Narrow sieves for parameterized paths and packings

Björklund, Andreas

2017-08

Björklund , A , Husfeldt , T , Kaski , P & Koivisto , M 2017 , ' Narrow sieves for parameterized paths and packings ' , Journal of Computer and System Sciences , vol. 87 , pp. 119-139 . https://doi.org/10.1016/j.jcss.2017.03.003

http://hdl.handle.net/10138/288208
https://doi.org/10.1016/j.jcss.2017.03.003

unspecified
acceptedVersion

Narrow sieves for parameterized paths and packings[☆]

Andreas Björklund[a], Thore Husfeldt[a,b], Petteri Kaski[c], Mikko Koivisto[d,*]

[a] *Lund University, Department of Computer Science, PO Box 118, SE-22100 Lund, Sweden*
[b] *IT University of Copenhagen, 2300 Copenhagen S, Denmark*
[c] *Helsinki Institute for Information Technology HIIT, Aalto University, Department of Information and Computer Science PO Box 15400, FI-00076 Aalto, Finland*
[d] *Helsinki Institute for Information Technology HIIT, Department of Computer Science, University of Helsinki, PO Box 68, FI-00014 University of Helsinki, Finland*

Abstract

We present parameterized algorithms for the k-path problem, the p-packing of q-sets problem, and the q-dimensional p-matching problem. Our algorithms solve these problems with high probability in time exponential only in the parameter (k, p, q) and using polynomial space. The constant bases of the exponentials are significantly smaller than in previous works; for example, for the k-path problem the improvement is from 2 to 1.66. We also show how to detect if a d-regular graph admits an edge coloring with d colors in time within a polynomial factor of $2^{(d-1)n/2}$. Our techniques generalize an algebraic approach studied in various recent works.

Keywords: determinant, edge coloring, graph algorithm, k-path, multidimensional matching, sieve, set packing, polynomial identity testing, randomized algorithm
2000 MSC: 68W20, 68W30, 05C38, 05C15, 05C85, 05C70

[☆]The work was supported by the Swedish Research Council, grant 2007-6595, and by the Academy of Finland, grants 117499 and 125637.

[*]Corresponding author; phone +358 9 191 51171, fax +358 9 191 51120.
Email addresses: andreas.bjorklund@yahoo.se (Andreas Björklund), thore@itu.dk (Thore Husfeldt), petteri.kaski@aalto.fi (Petteri Kaski), mikko.koivisto@cs.helsinki.fi (Mikko Koivisto)

1. Introduction

Combinatorial problems such as finding a long simple path in a graph or disjointly packing many members of a set of subsets are well-studied and hard. In fact, under standard complexity-theoretic assumptions, algorithms for these problems must either be inexact, or require running times of super-polynomial or maybe even exponential time.

It has been observed that the complexity of the problems studied in the present paper depends exponentially on the output size k instead of the input size n, i.e., they admit running times of the form $\exp(\text{poly}(k))\cdot\text{poly}(n)$ rather than, say, $\exp(O(n))$. A number of papers have improved the exponential dependencies dramatically over the past decade, arriving at exponential factors of size 2^k. We further improve this dependency, reducing the exponent base below the constant 2, sometimes significantly.

To express our results in the terminology of parameterized computational complexity theory, we improve the running time of several canonical fixed parameter tractable problems. In particular, we claim the ephemeral lead in the highly competitive "FPT races" for path finding, uniform set packing, and multidimensional matching.

Our techniques also allow us to report progress on an unparameterized problem: we show a nontrivial upper bound on the complexity of edge coloring for regular graphs.

We adopt the notational convention from parameterized and exponential time algorithms, letting $O^*(f(k))$ denote $f(k)n^{O(1)}$, where typically n is some aspect of the input size such as number of vertices, and k is a parameter such as path length. Our parameters are all polynomially bounded by n, so O^* also hides factors that are polynomial in the parameter size. We present our results in terms of decision problems, they can be turned into optimization or search problems by self-reductions in the obvious way.

All our algorithms are randomized. The error is one-sided in the sense that they never report a false positive. The error probability is constant and can be made exponentially small by a polynomial number of repetitions, which would again be hidden in the O^* notation.

1.1. Finding a path

Given an undirected graph G on n vertices, the *k-path problem* asks whether G contains a simple path on k vertices.

Table 1: k-path in time $O^*(f(k))$

$k!$		Monien [2]	4^k	r	Chen et al. [6]
$k!2^k$		Bodlaender [3]	2.86^k		Fomin et al. [7]
5.44^k	r	Alon et al. [5]	2.83^k	r	Koutis [8]
c^k		$c > 8000$, Alon et al. [5]	2^k	r	Williams [9]
16^k		Kneis et al. [10]	1.66^k	r	*this paper*
$4^{k+o(k)}$		Chen et al. [6]			

Theorem 1. *The undirected k-path problem can be solved in time $O^*(1.66^k)$ by a randomized algorithm with constant, one-sided error.*

The proof is in Section 2. For $k = n$, the result matches the running time of Björklund's [1] algorithm for Hamiltonian path.

Previous work. Naively, the k-path problem can be solved in time $O^*(n^k)$, but Monien [2] and Bodlaender [3] showed that the problem can be solved in time $O^*(f(k))$, leading Papadimitriou and Yannakakis [4] to conjecture that the problem was polynomial-time solvable for $k = O(\log n)$. This was confirmed in a strong sense by Alon et al. [5], with a beautiful $O^*(c^k)$ algorithm. A number of paper have since reduced the base c of the exponent using different techniques; see Table 1. In the tables here and henceforth, we mark randomized algorithms with 'r'.

Our result is yet another improvement of the exponent base, notable perhaps mostly because it breaks the psychological barrier of $c = 2$. We fully expect this development to continue. On the other hand, computational complexity informs us that an even more ambitious goal may be quixotic: An algorithm for k-path with running time $\exp(o(k))$ would solve the Hamiltonian path problem in time $\exp(o(n))$, which is known to contradict the exponential time hypothesis [11].

1.2. Packing disjoint triples

Let $\mathcal{F}$ be a family of subsets of an n-element ground set. A subset $\mathcal{A} \subseteq \mathcal{F}$ is a p-packing if $|\mathcal{A}| = p$ and the sets in $\mathcal{A}$ are pairwise disjoint. Given input set $\mathcal{F}$ of size-3 subsets, the *p-packing of 3-sets* problem asks whether $\mathcal{F}$ contains a p-packing. This problem includes a number of well-studied problems in which the ground set consists of the vertices of an input graph $G = (V, E)$.

In the *vertex-disjoint triangle p-packing problem*, the set $\mathcal{F}$ consists of the subsets of V that form a triangle K_3. In the *edge-disjoint triangle p-packing problem*, the set $\mathcal{F}$ consists of the subsets of E that form the edges of a triangle. In the *vertex-disjoint P_3 p-packing problem*, the set $\mathcal{F}$ consists of the vertex subsets $\{u, v, w\} \subseteq V$ for which $uv, vw \in E$.

Theorem 2. *The p-packing of 3-sets problem can be solved by a randomized algorithm in time $O^*(1.493^{3p})$ with constant one-sided error.*

The proof is in Section 4. For $p = n$, the result matches the running time of the recent algorithm for exact cover by 3-sets of Björklund [12].

Previous work. The naive algorithm for p-packing considers all $\binom{|\mathcal{F}|}{p}$ ways of selecting p sets from $\mathcal{F}$. Results of the form $O^*(f(p))$ go back to Downey and Fellows [13], and the dependency on p has been improved dramatically in a series of papers, see Table 2. Remarkably, the best previous running time is given by Koutis's algorithm for the more general problem of p-packing sets of size q, specialized to the case $q = 3$. (We return to the performance of our own algorithm in the case $q > 3$ in Section 1.3.)

It is known that packing vertex-disjoint copies of H into G is NP-complete as soon as H is connected and has more than 2 vertices [23]. Fellows et al. [16] raised the question of how hard the parameterized problem is, and we observe here that this can be answered under the exponential time hypothesis [11], which is equivalent to the parameterized complexity hypothesis FPT $\neq$ M[1] [24, Chapter 16].

Proposition 3. *There is no algorithm for vertex-disjoint triangle p-packing in time $\exp(o(p))$ unless the satisfiability of 3-CNF formulas in n variables can be decided in time $\exp(o(n))$.*

The proof of the proposition is routine and we omit a detailed presentation. Briefly, a 3-CNF formula with n variables and m clauses is transformed to a 3-dimensional matching instance of size $O(n+m)$, and the sparsification lemma of Impagliazzo et al. [11] is used to remove the dependency on m. The needed linear-time reduction can be found in the textbooks of Kleinberg and Tardos [25, p. 483] and Dasgupta et al. [26, p. 267]: For each variable we have a gadget that consists of $2k$ overlapping triangles, where k is the total number of negated and unnegated occurrences of the variable; the gadget has two 3-matchings, corresponding to the two possible assignments of the variable.

Table 2: p-packings of 3-sets in time $O^*(f(p))$

$2^{O(p)}(3p)!$		Downey and Fellows [13]
$(5.7p)^p$	r	Jia et al. [14]
$2^{O(p)}$		Koutis [15]
$(12.7c)^{3p}$		$c > 10.4$, Fellows et al. [16]
10.88^{3p}	r	Koutis [15]
4.68^{3p}		Kneis et al. [10]
$4^{3p+o(p)}$		Chen et al. [6]
2.52^{3p}	r	Chen et al. [6]
$2^{2p\log p+1.869p}$		vertex-disjoint triangles K_3, Fellows et al. [16]†
$22.628^{p\log p+p}$		edge-disjoint triangles K_3, Mathieson et al. [17]
4.61^{3p}		Liu et al. [18]
3.523^{3p}		Wang and Feng [19]
3.404^{3p}		vertex-disjoint paths P_3, Prieto and Sloper [20]
2.604^{3p}		vertex-disjoint paths P_3, Wang et al. [21]
2.482^{3p}		vertex-disjoint paths P_3, Fernau and Raible [22]
2^{3p}	r	Koutis [8]
1.493^{3p}	r	*this paper*

† The precise time bound can be seen to be $O(2^{2p\log p+1.869p}n^3)$.

In addition, we have a constant-size gadget for each clause and "cleanup" gadgets, which at most double the size of the construction. A subexponential time (in $p \leq n$) algorithm for 3-dimensional matching would thus solve the satisfiability of the 3-CNF formula in time $\exp(o(p)) = \exp(o(n))$.

1.3. Uniform set packing

As before, let $\mathcal{F}$ be a set of subsets of an n-element ground set. A subset $\mathcal{A} \subseteq \mathcal{F}$ is a *p-packing* if $|\mathcal{A}| = p$ and the sets in $\mathcal{A}$ are pairwise disjoint. Given as input a family $\mathcal{F}$ of size-q subsets, the *p-packing of q-sets* problem asks us to determine whether $\mathcal{F}$ contains a p-packing.

This is a generalization of the triple p-packing problem described in Section 1.2, and the algorithm we advertised in that section is merely a specialization of a more general result.

Theorem 4. *The p-packing of q-sets problem can be solved by a randomized*

Table 3: p-packing of q-sets in time $O^*(f(p,q))$

$\exp(O(pq))$		Fellows et al. [27]
5.44^{qp}	r	Koutis [15]
2^{qp}	r	Koutis [8]

algorithm in time $O^*(f(p,q))$ with constant, one-sided error, where

$$f(p,q) = \left\{ \frac{0.108157 \cdot 2^q (1 - 1.64074/q)^{1.64076-q} q^{0.679625}}{(q-1)^{0.679623}} \right\}^p .$$

Potentially, $\mathcal{F}$ can have size $\binom{n}{q}$, so reading in the input alone can take super-polynomial time in n. Thus we adopt the convention that $|\mathcal{F}|$ is polynomial in n. Thus, the O^* notation suppresses polynomial factors in n (and hence in $p, q \leq n$) and also in $|\mathcal{F}|$; a more careful (and even less readable) bound on the running time is given in Section 4.7.

Still, the above expression is difficult to parse. The previous best bound is Koutis's [8] much cleaner $O^*(2^{qp})$, and our bound is not $O^*((2-\epsilon)^{qp})$ for any ϵ. Instead, our algorithm behaves well on small q; for comparison, we can express bounds on f of the form $O^*(c^{qp})$ for small $q \geq 3$:

q	3	4	5	6	7	8
$f(p,q)$	1.4953^{3p}	1.6413^{4p}	1.7205^{5p}	1.7707^{6p}	1.8055^{7p}	1.8311^{8p}

,

q	10	20	50	100	500
$f(p,q)$	1.8663^{10p}	1.9345^{20p}	1.9741^{50p}	1.9871^{100p}	1.9975^{500p}

.

The proof is in Section 4. For $p = n$, the result matches the running time of the recent algorithm for exact cover by q-sets of Björklund [12].

Previous work. Most of the work on p-packing of q-sets has been done for the special case $q = 3$, described in Section 1.2. For general q, the first algorithm with running time of the form $O^*(f(p,q))$ is due to Jia et al. [14]; the subsequent improvements are shown in Table 3.

For the special graph packing case where $\mathcal{F}$ consists of isomorphic copies of a fixed graph H of size q, an earlier result established $O(2^{pq \log p + 2pq \log q} n^q)$ [16]. The only specific packing problem we are aware of that our general algorithm does not seem to improve is the problem of packing vertex-disjoint

stars into a given graph. A star $K_{1,q-1}$ consists of a center vertex connected to $q-1$ other vertices. Prieto and Sloper [20] exhibit a kernel of polynomial size $s(p,q) = p(q^3 + pq^2 + pq + 1)$ for this problem. They do not express running times in terms of q, but their "brute force" algorithm can be seen to run in time within a polynomial factor of $\binom{s(p,q)}{p} = \exp(O(p \log pq))$.

