

Shortest paths in one-counter systems^{*}

Dmitry Chistikov¹, Wojciech Czerwiński², Piotr Hofman³, Michał Pilipczuk²,
and Michael Wehar⁴

¹ Max Planck Institute for Software Systems (MPI-SWS),
Kaiserslautern and Saarbrücken, Germany, dch@mpi-sws.org

² Institute of Informatics, University of Warsaw, Poland,
{wczwin,michal.pilipczuk}@mimuw.edu.pl

³ Laboratoire Spécification et Vérification (LSV), ENS Cachan & CNRS, France,
piotr.hofman@lsv.ens-cachan.fr

⁴ Department of Computer Science and Engineering, University at Buffalo, USA,
mwehar@buffalo.edu

Abstract. We show that any one-counter automaton with n states, if its language is non-empty, accepts some word of length at most $O(n^2)$. This closes the gap between the previously known upper bound of $O(n^3)$ and lower bound of $\Omega(n^2)$. More generally, we prove a tight upper bound on the length of shortest paths between arbitrary configurations in one-counter transition systems. Weaker bounds have previously appeared in the literature, and our result offers an improvement.

1 Introduction

Extremal combinatorial questions are a routine in today's theory of computing: How many steps does an algorithm take in the worst case when traversing a data structure? How large is the most compact automaton for a formal language? While some specific questions of this form are best seen as standalone puzzles, only interesting for their own sake, others can be used as basic building blocks for more involved arguments.

We look into the following extremal problem: Given a one-counter automaton \mathcal{A} with n states, how long can the shortest word accepted by \mathcal{A} be? It is folklore that, unless the language of \mathcal{A} is empty, \mathcal{A} accepts some word of length at most polynomial in n . This fact and a number of related results of similar form have appeared as auxiliary lemmas in the literature on formal languages, analysis of infinite-state systems, and applications of formal methods [11, Lemma 6; 10, Section 8.1; 7, Lemma 5; 1, Lemma 11; 8, Lemmas 28 and 29]. A closer inspection

^{*} The main part of these results was obtained during Autobóz'15, the annual research camp of Warsaw automata group. During the work on these results, Mi. Pilipczuk held a post-doc position at Warsaw Centre of Mathematics and Computer Science and was supported by the Foundation for Polish Science via the START stipend programme. P. Hofman was supported by Labex Digicosme, Univ. Paris-Saclay, project VERICONISS. W. Czerwiński acknowledges a partial support by the Polish National Science Centre grant 2013/09/B/ST6/01575.

reveals that the arguments behind these results deliver (or can deliver) an upper bound of $O(n^3)$, while the best known lower bound comes from examples of one-counter automata with shortest accepted words of length $\Theta(n^2)$. In other words, the true value is at least quadratic and at most cubic.

The main result of the present paper is that we close this gap by showing a quadratic upper bound, $O(n^2)$. We also extend this result to a more general reachability setting: in any one-counter (transition) system with n control states, whenever there is a path from a configuration α to a configuration β —recall that configurations are pairs of the form (q, c) where $q \in Q$ is the control state, $|Q| = n$, and c is a counter value, a nonnegative integer—there is also a path from α to β that has length at most $O(n^2 + n \cdot \max(c_\alpha, c_\beta))$ where c_α and c_β are the counter values of α and β . We discuss our contribution in more detail in the following Section 2.

Related work and motivation. Reachability is a fundamental problem in theoretical computer science and in its applications in verification, notably via analysis of infinite-state systems [4; 16; 2; 13]. Among such systems, counter-based models of computation are a standard abstraction in many domains and have attracted a lot of attention [3]; machines with a single counter are, of course, the most basic. Nevertheless, while our main motivation has been purely theoretical, we find it interesting to note that bounds on the length of shortest paths in one-counter systems have appeared as building blocks in the literature on rather diverse topics.

More specifically, a polynomial upper bound is used by Etessami et al. [7] in an analysis of *probabilistic* one-counter systems (which are equivalent to so-called discrete-time quasi-birth-death processes, QBDs). They prove that in the $(q, 1) \rightsquigarrow (q', 0)$ -reachability setting the counter does not need to grow higher than n^2 and provide examples showing that this bound is tight. However, they only deduce upper bounds of n^3 and n^4 on the length of shortest paths without and with zero tests, respectively. A simple corollary shows that if a state q can eventually reach a state q' with a non-zero probability, then this probability is lower-bounded by $p^{\text{poly}(n)}$ where p is the smallest among positive probabilities associated with transitions. This becomes a step in the proof that a (decomposed) Newton’s method approximates *termination probabilities* of the system in time polynomial in its size, n ; the results of the present paper reduce the (theoretical) worst-case upper bounds on the number of steps roughly by a factor of n .

In a subsequent work, Hofman et al. [8] reuse the auxiliary lemmas on the length of shortest paths from [7] and show that (strong and weak) trace inclusion for a one-counter system and a finite-state process is decidable in PSPACE (and is, in fact, PSPACE-complete).

One may note that a stronger upper bound of $O(n^3)$ on the length of shortest paths can be derived from the above bound on the largest needed counter value even in the presence of zero tests. This value, $O(n^3)$, seems to be a recurring theme in the literature on one-counter systems; it already appears in the pumping lemma for one-counter languages due to Latteux [12] as the *pumping constant*: a number N such that any accepted word longer than N can be pumped. In fact,

the formulation in [12] does not permit *removals* of factors from an accepted word, but even such a version would only yield the same upper bound of $O(n^3)$ on the length of shortest paths. While the arguments of the present paper do not lead to an improvement in the pumping constant for one-counter languages, we nevertheless show that in the reachability setting the optimal value (the length of the shortest path) is actually $O(n^2)$.

A cubic upper bound on the largest needed counter value (for the reachability setting) in one-counter systems without zero tests, also known as one-counter nets, appears in the work of Lafourcade et al. [11; 10]. This result is applied in the context of the Dolev-Yao intruder model, where the question of whether a passive eavesdropper (an intruder) can obtain a piece of information is reduced to the decision problem for a deduction system. For several such systems, Lafourcade et al. show that, under certain assumptions, the problem is decidable in polynomial time. They construct a one-counter system where states represent terms from a finite set and the counter value corresponds to the number of applications of a free unary function symbol to a term. After this, the upper bound on counter values along shortest paths is extended to an upper bound on the size of terms that can be used in a minimal deductive proof; needless to say, an improvement in the upper bound extends in a natural way.

Finally, we would like to mention the work of Alur and Černý [1], who use a related model of one-counter systems with counter values in \mathbb{Z} and without zero tests. They reduce the equivalence problem for so-called streaming data-string transducers to $(q, 0) \rightsquigarrow (q', 0)$ -reachability in such counter systems: the transducers produce output at the end of the computation, and the counter is used to track the accumulated distance between a distinguished pair of symbols in the output. Since these transducers are designed to model list-manipulating programs (in two syntactically restricted models), decision procedures for equivalence of such programs can rely on the upper bounds for shortest paths to efficiently prune the search space. In [1], the upper bound on the path length is the familiar $O(n^3)$; this gives an upper bound on the length of smallest counterexamples to equivalence. Our upper bound of $O(n^2)$ extends to this model of counter systems too. The reduction to reachability in one-counter systems was recently implemented by Thakkar et al. [15] on top of ARMC, an abstraction-refinement model checker [14], for the purpose of verifying retransmission protocols over noisy channels.

2 Summary

One-counter systems. In this paper we work in the framework of one-counter systems, which are an abstract version of one-counter automata. More precisely, they are one-counter automata without input alphabet (see below).

Let us first introduce formal definitions. A *one-counter system (OCS)* \mathcal{O} consists of a finite set of states Q , a set of non-zero transitions $T_{>0} \subseteq Q \times \{-1, 0, 1\} \times Q$, and a set of zero tests $T_{=0} \subseteq Q \times \{0, 1\} \times Q$. A *configuration* of the OCS \mathcal{O} is a pair in $Q \times \mathbb{N}$. We define a binary relation \longrightarrow on the set $Q \times \mathbb{N}$

as follows: $(p, c) \longrightarrow (q, c + d)$ whenever (i) $c \geq 1$ and $(p, d, q) \in T_{>0}$ or (ii) $c = 0$ and $(p, d, q) \in T_{=0}$. The reflexive transitive closure of \longrightarrow is denoted by \longrightarrow^* ; we say that a configuration β is *reachable* from α if $\alpha \longrightarrow^* \beta$. This reachability is witnessed by a *path* in OCS \mathcal{O} , which is simply a path in the infinite directed graph with vertices $Q \times \mathbb{N}$ and edge relation \longrightarrow ; vertices and edges along the path can be repeated. The *length* of the path is the number of (not necessarily distinct) edges that occur on it.

Our contribution. We first formulate our results in terms of one-counter systems. Our first result is on reachability between configurations of OCS with zero counter values.

Theorem 1. *Let \mathcal{O} be a one-counter system with n states. Suppose a configuration $\beta = (p_\beta, 0)$ is reachable from a configuration $\alpha = (p_\alpha, 0)$ in \mathcal{O} . Then \mathcal{O} has a path from α to β of length at most $14n^2$.*

Using Theorem 1 as a black-box, we generalize it to the case where source and target configurations have arbitrary counter values.

Theorem 2. *Let \mathcal{O} be a one-counter system with n states. Suppose a configuration $\beta = (p_\beta, c_\beta)$ is reachable from a configuration $\alpha = (p_\alpha, c_\alpha)$ in \mathcal{O} . Then \mathcal{O} has a path from α to β of length at most $14n^2 + n \cdot \max(c_\alpha, c_\beta)$.*

The proof of Theorem 1 is the main technical contribution of this work. We prove Theorem 1 in Section 5 and Theorem 2, as well as an extension to OCS with negative counter values, in Section 6.

One-counter automata. We now restate our contribution in terms of one-counter automata (which are the original motivation for this work).

Take any finite set Σ . The set of all finite words over Σ is denoted by Σ^* , and the empty word by ε . A (*nondeterministic*) *one-counter automaton* \mathcal{A} over the input alphabet Σ is a one-counter system where every transition $t \in T_{>0} \cup T_{=0}$ is associated with a label, $\lambda(t) \in \Sigma \cup \{\varepsilon\}$, and where some subsets $I \subseteq Q$ and $F \subseteq Q$ are distinguished as sets of initial and final states respectively. The labeling function λ is extended from transitions to edges \longrightarrow and to paths in a natural way; the automaton *accepts* all words that are labels of paths from $I \times \{0\}$ to $F \times \mathbb{N}$. The *language* of a one-counter automaton \mathcal{A} is the set $L \subseteq \Sigma^*$ of all words accepted by \mathcal{A} .

Corollary 1. *Let \mathcal{A} be a nondeterministic one-counter automaton with n states. If the language of \mathcal{A} is non-empty, then \mathcal{A} accepts some word of length at most $14n^2$.*

Proof. Take \mathcal{A} with a non-empty language and add self-loops with decrements to all final states: $(p, -1, p) \in T_{>0}$ for $p \in F$. Since the language of \mathcal{A} is non-empty, some final configuration (in $F \times \mathbb{N}$) is reachable from some initial configuration (from $I \times \{0\}$); this implies that in the modified automaton, denoted by \mathcal{A}' , a configuration $\beta = (p_\beta, 0)$, $p_\beta \in F$, is reachable from a configuration $\alpha = (p_\alpha, 0)$, $p_\alpha \in I$. Consider the shortest path in \mathcal{A}' between α and β : by Theorem 1, its

length is at most $14n^2$. Take the shortest prefix of this path whose target has state from F ; this path is a path in \mathcal{A} . Since the label of the path cannot be longer than the path itself, the result follows. \square

As a concrete example, one can deduce from Corollary 1 that any nondeterministic one-counter automaton that accepts the singleton unary language $\{a^n\}$ —a basic version of counting to n —must have at least $\Omega(\sqrt{n})$ states. This lower bound is tight and shows that nondeterminism does not help to “count to n ”, because *deterministic* one-counter automata can also do this using $\Theta(\sqrt{n})$ states [5].

Lower bounds. As we already said, the lower bound on the length of the shortest path is $\Omega(n^2)$. We present constructions of OCS that match the upper bounds of Theorems 1 and 2. Note that Examples 1 and 2 seem to use different phenomena; they previously appeared in [7; 5].

Example 1. Consider an OCS \mathcal{O}_1 with $2n$ states: p_1, \dots, p_n and q_1, \dots, q_n . Let \mathcal{O}_1 have, for $1 \leq i < n$, transitions $(p_i, +1, p_{i+1})$ and $(q_i, 0, q_{i+1})$, as well as $(q_n, -1, q_1)$ and $(p_n, 0, q_1)$. All the transitions are non-zero, except for transition $(p_1, +1, p_2)$, which is a zero test. This OCS is deterministic: every configuration has at most one outgoing transition. The only path from $(p_1, 0)$ to $(q_1, 0)$ has length n^2 .

