Contributions to the study of probabilistic communication complexity classes

Bottesch, Ralph Christian

2016

Bottesch, R. C. (2016). Contributions to the study of probabilistic communication complexity classes. Doctoral thesis, Nanyang Technological University, Singapore.

<https://hdl.handle.net/10356/66033>

<https://doi.org/10.32657/10356/66033>

NANYANG
TECHNOLOGICAL
UNIVERSITY

**Contributions to the Study of Probabilistic
Communication Complexity Classes**

Ralph Christian Bottesch
School of Physical and Mathematical Sciences
2015

Contributions to the Study of Probabilistic Communication Complexity Classes

Ralph Christian Bottesch

School of Physical and Mathematical Sciences

A thesis submitted to the Nanyang Technological University
in partial fulfilment of the requirement for the degree of
Doctor of Philosophy

2015

Acknowledgements

I thank first and foremost my academic supervisor, Assistant Professor Dr. Hartmut Klauck, for his guidance throughout my doctoral studies in general, and in particular for introducing me to difficult and very interesting problems in Communication Complexity, and for encouraging me to focus my research efforts on some of them.

I am very grateful for the financial support I have received for my studies in Singapore: The first four years of my doctoral studies were funded by A*STAR through the SINGA (Singapore International Graduate Award) scholarship, and the final year was funded through a research grant of Associate Professor Dr. Troy Lee, for which I want to express my deepest gratitude.

Also, I want to thank Associate Professor Dr. Wu Guohua for helping me get in touch and collaborate with a number of other complexity theorists. And finally, I thank my father, Martin Bottesch, for taking the time to carefully read through different drafts of this thesis, and making helpful suggestions for improvements.

Contents

1	Introduction	1
2	Preliminaries	7
2.1	Communication complexity	9
2.2	Arrangements of hyperplanes	20
2.3	Information theory	27
3	Complexity Bounds for Disjointness	29
3.1	Overview	29
3.2	The distributional complexity of Disjointness under distributions with bounded information	30
3.3	A non-deterministic SMP protocol for the Disjointness problem	50
4	The UPP^{cc} vs. $PSPACE^{cc}$ Question	54
4.1	Overview	54
4.2	A new approach to the UPP^{cc} vs. $PSPACE^{cc}$ question	57
4.3	Signature-invariant operations	74
4.4	The geometry of extension arrangements	89
4.5	An algorithm for balancing subspaces	106
	Conclusion	113

Abstract

The focus of this work is on two problems in Communication Complexity Theory, both related to notions of communication complexity involving randomisation.

First, we investigate the effect of the amount of correlation between the marginals of a bipartite distribution on the distributional complexity of a problem under that distribution. In this context we prove tight bounds on the *distributional complexity under distributions with bounded mutual information* in the case of the Disjointness problem.

Second, we present new results regarding a three-decades-old open problem of Babai *et al.* [3], who conjectured that the complexity classes UPP^{cc} and $PSPACE^{cc}$ are incomparable. We prove that if a certain conjecture about hyperplane arrangements holds, then $UPP^{cc} \subset PSPACE^{cc}$. To support our geometric conjecture, and with the aim of proving it, we develop a theory of operations which are invariant with respect to the combinatorial structure of hyperplane arrangements, and use it to devise an algorithm which provides numerical evidence supporting our conjecture.

Chapter 1

Introduction

Computational Complexity Theory classifies computational problems in terms of the amount of some resource (such as time or memory) required in order to solve them. The typical question in this field is of the form “What is the minimum amount of the given resource that is sufficient for completing this computational task?”, and an answer to such a question constitutes the *complexity* of the problem in terms of the resource that is being considered.

This work focuses on questions in *Communication Complexity Theory*, the sub-field of Complexity Theory in which the resource being considered is the amount of *communication* required to perform a computation. This framework was formalised in 1979 by Yao [40] as follows: Two processors (often personified as human players in the literature) are to solve an instance of a problem, but each receives as input only part of the description of the instance. They can communicate by exchanging messages according to a set of pre-established rules, called a *protocol*. The cost of the computation performed by the two processors together, is measured as the number of bits that are communicated throughout the execution of the protocol, discarding any other costs of computations performed by either of one the processors individually. The communication complexity of the problem is then taken as the smallest num-

ber of bits communicated by any protocol which produces a correct solution. Thus the communication complexity of a problem is the amount of communication that needs to occur within a computational system when the input is stored piece-wise in different parts of the system.

As an example, consider the problem of determining whether two subsets of $\{1, \dots, n\}$, with $n \in \mathbb{N}$ fixed, are disjoint (this is the *Disjointness problem*). The two players (processors) each receive one of the subsets, encoded as an n -bit string (with the i -th bit being 1 if and only if i is in the set). Evidently there is a trivial protocol, where one player sends the other his entire input (in this case his n -bit string), and thus the communication complexity of a problem is never much larger than the size of the smallest part of the input. However, in many cases the problem can be solved much more efficiently. For example, if, instead of determining whether the sets are disjoint, the task were to determine the maximum element of their union, it would suffice for one player to send the maximum of his own set, using only $\lceil \log n \rceil$ bits, which would enable the other player to compute and output the maximum of the union.

In practical situations it is rare for communication to be the only resource that needs to be considered. Nevertheless, results from Communication Complexity have been applied in other areas, such as VLSI circuit design (see [40] or [24]), the design of data structures (see [11]), and Turing Complexity Theory (see [1]).

In Communication Complexity there has been considerably more success in proving lower bounds on the complexity of problems than in the more classical Turing Complexity Theory. This is at least in part due to the fact that to any problem in the communication framework, and to any communication protocol, there corresponds in a natural way a matrix, and that most questions in this field can therefore be phrased as questions about matrices. Since matrices are some of the most ubiquitous and

well studied objects in mathematics, it is often possible to use a variety of tools from different fields (such as Algebra and Combinatorics, among many others) to attack questions in Communication Complexity.

A number of variations of the basic communication model have been introduced and studied. These are usually defined by enhancing the protocols (for example by allowing the use of randomness, or an outside source of information), or by relaxing the conditions for the protocol to be successful (for example by allowing for some wrong outputs). Consequently, the complexity of a communication problem with respect to such enhanced protocols is either less than or the same as its complexity under standard protocols. However, comparing the power of different types of protocols is often difficult.

In this work we study two topics, both related to *probabilistic communication models*, in which the processors are given access to a source of randomness on which their messages can depend (in other words, they are able to make guesses), and furthermore they are allowed to give the wrong answer with some probability on every input.

The first topic is the complexity of the Disjointness problem in different communication models involving randomisation. We investigate a complexity measure which is related to the *distributional communication model*, in which the input is chosen according to a fixed probability distribution. In this model, protocols may be tailored to the distribution, and are allowed to output the wrong answer on some small portion of the inputs. It is known that distributions which choose the inputs in such a way that the two parts of the input (which are fed separately to the processors) are highly correlated, can force the processors to communicate more in order to achieve small error. We study how the necessary amount of communication increases as a function of the amount of correlation permitted between the parts of the input, specifically in

the case of Disjointness. We also give an upper bound on the complexity of Disjointness in a model of communication where the messages of the two processors are only sent to a third processor (a *referee*) who has to decide the output using only these messages, without knowledge of the input.

Our second topic is a question regarding the relation between two communication complexity classes, both defined in terms of powerful protocols: UPP^{cc} and $PSPACE^{cc}$. The former is the class of problems which can be efficiently solved by probabilistic protocols with error arbitrarily close to $\frac{1}{2}$ on every input, in other words, protocols whose probability of success is almost as low as that of a fair coin flip. The latter class is the communication-equivalent of the Turing complexity class $PSPACE$ (the class of problems which can be solved by algorithms using memory (“space”) polynomially bounded in the size of the instance), but since there is no equivalent notion of space in the communication framework, $PSPACE^{cc}$ is defined as the class of problems which can be efficiently solved by *alternating protocols* (matching a similar characterisation of $PSPACE$ in terms of *alternating Turing machines*).

Babai, Frankl, and Simon originally defined UPP^{cc} and $PSPACE^{cc}$ in [3], and conjectured that they are incomparable, meaning that neither class is included in the other. The question of whether this conjecture is true has been open for nearly three decades. However, in 2001, 15 years after the question was posed, Forster [12] made a major breakthrough by proving that $PSPACE^{cc} \not\subset UPP^{cc}$, confirming “half” of the conjecture.

We present a new approach to this question, whereby we attempt to disprove Babai *et al.*’s conjecture by showing that $UPP^{cc} \subset PSPACE^{cc}$. Specifically, we prove that this inclusion holds, depending on a geometric conjecture which states that, given any arrangement of linear hyperplanes in a Euclidean space, the regions formed by the hyperplanes can be made sufficiently large without changing the combinatorial struc-

ture of the arrangement (its *face poset* or *signature set*). Towards this goal we study operations on hyperplane arrangements which allow for controlled changes of the sizes of the regions of an arrangement while leaving its structure unchanged. Since we have not yet obtained the full proof of the above geometric conjecture, we strengthen our case for it with numerical evidence, by designing an algorithm based on the above-mentioned operations, and analysing its performance on randomly chosen hyperplane arrangements.

Structure of the thesis.

Our results are presented in two chapters (Chapter 3 and 4), which can be read independently.

