Electronic version of an article published as International Journal of Foundations of Computer Science, 24(07), 2013, 1009–1027.
http://dx.doi.org/10.1142/S0129054113400285 © World Scientific Publishing Company http://www.worldscientific.com/

COMPLEXITY OF ATOMS OF REGULAR LANGUAGES

JANUSZ BRZOZOWSKI
David R. Cheriton School of Computer Science,
University of Waterloo, Waterloo, ON, Canada N2L 3G1
brzozo@uwaterloo.ca

HELLIS TAMM
Institute of Cybernetics, Tallinn University of Technology,
Akadeemia tee 21, 12618 Tallinn, Estonia
hellis@cs.ioc.ee

Received (Day Month Year)
Accepted (Day Month Year)
Communicated by (xxxxxxxxxx)

The quotient complexity of a regular language L, which is the same as its state complexity, is the number of left quotients of L. An atom of a non-empty regular language L with n quotients is a non-empty intersection of the n quotients, which can be uncomplemented or complemented. An NFA is atomic if the right language of every state is a union of atoms. We characterize all reduced atomic NFAs of a given language, *i.e.*, those NFAs that have no equivalent states. We prove that, for any language L with quotient complexity n, the quotient complexity of any atom of L with r complemented quotients has an upper bound of $2^n - 1$ if $r = 0$ or $r = n$; for $1 \leqslant r \leqslant n-1$ the bound is

$$1 + \sum_{k=1}^{r} \sum_{h=k+1}^{k+n-r} \binom{n}{h}\binom{h}{k}.$$

For each $n \geqslant 2$, we exhibit a language with 2^n atoms which meet these bounds.

Keywords: atoms, finite automaton, atomic NFA, quotient complexity, regular language, state complexity, syntactic semigroup, witness

2000 Mathematics Subject Classification: 68Q45, 68Q19, 68Q70

1. Terminology and Notation

In this section we provide some background information, introduce atoms of regular languages, and state our reasons for studying them. For basic properties of regular languages and finite automata see [7, 9].

If Σ is a non-empty finite alphabet, then Σ^* is the free monoid generated by Σ. A *word* is any element of Σ^*, and the empty word is ε. A *language* over Σ is any subset of Σ^*. The *reverse* of a language L is denoted by L^R and defined as $L^R = \{w^R \mid w \in L\}$, where w^R is w spelled backwards.

The *(left) quotient* of a regular language L over an alphabet Σ by a word $w \in \Sigma^*$ is the language $w^{-1}L = \{x \in \Sigma^* \mid wx \in L\}$. It is well known that the quotient of a regular language is itself regular, and that a language is regular if and only if it has a finite number of distinct quotients. Also, L is its own quotient by the empty word ε, that is $\varepsilon^{-1}L = L$. Note too that the quotient by $u \in \Sigma^*$ of the quotient by $w \in \Sigma^*$ of L is the quotient by wu of L, that is, $u^{-1}(w^{-1}L) = (wu)^{-1}L$.

Although quotients have been known for over half a century, atoms were introduced only in 2011 by Brzozowski and Tamm [4]; here we use a slightly different definition for reasons explained at the end of Section 2. An *atom* of a regular language L with quotients $K_0, \dots, K_{n-1}$ is any non-empty language of the form $\widetilde{K_0} \cap \dots \cap \widetilde{K_{n-1}}$, where $\widetilde{K_i}$ is either K_i or $\overline{K_i}$, and $\overline{K_i}$ is the complement of K_i with respect to Σ^*. The atoms of any regular language L have the following properties, which have either been shown in [4] or are easily verified for the new definition:

(1) Atoms are regular because quotients are regular and regularity is preserved under complement and intersection.
(2) If L has n quotients, it has at most 2^n atoms, by the definition of atoms.
(3) Atoms are pairwise disjoint because, if two intersections differ, there must be a quotient that is complemented in one intersection but not in the other.
(4) The atoms of L partition Σ^*, since the union of all the intersections is Σ^*.
(5) Every quotient K of L is a (possibly empty) union of atoms, namely all those atoms in which K is not complemented.
(6) Every quotient of an atom is a (possibly empty) union of atoms, because the quotient of an intersection of quotients of L is an intersection of quotients of L.
(7) The complement of L is a union of atoms of L, namely all those atoms in which L is complemented.

In summary, the atoms of a regular language are its basic building blocks.

A *nondeterministic finite automaton (NFA)* is a quintuple $\mathfrak{N} = (Q, \Sigma, \eta, I, F)$, where Q is a finite, non-empty set of *states*, Σ is a finite non-empty *alphabet*, $\eta\colon Q \times \Sigma \to 2^Q$ is the *transition function*, $I \subseteq Q$ is the set of *initial states*, and $F \subseteq Q$ is the set of *final states*. As usual, we extend the transition function to functions $\eta'\colon Q \times \Sigma^* \to 2^Q$, and $\eta''\colon 2^Q \times \Sigma^* \to 2^Q$, but we use η for all three functions.

The *language accepted* by an NFA $\mathfrak{N}$ is $L(\mathfrak{N}) = \{w \in \Sigma^* \mid \eta(I, w) \cap F \neq \emptyset\}$. Two NFAs are *equivalent* if they accept the same language. The *right language* of a state q of $\mathfrak{N}$ is $L_{q,F}(\mathfrak{N}) = \{w \in \Sigma^* \mid \eta(q, w) \cap F \neq \emptyset\}$. The *right language* of a set S of states of $\mathfrak{N}$ is $L_{S,F}(\mathfrak{N}) = \bigcup_{q \in S} L_{q,F}(\mathfrak{N})$; hence $L(\mathfrak{N}) = L_{I,F}(\mathfrak{N})$. A state is *empty* if its right language is empty. Two states of an NFA are *equivalent* if their right languages are equal. The *left language* of a state q of $\mathfrak{N}$ is $L_{I,q} = \{w \in \Sigma^* \mid q \in \eta(I, w)\}$. A state is *unreachable* if its left language is empty. An NFA is *trim* if it has no empty or unreachable states. An NFA is *reduced* if it has no equivalent states. An NFA is *minimal* if it has the minimal number of states among all the equivalent NFAs.

A *deterministic finite automaton (DFA)* is a quintuple $\mathfrak{D} = (Q, \Sigma, \delta, q_0, F)$, where Q, Σ, and F are as in an NFA, $\delta\colon Q \times \Sigma \to Q$ is the transition function, and $q_0 \in Q$ is the initial state. It is evident that a DFA is a special type of NFA.

A DFA is *minimal* if all of its states are reachable, and no two states are equivalent. It is well-known that for every regular language L there exists a unique (up to isomorphism) minimal DFA.

We use the following operations on automata:

(1) The *determinization* operation D applied to an NFA $\mathfrak{N}$ yields a DFA $\mathfrak{N}^D$ obtained by the subset construction, where only subsets reachable from the initial subset of $\mathfrak{N}^D$ are used and the empty subset, if present, is included.
(2) The *reversal* operation R applied to an NFA $\mathfrak{N}$ yields an NFA $\mathfrak{N}^R$, where sets of initial and final states of $\mathfrak{N}$ are interchanged and each transition is reversed.

2. Quotient DFAs and Átomata

Let L be any non-empty regular language, and let its set of quotients be $\mathcal{K} = \{K_0, \dots, K_{n-1}\}$. The quotient of $\varepsilon^{-1}L = L$ is called *initial* and is denoted by K_{in}. The set of *final* quotients is $\mathcal{F} = \{K_i \mid \varepsilon \in K_i\}$. In the following definition we use a 1-1 correspondence $K_i \leftrightarrow \mathbf{K}_i$ between quotients K_i of a language L and states $\mathbf{K}_i$ of the *quotient DFA* $\mathfrak{D}$ defined below. We refer to the $\mathbf{K}_i$ as *quotient symbols*.

Definition 1. *The* quotient DFA *of L is $\mathfrak{D} = (\boldsymbol{K}, \Sigma, \delta, \mathbf{K}_{in}, \boldsymbol{F})$, where $\boldsymbol{K} = \{\mathbf{K}_0, \dots, \mathbf{K}_{n-1}\}$, $\delta(\mathbf{K}_i, a) = \mathbf{K}_j$ if and only if $a^{-1}K_i = K_j$ for all $\mathbf{K}_i, \mathbf{K}_j \in \boldsymbol{K}$ and $a \in \Sigma$, $\mathbf{K}_{in}$ corresponds to K_{in}, and $\boldsymbol{F} = \{\mathbf{K}_i \mid K_i \in \mathcal{F}\}$.*

The quotient DFA has the following properties:

(1) The right language of state $\mathbf{K}_i$ is the quotient K_i.
(2) The left language of state $\mathbf{K}_i$ is $\{w \in \Sigma^* \mid w^{-1}L = K_i\}$.
(3) $L(\mathfrak{D})$ is the right language of state $\mathbf{K}_{in}$, and hence $L(\mathfrak{D}) = L$.
(4) $\mathfrak{D}$ is minimal, since all the quotients in $\mathcal{K}$ are distinct.
(5) The complement $\overline{L}$ of L is accepted by the DFA $\mathfrak{D}' = (\boldsymbol{K}, \Sigma, \delta, \mathbf{K}_{in}, \boldsymbol{K} \setminus \boldsymbol{F})$, obtained from $\mathfrak{D}$ by changing the *final* states.

