

Deleting and Testing Forbidden Patterns in Multi-Dimensional Arrays

Omri Ben-Eliezer* Simon Korman† Daniel Reichman‡

March 27, 2017

Abstract

Understanding the local behaviour of structured multi-dimensional data is a fundamental problem in various areas of computer science. As the amount of data is often huge, it is desirable to obtain sublinear time algorithms, and specifically property testers, to understand local properties of the data.

We focus on the natural local problem of testing *pattern freeness*: given a large d -dimensional array A and a fixed d -dimensional pattern P over a finite alphabet Γ , we say that A is P -free if it does not contain a copy of the forbidden pattern P as a consecutive subarray. The distance of A to P -freeness is the fraction its entries that need to be modified to make it P -free. For any $\epsilon \in [0, 1]$ and any large enough pattern P – other than a very small set of exceptional patterns – we design a *tolerant tester* that distinguishes between the case that the distance is at least ϵ and the case that it is at most $a_d\epsilon$, with query complexity and running time $c_d\epsilon^{-1}$, where $a_d < 1$ and c_d depend only on the dimension d . For the 1-dimensional case, we provide a linear time algorithm for computing the distance from P -freeness.

To analyze the testers we establish several combinatorial results, including the following d -dimensional *modification lemma*, which might be of independent interest: for any large enough d -dimensional pattern P over any alphabet (excluding a small set of exceptional patterns for the binary case), and any d -dimensional array A containing a copy of P , one can delete this copy by modifying one of its locations without creating new P -copies in A .

Our results address an open question of Fischer and Newman, who asked whether there exist efficient testers for properties related to tight substructures in multi-dimensional structured data. They serve as a first step towards a general understanding of local properties of multi-dimensional arrays.

1 Introduction

Pattern matching is the algorithmic problem of finding occurrences of a fixed pattern in a given string. This problem appears in many settings and has applications in diverse domains such as computational biology, computer vision, natural language processing and web search. There has been extensive research concerned with developing algorithms that search for patterns in strings, resulting with a wide range of efficient algorithms [13, 26, 21, 15, 28, 27]. Higher-dimensional analogues where one searches for a d -dimensional pattern in a d -dimensional array have received

*Blavatnik School of Computer Science, Tel Aviv University, Tel Aviv, Israel omrib@mail.tau.ac.il

†Computer Science Department, University of California at Los-Angeles simon.korman@gmail.com

‡Electrical Engineering and Computer Science, University of California at Berkeley daniel.reichman@gmail.com

attention as well. For example, the 2D case arises in analyzing aerial photographs [7, 8] and the 3D case has applications in medical imaging.

Given a string S of length n and a pattern P of length $k \leq n$, any algorithm which determines whether P occurs in S has running time $\Omega(n)$ [14, 31] and a linear lower bound carries over to higher dimensions. For the 2D and 3D case, when an $n \times n$ image is concerned, algorithms whose run time is $O(n^2)$ are known [8]. These algorithms have been generalized to the 3D case to yield running time of $O(n^3)$ [20]. Finally it is also known (e.g., [23]) that for the d -dimensional case it is possible to solve the pattern matching problem in time $O(d^2 n^d \log m)$ (where the pattern is an array of size m^d). It is natural to ask which tasks of this type can be performed in sublinear (namely $o(n^d)$) time for d -dimensional arrays.

The field of *property testing* [22, 32] deals with decision problems regarding discrete objects (e.g., graphs, functions, images) that either have a certain property P or are *far* from satisfying P . Here, we are interested in deciding quickly whether a given d -dimensional array A is far from *not* containing a fixed d -dimensional pattern P . *Tolerant* property testing [29] is a useful extension of the standard notion, in which the tester needs to distinguish between objects that are *close* to satisfying the property to those that are *far* from satisfying it.

A d -dimensional $k_1 \times \dots \times k_d$ array A over an alphabet Γ is a function from $[k_1] \times \dots \times [k_d]$ to Γ . For simplicity of presentation, all results in this paper will be presented for cubic arrays in which $k_1 = \dots = k_d$, but they generalize to non-cubic arrays in a straightforward manner. We consider the (tolerant) *pattern-freeness problem* where one needs to distinguish between the case that a given d -dimensional array A is ϵ_1 -close to being P -free for a fixed pattern P , and the case that A is ϵ_2 -far from being P -free, where $\epsilon_1 < \epsilon_2$. An (ϵ_1, ϵ_2) -tester Q for this problem is a randomized algorithm that is given access to an array A , as well as its size and proximity parameters $0 \leq \epsilon_1 < \epsilon_2 < 1$. Q needs to distinguish with probability at least $2/3$ between the case that A is ϵ_1 -close to being P -free and the case that A is ϵ_2 -far from being P -free. The query complexity of Q is the number of queries it makes in A .

Our interest in the pattern-freeness problem stems from several applications. In certain scenarios of interest, we might be interested in identifying quickly that an array is far from not containing a given pattern. For the one dimensional case, being far from not containing a given text may indicate a potential anomaly which requires attention (e.g., an offensive word in social network media), hence such testing algorithms may provide useful in anomaly detection. Many computer vision methods for classifying images are feature based: hence being far from containing a certain pattern associated with a feature may be useful in rejection methods that enable to quickly discard images that do not possess a certain visual property.

Beyond practical applications, devising property testing algorithms for the pattern freeness problem is of theoretical interest. In the first place, it leads to a combinatorial characterization of the distance from being P -free. Such a characterization has proved fruitful in graph property testing [3, 4] where celebrated graph removal lemmas were developed en route of devising algorithms for testing subgraph freeness. We encounter a similar phenomena in studying patterns and arrays: at the core of our approach for testing pattern freeness lies a *modification lemma* for patterns which we state next. We believe that this Lemma may be of independent interest and find applications beyond testing algorithms. Later we show one such application: computing the exact distance of a (one dimensional) string from being P -free can be done in linear time.

For a pattern P of size $k \times k \times \dots \times k$, any of its entries that is in $\{0, k-1\} \times \dots \times \{0, k-1\}$ is said to be a *corner* of P . We say that P is *almost homogeneous* if all of its entries but one are

equal, and the different entry lies in a corner of P . Finally, P is *removable* (with respect to the alphabet Γ) if for any d -dimensional array A over Γ and any copy of P in A , one can destroy the copy by modifying one of its entries without creating new P -copies in A . The modification lemma states that for any d , and any large enough pattern P , when the alphabet is binary it holds that P is removable *if and only if* it is not almost homogeneous, and when the alphabet is not binary, P is removable provided that it is large enough.

Recent works [11, 12] have obtained tolerant testers for visual properties. As observed in [11, 12], tolerance is an attractive property for testing visual properties as real-world images are often noisy. With the modification Lemma at hand, we show that when P is removable, the (relative) *hitting number* of P in A , which is the minimal size of a set of entries that intersects all P -copies in A divided by $|A|$, differs from the distance of A from P -freeness by a multiplicative factor that depends only on d (and not on P or A). This relation allows us to devise very fast $(5^{-d}\epsilon, \epsilon)$ -tolerant testers for P -freeness, as the hitting number of P in A can be well approximated using only a very small sample of blocks of entries from A . The query complexity of our tester is $O(C_d/\epsilon)$, where C_d is a positive constant depending only on the dimension d of the array. Note that our characterization in terms of the hitting number is crucial: merely building on the fact that A contains many occurrences of P (as can be derived directly from the modification lemma) and randomly sampling $O(1/\epsilon)$ possible locations in A , checking whether the sub-array starting at these locations equals P would lead to query complexity of $O(k^d/\epsilon)$. Note that our tester is optimal (up to a multiplicative factor that depends on d), as any tester for this problem makes $\Omega(1/\epsilon)$ queries.

The one dimensional setting, where one seeks to determine quickly whether a string S is ϵ -far from being P -free is of particular interest. We are able to leverage the modification Lemma and show that the distance of a string S from being P -free for a fixed pattern P (that is not almost homogeneous) is *exactly equal* to the hitting number of P in A . For an arbitrary constant $0 < c < 1$, this characterization allows us to devise a $((1 - c)\epsilon, \epsilon)$ -tolerant tester making $O_c(\epsilon^{-1})$ queries for this case. For the case of almost homogeneous patterns, and an arbitrary constant $c > 0$, we devise a $((1/16 + c)\epsilon, \epsilon)$ -tolerant tester that makes $O_c(1/\epsilon)$ queries. Whether tolerant testers exist for almost homogeneous patterns of dimension larger than 1 is an open question.

Moreover, the characterization via the hitting number implies an $O(n + k)$ algorithm that calculates (exactly) the distance of A from being P -free where P is an *arbitrary* pattern (that may be almost homogeneous). We are not aware of a previous algorithm for the distance computation problem. Unlike the one-dimensional case, in d dimensions we do not know of a clean combinatorial description of the distance to being P -free for higher dimension. Furthermore, it can be shown via a direct reduction from covering problems in the plane [19], that for dimension $d > 1$ there exists patterns P for which calculating the distance to P -freeness is NP-hard.

2 Related Work

The problem of testing pattern freeness is related to the study of testing subgraph-freeness (see, for example, [1, 4]). This line of work examines how one can test quickly whether a given graph G is H -free or ϵ -far from being H -free, where H is a fixed subgraph. In this problem, a graph is ϵ -far from being H -free if at least an ϵ -fraction of its edges and non-edges need to be altered in order to ensure that the resulting graph does not contain H as a (not necessarily induced) subgraph. A key component in these works are removal lemmas: typically such lemmas imply that if G is ϵ -far from being H -free, it necessarily contains a “large” number of copies of H . Perhaps the best example for

this phenomena is the triangle removal lemma which asserts that for every $\epsilon \in (0, 1)$, there exists $\delta = \delta(\epsilon) > 0$ such that if an n -vertex graph G is ϵ -far from being triangle free, then G contains at least δn^3 triangles (see e.g., [6] and the reference within).

Alon *et. al.* showed [3] that regular languages over $\{0, 1\}$ are strongly testable. Testing pattern-freeness (1-dimensional, binary alphabet, constant pattern length k) is a special case of the former, since the language of all strings avoiding a fixed pattern is regular. The query complexity of their tester is $O\left(\frac{c}{\epsilon} \cdot \ln^3\left(\frac{1}{\epsilon}\right)\right)$, where c is a constant that depends on the minimal size of a DFA A_L , that accepts the regular language L . It is shown in [3] that c can be taken to be $O(s^3)$ where s is the size of A_L . In the case of the regular language considered here a simple pumping-lemma inspired argument shows that $s \geq \Omega(k)$. Hence the upper bound on testing pattern freeness implied by their algorithm is $O\left(\frac{k^3}{\epsilon} \cdot \ln^3\left(\frac{1}{\epsilon}\right)\right)$. Our 1D tester solves a very restricted case of the problem the tester of [3] deals with, but it achieves a better query complexity of $O(1/\epsilon)$ in this setting. Moreover, our tester is much simpler and can be applied in the more general high dimensional setting, or when the pattern length k is allowed to grow as a function of the string length n .

The problem of testing *submatrix freeness* was investigated in [16, 17, 5, 18, 2]. As opposed to our case, which is concerned with *tight* submatrices, all of these results deal with submatrices that are not necessarily tight (i.e. the rows and the columns need not be consecutive). Quantitatively, the submatrix case is very different from our case: in our case P -freeness can be testable using $O(\epsilon^{-1})$ queries, while in the submatrix case, for a binary submatrix of size $k \times k$ a lower bound of $\epsilon^{-\Omega(k^2)}$ on the needed number of queries is easy to obtain, and in the non-binary case there exist 2×2 matrices for which there exists a super polynomial lower bound of $\epsilon^{\Omega(\log 1/\epsilon)}$.