Example 2. Let k and m be coprime and let OCS \mathcal{O}'_2 have states p_0, \dots, p_{k-1} , q_0, \dots, q_{m-1} , and s_1, s_2 . Let \mathcal{O}'_2 have, for all $0 \leq i < k$ and $0 \leq j < m$, non-zero transitions $(p_i, +1, p_{i+1 \bmod k})$ and $(q_j, -1, q_{j+1 \bmod m})$, non-zero transitions $(p_0, 0, q_0)$, $(q_{m-1}, -1, s_1)$, and zero tests $(p_0, +1, p_1)$ and $(s_1, 0, s_2)$. Now paths from $(p_1, 0)$ to $(s_2, 0)$ correspond to solutions of $x \cdot k - (y \cdot m - 1) = 1$; the shortest path takes the first cycle $x = m$ times and the second cycle $y = k$ times and thus has length $2km + 2$. Setting $k = n$ and $m = n - 1$ gives an OCS \mathcal{O}_2 with $2n + 1$ states where not only does the shortest path have quadratic length, but all such paths also need to use quadratic counter values.

Example 3. This example justifies the need for the term $n \cdot \max(c_\alpha, c_\beta)$ in Theorem 2. Modify \mathcal{O}_1 from Example 1 as follows. Add states a_1, \dots, a_n , b_1, \dots, b_n and the following non-zero transitions: $(a_n, -1, a_1)$, $(b_n, +1, b_1)$, and, for all $0 \leq i < n$, $(a_i, 0, a_{i+1})$ and $(b_i, 0, b_{i+1})$. For each of these non-zero transition, apart from $(a_n, -1, a_1)$, introduce also the same transition as a zero test. Finally, add two more zero tests: $(a_n, 0, p_1)$ and $(q_1, 0, b_1)$. Thus, the obtained OCS \mathcal{O}_3 has $4n$ states. Observe that every path in \mathcal{O}_3 from (a_1, c_α) to (b_n, c_β) has to go through $(a_n, 0)$ and $(b_1, 0)$ and thus has length at least $n^2 + n(c_\alpha + c_\beta + 2)$.

3 Challenges and techniques

We now discuss shortly the intuition behind our approach to proving Theorem 1, and where the main challenges lie.

The first, obvious observation is as follows: if some configuration appears more than once on a path, then the segment between any two appearances of this

configuration can safely be removed. If we apply this modification exhaustively, then on each “level” — a set of configurations with the same counter value — we cannot see more than n configurations. If the maximum counter value observed on some path were bounded by $O(n)$, then we would immediately obtain a quadratic upper bound on its length. Unfortunately, this is not the case: as Example 2 shows, the counter values in the shortest accepting path can be as large as quadratic. Hence, applying this observation in a straightforward manner cannot lead to any upper bound better than cubic.

Instead, we perform an involved surgery on the path. The first idea is to start with a path ρ_\circ that is not the shortest, but uses the fewest zero tests; the observation above shows that their number is bounded by n . Each subpath between two consecutive zero tests is called an *arc*, and we aim at modifying each arc separately to make it short. An arc is called *low* if it contains only configurations with counter values at most $5n$, and *high* otherwise. The total length of low arcs can again be bounded by $O(n^2)$ by just excluding repeated configurations, so it suffices to focus on high arcs.

Suppose ρ is a high arc. Since we observe high counter values on ρ , one can easily find a *positive cycle* σ^+ in the early parts of ρ , and a *negative cycle* σ^- in the late parts of ρ . Here by a cycle we mean a sequence of transitions that starts and ends in the same state, and the cycle is positive/negative if the total effect it has on the counter during its traversal is positive/negative. Let A be the (positive) effect of σ^+ on the counter, and $-B$ be the (negative) effect of σ^- .

Now comes the crucial idea of the proof: we can modify ρ by pumping σ^+ and σ^- up many times, thus effectively “lifting” the central part of the path (called *cap*) to counter levels where there is no threat of hitting counter value zero while performing modifications (see Figure 1, p. 11). More importantly, the cap can now be unpumped “modulo $\gcd(A, B)$ ” in the following sense: we can exhaustively remove subpaths between configurations that have the same state and whose counter values are congruent modulo $\gcd(A, B)$. The reason is that any change in the total effect of the cap on the counter that is divisible by $\gcd(A, B)$ can be compensated by adjusting the number of times we pump cycles σ^+ and σ^- . In particular, the length of the cap becomes reduced to at most $\gcd(A, B) \cdot n$, at the cost of pumping σ^+ and σ^- several times.

By performing this operation (we call it *normalization*) on all high arcs, we make them *normal*. After this, we apply an involved amortization scheme to show that the total length of normal arcs is at most quadratic in n . This requires very delicate arguments for bounding (i) the total length of the caps and (ii) the total length of the pumped cycles σ^+ and σ^- throughout all the normal arcs. In particular, for this part of the proof to work we need to assert a number of technical properties of normal arcs; we ensure that these properties hold when we perform the normalization. Most importantly, whenever for two arcs the corresponding cycles σ^+ (or σ^-) lie in the same strongly connected component of the system (looking at the graph induced only by non-zero transitions), we stipulate that in both arcs σ^+ (or σ^-) actually refer to the same cycle. The final

amortization is based on the analysis of pairs of strongly connected components to which σ^+ and σ^- belong, for all normal arcs.

At least as of now, arguments of this flavor (inspired by *amortized analysis* reasoning) are not typical for formal language theory and are more characteristic of the body of work on algorithms and data structures; see, e.g., [9; 6].

4 Preliminaries

In this paper \mathbb{N} stands for the set of nonnegative integers. For any set X and a word $w \in X^*$, the *length* of $w = x_1 \dots x_n$, denoted $\text{LEN}(w)$, is the number n of symbols in w . For $k \in \mathbb{N}$ and a word w , by w^k we denote the word w repeated k times. For two positive integers x, y , by $\text{gcd}(x, y)$ and $\text{lcm}(x, y)$ we denote the greatest common divisor and the least common multiple of x and y , respectively. Recall that $x \cdot y = \text{gcd}(x, y) \cdot \text{lcm}(x, y)$.

We now give all definitions related to one-counter systems that we need in the sequel. For the reader's convenience, concepts that appeared in Section 2 are defined anew.

A *one-counter system* (OCS) \mathcal{O} consists of a finite set of *states* Q , a set of *non-zero transitions* $T_{>0} \subseteq Q \times \{-1, 0, 1\} \times Q$, and a set of *zero tests* $T_{=0} \subseteq Q \times \{0, 1\} \times Q$. The set of *transitions* is $T = T_{>0} \cup T_{=0}$. For a transition $t = (q, d, q') \in T$, by $\text{SRC}(t)$ and $\text{TARG}(t)$ we denote q and q' , i.e., the source and the target state of t respectively. Further, the *effect* of the transition $t = (q, d, q')$ is the number d ; we write $\text{EFF}(t) = d$. We extend this notion to sequences of transitions: $\text{EFF}(t_1 \dots t_m) = \sum_{i=1}^m \text{EFF}(t_i)$.

A *configuration* of the OCS \mathcal{O} is a pair in $Q \times \mathbb{N}$. The *state of a configuration* (q, c) is the state q ; we also say that configuration (q, c) *has state* q , and write $\text{ST}((q, c)) = q$. The *counter value* of configuration (q, c) is the number c ; we write $\text{CNT}((q, c)) = c$.

A transition $t = (q, d, q') \in T$ can be *fired* in a configuration $\gamma = (q, c)$ if either $t \in T_{>0}$ and $c > 0$ or $t \in T_{=0}$ and $c = 0$. In other words, zero tests can be fired only if the counter value is zero, and non-zero transitions can be fired only if the counter value is positive. The *result* of firing (q, d, q') in (q, c) is the configuration $\gamma' = (q', c + d)$. We then write $\gamma \xrightarrow{t} \gamma'$.

A *path* ρ of the OCS \mathcal{O} is a sequence of pairs

$$(\gamma_1, t_1)(\gamma_2, t_2) \dots (\gamma_m, t_m) \in ((Q \times \mathbb{N}) \times T)^*$$

such that for every $i \in \{1, \dots, m-1\}$ we have $\gamma_i \xrightarrow{t_i} \gamma_{i+1}$ and there exists a configuration γ_{m+1} such that $\gamma_m \xrightarrow{t_m} \gamma_{m+1}$. The *length* of this path is m . The *source* of ρ , denoted by $\text{SRC}(\rho)$, is γ_1 ; we also say that ρ *starts* in its source. Similarly, the *target* of ρ , denoted by $\text{TARG}(\rho)$, is γ_{m+1} ; we say that ρ *finishes* in its target. Note that now the source and target are configurations, rather than individual states; the path is *from* its source *to* its target. All $\gamma_2, \dots, \gamma_m$ are called *intermediate configurations*. We also say that configurations $\gamma_1, \gamma_2, \dots, \gamma_{m+1}$ *appear* on

ρ ; note that the target of ρ also appears on ρ . Finally, when such a path exists, the configuration γ_{m+1} is said to be *reachable* from the configuration γ_1 .

The *projection* of a path ρ is the sequence of its transitions $t_1 t_2 \dots t_m$; we write $\text{PROJ}(\rho) = t_1 t_2 \dots t_m$. We follow the convention of denoting paths by ρ and sequences of transitions by σ . The *effect* of a path ρ is $\text{EFF}(\rho) = \text{EFF}(\text{PROJ}(\rho))$. A sequence of transitions $\sigma = t_1 t_2 \dots t_m$ is *fireable* in a configuration γ_1 if there exists a path $\rho = (\gamma_1, t_1)(\gamma_2, t_2) \dots (\gamma_m, t_m)$. This path ρ is called the *fastening* of σ at γ_1 , denoted $\rho = \text{FASTEN}(\gamma_1, \sigma)$. Note that in particular $\text{PROJ}(\text{FASTEN}(\gamma, \sigma)) = \sigma$ for every γ in which σ is fireable.

A sequence of transitions $t_1 t_2 \dots t_m$ is *consistent* if for all $i \in \{1, \dots, m-1\}$ it holds that $\text{TARG}(t_i) = \text{SRC}(t_{i+1})$. Note that a sequence of transitions fireable in some configuration is always consistent, but the other implication does not hold in general. We extend the notation $\text{SRC}(\cdot)$ and $\text{TARG}(\cdot)$ to consistent sequences of transitions: $\text{SRC}(t_1 t_2 \dots t_m) = \text{SRC}(t_1)$ and $\text{TARG}(t_1 t_2 \dots t_m) = \text{TARG}(t_m)$. The sources and targets of the transitions of $t_1 t_2 \dots t_m$ are *visited* on $t_1 t_2 \dots t_m$.

A *cycle* σ is a consistent sequence of non-zero transitions that starts and finishes in the same state q . This q is called the *base state of the cycle* σ . If the effect of σ is positive (resp. negative), then it is a *positive* (resp. *negative*) cycle. A cycle σ is called *simple* if every state is visited at most once on σ , except for the base of σ , which is visited only at the start and at the end.

5 Proof of Theorem 1

5.1 Proof overview and notation

Let us fix the OCS \mathcal{O} we work with; let Q be its state set and let $n = |Q|$. Suppose ρ_0 is a path from α to β , and let ρ_0 be chosen such that it has the smallest possible number of configurations with counter value zero. Note that ρ_0 does not have to be the shortest path between α and β . The first step is to divide ρ_0 into subpaths, called *arcs*, between consecutive configurations with counter value zero. Then we modify the arcs separately. If a counter value in an arc does not exceed $5n$, then we say that the arc is *low*, otherwise it is *high*. The low arcs will not be changed at all, and the reason is that we can bound quadratically the total number of configurations with counter value at most $5n$ using the following straightforward proposition. It is similar, in the spirit, to pumping lemmas, but simply removes a part of the path.

Proposition 1. *Suppose $\rho = (\gamma_1, t_1)(\gamma_2, t_2) \dots (\gamma_m, t_m)$ is a path from α to β . Suppose further that for some i and j with $1 \leq i < j \leq m+1$ it holds that $\gamma_i = \gamma_j$, where γ_{m+1} is such that $\gamma_m \xrightarrow{t_m} \gamma_{m+1}$. Consider*

$$\rho' = (\gamma_1, t_1)(\gamma_2, t_2) \dots (\gamma_{i-1}, t_{i-1})(\gamma_j, t_j)(\gamma_{j+1}, t_{j+1}) \dots (\gamma_m, t_m).$$

Then ρ' is also a path from α to β .

However, the high arcs will be heavily modified. Roughly speaking, if an arc is high, then it contains both a positive cycle near its beginning and a negative

cycle near its end. We can use these cycles to pump the middle part of the path as much up as we like. Thus, the modified path will consist of a short prefix; then several iterations of a positive cycle pumping it up; then a so called *cap*: a part of the path with only high counter values; then several iterations of a negative cycle pumping it down; and finally a short suffix. We show in the sequel how to perform this construction in such a way that the total length of pumping cycles, short prefixes and suffixes, and caps is quadratic. The construction itself (with arc-level length estimates) is presented in the following subsection 5.2, and the upper bound on the length of the entire path is given in subsection 5.3.

Transition multigraph. One can view a transition $(p, c, q) \in Q \times \{-1, 0, 1\} \times Q$ also as an edge $(p, q) \in Q \times Q$ labelled by a number $c \in \{-1, 0, 1\}$. In the proof we will many times switch back and forth between these two perspectives. In order to keep the mathematical precision we introduce a bit of notation.