In summary, the quotients of a regular language define its minimal DFA.

Next, we use atoms instead of quotients as states of an automaton. An atom is *initial* if it has L (rather than $\overline{L}$) as a term; it is *final* if it contains ε. Since L is non-empty, it has at least one quotient containing ε. Hence it has exactly one final atom, the atom $\widehat{K_0} \cap \dots \cap \widehat{K_{n-1}}$, where $\widehat{K_i} = K_i$ if $\varepsilon \in K_i$, and $\widehat{K_i} = \overline{K_i}$ otherwise. If the intersection $\overline{K_0} \cap \dots \cap \overline{K_{n-1}}$ is non-empty, then we call it the *negative* atom; all the other atoms are *positive*. Let the number of atoms be m and let the number of positive atoms be p. Let $\mathcal{A} = \{A_0, \dots, A_{m-1}\}$ be the set of atoms of L. By convention, $\mathcal{I}$ is the set of initial atoms, A_{p-1} is the final atom, and the negative atom, if present, is A_{m-1}. The negative atom can never be final, since there must be at least one complemented final quotient in its intersection.

As above, we use a 1-1 correspondence $A_i \leftrightarrow \mathbf{A}_i$ between atoms A_i of a language L and the states $\mathbf{A}_i$ of the NFA $\mathfrak{A}$ defined below. We refer to the $\mathbf{A}_i$ as *atom symbols*.

Definition 2. *The* átomaton *of L is the NFA $\mathfrak{A} = (\boldsymbol{A}, \Sigma, \alpha, \boldsymbol{A}_I, \{\mathbf{A}_{p-1}\})$, where $\boldsymbol{A} = \{\mathbf{A}_i \mid A_i \in \mathcal{A}\}$, $\mathbf{A}_j \in \alpha(\mathbf{A}_i, a)$ if and only if $aA_j \subseteq A_i$, for all $\mathbf{A}_i, \mathbf{A}_j \in \boldsymbol{A}$ and $a \in \Sigma$, $\boldsymbol{A}_I = \{\mathbf{A}_i \mid A_i \in \mathcal{I}\}$, and $\mathbf{A}_{p-1}$ corresponds to A_{p-1}.*

The átomaton of any regular language L has the properties listed below; these results are either proved in [4], or easily verified for the new definition of átomaton.

(1) The right language of state $\mathbf{A}_i$ is the atom A_i.
(2) If A_i is not the negative atom, the left language of state $\mathbf{A}_i$ is $L_{\boldsymbol{A}_I, \mathbf{A}_i}(\mathfrak{A}) = ((x^R)^{-1}L^R)^R$, for $i \in \{0, \dots, p-1\}$ and $x \in A_i$, and this left language is non-empty. The left language of the negative atom is empty.
(3) The language accepted by the átomaton is $L(\mathfrak{A}) = L_{\boldsymbol{A}_I, \{\mathbf{A}_{p-1}\}} = L$.
(4) $\mathfrak{A}$ is reduced, since all the atoms in $\mathcal{A}$ are distinct.
(5) The complement $\overline{L}$ of L is accepted by $\mathfrak{A}' = (\boldsymbol{A}, \Sigma, \alpha, \boldsymbol{A}'_I, \{\mathbf{A}_{p-1}\})$, obtained from $\mathfrak{A}$ by changing the *initial* states to states whose atoms have $\overline{L}$ as a term.
(6) The determinized version $\mathfrak{A}^D$ of $\mathfrak{A}$ is isomorphic to the quotient DFA of L.
(7) The reverse $\mathfrak{A}^R$ of $\mathfrak{A}$ is isomorphic to the quotient DFA of L^R.

In summary, the atoms of a regular language define a unique reduced NFA of the language, and this NFA has some remarkable properties.

Example 3. *Let $L_2 \subseteq \{a, c\}^*$ be defined by the quotient equations below (left) and recognized by the DFA $\mathfrak{D}_2$ of Fig. 1 (a), where the initial state is indicated by an arrow and the final state, by a double circle.*

$$
\begin{aligned}
K_0 &= aK_1 \cup cK_0, & K_0 \cap K_1 &= a(K_0 \cap K_1) \cup c[(K_0 \cap K_1) \cup (K_0 \cap \overline{K_1})],\\
K_1 &= aK_0 \cup cK_0 \cup \varepsilon, & K_0 \cap \overline{K_1} &= a(\overline{K_0} \cap K_1),\\
& & \overline{K_0} \cap K_1 &= a(K_0 \cap \overline{K_1}) \cup \varepsilon,\\
& & \overline{K_0} \cap \overline{K_1} &= a(\overline{K_0} \cap \overline{K_1}) \cup c[(\overline{K_0} \cap \overline{K_1}) \cup (\overline{K_0} \cap K_1)].
\end{aligned}
$$

The equations for the atoms of L_2 are above on the right; they are obtained directly from the quotient equations [4]. *For example,*

$$
\begin{aligned}
K_0 \cap \overline{K_1} &= (aK_1 \cup cK_0) \cap \overline{(aK_0 \cup cK_0 \cup \varepsilon)}\\
&= (aK_1 \cup cK_0) \cap (a\overline{K_0} \cup c\overline{K_0})\\
&= a(K_1 \cap \overline{K_0}) \cup c(K_0 \cap \overline{K_0}) = a(\overline{K_0} \cap K_1).
\end{aligned}
$$

The átomaton $\mathfrak{A}_2$ is in Fig. 1 (b); here each atom is denoted by A_P, where P is the set of uncomplemented quotients. Thus $K_0 \cap \overline{K_1}$ becomes $A_{\{0\}}$, etc., and we represent the sets in the subscripts without brackets and commas. The reverse $\mathfrak{D}_2^R$ of $\mathfrak{D}_2$ is in Fig. 1 (c). The determinized reverse $\mathfrak{D}_2^{RD}$ is in Fig. 1 (d); this is the minimal DFA for L_2^R, the reverse of L_2. The reverse $\mathfrak{A}_2^R$ of the átomaton is in Fig. 1 (e). Note that $\mathfrak{D}_2^{RD}$ and $\mathfrak{A}_2^R$ are isomorphic.

We are now in a position to explain the differences between our present definition of an atom and that of [4]. The definition in [4] did not consider the intersection

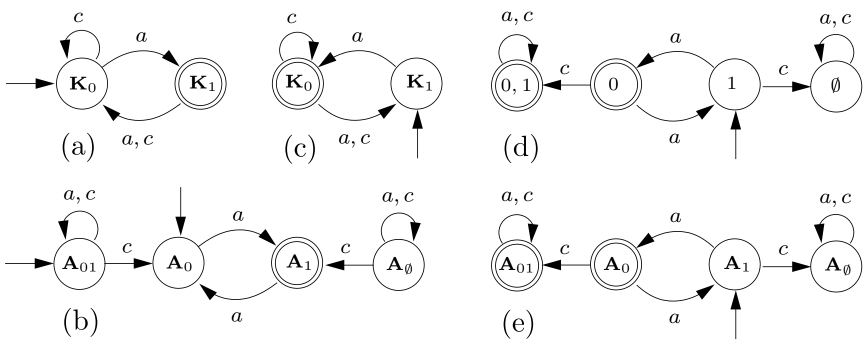

Fig. 1. (a) DFA $\mathfrak{D}_2$; (b) Átomaton $\mathfrak{A}_2$; (c) NFA $\mathfrak{D}_2^R$; (d) DFA $\mathfrak{D}_2^{RD}$; (e) DFA $\mathfrak{A}_2^R$.

of all the complemented quotients to be an atom, and so all atoms were positive. It was shown in [4] that the reverse of the átomaton with only positive atoms is the trim version of the minimal DFA of L^R. With the negative atom, we avoid the trimming operation; so the reverse of the átomaton is the minimal DFA of L^R. Also, with the negative atom, a language L and its complement language $\overline{L}$ have the same atoms. In addition, we have symmetry between the atoms with 0 and n complemented quotients, and the same upper bounds on quotient complexity for both, as will be shown in Section 5.

One might also consider a model in which there is an empty atom. Then there would be unnecessary transitions from every atom under every input to the empty atom. If this átomaton were reversed, the DFA for L^R would have an unreachable state. For this reason we avoided this definition.

3. Atomic NFAs

We show now that atoms lead naturally to a new class of NFAs: DFAs and átomata are special cases of atomic NFAs introduced in [4] and studied further in [5].

In this section we deal only with trim NFAs; thus we do not include the negative atom in the átomaton, if present. This also implies that we do not include the empty state when we determinize.

Definition 4. *An NFA* $\mathfrak{N} = (Q, \Sigma, \eta, I, F)$ *is* atomic *if for every* $q \in Q$*, the right language* $L_{q,F}(\mathfrak{N})$ *of* q *is a union of some atoms of* $L(\mathfrak{N})$.

The following theorem, slightly restated, was proved in [4]:

Theorem 5 (Atomicity) *A trim NFA* $\mathfrak{N}$ *is atomic if and only if* $\mathfrak{N}^{RD}$ *is minimal.*

This theorem allows us to test whether an NFA $\mathfrak{N}$ accepting a language L is atomic. To do this, reverse $\mathfrak{N}$ and apply the subset construction. Then $\mathfrak{N}$ is atomic if and only if $\mathfrak{N}^{RD}$ is isomorphic to the minimal DFA of L^R.

