

Fast Algorithms for Parameterized Problems with Relaxed Disjointness Constraints

Ariel Gabizon¹, Daniel Lokshтанov², and Michał Pilipczuk³

¹ Department of Computer Science, Technion, Israel.
ariel.gabizon@gmail.com

² Department of Informatics, University of Bergen, Norway.
daniello@ii.uib.no

³ Institute of Informatics, University of Warsaw, Poland.
michal.pilipczuk@mimuw.edu.pl

Abstract. In this paper we consider generalized versions of four well-studied problems in parameterized complexity and exact exponential time algorithms: k -PATH, SET PACKING, MULTILINEAR MONOMIAL TESTING and HAMILTONIAN PATH. The generalization is in every case obtained by introducing a *relaxation parameter*, which relaxes the constraints on feasible solutions. For example, the k -PATH problem is generalized to r -SIMPLE k -PATH where the task is to find a walk of length k that never visits any vertex more than r times. This problem was first considered by Abasi et al. [1]. HAMILTONIAN PATH is generalized to DEGREE BOUNDED SPANNING TREE, where the input is a graph G and integer d , and the task is to find a spanning tree T of G such that every vertex has degree at most d in T .

The generalized problems can easily be shown to be NP-complete for every fixed value of the relaxation parameter. On the other hand, we give algorithms for the generalized problems whose worst-case running time (a) *matches* the running time of the best algorithms for the original problems up to constants in the exponent, and (b) *improves* significantly as the relaxation parameter increases. For example, we give a deterministic algorithm with running time $O^*(2^{O(k \frac{\log r}{r})})$ for r -SIMPLE k -PATH matching up to a constant in the exponent the randomized algorithm of Abasi et al. [1], and a randomized algorithm with running time $O^*(2^{O(n \frac{\log d}{d})})$ for DEGREE BOUNDED SPANNING TREE improving upon an $O(2^{n+o(n)})$ algorithm of Fomin et al. [10].

On the way to obtain our results we generalize the notion of *representative sets* to multisets, and give an efficient algorithm to compute such representative sets. Both the generalization of representative sets to multisets and the algorithm to compute them may be of independent interest.

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 257575 and 240258. Mi. Pilipczuk is currently holding a post-doc position at Warsaw Center of Mathematics and Computer Science and is supported by Polish National Science Centre grant DEC-2013/11/D/ST6/03073.

1 Introduction

Many of the combinatorial optimization problems studied in theoretical computer science are idealized mathematical models of real-world problems. When the simplest model is well-understood, it can be enriched to better capture the real-world problem one actually wants to solve. Thus it comes as no surprise that many of the well-studied computational problems generalize each other: the CONSTRAINT SATISFACTION PROBLEM generalizes SATISFIABILITY, the problem of finding a spanning tree of maximum degree at most d generalizes HAMILTONIAN PATH, while the problem of packing sets of size 3 generalizes packing sets of size 2, also known as the MAXIMUM MATCHING problem.

By definition, the generalized problem is computationally harder than the original. However it is sometimes the case that the most difficult instances of the generalized problem are actually instances of the original problem. In other words, the “further away” an instance of the generalized problem is from being an instance of the original, the easier the instance is. Abasi et. al [1] initiated the study of this phenomenon in parameterized complexity (we refer the reader to the textbooks [6, 7, 9, 21] for an introduction to parameterized complexity). In particular, they study the r -SIMPLE k -PATH problem. Here the input is a graph G , and integers k and r , and the objective is to determine whether there is an r -simple k -path in G , where an r -simple k -path is a sequence v_1, v_2, \dots, v_k of vertices such that every pair of consecutive vertices is adjacent and no vertex of G is repeated more than r times in the sequence. Observe that for $r = 1$ the problem is exactly the problem of finding a simple path of length k in G . On the other hand, for $r = k$ the problem is easily solvable in polynomial time, as one just has to look for a walk in G of length k . Thus, gradually increasing r from 1 to k should provide a sequence of computational problems that become easier as r increases. Abasi et al. [1] confirm this intuition by giving a randomized algorithm for r -SIMPLE k -PATH with running time $O(r^{2k/r} n^{O(1)})$ for any $2 \leq r \leq k$.

In this paper we continue the investigation of algorithms for problems with a *relaxation parameter* r that interpolates between an NP-hard and a polynomial time solvable problem. We show that in several interesting cases one can get a sequence of algorithms with better and better running times as the relaxation parameter r increases, essentially providing a smooth transition from the NP hard to the polynomial time solvable case.

Our main technical contribution is a new algorithm for the (r, k) -MONOMIAL DETECTION problem. Here the input is an arithmetic circuit C that computes a polynomial f of degree k in n variables x_1, \dots, x_n . The task is to determine whether the polynomial f has a monomial $\prod_{i=1}^n x_i^{a_i}$, such that $0 \leq a_i \leq r$ for every $i \leq n$. The main result of Abasi et al. [1] is a randomized algorithm for (r, k) -MONOMIAL DETECTION with running time $O(r^{2k/r} |C| n^{O(1)})$, and their algorithm for r -SIMPLE k -PATH is obtained using a reduction to (r, k) -MONOMIAL DETECTION. We give a *deterministic* algorithm for the problem with running time $r^{O(k/r)} |C| n^{O(1)}$ in the case when the circuit C is *non-canceling*. Formally, this means that the circuit contains only variables at its leaves (i.e., no constants) and only addition and multiplication gates (i.e., no subtraction gates).

Informally, all monomials of the polynomials computed at intermediate gates of C contribute to the polynomial computed by C . Before stating our theorem, we remark that the formal definitions of problems and notations involved in the algorithmic results presented in the introduction can be found in Appendix B. We use the notation O_k in our theorems to hide $k^{O(1)}$ terms.

Theorem 1 *Given a non-canceling circuit C computing a polynomial $f(X_1, \dots, X_n) \in \mathbb{Z}[X_1, \dots, X_n]$, (r, k) -MONOMIAL DETECTION can be solved in deterministic time $O_k(|C| \cdot r^{18k/r} \cdot 2^{O(k/r)} \cdot n \log^3 n)$.*

Comparing our algorithm with the algorithm of Abasi et al. [1], our algorithm is slower by a constant factor in the exponent of r , and only works for non-canceling circuits. However our algorithm is *deterministic* (while the one by Abasi et al. is randomized) and also works for the *weighted* variant of the problem, while the one by Abasi et al. does not. In the weighted variant each variable x_i has a non-negative integer weight w_i , and the weight of a monomial $\prod_{i=1}^n x_i^{a_i}$ is defined as $\sum_{i=1}^n w_i a_i$. The task is to determine whether there exists a monomial $\prod_{i=1}^n x_i^{a_i}$, such that $0 \leq a_i \leq r$ for every $i \leq n$, and if so, to return one of minimum weight. As a direct consequence we obtain the first deterministic algorithm for r -SIMPLE k -PATH with running time $r^{O(k/r)} |C| n^{O(1)}$, and the first algorithm with such a running time for weighted r -SIMPLE k -PATH.

Theorem 2 (WEIGHTED) *r -SIMPLE k -PATH can be solved in deterministic time $O_k(r^{12k/r} \cdot 2^{O(k/r)} \cdot n^3 \cdot \log n)$.*

Here, by WEIGHTED r -SIMPLE k -PATH we mean the variant where the weights are on vertices, and the weight of an r -simple k -path is the sum of the weights of traversed vertices, including multiplicities. However, we can also solve the edge-weighted variant at a cost of a larger constant in the exponent.

The significance of an in-depth study of (r, k) -MONOMIAL DETECTION, is that it is the natural “relaxation parameter”-based generalization of the fundamental MULTI-LINEAR MONOMIAL DETECTION problem. The MULTI-LINEAR MONOMIAL DETECTION problem is simply (r, k) -MONOMIAL DETECTION with $r = 1$. A multitude of parameterized problems reduce to MULTI-LINEAR MONOMIAL DETECTION [4, 6, 11, 16, 26]. Thus, obtaining good algorithms (r, k) -MONOMIAL DETECTION is an important step towards efficient algorithms for relaxation parameter-variants of these problems. For some problems, such as k -PATH, efficient algorithms for the relaxation parameter variant (i.e r -SIMPLE k -PATH) follow directly from the algorithms for (r, k) -MONOMIAL DETECTION. In this paper we give two more examples of fundamental problems for which efficient algorithms for (r, k) -MONOMIAL DETECTION lead to efficient algorithms for their “relaxation parameter”-variant.

Our first example is the (r, p, q) -PACKING problem. Here the input is a family \mathcal{F} of sets of size q over a universe of size n , together with integers r and p . The task is to find a subfamily $\mathcal{A} \subseteq \mathcal{F}$ of size at least p such that every element of the universe is contained in at most r sets in \mathcal{A} . Observe that (r, p, q) -PACKING is the relaxation parameter variant of the classic SET PACKING problem ((r, p, q) -PACKING with $r = 1$). We give an algorithm for (r, p, q) -PACKING with running

time $2^{O(pq \cdot \frac{\log r}{r})} |\mathcal{F}| n^{O(1)}$. For $r = 1$ this matches the best known algorithm [4] for SET PACKING, up to constants in the exponent, and when r grows our algorithm is significantly faster than $2^{pq} |\mathcal{F}| n^{O(1)}$. Just as for r -SIMPLE k -PATH, our algorithm also works for weighted variants. We remark that (r, p, q) -PACKING was also studied by Fernau et al. [8] from the perspective of kernelization.

Theorem 3 *Let $k \triangleq p \cdot q$. Then (WEIGHTED) (r, p, q) -PACKING can be solved in deterministic time $O_k(|\mathcal{F}| \cdot r^{12k/r} \cdot 2^{O(k/r)} \cdot n \log^2 n)$.*

Again, by WEIGHTED (r, p, q) -PACKING we mean a variant where elements are assigned weights and the weight of a set is equal to the sum of the weights of its elements. However, at a cost of a larger constant in the exponent we can also solve the variant where each set is assigned its own weight.

