

GROUP, GRAPHS, ALGORITHMS: THE GRAPH ISOMORPHISM PROBLEM

László Babai

Abstract

Graph Isomorphism (GI) is one of a small number of natural algorithmic problems with unsettled complexity status in the P / NP theory: not expected to be NP-complete, yet not known to be solvable in polynomial time.

Arguably, the GI problem boils down to filling the gap between *symmetry* and *regularity*, the former being defined in terms of automorphisms, the latter in terms of equations satisfied by numerical parameters.

Recent progress on the complexity of GI relies on a combination of the *asymptotic theory of permutation groups* and asymptotic properties of highly regular combinatorial structures called *coherent configurations*. Group theory provides the tools to infer either global symmetry or global irregularity from local information, eliminating the symmetry/regularity gap in the relevant scenario; the resulting global structure is the subject of combinatorial analysis. These structural studies are melded in a *divide-and-conquer* algorithmic framework pioneered in the GI context by Eugene M. Luks (1980).

1 Introduction

We shall consider finite structures only; so the terms “graph” and “group” will refer to finite graphs and groups, respectively.

1.1 Graphs, isomorphism, NP-intermediate status. A *graph* is a set (the set of *vertices*) endowed with an irreflexive, symmetric binary relation called *adjacency*. Isomorphisms are adjacency-preserving bijections between the sets of vertices. The Graph Isomorphism (GI) problem asks to determine whether two given graphs are isomorphic.

It is known that graphs are *universal* among explicit finite structures in the sense that the isomorphism problem for explicit structures can be reduced in polynomial time to GI (in the sense of *Karp-reductions*¹) [Hedrlín and Pultr \[1966\]](#) and [Miller \[1979\]](#). This

MSC2010: primary 68Q25; secondary 20B05, 05B30, 05C68, 05C60, 05E30.

¹For basic concepts of complexity theory we refer to [Garey and Johnson \[1979\]](#).

makes GI a natural algorithmic problem. It is a *polynomial-time verifiable* problem: a candidate isomorphism is easily verified. This puts GI in the complexity class NP. Over time, increasingly strong conjectural evidence has been found that GI is not NP-complete, yet no polynomial-time algorithm is known to solve GI. This puts GI among the small number of natural NP-problems of potentially *intermediate complexity* (neither in P, nor NP-complete). (Another such problem is that of *factoring integers*, cf. [Section 11](#).) The interest in this status of GI was recognized at the dawn of the P / NP theory [Karp \[1972\]](#) and [Garey and Johnson \[1979\]](#).

1.2 Brief history of the GI problem. Combinatorial heuristics such as individualization and refinement (I/R) (see [Section 8](#)) have been used for the longest time to reduce the GI search space. It was shown that the “naive refinement” algorithm solves GI for almost all graphs in linear time [Babai, Erdős, and Selkow \[1980\]](#) and [Babai and Kucera \[1979\]](#). Efficient algorithms were found for special classes such as planar graphs [J. E. Hopcroft and Tarjan \[1972\]](#) and [J. E. Hopcroft and Wong \[1974\]](#). These algorithms exploited the combinatorial structure of the graphs concerned. However, combinatorial refinement methods alone cannot succeed in less than exponential time for the general GI problem, as shown in a seminal 1992 paper by [Cai, Fürer, and Immerman \[1992\]](#).

It has long been known that GI is equivalent to determining whether two vertices of a given graph belong to the same orbit of the automorphism group. Refinement procedures have been used to distinguish vertices, trying to refute *symmetry* by discovering *irregularity*. While this gives a first indication of the critical role of the gap between symmetry and regularity to GI, the CFI result shows the futility of trying to close this gap using combinatorial refinement heuristics alone. We use group theory to close a gap of this nature under particular circumstances (see [Theorem 5.3](#) and the paragraph preceding it). The relevant new group theoretic result, the “Unaffected Stabilizers Lemma,” is stated in [Theorem 6.2](#).

Elements of group theory were first introduced into the design of GI algorithms in 1979 [Babai \[1979\]](#). The *tower of groups* method described in that paper produced the following results. A *vertex-colored* graph has a “color” assigned to each vertex; isomorphisms preserve the colors by definition. The *multiplicity* of a color is the number of vertices of that color. The *adjacency matrix* of a graph with n vertices is the $n \times n$ $(0, 1)$ -matrix whose (i, j) -entry is 1 if vertex i is adjacent to vertex j , and 0 otherwise. By the *eigenvalues of a graph* we mean the eigenvalues of its adjacency matrix.

Theorem 1.1. (a) [Babai \[1979\]](#) and [Furst, J. Hopcroft, and E. Luks \[1980\]](#) *Isomorphism of vertex-colored graphs of bounded color multiplicities can be tested in polynomial time.*
 (b) [Babai, Grigoryev, and Mount \[1982\]](#) *Isomorphism of graphs with bounded eigenvalue multiplicities can be tested in polynomial time.*

It turns out that the CFI pairs of graphs, i. e., the pairs of graphs shown in [Cai, Fürer, and Immerman \[1992\]](#) to be hard to separate by combinatorial refinement, can be viewed as vertex-colored graphs with color multiplicity 4. This shows that elementary group theory (hardly more than the concept of cosets was used) was already capable of overcoming exponential barriers to combinatorial refinement methods. Modern extensions of the CFI result show that GI is hard for several more general refutation systems (see [Section 11](#)), putting GI in a somewhat paradoxical position in complexity theory (cf. [Section 11](#)).

In-depth use of group theory in the design of GI algorithms arrived with Luks's ground-breaking 1980 paper [E. M. Luks \[1982\]](#). We state the main result of that paper. Adjacent vertices of a graph are called *neighbors*; the *degree* of a vertex is the number of its neighbors.

Theorem 1.2 (Luks, 1980). *Isomorphism of graphs of bounded degree can be tested in polynomial time.*

Luks's group theoretic method, combined with a combinatorial refinement result by [Zemlyachenko, Korneenko, and Tyshkevich \[1982\]](#), have lead to the *moderately exponential* complexity bound of

$$(1) \quad \exp(O(\sqrt{n \log n})),$$

where n denotes the number of vertices (Luks, 1983, cf. [Babai and E. M. Luks \[1983\]](#) and [Babai, Kantor, and E. M. Luks \[1983\]](#)). In spite of intermittent progress on important special cases, notably for *strongly regular graphs* [Spielman \[1996\]](#), [Chen, Sun, and Teng \[2013\]](#), [Babai and Wilmes \[2013\]](#), and [Babai, Chen, Sun, Teng, and Wilmes \[2013\]](#) and for *primitive coherent configurations* [Sun and Wilmes \[2015\]](#), Luks's bound (1) for the general case had not been improved until this author's recent announcement [Babai \[2016\]](#) of a quasipolynomial-time algorithm. A *quasipolynomial* function is a function of the form $\exp(p(\log n))$ for some polynomial p . A quasipolynomial time bound is a bound of this form where n is the bit-length of the input; but if we take n to be the number of vertices of an input graph, the form of the bound will not be affected.

