THEORY OF COMPUTING, Volume 18 (8), 2022, pp. 1–18
www.theoryofcomputing.org

SPECIAL ISSUE: CCC 2018

The Cayley Semigroup Membership Problem

Lukas Fleischer*

Received June 28, 2018; Revised July 25, 2021; Published April 25, 2022

Abstract

The *Cayley semigroup membership problem* asks, given a multiplication table representing a semigroup S, a subset X of S and an element t of S, whether t can be expressed as a product of elements of X. It is well-known that this problem is NL-complete under AC^0-reductions. For groups, the problem can be solved in deterministic $\mathsf{Logspace}$. This raised the question of determining the exact complexity of this variant. Barrington, Kadau, Lange and McKenzie showed that for Abelian groups and for certain solvable groups, the problem is contained in the complexity class FOLL (polynomial-size, $O(\log\log n)$-depth circuits) and they concluded that these variants are not hard, under AC^0 reductions, for any complexity class containing the PARITY language. The more general case of arbitrary groups remained open. In this article, we apply results by Babai and Szemerédi directly to this setting to show that the problem is solvable in qAC^0 (quasi-polynomial size circuits of constant depth with unbounded fan-in). We prove a similar result for commutative semigroups. Combined with the Yao–Håstad circuit lower bound, it follows immediately that Cayley semigroup membership for groups and Cayley semigroup membership for commutative semigroups are not hard, under AC^0

A preliminary version of this paper appeared in the Proceedings of the 33rd Computational Complexity Conference (CCC'18) [17].

*This work was supported by the DFG grant DI 435/5–2.

ACM Classification: F.2.2, F.4.3

AMS Classification: 20M35, 68Q17, 68Q25, 68Q45, 68Q70

Key words and phrases: subsemigroup, multiplication table, generators, completeness, quasi-polynomial-size circuits, FOLL

 DOI: 10.4086/toc.2022.v018a008

reductions, for any class containing PARITY. Moreover, we prove that NL-completeness already holds for the classes of 0-simple semigroups and nilpotent semigroups. Together with our results on groups and commutative semigroups, we prove the existence of a natural class of finite semigroups that generates a variety of finite semigroups with NL-complete Cayley semigroup membership, while the Cayley semigroup membership problem for the class itself is not NL-hard. We also discuss applications of our technique to FOLL and describe varieties for which the Cayley semigroup membership problem is in AC^0.

1 Introduction

Back in 1976, Jones and Laaser studied the complexity of the *generation problem* which is formally defined as follows.

GEN	
Input:	A set G, a binary operation $\circ\colon G \times G \to G$, a set $X \subseteq G$ and an element $t \in G$
Question:	Is t contained in the smallest superset of X closed under $\circ$?

They showed that this problem is P-complete[1] [21], an observation which has since been used in many other P-completeness results. Barrington and McKenzie later studied natural subproblems and connected them to standard subclasses of P [8]. Following [10], the generation problem is also referred to as *Cayley groupoid membership problem*. This terminology stems from the fact that the set G forms a groupoid when equipped with the operation $\circ$ and the objective is to decide whether t belongs to the subgroupoid generated by X. The prefix *Cayley* is due to the representation of finite groupoid by its multiplication table, often also called *Cayley table*.

It is not surprising that imposing further structural properties on the multiplication table affects the complexity of the Cayley groupoid membership problem. For example, if the multiplication table is required to be associative, one obtains the *associative generation problem*, henceforth referred to as *Cayley semigroup membership problem*. This decision problem is NL-complete [22]. We will analyze its complexity when further restricting the semigroups encoded by the input. For a class of finite semigroups $\mathbf{C}$, the *Cayley semigroup membership problem for* $\mathbf{C}$ is defined as follows.

CSM($\mathbf{C}$)	
Input:	The Cayley table of a semigroup $S \in \mathbf{C}$, a set $X \subseteq S$ and an element $t \in S$
Question:	Is t in the subsemigroup of S generated by X?

The motivation for investigating this problem is two-fold. First, there is a direct connection between the Cayley semigroup membership problem and decision problems for regular languages: a language $L \subseteq \Sigma^+$ is regular if and only if there exist a finite semigroup S, a morphism $\varphi\colon \Sigma^+ \to S$ and a set $P \subseteq S$ such that $L = \varphi^{-1}(P)$. Thus, morphisms to finite

[1] In this paper, all completeness/hardness statements are with respect to AC^0 reductions. Even though some of the cited articles only claim completeness under Logspace or NC^1 reductions, their reductions can in fact be implemented in AC^0.

semigroups can be seen as a way of encoding regular languages. For encoding such a semigroup, specifying the multiplication table is a natural choice. Deciding emptiness of a regular language represented by a morphism $\varphi\colon \Sigma^+ \to S$ to a finite semigroup S and a set $P \subseteq S$ boils down to checking whether any of the elements from the set P is contained in the subsemigroup of S generated by the images of the letters of Σ under φ. Conversely, the Cayley semigroup membership problem is a special case of the emptiness problem for regular languages: an element $t \in S$ is contained in the subsemigroup generated by a set $X \subseteq S$ if and only if the language $\varphi^{-1}(P)$ with $\varphi\colon X^+ \to S, x \mapsto x$ and $P = \{t\}$ is non-empty.

Second, we hope to get a better understanding of the connection between algebra and low-level complexity classes included in NL in a fashion similar to the results of [8]. In the past, several intriguing links between so-called *varieties of finite semigroups* and the computational complexity of algebraic problems for such varieties were made. For example, the word problem for a fixed finite semigroup was shown to be in AC^0 if the semigroup is aperiodic, in ACC^0 if the semigroup is solvable and NC^1-complete otherwise [7, 9].

Related work. The first completeness results for the Cayley groupoid membership problem appeared in work by Jones and Laaser [21], and completeness results on the Cayley semigroup membership problem appeared in a paper by Jones, Lien and Laaser [22].

The semigroup membership problem and its restrictions to varieties of finite semigroups was also studied for other encodings of the input, such as matrix semigroups [2, 6, 4] or transformation semigroups [27, 18, 5, 12, 14, 13, 11]. In [6], Babai and Szemerédi introduced the Black Box Group model, and applied it to matrix groups over finite fields. The Black Box Group model also has direct applications in the Cayley table model — however, to the best of our knowledge, this connection has not been investigated prior to the present paper.

Further systematic study of the group membership problem in the Cayley model (CSM($\mathbf{G}$), using our notation) began with a paper by Barrington and McKenzie in 1991 [8]. They observed that the problem is in SymmetricLogspace which has been shown by Reingold in 2008 [26] to be the same as deterministic Logspace, and they suggested it might be complete for deterministic Logspace. However, all attempts to obtain a hardness proof failed (in fact, the conjecture is shown to be false in this paper). There was no progress in a long time until Barrington, Kadau, Lange and McKenzie showed in 2001 [10] that for Abelian groups and certain solvable groups, the problem lies in the complexity class FOLL (decidable by circuits of polynomial size and $O(\log\log n)$ depth) and thus cannot be hard for any complexity class containing PARITY.

