

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, Warsaw, Poland
`{wczerwin,michal.pilipczuk}@mimuw.edu.pl`

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

⁴ Department of Computer Science and Engineering,
University at Buffalo, 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).

1 Introduction

Extremal combinatorial questions are ubiquitous 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

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.

appeared as auxiliary lemmas in the literature on formal languages, analysis of infinite-state systems, and applications of formal methods [12, Lemma 6], [11, Sect. 8.1], [8, Lemma 5], [1, Lemma 11], [9, Lemmas 28 and 29], [7, Sect. 5].

A closer inspection 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 this 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 Sect. 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 [2, 4, 14, 17]. Among such systems, counter-based models of computation are a standard abstraction that has 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 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 Etesami et al. [8] 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. [9] reuse the auxiliary lemmas on the length of shortest paths from [8] 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 [13] as the *pumping constant*: a

number N such that any accepted word longer than N can be pumped. In fact, the formulation in [13] 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, 12]. 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; because of space constraints, details are only given in the full version of the paper. The reduction to reachability in one-counter systems was recently implemented by Thakkar et al. [16] on top of ARMC, an abstraction-refinement model checker [15], 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).

Formally, 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 paths between configurations 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 $14n^2$.*

We then generalize the result to arbitrary source and target configurations.

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. For this reason, in this extended abstract we focus on proving Theorem 1, while the reasoning leading to Theorem 2, as well as an extension to OCS with negative counter values, can be found in the full version of the paper. The full version contains also the proofs of all the statements marked with ♠.

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 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 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 $14n^2$.*

As a concrete example, from Corollary 1 it follows 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.

Example 1. [5, 8] 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. [8] 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 $(q_j, -1, q_{j+1 \bmod m})$, a non-zero $(p_0, -1, q_1)$, and zero tests $(s_1, +1, p_1)$, $(q_0, 0, s_2)$. Now paths from $(s_1, 0)$ to $(s_2, 0)$ correspond to solutions of $x \cdot k - y \cdot m = 0$; the shortest path takes the first cycle $x = m$ times and the second cycle $y = k$ times. Exiting the second cycle uses an additional transition, making the length $2km + 1$. 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 Fig. 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., [6, 10].

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 $\gcd(x, y)$ and $\text{lcm}(x, y)$ we denote the greatest common divisor and the least common multiple of x and y , respectively. Recall that $x \cdot y = \gcd(x, y) \cdot \text{lcm}(x, y)$.

We now give all definitions related to one-counter systems that we need later. For the reader's convenience, concepts from Sect. 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 Subsect. 5.2, and the upper bound on the length of the entire path is given in Subsect. 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 \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 (\spadesuit). *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.*

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 σ_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 Fig. 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 $S, T \in \mathfrak{G}$. See Fig. 1 for an illustration. 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 (\spadesuit). *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, $\text{LEN}(\rho_{\text{PREF}}) + \text{LEN}(\rho_{\text{SUFF}}) \leq 5n^2$. 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 Subsect. 5.3; the remainder of the present subsection provides an intuitive sketch of the proof of Lemma 2.

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.