In Chapter 2 we gather the essential background knowledge necessary throughout the next chapters. Section 2.1 lists the relevant definitions and facts from Communication Complexity. Section 2.2 lists the relevant definitions and facts related to the topic of hyperplane arrangements. Section 2.3 lists the relevant definitions and facts related to Information Theory.

In Chapter 3 we present new bounds on the complexity of the Disjointness problem in different models of communication. Section 3.1 gives an overview of the chapter. In Section 3.2 we introduce the distributional complexity under distributions with bounded information, and give tight bounds on this measure for the Disjointness function. In Section 3.3 we review the SMP communication model and give an upper bound on the non-deterministic SMP complexity of Disjointness.

In Chapter 4 we present our approach the “ UPP^{cc} vs. $PSPACE^{cc}$ ” question. Section 4.1 gives an overview of the chapter. In Section 4.2 we review some results related to the class UPP^{cc} , and translate the main question into a geometric conjecture. Section 4.3 and Section 4.4 contain the bulk of our results regarding this conjecture. In

Section 4.5 we present an algorithm for increasing the smallest size of any region of a given hyperplane arrangement while leaving its structure unchanged.

The content of Chapter 3 is based on joint research with Dr. Dmitry Gavinsky and Dr. Hartmut Klauck. The results in Section 3.2 also appear in [6] (R.C. Bottesch, D. Gavinsky, and H. Klauck. Correlation in Hard Distributions in Communication Complexity. To be presented at *RANDOM 2015*.). The result in Section 3.3 also appears in [7] (R.C. Bottesch, D. Gavinsky, and H. Klauck. Equality, Revisited. To be presented at *MFCS 2015*.). The content of Chapter 4 is based on joint research with Dr. Hartmut Klauck.

Chapter 2

Preliminaries

For reference we list here some notations to be used in this and the coming chapters.

We will use the common notations $O()$, $o()$, $\Omega()$, $\omega()$ and $\Theta()$ to describe the constant-independent asymptotic behaviour of a function.

Let $k, l \in \mathbb{N} = \{1, 2, 3, \dots\}$. Let S be a set.

- $[k]$ stands for $\{1, \dots, k\}$.
- $\mathcal{P}(S)$ is the power set of S , the set of subsets of S .
- $\text{poly}(k)$ is the class of polynomial functions with variable k ; $\text{polylog}(k)$ is the class of polynomial functions with variable $\log k$. As is common, we will abuse notation by saying “ $\text{polylog}(k)$ ” when we actually mean “some fixed element of $\text{polylog}(k)$ ”.
- $\mathcal{M}_{k,l}(S)$ denotes the set of matrices with k rows, l columns and entries from S .
 $\mathcal{M}_k(S) = \mathcal{M}_{k,k}(S)$.

Let $v, w \in \mathbb{R}^k$ and $A \in \mathcal{M}_{k,l}(\mathbb{R})$ be given.

- v_i denotes the i -th entry of v , unless vectors v_j (for $j \in J \subseteq \mathbb{N}$) are defined, in which case $v_{j,i}$ refers to the i -th entry of the vector v_j .
- Vectors can be represented either as horizontal or as vertical matrices. By default we consider them as vertical ones (column matrices), although we may abuse notation, for example by referring to a row vector of A as a vector in $\mathbb{R}^l$. However, we adhere to this convention strictly when multiplying matrices.
- $\langle v, w \rangle = \sum_{i \in [k]} v_i w_i$ (which is well-defined regardless of how v and w are represented as matrices). $\|v\| = \sqrt{\langle v, v \rangle}$ (by $\|\cdot\|$ we mean the Euclidean norm).
- A_i is the i -th row vector of A . $A_{\cdot,i}$ is the i -th column vector of A . $A[I, J]$ denotes the submatrix of A with row set I and column set J , where $I \subseteq [k]$ and $J \subseteq [l]$. $A_{i,j}$ denotes the element on row i and column j of A , but note that the rows and columns of matrices may occasionally be indexed by elements from sets other than $[k]$ (for $k \in \mathbb{N}$).

Let $S \subset \mathbb{R}^m$ be given. Let L be a subspace of $\mathbb{R}^m$.

- Unless stated otherwise, all subspaces are *linear subspaces*. Thus lines, planes and hyperplanes all pass through 0.
- $\langle S \rangle$ represents the subspace of $\mathbb{R}^m$ spanned by the vectors in S , the smallest set which is a subspace and contains all vectors in S . In particular, $\langle \emptyset \rangle = \{0\}$.
- The *canonical basis* of $\mathbb{R}^m$, also known as the *standard basis* of $\mathbb{R}^m$, is the set $\{e_i \mid i \in [m]\}$, where for all $i \in [m]$: $e_{i,j}$ is 1 if $j = i$, and 0 otherwise. We call a subspace of $\mathbb{R}^m$ a *canonical subspace* if it is spanned by a subset of the canonical basis, and denote by $\mathcal{E}_{m,k}$ the set of k -dimensional canonical subspaces of $\mathbb{R}^m$.
- $-S = \{-v \mid v \in S\}$.

- $C_L^\perp(S) = \{w \in L \mid \forall v \in S : v \perp w\}$ denotes the *orthogonal complement of S in L* ; $C^\perp(S) = C_{\mathbb{R}^m}^\perp(S)$ (when it is clear which Euclidean space is meant). We also write $C_L^\perp(v)$ to mean $C_L^\perp(\{v\})$ when v is a vector.
- $P_L(S) = \{P_L(v) \mid v \in S\}$, where $P_L(v)$ stands for the orthogonal projection of the vector v onto the subspace L .
- $\text{dist}(u, L) = \|u - P_L(u)\|$, for all $u \in \mathbb{R}^m$, is the *distance of u from the subspace L* .

2.1 Communication complexity

In this section we gather common definitions and facts related to Communication Complexity which will be needed in Chapter 3 and in Section 4.2. For statements and facts for which no proof is given here, we refer the reader to [26], the standard reference text in this field.

In the setting of Communication Complexity, two players, usually designated Alice and Bob, aim to compute the value of a function $f : X \times Y \rightarrow Z$ at some $(x, y) \in X \times Y$, with Alice having knowledge of only x , and Bob having knowledge of only y . The function f is known to both players, who are thought of as being computationally unbounded (meaning that they are able to compute arbitrary functions at no cost), but since neither of them knows the whole input, they must communicate with each other in turns in order to be able to compute $f(x, y)$, and they do this according to a *communication protocol*. Loosely speaking, a communication protocol is a set of rules according to which each of the players decides which message to send when it is their turn to communicate, based on the knowledge they possess at that point. This knowledge includes the part of the input that is known to the player, the messages

exchanged up to that point, and in some models of communication it can also include information from other sources, such as a random string or advice.

Formally, a communication protocol can be defined as follows: Consider a rooted tree whose edges are labelled with messages. We think of the nodes with even distance from the root (including the root) as belonging to Alice, and the remaining nodes as belonging to Bob, and we call such a tree a *protocol tree*. Note that to every node there corresponds uniquely a string of messages on the path from the root to that node. Now we can define a communication protocol as a function taking as input a node of the tree, the input of the player who owns the node, and any other knowledge available to the player, and returning as output a child of that node in the protocol tree. At every node, the player who owns it can be thought of as computing this function and sending the message which leads to the node returned by the function. The other player then computes the function at that node, and so on.

A protocol terminates when a leaf (child-less node) of the tree is reached, and the outcome of the protocol is the last communicated message. Note that this means that we require both players to have knowledge of the outcome at the end of the protocol. We consider all of the messages sent in a protocol as being binary strings. If the outcome of a protocol is a 1-bit message with value 1, we say that the protocol *accepts*, otherwise we say that it *rejects*. Note, however, that in models of communication where the messages are allowed to depend on additional knowledge available to the players, different outcomes may be possible for a single input. In such models of communication, a notion of the outcome of a protocol which depends only on the input is defined separately.

We will work only with *Boolean functions*, meaning functions with range $\{0, 1\}$. Furthermore, we restrict the domain by taking $X = Y = \{0, 1\}^n$, with $n \in \mathbb{N}$. All of the functions we consider will be defined for every $n \in \mathbb{N}$, and thus, as is common

in Communication Complexity, we will speak of “a” function when we actually mean an infinite sequence of functions. All statements which depend in some way on n are meant to hold for every sufficiently large value of n .

With every Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ we associate the *communication matrix* of f , $M_f \in \mathcal{M}_{2^n}(\{0, 1\})$, whose rows and columns are indexed by elements of $\{0, 1\}^n$ and which is defined such that $(M_f)_{x,y} = f(x, y)$ for all $x, y \in \{0, 1\}^n$. Many of the measures of communication complexity that we work with can be expressed in terms of combinatorial properties of the communication matrix. It will also be convenient occasionally to use a version of the communication matrix which is a *sign matrix* (a matrix whose entries are all ± 1), and we define $M_f^- \in \mathcal{M}_{2^n}(\{-1, +1\})$ such that for all i and j , $(M_f^-)_{i,j} = -1$ if $(M_f)_{i,j} = 1$, and is $+1$ otherwise.

Definition 2.1.1. A *deterministic communication protocol* is a communication protocol in which the messages of each player depend only on their knowledge of the input and the messages exchanged up to the point when a new message is sent. The *cost* of a deterministic communication protocol is the maximum number of bits communicated on any input (x, y) , taken over all inputs $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$.

Given a Boolean function f , we say that a deterministic communication protocol is a protocol *for* f (or that it *computes* f) if and only if for each $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$, the outcome of the protocol is $f(x, y)$. The *deterministic communication complexity* of f , denoted by $D(f)$, is the minimum cost of a deterministic communication protocol for f , taken over all such protocols.