If we allow equivalent states, there is an infinite number of atomic NFAs, but their behaviours are not distinct; hence we consider only reduced NFAs. Suppose $\mathfrak{B} = (\mathcal{B}, \Sigma, \beta, \mathcal{B}_I, \mathcal{B}_F)$ is any trim reduced atomic NFA accepting L. Since $\mathfrak{B}$ is atomic, the right language of any state of $\mathfrak{B}$ is a union of positive atoms of L; hence the states of $\mathfrak{B}$ can be represented by sets of positive atom symbols. Because $\mathfrak{B}$ is trim, it does not have a state with the empty set of atom symbols. Since $\mathfrak{B}$ is reduced, no set of atom symbols appears twice. Thus the state set $\mathcal{B}$ is a collection of non-empty sets of positive atom symbols.

Theorem 6 (Legality) *Suppose L is a regular language, its átomaton is $\mathfrak{A} = (\boldsymbol{A}, \Sigma, \alpha, \boldsymbol{A}_I, \{\mathbf{A}_{p-1}\})$, and $\mathfrak{B} = (\mathcal{B}, \Sigma, \beta, \mathcal{B}_I, \mathcal{B}_F)$ is a trim NFA, where $\mathcal{B} = \{\boldsymbol{B}_1, \dots, \boldsymbol{B}_r\}$ is a collection of sets of positive atom symbols and $\mathcal{B}_I, \mathcal{B}_F \subseteq \mathcal{B}$. If $\mathcal{B}_k \subseteq \mathcal{B}$, define $S(\mathcal{B}_k) = \bigcup_{\boldsymbol{B}_i \in \mathcal{B}_k} \boldsymbol{B}_i$ to be the set of atom symbols appearing in the sets $\boldsymbol{B}_i$ of $\mathcal{B}_k$. Then $\mathfrak{B}$ is a reduced atomic NFA of L if and only if it satisfies the following conditions:*

(1) $S(\mathcal{B}_I) = \boldsymbol{A}_I$.
(2) For all $\boldsymbol{B}_i \in \mathcal{B}$, $S(\beta(\boldsymbol{B}_i, a)) = \alpha(\boldsymbol{B}_i, a)$.
(3) For all $\boldsymbol{B}_i \in \mathcal{B}$, we have $\boldsymbol{B}_i \in \mathcal{B}_F$ if and only if $\mathbf{A}_{p-1} \in \boldsymbol{B}_i$.

Before proving the theorem, we require the following lemma:

Lemma 7. *If $\mathfrak{B}$ satisfies Condition 2 of Theorem 6, then $S(\beta(\boldsymbol{B}_i, w)) = \alpha(\boldsymbol{B}_i, w)$ for every $\boldsymbol{B}_i \in \mathcal{B}$ and $w \in \Sigma^*$.*

Proof. For $w = \varepsilon$, we have $S(\beta(\boldsymbol{B}_i, \varepsilon)) = S(\boldsymbol{B}_i) = \boldsymbol{B}_i$, and $\alpha(\boldsymbol{B}_i, \varepsilon) = \boldsymbol{B}_i$; so the claim holds for this case.

Assume that $S(\beta(\boldsymbol{B}_i, w)) = \alpha(\boldsymbol{B}_i, w)$ for all $\boldsymbol{B}_i \in \mathcal{B}$ and all $w \in \Sigma^*$ with length less than or equal to $l \geqslant 0$. We prove that $S(\beta(\boldsymbol{B}_i, wa)) = \alpha(\boldsymbol{B}_i, wa)$ for every $a \in \Sigma$. Let $\beta(\boldsymbol{B}_i, w) = \{\boldsymbol{B}_{i_1}, \dots, \boldsymbol{B}_{i_h}\}$ for some $\boldsymbol{B}_{i_1}, \dots, \boldsymbol{B}_{i_h} \in \mathcal{B}$. Since $\beta(\boldsymbol{B}_i, wa) = \beta(\beta(\boldsymbol{B}_i, w), a) = \beta(\boldsymbol{B}_{i_1}, a) \cup \cdots \cup \beta(\boldsymbol{B}_{i_h}, a)$, we have $S(\beta(\boldsymbol{B}_i, wa)) = S(\beta(\boldsymbol{B}_{i_1}, a) \cup \cdots \cup \beta(\boldsymbol{B}_{i_h}, a)) = S(\beta(\boldsymbol{B}_{i_1}, a)) \cup \cdots \cup S(\beta(\boldsymbol{B}_{i_h}, a))$. By Condition 2, the latter is equal to $\alpha(\boldsymbol{B}_{i_1}, a) \cup \cdots \cup \alpha(\boldsymbol{B}_{i_h}, a) = \alpha(\boldsymbol{B}_{i_1} \cup \cdots \cup \boldsymbol{B}_{i_h}, a) = \alpha(S(\beta(\boldsymbol{B}_i, w)), a)$. By the inductive assumption, we get $\alpha(S(\beta(\boldsymbol{B}_i, w)), a) = \alpha(\alpha(\boldsymbol{B}_i, w), a) = \alpha(\boldsymbol{B}_i, wa)$, which proves our claim. □

Proof. (Theorem 6) First we prove that any NFA $\mathfrak{B}$ satisfying Conditions 1–3 is an atomic NFA of L. Let $\boldsymbol{B}_i \in \mathcal{B}$ be a state of $\mathfrak{B}$. If $w \in L_{\boldsymbol{B}_i, \mathcal{B}_F}(\mathfrak{B})$, then by Condition 3, there exists $\boldsymbol{B}_j \in \beta(\boldsymbol{B}_i, w)$ such that $\mathbf{A}_{p-1} \in \boldsymbol{B}_j$, and we have $\mathbf{A}_{p-1} \in S(\beta(\boldsymbol{B}_i, w))$. By Lemma 7, we get $\mathbf{A}_{p-1} \in \alpha(\boldsymbol{B}_i, w)$, implying that there is some $\mathbf{A}_k \in \boldsymbol{B}_i$ such that $w \in L_{\mathbf{A}_k, \{\mathbf{A}_{p-1}\}}(\mathfrak{A})$. Conversely, if $w \in L_{\mathbf{A}_k, \{\mathbf{A}_{p-1}\}}(\mathfrak{A})$ and $\mathbf{A}_k \in \boldsymbol{B}_i$, then $\mathbf{A}_{p-1} \in \alpha(\boldsymbol{B}_i, w) = S(\beta(\boldsymbol{B}_i, w))$. Hence there exists $\boldsymbol{B}_j \in \beta(\boldsymbol{B}_i, w)$ such that $\mathbf{A}_{p-1} \in \boldsymbol{B}_j$. Consequently, every word accepted in $\mathfrak{B}$ from state $\boldsymbol{B}_i$ is in some atom A_k such that $\mathbf{A}_k \in \boldsymbol{B}_i$, and every word in an atom A_k such that $\mathbf{A}_k \in \boldsymbol{B}_i$, is also in $L_{\boldsymbol{B}_i, \mathcal{B}_F}(\mathfrak{B})$. Therefore the right language of $\boldsymbol{B}_i$ in $\mathfrak{B}$ is equal to the union of atoms A_k such that $\mathbf{A}_k \in \boldsymbol{B}_i$. In particular, $L_{\mathcal{B}_I,\mathcal{B}_F}(\mathfrak{B})$ is the union of atoms whose atom symbols appear in the initial collection of $\mathfrak{B}$ which, by Condition 1, is the same as the union of atoms whose atom symbols are initial in $\mathfrak{A}$. But that last union is precisely $L_{\boldsymbol{A}_I,\{\mathbf{A}_{p-1}\}}(\mathfrak{A}) = L$. Since any two sets $\boldsymbol{B}_i$ and $\boldsymbol{B}_j$ are different, and atoms are disjoint, $\mathfrak{B}$ is reduced. Hence $\mathfrak{B}$ is a reduced atomic NFA of L.

Conversely, we show that if $\mathfrak{B}$ is a reduced atomic NFA of L, then it must satisfy Conditions 1–3. We assume that $\mathfrak{B}$ is atomic, that is, for every state $\boldsymbol{B}_i$ of $\mathfrak{B}$, the right language of $\boldsymbol{B}_i$ is equal to the union of atoms A_k such that $\mathbf{A}_k \in \boldsymbol{B}_i$.

For Condition 1, let $\mathbf{A}_j \in S(\mathcal{B}_I)$. Then there is a state $\boldsymbol{B}_k \in \mathcal{B}_I$ such that $\mathbf{A}_j \in \boldsymbol{B}_k$. So for any $w \in A_j$, $w \in L(\mathfrak{B})$. Since $L(\mathfrak{B}) = L(\mathfrak{A})$, we have $w \in L(\mathfrak{A})$ for all $w \in A_j$. Thus $\mathbf{A}_j \in \boldsymbol{A}_I$. Conversely, if $\mathbf{A}_j \in \boldsymbol{A}_I$, then for all $w \in A_j$, $w \in L(\mathfrak{A}) = L(\mathfrak{B})$. Since $\mathfrak{B}$ is atomic, there is an initial state $\boldsymbol{B}_k$ such that $A_j \subseteq L_{\boldsymbol{B}_k,\mathcal{B}_F}(\mathfrak{B})$. Hence $\mathbf{A}_j \in S(\mathcal{B}_I)$.