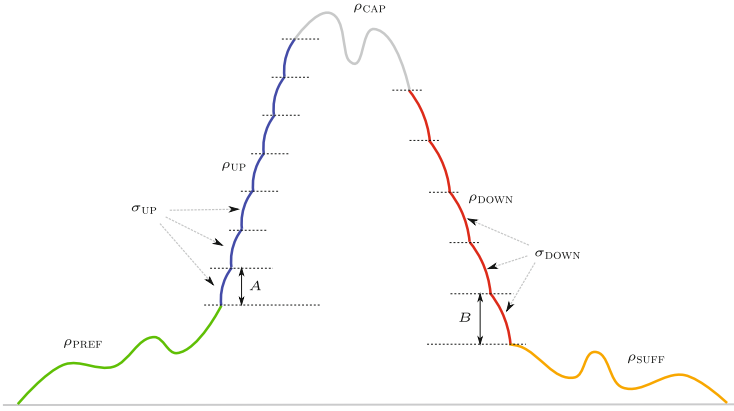


Fig. 1. A normal decomposition of an arc

Since ρ_o is not low, for each $k = 0, 1, 2, \dots, 5n$ we can distinguish the first configuration γ_{i_k} on ρ_o that has counter value k . Consider configurations γ_{i_k} for $2n \leq k \leq 3n$. Among these configurations, some state p repeats in two configurations γ_{i_k} and $\gamma_{i_{k'}}$, for some $2n \leq k < k' \leq 3n$. This means that the part of ρ_o between γ_{i_k} and $\gamma_{i_{k'}}$ corresponds to a cycle σ in the transition multigraph. This cycle has a positive effect on the counter, and hence it is contained in a positively enabled strongly connected component $S \in \mathfrak{S}$. Recall that we cannot simply pump the cycle σ in order to create ρ_{UP} , because by the definition of a normal arc, the cycle that creates ρ_{UP} has to be the pre-chosen cycle σ_S^+ assigned to S . This, however, poses no real difficulty for the following reason. Since p is contained in the same strongly connected component as σ_S^+ , we can travel through S from p to the base of σ_S^+ , then pump it arbitrarily many times, and then go back to p . The part ρ_{DOWN} is defined in a symmetric manner.

More precisely, the consecutive parts ρ_{PREF} , ρ_{UP} , ρ_{CAP} , ρ_{DOWN} and ρ_{SUFF} are defined as follows; the reader is advised to check the description against Fig. 2 while reading. First, ρ_{PREF} is constructed by taking the prefix of ρ_o up to configuration γ_{i_k} , and then traveling along a path σ_{pq} within S from p to q , the base state of σ_S^+ . Then we repeat cycle σ_S^+ a number of times, say a , thus creating ρ_{UP} . The reader should think of a as of a variable, because the possibility of changing this number will be essential for the constructions to follow. Note that we chose k so that $k \geq 2n$ in order to make sure that during these manipulations we never hit nonpositive counter values.

In a symmetrical manner we define ρ_{SUFF} and ρ_{DOWN} . First, we find a configuration $\gamma_{j_{\bar{k}}}$ on ρ_o such that it has some counter value \bar{k} with $2n \leq \bar{k} \leq 3n$, all the configurations later on ρ_o have smaller counter values, and its state \bar{p} belongs to a negatively enabled strongly connected component T . Let \bar{q} be the base state of σ_T^- . Then ρ_{SUFF} is the suffix of ρ_o starting from $\gamma_{j_{\bar{k}}}$, with a path $\sigma_{\bar{q}\bar{p}}$ appended in the front, where $\sigma_{\bar{q}\bar{p}}$ leads from \bar{q} to \bar{p} within T . Also, ρ_{DOWN} is constructed by repeating σ_T^- a number of times, say b . We denote $A = \text{EFF}(\sigma_S^+)$ and $-B = \text{EFF}(\sigma_T^-)$.




Fig. 2. The normalization procedure applied to an arc. Parts ρ_{PREF} , ρ_{UP} , ρ_{CAP} , ρ_{DOWN} , ρ_{SUFF} are depicted in green, blue, grey, red, and orange, respectively, as in Fig. 1 (Color figure online)

The middle part of the path, i.e. ρ_{CAP} , is constructed as follows. First, we follow a path σ_{qp} from q to p within T in order to get back to the original path. Then we repeat the cycle $(\sigma_{pq}\sigma_{qp})$ exactly $\text{gcd}(A, B) - 1$ times, so that the total number of times paths σ_{pq} and σ_{qp} are traversed is $\text{gcd}(A, B)$. It is necessary that the total effect on the counter value from traversing these paths is divisible by $\text{gcd}(A, B)$, so that this effect can be compensated for by selecting a and b appropriately. Then we follow the infix of the original path up to configuration γ_{jk} . Afterwards, a symmetric construction follows: we repeat the cycle $(\sigma_{\bar{p}\bar{q}}\sigma_{\bar{q}\bar{p}})$ exactly $\text{gcd}(A, B) - 1$ times, and then travel from \bar{p} to \bar{q} using $\sigma_{\bar{p}\bar{q}}$. It can be easily seen that one can choose numbers a and b so that the whole construction indeed yields a path from α to β . Here we use the fact that all numbers divisible by $\text{gcd}(A, B)$ can be represented as nonnegative combinations of A and $-B$, i.e., $a \cdot A - b \cdot B$ for some $a, b \in \mathbb{N}$.

