A Tight Lower Bound for Counting Hamiltonian Cycles via Matrix Rank[*]

Radu Curticapean[†]   Nathan Lindzey[‡]   Jesper Nederlof[§]

Abstract

For even $k \in \mathbb{N}$, the matchings connectivity matrix $\mathbf{M}_k$ is a binary matrix indexed by perfect matchings on k vertices; the entry at (M, M') is 1 iff $M \cup M'$ forms a single cycle. Cygan et al. (STOC 2013) showed that the rank of $\mathbf{M}_k$ over $\mathbb{Z}_2$ is $\Theta(\sqrt{2}^k)$ and used this to give an $O^*((2 + \sqrt{2})^{\mathsf{pw}})$ time algorithm for counting Hamiltonian cycles modulo 2 on graphs of pathwidth pw. The algorithm carries over to the decision problem via witness isolation. The same authors complemented their algorithm by an essentially tight lower bound under the Strong Exponential Time Hypothesis (SETH). This bound crucially relied on a large permutation submatrix within $\mathbf{M}_k$, which enabled a "pattern propagation" commonly used in previous related lower bounds, as initiated by Lokshtanov et al. (SODA 2011).

We present a new technique for a similar pattern propagation when only a *black-box* lower bound on the asymptotic rank of $\mathbf{M}_k$ is given; no stronger structural insights such as the existence of large permutation submatrices in $\mathbf{M}_k$ are needed. Given appropriate rank bounds, our technique yields lower bounds for counting Hamiltonian cycles (also modulo fixed primes p) parameterized by pathwidth.

To apply this technique, we prove that the rank of $\mathbf{M}_k$ over the rationals is $4^k/\mathrm{poly}(k)$, using the representation theory of the symmetric group and various insights from algebraic combinatorics. We also show that the rank of $\mathbf{M}_k$ over $\mathbb{Z}_p$ is $\Omega(1.97^k)$ for any prime $p \neq 2$ and even $\Omega(2.15^k)$ for some primes.

Combining our rank bounds with the new pattern propagation technique, we show that Hamiltonian cycles cannot be counted in time $O^*((6 - \varepsilon)^{\mathsf{pw}})$ for any $\varepsilon > 0$ unless SETH fails. This bound is tight due to a $O^*(6^{\mathsf{pw}})$ time algorithm by Bodlaender et al. (ICALP 2013). Under SETH, we also obtain that Hamiltonian cycles cannot be counted modulo primes $p \neq 2$ in time $O^*(3.97^{\mathsf{pw}})$ and, for some primes, not even in time $O^*(4.15^{\mathsf{pw}})$, indicating that the modulus can affect the complexity in intricate ways.

1 Introduction

Rank is a fundamental concept in linear algebra and has numerous applications in diverse areas of discrete mathematics and theoretical computer science, such as algebraic complexity [8], communication complexity [27], and extremal combinatorics [30], to name only a few. A common phenomenon is that low rank often helps in proving combinatorial upper bounds or designing algorithms, e.g., through representative sets [7, 13, 21] or the polynomial method (which ultimately relies on fast rectangular matrix multiplication, enabled through low-rank factorizations of problem-related matrices [37]). In particular, rank has recently found applications in *fine-grained complexity* (see [1] and the references therein) and *parameterized complexity*. In the latter, several influential results, such as algorithms for kernelization [21], the longest path problem [31], and connectivity problems parameterized by treewidth [12, 11, 7], rely crucially on low-rank factorizations.

In view of the utility of low rank in proving upper bounds, it is natural to ask, conversely, whether high rank translates into lower bounds. Indeed, examples for this connection can be found in communication complexity [22, Section 1.4] and circuit complexity [16]. In the present paper, we find such applications also in fine-grained and parameterized complexity: We develop a technique that allows us to transform rank lower bounds into conditional lower bounds for the problem #HC of counting Hamiltonian cycles. The decision version HC of #HC, which asks for the *existence* of a Hamiltonian cycle, is a classical subject of algorithmic research. For decades, the well-known $O^*(2^n)$ time dynamic programming algorithm [18] was essentially the fastest known algorithm for HC, until a breakthrough result [4] showed that HC can actually be solved in $O^*(1.657^n)$ randomized time. This result spawned several novel algorithmic insights into HC, but also showed that we still do not understand this problem in a satisfactory way: No *deterministic $O^*((2 - \varepsilon)^n)$* time algorithm for HC is known, and even no randomized $O^*((2-\varepsilon)^n)$ time algorithms are known for the more general traveling salesman problem, the directed Hamiltonian cycle problem, or the counting version #HC.

One of the novel algorithmic techniques for HC following in the wake of [4] is closely tied to the rank of the so-called *matchings connectivity matrix* [11]. For even k, the matchings connectivity matrix $\mathbf{M}_k$ is indexed by the perfect matchings of the complete graph K_k, and the entry $\mathbf{M}_k[M, M']$ for perfect matchings M and M' is defined as 1 if the union $M \cup M'$ is a single cycle, and 0 otherwise. See Figure 1 for an example. The matchings connectivity matrix can be seen as a description of the behavior of Hamiltonian cycles under graph separators, an interpretation that proved useful for algorithmic applications. For instance, the authors of [11] show that the rank of $\mathbf{M}_k$ over $\mathbb{Z}_2$ is precisely $2^{k/2-1}$ and use this surprisingly low rank to count Hamiltonian cycles modulo 2 in bipartite directed graphs in $O(1.888^n)$ time, which was recently improved to $O^*(3^{n/2})$ time in [6]. A randomized algorithm for the decision version follows from witness isolation.

The low rank of $\mathbf{M}_k$ also enabled an $O^*((2+\sqrt{2})^{\mathrm{pw}})$ time algorithm for HC on graphs with a given path decomposition of width pw. For this problem, the standard dynamic programming approach would require to keep track of all partitions of separators, resulting in a running time of $O^*(2^{\mathrm{pw}\log\mathrm{pw}})$; it is thus somewhat remarkable that the single-exponential running time of $O^*((2 + \sqrt{2})^{\mathrm{pw}})$ can be achieved. Even more surprisingly, the base $2 + \sqrt{2}$ appears to be optimal, as it is known that any $O((2 + \sqrt{2} - \varepsilon)^{\mathrm{pw}})$ time algorithm would violate the Strong Exponential Time Hypothesis (SETH) [11]. This was proven by combining a general reduction technique for SETH-based lower bounds from [24] with a special property of $\mathbf{M}_k$, namely, that $\mathbf{M}_k$ contains a principal minor of size $2^{k/2-1}$ that is a permutation matrix. In other words, there is a collection of $2^{k/2-1}$ perfect matchings such that every perfect matching in this collection can be extended to a Hamiltonian cycle by precisely one other member.

Nr		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1		0	0	0	0	1	1	0	1	1	1	1	0	1	1	0
2		0	0	0	1	0	1	1	1	0	0	1	1	1	0	1
3		0	0	0	1	1	0	1	0	1	1	0	1	0	1	1
4		0	1	1	0	0	0	0	1	1	1	0	1	1	0	1
5		1	0	1	0	0	0	1	0	1	0	1	1	1	1	0
6		1	1	0	0	0	0	1	1	0	1	1	0	0	1	1
7		0	1	1	0	1	1	0	0	0	0	1	1	0	1	1
8		1	1	0	1	0	1	0	0	0	1	0	1	1	1	0
9		1	0	1	1	1	0	0	0	0	1	1	0	1	0	1
10		1	0	1	1	0	1	0	1	1	0	0	0	0	1	1
11		1	1	0	0	1	1	1	0	1	0	0	0	1	0	1
12		0	1	1	1	1	0	1	1	0	0	0	0	1	1	0
13		1	1	0	1	1	0	0	1	1	0	1	1	0	0	0
14		1	0	1	0	1	1	1	1	0	1	0	1	0	0	0
15		0	1	1	1	0	1	1	0	1	1	1	0	0	0	0

Fig. 1: Originally from [11], this figure displays the matchings connectivity matrix $\mathbf{M}_6$, which indicates which pairs of the 15 perfect matchings on 6 vertices form a Hamiltonian cycle.

The general technique for SETH-based lower bounds from [24] was successfully applied to various problems parameterized by pathwidth: As a result, the optimal base in the exponential dependence on the pathwidth has been identified for many problems, assuming SETH. However, there are still natural open problems left, such as #HC: It has been shown that this problem can be solved in $O^*(6^{\mathrm{pw}})$ time [7], later extended to $O^*((2^\omega + 2)^{\mathrm{tw}})$ time when parameterized by treewidth [38], where $2 \le \omega < 2.371$ denotes the matrix multiplication constant. A tight lower bound however remained elusive, and this might justify optimism towards improved algorithms: For example, if we could lift the $O^*((2+\sqrt{2})^{\mathrm{pw}})$ time algorithm for #HC modulo 2 to an $O^*((4-\varepsilon)^{\mathrm{pw}})$ time algorithm for #HC, we could solve #HC on bipartite graphs in $O^*((2-\varepsilon')^n)$ time, since bipartite graphs have pathwidth at most $n/2$.

Our main results

We strike out the route towards faster algorithms for #HC sketched above: We show that the current pathwidth (and, assuming $\omega = 2$, treewidth) based algorithms are optimal assuming SETH.

Theorem 1.1. *Assuming SETH, there is no $\varepsilon > 0$ such that #HC can be solved in $O^*((6 - \varepsilon)^{\mathrm{pw}})$ time on graphs with a given path decomposition of width* pw.

This theorem gives a natural example for an NP-hard problem whose decision version (with base $2 + \sqrt{2}$) and counting version (with base 6) differ provably under SETH.

We prove Theorem 1.1 by starting from the general reduction technique in [24], augmented with a novel idea: We extend the technique in such a way that it can exploit arbitrary lower bounds on the matrix rank of $\mathbf{M}_k$, without further insights into the particular structure of basis vectors. That is, we derive Theorem 1.1 as a consequence of the following more general "black-box" connection

3

between the rank of $\mathbf{M}_k$ and the running time for #HC: If the exponential base of the rank can be lower-bounded by r, then we do not expect $O^*((2 + r - \varepsilon)^{\mathrm{pw}})$ time algorithms.

Theorem 1.2. *Let $r \in \mathbb{R}$ be such that $\log_r(\mathrm{rank}(\mathbf{M}_k))/k \to c$, where $c \geq 1$, as even k tends to infinity.*[1] *Assuming SETH, there is no $\varepsilon > 0$ such that #HC can be solved in $O^*((2 + r - \varepsilon)^{\mathrm{pw}})$ time on graphs with a given path decomposition of width* pw.

For prime numbers p, the same applies to #HC modulo p when replacing r by r_p, which is defined analogously to r by taking $\mathrm{rank}(\mathbf{M}_k)$ over $\mathbb{Z}_p$.

To prove Theorem 1.1, we then combine Theorem 1.2 with our second main contribution: We determine the rank of $\mathbf{M}_k$ over $\mathbb{Q}$ up to polynomial factors, and for primes $p \neq 2$, we additionally give lower bounds on the rank over $\mathbb{Z}_p$ that are higher than the rank over $\mathbb{Z}_2$.

Theorem 1.3. *The rank of $\mathbf{M}_k$ over the rational numbers is at least $\Omega(4^k/k^3)$. For any prime $p \neq 2$, the rank of $\mathbf{M}_k$ over $\mathbb{Z}_p$ is at least $\Omega(1.979^k)$, and for prime $5 \leq p \leq 13$, it is at least $\Omega(2.152^k)$. See Theorem 3.7 for a full list.*

The bound over $\mathbb{Q}$ is obtained by a novel application of representation theory, inspired by a previous approach from [33], where the rank of a bipartite version of $\mathbf{M}_k$ over $\mathbb{Q}$ was found to be $\Theta(2^k)$ up to polynomial factors. In the bipartite version of $\mathbf{M}_k$, only perfect matchings contained in the complete bipartite graph $K_{k/2,k/2}$ are considered. In our non-bipartite version, any perfect matching in the complete graph K_k is allowed; this is more appropriate for algorithmic applications, and our new bound over $\mathbb{Q}$ shows that going to the non-bipartite setting increases the rank significantly.

Combined with Theorem 1.2, our bound over $\mathbb{Z}_p$ suggests that #HC modulo prime $p \neq 2$ is harder than modulo 2: We can solve #HC modulo 2 in $O^*((2 + \sqrt{2})^{\mathrm{pw}})$ time, where $\sqrt{2} \leq 1.42$, but we cannot solve #HC modulo $p \neq 2$ in $O^*((2 + 1.97)^{\mathrm{pw}})$ time unless SETH fails. This connects to recent results [5, 6], which show that the counting Hamiltonian cycles modulo c^n (not parameterized by pathwidth) can be solved in time $O^*((2 - \varepsilon_c)^n)$, where $\varepsilon_c > 0$ depends on the constant c.

Connection matrices and fingerprints

The matchings connectivity matrix $\mathbf{M}_k$ fits into a bigger picture of so-called *connection matrices* for graph parameters, and our bounds on the rank of $\mathbf{M}_k$ translate into rank bounds in this framework.

The connection matrices of a graph parameter f are a sequence of matrices $\mathbf{C}_k$, for $k \in \mathbb{N}$, which describe the behavior of f under graph separators of size k. To define these matrices, say that a *k-boundaried graph*, for $k \in \mathbb{N}$, is a simple graph with k distinguished vertices that are labeled $1, \ldots, k$. Two k-boundaried graphs G and H can be glued together, yielding a graph $G \oplus H$, by taking the disjoint union of G and H and identifying vertices with the same label. The k-th connection matrix $\mathbf{C}_k$ of f then is an infinite matrix whose rows and columns are indexed by k-boundaried graphs such that the entry $\mathbf{C}_k[G, H]$ is $f(G \oplus H)$.

The ranks of connection matrices are closely related to graph-theoretic, algorithmic, and model-theoretic properties of graph parameters [20, 26, 25]. In particular, the connection matrices $\mathbf{C}_k$ for the number of Hamiltonian cycles were studied in [25, 26], where their rank was upper-bounded by $2^{O(k \log k)}$. As a consequence of Theorem 1.3, we can improve upon this and obtain the following essentially tight bounds.

---

[1] This implies that $\mathrm{rank}(\mathbf{M}_k)$ can be lower-bounded by $\Omega(r^k)$ up to sub-exponential factors.

Theorem 1.4. *For $k \in \mathbb{N}$, the rank of the connection matrix $\mathbf{C}_k$ for the number of Hamiltonian cycles is 6^k, up to polynomial factors.*

To prove this theorem, we use a third matrix, the fingerprint matrix $\mathbf{H}_k$ for Hamiltonian cycles, which will also play an important role in our main reduction.[2] A *fingerprint* of a k-boundaried graph is a pair (d, M), where $d \in \{0, 1, 2\}^k$ assigns 0, 1 or 2 to each boundary vertex, and M is a perfect matching on the boundary vertices to which d assigns 1. Fingerprints are essentially the states one would use in the natural dynamic programming routine for counting Hamiltonian cycles parameterized by pathwidth; they describe the behavior of a Hamiltonian cycle on a given side of a separation. A pair of fingerprints (d, M) and (d', M') on B *combines* if $d_v + d'_v = 2$ for every $v \in B$ and additionally $M \cup M'$ forms a single cycle. The fingerprint matrix $\mathbf{H}_k$ is a binary matrix, indexed by fingerprints, and the value at a pair of fingerprints is 1 iff the two fingerprints combine.

It can be derived easily from our rank bound for the matchings connectivity matrix $\mathbf{M}_k$ that the rank of $\mathbf{H}_k$ is 6^k up to polynomial factors, see Fact 2.2. To establish Theorem 1.4, we show in Fact 2.3 that $\mathbf{C}_k$ and $\mathbf{H}_k$ have the same rank.

Proof techniques

In the remainder of the introduction, we sketch the techniques used to obtain Theorems 1.3 and 1.2, which together imply Theorem 1.1.