Significant improvements for p-packing of q-sets, such as a general $\exp(q \cdot o(p))$ algorithm, are ruled out by Proposition 3. For the nonuniform p-packing problem, when there is no bound on the size of the packed sets, we are unlikely to find an algorithm with running time $O^*(f(p))$ for any function f; the main evidence is provided in terms of parameterized complexity, where the more specific problem of finding an independent set of size p (equivalently, packing p subsets of is the family $\mathcal{F}$ of closed vertex neighborhoods in a given graph) is $W[1]$-hard [13].

1.4. Multidimensional matching

Let $U_1, U_2, \ldots, U_q$ be pairwise disjoint sets, each of size r. Let $\mathcal{F}$ be a collection of subsets of $U_1 \cup U_2 \cup \cdots \cup U_q$ such that each $A \in \mathcal{F}$ satisfies $|A \cap U_j| = 1$ for each $j = 1, 2, \ldots, q$. Given $\mathcal{F}$ as input, the *q-dimensional p-packing problem* asks whether $\mathcal{F}$ contains a p-packing. One often views this problem as q-dimensional *p-matching*, in which case $\mathcal{F}$ is thought of as the edge set of a q-uniform q-partite hypergraph over $U_1 \times U_2 \times \cdots \times U_q$.

Again, $\mathcal{F}$ itself can have size up to r^q, so reading the input alone would take time super-polynomial in the size $n = qr$ of the universe already. Thus, we adopt the convention that $|\mathcal{F}|$ is polynomial in n. In particular, the O^* notation hides factors polynomial in both n (and hence $p, q \leq n$) and $|\mathcal{F}|$.

Theorem 5. *The q-dimensional p-packing problem can be solved in time $O^*(2^{(q-2)p})$ by a randomized algorithm with constant, one-sided error.*

The proof is in Section 3. For $p = r$, the result matches the running time of the recent algorithm for q-dimensional perfect matching of Björklund [12].

Previous work. The first parameterized multi-dimensional matching algorithm appears to be Downey, Fellows, and Koblitz's [28] application of the color-coding technique. Some of the ensuing improvements apply only to the 3-dimensional case. See Table 4.

Table 4: q-dimensional p-packing in time $O^*(f(p,q))$

$(qp)!(qp)^{3qp+1}$		Downey et al. [28]
$\exp(O(pq))$		Fellows et al. [27]
$\exp(O(pq))$		Koutis [15]
$(4e)^{pq}$	r	Koutis [15]
$2.5961^{(q-1)p}$		Zehavi [29]
2.80^{3p}		$q = 3$, Chen et al. [6]
2.77^{3p}		$q = 3$, Liu et al. [18]
2.52^{3p}	r	$q = 3$, Chen et al. [6]
2^{qp}	r	Koutis [8]
$2^{(q-1)p}$	r	Koutis and Williams [30]
$2^{(q-2)p}$	r	*this paper*

1.5. Edge coloring

Finally, we turn to an exact algorithm for edge coloring.

Let G be an undirected loopless n-vertex d-regular graph without parallel edges. The *edge-coloring problem* asks whether the edges of G can be colored so that no two edges that share an endvertex have same color. By Vizing's theorem [31] the number of colors required is either d or $d + 1$.

Theorem 6. *The d-edge coloring problem for d-regular graphs can be solved in time $O^*(2^{(d-1)n/2})$ and polynomial space by a randomized algorithm with constant, one-sided error.*

The proof is in Section 5.

Previous work. The naive algorithm for edge coloring tests all d^m assignments a d colors to m edges. To the best of our knowledge, the only other known exact algorithm is to look for a *vertex* d-coloring of the line graph $L(G)$ of G. The line graph $L(G)$ has m vertices, so the algorithm of Björklund, Husfeldt, and Koivisto Björklund et al. [32] solves the problem in time (and space) $O^*(2^m)$, which is $O^*(2^{dn/2})$ for d-regular graphs. For small numbers of colors, faster algorithms are given by Couturier et al. [33, Theorems 5 and 6]: 3-edge and 4-edge colorings can be counted in time $O(1.201^n)$ and $O(1.818^n)$, respectively, using exponential space.

It is known that for each $d \geq 3$ it is an NP-complete problem to decide whether d colors suffice [34, 35]. However, the exponential time complexity

of edge coloring remains wide open [36]. Curiously, we do not know how to apply these ideas to *vertex* coloring; a polynomial space algorithm with running time $O^*(2^n)$ has not yet been found.

1.6. Methods

Our methods follow the idea introduced by Koutis [8] of expressing a parameterized problem in an algebraic framework by associating multilinear monomials with the combinatorial structures we are looking for, ultimately arriving at a polynomial identity testing problem. We develop various ideas for sieving through these monomials by canceling unwanted contributions.

Some of our results are the parameterized analogues of recent work by the first author [12, 1], all using a determinant summation idea that essentially goes back to Tutte. To make these ideas work in the parameterized setting is not straightforward. In particular, while the Hamiltonian path algorithm [1] uses determinants to cancel the contribution of unwanted labeled cycle covers, our k-path algorithm from Theorem 1 uses a combinatorial argument to pair unwanted labeled walks of k vertices. With $k = n$ we recover a $O^*(1.66^n)$ Hamiltonian path algorithm, but using a different (and arguably more natural) approach. On the other hand, the parameterized packing and matching results of Theorems 2, 4, and 5, and also the edge coloring result in Theorem 6 all use determinants.

Our algorithm for k-path seems to be subtle in the sense that we see no way of extending it to other natural combinatorial structures, even directed paths. In contrast, the ideas of Koutis [8] and Williams [9] do work for directed graphs, and for detecting k-vertex trees and k-leaf spanning trees [30] in time $O^*(2^k)$.

It seems to be difficult to achieve our results using previous techniques. In particular, the limitations of the group algebra framework [8, 9] were studied by Koutis and Williams [30]; they show that multilinear polynomials of degree k cannot be detected in time faster than $O^*(2^k)$ in their model. Koutis [8] has argued that the color coding method [5] and the randomized divide-and-conquer approach [6]) also cannot achieve running times whose exponent base is better than $O^*(2^k)$.

1.7. Extensions and related work

The presented algorithms can be extended to natural weighted problem variants by augmenting the monomials by a new indeterminate whose degree accumulates the total weight of the respective object, like a k-path. For

the traveling salesman problem this approach has been adopted already by Kohn et al. [37] and more recently by Björklund [1]. Accommodating weights increases the running time of the algorithms by a factor that is nearly linear in the maximum weight over the objects, M, namely $M\text{polylog}(M)$.

While our randomized algorithms are the fastest we know to date for the problems in question, there has been recent progress in deterministic algorithms. Specifically, the method of representative sets has resulted in currently the fastest deterministic algorithms for the k-path problem [7] and the q-dimensional p-packing problem [29]; see Tables 1 and 4. The method applies also to weighted problem variants, increasing the running time by a factor that is only polylogarithmic in the maximum weight.

Since the publication of a preliminary version of the present work [38], several works have applied similar algebraic techniques to other combinatorial problems [39, 40]. Koutis and Williams [41] give a gentle review of the broader framework of "algebraic fingerprints."

2. A Projection Sieve for k-Paths

This section establishes Theorem 1.

2.1. Overview

In this section, we develop an inclusion–exclusion sieve over multivariate polynomials for the k-path problem.

The vertices of the input graph are randomly partitioned into two sets V_1, V_2 of roughly equal size. Central to our analysis is the family of *labeled walks*, defined relative to such a random partition as follows: each occurrence on the walk of a vertex in $G[V_1]$ and of an edge in $G[V_2]$ receives a unique label. We call this a *bijective* labeling. The labels need not correspond to the order in which the objects are visited, and the same object can incur more than one label. Examples of labeled walks on a graph whose vertices are partitioned into $V_1 = \{a, b, c\}$, $V_2 = \{d, e, f\}$ are

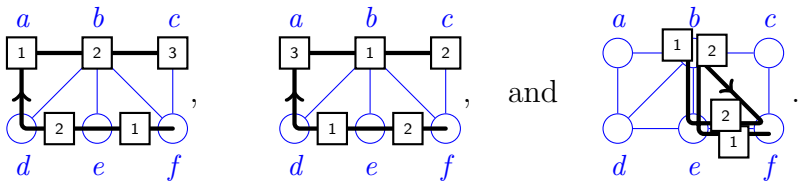

Note the asymmetry between what is labeled in $G[V_1]$ and $G[V_2]$; indeed, with good probability a path of length k has roughly $k/2$ vertex labels in $G[V_1]$ but only $k/4$ edge labels in $G[V_2]$. Very roughly speaking, the running time of our algorithm is around $2^{k/2+k/4}$ for that reason.

With each such labeled walk we associate a monomial consisting of variables x_e for every edge e on the walk, variables $y_{v,i}$ for every vertex v of $G[V_1]$ labeled i, and variables $z_{e,i}$ for every edge e of $G[V_2]$ labeled i. For example, the monomial associated with leftmost labeled walk above is

$$x_{ab} \cdot x_{ad} \cdot x_{bc} \cdot x_{de} \cdot x_{ef} \cdot y_{a,1} \cdot y_{b,2} \cdot y_{c,3} \cdot z_{de,2} \cdot z_{ef,1} \, ;$$

the monomial associated with the middle labeled walk contains the same xs, but the remaining factors are $y_{a,3} \cdot y_{b,1} \cdot y_{c,2} \cdot z_{de,1} \cdot z_{ef,2}$.

If the labeled walk is a path (i.e., it has no repeated vertices), then it can be uniquely recovered, including the labels, from its associated monomial and knowledge of the source vertex. On the other hand, two different labeled walks that are not paths can have the same monomial, for example,

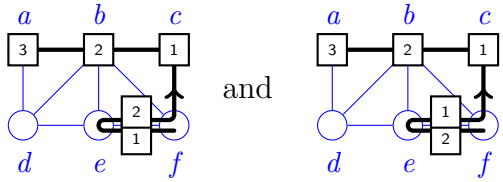

are both associated with

$$x_{ab} \cdot x_{bc} \cdot x_{cf} \cdot x_{ef}^2 \cdot y_{a,3} \cdot y_{b,2} \cdot y_{c,1} \cdot z_{ef,1} \cdot z_{ef,2} \, .$$

In fact, we will find a pairing such that every labeled non-path has exactly one such partner. In particular, their (identical) monomials will cancel each other when added in a field of characteristic 2, and only the monomials corresponding to labeled paths will remain. Representative examples of this pairing are given in Figure 1. A good part of our exposition is devoted to a very careful description of this pairing; to appreciate why such caution is necessary, note that walks like

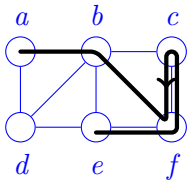

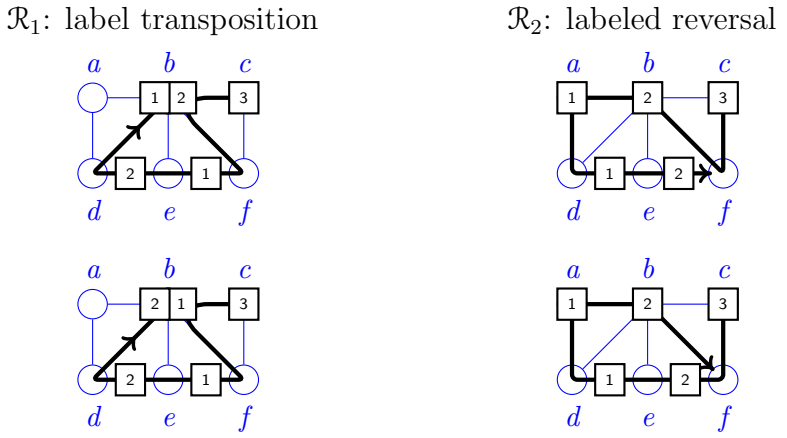

Figure 1: Representative examples of the pairing of labeled walks. Left: if the walk's first closed subwalk starts at a vertex in $V_1 = \{a, b, c\}$ (here, vertex c), then the labels at this vertex are transposed, while the walk remains unchanged. Right: if the walk's first closed subwalk starts at a vertex in $V_2 = \{d, e, f\}$ (here, vertex f), then the subwalk and its labels are reversed.

are not correctly handled by our set-up and indeed will require an exception in the definition.

In summary, there are four key ingredients. First, the labeled k-walks that are paths will be associated with monomials that have distinct variable supports. Second, the monomials associated with *bijectively* labeled k-walks that are not paths will cancel over a field of characteristic 2, established by a pairing argument. Third, an inclusion–exclusion sieve will be used to cancel all walks that are not bijectively labeled. The sieve requires at most $3k/4$ labels (with high probability) if we are careful about the walks we consider. Fourth, the polynomial of k-walks that *avoid* a given set of labels can be evaluated in time polynomial in n.

2.2. Strings

From a technical perspective it will be convenient to view a walk in a graph as a string. To this end, let us review some basic terminology on strings.

Let A be a set whose elements we view as symbols of an alphabet. A *string* of *length* ℓ over A is a sequence $S = s_1 s_2 \cdots s_\ell$ with $s_i \in A$ for each

$i = 1, 2, \ldots, \ell$. We say that s_i is the symbol at *position* i of the string. The *reverse* of a string $S = s_1 s_2 \cdots s_\ell$ is $\overleftarrow{S} = s_\ell s_{\ell-1} \cdots s_1$. The *concatenation* of two strings $S = s_1 s_2 \cdots s_\ell$ and $T = t_1 t_2 \cdots t_k$ is $ST = s_1 s_2 \cdots s_\ell t_1 t_2 \cdots t_k$. A string $T = t_1 t_2 \ldots t_k$ is a *substring* of a string $S = s_1 s_2 \cdots s_\ell$ if there exists a $j = 1, 2, \ldots, \ell - k + 1$ such that $t_i = s_{i+j-1}$ for all $i = 1, 2, \ldots, k$. A *palindrome* is a string that is identical to its reverse and has length at least 2. For $A_1, A_2, \ldots, A_\ell \subseteq A$, we say that a string $s_1 s_2 \cdots s_\ell$ is an $A_1 A_2 \cdots A_\ell$-*string* if $s_i \in A_i$ holds for every $i = 1, 2, \ldots, \ell$.

