Non-deterministic width of regular languages: the unary case

Giovanna D'Agostino[1], Alessandro Minisini[1] and Alberto Policriti[1]

[1]Dipartimento di Matematica, Informatica e Fisica, Università di Udine, Via delle Scienze, 206, Udine, 33100, Italy

Abstract

The (deterministic/non-deterministic) *width* of a regular language $\mathscr{L}$ is the minimum number of chains necessary to decompose a partial order of the states of an automaton (deterministic/non-deterministic) accepting $\mathscr{L}$. We begin the study of the non-deterministic *width* of a regular language, with a special focus on unary languages. We review preliminary results and introduce two possible definitions of co-lexicographic (co-lex) order on non-deterministic finite automata (NFAs)—namely, a global and a local order—along with their corresponding widths. We illustrate that, in general, these two widths do not coincide on NFAs. We then analyze the unary case, proving that the two language-level widths—defined as the minimum width among equivalent NFAs—do coincide. We establish this by showing that the minimum width is realised by automata in Chrobak normal form. Finally, we provide an operational definition for the non-deterministic width of a unary language and we analyze its computational complexity, showing that the problem of computing it is in NP. We further ask if the width of a language is a good indicator for its degree of non-determinism and show that, for unary languages, a constant degree of non-determinism is always achievable, regardless of the width.

Keywords

Regular languages, Automata theory, Co-lexicographical sorting

1. Introduction

In [1], *Wheeler automata* were introduced: automata whose states are endowed with a *total* order satisfying two simple axioms, thereby inducing an order—the co-lex one—on the prefixes of their accepted strings. Intuitively, the axioms generalize to the entire accepted language the rearrangement performed by the Burrows–Wheeler transform on a single string [2]. The benefit is concrete: a Wheeler automaton can be stored in little space and supports fast pattern-matching queries [3].

Not every automaton is Wheeler, yet much of the same algorithmic efficiency and succinctness still holds *in proportion to the width* of a co-lex partial order on the automaton—a measure of how far it is from being totally ordered, with width 1 corresponding to the Wheeler case. Since a language can be recognized by several equivalent automata, it is natural to ask which one minimizes the width, and how this minimum can be characterized. The deterministic case is well understood: the minimum width of a language is characterized via structural properties of its minimum-size automaton [3].

The non-deterministic case, by contrast, remains unexplored. NFAs can be exponentially smaller parametrically in their width than their deterministic counterparts—this is precisely why they are preferred in many applications—so a characterization of the width for NFAs would let the same indexing techniques operate on much smaller automata. Unlike the deterministic case, the two given notions of width do not coincide for NFAs, and whether this also holds at language level remains unknown. Moreover, the decidability of the non-deterministic width of a language is still an open problem.

In this work, we begin this study. We first revisit co-lex orders on NFAs and the two notions of width of an automaton and of a language, already presented in [3]. We then focus on the unary case, showing that automata in Chrobak normal form realise the minimum width of the language and we give an algebraic characterization of the non-deterministic width. Finally, we study the link between width and non-determinism and show that, in the unary case, every width admits a witness automaton with fan-out at most 2.

ICTCS 2026: Italian Conference on Theoretical Computer Science, September 7–9, 2026, Udine, Italy
giovanna.dagostino@uniud.it (G. D'Agostino); minisini.alessandro@spes.uniud.it (A. Minisini); alberto.policriti@uniud.it (A. Policriti)

2. Preliminaries

A *partial order* $(V, \leq)$ is a set V with a reflexive, antisymmetric, and transitive binary relation $\leq$. Elements u, v are *comparable* if $u \leq v$ or $v \leq u$; otherwise they are *incomparable*, written $u \parallel v$. A subset is a *chain* if it is totally ordered, and an *antichain* if every two elements are incomparable. The *width* $\text{width}(V, \leq)$ is the minimum number of chains partitioning V. By Dilworth's theorem [4], $\text{width}(V, \leq)$ equals the maximum size of an antichain.

Fix a totally ordered finite alphabet Σ. The *co-lexicographic (co-lex) order*, denoted $\preceq$, on Σ^* compares strings by reversing them and applying the lexicographic order: for $\alpha, \beta \in \Sigma^*$, we declare $\alpha \preceq \beta$ if and only if the reversed string α^R is lexicographically less than or equal to β^R. This is a total order on Σ^*; as usual, we write $\alpha \prec \beta$ for its strict part, i.e. $\alpha \preceq \beta$ and $\alpha \neq \beta$.

