# Ordering Regular Languages and Automata: Complexity 

Giovanna D'Agostino, Davide Martincigh, and Alberto Policriti<br>University of Udine, Italy


#### Abstract

Given an order of the underlying alphabet we can lift it to the states of a finite deterministic automaton: to compare states we use the order of the strings reaching them. When the order on strings is the co-lexicographic one and this order turns out to be total, the DFA is called Wheeler. This recently introduced class of automata-the Wheeler automata - constitute an important data-structure for languages, since it allows the design and implementation of a very efficient tool-set of storage mechanisms for the transition function, supporting a large variety of substring queries. In this context it is natural to consider the class of regular languages accepted by Wheeler automata, i.e. the Wheeler languages. An inspiring result in this area is the following: it has been shown that, as opposed to the general case, the classic determinization by powerset construction is polynomial on Wheeler automata. As a consequence, most classical problems, when considered on this class of automata, turn out to be "easy"-that is, solvable in polynomial time. In this paper we consider computational problems related to Wheelerness, but starting from non-deterministic automata. We also consider the case of reduced non-deterministic ones - a class of NFA where recognizing Wheelerness is still polynomial, as for DFA's. Our collection of results shows that moving towards non-determinism is, in most cases, a dangerous path leading quickly to intractability. Moreover, we start a study of "state complexity" related to Wheeler DFA and languages, proving that the classic construction for the intersection of languages turns out to be computationally simpler on Wheeler DFA than in the general case. We also provide a construction for the minimum Wheeler DFA recognizing a given Wheeler language.


Keywords: Regular languages • Finite Automata • Wheeler Automata
Ordering Languages.

## 1 Introduction

A simple and natural way of efficiently storing and composing regular languages presented by their accepting automata is by exploiting some kind of order imposed on their collection of states. After all, ordering a collection of objects is very often a way to shed light on their internal structure and ease their manipulation.

One way of ordering the states of a finite automaton is to consider their incoming languages-that is, the set of strings reaching the given states-and proposing a way to compare them. If we fix an order on the underlying alphabet $\Sigma$ and consider states as ending points of strings, we are naturally invited to start from their last character (the final one on the path reaching the state) and proceed backwards. This results in using the so-called co-lexicographic order over $\Sigma^{*}$. Since incoming languages of different states of a deterministic automaton $\mathcal{D}$ do not intersect, the co-lexicographic order can easily be lifted to the states of $\mathcal{D}: q \leq_{\mathcal{D}} q^{\prime}$ if all strings of the incoming language of $q$ are co-lexicographically smaller than any string of the incoming language of $q^{\prime}$. This order turns out to be very useful, allowing to store $\mathcal{D}$ using a succint index, that is, a space-saving data structure that supports fast matching queries [5]. It turns out that the complexity of constructing such an index depends on the width of the order $\leq_{\mathcal{D}}$ (see [5]), the best possible case being the one where $\leq_{\mathcal{D}}$ is a total order. In the latter case $\mathcal{D}$ is called a Wheeler automaton, and in [1] it has been proved that recognizing Wheelerness is an easy task over DFA's.

When moving from DFA's to NFA's things become more complicate and two possible approaches were considered:

- The first one consists in identifying some local properties of $\leq_{\mathcal{D}}$, used to define a general notion of a co-lex (possibly partial) order over the states of an NFA (see [2]). Turning back to DFA's, one can easily prove that $\leq_{\mathcal{D}}$ is the maximum co-lex (partial) order over $\mathcal{D}$. In general, co-lex orders over NFA's can still be used for indexing, with index-construction complexity parametric on the width (i.e. the maximum length of an anti-chain in $\leq_{\mathcal{D}}$ ) of the co-lex order. Unfortunately, co-lex orders are not as well behaving on NFA's as they are on DFA's: over an NFA we cannot guarantee the existence of a maximum co-lex order and also finding a maximal one turns out to be an NP-complete problem [6]. To overcome such difficulties, in [2] a new class of automata was introduced: the reduced NFA's. On reduced NFA's distinguished states have different incoming languages. While allowing non-determinism, the reduced NFA's share with DFA's the good behaviour of co-lex orders: any reduced NFA possesses a polynomial time computable, maximum co-lex order, so that recognizing Wheelerness is no longer an NP-complete problem over them.
- The second approach consists in generalizing the definition of $\leq_{\mathcal{D}}$ over NFA's states, by defining an order depending directly on the incoming languages. Such generalization must now take care of the fact that incoming languages may intersect. Actually, since in an NFA $\mathcal{A}$ there could be different states with the same incoming language, when lifting the order to the state of $\mathcal{A}$ we must be careful not to identify states with the same incoming languages.

As far as the first approach is concerned, in this paper we prove that deciding whether a language is Wheeler, i.e. whether it is recognized by an NFA with a total co-lex order, is PSPACE-complete. This remains the case even if we restrict to reduced NFA's. Note that the same problem using a recognizing DFA was proved to be easy (polynomially computable) in [1].

Regarding the second approach, even though the proposed partial order was shown to be useful for indexing [4], we first need to compute it. In this paper we prove that the task of computing $\leq_{\mathcal{A}}$ is difficult over NFA's, even on the class of reduced NFA's. Actually, as a corollary of this fact we also see that recognizing reduced-ness is a difficult task. The proof relies on the fact that the universality problem is PSPACE-complete over reduced NFA (as for the whole class of non-deterministic automata).

In the last part of the paper we go back to DFA's and tackle the problem of establishing the state complexity of the intersection of two Wheeler automata. We prove that equipping the input automata with an order on their collection of states allows us to do much better than in the general case: the standard procedure now turns out of a complexity proportional to the sum the sizes of the input automata. Our final result regards the (difficult) problem of computing the minimum-size Wheeler automaton, starting from the minimum automaton accepting a given Wheeler language.

### 1.1 Preliminaries

First of all, we fix some notation. Given a total order $(Z,<)$ we say that a subset $I \subseteq Z$ is an interval if, for any $x, y, z \in Z$ with $x<y<z$, if $x, z \in I$ then $y \in I$. Let $\Sigma$ denote a finite alphabet endowed with a total order $(\Sigma, \prec)$. We denote by $\Sigma^{*}$ the set of finite strings over $\Sigma$, with $\varepsilon$ being the empty string. We extend the order $\prec$ over $\Sigma$ to the co-lexicographic order $\left(\Sigma^{*}, \prec\right)$, where $\alpha \prec \beta$ if and only the reverse of $\alpha$, i.e. $\alpha$ read from the right to the left, precedes lexicographically the reverse of $\beta$. Given two strings $\alpha, \beta \in \Sigma^{*}$, we denote by $\alpha \dashv \beta$ the property that $\alpha$ is a suffix of $\beta$. For a language $\mathcal{L} \subseteq \Sigma^{*}$, we denote by $\operatorname{Pref}(\mathcal{L})$ the set of prefixes of strings in $\mathcal{L}$. We denote by $\mathcal{A}=\left(Q, q_{0}, \delta, F, \Sigma\right)$ a finite automaton (NFA), with $Q$ as set of states, $q_{0}$ initial state, $\delta: Q \times \Sigma \rightarrow 2^{Q}$ transition function, and $F \subseteq Q$ final states. The size of $\mathcal{A}$, denoted by $|\mathcal{A}|$, is defined to be $|Q|$. An automaton is deterministic (DFA) if $|\delta(q, a)| \leq 1$, for all $q \in Q$ and $a \in \Sigma$. As customary, we extend $\delta$ to operate on strings as follows: for all $q \in Q$, $a \in \Sigma$ and $\alpha \in \Sigma^{*}$

$$
\delta(q, \varepsilon)=\{q\}, \quad \delta(q, \alpha a)=\bigcup_{v \in \delta(q, \alpha)} \delta(v, a)
$$

We denote by $\mathcal{L}(\mathcal{A})=\left\{\alpha \in \Sigma^{*}: \delta\left(q_{0}, \alpha\right) \cap F \neq \emptyset\right\}$ the language accepted by the automaton $\mathcal{A}$. We assume that every automaton is trimmed, that is, every state is reachable from the initial state and every state can reach at least one final state. Note that this assumption is not restrictive, since removing every state not reachable from $q_{0}$ and every state from which is impossible to reach a final state from an NFA, can be done in linear time and does not change the accepted language. It immediately follows that:

- there might be only one state without incoming edges, namely $q_{0}$;
- every string that can be read starting from $q_{0}$ belongs to $\operatorname{Pref}(\mathcal{L})$.

We will often make use of the notion of the incoming language of a state of an NFA, defined as follows.

Definition 1 (Incoming language). Let $\mathcal{A}=\left(Q, q_{0}, \delta, F, \Sigma\right)$ be an NFA and let $q \in Q$. The incoming language of $q$, denoted by $I_{q}$, is the set of strings that can be read on $\mathcal{A}$ starting from $q_{0}$ and ending in $q$. In other words, $I_{q}$ is the language recognized by the automaton $\mathcal{A}_{q}=\left(Q, q_{0}, \delta,\{q\}, \Sigma\right)$.

The class of Wheeler automata has been recently introduced in [8]. An automaton in this class has the property that there exists a total order on its states that is propagated along equally labeled transition. Moreover, the order must be compatible with the underlying order of the alphabet:

Definition 2 (Wheeler Automaton). A Wheeler $N F A$ (WNFA) $\mathcal{A}$ is an $N F A$ $\left(Q, q_{0}, \delta, F, \Sigma\right)$ endowed with a binary relation $<$, such that: $(Q,<)$ is a linear order having the initial state $q_{0}$ as minimum, $q_{0}$ has no in-going edges, and the following two (Wheeler) properties are satisfied. Let $v_{1} \in \delta\left(u_{1}, a_{1}\right)$ and $v_{2} \in$ $\delta\left(u_{2}, a_{2}\right)$ :
(i) $a_{1} \prec a_{2} \rightarrow v_{1}<v_{2}$
(ii) $\left(a_{1}=a_{2} \wedge u_{1}<u_{2}\right) \rightarrow v_{1} \leq v_{2}$.

A Wheeler DFA (WDFA) is a deterministic WNFA.
Remark 1. A consequence of Wheeler property (i) is that $\mathcal{A}$ is input-consistent, that is all transitions entering a given state $u \in Q$ have the same label: if $u \in \delta(v, a)$ and $u \in \delta(w, b)$, then $a=b$. Therefore the function $\lambda: Q \backslash\left\{q_{0}\right\} \rightarrow \Sigma$ that associate to each state the unique label of its incoming edges is well defined. For the state $q_{0}$, the only one without incoming edges, we set $\lambda\left(q_{0}\right):=\#$.

In Figure 1 is depicted an example of a WDFA.
Remark 2. Note that, for a fixed (i.e. constant in size) alphabet, requiring an automaton to be input-consistent is not computationally demanding. In fact, given an NFA $\mathcal{A}=\left(Q, q_{0}, \delta, F, \Sigma\right)$ we can build an equivalent, input-consistent one just by creating, for each state $q \in Q$, at most $|\Sigma|$ copies of $q$, that is, one for each different incoming label of $q$. This operation can be performed in $O(|Q| \cdot|\Sigma|)$ time.

In [8] it is shown that WDFA's have a property called path coherence: let $\mathcal{A}=\left(Q, q_{0}, \delta, F, \Sigma\right)$ be a WDFA according to the order $(Q,<)$. Then for every interval of states $I=\left[q_{i}, q_{j}\right]$ and for all $\alpha \in \Sigma^{*}$, the set $J$ of states reachable starting from any state of $I$ by reading $\alpha$ is also an interval. Path coherence allows us to transfer the order $<$ over the states of $Q$ to the co-lexicographic order $\prec$ over the strings entering the states: two states $q$ and $p$ satisfy $q<p$ if and only if $\forall \alpha \in I_{q} \forall \beta \in I_{p}(\alpha \prec \beta)$ holds (again proved in [1]).

A consequence of this fact is that a WDFA admits an unique order of its states that makes it Wheeler and this order is univocally determined by the colexicographic order of any string entering its states (the order $\leq_{\mathcal{D}}$ mentioned in


Fig. 1. A WDFA's $\mathcal{A}$ recognizing the language $\mathcal{L}_{d}=a c^{*}+d c^{*} f$. Condition (i) of Definition 2 implies input consistency and induces the partial order $q_{1}<q_{2}, q_{3}<q_{4}<$ $q_{5}$. From condition (ii) it follows that $\delta\left(q_{1}, c\right) \leq \delta\left(q_{4}, c\right)$, thus $q_{2}<q_{3}$. Therefore, the only order that could make $\mathcal{A}$ Wheeler is $q_{0}<q_{1}<q_{2}<q_{3}<q_{4}<q_{5}$. The reader can verify that condition (ii) holds for each pair of equally labeled edges.
the introduction). This result is important for two different reasons. First of all, it makes possible to decide in polynomial time whether a DFA is Wheeler: for each state $q$, pick a string $\alpha_{q}$ entering it and order the states reflecting the colexicographic order of the strings $\left\{\alpha_{q}: q \in Q\right\}$; then check if the order satisfies the Wheeler conditions. Secondly, it is the key to adapt Myhill-Nerode Theorem to Wheeler automata. We recall the following definition.

Definition 3 (Myhill-Nerode equivalence). Let $\mathcal{L} \subseteq \Sigma^{*}$ be a language. Given a string $\alpha \in \Sigma^{*}$, we define the right context of $\alpha$ as

$$
\alpha^{-1} \mathcal{L}:=\left\{\gamma \in \Sigma^{*}: \alpha \gamma \in \mathcal{L}\right\}
$$

and we denote by $\equiv_{\mathcal{L}}$ the Myhill-Nerode equivalence on $\operatorname{Pref}(\mathcal{L})$ defined as

$$
\alpha \equiv_{\mathcal{L}} \beta \Longleftrightarrow \alpha^{-1} \mathcal{L}=\beta^{-1} \mathcal{L}
$$

The (classic) Myhill-Nerode Theorem, among many other things, establishes a bijection between equivalence classes of $\equiv_{\mathcal{L}}$ and the states of the minimum DFA recognizing $\mathcal{L}$. This minimum automaton is also unique up to isomorphism and a similar result, fully proved in [2], holds for Wheeler languages as well. In order to state such an analogous of Myhill-Nerode Theorem for Wheeler languages, the equivalence $\equiv_{\mathcal{L}}$ is replaced by the equivalence $\equiv_{\mathcal{L}}^{c}$ defined below.

Definition 4. The input consistent, convex refinement $\equiv_{\mathcal{L}}^{c}$ of $\equiv_{\mathcal{L}}$ is defined as follows. $\alpha \equiv_{\mathcal{L}}^{c} \beta$ if and only if
$-\alpha \equiv_{\mathcal{L}} \beta$,
$-\alpha$ and $\beta$ end with the same character,

- for all $\gamma \in \operatorname{Pref}(\mathcal{L})$, if $\min (\alpha, \beta) \preceq \gamma \preceq \max (\alpha, \beta)$, then $\alpha \equiv_{\mathcal{L}} \gamma \equiv_{\mathcal{L}} \beta$.