For Condition 2, if $\mathbf{A}_j \in S(\beta(\boldsymbol{B}_i, a))$, then $L_{\boldsymbol{B}_i,\mathcal{B}_F}(\mathfrak{B})$ must contain aA_j. So there exists some $\mathbf{A}_k \in \boldsymbol{B}_i$ such that $aA_j \subseteq A_k$. Thus $\mathbf{A}_j \in \alpha(\boldsymbol{B}_i, a)$. Conversely, if $\mathbf{A}_j \in \alpha(\boldsymbol{B}_i, a)$, then there is an atom $\mathbf{A}_k \in \boldsymbol{B}_i$ such that $\mathbf{A}_j \in \alpha(\mathbf{A}_k, a)$, implying $aA_j \subseteq A_k$. Since $\mathbf{A}_k \in \boldsymbol{B}_i$, $L_{\boldsymbol{B}_i,\mathcal{B}_F}(\mathfrak{B})$ must contain aA_j. Hence $\mathbf{A}_j \in S(\beta(\boldsymbol{B}_i, a))$.

For Condition 3, we first suppose that $\boldsymbol{B}_i \in \mathcal{B}_F$. Then ε is in the right language of $\boldsymbol{B}_i$. Since $\mathfrak{B}$ is atomic, ε must be in one of the atoms of $\boldsymbol{B}_i$. However, the only atom containing ε is A_{p-1}, so $\mathbf{A}_{p-1} \in \boldsymbol{B}_i$. Conversely, if $\mathbf{A}_{p-1} \in \boldsymbol{B}_i$, then ε is in the right language of $\boldsymbol{B}_i$, and $\boldsymbol{B}_i$ is a final state by definition of an NFA. □

The number of trim reduced atomic NFAs can be very large. There can be such NFAs with as many as $2^p - 1$ non-empty states, since there are that many non-empty sets of positive atoms. In the general case, however, not all sets of positive atom symbols can be states of an atomic NFA. The largest reduced atomic NFA is characterized in the following theorem.

Theorem 8 (Maximal Atomic NFA) *If $\mathcal{B}$ is the collection of all sets $\boldsymbol{B}_i$ such that $\boldsymbol{B}_i$ is a non-empty subset of the set of positive atom symbols $\{\mathbf{A}_h \mid A_h \subseteq K_j\}$ of some quotient K_j of L, then there exists a trim reduced atomic NFA of L with state set $\mathcal{B}$.*

Proof. Let $\mathfrak{B} = (\mathcal{B}, \Sigma, \beta, \mathcal{B}_I, \mathcal{B}_F)$ be an NFA in which the state set $\mathcal{B}$ is the collection of all sets $\boldsymbol{B}_i$ such that $\boldsymbol{B}_i$ is a non-empty subset of the set of atom symbols $\{\mathbf{A}_h \mid A_h \subseteq K_j\}$ of some quotient K_j of L, where $j \in \{0, \dots, n-1\}$, $\beta(\boldsymbol{B}_i, a) = \{\boldsymbol{B}_j \mid \boldsymbol{B}_j \subseteq \alpha(\boldsymbol{B}_i, a)\}$ for every $\boldsymbol{B}_i \in \mathcal{B}$ and $a \in \Sigma$, $\boldsymbol{B}_i \in \mathcal{B}_I$ if and only if $\boldsymbol{B}_i$ is a subset of the set of atom symbols of the initial quotient K_{in}, and $\boldsymbol{B}_i \in \mathcal{B}_F$ if and only if $\mathbf{A}_{p-1} \in \boldsymbol{B}_i$. We claim that $\mathfrak{B}$ is a trim reduced atomic NFA of L.

We show that $\mathfrak{B}$ is trim. Consider any state $\boldsymbol{B}_i$ of $\mathfrak{B}$. Let K_j be a quotient such that $\boldsymbol{B}_i$ is a subset of the set of atom symbols of K_j, and let $\boldsymbol{B}_j$ be the set of atom symbols corresponding to K_j. Let $\boldsymbol{B}_0$ be the set of atom symbols corresponding to the initial quotient K_{in} of L. Note that $\boldsymbol{B}_0 = \boldsymbol{A}_I$. Since every set of atom symbols corresponding to some quotient is reachable from the initial set of atom symbols in the átomaton $\mathfrak{A}$, there must be a word $w \in \Sigma^*$, such that $\boldsymbol{B}_j$ is reachable from $\boldsymbol{B}_0$ by w in $\mathfrak{A}$. We show that $\boldsymbol{B}_i$ is reachable from some initial state of $\mathfrak{B}$ by w. If $w = \varepsilon$, then $K_j = K_{in}$, and since $\boldsymbol{B}_i \subseteq \boldsymbol{B}_j$, it follows that $\boldsymbol{B}_i$ is an initial state of $\mathfrak{B}$ reachable from itself by ε. If $w = ua$ for some $u \in \Sigma^*$ and $a \in \Sigma$, then there is a state $\boldsymbol{B}_u$ of $\mathfrak{B}$, reachable from $\boldsymbol{B}_0$ by u, such that $\boldsymbol{B}_u$ corresponds to the quotient $u^{-1}L$ of L and $\boldsymbol{B}_j = \alpha(\boldsymbol{B}_u, a)$. Since $\boldsymbol{B}_i \subseteq \boldsymbol{B}_j$ and $\boldsymbol{B}_j = \alpha(\boldsymbol{B}_u, a)$, by the definition of β we have $\boldsymbol{B}_i \in \beta(\boldsymbol{B}_u, a)$. Thus, $\boldsymbol{B}_i$ is reachable from $\boldsymbol{B}_0$ in $\mathfrak{B}$ by ua.

We also have to show that there is a word $w \in \Sigma^*$, such that some final state of $\mathfrak{B}$ is reachable from $\boldsymbol{B}_i$ by w. If $\boldsymbol{B}_i$ is final, then it is reachable from itself by $w = \varepsilon$. If $\boldsymbol{B}_i$ is not final, then consider any $\mathbf{A}_k \in \boldsymbol{B}_i$. Since the right language of the state $\mathbf{A}_k$ in the átomaton $\mathfrak{A}$ is not empty, and $\mathbf{A}_k$ cannot be the final state of $\mathfrak{A}$, there must be some state $\mathbf{A}_l$ of $\mathfrak{A}$ and some $a \in \Sigma$, such that $\mathbf{A}_l \in \alpha(\mathbf{A}_k, a)$. Now we know that there is some $\boldsymbol{B}_j$ such that $\mathbf{A}_l \in \boldsymbol{B}_j$ and $\alpha(\boldsymbol{B}_i, a) = \boldsymbol{B}_j$. Since $\beta(\boldsymbol{B}_i, a)$ is the collection of all non-empty subsets of $\boldsymbol{B}_j$, it follows that $\{\mathbf{A}_l\} \in \beta(\boldsymbol{B}_i, a)$. Since the final state $\mathbf{A}_{p-1}$ of $\mathfrak{A}$ is reachable from $\mathbf{A}_l$ by any word $v \in A_l$, we get $\{\mathbf{A}_{p-1}\} \in \beta(\boldsymbol{B}_i, av)$ by the definition of β. So a final state $\{\mathbf{A}_{p-1}\}$ of $\mathfrak{B}$ is reachable from $\boldsymbol{B}_i$ by av. Thus, $\mathfrak{B}$ is trim.

To see that $\mathfrak{B}$ is a reduced atomic NFA, one verifies that Conditions 1–3 of Theorem 6 hold. Thus by Theorem 6, $\mathfrak{B}$ is a trim reduced atomic NFA of L. □

Theorem 9 (NFA with $2^p - 1$ States) *A regular language L has a trim reduced atomic NFA with $2^p - 1$ states if and only if there is some quotient K_i of L, such that $K_i = A_0 \cup \cdots \cup A_{p-1}$.*

Proof. Let $\mathfrak{B} = (\mathcal{B}, \Sigma, \beta, \mathcal{B}_I, \mathcal{B}_F)$ be a trim reduced atomic NFA of L with $2^p - 1$ states. Then there must be a state $\boldsymbol{B}_i$ of $\mathfrak{B}$ such that $\boldsymbol{B}_i = \{\mathbf{A}_0, \dots, \mathbf{A}_{p-1}\}$. Since the right language of any state of a trim NFA is a subset of some quotient, we have $L_{\boldsymbol{B}_i, \mathcal{B}_F}(\mathfrak{B}) = A_0 \cup \cdots \cup A_{p-1} \subseteq K_i$ for some quotient K_i of L. On the other hand, K_i must be a union of some atoms, so we get $K_i = A_0 \cup \cdots \cup A_{p-1}$.

Conversely, let $K_i = A_0 \cup \cdots \cup A_{p-1}$ be a quotient of L which includes all the positive atoms of L. Then by Theorem 8, there is a trim reduced atomic NFA of L in which the state set is the collection of all non-empty subsets of the set of positive atom symbols. This NFA has $2^p - 1$ states. □

A minimal atomic NFA of a language L can possibly be as small as a minimal NFA for L. An example of a language with this property is any language where a minimal DFA is also a minimal NFA as is the case, for instance, of the language L_2 with the minimal DFA $\mathcal{D}_2$ of Example 3.