Finally, we modify ρ_{CAP} by unpadding it “modulo $\text{gcd}(A, B)$ ” exhaustively: as long as there are two configurations that have the same state, and their counter values are congruent modulo $\text{gcd}(A, B)$, we remove the whole subpath between these configurations. When performing such an operation, the total effect of ρ_{CAP} on the counter changes by a number divisible by $\text{gcd}(A, B)$. Hence, by adjusting once more the numbers a and b we can compensate for this change.

From the description above, it should be relatively clear that the construction yields a normal path satisfying all the conditions apart from the quantitative

ones: (iii) and (iv). For this, some arithmetic calculations are needed to ensure that we can choose a and b small enough so that (iii) and (iv) hold, while all the necessary properties of a and b are satisfied. We remark that in this sketch we have glossed over some technicalities that are used to satisfy conditions (iii) and (iv).

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). If ρ^i is low, choose ρ^i to be the shortest possible low arc from γ_i to γ_{i+1} . If 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 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, of low arcs, of ρ_{PREF}^i and ρ_{SUFF}^i) is small. The following claim follows from a simple application of Proposition 1.

Lemma 3 (♠). *The following inequality holds:*

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

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 $\text{gcd}(\text{EFF}(\sigma_S^+), -\text{EFF}(\sigma_T^-))$.*

In the proof of Lemma 4 we observe that if such configurations δ_i and δ_j existed, then one could repeat σ_S^+ on ρ^i and σ_T^- on ρ_j more times so that the “lifted” configurations δ_i and δ_j would have the same counter value. Then we could cut the whole part of the path between them, thus reducing the number of configurations with counter value zero; this would be a contradiction with the choice of ρ_\circ . The following lemma is a simple corollary of Lemma 4.

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

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:*

$$\sum_{i \in \mathcal{N}_{(S,\cdot)}} \text{LEN}(\rho_{\text{CAP}}^i) \leq A_S \cdot n; \quad (1)$$

$$\sum_{i \in \mathcal{N}_{(\cdot,T)}} \text{LEN}(\rho_{\text{CAP}}^i) \leq n \cdot B_T; \quad (2)$$

$$\text{moreover, } \sum_{i \in \mathcal{N}} \text{LEN}(\rho_{\text{CAP}}^i) \leq n^2. \quad (3)$$

Proof. For (1), 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 (1) is proved, and (2) follows from a symmetric reasoning. The bound (3) follows by summing (1) 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:*

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

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,

$$\sum_{i \in \mathcal{N}} \text{LEN}(\rho_{\text{UP}}^i) \leq 2 \sum_{i \in \mathcal{N}} \frac{L_i n_{S_i}}{A_{S_i}} + 2 \sum_{i \in \mathcal{N}} \frac{\text{lcm}(A_{S_i}, B_{T_i}) \cdot n_{S_i}}{A_{S_i}}. \tag{4}$$

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(1). Thus

$$\sum_{i \in \mathcal{N}} \frac{L_i n_{S_i}}{A_{S_i}} = \sum_{S \in \mathfrak{S}} \sum_{i \in \mathcal{N}(S, \cdot)} \frac{L_i n_{S_i}}{A_{S_i}} \leq \sum_{S \in \mathfrak{S}} n_S \cdot n = n^2. \tag{5}$$

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,

$$\begin{aligned} \sum_{i \in \mathcal{N}} \frac{B_{T_i} \cdot n_{S_i}}{\text{gcd}(A_{S_i}, B_{T_i})} &= \sum_{S, T \in \mathfrak{S}} \sum_{i \in \mathcal{N}(S, T)} \frac{B_{T_i} \cdot n_{S_i}}{\text{gcd}(A_{S_i}, B_{T_i})} \\ &\leq \sum_{S, T \in \mathfrak{S}} n_T \cdot n_S = \sum_{S \in \mathfrak{S}} n_S \cdot \sum_{T \in \mathfrak{S}} n_T = n^2. \end{aligned} \tag{6}$$

By connecting equations (4), (5) and (6) 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(2) instead of Lemma 6(1). □

Combining bounds of Lemma 3, Lemma 6(3), 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.