Our second example is the DEGREE-BOUNDED SPANNING TREE problem. Here, we are given as input a graph G and integer d , and the task is to determine whether G has a spanning tree T whose maximum degree does not exceed d . For $d = 2$ this problem is equivalent to HAMILTONIAN PATH, and hence the problem is NP-complete in general, but for $d = n - 1$ it boils down to checking the connectedness of G . Thus, DEGREE-BOUNDED SPANNING TREE can be thought of as a relaxation parameter variant of HAMILTONIAN PATH. The problem has received significant attention in the field of approximation algorithms: there are classic results of Fürer and Raghavachari [13], Goemans [15], and of Singh and Lau [24] that give additive approximation algorithms for the problem and its weighted variant. From the point of view of exact algorithms, the currently fastest exact algorithm, working for any value of d , is due to Fomin et al. [10] and has running time $O(2^{n+o(n)})$. In this work, we give a randomized algorithm for DEGREE-BOUNDED SPANNING TREE with running time $2^{O(n \frac{\log d}{d})}$, by reducing the problem to an instance of (r, k) -MONOMIAL DETECTION. Thus, our algorithm significantly outperforms the algorithm of Fomin et al. [10] for all super-constant d , and runs in polynomial time for $d = \Omega(n)$. Interestingly, the instance of (r, k) -MONOMIAL DETECTION that we create crucially uses subtraction, since the constructed circuit computes the determinant of some matrix. Thus we are not able to apply our algorithm for non-canceling circuits, and have to resort to the randomized algorithm of Abasi et al. [1] instead. Obtaining a deterministic algorithm for DEGREE-BOUNDED SPANNING TREE that would match the running time of our algorithm, or extending the result to the weighted setting, remains as an interesting open problem.

Our methods. The starting point for our algorithms is the notion of *representative sets*. If \mathcal{A} is a family of sets, with all sets in \mathcal{A} having the same size p , we say that a subfamily $\mathcal{A}' \subseteq \mathcal{A}$ *q-represents* \mathcal{A} if for every set B of size q , whenever there exists a set $A \in \mathcal{A}$ such that A is disjoint from B , then there also exists a set $A' \in \mathcal{A}'$ such that A' is disjoint from B .

Representative sets were defined by Monien [18], and were recently used to give efficient parameterized algorithms for a number of problems [11, 12, 22, 23, 27, 28], including k -PATH [11, 12, 23], SET PACKING [23, 27, 28] and MULTI-LINEAR MONOMIAL DETECTION [11]. It is therefore very tempting to try to use

representative sets also for the relaxation parameter variants of these problems. However, it looks very hard to directly use representative sets in this setting. On a superficial level the difficulty lies in that representative sets are useful to guarantee *disjointness*, while the solutions to the relaxation parameter variants of the considered problems may self-intersect up to r times.

We overcome this difficulty by generalizing the notion of representative sets to multisets. When taking the union of two multisets A and B , an element that appears a times in A and b times in B will appear $a+b$ times in the union $A+B$. Thus, if two regular sets A and B are viewed as multisets, they are disjoint if and only if no element appears more than once in $A+B$. We can now relax the notion of disjointness and require that no element appears more than r times in $A+B$. Specifically, if \mathcal{A} is a family of multisets, with all multisets in \mathcal{A} having the same size p (counting duplicates), we say that a subfamily $\mathcal{A}' \subseteq \mathcal{A}$ *q-represents* \mathcal{A} if the following condition is satisfied. For every multiset B of size q , whenever there exists an $A \in \mathcal{A}$ such that no element appears more than r times in $A+B$, there also exists an $A' \in \mathcal{A}'$ such that no element appears more than r times in $A'+B$. The majority of the technical effort in the paper is spent on proving that every family \mathcal{A} of multisets has a relatively small q -representative family \mathcal{A}' in this new setting, and to give an efficient algorithm to compute \mathcal{A}' from \mathcal{A} . The formal statement of this result can be found in Corollary 2.

On the way to develop our algorithm for computing representative sets of multisets, we give a new construction of a pseudo-random object called *lopsided universal sets*. Informally speaking, an (n, p, q) -*lopsided universal set* is a set of strings such that, when focusing on any $k \triangleq p+q$ locations, we see all patterns of hamming weight p . These objects have been of interest for a while in mathematics and in theoretical computer science under the name *Cover Free Families* (Cf. [5]). We give, for the first time, an explicit construction of an (n, p, q) -lopsided universal set whose size is only polynomially larger than optimal for all p and q . See Theorem 13 in Appendix A for a formal statement.

Both our algorithm for computing representative sets of multisets, and the new construction of lopsided universal sets may be of independent interest.

Outline of the paper. In Section 2 we give the necessary definitions and set up notational conventions. In Section 3 we give our construction of representative sets for multisets. This construction requires an auxiliary tool called minimal separating families. The construction of Section 3 assumes that an appropriate construction of minimal separating families is given as a black box, and this construction is deferred to Appendix A. Our new construction of lopsided universal sets is a corollary of the construction of minimal separating families, and is also explained in Appendix A. In Section 4 we briefly discuss the applications of representative sets to (r, k) -MONOMIAL DETECTION, (r, p, q) -PACKING and r -SIMPLE k -PATH (details deferred to Appendix B). In Section 5 we present our algorithm for DEGREE-BOUNDED SPANNING TREE. Finally, we conclude by discussing open problems and directions for future research in Section 6.

The full version of this paper, which contains a natural order of introducing the consecutive concepts and results, is available on arxiv [14].

2 Preliminaries

Notation. Throughout the paper, we use the notation O_k to hide $k^{O(1)}$ terms. We denote $[n] = \{1, 2, \dots, n\}$. For sets A and B , by $\{A \rightarrow B\}$ we denote the set of all functions from A to B . The notation \triangleq is used to introduce new objects defined by formulas on the right hand side.

Separating families. Recall that, for an integer $t \geq 1$, we say that a family of functions $\mathcal{H} \subseteq \{[n] \rightarrow [m]\}$ is a t -*perfect hash family*, if for every $C \subseteq [n]$ of size $|C| = t$ there is $f \in \mathcal{H}$ that is injective on C . We will be interested in constructing families of perfect hash functions that, in addition to being injective on a set C , have the property of sending another large set D to an output *disjoint* from the image of C . We call such a family of functions a *separating family*.

Definition 1 (Separating family). *Fix integers t, k, s, n such that $1 \leq t \leq n$. For disjoint subsets $C, D \subseteq [n]$, we say that a function $h: [n] \rightarrow [s]$ separates C from D if*

- h is injective on C ; and
- there are no collisions between C and D . That is, $h(C) \cap h(D) = \emptyset$.

A family of functions $\mathcal{H} \subseteq \{[n] \rightarrow [s]\}$ is (t, k, s) -*separating* if for every disjoint subsets $C, D \subseteq [n]$ with $|C| = t$ and $|D| \leq k - t$, there is a function $h \in \mathcal{H}$ that separates C from D .

We say that \mathcal{H} is (t, k, s) -*separating* with probability γ if for any fixed C and D with sizes as above, a function h chosen uniformly at random from \mathcal{H} separates C from D with probability at least γ .

For us, the most important case of separating families is when the range size is $|C| + 1$. In this case we use the term *minimal separating family*. It will also be convenient to assume in the definition that C is mapped to the first $|C|$ elements in the range.

Definition 2 (Minimal separating family) *A family of functions $\mathcal{H} \subseteq \{[n] \rightarrow [t + 1]\}$ is (t, k) -minimal separating if for every disjoint subsets $C, D \subseteq [n]$ with $|C| = t$ and $|D| \leq k - t$, there is a function $h \in \mathcal{H}$ such that*

- $h(C) = [t]$.
- $h(D) \subseteq \{t + 1\}$.

3 Multiset separators and representative sets

The purpose of this section is to formally define and construct representative sets for multisets. We will use, as an auxiliary result, an efficient construction of a small separating family. This result is actually the most technical part of this paper, and its proof is deferred entirely to the appendix.

Theorem 4 Fix integers n, t, k such that $1 \leq t \leq \min(n, k)$. Then a (t, k) -minimal separating family of size $O_k((k/t)^{2t} \cdot 2^{O(t)} \cdot \log n)$ can be constructed in time $O_k((k/t)^{2t} \cdot 2^{O(t)} \cdot n \cdot \log n)$.

Proof. (sketch) We first hash the set C of size t injectively into t ‘buckets’ using known constructions. At this stage we have ‘taken care’ of C but the set D of size $\leq k - t$ is scattered somehow among the t buckets. The novel idea is that a combination of guesses as to how the D has been scattered, together with the use of hitting sets for combinatorial rectangles [17] can now be used to separate C from D in a way that is not too costly. See Appendix A for full details. \square

Our primary tool for the construction of representative sets for multisets is what we call a *multiset separator* (see Definition 3). Informally, a multiset separator is a not too large set of ‘witnesses’ for the fact that two multisets of bounded size do not jointly contain too many repetitions per element.

Notation for multisets. Fix integers $n, r, k \geq 1$. We use $[r]_0$ to denote $\{0, \dots, r\}$. An r -set is a multiset A where each element of $[n]$ appears at most r times. It will be convenient to think of A as a vector in $[r]_0^n$, where A_i denotes the number of times i appears in A . We denote by $|A|$ the number of elements in A counting repetitions. That is, $|A| = \sum_{i=1}^n A_i$. We refer to $|A|$ as the *size* of A . An (r, k) -set is an r -set $A \in [r]_0^n$, where the number of elements with repetitions is at most k . That is, $|A| \leq k$.

Fix r -sets $A, B \in [r]_0^n$. We say that $A \leq B$ when $A_i \leq B_i$ for all $i \in [n]$. By $\bar{A} \in [r]_0^n$ we denote the ‘complement’ of r -set A , that is, $\bar{A}_i = r - A_i$ for all $i \in [n]$. By $A + B$ we denote the ‘union’ of A and B , that is, $(A + B)_i = A_i + B_i$ for all $i \in [n]$. Suppose now that A and B are (r, k) -sets. We say that A and B are (r, k) -compatible if $A + B$ is also an (r, k) -set, and $|A + B| = k$. That is, the total number of elements with repetitions in A and B together is k and any specific element $i \in [n]$ appears in A and B together at most r times. With the notation above at hand, we can define the central object needed for our algorithms.

Definition 3 (Multiset separator) Let \mathcal{F} be a family of r -sets. We say that \mathcal{F} is an (r, k) -separator if for any (r, k) -sets $A, B \in [r]_0^n$ that are (r, k) -compatible, there exists $F \in \mathcal{F}$ such that $A \leq F \leq \bar{B}$.

Construction of multiset separators. The following theorem shows how an (r, k) -separator can be constructed from a minimal separating family.

