Canonization of a random circulant graph by counting walks*

Oleg Verbitsky[†] Maksim Zhukovskii[‡]

Abstract

It is well known that almost all graphs are canonizable by a simple combinatorial routine known as color refinement, also referred to as the 1-dimensional Weisfeiler-Leman algorithm. With high probability, this method assigns a unique label to each vertex of a random input graph and, hence, it is applicable only to asymmetric graphs. The strength of combinatorial refinement techniques becomes a subtle issue if the input graphs are highly symmetric. We prove that the combination of color refinement and vertex individualization yields a canonical labeling for almost all circulant digraphs (i.e., Cayley digraphs of a cyclic group). This result provides first evidence of good average-case performance of combinatorial refinement within the class of vertex-transitive graphs. Remarkably, we do not even need the full power of the color refinement algorithm. We show that the canonical label of a vertex v can be obtained just by counting walks of each length from v to an individualized vertex. Our analysis also implies that almost all circulant graphs are compact in the sense of Tinhofer, that is, their polytops of fractional automorphisms are integral. Finally, we show that a canonical Cayley representation can be constructed for almost all circulant graphs by the more powerful 2-dimensional Weisfeiler-Leman algorithm.

1 Introduction

1.1 Combinatorial refinement and canonization of random graphs

As is well known, the graph isomorphism problem is very efficiently solvable in the average case by a simple combinatorial method known as *color refinement* (also *degree refinement* or *naive vertex classification*). When a random graph $\mathbf{G}_n$ on n vertices is taken as an input, this algorithm produces a canonical labeling of all vertices in $\mathbf{G}_n$ by coloring them initially by their degrees and then by refining the

[*] A preliminary version of this paper was presented at WALCOM'24, the 18th International Conference and Workshops on Algorithms and Computation [45].

[†] Institut für Informatik, Humboldt-Universität zu Berlin, Germany. On leave from the IAPMM, Lviv, Ukraine.

[‡] School of Computer Science, University of Sheffield, UK.

color classes as follows: Two equally colored vertices u and v get new distinct colors if one of the initial colors occurs in the neighborhoods of u and v with different multiplicity. In this way, every vertex gets a unique color with probability $1 - O(n^{-1/7})$ (Babai, Erdős, and Selkow [5]). Thus, the method produces a canonical labeling for almost all graphs on a fixed set of n vertices.

This approach is not applicable to regular graphs, even with many refinement rounds, because if all vertices have the same degree, then the refinement step makes obviously no further vertex separation. Weisfeiler and Leman [46] came up with a more powerful refinement algorithm which colors pairs of vertices instead of single vertices (see Section 7.1 for a formal description). The idea can be lifted to k-tuples of vertices, for each integer parameter k, and the general approach is referred to as the k-dimensional Weisfeiler-Leman algorithm, abbreviated as k-WL. Thus, 2-WL is the original algorithm in [46], and 1-WL corresponds to color refinement. Remarkably, 2-WL is powerful enough to produce a canonical labeling for almost all regular graphs of a given vertex degree (Bollobás [9]; see also [31]).

Further restriction of regular input graphs to vertex-transitive graphs is challenging for combinatorial refinement because no vertex classification is at all possible in this case. Indeed, 2-WL assigns the same color to any two vertices u and v (or, more precisely, to the pairs (u, u) and (v, v)) because u is mapped to v by an automorphism of the graph. The same holds for any dimension k.

Moreover, the graph isomorphism problem for vertex-transitive graphs is provably unsolvable by k-WL for any fixed dimension k. Indeed, the Cai-Fürer-Immerman (CFI) construction [11] of non-isomorphic graphs X_k and Y_k indistinguishable by k-WL can be modified so that these graphs become vertex-transitive [23]. A natural way to enhance combinatorial refinement is to combine it with *vertex individualization*—that is, assigning a special color to one vertex in the graph [33]. While this algorithmic approach proves advantageous in many contexts (see, e.g., [4]), it nevertheless fails to overcome the obstacle posed by the CFI graphs. To see this, note that $(k + k')$-WL is more powerful than any combination of k-WL with prior individualization of k' vertices. As a consequence, for any pair of arbitrarily large integers k and k', isomorphism of vertex-transitive graphs cannot be solved by k-WL even under individualization of k' vertices.

Motivated by the question of whether or not these basic obstacles persist in the average case setting, we focus in this paper on Cayley graphs and, more specifically, on circulant graphs, that is, Cayley graphs of a cyclic group. While the canonization problem for this class of graphs is known to be solvable in polynomial time [19] by advanced algebraic methods, it is an open question whether this can be done by using k-WL for some fixed dimension k; see [38, 47]. This poses an ongoing challenge for the combinatorial refinement method, especially because the research on isomorphism of circulant graphs has a long history with many deep results (see [3, 19, 36, 35] and the references therein) and because k-WL with small dimension k is known to be applicable to many other natural graph classes (e.g., planar graphs [28]). The recent paper [29] investigates the round complexity of 2-WL on circulant graphs, exploiting the close connections of the subject with intricate mathematical concepts. Circulant graphs are also interesting on their own right as they naturally appear and are intensively investigated in many other theoretical and applied areas;

see, e.g., the books [10, 14, 16]. After all, our primary motivation for the study of circulant graphs is that this is the most natural first choice of a graph class for benchmarking of combinatorial refinement in the realm of vertex-transitive graphs.

1.2 Our contribution: Random circulant graphs

We begin by fixing the basic notation and terminology. Throughout the paper, the isomorphism relation $X \cong Y$ for graphs X and Y refers to the standard combinatorial notion of graph isomorphism, regardless of any underlying algebraic structure on the vertex sets of X and Y. In particular, when speaking of isomorphisms and automorphisms of graphs, we always mean the standard graph-theoretic concepts.

For a bijective vertex labeling $\lambda : V(X) \to \{0, 1, \ldots, n-1\}$ of an n-vertex digraph X, let X^λ denote the relabeled version of X, that is, the digraph on vertex set $\{0, 1, \ldots, n-1\}$ containing an edge from x to y whenever X contains an edge from $\lambda^{-1}(x)$ to $\lambda^{-1}(y)$. Given an input digraph X, a **canonical labeling algorithm** produces a *canonical labeling* λ_X of X, which satisfies the following properties:

- a labeling λ_Y is computed by the algorithm also for every $Y \cong X$, and
- $X^{\lambda_X} = Y^{\lambda_Y}$ for all such Y.

When applied to a randomly chosen X, such an algorithm may occasionally fail, i.e., terminate without producing any output (this failure occurs simultaneously for all isomorphic inputs). We say that the algorithms *succeeds* with probability at least $1 - \varepsilon(n)$, if the failure probability on the inputs with n vertices does not exceed $\varepsilon(n)$.

Canonical labeling of a random circulant. Our goal is to show that the individualization-refinement approach can be used to canonize almost all circulants at minimal computational costs. Our treatment covers also circulant directed graphs, which is advantageous for expository purposes as the case of digraphs is technically somewhat simpler. We show that the individualization of a single vertex suffices for random circulant digraphs, and that two individualized vertices are enough in the undirected case (in fact, we just perform color refinement twice, each time with a single individualized vertex). In both the directed and undirected cases, we maintain an overall running time of $O(n^2 \log n)$, which is the standard running time of color refinement [12]. This is possible because our input graphs are vertex-transitive, and hence, it does not actually matter which vertex is individualized — even though, in the undirected case, the choice of the second vertex to be individualized is not arbitrary. Here, n denotes the number of vertices. Thus, our time bound is actually linearithmic, that is, it is $O(N \log N)$ for the input length N where the input (di)graph is presented by the adjacency matrix and the cyclic structure is not explicitly given (see below the discussion of different representation concepts). Note also that, as one might expect, our average-case bound of $O(n^2 \log n)$ is substantially better than the worst-case bound resulting from [19].[1] We summarize our main result in a somewhat condensed form as follows.

[1] The algorithm in [19] involves the 2-dimensional Weisfeiler-Leman algorithm, which has time complexity $O(n^3 \log n)$. Its overall running time is stated as n^c, where the unspecified constant c depends on the complexity of several algorithms from computational group theory.

Theorem 1.1 (Main Theorem). *A uniformly random circulant (di)graph with n vertices is with probability at least $1 - n^{-1/2+o(1)}$ canonizable by color refinement combined with vertex individualization in running time $O(n^2 \log n)$.*

Theorem 1.1 includes two statements, one for undirected graphs (where all undirected circulant graphs are equiprobable) and the similar statement for directed graphs. Note that the concept of a uniform distribution of n-vertex circulants is somewhat ambiguous as the notion of an *n-vertex circulant* alone can be defined in three different natural ways:

- as a Cayley (di)graph of the cyclic group $\mathbb{Z}_n$,
- as an isomorphism class of Cayley (di)graphs of $\mathbb{Z}_n$ (which we call an *unlabeled circulant*),
- as a (di)graph on the vertex set $\{0, 1, \ldots, n-1\}$ isomorphic to a Cayley (di)graph of $\mathbb{Z}_n$ (which we call a *labeled circulant*).

More formally, a *connection set* $S \subseteq \mathbb{Z}_n \setminus \{0\}$ defines the *Cayley digraph* $\mathrm{Cay}(\mathbb{Z}_n, S)$, where two vertices $x, y \in \mathbb{Z}_n$ form a directed edge (x, y) if $y - x \in S$. If S is inverse-closed, that is, $S = -S$, then $\mathrm{Cay}(\mathbb{Z}_n, S)$ is undirected. When referring to a random Cayley (di)graph, we assume that all connection sets (with $S = -S$ in the undirected case) are equally likely. Detailed definitions of the other two distributions are provided in Section 6.

We first prove Theorem 1.1 for random Cayley digraphs and graphs (see the proof outline in Section 2), and then extend the result to the other two concepts. The transition from one distribution to another is quite general and is based on known results in algebraic graph theory [8, 15, 35].

Canonical Cayley representation. Our canonical labeling algorithm in Theorem 1.1 is based on individualization of a single vertex in the input digraph followed by color refinement (or 1-WL in other terminology). The aforementioned 2-dimensional Weisfeiler-Leman algorithm (2-WL) is strictly stronger than the combination of 1-WL with one-vertex individualization (cf. [40, Theorem 3.2]). It turns out that 2-WL can be used to solve an even more challenging algorithmic problem than computing a canonical labeling, which we describe below.

Theorem 1.1 provides an algorithm for producing a canonical labeling λ_X of a random circulant X. Note, however, that the proof does not guarantee that the canonical form X^{λ_X} is a Cayley digraph of $\mathbb{Z}_n$—or, equivalently, that the cycle $(\lambda_X^{-1}(0), \lambda_X^{-1}(1), \ldots, \lambda_X^{-1}(n-1))$ is an automorphism of X. In cases where this condition holds, i.e., when $X^{\lambda_X} = \mathrm{Cay}(\mathbb{Z}_n, S)$ for some connection set S, we say that the map λ_X is a *Cayley representation* of X. We are interested in an algorithm which, with high probability, succeeds in computing a map λ_X that is *both* a canonical labeling and a Cayley representation of X.

Theorem 1.2. *A uniformly random labeled circulant admits a canonical Cayley representation computable by means of* 2-WL *in time $O(n^3 \log n)$ with success probability $1 - O(n^2 2^{-n/8})$.*

Compared to Theorem 1.1, the proof of Theorem 1.2 requires less technical effort, primarily due to the greater expressive power of the 2-WL algorithm. Moreover, this algorithm is closely related to the notion of a coherent configuration in algebraic combinatorics [13]. This connection enables us to leverage strong results from algebraic graph theory [8, 15, 21], whose algorithmic interpretation is the core of Theorem 1.2.

2 Proof strategy and further implications

2.1 Proof overview of the Main Theorem: The power of walk counts

Remarkably, we show that canonization of a random circulant does not even need the full power of color refinement and can actually be accomplished by a weaker algorithmic tool. Let G be an arbitrary (di)graph on the vertex set $V = \{0, 1, \dots, n-1\}$, and let $T \subseteq V$. The *walk matrix* $W_T = (w_{ik})_{i,k \in V}$ is defined by setting w_{ik} to be the number of walks of length k from the vertex i to a vertex in T. That is, w_{ik} is the number of vertex sequences $x_0, x_1, \dots, x_k$ in G such that $x_0 = i$, $x_k \in T$, and the pair (x_j, x_{j+1}) is an edge of G for every $j < k$. If A is the adjacency matrix of G and χ_T denotes the characteristic vector of the subset T, then W_T is formed by the columns $\chi_T, A\chi_T, A^2\chi_T, \dots, A^{n-1}\chi_T$. The theory of walk matrices, including their applicability to isomorphism testing, has been developed by Godsil [24] and by Liu and Siemons [32]. Let G_T be obtained from G by coloring all vertices in T by the same color. We call G_T *walk-discrete* if the rows of W_T are pairwise different. For any walk-discrete G_T, the walk matrix W_T yields a canonical labeling of the vertices of G_T. This purely algebraic canonization method can be superseded by the purely combinatorial method of color refinement because if $w_{uk} \neq w_{vk}$ for some k, then color refinement assigns distinct colors to the vertices u and v in G_T (see Section 3.1 for details).

Let $W = W_V$ be the *standard* walk matrix of $G = G_V$. Obviously, G is walk-discrete whenever W is non-singular. Noteworthy, the rank of W for an undirected graph G is equal to the number of different *main* eigenvalues of the adjacency matrix A; see [25]. As shown by O'Rourke and Touri [37], a random undirected graph $\mathbf{G}_n$ has non-singular walk matrix with high probability. As a consequence, $\mathbf{G}_n$ is, with high probability, canonizable by computing its standard walk matrix.[2]

The above theory essentially exploits the fact that the adjacency matrices of undirected graphs are symmetric and, by this reason, does not apply to directed graphs. Nevertheless, we obtain the following spectral criterion for circulant digraphs.

Lemma 2.1. *Let X be a Cayley digraph of a cyclic group and $X_0 = X_{\{0\}}$ be its version with one individualized vertex. Let W_0 be the walk matrix of X_0. Then*

[2] In fact, only a small part of the walk matrix suffices for this purpose — as shown in [44], $\mathbf{G}_n$ is canonizable with high probability by assigning each vertex u the triple (w_{u1}, w_{u2}, w_{u3}). We also refer the interested reader to [30] for applications of color refinement and walk numbers in machine learning.

W_0 is non-singular (implying that X_0 is walk-discrete) if and only if X has simple spectrum, that is, all eigenvalues of X are pairwise distinct.

Suppose now that X is an undirected Cayley graph of $\mathbb{Z}_n$. In this case, the map $x \mapsto (-x) \bmod n$ is an automorphism of X_0, which implies that the walk matrix of X_0 has at most $\lceil (n+1)/2 \rceil$ different rows. If this bound is achieved, we call X_0 *walk-saturated*. On the other hand, the spectrum of X has at most $\lceil (n+1)/2 \rceil$ different eigenvalues. If there are exactly so many eigenvalues, we say that X has *saturated spectrum*. We have the following analog of Lemma 2.1 for the undirected case.