Theorem 1.3: Rank of the matchings connectivity matrix

To prove Theorem 1.3, we give two different lower bounds on the rank of $\mathbf{M}_k$: One is relatively simple and contained in Section 3. For this bound, we first explicitly compute the rank of small matching connectivity matrices and then use a product construction to give lower bounds for larger orders. While the resulting bound is loose, it also applies to the rank of $\mathbf{M}_k$ over $\mathbb{Z}_p$ for prime $p \neq 2$, whereas our more sophisticated main bound does not. In particular, we can use the bound to show that the rank of $\mathbf{M}_k$ over $\mathbb{Z}_3$ and other primes is asymptotically larger than the rank over $\mathbb{Z}_2$.

Our main result however concerns the rank of $\mathbf{M}_k$ over $\mathbb{Q}$, which we establish to be 4^k up to polynomial factors in Section 4. To this end, we build upon representation-theoretic techniques that were also used in Raz and Spieker's bound [33] for the bipartite version of $\mathbf{M}_k$, and which we first survey briefly: A *hook partition* λ of some number $k \in \mathbb{N}$ is a number partition with the particular form $(t, 1, \ldots, 1)$ for some $t \leq k$. One can view λ as a *Ferrers diagram*, which is a left-adjusted diagram made of cells, as shown in Figure 4a. A *standard Young tableau* of shape λ is a labeling of this diagram with numbers from $[k]$ such that the numbers are strictly increasing in each row and each column, see Figure 4b. Raz and Spieker showed that the rank of the bipartite variant of $\mathbf{M}_{2k}$ can be expressed as a weighted sum over all hook partitions λ of k, where each λ is weighted by the squared number of Young tableaux of shape λ. This sum simplifies to the central binomial coefficient $\binom{2k-2}{k-1}$, showing that the bipartite variant of $\mathbf{M}_k$ has rank $\Theta(2^k)$, up to polynomial factors.

To address the non-bipartite setting, we found ourselves in need of additional techniques from algebraic combinatorics that were not present in Raz and Spieker's original bound, such as the perfect matching association scheme [14] and zonal spherical functions [29]. With these at hand, we

---

[2] To avoid (or add) confusion, let us stress that we consider three matrices related to the Hamiltonian cycle problem: The matchings connectivity matrix $\mathbf{M}_k$, the connection matrix $\mathbf{C}_k$ for the number of Hamiltonian cycles, and the fingerprint matrix $\mathbf{H}_k$ for Hamiltonian cycles. We will revisit their differences in Section 2. While these matrices are closely related, our arguments benefit from using different matrices for different proofs.
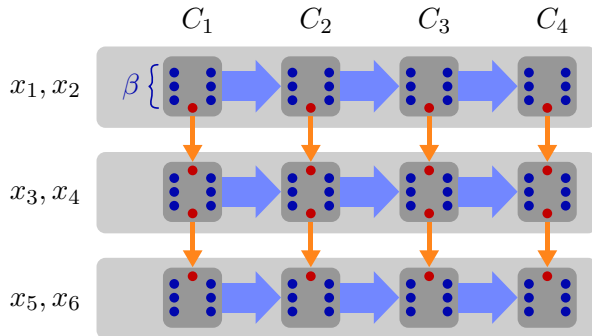
Fig. 2: A block propagation scheme for $n = 6$ variables, $m = 4$ clauses, variable block size $\gamma = 2$, and row pathwidth $\beta = 3$. The $q = n/\gamma = 3$ variable blocks correspond to rows. Assignments to blocks are propagated from left to right. Clauses correspond to columns and clause satisfaction is propagated downwards.

prove that the rank of $\mathbf{M}_{2k}$ can be lower-bounded by a similar sum over number partitions λ as in the bipartite case, this time however ranging over *domino hook partitions* λ, which have the form $(2t, 2t, 2, \ldots, 2)$ for some $t \leq k$. As in the bipartite case, we then observe that this sum simplifies significantly, this time however to (essentially) a product of two consecutive Catalan numbers. This entails a lower bound of 4^k for the rank of $\mathbf{M}_k$, up to polynomial factors. It then follows easily from the upper bound in the bipartite setting that this bound is tight up to polynomial factors.

Theorem 1.2: SETH-hardness via assignment propagation

To describe how we turn lower bounds on the rank of $\mathbf{M}_k$ into algorithmic intractability results for #HC under SETH, let us first survey the general construction from [24], which we dub a *block propagation scheme*: Given a CNF-formula φ with n variables, such a scheme produces an equivalent instance I of the target problem with parameter value $k \leq cn$ for some constant $c \leq 1$. An algorithm with running time $O^*((2^{1/c} - \varepsilon)^k)$ for the target problem would then refute SETH, as it would imply a $O^*((2 - \varepsilon)^n)$ time algorithm for CNF-SAT.

The constructed target instance I has the outline sketched in Figure 2: The n variables of φ are grouped into $q = \lceil n/\gamma \rceil$ blocks of constant size $\gamma \in \mathbb{N}$, where γ depends only on the ϵ in the running time we wish to rule out. The variable blocks are represented as rows, each propagating an assignment of type $\{0,1\}^\gamma$ using a thin graph of pathwidth $\beta \in \mathbb{N}$. Specifically, an assignment $\{0,1\}^\gamma$ is represented as the type of a partial solution for the target problem on a β-boundaried graph (e.g., as a partial coloring of the boundary, or in our case, as a fingerprint of a Hamiltonian cycle.) The relationship between β and γ is important in this construction: Intuitively, if we can choose β small for large γ, then the target problem has a large "combinatorial capacity" in the sense that it allows us to pack assignments to large blocks into thin wires.

In a block propagation scheme, the clauses of φ are then represented as columns; the column corresponding to clause C checks whether the overall assignment of type $\{0,1\}^n$ propagated by the rows satisfies C. To this end, one can use *cell gadgets*, which are graphs with β left and β right interface vertices, and c additional top/bottom interface vertices, where $c \in \mathbb{N}$ depends only on the target problem. The cell gadget is placed at the intersection of a row and a column, and it needs to "decode" an assignment $x \in \{0,1\}^\gamma$ from the state of the left β interface vertices, decide

whether x satisfies the clause, and "encode" x back into the state of right β vertices. The top and bottom interface is used to propagate, from the top of a column downwards, whether the respective clause is already satisfied by the partial assignments to the blocks above a given cell. Due to the grid-like construction, the overall pathwidth of the instance I is usually easily seen to be bounded by $\beta q + O(1)$, where the additive constant accounts for the size of cell gadgets.

The main technical effort in these reductions lies in constructing the cell gadget, and this usually subsumes constructing a "state tester", a gadget that tests whether, in a solution to I, the β left/right interface vertices are in a particular state S (say, a particular partial coloring, or a particular fingerprint of a Hamiltonian cycle). This requires constructing a graph that can be extended to a solution of the target instance iff the β relevant vertices are in state S, and for various problems, such constructions can be achieved with some effort. In the case of #HC, we face the situation that testers for fingerprints of Hamiltonian cycles *do not exist*: There are fingerprints S such that any graph that extends S to a Hamiltonian cycle also extends some unwanted fingerprints $S' \neq S$. Our main insight here is that this problem can be solved by firstly restricting to a set of 6^β good fingerprints that induce a full-rank submatrix $\mathbf{F}$ of the fingerprint matrix $\mathbf{H}_\beta$, and secondly simulating a "linear combination" of testers, with coefficients obtained from the inverse of $\mathbf{F}$. In the fingerprint tester for S, other fingerprints $S' \neq S$ *will* have extensions of non-zero weight, but the weights of these extensions are chosen in such a way (depending on $\mathbf{F}^{-1}$) that extensions of $S' \neq S$ cancel out. (A similar idea was used before to obtain conditional lower bounds for the complexity of permanents [10].) This allows us to simulate a state tester for fingerprints, and the use of cancellations also shows why this lower bound works only for #HC and not for HC—which is fortunate, since HC does admit an $O^*((2+\sqrt{2})^{\mathrm{pw}})$ time algorithm.

2 Preliminaries and notation

If φ is a CNF-formula, and x is a (partial) assignment to its variables, we write $x \models \varphi$ to denote that x satisfies φ. For integers n, we let $[n] = \{1, \ldots, n\}$. All graphs in this paper will be undirected. If $G = (V, E)$ is a graph, $v \in V$ is a vertex and $X \subseteq E$ is an edge set, we let $d_X(v)$ denote the number of edges in X that are incident to v. We write $\mathcal{M}_{2n}$ for the set of all perfect matchings of the complete graph K_{2n}.

Strong Exponential Time Hypothesis: As formulated by Impagliazzo and Paturi [19], the complexity assumption SETH states that for every $\varepsilon > 0$, there is a constant k such that k-CNF-SAT (the satisfiability problem for k-CNF formulas on n variables) cannot be solved in time $O^*((2-\varepsilon)^n)$. It is common to state this hypothesis as ruling out even randomized algorithms, making it slightly stronger. A result from Calabro et al. [9, Theorem 1] gives a randomized reduction from k-CNF-SAT to the problem UNIQUE-k-CNF-SAT, where the k-CNF formula is guaranteed to have at most one satisfying assignment. This allows us to assume that the k-CNF formulas in the statement of (randomized) SETH have at most one satisfying assignment.

Pathwidth: A *path decomposition* of a graph $G = (V, E)$ is a path $\mathbb{P}$ in which each node x has an associated set of vertices $B_x \subseteq V$ (called a *bag*) such that $\bigcup_x B_x = V$ and the following holds:

1. For each edge $\{u, v\} \in E$ there is a node x in $\mathbb{P}$ such that $u, v \in B_x$.

2. If $v \in B_x \cap B_y$ then $v \in B_z$ for all nodes z on the (unique) path from x to y in $\mathbb{P}$.

The *width* of $\mathbb{P}$ is the size of the largest bag minus one, and the pathwidth of a graph G is the minimum width over all possible path decompositions of G. A path decomposition *starts in L* if the first bag contains L and *ends in R* if the last bag contains R.

Since our focus here is on path decompositions, we only mention in passing that the related notion of treewidth can be defined in the same way, except for letting the nodes of the decomposition form a tree instead of a path.

Kronecker products: Given a field $\mathbb{F}$ and two matrices $\mathbf{A} \in \mathbb{F}^{n \times m}$ and $\mathbf{B} \in \mathbb{F}^{n' \times m'}$, the Kronecker product $\mathbf{A} \otimes \mathbf{B}$ is a matrix in $\mathbb{F}^{n \cdot n' \times m \cdot m'}$. Its rows can be indexed by pairs $(i, i') \in [n] \times [n']$, and similarly for columns. The entry of $\mathbf{A} \otimes \mathbf{B}$ at row (i, i') and column (j, j') is defined as $\mathbf{A}[i, j] \cdot \mathbf{B}[i', j']$. For $t \in \mathbb{N}$, the t-th Kronecker power of $\mathbf{A}$ is the t-fold product $\mathbf{A}^{\otimes t} = \mathbf{A} \otimes \ldots \otimes \mathbf{A}$, and we consider its rows and columns to be indexed by $[n]^t$ and $[m]^t$ respectively.

If $\mathbf{A}$ and $\mathbf{B}$ each have full rank over $\mathbb{F}$, then so does $\mathbf{A} \otimes \mathbf{B}$. Note that this requires $\mathbb{F}$ to be a field; it would fail if $\mathbb{F}$ contained zero divisors. For our purposes of computing the rank of matrices over $\mathbb{Z}_p$, this means we require p to be prime.

Fingerprints: The *HC-fingerprints* (which we often abbreviate as *fingerprints*) capture the states of the natural dynamic program for Hamiltonian cycles:

Definition 2.1 (Fingerprint, Partial Solutions). *Let $G = (V, E)$ be a graph and let $B \subseteq V$, where B is the set of 'boundary vertices'. A fingerprint on B is a pair (d, M) where $d \in \{0, 1, 2\}^B$ and M is a perfect matching on $d^{-1}(1)$. A partial solution in G for (d, M) is an edge set $H \subseteq E$ such that (i) $d_H(v) = 2$ for every $v \in V \setminus B$, (ii) $d_H(v) = d_v$ for every $v \in B$, and (iii) if $(u, v) \in M$ then u and v are the endpoints of the same path of H.*

Two fingerprints (d, M) and (d', M') on B combine (or match) if $d_v + d'_v = 2$ for every $v \in B$ and $M \cup M'$ forms a single cycle or is empty.

Variants of connection matrices: Our paper studies three related matrices that describe the behavior of Hamiltonian cycles under separators. We recall their definitions here for reference:

$\mathbf{M}_k$, the k-th matchings connectivity matrix: This binary matrix is indexed by perfect matchings $M, M' \in \mathcal{M}_k$, and $\mathbf{M}_k[M, M']$ is 1 iff $M \cup M'$ is a single cycle. This matrix appears naturally in our rank lower bounds.

$\mathbf{H}_k$, the k-th fingerprint matrix (for Hamiltonian cycles): This binary matrix is indexed by all fingerprints (d, M) on a fixed set B of size k. An entry $\mathbf{H}_k[f, f']$ equals 1 iff f and f' combine. This matrix will be crucially used in the algorithmic lower bound.

$\mathbf{C}_k$, the k-th connection matrix (for the number of Hamiltonian cycles): This integer-valued matrix is indexed by all k-boundaried graphs, and hence infinite. The entry $\mathbf{C}_k[G, G']$ counts the Hamiltonian cycles in the graph $G \oplus G'$. (If an edge between boundary vertices is present in both G and G', we count it twice in $G \oplus G'$.) This matrix will not be used in later sections, but we still mentioned it to *connect* to the established literature on connection matrices.

The subscript k is omitted when clear from the context. As we show below, we can easily transform rank bounds for $\mathbf{M}_k$ into bounds for $\mathbf{H}_k$ and $\mathbf{C}_k$, thus justifying our focus on the rank of $\mathbf{M}_k$:

Fact 2.2. *If p is a polynomial such that $\mathrm{rank}(\mathbf{M}_k) \geq c^k/p(k)$, then $\mathrm{rank}(\mathbf{H}_k) \geq (2+c)^k/q(k)$ for some polynomial q.*

Proof. Subject to proper indexing, the matrix $\mathbf{H}_k$ is a block-antidiagonal matrix that has a block for every vector $d \in \{0,1,2\}^k$, since fingerprints with degree functions d, d' not satisfying $d(v)+d'(v) = 2$ for all v cannot match. Therefore, we obtain:

$$\mathrm{rank}(\mathbf{H}_k) \;\geq\; \sum_{\substack{|d^{-1}(1)|=i \\ i \text{ even}}} \binom{k}{i} \cdot 2^{k-i} \cdot \mathrm{rank}(\mathbf{M}_i) \;\geq\; \frac{1}{k}\sum_{i=1}^{k} \binom{k}{i} \cdot 2^{k-i} \cdot c^i/p(i) \;\geq\; \frac{1}{k}(2+c)^k/q(k),$$

where q is a polynomial satisfying $p(i) \geq q(i)$ for $i = 1, \ldots, k$, and the last inequality follows from the binomial theorem. The upper bound follows similarly. $\qquad\square$

Furthermore, a simple argument shows that the fingerprint matrix $\mathbf{H}_k$ and the connection matrix $\mathbf{C}_k$ actually have the same rank.