Theorem 1.3 (B 2015). *Isomorphism of graphs can be tested in quasipolynomial time.*

In this paper we outline the main components of this result. For an introduction to the algorithmic theory of permutation groups we refer to the monograph [Seress \[2003\]](#).

Disclaimer. I should emphasize that the results discussed in this paper address the mathematical problem of the *asymptotic worst-case complexity* of GI and have little relevance to practical computation. A suite of remarkably efficient GI packages is available for practical GI testing; [McKay and Piperno \[2014\]](#) give a detailed comparison of methods and performance. These algorithms employ ingenious shortcuts to backtrack search. While

the worst-case performance of these heuristics seems to be exponential, this is increasingly difficult to demonstrate, cf. [Cai, Fürer, and Immerman \[1992\]](#), [Miyazaki \[1996\]](#), and [Neuen and Schweitzer \[2017\]](#).

2 The string isomorphism problem

We now define a generalization of the GI problem, introduced by [E. M. Luks \[1982\]](#).

Let Ω be a finite set; $\text{Sym}(\Omega)$ denotes the symmetric group acting on Ω . Let Σ be finite alphabet. An Ω -string (or just “string”) over Σ is a function $x : \Omega \rightarrow \Sigma$. There is a natural action $x \mapsto x^\sigma$ of $\text{Sym}(\Omega)$ on the set Σ^Ω of strings ($\sigma \in \text{Sym}(\Omega)$, $x \in \Sigma^\Omega$). We say that $\sigma \in \text{Sym}(\Omega)$ is a *G-isomorphism* between the strings x and y if $\sigma \in G$ and $x^\sigma = y$. The strings x and y are *G-isomorphic*, denoted $x \cong_G y$, if such a σ exists. The *String Isomorphism (SI) problem* asks, given G , x , and y , does $x \cong_G y$ hold? We refer to G as the *ambient group*; it is given by a list of generators.

Luks pointed out [E. M. Luks \[ibid.\]](#) that GI reduces to SI by encoding each graph X by the characteristic function f_X of its adjacency relation, $f_X : \binom{\Omega}{2} \rightarrow \{0, 1\}$, where $\binom{\Omega}{2}$ denotes the set of unordered pairs of elements of Ω . So f_X is an $\binom{\Omega}{2}$ -string over the alphabet $\{0, 1\}$. The pertinent ambient group is $\text{Sym}(\Omega)^{(2)}$, the induced action of $\text{Sym}(\Omega)$ on the set $\binom{\Omega}{2}$. It is easy to see that two graphs are isomorphic if and only if the corresponding $\binom{\Omega}{2}$ -strings are $\text{Sym}(\Omega)^{(2)}$ -isomorphic. The actual result we shall discuss concerns the complexity of SI [Babai \[2016\]](#).

Theorem 2.1 (B 2015). *String isomorphism can be tested in quasipolynomial time.*

[Theorem 1.3](#) is then a corollary. The previous best bound for SI was $\exp(\tilde{O}(n^{1/2}))$, where $n = |\Omega|$ is the length of the strings in question [Babai \[1983\]](#) (cf. [Babai, Kantor, and E. M. Luks \[1983\]](#)). (The tilde hides a polylogarithmic factor.)

Luks also observed that several other problems of computational group theory are polynomial-time equivalent to SI (under Karp-reductions), including the coset intersection, double coset membership, and ‘centralizer in coset’ problems. Given two subgroups G, H of the symmetric group S_n and two elements $\sigma, \pi \in S_n$, the *Coset Intersection* problem asks whether $G\sigma \cap H\pi \neq \emptyset$; the *double coset membership* problem asks whether $\sigma \in G\pi H$, and the *centralizer in coset* problem asks whether there exists an element in the coset $G\sigma$ that commutes with π . As a consequence, these problems, too, can be solved in quasipolynomial time.

The advantage of approaching GI through the SI problem is that SI permits recursion on the ambient group. This was Luks’s core idea.

3 Divide-and-Conquer

In the theory of algorithms, the term “Divide-and-Conquer” refers to recursive procedures that reduce an instance of a computational problem to a moderate number of significantly smaller instances. If our input has size n , we shall consider instances of size $\leq 0.9n$ to be “significantly smaller.” Let $q(n)$ be the number of such smaller instances to which our input is reduced; we refer to $q(n)$ as the *multiplicative cost* of the reduction. If $f(n)$ denotes the worst-case cost of processing an input of size n , this leads to the following recurrence (ignoring the additive cost of assembling all information from the smaller instances, which will typically not affect the cost estimate).

$$(2) \quad f(n) \leq q(n)f(0.9n)$$

Assuming that $q(n)$ is monotone, this gives the bound $f(n) \leq q(n)^{O(\log n)}$, so if $q(n)$ is quasipolynomially bounded then so is $f(n)$. Therefore our goal will be to significantly reduce the problem size at a quasipolynomial multiplicative cost.

4 Large primitive permutation groups

Not only did Luks point out that GI reduces to SI, but he also showed that (i) the SI problem for groups with restricted structure can be used to solve the GI problem for certain classes of graphs; and that (ii) SI can be solved efficiently under such structural constraints. The issue of relevance here is bounding the order of primitive permutation groups under structural constraints.

A *permutation group* acting on the set Ω (the *permutation domain*) is a subgroup $G \leq \text{Sym}(\Omega)$. (The “ \leq ” sign stands for “subgroup.”) The *degree* of G is $|\Omega|$. The set $x^G = \{x^\sigma \mid \sigma \in G\}$ is the G -*orbit* of x ; the orbit has *length* $|x^G|$. We say that G is *transitive* if $x^G = \Omega$ for some (and therefore any) $x \in \Omega$. A transitive group $G \leq \text{Sym}(\Omega)$ is *primitive* if $|\Omega| \geq 2$ and there is no nontrivial G -invariant equivalence relation on Ω .

In 1982, Pálffy [1982] and Wolf [1982] showed that primitive *solvable* groups of degree n have order $\leq n^c$ where $c \approx 3.243$. It turns out that the critical structural parameter of a group for polynomial bounds on the order of its primitive permutation representations is its “thickness.”

Definition 4.1. The *thickness*² $\theta(G)$ of a group G is the largest t such that the alternating group A_t is involved in G as a quotient of a subgroup.

The following result characterizes those hereditary classes of groups (classes that are closed under subgroups and quotients) which have only small primitive permutation representations.

²The term “thickness” was coined in Babai [2014].

Theorem 4.2 (B, Cameron, Pálffy, 1982). *If G is a primitive permutation group of degree n and thickness t then $|G| = n^{O(t)}$.*

This result first appeared in Babai, Cameron, and Pálffy [1982]; here it is stated with an improved exponent due to Pyber [1990]. Refined versions were subsequently obtained by Liebeck, Shalev, Maróti; see Liebeck and Shalev [2003, Sec. 3] for a survey of those developments. We note that while the initial motivation for Theorem 4.2 came from the GI problem, the result also found applications in other areas, such as the theory of profinite groups Borovik, Pyber, and Shalev [1996].