The 2D part of our work adds to a growing literature concerned with testing properties of images [30, 33, 11]. Ideas and techniques from the property testing literature have recently been used in the fields of computer vision and pattern recognition [24, 25].

3 Notation and definitions

With slight abuse of notation, for a positive integer n we let $[n]$ denote the set $\{0, \dots, n - 1\}$ and we write $[n]^d = [n] \times \dots \times [n]$.

Recall that a d -dimensional (cubic) array A over an alphabet Γ is a function from $[k]^d$ to Γ . The $x = (x_1, \dots, x_d)$ *entry* of A , denoted by A_x , is the value of the function A at location x . Let P be a (k, d) -array over an alphabet Γ of size at least two. We say that a d -dimensional array A *contains* a copy of P (or a P -copy) *starting in location* $x = (x_1, \dots, x_d)$ if for any $y \in [k]^d$ we have $A_{x+y} = P_y$. Finally, A is P -free if it does not contain copies of P .

A property \mathcal{P} of d -dimensional arrays is simply a family of such arrays over an alphabet Γ . For an array A and a property \mathcal{P} , the *absolute distance* $d_{\mathcal{P}}(A)$ of A to \mathcal{P} is the minimal number of entries that one needs to change in A to get an array from \mathcal{P} . The *relative distance* of A to \mathcal{P} is $\delta_{\mathcal{P}}(A) = d_{\mathcal{P}}(A)/|A|$, where clearly $0 \leq \delta_{\mathcal{P}}(A) \leq 1$ for any nontrivial \mathcal{P} and A . We say that A is ϵ -close (ϵ -far) to \mathcal{P} if $\delta_{\mathcal{P}}(A) \leq \epsilon$ ($\geq \epsilon$).

In this paper we will consider the property of P -freeness, which consists of all P -free arrays. The absolute and relative distance to P -freeness will be denoted by $d_P(A)$ and $\delta_P(A)$, respectively.

For an array A and a pattern P we will call a set of entries in A whose modification can turn it to be P -free a *deletion set* and therefore it is natural to call $d_P(A)$ (the absolute distance of A to P -freeness) the *deletion number*, since it is the size of a minimal deletion set. In a similar manner, for a given set of entries in A , if every P -copy in A contains at least one of these entries, we call

it a *hitting set* and we call the size of a minimal hitting set the *hitting number*, denoted by $h_{\mathcal{P}}(A)$. For all notations here and above, in the 1-dimensional case we will replace A by S (for String).

We will need several definitions from [29] as well. Let \mathcal{P} be a property of arrays and let $h_1, h_2 : [0, 1] \rightarrow [0, 1]$ be two monotone increasing functions. An (h_1, h_2) -*distance approximation* algorithm for \mathcal{P} is given query access to an unknown array A . The algorithm outputs an estimate $\hat{\delta}$ to $\delta_{\mathcal{P}}(A)$, such that with probability at least $2/3$ it holds that $h_1(\delta_{\mathcal{P}}(A)) \leq \hat{\delta} \leq h_2(\delta_{\mathcal{P}}(A))$. Finally, for a property \mathcal{P} and for $0 \leq \epsilon_1 < \epsilon_2 \leq 1$, an (ϵ_1, ϵ_2) -*tolerant tester* for \mathcal{P} is given query access to an array A . The tester *accepts* with probability at least $2/3$ if A is ϵ_1 -close to \mathcal{P} , and *rejects* with probability at least $2/3$ if A is ϵ_2 -far from \mathcal{P} . In the ‘standard’ notion of property testing, $\epsilon_1 = 0$. Thus, any tolerant tester is also a tester in the standard notion. Finally, we define the additive (multiplicative) *tolerance* of the tester above as $\epsilon_2 - \epsilon_1$ (ϵ_2/ϵ_1 respectively).

4 Main Results

The modification lemma result is central in the study of minimal deletion sets. It classifies the possible patterns into ones that are removable and ones that are not. The result that the vast majority of patterns are removable is used extensively throughout the paper in the design and proofs of algorithms for efficient testing of pattern freeness (in 1 and higher dimensions) as well as for the exact computation of the deletion number in 1-dimension.

Our 1-dimensional modification lemma (Lemma 9) gives the following full characterization of 1-dimensional patterns (i.e. strings). A binary pattern is removable *if and only if* it not almost homogeneous, while *any* pattern over a larger alphabet is removable. The multidimensional version of the lemma (Lemma 9) makes the exact same classification, but for (k, d) -arrays for which $k \geq 3 \cdot 2^d$.

The fact that most patterns are removable is very important for analyzing the deletion number (which is the distance to pattern freeness). As an example, a simple observation is that a removable pattern appears at least $d_{\mathcal{P}}(A)$ times (possibly with overlaps) in the array A , which implies an ϵ -tester that can simply check for the presence of the pattern in $1/\epsilon$ random locations in the array at a sample complexity of $O(k/\epsilon)$.

Another important part of our work makes explicit connections between the deletion number and the hitting number for both 1 and higher dimensions. These are needed in order to get improved testers (e.g. for getting rid of k in the sample complexity) in d -dimensions as well as for linear time computation of the distance (deletion number) in 1-dimension.

For the 1-dimensional case we show that the deletion number $d_{\mathcal{P}}(S)$ equals the hitting number $h_{\mathcal{P}}(S)$, which leads to an exact computation of $d_{\mathcal{P}}(S)$ in time $O(n + k)$ (Theorem 12) as well as a tolerant testers for Pattern Freeness: An (ϵ_1, ϵ_2) -tolerant tester for any $0 \leq \epsilon_1 < \epsilon_2 \leq 1$ at a complexity of $O(\epsilon_2^2/(\epsilon_2 - \epsilon_1)^3)$ (Theorem 22) as well as an $((1 - \tau)\epsilon, \epsilon)$ -tolerant tester for a fixed $\tau > 0$ and any $0 < \epsilon \leq 1$ at a complexity of $O(\epsilon^{-1}\tau^{-3})$ (Corollary 13).

For higher dimensions, we show (Lemma 11) that $h_{\mathcal{P}}(A) \leq d_{\mathcal{P}}(A) \leq \alpha_d h_{\mathcal{P}}(A) \leq \alpha_d k^{-d}$, a bound that relates the hitting number $h_{\mathcal{P}}(A)$ and the deletion number $d_{\mathcal{P}}(A)$ through a constant $\alpha_d = 4^d + 2^d$ that depends only on the dimension d . This bound enables a $((1 - \tau)^d \alpha_d^{-1} \epsilon, \epsilon)$ -tolerant tester making $C_{\tau} \epsilon^{-1}$ queries, where $C_{\tau} = O(1/\tau^d (1 - (1 - \tau)^d)^2)$ (Theorem 15).

In the 1-dimensional setting we also provide dedicated algorithms to handle the almost homogeneous (non-removable) patterns, achieving an $O(n)$ algorithm for computing the deletion number (Theorem 23) as well as a $(\epsilon/(16 + c), \epsilon)$ -tolerant tester, for any constant $c > 0$, at a complexity of $\alpha_c \epsilon^{-1}$ queries, where α_c depends only on c (Theorem 24).

Finally, we provide a lower bound of $\Omega(1/\epsilon)$ (Theorem 27, Appendix D) for any tester of pattern freeness. Unlike the previous lower bound of $\Omega(1/\epsilon)$ [3] on testing regular languages, ours extends to dimensions higher than 1 and applies to the case where k may depend on n .

Our main results are summarized in Table 1.

dim.	template type	deletion number computation	modification lemma	tester tolerance	query complexity
1D	general	$O(n+k)$	removable for any k	$1/(1-\tau)$	$O(1/\epsilon\tau^3)$
	almost homog.	$O(n+k)$	not removable for any k	$(16+c)$	α_c/ϵ
2+D	general	NP-Hard	removable for $k > 3 \cdot 2^d$	$(1-\tau)^{-d}\alpha_d$	$\beta_{d,\tau}/\epsilon$
	almost homog.	—	not removable for any k	—	—

Table 1: **Summary of results.** $0 < \tau < 1$ and $c > 0$ are arbitrary constants. α_c is a constant that depends only on c . $\beta_{d,\tau}$ is a constant that depends only on d and τ . 'modification lemma' specifies if patterns are classified as removable or not. the 'tester tolerance' is multiplicative

5 Modification Lemma

Theorem 1 (Modification Lemma). *Let $d > 1$ and let P be a (k, d) -array over the alphabet Γ where $k \geq 3 \cdot 2^d$.*

1. *If $|\Gamma| = 2$ then P is removable if and only if it is not almost homogeneous.*
2. *If $|\Gamma| \geq 3$ then P is removable.*

Remark. Theorem 1 states that any large enough binary pattern which is not almost homogeneous is removable. The requirement that the pattern is large enough is crucial, as the $2 \times \dots \times 2$ pattern P satisfying $P_x = 0$ for any $x = (x_1, \dots, x_d)$ with $x_1 = 0$ and $P_x = 1$ otherwise is not removable even though it is not almost homogeneous. To see this, consider the following $4 \times \dots \times 4$ array A : $M_x = 0$ if either $x_1 = 0$, or $x_1 = 1$ and $x_i \in \{1, 2\}$ for any $2 \leq i \leq d$, or $x_1 = 2$ and $x_i \in \{0, 3\}$ for some $2 \leq i \leq d$. For any other value of x , $M_x = 1$. Note that A contains a copy of P starting at $(1, \dots, 1)$, but flipping any bit in this copy creates a new P -copy in A . Still, the size of the counterexample is only $2 \times \dots \times 2$ while in the statement of Theorem 1, the dependence is exponential in d . It will be interesting to understand what is the correct order of magnitude of the dependence of k on d .

Proof of Theorem 1. The second statement of the theorem can be easily derived from the first statement; if P does not contain all letters in Γ then it is clearly removable, as changing any of its entries to any of the missing letters cannot create new P -copies. Otherwise, we can reduce the problem to the binary case: let σ_1, σ_2 be the letters in Γ that appear the smallest number of times in P . Consider the following (k, d) -array P' over $\{0, 1\}$: $P'_x = 0$ if $P_x \in \{\sigma_1, \sigma_2\}$ and $P'_x = 1$ otherwise. Observe that P' is not almost homogeneous, implying that it is removable. It is not hard to verify now that P is removable as well.

In what follows, we will prove the first statement. If P is binary and almost homogeneous then it is not removable: Without loss of generality $P_{(0, \dots, 0)} = 1$ and $P_x = 0$ for any $x \neq (0, \dots, 0)$. Consider a $(2k, d)$ -array A such that $M_{(0, \dots, 0)} = M_{(1, \dots, 1)} = 1$ and $A = 0$ elsewhere. Clearly, modifying any bit of the P -copy starting at $(1, \dots, 1)$ creates a new copy of P in A , so P is not removable.

The rest of the proof is dedicated to the other direction. Suppose that P is a binary (k, d) -array that is not removable. We would like to show that P must be almost homogeneous. As P is not removable, there exists a binary array A containing a copy of P such that flipping any single bit in this copy creates a new copy of P in A . This copy of P will be called the *template* of P in A .

Clearly, all of the new copies created by flipping bits in the template must intersect the template, so we may assume that A is of size $(3k - 2)^d$ and that the template starts in location $k = (k - 1, \dots, k - 1)$.


For convenience, let $I = [k]^d$ denote the set of indices of P . For any $x \in I$ let $\bar{x} = x + k$; \bar{x} is the location in A of bit x of the template.

