



Constant Depth Circuit Complexity for Generating Quasigroups

Nathaniel A. Collins
naco3124@colostate.edu

Colorado State University, Department of Mathematics
Fort Collins, Colorado, USA

Michael Levet
levetm@cofc.edu

College of Charleston, Department of Computer Science
Charleston, SC, USA

Joshua A. Grochow
jgrochow@colorado.edu

University of Colorado Boulder, Departments of Computer
Science and Mathematics
Boulder, Colorado, USA

Armin Weiß
armin.weiss@fmi.uni-stuttgart.de
Universität Stuttgart, FMI
Germany

ABSTRACT

We investigate the constant-depth circuit complexity of the ISOMORPHISM PROBLEM, MINIMUM GENERATING SET PROBLEM (MGS), and SUB(QUASI)GROUP MEMBERSHIP PROBLEM (MEMBERSHIP) for groups and quasigroups (=Latin squares), given as input in terms of their multiplication (Cayley) tables. Despite decades of research on these problems, lower bounds for these problems even against depth-2 AC circuits remain unknown. Perhaps surprisingly, Chattopadhyay, Torán, and Wagner (FSTTCS 2010; *ACM Trans. Comput. Theory*, 2013) showed that QUASIGROUP ISOMORPHISM could be solved by AC circuits of depth $O(\log \log n)$ using $O(\log^2 n)$ nondeterministic bits, a class we denote $\exists^{\log^2} \text{FOLL}$. We narrow this gap by improving the upper bound for these problems to quasiAC^0 , thus decreasing the depth to constant.

In particular, we show that MEMBERSHIP can be solved in $\text{NTIME}(\text{polylog}(n))$ and use this to prove the following:

- MGS for quasigroups belongs to $\exists^{\log^2} \forall^{\log} \text{NTIME}(\text{polylog}(n)) \subseteq \text{quasiAC}^0$. Papadimitriou and Yannakakis (*J. Comput. Syst. Sci.*, 1996) conjectured that this problem was $\exists^{\log^2} \text{NP}$ -complete; our results refute a version of that conjecture for completeness under quasiAC^0 reductions unconditionally, and under polylog -space reductions assuming $\text{EXP} \neq \text{PSPACE}$. It furthermore implies that this problem is not hard for any class containing PARITY . The analogous results concerning PARITY were known for QUASIGROUP ISOMORPHISM (Chattopadhyay, Torán, & Wagner, *ibid.*) and MEMBERSHIP for groups (Fleischer, *Theory Comput.* 2022), though not for MGS.

- MGS for groups belongs to $\text{AC}^1(\text{L})$. Our $\text{AC}^1(\text{L})$ bound improves on the previous, very recent, upper bound of P (Lucchini & Thakkar, *J. Algebra*, 2024). Our quasiAC^0 upper bound is incomparable to P, but has similar consequences to the above result for quasigroups.
- QUASIGROUP ISOMORPHISM $\in \exists^{\log^2} \text{AC}^0(\text{DTISP}(\text{polylog}(n), \log(n)))$, which is contained in quasiAC^0 . As a consequence of this result and previously known AC^0 reductions, this implies the same upper bound for the ISOMORPHISM PROBLEMS for: Steiner triple systems, pseudo-STS graphs, Latin square graphs, and Steiner $(t, t + 1)$ -designs. This improves upon the previous upper bound for these problems, which was $\exists^{\log^2} \text{NL} \cap \exists^{\log^2} \text{FOLL} \subseteq \text{quasiFOLL}$ (Chattopadhyay, Torán, & Wagner, *ibid.*; Levet, *Australas. J. Combin.* 2023).
- As a strong contrast, we show that MGS for arbitrary magmas is NP-complete.

Our results suggest that understanding the constant-depth circuit complexity may be key to resolving the complexity of problems concerning (quasi)groups in the multiplication table model.

CCS CONCEPTS

• Theory of computation \rightarrow Circuit complexity; Algebraic complexity theory; • Mathematics of computing \rightarrow Discrete mathematics.

KEYWORDS

Group Isomorphism, Quasigroup Isomorphism, Minimum Generating Set, Membership Testing, Constant-Depth Circuits, quasiAC^0 , Circuit Complexity

ACM Reference Format:

Nathaniel A. Collins, Joshua A. Grochow, Michael Levet, and Armin Weiß. 2024. Constant Depth Circuit Complexity for Generating Quasigroups. In *International Symposium on Symbolic and Algebraic Computation (ISSAC '24)*, July 16–19, 2024, Raleigh, NC, USA. ACM, New York, NY, USA, Article 111, 10 pages. <https://doi.org/10.1145/3666000.3669691>

1 INTRODUCTION

¹The GROUP ISOMORPHISM (GI) problem is a central problem in computational complexity and computer algebra. When the groups are given as input by their multiplication (a.k.a. Cayley) tables,

¹Some proofs are shortened, sketched, or omitted in this conference submission to save space; see the preprint of the full version [12] for full details.

Authors' Contact Information: Nathaniel A. Collins, naco3124@colostate.edu, Colorado State University, Department of Mathematics, Fort Collins, Colorado, USA; Joshua A. Grochow, jgrochow@colorado.edu, University of Colorado Boulder, Departments of Computer Science and Mathematics, Boulder, Colorado, USA; Michael Levet, levetm@cofc.edu, College of Charleston, Department of Computer Science, Charleston, SC, USA; Armin Weiß, armin.weiss@fmi.uni-stuttgart.de, Universität Stuttgart, FMI, Germany.



This work is licensed under a Creative Commons Attribution International 4.0 License.

ISSAC '24, July 16–19, 2024, Raleigh, NC, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0696-7/24/07
<https://doi.org/10.1145/3666000.3669691>

the problem reduces to GRAPH ISOMORPHISM (GI), and because the best-known runtimes for the two are quite close ($n^{O(\log n)}$ [29]² vs. $n^{O(\log^2 n)}$ [4]³), the former stands as a key bottleneck towards further improvements in the latter.

Despite this, GPI seems quite a bit easier than GI. For example, Tarjan's $n^{\log n + O(1)}$ algorithm for groups [29] can now be given as an exercise to undergraduates: every group is generated by at most $\lceil \log_2 |G| \rceil$ elements, so the algorithm is to select a generating set S for G and try all possible injections from S into H , checking whether any such map extends to an isomorphism. There are at most $n^{\log n}$ such maps, yielding a runtime of $n^{\log(n) + O(1)}$. In contrast, the quasipolynomial time algorithm for graphs was a tour de force that built on decades of cutting-edge research into algorithms and the structure of permutation groups. Nonetheless, it remains unknown whether the problem for groups is actually easier than that for graphs, or even whether both problems are in P!

Using a finer notion of reduction, Chattopadhyay, Torán, and Wagner [10] proved that there was no AC^0 reduction from GI to GPI. This gave the first (and still only known) unconditional evidence that there is *some formal sense* (namely, the AC^0 sense) in which GPI really is easier than GI. The key to their result was that the generator-enumeration technique described above can be implemented by non-deterministically guessing $\log^2 n$ bits (for the $\log n$ generators, each of $\log n$ bits), and then verifying an isomorphism by a circuit of depth only $O(\log \log n)$, a class we denote $\exists^{\log^2 n} \text{FOLL}$. Note that $\exists^{\log^2 n} \text{FOLL} \subseteq \text{quasiFOLL}$, which cannot compute PARITY [10, 32, 37]. As GI is DET-hard [40]—and hence can compute PARITY—there can be no AC^0 reduction from GI to GPI.

Such a low-depth circuit was quite surprising, although that surprise is perhaps tempered by the use of non-determinism. Nonetheless, it raises the question:

Is it possible that GROUP ISOMORPHISM is in AC^0 ?

The authors would be shocked if the answer were “yes,” and yet we do not even have results showing that GROUP ISOMORPHISM cannot be computed by polynomial-size circuits of (!) depth 2. The upper bound of $\exists^{\log^2 n} \text{FOLL}$ rules out most existing lower bound techniques against AC^0 , as most such techniques also yield similar lower bounds against $\exists^{\log^2 n} \text{FOLL}$.

In this paper, we aim to close the gap between AC^0 and $\exists^{\log^2 n} \text{FOLL}$ in the complexity of GROUP ISOMORPHISM and related problems. Our goal is to obtain constant-depth circuits of quasipolynomial⁴ size, a natural benchmark in circuit complexity [6]. Getting such circuits of polynomial size would resolve the long-standing question of putting these problems into P; in contrast, our current target is to reduce their *parallel complexity*. Our first main result along these lines is:

Theorem A. *(QUASI)GROUP ISOMORPHISM is in quasiAC^0 .*

(We discuss quasigroups in more detail below.) The prior best depth bounds on the general problem were all super-constant

(with quasipolynomial size): $\exists^{\log^2 n} \text{SC}^2$ by Tang [38] yields (by a standard simulation argument) circuits of depth $\log^2 n$, while $\exists^{\log^2 n} \text{L} \cap \exists^{\log^2 n} \text{FOLL}$ [10] has depth $\log \log n$. In particular, prior to our work, QUASIGROUP ISOMORPHISM was not known to be solvable using quasiAC circuits of depth $o(\log \log n)$. The special case of ABELIAN GPI was previously known to belong to quasiAC^0 [23].