Let $a_1 a_2 \cdots a_\ell$ be a string over an alphabet A. It will be convenient to view a string as a set consisting of pairs $(a, i) \in A \times \{1, 2, \ldots, \ell\}$, where the pair (a, i) indicates that the symbol a occurs at position i, that is, $a_i = a$.

For a subset $B \subseteq A$ and a string $a_1 a_2 \cdots a_\ell$, introduce the notation

$$B\{a_1 a_2 \cdots a_\ell\} = \{(a_i, i) : a_i \in B\} \subseteq A \times \{1, 2, \ldots, \ell\}.$$

In particular, we can recover the string $a_1 a_2 \cdots a_\ell$ from the set $A\{a_1 a_2 \cdots a_\ell\}$.

2.3. Walks

We assume that all graphs are undirected and contain neither loops nor parallel edges. For a graph G, we denote the vertex set of G by $V = V(G)$, and the edge set of G by $E = E(G)$. For convenience, we assume that V and E are disjoint sets.

A *k-walk* in G is a string of length $2k - 1$ such that

(a) each odd position contains a vertex of G;

(b) each even position contains an edge of G; and

(c) for every $i = 1, 2, \ldots, k - 1$, the edge at position $2i$ joins in G the vertices at positions $2i - 1$ and $2i + 1$.

The first and last positions of a walk are the *ends* of the walk.

A walk is a *path* if each vertex of G appears in at most one position of the walk. A k-walk is *closed* if its ends are identical and $k \geq 2$. A walk that is not a path always contains at least one closed subwalk. A *subwalk* of a walk W is a substring of W that is in itself a walk. Put otherwise, a subwalk of W is a substring with ends at odd positions of W.

To reduce the number of labels in the sieve, we will focus on a somewhat technical subset of k-walks that we will call "admissible" walks.

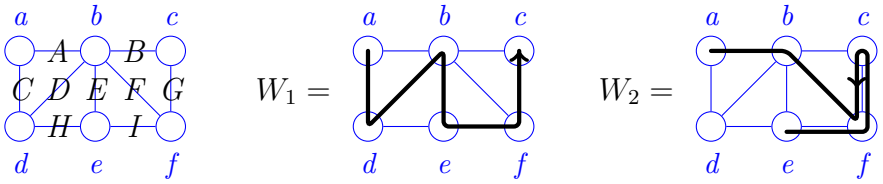

Figure 2: Terminology and notation for walks introduced in Section 2.2–2.3. The 6-vertex walk W_1 is encoded by the 11-letter string *aCdDbEeIjGc*. With $V_1 = \{a, b, c\}$ we have $E_2 = \{H, I\}$ and the "projections" $V_1\{W_1\} = \{(a, 1), (b, 5), (c, 11)\}$ and $E_2\{W_1\} = \{(I, 8)\}$. The walk W_1 is admissible with the following parameters: starting vertex $s = a$, number of vertices $k = 6$, number of V_1-vertices $k_1 = 3$, and number of E_2-edges $\ell_2 = 1$. The walk $W_2 = eIfGcGfFbAa$ is not admissible, because it contains the $V_2EV_1EV_2$-palindrome *fGcGf*. For readability, in all other examples we denote edges by node pairs, for instance writing *ef* instead of I, with the understanding that *ef* is viewed as a single symbol for purposes of indexing W.

Let G be a graph with vertex set V and let s be a fixed vertex of G. Partition the vertex set into two disjoint sets $V = V_1 \cup V_2$. Denote by E_1 the set of edges of G with both ends in V_1. Denote by E_2 the set of edges of G with both ends in V_2. Let k, k_1, ℓ_2 be nonnegative integers.

Let us say that a k-walk W in G is *admissible* if

(a) W starts at s;

(b) $|V_1\{W\}| = k_1$;

(c) $|E_2\{W\}| = \ell_2$; and

(d) W is $V_2EV_1EV_2$-palindromeless.

Here the term "palindromeless" refers to the property that a string has no palindrome as a substring. By $V_2EV_1EV_2$-palindromeless we refer to the lack of palindromes that are also $V_2EV_1EV_2$-strings. Observe that paths are palindromeless and hence $V_2EV_1EV_2$-palindromeless.

2.4. Random projection

For a fixed ordered partition (V_1, V_2), every k-path P in G that starts at s is admissible for some parameters k_1, ℓ_2. Conversely, for fixed parameters k_1, ℓ_2 and a fixed k-path P that starts at s, if we select (V_1, V_2) uniformly at random, then P is admissible with probability given by the following lemma.

Lemma 7 (Admissibility)**.** *Let k_1, ℓ_2 be nonnegative integers and let P be a k-path in G. For (V_1, V_2) selected uniformly at random, we have*

$$\Pr\Big(|V_1\{P\}| = k_1 \text{ and } |E_2\{P\}| = \ell_2\Big) = 2^{-k}\binom{k_1+1}{k-k_1-\ell_2}\binom{k-k_1-1}{\ell_2}.$$

Proof. There are 2^k strings of length k over the alphabet $\{1,2\}$. The probability in question is exactly the fraction of such strings that have exactly k_1 1-positions and exactly ℓ_2 22-substrings. There are exactly k_1+1 positions where to interleave the k_1 1s with substrings of 2s. Each such substring of length j contributes exactly $j-1$ 22-substrings. The total number of 2s is $k-k_1$, so there must be $k-k_1-\ell_2$ substrings of 2s. The positions where the substrings interleave the 1s are allocated by the first binomial coefficient. It remains to allocate the lengths of the strings. The total length is $k-k_1$, and each of the $k-k_1-\ell_2$ strings must have length at least 1. Thus there are $k-k_1-(k-k_1-\ell_2) = \ell_2$ free 2s to allocate to $k-k_1-\ell_2$ distinct bins. The second binomial coefficient carries out this allocation, since N unlabeled balls can be allocated to K labeled bins in $\binom{N+K-1}{K-1}$ different ways. □

In particular, a fixed k-path starting at s is admissible with positive probability if and only if either $k_1 = k$ and $\ell_2 = 0$ or $k_1 < k$ and $k_1 + \ell_2 \le k-1 \le 2k_1 + \ell_2$.

Let us now derive an asymptotic approximation for the probability in Lemma 7. We employ the following variant of Stirling's formula due to Robbins [42]. For all $j = 1, 2, \ldots$ it holds that

$$j! = \sqrt{2\pi j}\left(\frac{j}{e}\right)^j e^{\epsilon_j} \quad \text{where} \quad \frac{1}{12j+1} < \epsilon_j < \frac{1}{12j}\,. \tag{1}$$

Let us abbreviate

$$\left\langle {a \atop b} \right\rangle = \left(\frac{b}{a}\right)^{-b}\left(1-\frac{b}{a}\right)^{-a+b}.$$

From Stirling's formula (1) it follows that $\binom{a}{b} = \Theta^*\big(\left\langle {a \atop b} \right\rangle\big)$ holds uniformly for all $0 < b < a \le n$. We can thus approximate the probability in Lemma 7, uniformly for $0 < \ell_2, k_1 < k$ such that $k_1 + \ell_2 \le k-1 \le 2k_1+\ell_2$, with

$$\begin{aligned}\Pr\Big(|V_1\{P\}| = k_1 \text{ and } |E_2\{P\}| = \ell_2\Big) = \\ = \Theta^*\left(2^{-k}\left\langle {k_1 \atop k-k_1-\ell_2}\right\rangle\left\langle {k-k_1 \atop \ell_2}\right\rangle\right).\end{aligned} \tag{2}$$

2.5. Labeled admissible walks

The following labeling scheme for admissible walks serves two purposes. First, labeling enables us to "decouple" the sieve from the graphical domain (that is, vertices and edges) into a set of abstract labels whose number depends only on the parameters k_1, ℓ_2 and not on the size of the graph. Second, the labeling facilitates cancellation of non-paths in the sieve.

Let K_1 be a set of k_1 labels. Let L_2 be a set of ℓ_2 labels. For example, let $K_1 = \{1, 2, \ldots, k_1\}$ and $L_2 = \{1, 2, \ldots, \ell_2\}$.

Let W be an admissible walk. Let $\kappa_1 : V_1\{W\} \to K_1$ and $\lambda_2 : E_2\{W\} \to L_2$ be arbitrary functions. The three-tuple (W, κ_1, λ_2) is a *labeled* admissible walk. Intuitively, each position in W that contains a vertex in V_1 gets assigned a label in K_1 by κ_1. Similarly, each position in W that contains an edge in E_2 gets assigned a label in L_2 by λ_2. Let us say that the labeling is *bijective* if both κ_1 and λ_2 are bijections.

Example. Consider two labelings of the same walk W,

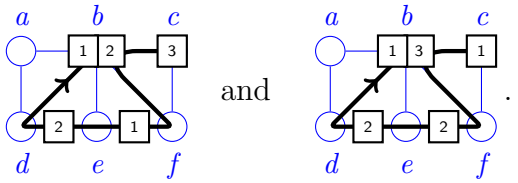

The walk is admissible with parameters $s = \mathrm{c}$, $k = 6$, $k_1 = 3$, and $\ell_2 = 2$. Both labelings associate a label with each position of W that contains a symbol from $V_1 = \{a, b, c\}$ or $E_2 = \{de, ef\}$. We have $V_1\{W\} = \{(b, 3), (b, 11), (c, 1)\}$, that is, there are three occurrences of symbols from V_1 in W; in particular the symbol b occurs at the 3rd and the 11th position. Similarly, we have $E_2\{W\} = \{(ef, 6), (de, 8)\}$. The labeling on the left is

$$\kappa_1(b, 3) = 2,\ \kappa_1(b, 11) = 1,\ \kappa_1(c, 1) = 3,\ \lambda_2(ef, 6) = 1,\ \lambda_2(de, 8) = 2.$$

We observe that this labeling is bijective since λ_1 and κ_2 are bijections.

The labeling on the right is

$$\kappa_1(b, 3) = 3,\ \kappa_1(b, 11) = \kappa_1(c, 1) = 1,\ \lambda_2(ef, 6) = \lambda_2(de, 8) = 2,$$

and not bijective. In fact, λ_1 avoids the label 2, and κ_2 avoids the label 1.

2.6. Fingerprinting and identifiability

We associate with each labeled admissible walk an algebraic object (or "fingerprint") that is used to represent the labeled admissible walk in sieving. Here it is important to observe that while we are careful to design the fingerprint so that each labeled path has a unique fingerprint, the fingerprints of labeled non-paths are by design *not* unique—we will explicitly take advantage of this property when canceling labeled non-paths in Section 2.8.

The sieve operates over a multivariate polynomial ring with the coefficient field $\mathbb{F}_{2^b}$ (the finite field of order 2^b) and the following indeterminates. Introduce one indeterminate x_e for each edge $e \in E$. Introduce one indeterminate $y_{v,i}$ for each pair $(v,i) \in V_1 \times K_1$. Introduce one indeterminate $z_{e,i}$ for each pair $(e,i) \in E_2 \times L_2$.

Let (W, κ_1, λ_2) be a labeled admissible walk. Associate with (W, κ_1, λ_2) the *monomial* (*fingerprint*)

$$m(W, \kappa_1, \lambda_2) = \prod_{(e,j) \in E\{W\}} x_e \prod_{(v,i) \in V_1\{W\}} y_{v,\kappa_1(v,i)} \prod_{(e,i) \in E_2\{W\}} z_{e,\lambda_2(e,i)} \,.$$

The following lemma is immediate.

Lemma 8 (Identifiability)**.** *The monomial $m(W, \kappa_1, \lambda_2)$ of a labeled admissible walk (W, κ_1, λ_2) uniquely determines the edges and their multiplicities of occurrence in W. In particular, any path is uniquely identified. Furthermore, if W is a path and κ_1, λ_2 are bijections, then $m(W, \kappa_1, \lambda_2)$ uniquely identifies (W, κ_1, λ_2).*

Example. We presented some example monomials already in Section 2.1. We can also consider a bijectively labeled walk that repeats an edge in E_2:

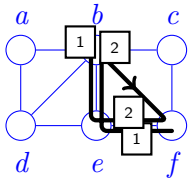

$$x_{be}^2 \cdot x_{bf} \cdot x_{ef}^2 \cdot y_{b,1} \cdot y_{b,2} \cdot z_{ef,1} \cdot z_{ef,2} \,,$$

and a non-bijectively labeled walk,

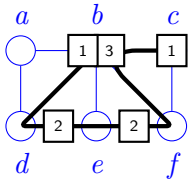

$$x_{bc} \cdot x_{bd} \cdot x_{bf} \cdot x_{de} \cdot x_{ef} \cdot y_{b,1} \cdot y_{b,3} \cdot y_{c,1} \cdot z_{de,2} \cdot z_{ef,2} \,.$$

In all these examples observe that if the walk is a path, we can reconstruct it from the x-variables and knowledge of the start vertex. Because a path has neither repeated vertices nor edges, the y- and z-variables in the monomial enable us to reconstruct the labeling.

2.7. Sieving for bijective labelings

Let us denote by $\mathcal{L}$ the set of all labeled admissible walks. For $I_1 \subseteq K_1$ and $J_2 \subseteq L_2$, denote by $\mathcal{L}[I_1, J_2]$ the set of all labeled admissible walks that *avoid* the labels in I_1 and J_2. Let us denote by $\mathcal{B}$ the set of all bijectively labeled admissible walks.