Lemma 2.2. *Let X be a Cayley graph of the cyclic group $\mathbb{Z}_n$. Then W_0 has the maximum possible rank $\lceil (n+1)/2 \rceil$ (implying that X_0 is walk-saturated) if and only if X has saturated spectrum.*

As an immediate consequence of Lemma 2.1, the rows of the walk matrix yield a canonical labeling of a circulant digraph whenever it has a simple spectrum. With only a small amount of additional technical effort (see Lemma 3.2 in the next section), Lemma 2.2 implies that the walk matrix can also be used to canonize a circulant graph whenever it has a saturated spectrum. The following theorem therefore estimates the success probability of these canonization methods on random circulants.

Theorem 2.3.

1. *If $S \subseteq \mathbb{Z}_n \setminus \{0\}$ is chosen uniformly at random, then the Cayley digraph* $\mathrm{Cay}(\mathbb{Z}_n, S)$ *has simple spectrum with probability at least* $1 - n^{-1/2+o(1)}$.

2. *If $S \subseteq \mathbb{Z}_n \setminus \{0\}$ is chosen uniformly at random among all inverse-closed sets, then the Cayley graph* $\mathrm{Cay}(\mathbb{Z}_n, S)$ *has saturated spectrum with probability at least* $1 - n^{-1/2+o(1)}$.

Thus, Theorem 1.1 for Cayley (di)graphs follows from Theorem 2.3 in view of Lemmas 2.1 and 2.2. Theorem 2.3—and hence, as already noted, also Theorem 1.1—extends to the other two models of a random circulant discussed in Section 1, namely, the random labeled and unlabeled circulants.

Theorem 2.3 is our main technical contribution and may be of independent interest in the context of research on random circulant matrices [10, 34]. Moreover, this result has further noteworthy consequences for properties of random circulant graphs, which we discuss in the next subsection.

2.2 Compactness and naïve canonization

Note that two n-vertex graphs G and H with adjacency matrices A and B, respectively, are isomorphic if and only if there is an $n \times n$ permutation matrix P such that $AP = PB$. This observation leads to a natural linear programming relaxation for the graph isomorphism problem. Recall that an $n \times n$ real matrix S is *doubly stochastic* if its elements are nonnegative and all its rows and columns sum up to 1.

A doubly stochastic matrix S satisfying the equation $AS = SB$ is called a *fractional isomorphism* from G to H. In particular, if $AS = SA$, then S is a *fractional automorphism* of G.

Tinhofer [42] calls a graph G *compact* if the polytope formed in $\mathbb{R}^{n^2}$ by the fractional automorphisms of G is integral, that is, the extreme points of this polytope have integer coordinates. If a compact graph G is isomorphic to another graph H, then the polytope of fractional isomorphisms from G to H is also integral. On the other hand, this polytope contains no integral point at all if G and H are non-isomorphic. This has an algorithmic consequence. If we know that a graph G is compact, then we can decide whether or not G is isomorphic to any other given graph H in polynomial time by using linear programming. It suffices to compute an extreme point of the polytope of fractional isomorphisms from G to H and check if it has integer coordinates.

While no efficient method is known in general for determining whether a graph G is compact, Schreck and Tinhofer [41] established a sufficient condition for the compactness of circulant graphs. In our terminology, they showed that circulant graphs with saturated spectrum are compact. Consequently, Part 2 of our Theorem 2.3 immediately implies the following result.

Corollary 2.4. *Almost all circulant graphs are compact. More precisely, for all three notions of a circulant graph, if X is a uniformly random circulant graph on n vertices, then X is compact with probability at least $1 - n^{-1/2+o(1)}$.*

In [43], Tinhofer presented another, fairly surprising combinatorial approach to testing whether two graphs are isomorphic provided one of them is compact. The approach can be recast as a canonization algorithm and was considered as such also by Immerman and Lander [26]. Specifically, they consider the following algorithmic procedure.

```
NAÏVE CANONIZATION
INPUT: a graph $X$.
  1. Set $\widetilde{X} = X$.
  2. Run color refinement on $\widetilde{X}$ and denote the resulting colored graph by $\widehat{X}$.
  3. If the vertex colors in $\widehat{X}$ are pairwise distinct, then output this coloring as a
     canonical labeling of $X$.
  4. Otherwise
     (a) choose an arbitrary vertex $v$ in the non-singleton color class of $\widehat{X}$ with
         least[3] color;
     (b) reset $\widetilde{X}$ to be $\widehat{X}$ with the vertex $v$ individualized;
     (c) repeat Step 2 again.
```

[3] The set of colors produced by color refinement is endowed with a natural order. Moreover, we can suppose that each vertex of $\widehat{X}$ is colored by an integer in $\{0, 1, \ldots, n-1\}$; see Section 3.1.

We say that the above procedure *works correctly* on input X if it produces a canonical labeling of X irrespectively of which vertex v is chosen in any execution of Step 4(a). Note that naïve canonization works correctly whenever it terminates after the first execution of Step 3, which occurs for almost all graphs X by [5]. In the general case, when Step 4 is executed at least once, it is clear that a non-backtracking refinement-individualization procedure like this cannot be expected to be correct. Indeed, naïve canonization obviously fails on any regular but non-vertex-transitive graph X. All the more surprising, then, is the result established in [43] that the approach still works correctly on such a broad and naturally defined class of graphs as the class of compact graphs.

Combining this fact with Corollary 2.4, we obtain the following result.

Corollary 2.5. *Naïve canonization works correctly for almost all circulant graphs. More precisely, this holds for all but a fraction of $n^{-1/2+o(1)}$ of circulant graphs on n vertices, under any of the three models of circulant graphs.*

Note that the class of graphs on which naïve canonization succeeds is strictly larger than the class of compact graphs; see [2]. This remains true also when we focus on circulant graphs. Indeed, Schreck and Tinhofer [41] proved that a non-empty and non-complete circulant graph with a prime number of vertices is compact if and only if it has saturated spectrum. Remarkably, Kluge [29] showed that naïve canonization works correctly on *every* circulant graph with a prime number of vertices.

Finally, we remark that naïve canonization makes sense only if we are a priori confident in its correctness. Corollary 2.5 is quite constructive in this regard, as the sufficient condition of Schreck and Tinhofer [41]—having a saturated spectrum—is verifiable in polynomial time. Moreover, it can be shown by a direct argument that naïve canonization works correctly on every walk-saturated circulant graph.

2.3 Organization of the rest of the paper

The proof of our Main Theorem (Theorem 1.1) spans Sections 3–6. Section 3 begins with preliminaries and then provides a detailed description of our canonical labeling algorithm for circulant graphs with saturated spectrum. Lemmas 2.1 and 2.2 are proved in Section 4, as special cases of a more general result stated there as Lemma 4.1. Theorem 2.3 is proved in Section 5. This establishes Theorem 1.1 for Cayley (di)graphs over $\mathbb{Z}_n$.

The proof of Theorem 1.1 is completed in Section 6, which is devoted to the other two models of a random circulant. In this section, we establish two "transition" lemmas. Specifically, Lemma 6.4 allows us to extend the probability bound of Theorem 2.3 from random Cayley (di)graphs to random *unlabeled* circulants, while Lemma 6.6 further extends this bound to random *labeled* circulants.

Finally, Theorem 1.2 is proved in Section 7.

3 The walk matrix and color refinement

3.1 Definitions and a relationship

Speaking of a directed graph or, for short, *digraph* $G = (V, E)$, we always assume that G is loopless, that is, the adjacency relation $E \subset V^2$ is irreflexive. Without loss of generality we suppose that G is defined on the vertex set $V = \{0, 1, \ldots, n-1\}$. If E is symmetric, G is referred to as an (undirected) *graph*. The definitions given below for digraphs apply, as a special case, also to graphs.

For $t \in V$, we write G_t to denote the digraph G with distinguished vertex t. The vertex t is called *terminal* or *individualized*. We consider G_t to be a vertex-colored digraph where all vertices are equally colored with the exception of t which has a special color. An isomorphism from G_t to another vertex-individualized digraph H_u is defined as a digraph isomorphism from G to H taking t to u.

A *walk* in G is a sequence of vertices $x_0x_1 \ldots x_k$ such that $(x_i, x_{i+1}) \in E$ for every $0 \le i < k$. We say that $x_0x_1 \ldots x_k$ is a walk of length k from x_0 to x_k. Note that a one-element sequence x_0 is a walk of length 0. Given a digraph G_t with terminal vertex t, we define its *walk matrix* $W_t = (w_{x,k})_{0 \le x,k<n}$ by setting $w_{x,k}$ to be the number of walks of length k from x to t. Let

$$W_t(x) = (w_{x,0}, w_{x,1}, \ldots w_{x,n-1})$$

be the row of W_t corresponding to the vertex x. If ϕ is an isomorphism from G_t to H_u, then clearly $W_u(\phi(x)) = W_t(x)$. This means that $W_t(x)$ can be used as a canonical label for a vertex x in G_t. We call G_t *walk-discrete* if $W_t(x) \neq W_t(x')$ for all $x \neq x'$. Thus, the walk matrix yields a canonical labeling for the class of walk-discrete digraphs with an individualized vertex.

As it was mentioned in Section 2, the walk matrix is efficiently computable by linear algebraic operations. For walk-discrete digraphs, the corresponding canonical labeling can also be obtained combinatorially via the *color refinement* algorithm (CR), which we now describe formally.

Given a vertex-colored digraph G with initial coloring C_0, CR iteratively computes new colorings

$$C_{i+1}(x) = \Big(C_i(x), \{\!\{C_i(y)\}\!\}_{y \in N(x)}\Big), \tag{1}$$

where $\{\!\{\}\!\}$ denotes a multiset and $N(x) = \{y : (x, y) \in E\}$ is the out-neighborhood of x. For each i, the coloring C_i is canonical, i.e., isomorphism-invariant. Indeed, an easy induction on i shows that if ϕ is a (color-preserving) isomorphism from G to a vertex-colored graph H, then

$$C_i(\phi(x)) = C_i(x). \tag{2}$$

The color classes of C_{i+1} refine the color classes of C_i: if $C_i(x) \neq C_i(x')$, then $C_{i+1}(x) \neq C_{i+1}(x')$. The algorithm terminates after performing n refinement rounds, where n is the number of vertices in G, and outputs the coloring C_n. Note that the color partition becomes stable by this point.

A relationship between CR and the walk counts was observed in [39]. We use the following adaptation of this result for our purposes.

Lemma 3.1. *Let $t \in V(G)$ and C_0 be the vertex coloring associated with G_t, that is, $C_0(x) = C_0(x')$ and $C_0(x) \neq C_0(t)$ for all $x \neq t$ and $x' \neq t$. If $w_{x,k} \neq w_{x',k}$, then $C_k(x) \neq C_k(x')$.*

Proof. Using the induction on k, we prove that $w_{x,k} = w_{x',k}$ whenever $C_k(x) = C_k(x')$. In the base case of $k = 0$, the equalities $w_{x,0} = w_{x',0}$ and $C_0(x) = C_0(x')$ are equivalent by definition. Assume that $C_k(y) = C_k(y')$ implies $w_{y,k} = w_{y',k}$ for all y and y'. Let $C_{k+1}(x) = C_{k+1}(x')$. By the definition of the refinement step, we have $\{\!\{C_k(y)\}\!\}_{y\in N(x)} = \{\!\{C_k(y)\}\!\}_{y\in N(x')}$. Using the induction assumption, from here we derive the equality $\{\!\{w_{y,k}\}\!\}_{y\in N(x)} = \{\!\{w_{y,k}\}\!\}_{y\in N(x')}$. The equality $w_{x,k+1} = w_{x',k+1}$ now follows by noting that $w_{x,k+1} = \sum_{y\in N(x)} w_{y,k}$. □

Lemma 3.1 shows that a canonical labeling of a walk-discrete digraph G_t can be obtained by running CR on G_t rather than by directly computing the walk matrix, and we conclude this subsection by commenting on the efficiency of CR.

The inductive definition (1) leads to an exponential increase in the size of the color encoding. To prevent this, colors are renamed after each refinement step. In this way, we never need more than n distinct color names. To encode the colors after each refinement round, we can use, for example, the first n non-negative integers in binary. This allows us, in the next subsection, to refer to the *least* color having a certain property. Note that, once the renaming rule is fixed, the modified coloring remains canonical in the sense of Equality (2).

Finally, recall that CR can be implemented in time $O(n^2 \log n)$ [7, 12, 26], while preserving the canonicity of the final coloring with respect to graph isomorphism; see [7].

3.2 Cayley graphs of a cyclic group

Let $\mathbb{Z}_n$ denote a cyclic group of order n. More specifically, we let $\mathbb{Z}_n = \{0, 1, \ldots, n-1\}$ and consider the addition modulo n on this set. The *Cayley digraph* $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ is defined by a *connection set* $S \subseteq \mathbb{Z}_n\setminus\{0\}$ as follows: $V(X) = \mathbb{Z}_n$ and $(x, y) \in E(X)$ if and only if $y - x \in S$. Note that $S = N(0)$, the out-neighborhood of 0. If S is *inverse-closed*, i.e., $S = -S$, then $E(X)$ is symmetric and we speak of a *Cayley graph*. Cayley (di)graphs of $\mathbb{Z}_n$ are also called *circulant (di)graphs* or *circulants*.

For $u \in \mathbb{Z}_n$, let X_u be the vertex-individualized version of X. Since X is vertex-transitive, all X_u are isomorphic to each other, and we can speak about X_0 without loss of generality. Clearly, in order to canonize X, it is sufficient to canonize X_0. Therefore, the canonization method in the preceding subsection applies to any Cayley digraph $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ provided that X_0 is walk-discrete. We just have to individualize an arbitrary vertex of X and then run CR.

This method does not work for circulant *graphs*. Indeed, define $\rho : \mathbb{Z}_n \to \mathbb{Z}_n$ by $\rho(x) = -x$. If $S = -S$, then ρ is an automorphism of $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ and, hence, $W_0(\rho(x)) = W_0(x)$. This implies that the walk matrix W_0 has at most $\lceil (n+1)/2 \rceil$ different rows, and X_0 cannot be walk-discrete. If this maximum is attained, we call X_0 *walk-saturated.*

Lemma 3.2. *Let $X = \mathrm{Cay}(\mathbb{Z}_n, S)$, where $S = -S$, and suppose that X_0 is walk-saturated. Fix $u \neq 0$ such that $u \neq n/2$ if n is even. Then*

$$(W_0(x), W_u(x)) \neq (W_0(y), W_u(y))$$

for any two different vertices x and y of X.

Proof. Since X_0 is walk-saturated, the equality $W_0(x) = W_0(y)$ for $x \neq y$ is possible only if $y = \rho(x)$, i.e., $y = -x$ in $\mathbb{Z}_n$. Note that $W_u(x) = W_0(x - u)$. Therefore, the equality $W_u(x) = W_u(y)$ implies that $y = 2u - x$. The equalities $y = -x$ and $y = 2u - x$ can be fulfilled simultaneously only if $2u = 0$, which is excluded. □

Lemma 3.2 justifies the correctness of the following algorithm for the class of walk-saturated circulant graphs.

Canonical labeling algorithm

Input: a circulant graph X.