The complexity class P^{cc} is defined as the class of Boolean functions which have a deterministic protocol of cost $\leq (\log n)^c$, for some constant $c > 0$.

The deterministic communication complexity of a Boolean function f is lower-bounded by the logarithm of the minimum number of *monochromatic combinatorial*

rectangles needed to partition M_f , where combinatorial rectangles are sets of the type $A \times B = \{(x, y) \mid x \in A, y \in B\}$, with $A, B \subseteq \{0, 1\}^n$, and “monochromatic” means that $(M_f)_{x,y}$ has the same value for all $(x, y) \in A \times B$. This is due to the fact that every leaf of the protocol tree of a deterministic protocol corresponds to a combinatorial rectangle on which the protocol returns the output value associated with that leaf (see [26], chap. 1). Thus every deterministic protocol induces a partition of $\{0, 1\}^n \times \{0, 1\}^n$ into monochromatic combinatorial rectangles, hence the above lower bound.

The minimum number of monochromatic combinatorial rectangles needed to partition M_f is, in turn, lower-bounded by the rank of the communication matrix, and thus we have that $D(f) \geq \log \text{rank}(M_f)$. Whether there is a constant $c \geq 1$ such that for all Boolean functions f , $D(f)$ is upper-bounded by $O((\log \text{rank}(M_f))^c)$, is one of the most famous open problems in this field, called the *log-rank conjecture* [25].

For an example of an application of the above lower bound, consider the *Equality function* $\text{EQ} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ defined as $\text{EQ}(x, y) = 1 \Leftrightarrow x = y$. It is easy to see that $M_{\text{EQ}} = I_{2^n}$ is the identity matrix of size 2^n , and hence $\text{rank}(M_{\text{EQ}}) = 2^n$. It follows that $D(\text{EQ}) = n + 1$, with the added 1 being due to the requirement that both players know the output (which forces the player who has learned the correct output first to send one more bit to the other player, informing him of the answer).

The next type of communication protocol is one where the players are given advice in addition to their part of the input. It is the Communication Complexity-equivalent of non-deterministic Turing machines.

Definition 2.1.2. Let l be a positive integer. A *non-deterministic communication protocol (with l bits of advice)* is a communication protocol in which the messages of each player can depend on their knowledge of the input, the messages exchanged up

to the point when a new messages is sent, and an *advice string* of l bits which they receive together with their inputs. Such a protocol *accepts* an input (x, y) if and only if

$$\exists w \in \{0, 1\}^l : \text{the protocol accepts on input } (x, y) \text{ with advice } w.$$

The *cost* of a non-deterministic protocol *on an accepted input* (x, y) is defined as l plus the maximum number of communicated bits, with the maximum being over all choices of w such that the protocol accepts on input (x, y) with advice w . The *cost* of a non-deterministic protocol is defined as the maximum cost on any accepted input, taken over all accepted inputs.

Given a Boolean function f , we say that a non-deterministic protocol is a protocol *for* f (or that it *computes* f) if for any input (x, y) we have that the protocol accepts (x, y) if and only if $f(x, y) = 1$. The *non-deterministic communication complexity* of f , denoted by $N(f)$, is the minimum cost of a non-deterministic protocol for f (with any number of advice bits), over all non-deterministic protocols for f .

The complexity class NP^{cc} is defined as the class of all Boolean functions for which there is a non-deterministic protocol of cost $\leq (\log n)^c$, for some constant $c > 0$.

The non-deterministic communication complexity of a function corresponds to the logarithm of the smallest number of *combinatorial rectangles* needed to cover (not partition) the 1-entries of M_f without covering any 0-entries.

Similarly to non-deterministic protocols, *co-non-deterministic protocols* can be defined, where the condition for accepting an input requires that the protocol accepts on that input with *every* possible advice string (of a prescribed length). Generalising such protocols, Babai, Frankl and Simon [3] defined *alternating protocols*, where more than one advice string is given to the players, and the condition for the protocol to accept an input applies alternating quantifiers to the advice strings.

Definition 2.1.3. Let $l, k \in \mathbb{N}$ be given, and let $l_1, \dots, l_k$ be positive integers such that $l = \sum_{i=1,k} l_i$. An *alternating communication protocol (with an existential leading quantifier, k alternations, and advice strings of lengths $l_1, \dots, l_k$)* is a communication protocol in which the messages of the players can depend on their knowledge of the input, the messages exchanged up to the point when a new message is sent, and k advice strings of lengths $l_1, \dots, l_k$ (which are received by the players at the start of the protocol). Such a protocol *accepts an input* (x, y) if and only if:

$\exists w_1 \forall w_2 \dots w_k$: the protocol accepts on input (x, y) with advice strings $w_1, \dots, w_k$.

Alternating communication protocols with a universal leading quantifier are defined similarly, except that the quantifiers in the acceptance condition are reversed.

The *cost* of an alternating protocol *on an accepted input* (x, y) is defined as l plus the maximum number of communicated bits, with the maximum being over all choices of advice strings $w_1, \dots, w_k$ such that the protocol accepts on input (x, y) with advice strings $w_1, \dots, w_k$. The *cost* of an alternating protocol is the maximum cost on any accepted input, taken over all accepted inputs.

Given a Boolean function f , we say that an alternating protocol is a protocol *for* f if for any input (x, y) , the protocol accepts (x, y) if and only if $f(x, y) = 1$.

For every fixed $k \in \mathbb{N}$, the complexity class Σ_k^{cc} is defined as the class of Boolean functions for which there is an alternating protocol with an existential leading quantifier, k alternations, and cost $\leq (\log n)^c$, for some constant $c > 0$. Similarly, Π_k^{cc} is defined as the class of Boolean functions which can be computed by alternating protocols with a universal leading quantifier, k alternations, and cost $\leq (\log n)^c$, for some constant $c > 0$. Finally, PH^{cc} is defined as $\bigcup_{k \in \mathbb{N}} \Sigma_k^{cc} (= \bigcup_{k \in \mathbb{N}} \Pi_k^{cc})$.

The complexity class $PSPACE^{cc}$ is the class of all Boolean functions for which there is an alternating protocol (with any number of alternations and advice strings

of any lengths) with cost $\leq (\log n)^c$, for some constant $c > 0$.

PH^{cc} , Σ_k^{cc} and Π_k^{cc} are the communication-equivalents of the Polynomial Hierarchy and its levels from Turing Complexity. While it is known that $\Sigma_1^{cc} \neq \Pi_1^{cc}$, the question of whether $\Sigma_2^{cc} \neq \Pi_2^{cc}$, originally posed in [3], is still open. It is also not known whether PH^{cc} “collapses” (is equal to the union of finitely many of its levels), and whether it is a proper subset of $PSPACE^{cc}$.

Note that whenever the advice strings are fixed, the communication between the two players only depends on their knowledge of the input and on the previous messages, just as it does in a deterministic communication protocol. Using this fact, any alternating communication protocol with cost $l + m$, where l is the combined length of the advice strings and m is the maximum amount of bits communicated on any accepted input, can be transformed into an alternating communication protocol with at most one more alternation, $l + m$ bits of advice, and only one bit of communication from each player, by providing the m -bit transcript of the communication between the two players as the last advice string, with an existential quantifier (see [3]). This leads to the following characterisation of the class $PSPACE^{cc}$:

Fact 2.1.4 (Proposition 4.4 in [3]). A Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ is in $PSPACE^{cc}$ if and only if for some constant $c > 0$, positive integers $l = l(n) \leq (\log n)^c$, $k = k(n) \leq (\log n)^c$, with k odd, and $l_1(n), \dots, l_k(n)$ such that $l = \sum_{i=1, k} l_i(n)$, and for two Boolean functions $g, h : \{0, 1\}^n \times \{0, 1\}^l \rightarrow \{0, 1\}$, we have that:

$$f(x, y) = 1 \Leftrightarrow \exists w_1 \in \{0, 1\}^{l_1(n)} \forall \dots \exists w_k \in \{0, 1\}^{l_k(n)} : g(x, w) \wedge h(y, w),$$

where w is the concatenation of $w_1, \dots, w_k$.

Probabilistic communication models.

Definition 2.1.5. A *(public coin) randomised protocol* is a communication protocol in which the messages of each player can depend on their knowledge of the input, the messages exchanged up to the point when a new message is sent, and a string of random bits to which the two players have shared access (it can be thought of as the outcome of a series of coin flips that can be seen by both players, hence a “public coin”). The *(worst-case) cost* of such a protocol is the maximum number of bits communicated, taken over all inputs (x, y) and all values of the random string.

Given a Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ and $\epsilon \in (0, 1/2)$, we say that a randomised protocol is an ϵ -error protocol for f if for each input (x, y) , the outcome of the protocol is $f(x, y)$ with probability at least $1 - \epsilon$ (where the probability is taken over all possible values of the random string). The *(public coin) randomised communication complexity of f with error at most ϵ* , denoted by $R_\epsilon(f)$ (usually $R_\epsilon^{\text{pub}}(f)$) is defined as the minimum cost of an ϵ -error protocol for f , taken over all such protocols. The *(public coin) randomised communication complexity of f* is defined as $R_{1/3}(f)$.