4. Atom Complexity

The *quotient complexity* [2] of L is the number of quotients of L, and this is the same number as the number of states in the minimal DFA recognizing L; the latter number is known as the *state complexity* [10] of L. Quotient complexity allows us to use language-theoretic methods, whereas state complexity is more amenable to automaton-theoretic techniques. We use one of these two points of view or the other, depending on convenience.

It has been suggested by Brzozowski and Ye [6] that syntactic complexity can be a useful measure of complexity. It has its roots in the *syntactic congruence* $\approx_L$ defined by a language $L \subseteq \Sigma^*$ as follows: For $x, y \in \Sigma^*$,

$$x \approx_L y \text{ if and only if } uxv \in L \Leftrightarrow uyv \in L \text{ for all } u, v \in \Sigma^*.$$

The *syntactic semigroup* of L is the quotient semigroup $\Sigma^+/\approx_L$. *Syntactic complexity* is the cardinality of the syntactic semigroup. This complexity may be able to distinguish two regular languages with the same quotient complexity. For example [6], a language with three quotients may have syntactic complexity as low as 3 or as high as 27. The syntactic semigroup is isomorphic to the semigroup of transformations of the set of states, called the *transition semigroup,* by non-empty words in the minimal DFA of L. The transition semigroup is often used to represent the syntactic semigroup.

Our main result concerns the quotient complexity of atoms of regular languages, which represents yet another complexity measure. We say that a language has *maximal atom complexity* if (a) it has all 2^n atoms, and (b) they all reach their maximal bounds, as stated below.

For $n = 1$, there is only one non-empty language, Σ^*; it has one atom, Σ^*, which has quotient complexity 1. From now on we consider only $n \geqslant 2$.

Theorem 10 (Atom Complexity) *Let $L \subseteq \Sigma^*$ be a non-empty regular language and let its set of quotients be $\mathcal{K} = \{K_0, \dots, K_{n-1}\}$. The quotient complexity of the atoms with 0 or n complemented quotients is less than or equal to $2^n - 1$. For r satisfying $1 \leqslant r \leqslant n-1$, the quotient complexity of any atom of L with r complemented quotients is less than or equal to*

$$f(n,r) = 1 + \sum_{k=1}^{r} \sum_{h=k+1}^{k+n-r} \binom{n}{h}\binom{h}{k}.$$

The atoms of the language L_n of the DFA $\mathfrak{D}_n$ of Fig. 2 meet the bounds given above.

The proof of this result is postponed to Sections 5–8.

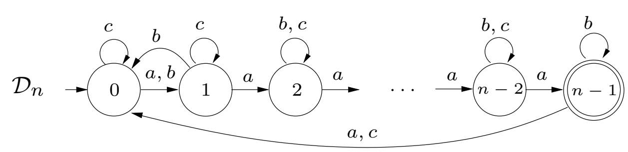

Fig. 2. DFA $\mathfrak{D}_n$ of language L_n whose atoms meet the bounds.

The following relations exist among the three complexity measures [3]:

Theorem 11 (Syntactic Semigroup, Atoms and Reversal) *Maximal syntactic complexity of a regular language implies maximal atom complexity, but the converse is false. Also, maximal atom complexity implies maximal complexity* (2^n) *of reversal, but the converse is false.*

Thus atom complexity defines a new complexity class of regular languages. These results provide additional motivation for studying the complexity of atoms.

5. Upper Bounds on the Quotient Complexities of Atoms

We now derive upper bounds on the quotient complexity of atoms. First we deal with the two atoms that have only uncomplemented or only complemented quotients.

Let $L \subseteq \Sigma^*$ be a non-empty regular language and let its set of quotients be $\mathcal{K} = \{K_0, \ldots, K_{n-1}\}$, with $n \geqslant 2$.

Proposition 12 (Atoms with 0 or n Complemented Quotients) *The quotient complexity of the two atoms* $A_{\mathcal{K}} = K_0 \cap \cdots \cap K_{n-1}$ *and* $A_{\emptyset} = \overline{K_0} \cap \cdots \cap \overline{K_{n-1}}$ *is less than or equal to* $2^n - 1$.

Proof. Each quotient $w^{-1}A_{\mathcal{K}}$ is the intersection of languages $w^{-1}K_i$, which are quotients of L: $w^{-1}A_{\mathcal{K}} = w^{-1}(K_0 \cap \cdots \cap K_{n-1}) = w^{-1}K_0 \cap \cdots \cap w^{-1}K_{n-1}$. Since these quotients of L need not be distinct, $w^{-1}A_{\mathcal{K}}$ may be the intersection of any non-empty subset of quotients of L. Hence $A_{\mathcal{K}}$ can have at most $2^n - 1$ quotients.

The argument for the atom $A_{\emptyset} = \overline{K_0} \cap \cdots \cap \overline{K_{n-1}}$ with n complemented quotients is similar, since $w^{-1}\overline{K_i} = \overline{w^{-1}K_i}$. □

Next, we present an upper bound on the quotient complexity of any atom with at least one and fewer than n complemented quotients.

Proposition 13 (Atoms with r Complemented Quotients, $1 \leqslant r \leqslant n-1$) *For* $1 \leqslant r \leqslant n-1$, *the quotient complexity of any atom with* r *complemented quotients is less than or equal to*

$$f(n,r) = 1 + \sum_{k=1}^{r} \sum_{h=k+1}^{k+n-r} \binom{n}{h}\binom{h}{k}. \tag{3}$$

Proof. Consider an intersection of complemented and uncomplemented quotients that constitutes an atom. Without loss of generality, we arrange the terms in the intersection in such a way that all complemented quotients appear on the right. Thus let $A_i = K_0 \cap \cdots \cap K_{n-r-1} \cap \overline{K_{n-r}} \cap \cdots \cap \overline{K_{n-1}}$ be an atom of L with r complemented quotients, where $1 \leqslant r \leqslant n-1$. The quotient of A_i by any $w \in \Sigma^*$ is

$$\begin{aligned} w^{-1}A_i &= w^{-1}(K_0 \cap \cdots \cap K_{n-r-1} \cap \overline{K_{n-r}} \cap \cdots \cap \overline{K_{n-1}}) \\ &= w^{-1}K_0 \cap \cdots \cap w^{-1}K_{n-r-1} \cap \overline{w^{-1}K_{n-r}} \cap \cdots \cap \overline{w^{-1}K_{n-1}}. \end{aligned}$$

Since each quotient $w^{-1}K_j$ is a quotient, say K_{i_j}, of L, we have

$$w^{-1}A_i = K_{i_0} \cap \cdots \cap K_{i_{n-r-1}} \cap \overline{K_{i_{n-r}}} \cap \cdots \cap \overline{K_{i_{n-1}}}.$$

The cardinality of a set S is denoted by $|S|$. Let the set of distinct quotients of L appearing in $w^{-1}A_i$ uncomplemented (respectively, complemented) be X (respectively, Y), where $1 \leqslant |X| \leqslant n-r$ and $1 \leqslant |Y| \leqslant r$. If $X \cap Y \neq \emptyset$, then $w^{-1}A_i = \emptyset$. Consider now the case where $X \cap Y = \emptyset$, and let $|X \cup Y| = h$, where $2 \leqslant h \leqslant n$; there are $\binom{n}{h}$ such sets $X \cup Y$. Suppose further that $|Y| = k$, where $1 \leqslant k \leqslant r$. There are $\binom{h}{k}$ ways of choosing Y. Hence there are at most $\sum_{h=k+1}^{k+n-r} \binom{n}{h}\binom{h}{k}$ distinct intersections with k complemented quotients. Thus, the total number of intersections of uncomplemented and complemented quotients can be at most $\sum_{k=1}^{r} \sum_{h=k+1}^{k+n-r} \binom{n}{h}\binom{h}{k}$. Adding 1 for the empty quotient of $w^{-1}A_i$, we get the required bound. □

We now consider the properties of the function $f(n,r)$.

Proposition 14 (Properties of Bounds) *For $1 \leqslant r \leqslant n-1$, the function $f(n,r)$ of Equation (3) satisfies the following properties:*

(1) $f(n,r) = f(n,n-r)$.
(2) For a fixed n, the maximal value of $f(n,r)$ occurs when $r = \lfloor n/2 \rfloor$.

Proof. Since $f(n,r) = 1 + \sum_{k=1}^{r} \sum_{h=k+1}^{k+n-r} \binom{n}{h}\binom{h}{k}$, and the following equations hold:

$$\sum_{k=1}^{r} \sum_{h=k+1}^{k+n-r} \binom{n}{h}\binom{h}{k} = \sum_{k=1}^{r} \sum_{l=1}^{n-r} \binom{n}{k+l}\binom{k+l}{k} = \sum_{l=1}^{n-r} \sum_{k=1}^{r} \binom{n}{k+l}\binom{k+l}{k}$$
$$= \sum_{l=1}^{n-r} \sum_{k=1}^{r} \binom{n}{k+l}\binom{k+l}{l} = \sum_{l=1}^{n-r} \sum_{m=l+1}^{l+r} \binom{n}{m}\binom{m}{l},$$

we have $f(n,r) = f(n,n-r)$.

