



HAL
open science

Invariants for One-Counter Automata with Disequality Tests

Dmitry Chistikov, Jérôme Leroux, Henry Sinclair-Banks, Nicolas Waldburger

► **To cite this version:**

Dmitry Chistikov, Jérôme Leroux, Henry Sinclair-Banks, Nicolas Waldburger. Invariants for One-Counter Automata with Disequality Tests. CONCUR2024 - International Conference on Concurrency Theory, Sep 2024, Calgary, Canada. pp.1-34, 10.4230/LIPICS.CONCUR.2024.17 . hal-04707413

HAL Id: hal-04707413

<https://hal.science/hal-04707413v1>

Submitted on 24 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Invariants for One-Counter Automata with Disequality Tests

Dmitry Chistikov¹ ✉ 

Centre for Discrete Mathematics and its Applications (DIMAP) &
Department of Computer Science, University of Warwick, Coventry, UK

Jérôme Leroux ✉

LaBRI, CNRS, Univ. Bordeaux, France

Henry Sinclair-Banks ✉  

Centre for Discrete Mathematics and its Applications (DIMAP) &
Department of Computer Science, University of Warwick, Coventry, UK

Nicolas Waldburger ✉ 

IRISA, Université de Rennes, France

Abstract

We study the reachability problem for one-counter automata in which transitions can carry disequality tests. A disequality test is a guard that prohibits a specified counter value. This reachability problem has been known to be NP-hard and in PSPACE, and characterising its computational complexity has been left as a challenging open question by Almagor, Cohen, Pérez, Shirmohammadi, and Worrell [1]. We reduce the complexity gap, placing the problem into the second level of the polynomial hierarchy, namely into the class coNP^{NP} . In the presence of both equality and disequality tests, our upper bound is at the third level, $\text{P}^{\text{NP}^{\text{NP}}}$.

To prove this result, we show that non-reachability can be witnessed by a pair of invariants (forward and backward). These invariants are almost inductive. They aim to over-approximate only a “core” of the reachability set instead of the entire set. The invariants are also leaky: it is possible to escape the set. We complement this with separate checks as the leaks can only occur in a controlled way.

2012 ACM Subject Classification Theory of computation → Models of computation

Keywords and phrases Inductive invariant, Vector addition system, One-counter automaton

Funding *Dmitry Chistikov*: Supported in part by the Engineering and Physical Sciences Research Council [EP/X03027X/1].

Henry Sinclair-Banks: Supported by EPSRC Standard Research Studentship (DTP), grant number EP/T51794X/1. Also supported in part by the International Emerging Actions grant (IEA’22) and by the ANR grant VeSyAM (ANR-22-CE48-0005).

Nicolas Waldburger: Supported in part by the International Emerging Actions grant (IEA’22) and by the ANR grant VeSyAM (ANR-22-CE48-0005).

Acknowledgements We would like to thank Mahsa Shirmohammadi, without whom this work would not have been possible, for many ideas and encouragement. We are also grateful to Thomas Colcombet, Rayna Dimitrova, and James Worrell for useful discussions.

1 Introduction

It is well known that the computational complexity of problems is often sensitive to seemingly minor adjustments in the problem setting. Consider, for example, vector addition systems

¹ During the work on this paper, DC was a visitor to the Max Planck Institute for Software Systems (MPI-SWS), Kaiserslautern and Saarbrücken, Germany, a visiting fellow at St Catherine’s College and a visitor to the Department of Computer Science at the University of Oxford, United Kingdom.

with states (VASS). Perhaps more commonly presented as Petri nets, VASS are a very simple yet powerful model of concurrency. Many important computational problems in logic, language theory, and formal verification reduce to or are even equivalent to the reachability problem in VASS (see, e.g., [29, 46]). However, a classical result due to Minsky shows that adding the capability to test counters for zero makes the problem undecidable [43].

At the same time, while reachability in VASS is known to be decidable [41], its computational complexity was recently shown to be extremely high, namely Ackermann-complete (see [39] for the upper bound and [19, 20, 38] for the lower bound), so from the practical point of view one might question the significance of the complexity jump arising from zero tests.

More recent, “down-to-earth”, and perhaps more striking is the following result on 1-dimensional VASS, which can be thought of as finite-state automata equipped with one counter (capable of storing a nonnegative integer). Reachability in these systems can be decided in NP [28] and is in fact NP-complete. It is not difficult to show, using the standard hill-cutting technique [50], that reachability can also always be witnessed by executions in which all values assumed by the counter are bounded from above by an exponential function in the size of the system and the bit length of counter values of the source and target configurations. Because of this, one might expect that placing an exponential bound on the counter values *upfront* does not change the problem much. But, in fact, the complexity jumps: the problem — which is equivalent to reachability in two-clock timed automata [15] — becomes PSPACE-complete [23]. One may say that, in this case, formal verification toolkit available for the reachability problem is not robust to this change in the problem setting.

In this paper we study a different seemingly benign variation of the standard reachability question. Consider one-counter automata in which transitions may test the value of the counter for *disequality* against a given integer (which depends on the transition). In other words, executions of the system can be blocked by disequality guards, which *prevent* the transition from being fired if the counter value is equal to a specified number. The initial motivation for studying this question comes from a model checking problem for flat Freeze LTL; see Demri and Sangnier [21] and Lechner, Mayr, Ouaknine, Pouly, and Worrell [35]. Additionally, recall that automata can be used for the modeling of imperative code; see, e.g., Hague and Lin [30, 31], as well as discussion in Section 2. Classical Minsky machines encode *if-then* conditionals with equality comparisons to constants, $x = k$. Simulating an *if-then-else* conditional of this type on a Minsky machine seems to require additional $O(k)$ states. If k is large, this growth in size may be exponential, even though *then* branches as well as increments $x += k$ and decrements $x -= k$ can be encoded directly. (If the machine model is extended with $x \leq k$ comparisons, then the asymmetry between *then* and *else* disappears, but reachability becomes PSPACE-complete [23].)

One might expect that, since only a small number of configurations are forbidden (by the disequality guards) in the infinitely large configuration space, the complexity of the problem should not change significantly and existing techniques should be applicable. This conclusion, however, has remained elusive. For the problem of reachability in one-counter automata with disequality tests, the exponential upper bound on counter values necessary for witnessing reachability carries over. But, despite progress on related problems [9, 11] (including the settling of the complexity of the above-mentioned flat Freeze LTL model checking problem [11]), it has not been possible to pin down the complexity of this problem, which has been known to be NP-hard (even without disequality tests) and to belong to PSPACE (thanks to the exponential bound on the counter) [35]. The apparent simplicity of the problem contrasts with the lack of robustness of the available toolbox. It was recently shown by Almagor, Cohen, Pérez, Shirmohammadi, and Worrell [1] that the coverability (or

state reachability) problem can be solved in polynomial time for this model, similarly to the standard 1-dimensional VASS without tests. The algorithm and its analysis, however, become sophisticated, and the complexity of reachability was left as a *challenging open problem*.

In the present paper, we make progress on this problem. Existing techniques need to be extended and developed significantly to handle seemingly benign disequality tests. We have been unable to find an easily verifiable witness for reachability, and instead show that *non-reachability* is witnessed by a form of invariants (or, more precisely, separators). The existence of counterexamples that violate such invariants can be checked in NP, thus placing the reachability problem for OCA with disequality tests into the second level of the polynomial hierarchy, namely in coNP^{NP} . This complexity class captures the complement of synthesis-type questions, which ask to find a single object (say, a circuit) that works correctly for all (exponentially many) inputs. In our problem, one configuration can reach another if and only if every potential invariant (of a form we describe) violates one of the invariance conditions; moreover, this violation can be checked in NP. In the presence of both equality and disequality tests, we need a slightly larger class $\text{P}^{\text{NP}^{\text{NP}}}$, at the third level of the hierarchy.

Our contributions. Traditionally, an invariant is an overapproximation of the set of reachable configurations which is inductive, i.e., closed under the transition relation. Our invariants are different in several ways:

1. We capture only some configurations within the reachability set, which form its *core*. Accordingly, we require a tailored notion of closure, namely closure under a restricted form of reachability relation.
2. Our invariants are leaky (*almost* inductive): an execution may escape the set. Allowing leaks is complemented by a separate controlling mechanism (check) that all leaks — which may occur at the interface between strongly connected components of the automaton — are safe.
3. To compose our local inductive invariants, i.e., those restricted to a single strongly connected component, the controlling mechanism for leaks relies on relaxed integer semantics for the execution. More precisely, we extend (to automata with disequality tests) a known technique [28] for lifting \mathbb{Z} -executions to actual executions.

Our notion of local invariants requires that we place a certain technical assumption at the interface (entry and exit points) of strongly connected components. To discharge this assumption, we use a combination of two invariants, one for the main (forward) VASS and another for the reverse VASS. Together, these two sets form a separator — a witness for non-reachability.

2 Related Work

Invariants. In formal verification, a forward exploration of countably infinite configuration spaces from the initial configuration, or a symmetrical backward exploration from the target, is a standard approach to reachability problems and targets bug finding. General heuristics can be used to improve such an exploration (see, for instance, the recent directed reachability algorithm [8]). However, in order to prove non-reachability, thus certifying the absence of bugs, an invariant-based approach is more popular.

Many techniques have been developed in the past for computing inductive invariants, depending on the structure of the underlying system based on counterexample-guided abstraction refinement [17], automata [32], property-directed reachability [3, 13, 14], and more generally in the abstract interpretation framework [18].

In vector addition systems, semilinear invariants [25] are sufficient for the general reachability problem [36]. Even if those invariants are intractable in general, for some instances, namely the control-state reachability problem, the implementation of efficient tools computing invariants (downward-closed sets in that case) is an active research area [7, 22] with implementation of tools [6, 24, 33].

In this paper we focus on 1-dimensional VASS in the presence of equality and disequality tests; we call them *one-counter automata (OCA) with tests*. The notion of local inductive invariants with leaks, which we propose, provides a way to reduce the search space of inductive invariants, by specifying the shape of the ‘core’ of the invariant (a union of arithmetic progressions within ‘bounded chains’), as well as restricting the problem to each strongly connected component one by one. We view this as a compositional approach for computing inductive invariants. As a theoretical application, we prove that the reachability problem for one-counter automata with tests is between NP and $P^{NP^{NP}}$, and in fact in $coNP^{NP}$ if only disequality tests are present.

The previous work on OCA with disequality tests by Almagor et al. [1] enables us to focus (subject to a technical assumption) on configurations in a small number of bounded chains (see Section 4). The structure of the set of reachable configurations in these chains admits a short description. At the core of our invariants are exactly such sets, and we need an appropriate notion of “inductiveness”, a condition to control “leaks” that violate the assumption above, and a verification mechanism for all these conditions.

One-counter automata. OCA can be seen as an abstraction of pushdown automata, a widely used model of recursive systems. Conceptually very simple, OCA are at the heart of a number of results in formal verification; see, e.g., [2, 4, 34]. Multi-counter automata are used to model imperative code with numerical data types [30, 31]; roughly speaking, a reachability query is expressed in logic, as a formula in existential linear integer arithmetic. In these two references an additional pushdown stack is available, capturing recursive function calls. We refer the reader to [16] for a retrospective on underlying ‘pumping’ results for OCA, crucial for many of the recent results. There is also a rich history of research on behavioural equivalences and model checking for a variety of one-counter processes and systems; see, e.g., [10, 27, 48, 50].

The above-mentioned result that reachability in OCA is NP-complete, by Haase, Kreutzer, Ouaknine, and Worrell [28], has recently been built upon to give a representation of the entire reachability relation in existential linear integer arithmetic, with an implementation available online, by Li, Chen, Wu, and Xia [40]. The idea of ‘lifting’ candidate runs to actual runs, which is shown in [28] and which we develop further by adding support for disequality tests, has been used in other settings as well [5, 37, 41]. For example, a construction similar to our Lemma 7.1 is an element of the proof of a tight upper bound on the length of shortest runs in OCA without disequality tests [16]. In comparison to the latter paper, our construction need not consider divisibility properties of run lengths, but at the same time applies in a more general scenario: the updates of our OCA are specified in binary notation (that is, succinctly); and, naturally, our OCA may have disequality tests.

We already mentioned above that, despite appearing atypical at first glance, the disequality tests do in fact contribute to the modeling power: namely, when modeling code, these tests enable the simulation of the *else* branch in conditional statements comparing an integer variable for equality with some constant. The framework of Hague and Lin [30, 31] assumes that each counter variable can undergo at most k reversals (i.e., changes between ‘increasing’ and ‘decreasing’), where k is fixed. This assumption is strong; without it, a reachability

instance would require a logical formula of exponential size. Results of Haase et al. [28] and Li et al. [40] avoid this assumption, but, for the standard syntax of one-counter automata, *if-then-else* conditionals remain out of reach — or rather require an exponential expansion of the automaton. Our leaky invariants technique allows us to handle such conditionals with equality tests on counters, without assuming any bound on the number of reversals.

3 OCA with Equality and Disequality Tests

We denote by \mathbb{Z} and \mathbb{N} the set of all integers and all nonnegative integers, respectively.

A *constraint* is either an *equality test* of the form $x = k$ with $k \in \mathbb{N}$, a *disequality test* of the form $x \neq k$ with $k \in \mathbb{N}$, or simply **true**; x denotes here our counter, which is a nonnegative integer variable. Let \mathcal{C} denote the set of all possible constraints. A *one-counter automaton (OCA) with equality and disequality tests* is a triplet $\mathcal{A} = (Q, \Delta, \tau)$, where Q is a finite set of *states*, $\Delta \subseteq Q \times \mathbb{Z} \times Q$ is a finite set of *transitions* and $\tau: Q \rightarrow \mathcal{C}$ is the *constraint function*. The automaton \mathcal{A} is an *OCA with disequality tests* if the constraint function τ does not have any equality tests. We sometimes refer to the constraints as *guards*.

Syntactically, \mathcal{A} can be seen as an integer-weighted graph with directed edges between states. Viewed this way, \mathcal{A} can be decomposed into a set of *strongly connected components (SCCs)*. The automaton \mathcal{A} is *strongly connected* when it has one strongly connected component only. A *path* π in \mathcal{A} is a sequence $\pi = (t_1, t_2, \dots, t_n)$ of transitions, where $t_i = (q_{i-1}, a_i, q_i)$ for each i and $n \geq 0$. We may refer to π as a q_0 - q_n path. The *length* of such a path is $\text{len}(\pi) \stackrel{\text{def}}{=} n$. The *effect* of π is $\text{eff}(\pi) \stackrel{\text{def}}{=} \sum_{i=1}^n a_i$. A *cycle* is a path starting and ending at the same state; for $q \in Q$, a *q-cycle* is a q - q path. A path or cycle is *simple* if it contains no repetition of states, except that a *simple cycle* has the same starting and ending state. Every simple cycle has length less than or equal to $|Q|$.

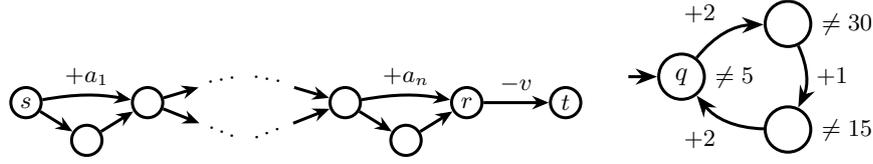
The size of an OCA \mathcal{A} is the bit size of its encoding, where all numbers are written in binary; we denote it by $|\mathcal{A}|$. We write $\|\Delta\|$ and $\|\tau\|$ to refer to the maximum absolute value of a transition update and test, respectively.

Configurations and runs. The semantics of \mathcal{A} is defined based on the set of valid configurations and the reachability relation, as follows.

A *configuration* is a pair (q, z) comprising a state $q \in Q$ and a nonnegative integer $z \in \mathbb{N}$; we may refer to z as the *counter value*. We say that (q, z) is a *valid configuration* if it respects the constraint $\tau(q)$. Write $\text{Conf} \stackrel{\text{def}}{=} Q \times \mathbb{N}$ for the set of all configurations. Given two configurations $(q, z), (q', z')$ and $t \in \Delta$, we write $(q, z) \xrightarrow{t} (q', z')$ when $t = (q, z' - z, q')$; we denote by $(q, z) \rightarrow (q', z')$ the existence of such a transition. A *run* of \mathcal{A} is a sequence $(q_0, z_0), \dots, (q_n, z_n)$ of valid configurations, for $n \geq 0$, such that there exists a path (t_1, \dots, t_n) with $(q_{i-1}, z_{i-1}) \xrightarrow{t_i} (q_i, z_i)$. We say that (q_n, z_n) is *reachable* from (q_0, z_0) if there exists a run from (q_0, z_0) to (q_n, z_n) . We write $(q_0, z_0) \xrightarrow{*} (q_n, z_n)$ to denote the existence of such a run. Given a path π , we write $(q_0, z_0) \xrightarrow{\pi} (q_n, z_n)$ if π *yields* a run from (q_0, z_0) to (q_n, z_n) .