By the principle of inclusion–exclusion, we have

$$\sum_{(W,\kappa_1,\lambda_2)\in\mathcal{B}} m(W,\kappa_1,\lambda_2) = \sum_{I_1\subseteq K_1} \sum_{J_2\subseteq L_2} (-1)^{|I_1|+|J_2|} \sum_{(W,\kappa_1,\lambda_2)\in\mathcal{L}[I_1,J_2]} m(W,\kappa_1,\lambda_2)\,. \qquad (3)$$

2.8. Bijectively labeled non-path fingerprints cancel

Let us partition $\mathcal{B}$ into $\mathcal{B} = \mathcal{P} \cup \mathcal{R}$, where $\mathcal{P}$ consists of bijectively labeled admissible paths, and $\mathcal{R}$ consists of bijectively labeled admissible non-paths. Accordingly, the left-hand side of (3) splits into

$$\sum_{(W,\kappa_1,\lambda_2)\in\mathcal{B}} m(W,\kappa_1,\lambda_2) = \sum_{(W,\kappa_1,\lambda_2)\in\mathcal{P}} m(W,\kappa_1,\lambda_2) + \sum_{(W,\kappa_1,\lambda_2)\in\mathcal{R}} m(W,\kappa_1,\lambda_2)\,.$$

We show that the rightmost sum vanishes. To this end, let us first recall that an *involution* is a permutation that is its own inverse. We claim that it suffices to construct a fixed-point-free involution $\phi : \mathcal{R} \to \mathcal{R}$ with $m(W,\kappa_1,\lambda_2) = m(\phi(W,\kappa_1,\lambda_2))$ for all $(W,\kappa_1,\lambda_2) \in \mathcal{R}$, Indeed, introduce an arbitrary total order to $\mathcal{R}$ and observe that in characteristic 2, we have

$$\begin{aligned}\sum_{(W,\kappa_1,\lambda_2)\in\mathcal{R}} m(W,\kappa_1,\lambda_2) &= \sum_{\substack{(W,\kappa_1,\lambda_2)\in\mathcal{R}\\ (W,\kappa_1,\lambda_2)<\phi(W,\kappa_1,\lambda_2)}} m(W,\kappa_1,\lambda_2) + m(\phi(W,\kappa_1,\lambda_2)) \\ &= \sum_{\substack{(W,\kappa_1,\lambda_2)\in\mathcal{R}\\ (W,\kappa_1,\lambda_2)<\phi(W,\kappa_1,\lambda_2)}} 2m(W,\kappa_1,\lambda_2) \quad = \quad 0\,.\end{aligned}$$

To construct a fixed-point-free involution $\phi : \mathcal{R} \to \mathcal{R}$ with $m(W,\kappa_1,\lambda_2) = m(\phi(W,\kappa_1,\lambda_2))$ for all $(W,\kappa_1,\lambda_2) \in \mathcal{R}$, we observe that every walk W that is not a path contains at least one closed subwalk. In particular, W contains

a *first* closed subwalk, that is, the closed subwalk C with the property that C is the unique closed subwalk in the prefix SC of $W = SCT$.

We denote the first closed subwalk of W by $C(W)$ and by $c(W)$ the first (and hence also the last) vertex of $C(W)$.

Let us partition $\mathcal{R}$ into two disjoint sets, $\mathcal{R}_1$ and $\mathcal{R}_2$, where

$$\mathcal{R}_1 = \{(W, \kappa_1, \lambda_2) \in \mathcal{R} : c(W) \in V_1\}\,,$$
$$\mathcal{R}_2 = \{(W, \kappa_1, \lambda_2) \in \mathcal{R} : c(W) \in V_2\}\,.$$

We proceed to construct the pairing ϕ on these two sets. See Figure 1 for examples.

2.9. The pairing on $\mathcal{R}_1$ – label transposition

Select an arbitrary $(W, \kappa_1, \lambda_2) \in \mathcal{R}_1$. Let j and ℓ be the positions of W that contain the symbol $c(W)$ and constitute the ends of $C(W)$. For brevity, let us write c for $c(W)$. Because $c \in V_1$, we have $(c, j), (c, \ell) \in V_1\{W\}$. Define κ_1' to be identical to κ_1 except that

$$\kappa_1'(c, j) = \kappa_1(c, \ell)\,, \qquad \kappa_1'(c, \ell) = \kappa_1(c, j)\,.$$

Observe that $\kappa_1(c, j) \neq \kappa_1(c, \ell)$ because (W, κ_1, λ_2) is bijectively labeled. Thus, $\kappa_1' \neq \kappa_1$ and $(W, \kappa_1', \lambda_2) \in \mathcal{R}_1$. Furthermore, we have $m(W, \kappa_1, \lambda_2) = m(W, \kappa_1', \lambda_2)$. Thus, we can set $\phi(W, \kappa_1, \lambda_2) = (W, \kappa_1', \lambda_2)$ to obtain the desired fixed-point-free involution on $\mathcal{R}_1$. Indeed, $\phi(W, \kappa_1, \lambda_2) = (W, \kappa_1', \lambda_2) \neq (W, \kappa_1, \lambda_2)$ and $\phi^2(W, \kappa_1, \lambda_2) = (W, \kappa_1, \lambda_2)$.

2.10. The pairing on $\mathcal{R}_2$ – labeled reversal of first closed subwalk

Select an arbitrary $(W, \kappa_1, \lambda_2) \in \mathcal{R}_2$. Let $C = C(W)$ and let S, T be strings such that

$$W = SCT\,.$$

Let us define the string W' by reversing C in W, that is,

$$W' = S\overleftarrow{C}T\,.$$

We observe that the strings C and $\overleftarrow{C}$ have identical ends because C is a closed walk, implying that W' is a walk in G. We also observe that W' is admissible. Indeed, because $c(W') = c(W) \in V_2$, any $V_2EV_1EV_2$-palindrome $ueveu$ in W' can either (a) occur as a subwalk of $\overleftarrow{C}$ in W', or (b) have at

most one position in W' common with $\overleftarrow{C}$. But in both cases we have that $ueveu$ occurs in W, which is a contradiction since W is admissible. Thus, W' is admissible.

We observe that $W = W'$ if and only if C is a palindrome. Furthermore, if we reverse $C(W') = \overleftarrow{C}$ in W', we obtain back W. That is, $W'' = W$.

In terms of string positions, we can characterize the reversal $W \mapsto W'$ using the following permutation of positions. Let j and ℓ be the positions of W that constitute the ends of C. Define the permutation $\rho : \{1, 2, \ldots, k\} \to \{1, 2, \ldots, k\}$ by

$$\rho(i) = \begin{cases} i & \text{if } i < j \text{ or } i > \ell; \text{ and} \\ \ell - i + j & \text{if } j \leq i \leq \ell. \end{cases}$$

Let us denote the symbol at the ith position of W by w_i. The reversal $W \mapsto W'$ can now be characterized by observing that $w'_{\rho(i)} = w_i$ holds for each $i = 1, 2, \ldots, k$.

We now introduce a labeling κ'_1, λ'_2 of W' using the labeling κ_1, λ_2 of W. In particular, let us label W' so that each position of W' is labeled using the label of the ρ-corresponding position in W, if any. In precise terms, using the labeling $\kappa_1 : V_1\{W\} \to K_1$, define the labeling $\kappa'_1 : V_1\{W'\} \to K_1$ for each $(w'_{\rho(i)}, \rho(i)) \in V_1\{W'\}$ by setting $\kappa'_1(w'_{\rho(i)}, \rho(i)) = \kappa_1(w_i, i)$. Similarly, using $\lambda_2 : E_2\{W\} \to L_2$, define $\lambda'_2 : E_2\{W'\} \to L_2$ by setting $\lambda'_2(w'_{\rho(i)}, \rho(i)) = \lambda_2(w_i, i)$ for each $(w'_{\rho(i)}, \rho(i)) \in E_2\{W'\}$.

Now set $\phi(W, \kappa_1, \lambda_2) = (W', \kappa'_1, \lambda'_2)$ and observe that $\phi(W, \kappa_1, \lambda_2) \in \mathcal{R}_2$, $\phi^2(W, \kappa_1, \lambda_2) = (W, \kappa_1, \lambda_2)$, and $m(W, \kappa_1, \lambda_2) = m(\phi(W, \kappa_1, \lambda_2))$.

What is not immediate, however, is that $\phi(W, \kappa_1, \lambda_2) \neq (W, \kappa_1, \lambda_2)$. There are two cases to consider, depending on C.

In the first case, C is not a palindrome, that is, $C \neq \overleftarrow{C}$. Thus, $W' \neq W$ and hence $(W', \kappa'_1, \lambda'_2) \neq (W, \kappa_1, \lambda_2)$.

In the second case, C is a palindrome. Since C is a closed walk, the string C has odd length at least 3. In particular, the length 3 (that is, a palindrome of the form ueu with $u \in V_2$ and $e \in E$) cannot occur because G has no loop edges. For palindromes of length 5, the only possibility is that C is a $V_2E_2V_2E_2V_2$-palindrome. Indeed, C can neither be a $V_1EV_1EV_1$-palindrome nor a $V_1EV_2EV_1$-palindrome because $c(W) \in V_2$. Furthermore, C cannot be a $V_2EV_1EV_2$-palindrome because such palindromes by definition do not occur in the admissible W. Thus, for length 5 the only possibility is

a $V_2EV_2EV_2$-palindrome, that is, a $V_2E_2V_2E_2V_2$-palindrome. Such a palindrome contains two occurrences of an edge in E_2 that are in ρ-corresponding positions. These occurrences get different labels under λ_2 and λ_2'. Thus, $(W', \kappa_1', \lambda_2') \neq (W, \kappa_1, \lambda_2)$. Finally, we observe that C cannot have length more than 5, because an odd-length palindrome of length 7 or more must include a palindrome of length 5, which would contradict the assumption that C is the first closed subwalk in W.

2.11. The algorithm

First, we recall the following result:

Lemma 9 (DeMillo–Lipton–Schwartz–Zippel [43, 44])**.** *Let $p(x_1, x_2, \ldots, x_n)$ be a nonzero polynomial of total degree at most d over the finite field $\mathbb{F}_q$. Then, for $a_1, a_2, \ldots, a_n \in \mathbb{F}_q$ selected independently and uniformly at random,*

$$\Pr(p(a_1, a_2, \ldots, a_n) \neq 0) \geq 1 - \frac{d}{q}.$$

Let us assume the parameters k, k_1, ℓ_2 are fixed so that $k_1 + \ell_2 \leq k - 1 \leq 2k_1 + \ell_2$. (We will set the precise values of k_1, ℓ_2 in what follows.) To decide the existence of a k-path starting at s, we repeat the following randomized procedure.

First, the procedure selects an ordered partition (V_1, V_2) uniformly at random among all the 2^n such partitions. Lemma 7 implies that a fixed k-path P that starts at s is admissible with positive probability.

Next, the procedure makes use of Lemma 9 to witness a nonzero evaluation of a multivariate generating function for labeled admissible k-paths starting at s. In particular, from (3) and Section 2.8 we have that

$$\sum_{(W,\kappa_1,\lambda_2)\in\mathcal{P}} m(W, \kappa_1, \lambda_2) = \sum_{I_1\subseteq K_1} \sum_{J_2\subseteq L_2} (-1)^{|I_1|+|J_2|} \sum_{(W,\kappa_1,\lambda_2)\in\mathcal{L}[I_1,J_2]} m(W, \kappa_1, \lambda_2)\,. \qquad (4)$$

The left-hand side of (4) is a multivariate polynomial of degree at most $k-1+k_1+\ell_2$. It follows from Lemma 8 that the polynomial is not identically zero if and only if G has an admissible k-path starting at s. Namely, if the polynomial is not identically zero, then G has at least one bijectively labeled admissible path (W, κ_1, λ_2), that is, an admissible k-path starting at s; and vice versa, if G has an admissible k-path W starting at s, then we can augment it by labelings κ_1, λ_2 into a bijectively labeled admissible path (W, κ_1, λ_2), rendering the polynomial nonzero.

It remains to evaluate the right-hand side of (4) for a random assignment of values in $\mathbb{F}_{2^b}$ to the indeterminates. To this end, the procedure iterates over each $I_1 \subseteq K_1$ and $J_2 \subseteq L_2$ and employs dynamic programming to evaluate the rightmost sum in (4).

Without loss of generality we can assume $k \geq 3$. For parameters k, k_1, ℓ_2 and a string $T = t_1t_2t_3t_4t_5$ over the alphabet $V \cup E$, our objective is to compute

$$M(k, k_1, \ell_2, T) = \sum_{\substack{(W,\kappa_1,\lambda_2)\in\mathcal{L}[I_1,J_2] \\ T \text{ is a suffix of } W}} m(W, \kappa_1, \lambda_2)\,. \tag{5}$$

In particular, taking the sum over all T, we obtain the rightmost sum in (4).

The recursion for (5) is as follows. For a logical proposition P, let us define $[P]$ to be 1 if P is true and 0 otherwise. For $k > 3$, we observe by induction on k that

$$\begin{aligned} &M(k, k_1, \ell_2, t_1t_2t_3t_4t_5) = \\ &\quad = [t_1t_2t_3t_4t_5 \text{ is not a } V_2EV_1EV_2\text{-palindrome}] \\ &\quad \times \sum_{\substack{e\in E \\ e\,=\,\{v,\,t_1\}}} x_e\big([v \notin V_1] + [v \in V_1] \sum_{j\in K_1\setminus I_1} y_{v,j}\big)\big([e \notin E_2] + [e \in E_2] \sum_{j\in L_2\setminus J_2} z_{e,j}\big) \\ &\qquad\qquad \times M(k-1, k_1 - [v \in V_1], \ell_2 - [e \in E_2], vet_1t_2t_3)\,. \end{aligned} \tag{6}$$

To set up the base cases for the recursion, we observe that $M(k, k_1, \ell_2, T)$ can be computed for all $0 \leq k_1, \ell_2 \leq k = 3$ and all $T = t_1t_2t_3t_4t_5$ in time polynomial in n. Furthermore, $M(k, k_1, \ell_2, T) = 0$ whenever $k_1 > k$ or $\ell_2 \geq k$ or $k_1 < 0$ or $\ell_2 < 0$.