We assume that the notions of deterministic/non-deterministic automata are known to the reader. Given a regular language $\mathscr{L}$, we denote by $\mathscr{D}_{\mathscr{L}}$ the (unique) minimum-size DFA recognizing $\mathscr{L}$. Throughout this paper we assume that every NFA $\mathscr{A}$ satisfies the following properties: (1) every state is reachable from the initial state; (2) every state can reach a final state; (3) the initial state is not reachable from any other state; and (4) all edges entering the same state carry the same label (input-consistency). We denote by $\lambda(u)$ the (uniquely determined) label of all incoming edges of state u. For the initial state s, we set $\lambda(s) = \# \notin \Sigma$ and assume $\# \prec c$ for all $c \in \Sigma$.

If $\mathscr{A} = (Q, s, \delta, F)$ is an NFA, we denote by $\text{Pref}(\mathscr{L}(\mathscr{A}))$ the set of prefixes of words in $\mathscr{L}(\mathscr{A})$. For each state $q \in Q$, let $I_q = \{\alpha \in \text{Pref}(\mathscr{L}(\mathscr{A})) \mid q \in \delta(s, \alpha)\}$ be the set of prefixes of $\mathscr{L}(\mathscr{A})$ reaching q. On the family $\{I_u \mid u \in Q\}$, define:

$$I_u \preceq I_v \iff (\forall \alpha \in I_u)(\forall \beta \in I_v)\,(\{\alpha, \beta\} \not\subseteq I_u \cap I_v \Rightarrow \alpha \preceq \beta)\,.$$

In [5] it was proven that the relation above is a partial order.

3. Co-lex orders and width

Given an NFA $\mathscr{A}$, one can define a partial order $<_{\mathscr{A}}$ on its states that allows $\mathscr{A}$ to be represented using an index, i.e., a succinct structure supporting fast pattern-matching queries [3].

Definition 3.1. Let $\mathscr{A} = (Q, s, \delta, F)$ be an NFA. Then, we define the partial order $<_{\mathscr{A}}$ on the states Q as:

$$u <_{\mathscr{A}} v \iff I_u \preceq I_v \quad \text{and} \quad I_u \neq I_v.$$

This is a global definition, and computing it is far from straightforward: indeed, Theorem 6 in [6] proves that deciding $u <_{\mathscr{A}} v$ is a PSPACE-complete problem.

In the literature, co-lex orders admit instead a local characterization.

Definition 3.2 ([5], Def. 3.1)**.** A *co-lex order* on $\mathscr{A}$ is a partial order $\leq$ on Q satisfying:

1. (Axiom 1) $\lambda(u) \prec \lambda(v) \Rightarrow u < v$ for every $u, v \in Q$.
2. (Axiom 2) If $(u', u), (v', v)$ are edges with $\lambda(u) = \lambda(v)$ and $u < v$, then $u' \leq v'$.

A natural question concerns the relationship between co-lex orders and $<_{\mathscr{A}}$. It is shown in [3] that for DFAs the latter is itself a co-lex order.

Lemma 3.3 ([5], Cor. 1)**.** *Let $\leq$ be a co-lex order on an NFA $\mathscr{A}$. If $u < v$, then $I_u \preceq I_v$.*

Since we are dealing with partial orders, it is natural to consider their width. The complexity of indexing algorithms for $\mathscr{A}$ is parameterized by the width [3], so finding co-lex orders of small width is crucial. We can also extend the notion of width to languages.

Definition 3.4. The *width* of an NFA $\mathscr{A}$, denoted $\text{width}(\mathscr{A})$, is the minimum width of a co-lex order on $\mathscr{A}$:

$$\text{width}(\mathscr{A}) = \min\{\text{width}(\leq) \mid \leq \text{ is a co-lex order on } \mathscr{A}\}.$$

Given a language $\mathscr{L}$, we define $\text{width}^N(\mathscr{L})$ as the minimum width of an NFA recognizing it, in other words:

$$\text{width}^N(\mathscr{L}) = \min\{\text{width}(\mathscr{A}) \mid \mathscr{L}(\mathscr{A}) = \mathscr{L}\}.$$