Definition 2.1.6. Given a Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ and a probability distribution μ on $\{0, 1\}^n \times \{0, 1\}^n$, the *distributional communication complexity of f with respect to μ and with error at most ϵ* , denoted by $D_\epsilon^\mu(f)$, is defined as the minimum cost of a deterministic protocol which on input (x, y) has outcome $f(x, y)$ for at least a $(1 - \epsilon)$ -fraction of the inputs, weighted by μ , where the minimum is taken over all such protocols. The maximum ϵ -error distributional complexity of a function with respect to any product distribution is denoted $D_\epsilon^\times(f)$.

Fact 2.1.7 ([39]). For every Boolean function f and every constant $\epsilon \in (0, 1/2)$,

$$R_\epsilon(f) = \max_\mu D_\epsilon^\mu(f).$$

Definition 2.1.8. An *unbounded-error probabilistic protocol* is a communication protocol in which the messages of each player can depend on their knowledge of the input, the messages exchanged up to the point when a new message is sent, and a private string of random bits of arbitrary length (“private” meaning unknown to the other player). The (*worst-case*) *cost* of such a protocol is defined as the maximum number of communicated bits, with the maximum being over all inputs and all possible values of the two random strings.

Given a Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, we say that an unbounded-error probabilistic protocol is a protocol *for* f if and only if for every input (x, y) , the outcome of the protocol is $f(x, y)$ with probability strictly greater than $\frac{1}{2}$ (with the probability being taken over all possible values of the random strings). The *unbounded-error probabilistic communication complexity of* f , denoted by $U(f)$, is defined as the smallest cost of any unbounded-error probabilistic protocol for f .

The communication complexity class UPP^c is defined as the class of all Boolean functions which have unbounded-error probabilistic protocols of cost $\leq (\log n)^c$, for some constant $c > 0$.

The unbounded-error probabilistic communication complexity of a Boolean function corresponds to an important combinatorial property of the communication matrix.

Definition 2.1.9. Given $m_1, m_2 \in \mathbb{N}$ and a matrix $M \in \mathcal{M}_{m_1, m_2}(\mathbb{R} \setminus \{0\})$, the *sign-rank of* M is defined as the smallest rank of any matrix $M' \in \mathcal{M}_{m_1, m_2}(\mathbb{R} \setminus \{0\})$ with the property that $M_{i,j}M'_{i,j} > 0$ for all $i \in [m_1], j \in [m_2]$. In other words, the sign-rank of M is the smallest value $r \in \mathbb{N}$ such that there exist two matrices $A \in \mathcal{M}_{m_1, r}(\mathbb{R})$ and $B \in \mathcal{M}_{r, m_2}(\mathbb{R})$ such that $(AB)_{i,j}M_{i,j} > 0$ for all $i \in [m_1], j \in [m_2]$.

Fact 2.1.10 (Theorem 3 in [28]). For any Boolean function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ we have that $U(f) = \log(\text{sign-rank}(M_f^-)) \pm O(1)$.

The next notion will also be important in the context of unbounded-error probabilistic protocols.

Definition 2.1.11. Let $m_1, m_2 \in \mathbb{N}$ and a matrix $M \in \mathcal{M}_{m_1, m_2}(\mathbb{R} \setminus \{0\})$ be given. Let $r \in \mathbb{N}$ be such that there exist matrices $A \in \mathcal{M}_{m_1, r}(\mathbb{R})$ and $B \in \mathcal{M}_{r, m_2}(\mathbb{R})$ such that $(AB)_{i,j} M_{i,j} > 0$ for all $i \in [m_1], j \in [m_2]$.

If $\gamma = \min_{i,j} \frac{|(AB)_{i,j}|}{\|A_i\| \|B_{\cdot,j}\|}$, then we say that the pair of matrices (A, B) *realises the matrix M with margin γ* . The *margin of M* is then defined as the supremum of the set of values of γ for which there exists a pair of matrices which realises M with margin γ .

The following are functions which will be relevant throughout the work.

Definition 2.1.12. The *Disjointness function* $\text{DISJ} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ is defined such that $\text{DISJ}(x, y) = 1$ if and only if $\sum_i x_i y_i = 0$.

The *Inner Product modulo 2* function $\text{IP}_2 : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ is defined such that $\text{IP}_2(x, y) = (\sum_i x_i y_i) \bmod 2$.

$M_{\text{IP}_2}^- \in \mathcal{M}_{2^n}(\{-1, +1\})$ is a *Sylvester's Hadamard matrix* for every $n \in \mathbb{N}$.

Boolean circuits.

We mention here only the definition of Boolean circuits and two related facts which will be used in the proof of Theorem 4.2.2.

Definition 2.1.13. A *Boolean circuit* is a finite directed acyclic graph such that every node is marked with either a Boolean function or a Boolean variable, as follows: Nodes with in-degree 0 are called *inputs*, and are marked with a variable from a finite

set $\{x_1, \dots, x_n\}$ (for some $n \in \mathbb{N}$). All other nodes are called *gates*, and each gate is marked with a Boolean function from a predetermined set of *basis functions*, in such a way that the in-degree of the node (also called the *fan-in* of the gate) matches the number of variables taken by the function it is marked with.

The Boolean function computed at a gate can be defined recursively as the composition of the function which marks the given gate, with the functions computed at the gates from which there are incoming edges.

A Boolean circuit has an *output vector*, whose elements are gates or inputs. The output of a Boolean circuit on a given input is then defined in the obvious way.

We are interested only in Boolean circuits for which the set of basis functions is $\{\vee, \wedge, \neg\}$, where the ‘ $\vee$ ’ and ‘ $\wedge$ ’ operations have fan-in 2 and ‘ $\neg$ ’ has fan-in 1.

See [38], chap. 3, for proofs of the following statements:

Fact 2.1.14. Let $k \in \mathbb{N}$ be given. The sum of k numbers of k bits each can be computed by $(\log k)^2$ -depth circuits using only the ‘ $\vee$ ’ and ‘ $\wedge$ ’ gates with fan-in 2, and with variables or negations of variables as inputs. The product of two numbers of k bits, as well as the comparison of a k -bit number with 0, can be computed by $\log(k)$ -depth circuits of the same type.

Finally, we mention a characterisation of the class $PSPACE^{cc}$ in terms of Boolean circuits, which is a straight-forward consequence of the characterisation of this complexity class given in Fact 2.1.4. This fact is used in [23] (in the proof of Theorem 5.3), without being stated explicitly.

Fact 2.1.15. A function $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ is in $PSPACE^{cc}$ if and only if for all $x, y \in \{0, 1\}^n$ we have that $f(x, y)$ can be computed by a circuit of depth $\text{polylog}(n)$ using only the ‘ $\vee$ ’ and ‘ $\wedge$ ’ gates with fan-in 2, and *with arbitrary functions of x and arbitrary functions of y as inputs*.

2.2 Arrangements of hyperplanes

In this section we review some relevant definitions and facts regarding arrangements of hyperplanes, which are needed throughout Chapter 4. See [27] for a comprehensive overview of the field, or [36] for a good introduction to the combinatorial theory of hyperplane arrangements.

Definition 2.2.1. Given $d \in \mathbb{N}$, $d \geq 1$, we say that a set H is a *(linear) hyperplane in $\mathbb{R}^d$* if for some non-zero vector $v \in \mathbb{R}^d$ we have that $H = \{w \in \mathbb{R}^d \mid \langle v, w \rangle = 0\}$ (in which case we say that v is a *normal vector of H*). Given $m \in \mathbb{N}$, $m \geq 1$, an ordered multi-set $\mathcal{A} = \{H_1, \dots, H_m\}$ is called an *arrangement of m hyperplanes in $\mathbb{R}^d$* if H_i is a hyperplane in $\mathbb{R}^d$ for each $i \in [m]$.

Throughout this work we take $m, d \in \mathbb{N}$, $d \geq 2$ and $m > d$, unless stated otherwise.

We always mean linear (as opposed to affine) real hyperplanes. With every hyperplane we associate two half-spaces (assume that its normal vector is $v \in \mathbb{R}^d \setminus \{0\}$): $\{w \in \mathbb{R}^d \mid \langle v, w \rangle > 0\}$ and $\{w \in \mathbb{R}^d \mid \langle v, w \rangle < 0\}$.

The reason for defining hyperplane arrangements as ordered multi-sets (as opposed to ordered or un-ordered sets) is that we wish to identify arrangements of m hyperplanes in $\mathbb{R}^d$ with $(m \times d)$ -matrices that have no rows equal to the zero vector. Thus we allow multiple copies of the same hyperplane in an arrangement, since a matrix can have two identical rows, and we consider the swapping of the positions of two distinct hyperplanes of an arrangement as resulting in a new arrangement, because the same is true when swapping distinct rows of a matrix. If A is an $(m \times d)$ -matrix with no rows equal to zero, we denote by $\mathcal{A}(A)$ the arrangement of hyperplanes whose normal vectors are the rows of A (in the same order).

Definition 2.2.2. Let $\mathcal{A}$ be an arrangement of m hyperplanes in $\mathbb{R}^d$. Let K be the subspace of $\mathbb{R}^d$ spanned by the normal vectors of the hyperplanes of $\mathcal{A}$, and

let $k := \dim(K)$ (= the rank of the matrix whose rows are the normal vectors). By identifying K with $\mathbb{R}^k$ (by fixing an orthonormal basis of K and mapping it bijectively to the canonical basis of $\mathbb{R}^k$), we then have an arrangement of m hyperplanes in $\mathbb{R}^k$, called the *essentialisation of $\mathcal{A}$* and denoted $ess(\mathcal{A})$, whose hyperplanes have the same normal vectors as $\mathcal{A}$, viewed as vectors in $\mathbb{R}^k$.