Consequently, for any given assignment of values in $\mathbb{F}_{2^b}$ to the indeterminates x_e, $y_{v,\ell}$, and $z_{e,\ell}$, the procedure evaluates the right-hand side of (4) via (6) in $O(2^{k_1+\ell_2}k^3n^4)$ arithmetic operations over $\mathbb{F}_{2^b}$.

Let us now complete the algorithm by optimizing the parameters for running time and $\Omega(1)$ probability of success. Denoting the probability that a k-path P starting at s is admissible by $P(k, k_1, \ell_2)$, we have that in r repetitions of the procedure at least one repetition finds P admissible with probability $1 - (1 - P(k, k_1, \ell_2))^r \geq 1 - e^{-P(k,k_1,\ell_2)r}$. Setting $r = \lceil 1/P(k, k_1, \ell_2)\rceil$ and $b = \lceil \log_2 6k\rceil$, it follows from Lemma 9 that any fixed k-path starting at s in G is witnessed with probability at least $(1 - e^{-1})/2$ in time $O\big(2^{k_1+\ell_2}k^5n^4/P(k, k_1, \ell)\big)$. Setting $k_1 = \lfloor\gamma_1 k\rfloor$, $\ell_2 = \lfloor\gamma_2 k\rfloor$, and employing (2) to approximate $P(k, k_1, \ell_2)$, we obtain $O^*(1.6569^k)$ time for $\gamma_1 = 0.5$ and

$\gamma_2 = 0.207107$. We found these values of γ_1 and γ_2 by numerical computations; note that any values of these parameters result in a valid upper bound for the running time.

We summarize the algorithm in the following pseudocode:

Algorithm 1 (Detect k-path in graph (V, E) starting at vertex s)**.**

$k_1 \leftarrow \lfloor 0.5k \rfloor$, $\ell_2 \leftarrow \lfloor 0.2017107k \rfloor$
$r \leftarrow \lceil 1/P(k, k_1, \ell_2) \rceil$, $b \leftarrow \lceil \log_2 6k \rceil$
$K_1 \leftarrow \{1, 2, \ldots, k_1\}$, $L_2 \leftarrow \{1, 2, \ldots, \ell_2\}$
repeat r times:
    Let (V_1, V_2) be a random partition of V
    Assign $(x_e), (y_{v,j}), (z_{e,j})$ values from $\mathbb{F}_{2^b}$ independently u.a.r.
    $value \leftarrow 0$
    for all $I_1 \subseteq K_1$ and $J_2 \subseteq L_2$:
        Compute the array M as described in (5–6)
        $value \leftarrow value + (-1)^{|I_1|+|J_2|} \sum_{T \in (V \cup E)^5} M(k, k_1, \ell_2, T)$
    if $value \neq 0$ **then return** "*A k-path detected*"
return "*No k-path detected*"

3. A Determinant Sieve for q-Dimensional p-Packings

This section establishes Theorem 5.

3.1. Prepackings and Edmonds's symbolic determinant

Recall that we are given pairwise disjoint sets $U_1, U_2, \ldots, U_q$, each of size r, and a family $\mathcal{F}$ of subsets of $U_1 \cup U_2 \cup \cdots \cup U_q$, each containing exactly one element from each U_i. The task is to find a p-packing, that is, p pairwise disjoint members of $\mathcal{F}$

Let us say that a subset $\mathcal{A} \subseteq \mathcal{F}$ is a *j-prepacking* if $|\mathcal{A}| = j$ and the sets in $\mathcal{A}$ are pairwise disjoint when projected to $U_1 \cup U_2$. Observe that each $A \in \mathcal{A}$ in a j-prepacking identifies both a unique $u_1(A) \in A \cap U_1$ and a unique $u_2(A) \in A \cap U_2$.

For a bijection $\sigma : U_1 \to U_2$, let us say that a j-prepacking $\mathcal{A}$ is *compatible* with σ if for all $A \in \mathcal{A}$ it holds that $\sigma(u_1(A)) = u_2(A)$. Note that each j-prepacking is compatible with at least one σ.

Edmonds [45] made the algorithmically seminal observation that the determinant of a symbolic $r \times r$ matrix $E = (e_{u_1,u_2})_{u_1 \in U_1, u_2 \in U_2}$ is a signed

generating function over partitions of $U_1 \cup U_2$ into 2-subsets with exactly one element from U_1 and exactly one element from U_2; the matrix is symbolic in the sense that each entry of the matrix is treated as a distinct indeterminate. Indeed, identifying each such partition with a bijection $\sigma : U_1 \to U_2$, we have

$$\det E = \sum_{\substack{\sigma : U_1 \to U_2 \\ \sigma \text{ bijective}}} \operatorname{sgn}_\tau(\sigma) \prod_{u_1 \in U_1} e_{u_1,\sigma(u_1)}\,, \tag{7}$$

where the sign $\operatorname{sgn}_\tau(\sigma)$ is the sign of the permutation $\sigma\tau$ for an arbitrary fixed bijection $\tau : U_2 \to U_1$.

Our strategy is to leverage Edmonds's observation from the dimensions U_1 and U_2 into q dimensions $U_1, U_2, \ldots, U_q$ with sieving. In particular, Edmonds's observation forces the packing constraint in the first two dimensions, which allows us to restrict the sieve to the remaining $q-2$ dimensions.

3.2. Fingerprinting and identifiability

Consider a j-prepacking $\mathcal{A} \subseteq \mathcal{F}$. The *domain* of the prepacking is the set

$$d(\mathcal{A}) = \{(u, A) : u \in A \in \mathcal{A}\} \subseteq (U_3 \cup U_4 \cup \cdots \cup U_q) \times \mathcal{F}\,. \tag{8}$$

Observe that $|d(\mathcal{A})| = j(q-2)$.

Let L be a set of $p(q-2)$ labels. A *labeling* of $\mathcal{A}$ is a pair (σ, λ), where $\sigma : U_1 \to U_2$ is a bijection compatible with $\mathcal{A}$ and $\lambda : d(\mathcal{A}) \to L$ is an arbitrary mapping. The labeling is *bijective* if λ is a bijection. We say that a triple $(\mathcal{A}, \sigma, \lambda)$ is a *labeled* j-prepacking.

The sieve operates over a multivariate polynomial ring with the coefficient field $\mathbb{F}_{2^b}$, with a sufficiently large integer b we fix later, and the following indeterminates. Introduce the indeterminate w for tracking the weight j of a j-prepacking. Associate with each $A \in \mathcal{F}$ an indeterminate x_A. Associate with each pair $(u_1, u_2) \in U_1 \times U_2$ an indeterminate y_{u_1,u_2}. Associate with each pair $(u, \ell) \in (U_3 \cup U_4 \cup \cdots \cup U_q) \times L$ an indeterminate $z_{u,\ell}$.

The *signed monomial* of a labeled j-prepacking $(\mathcal{A}, \sigma, \lambda)$ is

$$m(\mathcal{A}, \sigma, \lambda) = \operatorname{sgn}_\tau(\sigma) w^j \prod_{A \in \mathcal{A}} x_A \prod_{u_1 \in U_1} y_{u_1,\sigma(u_1)} \prod_{(u,A) \in d(\mathcal{A})} z_{u,\lambda(u,A)}\,. \tag{9}$$

For a later use, observe the degree of the monomial is $j + |\mathcal{A}| + |U_1| + |d(\mathcal{A})|$, which is equal to $j + j + r + j(q-2) = r + jq$.

Lemma 10 (Identifiability)**.** *The monomial $m(\mathcal{A}, \sigma, \lambda)$ uniquely determines both $\mathcal{A}$ and σ. Furthermore, if $\mathcal{A}$ is a p-packing and λ is bijective, then $m(\mathcal{A}, \sigma, \lambda)$ uniquely determines λ.*

3.3. Sieving for bijective labelings

Denote by $\mathcal{L}$ the set of all labeled p-prepackings. For $J \subseteq L$, denote by $\mathcal{L}[J]$ the subset of labeled p-prepackings whose labeling *avoids* each label in J. Denote by $\mathcal{B}$ the set of all bijectively labeled p-prepackings.

By the principle of inclusion-exclusion,

$$\sum_{(\mathcal{A},\sigma,\lambda)\in\mathcal{B}} m(\mathcal{A},\sigma,\lambda) = \sum_{J\subseteq L}(-1)^{|J|} \sum_{(\mathcal{A},\sigma,\lambda)\in\mathcal{L}[J]} m(\mathcal{A},\sigma,\lambda)\,. \tag{10}$$

3.4. Fingerprints of bijectively labeled non-p-packings cancel

Partition $\mathcal{B}$ into $\mathcal{B} = \mathcal{P} \cup \mathcal{R}$, where $\mathcal{P}$ is the set of bijectively labeled p-packings, and $\mathcal{R}$ is the set of bijectively labeled p-prepackings that are not packings. Accordingly, the left-hand side of (10) splits into

$$\sum_{(\mathcal{A},\sigma,\lambda)\in\mathcal{B}} m(\mathcal{A},\sigma,\lambda) = \sum_{(\mathcal{A},\sigma,\lambda)\in\mathcal{P}} m(\mathcal{A},\sigma,\lambda) + \sum_{(\mathcal{A},\sigma,\lambda)\in\mathcal{R}} m(\mathcal{A},\sigma,\lambda)\,.$$

We show that the rightmost sum vanishes in characteristic 2. To this end, it suffices to construct a fixed-point-free involution $\phi : \mathcal{R} \to \mathcal{R}$ such that $m(\phi(\mathcal{A},\sigma,\lambda)) = m(\mathcal{A},\sigma,\lambda)$ holds for all $(\mathcal{A},\sigma,\lambda) \in \mathcal{R}$. Consider an arbitrary $(\mathcal{A},\sigma,\lambda) \in \mathcal{R}$. Since $\mathcal{A}$ is a p-prepacking but not a packing, there is a minimum (with respect to e.g. lexicographic order) three-tuple $(u_0, A_1, A_2) \in (U_3 \cup U_4 \cup \cdots \cup U_q) \times \mathcal{A} \times \mathcal{A}$ such that $u_0 \in A_1 \cap A_2$ and $A_1 \neq A_2$. Define a labeling $\lambda' : d(\mathcal{A}) \to L$ of $\mathcal{A}$ by setting, for each $(u, A) \in d(\mathcal{A})$,

$$\lambda'(u,A) = \begin{cases} \lambda(u,A) & \text{if } u \neq u_0 \text{ or } A \notin \{A_1, A_2\}; \\ \lambda(u_0,A_2) & \text{if } u = u_0 \text{ and } A = A_1; \text{ and} \\ \lambda(u_0,A_1) & \text{if } u = u_0 \text{ and } A = A_2. \end{cases} \tag{11}$$

Note that λ' is bijective and that $\lambda' \neq \lambda$. From (18) and (11) it follows that $m(\mathcal{A},\sigma,\lambda') = m(\mathcal{A},\sigma,\lambda)$ holds for all $(\mathcal{A},\sigma,\lambda) \in \mathcal{R}$. We can now set $\phi(\mathcal{A},\sigma,\lambda) = (\mathcal{A},\sigma,\lambda')$ and observe that $\phi(\mathcal{A},\sigma,\lambda) \in \mathcal{R}$, $\phi(\mathcal{A},\sigma,\lambda) \neq (\mathcal{A},\sigma,\lambda)$, and $\phi^2(\mathcal{A},\sigma,\lambda) = (\mathcal{A},\sigma,\lambda)$ for all $(\mathcal{A},\sigma,\lambda) \in \mathcal{R}$. Thus, ϕ is a fixed-point-free involution on $\mathcal{R}$.

3.5. The algorithm

From (10) and Section 3.4 we have

$$\sum_{(\mathcal{A},\sigma,\lambda)\in\mathcal{P}} m(\mathcal{A},\sigma,\lambda) = \sum_{J\subseteq L}(-1)^{|J|} \sum_{(\mathcal{A},\sigma,\lambda)\in\mathcal{L}[J]} m(\mathcal{A},\sigma,\lambda)\,. \tag{12}$$

From (9) we observe that the left-hand side of (12) is a multivariate polynomial of degree at most $pq + r$. It follows from Lemma 10 that the polynomial is not identically zero if and only if $\mathcal{F}$ contains a p-packing.