Roughly speaking, our general strategy for the proof would be show that there exists at most two "special" entries in P such that when we flip a bit in the template, creating a new copy of P in A , the flipped bit usually plays the role of one of the special entries in the new copy. We will then show that in fact, there must be exactly one special entry, which must lie in a corner of P , and that all non-special entries are equal while the special entry is equal to their negation. This will finish the proof that P is almost homogeneous.

Definition 2. Let $i \leq d$ and let δ be positive integers. Let $x = (x_1, \dots, x_d)$ and $y = (y_1, \dots, y_d)$ be d -dimensional points. The pair (x, y) is (i, δ) -related if $y_i - x_i = \delta$ and $y_j = x_j$ for any $j \neq i$. An (i, δ) -related pair (x, y) is said to be an (i, δ) -jump in P if $P_x \neq P_y$.

Figure 1: **Illustration for Lemma 3.**

A 2-dimensional example, where i is the vertical coordinate: Flipping the bit (of the template P) at location \bar{a} creates the P -copy Q^a at location $m(a)$. Similarly, the copy Q^b is created at location $m(b)$. Note that the pair of points (\bar{x}, \bar{y}) (which is (x, y) in P) and the copy locations pair $(m(a), m(b))$ are both (i, Δ_i) -related. The values P_x and P_y ($M_{\bar{x}}$ and $M_{\bar{y}}$) must be equal.



Lemma 3. For any $1 \leq i \leq d$ there exists $0 < \Delta_i < k/3$ such that at most two of the (i, Δ_i) -related pairs of points from I are (i, Δ_i) -jumps in P .

Proof. Recall that, by our assumption, flipping any of the $K = k^d$ bits of the template creates a new copy of P in A . Consider the following mapping $m : I \rightarrow [2k - 1]^d$. $m(x_1, \dots, x_d)$ is the starting location of a new copy of P created in A as a result of flipping bit $x = (x_1, \dots, x_d)$ of the template (which is bit \bar{x} of A). If more than one copy is created by this flip, then we choose the starting location of one of the copies arbitrarily.

Observe that m is injective, and let S be the image of m , where $|S| = K$. Let $1 \leq i \leq d$ and consider the collection of (one-dimensional) lines

$$\mathcal{L}_i = \left\{ \{x_1\} \times \dots \times \{x_{i-1}\} \times [2k - 1] \times \{x_{i+1}\} \times \dots \times \{x_d\} \mid \forall j \neq i : x_j \in [2k - 1] \right\}.$$

Clearly $\sum_{\ell \in \mathcal{L}_i} |S \cap \ell| = K$. On the other hand, $|\mathcal{L}_i| = \prod_{j \neq i} (2k - 1) < 2^{d-1} \prod_{j \neq i} k = 2^{d-1} K/k$, so there exists a line $\ell \in \mathcal{L}_i$ for which $|S \cap \ell| > k/2^{d-1} \geq 6$. Hence $|S \cap \ell| \geq 7$. Let $\alpha_1 < \dots < \alpha_7$ be the smallest i -indices of elements in $S \cap \ell$. Since $\alpha_7 - \alpha_1 < 2k - 1$ there exists some $1 \leq l \leq 6$ such that $\alpha_{l+1} - \alpha_l < k/3$. That is, S contains an (i, Δ_i) -related pair with $0 < \Delta_i < k/3$. In other words, there are two points $a, b \in I$ such that flipping \bar{a} (\bar{b}) would create a new P -copy, denoted by Q^a (Q^b respectively), which starts in location $m(a)$ ($m(b)$ respectively) in A , and $(m(a), m(b))$ is an (i, Δ_i) -related pair.

The following claim finishes the proof of the lemma and will also be useful later on.

Claim 4. *For a and b as above, let (x, y) be a pair of points from I that are (i, Δ_i) -related and suppose that $y \neq \bar{a} - m(a)$ and that $x \neq \bar{b} - m(b)$. Then $P_x = P_y$.*

Proof. The bits that were flipped in A to create Q^a and Q^b are \bar{a}, \bar{b} respectively. Since $y + m(a) \neq \bar{a}$, the copy Q_a contains the original entry of A in location $y + m(a)$. Therefore, $P_y = M_{y+m(a)}$ (as $M_{y+m(a)}$ is bit y of Q^a , which is a copy of P). Similarly, since $x + m(b) \neq \bar{b}$, we have $P_x = M_{x+m(b)}$. But since both pairs (x, y) and $(m(a), m(b))$ are (i, Δ_i) -related, we get that $m(b) - m(a) = y - x$, implying that $x + m(b) = y + m(a)$, and therefore $P_x = M_{x+m(b)} = M_{y+m(a)} = P_y$, as desired. \square

Clearly, the number of (i, Δ_i) -related pairs that do not satisfy the conditions of the claim is at most two, finishing the proof of Lemma 3. \square

Let $\Delta = (\Delta_1, \dots, \Delta_d)$ where for any $1 \leq i \leq d$, we take Δ_i that satisfies the statement of Lemma 3 (its specific value will be determined later).

Definition 5. Let $x \in I$. The set of Δ -neighbours of x is

$$N_x = \{y \in I \mid \exists i : (x, y) \text{ is } (i, \Delta_i)\text{-related or } (y, x) \text{ is } (i, \Delta_i)\text{-related}\}$$

and the number of Δ -neighbours of x is $n_x = |N_x|$, where $d \leq n_x \leq 2d$. We say that x is a Δ -corner if $n_x(\Delta) = d$ and that it is Δ -internal if $n_x(\Delta) = 2d$. Furthermore, x is (Δ, P) -isolated if $P_x \neq P_y$ for any $y \in N_x$, while it is (Δ, P) -generic if $P_x = P_y$ for any $y \in N_x$.

When using the above notation, we will sometimes omit the parameters (e.g. simply writing *isolated* instead of (Δ, P) -isolated) as the context is usually clear.

The definition imposes a symmetric neighborhood relation, that is, $x \in N_y$ holds if and only if $y \in N_x$. If $x \in N_y$ we say that x and y are Δ -neighbours. Note that a point $x = (x_1, \dots, x_d) \in I$ is a Δ -corner if $x_i < \Delta_i$ or $x_i \geq k - \Delta_i$ for any $1 \leq i \leq d$, and that x is Δ -internal if $\Delta_i \leq x_i < k - \Delta_i$ for any $1 \leq i \leq d$.


Claim 6. *Two (Δ, P) -isolated points in I cannot be Δ -neighbors.*

Proof. Suppose towards contradiction that $x = (x_1, \dots, x_d)$ and $y = (y_1, \dots, y_d)$ are two distinct (Δ, P) -isolated points and that (x, y) is (i, Δ_i) -related for some $1 \leq i \leq d$. Since $\Delta_i < k/3$, at least one of x or y participates in two different (i, Δ_i) -related pairs: if $x_i < k/3$ then $y_i + \Delta_i = x_i + 2\Delta_i < k$ so y is in two such pairs, and otherwise $x_i \geq \Delta_i$, meaning that x participates in two such pairs. Assume without loss of generality that the two (i, Δ_i) -related pairs are (t, x) and (x, y) , then $P_t \neq P_x$ and $P_x \neq P_y$ as x is isolated. By Lemma 3, these are the only (i, Δ_i) -jumps in P .

Choose an arbitrary $j \neq i$ and take $v = (v_1, \dots, v_d)$ where $v_j = \Delta_j$ and $v_l = 0$ for any $l \neq j$. Recall that $\Delta_j < k/3$, implying that either $x_j + v_j < k$ or $x_j - v_j \geq 0$. Without loss of generality assume the former, and let $x' = x + v$ and $y' = y + v$. Since x and y are (Δ, P) -isolated, and since

$x' \in N_x$ and $y' \in N_y$, we get that $P_{x'} \neq P_x \neq P_y \neq P_{y'}$, and thus $P_{x'} \neq P_{y'}$ (as the alphabet is binary). Therefore, (x', y') is also an (i, Δ_i) -jump in P , a contradiction. \square

Illustration for Definition 7. Recall that flipping a bit \bar{a} in A creates a new P -copy Q^a (which contains \bar{a}), located at the point $m(a)$ in the coordinates of A . The bits x and a are mapped to y and $f(a)$ respectively.



Definition 7. For three points $x, y, a \in I$, we say that x is mapped to y as a result of the flipping of a if $\bar{x} = m(a) + y$. Moreover, define the function $f : I \rightarrow I$ as follows: $f(x) = \bar{x} - m(x)$ is the location to which x is mapped as a result of flipping x .

In other words, x is mapped to y as a result of flipping the bit a if bit \bar{x} of A "plays the role" of bit y in the new P -copy Q_a that is created by flipping a . Note that

- If $\bar{x} - m(a) \notin I$ then x is not mapped to any point. However, this cannot hold when $x = a$, so the function f is well defined.
- For a fixed a , the mapping as a result of flipping a is linear: if x and y are mapped to x' and y' respectively, then $y - x = y' - x'$. In particular, if (x, y) is (i, Δ_i) -related for some $1 \leq i \leq d$ then (x', y') is also (i, Δ_i) -related.
- If x is mapped to y as a result of flipping a and $x \neq a$, then $P_x = P_y$.
- On the other hand, we always have $P_x \neq P_{f(x)}$.
- If x is Δ -internal and (Δ, P) -generic, then $f(x)$ must be (Δ, P) -isolated.

The first four statements are easy to verify. To verify the last one, suppose that x is internal and generic and let $z \in N_{f(x)}$; we will show that $P_{f(x)} \neq P_z$. Since x is internal, there exists $y \in N_x$ such that $y - x = z - f(x)$. Then y is mapped to z as a result of flipping x , since $\bar{y} = y + k = z + (x + k) - f(x) = z + \bar{x} - f(x) = z + m(x)$. Therefore $P_y = P_z$. On the other hand, $P_x = P_y$ as x is generic and $P_x \neq P_{f(x)}$, and we conclude that $P_z \neq P_{f(x)}$.

Lemma 8. *There is exactly one (Δ, P) -isolated point in I .*

Proof. Let \mathcal{S} be the set of isolated points; our goal is to show that $|\mathcal{S}| = 1$. Consider the set

$$C = \{(x, y) : x, y \in I, (x, y) \text{ is an } (i, \Delta_i)\text{-jump for some } 1 \leq i \leq d\}.$$

Clearly, each point in \mathcal{S} is contained in at least d pairs from C . By claim 6 no pair of isolated points are Δ -neighbours and therefore every pair in C contains at most one point from \mathcal{S} . By Lemma 3, $|C| \leq 2d$ which implies that $|\mathcal{S}| \leq 2$. On the other hand we have $|\mathcal{S}| \geq 1$. To see this, observe that the number of (Δ, P) -internal points in I is greater than $\prod_{i=1}^d k/3 \geq 2^{d^2}$, while the number of non- Δ -generic points is at most $2|C| \leq 4d$, implying that at least $2^{d^2} - 4d > 0$ of the internal

points are generic. Therefore, pick an internal generic point $z \in I$. As we have seen before, $f(z)$ must be isolated.

To complete the proof it remains to rule out the possibility that $|\mathcal{S}| = 2$. If two different (Δ, P) -isolated points $a = (a_1, \dots, a_d)$ and $b = (b_1, \dots, b_d)$ exist, each of them must participate in exactly d pairs in C . This implies that both of them are Δ -corners with d neighbors. It follows that every Δ -internal point z must be generic (since an internal point and a corner point cannot be neighbours), implying that either $f(z) = a$ or $f(z) = b$.