This is essentially the same arrangement as $\mathcal{A}$, because $\mathcal{A}$ can be recovered from $ess(\mathcal{A})$ by extending each of its hyperplanes by the orthogonal complement of K in $\mathbb{R}^d$ (see [36]).

Definition 2.2.3. Let $\mathcal{A}$ be an arrangement of hyperplanes with normal vectors $n_1, \dots, n_m$. We say that a set $R \subseteq \mathbb{R}^d$ is a *region of $\mathcal{A}$* if for some vector $s \in \{-1, +1\}^m$ we have that $R = \{w \in \mathbb{R}^d \mid \forall i \in [m] : \text{sgn}\langle w, n_i \rangle = s_i\}$ and $R \neq \emptyset$. In this case we say that s is the *signature* of the region R , and we call the set of all signatures of regions of $\mathcal{A}$ the *signature set of $\mathcal{A}$* .

We have defined the regions of an arrangement as the intersection of m open half-spaces, each corresponding to one of the m hyperplanes of the arrangement. Alternatively we can define regions as maximal non-empty subsets which are disjoint from any of the hyperplanes of the arrangement, and it can be shown that the two definitions are equivalent, and that regions are convex, open and connected, and in our case always unbounded (see [36]).

It should be noted that the definition of the signature of a region depends on the directions of the chosen normal vectors, and that choosing normal vectors with different directions for the same arrangement of hyperplanes can result in a different signature set. Whenever we speak of the signature of an arrangement without specifying the directions of the corresponding normal vectors, we assume that the directions are fixed in some arbitrary way.

Definition 2.2.4. We say that an arrangement of hyperplanes $\mathcal{A} = \{H_1, \dots, H_m\}$ in $\mathbb{R}^d$ is *in general position* if for any $k \in [m]$ and $\{i_1, \dots, i_k\} \subseteq [m]$ with $i_1 < \dots < i_k$ we have that $\dim(H_{i_1} \cap \dots \cap H_{i_k}) = \max\{d - k, 0\}$.

Our definition of “general position” differs slightly from the one used in other sources (for example [36] or [27]), where it is required that $H_{i_1} \cap \dots \cap H_{i_k} = \emptyset$ whenever $k > d$. The change is necessary because we work with linear hyperplanes and hence with arrangements that are *centred* (the intersection of all hyperplanes is non-empty), and centred arrangements can be in general position in the traditional sense only if $m \leq d$ (so that k is always $\leq d$). Therefore, it makes sense, in our setting, to say that an arrangement is in general position if the intersection of any number of its hyperplanes has *the smallest possible dimension* (as a linear subspace).

It follows directly from the above definition, and basic linear algebra, that an arrangement of hyperplanes whose normal vectors are the rows of an $(m \times d)$ -matrix A is in general position if and only if every d -minor of A is non-zero.

We will study operations on hyperplane arrangements which leave their signature set unchanged. The following fact allows us to assume that any given arrangement is in general position.

Fact 2.2.5. For every hyperplane arrangement $\mathcal{A}$ in $\mathbb{R}^d$ there is a hyperplane arrangement $\mathcal{A}'$ in general position in $\mathbb{R}^d$ such that the signature set of $\mathcal{A}$ is contained in that of $\mathcal{A}'$.

Proof. Let $A \in \mathcal{M}_{m,d}(\mathbb{R})$ be such that its rows are the normal vectors of the hyperplanes of an arrangement $\mathcal{A}$. From the definition it follows that $\mathcal{A}$ is in general position if and only if every $(d \times d)$ -submatrix of A is non-singular.

Assume that $\mathcal{A}$ is not in general position, and hence that some number of d -minors

of A are 0. Let $I \subseteq [m]$ with $|I| = d$ be the set of row indices of a singular $(d \times d)$ -submatrix of A . Let k be the number of regions of $\mathcal{A}$ and let $x_1, \dots, x_k$ each lie in a different region. The non-singular $(d \times d)$ -matrices form a dense subset of $\mathcal{M}_d(\mathbb{R})$ (with respect to the l_2 norm), and thus any singular $(d \times d)$ -matrix can be made non-singular by an arbitrarily small perturbation of its values. On the other hand we have that for any non-singular $(d \times d)$ -matrix there is an $\epsilon > 0$ such that no perturbation of its values by at most ϵ will result in a singular matrix. Since there are finitely many submatrices of A and finitely many points x_i , it follows that we can choose an $\epsilon > 0$ and a matrix $P \in \mathcal{M}_{m,d}(\mathbb{R})$ with values in $(-\epsilon, \epsilon)$, such that: $(A + P)[[d], I]$ is non-singular, no non-zero d -minor of A is zero in $A + P$, and such that Ax_i and $(A + P)x_i$ have the same sign for all coordinates, for all $i \in [k]$.

Now let $\mathcal{A}'$ be the arrangement of hyperplanes whose normal vectors are the row vectors of $A + P$. Then $\mathcal{A}'$ has all of the regions that $\mathcal{A}$ has (because the points $x_1, \dots, x_k$ still lie in regions with the same signatures), but its matrix of normal vectors has at least one fewer d -minor equal to zero compared to A . If $\mathcal{A}'$ is not in general position, we replace $\mathcal{A}$ with $\mathcal{A}'$ and repeat the above process, and do so until we obtain an arrangement whose matrix of normal vectors has all d -minors $\neq 0$, at which point the arrangement is in general position. $\square$

From the proof of Fact 2.2.5 it is easy to see the following

Fact 2.2.6. An arrangement $\mathcal{A}$ of m hyperplanes in $\mathbb{R}^d$ has the maximal number of regions (in the sense that no other hyperplane arrangement of the same dimensions has a signature set which strictly contains that of $\mathcal{A}$) if and only if it is in general position.

The number of regions can be computed exactly for many types of hyperplane arrangements, using an important combinatorial result, due to Zaslavsky [41], which

expresses the number of regions of an arrangement (which can be affine and not in general position (in the traditional sense)) in terms of the *characteristic polynomial* associated with the arrangement (not to be confused with the characteristic polynomial of a square matrix) (see also [36]). For arrangements of (linear) hyperplanes in general position, it is also possible to compute this number directly.

Fact 2.2.7. For $m \geq d \geq 1$, all arrangements of m hyperplanes in general position in $\mathbb{R}^d$ have the same number of regions, and that if we denote this number by $f(m, d)$, then for $m > d$ we have that

$$f(m, d) = f(m - 1, d - 1) + f(m - 1, d).$$

Solving this recursion we then get that $\forall m \geq 2 : f(m, 2) = 2m$, $\forall m \geq 3 : f(m, 3) = m^2 - m + 2$, and in general that $f(m, d) = O(m^{d-1})$ whenever $m \geq d$.

Proof. The proof is by induction. For the base case ($m = d$ or $d = 1$) we have that all arrangements of m hyperplanes in general position in $\mathbb{R}^m$, and all arrangements in $\mathbb{R}$, have $f(m, m) = 2^m$ and $f(m, 1) = 2$ regions, respectively, for all $m \in \mathbb{N}$. Note that $d = 1$ is the only case where the arrangement is in general position even though its hyperplanes coincide (since all of them are $\{0\}$). Next we fix $m > d > 1$ and assume that the statement holds for smaller values of d . Let $\mathcal{A} = \{H_1, \dots, H_m\}$ be an arrangement in general position in $\mathbb{R}^d$. It is easy to see (by applying the definition directly) that the arrangement $\{H_1, \dots, H_{m-1}\}$, as well as $\{H_i \cap H_m \mid i \in [m-1]\}$, regarded as an arrangement in $\mathbb{R}^{d-1}$ (by identifying H_m with $\mathbb{R}^{d-1}$), are both in general position if $\mathcal{A}$ is. We can then state that the number of regions of $\mathcal{A}$ is the number of regions of $\mathcal{A} \setminus \{H_m\}$, $f(m-1, d)$, plus the number of regions (of $\mathcal{A} \setminus \{H_m\}$) that are further divided into two by H_m (because these must be counted twice). But this number is exactly $f(m-1, d-1)$, because every region of $\{H_m \cap H_j \mid j \in [m-1]\}$ is

the intersection of H_m with a region of $\mathcal{A} \setminus \{H_m\}$. This establishes by induction that $f(m, d)$ is well-defined and that it satisfies the above recursive relation. $\square$

Definition 2.2.8. Let $\mathcal{A}$ be an arrangement of hyperplanes (not necessarily in general position), and let R be a region of $\mathcal{A}$. If a hyperplane H of $\mathcal{A}$ does not coincide with any other hyperplane of the arrangement, we say that H (*uniquely*) *delimits* the region R if and only if R is a proper subset of a region of $\mathcal{A} \setminus \{H\}$. Otherwise (if H coincides with other hyperplanes of $\mathcal{A}$), we say that H *delimits* R if and only if H delimits R as a region of $\mathcal{A}'$, the arrangement obtained by eliminating all but one of the copies of H from $\mathcal{A}$.

Note that if H is a hyperplane that is unique in the arrangement, then H delimits a region R if and only if the arrangement also has a region that “neighbours” R , meaning a region whose signature differs from that of R only on the position corresponding to H .