Fact 2.3. *For every $k \in \mathbb{N}$, the matrices $\mathbf{H}_k$ and $\mathbf{C}_k$ have the same rank.*

Proof. We first show that $\mathrm{rank}(\mathbf{C}_k) \geq \mathrm{rank}(\mathbf{H}_k)$ by finding $\mathbf{H}_k$ as a submatrix of $\mathbf{C}_k$. To this end, we construct a k-boundaried graph G_F for every k-fingerprint F and then find $\mathbf{H}_k$ as the submatrix induced by these graphs. Given F, the graph G_F is constructed as follows: At first, it contains only the boundary vertices $1, \ldots, k$. Then we add an arbitrary partial solution for F to G_F. For instance, if $F = (d, M)$ and M is non-empty, pick the lexicographically first edge of the matching M, say $ij \in M$, and connect i to j in G_F with a path that passes through all vertices in $d^{-1}(2)$ in an arbitrary order. Then add all edges in $M \setminus \{ij\}$ to G_F as edges. If M is empty, add a Hamiltonian cycle on $d^{-1}(2)$. Finally, subdivide all edges of the graph; this adds some number of subdivision vertices to G_F, which we consider not to be part of the boundary. Note that the degree of boundary vertex $i \in [k]$ in G_F is precisely $d(i)$.

Given two k-fingerprints F, F', we observe that any Hamiltonian cycle in $G_F \oplus G_{F'}$ uses all edges of the graph, as every edge is incident to a (subdivision) vertex of degree 2. This implies, firstly, that the number of Hamiltonian cycles in $G_F \oplus G_{F'}$ is either 0 or 1. Secondly, it implies that $d_F + d_{F'}$ needs to be the constant 2-function for $G_F \oplus G_{F'}$ to have a Hamiltonian cycle. If this condition is fulfilled, then by construction, $G_F \oplus G_{F'}$ has a Hamiltonian cycle iff $M \cup M'$ forms a single cycle. Summarizing, we have that $\mathbf{C}_k[G_F, G_{F'}] \in \{0,1\}$ and that $\mathbf{C}_k[G_F, G_{F'}] > 0$ iff F and F' match. This shows that the set of k-boundaried graphs G_F, for k-fingerprints F, induce the fingerprint matrix $\mathbf{H}_k$ as a submatrix in $\mathbf{C}_k$, and the lower bound on the rank of $\mathbf{C}_k$ follows.

For the upper bound of $\mathrm{rank}(\mathbf{C}_k) \leq \mathrm{rank}(\mathbf{H}_k)$, we find a matrix $\mathbf{A}_k$ such that $\mathbf{C}_k = \mathbf{A}_k \cdot \mathbf{H}_k \cdot \mathbf{A}_k^T$. The rows of $\mathbf{A}_k$ are indexed by k-boundaried graphs G, the columns are indexed by fingerprints F, and we define $\mathbf{A}_k[G, F]$ to count the partial solutions in G for the fingerprint F.

Given two k-boundaried graphs G and G', every Hamiltonian cycle C in $G \oplus G'$ induces a partial solution in each of G and G', for fingerprints F and F', respectively. The pair of fingerprints $P_C = (F, F')$ can be determined uniquely from the partial solutions of C, and since C is a Hamiltonian cycle, it follows that F and F' match. Given a matching pair of fingerprints (F, F'), the number of Hamiltonian cycles of $G \oplus G'$ with $P_C = (F, F')$ is precisely $\mathbf{A}_k[G, F] \cdot \mathbf{A}_k[G', F']$, as the extensions in each of G and G' can be chosen independently, provided they agree with F and F' respectively.

We conclude that the number of Hamiltonian cycles in $G \oplus G'$ can be expressed as

$$\sum_{\substack{F,F' \\ \text{match}}} \mathbf{A}_k[G,F] \cdot \mathbf{A}_k[G',F'] = (\mathbf{A}_k \cdot \mathbf{H}_k \cdot \mathbf{A}_k^T)[G,G'].$$

It follows that $\mathbf{C}_k = \mathbf{A}_k \cdot \mathbf{H}_k \cdot \mathbf{A}_k^T$ as claimed, establishing the upper bound on the rank. $\qquad \square$

3   A simple rank lower bound

Before showing our main lower bound on the rank of the matchings connectivity matrix $\mathbf{M}_k$ in Section 4, we first establish the second part of Theorem 1.3; this turns out to be somewhat simpler.

Let $p \neq 2$ be a fixed prime. To obtain the lower bound on the rank of $\mathbf{M}_k$ over $\mathbb{Z}_p$, we proceed in two steps: First, we use a computer program to compute, for a small constant $B \in \mathbb{N}$, the rank of $\mathbf{M}_B$ over $\mathbb{Z}_p$. Then we use a product construction to amplify this initial rank to a lower bound on the rank of $\mathbf{M}_{tB}$ for $t \in \mathbb{N}$.

3.1   The initial matrix

We choose $B \in \mathbb{N}$ maximally such that $\mathbf{M}_B$ can still be computed, e.g., by the MATLAB script provided in the ancillary files. If the rank of $\mathbf{M}_B$ over $\mathbb{Z}_p$ is r, then the symmetry of $\mathbf{M}_B$ implies the existence of a set $\mathcal{I}$ of perfect matchings in K_B such that the submatrix $\mathbf{F} = \mathbf{M}_B[\mathcal{I}, \mathcal{I}]$ has full rank over $\mathbb{Z}_p$. Our computations enable the following choices:

Lemma 3.1. *For any prime $p \notin \{2,3\}$, the matrix $\mathbf{M}_{10}$ has (full) rank 945 over $\mathbb{Z}_p$. Furthermore, the matrix $\mathbf{M}_{12}$ has rank 3618 over $\mathbb{Z}_3$, rank 9890 over $\mathbb{Z}_5$, and rank 9933 over $\mathbb{Z}_p$ for $p \in \{7, 11, 13\}$.*

Proof. The dimensions of $\mathbf{M}_{10}$ are 945×945. Our calculations show that the determinant of $\mathbf{M}_{10}$ is non-zero and contains only the prime factors 2 and 3. It follows that $\mathbf{M}_{10}$ has full rank over $\mathbb{Z}_p$ for any prime $p \notin \{2,3\}$. Over $\mathbb{Z}_3$, the rank of $\mathbf{M}_{10}$ is found to be 567, but we will obtain a better bound by going to $\mathbf{M}_{12}$, a matrix of dimensions 10395×10395, where the claimed rank bounds can be obtained by calculation. $\qquad \square$

For concreteness, we illustrate our approach in the next subsections with $\mathbf{M}_6$ as initial matrix, see Figure 1, and revisit the better choices provided in Lemma 3.1 at the end of the proof. Calculation shows that $\det(\mathbf{M}_6) = -2^{17}$, so $\mathbf{M}_6$ has full rank over $\mathbb{Z}_p$ for primes $p \neq 2$, and we can choose a set $\mathcal{I}$ of size 15 to get a full-rank matrix $\mathbf{F} = \mathbf{M}_B[\mathcal{I}, \mathcal{I}]$. Already $B = 6$ gives a lower bound on the rank of $\mathbf{M}_n$ over $\mathbb{Z}_p$ with $p \neq 2$ that is higher than the rank over $\mathbb{Z}_2$.

3.2   Amplification via Kronecker products

After having obtained $\mathbf{F} = \mathbf{M}_B[\mathcal{I}, \mathcal{I}]$, we then "tensor up" this matrix to obtain rank lower bounds on $\mathbf{M}_{tB}$ for $t \in \mathbb{N}$: To this end, we find the Kronecker power $\mathbf{F}^{\otimes t}$, which is a full-rank matrix of dimensions $|\mathcal{I}|^t \times |\mathcal{I}|^t$, as a submatrix of $\mathbf{M}_{tB}$. It follows that $\operatorname{rank}(\mathbf{M}_n) \geq \Omega(|\mathcal{I}|^{n/B})$ whenever n is divisible by B. For $B = 6$, this yields $\operatorname{rank}(\mathbf{M}_n) \geq \Omega(15^{n/6}) = \Omega(1.57^n)$ over $\mathbb{Z}_p$ when $p \neq 2$.

To proceed, it will be useful to define a particular graph $K_B^{(t)}$ on tB vertices for each $t \in \mathbb{N}$, which can be viewed as a subgraph of K_{tB}. Only perfect matchings contained in $K_B^{(t)}$ will be relevant. The graph consists of t disjoint copies of K_B and "patch edges" between adjacent K_B-copies that will be used to combine solutions of the individual K_B-copies to a global solution.
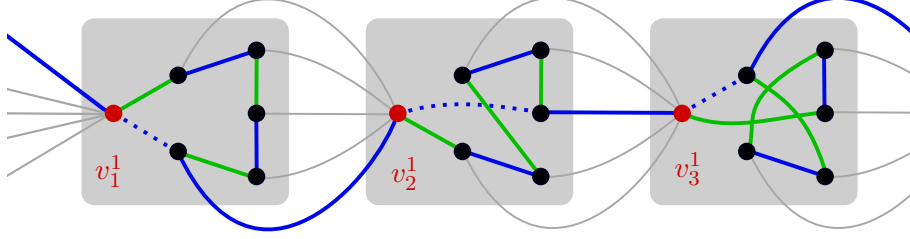
Fig. 3: The graph $K_B^{(t)}$ for $B = 6$ and $t = 3$. Each gray block represents a copy of K_6. Patch edges are drawn with light gray lines; the figure wraps around. Each K_6-copy shows a row matching (green) and a column matching (blue). The special vertex v_i^1 in the i-th copy is drawn red, and the edge incident with v_i^1 in the column matching is dotted. After deleting the dotted edges, we can choose patch edges, depending only on the column matching, so as to obtain a Hamiltonian cycle of $K_B^{(t)}$.

Definition 3.2. *Let $B \in \mathbb{N}$ be fixed. For $t \in \mathbb{N}$, let $K_B^{(t)}$ be obtained as follows, see also Figure 3:*

1. *Take t disjoint copies of K_B and denote the vertices of copy $i \in [t]$ by $v_i^1, \ldots, v_i^B$.*

2. *For each $i \in [t]$ and $j \in [t] \setminus \{1\}$, add an edge from v_i^j to v_{i+1}^1, interpreting $t+1$ as 1. These are the* patch edges.

The perfect matchings of $K_B^{(t)}$ contain a particular subset $\mathcal{I}^{\ominus t}$ of size $|\mathcal{I}|^t$ that is essentially the t-th power of $\mathcal{I}$; this set will be the row set of the full-rank submatrix we wish to find in $\mathbf{M}_{tB}$. The elements of $\mathcal{I}^{\ominus t}$ are disjoint unions of perfect matchings, one for each K_B-copy.

Definition 3.3. *Given a tuple $N = (N_1, \ldots, N_t) \in \mathcal{I}^t$, for $t \in \mathbb{N}$, we define a perfect matching $M_N^{\ominus}$ of the graph $K_B^{(t)}$ by*

$$M_N^{\ominus} = \{\{v_i^a, v_i^b\} \mid i \in [t] \text{ and } \{a, b\} \in N_i\}.$$

We write $\mathcal{I}^{\ominus t} = \{M_N^{\ominus} \mid N \in \mathcal{I}^t\}$ for the perfect matchings that can be obtained from $\mathcal{I}^t$ this way.

It remains to find an appropriate column set of perfect matchings. Note that we cannot reuse $\mathcal{I}^{\ominus t}$ for this purpose: If $t \geq 2$, the union of any two perfect matchings in $\mathcal{I}^{\ominus t}$ is disconnected, and therefore $\mathbf{M}_{tB}[\mathcal{I}^{\ominus t}, \mathcal{I}^{\ominus t}]$ contains only zeroes.

We do however obtain a suitable column set, which we denote by $\mathcal{I}^{\oplus t}$, by using the patch edges of $K_B^{(t)}$: Each perfect matching in $\mathcal{I}^{\oplus t}$ is obtained from some $M \in \mathcal{I}^{\ominus t}$ by deleting one particular edge from each K_B-copy, and patching the resulting isolated vertices to adjacent K_B-copies.

Definition 3.4. *Given a tuple $N = (N_1, \ldots, N_t) \in \mathcal{I}^t$, for $t \in \mathbb{N}$, we define a perfect matching $M_N^{\oplus}$ of the graph $K_B^{(t)}$:*

1. *Start with the perfect matching $M_N^{\ominus}$.*

2. *For $i \in [t]$, let $r(i)$ denote the neighbor of 1 in N_i. Delete the edge from v_i^1 to $v_i^{r(i)}$ in $M_N^{\ominus}$, rendering these two vertices isolated.*

3. *For $i \in [t]$, include the patch edge from $v_i^{r(i)}$ to v_{i+1}^1. (Consider $t+1 = 1$ here.)*

11

We then define $\mathcal{I}^{\oplus t} = \{M_N^{\oplus} \mid N \in \mathcal{I}^t\}$.

It is easily seen that $M_N^{\oplus}$ is indeed a perfect matching of the graph $K_B^{(t)}$, for each $N \in \mathcal{I}^t$: We started with the perfect matching $M_N^{\ominus}$, then reduced the degree of v_i^1 and $v_i^{r(i)}$ to 0 for all $i \in [t]$, and then increased these degrees back to 1 in the third step. No other degrees were affected.

Having defined our rows $\mathcal{I}^{\ominus t}$ and columns $\mathcal{I}^{\oplus t}$, we proceed to study the submatrix $\mathbf{M}_{tB}[\mathcal{I}^{\ominus t}, \mathcal{I}^{\oplus t}]$. Note that both $\mathcal{I}^{\ominus t}$ and $\mathcal{I}^{\oplus t}$ correspond bijectively to $\mathcal{I}^t$, so the indexing of $\mathbf{M}_{tB}[\mathcal{I}^{\ominus t}, \mathcal{I}^{\oplus t}]$ already puts this matrix close to the t-th Kronecker power of $\mathbf{F}$. Its content also does not fail us:

Lemma 3.5. *Identifying $\mathcal{I}^{\ominus t}$ and $\mathcal{I}^{\oplus t}$ each with $\mathcal{I}^t$ in the natural way, we have* $\mathbf{M}_{tB}[\mathcal{I}^{\ominus t}, \mathcal{I}^{\oplus t}] = \mathbf{F}^{\otimes t}$.

Proof. Given $R, C \in \mathcal{I}^t$ with $R = (R_1, \ldots, R_t)$ and $C = (C_1, \ldots, C_t)$, let $H = M_R^{\ominus} \cup M_C^{\oplus}$ be the union of its corresponding perfect matchings. We observe that H is a Hamiltonian cycle (in $K_B^{(t)}$) if and only if $R_i \cup C_i$ is a Hamiltonian cycle (in K_B) for each $i \in [t]$:

In the "if" direction, note that H is the result of deleting one edge each from t Hamiltonian cycles, then adding edges between the endpoints of the resulting Hamiltonian paths so as to obtain a Hamiltonian cycle in $K_B^{(t)}$.

In the "only if" direction, note that the restriction of H to the i-th K_B-copy for $i \in [t]$ is a Hamiltonian path between the v_i^1 and some neighbor. By adding back the edge between v_i^1 and its neighbor and deleting the patch edges, we obtain a Hamiltonian cycle in each copy of K_B.

The claim then follows from the definition of $\mathbf{F} = \mathbf{M}_B[\mathcal{I}, \mathcal{I}]$ and the Kronecker product. $\qquad\square$

Since $\mathbf{F}$ has full rank over $\mathbb{Z}_p$ and p was required to be prime, the Kronecker power $\mathbf{F}^{\otimes t}$ also has full rank, so we obtain:

Corollary 3.6. *The matrix $\mathbf{M}_{tB}[\mathcal{I}^{\ominus t}, \mathcal{I}^{\oplus t}]$ has full rank over $\mathbb{Z}_p$. Consequently, the rank of $\mathbf{M}_{tB}$ over $\mathbb{Z}_p$ is at least $|\mathcal{I}|^t$.*

In conclusion, by using $\mathbf{F} = \mathbf{M}_6$, we obtain that, for prime $p \neq 2$, the rank of $\mathbf{M}_n$ over $\mathbb{Z}_p$ is at least $\Omega(15^{n/6}) = \Omega(1.57^n)$. Using the larger initial matrices provided by Lemma 3.1, we obtain the following stronger bounds:

Theorem 3.7. *For prime p, the rank of $\mathbf{M}_n$ over $\mathbb{Z}_p$ is at least*

$$
\begin{aligned}
\Omega(945^{n/10}) = \Omega(1.984^n) &\quad \text{if } p \notin \{2, 3\}, \\
\Omega(3618^{n/12}) = \Omega(1.979^n) &\quad \text{if } p = 3, \\
\Omega(9890^{n/12}) = \Omega(2.152^n) &\quad \text{if } p = 5, \\
\Omega(9933^{n/12}) = \Omega(2.153^n) &\quad \text{if } p \in \{7, 11, 13\}.
\end{aligned}
$$

The bounds can be improved by using larger initial matrices $\mathbf{F}$, but we hit our computational limit with the 10395×10395 matrix $\mathbf{M}_{12}$. For this matrix, we could no longer compute determinants of the relevant submatrices to determine their prime factors, but we could still compute the rank of $\mathbf{M}_{12}$ for primes up to 13, thus obtaining the last three entries in Theorem 3.7.