The *transition multigraph* $G = (V, E, \ell)$ of an OCS consists of a set of nodes V , a multiset of directed edges E , and a labeling $\ell : E \rightarrow \{-1, 0, 1\}$. Set V equals the set of states Q . Every non-zero transition $t = (p, c, q) \in T_{>0}$ in the OCS \mathcal{O} gives rise to an edge $e = (u, v) \in E$ with $\ell(e) = c$. Note that the definition of the transition multigraph does not take into account the zero transitions.

In the proof we pay a special attention to strongly connected components (SCCs) of G . Recall that two vertices $p, q \in V$ are said to *communicate* if G has a walk from p to q and a walk from q to p . Communication is an equivalence relation, and its equivalence classes are called the *strongly connected components* of G . Let \mathfrak{S} be the set of all strongly connected components of G . For a strongly connected component $S \in \mathfrak{S}$, by n_S we denote the number of vertices in S . We say that a cycle σ is *contained* in S if each state appearing on σ belongs to S . Note that every cycle is contained in some SCC, and a simple cycle contained in S has length at most n_S . We say that an SCC S is *positively enabled* if it contains a cycle that has a positive effect. Similarly, S is *negatively enabled* if it contains a cycle that has a negative effect. Note that an SCC S can be both positively and negatively enabled.

Lemma 1. *Let G be a transition multigraph of an OCS and S a positively (respectively, negatively) enabled SCC. Then there exists a positive (respectively, negative) cycle σ contained in S that is simple.*

Proof. We prove the lemma for positively enabled SCCs; the proof for negatively enabled SCCs is symmetric. By definition, S contains a positive cycle σ . Choose σ to be the shortest such cycle; we claim that then σ is simple. Aiming towards a contradiction, suppose that some state repeats on σ . Then σ can be decomposed into two cycles σ_1, σ_2 that are strictly shorter than σ . Since σ is positive and $\text{EFF}(\sigma) = \text{EFF}(\sigma_1) + \text{EFF}(\sigma_2)$, we infer that either σ_1 or σ_2 is positive. This is a contradiction with the minimality of σ . \square

For every positively enabled SCC S we distinguish one, arbitrarily chosen, simple cycle with positive effect contained in S ; we denote it by σ_S^+ . Its existence is guaranteed by Lemma 1. Similarly, for every negatively enabled S we

distinguish one simple cycle with negative effect contained in S , and we denote it by $\sigma_{\bar{S}}$. The base states of these cycles are chosen arbitrarily.

5.2 Normal paths

A path is an *arc* if both its source and target have counter value zero, but all its intermediate configurations have counter values strictly larger than zero. An arc (or a path) is *low* if all its configurations (including the target) have counter values strictly smaller than $5n$. An arc ρ is (S, T) -*normal*, where S and T are some SCCs of the transition multigraph, if it admits the following *normal decomposition* (see Figure 1, p. 11):

$$\rho = \rho_{\text{PREF}} \rho_{\text{UP}} \rho_{\text{CAP}} \rho_{\text{DOWN}} \rho_{\text{SUFF}},$$

where

- ρ_{PREF} and ρ_{SUFF} are low;
- $\text{PROJ}(\rho_{\text{UP}}) = (\sigma_{\text{UP}})^k$ for some $k \in \mathbb{N}$, where $\sigma_{\text{UP}} = \sigma_S^+$;
- $\text{PROJ}(\rho_{\text{DOWN}}) = (\sigma_{\text{DOWN}})^\ell$ for some $\ell \in \mathbb{N}$, where $\sigma_{\text{DOWN}} = \sigma_T^-$;
- $\text{ST}(\text{SRC}(\rho_{\text{CAP}}))$ is the base state of σ_{UP} ; and
- $\text{ST}(\text{TARG}(\rho_{\text{CAP}}))$ is the base state of σ_{DOWN} .

We say that an arc ρ is *normal* if it is (S, T) -normal for some $S, T \in \mathfrak{S}$. Then a path ρ' is *normal* if it is a concatenation of normal arcs (possibly for different pairs (S, T)) and low arcs.

In the remaining part of the proof we will show that if β is reachable from α , where $\text{CNT}(\alpha) = \text{CNT}(\beta) = 0$, then there exists a short normal path from α to β . We start by analyzing a single arc. The following lemma, which is the most technically involved step in this paper, shows that we can restrict ourselves to normal arcs that have a very special structure.

Lemma 2. *If $\text{CNT}(\alpha) = \text{CNT}(\beta) = 0$ and there exists an arc from α to β , then there exists an arc ρ from α to β which is either low or normal. Moreover, in the case when ρ is normal, a normal decomposition $\rho = \rho_{\text{PREF}} \rho_{\text{UP}} \rho_{\text{CAP}} \rho_{\text{DOWN}} \rho_{\text{SUFF}}$ can be chosen such that:*

- (i) $\text{PROJ}(\rho_{\text{UP}}) = (\sigma_{\text{UP}})^a$, $\text{EFF}(\sigma_{\text{UP}}) = A$ for some $a, A \in \mathbb{N}$;
- (ii) $\text{PROJ}(\rho_{\text{DOWN}}) = (\sigma_{\text{DOWN}})^b$, $\text{EFF}(\sigma_{\text{DOWN}}) = -B$ for some $b, B \in \mathbb{N}$;
- (iii) $a \cdot A \leq 2 \cdot \text{LEN}(\rho_{\text{CAP}}) + 2 \cdot \text{lcm}(A, B)$;
- (iv) $b \cdot B \leq 2 \cdot \text{LEN}(\rho_{\text{CAP}}) + 2 \cdot \text{lcm}(A, B)$;
- (v) no infix of $\text{PROJ}(\rho_{\text{CAP}})$ is a cycle with effect divisible by $\text{gcd}(A, B)$;
- (vi) $\text{CNT}(\text{TARG}(\rho_{\text{UP}})), \text{CNT}(\text{SRC}(\rho_{\text{DOWN}})) > n$; and
- (vii) all configurations appearing on ρ_{PREF} and ρ_{SUFF} are pairwise different.

We now explain some intuition behind this statement. First note that, by condition (vii), the total number of configurations appearing on ρ_{PREF} and ρ_{SUFF} is at most $5n \cdot n$, since n is the number of states of the OCS \mathcal{O} and both of these paths are low (so counter values $5n$ and above do not occur). Thus,

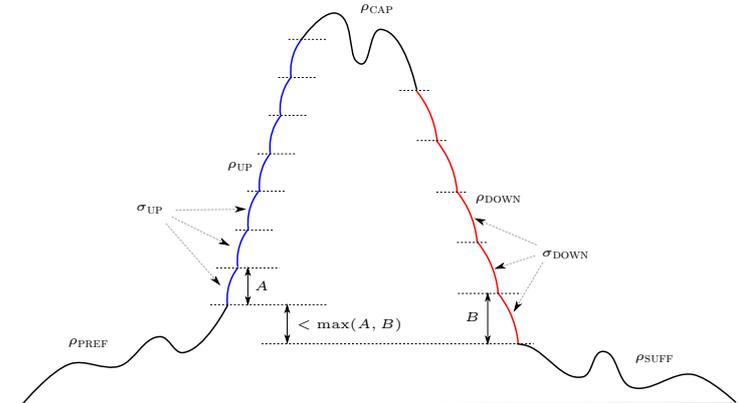


Fig. 1. The normal decomposition of an arc together with some notation from the proof of Lemma 2.

$\text{LEN}(\rho_{\text{PREF}}) + \text{LEN}(\rho_{\text{SUFF}}) \leq 5n^2$. Second, we can conclude from condition (v) that every state $q \in Q$ can occur in configurations appearing in ρ_{CAP} at most $\text{gcd}(A, B)$ times; hence, $\text{LEN}(\rho_{\text{CAP}}) \leq n \cdot \text{gcd}(A, B) \leq n^2$. Finally, condition (i) implies $\text{LEN}(\rho_{\text{UP}}) \leq a \cdot n$; if, for instance, $a \leq \text{const} \cdot n$, then $\text{LEN}(\rho_{\text{UP}}) \leq \text{const} \cdot n^2$; similarly, $\text{LEN}(\rho_{\text{DOWN}}) \leq \text{const} \cdot n^2$. Combined together, these bounds would in this case show that $\text{LEN}(\rho)$ is at most quadratic in n .

However, this reasoning would be insufficient for our purposes, since the number of normal arcs itself can be linear in n . This motivates more subtle upper bounds (iii) and (iv) and the fine-grained choice of parameter in (v). We show how to use Lemma 2 to obtain a quadratic upper bound on the size of the *entire* path in the following subsection 5.3; the remainder of the present subsection proves Lemma 2.

Proof. Fix configurations α and β such that $\text{CNT}(\alpha) = \text{CNT}(\beta) = 0$ and there exists an arc from α to β . If there is a low arc from α to β , then there is nothing to prove, so assume that all the arcs from α to β are not low. Let ρ_{\circ} be such an arc of the shortest possible length; then ρ_{\circ} is not low. Let

$$\rho_{\circ} = (\gamma_1, t_1) \dots (\gamma_m, t_m),$$

where $\alpha = \gamma_1$ and $\gamma_m \xrightarrow{t_m} \gamma_{m+1} = \beta$. Since ρ_{\circ} is shortest possible, from Proposition 1 we infer that configurations $\gamma_1, \gamma_2, \dots, \gamma_{m+1}$ are pairwise different.

We start with a short overview. Based on ρ_{\circ} we construct a normal arc ρ from α to β satisfying the promised conditions. Roughly speaking we proceed as follows. First, we carefully define ρ_{PREF} and ρ_{SUFF} so that condition (vii) is satisfied; in this step we also fix the components $S, T \in \mathfrak{S}$ for which ρ will be (S, T) -normal. Then we construct a sequence of transitions σ_{CAP} that, after fastening it at some configuration, will form ρ_{CAP} that satisfies condition (v). Intuitively, σ_{CAP} is formed by exhaustively unpumping the middle part of ρ_{\circ} .

As S, T are already fixed, so are also cycles $\sigma_{\text{UP}} = \sigma_S^+$ and $\sigma_{\text{DOWN}} = \sigma_T^-$. Hence at this point to completely define ρ it remains to choose numbers a and b . At the end we show that a, b can be chosen so that ρ is indeed a valid path, and moreover conditions (iii) and (iv) are satisfied.

Let us carry out this plan. Consider any k with $2n \leq k \leq 3n$. Let i_k be the smallest index for which $\text{CNT}(\gamma_{i_k}) = k$, and let j_k be the largest index for which $\text{CNT}(\gamma_{j_k}) = k$. Clearly such configurations exist, because ρ_o is not low. It moreover holds that $i_{2n} < i_{2n+1} < \dots < i_{3n} < j_{3n} < \dots < j_{2n+1} < j_{2n}$. By the pigeonhole principle there exist indices k, ℓ , where $k < \ell$, such that the state of γ_{i_k} equals the state of γ_{i_ℓ} . Let this state be $p \in Q$. Consider the sequence of transitions

$$\sigma_{\text{CYC}} = t_{i_k} t_{i_{k+1}} \dots t_{i_\ell-1}.$$

It follows that σ_{CYC} is a positive cycle with effect $\ell - k$ and base state p . Let S be the SCC of G in which σ_{CYC} is contained; the existence of σ_{CYC} asserts that S is positively enabled. Let $\sigma_{\text{UP}} = \sigma_S^+$.

Let $\tilde{\rho}_{\text{PREF}}$ be the prefix of ρ_o up to configuration γ_{i_k} (i.e., with $\text{TARG}(\tilde{\rho}_{\text{PREF}}) = \gamma_{i_k}$). Note that we cannot simply put $\rho_{\text{PREF}} = \tilde{\rho}_{\text{PREF}}$, because the state p in which $\tilde{\rho}_{\text{PREF}}$ finishes does not have to be the base state of σ_{UP} , which is the cycle that is required to be the one used for constructing ρ_{UP} . This, however, poses no real difficulty, because p and σ_{UP} are contained in the same SCC S , so we can easily augment $\tilde{\rho}_{\text{PREF}}$ by a path to the base state of σ_{UP} as follows.

Precisely we do the following. Let q be the base state of σ_{UP} . As both p and q belong to S , there exist consistent sequences of non-zero transitions σ_{pq} and σ_{qp} , leading from p to q and from q to p , respectively, such that:

- (a) states visited on σ_{pq} are pairwise different, and the same holds also for σ_{qp} ; in particular $\text{LEN}(\sigma_{pq}), \text{LEN}(\sigma_{qp}) < n_S$;
- (b) σ_{pq} is fireable in any configuration (p, c) for any $c \geq n$; and
- (c) σ_{qp} is fireable in any configuration (q, c) for any $c \geq n$.

Assertion (a) follows from the fact that σ_{pq} and σ_{qp} can be chosen so that they correspond to simple paths in G , i.e., walks with no state repeated. Assertion (a) in particular implies that the effects of prefixes of σ_{pq} and σ_{qp} are strictly larger than $-n$. This implies assertions (b) and (c).