The Myhill-Nerode Theorem for Wheeler languages proves that there exists a minimum (in the number of states) WDFA recognizing $\mathcal{L}$. As in the classic case, states of the minimum automaton are, in fact, $\equiv_{\mathcal{L}^{\prime}}^{c}$-equivalence classes, this time consisting of intervals of strings. Also, such WDFA is unique up to isomorphism.

Theorem 1. (see [2]) If $\mathcal{A}=(Q, s, \delta,<, F)$ is a WNFA with $|Q|=n$ and $\mathcal{L}=\mathcal{L}(\mathcal{A})$, then there exists a unique minimum-size WDFA $\mathcal{B}$ with $2 n-1-|\Sigma|$ states such that $\mathcal{L}=\mathcal{L}(\mathcal{B})$.

Starting from the (possibly non Wheeler) minimum DFA of a Wheeler language $\mathcal{L}$, we will give an algorithm constructing the minimum Wheeler automaton for the language. This automaton can be described as follows (see [2]): $\mathcal{B}=\left(Q^{\prime}, \delta^{\prime}, q_{0}^{\prime}, F^{\prime}\right)$ where

- $Q^{\prime}=\left\{[\alpha]_{\equiv_{\mathcal{L}}^{c}}: \alpha \in \operatorname{Pref}(\mathcal{L})\right\} ;$
- $q_{0}=[\epsilon]_{\equiv_{\mathcal{L}}^{c}}$;
- $\delta^{\prime}\left([\alpha]_{\equiv_{\mathcal{L}}^{c}}, a\right)=[\alpha a]_{\equiv_{\mathcal{L}}^{c}}$, for all $\alpha \in \operatorname{Pref}(\mathcal{L}), a \in \Sigma$,
- $F^{\prime}=\left\{[\alpha]_{\equiv_{\mathcal{L}}^{c}}: \alpha \in \mathcal{L}\right\}$.


## 2 Reduced NFA's meets Wheelerness

### 2.1 Automata

Among the two possible ways of presenting regular languages by automata, that is DFA's or NFA's, in general, computational problems tend to be significantly harder when referred to the non-deterministic class. Typical examples are: checking emptiness, computing the intersection, checking universality and much more. In the realm of Wheeler automata and languages a new class emerges: the class of reduced automata, formally defined below.

Definition 5. An NFA $\mathcal{A}=(Q, S, \delta, F, \Sigma)$ is called reduced if $q \neq p$ implies $I_{q} \neq I_{p}$.

Clearly, the class of reduced NFA's contains properly the class of DFA's. When Wheelerness is concerned, the class of reduced NFA's is interesting because it has been proved that deciding whether an NFA is Wheeler is an NP-complete problem [7], whereas deciding whether a reduced NFA is Wheeler turns out to be in P [2] as it is for DFA's [1]. Clearly, any NFA can be turned into a reduced one simply by merging all the states that recognize the same incoming language. Finding states to be merged is complex: the language-equivalence problem for NFA's can easily be proved as complex as deciding whether two states of an NFA recognize the same incoming language and, therefore, the latter is PSPACEcomplete.

A natural question is now whether switching from NFA's to reduced NFA's simplifies some otherwise difficult problem. In this section we prove that this is not always the case: some problems remain hard even when restricted to the class of reduced NFA's.

Lemma 1. The universality problem for reduced NFA's is PSPACE-complete.
Proof. This problem belongs to PSPACE, since it is a restriction of the universality problem over generic NFA's. To prove the completeness, we show a reduction from the universality problem.

Given an NFA, we can assume w.l.o.g. that there is only one initial state without incoming edges, hence let $\mathcal{A}=\left(Q, q_{0}, \delta, F, \Sigma\right)$ be an NFA with $Q=$ $\left\{q_{0}, \ldots, q_{n}\right\}$ be such an NFA. We build a new automaton $\mathcal{A}^{\prime}=\left(Q \cup P, q_{0}, \delta^{\prime}, F, \Sigma \cup\right.$ $\{d\}$ ), where $P=\left\{p_{1}, \ldots, p_{n-1}\right\}$ is a set of $n-1$ new states and $d$ is a new character. For each $q \in Q$ we add the self loop $(q, d, q)$. If we add only these transitions, it holds that $\mathcal{L}(\mathcal{A})=\Sigma^{*}$ iff $\mathcal{L}\left(\mathcal{A}^{\prime}\right)=(\Sigma+d)^{*}$. We can now add to the automaton as many $d$-transitions as we please without violating the property $\mathcal{L}(\mathcal{A})=\Sigma^{*}$ iff $\mathcal{L}\left(\mathcal{A}^{\prime}\right)=(\Sigma+d)^{*}$ : the right-to-left implication still holds if we only add $d$ transitions, whereas the left-to-right implication holds since adding transitions may only expand the recognized language, but $(\Sigma+d)^{*}$ is already maximal (with respect to the inclusion). Therefore we add the transitions ( $q_{0}, d, q_{1}$ ) and $\left(q_{0}, d, p_{1}\right)$. Moreover, for each $1 \leq i \leq n-1$ we add the transitions $\left(p_{i}, d, q_{i+1}\right)$ and ( $p_{i}, d, p_{i+1}$ ) (see Figure 2).


Fig. 2. The automaton $\mathcal{A}^{\prime}$ with only $d$-transitions depicted.

To conclude the proof that the reduction is correct, we need to show that $\mathcal{A}^{\prime}$ is reduced. Since $q_{0}$ has no incoming edges, we have

$$
\begin{aligned}
I_{q_{0}} & =d^{*} \\
I_{p_{i}} & =d^{i} \cdot d^{*} \quad \text { for } 1 \leq i \leq n-1
\end{aligned}
$$

Since $\mathcal{A}$ was trimmed and since each $q \in Q \backslash\left\{q_{0}\right\}$ is not an initial state, we have $I_{q} \cap \Sigma^{+} \neq \emptyset$ for each $q \in Q \backslash\left\{q_{0}\right\}$. Thus $I_{q} \neq I_{p}$ for each $q \in Q \backslash\left\{q_{0}\right\}$ and for each $p \in P \cup\left\{q_{0}\right\}$. Moreover, for each $1 \leq i<j \leq n$ we have $d^{i} \in I_{q_{i}} \backslash I_{q_{j}}$, hence $I_{q_{i}} \neq I_{q_{j}}$.

We will use the previous lemma to solve a problem related to another interesting aspect of the relationship between DFA's, NFA's, and reduced NFA's: indexability. Given an NFA $\mathcal{A}$, it is possible to define a partial order $<_{\mathcal{A}}$ on its states that allows to represent $\mathcal{A}$ using an index, that is, a succint structure that supports fast matching queries [4]. The partial order $<_{\mathcal{A}}$ is defined using the family of incoming languages $\left\{I_{q}: q \in Q\right\}$. As opposed to the case of DFA's, over NFA's these languages may not be pairwise disjoint, and we can compare them as follows:

$$
I_{q} \preceq I_{p} \Longleftrightarrow \forall \alpha \in I_{q} \forall \beta \in I_{p}\left(\{\alpha, \beta\} \nsubseteq I_{q} \cap I_{p} \Rightarrow \alpha \prec \beta\right)
$$

The above partial order can be lifted to the collection of states of an NFA.
Definition 6. Given two states $q$ and $p$ of an $N F A \mathcal{A}$, we say that $q<_{\mathcal{A}} p$ iff $I_{q} \preceq I_{p}$ and $I_{q} \neq I_{p}$.

Note that if $\mathcal{D}$ is a DFA, then $<_{\mathcal{D}}$ simplifies:

$$
q<_{\mathcal{D}} p \Longleftrightarrow \forall \alpha \in I_{q} \forall \beta \in I_{p}(\alpha \prec \beta),
$$

and this order satisfies the properties of a Wheeler order, with the exception of not necessarily being total. As a matter of fact, it can be proved that the DFA $\mathcal{D}$ is Wheeler if and only if $<_{\mathcal{D}}$ is a total order. Remarkably, this partial order can be computed in polynomial time [4] on DFA's.

Proposition 1. Let $\mathcal{D}$ be a DFA with n states. Then, we can compute the order $<_{\mathcal{D}}$ in $O\left(n^{5}\right)$ time.

It follows that, given a DFA $\mathcal{D}$, we can compute $<_{\mathcal{D}}$ in polynomial time and use it to index $\mathcal{D}$ efficiently. Would it be possible to generalized this result to NFA's using the corresponding partial order $<_{\mathcal{A}}$ of Definition 6? In the following lemma we give a negative answer to this question, even when restricted to reduced automata, proving that a different approach is needed to index NFA's (see [4] for a positive solution to the problem).

Theorem 2. Given two states $q$ and $p$ of an NFA $\mathcal{A}$, deciding whether $q<{ }_{\mathcal{A}} p$ is PSPACE-complete. The same result holds even if $\mathcal{A}$ is reduced.

Proof. First of all we need to prove that the problem is in PSPACE. We will show instead that its complement is in PSPACE and the thesis follows from the fact that PSPACE is closed under complementation. The complement of our problem consist of answering to the question whether $q \nless p$. To do so, first we check whether $I_{q}=I_{p}$. As we have already mentioned, this problem is in PSPACE, so we can get the answer in polynomial space. If $I_{q}=I_{p}$, then $q \nless p$ and we answer "yes". Otherwise, we have

$$
q<_{\mathcal{A}} p \Longleftrightarrow \forall \alpha \in I_{q} \forall \beta \in I_{p}\left(\{\alpha, \beta\} \nsubseteq I_{q} \cap I_{p} \Rightarrow \alpha \prec \beta\right)
$$

or equivalently

$$
q \nless_{\mathcal{A}} p \Longleftrightarrow \exists \alpha \in I_{q} \exists \beta \in I_{p}\left(\{\alpha, \beta\} \nsubseteq I_{q} \cap I_{p} \wedge \beta \prec \alpha\right) .
$$

Let $d$ be the number of states of the DFA $\mathcal{D}$ generated by the determinization of $\mathcal{A}$; clearly it holds $d \leq 2^{n}$. We claim that if $q \nless p$, then there exist two strings $\alpha, \beta$ of length at most $d^{2}+d$ such that

$$
\begin{equation*}
\alpha \in I_{q} \wedge \beta \in I_{p} \wedge\{\alpha, \beta\} \nsubseteq I_{q} \cap I_{p} \wedge \beta \prec \alpha \tag{1}
\end{equation*}
$$

Assume that $\alpha, \beta$ satisfy (1), with either $|\alpha|$ or $|\beta|$ (possibly both) greater than $d^{2}+d$. We assume, w.l.o.g., that $|\alpha| \leq|\beta|$ and distinguish two cases.

1) The last $d^{2}$ characters of $\alpha$ and $\beta$ differs; this also includes the case where $|\alpha|$ is strictly less than $d^{2}$. Consider the $d+1$ states of $\mathcal{D}$ visited by reading the first $d$ characters of $\beta$. Since $\mathcal{D}$ has $d$ states, at least one of them appears twice, implying that we visited a cycle. By erasing from the first $d$ characters of $\beta$ the factor corresponding to such cycle, we obtain a string $\beta^{\prime}$ such that $\alpha$ and $\beta^{\prime}$ also satisfy (1).
2) The last $d^{2}$ characters of $\alpha$ and $\beta$ coincide; in particular $|\alpha|,|\beta| \geq d^{2}$. Consider the last $d^{2}+1$ states $r_{0}, \ldots, r_{d^{2}}$ of $\mathcal{D}$ visited by reading the string $\alpha$, and the last $d^{2}+1$ states $p_{0}, \ldots, p_{d^{2}}$ visited by reading the string $\beta$. Since $\mathcal{D}$ has only $d$ states, there must exist $0 \leq i, j \leq d^{2}$ with $i<j$ such that $\left(r_{i}, p_{i}\right)=\left(r_{j}, p_{j}\right)$, implying that $\alpha$ and $\beta$ visited two cycles labeled by the same string. By erasing from the last $d^{2}$ characters of $\alpha$ and $\beta$ the factor corresponding to such cycles, we obtain two strings $\alpha^{\prime}, \beta^{\prime}$ which also satisfy (11).
In both cases, we were able to shorten the length of the longest string. By repeating this process as many times as needed, we will eventually obtain two strings both shorter than $d^{2}+d$, with $d \leq 2^{n}$.

Now that we have bounded the length of $\alpha, \beta$ with the constant $2^{2 n}+2^{n}$, we can use non-determinism to guess, bit by bit, the length of $\alpha$ and $\beta$ and store this guessed information in two counters $a, b$ respectively, using $O\left(\log \left(2^{2 n}+2^{n}\right)\right)=$ $O(n)$ space for each. These counters determine which string among $\alpha$ and $\beta$ is longer and we start guessing the characters of such longest string from the left to the right, decreasing by one its counter whenever we guess a character. Note that we are not storing the guessed characters, since it would use too much space. When the counter reaches the same value of the other counter, we start guessing the characters of both the first and the second string at the same time and we carry on until both counters reach the value 0 . While guessing the characters of $\alpha$ (respectively, $\beta$ ) we update at each step the set of states of $\mathcal{A}$ reachable from $q_{0}$ by reading the currently guessed prefix of $\alpha(\beta)$, so that in the end we obtain the sets $\delta\left(q_{0}, \alpha\right)$ and $\delta\left(q_{0}, \beta\right)$. With this information, we can check whether $\alpha \in I_{q}$ and $\beta \in I_{p}$ and $\{\alpha, \beta\} \nsubseteq I_{q} \cap I_{p}$. To complete checking condition (1), we need to show how to decide whether $\beta \prec \alpha$.

To confront co-lexicographically $\alpha$ and $\beta$, we use a variable $\rho$ that indicates whether $\alpha$ is less, equal or greater than $\beta$. We initialize $\rho$ based on the counters
$a, b$ as follows:

$$
\rho:= \begin{cases}= & \text { if } a=b \\ \dashv & \text { if } a<b \\ \vdash & \text { if } b<a\end{cases}
$$

We leave $\rho$ unchanged until we start guessing simultaneously the characters of $\alpha$ and $\beta$. When we guess the character $c_{1}$ for $\alpha$ and the character $c_{2}$ for $\beta$, we set

$$
\rho:= \begin{cases}\prec & \text { if } c_{1} \prec c_{2} \\ \succ & \text { if } c_{1} \succ c_{2} \\ \rho & \text { if } c_{1}=c_{2}\end{cases}
$$

Note that if at the end $\rho$ has value $\dashv$, it means that $\alpha \dashv \beta$, thus $\alpha \prec \beta$. Similarly, if $\rho$ has value $\vdash$ then $\beta \prec \alpha$. Otherwise, we have $\alpha \rho \beta$. Thus we are always able to determine the co-lexicographic order of $\alpha$ and $\beta$. Therefore, deciding whether $q \nless p$ is a problem in PSPACE, and so it is its complement.