A path π has no hope to yield a run from (q, z) if $z + \text{eff}(\pi') < 0$ for some prefix π' of π . We denote by $\text{drop}(\pi)$ the maximum of $-\text{eff}(\pi')$ over all prefixes π' of π , and call it the *drop* of π . Intuitively, $\text{drop}(\pi)$ is the smallest counter value $z \in \mathbb{N}$ such that π , when applied from (q, z) , remains nonnegative; note that hitting a guard is not a consideration here.

We use the following standard operators: $\text{Post}(c) \stackrel{\text{def}}{=} \{c' \in \text{Conf} \mid c \rightarrow c'\}$ and $\text{Pre}(c) \stackrel{\text{def}}{=} \{c' \in \text{Conf} \mid c' \rightarrow c\}$. For $X \subseteq \text{Conf}$, we write $\text{Post}(X) \stackrel{\text{def}}{=} \bigcup_{c \in X} \text{Post}(c)$ and $\text{Pre}(X) \stackrel{\text{def}}{=} \bigcup_{c \in X} \text{Pre}(c)$.



■ **Figure 1**

Left. This OCA without tests is constructed from an instance of the subset sum problem (a_1, \dots, a_n, v) ; this is in fact how the reachability problem in OCA without tests is proved to be NP-hard in [28]. Configuration $(t, 0)$ can be reached from configuration $(s, 0)$ whenever there exists a subset of $\{a_1, \dots, a_n\}$ whose elements sum up to v . Note that all unlabelled transitions have update zero. The set of configurations reachable from $(s, 0)$ can have size exponential in n , and its structure is unwieldy.

Right. (For Section 4.) The named state q belongs to Q_+ since there is a simple q -cycle with positive effect. There are six bounded chains of configurations at q . The disequality test $\neq 5$ bounds the counter values with residue 0 modulo 5, so $\{(q, 0)\}$ and $\{(q, 5)\}$ are bounded chains. The disequality test $\neq 30$ bounds the counter values with residue 3 modulo 5, so $\{(q, 3), (q, 8), (q, 13), (q, 18), (q, 23), (q, 28)\}$ is a bounded chain. The disequality test $\neq 15$ bounds the counter values with residue 2 modulo 5, so $\{(q, 2), (q, 7), (q, 12)\}$ is a bounded chain.

$\bigcup_{c \in X} \text{Pre}(c)$. Also, $\text{Post}^*(X) \stackrel{\text{def}}{=} \{d \mid \exists c \in X : c \xrightarrow{*} d\}$ and $\text{Pre}^*(X) \stackrel{\text{def}}{=} \{c \mid \exists d \in X : c \xrightarrow{*} d\}$.

For an OCA $\mathcal{A} = (Q, \Delta, \tau)$, we define the *reverse of \mathcal{A}* as $\mathcal{A}^R \stackrel{\text{def}}{=} (Q, \Delta^R, \tau)$ where $(q, a, q') \in \Delta^R$ if and only if $(q', -a, q) \in \Delta$. Given configurations c and d and a path π in \mathcal{A} , we have $c \xrightarrow{*} d$ in \mathcal{A} if and only if $d \xrightarrow{*} c$ in \mathcal{A}^R .

The reachability problem. We consider the following decision problem.

REACHABILITY

Input: An OCA \mathcal{A} with equality and disequality tests, a valid initial configuration `src`, and a valid target configuration `trg`.

Output: Does `src` $\xrightarrow{*}$ `trg` hold?

The model of OCA with disequality tests has been studied in [35] and [1]. The latter paper provides polynomial-time algorithms for the *coverability problem*: “given `src` and a state q , does there exist z such that `src` $\xrightarrow{*}$ (q, z) ?” and the related *unboundedness problem*: “is the set of configurations reachable from `src` infinite?”. The reachability problem, however, is NP-hard even without tests, see Figure 1 (Left).

Equality tests. In the reachability problem in OCA with equality and disequality tests, the main technical challenge stems from disequality tests. Indeed, a state with an equality test only has one valid configuration hence need not be visited more than once.

We now work with OCA with disequality tests only. We will discuss in Section 7.5 how our techniques are affected by the addition of equality tests.

4 Getting Familiar with Disequality Tests

In this section, we fix an OCA $\mathcal{A} = (Q, \Delta, \tau)$ with disequality tests and two valid configurations `src` and `trg`.

A configuration c is *bounded* when $\text{Post}^*(c)$ is finite, and *unbounded* otherwise. It is known, although far from trivial, that one can decide boundedness in polynomial time.

► **Lemma 4.1** (see [1, Theorem 19]). *Given an OCA with disequality tests and a configuration c , it is decidable in polynomial time whether c is bounded or unbounded.*

A *candidate run* is a run except that neither the nonnegativity condition nor the disequality tests are necessarily respected. Formally, a candidate run is simply a sequence $(q_0, z_0), \dots, (q_n, z_n)$ where all $(q_i, z_i) \in Q \times \mathbb{Z}$, $n \geq 0$, and such that there exist transitions $(q_{i-1}, a_i, q_i) \in \Delta$ with $z_{i+1} = z_i + a_i$ for all $i \in \{1, \dots, n\}$. We write $(q_0, z_0) \xrightarrow[\mathbb{Z}]{*} (q_n, z_n)$.

An ingredient of the NP upper bound for reachability in OCA *without* disequality tests [28] is establishing conditions under which a candidate run can be *lifted* to a run. We adapt the argument to OCA with disequality tests.

► **Lemma 4.2.** *Let \mathcal{A} be a strongly connected OCA with disequality tests. If src is unbounded in \mathcal{A} and trg is unbounded in \mathcal{A}^R , then there is a run from src to trg in \mathcal{A} ($\text{src} \xrightarrow{*} \text{trg}$) if and only if there is a candidate run from src to trg in \mathcal{A} ($\text{src} \xrightarrow[\mathbb{Z}]{*} \text{trg}$).*

The hypothesis that \mathcal{A} is strongly connected is crucial. Indeed, if \mathcal{A} is strongly connected and $\text{src} = (s, v)$ is unbounded in \mathcal{A} , then there is a cycle of positive effect that, from src , can be applied infinitely often to reach (s, z) with z arbitrarily large. If \mathcal{A} is not strongly connected, it could be that the positive cycles that make src unbounded are in another SCC and that the set $\{z \mid (s, v) \xrightarrow{*} (s, z)\}$ is finite. See Appendix A for a proof of Lemma 4.2.

Let $Q_+ \subseteq Q$ be the set of states $q \in Q$ such that there exists a q -cycle γ with $\text{len}(\gamma) \leq |Q|$ and with $\text{eff}(\gamma) > 0$. For each $q \in Q_+$, let γ_q be such a q -cycle with minimal drop. **We fix this choice for the remainder of the paper.** We define

$$\text{Conf}_+ \stackrel{\text{def}}{=} \{(q, z) : q \in Q_+ \text{ and } z \geq \text{drop}(\gamma_q)\}.$$

► **Lemma 4.3.** *There is a polynomial-time algorithm to identify Q_+ and to choose cycles γ_q for all $q \in Q_+$. Moreover, membership in Conf_+ can be decided in polynomial time.*

► **Remark 4.4.** Our choice of Q_+ differs slightly from the definition found in [1]: we use short cycles ($\text{len}(\gamma) \leq |Q|$) rather than simple cycles. For simple cycles, the ability to compute, in polynomial time, the minimal drop of a positive-effect simple q -cycle (for each $q \in Q$) is not justified in [1]. In fact, in Appendix A.1, we prove that deciding, for a given OCA without tests \mathcal{A} and a given state q , whether there exists a positive-effect simple q -cycle in \mathcal{A} is an NP-complete problem. However, all constructions and arguments of [1] appear to be insensitive to the replacement of ‘simple cycles’ by ‘short cycles’. As a result, we can still use polynomial-time algorithms for coverability and for unboundedness in OCA with disequality tests (1-VASS with disequality tests).

The set of all $(q, z) \in \text{Conf}_+$ can be partitioned into q -chains. For each $q \in Q_+$, let $\text{Conf}_+(q) = (\{q\} \times \mathbb{N}) \cap \text{Conf}_+$. A q -chain C is a maximal non-empty subset $C \subseteq \text{Conf}_+(q)$ such that, for every two distinct $c, c' \in C$, either c is reachable from c' by iterating γ_q , or vice versa. In other words, C is a non-empty minimal subset of $\text{Conf}_+(q)$ (with respect to set inclusion) such that, for all $c \in C$ and all $c' \in \text{Conf}_+$, if $c \xrightarrow{\gamma_q} c'$ or $c' \xrightarrow{\gamma_q} c$ then $c' \in C$.

A q -chain is *bounded* if it is a finite set, otherwise it is *unbounded*. Note that configurations in unbounded chains are all themselves unbounded, but configurations in bounded chains need not be bounded (they may be unbounded). Because the number of disequality guards that a cycle γ_q may encounter is small, so is the total number of bounded chains.

► **Lemma 4.5** (see [1, Remark 6]). *There are at most $2|Q|^2$ bounded chains.*

Given a chain C , the counter values z of every $(q, z) \in C$ have the same remainder modulo $\text{eff}(\gamma_q)$. Henceforth, a bounded q -chain can be described as $[\ell, u] \cap (r + \text{eff}(\gamma_q) \cdot \mathbb{N})$ where $[\ell, u]$ is an interval of nonnegative integers and $r + \text{eff}(\gamma_q) \cdot \mathbb{N}$ is an arithmetic progression with initial term r and difference $\text{eff}(\gamma_q)$. Since the OCA \mathcal{A} is encoded in binary, the values of l, u, r , and $|\gamma_q|$ may be exponential in the size of \mathcal{A} . See Figure 1 (Left) for an example.

5 Pessimistic Reachability

In this section, we exhibit a family run of runs, namely pessimistic runs, that are guaranteed to admit an NP certificate. This will already enable us to prove, in Section 6, that the reachability problem is in coNP^{NP} in the special case where the OCA is strongly connected.

Let \mathcal{A} be an OCA with disequality tests. We call a run of \mathcal{A} *pessimistic* if none of its configurations are in Conf_+ , except possibly the first one. Of course, some pessimistic runs may be exponentially long relative to the size of \mathcal{A} ; however, we provide a way to handle them. For $S \subseteq \text{Conf}_+$, we write $\text{Post}_*(S)$ for the set of configurations reachable from S using only pessimistic runs. In particular, $S \subseteq \text{Post}_*(S)$.

Consider the following decision problem:

PESSIMISTIC REACHABILITY

Input: An OCA \mathcal{A} with disequality tests, and two configurations `src` and `trg`.

Output: Is there a pessimistic run from `src` to `trg` in \mathcal{A} ?

Pessimistic runs turn out to be very handy, not least because we can adapt an existing ‘flow’ technique [28] to decide pessimistic reachability.

► **Lemma 5.1.** *The pessimistic reachability problem is in NP.*

In a nutshell, the idea [28] is to guess how many times the run traverses each transition. The guessed numbers are subject to polynomial-time checkable balance and connectivity conditions, akin to, e.g., [47]. However, we cannot check whether the (possibly very long) run constructed from the flow violates disequality constraints, so the technique cannot be applied directly.

Our solution (detailed in Appendix B.3) uses the pessimism of the run. Let $x \neq g$ be a guard on state q . We split the run in two: in the first part, all visits to q are above g ; then the run *jumps* the guard so that, in the second part, all visits are below g . This way, with at most $|Q|$ splits, we can reduce the problem to the case in which the run does not jump *any* disequality guard (always staying above or below each of them).

6 Reachability in Strongly Connected OCA

In this section, for pedagogical purposes, we study the particular case where the OCA is strongly connected. The case with multiple SCCs presented in Section 7 is more technical but relies on the same key idea.

► **Theorem 6.1.** *The reachability problem for strongly connected OCA with disequality tests belongs to the complexity class coNP^{NP} .*

We sketch the proof of Theorem 6.1 below. Throughout the section, we fix a strongly connected OCA with disequality tests \mathcal{A} and two configurations `src` and `trg`, and we are interested in whether $\text{src} \xrightarrow{*} \text{trg}$.

6.1 Ruling out the Unbounded Case

By Lemma 4.1, given an instance of reachability, we can check in polynomial time whether src is unbounded in \mathcal{A} and trg is unbounded in \mathcal{A}^R . If both are true, then, by Lemma 4.2, it suffices to determine whether there exists a candidate run from src to trg . The existence of a candidate run can be decided in NP (e.g., using integer linear programming, see [12]). This case will not affect our complexity result because $\text{NP} \subseteq \text{coNP}^{\text{NP}}$. Thus, without loss of generality, we may assume that src is bounded in \mathcal{A} or trg is bounded in \mathcal{A}^R . Moreover, if trg is bounded in \mathcal{A}^R , we symmetrically work with \mathcal{A}^R instead of \mathcal{A} .

In the remainder of this section, we assume that src is bounded in \mathcal{A} .

6.2 Inductive Invariants in the Bounded Case

We will show that $\text{src} \not\rightarrow^* \text{trg}$ if and only if there exists a certificate of a particular shape witnessing this non-reachability. This certificate takes the form of an inductive invariant separating src and trg . The exact set of configurations comprising this inductive invariant is unwieldy, so we concentrate on its *core* instead. This set of core configurations admits a short representation, as follows.

We call an *arithmetic progression* on state $q \in Q$ a set of configurations $\{(q, v) \mid \ell \leq v \leq L \wedge \exists k \in \mathbb{N}, v = kp + s\}$ with $p, s, \ell, L \in \mathbb{N}$. An arithmetic progression can be specified by writing q and the numbers p, s, ℓ, L . A set of configurations *has a concise description* if it is a union of at most $2|Q|^2 + 1$ arithmetic progressions whose configurations have counter value bounded by $2|Q| \cdot \|\Delta\| \cdot \|\tau\|$. Such a set can be described in polynomial space.

The set of all configurations in bounded chains has a concise description. This also holds for the set R of all *reachable* configurations in bounded chains: indeed, if a configuration of a chain can be reached, the same is true for all configurations above in the same chain. Because src is bounded, unbounded chains cannot be reached, and this R is in fact the set of reachable configurations in all chains. (Observe that runs that reach configurations from R may well visit the complement of Conf_+ .)

► **Lemma 6.2.** *The set R of reachable configurations in Conf_+ has a concise description.*

Intuitively, R is our desired ‘core invariant’, and the desired invariant is the set $\text{Post}_*(R)$. However, when given a set I , it is not easy to check whether I is actually equal to R . Instead, the following theorem defines possible invariant cores by 3 conditions.

Conditions involving src and trg are self-explanatory. Set inclusion $\text{Post}(\text{Post}_*(I)) \subseteq \text{Post}_*(I)$ would express inductiveness (closure of the set under $\text{Post}(\cdot)$). However, verifying this condition is computationally expensive, and we replace it with a version that ‘focuses’ on the core only, and thus has I rather than $\text{Post}_*(I)$ on the right-hand side.

► **Theorem 6.3.** *Suppose src is bounded in \mathcal{A} . Then $\text{src} \not\rightarrow^* \text{trg}$ if and only if there exists a set $I \subseteq \text{Conf}_+ \cup \{\text{src}\}$ with concise description such that:*

- (Cond1) $\text{src} \in I$,
- (Cond2) $\text{trg} \notin \text{Post}_*(I)$, and
- (Cond3) $\text{Post}(\text{Post}_*(I)) \cap \text{Conf}_+ \subseteq I$.

Proof. First, assume that there is such a set I . Because $\text{trg} \notin \text{Post}_*(I)$ by (Cond2), it suffices to prove that $\text{Post}^*(\text{src}) \subseteq \text{Post}_*(I)$. We proceed by induction on the length of the run from src to $c \in \text{Post}^*(\text{src})$. The base of induction is (Cond1). Assume that we have $d \in \text{Post}_*(I)$ and $d \rightarrow c$. If $c \notin \text{Conf}_+$ then we have a pessimistic run from I to c ,

so $c \in \text{Post}_-^*(I)$. If $c \in \text{Conf}_+$ then $c \in \text{Post}(\text{Post}_-^*(I)) \cap \text{Conf}_+$, hence $c \in I$ by (Cond3). For the other direction, assume that trg is not reachable from src . Let $I \stackrel{\text{def}}{=} R \cup \{\text{src}\}$; by Lemma 6.2, I has a concise description. (Cond1) and (Cond2) are trivially satisfied. Moreover, $\text{Post}(\text{Post}_-^*(I)) \cap \text{Conf}_+ \subseteq \text{Post}^*(\text{src}) \cap \text{Conf}_+ \subseteq I$, hence (Cond3) is satisfied. \blacktriangleleft

6.3 The Complexity of Reachability in Strongly Connected OCA

We now prove that reachability is in coNP^{NP} by, equivalently, proving that *non*-reachability is in NP^{NP} . We refer the reader to Appendix C.2 for definitions of these complexity classes. For our purpose, the following informal description suffices. A problem is in NP^{NP} whenever this problem is solvable in non-deterministic polynomial time by a Turing machine which has access to an oracle for some NP-complete problem. The oracle is a black box that may provide the answer to any problem in NP (and therefore to any problem in coNP).