Now we construct a path ρ''_{PREF} as follows. Let ρ_{pq} be the fastening of σ_{pq} at the configuration γ_{i_k} . The state of γ_{i_k} is p and its counter value is not smaller than $2n$, so indeed σ_{pq} is fireable from γ_{i_k} ; even more, since $\text{LEN}(\sigma_{pq}) < n$ and $\text{CNT}(\gamma_{i_k}) \geq 2n$, all the counter values on ρ_{pq} are larger than n . We define then

$$\rho''_{\text{PREF}} = \tilde{\rho}_{\text{PREF}} \rho_{pq} = (\gamma_1, t_1) \dots (\gamma_{i_k-1}, t_{i_k-1}) \rho_{pq}.$$

We construct path ρ''_{SUFF} in a completely symmetric manner, so we only make a short summary in order to introduce the notation. By the pigeonhole principle, for some $\bar{\ell}, \bar{k}$ with $\bar{\ell} < \bar{k}$ the state of $\gamma_{j_{\bar{\ell}}}$ and $\gamma_{j_{\bar{k}}}$ is the same, let it be \bar{p} . The part of the path between $\gamma_{j_{\bar{\ell}}}$ and $\gamma_{j_{\bar{k}}}$ projects to a negative cycle, so it is contained in some negatively enabled SCC T , to which \bar{p} also belongs. Define $\sigma_{\text{DOWN}} = \sigma_T^-$,

and let \bar{q} be the base state of σ_{DOWN} . As \bar{p} and \bar{q} both belong to T , we have $\sigma_{\bar{q}\bar{p}}$ and $\sigma_{\bar{p}\bar{q}}$ with similar properties as σ_{pq} and σ_{qp} . Path $\rho_{\bar{q}\bar{p}}$ can be again defined as an appropriate fastening of $\sigma_{\bar{q}\bar{p}}$, so we define

$$\rho''_{\text{SUFF}} = \rho_{\bar{q}\bar{p}}(\gamma_{j_{\bar{k}}}, t_{j_{\bar{k}}}) \cdots (\gamma_m, t_m).$$

Let $A = \text{EFF}(\sigma_{\text{UP}})$ and $B = -\text{EFF}(\sigma_{\text{DOWN}})$. Now, based on ρ''_{PREF} and ρ''_{SUFF} we define ρ'_{PREF} and ρ'_{SUFF} as follows. Observe that $\text{CNT}(\text{TARG}(\rho''_{\text{PREF}})) = k + \text{EFF}(\sigma_{pq}) > 2n - n = n$, and similarly $\text{CNT}(\text{SRC}(\rho''_{\text{SUFF}})) > n$. Suppose first that $\text{CNT}(\text{TARG}(\rho''_{\text{PREF}})) \leq \text{CNT}(\text{SRC}(\rho''_{\text{SUFF}})) - A$. Then we take $\rho'_{\text{SUFF}} = \rho''_{\text{SUFF}}$, whereas ρ'_{PREF} is obtained from ρ''_{PREF} by appending the cycle σ_{UP} a number of times so that $\text{CNT}(\text{SRC}(\rho'_{\text{SUFF}})) - A < \text{CNT}(\text{TARG}(\rho'_{\text{PREF}})) \leq \text{CNT}(\text{SRC}(\rho'_{\text{SUFF}}))$. Similarly, if $\text{CNT}(\text{TARG}(\rho''_{\text{PREF}})) \geq \text{CNT}(\text{SRC}(\rho''_{\text{SUFF}})) + B$, then we take $\rho'_{\text{PREF}} = \rho''_{\text{PREF}}$ whereas ρ'_{SUFF} is constructed from ρ''_{SUFF} by appending σ_{DOWN} a number of times in the front so that $\text{CNT}(\text{TARG}(\rho'_{\text{PREF}})) - B < \text{CNT}(\text{SRC}(\rho'_{\text{SUFF}})) \leq \text{CNT}(\text{TARG}(\rho'_{\text{PREF}}))$. If none of these cases holds, we simply take $\rho'_{\text{PREF}} = \rho''_{\text{PREF}}$ and $\rho'_{\text{SUFF}} = \rho''_{\text{SUFF}}$. Since $\sigma_{\text{UP}}, \sigma_{\text{DOWN}}$ have lengths at most n , and the first one is a positive cycle whereas the second one is a negative cycle, it can be easily verified that ρ'_{PREF} and ρ'_{SUFF} are indeed valid paths; here we use the property that $\text{CNT}(\text{TARG}(\rho''_{\text{PREF}})) > n$ and $\text{CNT}(\text{SRC}(\rho''_{\text{SUFF}})) > n$ in order to make sure that appending the cycles does not create nonpositive counter values on the path. Moreover, we achieved the property that $|\text{CNT}(\text{TARG}(\rho'_{\text{PREF}})) - \text{CNT}(\text{SRC}(\rho'_{\text{SUFF}}))| < \max(A, B)$.

Finally, we obtain ρ_{PREF} by applying Proposition 1 to ρ'_{PREF} exhaustively. In this manner ρ_{PREF} has still the same source and target as ρ'_{PREF} , but no configuration repeats on ρ_{PREF} . Similarly, ρ_{SUFF} is obtained from ρ'_{SUFF} by applying Proposition 1 exhaustively, so that no configuration repeats on ρ_{SUFF} .

Let $\zeta = \text{TARG}(\rho_{\text{PREF}}) = \text{TARG}(\rho'_{\text{PREF}})$ and $\bar{\zeta} = \text{SRC}(\rho_{\text{SUFF}}) = \text{SRC}(\rho'_{\text{SUFF}})$. We now verify that ρ_{PREF} and ρ_{SUFF} are as required.

Claim 1. *The following conditions hold:*

- (1.a) *paths ρ_{PREF} and ρ_{SUFF} are low;*
- (1.b) *the counter values in configurations appearing on ρ_{PREF} and ρ_{SUFF} are always positive, apart from the source of ρ_{PREF} (which is α) and the target of ρ_{SUFF} (which is β);*
- (1.c) $\text{CNT}(\zeta), \text{CNT}(\bar{\zeta}) > n$;
- (1.d) $|\text{EFF}(\rho_{\text{PREF}}) + \text{EFF}(\rho_{\text{SUFF}})| < \max(A, B)$;
- (1.e) *property (vii) is satisfied.*

Proof. By the definition of k , all the configurations appearing on $\tilde{\rho}_{\text{PREF}}$ have counter values at most $3n$. Since the counter value of configuration γ_{i_k} is not smaller than $2n$ and not larger than $3n$, and $|\text{EFF}(\sigma_{pq})| \leq \text{LEN}(\sigma_{pq}) < n$, we infer that the counter value on the path ρ''_{PREF} is always strictly smaller than $4n$. As $\text{LEN}(\sigma_{\text{UP}}) < n$, it can be easily seen that appending the cycles during the construction of ρ'_{PREF} cannot create counter values larger than $5n - 1$. Hence ρ'_{PREF} is low, and consequently ρ_{PREF} is also low. A symmetric reasoning shows the same conclusions for ρ_{SUFF} , and thus condition (1.a) is satisfied.

Since ρ_\circ is an arc, no configuration on $\tilde{\rho}_{\text{PREF}}$ apart from α has nonpositive counter value. As $\text{CNT}(\gamma_{i_k}) \geq 2n$, we have already argued that both after adding σ_{pq} when constructing ρ''_{PREF} , and after adding cycles σ_{UP} when constructing ρ'_{PREF} , we could not obtain a configuration with a nonpositive counter value. Hence the only configuration on ρ'_{PREF} that has zero counter value is α , and the same holds also for ρ_{PREF} . A symmetric reasoning yields a symmetric conclusion for ρ_{SUFF} , which proves condition (1.b).

From the construction we have $\text{CNT}(\zeta) \geq \text{CNT}(\text{TARG}(\rho''_{\text{PREF}}))$, and we already argued that $\text{CNT}(\text{TARG}(\rho''_{\text{PREF}})) > n$. Hence $\text{CNT}(\zeta) > n$, and a symmetric reasoning shows that $\text{CNT}(\bar{\zeta}) > n$. This proves condition (1.c).

Observe that $\text{EFF}(\rho_{\text{PREF}}) = \text{CNT}(\zeta)$ and $\text{EFF}(\rho_{\text{SUFF}}) = -\text{CNT}(\bar{\zeta})$. Hence, condition (1.d) follows from $|\text{CNT}(\zeta) - \text{CNT}(\bar{\zeta})| < \max(A, B)$.

For condition (1.e), aiming towards a contradiction suppose that some configuration γ appears more than once on ρ_{PREF} and ρ_{SUFF} . By construction, no configuration repeats on ρ_{PREF} and on ρ_{SUFF} individually, so one of the appearances of γ is on ρ_{PREF} and the second is on ρ_{SUFF} . Define a path ρ_1 by concatenating the prefix of ρ_{PREF} up to the appearance of γ together with the suffix of ρ_{SUFF} beginning from the appearance of γ . Clearly, ρ_1 is an arc from α to β , and moreover it is low because both ρ_{PREF} and ρ_{SUFF} are low. This contradicts the assumption that there is no low arc from α to β . \perp

The intuition now is that by repeating σ_{UP} and σ_{DOWN} appropriately many times (i.e., selecting numbers a and b) we can choose any difference of effects of $\rho_{\text{PREF}}\rho_{\text{UP}}$ and $\rho_{\text{DOWN}}\rho_{\text{SUFF}}$, as long as this difference belongs to a fixed congruence class modulo $\text{gcd}(A, B)$. This means that the middle part of the path ρ_\circ can be unpumped “modulo $\text{gcd}(A, B)$ ”: even if we change its effect by a multiple of $\text{gcd}(A, B)$, we will be able to compensate for this change by adjusting a and b .

We now proceed to showing how the middle part of the path, i.e., ρ_{CAP} , will be constructed. Intuitively, the idea is to take the part from ρ_\circ between indices k and \bar{k} , augment it with short connectives σ_{qp} and $\sigma_{\bar{p}\bar{q}}$ to link it with the cycles σ_{UP} and σ_{DOWN} , and unpump it “modulo $\text{gcd}(A, B)$ ” exhaustively. However, during further constructions we need certain divisibility properties of $\text{EFF}(\sigma_{\text{CAP}})$, and hence the construction of the connections to σ_{UP} and σ_{DOWN} is more complicated.

Claim 2. *There exists a sequence of transitions σ_{CAP} such that*

- (2.a) σ_{CAP} starts in q and finishes in \bar{q} (base states of σ_{UP} and σ_{DOWN} respectively);
- (2.b) no infix of σ_{CAP} is a cycle with effect divisible by $\text{gcd}(A, B)$; and
- (2.c) $\text{EFF}(\rho_{\text{PREF}}) + \text{EFF}(\sigma_{\text{CAP}}) + \text{EFF}(\rho_{\text{SUFF}})$ is divisible by $\text{gcd}(A, B)$.

Proof. The construction is depicted in Figure 2. Let $\sigma_{\text{PRE-CONN}} = \sigma_{pq}\sigma_{qp}$ and $\sigma_{\text{POST-CONN}} = \sigma_{\bar{p}\bar{q}}\sigma_{\bar{q}\bar{p}}$. We set

$$\begin{aligned}\sigma_{\text{PRE}} &= \sigma_{qp} (\sigma_{\text{PRE-CONN}})^c \\ \sigma_{\text{POST}} &= (\sigma_{\text{POST-CONN}})^c \sigma_{\bar{p}\bar{q}},\end{aligned}$$

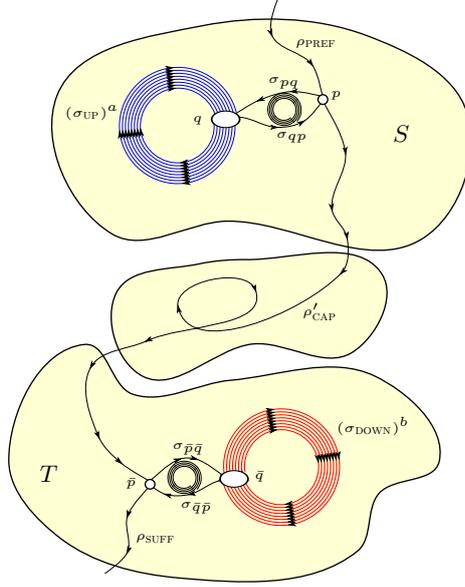


Fig. 2. The construction of ρ'_{CAP} in the first part of the proof of Claim 2. Note that cycles $(\sigma_{\text{PRE-CONN}})^c$ and $(\sigma_{\text{POST-CONN}})^c$ are depicted symbolically as small spirals.

where $c = \gcd(A, B) - 1$. It is easy to verify that $\sigma_{\text{PRE}}, \sigma_{\text{POST}}$ are consistent and

$$\text{EFF}(\sigma_{\text{PRE}}) + \text{EFF}(\sigma_{pq}) = (c + 1) \cdot \text{EFF}(\sigma_{\text{PRE-CONN}}) \equiv 0 \pmod{\gcd(A, B)}, \quad (1)$$

$$\text{EFF}(\sigma_{\text{POST}}) + \text{EFF}(\sigma_{\bar{q}\bar{p}}) = (c + 1) \cdot \text{EFF}(\sigma_{\text{POST-CONN}}) \equiv 0 \pmod{\gcd(A, B)}. \quad (2)$$

Consider the sequence of transitions

$$\sigma_{\text{MIDD}} = t_{i_k} t_{i_k+1} \cdots t_{j_{\bar{k}}-1}.$$

and define $\sigma'_{\text{CAP}} = \sigma_{\text{PRE}} \sigma_{\text{MIDD}} \sigma_{\text{POST}}$. We now verify that conditions (2.a) and (2.c) are satisfied for σ'_{CAP} . Condition (2.a) follows directly from the construction. For condition (2.c), observe that from the construction of ρ_{PREF} and ρ_{SUFF} we have

$$\begin{aligned} \text{EFF}(\rho_{\text{PREF}}) &= \text{EFF}(\rho'_{\text{PREF}}) = \text{EFF}(t_1 t_2 \cdots t_{i_k-1}) + \text{EFF}(\sigma_{pq}) + x \cdot A \\ \text{EFF}(\rho_{\text{SUFF}}) &= \text{EFF}(\rho'_{\text{SUFF}}) = \text{EFF}(t_{j_{\bar{k}}} t_{j_{\bar{k}}+1} \cdots t_m) + \text{EFF}(\sigma_{\bar{q}\bar{p}}) - y \cdot B \end{aligned}$$

for some $x, y \in \mathbb{N}$. Hence, from (1), (2) and the fact that $\text{EFF}(\rho_o) = 0$, it follows that

$$\begin{aligned} &\text{EFF}(\rho_{\text{PREF}}) + \text{EFF}(\sigma'_{\text{CAP}}) + \text{EFF}(\rho_{\text{SUFF}}) \equiv \\ &\text{EFF}(\rho_o) + \text{EFF}(\sigma_{pq}) + \text{EFF}(\sigma_{\bar{q}\bar{p}}) + \text{EFF}(\sigma_{\text{PRE}}) + \text{EFF}(\sigma_{\text{POST}}) \equiv 0 \pmod{\gcd(A, B)}. \end{aligned}$$

So condition (2.c) indeed holds for σ'_{CAP} .