E. M. Luks [1982] introduced a group theoretic divide-and-conquer technique to attack the SI problem. Luks's method, combined with the above bounds, yields the following.

Corollary 4.3. *The SI problem can be solved in polynomial time if the ambient group is solvable or more generally, if it has bounded thickness.*

Let G be the stabilizer of an edge in the automorphism group of a connected graph in which every vertex has degree $\leq k$. It is easy to see that every composition factor of G is a subgroup of the symmetric group S_{k-1} . In particular, $\theta(G) \leq k-1$ and therefore the SI problem can be solved in polynomial time for such G as the ambient group. This fact is at the heart³ of the proof of Theorem 1.2.

While Theorem 4.2 is helpful for groups with small thickness, our interest is in the general case. Luks's technique for SI works in quasipolynomial time as long as the primitive groups involved in the ambient group have quasipolynomially bounded orders. In 1981, building on the then expected completion of the classification of the finite simple groups (CFSG), Cameron [1981] gave a precise characterization of primitive groups of large order. The socle of a group is the product of its minimal normal subgroups. It is known that the socle of a primitive permutation group is a direct product of isomorphic simple groups. For a permutation group $T \leq \text{Sym}(\Delta)$, the product action of the direct power T^k on the Cartesian power Δ^k is the independent action of each copy of T on the corresponding coordinate. Wreath product in addition permutes the coordinates by some group “on the top.” For a permutation group $G \leq \text{Sym}(\Omega)$ we denote by $G^{(t)}$ the induced action of G on the set $\binom{\Omega}{t}$ of unordered t -tuples of elements of Ω .

Definition 4.4. $G \leq S_n$ is a *Cameron group* with parameters $s, t \geq 1$ and $k \geq \max(2t + 1, 5)$ if we have $n = \binom{k}{t}^s$, the socle of G is isomorphic to A_k^s and acts as $(A_k^{(t)})^s$ in the product action, and $(A_k^{(t)})^s \leq G \leq S_k^{(t)} \wr S_s$ (wreath product, product action), moreover the induced action $G \rightarrow S_s$ on the direct factors of the socle is transitive.

³Theorem 4.2 was not available to Luks at the time; he used a further layer of recurrence so a weaker group-theoretic result was sufficient for his analysis E. M. Luks [1982].

Theorem 4.5 (Cameron 1981). *For $n \geq 25$, if $G \leq S_n$ has order $|G| \geq n^{1+\log_2 n}$ then G is a Cameron group.*

This sharp version of Cameron's theorem [Cameron \[ibid.\]](#) is due to [Maróti \[2002\]](#).

5 Luks's method and the bottleneck

In attacking the SI problem, Luks applies a combination of the following two types of recursive operations to the ambient group.

- Descend to a subgroup.
- Process orbits one by one.

Orbit-by-orbit processing leads to ultra-efficient (linear-time) recurrence. Descent to a subgroup $H \leq G$ incurs a heavy penalty, namely, a multiplicative cost of $|G : H|$, so this can only be used to replace the ambient group with a subgroup of small index, and to compensate for the multiplicative cost, such a step needs to lead to significantly reduced problem size. Small primitive groups acting on a minimal system of imprimitivity (system of maximal blocks of a G -invariant equivalence relation) provide such an opportunity; the orbits of the kernel of the action of such a primitive group have length $\leq n/2$, hence orbit-by-orbit processing reduces the problem to significantly smaller instances.

Using [Theorem 4.5](#) we can identify the bottleneck for Luks's method.

Definition 5.1. We say that a group G has a *giant quotient of degree m* if G has an epimorphism onto S_m or A_m .

Proposition 5.2. *For any constant $C \geq 1$ one can use Luks recurrence for the SI problem to achieve one of the following at a multiplicative cost of $n^{O(\log n)}$.*

- (a) *Significantly reduce the problem size.*
- (b) *Reduce the ambient group to a transitive group with a giant quotient of degree $\geq C \log n$.*

Our work addresses case (b), the bottleneck situation. The goal is to either confirm or effectively break the symmetry represented by the giant quotient. This inserts another layer of recurrence into Luks's framework: significant reduction of m , the degree of the giant quotient.

More specifically, let $G \leq \text{Sym}(\Omega)$ be our ambient group and $x, y : \Omega \rightarrow \Sigma$ be two strings of which we wish to determine the G -isomorphisms. Let, further, $\varphi : G \rightarrow H$ be an epimorphism where $\text{Alt}(\Gamma) \leq H \leq \text{Sym}(\Gamma)$ for some large set Γ , where $\text{Alt}(\Gamma)$ denotes the alternating group (even permutations of Γ). Let $m = |\Gamma|$ and let $P(x) =$

$\varphi(\text{Aut}_G(x)) \leq \text{Sym}(\Gamma)$; define $P(y)$ analogously. We say that a group $K \leq \text{Sym}(\Psi)$ is a *giant* on Ψ if $\text{Alt}(\Psi) \leq K \leq \text{Sym}(\Psi)$.

Theorem 5.3 (Canonical obstruction to symmetry). *Either $P(x)$ acts as a giant on a P -orbit of length $\geq 0.9m$, or there exists a $P(x)$ -invariant canonical k -ary relational structure $\mathcal{X}(x)$ on Γ with $k = O(\log n)$ such that $\mathcal{X}(x)$ has symmetry defect > 0.1 . Moreover, in each case, we can find, via efficient Luks recurrences, an effective representation of the stated objects.*

We explain the concepts involved in this statement.

By ‘efficient Luks recurrence’ we mean a sequence of Luks operations that significantly reduces the problem size at a multiplicative cost of $n^{O(\log n)}$.

In the first case, ‘effective representation’ means we can find a subgroup $M \leq \text{Aut}_G(x)$ such that $\varphi(M)$ has a large orbit on which it acts as a giant. Note that $\text{Aut}_G(x)$ is not known; in fact, determining $\text{Aut}_G(x)$ is equivalent to the SI problem.

We need to explain the second case. A k -ary relation on a set Γ is a subset of the Cartesian power Γ^k . A k -ary relational structure on Γ is a pair $\mathcal{X} = (\Gamma, \mathcal{R})$ where $\mathcal{R} = (R_1, \dots, R_r)$ is a list of k -ary relations R_i on Γ . ‘Effective representation’ of \mathcal{X} simply means listing each R_i . We may assume the R_i are disjoint, so the total length of the lists is $\leq m^k$.

We say that the *symmetry defect* of \mathcal{X} is $\geq \alpha$ if every orbit of $\text{Aut}(\mathcal{X})$ on which $\text{Aut}(\mathcal{X})$ acts as a giant has size $\leq (1 - \alpha)m$.