To prove completeness, we show a reduction from the universality problem over generic, respectively reduced, NFA's.


Fig. 3. The automaton $\mathcal{A}^{\prime}$ built starting from the automaton $\mathcal{A}$ with $S=\left\{s_{1}, s_{2}\right\}$ recognizing the language $\mathcal{L}=\left\{\varepsilon, a_{2}, a_{1} a_{3}\right\}$. Edges entering a final state in $\mathcal{A}$ have been duplicated and redirected to $q_{e}$. Green edges are labeled $\Sigma=\left\{a_{1}, a_{2}, a_{3}\right\}$.

Let $\mathcal{A}=(Q, S, \delta, F, \Sigma)$ be an NFA with $Q=\left\{q_{1}, \ldots, q_{n}\right\}$ and $\Sigma=\left\{a_{1}, \ldots, a_{\sigma}\right\}$ recognizing the language $\mathcal{L}=\mathcal{L}(\mathcal{A})$, we build a new NFA $\mathcal{A}^{\prime}=\left(Q^{\prime}, q_{0}, \delta^{\prime}, F \cup\right.$ $\left.\left\{q_{e}, q_{f}\right\}, \Sigma^{\prime}\right)$ by adding a new initial state $q_{0}$ and two final states $\left\{q_{e}, q_{f}\right\}$ (see Figure 31). The new alphabet is $\Sigma^{\prime}=\Sigma \cup\{y, z\}$, where $a_{j} \prec y \prec z$ for each $1 \leq j \leq \sigma$. For each $q_{i} \in S$, we add a transition from $q_{0}$ to $q_{i}$ labeled $a_{1}$. Adding $q_{0}$ has the sole purpose of having an initial state without incoming edges. Note that we can not make the usual assumption that $\mathcal{A}$ has only one initial states without incoming edges: if we start from a reduced NFA and we build an equivalent NFA with the required property, there is no guarantee that the new automaton will still be reduced. The state $q_{e}$ represents the new final state that gathers all the strings in $a_{1} \cdot(\mathcal{L} \backslash\{\varepsilon\})$. To achieve this goal, for each transition $\left(q_{i}, a_{j}, q_{i^{\prime}}\right)$ of $\delta$ such that $q_{i^{\prime}} \in F$ we add a new transition $\left(q_{i}, a_{j}, q_{e}\right)$. The state $q_{f}$ gathers all the strings in $a_{1} \cdot \operatorname{Pref}(\mathcal{L}) \cdot \Sigma$, and this can be easily achieved by adding a transition $\left(q_{i}, a_{j}, q_{f}\right)$ for each $i \geq 1$ and $j \geq 1$. Lastly, we add the transitions $\left(q_{0}, y, q_{e}\right)$, $\left(q_{0}, y, q_{f}\right)$ and $\left(q_{0}, z, q_{f}\right)$. This way, if $\mathcal{A}$ is reduced then $\mathcal{A}^{\prime}$ is also reduced: note that $I_{q_{0}}=\{\varepsilon\}$, for each $i \geq 1$ it holds $I_{q_{i}}^{\mathcal{A}^{\prime}}=a_{1} \cdot I_{q_{i}}^{\mathcal{A}}$, the states $q_{e}, q_{f}$ are the only that can read the string $y$ and $q_{f}$ is the only state that can read the string $z$.

Let $\mathcal{L}_{\varepsilon}$ denote the language $\mathcal{L} \backslash\{\varepsilon\}$. By construction, we have

$$
\begin{aligned}
& I_{q_{e}}=a_{1} \cdot \mathcal{L}_{\varepsilon}+y \\
& I_{q_{f}}=a_{1} \cdot \operatorname{Pref}(\mathcal{L}) \cdot \Sigma+y+z
\end{aligned}
$$

We want to show that $\mathcal{L}=\Sigma^{*}$ iff $q_{e}<q_{f} \wedge \Sigma \subseteq \mathcal{L}_{\varepsilon}$. Note that $\Sigma \subseteq \mathcal{L}_{\varepsilon}$ is a necessary condition for $\mathcal{L}$ to be universal, and such condition can be checked in polynomial time using reachability on $\mathcal{A}$, therefore the reduction is still polynomial.
$(\Rightarrow)$ If $\mathcal{L}=\Sigma^{*}$, it clearly follows that $\Sigma \subseteq \mathcal{L}_{\varepsilon}$. Moreover we have $\operatorname{Pref}(\mathcal{L}) \cdot \Sigma=$ $\Sigma^{+}$and we obtain

$$
\begin{aligned}
& I_{q_{e}}=a_{1} \cdot \Sigma^{+}+y \\
& I_{q_{f}}=a_{1} \cdot \Sigma^{+}+y+z
\end{aligned}
$$

It follows immediately that $q_{e}<q_{f}$.
$(\Leftarrow)$ Note that $\mathcal{L}_{\varepsilon} \subseteq \operatorname{Pref}(\mathcal{L}) \cdot \Sigma$. We first prove that from the hypothesis it follows $\mathcal{L}_{\varepsilon}=\operatorname{Pref}(\mathcal{L}) \cdot \Sigma$. Assume by contradiction that $\mathcal{L}_{\varepsilon} \neq \operatorname{Pref}(\mathcal{L}) \cdot \Sigma$ and let $\beta$ be a string in $\operatorname{Pref}(\mathcal{L}) \cdot \Sigma \backslash \mathcal{L}_{\varepsilon}$. Then we have

$$
y \in I_{q_{e}}, \quad a_{1} \cdot \beta \in I_{q_{f}}, \quad\left\{y, a_{1} \cdot \beta\right\} \nsubseteq I_{q_{e}} \cap I_{q_{f}}
$$

but $y \succ a_{1} \cdot \beta$, a contradiction. Thus $\mathcal{L}_{\varepsilon}=\operatorname{Pref}(\mathcal{L}) \cdot \Sigma$.
We can then prove by induction on $|\alpha|$ that $\alpha \in \Sigma^{+}$implies $\alpha \in \mathcal{L}_{\varepsilon}$. If $|\alpha|=1$ then $\alpha \in \Sigma$ and by hypothesis we have $\Sigma \subseteq \mathcal{L}_{\varepsilon}$. If $|\alpha|=n+1>1$, then $\alpha=\alpha^{\prime} \cdot a_{j}$ for some $\alpha^{\prime} \in \Sigma^{+}$and some $a_{j} \in \Sigma$. By induction hypothesis we have $\alpha^{\prime} \in \mathcal{L}_{\varepsilon} \subseteq \operatorname{Pref}(\mathcal{L})$, and from $\mathcal{L}_{\varepsilon}=\operatorname{Pref}(\mathcal{L}) \cdot \Sigma$ it follows $\alpha \in \mathcal{L}_{\varepsilon}$.

This concludes the reduction from the universality problem to our problem over general NFA's. Since the construction described preserves the reduced-ness
of the starting automaton, it also works as a reduction from the universality problem over reduced NFA's to our problem over reduced NFA's. In Lemma 1 we proved that the former problem is PSPACE-complete, thus proving that the latter is also PSPACE- complete.

We can use the previous results to prove another complexity result over reduced NFA's.

Corollary 1. Deciding whether an NFA $\mathcal{A}$ is reduced is PSPACE-complete.
Proof. To prove that the problem is in PSPACE, note that $\mathcal{A}=\left(Q, q_{0}, \delta, F, \Sigma\right)$ is reduced iff, for all $q, p \in Q, q \neq p$ implies $I_{q} \neq I_{p}$. Therefore, it is sufficient to check $O\left(n^{2}\right)$ times whether $I_{q}=I_{p}$, where $n=|Q|$. As we have already mentioned, the problem of deciding whether $I_{q}=I_{p}$ belongs to PSPACE, thus the thesis follows.

To prove completeness, we combine the reductions shown in Lemma 1 and Theorem 2. Let $\Sigma_{d}=\Sigma \cup\{d\}$. We first apply the reduction shown in Lemma 1 to build a reduced automaton $\mathcal{A}^{\prime}$ such that $\mathcal{L}(\mathcal{A})=\Sigma^{*}$ iff $\mathcal{L}\left(\mathcal{A}^{\prime}\right)=\Sigma_{d}^{*}$. We set $\mathcal{L}^{\prime}:=\mathcal{L}\left(\mathcal{A}^{\prime}\right)$. Then, we apply the reduction showed in Theorem 2 to the automaton $\mathcal{A}^{\prime}$, but we remove the edge $\left(q_{0}, z, q_{f}\right)$; we call this new automaton $\mathcal{A}^{\prime \prime}$. The languages recognized by $q_{e}$ and $q_{f}$ change as follow:

$$
\begin{aligned}
& I_{q_{e}}=a_{1} \cdot \mathcal{L}_{\varepsilon}^{\prime}+y \\
& I_{q_{f}}=a_{1} \cdot \operatorname{Pref}\left(\mathcal{L}^{\prime}\right) \cdot \Sigma+y
\end{aligned}
$$

Since $\mathcal{A}^{\prime}$ is a reduced automaton and the states $q_{e}$ and $q_{f}$ are the only ones with an incoming edge labeled $y$, it immediately follows that $\mathcal{A}^{\prime \prime}$ is not reduced iff $I_{q_{e}}=I_{q_{f}}$. Applying the same argument we used in Theorem 2, we can conclude that $\mathcal{L}(\mathcal{A})^{\prime}=\Sigma_{d}^{*}$ iff $I_{q_{e}}=I_{q_{f}}$-again, we assumed that $\Sigma_{d} \subseteq \mathcal{L}\left(\mathcal{A}^{\prime}\right)$, since this condition can be checked in polynomial time. Summarizing we have that $\mathcal{A}^{\prime \prime}$ is not reduced iff $\mathcal{L}(\mathcal{A})=\Sigma^{*}$. Our claim follows from the equality PSPACE=NPSPACE.

Note that, as proved in [2], deciding whether a Wheeler NFA is reduced is a simpler problem, being in P .

### 2.2 Languages

In this section we switch our focus from automata to languages. An important consequence of the Myhill-Nerode Theorem for Wheeler languages is stated in the following Lemma (proved in [2]).

Lemma 2. A regular language $\mathcal{L}$ is Wheeler if and only if all monotone sequences in $(\operatorname{Pref}(\mathcal{L}), \prec)$ become eventually constant modulo $\equiv_{\mathcal{L}}$. In other words, for all sequences $\left(\alpha_{i}\right)_{i \geq 0}$ in $\operatorname{Pref}(\mathcal{L})$ with

$$
\alpha_{1} \preceq \alpha_{2} \preceq \ldots \alpha_{i} \preceq \ldots \quad \text { or } \quad \alpha_{1} \succeq \alpha_{2} \succeq \cdots \succeq \alpha_{i} \succeq \ldots
$$

there exists an $n$ such that $\alpha_{h} \equiv_{\mathcal{L}} \alpha_{k}$, for all $h, k \geq n$.

Lemma 2 shows how it is possible to recognize whether a language $\mathcal{L}$ is Wheeler simply by verifying a property on elements of $\operatorname{Pref}(\mathcal{L})$ : trying to find a WDFA that recognizes $\mathcal{L}$ is no longer needed to decide Wheelerness of $\mathcal{L}$. As shown in Theorem 3 (see [2]), we can verify whether the property mentioned in Lemma 2 is satisfied just analysing the structure of the minimum DFA recognizing $\mathcal{L}$.

Theorem 3. Let $\mathcal{D}_{\mathcal{L}}$ be the minimum DFA that recognizes the language $\mathcal{L}$, with initial state $q_{0}$ and dimension $n=\left|\mathcal{D}_{\mathcal{L}}\right|$.
$\mathcal{L}$ is not Wheeler if and only if there exist $\mu, \nu$ and $\gamma$ in $\Sigma^{*}$, with $\gamma \notin \mu, \nu$, such that:

1. $\mu \not 三_{\mathcal{L}} \nu$ and they label paths from $q_{0}$ to states $u$ and $v$, respectively;
2. $\gamma$ labels two cycles, one starting from $u$ and one starting from $v$;
3. $\mu, \nu \prec \gamma$ or $\gamma \prec \mu, \nu$.

The length of the strings $\mu, \nu$ and $\gamma$ satisfying the above can be bounded:
4. $|\mu|,|\nu| \leq|\gamma| \leq n^{3}+2 n^{2}+n+2$.

The proof of Theorem 3] [2] can be adapted to work on generic DFA's. Since such proof is both long and technical, we will prove instead (in the Appendix) the following proposition, where we worsen the bound given in condition 4. This is not a problem, since we will only use the fact that this bound is polynomial in $n$.

Proposition 2. Let $\mathcal{D}=\left(Q, q_{0}, \delta, F, \Sigma\right)$ be a $D F A$ recognizing the language $\mathcal{L}$, with $n=|\mathcal{D}|$.
$\mathcal{L}$ is not Wheeler if and only if there exist $\mu, \nu$ and $\gamma$ in $\Sigma^{*}$, with $\gamma \notin \mu, \nu$, such that:

1. $\mu \not 三_{\mathcal{L}} \nu$ and they label paths from $q_{0}$ to states $u$ and $v$, respectively;
2. $\gamma$ labels two cycles, one starting from $u$ and one starting from $v$;
3. $\mu, \nu \prec \gamma$ or $\gamma \prec \mu, \nu$.

The length of the strings $\mu, \nu$ and $\gamma$ satisfying the above can be bounded:
4. $|\mu|,|\nu| \leq|\gamma| \leq\left(n^{3}+2 n^{2}+n+2\right) \cdot n^{2}$.

The polynomial bound given by condition 4 of Theorem 3 allows us to design an algorithm that decides whether a given DFA recognizes a Wheeler language: using dynamic programming (see [1]) it is possible to keep track of all the relevant paths and cycles inside the DFA and check, in polynomial time, whether there exists three strings satisfying the conditions of the theorem.

Things change if, instead of a DFA, we are given an NFA. Trying to exploit the same idea used for DFA's does not work: the problem of deciding whether two strings $\mu$ and $\nu$ read by an NFA are Myhill-Nerode equivalent is PSPACEcomplete. Even worse, a straightforward attempt of building the minimum DFA recognizing the NFA's language might lead to a blow-up of the sates, resulting in a exponential time (and exponential space) algorithm.

We show that the problem of deciding whether an NFA recognizes a Wheeler language is indeed hard, but does not necessarily require exponential time to be solved. Instead, the problem turns out to be PSPACE-complete. To show this, we first show how to adapt Theorem 3 to work on NFA's, as described in the following corollary.