4 The rank of the matchings connectivity matrix over the rational numbers

In this section we establish the first part of Theorem 1.3. For this we need some basics on the representation theory of the symmetric group which we first briefly outline.
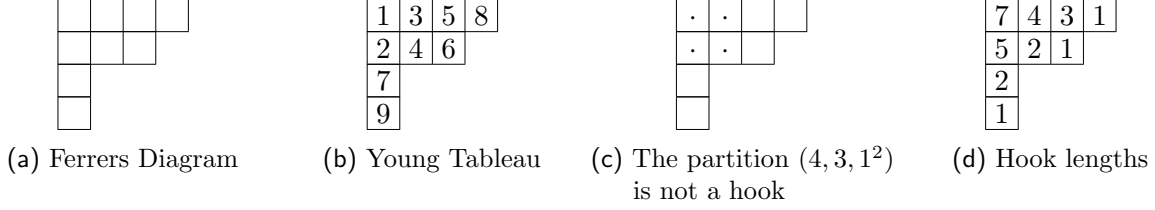
(a) Ferrers Diagram   (b) Young Tableau   (c) The partition $(4, 3, 1^2)$ is not a hook   (d) Hook lengths

Fig. 4: Illustrations of basic notions from the representation theory of the symmetric group

4.1 The Representation Theory of the Symmetric Group

The representation theory of the symmetric group S_n is remarkable, as much of it may be explained via the combinatorics of *integer partitions* and *tableaux*. We outline the relevant combinatorial aspects of the theory, leaving the algebraic basics of finite group representation theory for Appendix B. The reader is referred to [34] for a gentle but more thorough introduction.

Let $\lambda = (\lambda_1, \lambda_2, \cdots, \lambda_k) \vdash n$ such that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k$ and $\sum_{i=1}^{k} \lambda_i = n$ denote an *integer partition* of n. If j parts of the integer partition have the same size m, then we express them by the shorthand m^j. Let $\lambda(n)$ denote the number of integer partitions of n. It is well-known that there is a one-to-one correspondence between the irreducible representations of S_n and the integer partitions of n. We let $[\![\lambda]\!]$ denote the irreducible representation of S_n corresponding to $\lambda \vdash n$.

For an integer partition $\lambda \vdash n$, the *Ferrers diagram* of λ is an associated left-justified tableau that has λ_i cells in the ith row. Abusing notation, we let λ also refer the Ferrers diagram of $\lambda \vdash n$. In Figure 4a the Ferrers diagram for $(4, 3, 1^2) \vdash 9$ is illustrated.

We obtain a *standard Young tableau* from a Ferrers diagram by labeling its cells with numbers such that the numbers along each row are strictly increasing, and the numbers along each column are strictly increasing. In Figure 4b a standard Young tableau of shape $(4, 3, 1^2)$ is shown.

Let f^λ denote the number of standard Young tableaux of shape $\lambda \vdash n$. There is an elegant combinatorial formula for expressing f^λ.

We say that a tableau λ *covers* a tableau μ if the cells of μ are contained in the cells of λ. A *hook* is a tableau of shape $(k, 1^\ell)$, equivalently, a tableau that does not cover the shape (2^2). The partition $(4, 3, 1^2)$ is not a hook, as it covers (2^2), illustrated in Figure 4c.

For each cell $c \in \lambda$ of a Ferrers diagram, say at row i and column j, if we take c along with all cells in row i to the right of c, and all cells in column j that lie below c, we obtain a hook $(k, 1^\ell)$ for some $k, \ell \in \mathbb{N}$. Let $h(c) := k + \ell$ denote the number of cells in this hook, the so-called *hook length*. In Figure 4d, we have annotated each cell with its corresponding hook length. The following result connects hook lengths with enumerating standard Young tableaux.

Theorem 4.1 (Hook Theorem [34]). $f^\lambda = \frac{n!}{\prod_{c \in \lambda} h(c)}$.

For instance, it is easy to see using the hook formula that $f^\lambda = \binom{n-1}{\ell}$ for any hook $\lambda = (n-\ell, 1^\ell) \vdash n$. A classic result in the representation theory of the symmetric group S_n is that f^λ equals the dimension of the irreducible $[\![\lambda]\!]$ corresponding to λ.

Proposition 4.2 (Dimensions of Irreducibles of S_n [34]). $f^\lambda = \dim[\![\lambda]\!]$.

13

4.2 Relating rank to the number of Young tableaux

We proceed by studying $\mathbf{M} := \mathbf{M}_k$ where we let $k = 2n$. Let A, B be a partition of the vertices of K_{2n} into two parts of size n. Consider the sub-matrix $\mathbf{M}'$ of $\mathbf{M}$ induced by the perfect matchings of K_{2n} that are also bipartite perfect matchings with respect to the bipartition A, B. In [33], Raz and Spieker show that the eigenspaces of $\mathbf{M}'$ are in fact irreducible representations of S_n, and that the eigenspaces corresponding to nonzero eigenvalues of $\mathbf{M}'$ correspond to the hooks of length n. This result paired with some elementary combinatorics implies the following theorem.

Theorem 4.3 (Raz & Spieker [33])**.**

$$\mathrm{rank}(\mathbf{M}') = \sum_{\substack{\lambda \vdash n \\ \lambda \text{ does not cover } (2^2)}} (f^\lambda)^2 = \binom{2n-2}{n-1}.$$

Since there are $\frac{1}{2}\binom{2n}{2}$ ways to partition $V(K_{2n})$ into two parts A, B of size n, this already gives an upper bound of $\frac{1}{2}\binom{2n-2}{n-1}\binom{2n}{n}$ on the rank of $\mathbf{M}$. We will show this is almost tight.[3] One of our key technical theorems is the following exact formula for the rank of $\mathbf{M}$.

Theorem 4.4. *For any $\lambda \vdash n$, let us write $2\lambda = (2\lambda_1, 2\lambda_2, \cdots, 2\lambda_k) \vdash 2n$. Then*

$$\mathrm{rank}(\mathbf{M}) = \sum_{\substack{\lambda \vdash n \\ \lambda \text{ does not cover } (2^3)}} f^{2\lambda}.$$

This result can be seen as the non-bipartite analogue of Theorem 4.3. To prove it, we determine the nonzero eigenvalues of $\mathbf{M}$; however, this will require a fair amount of algebraic combinatorics, which we now develop.

Let H_n denote the *hyperoctahedral group* of order n, equivalently, the group of permutations $\sigma \in S_{2n}$ such that

$$\{\{\sigma(1), \sigma(2)\}, \{\sigma(3), \sigma(4)\}, \cdots, \{\sigma(2n-1), \sigma(2n)\}\} = \{\{1, 2\}, \{3, 4\}, \cdots, \{2n-1, 2n\}\}.$$

It is well-known that the set of perfect matchings of K_{2n} can be written as $\mathcal{M}_{2n} := S_{2n}/H_n$. Even though S_{2n}/H_n is not a group, these cosets possess a remarkable amount of algebraic structure.

Let $\mathbb{R}[\mathcal{M}_{2n}]$ be the vector space of real-valued functions over perfect matchings, equivalently, the space of real-valued functions over S_{2n} that are H_n-*invariant*, that is, $f(\sigma) = f(\sigma h) \; \forall \sigma \in S_{2n}, h \in H_n$. In [36], Thrall showed this vector space admits the following decomposition into irreducible representations of S_{2n}.

Theorem 4.5 (Thrall '42)**.** $\mathbb{R}[\mathcal{M}_{2n}] = \bigoplus_{\lambda \vdash n} [\![2\lambda]\!]$.

A consequence of Thrall's result is that $\mathcal{M}_{2n}$ admits a *symmetric association scheme*, the so-called *perfect matching association scheme* [14, Section 15.4].

Definition 4.6 (Symmetric Association Scheme)**.** *A collection of binary $m \times m$ matrices $A_0, A_1, \cdots, A_d$ is a symmetric association scheme if the following axioms are satisfied.*

---

[3] Moreover, it can be verified experimentally that for some constant n the rank of $\mathbf{M}$ is strictly smaller than this bound, but lower order terms will not be relevant for us.

1. $A_0 = I$ where I is the identity matrix.

2. $\sum_{i=0}^{d} A_i = J$ where J is the all-ones matrix.

3. $A_i = A_i^T$ for all $i \in \{1, 2, \cdots, d\}$.

4. $A_i A_j$ is a linear combination of $A_0, A_1, \cdots, A_d$ for all $0 \leq i, j \leq d$.

5. $A_i A_j = A_j A_i$ for all $0 \leq i, j \leq d$.

We refer the reader to [3, 14] for a more thorough treatment of association schemes.

Recall that the union of any two perfect matchings is a disjoint union of cycles, which can be represented by an integer partition of the form $2\lambda \vdash 2n$ where $2\lambda = (2\lambda_1, 2\lambda_2, \cdots, 2\lambda_k)$ for some $\lambda \vdash n$. The perfect matching association scheme is simply the collection of $\mathcal{M}_{2n} \times \mathcal{M}_{2n}$ matrices $\mathcal{A} := \{A_\lambda\}_{\lambda \vdash n}$ defined such that $(A_\lambda)_{ij} = 1$ if $i \cup j \cong 2\lambda$, and 0 otherwise.

Proposition 4.7. $\mathbf{M} \cong A_{(n)}$.

Since $\mathcal{A}$ is a symmetric association scheme, the eigenspaces of the A_λ's coincide, and are precisely the irreducibles in the decomposition given by Thrall [14, Section 15.4]. In light of this, we can take the distinct eigenvalues of these matrices as column vectors and collect them in a $\lambda(n) \times \lambda(n)$ matrix P. For example, when $n = 4$, we have $\mathcal{A} = \{A_{(4)}, A_{(3,1)}, A_{(2,2)}, A_{(2,1^2)}, A_{(1^4)}\}$, and

$$
P = \begin{array}{c}
\\ (4) \\ (3,1) \\ (2,2) \\ (2,1^2) \\ (1^4)
\end{array}
\begin{array}{c}
\begin{array}{ccccc} (4) & (3,1) & (2,2) & (2,1^2) & (1^4) \end{array} \\
\left[\begin{array}{ccccc}
48 & 32 & 12 & 12 & 1 \\
-8 & 4 & -2 & 5 & 1 \\
-2 & -8 & 7 & 2 & 1 \\
4 & -2 & -2 & -1 & 1 \\
-6 & 8 & 3 & -6 & 1
\end{array} \right]
\end{array}.
$$

For any $\lambda \vdash n$, we call $\Omega_\lambda := \{m \in \mathcal{M}_{2n} : m \cup m^* \cong 2\lambda\}$ the λ-*sphere*, where

$$m^* = \{\{1, 2\}, \{3, 4\}, \cdots, \{2n-1, 2n\}\}.$$

The following lemma gives a simple way to determine their size.

Lemma 4.8 ([23]). *Let $l(\lambda)$ denote the number of parts of $\lambda \vdash n$, m_i denote the number of parts of λ that equal i, and set $z_\lambda := \prod_{i \geq 1} i^{m_i} m_i!$. Then*

$$|\Omega_\lambda| = \frac{2^n n!}{2^{l(\lambda)} z_\lambda}.$$

Note that the first row of P lists the sizes of the respective spheres. This is no coincidence, as each A_λ has constant row sum $|\Omega_\lambda|$, and so its largest eigenvalue is $|\Omega_\lambda|$ respectively [15]. It is known that the entries of P are determined by the *zonal spherical functions* [29, Chapter VII], which can be thought of as an analogue of irreducible characters in our association scheme setting.

Theorem 4.9 ([29]). *Let η_λ be the eigenvalue of $A_{(n)}$ associated with the λ-eigenspace. Then*

$$\eta_\lambda = |\Omega_\lambda| \omega_\lambda^{(n)}.$$

A result of Diaconis and Lander is the following explicit formula for determining the value of the zonal spherical function ω_λ evaluated on perfect matchings $m \in \Omega_{(n)}$.

Lemma 4.10 (Diaconis & Lander [29])**.** *For any cell $c \in \lambda \vdash n$, let $w(c)$ denote the number of cells on the same row as c that lie strictly to the left of c, and let $n(c)$ denote the number of cells on the same column as c that lie strictly above c. Then*

$$\omega_\lambda^{(n)} = \frac{1}{2^{n-1}(n-1)!} \sum_{\substack{c \in \lambda \\ c \neq (1,1)}} (2w(c) - n(c)),$$

where the sum runs over all cells except for the northwest-most cell $(1,1)$. In particular, we have $\omega_\lambda^{(n)} = 0$ if and only if λ covers (2^3).

We are now in a position to prove Theorem 4.4.

Proof of Theorem 4.4. By Proposition 4.7, we have $\mathbf{M} = A_{(n)}$. Theorem 4.9 and Lemma 4.10 together imply that nonzero eigenvalues of $\mathbf{M}$ do not cover (2^3). Lemma 4.8 implies that spheres are nonempty, thus these are precisely the nonzero eigenvalues of $\mathbf{M}$. By Proposition 4.2 and Theorem 4.5, the dimension of the eigenspace corresponding to η_λ is $f^{2\lambda}$, completing the proof. $\qquad\square$

4.3 Counting Young tableaux

Our combinatorial formula for the rank of the matchings connectivity matrix does not seem to admit any particularly revealing closed-form; nevertheless, we can still get a good lower bound on its rank. We say that $2\lambda \vdash 2n$ such that $\lambda \vdash n$ is a *domino hook* if $\lambda = (k, k, 1^{n-2k})$ for some $0 \leq k \leq n$. For example,
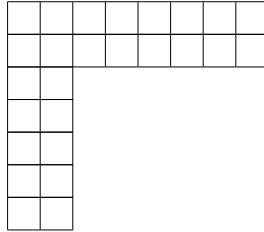
is the Ferrers diagram of the domino hook $2(4, 4, 1^5)$. Note that if $2\lambda \vdash 2n$ is a domino hook, then λ does not cover (2^3). Using WZ-theory [32], Regev showed that the number of standard Young tableaux of domino hook shape admits an elegant count. Recall that $C_n = \frac{1}{n+1}\binom{2n}{n}$ is the nth Catalan number, and that $\lim_{n \to \infty}((4^n/\sqrt{\pi}n^{3/2})/C_n) = 1$, see [35].

Theorem 4.11 (Regev [2])**.**
$$C_{n-1}C_n = \sum_{\substack{2\lambda \vdash 2n \\ 2\lambda \text{ domino hook}}} f^{2\lambda}.$$

Now we can combine all work and finish the proof:

Proof of Theorem 1.3, first part. We have by Theorem 4.4 that

$$\text{rank}(\mathbf{M}_n) = \sum_{\substack{\lambda \vdash n \\ \lambda \text{ does not cover } (2^3)}} f^{2\lambda} \geq \sum_{\substack{2\lambda \vdash 2n \\ 2\lambda \text{ domino hook}}} f^{2\lambda} = C_{n-1}C_n \geq 4^n/n^3,$$

where the inequality follows because if $2\lambda \vdash 2n$ is a domino hook, then λ does not cover (2^3). $\qquad\square$

Note that by the aforementioned upper bound of $\frac{1}{2}\binom{2n-2}{n-1}\binom{2n}{n}$, the lower bound is almost tight.