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

References

1. Alur, R., Černý, P.: 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, 26–28 January 2011, pp. 599–610 (2011)
2. Atig, M.F., Bouajjani, A., Kumar, K.N., Saivasan, P.: On bounded reachability analysis of shared memory systems. In: 34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, 15–17 December 2014, New Delhi, India, pp. 611–623 (2014)
3. Barrett, C., Demri, S., Deters, M.: Witness runs for counter machines. In: Fontaine, P., Ringeissen, C., Schmidt, R.A. (eds.) FroCoS 2013. LNCS, vol. 8152, pp. 120–150. Springer, Heidelberg (2013)
4. Bouajjani, A., Esparza, J., Maler, O.: Reachability analysis of pushdown automata: application to model-checking. In: Mazurkiewicz, A., Winkowski, J. (eds.) CONCUR 1997. LNCS, vol. 1243, pp. 135–150. Springer, Heidelberg (1997)
5. Chistikov, D.: Notes on counting with finite machines. In: 34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, 15–17 December 2014, New Delhi, India, pp. 339–350 (2014)
6. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: Introduction to Algorithms 2nd (edn.) The MIT Press and McGraw-Hill Book Company. ISBN 0-262-03293-7 (2001)
7. Demri, S., Gascon, R.: The effects of bounding syntactic resources on Presburger LTL. *J. Log. Comput.* **19**(6), 1541–1575 (2009)
8. Etesami, K., Wojtczak, D., Yannakakis, M.: Quasi-birth-death processes, tree-like QBDS, probabilistic 1-counter automata, and pushdown systems. *Perform. Eval.* **67**(9), 837–857 (2010)
9. Hofman, P., Mayr, R., Totzke, P.: 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, 25–28 June 2013, pp. 203–212 (2013)
10. Kozen, D.C.: Design and Analysis of Algorithms. Texts and Monographs in Computer Science. Springer, New York (1992). ISBN: 978-3-540-97687-5
11. Lafourcade, P., Lugiez, D., Treinen, R.: Intruder deduction for AC-like equational theories with homomorphisms. Research report LSV-04-16, Laboratoire Spécification et Vérification, ENS Cachan, France, p. 69, November 2004. http://www.lsv.ens-cachan.fr/Publis/RAPPORTS_LSV/PS/rr-lsv-2004-16.rr.ps
12. Lafourcade, P., Lugiez, D., Treinen, R.: Intruder deduction for AC-like equational theories with homomorphisms. In: Giesl, J. (ed.) RTA 2005. LNCS, vol. 3467, pp. 308–322. Springer, Heidelberg (2005)

13. Latteux, M.: Langages à un compteur. *J. Comput. Syst. Sci.* **26**(1), 14–33 (1983)
14. Leroux, J., Schmitz, S.: Demystifying reachability in vector addition systems. In: 30th Annual ACM/IEEE Symposium on Logic in Computer Science. LICS 2015, 6–10 July 2015, Kyoto, Japan, pp. 56–67 (2015)
15. Podelski, A., Rybalchenko, A.: ARMC: the logical choice for software model checking with abstraction refinement. In: Hanus, M. (ed.) PADL 2007. LNCS, vol. 4354, pp. 245–259. Springer, Heidelberg (2007)
16. Thakkar, J., Kanade, A., Alur, R.: Transducer-based algorithmic verification of retransmission protocols over noisy channels. In: Beyer, D., Boreale, M. (eds.) FORTE 2013 and FMOODS 2013. LNCS, vol. 7892, pp. 209–224. Springer, Heidelberg (2013)
17. Thomas, W.: The reachability problem over infinite graphs. In: Frid, A., Morozov, A., Rybalchenko, A., Wagner, K.W. (eds.) CSR 2009. LNCS, vol. 5675, pp. 12–18. Springer, Heidelberg (2009)