As argued in Section 6.1, we assume with no loss of generality that src is bounded. By Theorem 6.3, we have $\text{src} \xrightarrow{*} \text{trg}$ if and only if there exists I satisfying the three conditions (Cond1), (Cond2), and (Cond3). Moreover, by the same theorem, I can be assumed to have a concise description. Thus, we can guess such a set I in non-deterministic polynomial time. It remains to prove that the verification that a set I satisfies the three conditions can be performed using an NP oracle. To this end, we prove that this verification is a coNP problem. Indeed, I does *not* satisfy the three conditions when:

- either $\text{src} \notin I$ (which can be checked efficiently),
- or $\text{trg} \in \text{Post}_-^*(I)$ (this is when there is a *small* configuration c such that $c \in I$ and $\text{trg} \in \text{Post}_-^*(c)$),
- or there are some *small* configurations c and d such that $c \in I$, $d \in \text{Post}_-^*(c)$, and some successor of d belongs to Conf_+ but not to I .

The adjective *small* should here be understood as “bounded by an exponential in the size of \mathcal{A} , src , and trg ”. In fact, it is fairly easy to obtain an exponential bound on configurations to consider, see Appendix C.3. Thanks to Lemma 5.1, verification of both whether there is a $c \in I$, and $e \in \text{Post}(d)$ such that $\text{trg} \in \text{Post}_-^*(c)$ and whether there exist $c \in I$, $d \in \text{Post}_-^*(c)$ such that $e \notin I$ are in NP. Since membership in Conf_+ can be checked in polynomial time by Lemma 4.3, the entire third condition is also an NP condition. This completes the proof of Theorem 6.1.

7 Combining Strongly Connected Components

In this section, we extend the techniques from Section 6 to the general case in which the OCA is not assumed to be strongly connected. We fix an OCA \mathcal{A} with disequality tests and two configurations src and trg .

We first highlight why the techniques developed above do not apply to this general case. In Section 6, the hypothesis that \mathcal{A} is strongly connected was necessary for the application of Lemma 4.2. When \mathcal{A} is not strongly connected, knowing that src is unbounded is no longer satisfactory. Indeed, it no longer implies the existence of a positive cycle involving its state, as the positive cycle allowing us to pump up could be in another SCC. We need to be able to specify whether a configuration is unbounded *within its own SCC* or not.

7.1 Locally Bounded Configurations and Runs

Locally bounded configurations. Given a SCC S of \mathcal{A} , we denote by \mathcal{A}_S the automaton obtained when restricting \mathcal{A} to states and transitions within S . A configuration c is *locally bounded* if c is bounded in \mathcal{A}_S where S is the SCC of c . We denote by L the set of all locally bounded configurations (and by L^R in \mathcal{A}^R). Configurations that are not locally bounded are referred to as *locally unbounded*. We generalise the lifting technique from Lemma 4.2. The proof of the following lemma can be found in Appendix D.1.

► **Lemma 7.1** (Lifting). *For all $c \notin L$ and $d \notin L^R$, we have $c \xrightarrow{*} d$ if and only if $c \xrightarrow{Z}^* d$.*

Locally bounded runs. A run $c \xrightarrow{\pi} d$ is said to be *locally bounded* if all configurations visited by the run are locally bounded. We denote such a run by $c \xrightarrow{L}^{\pi} d$, and denote its existence by $c \xrightarrow{L}^* d$. Notice that a locally bounded run may go through several SCCs. Moreover, a run starting from a locally bounded configuration is not always locally bounded: once it goes to a new SCC, it may visit configurations that are not locally bounded. We define the locally bounded counterpart LPost_*^* of the pessimistic post-star operator: $d \in \text{LPost}_*^*(c)$ if there is a pessimistic and locally bounded run from c to d . We extend this definition to sets of configurations X in the usual way. Dually, we also define, for every set of configurations X , the set $\text{LPre}_+^*(X)$ as the same notion but in the reverse OCA \mathcal{A}^R . We extend Lemma 5.1 to these new operators.

► **Lemma 7.2.** *Given $c, d \in \text{Conf}$, deciding whether $d \in \text{LPost}_*^*(c)$ is in NP.*

Proof sketch. We split the run on its transitions between SCCs. We apply Lemma 5.1 on the portions remaining in one SCC. Since the run is pessimistic, we can bound all the counter values in it. The run is locally bounded when the first configuration visited in each SCC is locally bounded, which is checked using Lemma 4.1. The full proof can be found in Appendix D.1. ◀

7.2 Leaky Invariants

Unlike in the strongly-connected case, a single invariant construction is not sufficient for our needs. Indeed, if src is locally bounded but unbounded, then one could imagine the invariant technique from Theorem 6.3 applied to the SCC S of src , but then this invariant would not apply to other SCCs. For example, there could be runs that are locally bounded in the SCC S_{src} of src but not in the SCC S_{trg} of trg , making the invariant inapplicable. Instead, assuming that trg is locally bounded in \mathcal{A}^R , one may consider in the SCC of trg an invariant constructed in the reverse automaton \mathcal{A}^R . We therefore employ a pair of invariants, one for \mathcal{A} (the *forward invariant*) and another one for its reverse \mathcal{A}^R (the *backward invariant*). The two invariants will induce two sets of configurations that, in a negative instance of the reachability problem, separate the source and target.

The following lemma will allow us to avoid treating src and trg separately. The set Conf_+^R is defined as the counterpart of Conf_+ in \mathcal{A}^R .

► **Lemma 7.3.** *We may assume that $\text{src} \in \text{Conf}_+ \cap L$ and $\text{trg} \in \text{Conf}_+^R \cap L^R$.*

We now define our notion of a leaky invariant. As in Section 6, we represent the invariants using *core* sets of configurations that can be succinctly described, denoted by I and J . Our invariants must be inductive in the following weak sense:

► **Condition 7.4.** Let $I \subseteq \text{Conf}_+ \cap L$ and $J \subseteq \text{Conf}_+^R \cap L^R$ be sets of configurations. The pair (I, J) is *inductive* if

$$\begin{aligned} (\text{Ind}) \quad & \text{Post}(\text{LPost}_-^*(I)) \cap \text{Conf}_+ \cap L \subseteq I \quad \text{and} \\ & \text{Pre}(\text{LPre}_+^*(J)) \cap \text{Conf}_+^R \cap L^R \subseteq J. \end{aligned}$$

Notice that I and J play symmetric roles in \mathcal{A} and \mathcal{A}^R . We now provide some intuition for the (forward) inductive condition for I . The set I only contains configurations from $\text{Conf}_+ \cap L$, because the set $\text{Conf}_+ \cap L$ has a regular structure thanks to bounded chains. The set I is, again, only the *core* of the invariant. The full invariant² is $\text{Post}_-^*(I) \cup \text{Post}(\text{Post}_-^*(I))$, but this set is not easily described (see Remark 7.8). This explains why we use the composition $\text{Post}(\text{LPost}_-^*(\cdot))$ instead of the single-step $\text{Post}(\cdot)$ operator traditionally used to define inductiveness.

► **Remark 7.5.** We refer to our invariants as *leaky*, because they are not inductive in the traditional sense. Indeed, our invariants are ‘focused’ on locally bounded configurations, and can be escaped by transitions to locally unbounded configurations. This leak may, however, only happen with transitions going from one SCC to another.

► **Condition 7.6.** Let $\mathcal{I}, \mathcal{J} \subseteq \text{Conf}$ be sets of configurations. The pair $(\mathcal{I}, \mathcal{J})$ is a *separator* if, for all $c \in \mathcal{I}$ and $d \in \mathcal{J}$,

$$\begin{aligned} (\text{Sep1}) \quad & c \not\rightarrow d; \text{ and} \\ (\text{Sep2}) \quad & \text{if } c \notin L \text{ and } d \notin L^R, \text{ then } c \not\stackrel{*}{\rightarrow}_{\mathbb{Z}} d. \end{aligned}$$

Firstly, (Sep1) will forbid \mathcal{I} and \mathcal{J} from being connected by a single step. Secondly, (Sep2) will forbid connection between \mathcal{I} and \mathcal{J} using the lifting technique of Lemma 7.1. This does not, in general, prevent the existence of runs from \mathcal{I} to \mathcal{J} ; it will do so, however, for our leaky invariants that combine Conditions 7.4 and 7.6.

► **Definition 7.7.** Let $I \subseteq \text{Conf}_+ \cap L$ and $J \subseteq \text{Conf}_+^R \cap L^R$. Consider the sets

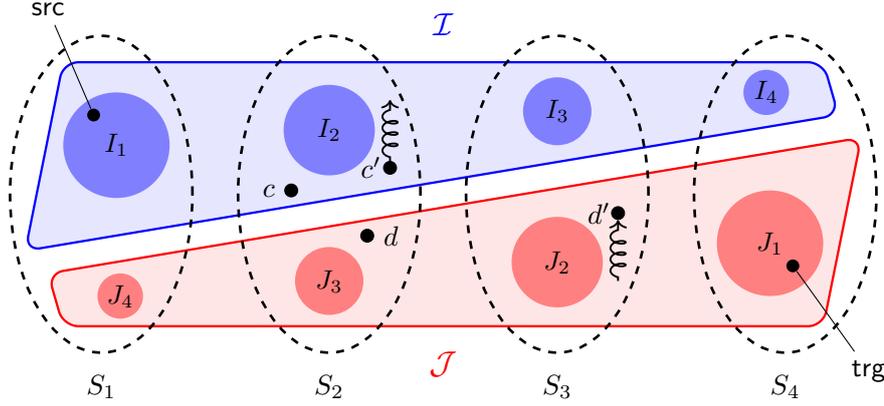
$$\begin{aligned} \mathcal{I} &:= \text{Post}_-^*(I) \cup \text{Post}(\text{Post}_-^*(I)) \text{ and} \\ \mathcal{J} &:= \text{Pre}_+^*(J) \cup \text{Pre}(\text{Pre}_+^*(J)). \end{aligned}$$

We call the pair (I, J) a non-reachability witness for src and trg if (I, J) is inductive, $(\mathcal{I}, \mathcal{J})$ is a separator, $\text{src} \in I$, and $\text{trg} \in J$.

The pair of sets (I, J) forms the *core* of the invariant, namely I represents the *forward leaky invariant* and J represents the *backward leaky invariant*. In this case we also say that (I, J) *induces* the separator $(\mathcal{I}, \mathcal{J})$. A visualisation of a pair (I, J) and its induced separator $(\mathcal{I}, \mathcal{J})$ can be seen in Figure 2. A helpful intuition is that \mathcal{I} is approximately $\text{Post}_-^*(I)$ (and similarly \mathcal{J} is approximately $\text{Pre}_+^*(J)$). One additional step of $\text{Post}(\cdot)$ (and $\text{Pre}(\cdot)$, respectively) ensures that the ‘outer boundary’ of this closure should also be included in the set.

► **Remark 7.8.** As in Section 6, our representation of invariants refers to their core only, i.e., the pair (I, J) . The example in Fig. 1 (left) demonstrates that the set $\text{Post}_-^*(I) \cup \text{Post}(\text{Post}_-^*(I))$ does not always have a tractable description. The set of all possible sums of subsets has no convenient description, therefore we want it to be captured by $\text{Post}_-^*(I) \cup \text{Post}(\text{Post}_-^*(I))$ only and not by I itself.

² Our invariant is $\text{Post}_-^*(I) \cup \text{Post}(\text{Post}_-^*(I))$ but the operator that appears in Condition 7.4 is $\text{LPost}_-^*(\cdot)$. In Appendix D.3, we discuss the issues encountered if $\text{Post}_-^*(\cdot) \cup \text{Post}(\text{Post}_-^*(\cdot))$ was used in Condition 7.4.



■ **Figure 2** Core inductive sets and the separator they induce. The core of the forward leaky invariant is $I = I_1 \cup I_2 \cup I_3 \cup I_4$ (the blue circular sets) and the core of the backward leaky invariant is $J = J_1 \cup J_2 \cup J_3 \cup J_4$ (the red circular sets). The induced separator $(\mathcal{I}, \mathcal{J})$ is shown as blue and red rounded quadrilaterals containing the core sets. Notice that $\text{src} \in I$ and $\text{trg} \in J$. The upwards coiled arrow from c' represents that c' is locally unbounded in the SCC S_2 and the downwards coiled arrow from d' represents that d' is locally unbounded in the SCC S_3^R . Note also the separator conditions: Condition 7.6 (Sep1) means that configurations $c \in \mathcal{I}$ and $d \in \mathcal{J}$ cannot reach one another by one transition, so $c \not\rightarrow d$; Condition 7.6 (Sep2) means that unbounded configurations $c' \in \mathcal{I}$ and $d' \in \mathcal{J}$ cannot reach one another via a candidate run, so $c' \not\stackrel{*}{\underset{Z}{\rightarrow}} d$.

► **Theorem 7.9.** *In an OCA \mathcal{A} with disequality tests, trg is not reachable from src if and only if there exists a non-reachability witness. Moreover, in this case, there is always a non-reachability witness with a concise description.*

In Section 7.3 we define ‘perfect cores’, which we use in Section 7.4 to sketch a proof of Theorem 7.9, details being relegated to Appendix D.5.

7.3 Perfect Cores

Condition (Ind) on the core of leaky invariants captures a weak inductiveness property, which is central to our approach. We will now discuss two features of this condition that are used in the proof of Theorem 7.9.

Our conditions capture a specific invariant, which we now define. Consider the set

$$B \stackrel{\text{def}}{=} \{c \in \text{Conf} : \text{src} \xrightarrow[L]{*} c\}$$

In words, B contains all configurations reachable from src using locally bounded runs. In line with ideas from Section 6, we do not want to store B entirely, so we will restrict the core to configurations in bounded chains. We call *perfect core* the set $B \cap \text{Conf}_+$; the term *perfect* is motivated by the fact that Definition 7.7 *aims* to capture this set exactly. Similarly, let $B^R := \{c \in \text{Conf} : c \xrightarrow[L^R]{*} \text{trg}\}$. The *perfect core* in the reverse automaton is $B^R \cap \text{Conf}_+^R$.

The two features of (Ind) are summarised in the following two lemmas. We use the words “sound” and “complete” to characterise the relationship between Condition 7.4 (as part of Definition 7.7) and the perfect cores defined above. Completeness expresses that in *every instance* of non-reachability, the perfect cores defined above induce a non-reachability invariant. Conversely, soundness states that *every invariant* must contain all configurations from the perfect cores. (Thus, the perfect core are the smallest possible invariants.) The full proofs of soundness and completeness can be found in Appendix D.4.

► **Lemma 7.10** (Soundness). *For all $I \subseteq \text{Conf}_+ \cap L$ and $J \subseteq \text{Conf}_+^R \cap L^R$ such that $\text{src} \in I$ and $\text{trg} \in J$, if (I, J) is inductive (Condition 7.4), then $B \cap \text{Conf}_+ \subseteq I$ and $B^R \cap \text{Conf}_+^R \subseteq J$.*

Proof sketch. Condition 7.4 for I gives $\text{Post}(\text{LPost}_-^*(I)) \cap \text{Conf}_+ \cap L \subseteq I$. Let $c \in B \cap \text{Conf}_+$ be a configuration of the perfect core. Thus, src reaches c by a locally bounded run. It is not always true that $c \in \text{Post}(\text{LPost}_-^*(\text{src}))$ because this run does not have to be pessimistic: it may observe configurations in Conf_+ . We prove by induction that all configurations in Conf_+ along the run are in I , using the property that $\text{Post}(\text{LPost}_-^*(I)) \cap \text{Conf}_+ \cap L \subseteq I$ once for each such configuration; this eventually proves that $c \in I$. The proof is analogous for J . ◀

► **Lemma 7.11** (Completeness). *If $I = B \cap \text{Conf}_+$ and $J = B^R \cap \text{Conf}_+^R$, then (I, J) is inductive (Condition 7.4).*

Proof sketch. We prove that $\text{Post}(\text{LPost}_-^*(B \cap \text{Conf}_+)) \cap \text{Conf}_+ \cap L \subseteq B \cap \text{Conf}_+$. Let $c \in \text{Conf}_+$ be locally bounded and belong to $\text{Post}(\text{LPost}_-^*(B \cap \text{Conf}_+))$. All configurations in $\text{LPost}_-^*(B \cap \text{Conf}_+)$ are in B by the definition of B , so c can be reached in one step from a configuration $d \in B$. By definition, d is reachable from src with a locally bounded run; since c is itself locally bounded, this is also true for c , and so $c \in B$. The case of J is similar. ◀

7.4 Non-reachability Witnesses and Their Complexity

We are ready to prove our main results. Here, we will provide proof sketches, the full proofs can be found in Appendix D.5.