Remark 1.1. We in fact get a more precise bound of

$$\exists^{\log^2 n} \sqrt{\log n} \exists^{\log n} \text{DTISP}(\text{polylog}(n), \log(n)).$$

This more precise bound is notable because it is contained in quasiAC^0 and $\exists^{\log^2 n} \text{FOLL} \cap \exists^{\log^2 n} \text{L}$, thus improving on [10]. We get similarly precise bounds with complicated-looking complexity classes for the other problems we study, but we omit the precise bounds in the introduction for readability.

Minimum generating set. Another very natural problem in computational algebra is the MIN GENERATING SET (MGS) problem. Given a group, this problem asks to find a generating set of the smallest possible size. Given that many algorithms on groups depend on the size of a generating set, finding a minimum generating set has the potential to be a widely applicable subroutine. The MGS problem for groups was shown to be in P by Lucchini & Thakkar very recently [28]. We improve their complexity bound:

Theorem B. *MGS for groups can be solved in quasiAC^0 and in $AC^1(\text{L})$ ($O(\log n)$ -depth, unbounded fan-in circuits with a logspace oracle).*

We note that, although quasiAC^0 is incomparable to P because of the quasipolynomial size (whereas $AC^1(\text{L}) \subseteq \text{P}$), the key we are focusing on here is reducing the depth. For nilpotent groups (widely believed to be the hardest cases of GPI), if we only wish to compute the minimum *number* of generators, we can further improve this complexity to a subclass of $\text{L} \cap \text{FOLL} \cap \text{quasiAC}^0$ (Prop. 7.3).

While our $AC^1(\text{L})$ bound above is essentially a careful complexity analysis of the polynomial-time algorithm of Lucchini & Thakkar [28], the quasiAC^0 upper bound is in fact a consequence of our next, more general result for *quasigroups*, which involves some new ingredients.

Enter quasigroups. Quasigroups can be defined in (at least) two equivalent ways: (1) an algebra whose multiplication table is a Latin square,⁵ or (2) a group-like algebra that need not have an identity nor be associative, but in which left and right division are uniquely defined, that is, for all a, b , there are unique x and y such that $ax = b$ and $ya = b$, denoted $x = a \backslash b$ and $y = b / a$.

In the paper in which they introduced $\log^2(n)$ -bounded nondeterminism, Papadimitriou and Yannakakis showed that for arbitrary magmas,⁶ testing whether the magma has $\log n$ generators was in fact *complete* for $\exists^{\log^2 n} \text{P}$, and conjectured:

Conjecture 1.2 (Papadimitriou & Yannakakis [31, p. 169]). *MGS for quasigroups is $\exists^{\log^2 n} \text{P}$ -complete.*

They explicitly did *not* conjecture the same for MGS for *groups*, writing:

²Miller [29] credits Tarjan for $n^{\log n + O(1)}$.

³Babai [4] proved quasipolynomial time, and the exponent of the exponent was analyzed and improved by Helfgott [20]

⁴In our setting, quasipolynomial means $n^{\text{poly}(\log n)} = 2^{\text{poly}(\log n)}$.

⁵A Latin square is an $n \times n$ matrix where for each row and each column, the elements of $[n]$ appear exactly once

⁶A magma is a set M together with a function $M \times M \rightarrow M$ that need not satisfy any additional axioms.

“We conjecture that this result [$\exists^{\log^2 n}$ P-completeness] also holds for the more structured MINIMUM GENERATOR SET OF A QUASIGROUP problem. In contrast, QUASIGROUP ISOMORPHISM was recently shown to be in $\text{DSPACE}(\log^2 n)$ [43]. Notice that the corresponding problems for groups were known to be in $\text{DSPACE}(\log^2 n)$ [26].”—Papadimitriou & Yannakakis [31, p. 169]

We thus turn our attention to the analogous problems for quasigroups: MGS, ISOMORPHISM, and the key subroutine, MEMBERSHIP.

Theorem C. *MGS for quasigroups is in $\text{quasiAC}^0 \cap \text{DSPACE}(\log^2 n)$.*

To the best of our knowledge, MGS FOR QUASIGROUPS has not been studied from the complexity-theoretic viewpoint previously. While a $\text{DSPACE}(\log^2 n)$ upper bound for MGS for *groups* follows from [1, 38], as far as we know it remained open for quasigroups prior to our work.

As with prior results on (QUASI)GROUP ISOMORPHISM [10] and other isomorphism problems (e.g., [25]), Thm. C shows that PARITY does not reduce to MGS FOR QUASIGROUPS, thus ruling out most known lower bound methods that might be used to prove that MGS FOR QUASIGROUPS is not in AC^0 . We observe a similar bound for MGS FOR GROUPS using Fleischer’s technique [17].

Papadimitriou & Yannakakis did not specify the type of reduction used in their conjecture, though their $\exists^{\log^2 n}$ P-completeness result for LOG GENERATING SET for magmas works in both logspace and AC^0 (under a suitable input encoding). Our two upper bounds rule out such reductions for MGS FOR QUASIGROUPS:

Corollary C. *Conj. 1.2 is false under quasiAC^0 reductions. It is also false under polylog-space reductions assuming $\text{EXP} \neq \text{PSPACE}$.*

In strong contrast, we show that MGS FOR MAGMAS is NP-complete (Thm. 7.10).

A key ingredient in our proof of Thm. C is an improvement in the complexity of another central problem in computational algebra: the SUB-QUASIGROUP MEMBERSHIP problem (MEMBERSHIP,⁷ for short):

Theorem D. *MEMBERSHIP for quasigroups is in $\text{NTIME}(\text{polylog}(n)) \subseteq \text{quasiAC}^0$.*

MEMBERSHIP for *groups* is well-known to belong to L, by reducing to the connectivity problem on the Cayley graph (cf. [8, 33]), but as L sits in between AC^0 and AC^1 , this is not low enough depth for us.

Additional results. We also obtain a number of additional new results on related problems, some of which we highlight here:

- By known AC^0 reductions (see, e.g., [25]), our quasiAC^0 analysis of Chattopadhyay, Torán, and Wagner’s algorithm for QUASIGROUP ISOMORPHISM yields the same upper bound for the isomorphism problems for Steiner triple systems, pseudo-STs graphs, Latin square graphs, and Steiner $(t, t+1)$ -designs, as well as LATIN SQUARE ISOTOPY.
- GPI for groups from a dense set of orders can be solved in a subclass of $\text{quasiAC}^0 \cap \text{FOLL} \cap \text{L}$. This improves the parallel complexity compared to the original result [15].
- ABELIAN GROUP ISOMORPHISM is in

$$\forall^{\log \log n} \text{MAC}^0(\text{DTISP}(\text{polylog}(n), \log(n))).$$

⁷In the literature, the analogous problem for groups is sometimes called CAYLEY GROUP MEMBERSHIP or CGM, to highlight that it is in the Cayley table model.

The key novelties here are (1) a new observation that allows us to reduce the number of co-nondeterministic bits from $\log n$ (as in [18]) down to $\log \log n$, and (2) using an $\text{AC}^0(\text{DTISP}(\text{polylog}(n), \log(n)))$ circuit for order finding, rather than FOLL as in [10].

- MEMBERSHIP for nilpotent groups is in $\text{NTISP}(\text{polylog}(n), \log(n))$, which is contained in $\text{FOLL} \cap \text{quasiAC}^0$.

1.1 Methods

Several of our results involve careful analysis of the low-level circuit complexity of extant algorithms, showing that they in fact lie in smaller complexity classes than previously known. We crucially use simultaneous time- and space-restricted computations. This not only facilitates several proofs and gives better complexity bounds, but also gives rise to new algorithms such as for MEMBERSHIP for nilpotent groups, which previously was not known to be in FOLL.

One such instance is in our improved bound for order-finding and exponentiation in a semigroup (Lem. 3.1). The previous proof [7] (still state of the art 23 years later) used a then-novel and clever “double-barrelled” recursive approach to compute these in FOLL. In contrast, our proof uses standard repeated doubling, noting that it can be done in $\text{DTISP}(\text{polylog}(n), \log(n)) \subseteq \text{FOLL} \cap \text{quasiAC}^0$, recovering their result with standard tools and reducing the depth. We use this improved bound on order-finding to improve the complexity of isomorphism testing of Abelian groups (Thm. 4.1), simple groups (Cor. 3.3), and groups of almost all orders (Thm. 5.1).