5   The reduction

This section is devoted to the proof of Theorem 1.2. We will show the following stronger lemma:

Lemma 5.1. *Let p be prime and let $r_p \in \mathbb{R}$ be such that $\log_{r_p}(\text{rank}_p(\mathbf{M}_k))/k \to c$, where $c \geq 1$, as even k tends to infinity. Suppose that $\#\text{HC}$ modulo p on graphs with given path decomposition of width pw can be solved in $O^*((2 + r_p - \varepsilon)^{\text{pw}})$ time, for some $\varepsilon > 0$. Then there is an $O^*((2 - \varepsilon')^n)$ time algorithm that counts the satisfying assignments of a given a CNF-formula on n variables modulo p, for some $\varepsilon' > 0$ depending on ε.*

Lemma 5.1 is a generalization of Theorem 1.2 as in SETH we can without loss of generality assume the number of satisfying solutions is at most one as mentioned in Section 2, so for the decision version we can simply check whether the modular count equals 1 or not. Lemma 5.1 will be used to prove Theorem 1.2. It should be noted that many of the gadgets used in the non-innovative parts of this section are heavily based on the lower bound for the decision version from [11].

5.1   Illustrated outline of proof.

Before describing the reduction in detail, we first give an illustrated outline. For this, a basic understanding of previous block propagation schemes as outlined in Section 1 will be advantageous. We start with a high level description of the statement of Lemma 5.2. The 2^n assignments of the variables of the given CNF-formula φ are encoded by fingerprints that form a basis in the matrix $\mathbf{H}_k$. The larger such a basis, the more assignments one can encode for fixed k. Lemma 5.2 asserts the existence of a certain graph G in which the number of partial solutions of a given fingerprint equals 0 if the fingerprint encodes an assignment not satisfying φ and a fixed positive quantity (depending on the fingerprint) otherwise. The boundary vertices L and R are partitioned into blocks and the fingerprints will also have some block-structure for technical reasons reminiscent to the block-propagation scheme. Confer the statement of Lemma 5.2 for the precise details.

In Figure 5a, we illustrate how the graph output by the reduction implied by Lemma 5.1 is obtained from the graph G obtained by Lemma 5.2: The *blocks* G_L^i and G_R^i (whose unions we jointly denote by G_L respectively G_R) are added to G, where G_R^i has vertex set R_i common with G, and G_L^i has vertex set L_i common with G. Additionally, graphs G_L^i and G_L^{i+1} share a common vertex b_i. By Lemma 5.2, the graph G has $\mathbf{A}[f_L, f_R]$ partial solutions for fingerprints f_L and f_R. The graph G_L has $\mathbf{l}_f$ fingerprints and $\mathbf{r}_{f'}$ fingerprints such that $\mathbf{l}^T \mathbf{A} \mathbf{r}$ equals the number of satisfying solutions of φ modulo p. To establish this, for each fingerprint, certain partial solutions are allowed such that two edgesets avoid creating subcycles only if they have combining fingerprints on L (or R). The structure of the partial solutions is tailored so that combinations of partial solutions with combining fingerprints can be extended in one (modulo p) way into a Hamiltonian cycle. For the latter part, the vertices b_i are used to connect parts from the subgraphs G_L^i. In Figure 5a, the edges present in all partial solutions are displayed everywhere (where dashed edges will actually be redirected to visit vertices with degree 0 or 2 in a fingerprint), and on the vertices $L_i \cup R_i$, a possible partial solution in G_L^i, G, G_R^i is displayed. The edges $l_{i,1}l_{i,2}$ and $r_{i,1}r_{i,3}$ in G are rerouted to avoid creating two subcycles if two matchings give exactly one cycle. If the fingerprint in both G_L^1 and G_R^1 matches with that in G, $l_{1,1}$ will be connected to $r_{1,1}$ and $l_{1,2}$ will be connected to $r_{1,2}$.

(a) From Lemma 5.2 to Theorem 1.2   (b) Lemma 5.2, inductive step   (c) Lemma 5.2, base case
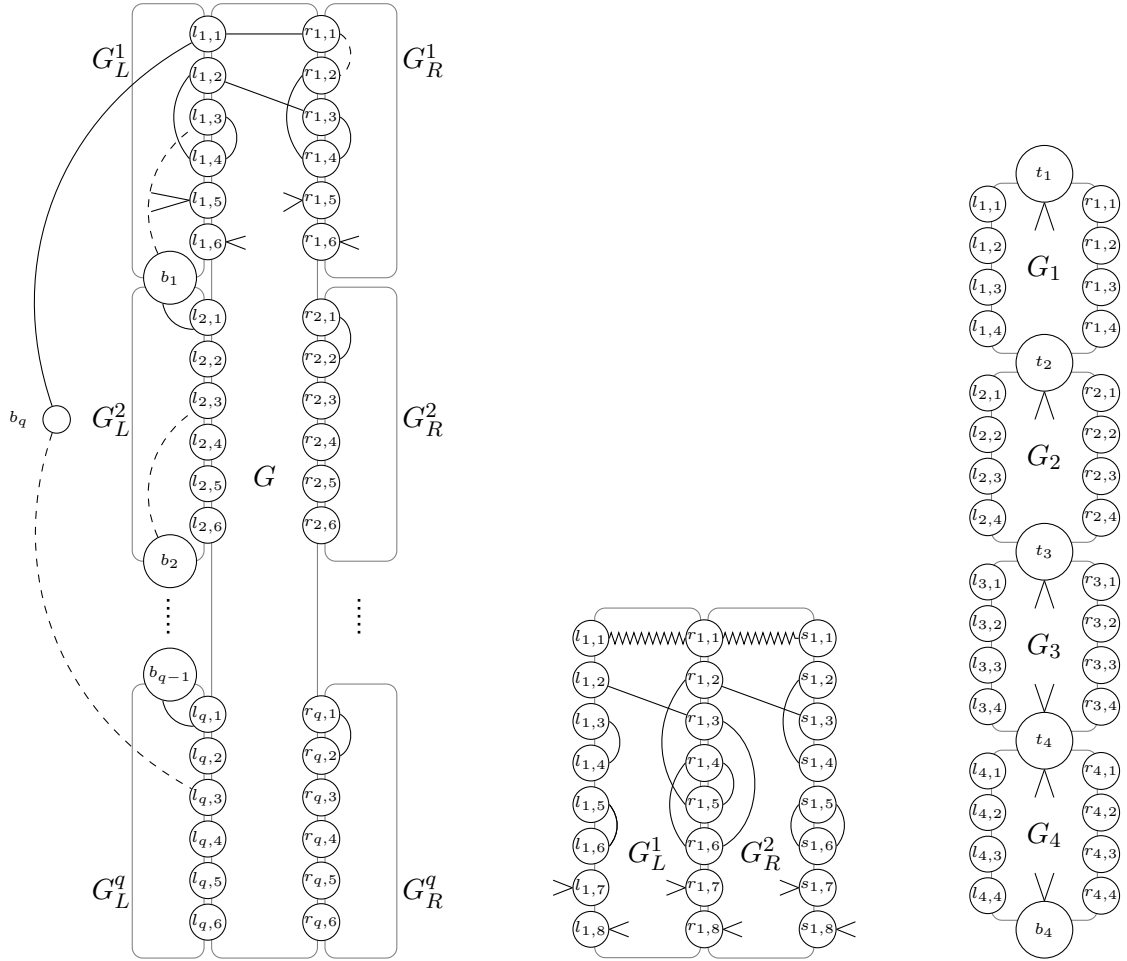
Fig. 5: Illustrations of parts of the proof of Theorem 1.2.

In Figure 5b, we illustrate the inductive step of the proof of Lemma 5.2. This lemma states that, for a CNF-formula $\varphi = C_1 \wedge \ldots \wedge C_m$ on m clauses, there exists a graph such that the number of partial solutions for certain fingerprints encodes the number of satisfying assignments of φ. The proof of this lemma is by induction on m, so in the inductive step, we assume a graph G_L is given for $\varphi' = C_1 \wedge \ldots \wedge C_{m-1}$, and a graph G_R for the CNF-formula C_m. The graphs G_L and G_R have a certain block structure consisting of smaller graphs G_L^i and G_R^i, which allows us to basically restrict attention to one block (i.e., a fingerprint of one block encodes a partial assignment of a block of variables).[4] Suppose the graphs G_L^i and G_R^i have boundary vertices $L \cup R$ and $R \cup S$ respectively. Then we can see these graphs as matrices $\mathbf{L}$ and $\mathbf{R}$ of which an entry indexed by f, f' equals the number of partial solutions on fingerprints f, f' in G_L^i (in which case f is on L and f' is on R), and respectively the number of partial solutions on fingerprints f, f' in G_R^i (in which case f is on R and f' is on S). Note the fingerprint of the partial solution in G_L^i on R and the fingerprint of G_R^i on R also need to match as otherwise the union cannot give a partial solution in the graph $G^i \cup G_R^i$, and by our construction two partial solutions with matching fingerprint on R are also a partial solution of $G^i \cup G_R^i$. It follows that if we let $\mathbf{A}$ be the number of partial solutions in $G^i \cup G_R^i$ indexed by fingerprint f_L on L and f_R on R, than this matrix can be computed as the matrix multiplication $\mathbf{LFR}$, where $\mathbf{F}$ is a full rank submatrix of the Hamiltonian fingerprint matrix by a set of fingerprints that we restrict ourselves to. By using gadgets to ensure that $\mathbf{L} = \mathbf{CF}^{-1}$ where $\mathbf{C}$ is a diagonal matrix checking whether the assignment satisfies φ' the inductive step can be carried out. In Figure 5b the showed partial solution in G_L^1 is derived from the fingerprint $f_L = (d_L, M_L)$ in G_L^1 with $d_L^{-1}(1) = \{l_{1,i}\}_{i=1}^6$, $d_L^{-1}(2) = \{l_{1,8}\}$, $d_L^{-1}(0) = \{l_{1,7}\}$ and $M_L = \{l_{1,1}l_{1,2}, l_{1,3}l_{1,4}, l_{1,5}l_{1,6}\}$ on L and from the fingerprint $f_R = (d_R, M_R)$ in G_L^1 with $d_R(v) = 2 - d_L(v)$, and $M_R = \{r_{1,1}r_{1,3}, r_{1,2}r_{1,5}, r_{1,4}l_{1,6}\}$ on L. If this (or a fingerprint of another block) encodes a partial assignment satisfying φ', there will be $\mathbf{F}^{-1}[f_L, f_R]$ partial solutions with fingerprint $f_L \cup f_R$ in G_L^1.

In Figure 5c we illustrate the base case of the proof of Lemma 5.2. The graph is partitioned into q blocks $G_1, \ldots, G_q$ (with $q = 4$) in the example, and a fingerprint on $L_i = \{l_{i,j}\}_{\{j \leq \beta\}}$ encodes a partial assignment of a block of variables of the CNF-formula, which we consider to be a single clause as $m = 1$ in the base case of the induction. A graphs G_i contains a top vertex t_i and bottom vertices b_i, and the consecutive graphs overlap in the sense that $b_{i+1} = t_i$. Partial solutions are locally (that is, per G_i) restricted such that i is the smallest integer with the property that the partial solution encodes a fingerprint satisfying the clause if and only if the partial solution in G_i has a fingerprint in which both t_i and b_i are of degree 2.

5.2 Pattern propagation using a rank lower bound

Let $\gamma := \gamma(\varepsilon) \geq 4$ be even. Assume we are given a CNF-formula $\varphi = C_1 \wedge \ldots \wedge C_m$ on variables $x_1, \ldots, x_n$ with n being a multiple of γ;[5] let $q = \frac{n}{\gamma}$. Partition the set $\{x_1, \ldots, x_n\}$ into n/γ blocks of size γ, denoted $X_1, \ldots, X_{n/\gamma}$. Intuitively, we will represent the 2^γ assignments of the block of variables X_i by HC-fingerprints on groups of vertices in a bag of the to be constructed path decomposition. We let $\mathcal{B}_l, \mathcal{B}_r$ be sets of HC-fingerprints on $[\beta]$ such that $\mathbf{H}[\mathcal{B}_l, \mathcal{B}_r]$ has full rank over p and if $(d_l, M_r) \in \mathcal{B}_l$ and $(d_r, M_r) \in \mathcal{B}_r$ then

- $d_l(i), d_r(i)$ equal 1 for $i = 1, 2, 3$, and

---

[4] This is not completely true, as partial assignments must interact to check whether any partial assignment satisfies a clause, but this does not turn out problematic.

[5] Note that since γ is a constant this is easily established by adding at most γ dummy variables.

- $\{1,3\} \in M_l$ and $\{1,2\} \in M_r$.

For $i = 1, \ldots, q$ we assume η_i is an injective function from $\mathcal{B}_r$ to $\{0,1\}^{X_i}$ which describes the encoded partial assignment to the variable set X_i. Note this is only possible when $|\mathcal{B}_r| \geq 2^\gamma$, which we will ensure in Subsection 5.5 when wrapping up the proof of Theorem 1.2. We let denote η denote the natural extension of η_i, i.e. $\eta(f_1, \ldots, f_{n/\gamma}) = (\eta_1(f_1), \ldots, \eta_{n/\gamma}(f_{n/\gamma}))$. In the following, we refer to $\mathbf{H}[\mathcal{B}_l, \mathcal{B}_r]$ as $\mathbf{F}$, and also frequently use the q-th Kronecker power $\mathbf{F}^{\otimes q}$. We also use $\mathbf{F}^{-1}$ (which is indexed by $\mathcal{B}_r$ and $\mathcal{B}_l$ respectively) and its Kronecker power $(\mathbf{F}^{-1})^{\otimes q} = (\mathbf{F}^{\otimes q})^{-1}$.

Lemma 5.2. *Let $L = L_1 \uplus \ldots \uplus L_q, R = R_1 \uplus \ldots \uplus R_q$ be disjoint sets with $L_i = \{l_{i,1}, \ldots, l_{i,\beta}\}$ and $R_i = \{r_{i,1}, \ldots, r_{i,\beta}\}$. There exists a graph $G(L, R, \varphi) = (V, E)$ with independent set $L \cup R \subseteq V$ such that for every sequence of fingerprints*

$$f_L = \left((d_L^i, M_R^i)\right) \text{ on } L_i\big)_{i \in [q]} \text{ and } f_R = \left((d_R^i, M_R^i) \text{ on } R_i\right)_{i \in [q]},$$

the number of partial solutions in G for fingerprint $f' = \left(\bigcup_{i=1}^q d_L^i \cup d_R^i, \cup_{i=1}^q M_{alt}^i\right)$ equals

$$\mathbf{A}[f_L, f_R] := \begin{cases} (\mathbf{F}^{\otimes q})^{-1}[f_L, f_R], & \text{if } f_L \in \mathcal{B}_r^q, f_R \in \mathcal{B}_l^q \text{ and } \eta(f_L) \text{ satisfies } \varphi, \\ 0, & \text{otherwise.} \end{cases}$$

Here M_{alt}^i is the altered matching defined by $M_{alt}^i = M_L^i \cup M_R^i \setminus \{l_{i,1}l_{i,2}, r_{i,1}r_{i,3}\} \cup \{l_{i,1}r_{i,1}, l_{i,2}r_{i,3}\}$. Moreover, G has a path decomposition $\mathbb{P}$ of width $\beta \cdot q + O(\beta)$ starting in L and ending in R, and G and $\mathbb{P}$ can be computed in polynomial time.

It may look counterintuitive (or like a typo) that $f_L \in \mathcal{B}_r^q$, but note that the subscripts l and r denote that fingerprints in $\mathcal{B}_l$ will be used for partial solutions connecting vertices 'to the left' of a certain vertex boundary, while fingerprints in $\mathcal{B}_r$ will be used for partial solution 'to the right' of a certain boundary (see Figure 5b). As L is the 'left boundary' of G connection made in G between vertices in L will be 'to the right' of L so the matchings will be in $\mathcal{B}_r$.

