The Complexity of Reversal-Bounded Model-Checking *

Marcello M. Bersani[1,2], Stéphane Demri[1]

[1] LSV, ENS Cachan, CNRS, INRIA, France [2] Politecnico di Milano, Italy

Abstract. We study model-checking problems on counter systems when guards are quantifier-free Presburger formulae, the specification languages are LTL-like dialects with arithmetical constraints and the runs are restricted to reversal-bounded ones. We introduce a generalization of reversal-boundedness and we show the NExpTime-completeness of the reversal-bounded model-checking problem as well as for related reversal-bounded reachability problems. As a by-product, we show the effective Presburger definability for sets of configurations for which there is a reversal-bounded run verifying a given temporal formula. Our results generalize existing results about reversal-bounded counter automata and provides a uniform and more general framework.

1 Introduction

Reversal-bounded model-checking. Given a counter system $\mathcal{S}$ and a linear-time property ϕ expressed in a logical formalism, a standard question in formal verification is to determine whether there is an infinite run ρ for $\mathcal{S}$ satisfying ϕ (written $\rho \models \phi$), or dually whether all runs satisfy ϕ. In full generality, existential model-checking problem is undecidable (as an immediate consequence of the undecidability of the halting problem for Minsky machines). A way to overcome this difficulty is to consider a subclass of runs for $\mathcal{S}$ for which decidability is regained; in that case, we answer a different question but in case of positive answer, starting from a subclass of runs does not harm. In the paper, we restrict the runs so that along an infinite run, for each counter the number of reversals is bounded by a given value r; a reversal is witnessed when a counter behaviour changes from increasing mode to decreasing mode, or vice-versa. We follow an approach similar to bounded model-checking (BMC), see e.g. [6], in which runs are built until positions of a bounded distance from the initial configuration. Analogously, in context-bounded model-checking, the number of segments of the computation during which only one thread is active is bounded in multithreaded programs, see e.g. [29]. As for bounded model-checking, in case of negative answer to the question, the value r can be incrementally augmented. Reversal-bounded counter systems have been first studied in [20] and several extensions have been considered in the literature, see e.g. [13]. A major property

* Work supported by Agence Nationale de la Recherche, grant ANR-06-SETIN-001 and by the European Commission, Project 227977-SMScom.

of such systems is that the sets of configurations reachable from a given initial configuration, are effectively Presburger-definable. However this does not entail that problems involving infinite runs are decidable, since infinite runs are not necessarily effectively representable in Presburger arithmetic, see e.g. [10]. In this paper, we study problems of the form: given a counter system $\mathcal{S}$, a bound $r \geq 0$ and a formula ϕ, is there an infinite r-reversal-bounded run ρ such that $\rho \models \phi$. To complete our analogy, it is fair to observe that BMC for finite systems benefits from nice properties on runs that allow the existence of an upper bound on the length of runs to be checked (a.k.a. *completeness threshold*). That is why, a finite amount of BMC instances needs to be checked in order to provide an answer to any instance of the model-checking problem. By contrast, since the reversal-boundedness detection problem on counter systems is undecidable [20], there is no guarantee that given an initialized counter system, there exists a r-reversal-bounded run, for some $r \geq 0$, satisfying a given temporal property.

Our motivations. In order to test whether there is an infinite run satisfying a temporal property, we restrict ourselves to r-reversal-bounded runs for some $r \geq 0$ so that for a fixed r, the problem is decidable. In case of positive answer, we stop the process, otherwise we increment r and perform again a test. This incremental approximation approach is applied to counter systems that are more general than Minsky machines (counter automata with increments, decrements and zero-tests), typically by considering guards definable in quantifier-free fragment of Presburger arithmetic and update vectors in $\mathbb{Z}^n$. Moreover, we aim at expressing the temporal property in a rich LTL-like dialect, including arithmetical constraints and past-time operators (i.e., not only restricted to Boolean combinations of **GF**-formulae). Finally, not only we characterize the computational complexity of the existence of r-reversal-bounded runs but also our goal is to effectively express the set of configurations admitting such runs in Presburger arithmetic, which will allow us to use SMT solvers to perform verification tasks on counter systems (see e.g. [2, 26]) or to take advantage of verification techniques developed in [5]. It is worth noting that the use of Presburger arithmetic for formal verification has been already advocated since the work [31].

Our contributions. As far as the methodology is concerned, we reduce model-checking problems to reachability problems (first, by synchronization of the counter system and the automaton representing the temporal formula and, then, we reduce the model-checking problems to reachability problems). Let us quote the major results of the paper. (i) We introduce a new concept for reversal-boundedness that makes explicit the role of arithmetical terms and it captures previous notions on reversal-boundedness (see Section 2). (ii) We show that the reversal-bounded model-checking problem for counter systems with guards in QFP (quantifier-free fragment of Presburger arithmetic) and temporal formulae with atomic formulae in QFP is decidable and NExpTime-complete (see Theo. 4). The same complexity applies to reversal-bounded control state repeated reachability problem and reversal-bounded reachability problem (see Corollary 6). (iii) We show that the existence of reversal-bounded runs satisfying a temporal property implies the existence of reversal-bounded runs that are

ultimately periodic, i.e. the sequences of transitions are of the form $\pi_1 \cdot (\pi_2)^\omega$ where π_1 and π_2 are finite sequences. This type of properties has been already shown useful to implement verification methods following the BMC paradigm (see Corollary 3). (iv) Besides, our complexity results provide as by-products that reachability sets for reversal-bounded counter systems are effectively Presburger definable (see Corollary 4) and sets of configurations for which there is a reversal-bounded run verifying a temporal formula are also effectively Presburger definable (Theorem 5).

Related works. Effective Presburger definability for reversal-bounded Minsky machines and more generally for reversal-bounded counter systems can be found in [20, 19, 13] whereas the NExpTime-completeness of the reversal-bounded reachability problem for Minsky machines has been shown in [17] (lower bound) and [14, 17] (upper bound). The NExpTime upper bounds established in this paper for several extensions with richer classes of counter systems or with richer concepts of reversal-boundedness build on [14] and on [30] with adaptations to handle more complicated technical features. Decidability results for reversal-bounded counter systems augmented with familiar data structures such as stacks or queues (also with restricted behaviours) can be found in [18]. Our temporal language is very expressive and includes control states as well as arithmetical constraints in QFP. Moreover, in the paper we deal with model-checking involving linear-time temporal logics with past-time and future temporal operators and with arithmetical constraints on counter values. In [10], it is shown that $\exists$-Presburger-infinitely often problem for reversal-bounded counter automata (with guards made of Boolean combinations of the form $\mathtt{x}_i \sim k$) is decidable. Moreover, $\exists$-Presburger-always problem for reversal-bounded counter automata is undecidable [10]. Our decidability results on model-checking refine these results in order to obtain new decidability results, by allowing a full LTL specification language with arithmetical constraints and by proposing new concepts for reversal-boundedness. Finally, in [21, Theorem 22], ExpTime upper bound for LTL model-checking over reversal-bounded counter automata is shown but the logical language has no arithmetical constraint and the number of reversals r is encoded in unary (see also [32]). In our setting, our complexity results deal with instances in which all the integers are encoded in binary.

The recent work [16] is also closely related to our paper. We are grateful to an anonymous referee for pointing it to us. Our work and [16] have been done independently but most of our results can be reproved by extending [16]. In [16], operational models extending pushdown systems with counters and clocks are considered; a version of reversal-bounded LTL model-checking is shown to be co-NExpTime [16, Theorem 2]. A prototypical implementation and experimental results are also presented in [16]. LTL dialect contains only control states and guards are Boolean combinations of constraints of the form $\mathtt{x} \sim k$. By contrast, models are more general than ours. Theorem 2 in [16] is based on [16, Theorem 1] that also implies that reversal-bounded reachability problem considered in our paper is in NExpTime (assuming atomic guards of the form $\mathtt{x} \sim k$). Unlike [16], our LTL dialect contains control states, past-time operators but also arithmetical

constraints in QFP allowing non-trivial arithmetical properties like $\mathbf{GF}(\mathtt{Xx} = \mathtt{x} + \mathtt{y})$ (which may lead to undecidability in the general case). Similarly, even though our paper deals only with counter systems (no stack, no clocks), we allow general guards from QFP; we also introduce a new concept for reversal-boundedness. The proof of [16, Theorem 1] share common features with our proof of Theorem 2, at least in the use of counter modes. In both cases Presburger formulae are built: our proof is based on a run analysis whereas the proof in [16, Theorem 1] builds directly the formula. We believe our treatment is more uniform and it generalizes notions presented in [19, 10]. Moreover, our run analysis for proving Theorem 2 is interesting for its own sake, see [4].

In general, omitted proofs can be found in [4] (submitted version).