1. Individualize an arbitrary vertex of X. By vertex-transitivity, we can without loss of generality assume that the individualized vertex is 0.
2. Run CR on X_0. Let C be the obtained coloring of the vertex set.
3. Let c be the least color such that there are exactly two vertices u_1 and u_2 with $C(u_1) = C(u_2) = c$. If such c does not exist, then give up. Let u be any of u_1 and u_2. Individualize u in X.
4. Run CR on X_u. Let C' be the obtained coloring.
5. To each vertex x, assign the label $L(x) = (C(x), C'(x))$.
6. Check that all labels $L(x)$ are pairwise distinct. If not, then give up.

For each circulant input graph, our canonization algorithm either produces a vertex labeling L or explicitly gives up (doing always the same for isomorphic inputs). The labeling L is canonical because it does not depend on which vertex is individualized in Step 1 (by vertex-transitivity) nor in Step 3 (because u_1 and u_2 are interchangeable by an automorphism of X_0). Lemma 3.2 ensures that the algorithm succeeds whenever X_0 is walk-saturated, and this will allow us to estimate the success probability based on Lemma 2.2 and Theorem 2.3.

Finally, we remark that if the algorithm is run on a non-circulant input graph and outputs a vertex labeling, then this labeling does not need to be canonical. To make it canonical in all cases, Steps 1 and 3 have to be performed for all possible individualized vertices, which can yield $2n$ different labelings $L_1, \ldots, L_{2n}$. Of all these candidate labelings, the algorithm chooses that which yields the isomorphic copy of X with lexicographically least adjacency matrix. The running time of this algorithm variant is $O(n^3 \log n)$. The similar modification is as well possible in the case of digraphs.

4 The walk matrix and the spectrum of circulants

4.1 The spectrum of a circulant

Let $A = (a_{ij})$ be the adjacency matrix of a circulant digraph X, that is, A is the 0–1 matrix whose rows and columns are indexed by the elements $0, 1, \ldots, n-1$ of $\mathbb{Z}_n$ such that $a_{ij} = 1$ exactly when $(i, j) \in E(X)$, that is, $j - i \in S$. Note that A is a *circulant matrix*, which means that the $(i+1)$-th row of A is obtained from its i-th row by the cyclic shift in one element to the right.

Let ω be an n-th root of unity, i.e., $\omega \in \mathbb{C}$ and $\omega^n = 1$. For the vector $V_\omega = (1, \omega, \omega^2, \ldots, \omega^{n-1})^\top$, the definition of a circulant matrix easily implies the equality

$$A\,V_\omega = \left(a_0 + a_1\omega + a_2\omega^2 + \cdots + a_{n-1}\omega^{n-1}\right)\,V_\omega = \left(\sum_{j\in S}\omega^j\right)V_\omega,$$

where $(a_0, a_1, a_2, \ldots, a_{n-1})$ is the first row of A, that is, the characteristic vector of $S \subset \mathbb{Z}_n$. We conclude that V_ω is an eigenvector of A corresponding to the eigenvalue $\lambda_{\omega,S} = \sum_{j\in S}\omega^j$.

Now, let $\omega = \zeta_n$ be a primitive n-th root of unity. To be specific, we fix $\zeta_n = e^{-2\pi\mathrm{i}/n}$. The n column vectors V_ω for $\omega = \zeta_n^0, \zeta_n^1, \zeta_n^2, \ldots, \zeta_n^{n-1}$ form a Vandermonde matrix with non-zero determinant. It follows that these n vectors are linearly independent and, therefore, $\lambda_{\zeta_n^0,S}, \lambda_{\zeta_n^1,S}, \ldots, \lambda_{\zeta_n^{n-1},S}$ is the full spectrum of A. The i-th eigenvalue in this sequence will be below denoted by

$$\lambda_{i,S} = \sum_{j\in S}\zeta_n^{ij}. \tag{3}$$

4.2 Discrete Fourier transform

Let $\mathbb{C}^{\mathbb{Z}_n}$ denote the vector space of all functions from $\mathbb{Z}_n$ to the field of complex numbers $\mathbb{C}$ with pointwise addition and pointwise scalar multiplication. The pointwise multiplication on $\mathbb{C}^{\mathbb{Z}_n}$ will be denoted by $\circ$. Another way to introduce a product on $\mathbb{C}^{\mathbb{Z}_n}$ is to consider the convolution $\alpha * \beta$ of two functions $\alpha, \beta : \mathbb{Z}_n \to \mathbb{C}$, which is defined by $(\alpha * \beta)(x) = \sum_{y\in\mathbb{Z}_n}\alpha(x-y)\beta(y)$ for each $x \in \mathbb{Z}_n$. Both $\circ$ and $*$ are bilinear and, therefore, both $(\mathbb{C}^{\mathbb{Z}_n}, \circ)$ and $(\mathbb{C}^{\mathbb{Z}_n}, *)$ are n-dimensional algebras over $\mathbb{C}$. The algebra $(\mathbb{C}^{\mathbb{Z}_n}, *)$ can alternatively be seen as the *group algebra* of $\mathbb{Z}_n$ over $\mathbb{C}$ and, as such, it is semisimple by Maschke's theorem; see, e.g., [17, Section 7.1]. Like any two n-dimensional commutative semisimple $\mathbb{C}$-algebras, the algebras $(\mathbb{C}^{\mathbb{Z}_n}, *)$ and $(\mathbb{C}^{\mathbb{Z}_n}, \circ)$ are isomorphic (see [17, Corollary 2.4.2]). We now describe an explicit algebra isomorphism from $(\mathbb{C}^{\mathbb{Z}_n}, *)$ to $(\mathbb{C}^{\mathbb{Z}_n}, \circ)$.

For $T \subseteq \mathbb{Z}_n$, let $\chi_T \in \mathbb{C}^{\mathbb{Z}_n}$ be the characteristic function of T. In particular, $\chi_{\mathbb{Z}_n}$ is the identically one function. For $x \in \mathbb{Z}_n$, we set $\delta_x = \chi_{\{x\}}$.

The *discrete Fourier transform (DFT)* is the linear operator $\mathcal{F} : \mathbb{C}^{\mathbb{Z}_n} \to \mathbb{C}^{\mathbb{Z}_n}$ mapping a function $\alpha : \mathbb{Z}_n \to \mathbb{C}$ into the function $\mathcal{F}(\alpha) : \mathbb{Z}_n \to \mathbb{C}$ defined by

$$\mathcal{F}(\alpha)(i) = \sum_{j=0}^{n-1}\zeta_n^{ij}\alpha(j). \tag{4}$$

In the standard basis $\delta_0, \delta_1, \ldots, \delta_{n-1}$, the DFT is represented by the matrix $F = (\zeta_n^{ij})_{i,j\in\mathbb{Z}_n}$. Since F is the familiar Vandermonde matrix with non-zero determinant, the map F is a linear isomorphism from $\mathbb{C}^{\mathbb{Z}_n}$ onto itself. It is well known and easy to derive from the definitions that

$$\mathcal{F}(\alpha * \beta) = \mathcal{F}(\alpha) \circ \mathcal{F}(\beta). \tag{5}$$

4.3 The rank of the walk matrix

We are now prepared to derive Lemmas 2.1 and 2.2 from the following more general fact.

Lemma 4.1. *$X = \mathrm{Cay}(\mathbb{Z}_n, S)$ has exactly $\mathrm{rk}\, W_0$ distinct eigenvalues, where W_0 is the walk matrix of X_0.*

Proof. Let $X' = \mathrm{Cay}(\mathbb{Z}_n, -S)$ be the transpose of the digraph X, i.e., the digraph obtained from X by reversing all arcs. Note that the adjacency matrices of X and X' are transposes of each other and therefore have the same spectrum. Denote the number of distinct eigenvalues of X', and hence also of X, by R.

A column vector $(a_0, a_1, \ldots, a_{n-1})^\top$ will be naturally identified with the function $\alpha \in \mathbb{C}^{\mathbb{Z}_n}$ defined by $\alpha(x) = a_x$ for all $x \in \mathbb{Z}_n$. In this way, the columns of the walk matrix W_0 correspond to the functions $\eta_0, \eta_1, \ldots, \eta_{n-1}$ where $\eta_k(x) = w_{x,k}$. Thus, the rank of W_0 is equal to the dimension of the linear subspace U of $\mathbb{C}^{\mathbb{Z}_n}$ spanned by these functions. We, therefore, have to prove that $\dim U = R$.

Note that

$$\eta_{k+1}(x) = \sum_{y\in N(x)} \eta_k(y) = \sum_{y\in\mathbb{Z}_n} \chi_S(y-x)\eta_k(y)$$
$$= \sum_{y\in\mathbb{Z}_n} \chi_{-S}(x-y)\eta_k(y) = (\chi_{-S} * \eta_k)(x).$$

It follows that $\eta_0 = \delta_0$, $\eta_1 = \chi_{-S}$, $\eta_2 = \chi_{-S} * \chi_{-S}$ and, generally, $\eta_k = \chi_{-S}^{*(k)}$ is the $(k-1)$-fold convolution of k copies of the characteristic function χ_{-S} of the set $-S$.

Let us apply the discrete Fourier transform $\mathcal{F}$. Note that $\mathcal{F}(\delta_0)$ is the all-ones vector. As easily seen from (3) and (4), $\mathcal{F}(\chi_{-S})$ is the vector whose entries are the eigenvalues $\lambda_{0,-S}, \lambda_{1,-S}, \ldots, \lambda_{n-1,-S}$ of the transpose $X' = \mathrm{Cay}(\mathbb{Z}_n, -S)$. Equality (5) readily implies that the matrix formed by the column vectors $\mathcal{F}(\eta_0), \mathcal{F}(\eta_1), \ldots, \mathcal{F}(\eta_{n-1})$ has exactly R distinct rows. Consequently, $\dim U = \dim \mathcal{F}(U) \le R$. In fact, equality holds because the aforementioned matrix contains a non-degenerate Vandermonde matrix of size $R \times R$ as a submatrix. □

Lemmas 2.1 and 2.2 correspond, respectively, to the special cases $\mathrm{rk}\, W_0 = n$ and $\mathrm{rk}\, W_0 = \lceil (n+1)/2 \rceil$ of Lemma 4.1. If $S = -S$, then $\mathrm{rk}\, W_0 \le \lceil (n+1)/2 \rceil$ due to the symmetry of the undirected graph $X = \mathrm{Cay}(\mathbb{Z}_n, S)$. Lemma 4.1 shows that X can have at most $\lceil (n+1)/2 \rceil$ distinct eigenvales[4], and that this maximum is attained exactly when $\mathrm{rk}\, W_0$ attains the same maximum value $\lceil (n+1)/2 \rceil$.

[4] This can be seen also directly as Equality (3) implies that $\lambda_{a,S} = \lambda_{b,S}$ for $a \neq b$ whenever $a + b = n$.

5 Proof of Theorem 2.3

We set $\zeta_n = e^{-2\pi \mathrm{i}/n}$. As discussed in Subsection 4.1, a circulant $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ has eigenvalues $\lambda_0, \lambda_1, \ldots, \lambda_{n-1}$ where

$$\lambda_a = \sum_{j \in S} \zeta_n^{aj} = \sum_{j=0}^{n-1} \chi_S(j)\zeta_n^{aj}.$$

Let $\sigma_j = \chi_S(j)$. If X is a random digraph, i.e., the connection set S is chosen uniformly at random among all subsets of $\mathbb{Z}_n \setminus \{0\}$, then $\sigma_1, \sigma_2, \ldots, \sigma_{n-1}$ is a Bernoulli process, that is, these $n-1$ random variables are independent and identically distributed with σ_j taking each of the values 0 and 1 with probability 1/2. If X is a random graph, i.e., the connection set $S = -S$ is chosen randomly among all inverse-closed subsets, then the values $\sigma_1, \sigma_2, \ldots, \sigma_{\lfloor n/2 \rfloor}$ form a Bernoulli process, and the remaining values are determined by the equality $\sigma_j = \sigma_{n-j}$. For each $a = 0, 1, \ldots, n-1$, the eigenvalue

$$\lambda_a = \sum_{j=1}^{n-1} \sigma_j \zeta_n^{aj} \tag{6}$$

becomes a random variable taking its values in the cyclotomic field $\mathbb{Q}(\zeta_n)$.

We will use the following observation. As usually, $\phi(n)$ stands for Euler's totient function.

Lemma 5.1. *No two different subsets of $\{\zeta_n^j : 1 \le j \le n/\ln n\}$ have equal sums of their elements.*

Proof. The known lower bounds for $\phi(n)$ (see, e.g., [6, Thm. 8.8.7]) imply that $\phi(n) > n/\ln n$ for $n \ge 3$. The existence of two different subsets with equal sums would therefore yield a non-trivial linear combinations with rational coefficient of $1, \zeta_n, \zeta_n^2, \ldots, \zeta_n^{\phi(n)-1}$, contradicting the fact that these numbers form a basis of $\mathbb{Q}(\zeta_n)$ considered as a vector space over $\mathbb{Q}$. □

Our overall strategy for proving Theorem 2.3 will be to use the union bound

$$\mathbb{P}[\lambda_a = \lambda_b \text{ for some } 0 \le a, b \le n-1] \le \sum_{0 \le a, b \le n-1} \mathbb{P}[\lambda_a = \lambda_b] \tag{7}$$

and to show that the right hand side is bounded by $n^{-1/2+o(1)}$. The summation ranges over unequal a and b; in the undirected case, we additionally require $a + b \neq n$. For a fixed pair $\{a, b\}$, we have to show that $\lambda_a = \lambda_b$ occurs with sufficiently small probability. Using (6), this equality can be rewritten as

$$\sum_{j=1}^{n-1} \sigma_j \zeta_n^{aj} = \sum_{j=1}^{n-1} \sigma_j \zeta_n^{bj} \tag{8}$$

The basic idea is to “distill” a sufficiently large set of indices $J \subset \{1, \ldots, n-1\}$ such that exposing all random variables σ_j for $j \notin J$ converts (8) into an equality

$$\sum_{j \in J} \sigma_j \zeta_n^{aj} = \sum_{j \in J} \sigma_j \zeta_n^{bj} + \text{const} \tag{9}$$

that can be satisfied by at most one assignment to the remaining random variables σ_j for $j \in J$. The last condition immediately implies the bound $\mathbb{P}[\lambda_a = \lambda_b] = 2^{-|J|}$, which will be strong enough for our purposes. To justify that (9) has at most one satisfying assignment, we will crucially rely on Lemma 5.1.

Let

$$\mathbb{U} = \{z \in \mathbb{C} \,:\, |z| = 1\}$$

denote the unit circle in the complex plane. Lemma 5.1 will be applicable not only when aj and bj (modulo n) do not exceed $n/\ln n$ for all $j \in J$—which would be technically difficult to ensure—but also whenever the set $\{\zeta_n^{aj}\}_{j\in J} \cup \{\zeta_n^{bj}\}_{j\in J}$ is contained within an arbitrary arc of $\mathbb{U}$ of length $2\pi/\ln n$. In such cases, we can “rotate” this set by multiplying both sides of (9) by a suitable root of unity ζ_n^t. This will transform (9) into a linear combination of the complex numbers ζ_n^{aj+t} and ζ_n^{bj+t} with exponents (modulo n) lying in the range covered by Lemma 5.1.

The choice of a suitable set J depends on specific properties of the pair $\{a, b\}$. We divide all such pairs into three categories and describe an appropriate “distillation” of J separately for each of the three cases. The relevant properties of a pair $\{a, b\}$ are expressed in terms of elementary number-theoretic parameters, which we now introduce.