Canonicity of the $x \mapsto \mathcal{X}(x)$ assignment means this construction is a *functor* from the category of G -isomorphisms of strings in the set $\{x, y\}$ (two objects) to the category of isomorphisms of k -ary relational structures on Γ , so every G -isomorphism $\beta_1 \rightarrow \beta_2$ ($\beta_i \in \{x, y\}$) induces an isomorphism $\mathcal{X}(\beta_1) \rightarrow \mathcal{X}(\beta_2)$.

The two cases listed in [Theorem 5.3](#) are mutually exclusive by the definition of symmetry defect. The result provides a *constructive obstruction* to certain type of very large symmetry (small symmetry defect); the structure \mathcal{X} has sufficient *irregularity* to preclude such large symmetry. This is the sense in which, under our special circumstances, we have been able to close a *symmetry vs. regularity gap* (see [Section 1](#)), a key step toward [Theorem 2.1](#).

6 Unaffected Stabilizers Lemma

In this section we state a group theoretic result, [Theorem 6.2](#) (a), that is our main mathematical (non-algorithmic) tool for the proof of [Theorem 5.3](#).

For a group $G \leq \text{Sym}(\Omega)$ and $x \in \Omega$, the *stabilizer* of x in G is the subgroup $G_x = \{\sigma \in G \mid x^\sigma = x\}$. For $\Delta \subseteq \Omega$, the *pointwise stabilizer* of Δ is the subgroup $G_{(\Delta)} = \bigcap_{x \in \Delta} G_x$.

For a group G and a set Γ we say that the action $\varphi : G \rightarrow \text{Sym}(\Gamma)$ is a *giant representation* of G (or a *giant homomorphism*) if the image $\varphi(G)$ is a giant, i. e., $\varphi(G) \geq \text{Alt}(\Omega)$. We now define our central new concept.

Definition 6.1 (Affected). Let Ω and Γ be sets, $G \leq \text{Sym}(\Omega)$, and let $\varphi : G \rightarrow \text{Sym}(\Gamma)$ be a giant representation. We say that $x \in \Omega$ is *affected* by φ if the φ -image of the stabilizer G_x is not a giant, i. e., $\varphi(G_x) \not\geq \text{Alt}(\Gamma)$.

We note that if $x \in \Omega$ is affected then every element of the orbit x^G is affected. So we can speak of *affected orbits*.

Theorem 6.2. Let $G \leq \text{Sym}(\Omega)$ be a permutation group of degree $n = |\Omega|$ and $\varphi : G \rightarrow S_k$ a giant representation, i. e., $\varphi(G) \geq A_k$. Let $U \subseteq \Omega$ denote the set of elements of Ω not affected by φ . Then the following hold.

- (a) (Unaffected Stabilizers Lemma) Assume $k > \max\{8, 2 + \log_2 n\}$. Then φ restricted to $G_{(U)}$, the pointwise stabilizer of U , is still a giant representation, i. e., $\varphi(G_{(U)}) \geq A_k$. In particular, $U \neq \Omega$ (at least one element is affected).
- (b) (Affected Orbit Lemma) Assume $k \geq 5$. If Δ is an affected G -orbit, i. e., $\Delta \cap U = \emptyset$, then $\ker(\varphi)$ is not transitive on Δ ; in fact, each orbit of $\ker(\varphi)$ in Δ has length $\leq |\Delta|/k$.

The affected/unaffected dichotomy underlies the core “local certificates” algorithm (Section 7).

Part (b) is an easy exercise; its significance is that it permits efficient Luks reductions on affected orbits.

Part (a) is the central result mentioned. The proof of part (a) builds on the O’Nan–Scott–Aschbacher characterization of primitive permutation groups (L. L. Scott [1980] and Aschbacher and L. Scott [1985], cf. Dixon and Mortimer [1996, Thm. 4.1A]) and depends on the classification of Finite Simple Groups (CFSG)⁴ through Schreier’s Hypothesis (a consequence of CFSG) that asserts that the outer automorphism group of every finite simple group is solvable.

Note that part (a) is counter-intuitive: it asserts that if the stabilizer of each $x \in U$ maps onto A_k or S_k then even the intersection of these stabilizers maps onto A_k or S_k .

The condition $k > 2 + \log_2 n$ in part (a) is tight. In fact, there are infinitely many examples with $k = 2 + \log_2 n$ which have *no affected points*, as shown by the example of a semidirect product $\mathbb{Z}_2^{k-2} \rtimes A_k \leq \text{AGL}(k-2, 2)$ for even k , acting on $n = 2^{k-2}$ elements.

⁴A less tight version of the lemma, still sufficient for the quasipolynomial claim, was recently proved by Pyber [2016] without the CFSG.

7 Local certificates

In this section we describe our core algorithmic result. The goal is to categorize ordered k -tuples of Γ , setting the stage for a combinatorial analysis of the resulting k -ary relational structure. The method requires the construction of global automorphisms from local information; our key tool is the Unaffected Stabilizers Lemma.

We consider the Luks bottleneck situation. The input is a transitive group $G \leq \text{Sym}(\Omega)$, a giant representation $\varphi : G \rightarrow \text{Alt}(\Gamma)$, and two strings $x, y : \Omega \rightarrow \Sigma$. We write $n = |\Omega|$ and $m = |\Gamma|$. We fix a number $k > 2 + \log_2 n$ (but not much greater, e. g., $k = 3 + \lfloor \log_2 n \rfloor$) and assume $m \geq 10k$. Subsets $T \subset \Gamma$ of size $|T| = k$ will be referred to as “test sets.”

If $L \leq G$ then L also acts on Γ via φ so for a test set T we can speak of the setwise stabilizer of T in L ; we write L_T for this subgroup.

We say that T is L -invariant if $L_T = L$. We write $\psi_T : G_T \rightarrow \text{Sym}(T)$ for the map that restricts the domain of φ to G_T and the codomain to $\text{Sym}(T)$. The group G_T can be computed in polynomial time as $G_T = \varphi^{-1}(\text{Sym}(\Gamma)_T)$. Our focus is the (unknown) group $P(T) := \psi_T(\text{Aut}_{G_T}(x))$.

Definition 7.1 (Fullness). Let T be a test set. We say that T is *full* with respect to the input string x if $P(T) \geq \text{Alt}(T)$, i. e., the G -automorphisms of x induce a giant on T .

We consider the problem of deciding whether a given test set is full and compute useful certificates of either outcome. We show that this question can efficiently (in time $k! n^{O(1)}$) be reduced to the String Isomorphism problem on inputs of size $\leq n/k$.

Certificate of non-fullness. We certify non-fullness of the test set T by computing a permutation group $M(T) \leq \text{Sym}(T)$ such that (i) $M(T) \not\geq \text{Alt}(T)$ and (ii) $M(T) \geq P(T)$ ($M(T)$ is guaranteed to contain the projection of the G -automorphism group of x). Such an “encasing group” $M(T)$ can be thought of as a constructive refutation of fullness.