The case of arbitrary finite groups remained open partly due to the lack of awareness of the relevance of the early work by Babai and Szemerédi [6]. With this paper we are closing this information gap. We give more details of the connection in Section 4.

Our contributions. We show that the Cayley semigroup membership problem for the variety $\mathbf{G}$ of finite groups and the variety $\mathbf{Com}$ of commutative semigroups are contained in qAC^0 and thus cannot be hard for any class containing PARITY. Our approach heavily relies on the application of techniques from [6] to the Cayley table setting. The key observation is that every element of a group G (or commutative semigroup S) can be computed by an algebraic circuit of size $O(\log^3 |G|)$ (size $O(\log^2 |S|)$, resp.) over any set of generators.

By means of a closer analysis of the technique used by Jones, Lien and Laaser in [22], we also show that the Cayley semigroup membership problem remains NL-complete when restricting the input to 0-simple semigroups or to nilpotent semigroups.

Combining our results, we obtain that the Cayley semigroup membership problem for the class $\mathbf{G} \cup \mathbf{Com}$ is decidable in qAC^0 (and thus not NL-hard) while the Cayley semigroup membership problem for the minimal variety of finite semigroups containing $\mathbf{G} \cup \mathbf{Com}$ is NL-complete.

We discuss the extent to which our approach can be used to establish membership of Cayley semigroup membership variants in the complexity class FOLL. Here, we use an idea based on repeated squaring. This technique generalizes some of the main concepts used in [10]. Finally, we give examples of varieties for which the Cayley semigroup membership problem is in AC^0.

2 Preliminaries

Algebra. A semigroup T is a *subsemigroup* of S if T is a subset of S closed under multiplication. The *direct product* of two semigroups S and T is the Cartesian product $S \times T$ equipped with componentwise multiplication. A semigroup T is a *quotient* of a semigroup S if there exists a surjective morphism $\varphi\colon S \to T$. A semigroup T *divides* a semigroup S if there exists a surjective morphism from a subsemigroup of S onto T.

For every element s of a finite semigroup S, there exist natural numbers $i, p > 0$ such that $s^{i+p} = s^i$. This implies $s^{j+p} = s^j$ for all $j \geqslant i$. In particular, we have $(s^{ip})^2 = s^{ip+ip} = s^{ip}$, which shows that in a finite semigroup, every element has an idempotent power. An element $z \in S$ is a *zero element* if $sz = z = zs$ for all $s \in S$. It is easy to see that every semigroup contains at most one zero element. It is usually denoted by 0.

A *variety of finite semigroups* is a class of finite semigroups that is closed under finite direct products and under taking divisors. Since we are only interested in finite semigroups, we will henceforth use the term *variety* for a variety of finite semigroups. Note that in the literature, such classes of semigroups are often called *pseudovarieties*, as opposed to Birkhoff varieties which are also closed under infinite direct products. The following varieties play an important role in this paper:

- **Ab**, the class of all finite Abelian groups,
- **Com**, the class of all finite commutative semigroups,
- **G**, the class of all finite groups,
- **N**, the class of all finite *nilpotent* semigroups, i. e., finite semigroups S with a zero element 0 such that for all $s \in S$, there exists an integer $e \in \mathbb{N}$ with $s^e = 0$,
- **LI**, the class of all finite *locally trivial* semigroups, i. e., finite semigroups S where $ese = e$ for all elements $s \in S$ and all idempotent elements $e \in S$,

- $\mathbf{LI}_k$, the class of all finite semigroups S that satisfy $x_1 \cdots x_k z y_k \cdots y_1 = x_1 \cdots x_k y_k \cdots y_1$ for all $x_1, \ldots, x_k, y_1, \ldots, y_k, z \in S$.

Note that by definition, the only idempotent element of a nilpotent semigroup is the zero element. For finite semigroups, having exactly one idempotent element which is a zero element actually characterizes nilpotency: for a finite semigroup S with this property and an element $s \in S$, choosing $e \in \mathbb{N}$ such that s^e is idempotent, we obtain $s^e = 0$. From this observation, it follows immediately that every finite nilpotent semigroup is locally trivial, i. e., $\mathbf{N} \subseteq \mathbf{LI}$.

It is easy to verify that every semigroup in $\mathbf{LI}_k$ is locally trivial. Moreover, if S is a finite locally trivial semigroup, then S belongs to $\mathbf{LI}_{|S|}$. Therefore, $\mathbf{LI} = \bigcup_{k \in \mathbb{N}} \mathbf{LI}_k$. The classes $\mathbf{LI}_k$ form an infinite strict hierarchy within $\mathbf{LI}$.

The *join* of two varieties $\mathbf{V}$ and $\mathbf{W}$, denoted by $\mathbf{V} \vee \mathbf{W}$, is the smallest variety containing both $\mathbf{V}$ and $\mathbf{W}$. A semigroup S is 0-*simple* if it contains a zero element 0 and if for each $s \in S \setminus \{0\}$, one has $SsS = S$. The class of finite 0-simple semigroups does not form a variety.

For a comprehensive introduction to the algebraic concepts used in this paper, we refer to the textbooks [20] and [25].

Complexity. We assume familiarity with standard definitions from circuit complexity; see, e. g., [28] for an introduction. We consider *unbounded fan-in Boolean circuits* which consist of unbounded fan-in AND gates, unbounded fan-in OR gates and fan-in-1 NOT gates. The *size* of such a circuit is the total number of AND and OR gates. The *length* of a path in the circuit is the total number of AND and OR gates occurring on the path. The length of the longest path from an input gate to the output gate is the *depth* of the circuit. Note that NOT gates are not counted when measuring the size or depth of a circuit. A function has *quasi-polynomial* growth if it is contained in $2^{O(\log^c n)}$ for some fixed $c \in \mathbb{N}$.

Throughout the paper, we will consider the following unbounded fan-in Boolean circuit families:

- AC^0, languages decidable by circuit families of depth $O(1)$ and polynomial size,
- qAC^0, languages decidable by circuit families of depth $O(1)$ and quasi-polynomial size,
- FOLL, languages decidable by circuit families of depth $O(\log \log n)$ and polynomial size,
- AC^1, languages decidable by circuit families of depth $O(\log n)$ and polynomial size,
- P/poly, languages decidable by circuit families of polynomial size (and unbounded depth).

We will also briefly refer to the complexity classes ACC^0, TC^0, NC^1, $\mathsf{Logspace}$ and NL. It is known that the PARITY function cannot be computed by AC^0, FOLL or qAC^0 circuits. This follows directly from Håstad's and Yao's famous lower bound results [19, 29], which state that the number of Boolean gates required for a depth-d circuit to compute PARITY is exponential in $n^{1/(d-1)}$.