2 Preliminaries

In this section, we introduce a language for arithmetical constraints, namely the quantifier-free fragment of Presburger arithmetic (over the set of natural numbers). This language serves two main purposes. Firstly, we define classes of operational models, namely counter systems, for which transitions are guarded by arithmetical constraints. Secondly, we introduce a version of linear-time temporal logic with past-time operators for which atomic formulae can state properties about counter values, i.e. arithmetical constraints.

Arithmetical constraints. We write $\mathbb{N}$ (resp. $\mathbb{Z}$) for the set of natural (resp. integers) numbers and $[m, m']$ with $m, m' \in \mathbb{Z}$ to denote the set $\{j \in \mathbb{Z} : m \leq j \leq m'\}$. For $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Z}^n$, we write $\boldsymbol{x}(1), \ldots, \boldsymbol{x}(n)$ for the entries of $\boldsymbol{x}$, $\boldsymbol{x} \preceq \boldsymbol{y} \stackrel{\text{def}}{\Leftrightarrow}$ for all $i \in [1, n]$, $\boldsymbol{x}(i) \leq \boldsymbol{y}(i)$ and $\boldsymbol{x} \prec \boldsymbol{y}$ when $\boldsymbol{x} \preceq \boldsymbol{y}$ and $\boldsymbol{x} \neq \boldsymbol{y}$.

Let $\mathrm{VAR} = \{\mathtt{x}_0, \mathtt{x}_1, \ldots\}$ be a countably infinite set of variables. We define below formulae from the quantifier-free theory of natural numbers with addition, also known as quantifier-free fragment of Presburger arithmetic. Terms $\mathtt{t}$ are defined from the grammar $\mathtt{t} := a\mathtt{x} \mid \mathtt{t} + \mathtt{t}$, where $\mathtt{x} \in \mathrm{VAR}$, $a \in \mathbb{Z}$ (encoded with a binary representation). A *valuation* $\mathbf{val}$ is a map $\mathbf{val} : \mathrm{VAR} \to \mathbb{N}$ and it can be extended to the set of all terms as follows: $\mathbf{val}(a\mathtt{x}) = a \times \mathbf{val}(\mathtt{x})$, $\mathbf{val}(\mathtt{t} + \mathtt{t}') = \mathbf{val}(\mathtt{t}) + \mathbf{val}(\mathtt{t}')$. It is worth noting that variables take values over $\mathbb{N}$ but terms take values over $\mathbb{Z}$. Formulae ξ of QFP are defined from the grammar $\xi ::= \top \mid \mathtt{t} \leq k \mid \mathtt{t} \geq k \mid \mathtt{t} \equiv_c k' \mid \xi \wedge \xi \mid \neg\xi$, where $\top$ is the truth constant, $c \in \mathbb{N} \setminus \{0, 1\}$, $k \in \mathbb{Z}$ and $k' \in \mathbb{N}$. The satisfaction relation $\models_{\mathrm{PA}}$ for QFP formulae is briefly recalled below:

- $\mathbf{val} \models_{\mathrm{PA}} \mathtt{t} \equiv_c k' \stackrel{\text{def}}{\Leftrightarrow}$ there is $n \in \mathbb{Z}$ such that $nc + \mathbf{val}(\mathtt{t}) = k'$,
- $\mathbf{val} \models_{\mathrm{PA}} \mathtt{t} \leq k \stackrel{\text{def}}{\Leftrightarrow} \mathbf{val}(\mathtt{t}) \leq k$ (and similarly with $\geq$),
- $\mathbf{val} \models_{\mathrm{PA}} \neg\phi \stackrel{\text{def}}{\Leftrightarrow} \mathbf{val} \not\models_{\mathrm{PA}} \phi$; $\mathbf{val} \models_{\mathrm{PA}} \phi \wedge \phi' \stackrel{\text{def}}{\Leftrightarrow} \mathbf{val} \models_{\mathrm{PA}} \phi$ and $\mathbf{val} \models_{\mathrm{PA}} \phi'$.

A valuation $\mathbf{val}$ restricted to variables in $V = \{\mathtt{x}_1, \ldots, \mathtt{x}_n\} \subseteq \mathrm{VAR}$ can be also represented by a vector $\boldsymbol{x} \in \mathbb{N}^n$, where $\mathbf{val}(\mathtt{x}_j) = \boldsymbol{x}(j)$ for $j \in [1, n]$. Hence, assuming that ϕ has n distinct variables, the satisfaction relation can be equivalently written with respect to a vector of values: $\boldsymbol{x} \models_{\mathrm{PA}} \phi$ (in place of

$\mathbf{val} \models_{\mathrm{PA}} \phi$ with $\mathbf{val}(\mathtt{x}_i) = \boldsymbol{x}(i)$). Full Presburger arithmetic (i.e., with first-order quantification over natural numbers) has been shown decidable in [28] by means of quantifier elimination. Moreover, the satisfiability problem for QFP is known to be NP-complete, see e.g. [27].

We present below a few notations used in the sequel: QFP is also denoted by QFP$(<, \equiv)$ whereas its restriction without periodicity constraints is denoted by QFP$(<)$. Similarly, we write QFP$(<_1, \equiv)$ to denote the restriction of QFP$(<, \equiv)$ with atomic formulae involving at most one variable; QFP$(<_1, \equiv)$ without periodicity constraints is denoted by QFP$(<_1)$. Wlog., we can assume that the atomic formulae of QFP$(<_1, \equiv)$ are of one of the forms below: $\mathtt{x} \sim k$ with $k \in \mathbb{N}$ and $\sim \in \{<, \leq, >, \geq\}$ or $\mathtt{t} \equiv_c k'$ with $c > 1$ and $k' \in [0, c-1]$.

Counter systems. In this paper, we consider *counter systems* to be finite-state automata equipped with a finite set of counters $\{1, \ldots, n\}$ with values over $\mathbb{N}$; a counter system is a tuple $\mathcal{S} = (Q, n, \delta)$ where Q is a finite set of control states, $n \geq 1$ is the number of counters and δ is a finite subset of $Q \times (\mathrm{QFP} \times \mathbb{Z}^n) \times Q$ such that whenever $(q, (\phi, \boldsymbol{v}), q') \in \delta$ (also written $q \xrightarrow{(\phi, \boldsymbol{v})} q'$), ϕ is a formula in QFP with variables among $\mathtt{x}_1, \ldots, \mathtt{x}_n$ (a guard on the n counters) and $\boldsymbol{v} \in \mathbb{Z}^n$ is the *update vector*. Elements of δ are called *transitions*, i.e. rules acting on counters. A *configuration* of $\mathcal{S}$ is defined as a pair $(q, \boldsymbol{x}) \in Q \times \mathbb{N}^n$, where $\boldsymbol{x}$ is the vector of values for counters. The *one-step transition relation* $\rightarrow \subseteq Q \times \mathbb{N}^n \times Q \times \mathbb{N}^n$ is defined between a pair of configurations such that $((q, \boldsymbol{x}), (q', \boldsymbol{x'})) \in \rightarrow \stackrel{\text{def}}{\Leftrightarrow}$ there is a transition $t = q \xrightarrow{(\phi, \boldsymbol{v})} q'$ in δ, $\boldsymbol{x} \models_{\mathrm{PA}} \phi$ and $\boldsymbol{x'} = \boldsymbol{x} + \boldsymbol{v}$ (in that case, we write $(q, \boldsymbol{x}) \xrightarrow{t} (q', \boldsymbol{x'})$). A *run* ρ is a (possibly infinite) sequence of configurations $(q_0, \boldsymbol{x}_0), (q_1, \boldsymbol{x}_1) \ldots$ such that two successive configurations agree with δ, i.e. for $i \geq 0$, we have $(q_i, \boldsymbol{x_i}) \xrightarrow{t} (q_{i+1}, \boldsymbol{x_{i+1}})$, for some $t \in \delta$. An *initialized* counter system is a pair $(\mathcal{S}, (q, \boldsymbol{x}))$ such that $\mathcal{S}$ is a counter system and $(q, \boldsymbol{x})$ is an *initial configuration* (with $\boldsymbol{x} \geq 0$).

Given a subset L of QFP, we write CS(L) to denote the class of counter systems for which transitions are restricted to guards in L. Clearly, Minsky machines (and also vector addition systems with states) are included in CS(QFP$(<_1)$). Then, most of all the reachability problems are already undecidable as soon as CS(L) contains CS(QFP$(<_1)$). For this reason, in order to get decidability for reachability and model-checking problems, some restrictions have to be imposed on the nature of the systems. The notion of reversal-boundedness introduced in [20] is based on a semantical restriction that entails the decidability of several reachability problems. Informally, a *reversal* for a counter occurs in a run when there is an alternation from nonincreasing to nondecreasing mode.