Given $z \in \mathbb{C}^\times$, we write $\langle z\rangle$ to denote the cyclic subgroup of the multiplicative group $\mathbb{C}^\times$ generated by z. For an integer c, let $g(c) = |\langle\zeta_n^c\rangle|$ denote the order of the element ζ_n^c in the group $\langle\zeta_n\rangle$. Note that $g(c) = n/\gcd(c, n)$. We also define

$$c' = c/\gcd(c, n),$$

noting that

$$\zeta_n^c = \zeta_{g(c)}^{c'}. \tag{10}$$

The above notation will be used both for $c = a$ and $c = b$. We set

$$g = g(a) \text{ and } h = g(b),$$

and suppose, without loss of generality, that

$$h \le g.$$

Since $\gcd(a', g) = 1$, the integer a' can be regarded as an invertible element of the ring $\mathbb{Z}_g$. Let r denote the inverse of a' in $\mathbb{Z}_g$, i.e., r is the smallest positive integer for which $ra' = 1 \pmod g$.

Note that $\langle\zeta_n^a\rangle = \langle\zeta_g\rangle$. Indeed, $\zeta_n^a \in \langle\zeta_g\rangle$ by (10). On the other hand, (10) also implies that $\zeta_g = \zeta_g^{ra'} = \zeta_n^{ar}$ belongs to $\langle\zeta_n^a\rangle$. Set

$$\xi = \zeta_g = \zeta_n^{ar} \text{ and } \eta = \zeta_n^{br}. \tag{11}$$

Thus, all $g-1$ non-unity elements of the multiplicative group $\langle \zeta_n^a \rangle = \langle \xi \rangle$ can be listed as

$$\xi = \zeta_n^{ar},\ \xi^2 = \zeta_n^{2ar},\ \dots\ , \xi^{g-1} = \zeta_n^{(g-1)ar}. \tag{12}$$

It is useful to note that they appear in the left hand side of Equality (8) for the indices $j = r, 2r, \dots, (g-1)r$, which are understood modulo n. The same positions in the right hand side of Equality (8) are occupied by

$$\eta = \zeta_n^{br},\ \eta^2 = \zeta_n^{2br},\ \dots\ , \eta^{g-1} = \zeta_n^{(g-1)br}. \tag{13}$$

The importance of the parameter g stems from the fact that the elements of (12), together with the unity $\xi^g = 1$, partition the unit circle $\mathbb{U}$ into g arcs, each of length $2\pi/g$. When g is large, this gives us better chances to obtain a linear combination (9) such that the degrees ζ_n^{aj} involved in (9) are sufficiently close to one another on $\mathbb{U}$, allowing us to apply Lemma 5.1 as described above.

The minimum distance in $\mathbb{U}$ between two distinct elements $\zeta_n^{brk_1} = \eta^{k_1}$ and $\zeta_n^{brk_2} = \eta^{k_2}$ is important by the same reason. This distance is equal to $2\pi/g'$ where

$$g' = |\langle \eta \rangle|$$

is the order of the element η in $\langle \zeta_n \rangle$. Since $\langle \eta \rangle \subseteq \langle \zeta_n^b \rangle$, we have

$$g' \le h \le g.$$

While all numbers in (12) are distinct, the number of distinct numbers in the sequence (13) is equal to g' if $g' < g$ and to $g'-1$ if $g' = g$.

The parameters g and g' are crucial for our analysis, and we partition all pairs $\{a, b\}$ into three groups based on the values of g and g'. Recall that we assume $g(b) \le g(a)$. In cases where $g(b) = g(a)$, we will further assume that $b < a$. This additional assumption does not play any substantive role in the analysis but allows us to treat a, b as an ordered pair.

We now outline our argument for each of the three cases—Cases A, B, and C—described below. After presenting the main ideas, we will proceed to a detailed proof, organized into Claims A, B, and C, respectively. Note that Case C is responsible for the probability bound $n^{-1/2+o(1)}$ in Theorem 2.3, as the estimates for the probability of multiple eigenvalues obtained in Cases A and B are stronger. We define Case C by assuming that the parameter g is relatively small in terms of n. Notably, this restriction simplifies the analysis, because both sides of Equality (8) can be viewed as linear combinations of a few independent binomial variables. We begin the outline with the most technically demanding case.

Case A: g and g' are large.

This case is further divided into three subcases. In the first two of them, we can find a sufficiently large set $K \subseteq \{1, \dots, n-1\}$, specifically of size $|K| = \Omega(\ln^2 n)$, such that:

(i) the subsequence $(\xi^k)_{k\in K}$ of the sequence (12) consists of distinct elements, and the same holds true for the corresponding subsequence $(\eta^k)_{k\in K}$ of (13);

(ii) the sets $\{\xi^k\}_{k\in K}$ and $\{\eta^k\}_{k\in K}$ are disjoint;

(iii) both sets are contained within a sufficiently short arc of $\mathbb{U}$, specifically of length less than $1/\ln n$.

We now argue that properties (i)–(iii) allow us to bound the probability of (8) based on Lemma 5.1.

Let us expose the values of all random variables σ_j excepting σ_{kr} with $k \in K$ (recall that the indices are considered modulo n). Equality (8) can now be written as

$$c_1 + \sum_{k\in K} \sigma_{kr}\xi^k = c_2 + \sum_{k\in K} \sigma_{kr}\eta^k \tag{14}$$

for some constants $c_1, c_2 \in \mathbb{C}$. In other words, we estimate the probability of the event that

$$c_1 + \sum_{k\in K'} \xi^k = c_2 + \sum_{k\in K'} \eta^k \tag{15}$$

for a random subset $K' \subseteq K$. We show that this equality can be true for at most one K'.

Indeed, assume that Equality (15) holds true for two different subsets $K' = K_1$ and $K' = K_2$ of K. This implies that $\sum_{k\in K'}(\xi^k - \eta^k) = c_2 - c_1$ for both $K' = K_1$ and $K' = K_2$ and, therefore,

$$\sum_{k\in K_1} \xi^k + \sum_{k\in K_2} \eta^k = \sum_{k\in K_2} \xi^k + \sum_{k\in K_1} \eta^k. \tag{16}$$

By Conditions (i)–(ii), both sides of (16) are sums of $|K_1| + |K_2|$ distinct numbers. Since $K_1 \neq K_2$, we can without loss of generality suppose that $K_1 \not\subseteq K_2$. Taking $k \in K_1 \backslash K_2$, we see that the number ξ^k occurs only in the left hand side of (16). Thus, we have equality of two sums of ζ_n^j over different sets of indices j. By multiplying both sides of Equation (16) by a suitable ζ_n^t, and using Condition (iii), we can "rotate" these index sets modulo n so that they lie within the interval $[1, n/\ln n]$. This leads to a contradiction with Lemma 5.1.

We conclude that Equality (15), and therefore also Equality (8), holds with probability at most $2^{-|K|} = n^{-\Omega(\ln n)}$.

The third subcase of Case A, which relies on the same idea and is somewhat simpler, is omitted from this outline (it appears as Case 3 in the proof of Claim A below).

Case B: g is large, and g' is small.

Since $\xi = \zeta_g = e^{-2\pi \mathrm{i}/g}$, the first m elements of the sequence (12) are contained in an arc of $\mathbb{U}$ of length $2\pi m/g$. The corresponding elements of the sequence (13) can take on at most g' different values. Therefore, there is a set $K \subseteq \{1, \dots, m\}$ of size

$$|K| \geq m/g'$$

such that if $k \in K$, then $\eta^k = \eta'$ for the same η'. Expose all σ_j excepting σ_{kr} for $k \in K$. Similarly to Case A, Equality (8) converts into Equality (14) for some constants c_1 and c_2. This event can be recast as

$$c_1 + \sum_{k \in K'} \xi^k = c_2 + |K'|\eta' \tag{17}$$

for a random subset $K' \subseteq K$. By Chernoff's bound, the size of K' is concentrated in the interval $\frac{1}{2}|K| - |K|^{2/3} < |K'| < \frac{1}{2}|K| + |K|^{2/3}$ with probability no less than $1 - 2\exp(-|K|^{1/3})$. Fix an integer m' such that $\frac{1}{2}|K| - |K|^{2/3} < m' < \frac{1}{2}|K| + |K|^{2/3}$ and consider the event (17) conditioned on $|K'| = m'$. Under this condition, Equality (17) reads

$$\sum_{k \in K'} \xi^k = c_2 + m'\eta' - c_1. \tag{18}$$

We choose m such that

$$2\pi m/g < 1/\ln n \tag{19}$$

while m/g', and hence $|K|$, is large. Condition (19) ensures that Lemma 5.1 applies, implying that Equality (18) holds for at most one subset $K' \subseteq K$. It follows that this equality holds with probability at most $1/\binom{|K|}{m'}$. We therefore conclude that Equality (8) holds with probability at most

$$1/\binom{|K|}{|K|/2 + |K|^{2/3}} + 2\exp\left(-|K|^{1/3}\right). \tag{20}$$

The specification of "large" g and "small" g' in the formal argument below ensures that this probability is subexponentially small.

Case C: g is small.

This case can be treated without invoking Lemma 5.1. Let $J = \{j < n : \zeta_n^{aj} = \xi\}$, and note that $|J| = n/g$. The set J contains a subset J' of size $|J'| \geq |J|/h \geq n/g^2$ such that the values ζ_n^{bj} are all equal for $j \in J'$. That is, for all $j \in J'$, we have $\zeta_n^{aj} = \xi$ and $\zeta_n^{bj} = \eta'$ where η' is the same h-th root of unity. Moreover, $J' \subseteq J$ can be chosen so that $\xi \neq \eta'$; see the proof of Claim C below for details.

Let us expose all random variables σ_j except those for $j \in J'$. Equality (8) then implies that

$$\sum_{j \in J'} \sigma_j(\xi - \eta') = c$$

for a constant $c \in \mathbb{C}$ and, therefore, the sum $\sum_{j \in J'} \sigma_j$ evaluates to a constant value. The probability of the last event is bounded by $\binom{|J'|}{\lfloor |J'|/2 \rfloor} 2^{-|J'|} \leq |J'|^{-1/2} \leq gn^{-1/2}$. Although this bound is weaker than those obtained in Cases A and B, it is still sufficient for use with the union bound (7) because, as we will see, the number of pairs a, b covered by Case C is relatively small.

After this outline, we proceed to the detailed proof. We begin by noting that the numbers ξ and η, defined by (11), are distinct. Indeed, the equality $\xi = \eta$ would imply that $\langle \zeta_n^b \rangle$ contains $\langle \eta \rangle = \langle \xi \rangle = \langle \zeta_n^a \rangle$, showing that $h \geq g$ and, hence, $h = g$ in this case. By (10), the last equality leads to

$$\zeta_g^{a'r} = \zeta_n^{ar} = \zeta_n^{br} = \zeta_h^{b'r} = \zeta_g^{b'r},$$

which holds only if $r(a'-b')$ is divisible by g. Since r and g are coprime, we conclude that $a'-b'$ must be divisible by g. Since a' and b' are strictly smaller than g, this is possible only when $a'=b'$, contradicting the assumption that $a \neq b$.

Let $\varepsilon > 0$ be an arbitrarily small constant. Once this parameter is fixed, we assume that n is sufficiently large. We divide the remainder of the proof into Claims A, B, and C, corresponding to Cases A, B, and C discussed above. We present a detailed argument for digraphs (Part 1 of the theorem), which, with minor modifications, also applies to graphs (Part 2). We comment on these modifications at the end of the proof.

For the argument $\arg(z)$ of a complex number z, we suppose that $\arg(z) \in [0, 2\pi)$.

Claim A. Let $P_1 = \{(a,b) : g \geq n^{6\varepsilon} \text{ and } g' \geq n^{\varepsilon}\}$. Then

$$\sum_{(a,b)\in P_1} \mathbb{P}[\lambda_a = \lambda_b] \leq n^{-0.5\ln n}.$$

Proof. We claim that there exists s such that

$$1 \leq s \leq \lceil \ln^4 n \rceil \tag{21}$$

and

$$\text{either } 0 < \arg(\eta^s) \leq 2\pi/\ln^4 n \text{ or } 0 < \arg(\eta^{-s}) \leq 2\pi/\ln^4 n.$$

Indeed, consider the set $\mathcal{S} = \{\eta^s : 1 \leq s \leq \lceil \ln^4 n \rceil\}$. All elements of $\mathcal{S}$ are pairwise distinct. Indeed, suppose that $\eta^{s_1} = \eta^{s_2}$ for $1 \leq s_1 < s_2 \leq \lceil \ln^4 n \rceil$. Then $\eta^{s_2 - s_1} = 1$, so $s_2 - s_1$ must be a multiple of the order of η, which is g'. This yields a contradiction because, since n is assumed to be sufficiently large, we have

$$s_2 - s_1 < \ln^4 n < n^{\varepsilon} \leq g'.$$

Let η^{s_1} and η^{s_2}, with $s_1 < s_2$, be two elements of $\mathcal{S}$ such that the distance between them on the unit circle $\mathbb{U}$ is minimal among all such pairs. Since this distance is at most $2\pi/\ln^4 n$, we can set $s = s_2 - s_1$.

We now consider three cases and show that, in each of them,

$$\mathbb{P}[\lambda_a = \lambda_b] \leq n^{-0.6\ln n}. \tag{22}$$

Case 1: $\arg(\eta^s) \leq 2\pi/\ln^4 n$.

We follow the strategy presented in Case A of the outline above. To establish the upper bound (22), it suffices to find a set $K \subseteq \{1, \dots, n-1\}$ of size $|K| = \lceil \ln^2 n \rceil$ that satisfies Conditions (i)–(iii) stated above. We set $K = \big\{s, 2s, \dots, \lceil \ln^2 n \rceil s\big\}$.

Condition (i) for this set is ensured by the estimate $\arg(\xi) = 2\pi - 2\pi/g \geq 2\pi - 2\pi/n^{6\varepsilon}$ and by the definition of Case 1. Since n is large enough, due to (21) we have $\arg(\xi^{\ell s}) > 2\pi - 2\pi/n^{5\varepsilon}$ for all $\ell \leq \lceil \ln^2 n \rceil$. In the case under consideration we also have $\arg(\eta^{\ell s}) < \pi/\ln n$ for all $\ell \leq \lceil \ln^2 n \rceil$. Consequently,

$$\big\{\arg(\xi^k)\big\}_{k\in K} \subset (2\pi - \pi/\ln n, 2\pi) \text{ and } \big\{\arg(\eta^k)\big\}_{k\in K} \subset (0, \pi/\ln n).$$

This implies the other two Conditions (ii) and (iii), yielding the upper bound

$$\mathbb{P}[\lambda_a = \lambda_b] \leq 2^{-|K|} \leq 2^{-\ln^2 n} \leq n^{-0.6 \ln n}. \tag{23}$$

Case 2: $\arg(\eta^{-s}) \leq 2\pi/\ln^4 n$ *and* $\eta^s \neq \xi^s$.