Remark 2.1. We use AC^0, ACC^0, TC^0, NC^1, AC^1, qAC^0, FOLL to refer to the *non-uniform* variants of these complexity classes, even though the same identifiers are sometimes also used to refer to uniform variants in related work. While our proofs also work in the uniform setting, our main results do not require uniformity. Proving that our algorithms can also be implemented as uniform circuits requires introducing the non-standard notion of $\mathsf{DPOLYLOGTIME}$-uniformity and some caution in the proofs. To avoid additional technical details and to keep the proofs short and self-contained, we refrain from doing so.

3 Hardness results

Before looking at parallel algorithms for the Cayley semigroup membership problem, we establish two new NL-hardness results. To this end, we first analyze the construction already used by Jones, Lien and Laaser [22]. It turns out that the semigroups used in their reductions are 0-simple which leads to the following result.

Theorem 3.1. *For a class containing all* 0*-simple semigroups, the Cayley semigroup membership problem is* NL*-complete (under* AC^0 *many-one reductions).*

Proof. To keep the proof self-contained, we briefly describe the reduction from the connectivity problem for directed graphs (henceforth called STConn) to the Cayley semigroup membership problem given in [22].

Let $G = (V, E)$ be a directed graph. We construct a semigroup on the set $S = V \times V \cup \{0\}$ where 0 is a zero element and the multiplication rule for the remaining elements is

$$(v, w) \cdot (x, y) = \begin{cases} (v, y) & \text{if } w = x, \\ 0 & \text{otherwise.} \end{cases}$$

By construction, the subsemigroup of S generated by $E \cup \{(v, v) \mid v \in V\}$ contains an element (s, t) if and only if t is reachable from s in G. To see that the semigroup S is 0-simple, note that for pairs of arbitrary elements $(v, w) \in V \times V$ and $(x, y) \in V \times V$, one has $(x, v)(v, w)(w, y) = (x, y)$, which implies $S(v, w)S = S$. ◻

In order to prove NL-completeness for another common class of semigroups, we use a construction reminiscent of the "layer technique", which is usually used to show that STConn remains NL-complete when the inputs are acyclic graphs.

Theorem 3.2. *CSM*($\mathbf{N}$) *is* NL*-complete (under* AC^0 *many-one reductions).*

Proof. Following the proof of Theorem 3.1, we describe an AC^0 reduction of STConn to CSM($\mathbf{N}$).

Let $G = (V, E)$ be a directed graph with n vertices. We construct a semigroup on the set $S = V \times \{1, \ldots, n-1\} \times V \cup \{0\}$ where 0 is a zero element and the multiplication rule for the remaining elements is

$$(v, i, w) \cdot (x, j, y) = \begin{cases} (v, i + j, y) & \text{if } w = x \text{ and } i + j < n, \\ 0 & \text{otherwise.} \end{cases}$$

The subsemigroup of S generated by the set $\{(v,1,w) \mid v = w \text{ or } (v,w) \in E\}$ contains the element $(s, n-1, t)$ if and only if t is reachable (in less than n steps) from s in G. Clearly, the zero element is the only idempotent in S, so S is nilpotent. Also, it is readily verified that the reduction can be performed by an AC^0 circuit family. □

4 Parallel algorithms for Cayley semigroup membership

In the Black Box Group model introduced by Babai and Szemerédi [6], group elements are encoded by bit strings of uniform length, and group operations (computing products and inverse elements[2]) are performed by an oracle. Babai and Szemerédi showed that subgroup membership is in NP relative to the group oracle. Together with the following observations, it follows that in the Cayley table setting, subgroup membership can be decided in non-deterministic polylogarithmic time in the random-access Turing machine model:

- Without loss of generality, we can assume that there are only logarithmically many generators in the input, since a generating subset of this size can be guessed in non-deterministic polylogarithmic time.
- A single oracle query can be simulated in non-deterministic logarithmic time (non-determinism is only required to compute the inverse of an element).

Since qAC^0 contains all languages decidable in non-deterministic polylogarithmic time, it follows that CSM($\mathbf{G}$) is in qAC^0.

In the remainder of this section, we will give a more self-contained proof of this result, and expand it to other classes of semigroups. We will use algebraic circuits as a succinct representation of elements in an algebraic structure, similarly to the approach taken in [6]. Unlike in usual algebraic circuits, in the context of the Cayley semigroup membership problem, the algebraic structure is not fixed but given as part of the input. We will introduce so-called Cayley circuits to deal with this setting. Since these circuits will be used for the Cayley semigroup membership problem only, we confine ourselves to cases where the algebraic structure is a finite semigroup.

4.1 Cayley circuits

A *Cayley circuit* is a directed acyclic graph with topologically ordered vertices such that each vertex has in-degree 0 or 2. In the following, to avoid technical subtleties when squaring an element, we allow multi-edges. The vertices of a Cayley circuit are called *gates*. The vertices with in-degree 0 are called *input gates* and vertices with in-degree 2 are called *product gates*. Each Cayley circuit also has a designated gate of out-degree 0, called the *output gate*. For simplicity, we assume that the output gate always corresponds to the maximal gate with regard to the

[2]Babai and Szemerédi [6] consider the more general model where multiple strings may encode the same group element; in this case, an oracle to recognize the identity element needs to be added. However, in the present paper we only consider the case of unique encoding.

vertex order. The *size* of a Cayley circuit C, denoted by $|C|$, is the number of gates of C. An *input* to a Cayley circuit C with k input gates consists of a finite semigroup S and elements $x_1, \dots, x_k$ of S. Given such an input, the *value* of the i-th input gate is x_i and the value of a product gate, whose predecessors have values x and y, is the product $x \cdot y$ in S. The *value of the circuit* C is the value of its output gate. We will denote the value of C under a finite semigroup S and elements $x_1, \dots, x_k \in S$ by $C(S, x_1, \dots, x_k)$.

A Cayley circuit can be seen as a circuit in the usual sense: the finite semigroup S and the input gate values are given as part of the input and the functions computed by product gates map a tuple, consisting of semigroup S and two elements of S, to another element of S. We say that a Cayley circuit with k input gates can be *simulated* by a family of unbounded fan-in Boolean circuits $(C_n)_{n\in\mathbb{N}}$ if, given the encodings of a finite semigroup S and of elements $x_1, \dots, x_k$ of S of total length n, the circuit C_n computes the encoding of $C(S, x_1, \dots, x_k)$. For a semigroup S with N elements, we assume that the elements of S are encoded by the integers $\{0, \dots, N-1\}$ such that the encoding of a single element uses $\lceil \log N \rceil$ bits. The semigroup itself is given as a multiplication table with N^2 entries of $\lceil \log N \rceil$ bits each.

Proposition 4.1. *Let C be a Cayley circuit of size m. Then, C can be simulated by a family of unbounded fan Boolean circuits $(C_n)_{n\in\mathbb{N}}$ of depth 2 and size at most n^m.*

Proof. Let C be a Cayley circuit with k input gates and $m-k$ product gates. We want to construct a Boolean circuit that can be used for all finite semigroups S with a fixed number of elements N. The input to such a circuit consists of $n = (N^2 + k)\lceil \log N \rceil$ bits.