Corollary 2. Let $\mathcal{A}=\left(Q, q_{0}, \delta, F, \Sigma\right)$ be an NFA of dimension $n:=|\mathcal{A}|$. Then $\mathcal{L}:=\mathcal{L}(\mathcal{A})$ is not Wheeler if and only if there exist three strings $\mu, \nu, \gamma$ such that $\gamma \notin \mu, \nu$ and

1. $\mu \gamma^{i} \not \equiv \mathcal{L} \nu \gamma^{j}$ for all $0 \leq i, j \leq 2^{n}$;
2. $\gamma$ labels two cycles, one starting from a state $p \in \delta\left(q_{0}, \mu\right)$ and one from a state $r \in \delta\left(q_{0}, \nu\right)$;
3. $\mu, \nu \prec \gamma$ or $\gamma \prec \mu, \nu$.

Moreover, the length of the strings $\mu, \nu$ and $\gamma$ satisfying the above can be bounded:
4. $|\mu|,|\nu|<|\gamma|<n^{3} \cdot\left(2^{3 n}+2 \cdot 2^{2 n}+2^{n}+2\right) \in O\left(2^{3 n}\right)$.

Proof. Let $\mathcal{D}=\left(\hat{Q}, \hat{q}_{0}, \hat{\delta}, \hat{F}, \Sigma\right)$ be the minimum DFA recognizing $\mathcal{L}$. Clearly $\mathcal{D}$ has at most $2^{n}$ states.
$(\Longleftarrow)$ From condition 2 it follows that $\mu \gamma^{*} \subseteq \operatorname{Pref}(\mathcal{L})$, so consider the following list of $2^{n}+1$ states of $\mathcal{D}$ :

$$
\hat{\delta}\left(\hat{q}_{0}, \mu \gamma^{0}\right), \hat{\delta}\left(\hat{q}_{0}, \mu \gamma^{1}\right), \ldots, \hat{\delta}\left(\hat{q}_{0}, \mu \gamma^{2^{n}}\right)
$$

Since $\mathcal{D}$ has at most $2^{n}$ states, there must exist two integers $0 \leq h<k \leq 2^{n}$ such that $\hat{\delta}\left(\hat{q}_{0}, \mu \gamma^{h}\right)=\hat{\delta}\left(\hat{q}_{0}, \mu \gamma^{k}\right)$. Therefore $\gamma^{k-h}$ labels a cycle starting from $\hat{\delta}\left(\hat{q}_{0}, \mu \gamma^{h}\right)$. Similarly, there exist $0 \leq h^{\prime}<k^{\prime} \leq 2^{n}$ such that $\gamma^{k^{\prime}-h^{\prime}}$ labels a cycle starting from $\hat{\delta}\left(\hat{q}_{0}, \nu \gamma^{h^{\prime}}\right)$. The strings

$$
\begin{aligned}
& \hat{\mu}:=\mu \gamma^{h} \\
& \hat{\nu}:=\nu \gamma^{h^{\prime}} \\
& \hat{\gamma}:=\gamma^{\operatorname{lcm}\left(k-h, k^{\prime}-h^{\prime}\right) \cdot 2^{n}},
\end{aligned}
$$

where the factor $2^{n}$ in the definition of $\hat{\gamma}$ ensures that $|\hat{\mu}|,|\hat{\nu}|<|\hat{\gamma}|$, so that $\hat{\gamma} \nrightarrow \hat{\mu}, \hat{\nu}$ and the strings $\hat{\mu}, \hat{\nu}, \hat{\gamma}$ satisfy condition 2 of Theorem 3. Condition 1 of Theorem 3 follows automatically from conditions 1 of this corollary. Lastly, condition 3 of Theorem 3 follows from conditions 3 of this corollary and the fact that $\gamma \not A \mu, \nu$. Thus we can apply Theorem 3 to conclude that $\mathcal{L}$ is not Wheeler. $(\Longrightarrow)$ Since $\mathcal{L}=\mathcal{L}(\mathcal{D})$ is not Wheeler, let $\hat{\mu}, \hat{\nu}, \hat{\gamma}$ be strings satisfying Theorem 3 The DFA $\mathcal{D}$ has at most $2^{n}$ states, hence the length of $\hat{\gamma}$ is bounded by the constant $2^{3 n}+2 \cdot 2^{2 n}+2^{n}+2$. We have $\hat{\mu} \hat{\gamma}^{*} \subseteq \operatorname{Pref}(\mathcal{L})$, so let $t_{0}=q_{0}, t_{1}, \ldots, t_{m}$ be a run of $\hat{\mu} \hat{\gamma}^{n}$ over $\mathcal{A}$. We set $u:=|\hat{\mu}|$ and $g:=|\hat{\gamma}|$, and consider the list of $n+1$ states

$$
t_{u}, t_{u+g}, t_{u+2 g}, \ldots, t_{u+n g}=t_{m}
$$

Since $\mathcal{A}$ has $n$ states, there must exist two integers $0 \leq h<k \leq n$ such that $t_{u+h g}=t_{u+k g}$. That is, there exists a state $p:=t_{u+h g}$ such that $p \in \delta\left(q_{0}, \hat{\mu} \hat{\gamma}^{h}\right)$ and $\hat{\gamma}^{k-h}$ labels a cycle starting from $p$. We can repeat the same argument for a run of $\hat{\nu} \hat{\gamma}^{n}$ over $\mathcal{A}$ to find a state $r$ and two integers $h^{\prime}, k^{\prime}$ such that $r \in \delta\left(q_{0}, \hat{\nu} \hat{\gamma}^{h^{\prime}}\right)$ and $\hat{\gamma}^{k^{\prime}-h^{\prime}}$ labels a cycle starting from $r$. We can then define the strings

$$
\begin{aligned}
\mu & :=\hat{\mu} \hat{\gamma}^{h} \\
\nu & :=\hat{\nu} \hat{\gamma}^{h^{\prime}} \\
\gamma & :=\hat{\gamma}^{\operatorname{lcm}\left(k-h, k^{\prime}-h^{\prime}\right) \cdot n}
\end{aligned}
$$

which satisfy the conditions 2 and 3 .
Condition 4 is satisfied since $|\hat{\gamma}| \leq 2^{3 n}+2 \cdot 2^{2 n}+2^{n}+2$ and $\operatorname{lcm}\left(k-h, k^{\prime}-h^{\prime}\right)<n^{2}$. Finally, condition 1 is satisfied for all $i, j \geq 0$. Indeed, for all $l$ the strings $\hat{\mu}$ and $\hat{\mu} \hat{\gamma}^{l}$ lead to the same state of $\mathcal{D}$, thus $\hat{\mu} \equiv_{\mathcal{L}} \hat{\mu} \hat{\gamma}^{l}$. Similarly, for all $l$ we also have $\hat{\nu} \equiv_{\mathcal{L}} \hat{\nu} \hat{\gamma}^{l}$. Since $\forall i \exists s_{i}$ such that $\mu \gamma^{i}=\hat{\mu} \hat{\gamma}^{s_{i}}$, and similarly, $\forall j \exists s_{j}$ such that $\nu \gamma^{j}=\hat{\nu} \hat{\gamma}^{s_{j}}$, the thesis follows from $\hat{\mu} \not \equiv \mathcal{L} \hat{\nu}$.

Despite the fact the the bound in condition 4 has become exponential by switching to NFA's, it is still possible to check in polynomial space (but exponential time) whether there are three strings $\mu, \nu$ and $\gamma$ satisfying the conditions of Proposition 2. Thus we can prove the following:

Theorem 4. Given an NFA $\mathcal{A}=\left(Q, q_{0}, \delta, F, \Sigma\right)$, deciding whether the language $\mathcal{L}:=\mathcal{L}(\mathcal{A})$ is Wheeler is PSPACE-complete. The same result holds even if $\mathcal{A}$ is reduced.

Proof. First of all we need to prove that the problem is in PSPACE. We will show instead that its complement is in NPSPACE, then the thesis follows from Savitch's Theorem, which states that NPSPACE $=$ PSPACE, and the fact that PSPACE is closed under complementation.

Let $\mathcal{D}$ be the automaton obtained by the determinization of $\mathcal{A}$ with dimension $d=|\mathcal{D}| \leq 2^{n}$. We prove that we can check the conditions in Proposition 2 for the automaton $\mathcal{D}$, without building it, using polynomial space. We use nondeterminism to guess, bit by bit, the length of $\mu, \nu$ and $\gamma$ and store this guessed information in three counters $u, v, g$ respectively, using $O\left(\log \left(d^{5}\right)\right)=O(n)$ space for each. These counters determine which string among $\mu, \nu$ and $\gamma$ is longer and we start guessing the characters of such string from the left to the right, decreasing by one its counter whenever we guess a character. When the counter reaches the same value of the second biggest counter, we start guessing the characters of both the first and the second string at the same time and we carry on until they reach the value of the last counter. Then, we guess simultaneously the characters of all three strings until all counters reach the value 0 . While guessing the characters of $\mu$ (respectively, $\nu$ ) we update at each step the set of states of $\mathcal{A}$ reachable from $q_{0}$ by reading the currently guessed prefix of $\mu(\nu)$, so that in the end we obtain the sets $\delta\left(q_{0}, \mu\right)$ and $\delta\left(q_{0}, \nu\right)$. We proceed similarly for $\gamma$, but this time we compute the set $\delta(q, \gamma)$ for each state $q \in Q$. Since $\mathcal{D}$
is the determinized version of $\mathcal{A}$, we can verify condition 2 of Proposition 2 by checking whether the set $\delta\left(q_{0}, \mu\right)$ and the set

$$
\delta\left(q_{0}, \mu \cdot \gamma\right)=\bigcup_{p \in \delta\left(q_{0}, \mu\right)} \delta(p, \gamma)
$$

are equal, and we do the same for $\delta\left(q_{0}, \nu\right)$ and $\delta\left(q_{0}, \nu \cdot \gamma\right)$. Condition 3 of Proposition [2 can be checked in constant space. To confront $\mu$ and $\gamma$, we use a variable $\rho$ that indicates whether $\mu$ is less, equal or greater than $\gamma$. We initialize $\rho$ based on the counters $u, g$ as follows:

$$
\rho:= \begin{cases}= & \text { if } u=g \\ \vdash & \text { if } u<g\end{cases}
$$

We leave $\rho$ unchanged until we start guessing simultaneously $\mu$ and $\gamma$. Then, when we guess simultaneously the character $c_{1}$ for $\mu$ and the character $c_{2}$ for $\gamma$, we set

$$
\rho:= \begin{cases}\prec & \text { if } c_{1} \prec c_{2} \\ \succ & \text { if } c_{1} \succ c_{2} \\ \rho & \text { if } c_{1}=c_{2}\end{cases}
$$

Note that if at the end $\rho$ has value $\vdash$, it means that $\mu \vdash \gamma$. Otherwise, we have $\mu \rho \gamma$. Therefore, we are always able to determine the co-lexicographic order of $\mu$ and $\gamma$. To check condition 1 of Proposition 2, consider the automata $A_{\mu}$ and $A_{\nu}$ obtained from the NFA $\mathcal{A}$ by considering as initial states the sets $\delta\left(q_{0}, \mu\right)$ and $\delta\left(q_{0}, \nu\right)$, respectively. We have that $\mu \not \equiv \mathcal{L} \nu$ if and only if $\mathcal{L}\left(\mathcal{A}_{\mu}\right) \neq \mathcal{L}\left(\mathcal{A}_{\nu}\right)$, and checking whether $\mathcal{L}\left(\mathcal{A}_{\mu}\right)=\mathcal{L}\left(\mathcal{A}_{\nu}\right)$ can be done in polynomial space, since deciding whether two NFA's recognize the same language is a well-known PSPACEcomplete problem.

To prove the completeness of the problem, we will show a polynomial reduction from the universality problem for NFA, i.e. the problem of deciding whether the language accepted by an NFA $\mathcal{A}$, over the alphabet $\Sigma$, is $\Sigma^{*}$.

Let $\mathcal{A}=\left(Q, q_{0}, \delta, F, \Sigma\right)$ be an NFA and let $\mathcal{L}=\mathcal{L}(\mathcal{A})$. We can assume without loss of generality that $q_{0} \in F$, otherwise $\mathcal{A}$ would not accept the empty string and we could immediately derive that $\mathcal{L} \neq \Sigma^{*}$. Let $a, b, c$ be three characters not in $\Sigma$ and such that $a \prec b \prec c$ with respect to the lexicographical order (the order of the characters of $\Sigma$ is irrelevant in this proof). First, we build the automaton $\mathcal{A}^{\prime}$ starting from $\mathcal{A}$ by adding an edge $\left(q_{f}, q_{0}, c\right)$ for each final state $q_{f} \in F$, see the top part of Figure 4. Notice that $\mathcal{A}^{\prime}$ recognizes the language $\mathcal{L}^{\prime}=\mathcal{L}\left(\mathcal{A}^{\prime}\right)=(\mathcal{L} c)^{*} \cdot \mathcal{L}$, and it is straightforward to prove that $\mathcal{L}=\Sigma^{*}$ if and only if $\mathcal{L}^{\prime}=(\Sigma+c)^{*}$ : if $\mathcal{L}=\Sigma^{*}$, let $\alpha$ be a string in $(\Sigma+c)^{*}$ containing $n$ occurrences of $c$. Then $\alpha=\alpha_{0} c \alpha_{2} c \ldots \alpha_{n-1} c \alpha_{n}$ for some $\alpha_{1}, \ldots, \alpha_{n} \in \Sigma^{*}$. Hence $\alpha \in\left(\Sigma^{*} c\right)^{*} \cdot \Sigma^{*}=\mathcal{L}^{\prime}$. On the other hand, if $\mathcal{L} \neq \Sigma^{*}$ let $\alpha$ be a string in $\Sigma^{*} \backslash \mathcal{L}$. Then $\alpha \cdot c \notin \mathcal{L}^{\prime}$.

We build a second automaton $\mathcal{A}^{\prime \prime}$ as depicted in Figure 4 Let $\mathcal{L}^{\prime \prime}=\mathcal{L}\left(\mathcal{A}^{\prime \prime}\right)$ be the language recognized by $\mathcal{A}^{\prime \prime}$. We claim that $\mathcal{L}=\Sigma^{*}$ if and only if $\mathcal{L}^{\prime \prime}$ is


Fig. 4. The automaton $\mathcal{A}^{\prime \prime}$. Every accepting state of $\mathcal{A}$, labeled A in the figure, has a back edge labeled $c$ connecting it to $q_{0}$. Conversely, non-accepting states of $\mathcal{A}$, labeled N in the figure, do not have such back edges.