Let $1 \leq i \leq d$ and define $\delta_i > 0$ to be the smallest integer such that there exists an (i, δ_i) -related pair (x, y) of generic internal points with $f(x) = f(y)$. For this choice of x and y we have $m(y) - m(x) = \bar{y} - f(y) - (\bar{x} - f(x)) = \bar{y} - \bar{x} = y - x$, so $(m(x), m(y))$ is also (i, δ_i) -related. In particular, we may take $\Delta_i = \delta_i$ (Recall that until now, we only used the fact that $\Delta_i < k/3$, without committing to a specific value). Without loss of generality we may assume that $f(x) = f(y) = a$. By Claim 4, any pair (s, t) of (i, Δ_i) -related points for which $s \neq \bar{y} - m(y) = f(y) = a$ and $t \neq \bar{x} - m(x) = f(x) = a$ is not an (i, Δ_i) -jump. Since b is not a Δ -neighbour of a , it does not participate in any (i, Δ_i) -jump, contradicting the fact that it is (Δ, P) -isolated. This finishes the proof of the lemma. \square

Finally, we are ready to show that P is almost homogeneous. Let $a = (a_1, \dots, a_d)$ be the single (Δ, P) -isolated point in I . Consider the set

$$J = \{x = (x_1, \dots, x_d) \in I : \Delta_i \leq x_i < \Delta_i + 2^d \text{ for any } 1 \leq i \leq d\}$$

and note that all points in J are Δ -internal. Let $1 \leq i \leq d$ and partition J into $(i, 1)$ -related pairs of points. There are $2^{d^2-1} \geq 4d$ pairs in the partition. On the other hand, the number of non-generic points in J is at most $2|C| - (d-1) < 4d$ (to see it, count the number of elements in pairs from C and recall that a is contained in at least d pairs). Therefore, there exists a pair (x, y) in the above partition such that x and y are both generic. As before, $f(x)$ and $f(y)$ must be isolated, and thus $f(x) = f(y) = a$, implying that $\Delta_i = \delta_i = 1$. We conclude that $\Delta = (1, \dots, 1)$.

Claim 4 now implies that any pair (s, t) of $(i, 1)$ -related points for which $s \neq \bar{y} - m(y) = f(y) = a$ and $t \neq \bar{x} - m(x) = f(x) = a$ is not an $(i, 1)$ -jump. That is, for any two neighbouring points $s, t \neq a$ in I , $P_s = P_t$, implying that $P_x = P_y$ for any $x, y \neq a$ (since $\Delta = (1, \dots, 1)$, a Δ -neighbour is a neighbour in the usual sense). To see this, observe that for any two points $x, y \neq a$ there exists a path $x_0 x_1 \dots x_t$ in I where x_j and x_{j+1} are neighbours for any $0 \leq j \leq t-1$, the endpoints are $x_0 = x$ and $x_t = y$, and $x_j \neq a$ for any $0 < j < t$. Since a is isolated, it is also true that $P_a \neq P_x$ for any $x \neq a$.

To finish the proof that P is almost homogeneous, it remains to show that a is a corner. Suppose to the contrary that $0 < a_i < k-1$ for some $1 \leq i \leq d$ and let $b, c \in I$ be the unique points such that (a, b) and (c, a) are $(i, 1)$ -related, respectively. Clearly $f(b) = a$, so a is mapped to $\bar{a} - m(b) = \bar{a} - \bar{b} + f(b) = c - a + a = c$ as a result of flipping b , which is a contradiction - as $P_a \neq P_c$ and $b \neq a, c$. This finishes the proof. \square

The above proof only works when the dimension is bigger than one, though it can be adapted to the one-dimensional case. However, we present here another proof for the one-dimensional case, which is simpler than the general proof above and works for any pattern which is not almost homogeneous (as opposed to the proof above, that required the forbidden pattern to also be large enough). The main strategy here is to consider the longest streaks of zeros and ones in the pattern - a strategy that cannot be used in higher dimensions.

Theorem 9 (1D Modification Lemma). *A one-dimensional pattern is removable if and only if it is almost homogeneous.*

Proof of Theorem 9. The reduction from a general alphabet to a binary one and the negative example for almost homogeneous patterns which were presented in the proof of Theorem 1 also hold here. It remains to prove that any 1-dimensional binary pattern that is not almost homogeneous is removable.

Let $P = P_0 \dots P_{k-1}$ be a binary pattern of length k , that is contained in an arbitrary binary string S . We need to show that one can flip one of the bits of P without creating a new P -copy in S . We assume that P contains both 0s and 1s (i.e. it is not homogeneous) otherwise flipping any bit would work. Therefore we can assume from now that $k \geq 3$ (since for $k = 1, 2$ all patterns are homogeneous or almost homogeneous).

Let us assume also that P starts with a 1, i.e. $P_0 = 1$ and let $t \leq k - 1$ be the length of the longest 0-streak (sub-string of consecutive 0s) in P . Let $i > 0$ be the leftmost index in which such a 0-streak of length t begins. Clearly, $P_{i-1} = 1$ and $P_i = \dots = P_{i+t-1} = 0$.

If $i + t \leq k$ (i.e. the streak is not at the end of P) then $P_{i+t} = 1$ and in such a case if we modify P_{i+t} to 0, the copy of P is removed without creating new P -copies in S . To see this, observe that a new copy cannot start at the bit flip location $i + t$ or within the 0-streak at any of its locations $i, \dots, i + t - 1$ since the bits in these locations are 0 while the starting bit of P is 1. On the other hand, a new copy cannot start after $i + t$ since it must include the bit flip location or anywhere before P_i since otherwise it would contain a 0-streak of length $t + 1$.

This implies that P contains exactly one 0-streak of length t at its last t locations. In particular, we have that at the last location $P_{k-1} = 1$, and if we denote by r the length of the longest 1-streak in P , a symmetric reasoning shows that P begins with its only longest 1-streak of length r .

If P is *not* of the form $1^s 0^t$, it can be verified that flipping P_s (the leftmost 0 in P) to 1 does not create any P -copy. The only case left is $P = 1^s 0^t$, where $s, t \geq 2$ since P is not almost homogeneous. Consider the bit of the string S that is to the left of P . If it is a 0 then we flip P_1 to 0 and otherwise, we flip P_0 to 0, where in both cases no new copy is created. \square

6 Characterizations of the Deletion Number

We use the modification lemmas of Section 5 to investigate several combinatorial characterizations of the deletion number, which will in turn allow exact (and efficient) computations of the deletion number in the 1-dimensional case, as well as efficient approximation and testing of pattern freeness for removable patterns in the d -dimensional case for any d .

In particular, we prove some surprising connections between minimal deletion sets and minimal hitting sets. The characterizations for almost homogeneous 1-dimensional patterns are given in Appendix C, along with an optimal algorithm to compute the exact deletion number and an optimal tester for pattern freeness in that case. The rest of this section deals with removable patterns, for both the 1-dimensional and multi-dimensional settings.

In the 1-dimensional case, we show that for any removable pattern there exist certain minimal hitting sets which are in fact minimal deletion sets. These are sets where none of the flips create new occurrences. Our constructive proof shows how to build such a set and allows for a linear time algorithm for finding the deletion number. The result is summarized in Theorem 10 and proved in Appendix A.

Theorem 10 ($d_P(S)$ equals $h_P(S)$; Linear time computation of $d_P(S)$). *For a binary string S of length n and a binary pattern P of length k that is removable, the deletion number $d_P(S)$ equals $h_P(S)$ and can be computed in time $O(n+k)$ and space $O(k)$.*

For the multidimensional case, we start by showing that when P is removable, the hitting number $h_P(A)$ of A approximates the deletion number up to a multiplicative constant that depends only on the dimension d . This is done in two stages, the first of which involves the analysis of a procedure that proves the existence of a large collection of P -copies with small pairwise overlaps, among the large set of at least $d_P(A)$ P -copies that exist in A . This procedure heavily relies on the fact that P is removable. The second stage shows the existence of a large hitting set of the collection with small pairwise overlaps. The result is summarized in Lemma 11 and fully proved in Appendix A.

Lemma 11 (relation between distance and hitting number). *Let P be a removable (k, d) -array over an alphabet Γ , and let A be an (n, d) -array over Γ . Let $\alpha_d = 4^d + 2^d$. It holds that: $h_P(A) \leq d_P(A) \leq \alpha_d h_P(A) \leq \alpha_d (n/k)^d$.*

7 Testers for Pattern Freeness

We describe efficient testers for both the one-dimensional and the d -dimensional removable patterns that have tolerance and query complexity that only depend on d (and not on k ; using a completely naive tester, it can be seen that the tolerance and the query complexity depend on k). The testers essentially approximate the hitting number, which is related to the deletion number by the characterizations that were shown in Section 6.

We start by presenting the distance approximation algorithm for P -freeness, which has both additive and multiplicative errors.

Theorem 12 (Approximating the deletion number in 1-dimension). *Let P be a removable string of length k and fix constants $0 < \tau < 1, 0 < \delta < 1/k$. Let $h_1, h_2 : [0, 1] \rightarrow [0, 1]$ be defined as $h_1(\epsilon) = (1 - \tau)\epsilon - \delta$ and $h_2(\epsilon) = \epsilon + \delta$. There exists an (h_1, h_2) -distance approximation algorithm for P -freeness with query complexity and running time of $O(1/k\tau\delta^2)$.*

Note that $d_P(S) = h_P(S) \leq n/k$ always holds, so having an additive error parameter of $\delta \geq 1/k$ is pointless. The proof of Theorem 12 can be adapted to derive (ϵ_1, ϵ_2) -tolerant testers for any $0 \leq \epsilon_1 < \epsilon_2 \leq 1$, which we describe in Theorem 22. An immediate corollary is the following multiplicative tester. The proofs for Theorems 12 and 22 can be found in Appendix B.

Corollary 13 (Multiplicative tolerant tester for pattern freeness in 1-dimension). *Fix $0 < \tau < 1$. For any $0 < \epsilon \leq 1$ there exists a $((1 - \tau)\epsilon, \epsilon)$ -tolerant tester whose number of queries and running time are $O(\epsilon^{-1}\tau^{-3})$.*

For the multidimensional case, our distance approximation algorithm and tolerant tester for P -freeness are given in Theorems 14 and 15. As their technical details are very similar to those in the 1D case, we provide in Appendix B only a sketch of the main ideas.

Theorem 14 (Approximating the deletion number in multidimensional arrays). *Let P be a removable (k, d) -array and fix constants $0 < \tau \leq 1, 0 \leq \delta \leq 1/k^d$. Let $h_1, h_2 : [0, 1] \rightarrow [0, 1]$ be defined as $h_1(\epsilon) = (1 - \tau)^d \alpha_d^{-1} \epsilon - \delta$ and $h_2(\epsilon) = \epsilon + \delta$. There exists an (h_1, h_2) -distance approximation*

algorithm for P -freeness making at most $\gamma/k^d\tau^d\delta^2$ queries, where $\gamma > 0$ is an absolute constant, and has running time $\zeta_\tau/k^d\delta^2$ where ζ_τ is a constant depending only on τ .