For a fixed vector $(y_1, \dots, y_m) \in S^m$, one can check using a single AND gate (and additional NOT gates at some of the incoming wires) whether $(y_1, \dots, y_m)$ corresponds to the sequence of values occurring at the gates of C under the given inputs. To this end, for each gate $i \in \{1, \dots, m\}$ of C, we add $\lceil \log N \rceil$ incoming wires to this AND gate: if the i-th gate of C is an input gate, we feed the bits of the corresponding input value into the AND gate, complementing the j-th bit if the j-th bit of y_i is zero. If the i-th gate is a product gate and has incoming wires from gates ℓ and r, we connect the entry (y_ℓ, y_r) of the multiplication table to the AND gate, again complementing bits corresponding to 0-bits of y_i.

To obtain a Boolean circuit simulating C, we put such AND gates for all vectors of the form $(y_1, \dots, y_m) \in S^m$ in parallel. In a second layer, we create $\lceil \log N \rceil$ OR gates and connect the AND gate for a vector $(y_1, \dots, y_m)$ to the j-th OR gate if and only if the j-th bit of y_m is one. The idea is that exactly one of the AND gates — the gate corresponding to the vector of correct guesses of the gate values of C — evaluates to 1 and the corresponding output value y_m then occurs as output value of the OR gates. This circuit has depth 2 and size $N^m + \lceil \log N \rceil \leqslant n^m$. □

4.2 The polylogarithmic circuits property

When analyzing the complexity of CSM(**Ab**), Barrington *et al.* introduced the so-called *logarithmic power basis property*. A class of semigroups has the logarithmic power basis property if every set X of generators for a semigroup S of cardinality N from the family has the property that every element of S can be written as a product of at most $\log(N)$ many powers of elements of X. In [10], it was shown that the class of Abelian groups has the logarithmic power basis property. Using a different technique, this result can easily be extended to arbitrary commutative semigroups.

Lemma 4.2. *The variety* **Com** *has the logarithmic power basis property.*

Proof. Suppose that S is a commutative semigroup of size N and let X be a set of generators for S. Let $y \in S$ be an arbitrary element. We choose $k \in \mathbb{N}$ to be the smallest value such that there exist elements $x_1, \dots, x_k \in X$ and integers $i_1, \dots, i_k \in \mathbb{N}$ with $y = x_1^{i_1} \cdots x_k^{i_k}$. Assume, for the sake of contradiction, that $k > \log(N)$.

The power set $\mathcal{P}(\{1, \dots, k\})$ forms a semigroup when equipped with set union as binary operation. Consider the morphism $h \colon \mathcal{P}(\{1, \dots, k\}) \to S$ defined by $h(\{j\}) = x_j^{i_j}$ for all $j \in \{1, \dots, k\}$. This morphism is well-defined because S is commutative.

Since $|\mathcal{P}(\{1, \dots, k\})| = 2^k > 2^{\log(N)} = |S|$, we know by the pigeonhole principle that there exist two sets $K_1, K_2 \subseteq \{1, \dots, k\}$ with $K_1 \neq K_2$ and $h(K_1) = h(K_2)$. We may assume, without loss of generality, that there exists some $j \in K_1 \setminus K_2$. Now, because

$$y = h(\{1, \dots, k\}) = h(K_1)h(\{1, \dots, k\} \setminus K_1) = h(K_2)h(\{1, \dots, k\} \setminus K_1)$$

and since neither K_2 nor $\{1, \dots, k\} \setminus K_1$ contain j, we know that y can be written as a product of powers of elements x_i with $1 \leqslant i \leqslant k$ and $i \neq j$, contradicting the choice of k. □

For the analysis of arbitrary groups, we introduce a more general concept. It is based on the idea that algebraic circuits (Cayley circuits with fixed inputs) can be used for succinct representations of semigroup elements.

Example 4.3 (repeated squaring)**.** Let $e \in \mathbb{N}$ be a positive integer. Then, one can construct a Cayley circuit of size at most $2\lceil \log e \rceil$ which computes, given a finite semigroup S and an element $x \in S$ as input, the power x^e in S. If $e = 1$, the circuit only consists of the input gate. If e is even, the circuit is obtained by taking the circuit for $e/2$, adding a product gate and creating two edges from the output gate of the circuit for $e/2$ to the new gate. If e is odd, the circuit is obtained by taking the circuit for $e - 1$ and connecting it to a new product gate. In this case, the second incoming edge for the new gate comes from the input gate.

A class of semigroups has the *polylogarithmic circuits property* if there exists a constant $c \in \mathbb{N}$ such that for each semigroup S of cardinality N from the class, for each subset X of S and for each y in the subsemigroup generated by X, there exists a Cayley circuit C of size $\log^c(N)$ with k input gates and there exist $x_1, \dots, x_k \in X$ such that $C(S, x_1, \dots, x_k) = y$.

For classes closed under taking subsemigroups, such as varieties of finite semigroups, this is equivalent to saying that each element y of a semigroup of cardinality N can be represented by a Cayley circuit of size $\log^c(N)$ over any set of generators. Alternatively, the polylogarithmic circuits property can be defined in terms of *straight-line programs*; this connection will be used further below.

Proposition 4.4. *Let* **C** *be a family of semigroups that is closed under subsemigroups and has the logarithmic power basis property. Then* **C** *has the polylogarithmic circuits property.*

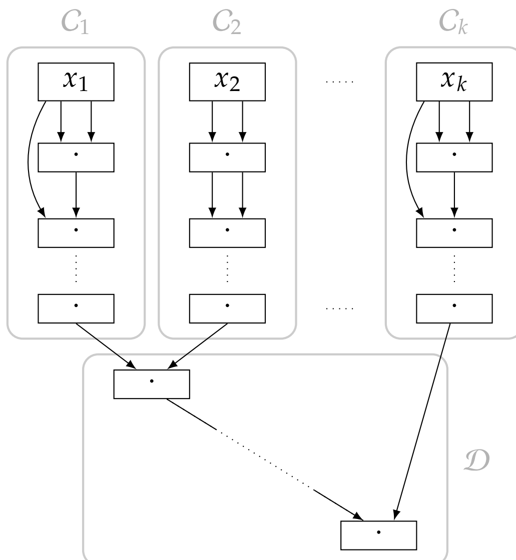

Figure 1: The Cayley circuit C from Proposition 4.4

Proof. Let X be a subset of a semigroup S of cardinality N. Let y be in the subsemigroup generated by X. Then, we have $y = x_1^{i_1} \cdots x_k^{i_k}$ for some $x_1, \ldots, x_k \in X$ with $k \leqslant \log(N)$ and $i_1, \ldots, i_k \in \mathbb{N}$. By the pigeonhole principle, we may assume without loss of generality that $1 \leqslant i_1, \ldots, i_k \leqslant N$. Using the method from Example 4.3, one can construct Cayley circuits $C_1, \ldots, C_k$ of size at most $2\lceil \log N \rceil$ such that $C_j(S, x) = x^{i_j}$ for all $j \in \{1, \ldots, k\}$ and $x \in S$. Using $k-1$ additional product gates $\mathcal{D}$, these circuits can be combined to a single circuit C with $C(S, x_1, \ldots, x_k) = x_1^{i_1} \cdots x_k^{i_k} = y$. The construction is depicted in Figure 1.