We will prove Lemma 5.2 by induction on the number of clauses m. The bulk of our technical efforts with gadgets will be to prove the following lemma for $m = 1$. On the other hand, for proving the inductive step our new key insight of applying matrix inversion is crucial. We postpone the technical proof of the base case $m = 1$ to Subsection 5.4 and first focus on the inductive step.

Proof of Lemma 5.2 (inductive step). Let $\varphi = C_1 \wedge \ldots \wedge C_m$ be a CNF formula on n variables and $\varphi' = C_1 \wedge \ldots \wedge C_{m-1}$. Let $L = L_1 \uplus \ldots \uplus L_q, R = R_1 \uplus \ldots \uplus R_q, S = S_1 \uplus \ldots \uplus S_q$ be disjoint sets with $L_i = \{l_{i,1}, \ldots, l_{i,\beta}\}, R_i = \{r_{i,1}, \ldots, r_{i,\beta}\}$ and $S_i = \{s_{i,1}, \ldots, s_{i,\beta}\}$. Let $\hat{G} = (\hat{V}, \hat{E}) = G(L, R, \varphi')$ and $\tilde{G} = (\tilde{V}, \tilde{E}) = G(R, S, C_m)$ be graphs as given by the induction hypothesis.

Define $G = (\hat{V} \cup \tilde{V}, \hat{E} \cup \tilde{E})$. We show G satisfies the conditions of Lemma 5.2 with φ, as required to prove the inductive step. Note that $\hat{E} \cap \tilde{E} = \emptyset$ as $\hat{V} \cap \tilde{V} = R$ is an independent set in both graphs. As $\hat{V} \cap \tilde{V} = R$, the path decomposition of $\hat{G}$ ending in R and the path decomposition of $\tilde{G}$ starting in R can clearly be combined into a path decomposition of G of the same width. It remains to show that for every sequence of fingerprints

$$f_L = \left((d_L^i, M_L^i) \text{ on } L_i\right)_{i \in [q]} \text{ and } f_S = \left((d_S^i, M_S^i) \text{ on } S_i\right)_{i \in [q]},$$

the number of partial solutions H in G for fingerprint $f = \left(\bigcup_{i=1}^q d_L^i \cup d_S^i, \cup_{i=1}^q M_{alt}^i\right)$ equals

$$\begin{cases} (\mathbf{F}^{\otimes q})^{-1}[f_L, f_S], & \text{if } f_L \in \mathcal{B}_r^q, f_S \in \mathcal{B}_l^q \text{ and } \eta(f_L) \text{ satisfies } \varphi, \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

20

Here $M_{alt}^i = M_L^i \cup M_S^i \setminus \{l_{i,1}l_{i,2}, s_{i,1}s_{i,3}\} \cup \{l_{i,1}s_{i,1}, l_{i,2}s_{i,3}\}$. To show this, we first show the following:

Claim 1. *The number of partial solutions in G for fingerprint f equals*

$$\sum_{\substack{\hat{f}_R \in \mathcal{B}_l^{\otimes q}, \tilde{f}_R \in \mathcal{B}_r^{\otimes q} \\ \eta(f_L) \models \varphi', \ \eta(\tilde{f}_R) \models C_m}} (\mathbf{F}^{\otimes q})^{-1}[f_L, \hat{f}_R] \cdot \mathbf{F}^{\otimes q}[\hat{f}_R, \tilde{f}_R] \cdot (\mathbf{F}^{\otimes q})^{-1}[\tilde{f}_R, f_S].$$

Proof. Let H be a partial solution in G for fingerprint f, let $\hat{H} = H \cap \hat{E}$ and $\tilde{H} = H \cap \tilde{E}$. By construction of $\hat{G}$ and $\tilde{G}$, there are fingerprints $\hat{f}$ and $\tilde{f}$ of the form

$$\hat{f} = \left(\bigcup_{i=1}^q d_L^i \cup \hat{d}_R^i, \cup_{i=1}^q \hat{M}_{alt}^i \right), \quad \tilde{f} = \left(\bigcup_{i=1}^q \tilde{d}_R^i \cup d_S^i, \cup_{i=1}^q \tilde{M}_{alt}^i \right),$$

such that $\hat{H}$ is a partial solution in $\hat{G}$ for $\hat{f}$ and $\tilde{H}$ is a partial solution in $\tilde{G}$ for $\tilde{f}$. Here the altered matchings are of the form

$$\hat{M}_{alt}^i = M_L^i \cup \hat{M}_R^i \setminus \{l_{i,1}l_{i,2}, r_{i,1}r_{i,3}\} \cup \{l_{i,1}r_{i,1}, l_{i,2}r_{i,3}\},$$
$$\tilde{M}_{alt}^i = M_R^i \cup M_S^i \setminus \{r_{i,1}r_{i,2}, s_{i,1}s_{i,3}\} \cup \{r_{i,1}s_{i,1}, r_{i,2}s_{i,3}\}.$$

As H is a partial solution in G we have $d_H(v) = 2$ for all vertices in $v \in R$ and thus $\hat{d}_R^i(j) + \tilde{d}_R^i(j) = 2$ for every i, j. Moreover, H cannot contain a cycle as it is a partial solution and therefore $\hat{M}_{alt}^i \cup \tilde{M}_{alt}^i$ cannot contain a cycle for every $i = 1, \ldots, q$. It follows that $\hat{M}_R^i \cup \tilde{M}_R^i$ must form a single cycle: if not, it contains at least two cycles as $\hat{M}_R^i \cup \tilde{M}_R^i$ are perfect matchings on the same set of vertices, and a cycle not containing the vertex $r_{i,1}$ will still be present in $\hat{M}_{alt}^i \cup \tilde{M}_{alt}^i$. Thus in summary we have that if $\hat{H}$ is a partial solution in $\hat{G}$ for fingerprint $\hat{f}_R = (\cup_{i=1}^q \hat{d}_R^i, \cup_{i=1}^q \hat{M}_R^i)$ on R and is $\tilde{H}$ is a partial solution in $\tilde{G}$ for fingerprint $\tilde{f}_R = (\cup_{i=1}^q \tilde{d}_R^i, \cup_{i=1}^q \tilde{M}_R^i)$ on R, then $\mathbf{F}^{\otimes q}[\hat{f}_R, \tilde{f}_R] = 1$.

For the reverse direction we have that if $\mathbf{F}^{\otimes q}[\hat{f}_R, \tilde{f}_R] = 1$, by the definition of the altered matchings H indeed has fingerprint $f = \left(\bigcup_{i=1}^q d_L^i \cup d_S^i, \cup_{i=1}^q M_{alt}^i \right)$ where

$$M_{alt}^i = M_L^i \cup M_S^i \setminus \{l_{i,1}l_{i,2}, s_{i,1}s_{i,3}\} \cup \{l_{i,1}s_{i,1}, l_{i,2}s_{i,3}\}.$$

To see this, note that in H, the vertex $l_{i,1}$ is connected to $r_{i,1}$ (directly via $\hat{M}_{alt}^i$), which is connected to $s_{i,1}$ (directly via $\tilde{M}_{alt}^i$), and $l_{i,2}$ is connected to $r_{i,3}$ (directly via $\hat{M}_{alt}^i$), which is connected to $r_{i,2}$ (indirectly via the path $(\hat{M}_{alt}^i \cup \tilde{M}_{alt}^i) \cap R \times R$), which is connected to $s_{i,3}$ (directly via $\tilde{M}_{alt}^i$).

The claim follows by summing over all $\hat{f}_R$ and $\tilde{f}_R$ such that $\eta(f_L) \models \varphi'$ and $\eta(\tilde{f}_R) \models C_m$ (the latter two properties follow by the properties of $\hat{G}$ and $\tilde{G}$). $\qquad \square$

By Claim 1 the number of partial solutions in G for fingerprint f equals $\mathbf{A}[f_L, f_S]$, where

$$\mathbf{A} = \mathbf{C}(\mathbf{F}^{\otimes q})^{-1}\mathbf{F}^{\otimes q}\mathbf{C}'(\mathbf{F}^{\otimes q})^{-1} = \mathbf{C}\mathbf{C}'(\mathbf{F}^{\otimes q})^{-1},$$

and $\mathbf{C}, \mathbf{C}'$ are diagonal matrices defined by

$$\mathbf{C}[f, f'] = \begin{cases} 1, & \text{if } f = f' \text{ and } \eta(f) \models \varphi' \\ 0, & \text{otherwise} \end{cases}$$

and

$$\mathbf{C}'[f, f'] = \begin{cases} 1, & \text{if } f = f' \text{ and } \eta(f') \models C_m \\ 0, & \text{otherwise.} \end{cases}$$

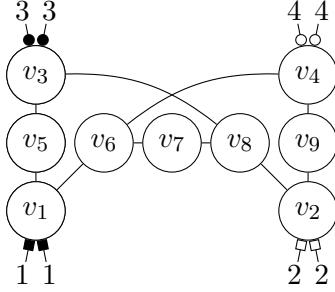and the inductive step follows as $\mathbf{A}[f_L, f_S]$ equals (1). $\qquad \square$

Fig. 6: Implementation of the label gadget. The four labels $\{1, 2, 3, 4\}$ are depicted symbolically with ■ □, •, and ○ respectively for reference in Figure 7.

5.3 Gadgets

In this section we introduce two general gadgets, adapted from previous constructions [11], that are used in the final construction to obtain strong control on the number of Hamiltonian cycles. Both gadgets accept parameters to be set in the final construction.

Label Gadget.   The following gadget allows us to label incident edges of a vertex v and control label combinations of the edges used in a Hamiltonian cycle.

Definition 5.3. *A* label gadget *is a pair* (v, λ_v) *where* $\lambda_v : I(v) \rightarrow \{1, 2, 3, 4\}$ *is a labeling of the edges* $I(v)$ *incident with* v. *A Hamiltonian cycle* C *is* consistent *with label gadget* (v, λ_v) *if* $\lambda_v(e) = 2i$ *and* $\lambda_v(e') = 2i - 1$ *for* $i \in \{1, 2\}$, *where* e, e' *are the two edges of* C *incident with* v.

When we use several label gadgets simultaneously, there will be several labelings and we say an e edge has label l with respect to v if $\lambda_v(e) = l$. We will now show how to replace a label gadget in a graph G with a certain graph to obtain G' such that the number of Hamiltonian cycles in G' equals the number of Hamiltonian cycles in G consistent with (v, λ_v). The graph is shown at the left-hand side in Figure 6. That is, the vertex v is replaced by the displayed graph on vertices $v_1, \ldots, v_9$ and each edge with label $i = 1, \ldots, 4$ is connected to v_i.

Any Hamiltonian cycle contains exactly two edges of the set of edges leaving the gadget and these have either labels 1 and 2 or labels 3 and 4. This follows from a simple case analysis: If the cycle enters the gadget in vertex v_1, it must continue with v_5, v_3. Then it cannot leave the gadget, because then it is impossible to visit all six remaining vertices. Hence it must continue with v_8 and then v_7, v_6, v_4, v_9, and v_2 are forced and the cycle uses edges labeled with 1 and 2. The cases where it enters at a different vertex are symmetric.

Fingerprint Gadget.   Now we present a general gadget allowing strong control on the fingerprints of partial solutions. If B is a set of vertices, we let $\mathcal{P}_B$ denote the set of all HC-fingerprints on B.

Definition 5.4 (Fingerprint Gadget)**.** *A* fingerprint gadget with boundary B *for a positive integer sequence* $\{m_f\}_{f \in \mathcal{P}_B}$ *is a graph* $G' = (V', E')$ *such that* $B \subseteq V$ *and for every HC-fingerprint* $f \in \mathcal{P}_B$ *the number of partial solutions in* G' *for* f *is exactly* m_f.

Lemma 5.5 (Fingerprint Gadget Implementation)**.** *Let $G, B, \{m_f\}_{f \in \mathcal{P}_B}$ as above. Assume there exist at least 2 distinct fingerprints f such that $m_f \neq 0$, and there exist $a, b \in B$ such that if $f = (d, M)$ and $m_f \neq 0$, we have $ab \in M$.[6] There is a fingerprint gadget G' for $\{m_f\}_{f \in \mathcal{P}_B}$ on $O(|B| \sum_{f \in \mathcal{P}_B} m_f)$ vertices with path decomposition of width $|B| + O(1)$ that starts and ends in B.*

Proof. We start with a formal definition of the graph $G' = (V', E')$, which we define using label gadgets. An illustrative example is provided in Figure 7.

1. Add to V' the set B and additionally two vertices a' and c. Add a subdivided edge aa'.[7]

2. Let $(f_i = (d_i, M_i))_{i=1}^{\ell}$ be a sequence of fingerprints that contains each $f \in \mathcal{P}_B$ exactly m_f times, so $\ell = \sum_{f \in \mathcal{P}_B} m_f$.

3. For $i = 1, \ldots, \ell$:

    (a) Let $d_i^{-1}(2) = \{t_i^1, \ldots, t_i^k\}$ be the vertices of degree 2 in d_i.

    (b) Define a sequence of edges $E_i = (e_i^j)_{j=i}^{\ell_i}$ to be an arbitrary ordering of the union of the edgeset of a path from a' through $d_i^{-1}(2)$ to c and the edgeset $M_i \setminus ab$.

        - More formally, let $e_i^1 = a't_i^1$; for $j = 1, \ldots, k - 1 = |d_i^{-1}(2)| - 1$ let $e_i^{j+1} = t_i^j t_i^{j+1}$; let $e_i^{k+1} = t_i^k c$, and let $\{e_i^{k+2}, \ldots, e_i^{\ell_i}\} = M_i \setminus ab$, where $\ell_i = k + 1 + |M_i|$.

    (c) For $j = 1, \ldots, \ell_i$:

        - Let $e_i^j = uv$. Add a label gadget p_i^j with incident edges $p_i^j u$ with label 1 and $p_i^j v$ with label 2.
        - If $j > 1$, add an edge $p_i^{j-1} p_i^j$ with $\lambda_{p_i^{j-1}}(e) = 4$ and $\lambda_{p_i^j}(e) = 3$.

    (d) If $i > 1$, add an edge $e = uv = p_{i-1}^{\ell_{i-1}} p_i^1$ with $\lambda_u(e) = 4$ and $\lambda_v(e) = 3$.

    (e) If $i > 2$, add an edge $e = uv = p_{i-2}^{\ell_{i-2}} p_i^1$ with $\lambda_u(e) = 4$ and $\lambda_v(e) = 3$.

4. For $i = 1, 2$, add an edge $e = uv = c p_1^i$ with $\lambda_v(e) = 3$.

5. For $i = \ell - 1, \ell$, add an edge $e = uv = p_i^{\ell_i} b$ with $\lambda_u(e) = 4$.

Correctness. Let H be a partial solution in G' consistent with all label gadgets. Note that H contains the edges aa'' and $a''a'$ as a'' has degree 2. Thus we see that a' needs to be adjacent in H to p_x^1 for exactly one chosen x, as this vertex has no other neighbors. We claim that for every x there is exactly one partial solution containing the edge $a' p_x^1$ and this partial solution has fingerprint f_x. Note that this is sufficient to prove the lemma, as it implies we have m_f choices for x that lead to fingerprint f.

To this end, suppose that H contains the edge $a' p_x^1$, and let $i \neq x$. Then H cannot contain an edge e incident to p_i^1 satisfying $\lambda_{p_i^1}(e) \in \{1, 2\}$ by the definition of a label gadget as it only has one edge not incident to a' with such a label. Thus H needs to contain edges incident to p_i^1 with labels 3 and 4 with respect to p_i^1. But the only edge with label 4 is to p_i^2 (if it exists) which has label 3

---

[6] These are just technical conditions to facilitate the implementation of this gadget.