For several results we additionally need to develop new tools to work with quasigroups. Notably, for the quasiAC^0 upper bound on MGS for quasigroups, we cannot directly adapt the technique of [10], as the existence alone of cube generating sets seems insufficient for MEMBERSHIP and MGS. The first key is Thm. D, putting MEMBERSHIP for quasigroups into $\text{NTIME}(\text{polylog}(n))$. Here, we avoid their use of cube generating sets and instead utilize straight-line programs (SLPs) directly. To show short SLPs exist we extend the Babai–Szemerédi Reachability Lemma [5, Thm. 3.1] from groups (its original setting) to quasigroups. As division in quasigroups is nuanced due to lack of associativity, our proof is a careful adaptation of the technique of [5, Thm. 3.1], with a few quasigroup twists.

1.2 Prior work

Isomorphism testing. The best known runtime bound for GPI is $n^{(1/4) \log_p(n) + O(1)}$, where p is the smallest prime dividing n , [35] (see [24, Sec. 2.2]), though this tells us little about parallel complexity. Lipton, Snyder, & Zalcstein [26] independently observed the generator-enumeration procedure and used it to give a bound of $\text{DSPACE}(\log^2 n)$. Miller [29] extended Tarjan’s result to quasigroups. There has been subsequent work on improving the parallel complexity of generator enumeration for quasigroups, resulting in bounds of $\exists^{\log^2 n} \text{AC}^1$ [43]⁸, $\exists^{\log^2 n} \text{SAC}^1$ [42], and $\exists^{\log^2 n} \text{L} \cap \exists^{\log^2 n} \text{FOLL}$ [10]. For groups, generator enumeration is also known to belong to $\exists^{\log^2 n} \text{SC}^2$ [38]. There has been considerable work on polynomial-time isomorphism tests for several families of groups,

⁸Wolf actually claims a bound of $\exists^{\log^2 n} \text{NC}^2$; however, he uses NC^1 circuits to multiply two elements of a quasigroup rather than AC^0 circuits.

and more recent work on NC isomorphism tests—see [15, 18, 19] for a survey. We are not aware of work on isomorphism testing for specific families of quasigroups that are not groups.

Min Generating Set. As every (quasi)group has a generating set of size $\leq \lceil \log n \rceil$, MGS admits an $n^{\log(n)+O(1)}$ -time solution for (quasi)groups. Arvind & Torán [1] improved the complexity to $\text{DSPACE}(\log^2 n)$ for groups. They also gave a polynomial-time algorithm in the special case of nilpotent groups. Tang further improved the general bound for MGS for groups to $\exists^{\log^2 n} \text{SC}^2$ [38]. We observe that Wolf’s technique for placing $\text{QUASIGROUP ISOMORPHISM}$ into $\text{DSPACE}(\log^2 n)$ also suffices to get MGS for quasigroups into the same class. Recently, MGS for groups was placed into P [28]. Prior to [28], MGS for groups was considered comparable to GROUP ISOMORPHISM in terms of difficulty [13]. Our $\text{AC}^1(\text{L})$ bound (Thm. C) further closes the gap between MEMBERSHIP and MGS for groups, and in particular suggests that MGS is of comparable difficulty to MEMBERSHIP rather than GPI . Note that MEMBERSHIP is known to belong to L [8, 33].

2 ADDITIONAL PRELIMINARIES

2.1 Algebra

A *magma* M is an algebraic structure together with a binary operation $\cdot : M \times M \rightarrow M$. We will frequently consider subclasses of finite magmas, such as groups, quasigroups, and semigroups (associative magmas).

Quasigroups. As quasigroups are non-associative, the parenthesization of a given expression may impact the resulting value. For a sequence $S := (s_0, s_1, \dots, s_k)$ and parenthesization P from a quasigroup G , define: $\text{Cube}(S) = \{P(s_0 s_1^{e_1} \dots s_k^{e_k}) : e_1, \dots, e_k \in \{0, 1\}\}$. We say that S is a *cube generating sequence* if each element $g \in G$ can be written as $g = P(s_0 s_1^{e_1} \dots s_k^{e_k})$, for $e_1, \dots, e_k \in \{0, 1\}$. Here, s_i^0 indicates that s_i is not being considered in the product. For every parenthesization, every quasigroup admits a cube generating sequence of size $O(\log n)$ [10].

Group Theory. For a standard reference, see [34]. A *chief series* of G is an ascending chain $(N_i)_{i=0}^k$ of normal subgroups of G , where $N_0 = 1$, $N_k = G$, and each N_{i+1}/N_i ($i = 0, \dots, k-1$) is minimal normal in G/N_i .

Algorithmic Problems. We will consider the following algorithmic problems, where the magmas are given by their multiplication tables. The $\text{QUASIGROUP ISOMORPHISM}$ problem takes as input two quasigroups Q_1, Q_2 and asks if there is an isomorphism $\varphi : Q_1 \cong Q_2$. The MEMBERSHIP problem for groups takes as input a group G , a set $S \subseteq G$, and an element $x \in G$, and asks if $x \in \langle S \rangle$ (the subgroup generated by S). We define the MEMBERSHIP problem analogously when the input is a semigroup or quasigroup, and $\langle S \rangle$ is considered as the sub-semigroup or sub-quasigroup generated by S , respectively. The $\text{MINIMUM GENERATING SET}$ (MGS) problem takes as input a magma M and asks for a generating set $S \subseteq M$ where $|S|$ is minimum. The *decision variant* of MGS additionally takes an integer k in the input and ask whether there exists a generating set of size at most k .

2.2 Computational Complexity

We assume familiarity with standard complexity classes such as L, NL, NP, and EXP. For a standard reference on circuit complexity, see [41]. We consider Boolean circuits using AND, OR, NOT, and Majority, where $\text{Majority}(x_1, \dots, x_n) = 1$ iff $\geq n/2$ of the inputs are 1. All our polynomial-size circuit families are DLOGTIME -uniform, see [12, §2.2] for details.

A language L belongs to NC^k if there is a family of circuits C_n with NOT gates, and 2-input AND and OR gates, of depth $O(\log^k n)$ and size $n^{O(1)}$, such that $x \in L \Leftrightarrow C_{|x|}(x) = 1$. The class SAC^k is defined analogously, but the OR gates may have unbounded fan-in, while the AND gates still must have fan-in 2. For AC^k also the AND gates are permitted to have unbounded fan-in. The complexity class TC^k is defined analogously as AC^k , except that our circuits are now also permitted Majority gates of unbounded fan-in.

$$\text{AC}^0 \subseteq \text{TC}^0 \subseteq \text{NC}^1 \subseteq \text{L} \subseteq \text{NL} \subseteq \text{SAC}^1 \subseteq \text{AC}^1 \subseteq \dots$$

For a language L the class $\text{AC}^k(L)$, apart from Boolean gates, also allows oracle gates for L .⁹ If $K \in \text{AC}^k(L)$, then K is said to be AC^k Turing reducible to L . For a complexity class C denote $\text{AC}^k(C)$ to be the set of decision problems that are AC^k -Turing reducible to problems in C .¹⁰

Further circuit classes. The complexity class MAC^0 is the set of languages decidable by TC^0 circuits with only a single Majority gate that must be the output gate. MAC^0 was introduced (but not so named) in [2], where it was shown that $\text{MAC}^0 \subseteq \text{TC}^0$, and later given the name MAC^0 in [21].

The complexity class FOLL is the set of languages decidable by uniform AC circuit families of depth $O(\log \log n)$ and polynomial size. It is known that $\text{AC}^0 \subseteq \text{FOLL} \subseteq \text{AC}^1$, and it is open as to whether FOLL is contained in NL [7].

We will be particularly interested in AC circuits of quasipolynomial ($n^{\text{poly} \log(n)}$) size. For a circuit class $C \subseteq \text{AC}$, the analogous class permitting a quasipolynomial number of gates is denoted $\text{quasi}C$. We will focus specifically on quasiAC^0 . Note that $\text{PARITY} \notin \text{quasiFOLL}$ [32, 37].

Bounded nondeterminism. For a complexity class C , define $\exists^{f(n)}C$ (resp. $\forall^{f(n)}C$) to be the set of languages L such that there exists an $L' \in C$ such that $x \in L$ iff there exists (resp., for all) y of length at most $O(f(|x|))$ such that $(x, y) \in L'$. For any $i \geq 0$, $\exists^{\log^i n} \text{FOLL} \cup \forall^{\log^i n} \text{FOLL} \subseteq \text{quasiFOLL}$, and so cannot compute PARITY [10, 37]. Note that $\forall^{\log n} C \cup \exists^{\log n} C \subseteq \text{AC}^0(C)$.