We now take condition (2.b) into consideration. Define σ_{CAP} to be any of the shortest possible consistent sequences of transitions satisfying conditions (2.a) and (2.c); the existence of σ'_{CAP} implies that there exists such a sequence, so σ_{CAP} is well-defined. Let $\sigma_{\text{CAP}} = t_1 \dots t_r$. Assume towards contradiction that σ_{CAP} does not satisfy condition (2.b). Then there is some infix $t_{i+1} \dots t_j$ that is a cycle and its effect is divisible by $\gcd(A, B)$. It is easy to observe that sequence $\sigma' = t_1 \dots t_i t_{j+1} \dots t_m$ is consistent, $\text{SRC}(\sigma') = q$, $\text{TARG}(\sigma') = \bar{q}$ and $\text{EFF}(\sigma') - \text{EFF}(\sigma) \equiv 0 \pmod{\gcd(A, B)}$. Hence σ' satisfies conditions (2.a) and (2.c) while being strictly shorter than σ_{CAP} . This is a contradiction with the minimality (shortness) of σ_{CAP} , which proves that σ_{CAP} satisfies condition (2.b). \dashv

Note that from condition (2.b) it follows that $\text{LEN}(\sigma_{\text{CAP}}) \leq \gcd(A, B) \cdot n$. In the final construction this condition will directly imply that ρ_{CAP} will satisfy property (v), since ρ_{CAP} will be simply σ_{CAP} fastened at some configuration.

We denote

$$\begin{aligned} L &= \text{LEN}(\sigma_{\text{CAP}}), \\ K &= \text{EFF}(\rho_{\text{PREF}}) + \text{EFF}(\sigma_{\text{CAP}}) + \text{EFF}(\rho_{\text{SUFF}}). \end{aligned}$$

Recall that $\gcd(A, B)$ divides K . Moreover, from condition (1.d) we know that

$$|K| \leq |\text{EFF}(\sigma_{\text{CAP}})| + |\text{EFF}(\rho_{\text{PREF}}) + \text{EFF}(\rho_{\text{SUFF}})| < L + \max(A, B). \quad (3)$$

Having defined ρ_{PREF} , ρ_{SUFF} and σ_{CAP} , we proceed to defining ρ_{UP} and ρ_{DOWN} . For this, we need to define $a, b \in \mathbb{N}$: the numbers of times the cycles σ_{UP} and σ_{DOWN} are repeated on ρ_{UP} and ρ_{DOWN} . As described earlier, they have to be chosen so that the resulting path ρ is valid and has zero effect, but they also need to be reasonably small so that conditions (iii) and (iv) are satisfied. We now prove that this is always possible.

Claim 3. *There exist $a, b \in \mathbb{N}$ such that the following conditions hold:*

- (3.a) $a \cdot A - b \cdot B = -K$;
- (3.b) $L \leq a \cdot A, b \cdot B \leq 2L + 2 \cdot \text{lcm}(A, B)$.

Proof. It is a well-known fact that there exist some integers x_0, y_0 such that $x_0 \cdot A - y_0 \cdot B = \gcd(A, B)$. Since K is divisible by $\gcd(A, B)$, we can take $x = x_0 \cdot (-K / \gcd(A, B))$ and $y = y_0 \cdot (-K / \gcd(A, B))$ so that $x \cdot A - y \cdot B = -K$. Moreover, by increasing x by $M \cdot B$ and increasing y by $M \cdot A$, for a sufficiently large integer M , we can further assume that $x, y \geq 0$. Suppose then that (x, y) is a pair of nonnegative integers satisfying $x \cdot A - y \cdot B = -K$ for which $x + y$ is the smallest possible. We claim that $x \cdot A, y \cdot B \leq |K| + \text{lcm}(A, B)$.

Aiming towards a contradiction, suppose that $x \cdot A > |K| + \text{lcm}(A, B)$. Then in particular $x > \text{lcm}(A, B) / A = B / \gcd(A, B)$. Also, $y \cdot B = x \cdot A + K > \text{lcm}(A, B)$, and hence $y > \text{lcm}(A, B) / B = A / \gcd(A, B)$. Consider $x' = x - B / \gcd(A, B)$ and $y = y' - A / \gcd(A, B)$. Then we have $x', y' \geq 0$ and it can be easily verified that $x' \cdot A - y' \cdot B = -K$. As $x' + y' < x + y$, this is a contradiction with the minimality of $x + y$. Hence indeed $x \cdot A \leq |K| + \text{lcm}(A, B)$, and symmetrically

we also prove that $y \cdot B \leq |K| + \text{lcm}(A, B)$. Note that by (3) the inequalities $|K| + \text{lcm}(A, B) \leq L + \max(A, B) + \text{lcm}(A, B) \leq L + 2 \cdot \text{lcm}(A, B)$ hold.

It remains to define a, b based on x, y so that the lower bound on $a \cdot A$ and $b \cdot B$ holds. If already $x \cdot A, y \cdot B \geq L$, then we can take $(a, b) = (x, y)$; assume therefore that this is not the case. For $i \in \mathbb{N}$, let

$$a_i = x + i \cdot (\text{lcm}(A, B)/A) \quad \text{and} \quad b_i = y + i \cdot (\text{lcm}(A, B)/B).$$

Clearly $a_i, b_i \geq 0$, and it is easy to verify that $a_i \cdot A - b_i \cdot B = -K$ for each $i \in \mathbb{N}$. It therefore only suffices to show that there exists i such that $L \leq a_i \cdot A, b_i \cdot B \leq 2L + 2 \cdot \text{lcm}(A, B)$. Let i be the smallest nonnegative integer so that $a_i \cdot A \geq L$ and $b_i \cdot B \geq L$; then $i > 0$. Suppose that $K \geq 0$; the other case is symmetric. As $a_i \cdot A - b_i \cdot B = -K \leq 0$, we have $a_i \cdot A \leq b_i \cdot B$; by our definition of a_i and b_i , similarly $a_{i-1} \cdot A \leq b_{i-1} \cdot B$. Since by the minimality of i either $a_{i-1} \cdot A < L$ or $b_{i-1} \cdot B < L$ holds, we deduce that $a_{i-1} \cdot A < L$ in both cases, and it follows that $a_i \cdot A < L + \text{lcm}(A, B)$. Now

$$b_i \cdot B = K + a_i \cdot A < (L + \max(A, B)) + (L + \text{lcm}(A, B)) \leq 2L + 2 \cdot \text{lcm}(A, B).$$

Thus we can take $(a, b) = (a_i, b_i)$. The case $K < 0$ is symmetric. \square

Let us fix the numbers $a, b \in \mathbb{N}$ given by Claim 3. We are finally ready to define the whole path ρ . Define ρ_{UP} as $(\sigma_{\text{UP}})^a$ fastened at configuration ζ . Symmetrically we define ρ_{DOWN} as $(\sigma_{\text{DOWN}})^b$ fastened at $(\bar{q}, b \cdot B + \text{CNT}(\bar{\zeta}))$, so that its target is $\bar{\zeta}$. Finally, let ρ_{CAP} be σ_{CAP} fastened at $\text{TARG}(\rho_{\text{UP}})$, and define

$$\rho = \rho_{\text{PREF}} \rho_{\text{UP}} \rho_{\text{CAP}} \rho_{\text{DOWN}} \rho_{\text{SUFF}}.$$

Note that in this definition we did not verify properly that appropriate sequences of transitions are fireable at certain configurations. We perform this check in the next claim.

Claim 4. ρ is a normal arc from α to β .

Proof. First, condition (1.b) ensures that on ρ_{PREF} all the configurations have positive counter values apart from the source configuration α . Similarly, on ρ_{SUFF} all the configurations have positive counter values apart from the target configuration β . Condition (1.c) asserts that $\text{CNT}(\zeta) > n$, so cycle σ_{UP} is fireable at ζ because $\text{LEN}(\sigma_{\text{UP}}) \leq n$. Since σ_{UP} is a positive cycle, it can be easily seen that also $(\sigma_{\text{UP}})^a$ is fireable at ζ , and moreover on ρ_{UP} we do not obtain any configuration with zero counter value. A symmetric reasoning shows that $(\sigma_{\text{DOWN}})^b$ is fireable at $(\bar{q}, b \cdot B + \text{CNT}(\bar{\zeta}))$ so that its target is $\bar{\zeta}$ and $\rho_{\text{DOWN}} \rho_{\text{SUFF}}$ is a valid path with β being the only configuration with zero counter value.

Now observe that $\text{CNT}(\text{TARG}(\rho_{\text{UP}})) = \text{EFF}(\rho_{\text{PREF}}) + a \cdot A$, which is strictly larger than L by condition (3.b). Since $\text{LEN}(\sigma_{\text{CAP}}) = L$ and $\text{SRC}(\sigma_{\text{CAP}}) = q = \text{ST}(\text{TARG}(\rho_{\text{UP}}))$, we see that indeed σ_{CAP} is fireable at $\text{TARG}(\rho_{\text{UP}})$, and moreover on ρ_{CAP} all the configurations have positive counter values.

To conclude that ρ is an arc from α to β , it remains to verify that $\text{TARG}(\rho_{\text{CAP}}) = \text{SRC}(\rho_{\text{DOWN}})$. Both these configurations have state \bar{q} , so we need to verify that their counter values are equal. However, using condition (3.a) we have the following:

$$\begin{aligned} \text{CNT}(\text{TARG}(\rho_{\text{CAP}})) &= \text{EFF}(\rho_{\text{PREF}}) + \text{EFF}(\rho_{\text{UP}}) + \text{EFF}(\rho_{\text{CAP}}) \\ &= a \cdot A + K - \text{EFF}(\rho_{\text{SUFF}}) = b \cdot B - \text{EFF}(\rho_{\text{SUFF}}) \\ &= -\text{EFF}(\rho_{\text{DOWN}}) - \text{EFF}(\rho_{\text{SUFF}}) = \text{CNT}(\text{SRC}(\rho_{\text{DOWN}})). \end{aligned}$$

Hence indeed ρ is an arc from α to β , normal by construction (as $\text{ST}(\text{SRC}(\rho_{\text{CAP}}))$ and $\text{ST}(\text{TARG}(\rho_{\text{CAP}}))$ are the base states of σ_{UP} and σ_{DOWN} respectively). \square

We summarize the properties of ρ that are required in the lemma statement. Properties (i) and (ii) follow directly from the construction. Properties (iii) and (iv) follow from our choice of a and b , in particular from condition (3.b). Property (v) follows from condition (2.b) and the fact that $\sigma_{\text{CAP}} = \text{PROJ}(\rho_{\text{CAP}})$. Property (vi) follows from condition (1.c) and the fact that σ_{UP} and σ_{DOWN} are a positive and a negative cycle, respectively. Finally, property (vii) follows from condition (1.e). This concludes the proof of Lemma 2. \square

5.3 Length of shortest paths

Let α and β be such as in the statement of Theorem 1. Let ρ_{\circ} be a path from α to β that has the minimum possible number of intermediate configurations with counter value zero. Let all these intermediate configurations with counter value zero be $\gamma_2, \dots, \gamma_k$, where $\gamma_1 = \alpha$ and $\gamma_{k+1} = \beta$. For $i = 1, 2, \dots, k$, let ρ_{\circ}^i be the subpath of ρ_{\circ} between configurations γ_i and γ_{i+1} . Then ρ_{\circ}^i is an arc from γ_i to γ_{i+1} . By Lemma 2, there exists also an arc ρ^i from γ_i to γ_{i+1} that is either low or is normal and admits a normal decomposition satisfying properties (i)–(vii). In case ρ^i is low, choose ρ^i to be the shortest possible low arc from γ_i to γ_{i+1} . In case ρ^i is normal, let

$$\rho^i = \rho_{\text{PREF}}^i \rho_{\text{UP}}^i \rho_{\text{CAP}}^i \rho_{\text{DOWN}}^i \rho_{\text{SUFF}}^i$$

be its normal decomposition. Our goal for the rest of the proof (i.e., for this subsection) is to show that $\rho = \rho^1 \dots \rho^k$, which is clearly a path from α to β , has length at most $14n^2$. Note that ρ has the same number of configurations with counter value zero as ρ_{\circ} . Let $\mathcal{N} \subseteq \{1, 2, \dots, k\}$ be the set of indices i for which ρ^i is normal, and let $\mathcal{L} = \{1, 2, \dots, k\} \setminus \mathcal{N}$ be the set of indices i for which ρ^i is low.