Below, we propose a slight generalization that captures the notion of reversal-boundedness from [20] and the notion of *strong* reversal-boundedness introduced in [19, Section 4.2.2]. In a few words, in our new definition below, reversal-boundedness applies to counters but *also* to terms occurring in guards. Let $\mathcal{S} = (Q, n, \delta)$ be a counter system and $\mathtt{T}$ be a finite set of terms including $\{\mathtt{x}_1, \ldots, \mathtt{x}_n\}$. Let us linearly order the terms in $\mathtt{T}$ with $\mathtt{x}_1, \ldots, \mathtt{x}_n, \mathtt{t}_1, \ldots, \mathtt{t}_{n'}$. So, $\mathrm{card}(\mathtt{T}) = n + n'$ (n' can possibly be equal to 0). From a run $\rho = (q_0, \boldsymbol{x_0}), (q_1, \boldsymbol{x_1}), \ldots$ of

$\mathcal{S}$, in order to describe the behavior of counters and terms varying along ρ, we define a sequence of *mode vectors* $\boldsymbol{m_0}, \boldsymbol{m_1}, \ldots$ (of the same length as ρ) such that each $\boldsymbol{m_i}$ belongs to $\{\nearrow, \searrow\}^{n+n'}$. Intuitively, each value in a mode vector records whether a term is currently in an increasing phasis or in an decreasing phasis (this includes the counters themselves as in standard reversal-boundedness). Given a term $\mathtt{t} = \sum_k a_k \mathtt{x}_k$ and a counter vector $\boldsymbol{x}$, we write $\boldsymbol{x}(\mathtt{t})$ to denote the integer $\sum a_k \boldsymbol{x}(k)$. We are now ready to define the sequence $\boldsymbol{m_0}, \boldsymbol{m_1}, \ldots$

- By convention, m_0 is the unique vector in $\{\nearrow\}^{n+n'}$.
- For $j \geq 0$ and $i \in [1, n+n']$ with the ith term in $\mathtt{T}$ equal to $\mathtt{t}$, we have $\boldsymbol{m}_{j+1}(i) \stackrel{\text{def}}{=} \boldsymbol{m}_j(i)$ when $\boldsymbol{x}_j(\mathtt{t}) = \boldsymbol{x}_{j+1}(\mathtt{t})$, $\boldsymbol{m}_{j+1}(i) \stackrel{\text{def}}{=} \nearrow$ when $\boldsymbol{x}_{j+1}(\mathtt{t}) - \boldsymbol{x}_j(\mathtt{t}) > 0$ and $\boldsymbol{m}_{j+1}(i) \stackrel{\text{def}}{=} \searrow$ when $\boldsymbol{x}_{j+1}(\mathtt{t}) - \boldsymbol{x}_j(\mathtt{t}) < 0$.

It is worth noting that if $(q_j, \boldsymbol{x_j}) \xrightarrow{t} (q_{j+1}, \boldsymbol{x_{j+1}})$ with $t = q_j \xrightarrow{(\phi, \boldsymbol{v})} q_{j+1}$, then $\boldsymbol{x}_{j+1}(\mathtt{t}) - \boldsymbol{x}_j(\mathtt{t}) = \sum_k a_k \boldsymbol{v}(k)$. Now, let $Rev_i = \{j \in \mathbb{N} : \boldsymbol{m}_j(i) \neq \boldsymbol{m}_{j+1}(i)\}$; we say that ρ is r-$\mathtt{T}$-*reversal-bounded* for some $r \geq 0 \stackrel{\text{def}}{\Leftrightarrow}$ for all $i \in [1, n+n']$, $\text{card}(Rev_i) \leq r$. Given a counter system $\mathcal{S}$, we write $\mathtt{T}_\mathcal{S}$ to denote the finite set of terms $\mathtt{t}$ occurring in atomic guards of the form $\mathtt{t} \sim k$ with $\sim \in \{\leq, \geq\}$ and $k \in \mathbb{Z}$, plus the distinguished terms (counters) from $\{\mathtt{x}_1, \ldots, \mathtt{x}_n\}$. Note that terms occurring only in periodicity constraints are not taken into account; we shall deal with them separately (see Section 3). An initialized counter system $(\mathcal{S}, (q, \boldsymbol{x}))$ is *reversal-bounded* $\stackrel{\text{def}}{\Leftrightarrow}$ there is $r \geq 0$ such that every run from $(q, \boldsymbol{x})$ is r-$\mathtt{T}_\mathcal{S}$-reversal-bounded.

When $\mathtt{T}$ is reduced to $\{\mathtt{x}_1, \ldots, \mathtt{x}_n\}$, $\mathtt{T}$-reversal-boundedness is equivalent to reversal-boundedness from [20]. Hence, for $\mathcal{S} \in \text{CS}(\text{QFP}(<_1))$ and initial configuration $(q, \boldsymbol{x})$, $(\mathcal{S}, (q, \boldsymbol{x}))$ is reversal-bounded in the sense herein iff $(\mathcal{S}, (q, \boldsymbol{x}))$ is reversal-bounded in the sense from [20]. In strong reversal-boundedness [19, Sect. 4.2.2], a phasis can be either strictly increasing, or strictly decreasing or constant (mode vectors belong to $\boldsymbol{m_i} \in \{\nearrow, \searrow, \rightarrow\}^{n+n'}$). This provides more constraints on runs: the guards are more general (typically in QFP($<$)) and the update vectors are in $\{-1, 0, +1\}^n$. Again, our notion of $\mathtt{T}$-reversal-boundedness allows us to provide a uniform and more general treatment. Indeed, when a sequence of transitions has a unique update vector, the mode vector remains constant. When an initialized counter system from CS(QFP), involving guards with terms in $\mathtt{T}'$, is strongly reversal-bounded in the sense of [19, Sect. 4.2.2], then it is $(\mathtt{T}' \cup \{\mathtt{x}_1, \ldots, \mathtt{x}_n\})$-reversal-bounded, too.

Given a class $\mathcal{C}$ of counter systems, the *reversal-bounded reachability problem for* $\mathcal{C}$, written RB-REACH($\mathcal{C}$), is defined as follows (all integers are encoded in binary): given a counter system $\mathcal{S} \in \mathcal{C}$, configurations $(q_0, \boldsymbol{x_0})$ and $(q_f, \boldsymbol{x_f})$, $r \geq 0$, is there an r-$\mathtt{T}_\mathcal{S}$-reversal-bounded run from $(q_0, \boldsymbol{x_0})$ to $(q_f, \boldsymbol{x_f})$? Clearly, when $(\mathcal{S}, (q_0, \boldsymbol{x_0}))$ is reversal-bounded, reversal-bounded reachability corresponds exactly to reachability. Similarly, the *reversal-bounded control state repeated reachability problem for* $\mathcal{C}$, written RB-REP-REACH($\mathcal{C}$), is defined as follows: given a counter system $\mathcal{S} \in \mathcal{C}$, a configuration $(q_0, \boldsymbol{x_0})$, a control state q_f and $r \geq 0$, is there an infinite r-$\mathtt{T}_\mathcal{S}$-reversal-bounded run from $(q_0, \boldsymbol{x_0})$ such that q_f is repeated infinitely often? Both problems RB-REACH($\mathcal{C}$) and RB-REP-REACH($\mathcal{C}$)

restrict the set of runs witnessing a simple property (reaching $(q_f, \boldsymbol{x_f})$ or repeating infinitely often q_f). This makes sense in our incremental approximation approach, since removing the restriction leads to undecidability. However, it is worth noting that our new notion of $\mathtt{T}$-reversal-boundedness is rich enough so that witness runs include standard reversal-bounded runs.

In the sequel, we show that RB-REACH(CS(QFP)) is NExpTime-complete. It is worth explaining why this is consistent with the fact that the reachability problem for (standard) reversal-bounded counter automata augmented with guards of the form $\mathtt{x}_i = \mathtt{x}_{i'}$ or $\mathtt{x}_i \neq \mathtt{x}_{i'}$ is undecidable [19]. Indeed, the presence of such guards entails the presence of terms of the form $\mathtt{x}_i - \mathtt{x}_{i'}$, that have to be reversal-bounded by definition of RB-REACH(CS(QFP)). However, it is not difficult to show that the undecidability proof in [19] produces enriched counter automata for which some terms of the form $\mathtt{x}_i - \mathtt{x}_{i'}$ are not reversal-bounded.