In total, the resulting circuit consists of $k \cdot 2\lceil \log N \rceil + k - 1 < 5 \log^2(N)$ gates. □

Let G be a finite group and let X be a subset of G. A sequence $(g_1, \ldots, g_\ell)$ of elements of G is a *straight-line program over* X if for each $i \in \{1, \ldots, \ell\}$, we have $g_i \in X$ or $g_i = g_p^{-1}$ or $g_i = g_p g_q$ for some $p, q < i$. The number ℓ is the *length* of the straight-line program and the elements of the sequence are said to be *generated* by the straight-line program. The following result by Babai and Szemerédi [6] is commonly known as the *Reachability Lemma*.

Lemma 4.5 (Reachability Lemma)**.** *Let G be a finite group and let X be a set of generators of G. Then, for each element $t \in G$, there exists a straight-line program over X generating t which has length at most* $(\log |G| + 1)^2$.

The proof of this lemma is based on a technique called "cube doubling". For details, we refer to [3]. It is now easy to see that groups admit polylogarithmic circuits.

Lemma 4.6. *The variety* **G** *has the polylogarithmic circuits property.*

Proof. Let G be a group of order N, let X be a subset of G and let y be an element in the subgroup of G generated by X. By Lemma 4.5, we know that there exists a straight-line program

$(g_1, \dots, g_\ell)$ over X with $\ell \leqslant (\log(N) + 1)^2$ and $g_\ell = y$. We may assume that the elements $g_1, \dots, g_\ell$ are pairwise distinct. It suffices to describe how to convert this straight-line program into a Cayley circuit C and values $x_1, \dots, x_k \in X$ such that $C(S, x_1, \dots, x_k) = y$.

We start with an empty circuit and with $k = 0$ and process the elements of the straight-line program left to right. For each element g_i, we add gates to the circuit. The output gate of the circuit obtained after processing the element g_i will be called the *g_i-gate*.

If the current element g_i is contained in X, we increment k, add a new input gate to the circuit and let $x_k = g_i$. If the current element g_i can be written as a product $g_p g_q$ with $p, q < i$, we add a new product gate to the circuit and connect the g_p-gate as well as the g_q-gate to this new gate. If the current element g_i is an inverse g_p^{-1} with $p < i$, we take a circuit C' with $2\lceil \log N \rceil$ gates and with $C'(G, x) = x^{N-1}$ for all $x \in S$. Such a circuit can be built by using the powering technique illustrated in Example 4.3. We add C' to C, replacing its input gate by an edge coming from the g_p-gate.

The resulting circuit has size at most $(\log(N) + 1)^2 \cdot 2\lceil \log N \rceil \leqslant 2(\log(N) + 1)^3$. □

We will now show that for classes of semigroups with the polylogarithmic circuits property, one can solve the Cayley semigroup membership problem in qAC^0.

Theorem 4.7. *Let* $\mathbf{C}$ *be a class of semigroups with the polylogarithmic circuits property. Then CSM*$(\mathbf{C})$ *is in* qAC^0*.*

Proof. We construct a family of unbounded fan-in constant-depth Boolean circuits with quasi-polynomial size, deciding, given the multiplication table of a semigroup $S \in \mathbf{C}$, a set $X \subseteq S$ and an element $t \in S$ as inputs, whether t is in the subsemigroup generated by X.

Since $\mathbf{C}$ has the polylogarithmic circuits property, we know that, for some constant $c \in \mathbb{N}$, the element t is in the subsemigroup generated by X if and only if there exist a Cayley circuit C of size $\log^c(n)$ and inputs $x_1, \dots, x_k \in X$ such that $C(S, x_1, \dots, x_k) = t$. There are at most $(\log^c(n) \cdot \log^c(n))^{\log^c(n)} = 2^{2c \log^c(n) \log\log(n)}$ different Cayley circuits of this size. Let us consider one of these Cayley circuits C. Suppose that C has k input gates. By Proposition 4.1, there exists an unbounded fan-in constant-depth Boolean circuit of size $n^{\log^c n} = 2^{\log^{c+1} n}$ deciding on input S and elements $x_1, \dots, x_k \in S$ whether $C(S, x_1, \dots, x_k) = t$. There are at most $n^k \leqslant n^{\log^c n} = 2^{\log^{c+1} n}$ possibilities of connecting (not necessarily all) input gates corresponding to the elements of X to this simulation circuit.

Thus, we can check for all Cayley circuits of the given size and all possible input assignments in parallel, whether the value of the corresponding circuit is t, and feed the results of all these checks into a single OR gate to obtain a quasi-polynomial-size Boolean circuit. □

In conjunction with Lemma 4.2 and Lemma 4.6, we immediately obtain the following corollary.

Corollary 4.8. *Both CSM*$(\mathbf{G})$ *and CSM*$(\mathbf{Com})$ *are contained in* qAC^0*.*

As stated in the preliminaries, problems in qAC^0 cannot be hard for any complexity class containing PARITY. Thus, we also obtain the following statement.

Corollary 4.9. *Let* **C** *be a class of semigroups with the polylogarithmic circuits property, such as the variety of finite groups* **G** *or the variety of finite commutative semigroups* **Com**. *Then CSM*(**C**) *is not hard, under* AC^0 *reductions, for any complexity class containing* PARITY, *such as* ACC^0, TC^0, NC^1, Logspace *or* NL.

4.3 Connections to FOLL

In a first attempt to solve outstanding complexity questions related to the Cayley semigroup membership problem, Barrington *et al.* introduced the complexity class FOLL. The approach presented in the present paper is quite different. This raises the question of whether our techniques can be used to design FOLL-algorithms for Cayley semigroup membership. Note that FOLL and qAC^0 are known to be incomparable, so we cannot use generic results from complexity theory to simulate qAC^0 circuits using families of FOLL circuits or vice versa. The direction $\mathsf{FOLL} \not\subseteq \mathsf{qAC}^0$ follows from bounds on the average sensitivity of bounded-depth circuits (Boppana [15]); using these bounds, one can show that there exists a padded version of the PARITY function that can be computed by a FOLL circuit family and cannot be computed by any qAC^0 circuit family. Conversely, each subset of $\{0,1\}^n$ of cardinality at most $n^{\log n}$ is decidable by a depth-2 circuit of size $n^{1+\log n}+1$, but for each fixed $k \in \mathbb{N}$, there is some large value $n \geqslant 1$ such that the number of such subsets exceeds the number of different circuits of size n^k. This shows that there exist languages in qAC^0 that are not contained in $\mathsf{P/poly} \supseteq \mathsf{FOLL}$.