An $(m \times d)$ -matrix A with no rows equal to zero can also be identified with a subspace D of $\mathbb{R}^m$ which is not orthogonal to any coordinate axis, by taking D as the subspace generated by the column vectors of A . This induces an equivalence between hyperplane arrangements and subspaces, and it will be advantageous to switch between the two points of view, as certain facts pertaining to the underlying matrix will be easier to visualise in one context or the other. Naturally, all of the notions and facts that we have seen so far related to hyperplane arrangements also transfer to subspaces.

Definition 2.2.9. Let D be a d -dimensional subspace of $\mathbb{R}^m$. We say that D is *in general position* if any matrix whose column vectors form a basis of D has no d -minors equal to 0.

For every vector $s \in \{-1, +1\}^m$, the set of all vectors $v \in \mathbb{R}^m$ such that $v_i s_i > 0$

for all $i \in [m]$, is an *orthant* of $\mathbb{R}^m$, and we say that s is its *signature*. The *signature set* of the subspace D is the set of signatures of orthants which are intersected by D .

It is easy to verify that the signature set of the subspace D associated with an $(m \times d)$ -matrix A is the same as that of the hyperplane arrangement associated with A . Using the following definition, we can also give a geometric interpretation of a subspace being in general position.

Definition 2.2.10. We say that a subspace which is spanned by a subset of the canonical (or standard) basis, $\{e_1, \dots, e_m\}$, is a *canonical subspace*. In particular, $\mathbb{R}^m$ has m *canonical hyperplanes*, and we define for each $i \in [m]$ the hyperplane $H_i := \langle \{e_j \mid j \in [m] \setminus \{i\}\} \rangle$ of $\mathbb{R}^m$. We also denote by $\mathcal{E}_{m,k}$ the set of k -dimensional canonical subspaces of $\mathbb{R}^m$.

Fact 2.2.11. If D is a d -dimensional subspace of $\mathbb{R}^m$, and A is a matrix such whose column vectors form a basis of D , then D is in general position if and only if for any $E \in \mathcal{E}_{m,m-d}$, we have that $\dim(D \cap E) = 0$.

Finally, we note a useful geometric interpretation of the equivalence between subspaces and hyperplane arrangements, which does not involve the matrix they are associated with.

Definition 2.2.12. Let D be a d -dimensional subspace of $\mathbb{R}^m$ which is not contained in any canonical hyperplane. Identifying D with $\mathbb{R}^d$, we call the arrangement of m hyperplanes in $D = \mathbb{R}^d$ whose respective normal vectors are $P_D(e_i)$ (the orthogonal projection of e_i onto D), for $i \in [m]$, the *induced arrangement of D* .

As subsets of D , the m hyperplanes of the induced arrangement are in fact the sets $D \cap H_i$, with $i \in [m]$, where H_i are the canonical hyperplanes of $\mathbb{R}^m$. Whenever we speak of the induced arrangement of a subspace, we will denote its hyperplanes (viewed

as hyperplanes in $\mathbb{R}^d$) in this way. It is easy to verify that the induced arrangement of a subspace is the same, up to an rotation of $\mathbb{R}^d$, as the arrangement associated with any matrix whose columns form an orthonormal basis of the subspace.

2.3 Information theory

The following definitions and facts are used in Section 3.2. See also [9], chap. 2. We assume some familiarity with basic probability theory.

Definition 2.3.1. If μ is a probability distribution defined on a set of the form $A = A_1 \times \dots \times A_k$, for some $k \in \mathbb{N}$, then for each $i \in [k]$, the *i-th marginal of μ* , μ_i , is a distribution on A_i such that $\mu_i(x) = \sum_{y \in A, y_i = x} \mu(y)$. Denoting the product operation for distributions by ‘ $\otimes$ ’, we have that μ is a product distribution if and only if $\mu = \mu_1 \otimes \dots \otimes \mu_k$.

Definition 2.3.2. Let X be a random variable. The *entropy of X* is defined as

$$H(X) := - \sum_x \Pr[X = x] \log(\Pr[X = x]),$$

where the sum is over the support of X . (Note that for expressions such as the last one, the convention that $0 \log 0 = 0$ is used.) If X and Y are two random variables, then we define the *mutual information between X and Y* , as

$$I(X; Y) := H(X) + H(Y) - H(X, Y)$$

(where $H(X, Y)$ is the entropy of the vector random variable (X, Y)).

Definition 2.3.3. Let μ and ν be two probability distributions with the same support set. The *relative entropy (Kullback-Leibler distance) of μ and ν* is defined as $D(\mu||\nu) = \sum_x \mu(x) \log \frac{\mu(x)}{\nu(x)}$, where x ranges over the support of the distributions.

Fact 2.3.4. Let μ be a bipartite distribution on a set $A_1 \times A_2$, and let μ_X and μ_Y be the two marginal distributions of μ (on A_1 and A_2 , respectively). If X and Y are two random variables such that $X \sim \mu_X$ and $Y \sim \mu_Y$, then $I(X; Y) = D(\mu \| \mu_X \otimes \mu_Y)$.

The next fact states that the ability of two players to sample together elements from a distribution that is known to both (“together” meaning that both get to know the sampled element), allows them to also sample together elements from a distribution that is known to only one of them, with a communication cost proportional to the distance between the two distributions.

Fact 2.3.5 ([16]). Let μ and ν be two probability distributions, and let $k = D(\mu \| \nu)$. Assume that Alice and Bob have access to a public coin (shared randomness), and that they both know μ , meaning that they can get the same sample distributed according to μ . If Alice also knows ν , then she can send a message of expected length $k + 2 \log k + O(1)$ to Bob, such that both obtain the same sample element distributed according to ν (the expectation being over all possible outcomes of the random coin).

When the players sample elements using a public coin (as is the case in the protocol from Fact 2.3.5), they do so by observing a large number of coin flips (which they can do without cost, since the outcomes are seen by both), followed by one player sending a message which tells the other player which batch of observed outcomes to use as their sample.

Chapter 3

Complexity Bounds for Disjointness

3.1 Overview

In this chapter we give three new bounds on the complexity of the Disjointness problem in different models of probabilistic communication.

In Section 3.2 we consider the questions of how the correlation between the parts of the input can affect the distributional complexity of a function, and of how much correlation is necessary in order for the distributional and randomised complexities to coincide. This leads naturally to the definition of a new notion of complexity for Boolean functions, denoted by $R_\epsilon^{I \leq k}$, the study of which makes it possible to formulate precise answers to the above questions.

We then look at $R_\epsilon^{I \leq k}(\text{DISJ})$, the maximum distributional complexity of Disjointness with respect to distributions with mutual information $\leq k$ between the inputs, and prove matching upper and lower bounds of $\Theta(\sqrt{(k+1)n})$ for this measure, for any real $k \in [0, n]$ and any sufficiently small constant $\epsilon > 0$, and with an at most polynomial dependence on $1/\epsilon$ (Theorem 3.2.3 and Theorem 3.2.4).

In Section 3.3 we briefly review the *SMP* (*simultaneous message passing*) model of communication, and adapt a protocol of Aaronson and Wigderson [1] to get a non-trivial upper bound on the non-deterministic SMP communication complexity of Disjointness. We get:

Theorem 3.3.1. There is a non-deterministic SMP protocol with cost $O(n^{2/3} \log n)$ for DISJ.

3.2 The distributional complexity of Disjointness under distributions with bounded information

The connection between the distributional and the randomised complexity of a function (see Definitions 2.1.5 and 2.1.6) is expressed by a theorem of Yao [39] which is a consequence of von Neumann’s well-known minimax theorem, and which states that

$$R_\epsilon(f) = \max_{\mu} D_\epsilon^\mu(f).$$

This relationship between the two complexities makes it possible to prove lower bounds on the randomised complexity of a function by pitting *deterministic* protocols for the function against a well-chosen distribution on the inputs, and showing that, unless the cost of any protocol is large, its error (with respect to the chosen distribution) will be greater than ϵ . Distributions for which the maximum distributional complexity is attained can then be called “hard” distributions.

The intuitively “easiest” distributions are the product distributions, which are equal to the product of their marginals, or, in other words, distributions which choose the x and y parts of the input independently. The fact that product distributions are in general insufficient for obtaining tight lower bounds on the randomised complexity of

a function, was known early on: For the Disjointness function (see Definition 2.1.12), Babai, Frankl and Simon [3] showed that $\max_{\mu: \mu_X \otimes \mu_Y} D_\epsilon^\mu(\text{DISJ}) \leq O(\sqrt{n} \log n)$ (where μ_X and μ_Y denote the marginals of μ); later, Kalyanasundaram and Schnitger [19] proved that the randomised complexity of Disjointness is $\Theta(n)$, and shortly thereafter Razborov [29] gave an explicit example of a distribution with respect to which the distributional complexity of Disjointness is $\Omega(n)$. Note, however, that in this example the gap between the two measures of complexity is only quadratic. In 1997, Kushilevitz and Nisan [26] posed the question of how large the gap between the randomised communication complexity of a function f and its distributional complexity under product distributions (denoted by $D_\epsilon^\times(f)$ and defined as $\max_{\mu: \mu = \mu_X \otimes \mu_Y} D_\epsilon^\mu(f)$) can be. This question remained open until 2008, when Sherstov [32] proved that there exist functions f with $R_{1/3}(f) = \Omega(n)$ and $D_{1/3}^\times(f) = O(1)$.