Time and space-restricted Turing machines. For complexity classes defined by Turing machines with a time bound $t(n) \in o(n)$, we use Turing machine with random access and a separate address (or *index*) tape. After writing an address, the machine can go to a query state reading the symbol from the input at the location specified by the address tape. The contents of the address tape are not deleted. For functions $t(n), s(n) \in \Omega(\log n)$, the classes $\text{DTISP}(t(n), s(n))$ and $\text{NTISP}(t(n), s(n))$ are defined by deterministic (resp. nondeterministic) $t(n)$ time and $s(n)$ space bounded

⁹An oracle gate outputs 1 if and only if its input is in L .

¹⁰Be aware that here we follow the notation of [41], which is different from [18, 42] (where $\text{AC}^k(C)$ is used to denote composition of functions).

Turing machines. There must be one Turing machine that simultaneously satisfies the time and space bound. See [41, §2.6] for details. For connections to quasiAC^0 , see [6, 16].

Fact 2.1. $\text{NTISP}(\text{polylog}(n), \log(n)) \subseteq \text{NTIME}(\text{polylog}(n)) \subseteq \text{quasiAC}^0$.

PROOF SKETCH. Take the OR over all $2^{\text{polylog}(n)}$ possible computation histories, of the AC^0 circuit that verifies a computation history (the latter as in the proof of the Cook–Levin Theorem). \square

Lemma 2.2. $\text{NTISP}(\text{polylog}(n), \log(n)) \subseteq \text{FOLL}$.

PROOF SKETCH. Follows the proof of Savitch’s Theorem. The $\log n$ space bound implies the configuration space has only $n^{O(1)}$ vertices, while the $\text{polylog}(n)$ time bound means one need only find a path of length $\text{polylog}(n)$, which can be done by Savitch’s “repeated doubling” technique in recursion depth $O(\log \log n)$. \square

By the very definition we have $\text{DTISP}(\text{polylog}(n), \log(n)) \subseteq \text{L}$ and $\text{NTISP}(\text{polylog}(n), \log(n)) \subseteq \text{NL}$. Thus, we obtain

- $\text{AC}^0(\text{DTISP}(\text{polylog}(n), \log(n))) \subseteq \text{L} \cap \text{FOLL} \cap \text{quasiAC}^0$ and
- $\text{AC}^0(\text{NTISP}(\text{polylog}(n), \log(n))) \subseteq \text{NL} \cap \text{FOLL} \cap \text{quasiAC}^0$.

3 ORDER FINDING AND APPLICATIONS

In this section, we improve the parallel complexity of order finding with an easier and more general proof than [7]. We use this in several applications in this section, §4, and §5. The key lemma is:

Lemma 3.1. *The following problem is in $\text{DTISP}(\text{polylog}(n), \log(n))$: On input of a multiplication table of a semi-group S , an element $s \in S$, and a unary or binary number $k \in \mathbb{N}$ with $k \leq |S|$, compute s^k .*

PROOF. If k is given in unary, we first compute its binary representation using a binary search (we can write it on the work tape as it uses at most $\lceil \log |S| \rceil$ bits). We identify the semigroup elements with the natural numbers $0, \dots, |S| - 1$. Now, compute s^k using the standard fast exponentiation algorithm. Multiplying two semigroup elements can be done in $\text{DTIME}(\log n)$ as computing the address of their product only involves multiplying two $\log n$ -bit addresses. Finally, it is well-known that the fast exponentiation algorithm needs only $O(\log k)$ algebra multiplications and $O(\log k + \log n)$ space. \square

Corollary 3.2. *On input of a group G , an element $g \in G$, and $k \in \mathbb{N}$, we may decide whether $\text{ord}(g) = k$ in $\forall^{\log n} \text{DTISP}(\text{polylog}(n), \log(n))$.*

Application to isomorphism testing. Using Cor. 3.2, we can improve the upper bound for isomorphism testing of finite simple groups. Previously, this problem was known to be in L [38] and FOLL [18].

Corollary 3.3. *Let G be a finite simple group and H be arbitrary. We can decide whether $G \cong H$ in $\text{AC}^0(\text{DTISP}(\text{polylog}(n), \log(n)))$.*

Cor. 3.2 also lets us improve on the $\text{L} \cap \text{FOLL}$ bound [7, Thm. 3.2] for testing whether a group is nilpotent to:

Corollary 3.4. *Deciding whether a group G (given by its multiplication table) is nilpotent is in $\text{AC}^0(\text{DTISP}(\text{polylog}(n), \log(n)))$.*

Application to membership testing. A group G has the $\log n$ power basis property [7, 17] if for every $X \subseteq G$ every $g \in \langle X \rangle$ can be written as $g = g_1^{e_1} \cdots g_m^{e_m}$ with $m \leq \log n$ and suitable $g_i \in X$ and $e_i \in \mathbb{Z}$. Lem. 3.1 provides that MEMBERSHIP for semigroups with the $\log n$ power basis property is in $\text{NTISP}(\text{polylog}(n), \log(n))$. This observation allows us to improve some FOLL bounds from [7, 17] to $\text{NTISP}(\text{polylog}(n), \log(n))$.

Corollary 3.5. *MEMBERSHIP is in $\text{NTISP}(\text{polylog}(n), \log(n))$ for:*

- (a) commutative semigroups;
- (b) nilpotent groups;
- (c) solvable groups of class $O(1)$.

The above result for nilpotent groups involves no restriction on the nilpotency class, improving considerably over [7, Cor. 3.2].

PROOF SKETCH FOR NILPOTENT GROUPS. Given $g \in G, X \subseteq G$. The m -th term of the lower central series $\Gamma_0 = \langle X \rangle, \Gamma_{m+1} = [\Gamma_m, \langle X \rangle]$ is generated by $C_m = \{[[[x_1, x_2], x_3], \dots, x_k] : x_i \in X, k \geq m\}$ (e. g., [11, Lem. 2.6]). Define C' by choosing a minimal generating set for Γ_0/Γ_1 from among C_0 , then a minimal generating set of Γ_1/Γ_2 among C_1 , and so on. C' will be a so-called polycyclic generating set of $\langle X \rangle$, so every $g \in \langle X \rangle$ can be written $g = c_1^{e_1} \cdots c_m^{e_m}$ with $e_i \in \mathbb{Z}$ and $c_i \in C'$. We nondeterministically guess the latter expression for g . As $m \leq \log |G|$ for any nilpotent group, this gives an $\text{NTISP}(\text{polylog}(n), \log(n))$ algorithm. \square

4 ABELIAN GROUP ISOMORPHISM

Our next application of our order-finding Lem. 3.1 is:

Theorem 4.1. *Let G be an Abelian group, and let H be arbitrary. We can decide isomorphism between G and H in*

$$\forall^{\log \log n} \text{MAC}^0(\text{DTISP}(\text{polylog}(n), \log(n))).$$

Chattopadhyay, Torán, and Wagner [10] established a $\text{TC}^0(\text{FOLL})$ upper bound on this problem. Grochow & Levet [18, Thm. 5] gave a tighter analysis of their algorithm, placing it in the sub-class $\forall^{\log n} \text{MAC}^0(\text{FOLL})$.¹¹ Chattopadhyay, Torán, & Wagner also established an upper bound of L for this problem, which is incomparable to the result of Grochow & Levet (*ibid.*). We improve upon both these bounds by (i) showing that $O(\log \log n)$ non-deterministic bits suffice instead of $O(\log n)$ bits, and (ii) using our improved bound on order-finding (Lem. 3.1).

While $\forall^{\log n} \text{MAC}^0(\text{FOLL})$ is contained in $\text{TC}^0(\text{FOLL})$, it is open whether this containment is strict. In contrast, it follows from [9, Thm. 7] that our new bound is a class that is in fact *strictly* contained in $\text{L} \cap \text{TC}^0(\text{FOLL})$ (see [12, Cor. 5.3] for details).

PROOF OF THM. 4.1. Following the strategy of [18, Theorem 7.15], we show that *non*-isomorphism can be decided in the same class but with existentially quantified non-deterministic bits.

We may check in AC^0 whether a group is Abelian. So now assume that H is Abelian. If $G \not\cong H$, then there exists a prime power p^e such that there are more elements of order p^e in G

¹¹Grochow & Levet consider $\forall^{\log n} \text{MAC}^0 \circ \text{FOLL}$, where \circ denotes composition (see [18] for a precise formulation). Note that as $\text{AC}^0 \circ \text{FOLL} = \text{FOLL} = \text{AC}^0(\text{FOLL})$, we have $\forall^{\log n} \text{MAC}^0 \circ \text{FOLL} = \forall^{\log n} \text{MAC}^0(\text{FOLL})$. Thus, Thm. 4.1 improves upon the previous bound of $\forall^{\log n} \text{MAC}^0(\text{FOLL})$ obtained by Grochow & Levet.

than in H . We first identify the order of each element, which is $AC^0(\text{DTISP}(\text{polylog}(n), \log(n)))$ -computable by Cor. 3.2.