Designing an FOLL-algorithm that works for arbitrary classes of semigroups with the polylogarithmic circuits property seems difficult. However, for certain special cases, there is an interesting approach, based on the repeated squaring technique. We first give an interpretation of the *Double-Barrelled Recursive Strategy* from [10] in the Cayley circuit setting.

Suppose we are given a cyclic group of large order N, generated by the element x, and some integer $e \in \{1,\dots,N\}$. Let $\ell = \lceil \log e \rceil$ be the length of the binary representation of e. The element x^e can be computed by a repeated squaring Cayley circuit as described in Example 4.3. These circuits only use two different "types" of product gates: gates squaring the current intermediate result and gates multiplying the intermediate result by the generator x. When viewed as operations on the exponent of x, the first gate type performs a left shift of the exponent by 1 bit whereas the second gate type toggles the last bit of an even exponent. Clearly, the integer e can be generated by a sequence of at most 2ℓ of these operations. The idea of the Double-Barrelled Recursive Strategy is that, instead of performing these shift-toggle operations on the exponent e sequentially, we can split its binary representation into two parts of roughly the same size. This yields values e_1 and e_2 with $\lceil \ell/2 \rceil$ bits each such that $e = e_1 \cdot 2^{\lceil \ell/2 \rceil} + e_2$. The value e_1 can be guessed. Then, we recursively run the same procedure to confirm that e_1 can be obtained from x by a sequence of ℓ operations and that e can be obtained from e_1 in the same way. In each recursion step, the number of required operations is halved. Therefore, the recursion depth is $\log(2\ell) \in \mathcal{O}(\log\log N)$.

In the Cayley circuit, this strategy corresponds to dividing the circuit into two parts of roughly equal size and handling the two parts recursively. This idea also works whenever the gates of a Cayley circuit can be ordered in a way such that the number of gate values produced by the first i gates and reused by the remaining gates is bounded by a constant. This property is formalized and used to describe a more general FOLL algorithm below.

For a Cayley circuit, the *width* of a topological ordering $(v_1, \ldots, v_m)$ of the gates is the smallest number $w \in \mathbb{N}$ such that for each $i \in \{1, \ldots, m-1\}$, at most w product gates from the set $A_i = \{v_1, \ldots, v_i\}$ are connected to gates in $B_i = \{v_{i+1}, \ldots, v_m\}$. Let C_i be the set of product gates belonging to A_i that are connected to gates in B_i. The subcircuit induced by A_i can be interpreted as a Cayley circuit computing multiple output values C_i. The subcircuit induced by B_i can be seen as a circuit which, in addition to the input gates of the original circuit, uses the gates from C_i as input gates. The *width* of a Cayley circuit is the smallest width of a topological ordering of its gates. Let us fix some width $w \in \mathbb{N}$.

We introduce a predicate $P(z_1, \ldots, z_w, y_1, \ldots, y_w, i)$ which is true if there exists a Cayley circuit of width at most w and size at most 2^i with w additional input gates and w additional *passthrough gates* (which have in-degree 1 and replicate the value of their predecessors), such that the elements $y_1, \ldots, y_w \in S$ occur as values of the passthrough gates when using $z_1, \ldots, z_w \in S$ as values for the additional input gates and using any subset of the original inputs X as values for the remaining input gates. The additional input gates (or passthrough gates) are not counted when measuring the circuit size but are considered as product gates when measuring width and they have to be the first (last, resp.) gates in all topological orderings considered for width measurement. For each fixed i, there are only n^{2w} such predicates.

The truth value of a predicate with $i = 0$ can be computed by a constant-depth unbounded fan-in Boolean circuit of polynomial size. This is achieved by computing all binary products of the elements $z_1, \ldots, z_w$ and elements of the input set X. For $i \geqslant 1$, the predicate $P(z_1, \ldots, z_w, y_1, \ldots, y_w, i)$ is true if and only if there exist $z'_1, \ldots, z'_w \in S$ such that both $P(z_1, \ldots, z_w, z'_1, \ldots, z'_w, i-1)$ and $P(z'_1, \ldots, z'_w, y_1, \ldots, y_w, i-1)$ are true. Having the truth values of all tuples for $i-1$ at hand, this can be checked with a polynomial number of gates in constant depth because there are only n^w different vectors $(z'_1, \ldots, z'_w) \in S^w$.

For a class of semigroups $\mathbf{C}$ with Cayley circuits of bounded width and polylogarithmic size, we obtain a circuit family of depth $\mathcal{O}(\log\log n)$ deciding CSM($\mathbf{C}$): the predicates are computed for increasing values of i, until i exceeds the logarithm of an upper bound for the Cayley circuit size and then, we return $P(x, \ldots, x, t, \ldots, t, i)$ for the element t given in the input and for an arbitrary element $x \in X$. The number of repetitions of both x and t in $P(x, \ldots, x, t, \ldots, t, i)$ is w.

One example of Cayley circuits of bounded width are the circuits constructed in the proof of Proposition 4.4. Recall that those circuits consist of subcircuits $C_1, \ldots, C_k$ and additional product gates $\mathcal{D}$. Let d_2 denote the gate computing the product of the output values of C_1 and C_2. For $j \in \{3, \ldots, k\}$, let d_j denote the gate computing the product of d_{j-1} and the output value of C_j. Now consider the topological ordering with all gates from C_i preceding all gates from C_j for $i < j$ and with each of the additional multiplication gates from $\mathcal{D}$ as early as possible, i. e., the sequence starts with the gates from C_1, followed by $C_2, d_2, \ldots, C_k, d_k$. This ordering has width at most 2. In particular, we obtain a self-contained proof of the following result.

Theorem 4.10. *Let $\mathbf{C}$ be a class of semigroups that is closed under taking subsemigroups and has the logarithmic power basis property. Then CSM($\mathbf{C}$) is in* FOLL.

By Lemma 4.2, we obtain the following corollary.

Corollary 4.11. *CSM*(**Com**) *is contained in* FOLL.

4.4 The complexity landscape of Cayley semigroup membership

Little is known about when there are algorithms more efficient than the qAC^0 or FOLL upper bounds given in the previous sections. We will now describe an interesting special case for which the Cayley semigroup membership problem is in AC^0.

Theorem 4.12. *For each $k \geqslant 1$, CSM($\mathbf{LI}_k$) is contained in* AC^0.