► **Theorem 7.9.** *In an OCA \mathcal{A} with disequality tests, trg is not reachable from src if and only if there exists a non-reachability witness. Moreover, in this case, there is always a non-reachability witness with a concise description.*

Proof sketch. First, if $\text{src} \not\rightarrow^* \text{trg}$ then the perfect cores $I = B \cap \text{Conf}_+$ and $J = B^R \cap \text{Conf}_+^R$ form a non-reachability witness. Indeed, by Lemma 7.11, (I, J) is inductive. Moreover, the induced $\mathcal{I} = \text{Post}_-^*(I) \cup \text{Post}(\text{Post}_-^*(I))$ and $\mathcal{J} = \text{Pre}_+^*(J) \cup \text{Pre}(\text{Pre}_+^*(J))$ form a separator. We have $\mathcal{I} \subseteq \text{Post}^*(\text{src})$ and $\mathcal{J} \subseteq \text{Pre}^*(\text{trg})$, proving Condition 7.6 (Sep1). If Condition 7.6 (Sep2) fails, Lemma 7.1 yields a contradiction. Moreover, I and J have a concise description thanks to bounded chains.

Conversely, suppose there is a non-reachability witness (I, J) . Assume for the sake of contradiction that $\text{src} \rightarrow^* \text{trg}$. By Lemma 7.10, since (I, J) is inductive, $B \cap \text{Conf}_+ \subseteq I$ and $B^R \cap \text{Conf}_+^R \subseteq J$. Consider a run from src to trg . It must leave $\mathcal{I} = \text{Post}_-^*(I) \cup \text{Post}(\text{Post}_-^*(I))$ therefore it visits locally unbounded configurations. Let c be the first such configuration visited. Similarly, let d be the last visited configuration that is locally unbounded in \mathcal{A}^R . First, if c occurs before d , then Condition 7.6 (Sep2) is violated. Second, if c occurs after d , then there is an overlap in the runs from src to c and from d to trg . The overlap must be in $\mathcal{I} \cap \mathcal{J}$, leading to a violation of Condition 7.6 (Sep1). ◀

► **Theorem 7.12.** *The reachability problem for OCA with disequality tests belongs to the complexity class coNP^{NP} .*

Proof sketch. We use Theorem 7.9, deciding the existence of a non-reachability witness in NP^{NP} . Recall that NP^{NP} is introduced in Section 6.3. Let a pair (I, J) be given; we show that a violation of the conditions for being a non-reachability witness can be checked in NP.

- One can check in polynomial time whether $\text{src} \in I$ and $\text{trg} \in J$.

- Violation of Condition 7.4 (Ind) is an NP property. Indeed, this follows because membership in $L\text{Post}_-^*(\cdot)$ is in NP (by Lemma 7.2) and membership in Conf_+ and L is polynomial-time checkable (by Lemma 4.3 and Lemma 4.1, respectively).
- Violation of Condition 7.6 (Sep1) is in NP, because $\mathcal{I} := \text{Post}_-(I) \cup \text{Post}(\text{Post}_-(I))$, $\mathcal{J} := \text{Pre}_+(J) \cup \text{Pre}(\text{Pre}_+(J))$, and membership of given configurations in the pessimistic post-star (optimistic pre-star, respectively) is in NP by Lemma 5.1. This assumes that we have an exponential bound on relevant configurations, which is guaranteed by Lemma C.2 in the appendix.
- To check violation of Condition 7.6 (Sep2) in NP, we again use Lemma 4.1 for L , as well as the fact that the existence of a candidate run is in NP (by integer programming). ◀

7.5 Adding Equality Tests

The previous techniques have been developed for OCA with disequality tests only. In particular, the lifting argument of Lemma 7.1 does not hold in the presence of equality tests: candidate runs that visit a state with an equality test cannot be lifted to greater counter values. However, at the cost of increasing the complexity, one can handle equality tests.

Complexity class $\text{P}^{\text{NP}^{\text{NP}}}$ consists of decision problems solvable in polynomial time with access to an NP^{NP} oracle (which can solve NP^{NP} problems in one step).

► **Corollary 7.13.** *The reachability problem for OCA with equality and disequality tests belongs to the complexity class $\text{P}^{\text{NP}^{\text{NP}}}$.*

Proof. Let \mathcal{A} be such an OCA with tests, and src and trg two configurations. Denote by $\text{Conf}_=$ the set of valid configurations at states with equality tests; $|\text{Conf}_=|$ does not exceed the number of states in \mathcal{A} . Consider the OCA with disequality tests \mathcal{A}' that is obtained by deleting all states with equality tests (and incident transitions) from \mathcal{A} . By Theorem 7.12, with an NP^{NP} oracle we can build a graph with vertex set $\text{Conf}_= \cup \{\text{src}, \text{trg}\}$ and edge set $\{(c, d) \mid c \xrightarrow{*} d \text{ in } \mathcal{A}'\}$. Depth-first search in this graph for a path from src to trg takes polynomial time. ◀

8 Conclusions

We have looked at the reachability problem for one-counter automata with equality and disequality tests. We have proposed the idea of local inductive invariants and combined them with the notion of unboundedness within an SCC. Our construction circumvents the lack of computationally tractable descriptions: indeed, in the subset sum example (Fig. 1 (left) and Remark 7.8) the reachability set has exponential size, depending on a_1, \dots, a_n . There is no obvious means of compression available, and guessing/storing a traditional invariant is prohibitively expensive even for moderate n .

An outstanding theoretical question is characterisation of complexity of reachability in OCA with disequality tests. We have placed the problem in coNP^{NP} and, in the presence of equality tests, in $\text{P}^{\text{NP}^{\text{NP}}}$. Both problems have already been known to be NP-hard. Are they NP-complete or coNP -hard too? We also leave it open whether our technique can be extended to other systems and settings, e.g., to parameter synthesis questions (see, e.g., [26, 35, 44]).

In a more practical direction, while the general invariant-based effective procedure for (non-)reachability in vector addition systems [36] has not, to the best of our knowledge, been implemented, our work identifies these potentially practical ways to reduce the search space for invariants in VASS. The idea of restricting invariant sets to just a small ‘core’ (in our

case: a union of arithmetic progressions), combined with the compositionality of invariants, can help to direct an exploration of the search space, or assist a learning algorithm.

References

- 1 Shaull Almagor, Nathann Cohen, Guillermo A. Pérez, Mahsa Shirmohammadi, and James Worrell. Coverability in 1-VASS with disequality tests. In Igor Konnov and Laura Kovács, editors, *31st International Conference on Concurrency Theory, CONCUR 2020, September 1-4, 2020, Vienna, Austria (Virtual Conference)*, volume 171 of *LIPICs*, pages 38:1–38:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- 2 Rajeev Alur and Pavol Černý. Streaming transducers for algorithmic verification of single-pass list-processing programs. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, pages 599–610, 2011.
- 3 Nicolas Amat, Silvano Dal-Zilio, and Thomas Hujsa. Property directed reachability for generalized Petri nets. In Dana Fisman and Grigore Rosu, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 28th International Conference, TACAS 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022, Proceedings, Part I*, volume 13243 of *Lecture Notes in Computer Science*, pages 505–523. Springer, 2022.
- 4 Pascal Baumann, Moses Ganardi, Rupak Majumdar, Ramanathan S. Thinniyam, and Georg Zetsche. Context-bounded verification of context-free specifications. *Proc. ACM Program. Lang.*, 7(POPL):2141–2170, 2023.
- 5 Michael Blondin, Matthias Englert, Alain Finkel, Stefan Göller, Christoph Haase, Ranko Lazic, Pierre McKenzie, and Patrick Totzke. The reachability problem for two-dimensional vector addition systems with states. *J. ACM*, 68(5):34:1–34:43, 2021.
- 6 Michael Blondin, Alain Finkel, Christoph Haase, and Serge Haddad. Approaching the coverability problem continuously. In Marsha Chechik and Jean-François Raskin, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 22nd International Conference, TACAS 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven, The Netherlands, April 2-8, 2016, Proceedings*, volume 9636 of *Lecture Notes in Computer Science*, pages 480–496. Springer, 2016.
- 7 Michael Blondin, Alain Finkel, Christoph Haase, and Serge Haddad. The logical view on continuous Petri nets. *ACM Trans. Comput. Log.*, 18(3):24:1–24:28, 2017.
- 8 Michael Blondin, Christoph Haase, and Philip Offtermatt. Directed reachability for infinite-state systems. In Jan Friso Groote and Kim Guldstrand Larsen, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 27th International Conference, TACAS 2021, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2021, Luxembourg City, Luxembourg, March 27 - April 1, 2021, Proceedings, Part II*, volume 12652 of *Lecture Notes in Computer Science*, pages 3–23. Springer, 2021.
- 9 Michael Blondin, Tim Leys, Filip Mazowiecki, Philip Offtermatt, and Guillermo A. Pérez. Continuous one-counter automata. *ACM Trans. Comput. Log.*, 24(1):3:1–3:31, 2023.
- 10 Stanislav Böhm, Stefan Göller, and Petr Jancar. Equivalence of deterministic one-counter automata is NL-complete. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 131–140. ACM, 2013.
- 11 Benedikt Bollig, Karin Quaas, and Arnaud Sangnier. The complexity of flat freeze LTL. *Log. Methods Comput. Sci.*, 15(3), 2019.
- 12 Itshak Borosh and Leon Bruce Treybig. Bounds on positive integral solutions of linear Diophantine equations. *Proceedings of the American Mathematical Society*, 55(2):299–304, 1976.

- 13 Aaron R. Bradley. SAT-based model checking without unrolling. In Ranjit Jhala and David A. Schmidt, editors, *Verification, Model Checking, and Abstract Interpretation - 12th International Conference, VMCAI 2011, Austin, TX, USA, January 23-25, 2011. Proceedings*, volume 6538 of *Lecture Notes in Computer Science*, pages 70–87. Springer, 2011.
- 14 Aaron R. Bradley. Understanding IC3. In Alessandro Cimatti and Roberto Sebastiani, editors, *Theory and Applications of Satisfiability Testing - SAT 2012 - 15th International Conference, Trento, Italy, June 17-20, 2012. Proceedings*, volume 7317 of *Lecture Notes in Computer Science*, pages 1–14. Springer, 2012.
- 15 Daniel Bundala and Joël Ouaknine. On parametric timed automata and one-counter machines. *Inf. Comput.*, 253:272–303, 2017.
- 16 Dmitry Chistikov, Wojciech Czerwinski, Piotr Hofman, Michal Pilipczuk, and Michael Wehar. Shortest paths in one-counter systems. *Log. Methods Comput. Sci.*, 15(1), 2019.
- 17 Edmund M. Clarke, Orna Grumberg, Somesh Jha, Yuan Lu, and Helmut Veith. Counterexample-guided abstraction refinement. In E. Allen Emerson and A. Prasad Sistla, editors, *Computer Aided Verification, 12th International Conference, CAV 2000, Chicago, IL, USA, July 15-19, 2000. Proceedings*, volume 1855 of *Lecture Notes in Computer Science*, pages 154–169. Springer, 2000.
- 18 Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In Robert M. Graham, Michael A. Harrison, and Ravi Sethi, editors, *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages, Los Angeles, California, USA, January 1977*, pages 238–252. ACM, 1977.
- 19 Wojciech Czerwinski, Slawomir Lasota, Ranko Lazic, Jérôme Leroux, and Filip Mazowiecki. The reachability problem for Petri nets is not elementary. *J. ACM*, 68(1):7:1–7:28, 2021.
- 20 Wojciech Czerwinski and Lukasz Orlikowski. Reachability in vector addition systems is Ackermann-complete. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1229–1240. IEEE, 2021.
- 21 Stéphane Demri and Arnaud Sangnier. When model-checking freeze LTL over counter machines becomes decidable. In C.-H. Luke Ong, editor, *Foundations of Software Science and Computational Structures, 13th International Conference, FOSSACS 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings*, volume 6014 of *Lecture Notes in Computer Science*, pages 176–190. Springer, 2010.
- 22 Javier Esparza, Ruslán Ledesma-Garza, Rupak Majumdar, Philipp J. Meyer, and Filip Nikić. An SMT-based approach to coverability analysis. In Armin Biere and Roderick Bloem, editors, *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings*, volume 8559 of *Lecture Notes in Computer Science*, pages 603–619. Springer, 2014.
- 23 John Fearnley and Marcin Jurdzinski. Reachability in two-clock timed automata is PSPACE-complete. *Inf. Comput.*, 243:26–36, 2015.
- 24 Alain Finkel, Serge Haddad, and Igor Khmelnitsky. Minimal coverability tree construction made complete and efficient. In Jean Goubault-Larrecq and Barbara König, editors, *Foundations of Software Science and Computation Structures - 23rd International Conference, FOSSACS 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020. Proceedings*, volume 12077 of *Lecture Notes in Computer Science*, pages 237–256. Springer, 2020.
- 25 Seymour Ginsburg and Edwin H. Spanier. Semigroups, Presburger formulas and languages. *Pacific Journal of Mathematics*, 16(2):285–296, 1966. doi:10.2140/pjm.1966.16.285.
- 26 Stefan Göller, Christoph Haase, Joël Ouaknine, and James Worrell. Model checking succinct and parametric one-counter automata. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010, Bordeaux, France, July 6-10,*

- 2010, *Proceedings, Part II*, volume 6199 of *Lecture Notes in Computer Science*, pages 575–586. Springer, 2010.
- 27 Stefan Göller and Markus Lohrey. Branching-time model checking of one-counter processes and timed automata. *SIAM J. Comput.*, 42(3):884–923, 2013.
 - 28 Christoph Haase, Stephan Kreutzer, Joël Ouaknine, and James Worrell. Reachability in succinct and parametric one-counter automata. In Mario Bravetti and Gianluigi Zavattaro, editors, *CONCUR 2009 - Concurrency Theory, 20th International Conference, CONCUR 2009, Bologna, Italy, September 1-4, 2009. Proceedings*, volume 5710 of *Lecture Notes in Computer Science*, pages 369–383. Springer, 2009.
 - 29 Michel Hack. *Decidability questions for Petri nets*. PhD thesis, MIT, 1975. URL: <http://publications.csail.mit.edu/lcs/pubs/pdf/MIT-LCS-TR-161.pdf>.
 - 30 Matthew Hague and Anthony Widjaja Lin. Model checking recursive programs with numeric data types. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, volume 6806 of *Lecture Notes in Computer Science*, pages 743–759. Springer, 2011.
 - 31 Matthew Hague and Anthony Widjaja Lin. Synchronisation- and reversal-bounded analysis of multithreaded programs with counters. In P. Madhusudan and Sanjit A. Seshia, editors, *Computer Aided Verification - 24th International Conference, CAV 2012, Berkeley, CA, USA, July 7-13, 2012 Proceedings*, volume 7358 of *Lecture Notes in Computer Science*, pages 260–276. Springer, 2012.
 - 32 Matthias Heizmann, Jochen Hoenicke, and Andreas Podelski. Software model checking for people who love automata. In Natasha Sharygina and Helmut Veith, editors, *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *Lecture Notes in Computer Science*, pages 36–52. Springer, 2013. doi:10.1007/978-3-642-39799-8_2.
 - 33 Alexander Kaiser, Daniel Kroening, and Thomas Wahl. Efficient coverability analysis by proof minimization. In Maciej Koutny and Irek Ulidowski, editors, *CONCUR 2012 - Concurrency Theory - 23rd International Conference, CONCUR 2012, Newcastle upon Tyne, UK, September 4-7, 2012. Proceedings*, volume 7454 of *Lecture Notes in Computer Science*, pages 500–515. Springer, 2012.
 - 34 Pascal Lafourcade, Denis Lugiez, and Ralf Treinen. Intruder deduction for AC-like equational theories with homomorphisms. In Jürgen Giesl, editor, *Term Rewriting and Applications, 16th International Conference, RTA 2005, Nara, Japan, April 19-21, 2005, Proceedings*, volume 3467 of *Lecture Notes in Computer Science*, pages 308–322. Springer, 2005.
 - 35 Antonia Lechner, Richard Mayr, Joël Ouaknine, Amaury Pouly, and James Worrell. Model checking flat freeze LTL on one-counter automata. *Log. Methods Comput. Sci.*, 14(4), 2018.
 - 36 Jérôme Leroux. The general vector addition system reachability problem by Presburger inductive invariants. *Log. Methods Comput. Sci.*, 6(3), 2010.
 - 37 Jérôme Leroux. Distance between mutually reachable Petri net configurations. In Arkadev Chattopadhyay and Paul Gastin, editors, *39th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2019, December 11-13, 2019, Bombay, India*, volume 150 of *LIPICs*, pages 47:1–47:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
 - 38 Jérôme Leroux. The reachability problem for Petri nets is not primitive recursive. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 1241–1252. IEEE, 2021.
 - 39 Jérôme Leroux and Sylvain Schmitz. Reachability in vector addition systems is primitive-recursive in fixed dimension. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*, pages 1–13. IEEE, 2019.
 - 40 Xie Li, Taolue Chen, Zhilin Wu, and Mingji Xia. Computing linear arithmetic representation of reachability relation of one-counter automata. In Jun Pang and Lijun Zhang, editors,