It can be shown that, on NFAs, maximal co-lex orders can have different widths, which are typically greater than the width of $\leq_{\mathscr{A}}$. On the other hand, it is not known whether, for a given language $\mathscr{L}$, the quantity

$$\text{width}_{glob}(\mathscr{L}) = \min\{\text{width}(\leq_{\mathscr{A}}) \mid \mathscr{L}(\mathscr{A}) = \mathscr{L}\}$$

coincides with $\text{width}^N(\mathscr{L})$. What is known is that $\text{width}_{glob}(\mathscr{L}) \leq \text{width}^N(\mathscr{L})$: indeed, for any $\mathscr{A}$ such that $\text{width}(\mathscr{A}) = \text{width}^N(\mathscr{L})$, we can always merge states u, v with $I_u = I_v$, preserving the language and the co-lex width. After merging, the new automaton $\mathscr{A}'$ is *reduced*, and Lemma 3.3 implies that any co-lex order on $\mathscr{A}'$ is refined by $\leq_{\mathscr{A}'}$.

4. Co-lex order in the unary case

In this section, we focus on the unary case, showing that the highly constrained structure of automata in Chrobak normal form causes the two definitions of width to collapse into a single notion.

Let $\Sigma = \{a\}$. From now on, we consider languages which are both infinite and co-infinite, since the remaining cases have already been studied. In this setting, deterministic automata consist of an acyclic component, usually called the tail, followed by a simple cycle; the resulting structure is often called a *lollipop*. Non-deterministic automata, by contrast, have in general a more complex structure.

Fortunately, Chrobak [7] showed that every unary NFA admits an equivalent automaton with a strict structure.

Definition 4.1 (Chrobak normal form)**.** An NFA $\mathscr{A} = (Q, s, \delta, F)$ over a unary alphabet is in *Chrobak Normal Form (CNF)* if it consists of a single deterministic path of some length $t > 0$ starting from the initial state s (the *tail*), and from the end of the tail, non-deterministic transitions branch off into $m \geq 1$ pairwise disjoint deterministic cycles $C_1, \dots, C_m$ of different lengths.

We show that CNF automata are the natural candidates for realising the minimum width.

Lemma 4.2 (width of CNF automata)**.** *Let $\mathscr{A}$ be an NFA in CNF with tail length t and cycles $C_1, \dots, C_m$ of length, respectively, $c_1, \dots, c_m$. Then*

$$\text{width}(\mathscr{A}) = \text{width}(\leq_{\mathscr{A}}) = \sum_{i=1}^{m} c_i.$$

Lemma 4.3. *For every NFA $\mathscr{A}$ there exists an equivalent automaton $\mathscr{A}_{chr}$ in CNF such that*

$$\text{width}(\mathscr{A}_{chr}) \leq \text{width}(\mathscr{A}).$$

Theorem 4.4. *Let $\mathscr{L}$ be a unary regular language. Then,* $\text{width}_{glob}(\mathscr{L}) = \text{width}^N(\mathscr{L})$.

5. Non-deterministic width of unary languages

In this section, we provide an algebraic characterization of the non-deterministic width of a unary language—which coicides with both the notions of width by Theorem 4.4—and we notice that such measure is related only to the divisors of length of the cycle of the minimum-size automaton.

First, notice that the infiniteness of $\mathscr{L}$ is given by the final states in the cycle of $\mathscr{D}_{\mathscr{L}}$, and thus this component must be periodic. In particular, we can identify representatives for each periodic component of the language.

Definition 5.1. Let $\mathscr{L}$ be a unary language over $\Sigma = \{a\}$, and suppose $\mathscr{D}_{\mathscr{L}}$ consists of a tail of length t and a cycle of length n (so n is the minimal period of $\mathscr{L}$).

$$\text{Res}(\mathscr{L}) = \{(x - t) \bmod n \mid x \geq t,\ a^x \in \mathscr{L}\} \subseteq \mathbb{Z}_n.$$

Consequently, $\mathscr{L}$ splits into its *initial part* $\{a^x \mid x < t,\ a^x \in \mathscr{L}\}$ and its *periodic part* $\{a^{t+x+kn} \mid x \in \mathrm{Res}(\mathscr{L}),\ k \in \mathbb{N}\}$.