Reversal-bounded model-checking problems. We define below a linear-time temporal logic with future-time and past-time operators. Atomic formulae are either control states or arithmetical constraints about counter values at the current position and next position. *Counter variables* in $\mathrm{VAR} = \{\mathtt{x}_1, \mathtt{x}_2, \ldots\}$ are free variables, only interpreted by the counter values on configurations. As for defining QFP, *arithmetical terms* are defined by the grammar $\mathtt{t} ::= a\mathtt{x} \mid a\mathrm{X}\mathtt{x} \mid \mathtt{t} + \mathtt{t}$ with $\mathtt{x} \in \mathrm{VAR}$ and $a \in \mathbb{Z}$. Intuitively, $\mathtt{x}$ refers to the current value for counter $\mathtt{x}$, $\mathrm{X}\mathtt{x}$ refers to the counter value for $\mathtt{x}$ at the next position from the current one. Formulae of CLTL(QFP) are defined as follows:

$$\phi ::= \top \mid q \mid \mathtt{t} \sim k \mid \mathtt{t} \equiv_c k' \mid \neg\phi \mid \phi \wedge \phi \mid \mathbf{X}\phi \mid \phi\mathbf{U}\phi \mid \mathbf{Y}\phi \mid \phi\mathbf{S}\phi$$

with $q \in Q$, $\sim \in \{<, \leq, >, \geq, =\}$, $k \in \mathbb{Z}$, $c \in \mathbb{N} \setminus \{0, 1\}$ and $k' \in \mathbb{N}$. As usual, we pose $\mathbf{F}\phi \stackrel{\text{def}}{=} \top\mathbf{U}\phi$ and $\mathbf{G}\phi \stackrel{\text{def}}{=} \neg\mathbf{F}\neg\phi$. The formula $\mathbf{GF}(\mathtt{x}_1 - \mathtt{x}_2 = 3)$ states that infinitely often the value for counter 1 is equal to the value for counter 2 plus 3. Given a fragment $\mathrm{L} \subseteq \mathrm{QFP}$, we write CLTL(L) to denote the restriction of CLTL(QFP) with arithmetical constraints built from L.

Models of CLTL(QFP) are intended to be infinite runs of counter systems; hence they are of the form $\rho = (q_0, \boldsymbol{x_0}), (q_1, \boldsymbol{x_1}), (q_2, \boldsymbol{x_2}), \ldots$ with $\rho \in (Q \times \mathbb{N}^n)^\omega$. In order to deal with arithmetical constraints, we need to introduce a few notations. Given a term $\mathtt{t}$ from CLTL(QFP), we write $\tilde{\mathtt{t}}$ to denote the term in QFP obtained from $\mathtt{t}$ by replacing $\mathrm{X}\mathtt{x}_i$ by a fresh variable $\mathtt{x}'_i$. Then, satisfaction relation $\models$ is defined as follows (we omit obvious Boolean clauses):

- $\rho, i \models q \stackrel{\text{def}}{\Leftrightarrow} q = q_i$.
- $\rho, i \models \mathtt{t} \sim k \stackrel{\text{def}}{\Leftrightarrow} \mathbf{val} \models_{\mathrm{PA}} \tilde{\mathtt{t}} \sim k$ where for $j \in [1, n]$, $\mathbf{val}(\mathtt{x}_j) = \boldsymbol{x_i}(j)$ and $\mathbf{val}(\mathtt{x}'_j) = \boldsymbol{x_{i+1}}(j)$. Similarly, $\rho, i \models \mathtt{t} \equiv_c k' \stackrel{\text{def}}{\Leftrightarrow} \mathbf{val} \models_{\mathrm{PA}} \tilde{\mathtt{t}} \equiv_c k'$.
- $\rho, i \models \mathbf{X}\phi \stackrel{\text{def}}{\Leftrightarrow} \rho, i+1 \models \phi$; $\rho, i \models \mathbf{Y}\phi \stackrel{\text{def}}{\Leftrightarrow} \rho, i-1 \models \phi$ and $i \geq 1$.
- $\rho, i \models \phi\mathbf{U}\phi' \stackrel{\text{def}}{\Leftrightarrow}$ there is $j \geq i$ s.t. $\rho, j \models \phi'$ and for all $h \in [i, j-1]$, $\rho, h \models \phi$.
- $\rho, i \models \phi\mathbf{S}\phi' \stackrel{\text{def}}{\Leftrightarrow}$ there is $j \leq i$ s.t. $\rho, j \models \phi'$ and for all $h \in [j-1, i]$, $\rho, h \models \phi$.

Observe that $\mathbf{X}$ is a temporal operator whereas X is used to refer to next counter values and it does not admit nesting. Moreover, the syntax of CLTL(QFP) does

not allow terms that refer to counter values at the previous position; again, this can be easily simulated. For instance, current value for counter 1 is equal to the value of counter 2 at the previous position can be encoded by the formula $\mathbf{Y}(\mathtt{x}_2 = \mathrm{X}\mathtt{x}_1)$.

The basic idea behind the design of CLTL(QFP) is to allow comparisons between counter values at successive positions of the runs. Similar motivations can be found in the introduction of concrete domains in description logics, that are logic-based formalisms for knowledge representation [25]. Temporal logics with Presburger constraints have been developed, for instance, in [9, 8, 22]. Some of them have quite expressive decidable fragments. Undecidability of the existential model-checking problem for CLTL(QFP) can be shown using the undecidability of the halting problem for Minsky machines. SMT solvers can be used for checking bounded reachability problems, see e.g., [5].

Given an CLTL(QFP) formula ϕ, we write $\mathtt{T}_\phi$ to denote the finite set of terms of the form $\sum_k (a_k + b_k)\mathtt{x}_k$ when $\mathtt{t} = (\sum_k a_k \mathrm{X}\mathtt{x}_k) + (\sum_k b_k \mathtt{x}_k)$ is a term occurring in ϕ (modulo AC for the operator $+$) in an atomic formula of the form $\mathtt{t} \sim k$ with $\sim \in \{\leq, \geq, <, >, =\}$ and $k \in \mathbb{Z}$. Since the next value of counter k (denoted by $\mathrm{X}\mathtt{x}_k$) is equal to the current value of the counter plus some $b \in \mathbb{Z}$ (depending on the update vectors of the transitions), the value of the term $(\sum_k a_k \mathrm{X}\mathtt{x}_k) + (\sum_k b_k \mathtt{x}_k)$ is equal to the current value of $\sum_k (a_k + b_k)\mathtt{x}_k$ plus some constant depending on the next transition. This explains the current definition of $\mathtt{T}_\phi$ and more justifications can be found in Section 3.

Reversal-bounded model-checking problem RBMC is defined as follows: given a counter system $\mathcal{S} \in \mathrm{CS(QFP)}$, a configuration $(q, \boldsymbol{x})$, a formula $\phi \in \mathrm{CLTL(QFP)}$ and bound $r \in \mathbb{N}$, is there an infinite run ρ from $(q, \boldsymbol{x})$ such that $\rho, 0 \models \phi$ and ρ is r-$\mathtt{T}$-reversal-bounded with $\mathtt{T} = \mathtt{T}_\mathcal{S} \cup \mathtt{T}_\phi$? The restriction of RBMC to counter systems in the class $\mathrm{CS(L_1)}$ and to formulae in $\mathrm{CLTL(L_2)}$ is denoted by $\mathrm{RBMC(CS(L_1),CLTL(L_2))}$ with $\mathrm{L_1, L_2} \subseteq \mathrm{QFP}$. If $\mathrm{L_1 = L_2 = QFP}(<_1)$, the witness run ρ should simply be reversal-bounded in the sense of [20] ($\mathtt{T}_\phi = \mathtt{T}_\mathcal{S} = \{\mathtt{x}_1, \ldots, \mathtt{x}_n\}$). Similarly, if $\mathrm{L_1 = L_2 = QFP}$, then the set of witness runs include the set of strongly reversal-bounded runs from $(q, \boldsymbol{x})$ in the sense of [19, Section 4.2.2]. We can impose that witness runs are exactly strongly reversal-bounded by adding the subformula $\bigvee_{q \xrightarrow{(\xi, \upsilon)} q' \in \delta} \mathbf{FG}(\bigwedge_{i \in [1,n]}((\mathrm{X}\mathtt{x}_i - \mathtt{x}_i) = \boldsymbol{v}(i)))$. Do note that a richer class of witness runs is allowed by our definition. The main result of the paper is the NExpTime-completeness of RBMC (with all integers admitting a binary representation). Observe also that both RB-REACH(CS(QFP)) and RB-REP-REACH(CS(QFP)) can be easily reduced to RBMC.

3 From reversal-bounded model-checking to reachability

Herein, we show how to reduce RMBC into RB-REP-REACH(QFP), RB-REP-REACH(QFP) into RB-REP-REACH(QFP(<)) and RB-REP-REACH(QFP(<)) into RB-REACH(QFP(<)). In Section 4, we deal with RB-REACH(QFP(<)) complexity as well as with RMBC and RB-REP-REACH(QFP) complexity. The two first reductions presented below use quite standard proof techniques but we

have to perform them carefully since we shall reuse their complexity functions to establish the final complexity upper bound for RMBC, see e.g. [12] for the first reduction (see also [33]). It is worth nothing that each reduction below produces an exponential blow-up.

Towards control state repeated reachability. In this section, we show how to reduce RBMC to RB-REP-REACH(QFP) by synchronizing counter systems with Büchi automata for temporal formulae, as done for LTL model-checking [34], see also developments for Petri nets in [12]. The definition of a synchronized product is motivated by the design of a unique counter system that captures the Büchi acceptance condition and the update of counters following the transitions of $\mathcal{S}$.