Certificate of fullness. We certify fullness of the test set T by computing a permutation group $K(T) \leq \text{Sym}(\Omega)$ such that (i) $K(T) \leq \text{Aut}_{G_T}(x)$ and (ii) $\psi_T(K(T)) \geq \text{Alt}(T)$. Note that $K(T) \leq P(T)$, so $K(T)$ represents a polynomial-time verifiable proof of fullness of T .

Our ability to find $K(T)$, the certificate of fullness, may be surprising because it means that from a local start (that may take only a small segment of x into account), we have to build up global automorphisms (automorphisms of the full string x). Our ability to do so critically depends on the “Unaffected Stabilizers Lemma” (Theorem 6.2 (a)).

Theorem 7.2 (Local certificates). *Let $T \subseteq \Gamma$ where $|T| = k$ is a test set. Assume $\max\{8, 2 + \log_2 n\} < k \leq m/10$ (where $m = |\Gamma|$). By making $\leq k! n^2$ calls to SI problems on domains of size $\leq n/k$ and performing $k! n^{O(1)}$ computation we can decide whether T is full and*

- (a) if T is full, find a certificate $K(T) \leq \text{Aut}_G(\mathfrak{x})$ of fullness
 (b) if T is not full, find a certificate $M(T) \leq \text{Sym}(T)$ of non-fullness.

To aggregate the local certificates, first we consider the group F generated by the fullness certificates. If the support of $\varphi(F) \leq \text{Sym}(\Gamma)$ has at least $m/10$ elements then the structure of $\varphi(F)$ suffices for the proof of [Theorem 5.3](#). In the alternative, non-fullness certificates dominate. In this case a slight extension of [Theorem 7.2](#) is needed, to encase not only the group $\psi_T(\text{Aut}_{G_T}(\mathfrak{x}))$ but also the images of the cosets $\text{Iso}_{G_{T,T'}}(\mathfrak{z}_1, \mathfrak{z}_2)$ for all pairs T, T' of test sets and all choices of $\mathfrak{z}_1, \mathfrak{z}_2 \in \{\mathfrak{x}, \mathfrak{y}\}$. The result will be two classifications of the ordered k -tuples of Γ , one associated with \mathfrak{x} , the other with \mathfrak{y} , yielding the canonical assignment $\mathfrak{x} \mapsto \mathfrak{X}(\mathfrak{x})$ and $\mathfrak{y} \mapsto \mathfrak{X}(\mathfrak{y})$.

8 Individualization and refinement

We consider k -ary *partition structures* $\mathfrak{X} = (\Gamma, \mathfrak{R})$ where $\mathfrak{R} = (R_1, \dots, R_r)$ is a partition of Γ^k . We think of such a structure as a *coloring* $c : \Gamma^k \rightarrow \{1, \dots, r\}$ where $c(\vec{x}) = i$ if $\vec{x} \in R_i$ ($\vec{x} \in \Gamma^k$). We also write $\mathfrak{X} = (\Gamma, c)$ instead of $\mathfrak{X} = (\Gamma, \mathfrak{R})$. A *refinement* of a coloring c is a coloring c' such that $(\forall \vec{x}, \vec{y} \in \Gamma^k)(c'(\vec{x}) = c'(\vec{y}) \implies c(\vec{x}) = c(\vec{y}))$.

An assignment $\mathfrak{X} \mapsto \mathfrak{X}'$ is *canonical* if it is defined by a functor between categories of isomorphisms of structures.

By a *binary configuration* we mean a binary partition structure $\mathfrak{X} = (\Gamma, c)$ such that

- (i) $(\forall x, y, z \in \Gamma)(c(x, y) = c(z, z) \implies x = y)$ and
 (ii) $(\forall x, y \in \Gamma)(c(x, y) \text{ determines } c(y, x))$.

The *Weisfeiler–Leman* canonical refinement process (WL) [Weisfeiler \[1968\]](#) and *On construction and identification of graphs* [\[1976\]](#) takes a binary configuration and with every pair $(x, y) \in \Gamma^2$ associates the list $c'(x, y) = (c(x, y), d_{i,j}(x, y) \mid i, j = 1, \dots, r)$ where $d_i(x, y) = |\{z \in \Gamma \mid c(x, z) = i, c(z, y) = j\}|$. This is clearly a canonical refinement.

Let $\mathfrak{X} = (\Gamma, c)$ be a k -ary partition structure. We assign colors to the elements by setting $c(x) = c(x, \dots, x)$. *Individualizing* an element $x \in \Gamma$ means assigning it a special color, thereby introducing irregularity. This irregularity propagates via canonical refinement, reducing the isomorphism search space. Let \mathfrak{X}_x denote \mathfrak{X} with $x \in \Gamma$ individualized. Then $\mathfrak{X} \cong \mathfrak{Y} \iff (\exists y \in \Gamma)(\mathfrak{X}_x \cong \mathfrak{Y}_y)$. So progress comes at a multiplicative cost of $m = |\Gamma|$. The multiplicative cost of individualizing t points is n^t , so we need $t \leq \text{polylog}$ for a quasipolynomial complexity bound.

9 Coherent configurations

The stable configurations of the WL process (where no proper refinement is obtained) are called *coherent configurations*. This concept goes back to [Schur \[1933\]](#) who abstracted its axiom from the *orbital configurations* of permutation groups. An *orbital* of $G \leq \text{Sym}(\Omega)$ is an orbit of the induced action of G on $\Omega \times \Omega$. Let $\mathfrak{X}(G)$ denote the configuration on Ω with the orbitals as the relations. This configuration is clearly coherent, but there are many coherent configurations that do not arise this way. For $v \geq 2k + 1$, the *Johnson scheme* $\mathfrak{J}(v, k)$ has $\binom{v}{k}$ vertices; it is defined as the orbital configuration of the group $S_v^{(k)}$ (induced action of S_v on unordered k -tuples).

A coherent configuration is *homogeneous* if every point has the same color. A homogeneous configuration is *primitive* if $|\Gamma| \geq 2$ and each off-diagonal color (relation) is a (strongly) connected (directed) graph. We note that the orbital configuration $\mathfrak{X}(G)$ of a permutation group G is homogeneous iff G is transitive and $\mathfrak{X}(G)$ is primitive iff G is primitive. The *rank* of a configuration is the number of colors, so for $|\Gamma| \geq 2$ the rank is at least 2. The only rank-2 configuration is the *clique*; its automorphism group is $\text{Sym}(\Gamma)$. The Johnson scheme $\mathfrak{J}(v, k)$ has rank $k + 1$.

The WL process and its natural k -ary generalization play a key role in the combinatorial analysis of the k -ary relational structures handed down by the Local Certificates algorithm.