[7] Formally, add another vertex a'' with only a and a' as neighbors. Here a'' is a vertex of degree 2 with the sole purpose of enforcing the edges $a'a''$ and $a''a'$ to be in a Hamiltonian cycle.
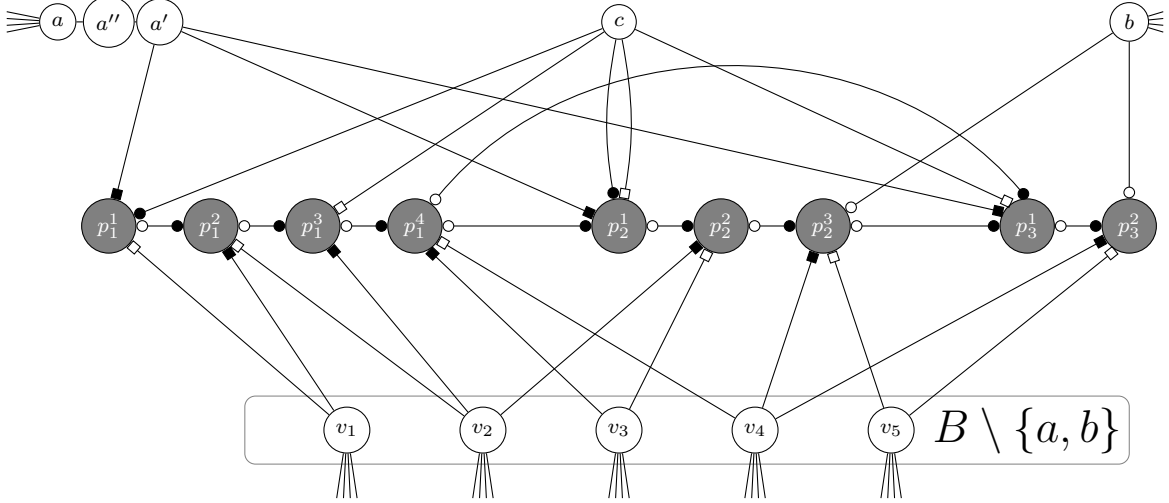
Fig. 7: Example of the fingerprint gadget. Here $B = \{v_1, \ldots, v_5, a, b\}$ and m_f equals one for $f = f_i = (d_i, M_i)$ (described in the next line), and zero otherwise. The vertices p_i^1 represent fingerprint f_i. Here $d_1^{-1}(2) = \{v_1, v_2\}$, $M_1 = \{ab, v_3v_4\}$; $d_2^{-1}(2) = \emptyset$, $M_2 = \{ab, v_2v_3, v_4v_5\}$; $d_3^{-1}(2) = \emptyset$, $M_3 = \{ab, v_4v_5\}$. The vertices p_i^j are label gadgets, and the symbols denote the labels $1, 2, 3, 4$ as in Figure 6.

with respect to p_i^2. By propagation it follows that in H we have that for every p_i^j the two edges incident to p_i^j must have label 3 and 4 with respect to p_i^j.

Now we focus on $p_x^1, \ldots, p_x^{\ell_x}$. We see that p_x^2 (if it exists) has only one incident edge with label 3 (with respect to p_x^2) which is to p_x^1. Therefore the edges of H incident to p_x^2 must have labels 1 and 2, and the same holds by propagation for all p_x^j. It follows that H has fingerprint f_x: every vertex in $d_x^{-1}(2)$ has indeed degree 2 as it is incident to two edges in the edges created in Step 3c from the set E_x; no edges incident to vertices in $d_x^{-1}(0)$ occur in H as they do not occur in E_x so they are not adjacent to the vertices p_x^i, which are the only vertices with incident edges with label 1 or 2. Moreover, for all edges in $uv \in M_i \setminus ab$ we see H has the path $up_x^j, p_x^j v$ as the edges incident to p_x^j must have labels 1 and 2.

Summarizing, we saw that H must contain the paths $p_i^1, \ldots, p_i^{\ell_i}$ for $i \neq x$ and all edges with labels 1 and 2 with respect to vertices p_x^j, and if it does it has the correct fingerprint if a and b are connected to each other. It remains to show that (without creating subcycles) the paths can be connected to one path from a to b visiting all vertices $V' \setminus (B \cup \{p_x^{1+|d_x^{-1}(2)|}, \ldots, p_x^{\ell_i})$ in a unique way. To see that this is the case, first note that a' is connected to c via the edges $e_x^1, \ldots, e_x^{k+1} \in E_i$ used in step 3b of the construction of G' where $k = |d_x^{-1}(2)|$, so they are also connected in G' via the vertices $p_x^1, \ldots, p_x^{k+1}$. To connect the paths $p_i^1, \ldots, p_i^{\ell_i}$ for $i \neq x$, note that $p_i^{\ell_i}$ can only be connected to vertices $p_{i'}^j$ with $i' > i$ as an incident edge of label 4 must be chosen. It follows that the only way to complete the paths is to connect

- c to p_1^1 if $x \neq 1$ or to p_2^1 if $x = 1$,

- $p_i^{\ell_i}$ to

    - p_{i+1}^1, if $x \neq i+1$ and $i < \ell$,

24

- p^1_{i+2}, if $x = i+1$ and $i < \ell - 1$,
- b, if $x = i+1$ and $i = \ell - 1$,
- to b if $i = \ell$,

which connects a to b, as required.

Path Decomposition and Size. The claimed size bound holds trivially, as all label gadgets have constant size. For the pathwidth bound, note that after removing $B \cup \{a, b\}$ the graph induced on the vertices p^j_i has constant pathwidth, as for every i, we have that $\{p^{\ell_i}_i, p^1_{i+1}\}$ forms a separator separating $p^j_{i'}$ with $i' \leq i$ from $p^j_{i'}$ with $i > i'$, and the graph between separators $\{p^{\ell_{i-1}}_i, p^1_i\}$ and $\{p^{\ell_i}_i, p^1_{i+1}\}$ has constant pathwidth, as it is a path of label gadgets, each of constant size. The required path decomposition can thus be obtained by including B in every bag. $\qquad\square$

5.4 The base case of Lemma 5.2

We prove Lemma 5.2 for $m = 1$, so the CNF-formula is a single clause C_1. The graph output by the reduction will consist of a graph G_i for $1 \leq i \leq q$ (recall $q = n/\gamma$). Each G_i contains 'left boundary vertices' L_i, and 'right boundary vertices' R_i, and additionally a 'top vertex' t_i and 'bottom vertex' b_i. The graphs G_i are glued together by unifying $b_i = t_{i-1}$ to get the graph G. The vertices t_i and b_i are used to propagate whether an encoded partial assignment has satisfied the clause. We now define the graph G_i. Recall from Subsection 5.2 that η_i is an injective function from $\mathcal{B}_r$ to $\{0, 1\}^{X_i}$.

Definition 5.6. *Let b_i, t_i be two vertices. The graph G_i is the following instantiation of a fingerprint gadget as implemented by Lemma 5.5 with boundary $B = L \cup R \cup \{b, t\}$, where we shorthand $L = L_i, R = R_i, t = t_i, b = b_i$. For every fingerprint $f_L = (d_L, M_L) \in \mathcal{B}_r$ on L, fingerprint $f_R = (d_R, M_R) \in \mathcal{B}_l$ on R and $d_b, d_t \in \{0, 2\}$, denote $f = (d_L \cup d_R \cup d_b \cup d_t, M^i_{alt})$, where $M^i_{alt} = M^i_L \cup M^i_S \setminus \{l_{i,1}l_{i,2}, r_{i,1}r_{i,3}\} \cup \{l_{i,1}r_{i,1}, l_{i,2}r_{i,3}\}$. For each such combinations we define $m_f = \mathbf{F}^{-1}[f_L, f_R]$ if at least one of the following conditions holds:*

1. *$\eta_i(f_L)$ is an assignment of X_i satisfying clause C_1, and*

   (a) *$d_t = d_b = 2$, or*
   (b) *both $d_b = 2$ and $d_t = 0$ hold.*

2. *$\eta_i(f_L)$ is an assignment of X_i not satisfying clause C_1, and*

   (a) *exactly one the propositions $d_t = 2$ and $d_b = 2$ holds.*

For other $f \in \mathcal{P}(B)$, set $m_f = 0$. Let G_i be the fingerprint gadget with boundary B for $(m_f)_{f \in \mathcal{P}_B}$.

Proof of Lemma 5.2 (base case). We claim that the graph G has the properties of Lemma 5.2. Let H be a partial solution which has fingerprint f^i_L on L_i and f^i_R on R_i. Suppose that for every $i = 1, \ldots, q$, the assignment $\eta_i(f^i_L)$ to X_i does not satisfy C_1. Then exactly q of the vertices $t_1, \ldots, t_q, b_q$ will have 2 edges from H incident to it so one vertex will have no incident edges in H and therefore H cannot be a partial solution and does not contribute to the count.

Otherwise, let i be the smallest integer such that the assignment $\eta_i(f^i_L)$ to X_i satisfies C_1. It follows that in fingerprint gadget G_i, H induces a partial solution for a fingerprint in which t_i and

b_i both have degree 2. Therefore, in the fingerprint gadget $G_{i'}$ with $i' < i$, H induces a partial solution for a fingerprint in which t_i has degree 2 and b_i has degree 0. Similarly in the fingerprint gadget $G_{i'}$ with $i' > i$, H induces a partial solution for a fingerprint in which t_i has degree 0 and b_i has degree 2. Therefore, in this case the number of combinations of partial solutions of $G_1, \ldots, G_q$ with the combined fingerprint is thus $(\mathbf{F}^{-1})^{\otimes q}[f_L, f_R] = (\mathbf{F}^{\otimes q})^{-1}[f_L, f_R]$ by the constructions of the graphs G_i, as required.

Pathwidth and Size. Let $L_0 = R_{q+1} = \emptyset$ for notational convenience. For $a = 0, \ldots, q+1$, define

$$S_{<a} = \bigcup_{i=1}^{a-1} L_i \cup R_i \cup \{b_i, t_i\}, \quad S_a = b_a \cup t_a \cup \bigcup_{i=a}^{q} L_i \cup \bigcup_{i=1}^{a} R_i, \quad S_{>a} = \bigcup_{i=a+1}^{q} L_i \cup R_i \cup \{b_i, t_i\}.$$

Note that S_a is a separator separating all vertices from $S_{<a}$ from all vertices from $S_{>a}$. Moreover, S_0 contains L and S_{q+1} contains R. We construct a path decomposition containing the bags $S_0, \ldots, S_{q+1}$ in this order. It remains to show that this can be completed into a efficient path decomposition by adding an appropriate path decompositions between bags S_a and S_{a+1}. To see this note that the vertices not in $S_{<a}$ and $S_{>a}$ must be in G_a and G_a admits a path decomposition of width $|B| + O(1) = \beta + O(1)$ starting in L_a and ending in R_a by Lemma 5.5. Thus in between bags S_a and S_{a+1} we can add bags with S_a and the path decomposition of G_a; after the last bag of this path decomposition we can forget all vertices of G_a except b_a and R_a which are contained in S_{a+1}. We obtain a path decomposition of width $q\beta + O(\beta)$, as required. $\qquad\square$

5.5 Putting things together to prove Lemma 5.1

Proof of Lemma 5.1. We first finish off the construction of G. Let $G' = G(L, R, \varphi)$ as in Lemma 5.2. Then do the following for $i = 1, \ldots, q$:

1. Add a fingerprint gadget G_i^R with boundary R_i to G' that has one partial solution for every fingerprint from $\mathcal{B}_r$ on R_i.

2. Add vertices b_i, t_i to G', where $b_i = t_{i+1}$ for $i < q$ and $t_q = b_1$.

3. Add a fingerprint gadget G_i^L with boundary $L_i \cup b_i \cup t_i$ to G' such that for every fingerprint $f_i^a = (d, M) \in \mathcal{B}_l$ on L_i, G' has $\sum_{f^l \in \mathcal{B}_r} (\mathbf{F}^{\otimes q})^{-1}[f^l, f^a]$ partial solutions for the fingerpint

$$f = (d', M \setminus \{l_{i,1}l_{i,3}\} \cup \{l_{i,1}t_i, l_{i,3}b_i\}),$$

on $L_i \cup b_i \cup t_i$. Here d' equals d with the addition that $d'(b_i) = d'(t_i) = 1$.

Number of solutions equals number of Hamiltonian cycles. Analogously to Claim 1, we first show the following:

Claim 2. *The number of Hamiltonian cycles of G equals*

$$\sum_{f^a, f^c \in \mathcal{B}_l^q, f^l, f^b, f^d \in \mathcal{B}_r^q} (\mathbf{F}^{\otimes q})^{-1}[f^l, f^a]\mathbf{F}^{\otimes q}[f^a, f^b]\mathbf{A}[f^b, f^c]\mathbf{F}^{\otimes q}[f^c, f^d]. \tag{2}$$

Proof. Denote G^L and G^R for the union of the graphs G_i^L and G_i^R, respectively. Let H be an edgeset with fingerprint constructed from $f^a = (f_1^a, \ldots, f_q^a)$ in Step 3 on L in G^L, and fingerprint f_L on L in G'. By the construction of G', the fingerprint f_L must be obtained from a fingerprint f_a by altering the matchings such that f_a and f_a match (indeed, as before, otherwise in one block two cycles will be created and one of them will contain the edge $\{l_{i,1}, l_{i,3}\}$ and thus not be altered so it remains a subcycle). Similarly, if H has fingerprint f_d on R in G^R and fingerprint f_R on R in G', f_R must be obtained from a fingerprint f_c by altering the matchings such that f_c and f_d match.

Conversely we claim that, if H is an edgeset with fingerprints constructed from f^a in Step 3 on L in G^L, f_L on L in G', f_d on R in G^R and a fingerprint obtained from f^c on R in G' by altering the matching, and both f^a and f^b as well as f^c and f^d match, then H is automatically a Hamiltonian cycle. To see this, first take into account the partial solution in G' and G^R. Similarly as in the proof of Claim 1, it is easily seen that this gives a set of paths that connect $l_{i,1}$ with $l_{i,2}$ for every i as the fingerprints f^c and f^d match. Taking also the partial solutions in G^L into account and that f^a and f^b match, we see that within each block G_L^i the partial solutions connect t_i to b_i (denoting $t_i = b_q$) and therefore the partial solutions combined give a Hamiltonian cycle.

By the way we constructed G_i^L the number of partial solutions that have fingerprint f^a on L is multiplied with $\sum_{f^l \in \mathcal{B}_r} (\mathbf{F}^{\otimes q})^{-1} [f^l, f^a]$. Therefore, by summing over all fingerprints f^a, f^b, f^c, f^d and counting number of partial solutions with these fingerprints as described in the fingerprint gadgets we obtain that the number of Hamiltonian cycles equals the claimed quantity. $\qquad\square$

By the definition of $\mathbf{A}$ from Lemma 5.2, $\mathbf{A} = \mathbf{C}(\mathbf{F}^{\otimes q})^{-1}$ with

$$\mathbf{C}[f, f] = \begin{cases} 1, & \text{if } \eta(f) \models \varphi, \\ 0, & \text{otherwise.} \end{cases}$$

Thus, in matrix terms, (2) can be rewritten into

$$\mathbf{1}^T (\mathbf{F}^{\otimes q})^{-1} \mathbf{F}^{\otimes q} \mathbf{C} (\mathbf{F}^{\otimes q})^{-1} \mathbf{F}^{\otimes q} \mathbf{1} = \mathbf{1}^T \mathbf{C} \mathbf{1},$$

which is easily seen to be the number of assignments of X that satisfy φ modulo p, as required.