Since distributions which choose the x and y parts of the input independently can not generally be used to tightly lower-bound $R_\epsilon(f)$, it is natural to ask: *How much correlation between x and y is necessary in order for a distribution to be “hard”?* To make this question precise, we first regard inputs chosen according to a distribution μ as values of a pair of random variables (X, Y) such that $(X, Y) \sim \mu$, meaning that $X \sim \mu_X$ and $Y \sim \mu_Y$, and then consider the *mutual information between X and Y* , $I(X; Y) = H(X) + H(Y) - H(X, Y)$, which is also the relative entropy (Kullback-Leibler distance) between the distribution μ and the product of its marginals, $D(\mu || \mu_X \otimes \mu_Y) = \sum_x \sum_y \mu(x, y) \log \frac{\mu(x, y)}{\mu_X(x) \mu_Y(y)}$ (see Section 2.3 or [9], chap. 2). Intuitively, this is a measure of how far μ is from being a product distribution, and we will refer to distributions for which this measure is upper-bounded as “distributions with bounded (mutual) information”. It is then natural to define the following class of communication complexity measures: For any $\epsilon \in (0, 1/3)$, any function

$f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$, and any (real) $k \in [0, n]$, let

$$R_\epsilon^{I \leq k}(f) := \max_{\mu: I(X; Y) \leq k} D_\epsilon^\mu(f).$$

Evidently this quantity is increasing as a function of k . Our question can now be restated as follows: Given a function f and a small constant ϵ , for which values of k (as a function of n) does it hold that $R_\epsilon^{I \leq k}(f) = R_\epsilon(f)$? (Note that for distributions on $\{0, 1\}^n \times \{0, 1\}^n$, the mutual information between the two parts of the input is always $\leq n$, and thus $R_\epsilon^{I \leq n}(f) = R_\epsilon(f)$ whenever f is defined on $\{0, 1\}^n \times \{0, 1\}^n$.) More generally, one can ask how $R_\epsilon^{I \leq k}(f)$, regarded as a function of k and n , changes from $R_\epsilon^{I \leq 0}(f) = D_\epsilon^\times(f)$ to $R_\epsilon^{I \leq n}(f) = R_\epsilon(f)$ as k goes from 0 to n . This type of complexity measure was first investigated by Jain and Zhang [18] in the setting of one-way communication complexity (where only one of the two players is allowed to send messages).

The question regarding the largest possible gap between $D_\epsilon^\times(f)$ and $R_\epsilon(f)$ can also be extended to $R_\epsilon^{I \leq k}(f)$, and we ask: What is the largest value of k (relative to n) for which the gap between $R_\epsilon^{I \leq k}(f)$ and $R_\epsilon(f)$ can be large? The strongest possible answer to this question would be that for some f it holds that $R_\epsilon^{I \leq k}(f) = O(1)$ for all $k \leq o(n)$, while $R_\epsilon(f) = \Omega(n)$ (which would extend Sherstov's result from [32]). A result that is almost as strong can actually be obtained via a probabilistic argument which shows that for every d with $6n \leq d \leq 2^n$, there exists a function f_d such that $R_{1/10}^{I \leq n/1000}(f_d) = O(\log n)$ but $R_{1/3}(f_d) = \Theta(\log d)$. For $d = \Omega(2^n)$ we get an $O(\log n)$ vs. $\Omega(n)$ gap between the distributional complexity under distributions with mutual information $\leq o(n)$, and the randomised complexity (see [6] for the detailed proof).

When studying the distributional complexity of a particular function under distributions with bounded mutual information, a possible approach is to try to adapt existing bounds. To obtain protocols in this situation, one can try to modify protocols

which work for product distributions, so that they still work when the parts of the input are correlated, but in such a way that the cost of the protocol increases gradually with the amount of correlation. Likewise, for lower bounds on the complexity under distributions with bounded information, one can look at proofs of lower bounds on the randomised complexity which exhibit explicit “hard” distributions, try to modify the distributions so that the mutual information between the parts of the input can be “dialed down”, and investigate how the original proof works with the modified distribution. As we will see below, this approach can work, but in some cases, entirely new ideas may be required.

When designing protocols to work when the parts of the input are correlated, the following fact can be very helpful.

Fact 3.2.1. Recall that the distributional complexity of f under μ and with error $\leq \epsilon$ is the smallest cost of any *deterministic* protocol for f with error $\leq \epsilon$ under μ . We may in fact take the minimum cost of ϵ -error *randomised* protocols to compute $D_\epsilon^\mu(f)$. This is due to the fact that fixing the random string turns randomised protocols into deterministic ones. If $P(r, x, y)$ denotes the deterministic protocol with input (x, y) obtained when the random string is r , we have that $\mathbb{E}_r\{P(r, x, y) = f(x)\} \geq 1 - \epsilon$ for *any* $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$, and hence that $1 - \epsilon \leq \mathbb{E}_{(x, y) \sim \mu} \mathbb{E}_r\{P(r, x, y) = f(x, y)\} = \mathbb{E}_r \mathbb{E}_{(x, y) \sim \mu}\{P(r, x, y) = f(x, y)\}$. We may conclude that for some r , $\mathbb{E}_{(x, y) \sim \mu}\{P(r, x, y) = f(x, y)\}$, meaning that the protocol $P(r, \cdot)$ has error $\leq \epsilon$ *under* μ . In other words, we can always fix the randomness of a protocol in such a way that the error of the resulting deterministic protocol with respect to a given distribution does not increase.

We now turn to the question of how $R_\epsilon^{I \leq k}(\text{DISJ})$ behaves as k goes from 0 to n . As

mentioned before, Babai *et al.* [3] gave a protocol which solves Disjointness with cost $O(\sqrt{n} \log n)$ for any product distribution on the inputs, and an almost matching lower bound of $\Omega(\sqrt{n})$ for the complexity under such distributions; Razborov [29] proved for an explicit distribution μ on $\{0, 1\}^n \times \{0, 1\}^n$ that $D_\epsilon^\mu(\text{DISJ}) = \Omega(n)$ for any sufficiently small ϵ . It turns out, as we show next, that the distributional complexity of Disjointness under distributions with bounded mutual information nicely interpolates between the two extremal values, as $R_\epsilon^{I \leq k}(\text{DISJ}) = \Theta(\sqrt{(k+1)n})$, with an at most polynomial dependence on $1/\epsilon$, for any sufficiently small constant $\epsilon > 0$. Note that the tight upper bound implies that the only “hard” distributions for Disjointness are the ones with $\Omega(n)$ mutual information between the parts of the input.

Upper bound on $R_\epsilon^{I \leq k}(\text{Disj})$.

In this subsection we speak of the *expected* cost of a protocol on an input (x, y) , the expectation being over the values of the shared random string. Generally, upper-bounding the expected cost of a protocol is a weaker statement than upper-bounding its worst-case cost: If with some very small probability the amount of communication far exceeds the average, then this might not affect the expected cost, while greatly increasing the worst-case cost of the protocol. However, as the next Fact shows, it is always possible to turn bounds on the expected cost into worst-case bounds, with only a small increase in the error of the protocol and a constant-factor increase of its cost.

Fact 3.2.2. Given a protocol P which computes some Boolean function with error $\leq \frac{2}{3}\epsilon$ and expected cost $\leq c$ on every input (x, y) , there is a protocol P' computing the same function with error $\leq \epsilon$ and worst-case cost $\leq O(c)$.

Proof. For P' , the players simply execute P while also keeping track of the number

of bits they communicate. If this number reaches $3c/\epsilon = O(c)$, they reject. While this may lead to them giving wrong answers, we have by Markov's inequality and the upper-bound of c on the expected cost, that such an event can occur with probability at most $\epsilon/3$. Adding this to the total error of P , we get that the protocol P' has error $\leq \epsilon$ and cost $O(c)$. $\square$

For the upper bound on $R_\epsilon^{I \leq k}(\text{DISJ})$ we first review the protocol from [3] for the case of product distributions, and then show how it can be generalised to work in the presence of correlation, and improved so as to remove the log-factor in the cost.

In the original protocol of Babai *et al.*, the two players keep track, throughout the protocol, of a set $V \subset [n]$ where the two input sets might still intersect. At the beginning, $V = [n]$. They then communicate for a number of rounds, reducing the size of V by at least $\sqrt{n}$ after each round. Once the intersection of x or y with V has size $\leq \sqrt{n}$, the corresponding player communicates this intersection set, using $O(\sqrt{n} \log n)$ bits, which allows the other player to produce the correct output. In each round of the protocol, the players first compare $|x \cap V|$ with $|y \cap V|$, and the one with the smaller set speaks. Assume that it is Alice, and hence that $\sqrt{n} < |x \cap V| \leq |y \cap V|$. She can then compute the probability that a set y' chosen according to the marginal distribution on Bob's inputs can be such that $|y' \cap V| = |y \cap V|$, and still be disjoint from x . If this probability is small, Alice rejects and terminates the protocol, otherwise the two use the public coin to obtain a shared sample of such a set, and take $V = V \setminus y'$. We note here that in the original proof, Babai *et al.* do not make use of Fact 3.2.1, and instead give a deterministic protocol in which the players have knowledge of a fixed database from which to pick sets which they subtract from V ; this introduces another source of error, and the existence of a suitable database for every product distribution must be shown via a probabilistic argument. Either way, the expected cost of communicating

a chosen set to the other player is $O(\log(1/\epsilon))$, and since the protocol has at most $O(\sqrt{n})$ rounds, the total expected cost is $O(\sqrt{n} \log n)$.