It remains to evaluate (12) for an assignment of values to the indeterminates. Let $J \subseteq L$ be fixed. We will show that the inner sum in (12) is obtained as the determinant of an $r \times r$ matrix $E(J)$ defined as follows. Index the rows with U_1 and the columns with U_2. Define the entry at row $u_1 \in U_1$, column $u_2 \in U_2$ by

$$e_{u_1,u_2}(J) = y_{u_1,u_2}\left(1 + w \sum_{A \in \mathcal{F} : \{u_1,u_2\} \subseteq A} x_A \prod_{u \in A \setminus \{u_1,u_2\}} \sum_{\ell \in L \setminus J} z_{u,\ell}\right).$$

From (7) it follows that

$$\det E(J) = \sum_{\substack{\sigma : U_1 \to U_2 \\ \sigma \text{ bijective}}} \operatorname{sgn}_\tau(\sigma) \prod_{u_1 \in U_1} e_{u_1,\sigma(u_1)}(J)\,.$$

After expanding the product into a sum of monomials, we observe the sum contains exactly one monomial

$$m(\mathcal{A}, \sigma, \lambda) = \operatorname{sgn}_\tau(\sigma) w^j \prod_{A \in \mathcal{A}} x_A \prod_{u_1 \in U_1} y_{u_1,\sigma(u_1)} \prod_{(u,A) \in d(\mathcal{A})} z_{u,\lambda(u,A)}\,.$$

for each j-prepacking $(\mathcal{A}, \sigma, \lambda)$ whose labeling avoids each label in J and $j = 0, 1, \ldots, r$. Thus

$$\det E(J) = \sum_{(\mathcal{A},\sigma,\lambda) \in \mathcal{L}[J]} m(\mathcal{A}, \sigma, \lambda)\,. \tag{13}$$

Consequently, for any given assignment of values in $\mathbb{F}_{2^b}$ to the indeterminates x_A, y_{u_1,u_2}, and $z_{u,\ell}$, we can evaluate the left-hand side of (13) *as a polynomial in the indeterminate* w via (13). Taking the sum over all $J \subseteq L$ and extracting the coefficient of the monomial w^p, we obtain an evaluation of the right-hand side of (12) in total $O(2^{p(q-2)}|\mathcal{F}|pq^2r^4)$ arithmetic operations in $\mathbb{F}_{2^b}$. Taking $b = \lceil \log_2 2(pq + r) \rceil$, Lemma 9 implies that we witness a p-packing in $\mathcal{F}$ (as a nonzero evaluation of the left-hand side of (12)) with probability at least $1/2$ in time $O^*(2^{p(q-2)})$ for polynomial size families $\mathcal{F}$.

4. A Projection–Determinant Sieve for p-Packings of q-Sets

This section establishes Theorems 2 and 4.

4.1. Tutte's observation

Let us recall that an involution is a permutation that is identical to its inverse. In particular, the cycle decomposition of an involution consists of fixed points and transpositions (cycles of length 2). It follows that the involutions on a set I are in a one-to-one correspondence with the set partitions of I into sets of cardinality 1 and 2. The following lemma is essentially due to Tutte [46].

Lemma 11 (Tutte's Determinant–Partition Lemma)**.** *Let T be an $m \times m$ matrix with entries in a multivariate polynomial ring over a field of characteristic 2. Index the rows and columns of T by the elements of a set I and suppose that T is symmetric so that*

$$t_{ij} = \begin{cases} \sum_k s^2_{\{i\},k} & \text{if } i = j; \\ \sum_k s_{\{i,j\},k} & \text{if } i \neq j \end{cases}$$

holds for all $i, j \in I$, with some monomials $s_{\{i,j\},k}$. Then,

$$\det T \;=\; \sum_{\substack{\iota : I \to I \\ \iota \text{ involution}}} \prod_{\substack{i \in I \\ i \leq \iota(i)}} \sum_k s^2_{\{i,\iota(i)\},k}\,. \tag{14}$$

Proof. Denote the set of all permutations of I by P_I. Observe that a permutation $\sigma \in P_I$ is not an involution if and only if the cycle decomposition of σ contains a cycle of length at least 3. Suppose that $\nu \in P_I$ is a permutation that is not an involution. Introduce an arbitrary total order on I, and order the cycles of length at least 3 in ν based on the least point in I moved by each such cycle. Denote by ν' the permutation obtained from ν by inverting the first cycle of length at least 3 in ν. Clearly, $\nu' \neq \nu$ and $(\nu')' = \nu$. Now observe that because T is a symmetric matrix, for a cyclic permutation $(i_1\ i_2\ \cdots\ i_j)$ and its inverse $(i_j\ i_{j-1}\ \cdots\ i_1)$, we have

$$t_{i_1,i_2} t_{i_2,i_3} \cdots t_{i_{j-1},i_j} t_{i_j,i_1} = t_{i_j,i_{j-1}} t_{i_{j-1},i_{j-2}} \cdots t_{i_2,i_1} t_{i_1,i_j}$$

It follows that

$$\prod_{i \in I} t_{i,\nu(i)} = \prod_{i \in I} t_{i,\nu'(i)}.$$

Partition P_I into $P_I = Q_I \cup R_I$, where Q_I consists of the involutions and R_I consists of the non-involutions. Introduce an arbitrary total order on R_I.

Because the determinant of T is equal to the permanent of T in characteristic 2, we have

$$\begin{aligned}
\det T &= \sum_{\sigma \in P_I} \prod_{i \in I} t_{i,\sigma(i)} \\
&= \sum_{\iota \in Q_I} \prod_{i \in I} t_{i,\iota(i)} + \sum_{\nu \in R_I} \prod_{i \in I} t_{i,\nu(i)} \\
&= \sum_{\iota \in Q_I} \prod_{i \in I} t_{i,\iota(i)} + \sum_{\substack{\nu \in R_I \\ \nu < \nu'}} \left(\prod_{i \in I} t_{i,\nu(i)} + \prod_{i \in I} t_{i,\nu'(i)} \right) \\
&= \sum_{\iota \in Q_I} \prod_{i \in I} t_{i,\iota(i)} + \sum_{\substack{\nu \in R_I \\ \nu < \nu'}} 2 \prod_{i \in I} t_{i,\nu(i)} \\
&= \sum_{\iota \in Q_I} \prod_{i \in I} t_{i,\iota(i)} \,.
\end{aligned}$$

Thus, splitting the product over fixed and moved points of ι, and using the symmetry of T, we have

$$\begin{aligned}
\det T &= \sum_{\iota \in Q_I} \prod_{\substack{i \in I \\ i = \iota(i)}} t_{i,\iota(i)} \prod_{\substack{i \in I \\ i \neq \iota(i)}} t_{i,\iota(i)} \\
&= \sum_{\iota \in Q_I} \prod_{\substack{i \in I \\ i = \iota(i)}} \sum_k s^2_{\{i\},k} \prod_{\substack{i \in I \\ i < \iota(i)}} \left(\sum_k s_{\{i,\iota(i)\},k} \right)^2 \\
&= \sum_{\iota \in Q_I} \prod_{\substack{i \in I \\ i = \iota(i)}} \sum_k s^2_{\{i\},k} \prod_{\substack{i \in I \\ i < \iota(i)}} \left(\sum_k s^2_{\{i,\iota(i)\},k} + 2 \sum_{k < k'} s_{\{i,\iota(i)\},k} s_{\{i,\iota(i)\},k'} \right) \\
&= \sum_{\iota \in Q_I} \prod_{\substack{i \in I \\ i = \iota(i)}} \sum_k s^2_{\{i\},k} \prod_{\substack{i \in I \\ i < \iota(i)}} \sum_k s^2_{\{i,\iota(i)\},k} \\
&= \sum_{\iota \in Q_I} \prod_{\substack{i \in I \\ i \leq \iota(i)}} \sum_k s^2_{\{i,\iota(i)\},k} \,.
\end{aligned}$$

The claim follows. □

Our strategy is to leverage Tutte's observation with random projection and sieving. In particular, we witness a p-packing by randomly projecting it

to a set $U_1 \subseteq U$ where Tutte's observation forces the packing constraint with positive probability, which allows us to restrict sieving to the complementary projection into $U_2 = U \setminus U_1$.

4.2. Admissible packings and prepackings

Let $\mathcal{F}$ be a set of q-subsets of an n-element universe U. Partition U into two disjoint sets $U = U_1 \cup U_2$ with $|U_1| = n_1$ and $|U_2| = n_2 = n - n_1$. We say that such an ordered partition (U_1, U_2) of U is an (n_1, n_2)-*partition*.

A subset $\mathcal{A} \subseteq \mathcal{F}$ is a p-*packing* if $|\mathcal{A}| = p$ and the sets in $\mathcal{A}$ are pairwise disjoint. We say that $\mathcal{A}$ is *admissible* if every set $A \in \mathcal{A}$ satisfies $|A \cap U_1| \leq 2$. We say that $\mathcal{A}$ is a (p_0, p_1, p_2)-*prepacking* if

(a) $|\mathcal{A}| = p_0 + p_1 + p_2$;

(b) $|\{A \in \mathcal{A} : |A \cap U_1| = j\}| = p_j$ for each $j = 0, 1, 2$; and

(c) the sets in $\mathcal{A}$ are pairwise disjoint when projected to U_1.

Note that a prepacking is by definition admissible. Also, every admissible p-packing is a (p_0, p_1, p_2)-prepacking for some parameters $p_0 + p_1 + p_2 = p$. In this case we say that the p-packing is a (p_0, p_1, p_2)-*packing*.

Let us say that a (p_0, p_1, p_2)-prepacking $\mathcal{A}$ is *compatible* with an involution $\iota : U_1 \to U_1$ if for every $A \in \mathcal{A}$ it holds that

(a) ι fixes the point in $A \cap U_1$ if $|A \cap U_1| = 1$; and

(b) ι transposes the two points in $A \cap U_1$ if $|A \cap U_1| = 2$.

Note that every prepacking is compatible with at least one involution.

4.3. Random projection

We analyze the probability that a given p-packing projects under a random (U_1, U_2) into a (p_0, p_1, p_2)-packing.

Lemma 12 (Admissibility)**.** *Let $\mathcal{A}$ be a p-packing. For an (n_1, n_2)-partition (U_1, U_2) of U selected uniformly at random, we have*

$$\begin{aligned}&\Pr\big(\mathcal{A} \text{ is a } (p_0, p_1, p_2)\text{-packing}\big) = \\ &\quad = \binom{p}{p_1 + p_2}\binom{p_1 + p_2}{p_2}\binom{q}{1}^{p_1}\binom{q}{2}^{p_2}\binom{n - pq}{n_1 - p_1 - 2p_2}\binom{n}{n_1}^{-1}. \qquad (15)\end{aligned}$$

Proof. Among the p pairwise disjoint sets in $\mathcal{A}$, there are $\binom{p}{p_1+p_2}$ ways to select the sets that intersect U_1 in 1 or 2 points, and $\binom{p_1+p_2}{p_2}$ ways to select among these the p_2 sets that intersect in 2 points. There are $\binom{q}{1}^{p_1}\binom{q}{2}^{p_2}$ possible intersection patterns with U_1 in these selected p_1+p_2 sets. There are $\binom{n-pq}{n_1-p_1-2p_2}$ ways to select the remaining $n_1-p_1-2p_2$ points of U_1 outside the pq points of $\mathcal{A}$. □

Remark. A nonzero probability is allocated if and only if $pq \le n$, $p_0 + p_1 + p_2 = p$, $p_1 + 2p_2 \le n_1$, and $n_1 - p_1 - 2p_2 \le n - pq$.

Using techniques similar to Section 2.4, let us derive an asymptotic approximation for (15). We begin by approximating the last two terms in (15). To this end, observe that for all $a, b, \delta > 0$ we have

$$\left\langle \begin{matrix} a-b \\ \delta a - \delta b \end{matrix} \right\rangle = \left[\delta^\delta (1-\delta)^{1-\delta}\right]^{-(a-b)} = \left\langle \begin{matrix} a \\ \delta a \end{matrix} \right\rangle \left\langle \begin{matrix} b \\ \delta b \end{matrix} \right\rangle^{-1}.$$

Setting $\delta = (p_1 + 2p_2)/(pq)$ and $n_1 = \lfloor \delta n \rfloor$ we obtain

$$\left\langle \begin{matrix} n-pq \\ \delta n - p_1 - 2p_2 \end{matrix} \right\rangle \left\langle \begin{matrix} n \\ \delta n \end{matrix} \right\rangle^{-1} = \left\langle \begin{matrix} pq \\ p_1 + 2p_2 \end{matrix} \right\rangle^{-1}.$$

Thus, uniformly for all $0 < p_0, p_1, p_2 < p$ with $p_0 + p_1 + p_2 = p$, we have

$$\begin{aligned} &\Pr\big(\mathcal{A} \text{ is a } (p_0,p_1,p_2)\text{-packing}\big) = \\ &\quad = \Theta^*\left(\left\langle \begin{matrix} p \\ p_1+p_2 \end{matrix} \right\rangle \left\langle \begin{matrix} p_1+p_2 \\ p_2 \end{matrix} \right\rangle \binom{q}{1}^{p_1} \binom{q}{2}^{p_2} \left\langle \begin{matrix} pq \\ p_1+2p_2 \end{matrix} \right\rangle^{-1}\right). \end{aligned} \tag{16}$$

4.4. Fingerprinting and identifiability

Let $\mathcal{A} \subseteq \mathcal{F}$ be a (p_0, p_1, p_2)-prepacking. The *domain* of the prepacking is the set

$$d(\mathcal{A}) = \{(u, A) : u \in A \in \mathcal{A}\} \subseteq U_2 \times \mathcal{F}. \tag{17}$$

Observe that $|d(\mathcal{A})| = qp_0 + (q-1)p_1 + (q-2)p_2$.

Let L be a set of $qp_0 + (q-1)p_1 + (q-2)p_2$ labels. A *labeling* of $\mathcal{A}$ is a pair (ι, λ), where $\iota : U_1 \to U_1$ is an involution compatible with $\mathcal{A}$ and $\lambda : d(\mathcal{A}) \to L$ is an arbitrary mapping. The labeling is *bijective* if λ is a bijection. We say that a triple $(\mathcal{A}, \iota, \lambda)$ is a *labeled* (p_0, p_1, p_2)-prepacking.

The sieve operates over a multivariate polynomial ring with the coefficient field $\mathbb{F}_{2^b}$ and the following indeterminates. Introduce the indeterminates

w_0, w_1, and w_2 for tracking the parameters p_0, p_1, p_2 of $\mathcal{A}$. Associate with each $A \in \mathcal{F}$ an indeterminate x_A. Associate with each set $K \subseteq U_1$ of size $1 \leq |K| \leq 2$ an indeterminate y_K. Associate with each pair $(u, \ell) \in U_2 \times L$ an indeterminate $z_{u,\ell}$.