Let $\mathcal{S} = (Q, n, \delta) \in \mathrm{CS(QFP)}$, $(q, \boldsymbol{x})$, $\phi \in \mathrm{CLTL(QFP)}$ and $r \in \mathbb{N}$ be an instance of RBMC. The formula ϕ can be viewed as a standard LTL formula in which the atomic formulae of the form q, $\mathtt{t} \sim k$ and $\mathtt{t} \equiv_c k'$ are viewed as propositional variables. From [34], we know that we can represent the symbolic models of ϕ by a Büchi automaton $\mathcal{A}_\phi$ whose size is exponential in the size of ϕ. At the symbolic level, the counter values are disgarded. The instance we shall build for RB-REP-REACH(QFP) is obtained by synchronizing $\mathcal{A}_\phi$ with $\mathcal{S}$, providing the counter system $\mathcal{S}'$ such that $\mathtt{T}_{\mathcal{S}'} = \mathtt{T}_{\mathcal{S}} \cup \mathtt{T}_\phi$.

Let us be a bit more precise in the construction of $\mathcal{A}_\phi$. We write A to denote the set of atomic formulae of the form either q, or $\mathtt{t} \sim k$ or $\mathtt{t} \equiv_c k'$ occurring in ϕ, as well as their negations. Similarly, we write $cl(\phi)$ to denote the *closure* of ϕ, defined as the smallest set of formulae closed under subformulae, closed under negations (double negations are eliminated) and containing ϕ. The set of *atoms* for ϕ, written $Atoms(\phi)$, contains the subsets of $cl(\phi)$ that are maximally consistent and such that for every formula $\xi \in A$ then either ξ or $\neg\xi$ belongs to the set (but not both). States of $\mathcal{A}_\phi$ are in $Atoms(\phi) \times [0, m]$ where ϕ has m **U**-formulae and its alphabet is a subset of $Q \times \mathcal{P}(A)$ (details can be found in [4] with the standard construction for the synchronized product $\mathcal{S}'$). An instance of RBMC can be reduced to several instances of RB-REP-REACH(QFP) with the synchronized product $\mathcal{S}'$. In particular, RMBC can be solved by checking a finite number of instances of RB-REP-REACH(QFP) depending which initial states and accepting states are considered.

Lemma 1. *Let $\mathcal{S} = (Q, n, \delta) \in$ CS(QFP), $(q, \boldsymbol{x})$, $\phi \in$ CLTL(QFP) and $r \in \mathbb{N}$ be an instance of RBMC and $\mathcal{S}'$ be the counter system in* CS(QFP) *obtained by synchronizing $\mathcal{S}$ with $\mathcal{A}_\phi$. The propositions below are equivalent:* (I) *there is an infinite r-($\mathtt{T}_{\mathcal{S}} \cup \mathtt{T}_\phi$)-reversal-bounded run ρ of $\mathcal{S}$ from $(q, \boldsymbol{x})$ such that $\rho, 0 \models \phi$;* (II) *there is an infinite r-$\mathtt{T}_{\mathcal{S}'}$-reversal-bounded run from $((q, X_0, 0), \boldsymbol{x})$ such that $(q_f, X_f, 0)$ is repeated infinitely often for some* initial *atom $X_0 \in Atoms(\phi)$ and for some $(q_f, X_f) \in Q \times Atoms(\phi)$.*

Actually, thanks to the previous lemma, the following corollary holds:

Corollary 1. *There is a polynomial-space reduction from* RMBC *into* RB-REP-REACH(CS(QFP)).

The next section is devoted to show how to reduce RB-REP-REACH(QFP($<$, $\equiv$)) to RB-REP-REACH(QFP($<$)).

Removing periodicity constraints. In this section, we show that given $\mathrm{L} \subseteq \mathrm{QFP}$ using periodicity constraints of the form $\mathtt{t} \equiv_c k$, the reversal-bounded reachability problem for counter systems in CS(L) can be reduced to the corresponding problem restricted to counter systems in CS(L$'$), where L$'$ is the restriction of L without periodicity constraints.

Reduction. Let us consider the class of counter systems CS(L). The underlying idea to remove periodicity constraints consists in defining a new counter system $\mathcal{S}' \in \mathrm{CS}(\mathrm{L}')$ from a given $\mathcal{S} \in \mathrm{CS}(\mathrm{L})$, whose control states store counter values modulo C, where C is the lcm of all the constants c appearing in atomic formulae of the form $\mathtt{t} \equiv_c k$ in guards of $\mathcal{S}$ (see [4] for standard justifications about the value C). The number of control states in $\mathcal{S}'$ is equal to number of control states in $\mathcal{S}$ multiplied by C, which is in $\mathcal{O}(2^{N^2})$ (N is the size of $\mathcal{S}$ with some reasonably succinct encoding). This construction entails an exponential blow-up of the number of control states of the new counter system $\mathcal{S}'$. The transitions of $\mathcal{S}'$ are defined accordingly to the update operations on them in order to correctly represent the classes of modulo for each counter. Let $\mathcal{S}' = (Q', n, \delta')$ be the counter system where $Q' = Q \times [0, C-1]^n$. Given $\boldsymbol{x} \in \mathbb{N}^n$, we write $\tilde{\boldsymbol{x}}$ to denote the unique tuple in $[0, C-1]^n$ such that for $i \in [1, n]$, we have $\boldsymbol{x}(i) \equiv_C \tilde{\boldsymbol{x}}(i)$. Let config_{ok} be the set of configurations for $\mathcal{S}'$ of the form $((q, \tilde{\boldsymbol{x}}), \boldsymbol{y})$ such that $\tilde{\boldsymbol{y}} = \tilde{\boldsymbol{x}}$. Let $f : (Q \times \mathbb{N}^n) \to \mathrm{config}_{ok}$ be the one-to-one map such that $f((q, \boldsymbol{x})) = ((q, \tilde{\boldsymbol{x}}), \boldsymbol{x})$. f and f^{-1} extend naturally to sequences (either finite or infinite ones). The transition relation δ' is defined as follows: if $q \xrightarrow{(\phi, \boldsymbol{b})} q' \in \delta$ then $(q, \tilde{\boldsymbol{x}}) \xrightarrow{(\phi', \boldsymbol{b})} (q', \tilde{\boldsymbol{y}}) \in \delta'$ for all tuples $\tilde{\boldsymbol{x}}$, $\tilde{\boldsymbol{y}}$, where ϕ' is defined from ϕ by substituting $\top$ in place of each $\mathtt{t} \equiv_c k$, with $\mathtt{t} = \sum_j a_j \mathtt{x}_j$, if $\sum_j a_j \tilde{\boldsymbol{x}}(j) \equiv_c k$; otherwise $\bot$. Moreover, we require that for $i \in [1, n]$, we have $\tilde{\boldsymbol{y}}(i) \equiv_C \tilde{\boldsymbol{x}}(i) + \boldsymbol{b}(i)$.

Lemma 2. *Let $\mathcal{S} = (Q, n, \delta)$ be in* CS(QFP) *and $\mathcal{S}' = (Q', n, \delta')$ be the counter system in* CS(QFP($<$)) *defined as above.* (I) *For every run ρ of $\mathcal{S}$, $f(\rho)$ is also a run of $\mathcal{S}'$.* (II) *For every run ρ of $\mathcal{S}'$ such that the first configuration belongs to* config_{ok}*, then all configurations in ρ belong to* config_{ok} *and $f^{-1}(\rho)$ is also a run of $\mathcal{S}$.*

From the previous result, the following corollary can be drawn.

Corollary 2. *Let* $\mathrm{L} = \mathrm{QFP}$ *[resp.* $\mathrm{L} = \mathrm{QFP}(<_1, \equiv)$*] and* $\mathrm{L}' = \mathrm{QFP}(<)$ *[resp.* $\mathrm{L}' = \mathrm{QFP}(<_1)$*].*

(I) *There is a polynomial-space reduction from* RB-REACH(CS(L)) *to* RB-REACH(CS(L$'$)). (II) *There is a polynomial-space reduction from* RB-REP-REACH(CS(L)) *to* RB-REP-REACH(CS(L$'$)).

Elimination of Büchi acceptance conditions. Let $\mathcal{S}$ be in CS(QFP($<$)), $(q_0, \boldsymbol{x_0})$ be an initial configuration, q_f be a control state and $r \geq 0$. We write $K_{min} \in \mathbb{Z}$ [resp. $K_{max} \in \mathbb{Z}$] to denote the minimal [resp. maximal] k occurring in atomic formulae of the form $\mathtt{t} \sim k$ in guards from $\mathcal{S}$. We show below how the existence of an *infinite* run can be characterized by the existence of a *finite*

run satisfying additional properties. The properties ($\star$) and ($\star\star$) below witness such an equivalence. This is comparable, but certainly a bit more technically involved, to the existence of infinite accepting runs in Büchi automata that is equivalent to conditions on finite runs. However, such a reduction is not possible with nondeterministic Minsky machines without the reversal-boundedness assumption. Indeed, the recurrence problem for nondeterministic Minsky machines is Σ_1^1-hard [1] whereas the halting problem for nondeterministic Minksy machines is in Σ_1^0. We show that the conditions below are equivalent.