Observation 5.2. The tail of a Chrobak automaton can be extended arbitrarily by rotating each cycle one step clockwise for each state added to the tail. Given the length of its cycles as $d_1, \dots, d_m$, it is also possible to extend the tail by $\mathrm{lcm}(d_1, \dots, d_m)$ without changing the cycles. This also applies to DFAs.

The following result helps in pruning the number of candidate lengths for the cycles in a minimum-width NFA.

Lemma 5.3. *Let $\mathscr{L}$ be a unary language whose minimum-size DFA $\mathscr{D}_{\mathscr{L}}$ has tail of length t and cycle of length n, and let $\mathscr{A}$ be a CNF automaton recognizing $\mathscr{L}$ with tail of length $\geq t$ and cycles of lengths $D = \{d_1, \dots, d_m\}$. For all $d \in D$, if $x \in \mathbb{Z}_d$ is such that $\left\{a^{t+x+kd} \mid k \in \mathbb{N}\right\} \subseteq \mathscr{L}$, then also $\left\{a^{t+x+k\,\mathrm{gcd}(n,d)} \mid k \in \mathbb{N}\right\} \subseteq \mathscr{L}$.*

Consequently, each cycle of length d_i can be safely replaced by one of length $\gcd(d_i, n) \leq d_i$, without increasing the width. Without loss of generality, we may therefore assume that the cycle lengths of any candidate automaton are divisors of n.

Theorem 5.4. *Let $\mathscr{L}$ be a unary language and $\mathscr{D}_{\mathscr{L}}$ its minimum-size automaton. Suppose $\mathscr{D}_{\mathscr{L}}$ is composed of a tail of length t and a cycle with length n. Then*

$$\mathrm{width}^N(\mathscr{L}) = \min\left\{\sum_{d\in D} d \;\middle|\; D \subseteq \mathrm{Div}(n),\ \forall r \in \mathrm{Res}(\mathscr{L})\ \exists\, d \in D \quad r + d\mathbb{Z}_n \subseteq \mathrm{Res}(\mathscr{L})\right\}.$$

The right-hand side is well defined since $D = \{n\}$ always satisfies the condition. Call $D \subseteq \mathrm{Div}(n)$ *admissible* if it satisfies this condition. Admissibility of D can be verified in polynomial time in n, so the decision problem $\mathrm{width}^N(\mathscr{L}) \leq k$ is in NP. Moreover, since $d(n) = |\mathrm{Div}(n)| = n^{o(1)}$, exhaustive search over the $2^{d(n)} = 2^{n^{o(1)}}$ subsets of $\mathrm{Div}(n)$ runs in subexponential time, so the problem cannot be NP-hard unless the Exponential Time Hypothesis fails. Finally, if $d, d' \in D$ with $d \mid d'$, then $r + d'\mathbb{Z}_n \subseteq r + d\mathbb{Z}_n$ for every $r \in \mathbb{Z}_n$, so d can be removed from D: the optimal D^* is an antichain in the divisor lattice $(\mathrm{Div}(n), |)$.

6. Relation between width and non-determinism

In [3] it is shown that the classical determinization blow-up is governed by the width of the input NFA. This suggests a further, more structural question: how much *non-determinism* is actually needed to realise a given width? We now make this question precise, and answer it for unary languages.

Definition 6.1. Let $\mathscr{A} = (Q, s, \delta, F)$ be an NFA. The *degree of non-determinism* of $\mathscr{A}$ is

$$\mathrm{nd}(\mathscr{A}) = \max_{u\in Q,\ a\in\Sigma} |\delta(u, a)|.$$

Clearly $\mathrm{nd}(\mathscr{A}) = 1$ if and only if $\mathscr{A}$ is a DFA. More generally, given a class $\mathscr{N}$ of NFAs, let $\mathscr{C} = \{\mathscr{L}(\mathscr{A}) \mid \mathscr{A} \in \mathscr{N}\}$ be the class of languages it describes. For $\mathscr{L} \in \mathscr{C}$, define the *degree of non-determinism of $\mathscr{L}$ relative to $\mathscr{N}$* as

$$\mathrm{nd}_{\mathscr{N}}(\mathscr{L}) = \min\{\mathrm{nd}(\mathscr{A}) \mid \mathscr{A} \in \mathscr{N},\ \mathscr{L}(\mathscr{A}) = \mathscr{L}\}.$$

Remark 6.2. If $\mathscr{N}$ is the class of all NFAs, then $\mathscr{C}$ is the class of all regular languages, and $\mathrm{nd}_{\mathscr{N}}(\mathscr{L}) = 1$ for every $\mathscr{L} \in \mathscr{C}$, witnessed by the minimal DFA of $\mathscr{L}$.