For the second part, we assume that $1 \leqslant r < \lfloor n/2 \rfloor$ holds. We will show that $f(n,r+1) > f(n,r)$ for this case. After some straightforward rewriting we find that $f(n,r+1) - f(n,r)$ is equal to

$$\sum_{h=r+2}^{n-r} \binom{n}{h}\binom{h}{r+1} + \sum_{k=1}^{r} \binom{n}{k+n-r}\binom{k+n-r}{r+1} - \sum_{k=1}^{r} \binom{n}{k+n-r}\binom{k+n-r}{k}.$$

Assuming $1 \leqslant k \leqslant r$, we will show that $\binom{k+n-r}{r+1} > \binom{k+n-r}{k}$. We consider

$$\frac{\binom{k+n-r}{r+1}}{\binom{k+n-r}{k}} = \frac{(k+n-r)!}{(r+1)!(k+n-2r-1)!} \div \frac{(k+n-r)!}{k!(n-r)!}$$
$$= \frac{k!(n-r)\cdots(n-2r+k)(n-2r+k-1)!}{(r+1)\cdots(k+1)k!(n-2r+k-1)!}$$
$$= \frac{n-r}{r+1} \cdot \frac{n-r-1}{r} \cdot \cdots \cdot \frac{n-2r+k}{k+1}.$$

The condition $1 \leqslant r < \lfloor n/2 \rfloor$ implies that $n > 2r+1$; consequently we have

$$n-r > r+1,\ n-r-1 > r,\ \ldots,\ n-2r+k > k+1.$$

Therefore $\binom{k+n-r}{r+1}/\binom{k+n-r}{k} > 1$, which implies that $\binom{k+n-r}{r+1} > \binom{k+n-r}{k}$.

It follows that $\sum_{k=1}^{r} \binom{n}{k+n-r}\binom{k+n-r}{r+1} > \sum_{k=1}^{r} \binom{n}{k+n-r}\binom{k+n-r}{k}$, and $f(n,r+1) - f(n,r) > 0$. So, if $1 \leqslant r < \lfloor n/2 \rfloor$, then $f(n,r+1) > f(n,r)$. Since $f(n,r) = f(n,n-r)$, the maximum of $f(n,r)$ occurs when $r = \lfloor n/2 \rfloor$. □

Some numerical values of $f(n,r)$ are shown in Table 1. The figures in boldface type are the maxima for a fixed n. The row marked *max* shows the maximal quotient complexity of the atoms of L. The row marked *ratio* shows the value of $f(n,\lfloor n/2 \rfloor)/f(n-1,\lfloor (n-1)/2 \rfloor)$, for $n \geqslant 2$.

The following observations are from Volker Diekert (personal communication). For $r = \lfloor n/2 \rfloor$ the difference $3^n - f(n,r)$ grows as $8^{n/2}$. Hence the ratio $f(n,r)/f(n-1,r)$ converges to 3. The ratio oscillates around 3: a combinatorial interpretation shows that for $n \geq 10$, we have $f(n,r)/f(n-1,r) > 3$ if n is even, and $f(n,r)/f(n-1,r) < 3$ if n is odd.

Table 1. Maximal quotient complexity of atoms.

n	1	2	3	4	5	6	7	8	9	10	⋯
r=0	**1**	**3**	7	15	31	63	127	255	511	1,023	⋯
r=1	∗	**3**	**10**	29	76	187	442	1,017	2,296	5,111	⋯
r=2	∗	**3**	**10**	**43**	**141**	406	1,086	2,773	6,859	16,576	⋯
r=3	∗	∗	7	29	**141**	**501**	**1,548**	4,425	12,043	31,681	⋯
r=4	∗	∗	∗	15	76	406	**1,548**	**5,083**	**15,361**	44,071	⋯
r=5	∗	∗	∗	∗	31	187	1,086	4,425	**15,361**	**48,733**	⋯
max	1	3	10	43	141	501	1,548	5,083	15,361	48,733	⋯
ratio	–	3	3.33	4.30	3.28	3.55	3.09	3.28	3.02	3.17	⋯

6. Another Representation of Átomata

The next theorem, a slightly modified version of a result from [1] also discussed in [4], will be used several times.

Theorem 15 (Determinization) *If an NFA $\mathfrak{N}$ has no empty states and $\mathfrak{N}^R$ is deterministic, then $\mathfrak{N}^D$ is minimal.*

We prove[a] that $\mathfrak{A}$ is isomorphic to $\mathfrak{D}^{RDR}$. We deal with the following automata:

(1) Quotient DFA $\mathfrak{D} = (\boldsymbol{K}, \Sigma, \delta, \mathbf{K}_{in}, \boldsymbol{F})$ of L whose states are *quotient symbols*.

[a]It was shown in [4] that the átomaton $\mathfrak{A}$ of L with reachable atoms only is isomorphic to the trim version of $\mathfrak{D}^{RDR}$, where $\mathfrak{D}$ is the quotient DFA of L.

(2) The reverse $\mathfrak{D}^R = (\boldsymbol{K}, \Sigma, \delta^R, \boldsymbol{F}, \{\mathbf{K}_{in}\})$ of $\mathfrak{D}$. The states in $\boldsymbol{K}$ are still *quotient symbols,* but their right languages are no longer quotients of L.
(3) The determinized reverse $\mathfrak{D}^{RD} = (S, \Sigma, \gamma, \boldsymbol{F}, G)$, where $S \subseteq 2^{\boldsymbol{K}}$ and $G = \{S_i \in S \mid \mathbf{K}_{in} \in S_i\}$. The states in S are *sets of quotient symbols*, *i.e.*, subsets of $\boldsymbol{K}$. Since $(\mathfrak{D}^R)^R = \mathfrak{D}$ is deterministic and all of its states are reachable, $\mathfrak{D}^R$ has no empty states. By Theorem 15, DFA $\mathfrak{D}^{RD}$ is minimal and accepts L^R; hence it is isomorphic to the quotient DFA of L^R.
(4) The reverse $\mathfrak{D}^{RDR} = (S, \Sigma, \gamma^R, G, \{\boldsymbol{F}\})$ of $\mathfrak{D}^{RD}$; here the states are still *sets of quotient symbols*.
(5) The átomaton $\mathfrak{A} = (\boldsymbol{A}, \Sigma, \alpha, \boldsymbol{A}_I, \{\mathbf{A}_{p-1}\})$, whose states are *atom symbols*.
(6) The reverse $\mathfrak{A}^R = (\boldsymbol{A}, \Sigma, \alpha^R, \mathbf{A}_{p-1}, \boldsymbol{A}_I)$ of $\mathfrak{A}$, whose states are still *atom symbols,* though their right languages are no longer atoms.

The results from [4] and our new definition of atoms imply that $\mathfrak{A}^R$ is a minimal DFA that accepts L^R. It follows that $\mathfrak{A}^R$ is isomorphic to $\mathfrak{D}^{RD}$. Our next result makes this isomorphism precise.

Proposition 16 (Isomorphism) *Let* $\varphi\colon \boldsymbol{A} \to S$ *be the mapping assigning to state* $\mathbf{A}_j$, *given by* $A_j = K_{i_0} \cap \cdots \cap K_{i_{n-r-1}} \cap \overline{K_{i_{n-r}}} \cap \cdots \cap \overline{K_{i_{n-1}}}$ *of* $\mathfrak{A}^R$, *the set* $\{K_{i_0}, \ldots, K_{i_{n-r-1}}\}$. *Then* φ *is a DFA isomorphism between* $\mathfrak{A}^R$ *and* $\mathfrak{D}^{RD}$.

Proof. The initial state $\mathbf{A}_{p-1}$ of $\mathfrak{A}^R$ is mapped to the set of all quotients containing ε, which is precisely the initial state $\boldsymbol{F}$ of $\mathfrak{D}^{RD}$. Since the quotient L appears uncomplemented in every initial atom $A_i \in \mathcal{I}$, the image $\varphi(\mathbf{A}_i)$ contains L. Thus the set of final states of $\mathfrak{A}^R$ is mapped to the set of final states of $\mathfrak{D}^{RD}$.

It remains to be shown that for all $\mathbf{A}_i, \mathbf{A}_j \in \boldsymbol{A}$ and $a \in \Sigma$, we have $\alpha^R(\mathbf{A}_j, a) = \mathbf{A}_i$ if and only if $\gamma(\varphi(\mathbf{A}_j), a) = \varphi(\mathbf{A}_i)$.

Consider atom A_i with P_i as the set of quotients that appear uncomplemented in A_i. Also define the corresponding set P_j for A_j. If there is a missing quotient K_h in the intersection $a^{-1}A_i$, we use $a^{-1}A_i \cap (K_h \cup \overline{K_h})$. We do this for all missing quotients until we obtain a union of atoms. Hence $\mathbf{A}_j \in \alpha(\mathbf{A}_i, a)$ can hold in $\mathfrak{A}$ if and only if $P_j \supseteq \delta(P_i, a)$ and $P_j \cap \delta(Q \setminus P_i, a) = \emptyset$. It follows that in $\mathfrak{A}^R$ we have $\alpha^R(\mathbf{A}_j, a) = \mathbf{A}_i$ if and only if $P_j \supseteq \delta(P_i, a)$ and $P_j \cap \delta(Q \setminus P_i, a) = \emptyset$.