Let $n = p_1^{e_1} \cdots p_\ell^{e_\ell}$ be the prime factorization of n . Using $O(\log \log n)$ non-deterministic bits, we can guess a pair (p, e) (p represented by its index in the prime decomposition), where G has more elements of p^e than H . We may then, in AC^0 , compute p^e from p . We use Cor. 3.2 to identify the elements of order p^e , and then an MAC^0 circuit to compare the number of such elements in G vs. H . See [12, Thm. 5.1] for full details. \square

5 GROUP ISOMORPHISM FOR ALMOST ALL ORDERS

We use our improved order-finding Lem. 3.1 to improve the parallel complexity of GpI for almost all orders. Dietrich & Wilson [15] proved that there is a dense set $\Upsilon \subseteq \mathbb{N}$ such that if $n \in \Upsilon$ and G_1, G_2 are magmas of order n given by their multiplication tables, we can (i) decide if G_1, G_2 are groups, and (ii) if so, decide whether $G_1 \cong G_2$ in time $O(n^2 \log^2 n)$, which is quasi-linear time relative to the input size. We improve the parallel complexity of their result:

Theorem 5.1. *Let $n \in \Upsilon$, and let G_1, G_2 be magmas of order n . We can decide whether the G_i are groups and whether $G_1 \cong G_2$ in $AC^0(\text{DTISP}(\text{polylog}(n), \log(n)))$.*

PROOF. Deciding whether a magma is a group is readily seen to be in AC^0 . Dietrich & Wilson showed [15, Thm. 2.5] that if G is a group of order $n \in \Upsilon$, then $G = H \ltimes B$, where:

- B is a cyclic group of order $p_1 \cdots p_\ell$, where for each $i \in [\ell]$, $p_i > \log \log n$ and p_i is the maximum power of p_i dividing n .
- $|H| = (\log n)^{\text{poly log log } n}$; and in particular, if a prime divisor p of n satisfies $p \leq \log \log n$, then p divides $|H|$.

Fix groups G_1, G_2 of order n . We may, in $AC^0(\text{DTISP}(\text{polylog}(n), \log(n)))$, do the following: (i) decompose $G_i = H_i \ltimes B_i$, as prescribed by [15, Thm. 2.5], (ii) select a generator b_i for B_i , and (iii) select an isomorphism for $H_1 \cong H_2$ by specifying cube generating sequences (or decide if $H_1 \not\cong H_2$). Step (iii) utilizes the proof of [12, Thm. 4.1] in the full version. As $|H_1| = |H_2| \leq (\log n)^{\text{poly log log } n}$, there are only $\text{poly}(n)$ such generating sequences to consider.

Suppose $H_1 \cong H_2, B_1 \cong B_2$, and $\gcd(|B_j|, |H_j|) = 1$ for $j = 1, 2$. We have by the Schur–Zassenhaus Theorem that $G_j = H_j \rtimes_{\theta_j} B_j$ ($j = 1, 2$). By Taunt’s Lemma [39], it remains to test whether the actions θ_1 and θ_2 are equivalent. This step is computable in $AC^0(\text{DTISP}(\text{polylog}(n), \log(n)))$; see [12, Thm. 6.1] for details. \square

6 QUASIGROUP ISOMORPHISM

The main results in the remainder of the paper do not rely on our order-finding Lem. 3.1, but rather on other techniques. The following proof is inspired by [14] where a similar problem is shown to be in the third level of the polynomial hierarchy using the same approach.

Theorem 6.1. *QUASIGROUP ISOMORPHISM belongs to*

$$\exists^{\log^2 n} \forall^{\log n} \exists^{\log n} \text{DTISP}(\text{polylog}(n), \log(n)).$$

Note that we have $\exists^{\log^2 n} \forall^{\log n} \exists^{\log n} \text{DTISP}(\text{polylog}(n), \log(n)) \subseteq \exists^{\log^2 n} AC^0(\text{DTISP}(\text{polylog}(n), \log(n))) \subseteq \text{quasi}AC^0 \cap \exists^{\log^2 n} L \cap \exists^{\log^2 n} \text{FOLL}$.

PROOF SKETCH. Careful analysis of the algorithm of [10, Thm. 3.4]. The key idea is to guess cube generating sequences for both input groups (in the $\exists^{\log^2 n}$ part) and then in the remaining computation verify whether these are in fact cube generating sequences that, indeed, induce an isomorphism. \square

We obtain several corollaries, that all essentially follow by known AC^0 -reductions; [25] is a convenient reference for definitions of the problems below and the reductions.

Corollary 6.2. *The following isomorphism problems are in*

$$\exists^{\log^2 n} AC^0(\text{DTISP}(\text{polylog}(n), \log(n))) :$$

- (1) *Isomorphism of Steiner triple systems*
- (2) *Isomorphism of pseudo-STS graphs*
- (3) *Isomorphism of Steiner $(t, t+1)$ -designs*
- (4) *Isomorphism of Latin square graphs*
- (5) *Latin Square Isotopy.*

Latin square graphs are one of the four families¹² of strongly regular graphs under Neumaier’s classification [30]. Levet [25] established an upper bound of $\exists^{\log^2 n} AC^0$ for isomorphism testing of conference graphs, which is a stronger upper bound than we obtain for Latin square graphs. In contrast, the best known algorithmic runtime for identifying conference graphs is $n^{2 \log(n) + O(1)}$ [3], whereas Latin square graphs admit an $n^{\log(n) + O(1)}$ -time solution [29].

7 MINIMUM GENERATING SET

In this section, we consider the MINIMUM GENERATING SET (MGS) problem for groups, quasigroups, and arbitrary magmas.

7.1 MGS for Groups in $AC^1(L)$

Theorem 7.1. *MGS for groups belongs to $AC^1(L)$.*

We begin with the following lemma from the full version [12].

Lemma 7.2. *Let G be a group. We can compute a chief series for G in $AC^1(L)$.*

PROOF OF THM. 7.1. By Lem. 7.2, we can compute a chief series $N_1 \triangleleft N_2 \triangleleft \cdots \triangleleft N_k = G$ in $AC^1(L)$. We proceed inductively down the N_i . As G/N_{k-1} is a finite simple group, and hence at most 2-generated, we can find a generating set in L [38].

Fix $i < k$. Given a minimum generating sequence $g_1, \dots, g_d \in G$ for G/N_i we construct a minimum generating sequence for G/N_{i-1} as follows. If $N = N_i/N_{i-1}$ is Abelian (case 1):

- **Case 1a:** We have $G/N_{i-1} = \langle g_1, \dots, g_i, g_j^n, g_{j+1}, \dots, g_d \rangle$ for some $j \in [d]$ and some $n \in \mathbb{N}$ (possibly $n = 1$). There are at most $d \cdot |N|$ generating sets to consider in this case and we can test each of them in L with MEMBERSHIP [8, 33].

¹²the other families being line graphs of Steiner 2-designs, conference graphs, and graphs whose eigenvalues satisfy the claw bound

- **Case 1b:** If Case 1a does not hold, then by [27, Thm. 4] we have that $G/N_{i-1} = \langle g_1, \dots, g_d, x \rangle$ for any non-identity element $x \in N$.

Otherwise, N is non-Abelian (case 2), and we have by [28, Cor. 13] the following holds. Let $\eta_G(N)$ denote the the number of factors in a chief series with order $|N|$. Let $u = \max\{d, 2\}$ and $t = \min\{u, \lceil \frac{8}{5} + \log_{|N|} \eta_G(N) \rceil\}$. Then there exist $n_1, \dots, n_t \in N_{i-1}$ (possibly $n_1 = \dots = n_t = 1$) such that $G/N_{i-1} = \langle g_1 n_1, \dots, g_t n_t, g_{t+1}, \dots, g_d \rangle$.

By [28, Cor. 13], there are at most $|N|^{\lceil \frac{8}{5} + \log_{|N|} \eta_G(N) \rceil}$ generating sets of this form. As $\log_{|N|} \eta_G(N) \in O(1)$, we may write down these generating sets in parallel with a single AC^0 circuit and test whether each generates G/N_{i-1} in L using MEMBERSHIP.

Since a chief series has $O(\log n)$ terms, this algorithm requires $O(\log n)$ iterations and each iteration is computable in L , resulting in an algorithm for MGS in $AC^1(L)$. \square

Improving upon the $AC^1(L)$ bound on MGS for groups appears daunting. We thus inquire as to families of groups where MGS is solvable in complexity classes contained within $AC^1(L)$. To this end, we examine the class of nilpotent groups. Arvind & Torán previously established a polynomial-time algorithm for nilpotent groups [1, Thm. 7]. We improve their bound as follows. Here, $d(G)$ denotes the minimum size of a generating set for G .