The first of the two assumptions ensures—similarly to Case 1—that Condition (i) is fulfilled for every set $K \subseteq \{s, 2s, \ldots, (2\lceil \ln^2 n \rceil)s\}$. Each such set K also satisfies Condition (iii) because, similarly to Case 1, we have

$$\{\arg(\xi^k)\}_{k\in K} \subset (2\pi - 2\pi/\ln n, 2\pi) \text{ and } \{\arg(\eta^k)\}_{k\in K} \subset (2\pi - 2\pi/\ln n, 2\pi).$$

These inclusions, however, do not guarantee Condition (ii). Nevertheless, we can show that Condition (ii) holds for at least one set $K \subseteq \{s, 2s, \ldots, (2\lceil \ln^2 n \rceil)s\}$ such that $|K| \geq \lceil \ln^2 n \rceil$.

Indeed, assume without loss of generality that $\arg(\xi^{-s}) < \arg(\eta^{-s})$. The existence of a suitable set K follows from the observation that, among any two consecutive values $\arg(\xi^{\ell s})$ and $\arg(\xi^{(\ell+1)s})$ for $\ell < 2\lceil \ln^2 n \rceil$, at most one can belong to the set $\{\arg(\eta^k)\}_{k\in\{s,2s,\ldots,(2\lceil \ln^2 n\rceil)s\}}$, because the distance between any two neighboring elements of this set on the unit circle $\mathbb{U}$ (which is equal to $\arg(\eta^{-s})$) is larger than the distance between $\xi^{\ell s}$ and $\xi^{(\ell+1)s}$ (which is equal to $\arg(\xi^{-s})$). It follows that the set K, consisting of all those ℓs with $\ell \leq 2\lceil \ln^2 n \rceil$ such that $\arg(\xi^{\ell s}) \notin \{\arg(\eta^k)\}_{k\in\{s,2s,\ldots,(2\lceil \ln^2 n\rceil)s\}}$, has the desired size. Condition (ii) is satisfied for this set by construction, and we obtain the upper bound (23) also in this case.

Case 3: $\arg(\eta^{-s}) \leq 2\pi/\ln^4 n$ *and* $\eta^s = \xi^s$.

Let $K = \{0, s, 2s, \ldots, \lceil \ln^2 n \rceil s\}$. Following our general strategy, let us expose all random variables σ_j excepting $\sigma_{(k+1)r}$ for $k \in K$. The event (8) can now be recast as the equality

$$\sum_{k\in K'} (\xi^{k+1} - \eta^{k+1}) = c \tag{24}$$

for a constant $c \in \mathbb{C}$ and a random subset $K' \subseteq K$. For $k = \ell s$, we have

$$\xi^{k+1} - \eta^{k+1} = \xi^{\ell s+1} - \eta^{\ell s+1} = \xi^{\ell s+1} - \xi^{\ell s}\eta = \xi^{\ell s}(\xi - \eta),$$

which allows us to rewrite Equality (24) as

$$\sum_{k\in K'} \xi^k = c/(\xi - \eta) \tag{25}$$

(recall that $\xi \neq \eta$). Since $\arg(\xi^s) = \arg(\eta^s) \geq 2\pi - 2\pi/\ln^4 n$, we have $\arg(\xi^{\ell s}) > 2\pi - 2\pi/\ln n$ for all $\ell \leq \lceil \ln^2 n \rceil$. We can, therefore, use Lemma 5.1 to conclude that Equality (25) can be true for at most one $K' \subseteq K$. It follows that Equality (24), and as well Equality (8), holds with probability at most $2^{-|K|} \leq n^{-0.6 \ln n}$.

Thus, Bound (22) is established in each of the three cases. Consequently,

$$\sum_{(a,b)\in P_1} \mathbb{P}[\lambda_a = \lambda_b] \le |P_1| \cdot n^{-0.6\ln n} \le n^{-0.5\ln n},$$

where the last inequality holds because n is assumed to be sufficiently large. The proof of the claim is complete. □

Claim B. Let $P_2 = \{(a,b) \,:\, g \ge n^{6\varepsilon} \text{ while } g' < n^{\varepsilon}\}$. Then

$$\sum_{(a,b)\in P_2} \mathbb{P}[\lambda_a = \lambda_b] \le \exp\left(-n^{\varepsilon}\right).$$

Proof. We closely follow the proof strategy presented (in a fairly complete form) in Case B of the outline above. Specifically, we set $m = n^{5\varepsilon}$. Since n is assumed to be sufficiently large, this implies $|K| \ge m/g' > n^{4\varepsilon}$ and also ensures Bound (19). This justifies the probability bound (20). Using the estimate $\binom{n}{k} \ge (n/k)^k$, we conclude from (20) that if $(a,b) \in P_2$, then

$$\mathbb{P}[\lambda_a = \lambda_b] < \left(\frac{1}{2} + |K|^{-1/3}\right)^{|K|/2} + 2\exp\left(-|K|^{1/3}\right).$$

Since n is sufficiently large, this probability is bounded by $3\exp\left(-n^{4\varepsilon/3}\right)$ and, therefore,

$$\sum_{(a,b)\in P_2} \mathbb{P}[\lambda_a = \lambda_b] \le |P_2| \cdot 3\exp\left(-n^{4\varepsilon/3}\right) \le \exp\left(-n^{\varepsilon}\right),$$

as required. □

Claim C. Let $P_3 = \{(a,b) \,:\, g < n^{6\varepsilon}\}$. Then

$$\sum_{(a,b)\in P_3} \mathbb{P}[\lambda_a = \lambda_b] \le n^{-1/2+20\varepsilon}.$$

Proof. As explained in Case C of the outline, we begin by considering the set $J = \{j < n \,:\, \zeta_n^{aj} = \xi\}$. This set can be described explicitly as $J = \{r + ig \,:\, 0 \le i < n/g\}$. Define the subset $J' \subseteq J$ by $J' = \{r + igh \,:\, 0 \le i < n/(gh)\}$ and note that $\zeta_n^{bj} = \eta$ for all $j \in J'$. The size of J' can be bounded below by $n/(gh) \ge n/g^2 > n^{1-12\varepsilon}$.

Recall that $\eta \ne \xi$. Exposing all σ_j except those for $j \in J'$, we obtain from Equality (8) that $\sum_{j\in J'} \sigma_j = c$ for a constant c. This occurs with probability at most

$$\binom{|J'|}{\lfloor |J'|/2 \rfloor} 2^{-|J'|} < |J'|^{-1/2} < n^{-1/2+6\varepsilon},$$

providing us with an upper bound for $\mathbb{P}[\lambda_a = \lambda_b]$ if $(a,b) \in P_3$.

We now estimate the number of pairs (a,b) in P_3. Recall that $a = \gcd(a,n) \cdot a'$, where $a' \le n/\gcd(a,n) = g < n^{6\varepsilon}$. The factor $\gcd(a,n)$ of a can take on at most

$d(n)$ different values, where $d(n)$ denotes the total number of divisors of n. It is known [1, Theorem 13.12] that $d(n) = n^{O(1/\ln\ln n)}$. Since there are less than $n^{6\varepsilon}$ possibilities to choose the factor a', the element a can be chosen in at most $n^{7\varepsilon}$ ways, and the same holds true as well for b. It follows that $|P_3| \leq n^{14\varepsilon}$, and we conclude that

$$\sum_{(a,b)\in P_3} \mathbb{P}[\lambda_a = \lambda_b] < |P_3| \cdot n^{-1/2+6\varepsilon} \leq n^{-1/2+20\varepsilon},$$

completing the proof of the claim. □

Claims A, B, and C readily imply that

$$\sum_{0\leq a,b\leq n-1} \mathbb{P}[\lambda_a = \lambda_b] \leq n^{-1/2+\varepsilon}$$

for each $\varepsilon > 0$ and sufficiently large n. This completes the proof of Part 1 of Theorem 2.3. The proof of Part 2 proceeds in essentially the same way. The only difference is that, for an unexposed random variable σ_j, instead of the term $\sigma_j\zeta_n^{aj}$ we now deal with $\sigma_j\zeta_n^{aj} + \sigma_{n-j}\zeta_n^{a(n-j)} = \sigma_j(\zeta_n^{aj} + \zeta_n^{-aj})$. Lemma 5.1 remains applicable after multiplying the entire sum by an additional factor of $\zeta_n^{t'}$ for a small value of t'. Finally, in Case 1 of the proof of Claim A, one has to address also the possibility that $\eta^{-s} = \xi^s$, which is handled similarly to Case 3 of the proof.

Remark 5.2. The probability bound in Theorem 2.3 is nearly optimal. This can be shown by noticing that a random digraph $\mathrm{Cay}(\mathbb{Z}_n, S)$ for $n = 3p$ with p prime has repeated eigenvalues with probability $\Omega(n^{-1/2})$. Indeed, for $r = 0, 1, 2$, let $S_r = \{s \in S : s = r \pmod 3\}$. Note that

$$\lambda_p = \sum_{j\in S} \zeta_n^{pj} = \sum_{j\in S} \zeta_3^j = \sum_{j\in S_0} 1 + \sum_{j\in S_1} \zeta_3 + \sum_{j\in S_2} \zeta_3^2.$$

Likewise,

$$\lambda_{2p} = \sum_{j\in S} \zeta_n^{2pj} = \sum_{j\in S} \zeta_3^{2j} = \sum_{j\in S_0} 1 + \sum_{j\in S_1} \zeta_3^2 + \sum_{j\in S_2} \zeta_3.$$

As a consequence, $\lambda_p = \lambda_{2p}$ whenever $|S_1| = |S_2|$. The last equality is true with probability $(1-o(1))/\sqrt{\pi n/3}$ because $|S_1|$ and $|S_2|$ are independent random variables with probability distribution $\mathrm{Bin}(p, 1/2)$.

It can be similarly shown that, for $n = 5p$, the spectrum of a random graph $\mathrm{Cay}(\mathbb{Z}_n, S)$ is not saturated with the same probability bound $\Omega(n^{-1/2})$.

6 Proof of Theorem 1.1

We are now ready to prove our main result. Specifically, we prove Theorem 1.1 for each of the three concepts of a circulant:

- A Cayley (di)graph $X = \mathrm{Cay}(\mathbb{Z}_n, S)$. The uniform probability distribution of X means that the connection set S is equiprobably chosen among all subsets of $\mathbb{Z}_n \setminus \{0\}$ in the case of digraphs and among all inverse-closed subsets in the case of graphs.

- An unlabeled circulant, i.e., an isomorphism class of Cayley (di)graphs $X = \mathrm{Cay}(\mathbb{Z}_n, S)$. The uniform distribution means that each isomorphism class on $\mathbb{Z}_n$ is chosen equiprobably. In the algorithmic setting, an isomorphism class is presented by its representative (a (di)graph from the class). Alternatively, we can think of the probability distribution on all Cayley (di)graphs $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ in which each X appears with probability $1/(\ell_n\, s(X))$, where ℓ_n is the total number of n-vertex unlabeled circulants and $s(X)$ is the number of connection sets S such that $\mathrm{Cay}(\mathbb{Z}_n, S)$ is isomorphic to X.

- A labeled circulant, i.e., an arbitrary (di)graph on the vertex set $\{0, 1, \ldots, n-1\}$ isomorphic to some Cayley (di)graph $X = \mathrm{Cay}(\mathbb{Z}_n, S)$. The uniform distribution is considered on all n-vertex (di)graphs in this class.

In each of the three cases, we use the same canonization algorithm presented in Subsection 3.2. For digraphs, the algorithm is extremely simple: We just individualize one vertex in an input digraph X and run CR on the obtained vertex-colored graph X_0. In this way either we get an individual label for each vertex of X or the algorithm gives up. The labeling is canonical for all circulants X, and it is successfully produced whenever X_0 is walk-discrete. For graphs, the algorithm is a little bit more complicated and is discussed in detail in Subsection 3.2. It succeeds whenever X_0 is walk-saturated.

Lemmas 2.1 and 2.2 provide us with two sufficient spectral conditions: X_0 is walk-discrete whenever X has simple spectrum, and X_0 is walk-saturated whenever X has saturated spectrum. This reduces our task to estimating the probability that the random digraph X has simple spectrum and, respectively, that the random graph X has saturated spectrum. In the case of Cayley (di)graphs, the proof is completed by applying Theorem 2.3.

It remains to show that the estimate of Theorem 2.3 stays as well true for the uniformly distributed labeled and unlabeled circulants. To this end, we present a general way to convert an estimate for one distribution into an estimate for another distribution with a small overhead cost.

6.1 Formal framework

We introduce the following notation simultaneously for graphs and digraphs. Let $\mathcal{C}$ be the set of all Cayley (di)graphs X of cyclic groups, that is, all (di)graphs $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ for any n and S. Recall that the notion of a Cayley (di)graph was formally defined in Section 3.2. According to this definition, two graphs $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ and $Y = \mathrm{Cay}(\mathbb{Z}_n, T)$ in $\mathcal{C}$ are different exactly when $S \neq T$.

In general, speaking of a (di)graph X, we always suppose that the vertex set of X is $\{0, 1, \ldots, n-1\}$, where n is the order of X. For a set of (di)graphs $\mathcal{Q}$, by $\mathcal{Q}_n$ we denote the set of the (di)graphs in $\mathcal{Q}$ that have order n. By $\mathcal{Q}^{\mathfrak{u}}$ we denote the unlabeled version of $\mathcal{Q}$ where isomorphic graphs are not distinguished. Formally, $\mathcal{Q}^{\mathfrak{u}}$ is the quotient set of $\mathcal{Q}$ by the isomorphism relation. In other words, $\mathcal{Q}^{\mathfrak{u}}$ consists of all unlabeled graphs whose representatives appear in $\mathcal{Q}$. Furthermore, $\mathcal{Q}^{\mathfrak{l}}$ is defined to be the closure of $\mathcal{Q}$ under isomorphism, that is, if $\mathcal{Q}$ contains a (di)graph X of order n, then $\mathcal{Q}^{\mathfrak{l}}$ contains all graphs on the vertex set $\{0, 1, \ldots, n-1\}$ isomorphic to

X. We write $\mathcal{Q}^{\mathfrak{u}}_n$ and $\mathcal{Q}^{\mathfrak{l}}_n$ to denote the restrictions of $\mathcal{Q}^{\mathfrak{u}}$ and $\mathcal{Q}^{\mathfrak{l}}$ to the (di)graphs of order n.

Note that $\mathcal{C}^{\mathfrak{u}}$ is precisely the set of unlabeled circulants, and $\mathcal{C}^{\mathfrak{l}}$ is the set of labeled circulants.

For an arbitrary set of (di)graphs $\mathcal{Q}$, let $a(\mathcal{Q}_n)$ denote the minimum number of automorphisms of a (di)graph in $\mathcal{Q}_n$. Note that

$$|\mathcal{Q}^{\mathfrak{l}}_n| \le |\mathcal{Q}^{\mathfrak{u}}_n|\, n!/a(\mathcal{Q}_n). \tag{26}$$

The following important fact is a consequence of the main result in [35]; see [35, Theorem 1.1] and the discussion right after its statement.

Proposition 6.1 (Muzychuk [35])**.** *For every $S \subseteq \mathbb{Z}_n \setminus \{0\}$ there are at most $\phi(n)$ sets $S' \subseteq \mathbb{Z}_n \setminus \{0\}$ such that* $\mathrm{Cay}(\mathbb{Z}_n, S') \cong \mathrm{Cay}(\mathbb{Z}_n, S)$.