Theorem 5 Fix integers n, r, k such that $1 < r \leq k$, and let $t \triangleq \lfloor 2k/r \rfloor$. Suppose a (t, k) -minimal separating family $\mathcal{H} \subseteq \{[n] \rightarrow [t+1]\}$ can be constructed in time $f(r, k, n)$. Then an (r, k) -separator \mathcal{F} of size $|\mathcal{H}| \cdot (r+1)^t$ can be constructed in time $O_k(f(r, k, \max(n, t)) \cdot (r+1)^t)$.

Proof. (sketch) Given (r, k) -compatible multisets A and B , we wish to ‘separate’ them by a multiset F such that $A \leq F \leq \bar{B}$. Because of the compatibility, there can only be $2k/r$ ‘large indices’ - where $A_i > r/2$ or $B_i > r/2$. If we knew these

indices in advance, we could afford to try all values F_i to find a value such that $A_i \leq F_i \leq \overline{B}_i$. For any other $i \in [n]$, defining $F_i = r/2$ is fine. The problem is that there are many options for the set of large indices. However, a minimal separating family allows us to separate the set of large indices from others while trying a small number of possibilities. See Appendix C for the full proof. \square

Combination of Theorems 4 and 5 immediately yields the following.

Corollary 1 $[\star]^4$ *Fix integers n, r, k such that $1 < r \leq k$. Then an (r, k) -separator \mathcal{F} of size $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$ can be constructed in time $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot n \cdot \log n)$*

Multisets over a weighted universe. Before proceeding, we discuss the issue of how the considered multisets will be equipped with *weights*. For simplicity, we assume that the universe $[n]$ is weighted, i.e., each element $i \in [n]$ is assigned an integer weight $\mathbf{w}(i)$. We define the weight of a multiset as the sum of the weights of its elements counting repetitions. Formally, for $A \in [r]_0^n$ we have

$$\mathbf{w}(A) = \sum_{i=1}^n A_i \cdot \mathbf{w}(i).$$

Whenever we talk about a *weighted family* of multisets, we mean that the universe $[n]$ is equipped with a weight function and the weights of the multisets are defined as in the formula above.

Representative sets for multisets. We are ready to define the notion of a representative set for a family of multisets.

Definition 4 (Representative sets for multisets) *Let \mathcal{P} be a weighted family of (r, k) -sets. We say that a subfamily $\hat{\mathcal{P}} \subseteq \mathcal{P}$ represents \mathcal{P} if for every (r, k) -set Q the following holds. If there exists some $P \in \mathcal{P}$ of weight w that is (r, k) -compatible with Q , then there also exists some $P' \in \hat{\mathcal{P}}$ of weight $w' \leq w$ that is (r, k) -compatible with Q .*

The following definition and lemma show that having an (r, k) -separator is sufficient for constructing representative sets.

Definition 5 *Let \mathcal{P} be a weighted family of r -sets and let \mathcal{F} be a family of (r, k) -sets. The weighted family $\text{Trim}_{\mathcal{F}}(\mathcal{P}) \subseteq \mathcal{P}$ is defined as follows: For each $F \in \mathcal{F}$, and for each $1 \leq i \leq k$, check if there exists some $P \in \mathcal{P}$ with $|P| = i$ and $P \leq F$. If so, insert into $\text{Trim}_{\mathcal{F}}(\mathcal{P})$ some $P \in \mathcal{P}$ that is of minimal weight among those with $|P| = i$ and $P \leq F$.*

Lemma 1. $[\star]$ *Let \mathcal{F} be an (r, k) -separator and let \mathcal{P} be a weighted family of (r, k) -sets. Then $\text{Trim}_{\mathcal{F}}(\mathcal{P})$ represents \mathcal{P} .*

⁴ Proof of statements marked with \star are omitted due to space constraints and may be found in the appendix.

We can now combine Lemma 1 with the construction of an (r, k) -separator from Corollary 1, and thus obtain a construction of a small representative family for a weighted family of multisets.

Corollary 2 [\star] *There exists a deterministic algorithm that, given a weighted family \mathcal{P} or (r, k) -sets, runs in time $O_k(|\mathcal{P}| \cdot r^{6k/r} \cdot 2^{O(k/r)} \cdot n \log n)$ and returns a family $\hat{\mathcal{P}} \subseteq \mathcal{P}$ that represents \mathcal{P} and has size $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$.*

4 Algorithmic applications

For lack of space, we defer all details to Appendix B, and now only sketch how representative sets for multisets can be applied to r -SIMPLE k -PATH, (r, p, q) -PACKING, and (r, k) -MONOMIAL DETECTION. In all cases, the algorithm applies the expand-and-shrink strategy that was used by Fomin et al. [12] for k -PATH and by Zehavi [27] for SET PACKING (see also the appropriate chapter of [6]). Basically, we iteratively expand the so far obtained family of partial solutions, e.g. by prolonging all the constructed paths by one vertex, and then trim the obtained expanded family by computing its representative subfamily. For (r, k) -MONOMIAL DETECTION, we compute for each gate a representative set of monomials that appear in the polynomial computed at this gate. This is done by “merging” the representative sets of monomials for the input gates, and computing the representative set of this merge. The assumption that the circuit is non-canceling is necessary to argue that the set of monomials appearing at each gate is indeed a proper merge of the sets appearing on its input gates.

5 Low degree monomials and low degree spanning trees

Let G be a simple, undirected graph with n vertices. Let \mathcal{T} be the family of spanning trees of G ; in particular, if G is not connected then $\mathcal{T} = \emptyset$. With every edge $e \in E(G)$ we associate a variable y_e . The *Kirchhoff's polynomial* of G is defined as:

$$K_G((y_e)_{e \in E(G)}) = \sum_{T \in \mathcal{T}} \prod_{e \in E(T)} y_e.$$

Thus, K_G is a polynomial in $\mathbb{Z}[(y_e)_{e \in E(G)}]$. Let v_1, v_2, \dots, v_n be an arbitrary ordering of $V(G)$. The *Laplacian* of G is an $n \times n$ matrix $L_G = [a_{ij}]$, where

$$a_{ij} = \begin{cases} \sum_{e \text{ incident to } v_i} y_e & \text{if } i = j, \\ -y_{v_i v_j} & \text{if } i \neq j \text{ and } v_i v_j \in E(G), \\ 0 & \text{if } i \neq j \text{ and } v_i v_j \notin E(G). \end{cases}$$

Observe that L_G is symmetric and the entries in every column and in every row of L_G sum up to zero. Then it can be shown that all the *first cofactors* of L_G , i.e., the determinants of matrix L_G after removing a row and a column with the same indices, are equal. Let N_G be this common value; then N_G is again a polynomial over variables $(y_e)_{e \in E(G)}$. The Kirchhoff's Matrix Tree Theorem, in its general form, states that these two polynomials coincide.

Theorem 6 (Matrix Tree Theorem) $K_G = N_G$.

We remark that the Matrix Tree Theorem is usually given in the more specific variant, where all variables y_e are replaced with 1; then the theorem expresses the number of spanning trees of G in terms of the first cofactors of L_G . However, the proof can be easily extended to the above, more general form; cf. [25].

Observe that Theorem 6 provides a polynomial-time algorithm for evaluating K_G over a vector of values of variables $(y_e)_{e \in E(G)}$. Indeed, we just need to construct matrix L_G , remove, say, the first row and the first column, and compute the determinant. We now present how this observation can be used to design a fast exact algorithm for the DEGREE-BOUNDED SPANNING TREE problem.

Theorem 7 *The DEGREE-BOUNDED SPANNING TREE problem can be solved in randomized time $O^*(d^{O(n/d)})$ with false negatives.*

Proof. Associate every vertex v of the given graph G with a distinct variable x_v . Let $\bar{K}_G \in \mathbb{Z}[(x_v)_{v \in V(G)}]$ be a polynomial defined as K_G with every variable y_{uv} , for $uv \in E(G)$, evaluated to $x_u x_v$. Then it follows that

$$\bar{K}_G((x_v)_{v \in V(G)}) = \sum_{T \in \mathcal{T}} \prod_{v \in V(G)} x_v^{\deg_T(v)}.$$

Observe also that \bar{K}_G is $2(n-1)$ -homogeneous, that is, all the monomials of \bar{K}_G have their total degrees equal to $2(n-1)$. Thus, graph G admits a spanning tree with maximum degree at most d if and only if polynomial \bar{K}_G contains a $(d, 2(n-1))$ -monomial. Using Theorem 6 we can construct a $n^{O(1)}$ -sized circuit evaluating \bar{K}_G . Hence, verifying whether \bar{K}_G contains a $(d, 2(n-1))$ -monomial boils down to applying the algorithm of Abasi et al. [1] for (r, k) -MONOMIAL DETECTION with $r = d$ and $k = 2(n-1)$. This algorithm runs in randomized time $O^*(d^{O(n/d)})$ and can only produce false negatives. \square

Let us repeat that in the proof of Theorem 7 we could not have used Theorem 16 instead of the result of Abasi et al. [1], because the constructed circuit is not non-canceling. Derandomizing the algorithm and extending it to the weighted setting remains hence open.

Interestingly, the running time of the algorithm of Theorem 7 is essentially optimal, up to the $\log d$ factor in the exponent. A similar lower bound for r -SIMPLE k -PATH was given by Abasi et al. [1].

Theorem 8 $[\star]$ *Unless ETH fails, there exists a constant $s > 0$ such that for no fixed integer $d \geq 2$ the DEGREE-BOUNDED SPANNING TREE problem with the degree bound d can be solved in time $O^*(2^{sn/d})$.*

6 Conclusions

In this paper we considered relaxation parameter variants of several well studied problems in parameterized complexity and exact algorithms. We proved, somewhat surprisingly, that instances with moderate values of the relaxation parameter are significantly easier than instances of the original problems. We hope that

our work, together with the result of Abasi et al. [1] breaks the ground for a systematic investigation of relaxation parameters in parameterized complexity and exact algorithms. We conclude with mentioning some of the most natural concrete follow up questions to our work.