Theorem 15 (Multiplicative tolerant tester for pattern freeness in multidimensional arrays). *Fix $0 < \tau \leq 1$ and let P be a removable (k, d) -array. For any $0 < \epsilon \leq 1$ there exists a $((1 - \tau)^d \alpha_d^{-1} \epsilon, \epsilon)$ -tolerant tester making $C_\tau \epsilon^{-1}$ queries, where $C_\tau = O(1/\tau^d(1 - (1 - \tau)^d)^2)$. The running time is $C'_\tau \epsilon^{-1}$ where C'_τ depends only on τ .*

8 Discussion and Open Questions

We have provided efficient algorithms for testing whether high-dimensional arrays do not contain a fixed pattern P for any removable pattern P . The results suggest several interesting open questions on the problem of pattern-freeness and more generally, on local properties - where we say that a property \mathcal{P} is k -local if any array A not satisfying \mathcal{P} , there exists a consecutive subarray of A of size at most $k \times \dots \times k$ which does not satisfy \mathcal{P} as well. That is, a property is local if any array not satisfying \mathcal{P} contains a small ‘proof’ for this fact. Note that P -freeness is indeed k -local where k is the side length of P , and that a property \mathcal{P} is k -local if and only if there exists a *family* \mathcal{F} of arrays of size at most $k \times \dots \times k$ each, such that A satisfies \mathcal{P} if and only if it does not contain any consecutive sub-array from \mathcal{F} . That is, to understand the general problem of testing local properties of arrays we will need to understand the testing of \mathcal{F} -freeness, where \mathcal{F} is a family of forbidden patterns (rather than a single forbidden pattern).

In particular, the problem of approximate pattern matching is of interest. The family of forbidden patterns for this problem might consist of a pattern and all patterns that are close enough to it, and the distance measures between patterns might also differ from the Hamming distance (e.g., ℓ_1 distance for grey-scale patterns).

Finally, it is desirable to settle the problem of testing pattern freeness for the almost homogenous case by either finding an efficient tester for the almost homogeneous multi dimensional case, or proving that an efficient tester cannot exist for such patterns. It is also of interest to examine which of the $[k]^d$ patterns with $k < 3 \cdot 2^d$ are removable.

Acknowledgements We are grateful to Swastik Kopparty for numerous useful comments. We are thankful to Sofya Raskhodnikova for her useful feedback.

References

- [1] N. Alon (2002). Testing subgraphs in large graphs, *Random structures and algorithms*, 21(34):359–370.
- [2] N. Alon and O. Ben Eliezer (2016). *Removal lemmas for matrices*, arXiv preprint 1609.04235.
- [3] N. Alon, M. Krivelevich, I. Newman and M. Szegedy (2001). Regular languages are testable with a constant number of queries, *SIAM Journal on Computing*, 30, 1842–1862.
- [4] N. Alon, E. Fischer, M. Krivelevich and M. Szegedy (2000). Efficient testing of large graphs, *Combinatorica*, 20, 451–476.
- [5] N. Alon, E. Fischer and I. Newman (2007). Efficient testing of bipartite graphs for forbidden induced subgraphs, *SIAM Journal on Computing*, 37.3, 959–976.

- [6] N. Alon and J. Spencer (2008). *The Probabilistic Method*. Wiley.
- [7] A. Amir, G. Benson (1998). Two-Dimensional Periodicity in Rectangular Arrays, *SIAM Journal on Computing*, 27, 90-106.
- [8] A. Amir, G. Benson, M. Farach (1994). An Alphabet Independent Approach to Two-Dimensional Pattern Matching, *SIAM Journal on Computing*, 23, 313-323.
- [9] P. Awasthi, M. Jha, M. Molinaro and S. Raskhodnikova (2016). Testing Lipschitz functions on hypergrid domains. *Algorithmica*, 74(3), 1055–1081.
- [10] Y. Bar-Hillel, M. Perles, and E. Shamir (1964). On formal properties of simple phrase structure grammars, In Y. Bar-Hillel, editor, *Language and Information: Selected Essays on Their Theory and Application*, 116–150. Addison-Wesley, Reading, Massachusetts.
- [11] P. Berman, M. Murzabulatov, S. Raskhodnikova (2015). Constant-Time Testing and Learning of Image Properties, arXiv preprint 1503.01363.
- [12] P. Berman, M. Murzabulatov and Sofya Raskhodnikova (2016). Tolerant Testers of Image Properties. *ICALP*, 1–90:14.
- [13] R.S. Boyer and J.S. Moore (1977). A fast string searching algorithm, *Comm. ACM*, 20(10), 762–772.
- [14] R. Cole (1991). Tight Bounds on the Complexity of the Boyer-Moore String Matching Algorithm, *SODA*, 224–233.
- [15] C. Maxime, A. Czumaj, L. Gasieniec, S. Jarominek, T. Lecroq, W. Plandowski, and W. Rytter (1994). Speeding up two string-matching algorithms, *Algorithmica*, 12, 247–267.
- [16] E. Fischer, and I. Newman (2001). Testing of matrix properties, *STOC*, 286–295
- [17] E. Fischer and I. Newman (2007). Testing of matrix-poset properties, *Combinatorica*, 27(3), 293–327.
- [18] E. Fischer, E. Rozenberg, Lower bounds for testing forbidden induced substructures in bipartite-graph-like combinatorial objects, Proc. RANDOM 2007, 464-478.
- [19] R. Fowler, M. S. Paterson and S. L. Tanimoto (1981). Optimal packing and covering in the plane are NP-complete. *Information processing letters*, 12(3), 133–137.
- [20] Z. Galil, J. G. Park and K. Park. (2004). Three-dimensional periodicity and its application to pattern matching. *SIAM Journal on Discrete Mathematics*, 18(2), 362–381.
- [21] Z. Galil and J. I. Seiferas (1983). Time-Space-Optimal String Matching, *J. Comput. Syst. Sci*, 26(3), 280–294.
- [22] O. Goldreich, S. Goldwasser and D. Ron (1998). Property testing and its connection to learning and approximation, *JACM*, 45, 653–750.
- [23] J. Krkkinen and E. Ukkonen (2008). Multidimensional string matching. In *Encyclopedia of Algorithms*, 559–562.

- [24] I. Kleiner, D. Keren, I. Newman, O. Ben-Zwi (2011). Applying Property Testing to an Image Partitioning Problem, *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(2), 256–265.
- [25] S. Korman, D. Reichman, G. Tsur and S. Avidan. Fast-Match: Fast Affine Template Matching, *International Journal of Computer Vision*, to appear.
- [26] D. E. Knuth, J. H. Morris Jr. and V. R. Pratt (1977). Fast Pattern Matching in Strings, *SIAM J. Comput.* 6(2): 323–350.
- [27] T. Lecroq (2007). Fast exact string matching algorithms, *Information Processing Letters*, 102(6), 229–235.
- [28] G. Navarro and M. Raffinot (2000). Fast and flexible string matching by combining bit-parallelism and suffix automata, *Journal of Experimental Algorithmics (JEA)* 5: 4.
- [29] M. Parnas, D. Ron and R. Rubinfeld (2006). Tolerant property testing and distance approximation. *Journal of Computer and System Sciences*, 72(6), 1012–1042.
- [30] S. Raskhodnikova (2003). Approximate testing of visual properties, *RANDOM*, 370–381.
- [31] R. L. Rivest (1977). On the Worst-Case Behavior of String-Searching Algorithms, *SIAM J. Comput.* 6(4): 669–674.
- [32] R. Rubinfeld and M. Sudan (1996). Robust characterization of polynomials with applications to program testing, *SIAM J. Comput.* 25, 252–271.
- [33] G. Tsur and D. Ron (2014). Testing properties of sparse images, *ACM Transactions on Algorithms* 4.
- [34] A. C. Yao (1977). Probabilistic computation, towards a unified measure of complexity, *FOCS*, 222–227.

A Characterizations of the Deletion Number: Proofs

Proof of Theorem 10. The main challenge is in proving that $d_P(S) = h_P(S)$, since then all we need is an algorithm that computes $h_P(S)$, which is relatively standard in template matching: Find the set \mathcal{O} of all P -copies in S ; Go through the P -copies in \mathcal{O} from left to right, repeating the following: (i) Let P^* be the leftmost P -copy in \mathcal{O} ; (ii) Increment the hitting set count by 1; (iii) Remove from \mathcal{O} all the (following) P -copies that intersect P^* (those whose starting location is not to the right of the rightmost location in P^*). Clearly, the complexity of the algorithm is dominated by the first step of finding \mathcal{O} , which can be done in $O(n + k)$ using, e.g., the KMP algorithm [26]. Taking the rightmost location in each of the visited P^* s creates a hitting set, which is minimal, due to the fact that the set of P^* s is independent.

It is trivial that $d_P(S) \geq h_P(S)$ and hence we have to show that $d_P(S) \leq h_P(S)$. Refer to Algorithm 1 below that constructs a set of bit flip locations. Note that the choice in Step 3 is possible using the modification lemma, while the choice in Step 4 is possible, since if h is contained in only one P -copy $P^0 \in \mathcal{D}$, by definition of \mathcal{D} there is some $P^1 \in \mathcal{D}$ such that P^0 and P^1 intersect at some location x (in particular one of the 2 endpoints of P^0 must be in the intersection). Simply replace h by x . It is easy to verify that the set of locations \mathcal{F} that it computes is a (particular)

minimal hitting set of \mathcal{O} , and hence $|\mathcal{F}| = h_P(S)$. It is therefore sufficient to show that flipping the bit locations in \mathcal{F} turns the string S to be P -free. This will be guaranteed, using the fact that \mathcal{F} is a hitting set of \mathcal{O} , by Lemma 16 that shows that no bit flip of a location in \mathcal{F} creates a new P -copy. Therefore, the proof of Lemma 16 will complete the proof of Theorem 10.

Algorithm 1

Input: Binary string S of length n and removable binary string P of length k

Output: Minimal set \mathcal{F} of flip locations in S that make it P -free ($|\mathcal{F}| = d_P(S)$)

1. Find the set \mathcal{O} of all P -copies in S
2. Divide \mathcal{O} into $\mathcal{I} \cup \mathcal{D}$, where \mathcal{I} is the subset of P -copies that do not intersect any other P -copy in \mathcal{O} , while \mathcal{D} is the subset of P -copies that intersect some other P -copy in \mathcal{O} .
3. For each P -copy $P^* \in \mathcal{I}$ add to \mathcal{F} a bit location whose flipping removes P^* without creating any other P -copy
4. Find a minimal hitting set \mathcal{H} of \mathcal{D} such that every location $h \in \mathcal{H}$ is contained in at least two P -copies in \mathcal{D} .
5. Add \mathcal{H} to \mathcal{F}

return \mathcal{F}

Lemma 16 (Flipping bits in F does not create new P -copies). *Let $f \in \mathcal{F}$. Flipping the bit at location f does not create any new P -copy in S .*

Proof. Recall that \mathcal{F} consisted of bits in \mathcal{I} as well as bits in \mathcal{D} . Each of the bit flips that are in \mathcal{I} was chosen (step 3 of Algorithm 1) using the modification lemma to be such that no new P -copy is created.

The main challenge is in showing that the remaining bit flips, i.e. at locations \mathcal{H} , do not create any new P -copies. Notice our requirement that any location $h \in \mathcal{H}$ is contained in at least two P -copies. By symmetry considerations, we have the following

Observation 17. *[Flipping an arbitrary bit in the intersection of 2 P -copies can create a new P -copy] \iff [Flipping an arbitrary bit in a P -copy can create 2 new P -copies]*

By Observation 17, in order to show that bit flips in \mathcal{H} do not create new P -copies, one can prove that an arbitrary bit-flip in a P -copy cannot create more than 1 P -copy, as is stated in the next lemma. □

Lemma 18 (Any bit flip in a pattern P cannot create more than 1 new P -copy). *Let $x \in [k]$. Flipping the bit P_x can create at most 1 new P -copy in S .*


Proof. The proof goes by contradiction, assuming that a bit flip in P has created two new P -copies P^1 and P^2 , and will analyze separately the two possible cases:

case 1: ‘ P^1 and P^2 intersect P from different sides’

In this case, flipping the bit location x of P creates a P -copy P^1 shifted t_1 locations to the left and

a P -copy P^2 shifted t_2 locations to the right, where we assume w.l.o.g. that $t_1 < t_2$. One can verify that $P_{x-t_2} = P_{x-t_2}^2 \neq P_x$ (and similarly that $P_x \neq P_{x+t_1}^1 = P_{x+t_1}$). We will assume that $P_x = 0$ and hence $P_{x-t_2} = 1$. We refer the reader to Figure 2 and its caption for the intuition of the proof.