- Dependable Software Engineering. Theories, Tools, and Applications - 6th International Symposium, SETTA 2020, Guangzhou, China, November 24-27, 2020, Proceedings*, volume 12153 of *Lecture Notes in Computer Science*, pages 89–107. Springer, 2020.
- 41 Ernst W. Mayr. An algorithm for the general Petri net reachability problem. *SIAM J. Comput.*, 13(3):441–460, 1984.
 - 42 Albert R. Meyer and Larry J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential space. In *13th Annual Symposium on Switching and Automata Theory, College Park, Maryland, USA, October 25-27, 1972*, pages 125–129. IEEE Computer Society, 1972.
 - 43 Marvin L. Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, Inc., USA, 1967.
 - 44 Guillermo A. Pérez and Ritam Raha. Revisiting parameter synthesis for one-counter automata. In Florin Manea and Alex Simpson, editors, *30th EACSL Annual Conference on Computer Science Logic, CSL 2022, February 14-19, 2022, Göttingen, Germany (Virtual Conference)*, volume 216 of *LIPICs*, pages 33:1–33:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
 - 45 Louis E. Rosier and Hsu-Chun Yen. A multiparameter analysis of the boundedness problem for vector addition systems. *J. Comput. Syst. Sci.*, 32(1):105–135, 1986.
 - 46 Sylvain Schmitz. The complexity of reachability in vector addition systems. *SIGLOG News*, 3(1):4–21, 2016.
 - 47 Helmut Seidl, Thomas Schwentick, Anca Muscholl, and Peter Habermehl. Counting in trees for free. In Josep Díaz, Juhani Karhumäki, Arto Lepistö, and Donald Sannella, editors, *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12-16, 2004. Proceedings*, volume 3142 of *Lecture Notes in Computer Science*, pages 1136–1149. Springer, 2004.
 - 48 Alistair Stewart, Kousha Etessami, and Mihalis Yannakakis. Upper bounds for Newton’s method on monotone polynomial systems, and P-time model checking of probabilistic one-counter automata. *J. ACM*, 62(4):30:1–30:33, 2015.
 - 49 Larry J. Stockmeyer. The polynomial-time hierarchy. *Theor. Comput. Sci.*, 3(1):1–22, 1976.
 - 50 Leslie G. Valiant and Mike Paterson. Deterministic one-counter automata. *J. Comput. Syst. Sci.*, 10(3):340–350, 1975.
 - 51 Celia Wrathall. Complete sets and the polynomial-time hierarchy. *Theor. Comput. Sci.*, 3(1):23–33, 1976.

A

 Details for Section 4

A.1 Finding Positive-Effect Simple Cycles is NP-hard

► **Proposition A.1.** *Deciding, for a given OCA without tests \mathcal{A} and a given state q , whether there exists a positive-effect simple q -cycle in \mathcal{A} is an NP-complete problem.*

Proof. Membership in NP is obtained by using the q -cycle itself as a certificate. To prove NP-hardness, we provide a reduction from the Hamiltonian path problem. Let $G = (V, E)$ be a directed graph and let $s, t \in V$ be two distinct vertices. A path from s to t is *Hamiltonian* if it is simple and visits every vertex in the graph. The Hamiltonian path problem takes as input a directed graph $G = (V, E)$ and two vertices $s, t \in V$ and asks whether there is a Hamiltonian path from s to t in G .

For the remainder of this proof, we fix an instance of this problem formed by $G = (V, E)$ and $s, t \in V$. Let $n = |V|$. We will now construct an OCA without tests (a 1-VASS) $\mathcal{A} = (Q, \Delta)$. Define $Q := V \cup \{q\}$, where $q \notin V$ is a new state, and

$$\Delta := \{(u, 1, v) : (u, v) \in E\} \cup \{(q, 0, s), (t, -(n-2), q)\}.$$

The construction of \mathcal{A} takes polynomial time. We claim that there exists a Hamiltonian path from s to t in G if and only if there exists a positive-effect simple q -cycle in \mathcal{A} .

Suppose there exists a Hamiltonian path π from s to t in G . Since π visits every vertex in G , we have $\text{len}(\pi) = n - 1$. Consider the path σ in \mathcal{A} that is obtained from π by replacing each edge $(u, v) \in E$ with the corresponding transition $(u, 1, v) \in \Delta$ as well as prepending the transition $(q, 0, s)$ and appending the transition $(t, -(n-2), q)$. Given that π is a simple path in G , we know that σ is a simple q -cycle in \mathcal{A} . Furthermore, given that $\text{len}(\pi) = n - 1$, we know that $\text{eff}(\sigma) = 0 + n - 1 - (n - 2) = 1$, so σ has positive effect.

Conversely, suppose there exists a positive-effect simple q -cycle σ in \mathcal{A} . This σ must begin with $(q, 0, s)$, the only outgoing transition from q , and end with $(t, -(n-2), q)$, the only transition leading back to q . Let $\sigma = (q, 0, s) \sigma' (t, -(n-2), q)$ for some σ' . Given that $\text{eff}(\sigma) \geq 1$ and all other transitions in \mathcal{A} have effect 1, we know that $\text{len}(\sigma') \geq n - 1$. Since the cycle σ is simple and $|Q \setminus \{q, s, t\}| = n - 2$, we conclude that σ' visits each of these $n - 2$ states exactly once. So the path π obtained from σ' by replacing each transition $(u, 1, v) \in \Delta$ with the corresponding edge $(u, v) \in E$ is a Hamiltonian path from s to t in G . ◀

A.2 Missing Proofs

► **Lemma 4.2.** *Let \mathcal{A} be a strongly connected OCA with disequality tests. If src is unbounded in \mathcal{A} and trg is unbounded in \mathcal{A}^R , then there is a run from src to trg in \mathcal{A} ($\text{src} \xrightarrow{*} \text{trg}$) if and only if there is a candidate run from src to trg in \mathcal{A} ($\text{src} \xrightarrow[*]{\mathbb{Z}} \text{trg}$).*

Proof. First, if $\text{src} \xrightarrow{*} \text{trg}$, then $\text{src} \xrightarrow[*]{\mathbb{Z}} \text{trg}$. Conversely, let π be a path such that $\text{src} \xrightarrow[\mathbb{Z}]{\pi} \text{trg}$. Let $(s, v) = \text{src}$. We claim that there is a positive cycle γ_+ such that, for all k , $(\gamma_+)^k$ yields a valid run from src (we say that γ_+ can be “pumped up” from src). Given that src is unbounded in \mathcal{A} , there exists a state $q \in Q$, a path π , and $z > v + \|\tau\| + \|\Delta\| \cdot |Q| + 1$ such that $\text{src} \xrightarrow{\pi} (q, z)$. Further, given that \mathcal{A} is strongly connected, there is a path σ , of length at most $|Q|$, from q back to s ; let $\gamma_+ \stackrel{\text{def}}{=} \pi \cdot \sigma$. We know that, because $z \geq \|\tau\| + \|\Delta\| \cdot |Q|$, following σ from (q, z) yields a valid run; suppose that run ends at the configuration (s, v') . Notice that $v' > v + \|\tau\|$, so the cycle γ_+ from (s, v') also yields a valid run, because all observed counter values are still greater than the disequality guards. Let $w_+ > 0$ be the effect

of γ_+ . Symmetrically, there is a cycle γ_- , with effect $w_- < 0$, that can be “pumped down” arbitrarily many times to trg . By letting $m > \text{drop}(\pi) + \|\tau\|$, the path $(\gamma_+)^{m|w_-|} \cdot \pi \cdot (\gamma_-)^{mw_+}$ yields a run from src to trg . ◀

► **Lemma 4.3.** *There is a polynomial-time algorithm to identify Q_+ and to choose cycles γ_q for all $q \in Q_+$. Moreover, membership in Conf_+ can be decided in polynomial time.*

Proof. The proof will use 1-dimensional vector addition systems with states (1-VASS). These are one-counter automata (as defined in this paper) without any tests.

- Syntactically, a 1-VASS is a pair (Q, T) , where Q is the set of states and $T \subseteq Q \times \mathbb{Z} \times Q$ is the set of transitions.
- The semantics is the same as that of an OCA with disequality tests (Q, T, true) , where the map true assigns true to all states in Q .

We will consider an auxiliary problem which takes as input a 1-VASS (Q, T) , a state $q \in Q$, and a natural number d (encoded in binary); the problem asks whether there is a positive-effect q -cycle γ of length at most $|Q|$ such that $\text{drop}(\gamma) \leq d$. We show that this problem can be solved in polynomial time, given that coverability in 1-VASS can be decided in polynomial time. Coverability is placed in P by a standard reduction from coverability to unboundedness (see, e.g., [1, Lemma 1]) and a polynomial-time algorithm for unboundedness by Rosier and Yen [45, Theorem 3.4]. Notice a stronger result from [1] is that coverability in 1-VASS is in fact in $\text{NC}^2 \subseteq \text{P}$.

Reduction of the auxiliary problem to coverability. Let $n = |Q|$. We can construct an instance of coverability as follows. Consider the unfolding (Q', T') where $Q' = \{q^{(i)} : q \in Q \text{ and } i \in [0, n]\}$ and $T' = \{(p^{(i-1)}, a, q^{(i)}) : (p, a, q) \in T \text{ and } i \in [1, n]\}$. Observe that there exists a positive-effect q -cycle whose length is at most $|Q|$ and with a drop bounded by d in (Q, T) if and only if $(q^{(i)}, d+1)$ can be covered from $(q^{(i)}, d)$ in (Q', T') for some $i \in [1, n]$. Moreover, in that case, such a cycle is obtained directly from a path in the unfolded 1-VASS that witnesses coverability.

Polynomial-time algorithms for minimum drop and membership in Conf_+ . We complete the proof of Lemma 4.3.

- First, observe that the minimum drop can be computed by a binary search for d . Let $m = \max\{|a| : (p, a, q) \in T\}$. By starting from an upper bound of $n(m+1)$, d can be computed using a polynomial number (at most $\lceil \log(n(m+1)) \rceil$) of coverability queries.
- Second, to decide membership of a configuration (q, v) in Conf_+ , it suffices to check that $q \in Q_+$, to compute $\text{drop}(\gamma_q)$, and to check that $v \geq \text{drop}(\gamma_q)$. ◀

► **Lemma 4.5** (see [1, Remark 6]). *There are at most $2|Q|^2$ bounded chains.*

Proof. As observed in [1, Remark 6], for each $q \in Q$, the number of bounded q -chain is at most $2|Q|$. Indeed, each guard seen in γ_q this can be seen by pulling forwards the disequality guards of each state along the simple cycle γ_q to q . Each of these guards, say $x \neq g$, induces two bounded chains, namely the singleton $\{(q, g)\}$ and the set of configurations below the guard with counter values congruent to g modulo $\text{eff}(\gamma_q)$. ◀

B

 Details for Section 5

B.1 Flows and Paths

We start by formally defining the concept of flow. Given a function $f : \Delta \rightarrow \mathbb{N}$ and $q \in Q$, the *flow entering* q is $\text{in}_q(f) := \sum_{t=(q',a,q) \in \Delta} f(t)$ and the *flow leaving* q is $\text{out}_q(f) := \sum_{t=(q,a,q') \in \Delta} f(t)$. The *underlying undirected graph* of f is the graph $\mathcal{G}(f) := (Q', E)$ such that:

- Q' contains all states $q \in Q$ such that either $\text{in}_q(f) \neq 0$ or $\text{out}_q(f) \neq 0$, and
- for all $q, q' \in Q$, $\{q, q'\} \in E$ whenever there exists $t \in \Delta$ such that $f(t) > 0$ and such that $t = (q, a, q')$ or $t = (q', a, q)$ for some $a \in \mathbb{Z}$.

In words, $\mathcal{G}(f)$ has as edges the undirected and unlabelled version of transitions in \mathcal{A} that have value greater than 0 for f , and as vertices all the states that are incident to these transitions. This graph is *connected* if there is a path between any two vertices.

► **Definition B.1.** A flow from $p \in Q$ to $q \in Q$ is a mapping $f : \Delta \rightarrow \mathbb{N}$ such that:

1. if $p = q$, then $\text{in}_p(f) = \text{out}_p(f)$;
2. if $p \neq q$, then $\text{out}_p(f) = \text{in}_p(f) + 1$ and $\text{in}_q(f) = \text{out}_q(f) + 1$;
3. for every $q' \in Q \setminus \{p, q\}$, $\text{in}_{q'}(f) = \text{out}_{q'}(f)$; and
4. $\mathcal{G}(f)$ is connected and contains p and q .

Let f be a flow from p to q . We may refer to p as the *start* of f and q as the *end* of f . The *size* of f is $|f| := \sum_{t \in \Delta} f(t)$. Its *effect* is $\text{eff}(f) := \sum_{t \in \Delta} \text{eff}(t) \cdot f(t)$.

A *decomposition* of a flow f from p to q is a finite sequence of flows f_1, \dots, f_m such that $\sum_{i=1}^m f_i = f$, the start of f_1 is p , the end of f_m is q , and, for all $i \in [1, m-1]$, the end of f_i is the start of f_{i+1} .

We can compare two flows f and g : we say that $f \leq g$ if $f(t) \leq g(t)$ for all $t \in \Delta$. A flow f *admits a positive cycle* if there exist $q \in Q$ and a flow $f_+ \leq f$ from q to q such that $\text{eff}(f_+) > 0$.

Let σ be a path from p to q . We define the function $\text{flow}(\sigma)$ that assigns, to each transition t , the number of occurrences of t in σ .

► **Lemma B.2.** Let σ be a path from p to q ; then $\text{flow}(\sigma)$ is a flow from p to q .

Proof. This follows immediately by induction on the prefixes of the path. ◀

We call $\text{flow}(\sigma)$ the *flow of* σ . Similarly, the *flow of a run* is the flow of the path underlying the run.

► **Lemma B.3.** For every flow f , there is a path σ such that $\text{flow}(\sigma) = f$.

Proof. We will prove, by induction, that for every $n \in \mathbb{N}$, for every flow f of size n , that there exists a path σ of length n such that $\text{flow}(\sigma) = f$.

Base case $n = 0$. The only flow f of size 0 has $f(t) = 0$ for all $t \in \Delta$. Let σ be the empty path; $\text{flow}(\sigma)[t] = 0$ for all $t \in \Delta$. Hence $\text{flow}(\sigma) = f$.

Inductive step $n > 0$. Let f be a flow from p to q of size n . For every $t \in \Delta$ such that $f(t) > 0$, we denote by $g_t : \Delta \rightarrow \mathbb{N}$ the function defined by $g_t(t') = f(t')$ for all $t' \in \Delta \setminus \{t\}$ and $g_t(t) = f(t) - 1$. First, we will justify that there is $t = (q', a, q) \in \Delta$ such that g_t is a flow from p to q' . There are two conditions to be a flow. These are the balance conditions (conditions 1, 2, and 3 of Definition B.1) and the connectivity condition (condition 4 of Definition B.1).

We first prove that, for every $t = (q', a, q)$, g_t satisfies the balance conditions. It suffices to prove the conditions at q and q' .