- We gave a deterministic algorithm for *non-canceling* (r, k) -MONOMIAL DETECTION with running time $2^{O(k \frac{\log r}{r})} |C| n^{O(1)}$, while Abasi et al. [1] gave a randomized algorithm with such a running time for (r, k) -MONOMIAL DETECTION without the non-cancellation restriction. Is there a deterministic algorithm for (r, k) -MONOMIAL DETECTION with running time $2^{O(k \frac{\log r}{r})} |C| n^{O(1)}$?
- Does there exist a deterministic algorithm with running time $2^{O(n \frac{\log r}{r})}$ for DEGREE-BOUNDED SPANNING TREE? Note that a deterministic algorithm with running time $2^{O(k \frac{\log r}{r})} |C| n^{O(1)}$ for (r, k) -MONOMIAL DETECTION would immediately imply such an algorithm for DEGREE-BOUNDED SPANNING TREE.
- Is there a $2^{O(k \frac{\log r}{r})} n^{O(1)}$ time algorithm for the problem where we are given as input a graph G , integers k and d , and asked whether G contains a subtree T on at least k vertices, such that the maximum degree of T is at most d ? Observe that for $k = n$ this is exactly the DEGREE-BOUNDED SPANNING TREE problem.
- Is it possible to obtain kernels for problems with relaxation parameters with smaller and smaller size bounds as the relaxation parameter increases?
- Is there an algorithm for r -SIMPLE k -PATH with running time $2^{O(k/r)} n^{O(1)}$? Or a $2^{O(n/d)}$ time algorithm for DEGREE BOUNDED SPANNING TREE?

Acknowledgements

The first author thanks Fedor V. Fomin for hosting him at a visit in the University of Bergen where this research was initiated. We thank the anonymous reviewers for helpful corrections.

References

1. H. Abasi, N. H. Bshouty, A. Gabizon, and E. Haramaty. On r -Simple k -Path. In *MFCS 2014 (Part II)*, pages 1–12, 2014.
2. N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple construction of almost k -wise independent random variables. *Random Struct. Algorithms*, 3(3):289–304, 1992.
3. N. Alon, R. Yuster, and U. Zwick. Color coding. In *Encyclopedia of Algorithms*. 2008.
4. A. Björklund, T. Husfeldt, P. Kaski, and M. Koivisto. Narrow sieves for parameterized paths and packings. *CoRR*, abs/1007.1161, 2010.
5. N. H. Bshouty. Testers and their applications. In *ITCS 2014*, pages 327–352, 2014.
6. M. Cygan, F. V. Fomin, D. Lokshtanov, L. Kowalik, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015. In press.

7. R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
8. H. Fernau, A. López-Ortiz, and J. Romero. Kernelization algorithms for packing problems allowing overlaps. In *TAMC 2015*, pages 415–427, 2015.
9. J. Flum and M. Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin, 2006.
10. F. V. Fomin, F. Grandoni, D. Lokshtanov, and S. Saurabh. Sharp separation and applications to exact and parameterized algorithms. *Algorithmica*, 63(3):692–706, 2012.
11. F. V. Fomin, D. Lokshtanov, F. Panolan, and S. Saurabh. Representative sets of product families. In *ESA 2014*, pages 443–454, 2014.
12. F. V. Fomin, D. Lokshtanov, and S. Saurabh. Efficient computation of representative sets with applications in parameterized and exact algorithms. In *SODA 2014*, pages 142–151, 2014.
13. M. Fürer and B. Raghavachari. Approximating the minimum-degree Steiner Tree to within one of optimal. *J. Algorithms*, 17(3):409–423, 1994.
14. A. Gabizon, D. Lokshtanov, and M. Pilipczuk. Representative sets for multisets. *CoRR*, abs/1411.6756, 2014.
15. M. X. Goemans. Minimum bounded degree spanning trees. In *FOCS 2006*, pages 273–282, 2006.
16. I. Koutis. Faster algebraic algorithms for path and packing problems. In *ICALP 2008*, pages 575–586, 2008.
17. N. Linial, M. Luby, M. E. Saks, and D. Zuckerman. Efficient construction of a small hitting set for combinatorial rectangles in high dimension. *Combinatorica*, 17(2):215–234, 1997.
18. B. Monien. How to find long paths efficiently. In *Analysis and design of algorithms for combinatorial problems (Udine, 1982)*, volume 109 of *North-Holland Math. Stud.*, pages 239–254. North-Holland, Amsterdam.
19. J. Naor and M. Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM J. Comput.*, 22(4):838–856, 1993.
20. M. Naor, L. J. Schulman, and A. Srinivasan. Splitters and near-optimal derandomization. In *FOCS 1995*, pages 182–191, 1995.
21. R. Niedermeier. *Invitation to fixed-parameter algorithms*, volume 31 of *Oxford Lecture Series in Mathematics and its Applications*. Oxford University Press, Oxford, 2006.
22. R. Y. Pinter, H. Shachnai, and M. Zehavi. Deterministic parameterized algorithms for the graph motif problem. In *MFCS 2014 (II)*, pages 589–600, 2014.
23. H. Shachnai and M. Zehavi. Representative families: A unified tradeoff-based approach. In *ESA 2014*, pages 786–797, 2014.
24. M. Singh and L. C. Lau. Approximating minimum bounded degree spanning trees to within one of optimal. *J. ACM*, 62(1):1:1–1:19, 2015.
25. W. T. Tutte. *Graph Theory*. Cambridge University Press, 2001.
26. R. Williams. Finding paths of length k in $O^*(2^k)$ time. *Inf. Process. Lett.*, 109(6):315–318, 2009.
27. M. Zehavi. Deterministic parameterized algorithms for matching and packing problems. *CoRR*, abs/1311.0484, 2013.
28. M. Zehavi. Solving parameterized problems by mixing color coding-related techniques. *CoRR*, abs/1410.5062, 2014.

A Construction of a separating family

The purpose of this section is to prove Theorem 4, that is, to construct a small separating family of functions. First, we need to introduce some auxiliary results.

Hashing families. Recall that, for an integer $t \geq 1$, we say that a family of functions $\mathcal{H} \subseteq \{[n] \rightarrow [m]\}$ is a t -perfect hash family, if for every $C \subseteq [n]$ of size $|C| = t$ there is $f \in \mathcal{H}$ that is injective on T . Alon, Yuster and Zwick [3] used a construction of Moni Naor (based on ideas from Naor et al. [20]) to hash a subset of size t into a world of size t^2 using a very small set of functions:

Theorem 9 ([3] based on Naor). *For integers $1 \leq t \leq n$, a t -perfect hash family $\mathcal{H} \subseteq \{[n] \rightarrow [t^2]\}$ of size $t^{O(1)} \cdot \log n$ can be constructed in time $O(t^{O(1)} \cdot n \cdot \log n)$.*

We will also use the following perfect hash family given by Naor, Schulman and Srinivasan [20].

Theorem 10 ([20]). *For integers $1 \leq t \leq n$, a t -perfect hash family $\mathcal{H} \subseteq \{[k^2] \rightarrow [t]\}$ of size $e^{t+O(\log^2 t)} \cdot \log k$ can be constructed in time $O(e^{t+O(\log^2 t)} \cdot k \cdot \log k)$.*

Hitting combinatorial rectangles. We first recall the notion of a hitting set for combinatorial rectangles. For a sequence of integers (m_1, m_2, \dots, m_t) , by $\prod_{i \in [t]} [m_i]$ we denote $[m_1] \times [m_2] \times \dots \times [m_t]$.

Definition 6 *Let $R \subseteq \prod_{i \in [t]} [m_i]$ be a set of the form $R_1 \times \dots \times R_t$, where $R_i \subseteq [m_i]$ for all $i \in [t]$. We say that R is a combinatorial rectangle with sidewise density γ , if for every $i \in [t]$ we have that $|R_i| \geq \gamma \cdot m_i$. A set $H \subseteq \prod_{i \in [t]} [m_i]$ is a hitting set for rectangles with sidewise density γ if for every set $R \subseteq \prod_{i \in [t]} [m_i]$ that is a combinatorial rectangle of sidewise density γ , it holds that $R \cap H \neq \emptyset$.*

Linial et al. [17] gave the following construction of a hitting set for combinatorial rectangles.

Theorem 11 ([17]). *A hitting set $H \subseteq [m]^t$ for rectangles with sidewise density $1/3$ of size $|H| = 2^{O(t)} \cdot m^{O(1)}$ can be constructed in time $2^{O(t)} \cdot m^{O(1)}$.*

We need a hitting set for combinatorial rectangles in a universe where the coordinates are from domains of different sizes. We show that Theorem 11 can be adapted to this setting.

Corollary 3 *Suppose $m_1, \dots, m_t \leq m$. Then a hitting set $H \subseteq \prod_{i \in [t]} [m_i]$ for rectangles with sidewise density $1/2$ of size $|H| = 2^{O(t)} \cdot m^{O(1)}$ can be constructed in time $2^{O(t)} \cdot m^{O(1)}$.*

Proof. For the purpose of the proof it will be convenient to redefine $[a]$ as $\{0, \dots, a-1\}$. Let $m' = 3m$. Define mapping $\pi: [m']^t \rightarrow \prod_{i \in [t]} [m_i]$ as follows:

$$\pi(a_1, a_2, \dots, a_t) = (a_1 \bmod m_1, a_2 \bmod m_2, \dots, a_t \bmod m_t).$$

Observe that if $R_i \subseteq [m_i]$ is such that $|R_i| \geq m_i/2$, and $R'_i \subseteq [m']$ is the set of all elements of $[m']$ whose remainders modulo m_i belong to R_i , then $|R'_i| \geq m'/3$. This follows from the fact that $m' = 3m \geq 3m_i$. Therefore, if $R \subseteq \prod_{i \in [t]} [m_i]$ is a combinatorial rectangle with sidewise density $1/2$, then $R' \triangleq \pi^{-1}(R) \subseteq [m']^t$ is a combinatorial rectangle with sidewise density $1/3$. Let $H' \subseteq [m']^t$ be the hitting set for combinatorial rectangles with sidewise density $1/3$ given by Theorem 11. Moreover, let $H \triangleq \pi(H')$; hence we have that $|H| \leq |H'| = 2^{O(t)} \cdot m^{O(1)}$. We have that $H' \cap R' \neq \emptyset$, so also $H \cap R = \pi(H') \cap \pi(R') \supseteq \pi(H' \cap R') \neq \emptyset$. Since R was chosen arbitrarily, we infer that H is a hitting set for combinatorial rectangles with sidewise density $1/2$. \square

Pairwise independent families. Another component in our construction is a family of ϵ -pairwise independent functions.

Definition 7 *A family of functions $\mathcal{H} \subseteq \{[n] \rightarrow [m]\}$ is ϵ -pairwise independent if for all $x, y \in [n]$ and $a, b \in [m]$, it holds that*

$$|\mathbb{P}_{f \leftarrow \mathcal{H}}(f(x) = a \wedge f(y) = b) - 1/m^2| \leq \epsilon.$$

The constructions of [19] and [2] imply the following construction of an ϵ -pairwise independent family of functions (cf. [3, Section 4]).