First we show that the sum of the lengths of low parts of ρ , more precisely, low arcs, ρ_{PREF}^i and ρ_{SUFF}^i , is small. For this, Proposition 1 will be very useful.

Lemma 3. *The following inequality holds:*

$$\sum_{i \in \mathcal{L}} \text{LEN}(\rho^i) + \sum_{i \in \mathcal{N}} (\text{LEN}(\rho_{\text{PREF}}^i) + \text{LEN}(\rho_{\text{SUFF}}^i)) \leq 5n^2.$$

Proof. For every $i \in \mathcal{L}$, no configuration appears on ρ^i more than once, because in such a case ρ^i could be made shorter using Proposition 1 without spoiling the property that it is low, which would contradict the assumption that ρ^i is the shortest possible. For every $i \in \mathcal{N}$, property (vii) of Lemma 2 ensures that no configuration appears more than once on the paths ρ_{PREF}^i and ρ_{SUFF}^i . Suppose that some configuration γ appears both in ρ^i and in ρ^j , for some $i < j$. Then by applying Proposition 1 to configuration γ in the path ρ , we would obtain a path from α to β with a strictly smaller number of intermediate configurations with counter value zero, a contradiction with the choice of ρ_\circ .

Hence, we conclude that among configurations appearing on paths from the set $\{\rho^i : i \in \mathcal{L}\} \cup \{\rho_{\text{PREF}}^i, \rho_{\text{SUFF}}^i : i \in \mathcal{N}\}$, no configuration appears more than once. Since all these paths are low, all these configurations have counter values between 0 and $5n - 1$. Hence, the total number of configurations appearing on these paths is at most $n \cdot 5n = 5n^2$, which concludes the proof. \square

Now we will estimate the length of the rest of the path ρ . First, however, we have to prepare a toolbox of lemmas. We introduce the following notation. For $S, T \in \mathfrak{S}$, let $\mathcal{N}_{(S,T)} \subseteq \mathcal{N}$ be the set of all those indices i for which ρ^i is (S, T) -normal. Moreover, let $\mathcal{N}_{(S,\cdot)} = \bigcup_{T' \in \mathfrak{S}} \mathcal{N}_{(S,T')}$ and $\mathcal{N}_{(\cdot,T)} = \bigcup_{S' \in \mathfrak{S}} \mathcal{N}_{(S',T)}$.

Lemma 4. *Let $S, T \in \mathfrak{S}$. Suppose $i \in \mathcal{N}_{(S,\cdot)}$ and $j \in \mathcal{N}_{(\cdot,T)}$ for some i, j with $1 \leq i < j \leq k$. Then there are no two configurations δ_i and δ_j appearing on ρ_{CAP}^i and ρ_{CAP}^j respectively, such that $\text{ST}(\delta_i) = \text{ST}(\delta_j)$ and $\text{CNT}(\delta_i) - \text{CNT}(\delta_j)$ is divisible by $\gcd(\text{EFF}(\sigma_S^+), -\text{EFF}(\sigma_T^-))$.*

Proof. Denote $A = \text{EFF}(\sigma_S^+)$ and $B = -\text{EFF}(\sigma_T^-)$. Assume towards contradiction that there exists δ_i on ρ_{CAP}^i and δ_j on ρ_{CAP}^j such that $\text{ST}(\delta_i) = \text{ST}(\delta_j)$ and $\text{CNT}(\delta_i) - \text{CNT}(\delta_j)$ is divisible by $\gcd(A, B)$. We will show that we can modify ρ^i and ρ^j so that the part between δ_i and δ_j can be cut off from ρ . This will be a contradiction with the assumption that ρ has the minimum possible number of intermediate configurations with counter value zero.

Let $\text{CNT}(\delta_j) - \text{CNT}(\delta_i) = Z$, where $Z = z \cdot \gcd(A, B)$ for some integer z . It is a known fact that there exist $a, b \in \mathbb{N}$ such that $a \cdot A - b \cdot B = \gcd(A, B)$; cf. the proof of Claim 3. If $z \geq 0$ then $az \cdot A - bz \cdot B = z \cdot \gcd(A, B) = Z$, where $az \geq 0$ and $bz \geq 0$. If $z < 0$ then $(MB + az) \cdot A - (MA + bz) \cdot B = z \cdot \gcd(A, B) = Z$, where M is large enough so that $MB + az \geq 0$ and $MA + bz \geq 0$. Therefore, there always exist numbers $a, b \geq 0$ such that $a \cdot A - b \cdot B = Z$.

We modify the path ρ as follows. In the path $\rho_{\text{UP}}^i \rho_{\text{CAP}}^i$ we insert a cycles σ_S^+ at the end of ρ_{UP}^i , and in $\rho_{\text{CAP}}^j \rho_{\text{DOWN}}^j$ we insert b cycles σ_T^- at the front of ρ_{DOWN}^j . By property (vi), this insertion does not introduce configurations with nonpositive counter values, since each of the cycles σ_S^+ and σ_T^- contains at most n edges. After this operation, the configuration δ_i that was originally on ρ_{CAP}^i becomes lifted to the configuration $(\text{ST}(\delta_i), \text{CNT}(\delta_i) + a \cdot A)$. On the other hand, the configuration δ_j that was originally on ρ_{CAP}^j becomes lifted to the configuration $(\text{ST}(\delta_j), \text{CNT}(\delta_j) + b \cdot B)$. However,

$$\text{CNT}(\delta_i) + a \cdot A = \text{CNT}(\delta_j) - Z + a \cdot A = \text{CNT}(\delta_j) + b \cdot B.$$

Since $\text{ST}(\delta_i) = \text{ST}(\delta_j)$, we conclude that these two lifted configurations are equal. Therefore, we can perform the following operation on ρ : insert the cycles σ_S^+ and σ_T^- as described above, and cut out the entire part of ρ between the lifted configurations originating in δ_i and δ_j by Proposition 1. In this manner we obtain a path from α to β that has strictly less intermediate configurations with counter value equal zero than ρ , which is a contradiction. \square

Lemma 5. *Let $S, T \in \mathfrak{S}$. Then $|\mathcal{N}_{(S,T)}| \leq \gcd(\text{EFF}(\sigma_S^+), -\text{EFF}(\sigma_T^-))$.*

Proof. Let us denote $A = \text{EFF}(\sigma_S^+)$ and $B = -\text{EFF}(\sigma_T^-)$. Assume towards a contradiction that $|\mathcal{N}_{(S,T)}| > \gcd(A, B)$. For $i \in \mathcal{N}_{(S,T)}$, let $\delta_i = \text{TARG}(\rho_{\text{CAP}}^i)$. By the pigeonhole principle, for some two indices $i < j$ configurations δ_i and δ_j have the same counter value modulo $\gcd(A, B)$. Moreover, δ_i and δ_j have the same state, which is the base state of σ_T^- by our definition of a normal arc. This contradicts Lemma 4. \square

Total length of caps. We have now all the necessary ingredients to establish the desired upper bounds on the lengths of caps. Recall that for a strongly connected component $S \in \mathfrak{S}$ we denote by n_S the number of vertices in S .

Lemma 6. *Let $S, T \in \mathfrak{S}$, let $A_S = \text{EFF}(\sigma_S^+)$ and let $B_T = -\text{EFF}(\sigma_T^-)$. Then:*

$$\sum_{i \in \mathcal{N}_{(S,\cdot)}} \text{LEN}(\rho_{\text{CAP}}^i) \leq A_S \cdot n; \quad (4)$$

$$\sum_{i \in \mathcal{N}_{(\cdot,T)}} \text{LEN}(\rho_{\text{CAP}}^i) \leq n \cdot B_T; \quad (5)$$

$$\text{moreover, } \sum_{i \in \mathcal{N}} \text{LEN}(\rho_{\text{CAP}}^i) \leq n^2. \quad (6)$$

Proof. For (4), assume towards a contradiction that $\sum_{i \in \mathcal{N}_{(S,\cdot)}} \text{LEN}(\rho_{\text{CAP}}^i) > A_S \cdot n$. Then by the pigeonhole principle there exists two configurations δ and δ' on the paths ρ_{CAP}^i for $i \in \mathcal{N}_{(S,\cdot)}$ which have the same state and the same counter value modulo A_S . Assume w.l.o.g. that δ is earlier in the path than δ' . By property (v) of Lemma 2, configurations δ and δ' cannot appear in the same path ρ_{CAP}^i . Indeed, otherwise the projection of the part of ρ_{CAP}^i between δ to δ' would be a cycle with effect divisible by A_S , so also by $\gcd(A_S, -\text{EFF}(\sigma_T^-))$, where T is the SCC for which ρ^i is (S, T) -normal. Therefore they have to belong to different arcs. Let δ belong to ρ^i and δ' belong to ρ^j , where $j \in \mathcal{N}_{(S,T)}$ for some $T \in \mathfrak{S}$. However, by Lemma 4, there are no two configurations δ and δ' on ρ^i and ρ^j , respectively, such that their states are the same and the difference in counter values is divisible by $\gcd(A_S, -\text{EFF}(\sigma_T^-))$. Contradiction, as δ and δ' are such configurations: the difference of its counter values is divisible by A_S , so also by $\gcd(A_S, -\text{EFF}(\sigma_T^-))$. Thus (4) is proved, and (5) follows from a symmetric reasoning. The bound (6) follows by summing (4) through all $S \in \mathfrak{S}$ and using the facts that $\text{EFF}(\sigma_S^+) \leq n_S$ and $\sum_{S \in \mathfrak{S}} n_S = n$. \square

Total length of positive and negative cycles. We now show that the total sum of the lengths of ρ_{UP}^i and ρ_{DOWN}^i is at most $8n^2$. This is the case where we need the key estimations (iii) and (iv) in Lemma 2.

Lemma 7. *The following inequalities hold:*

$$\begin{aligned} \sum_{i \in \mathcal{N}} \text{LEN}(\rho_{\text{UP}}^i) &\leq 4n^2, \\ \sum_{i \in \mathcal{N}} \text{LEN}(\rho_{\text{DOWN}}^i) &\leq 4n^2. \end{aligned}$$

Proof. We show how to bound the sum of lengths of paths ρ_{UP}^i . For any $S \in \mathfrak{S}$, let us denote $A_S = \text{EFF}(\sigma_S^+)$ and $B_S = -\text{EFF}(\sigma_S^-)$. For each $i \in \mathcal{N}$, let $S_i, T_i \in \mathfrak{S}$ be such that ρ^i is (S_i, T_i) -normal, and let $L_i = \text{LEN}(\rho_{\text{CAP}}^i)$. By Lemma 2 we know that $\text{EFF}(\rho_{\text{UP}}^i) \leq 2L_i + 2 \cdot \text{lcm}(A_{S_i}, B_{T_i})$. Since $\text{PROJ}(\rho_{\text{UP}}^i) = (\sigma_{S_i}^+)^a$ for some integer a , we have

$$\text{LEN}(\rho_{\text{UP}}^i) = \text{EFF}(\rho_{\text{UP}}^i) \cdot \frac{\text{LEN}(\sigma_{S_i}^+)}{\text{EFF}(\sigma_{S_i}^+)} \leq \text{EFF}(\rho_{\text{UP}}^i) \cdot \frac{n_{S_i}}{A_{S_i}} \leq (2L_i + 2 \cdot \text{lcm}(A_{S_i}, B_{T_i})) \cdot \frac{n_{S_i}}{A_{S_i}}.$$

Hence,

$$\sum_{i \in \mathcal{N}} \text{LEN}(\rho_{\text{UP}}^i) \leq 2 \sum_{i \in \mathcal{N}} \frac{L_i n_{S_i}}{A_{S_i}} + 2 \sum_{i \in \mathcal{N}} \frac{\text{lcm}(A_{S_i}, B_{T_i}) \cdot n_{S_i}}{A_{S_i}}. \quad (7)$$