The *monomial* of a labeled (p_0, p_1, p_2)-prepacking $(\mathcal{A}, \iota, \lambda)$ is

$$m(\mathcal{A}, \iota, \lambda) = w_0^{2p_0} w_1^{2p_1} w_2^{2p_2} \prod_{A \in \mathcal{A}} x_A^2 \prod_{\substack{i \in U_1 \\ i \leq \iota(i)}} y_{\{i,\iota(i)\}}^2 \prod_{(u,A) \in d(\mathcal{A})} z_{u,\lambda(u,A)}^2 \,. \tag{18}$$

Lemma 13 (Identifiability)**.** *The monomial $m(\mathcal{A}, \iota, \lambda)$ uniquely determines both $\mathcal{A}$ and ι. Furthermore, if $\mathcal{A}$ is a p-packing and λ is bijective, then $m(\mathcal{A}, \iota, \lambda)$ uniquely determines λ.*

Proof. The set family $\mathcal{A}$ can be read from the indeterminates x_A and the involution ι can be read from the 1- and 2-sets encoded by the indeterminates $y_{\{i,\iota(i)\}}$. For the second claim, let $(u, A) \in d(\mathcal{A})$. Observe that $\lambda(u, A) = \ell$ if the monomial includes the indeterminate $z_{u,\ell}$, for the set A is uniquely determined by the element u under the assumption that $\mathcal{A}$ is a p-packing. □

4.5. Sieving for bijective labelings

Denote by $\mathcal{L}_{p_0,p_1,p_2}$ the set of all labeled (p_0, p_1, p_2)-prepackings. For $J \subseteq L$, denote by $\mathcal{L}_{p_0,p_1,p_2}[J]$ the subset of labeled (p_0, p_1, p_2)-prepackings whose labeling *avoids* each label in J. Denote by $\mathcal{B}_{p_0,p_1,p_2}$ the set of all bijectively labeled (p_0, p_1, p_2)-prepackings.

By the principle of inclusion-exclusion,

$$\sum_{(\mathcal{A},\iota,\lambda) \in \mathcal{B}_{p_0,p_1,p_2}} m(\mathcal{A}, \iota, \lambda) = \sum_{J \subseteq L} (-1)^{|J|} \sum_{(\mathcal{A},\iota,\lambda) \in \mathcal{L}_{p_0,p_1,p_2}[J]} m(\mathcal{A}, \iota, \lambda) \,. \tag{19}$$

4.6. Fingerprints of bijectively labeled non-(p_0, p_1, p_2)-packings cancel

Let $\mathcal{B}_{p_0,p_1,p_2}$ be the set of bijectively labeled (p_0, p_1, p_2)-prepackings. Partition $\mathcal{B}_{p_0,p_1,p_2}$ into $\mathcal{B}_{p_0,p_1,p_2} = \mathcal{P}_{p_0,p_1,p_2} \cup \mathcal{R}_{p_0,p_1,p_2}$, where $\mathcal{P}_{p_0,p_1,p_2}$ is the set of bijectively labeled (p_0, p_1, p_2)-packings, and $\mathcal{R}_{p_0,p_1,p_2}$ is the set of bijectively labeled (p_0, p_1, p_2)-prepackings that are not (p_0, p_1, p_2)-packings. Accordingly, we have

$$\sum_{(\mathcal{A},\iota,\lambda) \in \mathcal{B}_{p_0,p_1,p_2}} m(\mathcal{A}, \iota, \lambda) = \sum_{(\mathcal{A},\iota,\lambda) \in \mathcal{P}_{p_0,p_1,p_2}} m(\mathcal{A}, \iota, \lambda) + \sum_{(\mathcal{A},\iota,\lambda) \in \mathcal{R}_{p_0,p_1,p_2}} m(\mathcal{A}, \iota, \lambda) \,.$$

By a pairing argument essentially identical to the one given in Section 3.4, the rightmost sum vanishes in characteristic 2.

4.7. The algorithm

Let us assume that the parameters $0 < p_0, p_1, p_2 < p$ and n_1, n_2 have been fixed so that any given p-packing is a (p_0, p_1, p_2)-packing with positive probability. (We will set the precise values in what follows.) The algorithm repeats the following randomized procedure.

First, the procedure selects an ordered (n_1, n_2)-partition (U_1, U_2) uniformly at random among all the $\binom{n}{n_1}$ such partitions.

Next, the procedure evaluates the following generating function for a random assignment of values to the indeterminates. From (19) and Section 4.6 we have

$$\sum_{(\mathcal{A},\iota,\lambda)\in\mathcal{P}_{p_0,p_1,p_2}} m(\mathcal{A},\iota,\lambda) = \sum_{J\subseteq L}(-1)^{|J|} \sum_{(\mathcal{A},\iota,\lambda)\in\mathcal{L}_{p_0,p_1,p_2}[J]} m(\mathcal{A},\iota,\lambda)\,. \tag{20}$$

The left-hand side of (20) is a multivariate polynomial of degree at most $2n_1 + (2q+4)p_0 + (2q+3)p_1 + (2q+2)p_2$. It follows from Lemma 13 that the polynomial is not identically zero if and only if $\mathcal{F}$ contains a (p_0, p_1, p_2)-packing.

It remains to evaluate the right-hand side of (20). Let $J \subseteq L$ be fixed. The procedure relies on the following observation that a sum over labeled prepackings factors into a product of two independent expressions. The first expression, $S(J)$, generates the sets that do not intersect U_1 with a simple product. The second expression, $\det T(J)$, generates the sets that intersect U_1 with Lemma 11.

In precise terms, let

$$S(J) = \prod_{\substack{A\in\mathcal{F}\\ A\cap U_1=\emptyset}} \left(1 + w_0^2 x_A^2 \prod_{u_2\in A} \sum_{\ell\in L\setminus J} z_{u_2,\ell}^2\right).$$

Expanding the products into sums yields a polynomial that consists of exactly one monomial

$$w_0^{2p_0} \prod_{A\in\mathcal{A}} x_A^2 \prod_{(u,A)\in d(\mathcal{A})} z_{u,\lambda(u,A)}^2$$

for each p_0-packing $\mathcal{A}$ whose members are disjoint from U_1, along with a function $\lambda : d(\mathcal{A}) \to L \setminus J$.

Next define the symmetric $n_1 \times n_1$ matrix $T(J)$ as follows. Index the rows and columns by elements of U_1. For $u_1 \in U_1$, define the diagonal entries by

$$t_{u_1,u_1}(J) = y_{\{u_1\}}^2 \left(1 + w_1^2 \sum_{\substack{A \in \mathcal{F} \\ A \cap U_1 = \{u_1\}}} x_A^2 \prod_{u_2 \in A \cap U_2} \sum_{\ell \in L \setminus J} z_{u_2,\ell}^2\right).$$

For $u_1, v_1 \in U_1$ with $u_1 \neq v_1$, define the off-diagonal entries by

$$t_{u_1,v_1}(J) = y_{\{u_1,v_1\}} \left(1 + w_2 \sum_{\substack{A \in \mathcal{F} \\ A \cap U_1 = \{u_1,v_1\}}} x_A \prod_{u_2 \in A \cap U_2} \sum_{\ell \in L \setminus J} z_{u_2,\ell}\right).$$

Using Lemma 11 we get

$$\det T(J) = \sum_{\substack{\iota : U_1 \to U_1 \\ \iota \text{ involution}}} \prod_{\substack{i \in U_1 \\ i \leq \iota(i)}} y_{\{i,\iota(i)\}}^2 \left(1 + w_{|\{i,\iota(i)\}|}^2 \sum_{\substack{A \in \mathcal{F} \\ A \cap U_1 = \{i,\iota(i)\}}} x_A^2 \prod_{u \in A \cap U_2} \sum_{\ell \in L \setminus J} z_{u,\ell}^2\right).$$

Expanding the products into sums yields a polynomial that consists of exactly one monomial

$$w_1^{2p_1} w_2^{2p_2} \prod_{A \in \mathcal{A}'} x_A^2 \prod_{\substack{i \in U_1 \\ i \leq \iota(i)}} y_{\{i,\iota(i)\}}^2 \prod_{(u,A) \in d(\mathcal{A}')} z_{u,\lambda'(u,A)}^2$$

for each labeled (p_0, p_1, p_2)-prepacking $(\mathcal{A}', \iota, \lambda')$, with $p_0 = 0$ and $0 \leq p_1 + 2p_2 \leq n_1$, whose labeling avoids each label in J.

Now, observe that multiplying a monomial of $S(J)$ by a monomial of $\det T(J)$ results in a monomial $m(\mathcal{A} \cup \mathcal{A}', \iota, \lambda \cup \lambda')$ of a labeled (p_0, p_1, p_2)-prepacking $(\mathcal{A}', \iota, \lambda')$ whose labeling avoids each label in J and, vice versa, the monomial of any such labeled prepacking factorizes uniquely into a monomial of $S(J)$ and a monomial of $\det T(J)$. Thus we have

$$\sum_{0 \leq p_0 \leq |\mathcal{F}|} \sum_{0 \leq p_1 + 2p_2 \leq n_1} \sum_{(\mathcal{A},\iota,\lambda) \in \mathcal{L}_{p_0,p_1,p_2}[J]} m(\mathcal{A}, \iota, \lambda) = S(J) \det T(J). \tag{21}$$

Consequently, for any given assignment of values in $\mathbb{F}_{2^b}$ to the indeterminates x_A, y_K, and $z_{u,\ell}$, the procedure evaluates the left-hand side of (21) *as a polynomial in the indeterminates* w_0, w_1, w_2 using a total of $O(|\mathcal{F}|^6 q |L| n_1^3)$ arithmetic operations in $\mathbb{F}_{2^b}$. From such an evaluation we can recover the

coefficient of the monomial $w_0^{p_0} w_1^{p_1} w_2^{p_2}$. This coefficient corresponds to an evaluation of the inner sum in the right-hand side of (20). Taking the sum over $J \subseteq L$ (and multiplying by $w_0^{p_0} w_1^{p_1} w_2^{p_2}$), we obtain an evaluation of the right-hand side of (20).

Denoting the probability that a p-packing $\mathcal{A}$ is a (p_0, p_1, p_2)-packing with $P(n, n_1, p_0, p_1, p_2)$, and taking $r = \lceil 1/P(n, n_1, p_0, p_1, p_2) \rceil$ repetitions of the procedure with $b = \lceil \log_2 16n \rceil$, Lemma 9 implies that at least one repetition of the procedure witnesses any fixed p-packing $\mathcal{A}$ (as a nonzero evaluation of (20)) with probability at least $(1 - e^{-1})/2$ in time

$$O\big(2^{qp_0+(q-1)p_1+(q-2)p_2} |\mathcal{F}|^6 pq^2 n_1 b^2 / P(n, n_1, p_0, p_1, p_2)\big).$$

Setting $n_1 = \lfloor \delta n \rfloor$, $p_1 = \lfloor \beta_1 p \rfloor$, $p_2 = \lfloor \beta_2 p \rfloor$, and $p_0 = p - p_1 - p_2$, we obtain from (16) the running time

$$O^*\left(\left(\frac{2^{q(1-\beta_1-\beta_2)+(q-1)\beta_1+(q-2)\beta_2} \left\langle {q \atop \beta_1+2\beta_2} \right\rangle}{\left\langle {1 \atop \beta_1+\beta_2} \right\rangle \left\langle {\beta_1+\beta_2 \atop \beta_2} \right\rangle \binom{q}{1}^{\beta_1} \binom{q}{2}^{\beta_2}}\right)^p\right)$$

for polynomial size families $\mathcal{F}$. In particular, we obtain time $O^*(3.3432^p)$ for $q = 3$ with $\beta_1 = 0.281509$ and $\beta_2 = 0.679622$, time $O^*(7.2562^p)$ for $q = 4$ with $\beta_1 = 0.323262$ and $\beta_2 = 0.612790$, time $O^*(15.072^p)$ for $q = 5$ with $\beta_1 = 0.338614$ and $\beta_2 = 0.582673$. We found these parameter values by numerical computations; note that any values results in a valid upper bound for the running time.

5. A Determinant Sieve for Edge-Coloring

This section establishes Theorem 6.

5.1. Tutte's observation revisited

The edges of G can be colored with d colors if and only if there exists a set of $d - 1$ pairwise edge-disjoint perfect matchings in G. Indeed, because the graph is d-regular, each color class must be a perfect matching.

Let us now return to Lemma 11. We observe that (14) in effect gives us a multivariate generating function for the perfect matchings in G. Our strategy is to introduce $d - 1$ independent copies of this generating function and sieve for edge-disjointness.

5.2. Fingerprinting and identifiability

Let $\vec{M} = (M_1, M_2, \ldots, M_p)$ be an ordered p-tuple of perfect matchings in G. The *domain* of $\vec{M}$ is the set

$$d(\vec{M}) = \{(e, i) : e \in M_i\} \subseteq E \times \{1, 2, \ldots, p\}\,. \tag{22}$$

Observe that $|d(\vec{M})| = pn/2$. Let L be a set of $pn/2$ labels. A *labeling* of $\vec{M}$ is a mapping $\lambda : d(\vec{M}) \to L$. The labeling is *bijective* if λ is a bijection.

The sieve operates over a multivariate polynomial ring with the coefficient field $\mathbb{F}_{2^b}$ and the following indeterminates. Associate with each pair $(e, i) \in E \times \{1, 2, \ldots, p\}$ an indeterminate $x_{e,i}$. Associate with each pair $(e, \ell) \in E \times L$ an indeterminate $y_{e,\ell}$.