Theorem 12 ([2, 19]) *Fix any $m \leq n$. Then a $1/m^2$ -pairwise independent family $\mathcal{H} \subseteq \{[n] \rightarrow [m]\}$ of size $m^{O(1)} \cdot \log n$ can be constructed in time $O(m^{O(1)} \cdot n \cdot \log n)$.*

We now use Theorem 12 to construct a small family that separates one element from k other elements with non-negligible probability.

Lemma 2. *Fix integers k, n with $1 \leq k \leq n$. Then a family $\mathcal{H}_{k,4k} \subseteq \{[n] \rightarrow [4 \cdot k]\}$ that is $(1, k, 4 \cdot k)$ -separating with probability $1/2$ and has size $k^{O(1)} \cdot \log n$ can be constructed in time $k^{O(1)} \cdot n \cdot \log n$.*

Proof. Let $\mathcal{H} \subseteq \{[n] \rightarrow [4k]\}$ be the family of $1/(4k)^2$ -pairwise independent functions, given by Theorem 12. Fix sets $C = \{a\}$ and $D = \{b_1, \dots, b_{k-1}\}$ we want to separate. For $j \in [k-1]$, let X_j be a random variable that is equal to one if $f(a) = f(b_j)$ and to zero otherwise, where f is chosen uniformly at random from \mathcal{H} . From the $1/4k$ -pairwise independence of \mathcal{H} , we have that $\mathbb{E}(X_j) = \mathbb{P}(X_j = 1) \leq 4k \cdot (1/(4k)^2 + 1/(4k)^2) \leq 1/(2k)$. Let us define $X = \sum_{j=1}^{k-1} X_j$, so that we have $\mathbb{E}(X) \leq (k-1)/(2k) < 1/2$. Note that X is precisely the number of collisions between C and D . From Markov's inequality, we have that $\mathbb{P}(X \geq 1) \leq 1/2$. So with probability at least $1/2$ over the choice of $f \in \mathcal{H}$, C and D are separated by f . \square

The next claim follows from the well-known theorem that the geometric mean of non-negative numbers is never larger than their arithmetic mean.

Proposition 1. *Let k_1, \dots, k_t be non-negative real numbers such that $k_1 + \dots + k_t \leq k$. Then $\prod_{i=1}^t k_i \leq (k/t)^t$.*

The main construction. We are ready to proceed to the main construction.

Proof (Proof of Theorem 4). Fix disjoint subsets $C, D \subseteq [n]$ with $|C| = t$ and $|D| \leq k - t$. Recall that we want to construct a family of functions $\mathcal{H} \subseteq \{[n] \rightarrow [t + 1]\}$, such that for any C and D as above, there exists $h \in \mathcal{H}$ that separates C from D . It will be convenient to present the family by constructing h adaptively given C and D . That is, for arbitrarily chosen C and D , we will adaptively construct a function h that separates C from D . Function h will be constructed by taking a number of *choices*, where each choice is taken among a number of possibilities. The final family \mathcal{H} will comprise all h that can be obtained using any such sequence of choices; thus, the product of the numbers of possibilities will limit the size of \mathcal{H} . As C and D are taken arbitrarily, it immediately follows that such \mathcal{H} separates every pair (C, D) .

1. Let $\mathcal{H}_0 \subseteq \{[n] \rightarrow [k^2]\}$ be the k -perfect hash family given by Theorem 9. Choose $f_0 \in \mathcal{H}_0$ that is injective on $C \cup D$ — there are $k^{O(1)} \cdot \log n$ choices for this stage.

From now on, we identify C and D with their images in $[k^2]$ under f_0 .

2. Let $\mathcal{H}_1 \subseteq \{[k^2] \rightarrow [t]\}$ be the t -perfect hash family given by Theorem 10. Choose a function $f_1 \in \mathcal{H}_1$ that is injective on C — there are $e^{t+O(\log^2 t)} \cdot \log k$ choices for this stage.

For $i \in [t]$, we denote from now on by c_i the (unique) element of C with $f_1(c_i) = i$. Moreover, elements mapped to i under f_1 will be denoted by \mathcal{B}_i , and we will think of them as the i -th bucket.

3. Choose non-negative integers k_1, \dots, k_t such that k_i is the number of elements $a \in D$ with $f_1(a) = i$. Note that $k_1 + \dots + k_t \leq k - t$, so the number of choices for this stage is at most $\binom{k}{t} \leq (ek/t)^t$.
4. For $i \in [t]$, let $\mathcal{H}_{k_i, 4k_i} \subseteq \{[k^2] \rightarrow [4 \cdot k_i]\}$ be family given by Lemma 2. That is, $\mathcal{H}_{k_i, 4k_i}$ is $(1, k_i, 4 \cdot k_i)$ -separating with probability $1/2$ and has size $m_i = k_i^{O(1)} \cdot \log k$.

Let R be the set of all vectors (h_1, \dots, h_t) such that for all $i \in [t]$, h_i is an element of $\mathcal{H}_{k_i, 4k_i}$ that separates c_i from D_i . Identify $[m_i]$ with $\mathcal{H}_{k_i, 4k_i}$ by ordering the functions in $\mathcal{H}_{k_i, 4k_i}$ arbitrarily. Observe that R is a combinatorial rectangle of sidewise density $1/2$ in $\prod_{i \in [t]} [m_i]$. Note that for all $i \in [t]$, $m_i \leq m$ for some $m = k^{O(1)}$. Let H be the hitting set for combinatorial rectangles with sidewise density $1/2$ given by Corollary 3. Choose an element $(h_1, \dots, h_t) \in H \cap R$. As $|H| = 2^{O(t)} \cdot k^{O(1)}$, there are $2^{O(t)} \cdot k^{O(1)}$ choices for this stage.

For $i \in [t]$, apply h_i to partition bucket \mathcal{B}_i into buckets $\mathcal{B}_{i,1}, \dots, \mathcal{B}_{i,4k_i}$; that is, element $a \in [n]$ is put into bucket $\mathcal{B}_{f_1(a), h_{f_1(a)}(a)}$. Since $(h_1, \dots, h_t) \in R$, we thus have that each c_i is mapped to a separate bucket from all elements of D_i .

5. For each $i \in [t]$, choose the unique index $j_i \in [4k_i]$ such that c_i was mapped to bucket \mathcal{B}_{i,j_i} . Construct the final function h by mapping all the elements of \mathcal{B}_{i,j_i} to i , and mapping all the elements of $\mathcal{B}_{i,j'}$ for $j' \neq j_i$ to $t + 1$, for all $i \in [t]$. The number of choices for this stage is $4^t \cdot \prod_{i=1}^t k_i \leq (4k/t)^t$, where the inequality follows from Proposition 1.

To sum up, the number of different functions enumerated in all the above stages is at most

$$\begin{aligned} k^{O(1)} \cdot \log n \cdot e^{t+O(\log^2 t)} \cdot \log k \cdot (ek/t)^t \cdot 2^{O(t)} \cdot k^{O(1)} \cdot (4k/t)^t &= (k/t)^{2t} \cdot 2^{O(t)} \cdot k^{O(1)} \cdot \log n \\ &= O_k((k/t)^{2t} \cdot 2^{O(t)} \cdot \log n). \end{aligned}$$

Similarly, going over the construction times of the different components used in the above stages, we get that the family can be constructed in time

$$\begin{aligned} O(k^{O(1)} \cdot n \cdot \log n \cdot e^{t+O(\log^2 t)} \cdot k \cdot \log k \cdot (ek/t)^t \cdot 2^{O(t)} \cdot k^{O(1)} \cdot (4k/t)^t) \\ = O((k/t)^{2t} \cdot 2^{O(t)} \cdot k^{O(1)} \cdot n \cdot \log n) = O_k((k/t)^{2t} \cdot 2^{O(t)} \cdot n \cdot \log n). \end{aligned}$$

□

Lopsided universal sets. Informally speaking, an (n, p, q) -lopsided universal set is a set of strings such that, when focusing on any $k \triangleq p + q$ locations, we see all patterns of hamming weight p . Universal set families, and in particular the lopsided ones, have important applications in the determinization of parameterized algorithms that use the technique of color coding; cf. [12].

Definition 8 *An (n, p, q) -lopsided universal set is a family of subsets $\mathcal{F} \subseteq [n]$ such that for every disjoint subsets $A, B \subseteq [n]$ with $|A| = p$ and $|B| = q$ there exists $F \in \mathcal{F}$ with $A \subseteq F$ and $F \cap B = \emptyset$.*

A probabilistic argument shows the existence of (n, p, q) -lopsided universal sets of size $\binom{k}{p} \cdot \sqrt{p \cdot q \cdot k} \cdot \log n$ (cf. Lemma 52 in [5]). Fomin, Lokshtanov and Saurabh [12] used the technique of [20] to construct an (n, p, q) -lopsided universal set of size $\binom{k}{p} \cdot 2^{O(\frac{k}{\log \log(k)})} \cdot \log n$. Consider the case $p = o(q)$: For simplicity of discussion, let us focus on the $\binom{k}{p} \leq (e \cdot k/p)^p$ term in the probabilistic construction and omit $\log n$ terms. In this case the $2^{O(\frac{k}{\log \log(k)})}$ term of [12] is much larger than $\binom{k}{p}$. Bshouty [5] gives a better construction for such p achieving size roughly k^{p+2} (see Lemma 53 in [5]). We remark that Bshouty [5] referred to this object as a cover free family. As a corollary of the last section we obtain an explicit construction of size $(k/p)^{O(p)}$.

Theorem 13 *Fix integers p, q and n such that $k = p + q \leq n$. Then an (n, p, q) -lopsided universal set of size $O_k((k/p)^{2 \cdot p} \cdot 2^{O(p)} \cdot \log n)$ can be constructed in time $O_k((k/p)^{2 \cdot p} \cdot 2^{O(p)} \cdot n \cdot \log n)$.*

Proof. Let $\mathcal{H} \subseteq \{[n] \rightarrow [p+1]\}$ be the (p, k) -minimal separating family of size $O_k((k/p)^{2p} \cdot 2^{O(p)} \cdot \log n)$ given by Theorem 4. For each $h \in \mathcal{H}$, construct the set $F_h = h^{-1}([p])$, and let $\mathcal{F} \triangleq \{F_h \mid h \in \mathcal{H}\}$. The definition of a minimal separating family immediately implies that \mathcal{F} is an (n, p, q) -lopsided universal set, and the time needed for the construction follows from Theorem 4. □

B Algorithmic applications

In this section we present applications of our construction of representative sets for multisets to the design of deterministic algorithms for problems with relaxed disjointness constraints. We first give some useful auxiliary tools, and then present algorithms for three problems: r -SIMPLE k -PATH, (r, p, q) -PACKING, (r, k) -MONOMIAL DETECTION. Let us remark that the algorithms for the first two problems follow from the algorithm for the (more general) last problem. Nevertheless, we find it more didactic to present the applications in increasing order of difficulty. Throughout this section, r always denotes an integer in the range $1 < r < k$.