Proposition 6.1 readily implies that if $\mathcal{Q} \subseteq \mathcal{C}$, then

$$|\mathcal{Q}^{\mathfrak{u}}_n| \le |\mathcal{Q}_n| \le \phi(n)|\mathcal{Q}^{\mathfrak{u}}_n|. \tag{27}$$

Seeing $\mathbb{Z}_n$ as a ring with addition and multiplication modulo n, we write $\mathbb{Z}_n^\times$ to denote the multiplicative group of order $\phi(n)$ consisting of the invertible elements of $\mathbb{Z}_n$. For a set $S \subseteq \mathbb{Z}_n \setminus \{0\}$, we define the subgroup $K(S) \le \mathbb{Z}_n^\times$ by $K(S) = \{k \in \mathbb{Z}_n^\times : kS = S\}$. In the case of digraphs, we call a connection set S *multiplier-free* if $K(S) = \{1\}$. In the case of graphs, an inverse-closed connection set $S = -S$ is called *multiplier-free* if $K(S) = \{1, -1\}$. The set of Cayley (di)graphs with multiplier-free connection sets is denoted by $\mathcal{A}$.

We say that a set $\mathcal{Q} \subseteq \mathcal{C}$ is *isomorphism-invariant within* $\mathcal{C}$ if for every $X \in \mathcal{Q}$ and $Y \in \mathcal{C}$, we have $Y \in \mathcal{Q}$ whenever $X \cong Y$.

Lemma 6.2. *The set $\mathcal{A}$ is isomorphism-invariant within $\mathcal{C}$.*

Proof. Assume that

$$\mathrm{Cay}(\mathbb{Z}_n, S) \cong \mathrm{Cay}(\mathbb{Z}_n, T). \tag{28}$$

Let $kS = S$ for $k \in \mathbb{Z}_n^\times$. To obtain the lemma, it is enough to prove that $kT = T$ as well.

We use the following fact, which was conjectured by Zibin and proved by Muzychuk, Klin, and Pöschel [36]. Let $d \mid n$, i.e., d is a divisor of n. For $S \subseteq \mathbb{Z}_n$, define $(S)_d = \{s \in S : \gcd(s, n) = d\}$. Then, according to [36, Theorem 5.1], the isomorphism (28) implies that for every $d \mid n$ there exists $m_d \in \mathbb{Z}_n^\times$ such that $(T)_d = m_d(S)_d$.

For $k \in \mathbb{Z}_n^\times$, let μ_k be the permutation of $\mathbb{Z}_n$ defined by $\mu_k(x) = kx$. The equality $kS = S$ means that S is a union of orbits of μ_k, that is, there is a set $H \subset \mathbb{Z}_n$ such that

$$S = \bigcup_{h\in H} \langle k\rangle h = \bigcup_{d|n} \bigcup_{h\in(H)_d} \langle k\rangle h,$$

where $\langle k\rangle$ denotes the subgroup of $\mathbb{Z}_n^\times$ generated by k. Note that $(S)_d = \bigcup_{h\in(H)_d} \langle k\rangle h$. It follows that

$$T = \bigcup_{d|n} (T)_d = \bigcup_{d|n} m_d(S)_d = \bigcup_{d|n} \bigcup_{h\in(H)_d} \langle k\rangle (hm_d).$$

This shows that T is also a union of orbits of μ_k and, as required, $kT = T$. □

If $\mathcal{Q} \subseteq \mathcal{A}$ is isomorphism-invariant within $\mathcal{C}$, then the lower bound in (27) can be improved as follows:

$$\begin{aligned} |\mathcal{Q}_n| &= \phi(n)|\mathcal{Q}_n^{\mathrm{u}}| \quad \text{for digraphs,} && (29)\\ \frac{\phi(n)}{2}|\mathcal{Q}_n^{\mathrm{u}}| \le |\mathcal{Q}_n| &\le \phi(n)|\mathcal{Q}_n^{\mathrm{u}}| \quad \text{for graphs.} && (30)\end{aligned}$$

Indeed, if $\mathrm{Cay}(\mathbb{Z}_n, S) \in \mathcal{A}$ and $k, k' \in \mathbb{Z}_n^\times$, then the isomorphic copies $\mathrm{Cay}(\mathbb{Z}_n, kS)$ and $\mathrm{Cay}(\mathbb{Z}_n, k'S)$ of this graph are distinct (i.e., $kS \neq k'S$) whenever $k' \neq k$ in the case of digraphs and $k' \neq \pm k$ in the case of graphs.

Lemma 6.3.

1 *A random connection set $S \subseteq \mathbb{Z}_n \setminus \{0\}$ is not multiplier-free with probability less than $n\,2^{-n/4}$.*

2 *A random inverse-closed connection set $S = -S$ is not multiplier-free with probability less than $2n\,2^{-n/8}$.*

Proof. Define the *annihilator* of $a \in \mathbb{Z}_n$ by $\mathrm{Ann}(a) = \{x \in \mathbb{Z}_n : xa = 0\}$. Since $\mathrm{Ann}(a)$ is a subgroup of $\mathbb{Z}_n$, we have $|\mathrm{Ann}(a)| \le n/2$ for every $a \neq 0$.

1. It suffices to prove that, for each $a \neq 1$ in $\mathbb{Z}_n^\times$, the equality $aS = S$ is fulfilled with probability at most $2^{-n/4}$. Let μ_a be the permutation of $\mathbb{Z}_n$ defined by $\mu_a(x) = ax$. We have $\mu_a(x) = x$ exactly when $x \in \mathrm{Ann}(a-1)$. Thus, μ_a is the identity on $\mathrm{Ann}(a-1)$ and a fixed-point-free permutation of the set $\mathbb{Z}_n \setminus \mathrm{Ann}(a-1)$. Denote the restriction of μ_a to the latter set by μ'_a. Let $c_1, \ldots, c_t$ be the cycle type of μ'_a. Note that $\sum_{i=1}^t c_i = n - |\mathrm{Ann}(a-1)| \ge n/2$. Note also that $t \le (\sum_{i=1}^t c_i)/2$ because $c_i \ge 2$ for all i. The equality $aS = S$ is true if and only if every cycle of μ'_a either is entirely in S or is disjoint from S. This happens with probability

$$\prod_{i=1}^t 2^{-c_i+1} = 2^{-(\sum_{i=1}^t c_i)+t} \le 2^{-(\sum_{i=1}^t c_i)/2} \le 2^{-n/4}.$$

2. Again, it is enough to prove that, for each $a \neq \pm 1$ in $\mathbb{Z}_n^\times$, the equality $aS = S$ is fulfilled with probability at most $2^{-n/8+2}$. Let Z be the set of all pairs $\{x, -x\}$ for $x \in \mathbb{Z}_n$ such that $x \neq -x$. Note that $|Z| = \lfloor (n-1)/2 \rfloor$. The permutation μ_a naturally acts on Z by $\mu_a(\{x, -x\}) = \{ax, -ax\}$. Let us estimate the number of fixed points under this action. We have $\mu_a(\{x, -x\}) = \{x, -x\}$ if and only if $ax = x$ or $ax = -x$, which happens exactly when $x \in \mathrm{Ann}(a-1) \cup \mathrm{Ann}(a+1)$. Note that

$$|\mathrm{Ann}(a-1) \cup \mathrm{Ann}(a+1)| \le \frac{n}{2} + 1. \tag{31}$$

Indeed, let $b = \gcd(a-1, n)$ and $c = \gcd(a+1, n)$, and note that $|\mathrm{Ann}(a-1)| = b$ and $|\mathrm{Ann}(a+1)| = c$. Since the sets $\mathrm{Ann}(a-1)$ and $\mathrm{Ann}(a+1)$ share at least one element, namely 0, it is enough to prove that $b + c \le n/2 + 2$. Since neither $a - 1 = 0$ nor $a + 1 = 0$, neither b nor c exceeds $n/2$. Hence, we are immediately done if $b \le 2$ or $c \le 2$. Suppose, therefore, that both $b \ge 3$ and $c \ge 3$.

Let $d = \gcd(b, c)$. Since d divides both $a - 1$ and $a + 1$, we have $d \le 2$. If $d = 1$, then b and c are coprime divisors of n and, therefore, $b + c \le b + n/b$.

Note that $3 \le b \le n/c \le n/3$. In particular, $n \ge 9$ in this case. It follows that $b + c \le n/3 + 3 < n/2 + 2$.

If $d = 2$, then $b/2$ and $c/2$ are coprime divisors of $n/2$. The similar argument yields

$$\frac{b}{2} + \frac{c}{2} \le \frac{b}{2} + \frac{n/2}{b/2} \le \frac{n}{3} + \frac{3}{2} < \frac{n}{2} + 2,$$

completing the proof of Bound (31).

Since x belongs to $\mathrm{Ann}(a-1) \cup \mathrm{Ann}(a+1)$ simultaneously with $-x$, the number of fixed points under the action of μ_a on Z is at most $n/4$. It follows that at least $n/4 - 1$ pairs in Z are non-fixed. Similarly to Part 1, we conclude that $aS = S$ with probability at most $2^{-n/8+1}$. □

To conclude the notational stuff, let $\mathcal{Q} \subseteq \mathcal{C}$. Then $\mathbb{P}[\mathcal{Q}_n] = |\mathcal{Q}_n|/|\mathcal{C}_n|$ is the probability that a random Cayley (di)graph $\mathrm{Cay}(\mathbb{Z}_n, S)$, where S is chosen equiprobably among all possible connection sets $S \subseteq \mathbb{Z}_n \setminus \{0\}$, belongs to $\mathcal{Q}$. Exactly this random Cayley (di)graph model is considered in Theorem 2.3 and studied in Section 5. In what follows, we also write $\mathbb{P}_{\mathfrak{l}}[\mathcal{Q}^{\mathfrak{l}}_n]$ to denote the probability that a random labeled circulant of order n, i.e., a (di)graph chosen randomly and uniformly in $\mathcal{C}^{\mathfrak{l}}_n$, belongs to $\mathcal{Q}^{\mathfrak{l}}_n$. Thus, $\mathbb{P}_{\mathfrak{l}}[\mathcal{Q}^{\mathfrak{l}}_n] = |\mathcal{Q}^{\mathfrak{l}}_n|/|\mathcal{C}^{\mathfrak{l}}_n|$. Similarly, $\mathbb{P}_{\mathfrak{u}}[\mathcal{Q}^{\mathfrak{u}}_n] = |\mathcal{Q}^{\mathfrak{u}}_n|/|\mathcal{C}^{\mathfrak{u}}_n|$ is the probability that a random unlabeled circulant of order n belongs to $\mathcal{Q}^{\mathfrak{u}}_n$.

6.2 From Cayley (di)graphs to unlabeled circulants

Lemma 6.4. *Let $\mathcal{R} \subseteq \mathcal{C}$. If $\mathcal{R}$ is isomorphism-invariant within $\mathcal{C}$, then*

$$\begin{array}{lcll} \mathbb{P}_{\mathfrak{u}}[\mathcal{R}^{\mathfrak{u}}_n] & \le & (1 + o(1))\, \mathbb{P}[\mathcal{R}_n] + n^2\, 2^{-n/4} & \textit{for digraphs,} \\ \mathbb{P}_{\mathfrak{u}}[\mathcal{R}^{\mathfrak{u}}_n] & \le & (2 + o(1))\, \mathbb{P}[\mathcal{R}_n] + 2n^2\, 2^{-n/8} & \textit{for graphs.} \end{array}$$

Proof. We prove the inequality for graphs; the case of digraphs is similar. Note that $\mathcal{C}^{\mathfrak{u}}_n \setminus \mathcal{A}^{\mathfrak{u}}_n = (\mathcal{C}_n \setminus \mathcal{A}_n)^{\mathfrak{u}}$ due to Lemma 6.2. Applying the inequalities (27) to $\mathcal{Q} = \mathcal{C} \setminus \mathcal{A}$ and to $\mathcal{Q} = \mathcal{C}$, we derive from Lemma 6.3

$$\mathbb{P}_{\mathfrak{u}}[\mathcal{C}^{\mathfrak{u}}_n \setminus \mathcal{A}^{\mathfrak{u}}_n] = \frac{|\mathcal{C}^{\mathfrak{u}}_n \setminus \mathcal{A}^{\mathfrak{u}}_n|}{|\mathcal{C}^{\mathfrak{u}}_n|} = \frac{|(\mathcal{C}_n \setminus \mathcal{A}_n)^{\mathfrak{u}}|}{|\mathcal{C}^{\mathfrak{u}}_n|} \le \frac{|\mathcal{C}_n \setminus \mathcal{A}_n|}{|\mathcal{C}_n|/\phi(n)} = \phi(n)\, \mathbb{P}[\mathcal{C}_n \setminus \mathcal{A}_n] \le 2n^2\, 2^{-n/8}. \tag{32}$$

Note that $\mathcal{R}^{\mathfrak{u}}_n \cap \mathcal{A}^{\mathfrak{u}}_n = (\mathcal{R}_n \cap \mathcal{A}_n)^{\mathfrak{u}}$ because $\mathcal{A}$ is isomorphism-invariant within $\mathcal{C}$ by Lemma 6.2 and $\mathcal{R}$ is isomorphism-invariant within $\mathcal{C}$ by assumption. Using Bound (32) and applying the inequalities (30) to $\mathcal{Q} = \mathcal{R} \cap \mathcal{A}$ and to $\mathcal{Q} = \mathcal{A}$, we obtain

$$\mathbb{P}_{\mathfrak{u}}[\mathcal{R}^{\mathfrak{u}}_n] \le \mathbb{P}_{\mathfrak{u}}[\mathcal{R}^{\mathfrak{u}}_n \cap \mathcal{A}^{\mathfrak{u}}_n] + \mathbb{P}_{\mathfrak{u}}[\mathcal{C}^{\mathfrak{u}}_n \setminus \mathcal{A}^{\mathfrak{u}}_n] \le \frac{\mathbb{P}_{\mathfrak{u}}[\mathcal{R}^{\mathfrak{u}}_n \cap \mathcal{A}^{\mathfrak{u}}_n]}{\mathbb{P}_{\mathfrak{u}}[\mathcal{A}^{\mathfrak{u}}_n]} + 2n^2\, 2^{-n/8} = \frac{|\mathcal{R}^{\mathfrak{u}}_n \cap \mathcal{A}^{\mathfrak{u}}_n|}{|\mathcal{A}^{\mathfrak{u}}_n|} + 2n^2\, 2^{-n/8}$$

$$= \frac{|(\mathcal{R}_n \cap \mathcal{A}_n)^{\mathfrak{u}}|}{|\mathcal{A}^{\mathfrak{u}}_n|} + 2n^2\, 2^{-n/8} \le \frac{2\, |\mathcal{R}_n \cap \mathcal{A}_n|/\phi(n)}{|\mathcal{A}_n|/\phi(n)} + 2n^2\, 2^{-n/8} \le \frac{2\, |\mathcal{R}_n|}{|\mathcal{A}_n|} + 2n^2\, 2^{-n/8}$$

$$= \frac{2\, \mathbb{P}[\mathcal{R}_n]}{\mathbb{P}[\mathcal{A}_n]} + 2n^2\, 2^{-n/8} = (2 + o(1))\, \mathbb{P}[\mathcal{R}_n] + 2n^2\, 2^{-n/8}.$$

For the last equality we used Lemma 6.3 once again. □

Let $\mathcal{R}$ be the set of all Cayley digraphs $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ whose spectrum is *not* simple. In the case of graphs, we set $\mathcal{R}$ to be the set of all $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ whose spectrum is *not* saturated. Lemma 6.4 implies that the bound of Theorem 2.3 holds true also for unlabeled circulants. This completes the proof of Theorem 1.1 in the unlabeled case.