Proof. Let $k \in \mathbb{N}$ be fixed. Then, for a given input set X of cardinality at most N, there are at most $|X|^{2k} \leqslant N^{2k}$ different products of the form $x_1 \cdots x_k y_k \cdots y_1$ with $x_1, \dots, x_k, y_1, \dots, y_k \in X$. By the definition of $\mathbf{LI}_k$, the element t belongs to the subsemigroup of S generated by X if and only if it is equal to one of these products. We can compute all these products with polynomially many gates in constant depth. Then, we compare each of the results with t. ◻

We recall that the union $\bigcup_{k \in \mathbb{N}} \mathbf{LI}_k$ is the variety of all locally trivial semigroups, which is known to properly contain $\mathbf{N}$. Thus, CSM($\bigcup_{k \in \mathbb{N}} \mathbf{LI}_k$) is NL-complete by Theorem 3.2. This implies that there is no class of finite semigroups that covers all (and only those) varieties of finite semigroups for which Cayley Semigroup Membership is in AC^0. If $\mathbf{C}$ is *any* class containing all varieties $\mathbf{V}$ with CSM($\mathbf{V}$) $\in \mathsf{AC}^0$, then CSM($\mathbf{C}$) is as hard as in the general case.

The previous construction strongly relies on the fact that the classes $\mathbf{LI}_k$ contain only semigroups without a neutral element. However, a slightly weaker statement also holds for varieties of finite monoids. It relies on the following result, which can be seen as a consequence of [1] and the fact that the zero element in a semigroup is always central. For completeness, we provide a short and self-contained proof.

Proposition 4.13. *The variety* $\mathbf{N}$ *is included in* $\mathbf{G} \vee \mathbf{Com}$.

Proof. We show that every finite nilpotent semigroup divides a direct product of a finite group and a finite commutative semigroup. Note that in a finite nilpotent semigroup S, there exists an integer $e \geqslant 0$ such that for each $x \in S$, the power x^e is the zero element. Let $T = \{1, \dots, e\}$ be the commutative semigroup with the product of two elements i and j defined as $\min\{i + j, e\}$.

Let X be the set of non-zero elements of S and let $F(X)$ be the free group over X. For an element $w \in F(X)$, we use $\overline{w}$ to denote its inverse. We use $|w|$ to denote the length of the (freely reduced) normal form of w. Since $F(X)$ is residually finite [23, 24], for each $w \in F(X) \setminus \{1\}$, there exist a finite group G_w and morphism $\psi_w \colon F(X) \to G_w$ such that $\psi_w(w) \neq 1$. Let G be the direct product of all groups G_w for $|w| < 2e - 1$ and let $\psi \colon F(X) \to G$ be the product morphism of the corresponding morphisms ψ_w. Note that for $u, v \in X^*$ with $|u|, |v| < e$, we have $\psi(u) \neq \psi(v)$: if $\psi(u)$ were equal to $\psi(v)$, we would have $\psi(u\overline{v}) = 1$, and thus $\psi_{u\overline{v}}(u\overline{v}) = 1$, contradicting the choice of $\psi_{u\overline{v}}$.

Let U be the subsemigroup of $G \times T$ generated by $\{(\psi(x), 1) \mid x \in X\}$. Now, we define a mapping $\varphi \colon U \to S$ as follows. Each element of the form (g, e) is mapped to zero. For every $(g, \ell) \in U$ with $\ell < e$, there exists, by choice of ψ and by the definition of U, a unique factorization

$g = \psi(x_1 \cdots x_\ell)$ with $x_1, \ldots, x_\ell \in X$. We map (g, ℓ) to the product $x_1 \cdots x_\ell$ evaluated in S. It is straightforward to verify that φ is a surjective morphism and thus, S is a quotient of U. $\square$

Corollary 4.14. *There exist two varieties of finite monoids* $\mathbf{V}$ *and* $\mathbf{W}$ *such that both CSM*$(\mathbf{V})$ *and CSM*$(\mathbf{W})$ *are contained in* qAC^0 *(and thus not hard for any class containing* PARITY*) but CSM*$(\mathbf{V} \vee \mathbf{W})$ *is* NL*-complete.*

The corollary is a direct consequence of the previous proposition, Corollary 4.8 and Theorem 3.2. As was observed in [10] already, Cayley semigroup problems seem to have "strange complexity". The results in this section make this intuition more concrete and suggest that it is difficult to find "nice" descriptions of maximal classes of semigroups for which the Cayley semigroup membership problem is easier than any NL-complete problem.

5 Summary and outlook

We provided new insights into the complexity of the Cayley semigroup membership problem for classes of finite semigroups, giving parallel algorithms for the variety of finite commutative semigroups and the variety of finite groups. We also showed that a maximal class of semigroups with Cayley semigroup membership decidable by qAC^0 circuits does not form a variety. Afterwards, we discussed applicability to FOLL and gave examples of classes for which the problem is in AC^0.

It is tempting to ask whether one can find nice connections between algebra and the complexity of the Cayley semigroup membership problem by conducting a more fine-grained analysis. Does the maximal class of finite semigroups, for which the Cayley semigroup membership problem is in AC^0, form a variety of finite semigroups? Is it possible to show that AC^0 does not contain CSM($\mathbf{G}$) or CSM($\mathbf{Com}$)? Potential approaches to tackling the latter question are reducing small distance connectivity for paths of non-constant length [16] to CSM($\mathbf{G}$) or developing a suitable switching lemma. Another related question is whether there exist classes of semigroups for which the Cayley semigroup membership problem cannot be NL-hard but, at the same time, is not contained within qAC^0.

Moreover, it would be interesting to see whether the Cayley semigroup membership problem can be shown to be in FOLL for all classes of semigroups with the polylogarithmic circuits property. More generally, investigating the relationship of FOLL and qAC^0 to other complexity classes remains an interesting subject for future research.

Finally, we wish to point out that while we have shown that the Cayley semigroup membership problem for groups is not hard for Logspace under AC^0 reductions, it remains open whether it is Logspace-complete under NC^1 reductions.

Acknowledgements. I would like to thank Armin Weiß, Samuel Schlesinger, and Madhur Tulsiani for discussions and comments that led to significant improvements of the presentation of the paper. Special thanks to Armin Weiß for pointing out that qAC^0 is not contained within P/poly, to Samuel Schlesinger for pointing out that results on the average sensitivity of bounded-depth circuits can be used to show that FOLL is not contained within qAC^0, and to Madhur

Tulsiani for pointing out a direct application of the Black Box Group model in the Cayley table setting. Moreover, I am grateful to the anonymous referees of both the conference and the full version of this paper for providing helpful comments.