($\star$) There is an infinite r-$\mathtt{T}_{\mathcal{S}}$-reversal-bounded run from $(q_0, \boldsymbol{x_0})$ such that q_f is repeated infinitely often.

($\star\star$) There exist a finite run $(q_0, \boldsymbol{x_0}), \ldots, (q_l, \boldsymbol{x_l})$, $l' < l$, $Z_{\rightarrow} \subseteq [1, n]$ and $\mathtt{T}_{\rightarrow}, \mathtt{T}_{\searrow}, \mathtt{T}_{\nearrow} \subseteq (\mathtt{T}_{\mathcal{S}} \setminus \{\mathtt{x}_1, \ldots, \mathtt{x}_n\})$ such that

1. $q_{l'} = q_l = q_f$ and $(q_0, \boldsymbol{x_0}), \ldots, (q_l, \boldsymbol{x_l})$ is r-$\mathtt{T}_{\mathcal{S}}$-reversal-bounded.
2. For $j \in [l'+1, l]$ and $i \in Z_{\rightarrow}$, $\boldsymbol{x_j}(i) - \boldsymbol{x_{j-1}}(i) = 0$.
3. For $j \in [l'+1, l]$ and $i \in [1, n] \setminus Z_{\rightarrow}$, $\boldsymbol{x_j}(i) - \boldsymbol{x_{j-1}}(i) \geq 0$.
4. For $i \in [1, n] \setminus Z_{\rightarrow}$, $\boldsymbol{x_{l'}}(i) \geq K_{max}$.
5. $\mathtt{T}_{\rightarrow}, \mathtt{T}_{\searrow}, \mathtt{T}_{\nearrow}$ is a partition of $(\mathtt{T}_{\mathcal{S}} \setminus \{\mathtt{x}_1, \ldots, \mathtt{x}_n\})$.
6. For $j \in [l'+1, l]$ and $\mathtt{t} \in \mathtt{T}_{\rightarrow}$, we have $\boldsymbol{x_j}(\mathtt{t}) - \boldsymbol{x_{j-1}}(\mathtt{t}) = 0$.
7. For $j \in [l'+1, l]$ and $\mathtt{t} \in \mathtt{T}_{\searrow}$, we have $\boldsymbol{x_j}(\mathtt{t}) - \boldsymbol{x_{j-1}}(\mathtt{t}) \leq 0$.
8. For $j \in [l'+1, l]$ and $\mathtt{t} \in \mathtt{T}_{\nearrow}$, $\boldsymbol{x_j}(\mathtt{t}) - \boldsymbol{x_{j-1}}(\mathtt{t}) \geq 0$.
9. For $\mathtt{t} \in \mathtt{T}_{\searrow}$, $\boldsymbol{x_{l'}}(\mathtt{t}) \leq K_{min}$; 10. For $\mathtt{t} \in \mathtt{T}_{\nearrow}$, $\boldsymbol{x_{l'}}(\mathtt{t}) \geq K_{max}$.

Lemma 3. *($\star$) is equivalent to ($\star\star$)*

Proof. ($\star$) implies ($\star\star$). Let $(q_0, \boldsymbol{x_0}), (q_1, \boldsymbol{x_1}), \ldots$ be an infinite r-$\mathtt{T}_{\mathcal{S}}$-reversal-bounded run from $(q_0, \boldsymbol{x_0})$ such that q_f is repeated infinitely often (with $\mathtt{T}_{\mathcal{S}} = \{\mathtt{x}_1, \ldots, \mathtt{x}_n\} \cup \{\mathtt{t}_1, \ldots, \mathtt{t}_{n'}\}$). All the atomic guards in $\mathcal{S}$ are of the form $\mathtt{t} \sim k$ with $\mathtt{t} \in \mathtt{T}_{\mathcal{S}}$ and $k \in [K_{min}, K_{max}]$. Let us make the following observations.

- Let $i \in [1, n]$. Because counter i has a bounded number of reversals, from some position, the value of counter i either remains constant or it is diverging to $+\infty$ and the update values (on counter i) are always greater than 0. Let $Z_{\rightarrow}$ be the subset of $[1, n]$ containing the counters whose values remain constant after some position. In the second case, there is a position j_1 such that for $j \geq j_1$, $\boldsymbol{x_j}(i) \geq K_{max}$, for all $i \in [1, n] \setminus Z_{\rightarrow}$.
- Let $i \in [1, n']$. Because the term $\mathtt{t}_i$ has a bounded number of reversals, one of the conditions below hold true (leading to the definition of $\mathtt{T}_{\rightarrow}, \mathtt{T}_{\searrow}, \mathtt{T}_{\nearrow}$).
  1. From some position, the value of the term $\mathtt{t}_i$ remains constant, i.e. there is $j_0 \in \mathbb{N}$, such that for $j \geq j_0$, $\boldsymbol{x_{j+1}}(\mathtt{t}_i) - \boldsymbol{x_j}(\mathtt{t}_i) = 0$.
  2. The value of the term $\mathtt{t}_i$ diverges to $-\infty$ and there is $j_0 \in \mathbb{N}$, such that for $j \geq j_0$, $\boldsymbol{x_{j+1}}(\mathtt{t}_i) - \boldsymbol{x_j}(\mathtt{t}_i) \leq 0$. In particular, there is a position $j_1 \geq j_0$ such that $\boldsymbol{x_{j_1}}(\mathtt{t}_i) \leq K_{min}$.
  3. The value of the term $\mathtt{t}_i$ diverges to $+\infty$ and there is $j_0 \in \mathbb{N}$, such that for $j \geq j_0$, $\boldsymbol{x_{j+1}}(\mathtt{t}_i) - \boldsymbol{x_j}(\mathtt{t}_i) \geq 0$. In particular, there is a position $j_1 \geq j_0$ such that $\boldsymbol{x_{j_1}}(\mathtt{t}_i) \geq K_{max}$.
- Since q_f is repeated infinitely often, there are two positions $l' < l$ satisfying the conditions (1)–(10).

($\star\star$) implies ($\star$). It remains to show that the existence of a finite run $(q_0, \boldsymbol{x_0})$, $(q_1, \boldsymbol{x_1})$, ..., $(q_l, \boldsymbol{x_l})$, $l' < l$, $Z_{\rightarrow} \subseteq [1, n]$ and $\mathtt{T}_{\rightarrow}, \mathtt{T}_{\searrow}, \mathtt{T}_{\nearrow} \subseteq (\mathtt{T}_{\mathcal{S}} \setminus \{\mathtt{x}_1, \ldots, \mathtt{x}_n\})$ such that (1)-(10) hold true implies that there is an infinite r-$\mathtt{T}_{\mathcal{S}}$-reversal-bounded run from $(q_0, \boldsymbol{x_0})$ such that q_f is repeated infinitely often. Let ρ be the run $(q_0, \boldsymbol{x_0}) \xrightarrow{t_1} (q_1, \boldsymbol{x_1}) \cdots \xrightarrow{t_{l'}} (q_{l'}, \boldsymbol{x_{l'}}) \cdots \xrightarrow{t_l} (q_l, \boldsymbol{x_l})$. For each transition t_i, we assume that the guard is ϕ_i and the update vector is $\boldsymbol{b_i}$. Let us consider the infinite sequence of configurations below

$$\rho' = (q_0, \boldsymbol{x_0}) \xrightarrow{t_1} (q_1, \boldsymbol{x_1}) \cdots \xrightarrow{t_{l'}} (q_{l'}, \boldsymbol{x_{l'}}) \cdots \xrightarrow{t_l} (q_l, \boldsymbol{x_l}) = (q_l, \boldsymbol{y_l}) \xrightarrow{t_{l'+1}} \cdots$$

$$\cdots (q_{l'+1}, \boldsymbol{y_{l'+(l-l')+1}}) \xrightarrow{t_l} (q_l, \boldsymbol{y_{l+(l-l')}}) \cdots$$

such that for $k \geq 0$ and $k' \in [0, l - l' - 1]$, we have $\boldsymbol{y_{l+k(l-l')+k'}} = \boldsymbol{x_{l+k'}} + k(\boldsymbol{x_l} - \boldsymbol{x_{l'}})$ and the sequence of transitions is $t_1 \cdots t_{l'} (t_{l'+1} \cdots t_l)^{\omega}$.

1. Obviously q_f is repeated infinitely often in ρ'.
2. ρ' is indeed a run as for $k \geq 0$ and $k' \in [0, l-l'-1]$, $\boldsymbol{y_{l+k(l-l')+k'}} \models \phi_{l'+1+k'}$ since $\boldsymbol{x_{l+k'}} \models \phi_{l'+1+k'}$ and after position l', atomic guards of the form $\mathtt{t} \sim k$ have a constant truth status. Indeed, $(\boldsymbol{x_l} - \boldsymbol{x_{l'}})$ is constant.
3. ρ' is r-$\mathtt{T}_{\mathcal{S}}$-reversal-bounded since after position l', no new reversal happens.