10 Combinatorial partitioning

Recall that we have a giant homomorphism $\varphi : G \rightarrow \text{Sym}(\Gamma)$ for some ‘ideal domain’ Γ and we are given a canonical k -ary partition structure $\mathfrak{X}(x) = (\Gamma, c_x)$ with symmetry defect ≥ 0.1 where x is the input string. Here $k = O(\log n)$ where $n = |\Omega|$ is the size of our original domain. Recall that our recursive goal is to significantly reduce the size of the ideal domain at moderate multiplicative cost. Ideally we would like to achieve this by finding a *good canonical coloring* of Γ (no color has multiplicity greater than $0.9m$) or a *good equipartition*, i. e., a nontrivial canonical equipartition of the dominant ($> 0.9m$) vertex-color class.

This goal cannot be achieved because of the resilience of the Johnson schemes to canonical partitioning.

Proposition 10.1 (Resilience of Johnson schemes). *The multiplicative cost of a good canonical coloring or a good canonical equipartition of the Johnson scheme $\mathfrak{J}(v, t)$ is $\geq (4t)^{v/(4t)}$.*

The proof shows that if we pay less than exponential multiplicative cost then our Johnson scheme is simply reduced to a slightly smaller Johnson scheme.

Note that $t = 2$ is an interesting case, largely responsible for the lack of progress over the $\exp(\tilde{O}(\sqrt{n}))$ bound for a long time.

The good news is that in a sense, the Johnson schemes are the only obstacles.

So our modified goal will be to find either (a) a good canonical coloring, or (b) a good canonical equipartition, or (c) a canonically embedded Johnson scheme on a dominant vertex-color class. In item (c), canonical embedding means a functor from the isomorphisms of the input structures \mathcal{X} to the isomorphisms of the secondary structures whose vertex set is a dominant vertex-color class in Γ (under a canonical coloring).

We achieve this goal in two stages: first we go from k -ary to binary (Design Lemma) and then from binary to the desired goal (Split-or-Johnson).

Theorem 10.2 (Design lemma). *Let $\mathcal{X} = (\Gamma, c)$ be a k -ary partition structure with $m = |\Gamma|$ elements, $2 \leq k \leq m/2$, and symmetry defect ≥ 0.1 . Then in time $m^{O(k)}$ we can find a sequence S of at most $k - 1$ vertices such that after individualizing each element of S we can either find*

- (a) *a good canonical coloring of Γ , or*
- (b) *a good canonical equipartition of Γ , or*
- (c) *a good canonically embedded primitive coherent configuration of rank ≥ 3 .*

Here canonicity is relative to the arbitrary choice of the sequence S .

Outcomes (a) and (b) allow for efficient Luks reduction. Case (c) requires further processing.

Theorem 10.3 (Split-or-Johnson). *Given a primitive coherent configuration $\mathcal{X} = (\Gamma, c)$ of rank ≥ 3 , at quasipolynomial multiplicative cost we can find either*

- (a) *a good canonical coloring of Γ , or*
- (b) *a good canonical equipartition of Γ , or*
- (c) *a good canonically embedded nontrivial Johnson scheme.*

Here canonicity is relative to the arbitrary choices made that resulted in the multiplicative cost. The trivial Johnson schemes are the cliques $\mathfrak{J}(v, 1)$.

Outcomes (a) and (b) again allow for efficient Luks reduction. Outcome (c) provides even greater efficiency. Assume the canonically embedded Johnson scheme is $\mathfrak{J}(m', t)$; so $m \geq \binom{m'}{t} \geq \binom{m'}{2}$ and therefore $m' < 1 + \sqrt{2m}$. Now $\text{Aut}(\mathfrak{J}(m', t)) \cong S_{m'}$, so we can replace Γ by a set Γ' of size $m' = O(\sqrt{m})$, a dramatic reduction of the size of the ideal domain.

Overall algorithm. We follow Luks's algorithm until we hit a bottleneck, at which time an "ideal domain" Γ arises and our recursive goal becomes to significantly reduce the size of

the ideal domain. First we use our central group theoretic algorithm (“Local certificates”), based on the “Unaffected Stabilizers Lemma,” to construct a canonical structure on Γ of logarithmic arity and with non-negligible symmetry defect. Then we use our combinatorial partitioning algorithms to achieve the desired reduction. Once Γ itself becomes very small (polylogarithmic), we can individualize all of its elements, yielding a significant reduction of n , the size of the input string.

11 Paradoxes of Graph Isomorphism

GI is perceived to be an “easy” computational problem. As discussed in the Introduction (see “Disclaimer”), it is efficiently solved in practice. It is also provably easy on average. Our result shows it has rather low worst-case time complexity. In comparison, the problem of factoring integers is perceived to be “hard” – the assumption that it is hard, not only in the worst case but even of average, is the basis of the RSA cryptosystem and many other cryptographic applications. Yet, by common measures used in structural complexity theory, GI seems harder than factoring. The decision version of the factorization problem is in $\text{NP} \cap \text{coNP}$; this is not known to be the case for GI. Factoring is solvable in polynomial time in the quantum computation model; no quantum advantage has been found (in spite of significant effort) for GI. Most remarkable is the series of recent hardness results for GI in proof complexity, inspired by the CFI result. It turns out that in commonly studied hierarchies of semialgebraic and algebraic proof systems, isomorphism of certain pairs of graphs cannot be refuted on levels lower than cn for some constant $c > 0$ (where n is the number of vertices), corresponding to refutation proofs of exponential length in these systems [Atserias and Maneva \[2013\]](#), [O’Donnell, Wright, Wu, and Zhou \[2014\]](#), and [Berkholz and Grohe \[2015\]](#). (Cf. [Atserias and Ochremiak \[2017\]](#) for an overview of these and related systems.)

12 Open problems

Complexity theory. It is not known whether GI belongs to coNP . On the other hand, it is also not known whether P has logspace reductions to GI. This is equivalent to a logspace reduction of the *circuit value problem* (CVP) to GI. The CVP takes a Boolean circuit and an input to the circuit and asks to evaluate the circuit. Such a reduction would be viewed as strong evidence against the existence of an efficient parallel algorithm for GI.

While GI is universal over isomorphism problems for explicit structures, there are interesting classes of isomorphism problems for non-explicit structures that are also not expected to be NP -complete (based on strong evidence from the theory of interactive proofs), yet cannot currently be solved in less than exponential time. Perhaps the simplest among

them is the *code equivalence problem* that asks, given two subspaces U and V of \mathbb{F}^n for some finite field \mathbb{F} , is there a permutation $\sigma \in S_n$ such that $U^\sigma = V$? Here σ acts on \mathbb{F}^n by permuting the coordinates.

Can GI be solved in quasipolynomial time and *polynomial space*? (Luks)

Can *canonical forms* of graphs be constructed in quasipolynomial time? (Cf. Babai and E. M. Luks [1983].)

Can isomorphism of hypergraphs be decided in time, quasipolynomial in the number of vertices and *polynomial in the number of edges*?

