaton, we build a one-counter automaton with parameterised tests (of exponential size). We then guess the normal form of the path shapes (of exponential size in the size of the automaton). We finally check the resulting existential Presburger formula. Since the Presburger formula has size double exponential in the size of the input, we get a naive upper bound of 2NEXPTIME for our algorithm. Improving this bound is a subject of ongoing work.

Another interesting complexity question concerns configuration reachability in one-counter automata with non-parameterised equality and disequality tests. For automata with only equality tests and with counter updates in binary, reachability is known to be NP-complete [9]. If inequality tests are allowed then reachability is PSPACE-complete [7]. Now automata with equality and disequality tests are intermediate in expressiveness between these two models and the complexity of reachability in this case is open as far as we know.

REFERENCES

- [1] P. Bouyer, N. Markey, J. Ouaknine, and J. Worrell. On expressiveness and complexity in real-time model checking. In *Proceedings of ICALP*, volume 5126 of *LNCS*, pages 124–135. Springer, 2008.
- [2] H. Comon and V. Cortier. Flatness is not a weakness. In *Proceedings of CSL*, volume 1862 of *LNCS*. Springer, 2000.
- [3] S. Demri and R. Lazic. LTL with the freeze quantifier and register automata. *ACM Trans. Comput. Log.*, 10(3):16:1–16:30, 2009.
- [4] S. Demri, R. Lazic, and D. Nowak. On the freeze quantifier in constraint LTL: Decidability and complexity. *Inf. Comput.*, 205(1):2–24, 2007.
- [5] S. Demri, R. Lazic, and A. Sangnier. Model checking memoryful linear-time logics over one-counter automata. *Theor. Comput. Sci.*, 411(22-24):2298–2316, 2010.
- [6] S. Demri and A. Sangnier. When model-checking freeze LTL over counter machines becomes decidable. In *Proceedings of FOSSACS*, volume 6014 of *LNCS*, pages 176–190, 2010.
- [7] John Fearnley and Marcin Jurdzinski. Reachability in two-clock timed automata is PSPACE-complete. *Inf. Comput.*, 243:26–36, 2015.
- [8] T. French. Quantified propositional temporal logic with repeating states. In *Proceedings of TIME-ICTL*, pages 155–165. IEEE Computer Society, 2003.
- [9] C. Haase, S. Kreutzer, J. Ouaknine, and J. Worrell. Reachability in succinct and parametric one-counter automata. In *Proceedings of CONCUR*, volume 5710 of *LNCS*, pages 369–383. Springer, 2009.
- [10] O. H. Ibarra, T. Jiang, N. Tran, and H. Wang. New decidability results concerning two-way counter machines and applications. In *Proceedings of ICALP*, volume 700 of *LNCS*. Springer, 1993.
- [11] J. Leroux and G. Sutre. Flat counter automata almost everywhere! volume 3707 of *Lecture Notes in Computer Science*. Springer, 2005.
- [12] A. Lisitsa and I. Potapov. Temporal logic with predicate lambda-abstraction. In *Proceedings of TIME*, pages 147–155. IEEE Computer Society, 2005.
- [13] M. Presburger. Über die Vollständigkeit eines gewissen Systems der Arithmetik ganzer Zahlen, in welchem die Addition als einzige Operation hervortritt. In *Comptes Rendus du I congrs de Mathmaticiens des Pays Slaves. Warsaw*, pages 92–101, 1929.