B.1 Algorithmic preliminaries

The following simple observation will be used in all of our algorithmic applications.

Lemma 3. *Let \mathcal{P} be a family of (r, k) -sets that contains a multiset P of size k and weight w . Then any $\hat{\mathcal{P}} \subseteq \mathcal{P}$ that represents \mathcal{P} contains a multiset P' of size k and weight $w' \leq w$.*

Proof. Take $B = \emptyset$. Then P and B are (r, k) -compatible. So there must exist $P' \in \hat{\mathcal{P}}$ that is (r, k) -compatible with B and has weight at most w . As $|P' + B| = k$, we have $|P'| = k$. \square

The following simple fact is immediate and will be used implicitly.

Proposition 2. *If \mathcal{A}'' represents \mathcal{A}' and \mathcal{A}' represents \mathcal{A} , then \mathcal{A}'' represents \mathcal{A} .*

We now show that representative sets behave robustly under union of families.

Lemma 4. *Suppose \mathcal{A} and \mathcal{B} are two weighted families of (r, k) -sets, and suppose further that $\hat{\mathcal{A}}$ represents \mathcal{A} and $\hat{\mathcal{B}}$ represents \mathcal{B} . Then $\hat{\mathcal{A}} \cup \hat{\mathcal{B}}$ represents $\mathcal{A} \cup \mathcal{B}$.*

Proof. Take any (r, k) -set D such that there exists $C \in \mathcal{A} \cup \mathcal{B}$ that is (r, k) -compatible with D . If $C \in \mathcal{A}$, then there exists $C' \in \hat{\mathcal{A}}$ with $\mathbf{w}(C') \leq \mathbf{w}(C)$ such that C' is (r, k) -compatible with D . Then also $C' \in \hat{\mathcal{A}} \cup \hat{\mathcal{B}}$ and we are done. The reasoning for the case when $C \in \mathcal{B}$ is symmetric. \square

Finally, we introduce operation \bullet that extends the considered weighted family of (r, k) -sets. This is the crucial ingredient in the classic technique of iteratively extending the current family by a new object, and then shrinking the size of the family by taking a representative subfamily.

Definition 9 *We say that (r, k) -sets A and B are (r, k) -consistent, if $|A+B| \leq k$ and $A_i + B_i \leq r$ for all $i \in [n]$. (Note that this is a weaker definition than being (r, k) -compatible, where it is required that $|A + B| = k$.)*

Let \mathcal{A} and \mathcal{B} be weighted families of (r, k) -sets. We define the weighted family $\mathcal{A} \bullet \mathcal{B}$ as follows:

$$\mathcal{A} \bullet \mathcal{B} \triangleq \{A + B \mid A \in \mathcal{A}, B \in \mathcal{B}, A \text{ and } B \text{ are } (r, k)\text{-consistent}\}.$$

In particular, for an element $i \in [n]$, we denote

$$\mathcal{A} \bullet i \triangleq \{A + \{i\} \mid A \in \mathcal{A}, |A| < k, A_i < r\}.$$

Lemma 5. *Suppose \mathcal{A} and \mathcal{B} are two weighted families of (r, k) -sets, and suppose further that $\hat{\mathcal{A}}$ represents \mathcal{A} and $\hat{\mathcal{B}}$ represents \mathcal{B} . Then $\hat{\mathcal{A}} \bullet \hat{\mathcal{B}}$ represents $\mathcal{A} \bullet \mathcal{B}$.*

Proof. Fix an (r, k) -set D such that there exists $C \in \mathcal{A} \bullet \mathcal{B}$ that is (r, k) -compatible with D . Hence, we have that $C = A + B$ for some $A \in \mathcal{A}$ and $B \in \mathcal{B}$ that are (r, k) -consistent. In particular, A is (r, k) -compatible with $B + D$. Hence, there exists $A' \in \hat{\mathcal{A}}$ with $\mathbf{w}(A') \leq \mathbf{w}(A)$ that is (r, k) -compatible with $B + D$. In other words, $A' + B + D$ is an (r, k) -set of size $|A' + B + D| = k$. It follows that B is (r, k) -compatible with $A' + D$. Hence, there exists $B' \in \hat{\mathcal{B}}$ with $\mathbf{w}(B') \leq \mathbf{w}(B)$ that is (r, k) -compatible with $A + D$. Again, for this B' we have that $A' + B' + D$ is an (r, k) -set of size k . Therefore, $A' + B'$ is an (r, k) -set of size at most k and weight not larger than the weight of $A + B$. In other words, A' and B' are (r, k) -consistent, and therefore $A' + B' \in \mathcal{A}' \bullet \mathcal{B}'$. As $A' + B'$ is (r, k) -compatible with D and $\mathbf{w}(A' + B') \leq \mathbf{w}(A + B)$, we are done. \square

B.2 r -Simple k -Path

We first introduce some notation for defining the r -SIMPLE k -PATH problem. Let G be a directed graph on n vertices. A k -path in G is a *simple* walk $P = v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_k$ in G , that is, all vertices traversed by the walk are pairwise different. An r -simple k -path is a walk $P = v_1 \rightarrow v_2 \rightarrow \dots \rightarrow v_k$ in G such that no vertex is repeated more than r times. In case the vertex set of G is equipped with a weight function, the weight of an r -simple k -path P is defined as the weight of the multiset of vertices traversed by P , where each vertex is taken the number of times equal to the number of its visits on P .

r -SIMPLE k -PATH

Input: Directed graph $G = (V, E)$, weight function $\mathbf{w}: V \rightarrow \mathbb{R}$, integers r and k .

Parameters: r, k

Question: Determine whether there exists an r -simple k -path in G , and if so then return one of minimal possible weight.

Theorem 14 r -SIMPLE k -PATH can be solved in deterministic time $O_k(r^{12k/r} \cdot 2^{O(k/r)} \cdot n^3 \cdot \log n)$.

Proof. Identify the vertex set V with $[n]$. For $u \in [n]$ and $i \in [k]$, let $\mathcal{P}_{i,u}$ denote the set of r -simple paths in G that have length i and end in u . Note that an element P of such a set $\mathcal{P}_{i,u}$ can be viewed as an (r, k) -set when we ignore the order of vertices in the path. We can thus view the sets $\mathcal{P}_{i,u}$ as families of (r, k) -sets. We will compute, using a dynamic programming algorithm, for each $u \in [n]$ and $i \in [k]$, a set $\hat{\mathcal{P}}_{i,u} \subseteq \mathcal{P}_{i,u}$ that represents $\mathcal{P}_{i,u}$ and has size at most $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$. By Lemma 3, a minimal-weight r -simple k -path in G can be recovered by taking a minimum-weight multiset among families $\hat{\mathcal{P}}_{k,u}$, for $u \in [n]$. In particular, if all these sets are empty, then no r -simple k -path exists in G . Thus, such an algorithm can be used to solve r -SIMPLE k -PATH.

We proceed with the algorithm description. For $i = 1$ and $u \in [n]$, family $\hat{\mathcal{P}}_{1,u}$ simply contains only the length-one path u . Assume that we have computed families $\hat{\mathcal{P}}_{i,u}$ for every $u \in [n]$. We describe the computation of $\hat{\mathcal{P}}_{i+1,v}$ for a fixed $v \in [n]$, together with its running time.

1. We compute $\mathcal{P}' \triangleq \bigcup_{(u,v) \in E} \hat{\mathcal{P}}_{i,u} \bullet v$. Clearly, we have that $|\mathcal{P}'| \leq n \cdot O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$. We now claim that \mathcal{P}' represents $\mathcal{P}_{i+1,v}$. Note first that $\mathcal{P}_{i+1,v} = \bigcup_{(u,v) \in E} \mathcal{P}_{i,u} \bullet v$. By Lemma 5, for any $u \in [n]$, $\hat{\mathcal{P}}_{i,u} \bullet v$ represents $\mathcal{P}_{i,u} \bullet v$. Therefore, $\mathcal{P}' = \bigcup_{(u,v) \in E} \hat{\mathcal{P}}_{i,u} \bullet v$ represents $\bigcup_{(u,v) \in E} \mathcal{P}_{i,u} \bullet v = \mathcal{P}_{i+1,v}$. Computing \mathcal{P}' directly from the definition requires time $O_k(\sum_{u \in [n]} |\hat{\mathcal{P}}_{i,u}|) = O_k(n \cdot r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$.
2. Now, using the algorithm of Corollary 2, compute a set $\hat{\mathcal{P}}_{i+1,v}$ that represents \mathcal{P}' and has size $|\hat{\mathcal{P}}_{i+1,v}| = O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$. This takes time $O_k(|\mathcal{P}'| \cdot r^{6k/r} \cdot 2^{O(k/r)} \cdot n \log n) = O_k(n^2 \cdot r^{12k/r} \cdot 2^{O(k/r)} \cdot \log n)$.

Since during each of k steps of the algorithm, this procedure is applied to every vertex $v \in [n]$, we conclude that the running time of the algorithm is $O_k(r^{12k/r} \cdot 2^{O(k/r)} \cdot n^3 \cdot \log n)$, as claimed. \square

We have considered the r -SIMPLE k -PATH problem where the weights are on the vertices. One could also consider the *edge weighted* variant where every edge has a weight and the weight of a walk is the sum of the weight of the edges on the walk (counting multiplicities of edges). One can reduce this variant to the vertex weighted variant as follows. Let G be the input graph, we make a new graph G' from G . Here each edge uv of G is replaced by a new vertex x_{uv} with u being the only in-neighbor and v being the only out-neighbor of x_{uv} . The weight of x_{uv} is set to the weight of the edge uv . The weight of the vertices of G' that correspond to the vertices of G is set to 0. An r -simple k -path in G corresponds to an r -simple $(2k+1)$ -path in G' of the same weight. On the other hand, a minimum weight r -simple $(2k+1)$ -path in G' must start and end in vertices corresponding to vertices of G , since these have weight 0. Such r -simple $(2k+1)$ -paths in G' correspond to an r -simple k -path in G of the same weight. Thus we can run the algorithm for Theorem 14 on G' , obtaining an algorithm for the edge weighted variant with running time $r^{O(k/r)} n^{O(1)}$.