Proposition 7.3. *For a nilpotent group G , we can compute $d(G)$ in $L \cap AC^0(NTISP(\text{polylog}(n), \log(n)))$.*

PROOF SKETCH. Use order-finding (Cor. 3.2) to compute the Sylow p -subgroups in $AC^0(DTISP(\text{polylog}(n), \log(n)))$; we have $d(G) = \max_{\text{Sylow } p} d(P)$. For each Sylow subgroup P , compute the Frattini subgroup $PP[P, P]$ based on the Burnside Basis Theorem (see [34, Thm. 5.3.2]) and Cor. 3.2. Then $d(P) = |P/PP[P, P]|$. A key fact we use along the way is that every element of $[P, P]$ is the product of at most $\log |P|$ commutators [36, I.§4 Ex. 5]. \square

Remark 7.4. While Prop. 7.3 allows us to compute $d(G)$ for a nilpotent group G , the algorithm is non-constructive. It is not clear how to find such a generating set in L . We can, however, provide such a generating set in $AC^1(NTISP(\text{polylog}(n), \log(n)))$. This bound is incomparable to $AC^1(L)$. See the full version [12, Rmk. 7.4] for details.

7.2 MGS for Quasigroups

In this section, we consider the MINIMUM GENERATING SET problem for quasigroups. Our goal is to establish the following.

Theorem 7.5. *For MGS for quasigroups,*

- The decision version belongs to $\exists^{\log^2 n} \forall^{\log n} NTIME(\text{polylog}(n))$;*
- The search version belongs to $\exists^{\log^2 n} SAC^1 \subseteq DSPACE(\log^2 n)$,¹³ and*
- The search version belongs to $quasiAC^0$.*

In the paper in which they introduced (polylog-)limited nondeterminism, Papadimitriou and Yannakakis conjectured that MGS for quasigroups was $\exists^{\log^2 n} P$ -complete [31, after Thm. 7]. While they did not specify the type of reductions used, it may be natural

¹³Wolf [43] showed the containment $\exists^{\log^2 n} SAC^1 \subseteq DSPACE(\log^2 n)$.

to consider polynomial-time many-one reductions. Thm. 7.5 refutes two versions of their conjecture under other kinds of reductions, that are incomparable to polynomial-time many-one reductions: $quasiAC^0$ reductions unconditionally and polylog-space reductions conditionally. We note that their other $\exists^{\log^2 n} P$ -completeness result in the same section produces a reduction that in fact can be done in logspace and (with a suitable, but natural, encoding of the gates in a circuit) also in AC^0 , so our result rules out any such reduction for MGS.

Corollary 7.6. *MGS for quasigroups and QUASIGROUP ISOMORPHISM are not $\exists^{\log^2 n} P$ -complete...*

- ...under $quasiAC^0$ Turing reductions;*
- ...under polylog-space Turing reductions unless $EXP = PSPACE$.*

PROOF. (a) As PARITY is in P but not in $quasiAC^0$, problems in $quasiAC^0$ can't be $\exists^{\log^2 n} P$ -complete under $quasiAC^0$ reductions.

(b) Both MGS for quasigroups and QUASIGROUP ISOMORPHISM are in $DSPACE(\log^2 n)$ by Thm. 7.5 (b), resp. [10]. The closure of $DSPACE(\log^2 n)$ under poly-log space reductions is contained in $\text{polyL} = \bigcup_{k \geq 0} DSPACE(\log^k n)$. If either of these two quasigroup problems were complete for $\exists^{\log^2 n} P$ under polylog-space Turing reductions, we would get $\exists^{\log^2 n} P \subseteq \text{polyL}$. Under the latter assumption, by a standard padding argument, we can show that $EXP = PSPACE$. See [12, Cor. 7.6] for full details. \square

Now we return to establishing the main result of this section, Thm. 7.5. To establish Thm. 7.5 (a) and (c), we will crucially leverage the MEMBERSHIP problem for quasigroups. To this end, we will first establish the following.

Theorem 7.7. *MEMBERSHIP for quasigroups belongs to $NTIME(\text{polylog}(n))$.*

Thm. 7.7 immediately yields the following corollary.

Corollary 7.8. *For quasigroups, MEMBERSHIP and MGS are not hard under AC^0 -reductions for any complexity class containing PARITY.*

The proofs of Thms. 7.7 and 7.5 rely crucially on the following adaption of the Babai–Szemerédi Reachability Lemma [5, Thm. 3.1] to quasigroups. We first generalize the notion of a straight-line program for groups [5] to SLPs for quasigroups. Let X be a set of generators for a quasigroup G . We call a sequence of elements $g_1, \dots, g_\ell \in G$ a *straight-line program* (SLP for short) if each g_i ($i \in [\ell]$) either belongs to X , or is of the form or $g_j g_k$, $g_j \setminus g_k$, or g_j / g_k for some $j, k < i$ (recall from §1 that $a \setminus b$ resp. b / a denotes the unique x such that $ax = b$ resp. $xa = b$). An SLP is said to *compute* or *generate* a set S (or an element g) if $S \subseteq \{g_1, \dots, g_\ell\}$ (resp. $g \in \{g_1, \dots, g_\ell\}$).

Lemma 7.9 (Reachability Lemma for quasigroups). *Let G be a finite quasigroup and let X be a set of generators for G . For each $g \in G$, there exists a straight-line program over X generating g which has length $O(\log^2 |G|)$.*

We follow the same strategy as in the proof of [5, Thm. 3.1], but there are some subtle, yet crucial, modifications due to the fact that quasigroups are non-associative and need not possess an identity element.

PROOF. For any sequence of elements z_1, \dots, z_k , let $P(z_1 z_2 \dots z_k)$ denote the left-to-right parenthesization, e.g., $P(z_1 z_2 z_3) = (z_1 z_2) z_3$. For some initial segment z_0, z_1, \dots, z_i define the cube

$$K(i) = \{P(z_0 z_1^{e_1} \dots z_i^{e_i}) : e_1, \dots, e_i \in \{0, 1\}\},$$

where $e_j = 0$ denotes omitting z_j from the product (since there need not be an identity element).

Define $L(i) = K(i) \setminus K(i) = \{g \setminus h : g, h \in K(i)\}$. We will construct a sequence z_0, z_1, \dots, z_t such that $t \leq \lceil \log_2(n) \rceil$ and $L(t) = G$. Moreover, we derive a bound on the straight-line cost $c(i)$ for $\{z_0, z_1, \dots, z_i\}$ ($1 \leq i \leq t$), which is defined as the length of the shortest SLP generating $\{z_0, z_1, \dots, z_i\}$.

We take z_0 as an arbitrary element from X . Hence, $K(0) = \{z_0\}$, and so $c(0) = 1$. Next, let us construct $K(i+1)$ from $K(i)$. If $L(i) \neq G$, we set z_{i+1} to be an element $z' \notin L(i)$ that minimizes $c(i+1) - c(i)$. We first claim that $|K(i+1)| = 2 \cdot |K(i)|$. Note that $K(i+1) = K(i) \cup K(i)z_{i+1}$ by definition. As right-multiplication by a fixed element is a bijection in a quasi-group, it suffices to show that $K(i) \cap K(i)z_{i+1} = \emptyset$. So, suppose that there exists some $a \in K(i) \cap K(i)z_{i+1}$. Then $a = gz_{i+1}$ for some $g \in K(i)$. Hence, $z_{i+1} = g \setminus a$, contradicting $z_{i+1} \notin L(i)$, since both a and g are in $K(i)$.

It now follows that $|K(i)| = 2^i$ and, hence, $t \leq \lceil \log_2(|G|) \rceil$ and $L(t) = G$. Moreover, for every $g \in G$ we obtain an SLP of length at most $c(t) + 2t + 1$: write $g = a \setminus b$ for $a, b \in K(t)$ and start with the SLP computing $\{z_0, z_1, \dots, z_t\}$. We obtain an SLP for $a = z_0 z_1^{e_1} \dots z_t^{e_t}$ by adding a new element for each $z_0 z_1^{e_1} \dots z_i^{e_i}$ with $e_i = 1$ ($1 \leq i \leq t$, thus, at most t new elements). Likewise, we get an SLP for b . Adding a last element $g = a \setminus b$, we obtain an SLP computing g .

It remains to bound the straight-line cost $c(i)$ of $\{z_0, z_1, \dots, z_i\}$. Here, we claim that that $c(i+1) - c(i) \leq 4i + 3$. It follows from this claim that $c(i) \in O(i^2)$. We will now turn to proving our claim:

If $G \neq L(i)$, then either $X \not\subseteq L(i)$ or $L(i)$ is not a sub-quasigroup. Hence, we have one of the following cases:

Case 1: Suppose that there is some $g \in X - L(i)$. In this case, there is an SLP of length one for g and we obtain $c(i+1) - c(i) \leq 1$.