Recall that an NFA $\mathscr{A}$ is a *Wheeler automaton* if $\mathrm{width}(\mathscr{A}) = 1$, i.e., if it admits a co-lex order that is total. Let $\mathscr{W}$ be the class of Wheeler NFAs, and let $\mathscr{C}_1 = \{\mathscr{L}(\mathscr{A}) \mid \mathscr{A} \in \mathscr{W}\}$ be the class of *Wheeler languages*. Since every Wheeler NFA admits an equivalent Wheeler DFA [8], we again have $\mathrm{nd}_{\mathscr{W}}(\mathscr{L}) = 1$ for every $\mathscr{L} \in \mathscr{C}_1$.

In both examples above, the degree of non-determinism collapses to 1: every language describable by some class of NFAs of width 1 is in fact recognized by an automaton that is simultaneously deterministic *and* of width 1. This motivates asking whether the same phenomenon, suitably relativized, holds at every width $k > 1$.

Definition 6.3. Let $\mathscr{L}$ be a regular language with $\text{width}^N(\mathscr{L}) = k$. The *degree of non-determinism of $\mathscr{L}$ (at its minimum width)* is

$$\text{nd}^N(\mathscr{L}) = \min\{\text{nd}(\mathscr{A}) \mid \mathscr{L}(\mathscr{A}) = \mathscr{L},\ \text{width}(\mathscr{A}) = k\}.$$

Problem 6.4. Does there exist a function $f\colon \mathbb{N} \to \mathbb{N}$ such that, for every regular language $\mathscr{L}$,

$$\text{nd}^N(\mathscr{L}) \leq f(\text{width}^N(\mathscr{L}))\,?$$

By Remark 6.2, $f(1) = 1$ is achievable. We know that, in the unary case, the degree of non-determinism is not greater than 2.

Theorem 6.5. *Let $\mathscr{L}$ be a unary regular language. Then* $\text{nd}^N(\mathscr{L}) \leq 2$.

Lemma 8 in [5] implies that there are unary languages for which no DFA has width equal to $\text{width}^N(\mathscr{L})$, thus for those languages $\text{nd}^N(\mathscr{L}) = 2$.

7. Conclusions

We initiated the study of the non-deterministic width of regular languages, focusing on the unary case. Starting from the two natural notions of co-lexicographic width for NFAs we showed that, although these notions may differ at the automaton level, they coincide for unary languages at the language level. The key structural ingredient is a "classic" for unary languages: Chrobak Normal Form, by which we proved that, for every unary NFA, the corresponding CNF automaton does not increase the width. As a consequence, the minimum non-deterministic width of a unary language is always realised by an automaton in CNF, where the width is exactly the sum of the lengths of its cycles. This allows us to give an algebraic characterization of the width in terms of the divisors of the period of the minimum DFA and of the residue classes defining the ultimately periodic part of the language. This characterization also yields algorithmic consequences: the problem of deciding whether the non-deterministic width of a unary language is at most a given value belongs to NP and the divisor-based formulation gives a subexponential exhaustive-search procedure. Moreover, the optimal solutions can be restricted to antichains in the divisor lattice, further clarifying the combinatorial structure of the problem.

Finally, we investigated the relation between width and the degree of non-determinism. For unary languages, we showed that minimum width can always be realised with degree of non-determinism at most two. Thus, in the unary setting, large width does not necessarily imply large branching. This separates width from a simple measure of non-deterministic fan-out and suggests that width captures a different structural feature of NFAs.

Several questions remain open. The most important one is whether the equality between global and local language width extends beyond unary languages. More generally, it remains to understand whether the non-deterministic width of an arbitrary regular language is decidable. Along this line, a further interesting question is if some analogue of Chrobak Normal Form can be put forward for non-unary alphabets.

Another natural direction is to investigate whether bounds on the degree of non-determinism in terms of width exist for larger classes of regular languages. These questions would ultimately contribute to a broader understanding of the role of co-lexicographic structure in succinct automata representations and indexing.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] J. Alanko, G. D'Agostino, A. Policriti, N. Prezza, Regular languages meet prefix sorting, in: Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '20, Society for Industrial and Applied Mathematics, USA, 2020, p. 911–930.