Wheeler.
$(\Longrightarrow)$ If $\mathcal{L}=\Sigma^{*}$, we have already proved that $\mathcal{L}^{\prime}=(\Sigma+c)^{*}$. Hence we have $\mathcal{L}^{\prime \prime}=(a+b) \cdot(\Sigma+c)^{*}$. The minimum DFA recognizing $\mathcal{L}^{\prime \prime}$ has only one loop, therefore by Theorem $3 \mathcal{L}^{\prime \prime}$ is Wheeler.
$(\Longleftarrow)$ If $\mathcal{L} \neq \Sigma^{*}$, let $\alpha$ be a string in $\Sigma^{*} \backslash \mathcal{L}$. Note that $\alpha \neq \varepsilon$ since we assumed that $\varepsilon \in \mathcal{L}$. Every possible run of $\alpha$ over $\mathcal{A}$ must lead to a non-accepting state, hence $\alpha \cdot c \notin \mathcal{L}^{\prime}$. This implies that for all $i \geq 0$ we have $a \cdot c^{i} \cdot \alpha \cdot c \notin \mathcal{L}^{\prime \prime}$ (notice that the only edge labeled $c$ leaving $q_{0}$ ends in $q_{0}$ ). On the other hand, for all $j \geq 0$ we have $b c^{j} \cdot \alpha \cdot c \in \mathcal{L}^{\prime \prime}$, hence for all $i, j \geq 0$ we have $a c^{i} \not \equiv \mathcal{L}^{\prime \prime} b c^{j}$. Thus the following monotone sequence in $\operatorname{Pref}\left(\mathcal{L}^{\prime \prime}\right)$

$$
a c \prec b c \prec a c c \prec b c c \prec \cdots \prec a c^{n} \prec b c^{n} \prec \ldots
$$

is not eventually constant modulo $\equiv \mathcal{L}^{\prime \prime}$. From Lemma 2 it follows that $\mathcal{L}^{\prime \prime}$ is not Wheeler.

Note that in the reduction described in Figure 4, if the starting NFA $\mathcal{A}$ was reduced, then also $\mathcal{A}^{\prime \prime}$ would be reduced. This means that the statement of the theorem holds even if restricted to reduced NFA's.

Remark 3. Note that the previous theorem is in contrast with what happens when we consider the problem of deciding whether an NFA is Wheeler, instead of whether it accepts a Wheeler language: in that case, restricting the problem to reduced NFA's makes it solvable in polynomial time.

## 3 State complexity

As already mentioned above, a significant property on the interplay between deterministic and non-deterministic Wheeler Automata is that given a size- $n$

WNFA $\mathcal{A}$, there always exists a WDFA that recognizes the same language whose size is at most $2 n$. The announced amount of states can be computed using the (classic) powerset construction. In other words, the blow-up of the number of states that we might observe when converting NFA's to DFA's, does not occur for Wheeler non-deterministic automata. This property is a direct consequence of an important feature of Wheeler automata: for any state $q$, the set of strings recognized by $q$-namely $I_{q}$-is an interval over $\operatorname{Pref}(\mathcal{L})$ with respect to the co-lexicographic order.

State complexity is also used to measure the complexity of operations on regular languages. In the next section we prove that the interval property of a Wheeler DFA can also be exploited to prove that the state complexity of the intersection of Wheeler languages is significantly better than the state complexity of the intersection of general regular languages.

### 3.1 Intersecting Wheeler languages

The state complexity of a regular language $\mathcal{L}$ is defined as the number of states of the minimum DFA $\mathcal{D}_{\mathcal{L}}$ recognizing $\mathcal{L}$. The state complexity of an operation on regular languages is a function that associates to the state complexities of the operand languages the worst-case state complexity of the language resulting from the operation. For instance, we say that the state complexity of the intersection of $\mathcal{L}_{1}$ and $\mathcal{L}_{2}$ is $m n$, where $m$ and $n$ are the number of states of $\mathcal{D}_{\mathcal{L}_{1}}$ and $\mathcal{D}_{\mathcal{L}_{2}}$ respectively. The bound $m n$ for the intersection can easily be proved using the state-product construction for $\mathcal{D}_{\mathcal{L}_{1}}$ and $\mathcal{D}_{\mathcal{L}_{2}}$, and it is a known fact that this bound is tight 9 .

It is natural to define the Wheeler state complexity of a Wheeler language $\mathcal{L}$ as the number of states of the minimum WDFA $\mathcal{D}_{\mathcal{L}}^{W}$ recognizing $\mathcal{L}$. In the following theorem, we show what it is the Wheeler state complexity of the intersection of two Wheeler languages $\mathcal{L}_{1}$ and $\mathcal{L}_{2}$.

Theorem 5. Let $\mathcal{D}_{\mathcal{L}_{1}}^{W}$ and $\mathcal{D}_{\mathcal{L}_{2}}^{W}$ be the minimum WDFA's recognizing the languages $\mathcal{L}_{1}$ and $\mathcal{L}_{2}$ respectively. Then, the minimum WDFA recognizing $\mathcal{L}:=$ $\mathcal{L}_{1} \cap \mathcal{L}_{2}$ has at most $\left|D_{\mathcal{L}_{1}}^{W}\right|+\left|D_{\mathcal{L}_{2}}^{W}\right|-|\Sigma|-1$ states.
This bound is tight.
Proof. First we prove that, given any two strings $\alpha, \beta \in \Sigma^{*}$, if $\alpha \equiv_{\mathcal{L}_{1}}^{c} \beta$ and $\alpha \equiv_{\mathcal{L}_{2}}^{c} \beta$ then $\alpha \equiv_{\mathcal{L}}^{c} \beta$. From $\alpha \equiv_{\mathcal{L}_{1}} \beta$ and $\alpha \equiv \mathcal{L}_{2} \beta$ it follows that $\alpha \equiv \equiv_{\mathcal{L}} \beta$. Moreover, from $\alpha \equiv{ }_{\mathcal{L}_{1}}^{c} \beta$ it follows that $\alpha$ and $\beta$ end with the same letter. What it is left to prove is that for any $\gamma \in \Sigma^{*}$ such that $\alpha \prec \gamma \prec \beta$ it holds $\alpha \equiv \mathcal{L} \gamma$. This follows immediately since $\alpha \equiv_{\mathcal{L}_{1}}^{c} \beta$ implies $\alpha \equiv_{\mathcal{L}_{1}} \gamma$ and $\alpha \equiv_{\mathcal{L}_{2}}^{c} \beta$ implies $\alpha \equiv \mathcal{L}_{1} \gamma$.

Let $C_{0}^{1}, \ldots C_{n-1}^{1}$ be the $\equiv_{\mathcal{L}_{1}}^{c}$-classes and let $C_{0}^{2}, \ldots C_{m-1}^{2}$ be the $\equiv_{\mathcal{L}_{2}}^{c}$-classes; we assume that both lists are ordered co-lexicographically. Since the $\equiv_{\mathcal{L}_{1}}^{c}$-classes are pairwise disjoint - and the same holds for the $\equiv_{\mathcal{L}_{2}}^{c}$-classes-the number of $\equiv_{\mathcal{L}_{1} \cap \mathcal{L}_{2}}^{c}$-classes is at most equal to the number of non-empty intersections of the form $C_{i}^{1} \cap C_{j}^{2}$, for $1 \leq i \leq n$ and $1 \leq j \leq m$. Classes that end with different
characters of the alphabet must have empty intersection; a particular case are the classes $C_{0}^{1}=C_{0}^{2}=\{\varepsilon\}$, which always lead to the non-empty intersection $C_{0}^{1} \cap C_{0}^{2}=\{\varepsilon\}$. We will focus on classes whose elements end with a specific character, say $a$. Let $C_{1}^{1 a}, \ldots, C_{n_{a}}^{1 a}$ be all the $\equiv_{\mathcal{L}_{1}}^{c}$-classes that end with $a$, colexicographically ordered, and let let $C_{1}^{2 a}, \ldots, C_{m_{a}}^{2 a}$ be all the $\equiv_{\mathcal{L}_{2}}^{c}$-classes that end with $a$. Let $k$ be the number of non-empty intersections of the form $C_{i}^{1 a} \cap C_{j}^{2 a}$, and let $\alpha_{1} \prec \cdots \prec \alpha_{k}$ be an ordered list containing one representatives for each non-empty intersection. For any $1 \leq s<k$, consider the strings $\alpha_{s}$ and $\alpha_{s+1}$. There must exist four unique indexes $i, j, i^{\prime}, j^{\prime}$ such that $\alpha_{s} \in C_{i}^{1 a} \cap C_{j}^{2 a}$ and $\alpha_{s+1} \in C_{i^{\prime}}^{1 a} \cap C_{j^{\prime}}^{2 a}$. From $\alpha_{s} \prec \alpha_{s+1}$ it follows that both $i \leq i^{\prime}$ and $j \leq j^{\prime}$ hold, since the $\equiv{ }_{\mathcal{L}_{1}}^{c}$-classes-and the $\equiv{ }_{\mathcal{L}_{2}}^{c}$-classes - are pairwise disjoint and colexicographically ordered. On the other hand, it can not be the case that both $i=i^{\prime}$ and $j=j^{\prime}$ hold, because $\alpha_{s}$ and $\alpha_{s+1}$ belong to different intersections. Therefore we have that $i^{\prime}+j^{\prime} \geq i+j+1$. The values of the function $f\left(\alpha_{s}\right)=i+j$ can range from 2 to $n_{a}+m_{a}$, hence there might be at most $n_{a}+m_{a}-1$ different representatives. Taking the sum over every possible characters of $\Sigma$ and adding the class $C_{0}^{1} \cap C_{j}^{2}=\{\varepsilon\}$, we get an upper bound of

$$
\begin{aligned}
1+\sum_{a \in \Sigma}\left(n_{a}+m_{a}-1\right) & =1+\sum_{a \in \Sigma} n_{a}+\sum_{a \in \Sigma} m_{a}-|\Sigma|= \\
& =1+(n-1)+(m-1)-|\Sigma|=n+m-|\Sigma|-1
\end{aligned}
$$

different possible representatives.

To show that the bound is tight (at least for $|\Sigma|=2$ ), consider the following families of languages over the alphabet $\Sigma=\{a, b\}$, with $a \prec b$ :

$$
\begin{aligned}
& A_{n}:=\left\{\alpha \in \Sigma^{*}: a^{n+1} \text { is not a factor of } \alpha\right\} \\
& B_{m}:=\left\{\beta \in \Sigma^{*}: b^{m+1} \text { is not a factor of } \beta\right\}
\end{aligned}
$$

We can easily prove that all these languages are Wheeler. The minimum DFA recognizing $B_{m}$ is already a WDFA, see Figure 5. with $m+2$ states. A list of representatives of such classes is

$$
\varepsilon, a, b, \ldots, b^{m}
$$



Fig. 5. The minimum WDFA recognizing $B_{3}$.


Fig. 6. The minimum WDFA recognizing $A_{3}$.

The minimum WDFA recognizing $A_{n}$ has more states than the minimum DFA: for $1 \leq i<n$ we have that $a^{i} \prec a^{n} \prec b a^{i}$, hence we have to split the $\equiv_{A_{n}}$-class containing both $a^{i}$ and $b a^{i}$ into two different $\equiv_{A_{n}}^{c}$-classes. The automaton has $2 n+1$ states, see Figure 6] A list of representatives of the $\equiv_{A_{n}}^{c}$-classes is

$$
\varepsilon, a, \ldots, a^{n}, b a^{n-1}, \ldots, b a, b
$$

We have already proved that the language $\mathcal{L}:=A_{n} \cap B_{m}$ might have at most $(2 n+1)+(m+2)-|\Sigma|-1=2 n+m$ different $\equiv_{\mathcal{L}^{-}}^{c}$-classes, hence it is sufficient to show that there are at least $2 n+m$ different ones. We claim that the $2 n+m$ strings

$$
\varepsilon, a, \ldots, a^{n}, b a^{n-1}, \ldots, b a, b, \ldots, b^{m}
$$

all belong to different $\equiv_{\mathcal{L}}^{c}$-classes. Strings that ends with a different amount of $a^{\prime}$ s (or $b^{\prime}$ s) belong to different $\equiv \mathcal{L}_{\mathcal{L}}$-classes, so there is nothing to prove. Therefore we only have to check, for each $1 \leq i<n$, that $a^{i}$ and $b a^{i}$ belong to different $\equiv_{\mathcal{L}^{c}}$-classes, and again this is true since $a^{i} \prec a^{n} \prec b a^{i}$.

Remark 4. Similarly to the case of determinizing a WNFA, where we can use the classic powerset construction without generating too many states, to compute a WDFA that recognizes the intersection of the languages accepted by two WDFA's $\mathcal{W}_{1}$ and $\mathcal{W}_{2}$ we can use the classic state-product construction with the certainty that it will not produce more states that necessary; that is, the number of states generated will be at most the sum of the number of states of $\mathcal{W}_{1}$ and $\mathcal{W}_{2}$.

Open Problem 1. Wheeler automata are closed under few operations: intersection and right-concatenation with a finite language, i.e. if $\mathcal{L}$ is a Wheeler language and $\mathcal{F}$ is a finite language, then also $\mathcal{L} \cdot \mathcal{F}$ is a Wheeler language. In general, the state complexity of the concatenation of $\mathcal{L}\left(\mathcal{D}_{1}\right) \cdot \mathcal{L}\left(\mathcal{D}_{2}\right)$ can result in an exponential blow-up in the number of states of $\mathcal{D}_{2}$ [9], even when restricted to finite languages [3]. It remains open the question whether it is possible to obtain a better-that is, sub-exponential-upper bound for Wheeler automata.

### 3.2 Computing the minimum WDFA

Despite of the good behaviour that Wheeler automata show regarding determinization and intersection, there are cases when the state complexity of a construction is exponential. In fact, it is known [1] that a blow-up of states can occur when switching from the minimum DFA recognizing a language $\mathcal{L}$ to its minimum WDFA. As a last contribution we provide an algorithm to compute the minimum WDFA starting from the minimum DFA $\mathcal{D}_{\mathcal{L}}$ of a Wheeler language $\mathcal{L}$, consisting in two steps: first, we describe an algorithm that extracts a fingerprint of $\mathcal{L}$ starting from $\mathcal{D}_{\mathcal{L}}$, that is, a set of string containing exactly one representative of each $\equiv_{\mathcal{L}^{-}}^{c}$-class of $\mathcal{L}$. Second, we provide an algorithm that builds the minimum WDFA recognizing $\mathcal{L}$ starting from any of its fingerprints.

Definition 7 (Fingeprint). Let $\mathcal{L}$ be a Wheeler language, and let $m$ be the number of equivalence classes of $\equiv_{\mathcal{L}}^{c}$. A set of strings $F=\left\{\alpha_{1}, \ldots, \alpha_{m}\right\} \subseteq \Sigma^{*}$ is called $a$ fingerprint of $\mathcal{L}$ if and only if for each $\equiv_{\mathcal{L}}^{c}$-class $C$ it holds $|F \cap C|=1$.

We start by proving that we can impose an upper bound to the length of the representative of a fingerprint.