Combinatorics. The author's decades-old project to find combinatorial relaxations of Cameron's [Theorem 4.5](#) has seen major progress recently, made by PhD students. *Cameron schemes* are the orbital configurations of Cameron groups ([Definition 4.4](#)). Let us say that a primitive coherent configuration is a *non-Cameron PCC* if it is not a Cameron scheme. The author has circulated various versions of the following conjectures for some time.

Conjecture 12.1. There exists a polynomial p such that the following hold. Let \mathcal{X} be a non-Cameron PCC with n vertices. Let $G = \text{Aut}(\mathcal{X})$. Then

- (a) $\theta(\text{Aut}(\mathcal{X})) \leq p(\log n)$ (where θ denotes the thickness, [Definition 4.1](#))
(polylogarithmically bounded thickness)
- (b) $|G| \leq \exp(p(\log n))$ (quasipolynomially bounded order)

Part (a) obviously follows from part (b). Regarding (b), for non-Cameron PCCs, an upper bound $|G| \leq \exp(\tilde{O}(\sqrt{n}))$ was proved in Babai [1981] in 1981. After no progress for three and a half decades, in a recent *tour de force* of combinatorial reasoning, Sun and Wilmes reduced this upper bound to $\exp(\tilde{O}(n^{1/3}))$, building a new combinatorial structure theory of primitive coherent configurations along the way. The weaker Conjecture (a) has been confirmed for rank-3 configurations (essentially, strongly regular graphs) in Babai [2014] (2014). Overcoming an array of technical obstacles through a powerful combination of structural and spectral theory, Kivva [2017] very recently confirmed (a) for rank-4 configurations. These are major steps, and raise the hope of further progress, although the technical challenges seem daunting.

References

- M. Aschbacher and L. Scott (1985). “Maximal subgroups of finite groups”. *J. Algebra* 92.1, pp. 44–80. MR: [772471](#) (cit. on p. 3311).
- Albert Atserias and Elitza Maneva (2013). “Sherali-Adams relaxations and indistinguishability in counting logics”. *SIAM J. Comput.* 42.1, pp. 112–137. MR: [3033123](#) (cit. on p. 3316).

- Albert Atserias and Joanna Ochremiak (2017). “Proof complexity meets algebra”. In: *44th International Colloquium on Automata, Languages, and Programming*. Vol. 80. LIPIcs. Leibniz Int. Proc. Inform. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, Art. No. 110, 14. arXiv: [1711.07320](#). MR: [3685850](#) (cit. on p. [3316](#)).
- L. Babai, P. J. Cameron, and P. P. Pálffy (1982). “On the orders of primitive groups with restricted nonabelian composition factors”. *J. Algebra* 79.1, pp. 161–168. MR: [679977](#) (cit. on p. [3308](#)).
- László Babai (1979). “Monte-Carlo algorithms in graph isomorphism testing”. *Université de Montréal Technical Report, DMS* 79-10, p. 42 (cit. on p. [3304](#)).
- (1981). “On the order of uniprimitive permutation groups”. *Ann. of Math.* (2) 113.3, pp. 553–568. MR: [621016](#) (cit. on p. [3317](#)).
 - (1983). “Permutation Groups, Coherent Configurations, and Graph Isomorphism”. PhD thesis. D. Sc. Thesis (Hungarian), Hungarian Academy of Sciences (cit. on p. [3306](#)).
 - (2014). “On the automorphism groups of strongly regular graphs I”. In: *ITCS’14—Proceedings of the 2014 Conference on Innovations in Theoretical Computer Science*. ACM, New York, pp. 359–368. MR: [3359489](#) (cit. on pp. [3307](#), [3317](#)).
 - (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, pp. 684–697. arXiv: [1512.03547](#). MR: [3536606](#) (cit. on pp. [3305](#), [3306](#)).
- László Babai, Xi Chen, Xiaorui Sun, Shang-Hua Teng, and John Wilmes (2013). “Faster canonical forms for strongly regular graphs”. In: *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*. IEEE, pp. 157–166 (cit. on p. [3305](#)).
- László Babai, Paul Erdős, and Stanley M. Selkow (1980). “Random graph isomorphism”. *SIAM J. Comput.* 9.3, pp. 628–635. MR: [584517](#) (cit. on p. [3304](#)).
- László Babai, D Yu Grigoryev, and David M Mount (1982). “Isomorphism of graphs with bounded eigenvalue multiplicity”. In: *Proceedings of the fourteenth annual ACM symposium on Theory of computing*. ACM, pp. 310–324 (cit. on p. [3304](#)).
- László Babai, William M Kantor, and Eugene M Luks (1983). “Computational complexity and the classification of finite simple groups”. In: *Foundations of Computer Science, 1983., 24th Annual Symposium on*. IEEE, pp. 162–171 (cit. on pp. [3305](#), [3306](#)).
- László Babai and Ludik Kucera (1979). “Canonical labelling of graphs in linear average time”. In: *Foundations of Computer Science, 1979., 20th Annual Symposium on*. IEEE, pp. 39–46 (cit. on p. [3304](#)).
- László Babai and Eugene M Luks (1983). “Canonical labeling of graphs”. In: *Proceedings of the fifteenth annual ACM symposium on Theory of computing*. ACM, pp. 171–183 (cit. on pp. [3305](#), [3317](#)).