First, suppose $p = q$. In this case, $\text{in}_q(f) = \text{out}_q(f)$ and $\text{in}_{q'}(f) = \text{out}_{q'}(f)$. By definition of g_t , $\text{in}_q(g_t) - 1 = \text{out}_q(g_t)$ and $\text{in}_{q'}(g_t) = \text{out}_{q'}(g_t) - 1$. Hence g_t satisfies the balance conditions to be a flow from p to q' . Second, suppose $p \neq q$. In this case, $\text{in}_q(f) = \text{out}_q(f) + 1$, and $\text{in}_{q'}(f) = \text{out}_{q'}(f)$. Here, there is an additional case distinction: whether $p = q'$ or not. Suppose $p = q'$. In this case, we know that $\text{in}_q(g_t) = \text{out}_q(g_t)$. Furthermore, $\text{in}_{q'}(g_t) = \text{in}_{q'}(f) = \text{in}_p(f) = \text{out}_p(f) - 1 = \text{out}_{q'}(f) - 1 = \text{out}_{q'}(g_t)$. Because $q' = p$, this also means that $\text{in}_p(g_t) = \text{out}_p(g_t)$, so that g_t satisfies the balance conditions to be a flow of size $n - 1$ from p to $q' = p$. Now, suppose $p \neq q'$. Since $g_t(t) = f(t) - 1$, we know that, again, $\text{in}_q(g_t) = \text{out}_q(g_t)$. However, this time, $\text{in}_p(g_t) + 1 = \text{out}_p(g_t)$ is unchanged but $\text{in}_{q'}(g_t) - 1 = \text{out}_{q'}(g_t)$. Hence g_t satisfies the balance conditions of a flow of size $n - 1$ from p to q' .

We have proven that, for every t leading to q , the function g_t satisfies the balance conditions of a flow from p to q' . It remains to prove that there exists a transition t such that g_t satisfies the connectivity condition. First, if there is a self-loop transition $t = (q, a, q)$ such that $f(t) > 0$, then g_t immediately satisfies the connectivity condition. Now, suppose that there is no such self-loop at q . Let T_q denote the set of transitions of Δ that lead to q . First, observe that $\Delta_q \neq \emptyset$. Since $\mathcal{G}(f)$ is connected, there is an edge of this graph that contains q . Because $\text{out}_q(f) \leq \text{in}_q(f)$, there is a transition $t \in \Delta_q$ such that $f(t) > 0$. If Δ_q only contains one transition t , then q only has one incident edge in $\mathcal{G}(f)$. The graph $\mathcal{G}(g_t)$ is the same as $\mathcal{G}(f)$ except potentially with q removed if $f(t) = 1$ (otherwise, if $f(t) > 1$, q remains in the graph). Therefore, g_t satisfies the connectivity condition.

Now, suppose that Δ_q contains at least two transitions. Because f is a flow, we have $\text{in}_q(f) \leq \text{out}_q(f) + 1$. There exists $t \in \Delta_q$ such that $f(t) > 0$. For all $r \in Q$ such that $r \neq q$, $\text{in}_r(f) \leq \text{out}_r(f)$, so we can build a path by starting with t and iteratively selecting new transitions until we arrive back at q . This results in a q -cycle α such that $\text{flow}(\alpha) \leq f$. Let t' denote the last transition of α and let $q' \in Q$ such that $t' = (q', a, q)$. Given that, for all transitions s in α , we know that $f(s) > 0$, the removal of the edge $\{q', q\}$ from $\mathcal{G}(f)$ does not change its connectivity. The graph $\mathcal{G}(g_{t'})$ is either the same as $\mathcal{G}(f)$ or it has the edge $\{q', q\}$ removed (depending on whether $f(t') > 1$ or $f(t') = 1$). Regardless, $\mathcal{G}(g_{t'})$ is connected.

We have proven that there is $t = (q', a, q)$ such that g_t is a flow from p to q' of size $n - 1$. By the inductive hypothesis, there exists a path σ' of length $n - 1$ from p to q' . Therefore, the path $\sigma = \sigma' t$ is from p to q , has length n , and satisfies $\text{flow}(\sigma) = f$. ◀

B.2 There are No Positive Cycles in the Flows of Pessimistic Runs

We now prove that pessimistic runs may have no cycle in their flow. To do so, we need the following technical result inspired by [28, Proposition 6].

► **Lemma B.4.** *Let γ_+ be a positive cycle of length ℓ (ℓ transitions, ℓ states) and, for all $i \in [1, \ell]$, let γ_i be the cycle obtained when concatenating the suffix of γ_+ of length $\ell - i + 1$ with the prefix of γ_+ of length $i - 1$ (γ_i corresponds to γ_+ but starting from and finishing at the i th state). Then there exists $k \in [1, \ell]$ such that $\text{drop}(\gamma_k) = 0$.*

Proof. For all $j \in [1, \ell]$, let $\gamma_{i,j}$ the path composed of the first j transitions of γ_i . Let $k \in [1, \ell]$ such that $\text{eff}(\gamma_{1,k}) = \min_i \text{eff}(\gamma_{1,i})$. We claim that $\text{drop}(\gamma_k) = 0$, *i.e.*, that, for all j , $\text{eff}(\gamma_{k,j}) \geq 0$. If $j \leq \ell - k + 1$ then $\gamma_{1,k} \cdot \gamma_{k,j} = \gamma_{1,k+j}$ hence $\text{eff}(\gamma_{1,k}) + \text{eff}(\gamma_{k,j}) = \text{eff}(\gamma_{1,k+j})$ but $\text{eff}(\gamma_{1,k+j}) \geq \text{eff}(\gamma_{1,k})$ so $\text{eff}(\gamma_{k,j}) \geq 0$. If $j > \ell - k + 1$, then $\gamma_{k,j} = \gamma_{k,\ell-k+1} \cdot \gamma_{1,j-\ell+k-1}$ so $\text{eff}(\gamma_{k,j}) = \text{eff}(\gamma_{k,\ell-k+1}) + \text{eff}(\gamma_{1,j-\ell+k-1}) \geq \text{eff}(\gamma_{k,\ell-k+1}) + \text{eff}(\gamma_{1,k}) = \text{eff}(\gamma_+) \geq 0$. We have proven that, for all j , $\text{eff}(\gamma_{k,j}) \geq 0$ therefore $\text{drop}(\gamma_k) \geq 0$. \blacktriangleleft

B.3 Proof of Lemma 5.1

Proof of Lemma 5.1. Combine the following two lemmas (Lemma B.5 and Lemma B.6). \blacktriangleleft

\blacktriangleright **Lemma B.5.** *The flow of a pessimistic run admits no positive cycle.*

Proof. Let ρ a pessimistic run, f its flow and assume that f admits a positive cycle; let $f_+ \leq f$ a flow from q to q such that $\text{eff}(f_+) > 0$. We build a positive cycle π_+ on q from f_+ thanks to Lemma B.3. By iteratively deleting negative simple cycles from π_+ , we obtain a positive simple cycle γ_+ . By Lemma B.4, one can reorder the transitions of γ_+ to obtain a positive cycle γ'_+ on some state q' such that $\text{drop}(\gamma'_+) = 0$. This means that any visit to q' is a visit to Conf_+ . Because γ_+ was built by deleting edges from γ_+ , there is a transition in f that ends on q' : q' is visited in ρ after its first configuration which contradicts the fact that ρ is pessimistic. \blacktriangleleft

For the next lemma, we recall that the size of an OCA \mathcal{A} with disequality tests, denoted $|\mathcal{A}|$, is defined as the bit size of its encoding, where all numbers are written in binary.

\blacktriangleright **Lemma B.6.** *If the flow of a run $\rho : (q, x) \xrightarrow{*} (q', y)$ admits no positive cycle, then $(q, x) \xrightarrow{*} (q', y)$ admits a certificate verifiable in time polynomial in $\log(x + y + 1) + |\mathcal{A}|$.*

Proof. Thanks to Lemma C.1, if such a ρ exists then we can assume that its length is bounded by $M := \psi(|\mathcal{A}|) + x + y$. Therefore, its flow f satisfies $f \leq \psi(|\mathcal{A}|) + x + y$ hence is storable in space $|\Delta| \log(\psi(|\mathcal{A}|) + x + y)$. We assume that all flows described below are bounded by $\psi(|\mathcal{A}|) + x + y$.

A *certificate* for the existence of a run from (q_s, x) to (q_f, y) consists in two things. First, a flow f from q_s to q_f of effect $y - x$ with no positive cycle and of largest value bounded by $\psi(|\mathcal{A}|) + x + y$; additionally, a decomposition of f into f_1, \dots, f_m such that $m \leq 4|Q| + 1$. The decomposition must moreover satisfy some conditions stated below. Let $q_s = q_0, q_1, \dots, q_m = q_f$ the corresponding states, so that f_i is a flow from q_{i-1} to q_i . For all i , let $w_i := \sum_{j=0}^i \text{eff}(f_j)$ and $x_i := w_i + x$. For all $q \in Q$, let $I(q) := \{i \in [0, m] \mid q_i = q\}$ and $X(q) := \{x_i \mid i \in I(q)\}$. Also, for all q such that $I(q) \neq \emptyset$, let $m_q := \min I(q)$ and $M_q := \max I(q)$. First, we impose that, for all i , $x_i \geq 0$ and $x_i \notin \tau(q_i)$ (all values x_i are non-negative and do not hit disequality guards). Moreover, for every state q appearing in the flow:

1. $I(q) \neq \emptyset$; and
2. for all $i \geq M_q$, q does not appear in f_i except in one transition leaving q for $i = M_q$; and
3. for all $i \leq m_q$, q does not appear in f_i except in one transition entering q for $i = m_q$.

The previous set of conditions guarantees that the counter value stays above 0.

Additionally, for every state q appearing in f that has some disequality guard $g \geq 0$, if there exist $x, y \in X(q)$ such that $x < g < y$ (*i.e.*, the flow crosses g), then there exists $i, j \in I(q)$ such that $i < j$, $I(q) \cap [i, j] = \{i, j\}$ and:

4. $x_i > g > x_j$;
5. q does not appear in flows f_{i+1} to f_j , except in one transition leaving q in f_{i+1} and one transition entering q in f_j .

These conditions guarantee that the disequality guards are not hit. Note, that if all values x_i are above (resp. below) g , then the last (resp. first) occurrence of q not hitting the guard witnesses that the guard is never hit.

We now prove that (q_f, y) is reachable from (q_s, x) if and only if there is a certificate for the existence of a run from (q_s, x) to (q_f, y) .

Suppose that there is a run ρ from (q_s, x) to (q_f, y) . Let f denote the flow of ρ . In order to build the required decomposition of f , we simply split ρ into successive runs. First, we split ρ on every last first and occurrence of a state. Moreover, for every state q that has a disequality guard g , we split run on the last visit to state q above g , if it exists, and on the first visit to state q below g , if it exists. Doing so, we have divided ρ into at most $4|Q| + 1$ consecutive runs ρ_1, \dots, ρ_m ; we simply let f_i the flow of ρ_i and all conditions above are trivially satisfied.

Conversely, assume that we have a flow f from q_s to q_f of effect $y - x$ and a decomposition $f_1 \dots f_m$ with $m \leq 4|Q| + 1$ that satisfies the conditions above. Let $q_s = q_0, q_1, \dots, q_m = q_f$ such that, for all i , f_i is a flow from q_{i-1} to q_i . For all i , let $w_i := \text{eff}(\sum_{j=1}^i f_j)$, and $x_i := w_i + x$. For every i , let σ_i a path obtained from f_i using Lemma B.3. By immediate induction, for all i , there is a candidate run ρ_i from (q_{i-1}, x_{i-1}) to (q_i, x_i) , which forms a candidate run ρ from (q_s, x) to (q_f, y) . We will now prove that all candidate runs are actual runs, *i.e.*, that all values remain non-negative and that no disequality guard is hit.

First, all values remain positive throughout ρ . Indeed, let $q \in Q$. Because f does not have negative cycles, if there is $q \in Q$ such that (q, x) and (q, y) are seen in ρ in this order, then $x \geq y$. Let $k := M_q$. Because ρ_k ends on (q, x_k) and $x_k \geq 0$, all values in ρ on state q until (q, x_k) have non-negative counter value. By Item 2, (q, x_k) is the last visit of ρ to q hence all counter values on q are positive.

Moreover, no disequality guard is hit. Let g be a disequality guard on some state q . If for all $x \in X(q)$, $x > g$, then in particular this is true for x_{M_q} which proves that g is not hit. Similarly, if for all $x \in X(q)$, $x < g$, then in particular this is true for x_{m_q} which is, by Item 3, the highest counter value on state q in ρ , so g is not hit. Suppose now that there exists $x, y \in X(q)$ such that $x < g < y$. Let $i, j \in I(q)$ satisfying the conditions above. Before the visit to (q, x_i) and after the visit to (q, x_j) , we know that g is not hit because f has no positive cycle. Thanks to Item 5, we have that there is no visit to q in ρ between the visit to (q, x_i) and the one to (q, x_j) , proving that g is not hit.

We finally argue that one can check in polynomial time whether a given flow along with a decomposition are a certificate; in particular, one can check in polynomial time whether a flow has no positive cycle. \blacktriangleleft

C Details for Section 6

C.1 Proof of Lemma 6.2

First, there are at most $2|Q|^2$ bounded chains by Lemma 4.5 above. It suffices to prove that the set of reachable configurations within a chain is an arithmetic progression of upper limit less than $\|\tau\| + (n+1)\|\Delta\|$. Let C a chain on state q , let $\gamma_+ := \gamma_q$ and $w := \text{eff}(\gamma_+)$. For every $x < y$ such that $(q, x), (q, y) \in C$, we have that there is $k \geq 0$ such that $(q, x) \xrightarrow{(\gamma_+)^k} (q, y)$; therefore, $(q, x) \equiv (q, y) \pmod{w}$. Moreover, if there is z such that $x < z < y$ and $z \equiv x$

(mod w) then there exists $\ell < k$ such that $(q, x) \xrightarrow{(\gamma_+)^{\ell}} (q, z)$ hence $(q, z) \in C$. This proves that C is an arithmetic progression; we now bound its largest configuration (q, x_{\max}) . Because γ_+ is a simple cycle, $w \leq (|Q| + 1)\|\Delta\|$. By maximality of x_{\max} , γ_+ cannot be taken from (q, x_{\max}) but $(q, x_{\max}) \in \text{Conf}_+$ therefore γ_+ must hit a disequality guard when applied from (q, x_{\max}) hence $x_{\max} \leq (|Q| + 1)\|\Delta\| + \|\tau\| \leq 2|Q| \cdot \|\Delta\| \cdot \|\tau\|$. \blacktriangleleft

C.2 The Polynomial Hierarchy

This hierarchy of computational complexity classes was first introduced by Meyer and Stockmeyer [42, 49], see also Wrathall [51]. Let $\Sigma_0^P \stackrel{\text{def}}{=} \Pi_0^P \stackrel{\text{def}}{=} P$, comprising all polynomial-time decision problems. For all $n \in \mathbb{N}$, let Σ_n^P be the set of all decision problems S such that there exists a predicate M decidable in polynomial time and a polynomial P satisfying that, for every instance x , we have that $x \in S$ if and only if

$$Q_1 y_1 \in \{0, 1\}^{P(|x|)} Q_2 y_2 \in \{0, 1\}^{P(|x|)} \dots Q_n y_n \in \{0, 1\}^{P(|x|)} : \\ M(x, y_1, \dots, y_n) = \text{YES},$$

where $Q_1 Q_2 \dots Q_n \in \{\exists, \forall\}^n$ is a (strictly) alternating sequence of quantifiers \exists and \forall that starts with \exists . The definition of Π_n^P is identical except that the sequence of quantifiers starts with \forall . In particular, $\Sigma_1^P = \text{NP}$ and $\Pi_1^P = \text{coNP}$. An equivalent definition uses oracles, by letting $\Sigma_0^P = P$ and for all $n \geq 0$, $\Sigma_{n+1}^P = \text{NP}^{\Sigma_n^P}$. Said otherwise, a problem is in Σ_{n+1}^P when it can be solved in non-deterministic polynomial time using a Σ_n^P oracle. Similarly, $\Pi_{n+1}^P = \text{coNP}^{\Pi_n^P}$.

The *polynomial hierarchy* is $\text{PH} \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} \Sigma_n^P$, or equivalently $\text{PH} \stackrel{\text{def}}{=} \bigcup_{n \in \mathbb{N}} \Pi_n^P$. In words, PH is the class of all problems that can be decided with a constant number of alternating quantifiers.

By applying the quantifier-based definition to NP^{NP} , we obtain that a decision problem S is in NP^{NP} such that there exists a predicate M decidable in polynomial time and a polynomial P satisfying that, for every instance x , we have that $x \in S$ if and only if

$$\exists y \in \{0, 1\}^{P(|x|)} \forall z \in \{0, 1\}^{P(|x|)} M(x, y, z) = \text{YES}.$$

The oracle-based definition is that the class NP^{NP} contains all decision problems that can be solved in non-deterministic polynomial time with access to an NP oracle. Note that, for our purposes, an NP oracle is also a coNP oracle, as it may also solve coNP problems.