Figure 2: **Illustration for case 1:** Our proof is based on 'skipping' along a 'path' from location x to location $x - t_2$ in P , while each skip is done between entries with equal values. A complete path from x to $x - t_2$ will give a contradiction, since $P_{x-t_2} \neq P_x$. The path starts at x and makes skips of size t_1 to the left as long as it does not pass $x - t_2$, then it makes a single skip to the right of size t_2 . It repeats this traversal until reaching $x - t_2$.



Since the P -copy P^1 was created from P at a left offset of t_1 by the flipping at location x , we can infer that $P_y = P_{y+t_1}$ for any $y \in [k - t_1]$, $y \neq x$ (or informally that " P is t_1 -cyclic except at x from the right"). Similarly, we know that " P is t_2 -cyclic except at x from the left".

We define a 'path' of skips that starts from location x , makes skips of size t_1 to the left as long as it does not pass $x - t_2$, then it makes a single skip to the right of size t_2 . Call this short path a traversal. The path repeats this traversal until reaching $x - t_2$. It is easy to verify that the path is always within the open range $(x - t_2, x + t_1)$ (except for the last step that reaches $x - t_2$). This implies in particular that the path does not go from $x + t_1$ to x or from x to $x - t_2$ (i.e. through the two only "value switching skips"), and hence the value of P along the path must be 0.

It remains to prove that the path eventually reaches $x - t_2$ and does not continue in some infinite loop. For each location y that the path goes through we can look at the value $y \pmod{t_1}$. Assume w.l.o.g. that for the 'target' location $x - t_2$ we have that $x - t_2 = 0 \pmod{t_1}$. This implies for the 'starting' location x that $x = t_2 \pmod{t_1} = 0$. Now, each skip by t_1 does not change the location $\pmod{t_1}$, while a skip by t_2 to the right increases the value by $t_2 \pmod{t_1}$. In other words, the sub-sequence of locations at the beginning of each traversal (before the first left skip) is of the form $\ell \cdot t_2 \pmod{t_1}$, for $\ell = 1, 2, 3, \dots$. This is exactly the subgroup of \mathbb{Z}_{t_1} (the additive group of integers modulo t_1) generated by the element t_2 and hence must contain the identity element $0 \pmod{t_1}$. This proves that the location $x - t_2$ will be reached.

case 2: ' P^1 and P^2 intersect P from one (the same) side'

Flipping a location x in a P -copy P creates two new P -copies P^i ($i = 1, 2$) that intersect P from the same side, w.l.o.g. right, at a shift of t_i , where $t_1 < t_2$. Refer to Figure 3 and its caption for the intuition of the proof. We 'follow' the two disjoint 'arrow paths' shown in the figure that lead from x in P to $x' := x - t_2$ in P^1 to reach a contradiction. Formally:


$$P_x = P_x^1 = P_{x-t_2+t_1}^2 = P_{x-t_2+t_1} = P_{x-t_2}^1$$

$$P_x \neq P_{x-t_2}^2 = P_{x-t_2} = P_{x-t_2}^1$$

□

□

Figure 3: **Illustration for case 2:** All arrows (ignoring directions) except the red one represent equality, while the red arrow represents inequality. The two disjoint 'arrow paths' from x in P to x' in P^1 imply that both $P_x = P_{x'}$ and $P_x \neq P_{x'}$, leading to contradiction.



Proof of Lemma 11. The first inequality follows from the fact that one needs to modify at least one entry in any P -copy in A . For the third inequality, note that the set $\{(x_1, \dots, x_d) \in [k]^d : \forall 1 \leq i \leq d, x_i \equiv k-1 \pmod{k}\}$ is a set of size $[n/k]^d$ that hits all $k \times \dots \times k$ consecutive subarrays of A , and in particular all P -copies. It remains to prove that $d_P(A) \leq \alpha_d h_P(A)$. We may assume that the alphabet Γ is binary by applying the standard reduction from non-binary to binary alphabets presented in Section 5.

We present a procedure on the array A that makes it P -free by sequentially flipping bits in it. In what follows, we will say that the *center* of a (k, d) matrix lies in location $(\lfloor k/2 \rfloor, \dots, \lfloor k/2 \rfloor)$ in the matrix. Let \mathcal{P} be the set of all P -copies *before* A is modified. In Phase 1, the procedure "destroys" all P -copies in \mathcal{P} by flipping central bits of a subset of the original P -copies in A , which is chosen in a greedy manner. However, these bit flips might create new P -copies in M , which are removed in Phase 2 using the modification lemma. The procedure maintains sets \mathcal{A}, \mathcal{B} that contain the bits flipped in phases 1,2 respectively.

- Let \mathcal{P} be the set of all P -copies in A , $\mathcal{N} \leftarrow \phi$, $\mathcal{A} \leftarrow \phi$, $\mathcal{B} \leftarrow \phi$.
- **Phase 1:** While $\mathcal{P} \neq \phi$
 - Pick $Q \in \mathcal{P}$ arbitrarily.
 - Flip A_x where x is the center of Q .
 - Add Q to \mathcal{A} and remove all P -copies containing x from \mathcal{P} .
 - Add all P -copies *created* by flipping A_x to \mathcal{N} .
- **Phase 2:** While $\mathcal{N} \neq \phi$
 - Pick $Q \in \mathcal{N}$ arbitrarily.
 - Pick a location x in Q whose flipping does not create new P -copies in A (exists by modification lemma).
 - Flip the bit A_x and add x to \mathcal{B} .

For the analysis of the procedure, we say that two P -copies Q, Q' in A whose starting points are $x = (x_1, \dots, x_d), y = (y_1, \dots, y_d) \in [n]^d$ respectively are *1/2-independent* if $|x_i - y_i| \geq k/2$ for some $1 \leq i \leq d$. Note that 1/2-independence is a symmetric relation. A set of P -copies is 1/2-independent if all pairs of copies in it are 1/2-independent. Denote by $i_P(A)$ the maximal size of a 1/2-independent set in A , divided by n^d .

For Q and Q' as above, if Q' does not contain the center of Q then Q, Q' are 1/2-independent, as there is some $1 \leq i \leq d$ for which either $y_i < x_i + \lfloor k/2 \rfloor - (k-1) \leq x_i - k/2 + 1$ or $y_i > x_i + \lfloor k/2 \rfloor \geq x_i + k/2 - 1$. In both cases $|y_i - x_i| \geq k/2$, implying the 1/2-independence. Therefore the set \mathcal{A}

generated by the procedure is 1/2-independent: if $Q, Q' \in \mathcal{A}$ are two different P -copies and Q was added to \mathcal{A} before Q' , then Q' does not contain the center of Q , so Q and Q' are 1/2-independent. Using the following claims, it is not hard finish the proof of the lemma.

Claim 19. $i_P(A) \leq 2^d h_P(A)$.

The proof of Claim 19 will be given later. For what follows, we say that P has a *cycle* of size $t = (t_1, \dots, t_d) \in \mathbb{Z}^d$, if $P_x = P_y$ for every pair of locations $x = (x_1, \dots, x_d), y = (y_1, \dots, y_d) \in [k]^d$ such that $x_i \equiv y_i \pmod{|t_i|} \forall i \in [d]$. The following claim is straightforward to verify.

Claim 20. [Shifted occurrences imply a cyclic pattern] *If M contains two overlapping occurrences of A , at a relative offset of $t \in \mathbb{Z}^d$, then P has a cycle of size t .*

Claim 21. [Central bit flip creates few new occurrences] *Flipping the central bit of a P -occurrence in A creates at most 2^d new occurrences of P in A .*

We first show how to use these claims to finish the proof. Consider the sets \mathcal{A}, \mathcal{B} after the procedure ends. The procedure flips $|\mathcal{A}| + |\mathcal{B}|$ bits in A , so $|\mathcal{A}| + |\mathcal{B}| \geq d_P(A)n^d$. On the other hand, $|\mathcal{A}| \leq i_P(A)n^d \leq (2n)^d h_P(A)$ as \mathcal{A} is 1/2-independent. Claim 21 now implies that $|\mathcal{B}| \leq 2^d |\mathcal{A}|$, and we get that

$$n^d d_P(A) \leq |\mathcal{A}| + |\mathcal{B}| \leq (2^d + 1)|\mathcal{A}| \leq \alpha_d n^d h_P(A)$$

Dividing by n^d yields the desired inequality. We now prove the claims.

Proof of Claim 19. Let \mathcal{S} be a 1/2-independent set of P -copies in A , which is of size $i_P(A)$. We will show that no point in $[n]^d$ is contained in more than 2^d copies from \mathcal{S} , implying that to hit all copies of P in A (and in particular, all copies of P in \mathcal{S}) we will need at least $|\mathcal{S}|/2^d = i_P(A)/2^d$ entries.

Suppose to the contrary that there are $2^d + 1$ copies from \mathcal{S} that contain the point $x = (x_1, \dots, x_d) \in [n]^d$. We will say that a copy from \mathcal{S} containing x is i -lower if $k/2 \leq x_i - y_i < k$ and i -higher if $0 \leq x_i - y_i < k/2$ (note that $0 \leq x_i - y_i < k$ must hold). Therefore, there exist two copies $Q, Q' \in \mathcal{S}$ containing x , starting at (y_1, \dots, y_d) and (y'_1, \dots, y'_d) respectively, such that for any $1 \leq i \leq d$, Q is i -higher (i -lower) if and only if Q' is i -higher (i -lower respectively). But then, for any i , either $0 \leq y_i, y'_i < k/2$ or $k/2 \leq y_i, y'_i < k$, implying that $|y_i - y'_i| < k/2$, thus contradicting the fact that \mathcal{S} is 1/2-independent. \square

Proof of Claim 21. Assume that more than 2^d new occurrences are created. Since these occurrences overlap (at the bit flip location), the same argument as in Claim 19 implies that there must be two of them, P_1 and P_2 , that are shifted (one from the other) by some vector $t \in \mathbb{Z}^d$, where $|t_i| < k/2 \forall i \in [d]$.

By Claim 20, P (and hence also P_1 and P_2) has a cycle of size t . Let x be the point in M of the (central) flipped bit in P_0 and consider the point $x' = x + t$, which is also in P_0 , since $|t_i| < k/2 \forall 1 \leq i \leq d$. The occurrence P_2 overlaps both locations x and x' (since both new occurrences P_1 and P_2 overlap the bit flip location x and P_2 is shifted by t from P_1 , which overlaps x).

On one hand we have $M_x = M_{x'}$ (before the bit flip), since both locations belong to P_0 , which has a cycle of size t . On the other hand, $M_x \neq M_{x'}$, since these locations both belong to P_2 and must be equal after flipping M_x as P_2 has a cycle length of t . This leads to a contradiction. \square

\square

B Testers for Pattern Freeness: Proofs

Theorem 22. *Let P be a removable string of length k and let $0 \leq \epsilon_1 < \epsilon_2 \leq 1$. There exists an (ϵ_1, ϵ_2) -tolerant tester whose number of queries and running time are $O(\epsilon_2^2/(\epsilon_2 - \epsilon_1)^3)$ where the constant term does not depend on k .*

It is not clear whether this upper bound is tight in general. However, for the important special case of tolerant testers with multiplicative tolerance of $1 + \tau$, where $\tau > 0$ is a constant, the above tester is optimal (up to a multiplicative constant that depends on τ), as is shown by taking $\epsilon_2 = \epsilon$ and $\epsilon_1 = (1 - \tau)\epsilon$ in Theorem 22, leading to the multiplicative tester given in Corollary 13.