B.3 (r, p, q) -Packing

We say that $P \subseteq [n]$ is a q -set if $|P| = q$. We say that a family of q -sets $\mathcal{A} = \{P_1, \dots, P_t\}$ is an r -packing if every element $j \in [n]$ appears in at most r of the q -sets in \mathcal{A} . The (r, p, q) -PACKING problem asks whether a given family of q -sets contains an r -packing. Again, in case universe $[n]$ is equipped with a weight function, we can say that a family of q -sets is *weighted* by setting the weight of a q -set to be the total weight of its elements. The weight of an r -packing is defined as the sum of weights of its sets.

(r, p, q) -PACKING

Input: Weighted family \mathcal{F} of q -sets, integers r, p, q

Parameters: r, p, q

Question: Determine whether there exists a subfamily $\mathcal{A} \subseteq \mathcal{F}$ with $|\mathcal{A}| = p$ that is an r -packing, and if so then return such \mathcal{A} of minimal total weight.

Theorem 15 *Let $k \triangleq p \cdot q$. Then (r, p, q) -PACKING can be solved in deterministic time $O_k(|\mathcal{F}| \cdot r^{12k/r} \cdot 2^{O(k/r)} \cdot n \log^2 n)$.*

Proof. Let \mathcal{F} be the input family of q -sets. For $0 \leq i \leq p$, denote by \mathcal{P}_i the set of r -packings of size i in \mathcal{F} . Note that an element $\mathcal{A} = \{P_1, \dots, P_i\}$ of \mathcal{P}_i can be viewed as an (r, k) -set: We think of \mathcal{A} as the multiset $\mathcal{A} = P_1 + \dots + P_i$. We know that this multiset has size at most $p \cdot q = k$ and every element $j \in [n]$ appears in \mathcal{A} at most r times. We will compute, for each $0 \leq i \leq p$, a subfamily $\hat{\mathcal{P}}_i \subseteq \mathcal{P}_i$ that represents \mathcal{P}_i and has size at most $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$. By Lemma 3, a minimum-weight r -packing in \mathcal{F} can be recovered by taking a minimum-weight element of $\hat{\mathcal{P}}_p$. In particular, if $\hat{\mathcal{P}}_p$ is empty, then no r -packing exists in \mathcal{F} . Thus, such an algorithm can be used to solve (r, p, q) -PACKING.

We proceed with the algorithm description. For $i = 0$, we take $\hat{\mathcal{P}}_0 = \mathcal{P}_0 = \{\emptyset\}$. Assume we have computed $\hat{\mathcal{P}}_i$ for some $0 \leq i < p$. We describe the computation of $\hat{\mathcal{P}}_{i+1}$.

1. We compute $\mathcal{P}' \triangleq \hat{\mathcal{P}}_i \bullet \mathcal{F}$. Clearly, we have that $|\mathcal{P}'| \leq |\mathcal{F}| \cdot |\hat{\mathcal{P}}_i| \leq |\mathcal{F}| \cdot O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$. We claim that \mathcal{P}' represents \mathcal{P}_{i+1} . Note first that $\mathcal{P}_{i+1} = \mathcal{P}_i \bullet \mathcal{F}$. Hence, by Lemma 5, $\mathcal{P}' = \hat{\mathcal{P}}_i \bullet \mathcal{F}$ represents $\mathcal{P}_i \bullet \mathcal{F} = \mathcal{P}_{i+1}$. Computing \mathcal{P}' directly from the definition requires time $O_k(|\mathcal{F}| \cdot |\hat{\mathcal{P}}_i| \cdot n) = O_k(|\mathcal{F}| \cdot r^{6k/r} \cdot 2^{O(k/r)} \cdot n \log n)$.
2. Now, using the algorithm of Corollary 2, compute a set $\hat{\mathcal{P}}_{i+1}$ that represents \mathcal{P}' and has size $|\hat{\mathcal{P}}_{i+1}| = O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$. This takes time $O_k(|\mathcal{P}'| \cdot r^{6k/r} \cdot 2^{O(k/r)} \cdot n \log n) = O_k(|\mathcal{F}| \cdot r^{12k/r} \cdot 2^{O(k/r)} \cdot n \log^2 n)$.

Since we repeat this procedure p times, the claimed running time follows \square

In Theorem 15 we gave an algorithm for (r, p, q) -PACKING where every element has a weight and the weight of a set is equal to the sum of the weights of the elements. A related variant is one where a weight function $w : \mathcal{F} \rightarrow \mathbb{N}$ is given as input. That is, every set $P \in \mathcal{F}$ has its own weight $w(P)$, and the

weight of a subfamily $\mathcal{A} \subseteq \mathcal{F}$ is the sum of the weights of the sets in \mathcal{A} . We will call this the *set weighted* (r, p, q) -PACKING problem. The set weighted problem can be reduced to the original by adding for every set $P \in \mathcal{F}$ a new element e_P of weight $w(P)$, inserting e_P into P and giving e_P weight $w(P)$. Note that P is the only set in the new instance that contains e_P . All elements corresponding to the elements of the instance of set weighted (r, p, q) -PACKING are given weight 0. All sets in the new instance have size $q + 1$, and r -packings in the new instance correspond to r -packings in the original instance with the same weight. Thus we can apply the algorithm of Theorem 15 on the new instance in order to solve the instance of set weighted (r, p, q) -PACKING. This yields an algorithm for set weighted (r, p, q) -PACKING with running time $|\mathcal{F}|r^{O(k/r)}n^{O(1)}$.

B.4 (r, k) -Monomial Detection

In this subsection we consider polynomials from ring $\mathbb{Z}[X_1, \dots, X_n]$. We say a monomial $X_1^{d_1} \cdots X_n^{d_n}$ is an r -*monomial* if for all $i \in [n]$, we have $d_i \leq r$. We say the monomial is an (r, k) -*monomial* if the above holds and the total degree of the monomial is k . Let C be an arithmetic circuit computing a polynomial $f(X_1, \dots, X_n) \in \mathbb{Z}[X_1, \dots, X_n]$. We say C is *non-canceling* if it contains only variables at its leaves (i.e., no constants), and only addition and multiplication gates (i.e., no subtractions). For a non-canceling circuit C , we define $|C|$ to be the number of multiplication and addition gates plus the number of leaves (each containing a variable). We assume the fan-in of a non-canceling circuit is two, i.e., each multiplication and addition gate has at most two wires coming in. (This will simply be convenient for bounding the running time as a function of $|C|$.)

(r, k) -MONOMIAL DETECTION

Input: A non-canceling circuit C computing a polynomial $f(X_1, \dots, X_n) \in \mathbb{Z}[X_1, \dots, X_n]$, integers r, k .

Parameters: r, k

Question: Determine if there exists an (r, k) -monomial in f with non-zero coefficient and if so, then return such a monomial.

Theorem 16 *Given a non-canceling circuit C computing a polynomial $f(X_1, \dots, X_n) \in \mathbb{Z}[X_1, \dots, X_n]$, (r, k) -MONOMIAL DETECTION can be solved in deterministic time $O_k(|C| \cdot r^{18k/r} \cdot 2^{O(k/r)} \cdot n \log^3 n)$.*

Proof. For each gate s of C , let f_s be the polynomial computed at s . Define \mathcal{P}_s to be the set of r -monomials of total degree at most k that appear in f_s with nonzero coefficient. We can view \mathcal{P}_s as a family of (r, k) -sets: An r -monomial $M = X_1^{d_1} \cdots X_n^{d_n}$ of total degree at most k corresponds to the (r, k) -set where each $j \in [n]$ appears d_j times (in this case, we put uniform weights on the elements of universe $[n]$). We present an algorithm that computes, for every gate s of C , a subfamily $\mathcal{P}_s \subseteq \mathcal{P}_s$ of size $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$ that represents \mathcal{P}_s . Let s_{out} be the output gate of C . By Lemma 3, if f contains an r -monomial of

total degree k , then $\hat{\mathcal{P}}_{s_{out}}$ will contain such a monomial. Hence, to solve (r, k) -MONOMIAL DETECTION it suffices to check whether $\hat{\mathcal{P}}_{s_{out}}$ is nonempty.

We compute the sets $\hat{\mathcal{P}}_s$ from bottom to top. For s being an input gate containing variable X_i , we simply put $\hat{\mathcal{P}}_s = \mathcal{P}_s = \{\{i\}\}$. Take then any non-input gate s , and suppose we have computed $\hat{\mathcal{P}}_{s_1}$ and $\hat{\mathcal{P}}_{s_2}$ for the gates s_1 and s_2 having wires into s .

1. If s is an addition gate, then we define $\mathcal{P}' \triangleq \hat{\mathcal{P}}_{s_1} \cup \hat{\mathcal{P}}_{s_2}$. We claim that \mathcal{P}' represents \mathcal{P}_s : Since C is non-canceling, the set of monomials that appear in f_s with a nonzero coefficient is simply the union of the sets of monomials of appearing in f_{s_1} and in f_{s_2} . In particular, $\mathcal{P}_s = \mathcal{P}_{s_1} \cup \mathcal{P}_{s_2}$. Therefore, by Lemma 4 we infer that $\mathcal{P}' = \hat{\mathcal{P}}_{s_1} \cup \hat{\mathcal{P}}_{s_2}$ represents \mathcal{P}_s . Note that $|\mathcal{P}'| \leq |\hat{\mathcal{P}}_{s_1}| + |\hat{\mathcal{P}}_{s_2}| = O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$ and \mathcal{P}' can be computed in time $O_k(|\hat{\mathcal{P}}_{s_1}| \cdot |\hat{\mathcal{P}}_{s_2}| \cdot n) = O_k(r^{12k/r} \cdot 2^{O(k/r)} \cdot n \log^2 n)$.
2. If s is a multiplication gate, then we define $\mathcal{P}' \triangleq \hat{\mathcal{P}}_{s_1} \bullet \hat{\mathcal{P}}_{s_2}$. Since C is non-canceling, the set of monomials appearing in f_s is exactly the set of all products of a monomial appearing in f_{s_1} and a monomial appearing in f_{s_2} . In particular, \mathcal{P}_s is exactly the set of these products that are also (r, k) -monomials. When viewed as a multiset, the product of monomials M_1 and M_2 is the multiset $M_1 + M_2$. Thus, we have that $\mathcal{P}_s = \mathcal{P}_{s_1} \bullet \mathcal{P}_{s_2}$ and therefore, using Lemma 5, we infer that $\mathcal{P}' = \hat{\mathcal{P}}_{s_1} \bullet \hat{\mathcal{P}}_{s_2}$ represents \mathcal{P}_s . Note that $|\mathcal{P}'| \leq |\hat{\mathcal{P}}_{s_1}| \cdot |\hat{\mathcal{P}}_{s_2}| = O_k(r^{12k/r} \cdot 2^{O(k/r)} \cdot \log^2 n)$ and \mathcal{P}' can be computed in time $O_k(|\hat{\mathcal{P}}_{s_1}| \cdot |\hat{\mathcal{P}}_{s_2}| \cdot n) = O_k(r^{12k/r} \cdot 2^{O(k/r)} \cdot n \log^2 n)$.
3. Now, using the algorithm of Corollary 2, compute subfamily $\hat{\mathcal{P}}_s$ of size $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$ that represents \mathcal{P}' . This takes time $O_k(|\mathcal{P}'| \cdot r^{6k/r} \cdot 2^{O(k/r)} \cdot n \log n) = O_k(r^{18k/r} \cdot 2^{O(k/r)} \cdot n \log^3 n)$.