6.3 From unlabeled to labeled circulants

For $a \in \mathbb{Z}_n$, define a bijection $\sigma_a : \mathbb{Z}_n \to \mathbb{Z}_n$ by $\sigma_a(x) = x + a$. For every Cayley digraph $X = \mathrm{Cay}(\mathbb{Z}_n, S)$, the map σ_a is an automorphism of X for all a. If X has no other automorphism, that is, $\mathrm{Aut}(X)$ is as small as possible (see [16, Section 8.1]), then we call X *firm*. In other words, X is firm exactly when $\mathrm{Aut}(X) \cong \mathbb{Z}_n$.

Now, let X be a Cayley graph. In this case there is also another automorphism ρ defined by $\rho(x) = -x$. The automorphisms σ_1 and ρ generate a subgroup of $\mathrm{Aut}(X)$ isomorphic to the dihedral group D_{2n}. If X has no other automorphisms, i.e., $\mathrm{Aut}(X) \cong D_{2n}$, then we say that X is a *firm Cayley graph* of $\mathbb{Z}_n$. Note that a firm graph is not firm as a digraph; this should not make any confusion because we treat random graphs and random digraphs separately (even when in parallel).

The set of firm (di)graphs is denoted by $\mathcal{F}$. The following equalities easily follow from the definitions: If $\mathcal{Q} \subseteq \mathcal{F}$, then

$$|\mathcal{Q}_n^{\mathfrak{l}}| = (n-1)!\,|\mathcal{Q}_n^{\mathfrak{u}}| \quad \text{for digraphs,} \tag{33}$$

$$|\mathcal{Q}_n^{\mathfrak{l}}| = \frac{(n-1)!}{2}\,|\mathcal{Q}_n^{\mathfrak{u}}| \quad \text{for graphs.} \tag{34}$$

We will need the following estimates obtained in [8, 15].[5]

Proposition 6.5.

1. $\mathbb{P}[\mathcal{C}_n \setminus \mathcal{F}_n] = 2^{-n/4+o(n)}$ *for digraphs* (Dobson, Spiga, and Verret [15, Theorem 1.6]).
2. $\mathbb{P}[\mathcal{C}_n \setminus \mathcal{F}_n] = O(n2^{-n/8})$ *for graphs* (Bhoumik, Dobson, and Morris [8, Theorem 3.2]).

Note that $\mathcal{C}_n^{\mathfrak{u}} \setminus \mathcal{F}_n^{\mathfrak{u}} = (\mathcal{C}_n \setminus \mathcal{F}_n)^{\mathfrak{u}}$ because $\mathcal{F}$ is obviously isomorphism-invariant within $\mathcal{C}$. Applying the inequalities (27) to $\mathcal{Q} = \mathcal{C} \setminus \mathcal{F}$ and to $\mathcal{Q} = \mathcal{C}$, similarly to (32) we get the relation

$$\mathbb{P}_{\mathfrak{u}}[\mathcal{C}_n^{\mathfrak{u}} \setminus \mathcal{F}_n^{\mathfrak{u}}] \le \phi(n)\,\mathbb{P}[\mathcal{C}_n \setminus \mathcal{F}_n].$$

By Proposition 6.5, this implies that

$$\mathbb{P}_{\mathfrak{u}}[\mathcal{C}_n^{\mathfrak{u}} \setminus \mathcal{F}_n^{\mathfrak{u}}] = 2^{-n/4+o(n)} \quad \text{for digraphs,} \tag{35}$$

$$\mathbb{P}_{\mathfrak{u}}[\mathcal{C}_n^{\mathfrak{u}} \setminus \mathcal{F}_n^{\mathfrak{u}}] = O(n^2\,2^{-n/8}) \quad \text{for graphs.} \tag{36}$$

[5] Note that, since $\mathcal{F} \subseteq \mathcal{A}$, we could use Proposition 6.5 instead of Lemma 6.3 in Subsection 6.2. However, Lemma 6.3 has the advantage of being proved by an elementary method.

Lemma 6.6. *Let $\mathcal{S} \subseteq \mathcal{C}$. If $\mathcal{S}$ is isomorphism-invariant within $\mathcal{C}$, then*

$$\begin{aligned} \mathbb{P}_\mathfrak{l}[\mathcal{S}_n^\mathfrak{l}] &\geq \mathbb{P}_\mathfrak{u}[\mathcal{S}_n^\mathfrak{u}] - 2^{-n/4+o(n)} \quad \textit{for digraphs,} \\ \mathbb{P}_\mathfrak{l}[\mathcal{S}_n^\mathfrak{l}] &\geq \mathbb{P}_\mathfrak{u}[\mathcal{S}_n^\mathfrak{u}] - O(n^2\, 2^{-n/8}) \quad \textit{for graphs.} \end{aligned}$$

Proof. We prove the inequality for graphs; the case of digraphs is similar. Note that $\mathcal{S}_n^\mathfrak{u} \cap \mathcal{F}_n^\mathfrak{u} = (\mathcal{S}_n \cap \mathcal{F}_n)^\mathfrak{u}$ and $\mathcal{S}_n^\mathfrak{l} \cap \mathcal{F}_n^\mathfrak{l} = (\mathcal{S}_n \cap \mathcal{F}_n)^\mathfrak{l}$ because both $\mathcal{S}$ and $\mathcal{F}$ are isomorphism-invariant within $\mathcal{C}$. Applying the inequality (26) to $Q = \mathcal{C}$ and the equality (34) to $\mathcal{Q} = \mathcal{S} \cap \mathcal{F}$, we derive

$$\begin{aligned} \mathbb{P}_\mathfrak{l}[\mathcal{S}_n^\mathfrak{l}] \geq \mathbb{P}_\mathfrak{l}[\mathcal{S}_n^\mathfrak{l} \cap \mathcal{F}_n^\mathfrak{l}] = \frac{|\mathcal{S}_n^\mathfrak{l} \cap \mathcal{F}_n^\mathfrak{l}|}{|\mathcal{C}_n^\mathfrak{l}|} &= \frac{|(\mathcal{S}_n \cap \mathcal{F}_n)^\mathfrak{l}|}{|\mathcal{C}_n^\mathfrak{l}|} \\ &\geq \frac{(n-1)!\,|(\mathcal{S}_n \cap \mathcal{F}_n)^\mathfrak{u}|/2}{(n-1)!\,|\mathcal{C}_n^\mathfrak{u}|/2} = \frac{|\mathcal{S}_n^\mathfrak{u} \cap \mathcal{F}_n^\mathfrak{u}|}{|\mathcal{C}_n^\mathfrak{u}|} = \mathbb{P}_\mathfrak{u}[\mathcal{S}_n^\mathfrak{u} \cap \mathcal{F}_n^\mathfrak{u}]. \quad (37) \end{aligned}$$

This implies that

$$\mathbb{P}_\mathfrak{l}[\mathcal{S}_n^\mathfrak{l}] \geq \mathbb{P}_\mathfrak{u}[\mathcal{S}_n^\mathfrak{u}] - \mathbb{P}_\mathfrak{u}[\mathcal{C}_n^\mathfrak{u} \setminus \mathcal{F}_n^\mathfrak{u}] \geq \mathbb{P}_\mathfrak{u}[\mathcal{S}_n^\mathfrak{u}] - O(n^2\, 2^{-n/8}).$$

The last inequality follows from the bound (36). □

Now, let $\mathcal{S}$ be the set of all Cayley digraphs $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ with simple spectrum. In the case of graphs, we set $\mathcal{S}$ to be the set of all $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ with saturated spectrum. We already know that the bound of Theorem 2.3 holds true for unlabeled circulants. Lemma 6.6 implies that it is as well true for labeled circulants.

Lemma 6.6 provides a rather general way of showing that if a property holds for almost all unlabeled circulants, then it also holds for almost all labeled circulants. We remark that that for the property $\mathcal{S}$ that a circulant digraph has simple spectrum (or that a circulant graph has saturated spectrum) this can be alternatively derived from Lemmas 2.1, 2.2, and 3.2. Indeed, these lemmas imply that $\mathcal{S} \subseteq \mathcal{F}$, which allows us to obtain the inequality $\mathbb{P}_\mathfrak{l}[\mathcal{S}_n^\mathfrak{l}] \geq \mathbb{P}_\mathfrak{u}[\mathcal{S}_n^\mathfrak{u}]$ directly from (37).

The proof of Theorem 1.1 is complete.

7 Canonical Cayley representations via 2-WL

This section is devoted to the proof of Theorem 1.2. Before presenting the proof in Subsection 7.4, we provide a formal description of the 2-WL algorithm in Subsection 7.1, and collect relevant preliminary results in Subsections 7.2 and 7.3.

7.1 Description of 2-WL

For notational simplicity, an ordered pair (a, b) will be denoted by ab. Given a loopless digraph $X = (V, E)$ as an input, 2-WL iteratively computes a sequence of colorings c_X^i of the Cartesian square $V \times V$. The initial coloring is defined by $c_X^0(uv) = (\mathit{type}(uv), \mathit{type}(vu))$ where $\mathit{type}(uv)$ takes on one of three values according to the type of an ordered vertex pair uv, namely *arc* if $uv \in E$, *nonarc* if $uv \notin E$ and $u \neq v$, and *loop* if $u = v$. The coloring is modified iteratively as follows:

$$c_X^{i+1}(uv) = \{\!\!\{ \left(c_X^i(uw), c_X^i(wv) \right) \}\!\!\}_{w \in V}.$$

In words, the new color of a pair uv is a "superposition" of all old color pairs observable along the extensions of uv to a triple uwv. Denote the partition of $V \times V$ into the color classes of c_X^i by $\mathcal{C}_X^i$. A simple inductive argument shows that

$$c_X^{i+1}(uv) \neq c_X^{i+1}(u'v') \text{ whenever } c_X^i(uv) \neq c_X^i(u'v')$$

which means that $\mathcal{C}_X^{i+1}$ is finer than or equal to $\mathcal{C}_X^i$. It follows that the partition stabilizes starting from some step $t = t(X)$, that is, $\mathcal{C}_X^t = \mathcal{C}_X^{t-1}$, which implies that $\mathcal{C}_X^i = \mathcal{C}_X^{t-1}$ for all $i \geq t$. Note that $t \leq |V|^2$. As soon as the stabilization is reached, the algorithm terminates and outputs the coloring c_X^t.

An easy induction on i shows that, if ϕ is an isomorphism from X to Y, then

$$c_X^i(uv) = c_Y^i(\phi(u)\phi(v)). \tag{38}$$

Note that the length of c_X^i-colors (in any natural encoding) grows exponentially with i increasing. Similarly to CR, the exponential blow-up is remedied by renaming the colors after each step. Finally, note that 2-WL can be implemented in time $O(n^3 \log n)$; see [26, 27].

Let $\mathrm{WL}_2(X) = \mathcal{C}_X^t$ denote the stabilized partition of V^2. It can be noticed that $\mathrm{WL}_2(X)$ forms a *coherent configuration* [13], but we will not use this fact directly. However, we will need another source of coherent configurations, which we introduce in the next subsection.

7.2 Orbitals of a permutation group

Let $G \leq \mathrm{Sym}(V)$ be a group of permutations of a set V. A natural action of G on V^2 is defined by $\alpha(x, y) = (\alpha(x), \alpha(y))$ for $x, y \in V$ and $\alpha \in G$. An orbit of this action is called *orbital*. We denote the partition of V^2 into the orbitals of G by $\mathrm{Orb}_2(G)$.

Equality (38) readily implies that the partition $\mathrm{WL}_2(X)$ is refined by the partition $\mathrm{Orb}_2(\mathrm{Aut}(X))$. Similarly to the concept of a Schurian coherent configuration [13], we call a digraph X *Schurian* if $\mathrm{WL}_2(X) = \mathrm{Orb}_2(\mathrm{Aut}(X))$.

For each $a \in \mathbb{Z}_n$, a digraph $X = \mathrm{Cay}(\mathbb{Z}_n, S)$ has an automorphism σ_a defined by $\sigma_a(x) = x + a$. These automorphisms form a subgroup of $\mathrm{Aut}(X)$. If this subgroup is normal, the circulant X is called *normal* (see [16, Chapter 8.1], where this concept is discussed in the more general setting of Cayley graphs for any groups). Note that if X is firm—according to the definition given in Section 6.3, then X is normal. This is true both for digraphs and graphs (in the case of graphs, recall that every subgroup of index 2 is normal).

Proposition 7.1 ([21, Theorem 6.1]). *Every normal circulant digraph is Schurian.*

Extending the definition of a firm Cayley (di)graph in Section 6.3, we also call an arbitrary labeled circulant (di)graph *firm* if it is isomorphic to a firm Cayley (di)graph $\mathrm{Cay}(\mathbb{Z}_n, S)$ or, equivalently, if $\mathrm{Aut}(X)$ is as small as possible, i.e., $\mathrm{Aut}(X)$ isomorphic to the cyclic group in the directed case and to the dihedral group in the undirected case.

The Schurity property of a digraph X is beneficial because it enables an efficient computation of the partition $\mathrm{Orb}_2(\mathrm{Aut}(X))$ just by running 2-WL on X. Moreover, if X is a firm circulant (di)graph, then the knowledge of $\mathrm{Orb}_2(\mathrm{Aut}(X))$ allows us to determine the automorphism group $\mathrm{Aut}(X)$ as a permutation group. We now state this fact formally in the form that will be useful in the next subsection.

Let C_n denote a cyclic permutation group on the n-element set V that acts on V transitively or, equivalently, contains a cycle of length n. Assume that $n \geq 3$ and note that C_n has a unique extention to a dihedral group (of permutations of V). We denote this dihedral permutation group by D_{2n}. Note that all $\phi(n)$ elements of order n in C_n are cycles of length n. The same holds true for D_{2n}. Indeed, every element in $D_{2n} \setminus C_n$ has degree 2. We now observe that, given the partitions $\mathrm{Orb}_2(C_n)$ and $\mathrm{Orb}_2(D_{2n})$, the elements of order n in C_n and D_{2n} can be explicitly constructed as permutations of the set V.

Lemma 7.2.

1. *The $\phi(n)$ generators of C_n are uniquely determined by* $\mathrm{Orb}_2(C_n)$ *and can be constructed from* $\mathrm{Orb}_2(C_n)$ *in time $O(n^2)$.*

2. *The $\phi(n)$ elements of order n of D_{2n} are uniquely determined by* $\mathrm{Orb}_2(D_{2n})$ *and can be constructed from* $\mathrm{Orb}_2(D_{2n})$ *in time $O(n^2)$.*

Proof. 1. Assume for a while that $V = \mathbb{Z}_n$ and that C_n consists of the permutations σ_a, $a \in \mathbb{Z}_n$, defined (like above) by $\sigma_a(x) = x + a$. Note that σ_a is a generator of C_n if and only if $a \in \mathbb{Z}_n^\times$. Each orbital of C_n will be regarded as a digraph. As easily seen, two pairs (x, y) and (x', y') are in the same orbital exactly when $y - x = y' - x'$. It follows that the orbitals of C_n are exactly the digraphs $\mathrm{Cay}(\mathbb{Z}_n, \{a\})$ for $a \in \mathbb{Z}_n$. Note that $\mathrm{Cay}(\mathbb{Z}_n, \{a\})$ has $\gcd(a, n)$ connected components, and each of them is isomorphic to the directed cycle of length $n/\gcd(a, n)$. This implies that, in general, the generating elements of C_n can be identified by finding all $\phi(n)$ orbitals of C_n that, viewed as digraphs, are directed cycles of length n. For each such cycle, one then forms a cyclic permutation of V along the cycle.