- László Babai and John Wilmes (2013). “Quasipolynomial-time canonical form for Steiner designs”. In: *STOC’13—Proceedings of the 2013 ACM Symposium on Theory of Computing*. ACM, New York, pp. 261–270. MR: [3210787](#) (cit. on p. [3305](#)).
- Christoph Berkholz and Martin Grohe (2015). “Limitations of algebraic approaches to graph isomorphism testing”. In: *Automata, languages, and programming. Part I*. Vol. 9134. Lecture Notes in Comput. Sci. Springer, Heidelberg, pp. 155–166. arXiv: [1502.05912](#). MR: [3382436](#) (cit. on p. [3316](#)).
- Alexandre V. Borovik, Laszlo Pyber, and Aner Shalev (1996). “Maximal subgroups in finite and profinite groups”. *Trans. Amer. Math. Soc.* 348.9, pp. 3745–3761. MR: [1360222](#) (cit. on p. [3308](#)).
- Jin-Yi Cai, Martin Fürer, and Neil Immerman (1992). “An optimal lower bound on the number of variables for graph identification”. *Combinatorica* 12.4, pp. 389–410. MR: [1194730](#) (cit. on pp. [3304](#)–[3306](#)).
- Peter J. Cameron (1981). “Finite permutation groups and finite simple groups”. *Bull. London Math. Soc.* 13.1, pp. 1–22. MR: [599634](#) (cit. on pp. [3308](#), [3309](#)).
- Xi Chen, Xiaorui Sun, and Shang-Hua Teng (2013). “Multi-stage design for quasipolynomial-time isomorphism testing of Steiner 2-systems”. In: *STOC’13—Proceedings of the 2013 ACM Symposium on Theory of Computing*. ACM, New York, pp. 271–280. MR: [3210788](#) (cit. on p. [3305](#)).
- John D. Dixon and Brian Mortimer (1996). *Permutation groups*. Vol. 163. Graduate Texts in Mathematics. Springer-Verlag, New York, pp. xii+346. MR: [1409812](#) (cit. on p. [3311](#)).
- Merrick Furst, John Hopcroft, and Eugene Luks (1980). “Polynomial-time algorithms for permutation groups”. In: *21st Annual Symposium on Foundations of Computer Science (Syracuse, N.Y., 1980)*. IEEE, New York, pp. 36–41. MR: [596045](#) (cit. on p. [3304](#)).
- Michael R. Garey and David S. Johnson (1979). *Computers and intractability. A guide to the theory of NP-completeness*, A Series of Books in the Mathematical Sciences. W. H. Freeman and Co., San Francisco, Calif., pp. x+338. MR: [519066](#) (cit. on pp. [3303](#), [3304](#)).
- Z. Hedrlín and A. Pultr (1966). “On full embeddings of categories of algebras”. *Illinois J. Math.* 10, pp. 392–406. MR: [0191858](#) (cit. on p. [3303](#)).
- J. E. Hopcroft and R. E. Tarjan (1972). “Isomorphism of planar graphs”, pp. 131–152, 187–212. MR: [0403302](#) (cit. on p. [3304](#)).
- J. E. Hopcroft and J. K. Wong (1974). “Linear time algorithm for isomorphism of planar graphs: preliminary report”, pp. 172–184. MR: [0433964](#) (cit. on p. [3304](#)).
- Richard M. Karp (1972). “Reducibility among combinatorial problems”, pp. 85–103. MR: [0378476](#) (cit. on p. [3304](#)).
- Bohdan Kivva (2017). “On the automorphism groups of distance-regular graphs and rank-4 primitive coherent configurations”. Manuscript, University of Chicago (cit. on p. [3317](#)).

- Martin W. Liebeck and Aner Shalev (2003). “Bases of primitive permutation groups”. In: *Groups, combinatorics & geometry (Durham, 2001)*. World Sci. Publ., River Edge, NJ, pp. 147–154. MR: [1994965](#) (cit. on p. [3308](#)).
- Eugene M. Luks (1982). “Isomorphism of graphs of bounded valence can be tested in polynomial time”. *J. Comput. System Sci.* 25.1, pp. 42–65. MR: [685360](#) (cit. on pp. [3305](#), [3306](#), [3308](#)).
- Attila Maróti (2002). “On the orders of primitive groups”. *J. Algebra* 258.2, pp. 631–640. MR: [1943938](#) (cit. on p. [3309](#)).
- Brendan D. McKay and Adolfo Piperno (2014). “Practical graph isomorphism, II”. *J. Symbolic Comput.* 60, pp. 94–112. arXiv: [1301.1493](#). MR: [3131381](#) (cit. on p. [3305](#)).
- Gary L. Miller (1979). “Graph isomorphism, general remarks”. *J. Comput. System Sci.* 18.2, pp. 128–142. MR: [532172](#) (cit. on p. [3303](#)).
- Takunari Miyazaki (1996). *Luks’s reduction of Graph isomorphism to code equivalence*. Comment on The Math Forum (cit. on p. [3306](#)).
- Daniel Neuen and Pascal Schweitzer (May 2017). “An exponential lower bound for Individualization Refinement algorithms for Graph Isomorphism”. arXiv: [1705.03283](#) (cit. on p. [3306](#)).
- Ryan O’Donnell, John Wright, Chenggang Wu, and Yuan Zhou (2014). “Hardness of robust graph isomorphism, Lasserre gaps, and asymmetry of random graphs”. In: *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*. ACM, New York, pp. 1659–1677. MR: [3376480](#) (cit. on p. [3316](#)).
- On construction and identification of graphs* (1976). Lecture Notes in Mathematics, Vol. 558. With contributions by A. Lehman, G. M. Adelson-Velsky, V. Arlazarov, I. Faragev, A. Uskov, I. Zuev, M. Rosenfeld and B. Weisfeiler, Edited by Boris Weisfeiler. Springer-Verlag, Berlin-New York, pp. xiv+237. MR: [0543783](#) (cit. on p. [3313](#)).
- P. P. Pálffy (1982). “A polynomial bound for the orders of primitive solvable groups”. *J. Algebra* 77.1, pp. 127–137. MR: [665168](#) (cit. on p. [3307](#)).
- László Pyber (1990). Unpublished (cit. on p. [3308](#)).
- (May 2016). “A CFSG-free analysis of Babai’s quasipolynomial GI-algorithm”. arXiv: [1605.08266](#) (cit. on p. [3311](#)).
- Issai Schur (1933). “Zur Theorie der einfach transitiven Permutationsgruppen”. *Sitzungsber. Preuss. Akad. Wiss.* Pp. 598–623 (cit. on p. [3314](#)).
- Leonard L. Scott (1980). “Representations in characteristic p ”. In: *The Santa Cruz Conference on Finite Groups (Univ. California, Santa Cruz, Calif., 1979)*. Vol. 37. Proc. Sympos. Pure Math. Amer. Math. Soc., Providence, R.I., pp. 319–331. MR: [604599](#) (cit. on p. [3311](#)).
- Ákos Seress (2003). *Permutation group algorithms*. Vol. 152. Cambridge Tracts in Mathematics. Cambridge University Press, Cambridge, pp. x+264. MR: [1970241](#) (cit. on p. [3305](#)).

- Daniel A. Spielman (1996). “Faster isomorphism testing of strongly regular graphs”. In: *Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (Philadelphia, PA, 1996)*. ACM, New York, pp. 576–584. MR: [1427556](#) (cit. on p. [3305](#)).
- Xiaorui Sun and John Wilmes (2015). “Faster canonical forms for primitive coherent configurations (extended abstract)”. In: *STOC’15—Proceedings of the 2015 ACM Symposium on Theory of Computing*. ACM, New York, pp. 693–702. MR: [3388249](#) (cit. on p. [3305](#)).
- Boris Weisfeiler (1968). “A reduction of a graph to a canonical form and an algebra arising during this reduction”. *Nauchno-Tekhnicheskaya Informatsiya* 9, pp. 12–16 (cit. on p. [3313](#)).
- Thomas R. Wolf (1982). “Solvable and nilpotent subgroups of $GL(n, q^m)$ ”. *Canad. J. Math.* 34.5, pp. 1097–1111. MR: [675682](#) (cit. on p. [3307](#)).
- V. N. Zemlyachenko, N. M. Korneenko, and R. I. Tyshkevich (1982). “The graph isomorphism problem”. *Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI)* 118. The theory of the complexity of computations, I, pp. 83–158, 215. MR: [659084](#) (cit. on p. [3305](#)).

Received 2017-12-11.

László Babai
University of Chicago
laci@cs.uchicago.edu