□

Theorem 1. *There is a polynomial-space many-one reduction from* RB-REP-REACH(CS(QFP(<))) *into* RB-REACH(CS(QFP(<))).

Proof. Let $\mathcal{S}$ be in CS(QFP(<)), $(q_0, \boldsymbol{x_0})$ be an initial configuration, q_f be a control state and $r \geq 0$. We write K_{min} [resp. K_{max}] to denote the minimal [resp. maximal] k occurring in atomic formulae of the form $\mathtt{t} \sim k$ in guards from $\mathcal{S}$. Let us build an instance of RB-REACH(CS(QFP(<))) which captures the condition ($\star\star$). We construct a counter automaton $\mathcal{S}' = (Q', n, \delta')$ such that ($\star\star$) iff there is an $(r+1)$-$\mathtt{T}_{\mathcal{S}'}$-reversal-bounded run from $(q_0, \boldsymbol{x_0})$ to $(q_{new}, \mathbf{0})$. $\mathcal{S}'$ is made of the original version of $\mathcal{S}$ (called below the *original copy*) augmented with copies of $\mathcal{S}$; each copy corresponds to a possible tuple $C = (Z_{\rightarrow}, \mathtt{T}_{\rightarrow}, \mathtt{T}_{\searrow}, \mathtt{T}_{\nearrow})$. By the *$C$-copy*, we mean the copy of $\mathcal{S}$ in which we keep only the transitions with update vector $\boldsymbol{b}$ such that for $i \in Z_{\rightarrow}$, $\boldsymbol{b}(i) = 0$; for $i \notin Z_{\rightarrow}$, $\boldsymbol{b}(i) \geq 0$. for $\mathtt{t} \in \mathtt{T}_{\rightarrow}$, $\boldsymbol{b}(\mathtt{t}) = 0$; for $\mathtt{t} \in \mathtt{T}_{\searrow}$, $\boldsymbol{b}(\mathtt{t}) \leq 0$; for $\mathtt{t} \in \mathtt{T}_{\nearrow}$, $\boldsymbol{b}(\mathtt{t}) \geq 0$.

In order to simulate the subrun $(q_{l'}, \boldsymbol{x_{l'}}) \cdots (q_l, \boldsymbol{x_l})$, from the original copy, nondeterministically we move from the original copy to some C-copy in $\mathcal{S}'$ (and therefore we choose the sets for C) and we test whether the counters in $[1, n] \setminus Z_{\rightarrow}$ have a value greater than K_{max} (with guards $\mathtt{x} \geq K_{max}$), the terms $\mathtt{t}$ in $\mathtt{T}_{\searrow}$ have a value smaller than K_{min} (with guards $\mathtt{t} \leq K_{min}$), the terms $\mathtt{t}$ in $\mathtt{T}_{\nearrow}$ have a value greater than K_{max} (with guards $\mathtt{t} \geq K_{max}$). Finally, in the C-copy, when q_f is reached again, nondeterministically we may jump to the new accepting control state q_{new}. Self-loops on q_{new} allows to decrement any counter. It is also worth noting that $\mathtt{T}_{\mathcal{S}'} = \mathtt{T}_{\mathcal{S}}$; $\mathcal{S}$ and $\mathcal{S}'$ have the same set of constants k occuring in atomic formulae of the form $\mathtt{t} \sim k$; the numbers of states of $\mathcal{S}'$ is bounded by $\mathrm{card}(Q) \times (1 + 2^n \times (2^{n'} \times 2^{n'}))) + 1$ (with $\mathrm{card}(\mathtt{T}_{\mathcal{S}}) = n + n'$). □

Given a counter system $\mathcal{S} = (Q, n, \delta)$ and an infinite run ρ, there exists at least one sequence of transitions $\pi \in \delta^\omega$ such that ρ is built from the successive firing of transitions from π. A sequence π is ultimately periodic if $\pi = \pi_1(\pi_2)^\omega$ for some finite sequences π_1 and π_2. The different reductions established in this section (see also their proofs) allow us to show the result below.

Corollary 3. *Let $\mathcal{S}$ be in* CS(QFP)*, $(q, \boldsymbol{x})$ be a configuration, ϕ be in* CLTL(QFP) *and $r \in \mathbb{N}$.* (I) *and* (II) *are equivalent:* (I) *there is an infinite run ρ from $(q, \boldsymbol{x})$ such that $\rho, 0 \models \phi$ and ρ is r-$\mathtt{T}$-reversal-bounded with $\mathtt{T} = \mathtt{T}_\mathcal{S} \cup \mathtt{T}_\phi$;* (II) *there exists an* ultimately periodic *run ρ satisfying the same properties as in (I).*

4 Complexity and effective Presburger-definability

In this section, we present the following results: RB-REACH(QFP), RB-REP-REACH(QFP) and RMBC are NExpTime-complete and the sets of initial configurations satisfying the properties related to these problems (witness run properties) are effectively definable in Presburger arithmetic, a key result for performing verification practically.

Theorem 2. RB-REACH(CS(QFP(<))) *is* NExpTime*-complete.*

The proof of Theorem 2 is the most involved part of the paper; it is presented in [4]. It generalizes the proof provided for [17, Theorem 3] and uses arguments that can be found also in [30] but in some other context (complexity upper bound for decision problems about Petri nets), see also [11]. It is essential to use the existence of small solutions for integer (inequality) systems [7]. Thanks to Theorem 2, we can improve [19, Theorem 4.4] by establishing that strong reversal-bounded reachability problem is in NExpTime (no complexity bound is provided in the proof of [19, Theorem 4.4]).

As a by-product of the previous result, we can show the following result.

Corollary 4. *Given $\mathcal{S}$ in* CS(QFP)*, $r \geq 0$ and control states q, q', one can effectively compute a Presburger formula $\phi_{q,q'}(\mathtt{x}_1, \ldots, \mathtt{x}_n, \mathtt{y}_1, \ldots, \mathtt{y}_n)$ such that for all valuations* $\mathbf{val}$*,* $\mathbf{val} \models_{\mathrm{PA}} \phi$ *iff there is an r-$\mathtt{T}_\mathcal{S}$-reversal-bounded run from $(q, (\mathbf{val}(\mathtt{x}_1), \ldots, \mathbf{val}(\mathtt{x}_n)))$ to $(q', (\mathbf{val}(\mathtt{y}_1), \ldots, \mathbf{val}(\mathtt{y}_n)))$.*

Consequently, when an initialized counter system is r-reversal-bounded for some $r \geq 0$, then the reachability set is effectively Presburger-definable. This captures the standard case when the counter system belongs to CS(QFP($<_1$)) [20, 24] but Corollary 4 goes much beyond.

Theorem 3. RB-REP-REACH(CS(QFP(<))) *is* NExpTime*-complete.*

Corollary 5. *Given $\mathcal{S}$ in* CS(QFP)*, $r \geq 0$ and control states q, q_f, one can effectively compute a Presburger formula $\phi_{q,q_f}(\mathtt{x}_1, \ldots, \mathtt{x}_n)$ such that for all valuations* $\mathbf{val}$*,* $\mathbf{val} \models_{\mathrm{PA}} \phi$ *iff there is an infinite r-$\mathtt{T}_\mathcal{S}$-reversal-bounded run from $(q, (\mathbf{val}(\mathtt{x}_1), \ldots, \mathbf{val}(\mathtt{x}_n)))$ such that q_f is repeated infinitely often.*

We are now ready to state our main results (Theorem 4 and Theorem 5).

Theorem 4. RBMC *is* NExpTime*-complete.*

As a consequence, we obtain the following results since RB-REACH(QFP) and RB-REP-REACH(QFP) can be reduced in logarithmic space to RBMC.

Corollary 6. RB-REACH(QFP) *and* RB-REP-REACH(QFP) *are* NExpTime*-complete.*

Interestingly, vector addition systems with states (VASS) are elements of CS(QFP($<_1$)) and therefore RBMC(VASS,CLTLQFP($<_1, \equiv$))) is in NExpTime, which contrasts with the ExpSpace-completeness of the model-checking problem with LTL (the only atomic formulae are control states) restricted to VASS [15]. Unlike LTL, CLTL(QFP($<_1, \equiv$)) admits arithmetical constraints.