2. Note that $\mathrm{Orb}_2(C_n)$ contains, along with each orbital $C = \mathrm{Cay}(\mathbb{Z}_n, \{a\})$, its transpose $C' = \mathrm{Cay}(\mathbb{Z}_n, \{-a\})$. As easily seen, the orbitals of D_{2n} are exactly the symmetric closures $C \cup C' = \mathrm{Cay}(\mathbb{Z}_n, \{-a, a\})$ of the orbitals of C_n. Therefore, the elements of order n of D_{2n} are identified by finding all $\phi(n)/2$ orbitals of D_{2n} that are (undirected) cycles of length n and by forming, along each of these cycles, two cyclic permutations of V in both directions. □

7.3 Cayley representations of firm circulants

Recall that a Cayley representation of a labeled circulant X on n vertices is a map $\lambda : V(X) \to \mathbb{Z}_n$ such that X^λ is a Cayley graph, i.e., $X^\lambda = \mathrm{Cay}(\mathbb{Z}_n, S)$ for some

S. We call two Cayley representation λ and λ' of X *equivalent* if $X^\lambda = X^{\lambda'}$. If $X^\lambda = \mathrm{Cay}(\mathbb{Z}_n, S)$, then λ and λ' are equivalent if and only if $\lambda' = \alpha\lambda$ for some automorphism α of $\mathrm{Cay}(\mathbb{Z}_n, S)$. Thus, if X is a firm digraph, then λ has exactly n equivalent Cayley representations, and if X is a firm graph, then λ has exactly $2n$ equivalent representations. It is useful to notice the following simple fact.

Lemma 7.3.

1. *Let X be a firm circulant digraph. There is a one-to-one correspondence between the equivalence classes of Cayley representations of X and the generators of* $\mathrm{Aut}(X)$.

2. *Let X be a firm circulant graph on n vertices. There is a one-to-one correspondence between the equivalence classes of Cayley representations of X and the pairs of mutually inverse cycles of length n in* $\mathrm{Aut}(X)$.

3. *In both cases, if a cycle of length n in* $\mathrm{Aut}(X)$ *is given, a Cayley representation of X from the corresponding equivalence class is constructible in linear time.*

Proof. Let λ be a Cayley representation of X. The cycle $(\lambda^{-1}(0), \lambda^{-1}(1), \ldots, \lambda^{-1}(n-1))$ is a generator of $\mathrm{Aut}(X)$. Every Cayley representation λ' equivalent to λ yields the same generator. Conversely, every cycle of length n in $\mathrm{Aut}(X)$ determines, in an explicit way, n equivalent Cayley representations of X.

The case of graphs is similar with the only difference that a cycle $(v_0, v_1, \ldots, v_{n-1})$ and its inverse $(v_{n-1}, \ldots, v_1, v_0)$ yield equivalent Cayley representations of X. □

Lemma 7.3 implies that every firm circulant digraph X on n vertices has, up to equivalence, exactly $\phi(n)$ Cayley representations. In the case that X is a graph, there are, up to equivalence, exactly $\phi(n)/2$ Cayley representations.

7.4 Proof of Theorem 1.2

We design an algorithm such that the following three conditions are fulfilled for a certain class of (di)graphs $\mathcal{C}$:

- the algorithm computes a canonical Cayley representation for all inputs in $\mathcal{C}$;
- the algorithm gives up on all inputs not in $\mathcal{C}$;
- $\mathcal{C}$ contains all firm circulant (di)graphs.

The last condition ensures the success probability bound stated in Theorem 1.2. Indeed, a random Cayley (di)graph $\mathrm{Cay}(\mathbb{Z}_n, S)$ is firm with high probability by Proposition 6.5, and this remains true for a random labeled circulant by the transition lemmas obtained in Section 6, i.e., by Lemmas 6.4 and 6.6.

Since firm circulant (di)graphs are normal, Proposition 7.1 shows that the orbital partition $\mathrm{Orb}_2(\mathrm{Aut}(X))$ for a firm circulant X can be computed just by running 2-WL on X. According to [26, 27], this takes time $O(n^3 \log n)$, where n is the number of vertices in X. Given $\mathrm{Orb}_2(\mathrm{Aut}(X))$, one can easily determine all cycles of length n in the permutation group $\mathrm{Aut}(X)$; see Lemma 7.2. By Lemma 7.3, these can be used to efficiently construct all Cayley representations of X. Note that it is enough to have one representation from each equivalence class. Summing up, we come to the following procedure.

Canonical Cayley representation algorithm

Input: a (di)graph X.

1. Run 2-WL on X.
2. Check whether the partition $\mathrm{WL}_2(X)$ contains a part that is isomorphic to a cycle (di)graph of length n. If not, terminate. Otherwise, let C be the part of $\mathrm{WL}_2(X)$ of this kind that has the lexicographically smallest 2-WL-color.
3. Choose an arbitrary vertex x_0 and enumerate the vertices of X along C starting from x_0 (in the case of graphs, choose any of the two directions in C). Let $x_0, x_1, \ldots, x_{n-1}$ be the obtained enumeration.
4. Check whether the cyclic permutation $(x_0 x_1 \ldots x_{n-1})$ is an automorphism of X. If not, then give up.
5. Output the labeling $\lambda_X : V(X) \to \mathbb{Z}_n$ where $\lambda_X(x_i) = i$.

The proof of Theorem 1.2 is complete.

Remark 7.4. Proposition 7.1 can also be used to show that 2-WL distinguishes a firm circulant (di)graph X from any other non-isomorphic (di)graph Y in the sense that the color palettes produced by 2-WL on X and Y are different, i.e., $\{\!\{c_X^{t(X)}(uv) \;:\; uv \in V(X)^2\}\!\} \neq \{\!\{c_Y^{t(Y)}(uv) \;:\; uv \in V(Y)^2\}\!\}$. The class of the firm circulant digraphs contains all circulant digraphs with simple spectrum [18, Theorem 3]. For this smaller class of circulant digraphs, the identifiability by 2-WL follows also from any of the results stated in [22, Corollary 4.5] or [20, Corollary of Theorem 1] (the latter results is stronger than the former in view of [13, Theorem 3.3.19]).

Competing interests: The authors declare none.

Data availability: This is a purely theoretical research, no new data were created or analyzed in the study.

Acknowledgements

The authors gratefully acknowledge the anonymous referee for their careful reading of the manuscript and numerous helpful comments, which significantly improved the presentation of the manuscript.

Oleg Verbitsky is supported by DFG grant KO 1053/8–2.

References

[1] T. M. Apostol. *Introduction to analytic number theory.* New York, NY: Springer, 1998.

[2] V. Arvind, J. Köbler, G. Rattan, and O. Verbitsky. Graph isomorphism, color refinement, and compactness. *Computational Complexity*, 26(3):627–685, 2017.

[3] L. Babai. Isomorphism problem for a class of point-symmetric structures. *Acta Math. Acad. Sci. Hungar.*, 29(3-4):329–336, 1977.

[4] L. Babai, X. Chen, X. Sun, S. Teng, and J. Wilmes. Faster canonical forms for strongly regular graphs. In *54th Annual IEEE Symposium on Foundations of Computer Science (FOCS'13)*, pages 157–166. IEEE Computer Society, 2013.

[5] L. Babai, P. Erdős, and S. M. Selkow. Random graph isomorphism. *SIAM Journal on Computing*, 9(3):628–635, 1980.

[6] E. Bach and J. Shallit. *Algorithmic number theory, Vol. 1: Efficient algorithms.* Cambridge, MA: The MIT Press, 1996.

[7] C. Berkholz, P. S. Bonsma, and M. Grohe. Tight lower and upper bounds for the complexity of canonical colour refinement. *Theory Comput. Syst.*, 60(4):581–614, 2017.

[8] S. Bhoumik, T. Dobson, and J. Morris. On the automorphism groups of almost all circulant graphs and digraphs. *Ars Math. Contemp.*, 7(2):499–518, 2014.

[9] B. Bollobás. Distinguishing vertices of random graphs. *Annals of Discrete Mathematics*, 13:33–49, 1982.

[10] A. Bose and K. Saha. *Random circulant matrices.* Boca Raton, FL: CRC Press, 2019.

[11] J. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identifications. *Combinatorica*, 12(4):389–410, 1992.

[12] A. Cardon and M. Crochemore. Partitioning a graph in $O(|A| \log_2 |V|)$. *Theor. Comput. Sci.*, 19:85–98, 1982.

[13] G. Chen and I. Ponomarenko. *Coherent configurations.* Wuhan: Central China Normal University Press, 2019. A draft version is available at `http://www.pdmi.ras.ru/~inp/ccNOTES.pdf`.

[14] P. J. Davis. *Circulant matrices.* New York, NY: AMS Chelsea Publishing, 2nd ed. edition, 1994.

[15] E. Dobson, P. Spiga, and G. Verret. Cayley graphs on abelian groups. *Combinatorica*, 36(4):371–393, 2016.

[16] T. Dobson, A. Malnič, and D. Marušič. *Symmetry in graphs*, volume 198 of *Camb. Stud. Adv. Math.* Cambridge: Cambridge University Press, 2022.

[17] Y. A. Drozd and V. V. Kirichenko. *Finite dimensional algebras.* Springer-Verlag, Berlin, 1994.

[18] B. Elspas and J. Turner. Graphs with circulant adjacency matrices. *Journal of Combinatorial Theory*, 9(3):297–307, 1970.

[19] S. Evdokimov and I. Ponomarenko. Circulant graphs: recognizing and isomorphism testing in polynomial time. *St. Petersbg. Math. J.*, 15(6):813–835, 2004.

[20] S. A. Evdokimov and I. N. Ponomarenko. Two inequalities for the parameters of a cellular algebra. *J. Math. Sci.*, 96(5):3496–3504, 1999.

[21] S. A. Evdokimov and I. N. Ponomarenko. Characterization of cyclotomic schemes and normal Schur rings over a cyclic group. *St. Petersbg. Math. J.*, 14(2):189–221, 2003.

[22] S. Friedland. Coherent algebras and the graph isomorphism problem. *Discret. Appl. Math.*, 25(1-2):73–98, 1989.

[23] F. Fuhlbrück, J. Köbler, I. Ponomarenko, and O. Verbitsky. The Weisfeiler-Leman algorithm and recognition of graph properties. *Theor. Comput. Sci.*, 895:96–114, 2021.

[24] C. Godsil. Controllable subsets in graphs. *Ann. Comb.*, 16(4):733–744, 2012.

[25] E. M. Hagos. Some results on graph spectra. *Linear Algebra Appl.*, 356(1-3):103–111, 2002.

[26] N. Immerman and E. Lander. Describing graphs: A first-order approach to graph canonization. In *Complexity Theory Retrospective*, pages 59–81. Springer, 1990.

[27] N. Immerman and R. Sengupta. The k-dimensional Weisfeiler-Leman algorithm. Technical report, `arxiv.org/abs/1907.09582`, 2019.

[28] S. Kiefer, I. Ponomarenko, and P. Schweitzer. The Weisfeiler-Leman dimension of planar graphs is at most 3. *J. ACM*, 66(6):44:1–44:31, 2019.

[29] L. Kluge. Combinatorial refinement on circulant graphs. *Computational Complexity*, 33(2):9, 2024.

[30] N. M. Kriege. Weisfeiler and Leman go walking: Random walk kernels revisited. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022 (NeurIPS'22)*, 2022.

[31] L. Kucera. Canonical labeling of regular graphs in linear average time. In *28th Annual Symposium on Foundations of Computer Science (FOCS'87)*, pages 271–279, 1987.

[32] F. Liu and J. Siemons. Unlocking the walk matrix of a graph. *J. Algebr. Comb.*, 55(3):663–690, 2022.

[33] B. D. McKay and A. Piperno. Practical graph isomorphism, ii. *Journal of Symbolic Computation*, 60:94–112, 2014.

[34] M. W. Meckes. Some results on random circulant matrices. In *High dimensional probability. V: The Luminy volume.*, pages 213–223. Beachwood, OH: IMS, Institute of Mathematical Statistics, 2009.

[35] M. Muzychuk. A solution of the isomorphism problem for circulant graphs. *Proc. Lond. Math. Soc. (3)*, 88(1):1–41, 2004.

[36] M. E. Muzychuk, M. H. Klin, and R. Pöschel. The isomorphism problem for circulant graphs via Schur ring theory. In *Codes and Association Schemes*, volume 56 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 241–264. DIMACS/AMS, 1999.

[37] S. O'Rourke and B. Touri. On a conjecture of Godsil concerning controllable random graphs. *SIAM J. Control. Optim.*, 54(6):3347–3378, 2016.

[38] I. Ponomarenko. On the WL-dimension of circulant graphs of prime power order. *Algebraic Combinatorics*, 6(6):1469–1490, 2023.

[39] D. L. Powers and M. M. Sulaiman. The walk partition and colorations of a graph. *Linear Algebra Appl.*, 48:145–159, 1982.

[40] G. Rattan and T. Seppelt. Weisfeiler-Leman and graph spectra. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms (SODA'23)*, pages 2268–2285. SIAM, 2023.

[41] H. Schreck and G. Tinhofer. A note on certain subpolytopes of the assignment polytope associated with circulant graphs. *Linear Algebra and its Applications*, 111:125–134, 1988.

[42] G. Tinhofer. Graph isomorphism and theorems of Birkhoff type. *Computing*, 36:285–300, 1986.

[43] G. Tinhofer. A note on compact graphs. *Discrete Applied Mathematics*, 30(2-3):253–264, 1991.

[44] O. Verbitsky and M. Zhukovskii. Canonization of a random graph by two matrix-vector multiplications. In *31st Annual European Symposium on Algorithms (ESA'23)*, volume 274 of *LIPIcs*, pages 100:1–100:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

[45] O. Verbitsky and M. Zhukovskii. Canonization of a random circulant graph by counting walks. In *WALCOM: Algorithms and Computation – 18th International Conference and Workshops on Algorithms and Computation (WALCOM'24)*, volume 14549 of *Lecture Notes in Computer Science*, pages 319–334. Springer, 2024.

[46] B. Weisfeiler and A. Leman. The reduction of a graph to canonical form and the algebra which appears therein. *NTI, Ser. 2*, 9:12–16, 1968. English translation is available at `https://www.iti.zcu.cz/wl2018/pdf/wl_paper_translation.pdf`.

[47] Y. Wu and I. Ponomarenko. On the Weisfeiler-Leman dimension of circulant graphs. Technical report, `arxiv.org/abs/2406.15822`, 2024.