Since the above procedure is applied to every gate of C , the claimed running time follows. \square

We remark that, after a trivial modification, the algorithm above can equally easily solve also a weighted variant of (r, k) -MONOMIAL DETECTION, where each variable is equipped with a weight and we are interested in extracting a monomial of minimum total weight, defined as the sum of the weights of variables times their degrees. It is not hard to reduce the problems r -SIMPLE k -PATH and (r, p, q) -PACKING, considered in the previous sections, to this variant; we leave the details to the reader.

C Proofs omitted from main text

Statement and proof of Theorem 5:

Fix integers n, r, k such that $1 < r \leq k$, and let $t \triangleq \lfloor 2k/r \rfloor$. Suppose a (t, k) -minimal separating family $\mathcal{H} \subseteq \{[n] \rightarrow [t+1]\}$ can be constructed in time $f(r, k, n)$. Then an (r, k) -separator \mathcal{F} of size $|\mathcal{H}| \cdot (r+1)^t$ can be constructed in time $O_k(f(r, k, \max(n, t)) \cdot (r+1)^t)$.

Proof. In the proof we will assume that $n \geq t$, since otherwise we can apply the same construction for n increased to t , and at the end remove from each multiset of the obtained \mathcal{F} all the elements from $[t] \setminus [n]$. Note that thus we apply the construction of a minimal separating family to the set of size $\max(n, t)$, rather than n . Also, observe that from the assumption that $r > 1$ it follows that $t \leq k$.

Let \mathcal{H} be the constructed (t, k) -minimal separating family of functions from $[n]$ to $[t+1]$. For each $h : [n] \rightarrow [t+1]$ in \mathcal{H} , and for each $w = (w_1, \dots, w_t) \in [r]_0^t$, we construct the following r -set $F^{h,w} \in [r]_0^n$. For all $j \in [t]$ and all $i \in [n]$ with $h(i) = j$, we put $F_i^{h,w} = w_j$. For all $i \in [n]$ with $h(i) = t+1$, we put $F_i^{h,w} = r/2$. Let \mathcal{F} consist of all the constructed r -sets $F_i^{h,w}$. Thus we have that

$$|\mathcal{F}| = |\mathcal{H}| \cdot (r+1)^t,$$

and \mathcal{F} clearly can be constructed in time as claimed in the theorem statement. We are left with proving that \mathcal{F} is indeed an (r, k) -separator.

Fix (r, k) -compatible (r, k) -sets $A, B \in [r]_0^n$. Let U be the set of all elements that appear in A or in B , that is, $U = \{i \in [n] \mid A_i > 0 \text{ or } B_i > 0\}$. Denote by $C_0 \subseteq U$ the sets of elements in A and B that appear more than $r/2$ times in one of the sets, that is, $C_0 = \{i \in [n] \mid A_i > r/2 \text{ or } B_i > r/2\}$. Note that since A and B are (r, k) -sets, we have that $|C_0| \leq \lfloor 2k/r \rfloor = t$. Let C be a superset of C_0 of size exactly t , constructed by augmenting C_0 with arbitrary elements of $U \setminus C_0$ up to size t , and if there is not enough of them, then by additionally augmenting it with the remaining number of arbitrary elements of $[n] \setminus U$. Note that this is always possible since $t \leq n$. Let $D = U \setminus C$. Since A and B are (r, k) -compatible, we have that $|U| \leq k$ and from the construction of C it follows that $|D| \leq k - t$.

Therefore, there exists some $h \in \mathcal{H}$ that separates C from D . For $j \in [t]$, let i_j be the unique element of C mapped to j under h . Choose $w_j \in [r]_0$ such that $A_{i_j} \leq w_j \leq r - B_{i_j}$; such w_j exists since A and B are (r, k) -compatible. Let $w = (w_1, \dots, w_t)$. We claim that $A \leq F^{h,w} \leq \bar{B}$. For $i \in C$, the choice of w guarantees that $A_i \leq F_i^{h,w} \leq \bar{B}_i$. For $i \in D$ we have $F_i^{h,w} = \lfloor r/2 \rfloor$, and from the definition of D we have that $D \cap C_0 = \emptyset$. So for such i it holds that $A_i \leq \lfloor r/2 \rfloor \leq \bar{B}_i$. Finally, for $i \notin C \cup D$ we have that $A_i = 0$ and $\bar{B}_i = r$, so surely $A_i \leq F_i^{h,w} \leq \bar{B}_i$. \square

Statement and proof of Corollary 1:

Fix integers n, r, k such that $1 < r \leq k$. Then an (r, k) -separator \mathcal{F} of size $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$ can be constructed in time $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot n \cdot \log n)$.

Proof. Let $t \triangleq \lfloor 2k/r \rfloor$. By Theorem 4, a (t, k) -minimal separating family $\mathcal{H} \subseteq \{[n] \rightarrow [t+1]\}$ of size $O_k((k/t)^{2t} \cdot 2^{O(t)} \cdot \log n)$ can be constructed in time $O_k((k/t)^{2t} \cdot 2^{O(t)} \cdot n \cdot \log n)$ from Theorem 4. Using this construction in Theorem 5, we obtain an (r, k) -separator \mathcal{F} such that

$$|\mathcal{F}| = |\mathcal{H}| \cdot (r+1)^t = O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n).$$

Moreover, from Theorem 5 it follows that \mathcal{F} can be constructed in time $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot n \cdot \log n)$. \square

Statement and proof of Lemma 1:

Let \mathcal{F} be an (r, k) -separator and let \mathcal{P} be a weighted family of (r, k) -sets. Then $\text{Trim}_{\mathcal{F}}(\mathcal{P})$ represents \mathcal{P} .

Proof. Fix an (r, k) -set Q and suppose there is an (r, k) -set $P \in \mathcal{P}$ that is (r, k) -compatible with Q . In particular, we have $|P| = k - |Q|$. Since \mathcal{F} is an (r, k) -separator, there exists $F \in \mathcal{F}$ such that $P \leq F \leq \overline{Q}$. As $P \leq F$, when constructing $\text{Trim}_{\mathcal{F}}(\mathcal{P})$ we must have inserted into it an (r, k) -set $P' \in \mathcal{P}$ of size $k - |Q|$ such that $P' \leq F$ and $\mathbf{w}(P') \leq \mathbf{w}(P)$. Therefore $P' \leq \overline{Q}$, which implies that $P' + Q$ is an (r, k) -set. As $|P' + Q| = k$, we have that P' and Q are (r, k) -compatible. \square

Statement and proof of Corollary 2:

There exists a deterministic algorithm that, given a weighted family \mathcal{P} of (r, k) -sets, runs in time $O_k(|\mathcal{P}| \cdot r^{6k/r} \cdot 2^{O(k/r)} \cdot n \log n)$ and returns a family $\hat{\mathcal{P}} \subseteq \mathcal{P}$ that represents \mathcal{P} and has size $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$.

Proof. Let \mathcal{F} be the (r, k) -separator of size $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot \log n)$ given by Corollary 1; recall that \mathcal{F} can be computed in time $O_k(r^{6k/r} \cdot 2^{O(k/r)} \cdot n \log n)$. We compute $\hat{\mathcal{P}} = \text{Trim}_{\mathcal{F}}(\mathcal{P})$ which represents \mathcal{P} by Lemma 1. The construction of $\text{Trim}_{\mathcal{F}}(\mathcal{P})$ amounts to going over all pairs of r -sets $P \in \mathcal{P}$ and $F \in \mathcal{F}$ and checking whether $P \leq F$. Thus, the computation takes time at most $O_k(|\mathcal{P}| \cdot |\mathcal{F}| \cdot n) = O_k(|\mathcal{P}| \cdot r^{6k/r} \cdot 2^{O(k/r)} \cdot n \log n)$. \square

Statement and proof of Theorem 8:

Unless ETH fails, there exists a constant $s > 0$ such that for no fixed integer $d \geq 2$ the DEGREE-BOUNDED SPANNING TREE problem with the degree bound d can be solved in time $O^*(2^{sn/d})$.

Proof. It is known (see e.g. [6]) that, assuming ETH, there exists a constant $s > 0$ such that the HAMILTONIAN PATH problem cannot be solved in time $O^*(2^{sn})$ on n -vertex graphs. Consider the following reduction from HAMILTONIAN PATH to DEGREE-BOUNDED SPANNING TREE with the degree bound d : given an instance G of HAMILTONIAN PATH, create G' by attaching to every vertex $v \in V(G)$ a set of $d - 2$ degree-1 vertices, adjacent only to v . Since the new vertices have to be leaves in every spanning tree of G' , it follows that every spanning tree T' of G' is in fact a spanning tree of G with all the vertices of $V(G') \setminus V(G)$ attached as leaves. In particular, G' admits a spanning tree with maximum degree d if and only if G admits a spanning tree with maximum degree 2, i.e., a hamiltonian path.

Observe that $|V(G')| = (d - 1) \cdot |V(G)|$. Hence, if there was an algorithm solving DEGREE-BOUNDED SPANNING TREE with the degree bound d in time $O^*(2^{sn/d})$, then by composing the reduction with the algorithm we would be able to solve HAMILTONIAN PATH on an n -vertex graph in time $O^*(2^{s(d-1)n/d}) \leq O^*(2^{sn})$, thus contradicting ETH. \square