Lemma 3. Let $\mathcal{D}_{\mathcal{L}}$ be the minimum DFA recognizing the Wheeler language $\mathcal{L}$ over the alphabet $\Sigma$, and let $C_{1}, \ldots, C_{m}$ be the pairwise distinct equivalence classes of $\equiv_{\mathcal{L}}^{c}$. Then, for each $1 \leq i \leq m$, there exists a string $\alpha_{i} \in C_{i}$ such that $\left|\alpha_{i}\right|<n+n^{2}$, where $n:=\left|\mathcal{D}_{\mathcal{L}}\right|$.

Proof. Suppose by contradiction that there exists a class $C_{i}$ such that for all $\alpha \in C_{i}$ it holds $|\alpha| \geq n+n^{2}$, and let $\alpha \in C_{i}$ be a string of minimum length. Consider the first $n+1$ states $q_{0}=t_{0}, \ldots, t_{n}$ of $\mathcal{D}_{\mathcal{L}}$ visited by reading the first $n$ characters of $\alpha$. Since $\mathcal{D}_{\mathcal{L}}$ has only $n$ states, there must exist $0 \leq i, j \leq n$ with $i<j$ such that $t_{i}=t_{j}$. Let $\alpha^{\prime}$ be the prefix of $\alpha$ of length $i$ (if $i=0$ then $\alpha^{\prime}=\varepsilon$ ), let $\delta$ be the factor of $\alpha$ of length $j-i$ labeling the path $t_{i}, \ldots, t_{j}$, and let $\zeta$ be the suffix of $\alpha$ such that $\alpha=\alpha^{\prime} \delta \zeta$. By construction, the strings $\alpha$ and $\beta:=\alpha^{\prime} \zeta$ end in the same state, hence $\alpha \equiv_{\mathcal{L}} \beta$. Moreover, from $|\beta|<|\alpha|$ and the minimality of $\alpha$ it follows that $\alpha \not \equiv_{\mathcal{L}}^{c} \beta$.
Suppose that $\alpha \prec \beta$, the other case being completely symmetrical. Since $\alpha$ and $\beta$ share the same suffix $\zeta$, they end with the same character. This means that the strings $\alpha$ and $\beta$, which are Myhill-Nerode equivalent but not $\equiv_{\mathcal{L}}^{c}$ equivalent, were not split into two distinct $\equiv_{\mathcal{L}}^{c}$-classes due to input-consistency, therefore there must exists a string $\eta$ such that $\alpha \prec \eta \prec \beta$ and $\eta \not \equiv \mathcal{L}_{\mathcal{L}} \alpha$. Formally, assume by contradiction that for all strings $\eta$ such that $\alpha \prec \eta \prec \beta$ it holds $\eta \equiv_{\mathcal{L}} \alpha$. Then, by definition of $\equiv_{\mathcal{L}}^{c}$, it would follow $\alpha \equiv_{\mathcal{L}}^{c} \beta$, a contradiction.
Let $\eta$ be a string such that $\alpha \prec \eta \prec \beta$ and $\eta \not \equiv \mathcal{L}_{\mathcal{L}} \alpha$. From $\zeta \dashv \alpha, \beta$ it follows that $\zeta \dashv \eta$, so we can write $\eta=\eta^{\prime} \zeta$ for some $\eta^{\prime} \in \Sigma^{*}$. Recall that by construction $\alpha=\alpha^{\prime} \delta \zeta$ with $\left|\alpha^{\prime} \delta\right| \leq n$, hence $|\zeta| \geq n^{2}$. Consider the last $n^{2}+1$ states $r_{0}, \ldots, r_{n^{2}}$ of $\mathcal{D}_{\mathcal{L}}$ visited by reading the string $\alpha$, and the last $n^{2}+1$ states $p_{0}, \ldots, p_{n^{2}}$ visited by reading the string $\eta$. Since $\mathcal{D}_{\mathcal{L}}$ has only $n$ states, there must exist $0 \leq i, j \leq n^{2}$ with $i<j$ such that $\left(r_{i}, p_{i}\right)=\left(r_{j}, p_{j}\right)$. Notice that it can't be $r_{i}=p_{i}$, otherwise from the determinism of $\mathcal{D}_{\mathcal{L}}$ it would follow $r_{n^{2}}=p_{n^{2}}$; from the minimality of $\mathcal{D}_{\mathcal{L}}$ it would then follow $\alpha \equiv_{\mathcal{L}} \eta$, a contradiction.
Let $\zeta^{\prime \prime}$ be the suffix of $\zeta$ of length $n^{2}-j$, and let $\gamma$ be the factor of $\zeta$ of length $j-i$ labeling the path $r_{i}, \ldots, r_{j}$. Since $|\zeta| \geq n^{2}$, there exists $\zeta^{\prime} \in \Sigma^{*}$ such that $\zeta=\zeta^{\prime} \gamma \zeta^{\prime \prime}$. We can then rewrite $\alpha, \eta$ and $\beta$ as

$$
\begin{aligned}
& \alpha=\alpha^{\prime} \delta \zeta=\alpha^{\prime} \delta \zeta^{\prime} \gamma \zeta^{\prime \prime} \\
& \eta=\eta^{\prime} \zeta=\eta^{\prime} \zeta^{\prime} \gamma \zeta^{\prime \prime} \\
& \beta=\alpha^{\prime} \zeta=\alpha^{\prime} \zeta^{\prime} \gamma \zeta^{\prime \prime}
\end{aligned}
$$

Let $k$ be an integer such that $\left|\gamma^{k}\right|$ is greater than $\left|\alpha^{\prime} \delta \zeta^{\prime}\right|$ and $\left|\eta^{\prime} \zeta^{\prime}\right|$. Set $\mu:=\eta^{\prime} \zeta^{\prime}$; from $\alpha \prec \eta \prec \beta$ it follows that $\alpha^{\prime} \delta \zeta^{\prime} \prec \mu \prec \alpha^{\prime} \zeta^{\prime}$. If $\gamma^{k} \prec \mu$ set $\nu:=\alpha^{\prime} \zeta^{\prime}$, otherwise set $\nu:=\alpha^{\prime} \delta \zeta^{\prime}$. In both cases, the hypothesis of Theorem3 are satisfied, since $\gamma^{k}$ labels two cycles starting from the states $r_{i}$ and $p_{i}$, that we have proved to be distinct. We can conclude that $\mathcal{L}$ is not Wheeler, a contradiction, and the thesis follows.

We show now how to compute the minimum WDFA recognizing a Wheeler language $\mathcal{L}$ if we are given its minimum $\operatorname{DFA} \mathcal{D}_{\mathcal{L}}$ and one of its fingeprints.

Proposition 3 (Fingerprint to min WDFA). Let $\mathcal{D}_{\mathcal{L}}$ be the minimum automaton recognizing the Wheeler language $\mathcal{L}$ with $\left|\mathcal{D}_{\mathcal{L}}\right|=n$ and let $C_{1}, \ldots, C_{m}$ be the pairwise distinct equivalence classes of $\equiv_{\mathcal{L}}^{c}$. Assume that we are given a fingerprint of $\mathcal{L}$, whose elements have length less than $n^{2}+n$. Then it is possible to build the minimum WDFA recognizing $\mathcal{L}$ in $O\left(n^{2} \cdot \sigma \cdot m \log m\right)$ time.

Proof. Let $\left\{\alpha_{1}, \ldots, \alpha_{m}\right\}$ be a fingerprint of $\mathcal{L}$ and let $\mathcal{D}_{\mathcal{L}}$ be the minimum DFA recognizing $\mathcal{L}$. We can assume without loss of generality that $\alpha_{1} \prec \ldots \prec \alpha_{m}$. We build the automaton $\mathcal{D}_{\mathcal{L}}^{W}=\left(Q, \alpha_{1}, \delta, F, \Sigma\right)$, where the set of states is $Q=$ $\left\{\alpha_{1}, \ldots, \alpha_{m}\right\}$ and the set of final states is $F=\left\{\alpha_{j}: \alpha_{j} \in \mathcal{L}\right\}$. The transition function $\delta$ can be computed as follow. For all $1 \leq j \leq m$ and for all $c \in \Sigma$, check whether $\alpha_{j} \cdot c \in \operatorname{Pref}(\mathcal{L})$. If $\alpha_{j} \cdot c \notin \operatorname{Pref}(\mathcal{L})$, there are no edges labeled $c$ that exit from $\alpha_{j}$. If instead $\alpha_{j} \cdot c \in \operatorname{Pref}(\mathcal{L})$, in order to define $\delta\left(\alpha_{j}, c\right)$ we just have to determine the $\equiv_{\mathcal{L}}^{c}$-class of the string $\alpha_{j} \cdot c$ (see Theorem 11). We first locate the position of $\alpha_{j} \cdot c$ in the intervals defined by $\alpha_{1} \prec \ldots \prec \alpha_{m}$ using a binary search. There are three possible cases.

1. $\alpha_{j} \cdot c \preceq \alpha_{1}$. Then by the properties of $\equiv_{\mathcal{L}}^{c}$ it easily follows $\alpha_{j} \cdot c \equiv_{\mathcal{L}}^{c} \alpha_{1}$ and we define $\delta\left(\alpha_{j}, c\right)=\alpha_{1}$.
2. $\alpha_{m} \preceq \alpha_{j} \cdot c$. Similarly to the previous case, we have $\alpha_{j} \cdot c \equiv_{\mathcal{L}}^{c} \alpha_{m}$ and we define $\delta\left(\alpha_{j}, c\right)=\alpha_{m}$.
3. There exists $s$ such that $\alpha_{s} \preceq \alpha_{j} \cdot c \preceq \alpha_{s+1}$. It can not be the case that both $\alpha_{j} c \not \equiv \mathcal{L}_{\mathcal{L}} \alpha_{s}$ and $\alpha_{j} c \not \mathcal{F}_{\mathcal{L}} \alpha_{s+1}$, since $\left\{\alpha_{1}, \ldots, \alpha_{m}\right\}$ is a fingerprint of $\mathcal{L}$ and $\equiv_{\mathcal{L}^{c}}^{c}$-classes are intervals in $\operatorname{Pref}(\mathcal{L})$. Hence we distinguish three cases.
(a) $\alpha_{s} \equiv_{\mathcal{L}} \alpha_{j} \cdot c \not \equiv_{\mathcal{L}} \alpha_{s+1}$. Then $\alpha_{j} \cdot c \equiv_{\mathcal{L}}^{c} \alpha_{s}$ and we define $\delta\left(\alpha_{j}, c\right)=\alpha_{s}$.
(b) $\alpha_{s} \not \equiv_{\mathcal{L}} \alpha_{j} \cdot c \equiv_{\mathcal{L}} \alpha_{s+1}$. Then $\delta\left(\alpha_{j}, c\right)=\alpha_{s+1}$.
(c) $\alpha_{s} \equiv_{\mathcal{L}} \alpha_{j} \cdot c \equiv_{\mathcal{L}} \alpha_{s+1}$. Since $\left\{\alpha_{1}, \ldots, \alpha_{m}\right\}$ is a fingerprint of $\mathcal{L}$, it is either $c=\operatorname{end}\left(\alpha_{j} c\right)=\operatorname{end}\left(\alpha_{s}\right)$, in which case $\alpha_{j} \cdot c \equiv_{\mathcal{L}}^{c} \alpha_{s}$ and we define $\delta\left(\alpha_{j}, c\right)=\alpha_{s}$, or $c=\operatorname{end}\left(\alpha_{s+1}\right)$, in which case $\alpha_{j} \cdot c \equiv_{\mathcal{L}}^{c} \alpha_{s+1}$ and we define $\delta\left(\alpha_{j}, c\right)=\alpha_{s+1}$ (where by end $(\beta)$ we denote the last letter of the string $\beta$, for $\beta \in \Sigma^{+}$).

To complete the construction, we show how to extract a fingerprint of a Wheeler language $\mathcal{L}$ starting from its minimum DFA. We first need to prove the following Lemma.

Lemma 4. Given a DFA $\mathcal{D}$ with $n$ states, a state $q$ and a string $\gamma \notin I_{q}$ with $|\gamma| \leq n^{2}+n$, we can find in polynomial time, if it exists, the greatest (smallest) string in $I_{q}$ that is smaller (greater) than $\gamma$ and has length at most $n^{2}+n$.

Proof. Let UB the the upper bound $\mathrm{UB}=n^{2}+n$. Using dynamic programming, we can extract a $n \times \mathrm{UB}$ table storing, for each $(i, j)$, the smallest and the greatest string in $I_{q_{i}}$ of length at most $j$ (see [ADPP]). Given a string $\alpha$, we use the notation $\alpha[i]$ to denote the $i$-th to last character of $\alpha$ (or $\varepsilon$ if $i>|\alpha|)$, and the notation $\alpha_{i}$ to denote the suffix of $\alpha$ of length $i$. In particular we have $\alpha_{i+1}=$ $\alpha[i+1] \cdot \alpha_{i}$. In this Lemma we are interested only in strings with length less than

UB, therefore every string (subset of strings) that will be mentioned has to be intended as an element (subset, respectively) of $\Sigma^{\leq \mathrm{UB}}=\left\{\alpha \in \Sigma^{*}:|\alpha| \leq \mathrm{UB}\right\}$.

We want to find the greatest string in $I_{q}$ that is smaller than $\gamma$. Note that if $\gamma$ is the suffix of a string $\alpha$, then $\gamma \prec \alpha$ so we do not have to worry about strings ending with $\gamma$. Note also that the greatest string smaller than $\gamma$ must maximize the length of the longest suffix it has in common with $\gamma$. Therefore, we look for all the states of $\mathcal{D}$ starting from which it is possible to read the longest proper suffix of $\gamma$ that ends in $q$. To do that, for each $1 \leq i<|\gamma|$ we build the set $S_{i}=\left\{p \in Q: p \stackrel{\gamma_{i}}{\rightsquigarrow} q\right\}$. We start from the set $S_{0}=\{q\}$ and to build $S_{i+1}$ from $S_{i}$ we simply follow backward the edges labeled $\gamma[i+1]$. Every time we determine a set $S_{i}$, we check if there exists at least one incoming edge with a label strictly less than $\gamma[i+1]$. If this is the case, we keep in memory $S_{i}$ as the last set we built with such property; previously stored sets can be overwritten. This procedure ends either when we find an $S_{i}$ that is empty or when we successfully build the last set $S_{|\gamma|-1}$. If we did not store any of the $S_{i}$ we have built, then there is no string in $I_{q}$ smaller than $\gamma$. If instead we have stored at least one $S_{i}$, we consider the last one stored (that is, the only one that has not been overwritten), say $S_{k}$. Clearly, the computation of any string in $I_{q}$ smaller than $\gamma$ that maximizes the length of the longest suffix it has in common with $\gamma$ must reach a state of $S_{k}$ at his $k$-th to last step. Therefore, let $c$ be the greatest label smaller than $\gamma[k+1]$ that enters $S_{k}$ (note that $c$ must exists since we stored $S_{k}$ ), and let $S$ be the set of states that can reach $S_{k}$ by an edge labeled $c$. Using the table computed at the very beginning of this lemma, we can easily find, if it exists, the greatest string $\bar{\alpha}$ of length at most $\mathrm{UB}-(k+1)$ that can reach a state of $S$. Then, the greatest string in $I_{q}$ that is smaller than $\gamma$ is $\bar{\alpha} \cdot c \cdot \gamma_{k}$.