C.3 Technical Lemmas for Bounding Counter Values

► **Lemma C.1.** *For every OCA with equality and disequality tests \mathcal{A} and every two configurations $\text{src} = (s, x)$ and $\text{trg} = (t, y)$, if $\text{src} \xrightarrow{*} \text{trg}$ then $\text{src} \xrightarrow{\pi} \text{trg}$ for some valid run ρ whose length and highest counter value are at most $\exp(|\mathcal{A}| + \log(x + 1) + \log(y + 1))$.*

Proof. We transform the given \mathcal{A} into another OCA, where updates are encoded in unary. Each transition are replaced by a simple path of length equal to the absolute value of its update value. An equality test $x = k$ can be first replaced with a sequence of three operations: decrement by k , equality test for 0, and increment by k . A disequality test $x \neq k$ can be directly converted into a sequence of $k + 1$ decrements and $k + 1$ increments, connected by a simple gadget comprising of k zero-test transitions and one ordinary transition. The result is a new OCA without disequality tests (but possibly with zero tests), which mimics the behaviour of the original one and in which all updates are encoded in unary.

We can also attach simple paths of length x and y to the obtained OCA so that the reachability query is for two configurations with counter value 0. The size of the new OCA is bounded from above by $n \stackrel{\text{def}}{=} \exp(|\mathcal{A}| + \log(x+1) + \log(y+1))$.

It is known that in every OCA with n states and unary updates (as well as possibly with zero tests) if there exists a run from the source configuration to the target configuration (with counter values 0), then there also exists such a run of length $O(n^2)$ (see, e.g., [16]). This run is mimicking a valid run in the original OCA with equality and disequality tests, so the upper bound on the length (and thus on the counter value) carries over. ◀

► **Lemma C.2.** *Suppose all configurations of a set $S \subseteq \text{Conf}$ have counter value bounded by $N \in \mathbb{N}$. Then all configurations in $\text{Post}_-^*(S)$ have counter value bounded by $N + (|Q| - 1)\|\Delta\|$.*

Proof. Let $\rho : (q, x) \xrightarrow{*} (q', y)$ be a pessimistic run, with $(q, x) \in S$. We prove, by induction on the number k of states visited in ρ , that the largest counter value in ρ is at most $N + (k - 1)\|\Delta\|$. For $k = 1$, because $(q, x) \in S$, we have $x \leq N$. Assume that ρ visits $k + 1$ states; let q_1, \dots, q_{k+1} the states appearing in ρ ordered by first visit in ρ . We decompose ρ on its first visit to q_{k+1} : $\rho : (q, x) \xrightarrow{\sigma_p} (q_p, x_p) \xrightarrow{t} (q_{k+1}, x) \xrightarrow{\sigma_s} (q', y)$. Let $\rho_p : \rho : (q, x) \xrightarrow{\sigma_p} (q_p, x_p)$ and $\rho_s : (q_{k+1}, x) \xrightarrow{\sigma_s} (q', y)$. By induction hypothesis, all counter values in ρ_p are less than $N + (k - 1)\|\Delta\|$ therefore $x \leq N + k\|\Delta\|$. We now prove it for ρ_s . Let (q_s, x_s) a configuration seen in ρ_s . If $q_s = q_i$ for $i \leq k$, then $x_s \leq x_i \leq N + (k - 1)\|\Delta\|$ as it otherwise would imply a positive cycle in the flow of ρ , contradicting Lemma B.5. For the same reason, if $q_s = q_{k+1}$ then $x_s \leq x \leq N + k\|\Delta\|$. This proves that all configurations in ρ are bounded by $N + k\|\Delta\|$, concluding the induction step. ◀

In particular, when checking the third condition of the invariant in Section 6.3, one already know that $c \in I$ can be considered to be at most exponential because the counter values I are bounded by an exponential in $|\mathcal{A}|$. Lemma C.2 therefore yields a bound on the counter value of the configuration d to consider, because there must be a pessimistic run from c to d . The same argument applies to the proof of Theorem 7.12.

D Details for Section 7

D.1 Proofs of Section 7.1

► **Lemma 7.1** (Lifting). *For all $c \notin L$ and $d \notin L^R$, we have $c \xrightarrow{*} d$ if and only if $c \xrightarrow{*} d$.*

Proof. First, if the instance is positive then there is a run and therefore a candidate run. In the other direction, let σ be a path corresponding to a candidate run from c to d . Recall that the configuration c is locally unbounded. Almagor et al. [1, Prop. 4] prove that a configuration in an OCA with disequality tests is unbounded if and only if it can reach an unbounded configuration in a set denoted Conf_+ . (The set Conf_+ in [1] is chosen based on a slightly different set of cycles compared to the present paper, but the difference does not matter to us here.) By this result, there exists an infinite lasso-shaped run starting in c , that is, whose path is of the form

$$\pi \lambda^\omega \stackrel{\text{def}}{=} \pi \lambda \lambda \lambda \dots,$$

where π and λ are finite sequences, and moreover the latter corresponds to some positive cycle (call it γ_+) within the SCC of c . It is no loss of generality to assume that λ has zero drop and starts in state q . (If this is not the case, the cycle can be ‘rotated’ appropriately;

see, e.g., Lemma B.4 in Appendix B.) Thus, from c it is possible to reach some configuration c' , pump the cycle γ_+ many times and get back to s using a path of less than $|Q|$ transitions. (Recall that π and λ stay in the SCC of c .) Therefore, we can build a positive s -cycle α_+ that can be pumped to infinity from c .

With a similar reasoning, we obtain in \mathcal{A}^R a positive t -cycle α_- that can be pumped down from infinity to d . Let σ be a candidate run from c to d . Let $M \stackrel{\text{def}}{=} \text{drop}(\sigma) + \|\tau\|$. We thus obtain a run from c to d by pumping the cycle α_+ up $M \cdot \text{eff}(\alpha_-)$ times, taking σ and pumping the cycle α_- down $M \cdot \text{eff}(\alpha_+)$ times. \blacktriangleleft

► **Lemma 7.2.** *Given $c, d \in \text{Conf}$, deciding whether $d \in \text{LPost}_*^*(c)$ is in NP.*

Proof. We want to provide a certificate the existence of a pessimistic and locally bounded run $c \xrightarrow[L]{\pi} d$. First, thanks to Lemma C.2, counter values in such a run are bounded by $(|Q| - 1)\|\Delta\| + |c|$. We are therefore able to store a given configuration observed by the run in space polynomial in the size of the system and the size of c . The certificate for $c \xrightarrow[L]{\pi} d$ consists in the following:

- the list S_1, \dots, S_m of SCCs visited by π in this order;
- for every SCC S_i , configurations e_i and e'_i where e_i is the first configuration of the run in S_i and e'_i is the last configuration of the run in S_i ,
- for every $i < m$, the transition $t \in \Delta$ from S_i to S_{i+1} that witnesses $e'_i \xrightarrow{t} e_{i+1}$;
- within a given SCC S_i , the certificate for the pessimistic run from e_i to e'_i as in Lemma 5.1.

To verify such a certification in polynomial time, we verify that the previous elements indeed define a run $c \xrightarrow{*} d$; moreover, we verify that all configurations e_i are locally bounded, which can be done in polynomial time thanks to Lemma 4.1. \blacktriangleleft

D.2 Proof of Lemma 7.3

Proof of Lemma 7.3. The assumption that $\text{src} \in \text{Conf}_+ \cap L$ and $\text{trg} \in \text{Conf}_+^R \cap L^R$ follows from Claim D.1. \blacktriangleleft

► **Claim D.1.** *Let \mathcal{A} be an OCA with disequality tests, and let (p, x) and (q, y) be two configurations. Then, there exists an OCA with disequality tests \mathcal{B} and two configurations (p', x) and (q', y) such that:*

1. $(p, x) \xrightarrow{*} (q, y)$ in \mathcal{A} if and only if $(p', x) \xrightarrow{*} (q', y)$ in \mathcal{B} ; and
2. $(p', x) \in \text{Conf}_+ \cap L$ and $(q', y) \in \text{Conf}_+^R \cap L^R$.

Proof. Given an OCA with disequality tests \mathcal{A} and two configurations (p, x) and (q, y) , we will construct an OCA with disequality tests \mathcal{B} by adding two states p' and q' , and four transitions $(p', 0, p)$; $\ell_1 = (p', 1, p')$; $(q, 0, q')$; and $\ell_2 = (q', -1, q')$. We also add disequality tests $\neq x + 1$ and $\neq y + 1$ at p' and q' , respectively.

First, observe that, in \mathcal{B} , indeed $(p', x) \in \text{Conf}_+$ since there is a positive-effect (single-transition) cycle (ℓ_1) with zero drop at p' (and so it can be taken from any counter value, including x). Also, $(q', y) \in \text{Conf}_+^R$ since, in reverse, the single-transition cycle (ℓ_2) has positive effect and zero drop at q' (and so it can be taken from any counter value, including y).

Second, it is also true that $(p', x) \in L$ in \mathcal{B} . This is true because p' belongs to its own SCC, and even though there is the positive-effect transition ℓ_1 , we know that (p', x) is locally bounded because of the disequality test added at p' that forbids the counter from being equal to $x + 1$. The same can be said for $(q', y) \in L^R$ in \mathcal{B}^R .

If $(p, x) \xrightarrow{*} (q, y)$ in \mathcal{A} , then trivially $(p', x) \xrightarrow{*} (q', y)$ in \mathcal{B} by appending and prepending the zero-effect transitions from p' to p and q to q' at the start and end of the run, respectively.

Conversely, suppose that $(p', x) \xrightarrow{\pi} (q', y)$ in \mathcal{B} , for some path π . It is not possible for π to use the transition ℓ_1 , otherwise the disequality test at p' is violated. Also, π does not use the transition ℓ_2 . Suppose, for a moment, that ℓ_2 is used. Since (ℓ_2) is the only cycle at q' and it has negative effect, then it must be the case that the run $(p', x) \xrightarrow{\pi} (q', y)$ first observes a configuration (q', y') for some $y' > y$. However, the only way to reach (q', y) from (q', y') is to use (ℓ_2) . This means that the configuration observed prior to reaching (q', y) would have to be $(q', y + 1)$ which, however, violates the disequality tests at q' . Now, given that $(p', x) \xrightarrow{\pi} (q', y)$ does not use transitions ℓ_1 and ℓ_2 , one can trim the zero-effect transitions from p' to p and q to q' away from the start and end of the run, respectively. This yields a run witnessing $(p, x) \xrightarrow{*} (q, y)$ in \mathcal{A} . ◀

D.3 Discussion of the Choice of Operators

In Condition 7.4, we made the choice of using operator $\text{LPost}_*(\cdot)$, and not $\text{Post}_*(\cdot) \cup \text{Post}(\text{Post}_*(\cdot))$ as in Section 6. Indeed, if we had used $\text{Post}_*(\cdot) \cup \text{Post}(\text{Post}_*(\cdot))$ instead, then in order to obtain completeness (Lemma 7.11), one would have to change the perfect core. We want the perfect core to be contained in $L \cap \text{Conf}_+$ so that it has a short representation; the natural candidate would be to take $\text{Post}^*(\text{src}) \cap L \cap \text{Conf}_+$ (and symmetrically in \mathcal{A}^R). This perfect core would satisfy the inductive property; however, this choice would break soundness (Lemma 7.10). Indeed, this invariant could contain a locally bounded configuration c that is reached from src using a run that visits many locally unbounded configurations in $\text{Conf} \setminus \text{Conf}_+$ before coming back in L . In this case, it could be that c is not captured by the inductive property, so one could find an inductive invariant I that does not contain c .

D.4 Proofs of Section 7.3

► **Lemma 7.10 (Soundness).** *For all $I \subseteq \text{Conf}_+ \cap L$ and $J \subseteq \text{Conf}_+^R \cap L^R$ such that $\text{src} \in I$ and $\text{trg} \in J$, if (I, J) is inductive (Condition 7.4), then $B \cap \text{Conf}_+ \subseteq I$ and $B^R \cap \text{Conf}_+^R \subseteq J$.*

Proof. Suppose (I, J) is an inductive core. We know therefore that

$$\text{Post}(\text{LPost}_*(I)) \cap \text{Conf}_+ \cap L \subseteq I.$$

Here, we will show that $B \cap \text{Conf}_+ \subseteq I$, showing $B^R \cap \text{Conf}_+^R \subseteq J$ follows in the same way.

For sake of contradiction, suppose there is a configuration in $B \cap \text{Conf}_+$ that does not belong to I . In fact, among all such configurations, we will have a configuration c that is witnessed by a locally bounded run $\text{src} \xrightarrow[L]{\pi} c$ observing the fewest configurations in Conf_+ . Note that since $\text{src} \in I$, we may assume that $\text{len}(\pi) \geq 1$. Suppose ℓ is the last configuration in the run that belongs to I and suppose p is the penultimate configuration in the run. Allow us to decompose the run into parts:

$$\text{src} \xrightarrow[L]{\theta} \ell \xrightarrow[L]{\sigma} p \xrightarrow[L]{t} c,$$

where $\theta \sigma t = \pi$.

We will now show that $\ell \xrightarrow[L]{\sigma} p$ is a pessimistic run. Indeed, if not, then this run observes a configuration $d \in \text{Conf}_+$ such that $d \neq p$. Suppose

$$\ell \xrightarrow[L]{\sigma_1} d \xrightarrow[L]{\sigma_2} p,$$

where $\sigma = \sigma_1 \sigma_2$. Then since $\text{src} \xrightarrow[L]{\theta \sigma_2} d$ is locally bounded and by assumption ℓ was the last configuration in I , we know that $d \in (B \cap \text{Conf}_+) \setminus I$. The run $\text{src} \xrightarrow[L]{\theta \sigma_2} d$ observes fewer configurations in Conf_+ than $\text{src} \xrightarrow[L]{\pi} d$, contradicting its minimality. Therefore, $\ell \xrightarrow[L]{\sigma} p$ is a pessimistic run for it does not observe any configurations in Conf_+ .

Now, since $\ell \in I$ and $\ell \xrightarrow[L]{\sigma} p$ is a pessimistic run, we know that $p \in \text{LPost}_-^*(I)$. Hence, since $p \xrightarrow[L]{t} c$, it is true that $c \in \text{Post}(\text{Post}_-^*(I))$. And so, since $B \subseteq L$, we know that $c \in \text{Post}(\text{LPost}_-^*(I)) \cap L \cap \text{Conf}_+ \subseteq I$, contradicting the assumption that $c \notin I$.

Therefore, for inductive cores (I, J) , $B \cap \text{Conf}_+ \subseteq I$ and $B^R \cap \text{Conf}_+^R \subseteq J$. \blacktriangleleft

► **Lemma 7.11 (Completeness).** *If $I = B \cap \text{Conf}_+$ and $J = B^R \cap \text{Conf}_+^R$, then (I, J) is inductive (Condition 7.4).*

Proof. Let $I = B \cap \text{Conf}_+$ and $J = B^R \cap \text{Conf}_+^R$. Here, we will indeed show that $\text{Post}(\text{LPost}_-^*(I)) \cap \text{Conf}_+ \cap L \subseteq I$, as per Condition 7.4. Showing that $\text{Pre}(\text{LPre}_+^*(J)) \cap \text{Conf}_+^R \cap L^R \subseteq J$ follows in the same way.

Suppose, for sake of contradiction, that $\text{Post}(\text{LPost}_-^*(I)) \cap \text{Conf}_+ \cap L \not\subseteq I$; therefore

$$\text{Post}(\text{LPost}_-^*(B \cap \text{Conf}_+)) \cap \text{Conf}_+ \cap L \not\subseteq B \cap \text{Conf}_+.$$

Thus, there exists a configuration c belonging to $\text{Post}(\text{LPost}_-^*(B \cap \text{Conf}_+))$, Conf_+ , and L such that c either does not belong to B or Conf_+ . Clearly, since $c \in \text{Conf}_+$, we may assume that $c \notin B$. Recall that B is the set of configurations that can be reached from src by a locally bounded run.