We will separately estimate the first and the second term. First we focus on $\sum_{i \in \mathcal{N}} \frac{L_i n_{S_i}}{A_{S_i}}$. Let us fix some specific $S \in \mathfrak{S}$. We have

$$\sum_{i \in \mathcal{N}(S, \cdot)} \frac{L_i n_{S_i}}{A_{S_i}} = \frac{n_S}{A_S} \cdot \sum_{i \in \mathcal{N}(S, \cdot)} L_i \leq \frac{n_S}{A_S} \cdot A_S \cdot n = n_S \cdot n,$$

where the inequality follows from Lemma 6(4). Thus

$$\sum_{i \in \mathcal{N}} \frac{L_i n_{S_i}}{A_{S_i}} = \sum_{S \in \mathfrak{S}} \sum_{i \in \mathcal{N}(S, \cdot)} \frac{L_i n_{S_i}}{A_{S_i}} \leq \sum_{S \in \mathfrak{S}} n_S \cdot n = n^2. \quad (8)$$

In order to estimate the second term, fix some $S, T \in \mathfrak{S}$. Note that $\frac{\text{lcm}(x, y)}{x} = \frac{xy}{\text{gcd}(x, y) \cdot x} = \frac{y}{\text{gcd}(x, y)}$ for all positive integers x, y . Now we have

$$\begin{aligned} \sum_{i \in \mathcal{N}(S, T)} \frac{B_{T_i} \cdot n_{S_i}}{\text{gcd}(A_{S_i}, B_{T_i})} &= \sum_{i \in \mathcal{N}(S, T)} \frac{B_T \cdot n_S}{\text{gcd}(A_S, B_T)} = |\mathcal{N}(S, T)| \cdot \frac{B_T \cdot n_S}{\text{gcd}(A_S, B_T)} \\ &\leq \text{gcd}(A_S, B_T) \cdot \frac{B_T \cdot n_S}{\text{gcd}(A_S, B_T)} = B_T \cdot n_S \leq n_T \cdot n_S, \end{aligned}$$

where the first inequality follows from Lemma 5 and the second one from the fact that the effect of a path is bounded by its length. Therefore,

$$\begin{aligned} \sum_{i \in \mathcal{N}} \frac{B_{T_i} \cdot n_{S_i}}{\gcd(A_{S_i}, B_{T_i})} &= \sum_{S, T \in \mathfrak{S}} \sum_{i \in \mathcal{N}(S, T)} \frac{B_{T_i} \cdot n_{S_i}}{\gcd(A_{S_i}, B_{T_i})} \\ &\leq \sum_{S, T \in \mathfrak{S}} n_T \cdot n_S = \sum_{S \in \mathfrak{S}} n_S \cdot \sum_{T \in \mathfrak{S}} n_T = n^2. \end{aligned} \quad (9)$$

By connecting equations (7), (8) and (9) we obtain

$$\sum_{i \in \mathcal{N}} \text{LEN}(\rho_{\text{UP}}^i) \leq 2n^2 + 2n^2 = 4n^2.$$

The upper bound on the sum of lengths of paths ρ_{DOWN}^i is obtained analogously, using Lemma 6(5) instead of Lemma 6(4). \square

Combining the bounds of Lemma 3, Lemma 6(6), and Lemma 7, we conclude that

$$\text{LEN}(\rho) \leq 5n^2 + n^2 + 4n^2 + 4n^2 \leq 14n^2,$$

which completes the proof of Theorem 1.

6 Generalizations

6.1 Proof of Theorem 2

In this section we prove Theorem 2, which provides an upper bound on the length of the shortest path between any pair of configurations. For convenience, we recall its statement.

Theorem 2. *Let \mathcal{O} be a one-counter system with n states. Suppose a configuration $\beta = (p_\beta, c_\beta)$ is reachable from a configuration $\alpha = (p_\alpha, c_\alpha)$ in \mathcal{O} . Then \mathcal{O} has a path from α to β of length at most $14n^2 + n \cdot \max(c_\alpha, c_\beta)$.*

Proof. Let $a = \text{CNT}(\alpha)$ and $b = \text{CNT}(\beta)$. Assume without loss of generality that $a \geq b$; the second case is symmetric. Let ρ be some path from α to β . We first formulate the following Claim.

Claim 5. *There exists a one counter system \mathcal{O}_a with the following property. For any $p, q \in Q$ there is a path from $(p, 0)$ to $(q, 0)$ of length exactly K (so with $K + 1$ configurations) in \mathcal{O}_a if and only if there is a path from (p, a) to (q, a) in \mathcal{O} which contains exactly $K + 1$ configurations of counter value at least a . Moreover \mathcal{O}_a and \mathcal{O} have the same number of states.*

Proof. We construct \mathcal{O}^a with the set of states Q^a , the set of non-zero transitions $T_{>0}^a$ and the set of zero tests $T_{=0}^a$ as follows. It has the same set of states as \mathcal{O} , so $Q^a = Q$, and the same set of non-zero transitions, so $T_{>0}^a = T_{>0}$. Only the set of zero tests is different. The set $T_{=0}^a$ contains only tuples of the form $(q, 0, q')$

for $q, q' \in Q^a$. Tuple $(q, 0, q')$ belongs to $T_{=0}^a$ if and only if there is a path in \mathcal{O} from configuration (q, a) to configuration (q', a) such that all the intermediate configurations have counter value smaller than a .

We now verify that for such \mathcal{O}^a , the claim indeed holds.

Assume that there is a path from $(p, 0)$ to $(q, 0)$ in \mathcal{O}^a . Observe that there is a corresponding path from (p, a) to (q, a) in \mathcal{O} . Every non-zero transition from (r, c) for $c > 0$ in \mathcal{O}^a is simulated by one transition from $(r, c + a)$ in \mathcal{O} , as $T_{>0}^a = T_{>0}$. Every zero-test from $(r, 0)$ to $(r', 0)$ in \mathcal{O}^a is simulated by a path from (r, a) to (r', a) in \mathcal{O} , where all the intermediate configurations have counter value smaller than a . Such a path exists by the definition of $T_{=0}^a$. Note that the corresponding path of \mathcal{O} indeed has exactly as many configurations with counter value at least a as there are configurations in the original path of \mathcal{O}^a .

Assume now that there is a path from (p, a) to (q, a) in \mathcal{O} . We construct the corresponding path of \mathcal{O}^a as follows. Every part of the path from some (r, a) to some (r', a) where all the configurations in between have smaller counter values is replaced by a zero-test of \mathcal{O}^a . The constructed path of \mathcal{O}^a indeed has as many configurations as there are configurations in the path of \mathcal{O} which have counter value at least a . \lrcorner

Let $\gamma = (q, a)$ be the last configuration in ρ which has counter value a . Assume that $\alpha = (p, a)$. Let us consider the OCS \mathcal{O}^a from Claim 5. As there is a path from (p, a) to (q, a) in \mathcal{O} then there is a path from $(p, 0)$ to $(q, 0)$ in \mathcal{O}^a . By Theorem 1 there is a path from $(p, 0)$ to $(q, 0)$ of length at most $14n^2$. Now one more time by Claim 5 there is a path from (p, a) to (q, a) which has at most $14n^2 + 1$ configurations of counter value at least a . Let us denote by ρ' concatenation of this path and the suffix of ρ starting in γ and finishing in β . Since γ is the last configuration in ρ which has counter value at least a then also in ρ' there are at most $14n^2 + 1$ configurations of counter value at least a .

Let ρ'' be a shortest path from α to β such that there are at most $14n^2 + 1$ configurations in this path with counter value at least a . There is at least one such path, namely ρ' , so a shortest one clearly exists. Let us define a set LOW as the set of all configurations appearing in ρ'' whose counter values are smaller than a . We claim that $|\text{LOW}| \leq an$. Indeed, assume otherwise, that $|\text{LOW}| > an$. Then, by pigeonhole principle, some two configurations appearing on ρ'' , say δ and δ' , are equal. Then by Proposition 1, cutting out the part of the path between δ and δ' would leave a strictly shorter path from α to β that would have not more configurations with counter value at least a . This would be a contradiction with the choice of ρ'' .

We therefore know that there are at most $14n^2 + 1$ configurations in ρ'' with counter value at least a and at most an configurations with counter value smaller than a . Thus altogether there are at most $14n^2 + 1 + an$ configurations on ρ'' , so $\text{LEN}(\rho'') \leq 14n^2 + an \leq 14n^2 + n \cdot \max(a, b)$. This finishes the proof of Theorem 2. \square

6.2 Generalization to counters with values in \mathbb{Z}

In this section we show how our results can be also used to give improved upper bounds on the length of the shortest path in the model considered by Alur and Černý [1]. Recall that in this model, the counter can take arbitrary values in \mathbb{Z} and there are no zero-tests. In fact, we can show that a quadratic upper bound holds in a much more general model, where zero tests are allowed and transitions fireable at positive counter values may differ from transitions fireable at negative counter values. We start with defining formally the model we are working with.

A *one- \mathbb{Z} -counter system* (\mathbb{Z} -OCS) \mathcal{O} consists of a finite set of states Q , a set of *positive transitions* $T_{>0} \subseteq Q \times \{-1, 0, 1\} \times Q$, a set of *negative transitions* $T_{<0} \subseteq Q \times \{-1, 0, 1\} \times Q$, and a set of *zero tests* $T_{=0} \subseteq Q \times \{-1, 0, 1\} \times Q$. The set of *transitions* is $T = T_{>0} \cup T_{<0} \cup T_{=0}$. The positive transitions are fireable in configurations where the counter value is positive, negative transitions are fireable whenever the counter value is negative, and zero tests are fireable whenever the counter value is equal to zero. We adopt all the notation from one-counter systems in a natural way. In particular, the configurations of a \mathbb{Z} -OCS \mathcal{O} are pairs (q, c) , where $q \in Q$ is the configuration's state, and $c \in \mathbb{Z}$ is the configuration's counter value. Observe that one- \mathbb{Z} -counter systems generalize standard one-counter systems, because we can take $T_{<0} = \emptyset$ and disallow zero tests having effect -1 on the counter.

Again, a path is a sequence of the form

$$(\gamma_1, t_1)(\gamma_2, t_2) \dots (\gamma_m, t_m) \in ((Q \times \mathbb{Z}) \times T)^*,$$

for which some final configuration γ_{m+1} exists, such that for each $i = 1, 2, \dots, m$ we have that $\gamma_i \xrightarrow{t_i} \gamma_{i+1}$, i.e., firing transition t_i at configuration γ_i results in configuration γ_{i+1} . We now state formally our result for one- \mathbb{Z} -counter systems.

Theorem 3. *Let \mathcal{O} be a one- \mathbb{Z} -counter system with n states. Suppose configuration β is reachable from configuration α in \mathcal{O} , where $\text{CNT}(\alpha) = c_\alpha$ and $\text{CNT}(\beta) = c_\beta$. Then \mathcal{O} has a path from α to β of length at most $56n^2 + n \cdot (|c_\alpha| + |c_\beta|)$.*

We remark that one can approach Theorem 3 by following the lines of the proofs of Theorems 1 and 2, and adjusting the argumentation to the setting of one- \mathbb{Z} -counter systems. The proof, however, would be even more technical, because having both positive and negative counter values requires performing the pumping arguments twice: both for very high (positive) values and for very low (negative) values. Instead, we show how Theorem 3 can be deduced from the proofs of Theorems 1 and 2 directly, using a similar approach as for Theorem 2.

Proof. Let \mathcal{O} be the given \mathbb{Z} -OCS, and let $Q, T_{>0}, T_{<0}, T_{=0}$ be its set of states, its sets of positive and negative transitions, and its set of zero tests, respectively. Denote $T = T_{>0} \cup T_{<0} \cup T_{=0}$. Let $a = \text{CNT}(\alpha)$ and $b = \text{CNT}(\beta)$.

For a transition $t = (q, c, q')$, its *negation* is transition $t^{\text{neg}} = (q, -c, q')$. For any set of transitions M , let M^{neg} denote the set of negations of transitions from M . For a configuration $\gamma = (q, c)$, the *negated* configuration γ^{neg} is equal

to $(q, -c)$. For a path ρ , by ρ^{neg} we mean a path obtained from ρ by negating all the configurations and all the transitions.

Let $T_{=0,+} = \{(q, c, q') \in T_{=0} \mid c \geq 0\}$ and $T_{=0,-} = \{(q, c, q') \in T_{=0} \mid c \leq 0\}$. Define the following (standard) one counter systems:

- \mathcal{O}_+ has state set Q , non-zero transition set $T_{>0}$, and zero tests $T_{=0,+}$; and
- \mathcal{O}_- has state set Q , non-zero transition set $(T_{<0})^{\text{neg}}$, and zero tests $(T_{=0,-})^{\text{neg}}$.

The following properties are immediate.

Claim 6. *The following holds for any sequence $\rho \in ((Q \times \mathbb{N}) \times T)^*$:*

- (a) ρ is a path of \mathcal{O}_+ if and only if ρ is a path of \mathcal{O} such that no configuration appearing on ρ has a negative counter value;
- (b) ρ^{neg} is a path of \mathcal{O}_- if and only if ρ is a path of \mathcal{O} such that no configuration appearing on ρ has a positive counter value.