To find the smallest string in $I_{q}$ that is greater than $\gamma$, we split the problem into two sub-problems: 1) find the smallest string in $I_{q}$ that is greater than $\gamma$ but has not $\gamma$ as a suffix and 2) find the smallest string in $I_{q}$ that has $\gamma$ as a suffix. The first problem is a symmetric version of the one discussed above, and can be solved in a similar way: we use exactly the same sets $S_{i}$, but this time we store a set $S_{i}$ if there exists at least one incoming edge with a label strictly greater than $\gamma[i+1]$. To also solve the second problem, instead of stopping when computing $S_{|\gamma|-1}$ we carry on and compute $S_{|\gamma|}$. We do this since the following implication holds: there exists at least one string in $I_{q}$ that has $\gamma$ as a suffix iff $S_{|\gamma|}$ is not empty and there is at least one string of length at most $U B-|\gamma|$ that can reach a state of $S_{|\gamma|}$. If $S_{|\gamma|} \neq \emptyset$, we use again the table to determine, if it exists, the smallest string $\bar{\beta}$ of length at most $U B-|\gamma|$ that can reach a state of $S_{|\gamma|}$. Lastly, we confront $\bar{\beta} \cdot \gamma$ with the string obtained by solving the first problem and we choose the smaller one.

As a last step, Algorithm 1 generates a fingerprint of a language $\mathcal{L}$ starting from the minimum DFA $\mathcal{D}_{\mathcal{L}}$. The algorithm uses the subroutines described in Lemma 4 given a DFA $\mathcal{D}$ with set of states $Q=\left\{q_{0}, \ldots, q_{n-1}\right\}$ and two strings $m, m^{\prime} \in \operatorname{Pref}(\mathcal{L}(\mathcal{D}))$ with $m \in I_{q_{k}}$ (for some $\left.0 \leq k \leq n-1\right)$,

- MinMaxPair $(\mathcal{D})$ returns the set of pairs $\left(m_{0}, M_{0}\right), \ldots,\left(m_{n-1}, M_{n-1}\right)$, where $m_{i}$ is the co-lexicographically smallest string in $I_{q_{i}}$ of length at most $n^{2}+n$, and $M_{i}$ is the greatest.
- GreatestSmaller $\left(m, m^{\prime}, \mathcal{D}\right)$ returns the greatest string in $I_{q_{k}}$ smaller than $m^{\prime}$ of length at most $n^{2}+n$.
- SmallestGreater $\left(m, m^{\prime}, \mathcal{D}\right)$ returns the smallest string in $I_{q_{k}}$ greater than $m^{\prime}$ of length at most $n^{2}+n$.

```
Algorithm 1 Min DFA to fingerprint
Input: The minimum DFA \(\mathcal{D}_{\mathcal{L}}\) recognizing \(\mathcal{L}\)
Output: A fingerprint of \(\mathcal{L}\)
    \(\tau \leftarrow \operatorname{MinMaxPairs}\left(\mathcal{D}_{\mathcal{L}}\right) \quad \triangleright\) We initialize a set of \(\left|\mathcal{D}_{\mathcal{L}}\right|\) pairs of strings
    while there exist \((m, M),\left(m^{\prime}, M^{\prime}\right) \in \tau\) such that \(m \prec m^{\prime} \prec M\) do
        \(M_{1} \leftarrow\) GreatestSmaller \(\left(m, m^{\prime}, \mathcal{D}_{\mathcal{L}}\right)\)
        \(m_{2} \leftarrow \operatorname{SmallestGreater}\left(m, m^{\prime}, \mathcal{D}_{\mathcal{L}}\right)\)
        \(\tau \leftarrow \tau \backslash\{(m, M)\}\)
        \(\tau \leftarrow \tau \cup\left\{\left(m, M_{1}\right),\left(m_{2}, M\right)\right\}\)
    end while
    \(\tau \leftarrow \operatorname{Expand}(\tau)\)
    return the first component of each element of \(\tau\)
```

At each iteration of the while cycle, we check the existence of two overlapping pairs $(m, M),\left(m^{\prime}, M^{\prime}\right)$ and replaces the first one with two new pairs $\left(m, M_{1}\right)$ and $\left(m_{2}, M\right)$. As we will prove in the Appendix, this cycle always ends. Clearly, when we exit the cycle $\tau$ can not contain overlapping pairs. We will also prove that, at this point, each pair $(m, M) \in \tau$ satisfy the following properties:

1. $m$ and $M$ belong to the same $\equiv_{\mathcal{L}^{-} \text {-class; }}$
2. if there exists a Wheeler class $C$ such that $m \prec C \prec M$, then $C \subseteq[m]_{\equiv_{\mathcal{L}}}$.

Lastly, we use the subroutine Expand to extract, from each pair $(m, M) \in \tau$, a representative of all the Wheeler classes $C$ (if any exists) such that $m \prec C \prec M$. Since property $1-2$ hold, if $\operatorname{end}(m)=\operatorname{end}(M)$ there are no Wheeler classes $C$ such that $m \prec C \prec M$; moreover, $m$ and $M$ belong to the same Wheeler class, so we leave the pair $(m, M)$ unchanged. Otherwise, if $\operatorname{end}(m) \neq \operatorname{end}(M)$ and there is a Wheeler class $C$ is such that $m \prec C \prec M$, it must be the case that the strings in $C$ end with a character that differs from both the last character of $m$ and the last one of $M$. For each character $c$ such that end $(m) \prec c \prec \operatorname{end}(M)$, we check whether there exists a string $\alpha_{c} \in I_{\delta\left(q_{0}, m\right)}$ such that end $\left(\alpha_{c}\right)=c$. Every time we find an $\alpha_{c}$ with such property, we add to $\tau$ the pair $\left(\alpha_{c}, \alpha_{c}\right)$. As a last step, we replace the pair $(m, M)$ with the pairs $(m, m)$ and $(M, M)$, since from $\operatorname{end}(m) \neq \operatorname{end}(M)$ it follows that $m$ and $M$ belong to different Wheeler classes.

After the Expand subroutine has been run, $\tau$ will contain exactly one pair for each Wheeler class $C$ of $\mathcal{L}$, whose components both belong to $C$. By extracting from each pair one of its components, e.g. the first one, we obtain a fingerprint of $\mathcal{L}$.

## 4 Conclusions

In this paper we considered a number of computational complexity problems related with the general idea of ordering states of a finite automaton. In general, ordering objects might lead to significant simplification of otherwise difficult storage and/or manipulation problems. In fact, ordered finite automata can ease such tasks as index construction, membership testing, and even determinization of NFA's accepting a given regular language. Clearly, a key point is the complexity of finding the right order from scratch. Even though this turned out to be simple on DFA's and, as opposed to the non-ordered case, turning a Wheeler NFA into a Wheeler DFA is polynomial, things become much more tricky when the input automaton is a non-deterministic one. This issue, together with some of its natural variants, were the main theme of this paper. We proved that a number of ordered-related results, ultimately guaranteeing the existence of polynomial time algorithms on DFA's, are much more complex if the starting automaton is an NFA- even in the case of a reduced NFA.

The complexity bounds we studied and presented here suggest the "dangerous" directions along which generalisations can be searched.

An interesting theme we did not explore here, is the possibility of exploiting order over more general classes of automata and languages. Can ordering states of a push-down (deterministic) automata or even a (deterministic) Turing Machine, be a way to obtain a simplification of interesting problems over the recognized languages? Can we define an order over the states of a DFA and use this order to simplify problems relative to language over infinite strings?

The order imposed on states of an accepting automaton is reflecting, in a variety of ways, the underlying properties of the ordering on strings reaching that state. The co-lexicographic order seems to be an especially effective one. However, exploring this relationship - and the corresponding complexity bounds-in interesting and expressive contexts, can be extremely stimulating in terms challenging formal language problems.

## References

1. Alanko, J., D'Agostino, G., Policriti, A., Prezza, N.: Regular languages meet prefix sorting. In: Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms. pp. 911-930 (2020). https://doi.org/10.1137/1.9781611975994.55, https://epubs.siam.org/doi/abs/10.1137/1.9781611975994.55
2. Alanko, J., D'Agostino, G., Policriti, A., Prezza, N.: Wheeler languages. Inf. Comput. 281(C) (dec 2021). https://doi.org/10.1016/j.ic.2021.104820, https://doi.org/10.1016/j.ic.2021.104820
3. Câmpeanu, C., Culik, K., Salomaa, K., Yu, S.: State complexity of basic operations on finite languages. In: Boldt, O., Jürgensen, H. (eds.) Automata Implementation. pp. 60-70. Springer Berlin Heidelberg, Berlin, Heidelberg (2001)
4. Cotumaccio, N., D'Agostino, G., Policriti, A., Prezza, N.: A theory of (co-lex) ordered regular languages. In preparation (2022)
5. Cotumaccio, N., Prezza, N.: On Indexing and Compressing Finite Automata, pp. 2585-2599. https://doi.org/10.1137/1.9781611976465.153, https://epubs.siam.org/doi/abs/10.1137/1.9781611976465.153
6. Gibney, D., Hoppenworth, G., Thankachan, S.V.: Simple reductions from formula-sat to pattern matching on labeled graphs and subtree isomorphism. In: Le, H.V., King, V. (eds.) 4th Symposium on Simplicity in Algorithms, SOSA 2021, Virtual Conference, January 11-12, 2021. pp. 232-242. SIAM (2021). https://doi.org/10.1137/1.9781611976496.26, https://doi.org/10.1137/1.9781611976496.26
7. Gibney, D., Thankachan, S.V.: On the hardness and inapproximability of recognizing wheeler graphs. In: 27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany. pp. 51:1-51:16 (2019). https://doi.org/10.4230/LIPIcs.ESA.2019.51, https://doi.org/10.4230/LIPIcs.ESA.2019.51
8. Travis Gagie, Giovanni Manzini e Sirén, J.: Wheeler graphs: A framework for bwtbased data structures. Theoretical computer science 698, 67-78 (2017)
9. Yu, S., Zhuang, Q., Salomaa, K.: The state complexities of some basic operations on regular languages. Theor. Comput. Sci. 125, 315-328 (1994)

## 5 Appendix

Proof of Proposition 2, Let $\mathcal{D}_{\mathcal{L}}=\left(\hat{Q}, \hat{q_{0}}, \hat{\delta}, \hat{F}, \Sigma\right)$ be the minimum DFA recognizing $\mathcal{L}$. Clearly $\mathcal{D}_{\mathcal{L}}$ has at most $n$ states.
$(\Longleftarrow)$ From condition 2 it follows that $\delta\left(q_{0}, \mu\right)=\delta\left(q_{0}, \mu \gamma\right)$, thus $\mu \equiv_{\mathcal{L}} \mu \gamma$. Therefore, in $\mathcal{D}_{\mathcal{L}}$ we also have $\delta\left(\hat{q}_{0}, \mu\right)=\delta\left(\hat{q}_{0}, \mu \gamma\right)$. Similarly, it holds $\delta\left(\hat{q}_{0}, \nu\right)=$ $\delta\left(\hat{q}_{0}, \nu \gamma\right)$. It follows that $\mu, \nu, \gamma$ satisfy condition 1-3 of Theorem 3, hence $\mathcal{L}$ is not Wheeler.
$(\Longrightarrow)$ Since $\mathcal{L}$ is not Wheeler, let $\hat{\mu}, \hat{\nu}, \hat{\gamma}$ be three strings satisfying conditions 1-4 of Theorem 3 The DFA $\mathcal{D}_{\mathcal{L}}$ has at most $n$ states, hence the length of $\hat{\mu}, \hat{\nu}$ and $\hat{\gamma}$ is bounded by $n^{3}+2 n^{2}+n+2$. We have $\hat{\mu} \hat{\gamma}^{*} \subseteq \operatorname{Pref}(\mathcal{L})$, so let $t_{0}=q_{0}, t_{1}, \ldots, t_{m}$ be a run of $\hat{\mu} \hat{\gamma}^{n}$ over $\mathcal{D}$. We set $u:=|\hat{\mu}|$ and $g:=|\hat{\gamma}|$, and consider the list of $n+1$ states

$$
t_{u}, t_{u+g}, t_{u+2 g}, \ldots, t_{u+n g}=t_{m}
$$

Since $\mathcal{D}$ has $n$ states, there must exist two integers $0 \leq h<k \leq n$ such that $t_{u+h g}=t_{u+k g}$. That is, there exists a state $p:=t_{u+h g}$ such that $p \in \delta\left(q_{0}, \hat{\mu} \hat{\gamma}^{h}\right)$ and $\hat{\gamma}^{k-h}$ labels a cycle starting from $p$. We can repeat the same argument for a run of $\hat{\nu} \hat{\gamma}^{n}$ over $\mathcal{D}$ to find a state $r$ and two integers $h^{\prime}, k^{\prime}$ such that $r \in \delta\left(q_{0}, \hat{\nu} \hat{\gamma}^{h^{\prime}}\right)$ and $\hat{\gamma}^{h^{\prime}-h^{\prime}}$ labels a cycle starting from $r$. We define the constant $h^{\prime \prime}$ as the minimum multiple of $(k-h) \cdot\left(k^{\prime}-h^{\prime}\right)$ greater than $\max \left\{h+1, h^{\prime}+1\right\}$; it can be proved that $h^{\prime \prime} \leq n^{2}$, and by construction $\hat{\gamma}^{h^{\prime \prime}}$ labels both a cycle starting from $p$ and one starting from $r$. We then define the strings

$$
\begin{aligned}
\mu & :=\hat{\mu} \hat{\gamma}^{h} \\
\nu & :=\hat{\nu} \hat{\gamma}^{h^{\prime}} \\
\gamma & :=\hat{\gamma}^{h^{\prime \prime}},
\end{aligned}
$$

which satisfy conditions 2 and 3 . Note that we have chosen a $h^{\prime \prime}$ such that $|\gamma|>|\mu|,|\nu|$, so that $\gamma A \mu, \nu$. Condition 4 is satisfied since $|\hat{\gamma}| \leq n^{3}+2 n^{2}+n+2$ and $h^{\prime \prime} \leq n^{2}$. Lastly, condition 1 is satisfied since the strings $\hat{\mu}$ and $\hat{\mu} \hat{\gamma}^{h}$ lead to the same state of $\mathcal{D}_{\mathcal{L}}$, thus $\hat{\mu} \equiv_{\mathcal{L}} \hat{\mu} \hat{\gamma}^{h}$. Similarly, we have $\hat{\nu} \equiv_{\mathcal{L}} \hat{\nu} \hat{\gamma}^{h^{\prime}}$. The thesis then follows from $\hat{\mu} \not \equiv \mathcal{L} \hat{\nu}$.