Proof of Theorems 12 and 22. Let S be a string of length $n \geq \beta k$, where $\beta = 3/\tau$. Write $\epsilon = \delta_P(S)$ and let $H \subseteq [n]$ be a hitting set for P in S whose size is ϵn . That is, H is a minimal set of locations that satisfies the following: if S contains a copy of P starting at location l , then $\{l, \dots, l + k - 1\} \cap H \neq \emptyset$.

For $i \in [n]$ let I_i denote the ‘‘cyclic interval’’ of length βk starting at i . That is, if $i + \beta k > n$ then $I_i = \{i, \dots, n\} \cup \{0, \dots, i + \beta k - n - 1\}$ and otherwise $I_i = \{i, \dots, i + \beta k - 1\}$.

Let the random variable X denote the size of the minimal hitting set H_i for P in the interval I_i , divided by βk , where $i \in [n]$ is chosen uniformly at random. Note that X is computable in time $O(\beta k)$, by Theorem 10. Let μ and σ^2 denote the expectation and the variance of X , respectively. By the minimality of H_i , we have that $|H_i| \leq |H \cap I_i|$ since the set in the RHS is a hitting set for P with respect to the interval I_i . Thus, $\mu \leq \mathbb{E}[|H \cap I_i|]/\beta k = \epsilon$.

Next we bound μ from below. Since H_i hits all P -copies that lie exclusively inside I_i , and by the minimality of H , we must have $|H_i| > |H \cap I'_i|$ where I'_i is the cyclic interval that starts in $i + k$ and ends in $(i + (\beta - 1)k - 1) \bmod n$. Therefore, $\mu \geq \mathbb{E}[|H \cap I'_i|]/\beta k = (1 - 2/\beta)\epsilon \geq (1 - \tau)\epsilon$. To conclude, we have seen that $(1 - \tau)\epsilon \leq \mu \leq \epsilon$.

To compute the variance of X , note that $0 \leq X \leq 1/k$, as there exist β entries in I_i such that any subinterval of length k in I_i contains at least one of them. By convexity, the variance satisfies $\sigma^2 \leq k\mu(1/k - \mu)^2 + (1 - k\mu)(0 - \mu)^2 = \mu(1/k - \mu) \leq \epsilon/k$.

Now let $Y = \frac{1}{t} \sum_{j=1}^t X_j$ where the X_j are independent copies of X and t will be determined later. Then $\mathbb{E}[Y] = \mu$ and $\text{Var}(Y) = \sigma^2/t \leq \epsilon/kt$.

Recall that $\mathbb{E}[Y] = \mu$, where $(1 - \tau)\epsilon \leq \mu \leq \epsilon$, so to get the desired approximation, it suffices to estimate Y with an additive error of no more than δ with constant probability. Chebyshev inequality implies that it suffices to have $\text{Var}(Y) = \Theta(\delta^2)$. In other words, it will be enough to sample $t = \Theta(\epsilon/k\delta^2)$ blocks, each of size $\beta k = \Theta(k/\tau)$. In total, it is enough to make $\Theta(k\epsilon(1/k - \epsilon)/\tau\delta^2) = O(\epsilon/\tau\delta^2)$ queries.

In the setting of approximation, ϵ is not known in advance, but $\epsilon \leq 1/k$ always holds, so sampling $t = \Theta(1/k^2\delta^2)$ blocks would suffice to get the desired additive error. The return value of the approximation algorithm will be its estimate of Y . The query complexity and running time are $\beta tk = \Theta(1/k\tau\delta^2)$. This finishes the proof of Theorem 12.

Now consider the setting of (ϵ_1, ϵ_2) -tolerant testing. By monotonicity of the tester, we can assume that we are given a string whose distance from P -freeness is either exactly ϵ_2 or exactly ϵ_1 . Pick $\epsilon = \epsilon_2$, $\delta = (\epsilon_2 - \epsilon_1)/4$, $\tau = (\epsilon_2 - \epsilon_1)/4\epsilon_2$, and sample $t = \Theta(\epsilon/k\delta^2)$ blocks, with query complexity and running time $\Theta(\epsilon/\tau\delta^2) = \Theta(\epsilon_2^2/(\epsilon_2 - \epsilon_1)^3)$, as was stated above. If the given string S is ϵ_2 -far from P -freeness, then with probability at least $2/3$, after sampling $t = \Theta(\epsilon/\tau\delta^2)$ samples, the value of Y will be bigger than $(\epsilon_2) * (1 - \tau) - \delta = (\epsilon_2 + \epsilon_1)/2$. On the other hand, if S is ϵ -close

then with probability at least $2/3$, $Y \leq \epsilon_1 + \delta < (\epsilon_2 + \epsilon_1)/2$. Therefore, the tester will answer that the input is ϵ_2 -far if and only if $Y \geq (\epsilon_2 + \epsilon_1)/2$. This finishes the proof of Theorem 22. \square

Proof sketch for Theorems 14 and 15. Take $\beta = 2/\tau$. Let A be an (n, d) -array where we may assume that $n \geq \beta k$ for a suitable choice of C . Again, the strategy is to take t (to be determined) independent samples of blocks of size $\beta k \times \dots \times \beta k$ and compute the hitting number of each sampled block. Note that (as opposed to the one-dimensional case), computing the minimal hitting set is generally an *NP*-complete problem, but since the hitting number of each of these blocks is at most $\beta^d = \Theta(\tau^{-d})$, here we may compute it with running time that depends only on τ and d . As in the 1D case, the expected relative hitting number μ of a sampled block satisfies $(1 - \tau)^d h_P(A) = (1 - 2/\beta)^d h_P(A) \leq \mu \leq h_P(A)$. The variance of the hitting number for a single sample is no bigger than $k^d(1/k^d - \mu)^2 + (1 - k^d\mu)\mu^2 = \mu(1/k^d - \mu) \leq \mu/k^d$, so for t samples it is $O(h_P(A)/k^d t)$. To get additive error of at most δ with constant probability, we may have (by Chebyshev inequality) $h_P(A)/k^d t = \Theta(\delta^2)$, or $t = \Theta(h_P(A)/k^d \delta^2)$.

Therefore, for an approximation algorithm (in which we don't know $h_P(A)$ in advance, though we have an upper bound of $h_P(A) \leq 1/k^d$), $t = \Theta(k^{-2d} \delta^{-2})$ sampled blocks are enough, and the total number of samples is $O(1/k^d \tau^d \delta^2)$. For a $((1 - \tau)^d \epsilon, \epsilon)$ -tolerant tester for the hitting number (which translates to a $((1 - \tau)^d \alpha_d^{-1} \epsilon, \epsilon)$ -tolerant tester for the deletion number), as observed in the 1D case, when deciding on the number of samples we may assume that $h_P(A) = \epsilon$ and pick $\delta = \Theta((1 - (1 - \tau)^d) \epsilon)$, so $t = \Theta(\epsilon/k^d \delta^2) = \Theta(1/k^d (1 - (1 - \tau)^d)^2 \epsilon)$ sampled block suffice. Since each block is of size $\Theta(k^d/\tau^d)$, the total number of queries is $O(C_\tau \epsilon^{-1})$ where $C_\tau = 1/\tau^d (1 - (1 - \tau)^d)^2$, while the running time is $C'_\tau \epsilon^{-1}$, where C'_τ depends on the time required to compute the hitting number in a single sampled block. \square

C Almost Homogeneous Patterns

The testers discussed above only consider removable patterns. However, Theorem 1 states that almost homogeneous patterns over a binary alphabet are not removable. As the above testers are not applicable for this case, it is natural to ask whether there exist efficient testers (and in particular tolerant testers) for pattern freeness when the pattern is almost homogeneous.

In this subsection we partially answer this question, addressing the one-dimensional case; we do so by describing a simple yet powerful characterization of the distance to P -freeness when P is almost homogeneous. The characterization is then utilized to get both an optimal $1/(16+c)$ -tolerant tester for P -freeness in the one-dimensional case and an efficient *exact* algorithm for computing the distance to P -freeness.

Unfortunately, this characterization does not hold in higher dimensions, and the question of testing and approximating the distance to P -freeness in this case is left open. It is not clear whether there exist efficient algorithms for computing the exact distance to P -freeness in high dimensions. It will be interesting to either find such an algorithm or show that the problem is *NP*-hard.

Denote the absolute distance of a string S to P -freeness by $d_P(S)$. Our main results here are as follows.

Theorem 23. *There exists an algorithm that, given an almost homogeneous binary string P and a binary string S of length n , compute $d_P(S)$ in time $O(n)$ and space $O(1)$ where the constants do*

not depend on the length of P .

The above algorithm is clearly optimal (up to a multiplicative constant that does not depend on k) with respect to time and space complexity. Our second result is a $1/(16+c)$ -tolerant tester for P -freeness when P is almost homogeneous.

Theorem 24. *Let P be an almost homogeneous binary string and fix a constant $c > 0$. Then there exists a $1/(16+c)$ -tolerant ϵ -tester for P -freeness making at most $\alpha_c \epsilon^{-1}$ queries, where α_c depends only on c .*

Again, the running time of the above tester is optimal (up to a multiplicative constant). In particular, it does not depend on the length of the forbidden pattern P or the tested string. We did not try to optimize the tolerance factor in the above statement.

In the rest of this subsection we will prove the above two theorems. We will start by describing a characterization of the distance to P -freeness which is easier to work with. Without loss of generality, in the rest of the subsection assume that $P = 10^{k-1}$ for some $k > 0$ and that S is a string of length n . A pair (i, j) with $0 \leq i < j \leq n - k + 1$ is said to be a P -evidence in S if $P_i = 1$ and $P_l = 0$ for any $j \leq l \leq j + k - 2$.

Observation 25. S is P -free if and only if there are no P -evidences in S .

Proof. If S contains a P -copy at location i then the pair $(i, i + 1)$ is a P -evidence. On the other hand, suppose that (i, j) is a P -evidence and let $i' < j$ be the maximal integer such that $P_{i'} = 1$. Then $P_l = 0$ for any $i' < l \leq j + k - 2$, implying that S contains a P -copy starting at i' . \square

The main technical result of this subsection relates the distance to P -freeness to the maximal number of non-overlapping P -evidences, where two P -evidences (i, j) and (i', j') are non-overlapping if $i \neq i'$ and $|j - j'| \geq k - 1$. Denote the maximal number of non-overlapping P -evidences in S by $d_E(S)$.

Lemma 26. $d_P(S) = d_E(S)$.

Proof of Lemma 26. It is clear that $d_P(S) \geq d_E(S)$. Indeed, take a set of non-overlapping P -evidences which is of size $d_E(S)$, then by Observation 25 we will need to modify at least one bit corresponding to any evidence, and in total at least $d_E(S)$ bits in S , to make it P -free. Next we show that $d_P(S) \leq d_E(S)$ by showing that there exists a set T of non-overlapping P -evidences in S with the following property: to make S P -free, it is enough to modify exactly one bit in any evidence in T .

The construction of T is carried by “marking” entries corresponding to P -evidences in S in the following manner. Initially, T is empty and no entry of S is marked. We repeat the following procedure as long as S contains a P -evidence whose entries are unmarked: in each round, we pick the smallest $0 < j \leq n - k + 1$ for which both S_l is unmarked for any $j \leq l \leq j + k - 2$ and the set $A_j = \{i < j : S_i \text{ is unmarked and equal to } 1\}$ is not empty. Pick $i = \max A_j$. Now add (i, j) to T and mark the entries of S in the k locations $i, j, j + 1, \dots, j + k - 2$. Consider the set $T = \{(i_1, j_1), \dots, (i_t, j_t)\}$ in the end of the process, where (i_s, j_s) was chosen in round s of the process, and let m be the largest integer such that S_m equals 1 and is not marked (if no such entry exists, we take $m = n$). Then T satisfies the following:

- $j_1 < \dots < j_t$. This is true by the minimality of j picked in any round. In particular, if S_i is marked before S_j and both entries are zero then $i < j$.