Case 2: Suppose there exist $g, h \in L(i)$ with one of $gh, g/h$, or $g \setminus h \notin L(i)$. For simplicity, suppose $gh \notin L(i)$. The argument is identical for g/h and $g \setminus h$. As above, given an SLP to compute $\{z_0, z_1, \dots, z_i\}$, we may construct SLPs for g and h each of additional length $2i + 1$. This yields an SLP for gh of total length at most $c(i) + 4i + 3$, and shows that $c(i+1) - c(i) \leq 4i + 3$. The result now follows. \square

For proving Thm. 7.7 we follow essentially the ideas of [17] (though we avoid introducing the notion of Cayley circuits). Fleischer obtained a quasiAC⁰ bound for MEMBERSHIP for group by then showing that the Cayley circuits for this problem can be simulated by a quasiAC⁰ circuit. We will instead directly analyze the straight-line programs using an NTIME(polylog(n)) algorithm.

PROOF OF THM. 7.7. To decide whether $g \in \langle X \rangle$, guess the sequence of operations for an SLP of length ℓ , for $\ell \in O(\log^2 |G|)$, which exists iff $g \in \langle X \rangle$ by Lem. 7.9. Computing the values in the SLP and verifying that $g_\ell = g$ can be done in time $\text{poly}(\ell, \log n) \leq \text{polylog}(n)$. \square

The proof of Thm. 7.5 (a) and (c) below is by describing a reduction from MGS to MEMBERSHIP. Thm. 7.5 (b) uses a result of Wagner.

PROOF OF THM. 7.5. (a) Let G denote the input quasigroup (of order n). First, observe that every quasigroup has a generating set of size $\leq \lceil \log n \rceil$ [29]. Therefore, we start by guessing a subset $X \subseteq G$ of size at most $\leq \lceil \log n \rceil$ (resp. the size bound given in the input). For the decision version, we use $O(\log^2 n)$ existentially quantified non-deterministic bits ($\exists^{\log^2 n}$) to guess a generating sequence. To find a minimum-sized generating sequence, we enumerate all possible generating sequences in quasiAC⁰. In the next step, we verify whether X actually generates G . This is done by checking for all $g \in G$ (universally verifying $O(\log n)$ bits, $\forall^{\log n}$) whether $g \in \langle X \rangle$, which can be done in NTIME(polylog(n)) \subseteq quasiAC⁰ by Thm. 7.7.

(c) Proceed as in (a), using the same technique to check that all $Y \subseteq G$ with $|Y| < |X|$ do not generate G .

(b) Existentially guess a generating set using $O(\log^2 n)$ bits, then use [42, Thm. 10.2.1] to compute $\langle X \rangle$ in SAC¹. Then we check, in AC⁰, whether each element of the quasigroup belongs to $\langle X \rangle$. \square

We may similarly reduce QUASIGROUP ISOMORPHISM to MEMBERSHIP for quasigroups. This formalizes the intuition that MEMBERSHIP is an essential subroutine for isomorphism testing and MGS. In particular, in the setting of quasiAC⁰, we have that isomorphism testing and MGS reduce to MEMBERSHIP. This latter consequence might seem surprising, as in the setting of groups, MEMBERSHIP belongs to L, while MGS belongs to AC¹(L) (Thm. 7.1), yet it is a longstanding open problem whether GROUP ISOMORPHISM is even in P.

7.3 MGS for Magmas

Theorem 7.10. *The decision variant of MINIMUM GENERATING SET for magmas is NP-complete.*

This NP-completeness result helps explain the use of Integer Linear Programming in practical heuristic algorithms for the search version of this problem, e.g., [22].

For the closely related problem LOG GENERATORS—given the multiplication table of a magma of order n , decide whether it has a generating set of size $\leq \lceil \log_2 n \rceil$ —LOG GENERATORS OF MAGMAS is $\exists^{\log^2 n}$ P-complete under polynomial-time reductions [31, Thm. 7].

PROOF SKETCH. The problem is in NP by guessing a suitable generating set. To show NP-hardness we reduce from 3SAT. Let $F = \bigwedge_{j=1}^m C_j$ with variables X_1, \dots, X_n be an instance of 3SAT. Our magma M consists of the following elements: an element C_j for $j \in [m]$, elements X_i, \bar{X}_i for $i \in [n]$, elements $T_{j,k}$ for $1 \leq j \leq k \leq m$, and a zero element 0. We define the multiplication as follows:

$$\begin{aligned} C_j X &= T_{j,j} \text{ if the literal } X \text{ appears in } C_j \\ T_{1,m} X_i &= \bar{X}_i, & T_{1,m} \bar{X}_i &= X_i, \\ T_{j,k} T_{k+1,\ell} &= T_{j,\ell}. \end{aligned}$$

all other products are defined as 0.

Any generating set for M must include all the C_j , since without them, there is no way to generate them from other elements. Similarly, any generating set must include, for each $i \in [n]$, at least one of X_i or \bar{X}_i , since they also can't be generated from other elements.

We show that F is satisfiable if and only if M can be generated by $n + m$ elements. When F is satisfiable, let φ be a satisfying assignment, then M is generated by $\{C_j : j \in [m]\} \cup \{X_i : \varphi(x_i) = 1\} \cup \{\bar{X}_i : \varphi(x_i) = 0\}$. Conversely, any generating set of size $n + m$ must consist of all C_j and exactly one of each $\{X_i, \bar{X}_i\}$. As the only way to generate $T_{j,j}$ is for one of the literals to satisfy C_j , the choice of X_i, \bar{X}_i in the generating set must be a satisfying assignment. \square

8 OPEN QUESTIONS

The biggest open question about constant-depth complexity on algebras given by multiplication tables is, in our opinion, still whether or not GROUP ISOMORPHISM is in AC^0 in the Cayley table model. Our results make salient some more specific, and perhaps more approachable, open questions that we now highlight.

Question 8.1. Does MGS for groups belong to L?

Question 8.2. Does MEMBERSHIP for quasigroups belong to L?

The analogous result is known for groups, by reducing to the connectivity problem on Cayley graphs. The best known bound for quasigroups is SAC^1 [42]. Improvements in this direction would yield improvements in MGS for quasigroups. A *constructive* membership test would also yield improvements for isomorphism testing of $O(1)$ -generated quasigroups. Note that isomorphism testing of $O(1)$ -generated groups is known to belong to L [38].