Termination and correctness of Algorithm 1, We start by analyzing the subroutines used by the algorithm. The subroutine MinMaxPairs can be computed simply by looking at the $n \times \mathrm{UB}$ table described in Lemma 4 Note that Lemma 3 ensures that $m_{q_{i}}$ (respectively, $M_{q_{i}}$ ) belongs to the smaller (greatest) Wheeler class contained in $I_{q_{i}}$. Subroutines GreatestSmaller and SmallestGreater are thoroughly described in Lemma 4

We now prove that Algorithm $\square$ always terminates. Let $\tau_{i}$ be the set $\tau$ at the end of the $i-t h$ iteration of the while cycle, and for any pair of strings $c=(m, M)$ let $w(c)$ denote the number of Wheeler classes $C$ such that $m \prec C \prec M$ and $C \nsubseteq[m]_{\equiv_{C}}$. Given a set of pairs $\tau$, let $w(\tau)$ denote the value

$$
w(\tau):=\sum_{c \in \tau} w(c) \geq 0 .
$$

We say that two pairs $c=(m, M)$ and $c^{\prime}=\left(m^{\prime}, M^{\prime}\right)$ are ordered if either $M \prec m^{\prime}$ or $M^{\prime} \prec m$. In order to prove that $w\left(\tau_{i+1}\right)<w\left(\tau_{i}\right)$, we will maintain the following invariants:
0 . if $c=(m, M)$ and $c^{\prime}=\left(m^{\prime}, M^{\prime}\right)$ are two distinct pairs in $\tau_{i}$, then $\{m, M\} \cap$ $\left\{m^{\prime}, M^{\prime}\right\}=\emptyset$;

1. if $c=(m, M) \in \tau_{i}$, then $m \preceq M, m \equiv_{\mathcal{L}} M$, and end $(m)=\operatorname{end}(M)$;
2. if $c=(m, M)$ and $c^{\prime}=\left(m^{\prime}, M^{\prime}\right)$ are two distinct pairs in $\tau_{i}$ such that $m \equiv_{\mathcal{L}}$ $m^{\prime}$, then $c$ and $c^{\prime}$ are ordered. Moreover, if $c$ and $c^{\prime}$ are also consecutive, that is if there is no $c^{*}=\left(m^{*}, M^{*}\right)$ in $\tau_{\mathcal{L}}$ such that $m \equiv_{\mathcal{L}} m^{*}$ and $m \prec m^{*} \prec m^{\prime}$, then there is no Wheeler class $C \subseteq[m]_{\equiv_{\mathcal{c}}}$ such that $m \prec C \prec m^{\prime}$;
3. let $x$ be the first or second component of any pair $c$ in $\tau_{i}$ and $y$ be the first or second component of any pair $c^{\prime}$ in $\tau_{i}$. If $x \equiv_{\mathcal{L}}^{c} y$, then $c=c^{\prime}$.
Note that invariant 0 implies that every time we have two distinct, not ordered pairs $c=(m, M)$ and $c^{\prime}=\left(m^{\prime}, M^{\prime}\right)$, the strict inequalities $m \prec M^{\prime}$ and $m^{\prime} \prec M$ hold. By construction, these invariants hold for $\tau_{0}$, which is the set returned by MinMaxPairs. For instance, invariant 2, 3 hold since in $\tau_{0}$ distinct pairs have components belonging to different $\equiv_{\mathcal{L}}$-classes.

Invariants $0 \mid 2$ can be easily proved by induction on $i$ just by looking at how new pairs are created. We prove by induction invariant 3) suppose that it holds for $\tau_{i}$. Let $c=(m, M), c^{\prime}=\left(m^{\prime}, M^{\prime}\right)$ be the pairs that meet the while condition on Line 2, and let $c_{1}=\left(m, M_{1}\right), c_{2}=\left(m_{2}, M\right)$ be the two pairs that replace $c$ on Line 5-6. Note that $c, c^{\prime}$ are not ordered, hence invariant 2 implies that $m$ and $m^{\prime}$ belong to different $\equiv \mathcal{L}$-classes. Suppose by contradiction that the invariant does not hold for $\tau_{i+1}$, that is, there exist two distinct pairs $d, d^{\prime} \in \tau_{i+1}$ such that a component $x$ of $d$ belongs to the same Wheeler class of a component $y$ of $d^{\prime}$. By induction, it can not be the case that $d, d^{\prime} \in \tau_{i}$. Therefore, at least one among $d$ and $d^{\prime}$ belongs to $\left\{c_{1}, c_{2}\right\}$. Moreover, $d$ and $d^{\prime}$ can not both belong to $\left\{c_{1}, c_{2}\right\}$ : we have by construction that the (possibly identical) Wheeler classes of $m$ and $M_{1}$ are different from the Wheeler classes of $m_{2}$ and $M$. We assume, w.l.o.g., that $d$ belongs to $\left\{c_{1}, c_{2}\right\}$ whereas $d^{\prime}$ doesn't; in particular $d^{\prime} \in \tau_{i}$ and since $d^{\prime} \in \tau_{i+1}$ we also have $d^{\prime} \neq c$. There are two possibilities. If $d=c_{1}$, it can not be the case that $x=m$, otherwise the pairs $c, d^{\prime} \in \tau_{i}$ would violate the inductive hypothesis. Thus $x=M_{1}$. From $y \equiv_{\mathcal{L}}^{c} M_{1} \equiv_{\mathcal{L}} m$ it follows that $y \equiv_{\mathcal{L}} m$ and invariants [ 2 applied to $\tau_{i}$ imply that $c$ and $d^{\prime}$ are ordered, that is, either $y \prec m$ or $M \prec y$ holds. If $y \prec m$ we get $y \prec m \preceq M_{1}$, thus $m$ belongs to the same Wheeler class of $y$. If $M \prec y$ we get $M_{1} \preceq M \prec y$, thus $M$ belongs to the same Wheeler class of $y$. In both cases, considering $c, d^{\prime} \in \tau_{i}$, we reach a contradiction with our inductive hypothesis.
If instead $d=c_{2}$, we use a similar argument to show that $x=m_{2}$ and that either $y \prec m \preceq m_{2}$ or $m_{2} \preceq M \prec y$ hold. Since both inequalities lead to a contradiction, we can conclude that invariant 3 holds. Hence we proved that invariant 3 holds for all $\tau_{i}$.

We can now prove that $w\left(\tau_{i+1}\right)<w\left(\tau_{i}\right)$. Let ( $m, M_{1}$ ), ( $m_{2}, M$ ) be the pairs added to $\tau_{i}$ on Line 6 of the Algorithm Note that if $C$ is a Wheeler class such that $m \prec C \prec M$ and $C \nsubseteq[m]_{\equiv_{\mathcal{L}}}$, it can not be the case that both $m \prec C \prec M_{1}$
and $m_{2} \prec C \prec M$ occurs, since $M_{1} \prec m_{2}$. Moreover, let $C^{\prime}$ be the Wheeler class containing $m^{\prime}$. From invariant 2 it follows that $C^{\prime} \nsubseteq[m]_{\equiv_{\mathcal{L}}}$, and from $M_{1} \prec m^{\prime} \prec m_{2}$ it follows that neither $m \prec C^{\prime} \prec M_{1}$ nor $m_{2} \prec C^{\prime} \prec M$ holds. Therefore we have $w\left(c_{1}\right)+w\left(c_{2}\right) \leq w(c)-1$, and $w\left(\tau_{i+1}\right)<w\left(\tau_{i}\right)$ follows.

We want to prove that if $w\left(\tau_{i}\right)>0$ then there exist two pairs $c=(m, M)$ and $c^{\prime}=\left(m^{\prime}, M^{\prime}\right)$ in $\tau_{i}$ such that $c$ and $c^{\prime}$ are not ordered, thus proving that $w\left(\tau_{i}\right)=0$ holds when we exit the while cycle. If $w\left(\tau_{i}\right)>0$, then there exists a pair $c=(m, M)$ in $\tau_{i}$ such that $w(c)>0$, that is, there exists a Wheeler class $C$ with $m \prec C \prec M$ and $C \nsubseteq[m]_{\equiv_{\mathcal{L}}}$; in particular we have $m \not \equiv_{\mathcal{L}}^{c} M$. By Lemma 3, there exists $\alpha \in C$ with $|\alpha| \leq n^{2}+n$. We want to prove, by induction on $j$, that for each $0 \leq j \leq i$ there exists a pair $c_{j}=\left(m_{j}, M_{j}\right) \in \tau_{j}$ such that $c$ and $c_{j}$ are not ordered and $m_{j} \equiv_{\mathcal{L}} \alpha$; note that $c$ may not belong to $\tau_{j}$ for $j<i$. If $j=0$, let $q_{k}$ be the state of $\mathcal{D}_{\mathcal{L}}$ such that $\alpha \in I_{q_{k}}$. The pairs $c$ and $c_{k}=\left(m_{k}, M_{k}\right) \in \tau_{0}$ are not ordered, since we have both $m \prec \alpha \prec M$ and $m_{k} \preceq \alpha \preceq M_{k}$, therefore we set $c_{0}:=c_{k}$.
If $0<j<i$, suppose the thesis holds for $\tau_{j}$, that is, there exists a pair $c_{j}=$ $\left(m_{j}, M_{j}\right)$ in $\tau_{j}$ such that $c$ and $c_{j}$ are not ordered and $m_{j} \equiv_{\mathcal{L}} \alpha$. If $c_{j} \in \tau_{i+1}$, we set $c_{j+1}:=c_{j}$. Otherwise, $c_{j}$ had been split into two pairs $c_{j 1}=\left(m_{j}, M_{j 1}\right)$ and $c_{j 2}=\left(m_{j 2}, M_{j}\right)$ with $m_{j} \preceq M_{j 1} \prec m_{j 2} \preceq M_{j}$. By construction it holds $m_{j 1} \equiv_{\mathcal{L}} m_{j 2} \equiv_{\mathcal{L}} m_{j} \equiv_{\mathcal{L}} \alpha$. Since $c$ and $c_{j}$ are not ordered, we have $m \prec M_{j}$ and $m_{j} \prec M$. If $m \prec M_{j 1}$ we have both $m \prec M_{j 1}$ and $m_{j} \prec M$, hence $c$ and $c_{j 1}$ are not ordered and we set $c_{j+1}:=c_{j 1}$; similarly, if $m_{j 2} \prec M$ we set $c_{j+1}:=c_{j 2}$. Otherwise we have $M_{j 1} \preceq m$ and $M \preceq m_{j 2}$ and strings have the following order:

$$
\begin{equation*}
m_{j} \preceq M_{j 1} \preceq m \prec \alpha \prec M \preceq m_{j 2} \preceq M_{j} . \tag{2}
\end{equation*}
$$

Let $m^{\prime}$ be the string used to split $c_{j}$ : by construction it holds $M_{j 1} \prec m^{\prime} \prec$ $m_{j 2}$, and it can not be $m^{\prime}=\alpha$ since $m_{j} \equiv \mathcal{L} \alpha$. By construction, the string $m_{j 2}$ is the smallest string $\gamma$ of length at most $n^{2}+n$ such that $m^{\prime} \prec \gamma$ and $\gamma \equiv \mathcal{L} m_{j}$. Thus, if $m^{\prime} \prec \alpha$, the string $\alpha$ would have all this properties, hence it would follow that $m_{j 2} \preceq \alpha$, which contradicts (22). Similarly, if $\alpha \prec m^{\prime}$ it would follow $\alpha \preceq M_{j 1}$, a contradiction. Therefore the condition depicted in (22) can not occur, ending the proof of the inductive step. Hence, if $c \in \tau_{i}$ and $w(c)>0$ then there exists $c^{\prime} \in \tau_{i}$ such that $c, c^{\prime}$ are not ordered and we iterate the while cycle.

Let $\tau_{p}$ be the last set built before exiting the while cycle. We need to prove that the collection of the first components of the pairs in $\operatorname{Expand}\left(\tau_{p}\right)$ is a fingerprint of $\mathcal{L}$. First we prove that all pairs in $\tau_{p}$ are ordered. Let $c_{i}=\left(m_{i}, M_{i}\right)$ and $c_{j}=\left(m_{j}, M_{j}\right)$ be two distinct pairs in $\tau_{p}$. If $m_{i} \equiv_{\mathcal{L}} m_{j}$, then $c_{i}$ and $c_{j}$ are ordered by invariant 2, If instead $m_{i} \not \equiv \mathcal{L} m_{j}$, suppose $c_{i}$ and $c_{j}$ are not ordered. Then either $m_{i} \prec m_{j} \prec M_{i}$ or $m_{j} \prec m_{i} \prec M_{j}$. In the first case, if $C$ is the Wheeler class containing $m_{j}$, we have $m_{i} \prec C \prec M_{j}$ and $w\left(c_{i}\right)>0$. Similarly, the second case implies $w\left(c_{j}\right)>0$ and in both cases we reach a contradiction with $w\left(\tau_{p}\right)=0$.

Second, we prove that if $C$ is a Wheeler class that is not represented by $\tau_{p}$, i.e. $C$ does not contain any components of any pair in $\tau_{p}$, then there exists a pair $(m, M) \in \tau_{p}$ such that $m \prec C \prec M$. Then the proof is complete, since the
subroutine Expand extracts a representative $\alpha_{c}$ of $C$ and adds to $\tau_{p}$ the pair $\left(\alpha_{c}, \alpha_{c}\right)$. Let $C$ be a Wheeler class not represented by $\tau_{p}$, and consider the state $q_{i}$ in $\mathcal{D}_{\mathcal{L}}$ such that $C \subseteq I_{q_{i}}$. By construction, the pair $\left(m_{q_{i}}, M_{q_{i}}\right) \in \tau_{0}$ is such that $m_{q_{i}}$ (respectively, $M_{q_{i}}$ ) belongs to the smaller (greatest) Wheeler class contained in $I_{q_{i}}$. Since when we build $\tau_{i+1}$ from $\tau_{i}$ we only add, and never delete, representatives of Wheeler classes, both $m_{q_{i}}$ and $M_{q_{i}}$ must appear as a component of some pair in $\tau_{p}$. Therefore, it is well defined the smallest representative $m$ in $\tau_{p}$ such that $m \prec C$, as well as the greatest representative $M$ in $\tau_{p}$ such that $C \prec M$. The second part of invariant 2 implies that $m$ and $M$ belong to the same pair, which completes the proof.