- Let $0 < s \leq t$. If $i_s < i_{s-1}$ then we must have $j_s = j_{s-1} + k - 1$: by the maximality of i_s , it must be true that $P_l = 0$ for any $j_{s-1} + k - 1 \leq l < j_s$. Thus, by the minimality of j_s , we must have that $j_s = j_{s-1} + k - 1$, as desired.
- On the other hand, if $i_s > i_{s-1}$ then $j_s = i_s + 1$. Again, this is true by the minimality of j_s and the maximality of i_s with respect to j_s .
- For any $1 \leq s \leq t$, either $j_s < m$ or $i_s > m$. This is true by the maximality of i_s with respect to j_s .

We now show how one can modify exactly t bits in S to make it P -free. For any $(i, j) \in T$ such that $j < m$, we modify P_i to zero. After the modification, $P_l = 0$ for any $l < m$, so there cannot be P -copies in S starting at $l < m$.

On the other hand, for any $(i, j) \in T$ such that $i > m$, we modify P_{j+k-2} to one. It remains to show that after the modification, S will not contain a P -copy starting at some $l \geq m$. By the third observation and the minimality of j in any round of the process, we cannot have a P -copy starting in an unmarked entry $S_l = 1$ - since otherwise, at some point in the above process, we would have to add the evidence $(l, l + 1)$ to T , while marking S_l , a contradiction. To complete the proof, suppose that S_l is marked and equals one after the modification; that is, $l = j_s + k - 2$ for some $(i_s, j_s) \in T$ with $i_s > m$. There are two cases to consider.

- If $i_{s+1} < i_s$ then, by the second observation above, $j_{s+1} = j_s + k - 1$. Therefore, the entry in location $j_{s+1} + k - 2 = l + k - 1$ was flipped to one during the modification, so it is not possible to have a P -copy starting in location l after the modification.
- On the other hand, if $i_{s+1} > i_s$ then it is not possible to have $P_{l+1} = P_{l+2} = \dots = P_{l+k-1} = 0$ before the modification, since otherwise we would choose the pair $(i'_{s+1}, l+1)$ satisfying $i'_{s+1} \geq m$ in round $s+1$ of the process instead of choosing (i_{s+1}, j_{s+1}) , as we have $l+1 < i_{s+1} < j_{s+1}$ and by the first observation above.

We conclude that after the modification, S is P -free. Hence $d_P(S) \leq |T|$. Since obviously $|T| \leq d_E(S)$, we get that $d_P(S) = d_E(S) = |T|$, as desired. \square

Proof of Theorem 23. Consider the following algorithm to compute the distance of a string S of length n to P -freeness, where $P = 10^{k-1}$. The idea of the proof is to read S from left to right, marking P -evidences along the way and keeping track of the following quantities.

- The number of 1's that were already observed but not yet marked, denoted by a .
- The length of the longest streak of unmarked 0's lying immediately to our left, denoted by b .
- The number of non-overlapping P -evidences found until now, denoted by c .

We initialize $a, b, c = 0$. For i running from 0 to $n - 1$ we do as follows:

- If $S_i = 1$ then a is increased by one.
- If $S_i = 0$ then
 - If $a = 0$ then we do nothing.

- if $a > 0$ then b is increased by one. If we now have $b = k - 1$ then a new P -evidence was found - and we update $a \leftarrow a - 1, b \leftarrow 0, c \leftarrow c + 1$.

The algorithm returns the value of c after the above loop stops. It is not hard to see that the algorithm indeed calculates $d_E(S) = d_P(S)$. Each round of the loop takes $O(1)$ times, so the total running time is $O(n)$ (with no dependence on k). The space complexity is clearly $O(1)$, as we only need to keep track of a, b, c . \square

Proof of Theorem 24. Let S be a string of length n and suppose that $P = 10^{k-1}$. Given $\epsilon > \delta > 0$ where $\delta \leq \epsilon/(16 + c)$, we will use the following simple tester to distinguish between the case that S is ϵ -far and the case that S is δ -close to P -freeness. Since any string S must be $1/k$ -close to P -freeness, we may assume that $\epsilon \leq 1/k \leq 1/2$. The tester makes at most m queries, where m will be determined later. First it chooses uniformly at random and independently (allowing repetitions) $i_1, \dots, i_{m/3}, j_1, \dots, j_{m/3k}$. Then, the tester samples single entries S_{i_l} for any $1 \leq l \leq m/3$ and blocks $S_{j_r}, \dots, S_{j_r+2k-2}$ for any $1 \leq r \leq m/3k$. The tester accepts the input if and only if there are no $i_l < j_r$ such that $S_{i_l} = 1$ and $S_j = \dots = S_{j+k-1} = 0$ for some $j_r \leq j \leq j_r + k - 1$.

Let $T = \{(x_1, y_1), \dots, (x_t, y_t)\}$ be a maximal set of non-overlapping P -evidences in S . Suppose first that S is ϵ -far from P -freeness; then $t \geq \epsilon n$. The probability that none of the locations $x_1, \dots, x_{t/2}$ was sampled by the tester as a single entry is at most $(1 - \epsilon/2)^{m/3} \leq e^{-\epsilon m/6}$. On the other hand, the probability that none of the 0^{k-1} -copies starting at locations $y_{t/2+1}, \dots, y_t$ is fully contained in one of the sampled blocks is at most $(1 - \epsilon k/2)^{m/3k} \leq e^{-\epsilon m/6}$. In total, the acceptance probability in this case, i.e. the probability that no P -evidence was found is at most $2e^{-\epsilon m/6}$.

On the other hand, if S is δ -close to P -freeness then $t \leq \delta n$. Since T is a maximal set of non-overlapping P -evidences, to find a P -evidence, the tester must either sample at least one of x_1, \dots, x_t as a single entry or one of the entries of the form $y_i + \delta$, with $1 \leq i \leq t$ and $0 \leq \delta k - 1$ as part of one of the sampled blocks. The probability that the first event does not occur is at least $(1 - \delta)^{m/3} > e^{-2\delta m/3}$. The probability that the second event does not occur is at least $(1 - 3k\delta)^{m/3k} > e^{-2\delta m}$. Hence the acceptance probability here, which is the probability that both events do not occur is at least $e^{-8\delta m/3} \geq e^{-\epsilon m/(6+c/8)}$. Clearly, if we pick $m = c'/\epsilon$ for a large enough constant c' that depends only on c , we will get that the quantity $e^{-\epsilon m/(6+c/8)} - 2e^{-\epsilon m/6}$ is positive and bounded away from zero; Therefore, with this choice of m the tester accepts δ -close inputs with probability that is bigger by an absolute constant than the acceptance probability for any ϵ -far input. This finishes the proof of the theorem, since one can build a tester that has the same query complexity (up to a multiplicative constant), in which δ -close inputs are accepted with probability at least $2/3$ while ϵ -far inputs are rejected with probability at least $2/3$, as desired. \square

D A Lower Bound for Testing Pattern Freeness

Here we show that any tester that makes $o_\epsilon(\frac{1}{\epsilon})$ queries in an attempt to test if a fixed pattern appears in a string, must err with probability greater than $1/3$. While the lower bound of [3] already entails a lower bound of $\Omega(1/\epsilon)$ for testing pattern-freeness (as for any fixed pattern J the language consisting of all strings not containing J is regular), we give here a self-contained proof for two reasons. First, our proof extends to the case where k , the length of J , is allowed to depend on n . Second, it is not hard to adapt the proof to the 2D (or higher dimensions) and demonstrate a lower bound of $\Omega(1/\epsilon)$ on the query complexity of testing pattern freeness of multidimensional arrays.

Theorem 27. [Lower bound of $\Omega(1/\epsilon)$ queries] Suppose k is even and consider the pattern $J := 0^{k/2-1}10^{k/2}$. Any tester that distinguishes with probability at least $2/3$ between the case that a string I is J -free and the case that I is ϵ -far from being J -free, makes at least $\frac{1}{13\epsilon}$ queries to I .

Note: we show that the lower bound above applies for wide ranges of possible values for ϵ : ranging from $\epsilon = O_n(1/n)$ to $\epsilon = \Omega_k(1/k)$. We did not attempt to optimize the constant $\frac{1}{13}$ in the theorem above. Furthermore, to ease readability we avoid using floor/ceiling signs.

Proof. By Yao's Principle [34], it suffices to construct two distributions B and C over length n strings where all strings in B are J -free and all strings in C are ϵ -far from being J -free, with the property that any deterministic algorithm D making less than $\frac{1}{13\epsilon}$ queries to I sampled from $A = \frac{1}{2}B + \frac{1}{2}C$ (namely we sample from B with probability $1/2$ and from C otherwise), must err with probability greater than $1/3$. We now explain how to construct these distributions.

The distribution B is just the single string 0^n sampled with probability 1. For the distribution C we take the string $I = 0^n$ and divide it to n/k disjoint intervals of length k : $I_1, \dots, I_{n/k}$. Then, we sample randomly a subset of $2\epsilon n$ of these intervals. For each length k interval in the subset, we choose one of the last (right side) $k/2$ locations and flip it to be a 1.

Clearly any string from B is J -free with probability 1, while every string in C contains exactly $2\epsilon n$ occurrences of J . Furthermore, as changing a single location in a string I sampled from C can remove at most 2 occurrences of J , such a string I is ϵ -far from being J -free.

Consider a deterministic tester D that performs at most $\frac{1}{13\epsilon}$ queries to the string I . Since D is deterministic, in the case that it encounters only zeros during its queries, it has to either declare the string I as being J -free or not. The first option is that D rejects (declares ' ϵ -far from J -free') when it sees only zeros. In this case D errs with probability 1 if I was chosen from B , hence it errs with overall probability of at least $1/2$. We next handle the second case, where D accepts (declares 'free') when it sees only zeros. We show that in this case D errs with probability greater than $1/3$, by focusing on the distribution C (i.e. the ϵ -far case).

Being deterministic, when D sees only zeros it queries a *fixed* set X of locations, where we chose $|X| = \frac{1}{13\epsilon}$. All we need to show is that under the random process of the distribution C , the fixed set of locations X does not contain a 1 w.p. $> 2/3$. This will suffice, since all strings from C (which is chosen w.p. $1/2$) are ϵ -far from being J -free, and therefore D will fail w.p. $> 1/2 \cdot 2/3 = 1/3$.

Let $X_i = X \cap I_i$ be the set of locations from X that are in interval I_i . An interval I_i has $k/2$ locations where a 1 could be placed (the right half of the interval) and we can assume that the set X_i is contained in the right half of I_i (since the left half is zero w.p. 1 hence would be pointless to query by any algorithm). Clearly, $Pr[X_i \text{ contains a 1}] = Pr[I_i \text{ contains a 1}] \cdot Pr[\text{the 1 is within } X_i] = \frac{2\epsilon n}{n/k} \cdot \frac{|X_i|}{k/2} = 4\epsilon|X_i|$, leading to $Pr[X \text{ contains a 1}] \leq \sum_{i=1}^{n/k} 4\epsilon|X_i| = 4\epsilon|X| < 1/3$, by union bound. \square

As noted, Theorem 27 can be generalized to the higher-dimensional settings (details omitted).