The *monomial* of a labeled p-tuple $(\vec{M}, \lambda)$ is

$$m(\vec{M}, \lambda) = \prod_{(e,i) \in d(\vec{M})} x_{e,i}^2 y_{e,\lambda(e,i)}^2\,. \tag{23}$$

Lemma 14 (Identifiability)**.** *The monomial $m(\vec{M}, \lambda)$ uniquely determines $\vec{M}$. Furthermore, if $\vec{M}$ consists of pairwise edge-disjoint perfect matchings and λ is bijective, then $m(\vec{M}, \lambda)$ uniquely determines λ.*

5.3. Sieving for bijective labelings

Denote by $\mathcal{L}$ the set of all labeled p-tuples of perfect matchings of G. For $J \subseteq L$, denote by $\mathcal{L}[J]$ the subset of labeled p-tuples of perfect matchings of G whose labeling *avoids* each label in J. Denote by $\mathcal{B}$ the set of all bijectively labeled p-tuples of perfect matchings of G.

By the principle of inclusion-exclusion,

$$\sum_{(\vec{M},\lambda) \in \mathcal{B}} m(\vec{M}, \lambda) = \sum_{J \subseteq L} (-1)^{|J|} \sum_{(\vec{M},\lambda) \in \mathcal{L}[J]} m(\vec{M}, \lambda)\,. \tag{24}$$

5.4. Fingerprints of bijectively labeled non-disjoint p-tuples cancel

Let $\mathcal{B}$ be the set of all bijectively labeled p-tuples of perfect matchings of G. Partition $\mathcal{B}$ into $\mathcal{B} = \mathcal{P} \cup \mathcal{R}$, where $\mathcal{P}$ is the set of bijectively labeled p-tuples of perfect matchings that are pairwise edge-disjoint, and $\mathcal{R}$ is the set of bijectively labeled p-tuples of perfect matchings for which there exists at

least one edge that occurs in at least two matchings in the tuple. Accordingly, we have

$$\sum_{(\vec{M},\lambda)\in\mathcal{B}} m(\mathcal{A},\iota,\lambda) = \sum_{(\vec{M},\lambda)\in\mathcal{P}} m(\vec{M},\lambda) + \sum_{(\vec{M},\lambda)\in\mathcal{R}} m(\vec{M},\lambda)\,.$$

By a pairing argument essentially identical to the one given in Section 3.4, the rightmost sum vanishes in characteristic 2.

5.5. The algorithm

First, the procedure evaluates the following generating function for a random assignment of values to the indeterminates. From (24) and Section 5.4 we have

$$\sum_{(\vec{M},\lambda)\in\mathcal{P}} m(\vec{M},\lambda) = \sum_{J\subseteq L}(-1)^{|J|} \sum_{(\vec{M},\lambda)\in\mathcal{L}[J]} m(\vec{M},\lambda)\,. \tag{25}$$

The left-hand side of (25) is a multivariate polynomial of degree at most $2pn$. It follows from Lemma 14 that the polynomial is not identically zero if and only if G has a set of p pairwise edge-disjoint perfect matchings.

It remains to evaluate the right-hand side of (25). Let $J \subseteq L$ be fixed. The procedure relies on Tutte's Lemma (Lemma 11). For $i = 1, 2, \ldots, p$ define the symmetric $n \times n$ matrix $T^{(i)}(J)$ as follows. Index the rows and columns by the vertices V of G. Define the entries of $T^{(i)}(J)$ for all $u, v \in V$ by

$$t^{(i)}_{u,v}(J) = \begin{cases} 0 & \text{if } u = v \text{ or } \{u,v\} \notin E; \\ x_{\{u,v\},i} \sum_{\ell\in L\setminus J} y_{\{u,v\},\ell} & \text{if } \{u,v\} \in E. \end{cases} \tag{26}$$

From Lemma 11 we have

$$\sum_{(\vec{M},\lambda)\in\mathcal{L}[J]} m(\vec{M},\lambda) = \prod_{i=1}^{p} \det T^{(i)}(J)\,. \tag{27}$$

Consequently, for any given assignment of values in $\mathbb{F}_{2^b}$ to the indeterminates $x_{e,i}$ and $y_{e,\ell}$, the procedure evaluates the left-hand side of (27) using a total of $O(pn^3)$ arithmetic operations in $\mathbb{F}_{2^b}$. Taking the sum over $J \subseteq L$, we obtain an evaluation of the right-hand side of (25). Taking $b = \lceil \log_2 4pn \rceil$, we witness a set of p pairwise edge-disjoint perfect matchings in G as a nonzero evaluation of the left-hand side of (25) with probability $\Omega(1)$ in time $O^*(2^{pn/2})$ and space polynomial in n. Taking $p = d - 1$, we obtain a polynomial-space randomized algorithm for deciding whether a d-regular graph admits a coloring of its edges with d colors in $O^*(2^{(d-1)n/2})$ time.

5.6. Graphs that are not regular

Let $m = |E|$. We can modify the previous algorithm to run in time $O^*(2^m)$ and space polynomial in n on graphs that are not regular. In particular, instead of perfect matchings consider matchings, set $|L| = m$, in (26) set the diagonal entries equal to 1, and set $p = \Delta$, where Δ is the maximum degree of a vertex in G.

References

[1] A. Björklund, Determinant sums for undirected Hamiltonicity, SIAM J. Comput. 43 (2014) 280–299.

[2] B. Monien, How to find long paths efficiently, Ann. Discrete Math. 25 (1985) 239–254.

[3] H. L. Bodlaender, On linear time minor tests with depth-first search, J. Algorithms 14 (1993) 1–23.

[4] C. H. Papadimitriou, M. Yannakakis, On limited nondeterminism and the complexity of the V-C dimension, J. Comput. Syst. Sci. 53 (1996) 161–170.

[5] N. Alon, R. Yuster, U. Zwick, Color-coding, J. ACM 42 (1995) 844–856.

[6] J. Chen, J. Kneis, S. Lu, D. Mlle, S. Richter, P. Rossmanith, S.-H. Sze, F. Zhang, Randomized divide-and-conquer: Improved path, matching, and packing algorithms, SIAM J. Comput. 38 (2009) 2526–2547.

[7] F. V. Fomin, D. Lokshtanov, F. Panolan, S. Saurabh, Efficient computation of representative families with applications in parameterized and exact algorithms, J. ACM 63 (2016) 29:1–29:60.

[8] I. Koutis, Faster algebraic algorithms for path and packing problems, in: L. Aceto, I. Damgård, L. A. Goldberg, M. M. Halldórsson, A. Ingólfsdóttir, I. Walukiewicz (Eds.), ICALP (1), volume 5125 of *Lecture Notes in Computer Science*, Springer, 2008, pp. 575–586.

[9] R. Williams, Finding paths of length k in $O^*(2^k)$ time, Inf. Process. Lett. 109 (2009) 315–318.

[10] J. Kneis, D. Mölle, S. Richter, P. Rossmanith, Divide-and-color, in: F. V. Fomin (Ed.), WG, volume 4271 of *Lecture Notes in Computer Science*, Springer, 2006, pp. 58–67.

[11] R. Impagliazzo, R. Paturi, F. Zane, Which problems have strongly exponential complexity?, J. Comput. Syst. Sci. 63 (2001) 512–530.

[12] A. Björklund, Exact covers via determinants, in: J.-Y. Marion, T. Schwentick (Eds.), STACS, volume 5 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010, pp. 95–106.

[13] R. G. Downey, M. R. Fellows, Parameterized Complexity, Springer, 1999.

[14] W. Jia, C. Zhang, J. Chen, An efficient parameterized algorithm for m-set packing, J. Algorithms 50 (2004) 106–117.

[15] I. Koutis, A faster parameterized algorithm for set packing, Inf. Process. Lett. 94 (2005) 7–9.

[16] M. Fellows, P. Heggernes, F. A. Rosamond, C. Sloper, J. A. Telle, Finding k disjoint triangles in an arbitrary graph, in: J. Hromkovic, M. Nagl, B. Westfechtel (Eds.), WG, volume 3353 of *Lecture Notes in Computer Science*, Springer, 2004, pp. 235–244.

[17] L. Mathieson, E. Prieto, P. Shaw, Packing edge disjoint triangles: A parameterized view, in: R. G. Downey, M. R. Fellows, F. K. H. A. Dehne (Eds.), IWPEC, volume 3162 of *Lecture Notes in Computer Science*, Springer, 2004, pp. 127–137.

[18] Y. Liu, S. Lu, J. Chen, S.-H. Sze, Greedy localization and color-coding: Improved matching and packing algorithms, in: H. L. Bodlaender, M. A. Langston (Eds.), IWPEC, volume 4169 of *Lecture Notes in Computer Science*, Springer, 2006, pp. 84–95.

[19] J. Wang, Q. Feng, An $O^*(3.523^k)$ parameterized algorithm for 3-set packing, in: M. Agrawal, D.-Z. Du, Z. Duan, A. Li (Eds.), TAMC, volume 4978 of *Lecture Notes in Computer Science*, Springer, 2008, pp. 82–93.

[20] E. Prieto, C. Sloper, Looking at the stars, Theor. Comput. Sci. 351 (2006) 437–445.

[21] J. Wang, D. Ning, Q. Feng, J. Chen, Improved parameterized algorithm for P_2-packing problem (in chinese), Journal of Software 19 (2008) 2879–2886.

[22] H. Fernau, D. Raible, A parameterized perspective on packing paths of length two, J. Comb. Optim. 18 (2009) 319–341.

[23] D. G. Kirkpatrick, P. Hell, On the completeness of a generalized matching problem, in: R. J. Lipton, W. A. Burkhard, W. J. Savitch, E. P. Friedman, A. V. Aho (Eds.), STOC, ACM, 1978, pp. 240–245.

[24] J. Flum, M. Grohe, Parameterized Complexity Theory, Texts in Theoretical Computer Science. An EATCS Series, Springer, 2006.

[25] J. Kleinberg, É. Tardos, Algorithm Design, Addison-Wesley, 2005.

[26] S. Dasgupta, C. H. Papadimitriou, U. Vazirani, Algorithms, McGraw-Hill, 2008.

[27] M. R. Fellows, C. Knauer, N. Nishimura, P. Ragde, F. A. Rosamond, U. Stege, D. M. Thilikos, S. Whitesides, Faster fixed-parameter tractable algorithms for matching and packing problems, Algorithmica 52 (2008) 167–176.

[28] R. G. Downey, M. R. Fellows, N. Koblitz, Techniques for exponential parameterized reductions in vertex set problems, 1996. Unpublished, reported in [13, §8.3].

[29] M. Zehavi, Mixing color coding-related techniques, in: N. Bansal, I. Finocchi (Eds.), Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings, volume 9294 of *Lecture Notes in Computer Science*, Springer, 2015, pp. 1037–1049.

[30] I. Koutis, R. Williams, Limits and applications of group algebras for parameterized problems, in: S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. E. Nikoletseas, W. Thomas (Eds.), ICALP (1), volume 5555 of *Lecture Notes in Computer Science*, Springer, 2009, pp. 653–664.

[31] V. G. Vizing, On an estimate of the chromatic class of a p-graph, Diskret. Analiz. 3 (1964) 25–30.

[32] A. Björklund, T. Husfeldt, M. Koivisto, Set partitioning via inclusion-exclusion, SIAM J. Comput. 39 (2009) 546–563.

[33] J. Couturier, P. A. Golovach, D. Kratsch, M. Liedloff, A. V. Pyatkin, Colorings with few colors: Counting, enumeration and combinatorial bounds, Theory Comput. Syst. 52 (2013) 645–667.

[34] I. Holyer, The NP-completeness of edge-coloring, SIAM J. Comput. 10 (1981) 718–720.

[35] D. Leven, Z. Galil, NP completeness of finding the chromatic index of regular graphs, J. Algorithms 4 (1983) 35–44.

[36] L. Kowalik, Edge colouring, in: F. D. et al. (Ed.), Open Problems: Moderately Exponential Time Algorithms, volume 08431 of *Dagstuhl Seminar Proceedings*, pp. 6–7.

[37] S. Kohn, A. Gottlieb, M. Kohn, A generating function approach to the traveling salesman problem, in: Proceedings of the 1977 Annual Conference, ACM '77, ACM, 1977, pp. 294–300.

[38] A. Björklund, T. Husfeldt, P. Kaski, M. Koivisto, Narrow sieves for parameterized paths and packings, CoRR abs/1007.1161 (2010).

[39] A. Björklund, V. Kamat, L. Kowalik, M. Zehavi, Spotting trees with few leaves, in: M. M. Halldórsson, K. Iwama, N. Kobayashi, B. Speckmann (Eds.), Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I, volume 9134 of *Lecture Notes in Computer Science*, Springer, 2015, pp. 243–255.

[40] A. Björklund, P. Kaski, L. Kowalik, Constrained multilinear detection and generalized graph motifs, Algorithmica 74 (2016) 947–967.

[41] I. Koutis, R. Williams, Algebraic fingerprints for faster algorithms, Commun. ACM 59 (2016) 98–105.

[42] H. Robbins, A remark on Stirling's formula, Amer. Math. Monthly 62 (1955) 26–29.

[43] R. A. DeMillo, R. J. Lipton, A probabilistic remark on algebraic program testing, Inf. Process. Lett. 7 (1978) 193–195.

[44] J. T. Schwartz, Fast probabilistic algorithms for verification of polynomial identities, J. ACM 27 (1980) 701–717.

[45] J. Edmonds, Systems of distinct representatives and linear algebra, J. Res. Nat. Bur. Standards Sect. B 71B (1967) 241–245.

[46] W. T. Tutte, The factorization of linear graphs, J. London Math. Soc. 22 (1947) 107–111.