Pathwidth bound. Recall the graph G' has pathwidth $q\beta + O(\beta)$. It is easy to see that the additions of the graphs G_i^L and G_i^R do not increase the pathwidth beyond this bound: We can simply introduce and forget each G_i^R separately at the end of the path decomposition. As similar approach can be used for G_i^L in the start of the path decomposition where we each time only forget the top vertex (except t_1). Recall that $q = n/\gamma$. We will now set the parameters β and γ. We first show that we can find the needed sufficiently large sets $\mathcal{B}_l$ and $\mathcal{B}_r$:

Lemma 5.7. *Let $rk_p \in \mathbb{R}$ be such that $\log_{rk_p}(\text{rank}(\mathbf{M}_k))/k \to c$, where $c \geq 1$, as even k tends to infinity. Then for any ε', there exists a large enough β and sets $\mathcal{B}_l, \mathcal{B}_r$ of fingerprints on $[\beta]$ of size at least $(2 + rk_p - \varepsilon')^\beta$ such that $\mathbf{H}_\beta[\mathcal{B}_l, \mathcal{B}_r]$ has full rank over p and if $(d_L, M_L) \in \mathcal{B}_l$ and $(d_R, M_R) \in \mathcal{B}_r$ then (i) $d_l(1), d_r(1), d_l(2)$ and $d_r(3)$ all equal 1, and (ii) $\{1,3\} \in M_l$ and $\{1,2\} \in M_r$.*

Proof. We have $\text{rank}(\mathbf{H}_\beta) \geq (2 + r_p - \varepsilon')^\beta$ for large enough β by Fact 2.2, and by the same binomial theorem argument as in Fact 2.2 we also have that sets of linearly independent rows and columns exist consisting of fingerprints (d, M) satisfying $|d^{-1}(1)| \geq 4$. Then there must be a vertex matched to the same vertices in the row index M_1 and column index M_2 in at least an $1/\beta^2$ fraction of both

basis matchings by the pigeon principle. Denoting this vertex and its two frequent neighbors with $1, 2, 3$ the claim follows. $\square$

Let $\hat{\varepsilon} = \varepsilon/2$, and pick β sufficiently large such that Lemma 5.7 ensures the sets $\mathcal{B}_l, \mathcal{B}_r$ of fingerprints on $[\beta]$ of size at least $(2 + \mathrm{rk}_p - \hat{\varepsilon})^\beta$ exist. To ensure the existence of the injective encoding functions $\eta_1, \ldots, \eta_q$, we pick γ such that $2^\gamma < (2 + \mathrm{rk}_p - \hat{\varepsilon})^\beta$; set γ as large as possible under this constraint so that $\gamma \geq \beta \log_2(2 + \mathrm{rk}_p - \hat{\varepsilon}) - 1$. The pathwidth of G will be at most

$$q\beta + O(\beta) = (n/\gamma)\beta + O(\beta) \leq n/(\log_2(2 + \mathrm{rk}_p - \hat{\varepsilon}) - 1/\beta) + O(\beta).$$

The running time of the assumed algorithm for counting the number of Hamiltonian cycles modulo p of the created instance will thus be $O^*(2^{\alpha n + O(\beta)}) = O^*(2^{\alpha n})$ time where

$$\alpha = \frac{\log_2(2 + \mathrm{rk}_p - \varepsilon)}{\log_2(2 + \mathrm{rk}_p - \hat{\varepsilon}) - 1/\beta},$$

which is smaller than 1 for sufficiently large β (which may depend on ε). Finally, it can be easily checked that the graph G can trivially be constructed in time polynomial in the size of φ and p for constant ε. $\square$

6 Conclusions

As future work, we suggest the problem of counting the connected induced subgraphs of a graph, where one could try to exclude $O^*(2^{o(\mathrm{pw} \log \mathrm{pw})})$ time algorithms. The connection matrix for this problem is the meet matrix of the partition lattice, ordered by refinement, so that the coarsest partition with one block is the smallest element. For this setting, the powerful technology of Möbius functions (see e.g. [28]) can readily give rank lower bounds, but it is not a priori clear how to construct the gadgets required for converting the rank bound into an algorithmic lower bound. Another example could be the problem of counting Steiner Trees, which has an $O^*(5^{\mathrm{pw}})$ time algorithm from [7], or the evaluation of graph polynomials such as the Tutte polynomial.

A further natural direction for future research is to find the optimal constant c_p such that #HC modulo p can be solved in $O^*(c_p^{\mathrm{pw}})$ time and not in $O^*((c_p - \varepsilon)^{\mathrm{pw}})$ time for $\varepsilon > 0$. It is natural to conjecture that $c_p = 2 + r_p$, where r_p is the exponential base of the rank of $\mathbf{M}_k$ over $\mathbb{Z}_p$. However, note that in [11], obtaining an algorithm from the rank upper bound was not trivial and, unlike the lower bound from Theorem 1.3, it is not a priori clear how to use rank *upper bounds* as a black box. The main reason is that we cannot seem to reduce the work related to $\mathbf{M}$ to constant-sized copies as done in the proof of Theorem 1.3. Additionally, we need better lower bounds for the rank of $\mathbf{M}_k$ over $\mathbb{Z}_p$, since the bounds from our paper are tight only over $\mathbb{Q}$.

Acknowledgments

References

[1] Josh Alman and Ryan Williams. Probabilistic rank and matrix rigidity. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pages 641–652, New York, NY, USA, 2017. ACM.

[2] Regev Amitai. Identities for the number of standard Young tableaux in some (k, l)-hooks. *Séminaire Lotharingien de Combinatoire*, 63(3), 2010.

[3] Eiichi Bannai and Tatsuro Ito. *Algebraic combinatorics I: association schemes*. Mathematics lecture note series. Benjamin/Cummings Pub. Co., 1984.

[4] Andreas Björklund. Determinant sums for undirected hamiltonicity. *SIAM J. Comput.*, 43(1):280–299, 2014.

[5] Andreas Björklund and Thore Husfeldt. The parity of directed hamiltonian cycles. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 727–735. IEEE Computer Society, 2013.

[6] Andreas Björklund, Petteri Kaski, and Ioannis Koutis. Directed hamiltonicity and out-branchings via generalized laplacians. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland*, pages 91:1–91:14, 2017.

[7] Hans L. Bodlaender, Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Inf. Comput.*, 243:86–111, 2015.

[8] Peter Bürgisser, Michael Clausen, and Mohammad Amin Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1997.

[9] Chris Calabro, Russell Impagliazzo, Valentine Kabanets, and Ramamohan Paturi. The complexity of unique k-SAT: An isolation lemma for k-CNFs. *J. Comput. Syst. Sci.*, 74(3):386–393, 2008.

[10] Radu Curticapean and Mingji Xia. Parameterizing the permanent: Genus, apices, minors, evaluation mod 2^k. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 994–1009, 2015.

[11] Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Fast hamiltonicity checking via bases of perfect matchings. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 301–310, 2013.

[12] Marek Cygan, Jesper Nederlof, Marcin Pilipczuk, Michał Pilipczuk, Johan M. M. van Rooij, and Jakub Onufry Wojtaszczyk. Solving connectivity problems parameterized by treewidth in single exponential time. In Rafail Ostrovsky, editor, *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 150–159. IEEE Computer Society, 2011.

[13] Fedor V. Fomin, Daniel Lokshtanov, Fahad Panolan, and Saket Saurabh. Efficient computation of representative families with applications in parameterized and exact algorithms. *J. ACM*, 63(4):29:1–29:60, 2016.

[14] Chris Godsil and K. Meagher. *Erdos-Ko-Rado Theorems: Algebraic Approaches*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2015.

[15] Chris Godsil and Gordon Royle. *Algebraic graph theory*, volume 207 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2001.

[16] Jozef Gruska, editor. *Mathematical Foundations of Computer Science 1977, 6th Symposium, Tatranska Lomnica, Czechoslovakia, September 5-9, 1977, Proceedings*, volume 53 of *Lecture Notes in Computer Science*. Springer, 1977.

[17] Godfrey H. Hardy and Edward M. Wright. *An Introduction to the Theory of Numbers*. Oxford University Press, Oxford, sixth edition, 2008. Revised by D. R. Heath-Brown and J. H. Silverman, With a foreword by Andrew Wiles.

[18] Michael Held and Richard M Karp. A dynamic programming approach to sequencing problems. *Journal of the Society for Industrial and Applied Mathematics*, 10(1):196–210, 1962.

[19] Russel Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.

[20] Tomer Kotek and Johann A. Makowsky. Connection matrices and the definability of graph parameters. In *Computer Science Logic (CSL'12) - 26th International Workshop/21st Annual Conference of the EACSL, CSL 2012, September 3-6, 2012, Fontainebleau, France*, pages 411–425, 2012.

[21] Stefan Kratsch and Magnus Wahlström. Compression via matroids: A randomized polynomial kernel for odd cycle transversal. *ACM Trans. Algorithms*, 10(4):20:1–20:15, 2014.

[22] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, New York, NY, USA, 1997.

[23] Nathan Lindzey. Erdős–Ko–Rado for perfect matchings. *European Journal of Combinatorics*, 65:130 – 142, 2017.

[24] Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Known algorithms on graphs on bounded treewidth are probably optimal. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 777–789. SIAM, 2011.

[25] László Lovász. The rank of connection matrices and the dimension of graph algebras. *Eur. J. Comb.*, 27(6):962–970, 2006.

[26] László Lovász. *Large Networks and Graph Limits*, volume 60 of *Colloquium Publications*. American Mathematical Society, 2012.

[27] László Lovász and Michael E. Saks. Lattices, Möbius functions and communication complexity. In *29th Annual Symposium on Foundations of Computer Science, White Plains, New York, USA, 24-26 October 1988*, pages 81–90. IEEE Computer Society, 1988.

[28] László Lovăsz and Michael Saks. Communication complexity and combinatorial lattice theory. *Journal of Computer and System Sciences*, 47(2):322 – 349, 1993.

[29] Ian G. Macdonald. *Symmetric functions and Hall polynomials*. Oxford mathematical monographs. Clarendon Press, 1995.

[30] Jiří Matoušek. *Thirty-three miniatures: Mathematical and Algorithmic applications of Linear Algebra*, volume 53. American Mathematical Soc., 2010.

[31] Burkhard Monien. How to find long paths efficiently. *North-Holland Mathematics Studies*, 109:239 – 254, 1985. Analysis and Design of Algorithms for Combinatorial Problems.

[32] Marko Petkovsek, Herbert S. Wilf, and Doron Zeilberger. *A = B*. A K Peters Series. Taylor & Francis, 1996.

[33] Ran Raz and Boris Spieker. On the "log rank"-conjecture in communication complexity. *Combinatorica*, 15(4):567–588, 1995.

[34] Bruce Sagan. *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*. Graduate Texts in Mathematics. Springer New York, 2001.

[35] Richard P. Stanley. *Enumerative Combinatorics:*, volume 2 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 2001.

[36] Robert M. Thrall. On symmetrized Kronecker powers and the structure of the free Lie ring. *American Journal of Mathematics*, 64(1):pp. 371–388, 1942.

[37] Ryan Williams. The polynomial method in circuit complexity applied to algorithm design (invited talk). In Venkatesh Raman and S. P. Suresh, editors, *34th International Conference on Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, December 15-17, 2014, New Delhi, India*, volume 29 of *LIPIcs*, pages 47–60. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014.

[38] Michał Włodarczyk. Clifford algebras meet tree decompositions. In Jiong Guo and Danny Hermelin, editors, *11th International Symposium on Parameterized and Exact Computation, IPEC 2016, August 24-26, 2016, Aarhus, Denmark*, volume 63 of *LIPIcs*, pages 29:1–29:18. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.

A   Proof of Theorem 1.1 from Theorem 1.2 and Theorem 1.3

We prove a slightly stronger consequence, namely, that there is an algorithm that counts the satisfying assignments of a given a CNF-formula on n variables in $O^*((2 - \varepsilon)^n)$ time for some $\varepsilon > 0$.

Let n be the number of variables of the given CNF-formula φ. The Chinese Remainder Theorem (CRT) tells us that given the number of satisfying assignments of φ modulo primes $p_1, \ldots, p_\ell$, we can compute the number of satisfying solutions of φ as long as $\prod_{i=1}^{\ell} p_i \geq 2^n$. By the Prime Number Theorem [17, p. 494, Eq. (22.19.3)], there are at least $r/\log_2 r$ primes between r and $2r$, and thus

$$\prod_{r \leq p \text{ prime} \leq 2r} p \geq r^{r/\log r} \geq 2^{\Omega(r)}.$$

It follows that for counting the number of satisfying assignments of a given CNF-formula, it is sufficient to count the number of satisfying assignments modulo p for any $p = \Theta(n)$. We do this using Lemma 5.1 combined with the algorithm for #HC. For fixed t we have that $\mathrm{rank}_p(\mathbf{M}_t) = \mathrm{rank}(\mathbf{M}_t)$ for large enough p (which can for example be shown by upper bounding the determinant of $\mathbf{M}_t$ by $t!$). The assumed algorithm for #HC also counts the number of Hamiltonian cycles modulo p. By Theorem 1.3 we have $\lim_{p \to \infty} r_p = 4$ and Theorem 1.1 follows.

B   Finite Group Representation Theory

For a set X, let $\mathbb{C}[X]$ denote the vector space of dimension $|X|$ of complex-valued functions from X. A *representation* (ϕ, V) of a finite group G is a homomorphism $\phi : G \to \mathbb{GL}(V)$ where $\mathbb{GL}(V)$ is the group of $\dim V \times \dim V$ invertible matrices. We refer to (ϕ, V) simply as ϕ when V is understood, or as V when ϕ is understood. For any representation ϕ, we define its *dimension* to be $\dim \phi := \dim V$. Two representations ρ, ϕ are *equivalent* if $\rho(g)$ and $\phi(g)$ are similar for all $g \in G$.

Let (ϕ, V) be a representation of a finite group G, and let $W \leq V$ be a *G-invariant* subspace, that is, $\phi(g)w \in W$ for all $w \in W$ and for all $g \in G$. We say that $(\phi|_W, W)$ is a *sub-representation* of ϕ where $\phi|_W$ is the restriction of ϕ to the subspace W. A representation (ϕ, V) is an *irreducible representation* (or simply, an *irreducible*) if it has no proper sub-representations.

It is well-known that there is a one-to-one correspondence between the set of inequivalent irreducibles of G and its conjugacy classes $\mathcal{C}$, and that any representation V of G uniquely decomposes as a finite direct sum of inequivalent irreducibles V_i of G:

$$V \cong \bigoplus_{i=1}^{|\mathcal{C}|} m_i \, V_i$$

where m_i is the *multiplicity* of V_i (the number of times that V_i occurs in the decomposition). Natural representations of groups can be obtained by letting them act on sets. In particular, for any group G acting on a set X, let $(\phi, \mathbb{C}[X])$ be the *permutation representation* of G on X defined such that

$$\phi(g)[f(x)] = f(g^{-1}x)$$

for all $g \in G$, $f \in \mathbb{C}[X]$, and $x \in X$. If we let G act on itself ($X = G$), then we obtain the *regular representation*, which admits the following decomposition into irreducibles:

$$\mathbb{C}[G] \cong \bigoplus_{i=1}^{|\mathcal{C}|} (\dim \phi_i) \, V_i$$

where (ϕ_i, V_i) is the ith irreducible of G. Letting $e_g e_h = e_{gh}$ over the standard basis $\{e_g\}_{g \in G}$ of $\mathbb{C}[G]$, we see that $\mathbb{C}[G]$ is an algebra, the so-called *group algebra* of G.

For any (irreducible) representation ϕ of G, the *(irreducible) character* χ_ϕ of ϕ is the map $\chi_\phi : G \to \mathbb{C}$ such that $\chi_\phi(g) := \mathrm{Tr}(\phi(g))$. Similar matrices have the same trace, thus the character of a representation is a *class function*, that is, they are constant on conjugacy classes. Furthermore, the characters of the set of all irreducible representations of a group G form an orthonormal basis for the space of all class functions of $\mathbb{C}[G]$.