[2] M. Burrows, D. J. Wheeler, A block-sorting lossless data compression scheme, Digital Equipment Corporation, Palo Alto, CA 124 (1994).

[3] N. Cotumaccio, G. D'Agostino, A. Policriti, N. Prezza, Co-lexicographically ordering automata and regular languages - part i, Journal of the ACM (JACM) 70 (2023) 1–36.

[4] R. P. Dilworth, A decomposition theorem for partially ordered sets, Annals of Mathematics 51 (1950) 161–166.

[5] N. Cotumaccio, G. D'Agostino, A. Policriti, N. Prezza, Co-lexicographically ordering automata and regular languages – part ii, 2023. URL: https://arxiv.org/abs/2102.06798. arXiv:2102.06798.

[6] G. D'Agostino, D. Martincigh, A. Policriti, Ordering regular languages and automata: Complexity, Theoretical Computer Science 949 (2023) 113709.

[7] M. Chrobak, Finite automata and unary languages, Theoretical Computer Science 47 (1986) 149–158.

[8] J. Alanko, G. D'Agostino, A. Policriti, N. Prezza, Wheeler languages, Information and Computation 281 (2021) 104820.

A. Proofs

A.1. Preliminary Lemma

We begin with a preliminary lemma that is used in the proofs.

Proposition A.1. *Let $A = \{r_1 + kp_1 \mid k \in \mathbb{N}\}$ and $B = \{r_2 + kp_2 \mid k \in \mathbb{N}\}$ with $p_1, p_2 \in \mathbb{N}_{>0}$, and assume that A and B are* not eventually equal, *i.e. it is not the case that $p_1 = p_2$ and $r_1 \equiv r_2 \pmod{p_1}$. Identifying unary strings with their length, so that for $X, Y \subseteq \mathbb{N}$*

$$X \preceq Y \iff \forall a \in X \forall b \in Y \left(\{a, b\} \not\subseteq X \cap Y \Rightarrow a \leq b\right),$$

the sets A and B are incomparable for $\preceq$, that is $A \parallel B$.

Proof. Let $g = \gcd(p_1, p_2)$ and $l = \operatorname{lcm}(p_1, p_2)$. If $r_1 \not\equiv r_2 \pmod{g}$, then $A \cap B = \emptyset$: a common element x would satisfy $x \equiv r_1 \pmod{p_1}$ and $x \equiv r_2 \pmod{p_2}$, hence $x \equiv r_1 \equiv r_2 \pmod{g}$, a contradiction. Thus $A \setminus B = A$ and $B \setminus A = B$ are both infinite (case 1).

If instead $r_1 \equiv r_2 \pmod{g}$, by the Chinese Remainder Theorem the system $x \equiv r_1\ (p_1),\ x \equiv r_2\ (p_2)$ is solvable and its solutions form a single class modulo l; thus $A \cap B$ is a progression of period l, in particular infinite. Since A has period p_1 and $A \cap B$ has period l, the set $A \setminus B$ is finite iff $l = p_1 \iff p_2 \mid p_1$; symmetrically $B \setminus A$ is finite if and only if $p_1 \mid p_2$. As A and B are not eventually equal we have $p_1 \neq p_2$, so at most one of these divisibilities holds. Hence either none of the two periods divides the other and both differences are infinite (case 1), or exactly one does (case 2).

In case 1, both $A \setminus B$ and $B \setminus A$ are infinite, so pick $a \in A \setminus B$ with $a > \min B$: so $A \not\preceq B$. Symmetrically, pick $b \in B \setminus A$ with $b > \min A$: then $B \not\preceq A$. Thus $A \parallel B$.