Theorem 5. *Let* $\mathcal{S}$ *be in* CS(QFP)*,* ϕ *be in* CLTL(QFP) $r \geq 0$ *and* q *be a control state. One can effectively build a Presburger formula* $\phi_q(\mathtt{x}_1, \ldots, \mathtt{x}_n)$ *such that for all* **val***,* **val** $\models_{\text{PA}} \phi_q$ *iff there is an infinite run* ρ *from* $(q, (\mathbf{val}(\mathtt{x}_1), \ldots, \mathbf{val}(\mathtt{x}_n)))$ *such that* $\rho, 0 \models \phi$ *and* ρ *is* r*-*$\mathtt{T}$*-reversal-bounded with* $\mathtt{T} = \mathtt{T}_{\mathcal{S}} \cup \mathtt{T}_{\phi}$*.*

We are also able to improve Corollary 7 since we also have bounds on the lenght of reversal-bounded runs (see the proof of Theorem 4).

Corollary 7. *Let* $\mathcal{S}$ *be in* CS(QFP)*,* $(q, \boldsymbol{x})$ *be an initial configuration,* ϕ *be in* CLTL(QFP) *and* $r \in \mathbb{N}$*. Condition* (I) *in Corollary 3 is equivalent to* (II) *in Corollary 3 with the following additional condition: the sequence of transitions* $\pi_1(\pi_2)^\omega$ *verifies that the length of* $\pi_1\pi_2$ *is bounded by* $2^{2^{p_0(N)}}$*, for some polynom* $p_0(\cdot)$ *and* N *is the size of the instance of* RBMC*.*

Let us explain the benefits of these results from a practical point of view. From Theorem 5, given the formula $\phi_q(\mathtt{x}_1, \ldots, \mathtt{x}_n)$, we can check if an initial configuration verifies the existence of an infinite run satisfying a temporal formula. This can be done with a solver for Presburger arithmetic (tools handling first-order logics with linear arithmetic are for instance LIRA [3], TAPAS [23], CVC3 [2] and Z3 [26]). Hence, Theorem 5 is the final step in our investigations since verification problems are then reduced effectively to satisfiability in Presburger arithmetic. Moreover, our results on the computational complexity guarantee that we are optimal. Another approach arises from Corollary 7 which takes advantage of the method for checking bounded reachability problems as developed in [5]. Since an instance of RBMC can be transformed into an instance of RB-REACH(QFP) and by Theorem 2, one could solve the reversal-bounded model checking problem by looking for finite runs of length at most doubly exponential.

5 Conclusion

We have studied the model-checking problem RBMC over counter systems when runs are reversal-bounded and the specification language is an LTL-like dialect

with arithmetical constraints, past-time and future-time operators. A major result is the NExpTime-completeness of the problem RBMC. Even more importantly, in order to implement decision procedures, we have shown that given a counter system, a temporal formula ϕ and $r \geq 0$, one can build effectively a Presburger formula encoding the set of configurations $(q, \boldsymbol{x})$ such that there is an r-$(\mathtt{T}_\phi \cup \mathtt{T}_\mathcal{S})$-reversal-bounded infinite run ρ from $(q, \boldsymbol{x})$ such that ϕ is satisfied by ρ. Finally, we have also characterized the complexity of several reversal-bounded reachability problems and control state repeated reachability problem (obtaining NExpTime-completeness). It is worth noting that our proofs for NExpTime-easiness are obtained by an explicit run analysis that shortens the runs, as in [16] but in a different way.

Acknowledgment: We would like to thank the anonymous referees for their suggestions and constructive remarks; a special thank is due to the referee that pointed us to [16].

References

1. Alur, R., Henzinger, T.: A really temporal logic. In: FOCS'89. pp. 164–169. IEEE (1989)
2. Barrett, C., Tinelli, C.: CVC3. In: CAV'07. LNCS, vol. 4590, pp. 298–302. Springer (2007)
3. Becker, B., Dax, C., Eisinger, J., Klaedtke, F.: LIRA: Handling Constraints of Linear Arithmetics over the Integers and the Reals. In: CAV'07. LNCS, vol. 4590, pp. 307–310. Springer (2007)
4. Bersani, M., Demri, S.: The complexity of reversal-bounded model checking. Tech. Rep. LSV-11-10, LSV, ENS Cachan, France (May 2011)
5. Bersani, M., Frigeri, A., Morzenti, A., Pradella, M., Rossi, M., San Pietro, P.: Bounded reachability for temporal logic over constraint systems. In: TIME'10. pp. 43–50. IEEE (2010)
6. Biere, A., Cimatti, A., Clarke, E.M., Strichman, O., Zhu, Y.: Bounded model checking. Advances in Computers 58, 118–149 (2003)
7. Borosh, I., Treybig, L.: Bounds on positive integral solutions of linear diophantine equations. Proocedings of The American Mathematical Society 55, 299–304 (1976)
8. Bouajjani, A., Echahed, R., Habermehl, P.: On the verification problem of nonregular properties for nonregular processes. In: LICS'95. pp. 123–133 (1995)
9. Čerāns, K.: Deciding properties of integral relational automata. In: ICALP'94. LNCS, vol. 820, pp. 35–46. Springer (1994)
10. Dang, Z., Ibarra, O., San Pietro, P.: Liveness verification of reversal-bounded multicounter machines with a free counter. In: FST&TCS'01. LNCS, vol. 2245, pp. 132–143. Springer (2001)
11. Demri, S.: On Selective Unboundedness of VASS. In: INFINITY'10. EPTCS, vol. 39, pp. 1–15 (2010)
12. Esparza, J.: Decidability and complexity of Petri net problems — an introduction. In: Advances in Petri Nets. LNCS, vol. 1491, pp. 374–428. Springer, Berlin (1998)
13. Finkel, A., Sangnier, A.: Reversal-bounded counter machines revisited. In: MFCS'08. LNCS, vol. 5162, pp. 323–334. Springer (2008)

14. Gurari, E., Ibarra, O.: The complexity of decision problems for finite-turn multicounter machines. In: ICALP'81. LNCS, vol. 115, pp. 495–505. Springer (1981)
15. Habermehl, P.: On the complexity of the linear-time mu-calculus for Petri nets. In: ICATPN'97. LNCS, vol. 1248, pp. 102–116. Springer (1997)
16. Hague, M., Lin, A.W.: Model checking recursive programs numeric data types. In: CAV'11. LNCS, Springer (2011), to appear
17. Howell, R., Rosier, L.: An analysis of the nonemptiness problem for classes of reversal-bounded multicounter machines. JCSS 34(1), 55–74 (1987)
18. Ibarra, O., Bultan, T., Su, J.: Reachability analysis for some models of infinite-state transition systems. In: CONCUR'00. LNCS, vol. 1877, pp. 183–198. Springer (2000)
19. Ibarra, O., Su, J., Dang, Z., Bultan, T., Kemmerer, R.: Counter Machines and Verification Problems. TCS 289(1), 165–189 (2002)
20. Ibarra, O.H.: Reversal-bounded multicounter machines and their decision problems. JACM 25(1), 116–133 (1978)
21. Kopczynski, E., To, A.: Parikh Images of Grammars: Complexity and Applications. In: LICS'10. pp. 80–89. IEEE (2010)
22. Laroussinie, F., Meyer, A., Petonnet, E.: Counting LTL. In: TIME'10. pp. 51–58. IEEE (2010)
23. Leroux, J., Point, G.: TaPAS: The Talence Presburger Arithmetic Suite. In: TACAS'09. LNCS, vol. 5505, pp. 182–185. Springer (2009)
24. Leroux, J., Sutre, G.: Flat counter systems are everywhere! In: ATVA'05. LNCS, vol. 3707, pp. 489–503. Springer (2005)
25. Lutz, C.: NEXPTIME-complete description logics with concrete domains. ACM ToCL 5(4), 669–705 (2004)
26. de Moura, L., Björner, N.: Z3: An Efficient SMT Solver. In: TACAS'08. LNCS, vol. 4963, pp. 337–340. Springer (2008)
27. Papadimitriou, C.: On the complexity of integer programming. JACM 28(4), 765–768 (1981)
28. Presburger, M.: Über die Vollständigkeit eines gewissen Systems der Arithmetik ganzer Zahlen, in welchem die Addition als einzige Operation hervortritt. In: Comptes Rendus du premier congrès de mathématiciens des Pays Slaves, Warszawa. pp. 92–101 (1930)
29. Qadeer, S., Rehof, J.: Context-bounded model checking of concurrent software. In: TACAS'05. LNCS, vol. 3440, pp. 93–107. Springer (2005)
30. Rackoff, C.: The covering and boundedness problems for vector addition systems. TCS 6(2), 223–231 (1978)
31. Suzuki, N., Jefferson, D.: Verification Decidability of Presburger Array Programs. JACM 27(1), 191–205 (1980)
32. To, A.: Model Checking Infinite-State Systems: Generic and Specific Approaches. Ph.D. thesis, School of Informatics, University of Edinburgh (2010)
33. To, A., Libkin, L.: Algorithmic metatheorems for decidable LTL model checking over infinite systems. In: FOSSACS'10. LNCS, vol. 6014, pp. 221–236. Springer (2010)
34. Vardi, M., Wolper, P.: Reasoning about infinite computations. I&C 115, 1–37 (1994)