Take a path ρ_\circ of \mathcal{O} that starts in α and finishes in β (by the assumption such a path exists), has the minimum possible number of configurations with counter value zero appearing on it, and among such paths has the smallest possible length. Assume first that there is no configuration appearing on ρ_\circ that has counter value zero. Since counter values in consecutive configurations of ρ_\circ change by at most one at a time, we infer that the counter values of configurations appearing on ρ_\circ are either always positive or always negative. Suppose first that they are always positive. By Claim 6(a), ρ_\circ is a path of \mathcal{O}_+ from α to β . Since ρ_\circ was chosen to be of the minimum possible length, by Theorem 2 applied to \mathcal{O}^+ we infer that $\text{LEN}(\rho_\circ) \leq 14n^2 + n \cdot \max(a, b) < 56n^2 + n \cdot (|a| + |b|)$, and we are done. Similarly, if all configurations appearing on ρ have negative counter values, then from Claim 6(b) we have that ρ_\circ^{neg} is a path from α^{neg} to β^{neg} in \mathcal{O}_- , and hence by Theorem 2 we infer that $\text{LEN}(\rho_\circ) = \text{LEN}(\rho_\circ^{\text{neg}}) \leq 14n^2 + n \cdot \max(-a, -b) < 56n^2 + n \cdot (|a| + |b|)$.

We are left with the case when at least one configuration on ρ_\circ has counter value zero. In the following we suppose that $a, b \geq 0$; the other three cases of the sign of a and b are symmetric. We denote

$$\rho_\circ = (\gamma_1, t_1)(\gamma_2, t_2) \dots (\gamma_m, t_m),$$

and $\gamma_m \xrightarrow{t_m} \gamma_{m+1}$, where $\gamma_1 = \alpha$ and $\gamma_{m+1} = \beta$. Let z and z' be respectively the indices of the first and the last configuration on ρ_\circ that has counter value zero; in particular $\text{CNT}(\gamma_z) = \text{CNT}(\gamma_{z'}) = 0$. Denote $\rho_\circ = \rho_{\text{PREF}}\rho_{\text{MIDD}}\rho_{\text{SUFF}}$, where

$$\begin{aligned} \rho_{\text{PREF}} &= (\gamma_1, t_1) \dots (\gamma_{z-1}, t_{z-1}), \\ \rho_{\text{MIDD}} &= (\gamma_z, t_z) \dots (\gamma_{z'-1}, t_{z'-1}), \\ \rho_{\text{SUFF}} &= (\gamma_{z'}, t_{z'}) \dots (\gamma_m, t_m). \end{aligned}$$

In case $z = 1$ or $z' = m + 1$, we take ρ_{PREF} or ρ_{SUFF} to be an empty path, respectively.

By the choice of z and the assumption that $a \geq 0$, we have that all the configurations appearing on ρ_{PREF} have nonnegative counter values. By Claim 6(a), we infer that ρ_{PREF} is also a path of \mathcal{O}_+ , starting from α and ending in γ_z . By Theorem 2, there is a path ρ'_{PREF} of \mathcal{O}_+ that starts in α and ends in γ_z , and for which we have $\text{LEN}(\rho'_{\text{PREF}}) \leq 14n^2 + na = 14n^2 + n \cdot |a|$. Note that in the case when $a \leq 0$, we could apply a symmetric reasoning to $\rho_{\text{PREF}}^{\text{neg}}$, which is a path in \mathcal{O}_- , and, by negating once more, again obtain a path ρ'_{PREF} from α to γ_z of length at most $14n^2 + n \cdot |a|$. A symmetric argument applied to ρ_{SUFF} shows that we can find a path ρ'_{SUFF} from $\gamma_{z'}$ to β with $\text{LEN}(\rho'_{\text{SUFF}}) \leq 14n^2 + n \cdot |b|$.

We are left with bounding the length of ρ_{MIDD} . Let z_1, z_2, \dots, z_k be the consecutive indices between $z = z_1$ and $z' = z_k$ for which the corresponding configurations γ_{z_i} have counter value zero. For each $i = 1, 2, \dots, k-1$, let $\rho^i = (\gamma_{z_i}, t_i) \dots (\gamma_{z_{i+1}-1}, t_{z_{i+1}-1})$. Clearly, for each $i = 1, 2, \dots, k-1$ we either have that all the intermediate configurations of ρ^i have positive counter values, in which case we call ρ^i a *positive arc*, or they all have negative counter values, in which case ρ^i is called a *negative arc*, or ρ^i has length 1, in which case it is called a *zero arc*.

Consider a (standard) one-counter system \mathcal{O}'_+ obtained from \mathcal{O}_+ as follows: we add a zero test $(q, 0, q')$ for every pair of configurations (q, q') for which there exists a path from $(q, 0)$ to $(q', 0)$ in \mathcal{O} such that all its intermediate configurations have negative counter values. Similarly, let \mathcal{O}'_- be a (standard) one-counter system obtained from \mathcal{O}_- by adding a zero test $(q, 0, q')$ for every pair of configurations (q, q') for which there exists a path from $(q, 0)$ to $(q', 0)$ in \mathcal{O} such that all its intermediate configurations have positive counter values. Observe that every positive arc ρ^i gives rise to a zero test in \mathcal{O}'_- , and every negative arc ρ^i gives rise to a zero test in \mathcal{O}'_+ . Let ρ_+ be a path obtained from ρ_{MIDD} by replacing every negative arc by a corresponding zero test in \mathcal{O}'_+ . Similarly, let ρ_- be a path obtained from ρ_{MIDD} by replacing every positive arc by a corresponding zero test in \mathcal{O}'_- , and negating the whole path. It is straightforward to verify that ρ_+ and ρ_- are paths in \mathcal{O}'_+ and \mathcal{O}'_- , respectively. Moreover, they both start in γ_z and end in $\gamma_{z'}$, and for both the number of appearing configurations with counter value zero is the same as on ρ_{\circ} .

We now claim that ρ_+ is a path that has the minimum possible number of intermediate configurations with counter value zero, among paths from γ_z to $\gamma_{z'}$ in \mathcal{O}'_+ . Indeed, if there was such a path ρ' with a smaller number of intermediate configurations with counter value zero, then we could modify ρ' to a path ρ'' in \mathcal{O} by replacing each usage of any zero test $(q, 0, q')$ added in the construction of \mathcal{O}'_+ by a corresponding negative arc, i.e., a path in \mathcal{O} from $(q, 0)$ to $(q', 0)$ in which all the intermediate configurations have negative counter values. Then ρ'' would have the same number of intermediate configurations with counter value zero as ρ' . Hence $\rho = \rho_{\text{PREF}}\rho''\rho_{\text{SUFF}}$ would be a path from α to β in \mathcal{O} that would have strictly fewer configurations with counter value zero appearing on it than ρ_{\circ} ; this is a contradiction with the choice of ρ_{\circ} . A symmetric argument shows that ρ_- has the minimum possible number of intermediate configurations with counter value zero, among paths in \mathcal{O}'_- from γ_z to $\gamma_{z'}$.

We now apply the same substitution scheme to ρ_+ as in Section 5.3. That is, we partition the arcs of ρ_+ (which correspond to positive arcs in the decomposition of ρ_{MIDD}) into *low arcs*, where the counter values are always smaller than $5n$, and *high arcs*, where some counter values are at least $5n$. Each low arc is replaced by the shortest possible low arc that starts and ends in the same configurations, whereas each high arc is replaced by a normal arc obtained using Lemma 2. Let ρ'_+ be the obtained path from γ_z to $\gamma_{z'}$. The discussion of the previous paragraph ensures that ρ_+ satisfies the same minimality property as the original arc ρ_o considered in Section 5.3, i.e., it is a path from γ_z to $\gamma_{z'}$ that has the minimum possible number of intermediate configurations with counter value zero. Hence, the same analysis as in Section 5.3 applies, yielding that $\text{LEN}(\rho'_+) \leq 14n^2$. In a symmetric manner we can modify ρ_- to a path ρ'_- from γ_z to $\gamma_{z'}$ in \mathcal{O}'_- with $\text{LEN}(\rho'_-) \leq 14n^2$.

Now we consider ρ'_{MIDD} obtained from ρ_{MIDD} in the same manner: every positive arc is replaced as in the construction of ρ'_+ from ρ_+ , and every negative arc is replaced as in the construction of ρ'_- from ρ_- (obviously, with the negation applied). It follows that ρ'_{MIDD} is a path in \mathcal{O} from γ_z to $\gamma_{z'}$ with $\text{LEN}(\rho'_{\text{MIDD}}) \leq \text{LEN}(\rho'_+) + \text{LEN}(\rho'_-) \leq 28n^2$. Then define

$$\rho = \rho'_{\text{PREF}} \rho'_{\text{MIDD}} \rho'_{\text{SUFF}}.$$

We therefore have that ρ is a path from α to β in \mathcal{O} , and

$$\text{LEN}(\rho) = \text{LEN}(\rho'_{\text{PREF}}) + \text{LEN}(\rho'_{\text{MIDD}}) + \text{LEN}(\rho'_{\text{SUFF}}) \leq 56n^2 + n \cdot (|a| + |b|).$$

This finishes the proof of Theorem 3. \square

Finally, we remark that the model used by Alur and Černý [1] corresponds to setting $T_{>0} = T_{<0} = T_{=0}$ in our definition of a one- \mathbb{Z} -counter system. In this case, one can very easily obtain a marginally better upper bound of $56n^2 + n \cdot |c_\alpha - c_\beta|$. Indeed, since the fireability of transitions of \mathcal{O} is independent of the sign of the counter, on any path in \mathcal{O} we can add an arbitrary integer to all the counter values throughout the path, and we still obtain a path of \mathcal{O} . Hence, by decrementing all the counter values by c_α , we can equivalently consider the problem of bounding the length of the shortest path from α to β when we know that $c_\alpha = 0$. Then an application of Theorem 3 yields a path from α to β with length at most $56n^2 + n \cdot |c_\beta|$, which translates to the bound $56n^2 + n \cdot |c_\alpha - c_\beta|$ in the general case before decrementing.

Acknowledgements. The authors are grateful to Christoph Haase and Aditya Kanade for discussions and comments.

References

- [1] R. Alur and P. 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.

- [2] M. F. Atig, A. Bouajjani, K. N. Kumar, and P. Saivasan. On bounded reachability analysis of shared memory systems. In *34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, December 15-17, 2014, New Delhi, India*, pages 611–623, 2014.
- [3] C. Barrett, S. Demri, and M. Deters. Witness runs for counter machines. In *Frontiers of Combining Systems - 9th International Symposium, FroCoS 2013, Nancy, France, September 18-20, 2013. Proceedings*, pages 120–150, 2013.
- [4] A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of push-down automata: Application to model-checking. In *CONCUR '97: Concurrency Theory, 8th International Conference, Warsaw, Poland, July 1-4, 1997, Proceedings*, pages 135–150, 1997.
- [5] D. Chistikov. Notes on counting with finite machines. In *34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, December 15-17, 2014, New Delhi, India*, pages 339–350, 2014.
- [6] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, Second Edition*. The MIT Press and McGraw-Hill Book Company, 2001. ISBN 0-262-03293-7.
- [7] K. Etessami, D. Wojtczak, and M. Yannakakis. Quasi-birth-death processes, tree-like QBDs, probabilistic 1-counter automata, and pushdown systems. *Perform. Eval.*, 67(9):837–857, 2010.
- [8] P. Hofman, R. Mayr, and P. Totzke. Decidability of weak simulation on one-counter nets. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25-28, 2013*, pages 203–212, 2013.
- [9] D. C. Kozen. *Design and Analysis of Algorithms*. Texts and Monographs in Computer Science. Springer, 1992. ISBN 978-3-540-97687-5.
- [10] P. Lafourcade, D. Lugiez, and R. Treinen. Intruder deduction for AC-like equational theories with homomorphisms. Research Report LSV-04-16, Laboratoire Spécification et Vérification, ENS Cachan, France, Nov. 2004. URL http://www.lsv.ens-cachan.fr/Publis/RAPPORTS_LSV/PS/rr-lsv-2004-16.rr.ps. 69 pages.
- [11] P. Lafourcade, D. Lugiez, and R. Treinen. Intruder deduction for AC-like equational theories with homomorphisms. In *Term Rewriting and Applications, 16th International Conference, RTA 2005, Nara, Japan, April 19-21, 2005, Proceedings*, pages 308–322, 2005.
- [12] M. Latteux. Langages à un compteur. *J. Comput. Syst. Sci.*, 26(1):14–33, 1983.
- [13] J. Leroux and S. Schmitz. Demystifying reachability in vector addition systems. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, July 6-10, 2015*, pages 56–67, 2015.
- [14] A. Podelski and A. Rybalchenko. ARMC: the logical choice for software model checking with abstraction refinement. In *Practical Aspects of Declar-*

- ative Languages, 9th International Symposium, PADL 2007, Nice, France, January 14-15, 2007.*, pages 245–259, 2007.
- [15] J. Thakkar, A. Kanade, and R. Alur. Transducer-based algorithmic verification of retransmission protocols over noisy channels. In *Formal Techniques for Distributed Systems - Joint IFIP WG 6.1 International Conference, FMOODS/FORTE 2013, Held as Part of the 8th International Federated Conference on Distributed Computing Techniques, DisCoTec 2013, Florence, Italy, June 3-5, 2013. Proceedings*, pages 209–224, 2013.
- [16] W. Thomas. The reachability problem over infinite graphs. In *Computer Science - Theory and Applications, Fourth International Computer Science Symposium in Russia, CSR 2009, Novosibirsk, Russia, August 18-23, 2009. Proceedings*, pages 12–18, 2009.