In case 2, we can assume without loss of generality that $p_1 \mid p_2$ with $p_1 < p_2$, so $A \setminus B$ is infinite while $B \setminus A$ is finite. For $A \not\preceq B$, pick $a \in A \setminus B$ with $a > \min B$, as in case 1. For $B \not\preceq A$ fix $a \in A \setminus B$ (which exists, $A \setminus B$ being infinite) and choose $b \in B$ with $b > a$ (possible since B is infinite). Then $a \in A$, $b \in B$, and $a \notin B$ gives $\{a, b\} \not\subseteq A \cap B$ while $b > a$; so $B \not\preceq A$. Hence $A \parallel B$. □

A.2. Main proofs

Proof of Lemma 4.2. We first show that $\mathrm{width}(\mathscr{A}) \geq \sum_{i=1}^{m} c_i$. For a fixed co-lex order $\leq$ on $\mathscr{A}$, let $u \in C_i$ and $v \in C_j$ with $u \neq v$. If $i = j$, then $u \parallel v$ by Lemma 5 in [5]; otherwise, I_u and I_v are not eventually equal (distinct cycles have distinct lengths), so Proposition A.1 gives $I_u \parallel I_v$; by Lemma 3.3 it follows that $u \parallel v$. Hence, all states in cycles are mutually incomparable. Notice that this property holds also for $\leq_{\mathscr{A}}$.

To prove equality, we check that $\leq_{\mathscr{A}}$ is a co-lex order: since Axiom 1 always holds, it suffices to verify Axiom 2. Consider $u <_{\mathscr{A}} v$ with $u \in \delta(u', a)$ and $v \in \delta(v', a)$; we must show $u' \leq_{\mathscr{A}} v'$. Note that u must be a state in the tail, since states in cycles accept strings longer than t, so u' is in the tail as well. Let $\{a^i\} = I_{u'}$ and $a^j \in I_{v'}$; notice that $a^i \notin I_{v'}$. Since $\{a^{i+1}\} = I_u$ and $a^{j+1} \in I_v$, $u <_{\mathscr{A}} v$ gives $a^{i+1} < a^{j+1}$, hence $i+1 < j+1 \implies i < j$. So $a^i \prec a^j$, and since $a^i \notin I_{v'}$, it follows that $I_{u'} \prec I_{v'}$. As the incomparable states are exactly those in the cycles, the width of the order is $\sum_{i=1}^{m} c_i$. □

Proof sketch of Lemma 4.3. We will use the notion of period of $A \subseteq \{a\}^*$: p is a period for A if $a^x \in A \implies a^{x+p} \in A$. For a non-singleton strongly connected component (scc) D of $\mathscr{A}$, let $\mathrm{period}(D)$ be the minimum p for which there exists $n_0 \in \mathbb{N}$ such that p is a period for $I_q \cap \{a^x \mid x > n_0\}$, for all $q \in D$. Let $D_1, \dots, D_k$ be the non-singleton sccs of $\mathscr{A}$ and let $p_1 < \cdots < p_m$ be the distinct values among $\mathrm{period}(D_1), \dots, \mathrm{period}(D_k)$. Let $\mathscr{A}_{chr}$ be an NFA in CNF with cycles $C_1, \dots, C_m$ of pairwise distinct lengths $|C_i| = p_i$, equivalent to $\mathscr{A}$. The existence of such an NFA equivalent to $\mathscr{A}$ can be proved in a way similar to the construction made in [7]; then $\mathrm{width}(\mathscr{A}_{chr}) = \sum_{i=1}^{m} p_i$ by Lemma 4.2.

For each $1 \leq i \leq m$, fix a component D_i with $\mathrm{period}(D_i) = p_i$ and a simple cycle $K_i \subseteq D_i$. By Lemma 5 in [5] the states of K_i are pairwise incomparable in every co-lex order on $\mathscr{A}$. Since $|K_i|$ is itself a period for all I_q with $q \in K_i$ and p_i is minimum, $|K_i|$ must be a multiple of p_i (otherwise $\gcd(|K_i|, p_i)$ is a smaller period for I_q with $q \in K_i$, in a way similar to Lemma 5.3), so the cycle K_i has at least p_i states; choose $S_i \subseteq K_i$ with $|S_i| = p_i$. By strong connectivity of D_i, each I_q with q in S_i is infinite and ultimately periodic of minimum period p_i, and no I_q is cofinite (since the state q reaches a final state, a cofinite I_q would make $\mathscr{L}$ cofinite).

We claim $S = \bigcup_i S_i$ is an antichain in every co-lex order $\leq$ on $\mathscr{A}$. States within the same S_i are incomparable by Lemma 5 in [5]. For $u \in S_i$, $v \in S_j$ with $i \neq j$, the sets I_u, I_v are infinite, ultimately periodic, with distinct minimum periods $p_i \neq p_j$; hence they are not eventually equal, so by a generalization of Proposition A.1, $I_u \parallel I_v$ and thus $u \parallel v$ by Lemma 3.3.