Now consider $\mathfrak{D}^{RD}$. Let P_i be any subset of Q; then the successor set of P_i in $\mathfrak{D}$ is $\delta(P_i, a)$. Let $\delta(P_i, a) = P_k$. So in $\mathfrak{D}^R$, we have $P_i \in \delta^R(P_k, a)$. But suppose that state q is not in $\delta(Q, a)$; then $\delta^R(q, a) = \emptyset$. Thus we also have $P_i \in \delta^R(P_k \cup \{q\}, a)$. So for any P_j containing $\delta(P_i, a)$ and satisfying $P_j \cap \delta(Q \setminus P_i, a) = \emptyset$, $\gamma(P_j, a) = P_i$.

We have now shown that $\alpha^R(\mathbf{A}_j, a) = \mathbf{A}_i$ if and only if $\gamma(P_j, a) = P_i$, for all subsets $P_i, P_j \in S$, that is, if and only if $\gamma(\varphi(\mathbf{A}_j), a) = \varphi(\mathbf{A}_i)$. □

Corollary 17. *The mapping* φ *is an NFA isomorphism between* $\mathfrak{A}$ *and* $\mathfrak{D}^{RDR}$.

In the remainder of the paper it is more convenient to use the $\mathfrak{D}^{RDR}$ representation of átomata, rather than that of Definition 2.

7. The Witness Languages and Automata

We now formally introduce a class $\{L_n \mid n \geqslant 2\}$ of regular languages defined by the quotient DFAs $\mathfrak{D}_n$ given below; we shall prove that the atoms of each language $L_n = L(\mathfrak{D}_n)$ in this class meet the worst-case quotient complexity bounds.

Definition 18 (Witness) *For $n \geqslant 2$, let $\mathfrak{D}_n = (Q, \Sigma, \delta, q_0, F)$, where $Q = \{0, \dots, n-1\}$, $\Sigma = \{a, b, c\}$, $\delta(i, a) = i+1 \bmod n$, $\delta(0, b) = 1$, $\delta(1, b) = 0$, $\delta(i, b) = i$ for $i > 1$, $\delta(i, c) = i$ for $0 \leqslant i \leqslant n-2$, and $\delta(n-1, c) = 0$, $q_0 = 0$, and $F = \{n-1\}$. Let L_n be the language accepted by $\mathfrak{D}_n$.*

For $n \geqslant 3$, the DFA $\mathcal{D}_n$ of Definition 18 is illustrated in Fig. 2, and $\mathfrak{D}_2$ is the DFA of Example 3 (a and b coincide). DFA $\mathfrak{D}_n$ is minimal, since for $0 \leqslant i \leqslant n-1$, state i accepts a^{n-1-i}, and no other state accepts this word.

A *transformation* of a set Q is a mapping of Q into itself. The set of all transformations of a finite set Q is a semigroup under composition, in fact, a monoid $\mathcal{T}_Q$ of n^n elements. A *permutation* of Q is a mapping of Q *onto* itself. A *transposition,* (i, j), interchanges i and j but does not affect any other element. A *unitary* transformation, $(i \to j)$, changes i to j but does not affect any other element.

The following is well known:

Theorem 19 (Transformations) *The transformation monoid $\mathcal{T}_Q$ can be generated by any cyclic permutation of n elements together with any transposition and any unitary transformation.*

In any DFA $\mathfrak{D} = (Q, \Sigma, \delta, q_0, F)$, each word w in Σ^+ performs a transformation on Q defined by $\delta(\cdot, w)$. The set of all these transformations is the transition semigroup of $\mathfrak{D}$. By Theorem 19, the transition semigroup of our witness $\mathfrak{D}_n$ has n^n elements, since a is a cyclic permutation, b is a transposition and c is a unitary transformation.

The following is a result of Salomaa, Wood and Yu [8] concerning reversal:

Theorem 20 (Transformations and Reversal) *Let $\mathfrak{D}$ be a minimal DFA with $n \geqslant 2$ states accepting a language L. If the transition semigroup of $\mathfrak{D}$ has n^n elements, then the quotient complexity of L^R is 2^n.*

Corollary 21 (Reversal) *For $n \geqslant 2$, the quotient complexity of L_n^R is 2^n.*

Corollary 22 (Number of Atoms of L_n) *The language L_n has 2^n atoms.*

Proof. By Corollary 17, the átomaton of L_n is isomorphic to the reversed quotient DFA of L_n^R. By Corollary 21, the quotient DFA of L_n^R has 2^n states, and so the empty set of states of L_n is reachable in L_n^R. Hence L_n^R has the empty quotient, implying that the intersection of all the complemented quotients of L_n is non-empty, and so L_n has 2^n atoms. □

Proposition 23 (Transitions of the Átomaton) *Let $\mathfrak{D}_n = (Q, \Sigma, \delta, q_0, F)$ be the DFA of Definition 18. The átomaton of $L_n = L(\mathfrak{D}_n)$ is the NFA $\mathfrak{A}_n = (2^Q, \Sigma, \alpha, I, \{n-1\})$, where*

(1) If $S = \{\emptyset\}$, then $\alpha(S, a) = \{\emptyset\}$. Otherwise,
$\alpha(\{s_1, \dots, s_k\}, a) = \{s_1 + 1, \dots, s_k + 1\}$, where the addition is modulo n.
(2) If $\{0, 1\} \cap S = \emptyset$, then

- *(a) $\alpha(S, b) = S$,*
- *(b) $\alpha(\{0\} \cup S, b) = \{1\} \cup S$,*
- *(c) $\alpha(\{1\} \cup S, b) = \{0\} \cup S$,*
- *(d) $\alpha(\{0, 1\} \cup S, b) = \{0, 1\} \cup S$.*

(3) If $\{0, n-1\} \cap S = \emptyset$, then

- *(a) $\alpha(S, c) = \{S, \{n-1\} \cup S\}$,*
- *(b) $\alpha(\{0, n-1\} \cup S, c) = \{\{0, n-1\} \cup S, \{0\} \cup S\}$,*
- *(c) $\alpha(\{0\} \cup S, c) = \emptyset$,*
- *(d) $\alpha(\{n-1\} \cup S, c) = \emptyset$.*

Proof. The reverse of DFA $\mathfrak{D}_n$ is the NFA $\mathfrak{D}_n^R = (Q, \Sigma, \delta^R, \{n-1\}, \{0\})$, where δ^R is defined by $\delta^R(i, a) = i - 1 \bmod n$, $\delta^R(i, b) = \delta(i, b)$, $\delta^R(0, c) = \{0, n-1\}$, $\delta^R(n-1, c) = \emptyset$, and $\delta^R(i, c) = i$, for $0 < i < n-1$. After applying determinization and reversal to $\mathfrak{D}_n^R$, the claims follow by Corollary 17. □

8. Tightness of the Upper Bounds

We now show that the upper bounds derived in Section 5 are tight by proving that the atoms of the language L_n of Definition 18 meet those bounds.

Since the states of any átomaton $\mathfrak{A}_n = (\boldsymbol{A}, \Sigma, \alpha, \boldsymbol{A}_I, \{\mathbf{A}_{p-1}\})$ are atom symbols $\mathbf{A}_i$, and the right language of each $\mathbf{A}_i$ is the atom A_i, the languages A_i are properly represented by the átomaton. Since, however, the átomaton is an NFA, to find the quotient complexity of A_i, we need the equivalent minimal DFA.

Let $\mathfrak{D}_n$ be the n-state quotient DFA of Definition 18 for $n \geqslant 2$, and recall that $L(\mathfrak{D}_n) = L_n$. In the sequel, using Corollary 17, we represent the átomaton $\mathfrak{A}_n$ of L_n by the isomorphic NFA $\mathfrak{D}_n^{RDR} = (S, \Sigma, \gamma^R, G, \{\boldsymbol{F}\})$, and identify the atoms by their sets of uncomplemented quotients. We represent atoms by the subscripts of the quotients, that is, by subsets of $Q = \{0, \dots, n-1\}$, as in Definition 18.

In this framework, to find the quotient complexity of an atom A_P, with $P \subseteq Q$, we start with the NFA $\mathfrak{A}_P = (S, \Sigma, \gamma^R, \{P\}, \{\boldsymbol{F}\})$, which has the same states, transitions, and final state as the átomaton, but has only one initial state P corresponding to the atom symbol $\mathbf{A}_P$. Because $\mathfrak{A}_P^R$ is deterministic and $\mathfrak{A}_P$ has no empty states, $\mathfrak{A}_P^D$ is minimal by Theorem 15. Therefore, $\mathfrak{A}_P^D$ is the quotient DFA of the atom A_P. The states of $\mathfrak{A}_P^D$ are certain *sets of sets* of quotient symbols; to reduce confusion we refer to them as *collections of sets*.

The particular collections appearing in $\mathfrak{A}_P^D$ will be called "intervals". Let U be a subset of Q with $|U| = u$, and let V be a subset of U with $|V| = v$. Define $[V, U]$ to be the collection of all 2^{u-v} subsets of U containing V. There are $\binom{n}{u}\binom{u}{v}$ collections of the form $[V, U]$, because there are $\binom{n}{u}$ ways of choosing U, and for each such choice there are $\binom{u}{v}$ ways of choosing V. The collection $[V, U]$ is called the *interval between* V *and* U. The *type* of an interval $[V, U]$ is the ordered pair (v, u).