References

[1] JORGE ALMEIDA: Some pseudovariety joins involving the pseudovariety of finite groups. *Semigroup Forum*, 37(1):53–57, 1988. [doi:10.1007/BF02573123] 14

[2] LÁSZLÓ BABAI: Trading group theory for randomness. In *Proc. 17th STOC*, pp. 421–429. ACM Press, 1985. [doi:10.1145/22145.22192] 3

[3] LÁSZLÓ BABAI: Local expansion of vertex-transitive graphs and random generation in finite groups. In *Proc. 23rd STOC*, pp. 164–174. ACM Press, 1991. [doi:10.1145/103418.103440] 10

[4] LÁSZLÓ BABAI, ROBERT BEALS, JIN-YI CAI, GÁBOR IVANYOS, AND EUGENE M. LUKS: Multiplicative equations over commuting matrices. In *Proc. 7th Ann. ACM–SIAM Symp. on Discrete Algorithms (SODA'96)*, pp. 498–507. SIAM, 1996. ACM DL. 3

[5] LÁSZLÓ BABAI, EUGENE M. LUKS, AND ÁKOS SERESS: Permutation groups in NC. In *Proc. 19th STOC*, pp. 409–420. ACM Press, 1987. [doi:10.1145/28395.28439] 3

[6] LÁSZLÓ BABAI AND ENDRE SZEMERÉDI: On the complexity of matrix group problems I. In *Proc. 25th FOCS*, pp. 229–240. IEEE Comp. Soc., 1984. [doi:10.1109/SFCS.1984.715919] 3, 7, 10

[7] DAVID A. MIX BARRINGTON: Bounded-width polynomial-size branching programs recognize exactly those languages in NC^1. *J. Comput. System Sci.*, 38(1):150–164, 1989. Preliminary version in STOC'86. [doi:10.1016/0022-0000(89)90037-8] 3

[8] DAVID A. MIX BARRINGTON AND PIERRE MCKENZIE: Oracle branching programs and Logspace versus P. *Inform. Comput.*, 95(1):96–115, 1991. Preliminary version in MFCS'89. [doi:10.1016/0890-5401(91)90017-V] 2, 3

[9] DAVID A. MIX BARRINGTON AND DENIS THÉRIEN: Finite monoids and the fine structure of NC^1. *J. ACM*, 35(4):941–952, 1988. Preliminary version in STOC'87. [doi:10.1145/48014.63138] 3

[10] DAVID MIX BARRINGTON, PETER KADAU, KLAUS-JÖRN LANGE, AND PIERRE MCKENZIE: On the complexity of some problems on groups input as multiplication tables. *J. Comput. System Sci.*, 63(2):186–200, 2001. Preliminary version in CCC'00. [doi:10.1006/jcss.2001.1764] 2, 3, 4, 8, 12, 15

[11] MARTIN BEAUDRY: *Membership Testing in Transformation Monoids*. Ph. D. thesis, McGill University, 1987. eScholarship@McGill. 3

[12] MARTIN BEAUDRY: Membership testing in commutative transformation semigroups. *Inform. Comput.*, 79(1):84–93, 1988. Preliminary version in ICALP'87. [doi:10.1016/0890-5401(88)90018-1] 3

[13] MARTIN BEAUDRY: Membership testing in threshold one transformation monoids. *Inform. Comput.*, 113(1):1–25, 1994. [doi:10.1006/inco.1994.1062] 3

[14] MARTIN BEAUDRY, PIERRE MCKENZIE, AND DENIS THÉRIEN: The membership problem in aperiodic transformation monoids. *J. ACM*, 39(3):599–616, 1992. [doi:10.1145/146637.146661] 3

[15] RAVI B. BOPPANA: The average sensitivity of bounded-depth circuits. *Inform. Process. Lett.*, 63(5):257–261, 1997. [doi:10.1016/S0020-0190(97)00131-2] 12

[16] XI CHEN, IGOR CARBONI OLIVEIRA, ROCCO A. SERVEDIO, AND LI-YANG TAN: Near-optimal small-depth lower bounds for small distance connectivity. In *Proc. 48th STOC*, pp. 612–625. ACM Press, 2016. [doi:10.1145/2897518.2897534, arXiv:1509.07476] 15

[17] LUKAS FLEISCHER: On the complexity of the Cayley semigroup membership problem. In *Proc. 33rd Comput. Complexity Conf. (CCC'18)*, pp. 25:1–12. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018. [doi:10.4230/LIPIcs.CCC.2018.25] 1

[18] MERRICK L. FURST, JOHN E. HOPCROFT, AND EUGENE M. LUKS: Polynomial-time algorithms for permutation groups. In *Proc. 21st FOCS*, pp. 36–41. IEEE Comp. Soc., 1980. [doi:10.1109/SFCS.1980.34] 3

[19] JOHAN HÅSTAD: Almost optimal lower bounds for small depth circuits. In *Proc. 18th STOC*, pp. 6–20. ACM Press, 1986. [doi:10.1145/12130.12132] 5

[20] JOHN M. HOWIE: *Fundamentals of Semigroup Theory*. Clarendon Press, 1995. 5

[21] NEIL D. JONES AND WILLIAM T. LAASER: Complete problems for deterministic polynomial time. *Theoret. Comput. Sci.*, 3(1):105–117, 1976. Preliminary version in STOC'74. [doi:10.1016/0304-3975(76)90068-2] 2, 3

[22] NEIL D. JONES, Y. EDMUND LIEN, AND WILLIAM T. LAASER: New problems complete for nondeterministic log space. *Math. Sys. Theory*, 10(1):1–17, 1976. [doi:10.1007/BF01683259] 2, 3, 4, 6

[23] FRIEDRICH LEVI: Über die Untergruppen der freien Gruppen. (2. Mitteilung). *Mathematische Zeitschrift*, 37:90–97, 1933. *Eu*DML. 14

[24] ROGER C. LYNDON AND PAUL E. SCHUPP: *Combinatorial Group Theory*. Springer, 2001. [doi:10.1007/978-3-642-61896-3] 14

[25] JEAN-ÉRIC PIN: *Varieties of Formal Languages*. North Oxford Academic, 1986. 5

[26] OMER REINGOLD: Undirected connectivity in Log-space. *J. ACM*, 55(4):17:1–24, 2008. Preliminary version in STOC'05. [doi:10.1145/1391289.1391291] 3

[27] CHARLES C. SIMS: Computational methods in the study of permutation groups. In *Computational Problems in Abstract Algebra*, pp. 169–183. Pergamon, 1970. See also SYMSAC'71. [doi:10.1016/B978-0-08-012975-4.50020-5] 3

[28] HERIBERT VOLLMER: *Introduction to Circuit Complexity*. Springer, 1999. [doi:10.1007/978-3-662-03927-4] 5

[29] ANDREW CHI-CHIH YAO: Separating the polynomial-time hierarchy by oracles. In *Proc. 26th FOCS*, pp. 1–10. IEEE Comp. Soc., 1985. [doi:10.1109/SFCS.1985.49] 5

AUTHOR

Lukas Fleischer
lfleischer@lfos.de
https://lfos.de/

ABOUT THE AUTHOR

LUKAS FLEISCHER received his Ph. D. at the University of Stuttgart, supervised by Volker Diekert, where his research focused on algorithmic and complexity aspects of finite semigroups. In 2019, he was a postdoctoral researcher in the Cheriton School of Computer Science at the University of Waterloo, working with Jeffrey Shallit. Lukas is now working as a software engineer in Kitchener–Waterloo. In his free time, he enjoys contributing to open source projects. He is also interested in technology, investing, entrepreneurship, sustainability, health, and fitness.