Therefore

$$\mathrm{width}(\mathscr{A}) \;\geq\; \sum_{i=1}^{m} |S_i| \;=\; \sum_{i=1}^{m} p_i \;=\; \mathrm{width}(\mathscr{A}_{chr}).$$

□

Proof of Lemma 5.3. By Bézout's theorem, we have that there exist $k, h \in \mathbb{Z}$ such that $kn + hd = \gcd(n, d)$. So, $hd \equiv_n \gcd(n, d)$ and

$$t + x + hd \equiv_n t + x + \gcd(n, d). \tag{1}$$

Since $a^{t+x+hd} \in \mathscr{L}$, the state that we reach reading a^{t+x+hd} in $\mathscr{D}_{\mathscr{L}}$ is final, and so also $a^{t+x+\gcd(n,d)} \in \mathscr{L}$. Then, adding l times respectively hd and $\gcd(n, d)$ to the left and right side of (1), we obtain that $a^{t+x+l\gcd(n,d)} \in \mathscr{L}$. □

Proof sketch of Theorem 5.4. Let $\mathscr{A}$ be any NFA recognizing $\mathscr{L}$. By Lemma 4.3 and Observation 5.2, we may assume $\mathscr{A}$ is in CNF with tail length $\geq t$ and cycle lengths $d_1, \dots, d_m$ dividing n. Then $D = \{d_1, \dots, d_m\}$ is admissible and $\mathrm{width}(\mathscr{A}) = \sum_{d \in D} d$.

Given any admissible D, we construct a CNF automaton recognizing $\mathscr{L}$: build an automaton in CNF with tail t and one cycle of length d for each $d \in D$. In cycle d, for every $r \in \mathrm{Res}(\mathscr{L})$ mark state $(r \bmod d)$ final if and only if $r + d\mathbb{Z}_n \subseteq \mathrm{Res}(\mathscr{L})$. Then, for $r = (r \bmod d) + hd$, the automaton accepts a^{t+r+kn} by traversing cycle d a total of $h + k(n/d)$ times from state $r \bmod d$, and does not accept other strings since $r + d\mathbb{Z}_n = (r \bmod d) + d\mathbb{Z}_n \subseteq \mathrm{Res}(\mathscr{L})$. □

Proof sketch of Theorem 6.5. Let $\mathscr{A}$ be an automaton in CNF such that $\text{width}(\mathscr{A}) = \text{width}^N(\mathscr{L})$. Then $\mathscr{A}$ consists of a tail of length t followed by cycles $C_1, \ldots, C_m$, so $\text{nd}(\mathscr{A}) = m$, witnessed by the unique non-deterministic state at the end of the tail, which we denote by u. We show that there is an equivalent automaton $\mathscr{A}'$ with $\text{width}(\mathscr{A}') = \text{width}(\mathscr{A})$ and $\text{nd}(\mathscr{A}') = 2$, obtained by replacing u with a suitable gadget. Let $\mathscr{T}$ be a complete binary tree with at least m leaves, and label all its edges with a, so that every node at level k is reached by the same string a^k. The partial order $\leq_{\mathscr{T}}$ that first compares nodes by level, and then by left-to-right sibling order, is a co-lex order on $\mathscr{T}$. We define $\mathscr{A}'$ as follows: start with a tail of length $t - 1$, attach the tree $\mathscr{T}$ to its last state, then choose m leaves of $\mathscr{T}$ and attach the cycles $C_1, \ldots, C_m$ to them, one per leaf. The final states of the tail are unchanged, while those of each cycle must be rotated as needed to preserve the language (Observation 5.2); a state of $\mathscr{T}$ at level k is marked final whenever $a^{t-1+k} \in \mathscr{L}$. A co-lex order on $\mathscr{A}'$ realising the minimum width is then obtained by extending $<_{\mathscr{A}} \cup <_{\mathscr{T}}$ so that the states of $\mathscr{T}$ lie above those of the tail and below those of the cycles. Indeed, all the nodes in the tree are comparable with any other states, so the longest antichain is again $\bigcup_{i=1}^{m} C_i$. $\square$