The following result is well-known:

Theorem 24 (Permutations) *The symmetric group of size* $n!$ *of all permutations of a set* $Q = \{0, \ldots, n-1\}$ *is generated by any cyclic permutation of* Q *together with any transposition.*

Lemma 25 (Strong-Connectedness of Intervals) *Intervals of the same type are strongly connected by words in* $\{a, b\}^*$.

Proof. Let $[V_1, U_1]$ and $[V_2, U_2]$ be any two intervals of the same type. Arrange the elements of V_1 in increasing order, and do the same for the elements of the sets V_2, $U_1 \setminus V_1$, $U_2 \setminus V_2$, $Q \setminus U_1$, and $Q \setminus U_2$. Let $\pi\colon Q \to Q$ be the mapping that assigns the ith element of V_2 to the ith element of V_1, the ith element of $U_2 \setminus V_2$ to the ith element of $U_1 \setminus V_1$, and the ith element of $Q \setminus U_2$ to the ith element of $Q \setminus U_1$. For any R_1 such that $V_1 \subseteq R_1 \subseteq U_1$, there is a corresponding subset $R_2 = \pi(R_1)$, where $V_2 \subseteq R_2 \subseteq U_2$. Thus π establishes a one-to-one correspondence between the elements of the intervals $[V_1, U_1]$ and $[V_2, U_2]$. Also, π is a permutation of Q and can be performed by a word $w \in \{a, b\}^*$ in $\mathfrak{D}_n$, by Theorem 24. Thus every set R_2 above is reachable from R_1 by w. So $[V_2, U_2]$ is reachable from $[V_1, U_1]$. □

Lemma 26 (Reachability) *Let* $[V, U]$ *be any interval of type* (v, u). *If* $v \geqslant 2$, *then from* $[V, U]$ *we can reach an interval of type* $(v-1, u)$. *If* $u \leqslant n-2$, *then from* $[V, U]$ *we can reach an interval of type* $(v, u+1)$.

Proof. If $v \geqslant 2$, then by Lemma 25, from $[V, U]$ we can reach an interval $[V', U']$ of type (v, u) such that $\{0, n-1\} \subseteq V'$. By input c we reach $[V' \setminus \{n-1\}, U']$ of type $(v-1, u)$. For the second claim, if $u \leqslant n-2$, then by Lemma 25, from $[V, U]$ we can reach an interval $[V', U']$ of type (v, u) such that $\{0, n-1\} \cap V' = \emptyset$. By input c we reach $[V', U' \cup \{n-1\}]$ of type $(v, u+1)$. □

Proposition 27 (Atoms with 0 or n Complemented Quotients) *The quotient complexity of the atoms* A_Q *and* $A_\emptyset$ *of* L_n *is* $2^n - 1$.

Proof. Let $\mathfrak{A}_Q$ ($\mathfrak{A}_\emptyset$) be the modified átomaton with only one initial state, Q ($\emptyset$). By the arguments above, $\mathfrak{A}_Q^D$ ($\mathfrak{A}_\emptyset^D$) is the quotient DFA of A_Q ($A_\emptyset$); hence it suffices to prove the reachability of $2^n - 1$ collections.

For A_Q, the initial state of $\mathfrak{A}_Q^D$ is the collection $\{Q\}$, which is the interval $[Q, Q]$. Now suppose that we have reached an interval of type (v, n). By Lemma 25, we can reach every other interval of type (v, n). If $v \geqslant 2$, then by Lemma 26 we can reach an interval of type $(v-1, n)$. Thus we can reach all intervals $[V, Q]$, one for each non-empty subset V of Q. Since there are at most $2^n - 1$ collections and that many can be reached, no other collection can be reached.

For $A_\emptyset$, the initial state of $\mathfrak{A}_\emptyset^D$ is the empty collection, which is the interval $[\emptyset, \emptyset]$. Now suppose we have reached an interval of type $(0, u)$. By Lemma 25, we can reach every other interval of type $(0, u)$. If $u \leqslant n-2$, then by Lemma 26 we can reach an interval of type $(0, u+1)$. Thus we can reach all intervals $[\emptyset, U]$, one for each non-empty subset U of Q. Since there are at most $2^n - 1$ collections and that many can be reached, no other collection can be reached. Hence the proposition holds. □

Proposition 28 (Tightness) *For $1 \leqslant r \leqslant n-1$, the quotient complexity of any atom of L_n with r complemented quotients is $f(n, r)$.*

Proof. Let A_P be an atom of L_n with $n-r$ uncomplemented quotients, where $1 \leqslant r \leqslant n-1$, that is, let P be the set of subscripts of the uncomplemented quotients. Let $\mathfrak{A}_P$ be the modified átomaton with the initial state P. As discussed above, $\mathfrak{A}_P^D$ is minimal; hence it suffices to prove the reachability of $f(n, r)$ collections.

We start with the interval $[P, P]$ of type $(n-r, n-r)$. By Lemmas 25 and 26, we can now reach all intervals of types

$$\begin{gathered}
(n-r, n-r), (n-r-1, n-r), \ldots, (1, n-r), \\
(n-r, n-r+1), (n-r-1, n-r+1), \ldots, (1, n-r+1), \\
\cdots \\
(n-r, n-1), (n-r-1, n-1), \ldots, (1, n-1).
\end{gathered}$$

Since the number of intervals of type (v, u) is $\binom{n}{u}\binom{u}{v}$, we can reach

$$g(n, r) = \sum_{u=n-r}^{n-1} \sum_{v=1}^{n-r} \binom{n}{u}\binom{u}{v}$$

intervals. Changing the first summation index to $k = n - u$, we get

$$g(n, r) = \sum_{k=1}^{r} \sum_{v=1}^{n-r} \binom{n}{n-k}\binom{n-k}{v}.$$

Note that $\binom{n}{n-k}\binom{n-k}{v} = \binom{n}{k+v}\binom{k+v}{k}$, because

$$\binom{n}{n-k}\binom{n-k}{v} = \frac{n!}{(n-k)!k!} \cdot \frac{(n-k)!}{v!(n-k-v)!} = \frac{n!}{k!v!(n-k-v)!}, \text{ and}$$

$$\binom{n}{k+v}\binom{k+v}{k} = \frac{n!}{(k+v)!(n-k-v)!} \cdot \frac{(k+v)!}{k!v!} = \frac{n!}{(n-k-v)!k!v!}.$$

Now, we can write $g(n,r) = \sum_{k=1}^{r} \sum_{v=1}^{n-r} \binom{n}{k+v}\binom{k+v}{k}$, and changing the second summation index to $h = k + v$, we have

$$g(n,r) = \sum_{k=1}^{r} \sum_{h=k+1}^{k+n-r} \binom{n}{h}\binom{h}{k}.$$

We notice that $g(n,r) = f(n,r) - 1$. From the interval $[V, V]$, where $V = \{0, 1, \ldots, n-r-1\}$, we reach the empty quotient by input c, since V contains 0, but not $n-1$. Since we can reach $f(n,r)$ intervals, no other collection can be reached, and the proposition holds. □

9. Conclusions

The atoms of a regular language L are its building blocks. We characterized atomic NFAs of L. We studied the quotient complexity of the atoms of L as a function of the quotient complexity of L. We computed an upper bound for the quotient complexity of any atom and exhibited languages $\{L_n\}$ whose atoms meet this bound.

Acknowledgments This work was supported by the Natural Sciences and Engineering Research Council of Canada under grant No. OGP0000871, the ERDF funded Estonian Center of Excellence in Computer Science, EXCS, the Estonian Science Foundation grant 7520, and the Estonian Ministry of Education and Research target-financed research theme no. 0140007s12.

References

[1] J. Brzozowski, Canonical regular expressions and minimal state graphs for definite events, *Proc. Symposium on Mathematical Theory of Automata*, *MRI Symposia* **12**, (Polytechnic Institute of Brooklyn, N.Y., 1963), pp. 529–561.

[2] J. Brzozowski, Quotient complexity of regular languages, *J. Autom. Lang. Comb.* **15**(1/2) (2010) 71–89.

[3] J. Brzozowski and G. Davies, Maximal syntactic complexity of regular languages implies maximal quotient complexities of atoms., `http://arxiv.org/abs/1302.3906` (2013).

[4] J. Brzozowski and H. Tamm, Theory of átomata, *DLT 2011*, eds. G. Mauri and A. Leporati *LNCS* **6795**, (Springer, 2011), pp. 105–116.

[5] J. Brzozowski and H. Tamm, Minimal nondeterministic finite automata and atoms of regular languages (2013), `http://arxiv.org/abs/1301.5585`.

[6] J. Brzozowski and Y. Ye, Syntactic complexity of ideal and closed languages, *DLT 2011*, eds. G. Mauri and A. Leporati *LNCS* **6795**, (Springer, 2011), pp. 117–128.

[7] D. Perrin, Finite automata, *Handbook of Theoretical Computer Science*, ed. J. van Leewen, **B** (Elsevier, 1990), pp. 1–57.

[8] A. Salomaa, D. Wood and S. Yu, On the state complexity of reversals of regular languages, *Theoret. Comput. Sci.* **320** (2004) 315–329.

[9] S. Yu, Regular languages, *Handbook of Formal Languages*, eds. G. Rozenberg and A. Salomaa, **1** (Springer, 1997), pp. 41–110.

[10] S. Yu, State complexity of regular languages, *J. Autom. Lang. Comb.* **6** (2001) 221–234.