REFERENCES

- [1] Vikraman Arvind and Jacobo Torán. 2006. The Complexity of Quasigroup Isomorphism and the Minimum Generating Set Problem. In *Algorithms and Computation, 17th International Symposium, ISAAC 2006, Kolkata, India, December 18–20, 2006, Proceedings (Lecture Notes in Computer Science, Vol. 4288)*, Tetsuo Asano (Ed.). Springer, 233–242. https://doi.org/10.1007/11940128_25
- [2] James Aspnes, Richard Beigel, Merrick Furst, and Steven Rudich. 1991. The Expressive Power of Voting Polynomials. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing (New Orleans, Louisiana, USA) (STOC '91)*. Association for Computing Machinery, New York, NY, USA, 402–409. <https://doi.org/10.1145/103418.103461>
- [3] László Babai. 1980. On the Complexity of Canonical Labeling of Strongly Regular Graphs. *SIAM J. Comput.* 9, 1 (1980), 212–216. <https://doi.org/10.1137/0209018>
- [4] László Babai. 2016. Graph isomorphism in quasipolynomial time [extended abstract]. In *STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, New York, 684–697. <https://doi.org/10.1145/2897518.2897542> Preprint of full version at arXiv:1512.03547v2 [cs.DS].
- [5] László Babai and Endre Szemerédi. 1984. On the Complexity of Matrix Group Problems I. In *25th Annual Symposium on Foundations of Computer Science, West Palm Beach, Florida, USA, 24–26 October 1984*. IEEE Computer Society, 229–240. <https://doi.org/10.1109/SFCS.1984.715919>
- [6] D.A.M. Barrington. 1992. Quasipolynomial size circuit classes. In *[1992] Proceedings of the Seventh Annual Structure in Complexity Theory Conference*. 86–93. <https://doi.org/10.1109/SCT.1992.215383>
- [7] David A. Mix Barrington, Peter Kaday, Klaus-Jörn Lange, and Pierre McKenzie. 2001. On the Complexity of Some Problems on Groups Input as Multiplication Tables. *J. Comput. Syst. Sci.* 63, 2 (2001), 186–200. <https://doi.org/10.1006/jcss.2001.1764>
- [8] David A. Mix Barrington and Pierre McKenzie. 1991. Oracle branching programs and Logspace versus P. *Inf. Comput.* 95, 1 (1991), 96–115. [https://doi.org/10.1016/0890-5401\(91\)90017-V](https://doi.org/10.1016/0890-5401(91)90017-V)
- [9] David A. Mix Barrington and Howard Straubing. 1994. Complex Polynomials and Circuit Lower Bounds for Modular Counting. *Comput. Complex.* 4 (1994), 325–338. <https://doi.org/10.1007/BF01263421>
- [10] Arkadev Chattopadhyay, Jacobo Torán, and Fabian Wagner. 2013. Graph isomorphism is not AC^0 -reducible to group isomorphism. *ACM Trans. Comput. Theory* 5, 4 (2013), Art. 13, 13. <https://doi.org/10.1145/2540088> Preliminary version appeared in FSTTCS '10; ECCC Tech. Report TR10-117.
- [11] Anthony E. Clement, Stephen Majewicz, and Marcos Zymam. 2017. *The theory of nilpotent groups*. Birkhäuser/Springer, Cham. xviii+307 pages.
- [12] Nathaniel A. Collins, Joshua A. Grochow, Michael Levet, and Armin Weiß. 2024. On the Constant-Depth Circuit Complexity of Generating Quasigroups. arXiv:2402.00133 [cs.CC].
- [13] Bireswar Das and Dhara Thakkar. 2023. Algorithms for the Minimum Generating Set Problem. arXiv:2305.08405 [math.GR]
- [14] Heiko Dietrich, Murray Elder, Adam Piggott, Youming Qiao, and Armin Weiß. 2022. The Isomorphism Problem for Plain Groups Is in Σ_3^P . In *39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022, March 15–18, 2022, Marseille, France (Virtual Conference) (LIPIcs, Vol. 219)*, Petra Berenbrink and Benjamin Monmege (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 26:1–26:14. <https://doi.org/10.4230/LIPIcs.STACS.2022.26>
- [15] Heiko Dietrich and James B. Wilson. 2022. Group isomorphism is nearly-linear time for most orders. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. 457–467. <https://doi.org/10.1109/FOCS52979.2021.00053>
- [16] Flavio Ferrarotti, Senén González, Klaus-Dieter Schewe, and José María Turull-Torres. 2020. Proper Hierarchies in Polylogarithmic Time and Absence of Complete Problems. In *Foundations of Information and Knowledge Systems*, Andreas Herzig and Juha Kontinen (Eds.). Springer International Publishing, Cham, 90–105.
- [17] Lukas Fleischer. 2022. The Cayley Semigroup Membership Problem. *Theory of Computing* 18, 8 (2022), 1–18. <https://doi.org/10.4086/toc.2022.v018a008>
- [18] Joshua A. Grochow and Michael Levet. 2023. On the Parallel Complexity of Group Isomorphism via Weisfeiler-Leman. In *Fundamentals of Computation Theory - 24th International Symposium, FCT 2023, Trier, Germany, September 18–21, 2023, Proceedings (Lecture Notes in Computer Science, Vol. 14292)*, Henning Fernau and Klaus Jansen (Eds.). Springer, 234–247. https://doi.org/10.1007/978-3-031-43587-4_17 Preprint of full version at arXiv:2112.11487 [cs.DS].
- [19] Joshua A. Grochow and Youming Qiao. 2017. Algorithms for group isomorphism via group extensions and cohomology. *SIAM J. Comput.* 46, 4 (2017), 1153–1216. <https://doi.org/10.1137/15M1009767> Preliminary version in IEEE Conference on Computational Complexity (CCC) 2014 (DOI:10.1109/CCC.2014.19). Also available as arXiv:1309.1776 [cs.DS] and ECCC Technical Report TR13-123..
- [20] Harald Andrés Helfgott, Jitendra Bajpai, and Daniele Dona. 2017. Graph isomorphisms in quasi-polynomial time. <https://doi.org/10.48550/ARXIV.1710.04574>
- [21] Jeffrey C. Jackson, Adam R. Klivans, and Rocco A. Servedio. 2002. Learnability beyond AC^0 . In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing (Montreal, Quebec, Canada) (STOC '02)*. Association for Computing Machinery, New York, NY, USA, 776–784. <https://doi.org/10.1145/509907.510018>
- [22] Mikolás Janota, António Morgado, and Petr Vojtechovský. 2023. Computing generating sets of minimal size in finite algebras. *J. Symb. Comput.* 119 (2023), 50–63. <https://doi.org/10.1016/j.jsc.2023.02.002>
- [23] Emil Jerábek. 2020. Answer to “Is Abelian group isomorphism in AC^0 ?”. Theoretical Computer Science Stack Exchange. <https://cstheory.stackexchange.com/q/40118> (version: 2020-06-17).
- [24] François Le Gall and David J. Rosenbaum. 2016. On the Group and Color Isomorphism Problems. arXiv:1609.08253 [cs.CC].
- [25] Michael Levet. 2023. On the Complexity of Identifying Strongly Regular Graphs. *Australasian Journal of Combinatorics* 87 (2023), 41–67. https://ajc.maths.uq.edu.au/pdf/87/ajc_v87_p041.pdf
- [26] R. J. Lipton, L. Snyder, and Y. Zalcstein. 1977. The Complexity of Word and Isomorphism Problems for Finite Groups. Yale University Dept. of Computer Science Research Report # 91. <https://apps.dtic.mil/dtic/tr/fulltext/u2/a053246.pdf>
- [27] Andrea Lucchini and Federico Menegazzo. 1994. Computing a Set of Generators of Minimal Cardinality in a Solvable Group. *Journal of Symbolic Computation* 17, 5 (1994), 409–420. <https://doi.org/10.1006/jsc.1994.1027>
- [28] Andrea Lucchini and Dhara Thakkar. 2024. The minimum generating set problem. *Journal of Algebra* 640 (2024), 117–128. <https://doi.org/10.1016/j.jalgebra.2023.11.012>
- [29] Gary L. Miller. 1978. On the $n^{\log n}$ Isomorphism Technique (A Preliminary Report). In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing (San Diego, California, USA) (STOC '78)*. Association for Computing Machinery, New York, NY, USA, 51–58. <https://doi.org/10.1145/800133.804331>
- [30] A. Neumaier. 1979. Strongly regular graphs with smallest eigenvalue $-m$. *Archiv der Mathematik* 33 (1979), 392–400. <https://doi.org/10.1007/BF01222774>
- [31] Christos H. Papadimitriou and Mihalis Yannakakis. 1996. On Limited Nondeterminism and the Complexity of the V-C Dimension. *J. Comput. Syst. Sci.* 53, 2 (1996), 161–170. <https://doi.org/10.1006/JCSS.1996.0058>
- [32] Alexander A Razborov. 1993. An Equivalence between Second Order Bounded Domain Bounded Arithmetic and First Order Bounded Arithmetic. In *Arithmetic, proof theory, and computational complexity*. Oxford University Press. <https://doi.org/10.1093/oso/9780198536901.003.0012>

- [33] Omer Reingold. 2008. Undirected Connectivity in Log-Space. *J. ACM* 55, 4, Article 17 (sep 2008), 24 pages. <https://doi.org/10.1145/1391289.1391291>
- [34] D. Robinson. 1982. *A Course in the Theory of Groups*. Springer.
- [35] David J. Rosenbaum. 2013. Bidirectional Collision Detection and Faster Deterministic Isomorphism Testing. arXiv:1304.3935 [cs.DS].
- [36] Jean-Pierre Serre. 1997. *Galois Cohomology* (1 ed.). Springer Berlin, Heidelberg. <https://doi.org/10.1007/978-3-642-59141-9>
- [37] Roman Smolensky. 1987. Algebraic Methods in the Theory of Lower Bounds for Boolean Circuit Complexity. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing, 1987, New York, New York, USA*, Alfred V. Aho (Ed.). ACM, 77–82. <https://doi.org/10.1145/28395.28404>
- [38] Bangsheng Tang. 2013. *Towards Understanding Satisfiability, Group Isomorphism and Their Connections*. Ph. D. Dissertation. Tsinghua University. http://papakonstantinou.org/periklis/pdfs/bangsheng_thesis.pdf
- [39] D. R. Taunt. 1955. Remarks on the isomorphism problem in theories of construction of finite groups. *Mathematical Proceedings of the Cambridge Philosophical Society* 51, 1 (1955), 16–24. <https://doi.org/10.1017/S030500410002987X>
- [40] Jacobo Torán. 2004. On the Hardness of Graph Isomorphism. *SIAM J. Comput.* 33, 5 (2004), 1093–1108. <https://doi.org/10.1137/S009753970241096X>
- [41] Heribert Vollmer. 1999. *Introduction to Circuit Complexity - A Uniform Approach*. Springer. <https://doi.org/10.1007/978-3-662-03927-4>
- [42] F. Wagner. 2010. *On the complexity of isomorphism testing for restricted classes of graphs*. Ph. D. Dissertation. Universität Ulm. https://oparu.uni-ulm.de/xmlui/bitstream/handle/123456789/3923/vts_7264_10267.pdf
- [43] Marty J. Wolf. 1994. Nondeterministic circuits, space complexity and quasigroups. *Theoretical Computer Science* 125, 2 (1994), 295–313. [https://doi.org/10.1016/0304-3975\(92\)00014-1](https://doi.org/10.1016/0304-3975(92)00014-1)

Received 6 February 2024; revised 12 March 2009; accepted 5 June 2009