Now, if $c \in \text{Post}(\text{LPost}_-^*(B \cap \text{Conf}_+))$, then there exists $d \in B \cap \text{Conf}_+$ and $p \in \text{LPost}_-^*((p, x))$ such that

$$\text{src} \xrightarrow[L]{\theta} d \xrightarrow[L]{\sigma} p \xrightarrow[t]{} c.$$

Where $\text{src} \xrightarrow[L]{\theta} d$ is the locally bounded run for $d \in B$; $d \xrightarrow[L]{\sigma} p$ is the locally bounded (and pessimistic) run for $p \in \text{LPost}_-^*(d)$; and there is the single transition run $p \xrightarrow[t]{} c$ for $c \in \text{Post}(\text{LPost}_-^*(B \cap \text{Conf}_+))$. Notice that the final single transition run is also locally bounded because there is a locally bounded run to p and $c \in L$ is assumed. Altogether, we observe that there is a locally bounded from from src to c :

$$\text{src} \xrightarrow[L]{\theta \sigma t} c.$$

Hence, $c \in B$, contradicting the opening assumption. We therefore conclude that $\text{Post}(\text{LPost}_-^*(I)) \cap \text{Conf}_+ \cap L \subseteq I$. \blacktriangleleft

D.5 Proofs of Section 7.4

► **Theorem 7.9.** *In an OCA \mathcal{A} with disequality tests, trg is not reachable from src if and only if there exists a non-reachability witness. Moreover, in this case, there is always a non-reachability witness with a concise description.*

Before we prove Theorem 7.9, we first provide an auxiliary statement that will be used a few times in the proof. In short, the following claim states (for the forwards direction) that all configurations that are reachable, from the initial configuration, via locally bounded runs are captured by $\text{Post}_-^*(I)$.

► **Claim D.2.** *Let I and J be sets of configurations such that (I, J) is inductive (Condition 7.4) and let c be a configuration.*

(1) *If $\text{src} \xrightarrow[L]{*} c$, then $c \in \text{Post}^*(I)$.*

(2) *If $c \xrightarrow[L^R]{*} \text{trg}$, then $c \in \text{Pre}_+^*(J)$.*

Proof. We will prove statement (1) and statement (2) follows symmetrically. Suppose $\text{src} \xrightarrow[L]{*} c$.

First, suppose $c \in \text{Conf}_+$. Since $\text{src} \xrightarrow[L]{*} c$ is a locally bounded run, we know that $c \in B$. Since $c \in B \cap \text{Conf}_+$ and give that (I, J) is inductive, by Lemma 7.10 (soundness), we conclude that $c \in I \subseteq \text{Post}^*(I)$.

Second, suppose $c \notin \text{Conf}_+$. Let ℓ be the last configuration in $\text{src} \xrightarrow[L]{*} c$ that belongs to Conf_+ . Such a configuration ℓ exists because $\text{src} \in \text{Conf}_+$ by Lemma 7.3. Suppose $\text{src} \xrightarrow[L]{\pi_1} \ell \xrightarrow[L]{\pi_2} c$, where $\pi_1 \pi_2 \pi$. Since $\text{src} \xrightarrow[L]{\pi_1} \ell$, we know that $\ell \in B$. Thus, $\ell \in B \cap \text{Conf}_+$, so given that (I, J) is inductive, by Lemma 7.10 (soundness), we find that $\ell \in I$. Furthermore, by choice of ℓ , we know that no configurations in the run $\ell \xrightarrow[L]{\pi_2} c$, except ℓ , belong to Conf_+ . This means that $\ell \xrightarrow[L]{\pi_2} c$ is a pessimistic run, so $c \in \text{Post}_-(\ell) \subseteq \text{Post}_-(I)$, as required. ◀

Proof of Theorem 7.9.

First direction. If src cannot reach trg in \mathcal{A} , one can construct a non-reachability witness (I, J) for src and trg .

Suppose, indeed, that src cannot reach trg in \mathcal{A} . Let $I = B \cap \text{Conf}_+$ and $J = B^R \cap \text{Conf}_+^R$. Recall that B is the set of all configurations that can be reached from src via locally bounded runs in \mathcal{A} and B^R is the set of all configurations that can be reached from trg via locally bounded runs in \mathcal{A}^R . By Lemma 7.11 (completeness), (I, J) is inductive, so the pair satisfies Condition 7.4.

Now, consider the induced $\mathcal{I} = \text{Post}_-(I) \cup \text{Post}(\text{Post}_-(I))$ and $\mathcal{J} = \text{Pre}_+^*(J) \cup \text{Pre}(\text{Pre}_+^*(J))$. We will prove that $(\mathcal{I}, \mathcal{J})$ form a separator, so the pair satisfies Condition 7.6. First, we will argue that all configurations in \mathcal{I} can be reached from src . Consider $c \in \mathcal{I}$, there must exist $d \in I$ such that $c \in \text{Post}_-(d)$ or $c \in \text{Post}(\text{Post}_-(I))$. Regardless of whether $c \in \text{Post}_-(I)$ or $c \in \text{Post}(\text{Post}_-(I))$, there is a run from d to c . Since $d \in I \subseteq B$, there is a locally bounded run from src to d . So, combining the two runs, one observes that $\text{src} \xrightarrow{*} c$ in \mathcal{A} . The same argument can be used to show that all configurations in $c \in \mathcal{J}$ can be reached from trg in \mathcal{A}^R ; that means that $c \xrightarrow{*} \text{trg}$ in \mathcal{A} .

Suppose that $(\mathcal{I}, \mathcal{J})$ does not satisfy Condition 7.6 (Sep1). There exists $c \in \mathcal{I}$ and $d \in \mathcal{J}$ such that $c \rightarrow d$. By the above, we know that $(s, x) \xrightarrow{*} c$ and $d \xrightarrow{*} (t, y)$;

$$(s, x) \xrightarrow{*} c \rightarrow d \xrightarrow{*} (t, y) \implies (s, x) \xrightarrow{*} (t, y).$$

This creates a contradiction, so $(\mathcal{I}, \mathcal{J})$ must satisfy Condition 7.6 (Sep1).

Now, suppose that $(\mathcal{I}, \mathcal{J})$ does not satisfy Condition 7.6 (Sep2). In this case, there exists $c \in \mathcal{I}$ and $d \in \mathcal{J}$ such that c is locally unbounded in \mathcal{A} , d is locally unbounded in \mathcal{A}^R , and $c \xrightarrow{z}^* d$. By Lemma 7.1, we know that actually $c \xrightarrow{*} d$ holds. Again, by the above, we know that $(s, x) \xrightarrow{*} c$ and $d \xrightarrow{*} (t, y)$;

$$(s, x) \xrightarrow{*} c \xrightarrow{*} d \xrightarrow{*} (t, y) \implies (s, x) \xrightarrow{*} (t, y).$$

This creates a contradiction, so $(\mathcal{I}, \mathcal{J})$ must satisfy Condition 7.6 (Sep2).

We now prove that I and J defined above have a concise description. The argument is very similar to the one for the proof of Theorem 6.3. Indeed, B cannot intersect unbounded chains, as all configurations in unbounded chains are locally unbounded. Let C be one of at most $2|Q|^2$ bounded chains. The intersection of B with C is an arithmetic progression whose largest value is bounded by an exponential in the size of the system. Here, we can use Lemma 6.2.

Second direction. If there is a non-reachability witness (I, J) for src and trg , then src cannot reach trg in \mathcal{A} .

Suppose that there exist sets of configurations I and J such that $\text{src} \in I$; $\text{trg} \in J$; (I, J) is inductive (Condition 7.4); and the induced $\mathcal{I} = \text{Post}^*(I) \cup \text{Post}(\text{Post}^*(I))$ and $\mathcal{J} = \text{Pre}_+^*(J) \cup \text{Pre}(\text{Pre}_+^*(J))$ form a separator $(\mathcal{I}, \mathcal{J})$ (Condition 7.6). We will argue that src cannot reach trg .

Assume, for sake of contradiction, that src can reach trg in \mathcal{A} , so there is a run $\text{src} \xrightarrow{\pi} \text{trg}$. We will decompose the run into parts about reach of the k SSCs that were visited. Precisely, we let e_1, \dots, e_k be the first, and f_1, \dots, f_k be the last, configurations observed in each SCC, respectively.

$$\text{src} = e_1 \xrightarrow{\pi_1} f_1 \xrightarrow{t_1} e_2 \xrightarrow{\pi_2} f_2 \xrightarrow{t_2} \dots \xrightarrow{t_{k-1}} e_k \xrightarrow{\pi_k} f_k = \text{trg}$$

Note that e_1, \dots, e_k are the first, and f_1, \dots, f_k are the last, configurations observed in each SCC, respectively.

Consider the case where all e_1, \dots, e_k are locally bounded. This means that the entire run is locally bounded $\text{src} \xrightarrow[L]{\pi} \text{trg}$. Since $\text{src} \neq \text{trg}$, it is true that $\text{len}(\pi) \geq 1$, so we can decompose the run $\text{src} \xrightarrow[L]{\pi'} p \xrightarrow[L]{t} \text{trg}$, where p is the penultimate configuration in the run. By Claim D.2, we know that $p \in \text{Post}^*(I) \subseteq \mathcal{I}$. We know that $\text{trg} \in J$, and since $J \subseteq \mathcal{J}$, we have $\text{trg} \in \mathcal{J}$. This is a contradiction because $p \rightarrow \text{trg}$ holds, but $(\mathcal{I}, \mathcal{J})$ was assumed to satisfy Condition 7.6 (Sep1).

It is therefore the case that not all e_1, \dots, e_k are locally bounded. In fact, the same argument can be used to show that not all f_1, \dots, f_k are locally bounded in \mathcal{A}^R . Let $i \in \{1, \dots, k\}$ be the first index such that e_i is not a locally bounded configuration in \mathcal{A} and let $j \in \{1, \dots, k\}$ be the last index such that f_j is not a locally bounded configuration in \mathcal{A}^R . We will consider two subcases: when $i > j$ and when $i \leq j$.

Before splitting into these the two subcases, given that $e_1 = \text{src}$ and

$$e_1 \xrightarrow[L]{\pi_1 t_1 \dots t_{i-2} \pi_{i-1}} f_{i-1},$$

we know that $f_{i-1} \in \text{Post}^*(I)$ by Claim D.2. Furthermore, since $f_{i-1} \xrightarrow{t_{i-1}} e_i$, we deduce that $e_i \in \text{Post}(\text{Post}^*(I))$. Together, and to summarise, we know that both f_{i-1} and e_i belong to \mathcal{I} . Likewise, given that $f_k = \text{trg}$ and again by Claim D.2, we can conclude that f_j and e_{j+1} belong to \mathcal{J} .

Case $i > j$. We know that f_j occurs before f_{i-1} ($j \leq i-1$) so there is a locally bounded run from src to f_j :

$$e_1 \xrightarrow[L]{\pi_1 t_1 \dots t_{j-1} \pi_j} f_j \xrightarrow[L]{t_j \pi_{j+1} \dots t_{i-2} \pi_{i-1} f_{i-1}} .$$

By Claim D.2, we conclude that $f_{i-1} \in \text{Post}^*(I) \subseteq \mathcal{I}$. Above, we justified that $e_{j+1} \in \mathcal{J}$. However, we know that $f_j \xrightarrow{t_j} e_{j+1}$. This creates a contradiction as $(\mathcal{I}, \mathcal{J})$ was assumed to be a separator (Condition 7.6 Item (Sep1)).

Case $i \leq j$. Above, we justified that $e_i \in \mathcal{I}$ and $f_j \in \mathcal{J}$. Furthermore, by assumption e_i is not locally bounded in \mathcal{A} , so $e_i \notin L$ and f_j is not locally bounded in \mathcal{A}^R , so $f_j \notin L^R$. However, there is a run from e_i to f_j :

$$e_i \xrightarrow[L]{\pi_i t_i \cdots t_{j-1} \pi_j} f_j.$$

This, of course, implies that $e_i \xrightarrow[\mathbb{Z}]{*} f_j$ holds; a contradiction is created because $(\mathcal{I}, \mathcal{J})$ were assumed to satisfy Condition 7.6 (Sep2).

Therefore, to conclude the second direction of the proof, we have shown that the existence of a non-reachability witness (I, J) for src and trg indeed implies that src cannot reach trg in \mathcal{A} . \blacktriangleleft

► **Theorem 7.12.** *The reachability problem for OCA with disequality tests belongs to the complexity class coNP^{NP} .*

Proof. We will use Theorem 7.9 to prove that non-reachability in OCA with disequality tests is in NP^{NP} . For an instance of (non-)reachability, let \mathcal{A} be an OCA with disequality tests, let src be the initial configuration, and let trg be the target configuration. Suppose that there is no run from src to trg in \mathcal{A} . By Theorem 7.9, there exists a non-reachability witness (I, J) for src to trg such that I and J have concise descriptions. Let $\mathcal{I} := \text{Post}_-(I) \cup \text{Post}(\text{Post}_-(I))$ and $\mathcal{J} := \text{Pre}_+(J) \cup \text{Pre}(\text{Pre}_+(J))$. We now argue that checking that (I, J) satisfies the conditions to be a non-reachability witness can be completed in non-deterministic polynomial time.

1. In polynomial time, one can check whether $\text{src} \in I$, $\text{trg} \in J$, $I \subseteq \text{Conf}_+ \cap L$, and $J \subseteq \text{Conf}_+^R \cap L^R$. Recall that I consists of at most $3|Q| + 1$ arithmetic progressions; suppose $I = I_1 \cup \dots \cup I_k$ where $k \leq 3|Q| + 1$ and, for each $i \in [1, k]$, I_i is an arithmetic progression on some state q . To check $I \subseteq \text{Conf}_+ \cap L$, we can check that $I_i \subseteq \text{Conf}_+$ and $I_i \subseteq L$ for every $i \in [1, k]$. For convenience, suppose (q, u) and (q, v) are the configurations with the smallest and largest counter values in I_i and suppose d is the difference of I_i . First, to check that $I_i \subseteq \text{Conf}_+$, it suffices to check that $(q, u) \in \text{Conf}_+$; this can be done in polynomial time via Lemma 4.3. Second, to check that $I_i \subseteq L$, we will construct an instance of unboundedness in a strongly connected 1-VASS with disequality tests. Let \mathcal{A}_S be the strongly connected component of \mathcal{A} that q belongs to. Add, to \mathcal{A}_S , an additional state q' with a disequality test $x \neq v + d$ and two transitions $(q', 0, q)$ and (q', d, q') , to obtain a new OCA with disequality tests \mathcal{B} . It follows that (q', u) is bounded in \mathcal{B} if and only if $I_i \subseteq L$. This is true because if (q', u) is unbounded in \mathcal{B} , then there is some $(q, x) \in I_i$ that is unbounded in \mathcal{A}_S which means that (q, x) is not locally bounded. Recall that (un)boundedness can be decided in polynomial time (see Lemma 4.1).
2. In non-deterministic polynomial time, one can check whether or not (I, J) are inductive (Condition 7.4). Verifying that $\text{Post}(\text{LPost}_-(I)) \cap \text{Conf}_+ \cap L \not\subseteq I$ requires checking membership in Conf_+ and L , which can be completed in polynomial time by Lemma 4.3 and Lemma 4.1, respectively. The same can be said for verifying $\text{Pre}(\text{LPre}_+(J)) \cap \text{Conf}_+^R \cap L^R \subseteq J$.
3. In non-deterministic polynomial time, one can check whether or not $(\mathcal{I}, \mathcal{J})$ satisfies the first condition of being a separator (Condition 7.6 (Sep1)). First, guess four configurations $c, c', d',$ and d . One can check whether $c \in I$, $d \in J$, whether $c' \in \text{Post}_-(c)$ or $c' \in \text{Post}(\text{Post}_-(c))$ and whether $d' \in \text{Pre}_+(d)$ or $d' \in \text{Pre}(\text{Pre}_+(d))$ in non-deterministic polynomial time (using Lemma 5.1). In which case, in linear time, one can see whether (Sep1) is violated by checking $c' \rightarrow d'$.

4. In non-deterministic polynomial time, one can check whether or not $(\mathcal{I}, \mathcal{J})$ satisfies the second condition of being a separator (Condition 7.6 (Sep2)). Similar to checking the first condition of being a separator, guess for configurations c, c', d' , and d . Like for condition 3, one can check whether $c \in I, d \in J$, whether $c' \in \text{Post}_*(c)$ or $c' \in \text{Post}(\text{Post}_*(c))$ and whether $d' \in \text{Pre}_*(d)$ or $d' \in \text{Pre}(\text{Pre}_*(d))$ in non-deterministic polynomial time (using Lemma 5.1). However, in this case, one can check whether (Sep2) is violated by checking if $c' \notin L$ and $d' \notin L^R$ in polynomial time using Lemma 4.1 as well as checking $c' \xrightarrow[\mathbb{Z}]{*} d'$ in non-deterministic polynomial time (say by using integer linear programming).

Note that checking conditions 3 and 4 requires exponential bounds on the counter values of configurations in \mathcal{I} and \mathcal{J} (see Lemma C.2). ◀