The above protocol can be generalised so as to obtain for any distribution on the inputs with mutual information $\leq k$, a protocol with cost $O(\sqrt{(k+1)n})$. The new protocol follows the same two-phase strategy in which the players first repeatedly reduce the universe of possible intersections, and then solve the problem when the sets are small. However, two technical issues arise when the parts of the input are allowed to be correlated.

First, as the players pass messages to one another, they learn more about the whole input, which restricts the support of the distribution on the inputs from which they sample elements. For deterministic and randomised protocols, this restriction is always to a combinatorial sub-rectangle of $\{0,1\}^n \times \{0,1\}^n$, but the mutual information between the random variables representing the parts of the input can increase when conditioning on smaller combinatorial rectangles, meaning that after one round of communication we no longer have the condition that the relative entropy between the *restricted* distribution and the product of its marginals is $\leq k$. This issue can be solved by observing that *on average* the relative entropy does not increase. Specifically, we have that $I(X;Y|R_i) \leq k$, where R_i is a random variable, depending on the input (x,y) and the outcomes of the public coin, which represents the combinatorial rectangle that has been reached at some point in round i (see [6] for the details of this computation).

The second issue is that whenever one of the players, for example Alice, is choosing a set y' with large intersection with the current set V and disjoint from x , the distribution according to which she is sampling now depends on her part of the input (unlike the case when the distribution is product), and thus is unknown to Bob. However, using the above-mentioned fact that $I(X;Y|R_i) \leq k$, we can show that

the expected relative entropy between the distribution according to which Alice must sample at round i , and a distribution which is known to both players, is $\leq k$ (the expectation is over both $x \in \{0, 1\}^n$ and the outcomes of the public coin): Let $\mu_{r_i}^x$ be the distribution μ conditioned on $R_i = r_i$ (the rectangle that has been reached before sampling in round i) and $X = x$, and let $\mu_{r_i, Y}^x$ be its marginal distribution on Bob's inputs (this is the distribution which is known only to Alice, and according to which she picks a set y' that is disjoint from x at round i). Let μ_{r_i} be the distribution μ conditioned on $R_i = r_i$, and let $\mu_{r_i, Y}$ be its marginal distribution on Bob's inputs (this is known to both players). Then we have:

$$\begin{aligned}
\mathbb{E}_{r_i, x} D(\mu_{r_i, Y}^x \| \mu_{r_i, Y}) &= \mathbb{E}_{r_i} \left[\sum_x \mu_{r_i, X}(x) \sum_y \mu_{r_i, Y}^x(y) \log \frac{\mu_{r_i, Y}^x(y)}{\mu_{r_i, Y}(y)} \right] \\
&= \mathbb{E}_{r_i} \left[\sum_x \mu_{r_i, X}(x) \sum_y \frac{\mu_{r_i}(x, y)}{\mu_{r_i, X}(x)} \log \frac{\mu_{r_i}(x, y)}{\mu_{r_i, X}(x) \mu_{r_i, Y}(y)} \right] \\
&= \mathbb{E}_{r_i} D(\mu_{r_i} \| \mu_{r_i, X} \otimes \mu_{r_i, Y}) = I(X; Y | R_i) \leq k.
\end{aligned}$$

Thus we may use Fact 2.3.5, and claim that using the public coin, and with a message of expected length $k + 2 \log k + O(1)$, Alice and Bob can obtain the same sample drawn according to Alice's distribution.

Finally, we discuss how to remove the log-factor from the cost. Notice that in the original protocol of Babai *et al.*, there are two sources for the factor of $\log n$ in the cost: the fact that each round the players inform each other of the size of a subset of $[n]$ to decide whether they can move to phase 2 ($O(\log n)$ bits times $O(\sqrt{n})$ rounds), and that in phase 2 one of the players communicates a subset of $[n]$ with $O(\sqrt{n})$ elements ($O(\sqrt{n} \log n)$ bits). The second factor can be eliminated by requiring that the players continue phase 1 until *both* $|x \cap V|$ and $|y \cap V|$ are small, and then use the protocol of Håstad and Wigderson [17] which solves Disjointness for sets of size s with cost $O(s)$. With this requirement on $|x \cap V|$ and $|y \cap V|$, it is also not necessary for

the players to communicate the sizes of sets at the start of each round: they can send one bit each to signify whether their sets are still large, and either continue phase 1 or move on to phase 2, accordingly.

Thus our protocol for Disjointness under bipartite distributions with bounded mutual information works as follows: In phase 1, the two players again keep track of a set V of possible intersections, which they gradually reduce. In each round of phase 1, the two players first communicate one bit each, to indicate whether the intersection of their set with V has size $\leq \sqrt{(k+1)n}$. If both intersection sets are smaller than this value, they move on to phase 2. Otherwise, Alice speaks, unless her set is large while Bob's is small (in which case Bob plays the same role as Alice does in what follows). At round i , if Alice speaks, she first computes the probability that a set y' drawn according to $\mu_{r_i, Y}^x$ is disjoint from x (note that such a y' is necessarily such that $|y' \cap V| > \sqrt{(k+1)n}$, since the fact that Alice speaks already means that Bob has communicated that $y \cap V$ is large, and this fact is reflected in r_i , on which we condition). If this probability is $< \epsilon/2$, she terminates the protocol with rejection. Otherwise she communicates a message of length $O(k+1)$ which allows the two players to sample an element distributed according to $\mu_{r_i, Y}^x$. By repeating this last process $O(\frac{1}{\epsilon} \log \frac{1}{\epsilon})$ times, they will have obtained with high probability ($> 1 - \epsilon/3$) an element y' that is disjoint with x . They set $V = V \setminus \{y'\}$ and move on to the next round. Since the size of V decreases by $\sqrt{(k+1)n}$ elements each round, phase 1 can have at most $n/\sqrt{(k+1)n} = \sqrt{\frac{n}{k+1}}$ rounds, each with an expected communication cost of $O(\frac{1}{\epsilon} \log(\frac{1}{\epsilon})(k+1))$ bits. Phase 2 consists of the above-mentioned protocol of Håstad and Wigderson for small-set Disjointness [17].

We then have

Theorem 3.2.3. For any $\epsilon \in (0, 1)$ and any distribution μ on $\{0, 1\}^n \times \{0, 1\}^n$

with mutual information $\leq k$ ($k \in [0, n]$), there is a protocol which solves DISJ with cost $O(\sqrt{(k+1)n})$ and error at most ϵ under μ . In other words, $R_\epsilon^{I \leq k}(\text{DISJ}) \leq O(\sqrt{(k+1)n})$. The error dependence of the complexity is $\leq \frac{1}{\epsilon} \log(\frac{1}{\epsilon})$.

We remark that the dependence of $R_\epsilon^{I \leq k}(\text{DISJ})$ on ϵ can actually be improved to $1/\epsilon$ (see [6]).

Lower bound on $R_\epsilon^{I \leq k}(\text{Disj})$.

For the lower bound we give a detailed proof in which we exhibit a distribution, depending on n and k , such that the mutual information of the two parts of the input, which are respectively distributed according to the marginals, is at most k ; we then prove an $\Omega(\sqrt{n(k+1)})$ lower-bound for the distributional complexity under this distribution. In what follows we consider $k = k(n)$ as being $\in o(n)$, since for $k = \Omega(n)$ the upper bound on the information is trivial and the lower bound on the communication is known.

Let $c = \frac{1}{\log e}$ and $m = c\sqrt{n(k+1)}$. Note that $m = o(n)$ as well. Now $\mu_{n,k}$ can be defined as the distribution obtained by mixing two distributions, one where a pair of disjoint subsets of $[n]$ of size m is chosen uniformly among all such pairs, and one where a pair of subsets of size m with intersection of size 1 is chosen uniformly among all such pairs. This is essentially the distribution used in the proof by Razborov [29], but with smaller sets. Letting X and Y follow the marginal distributions of $\mu_{n,k}$, respectively, we have:

$$\begin{aligned} I(X; Y) &= H(X) - H(X|Y) = \log \binom{n}{m} - \mathbb{E}_{y \in Y} (\mu_{n,k,Y}(y) H(X|Y = y)) \\ &= \log \binom{n}{m} - H(X|Y = y_0) \text{ (where } y_0 \text{ is any set with } \mu_{n,k,Y}(y_0) > 0) \\ &= \log \binom{n}{m} - \left(2 \log \binom{n-m}{m} + 2 \log \binom{n-m}{m-1} + 2 \right) \end{aligned}$$

$$\begin{aligned}
&\leq \log \binom{n}{m} - \log \binom{n-m}{m} = \log \frac{n(n-1) \cdots (n-m+1)}{(n-m)(n-m-1) \cdots (n-2m+1)} \\
&\leq \log \left(1 + \frac{m}{n-2m+1} \right)^m \leq (\log e) \frac{m^2}{n-2m+1} \\
&= c^2 (\log e) (1 + o(1)) (k+1) \leq k,
\end{aligned}$$

for any sufficiently large n .

Regarding the lower bound in the theorem below, we remark that while the structure of the proof is similar to that of the original proof of Razborov [29], two major difficulties arise when working with smaller sets: The first is that by mixing the two distributions with similar probabilities ($\frac{1}{4}$ and $\frac{3}{4}$, respectively), the weight of any pair of intersecting sets is much larger than that of a pair of disjoint sets. Since the proof relies heavily on the properties of the distribution when conditioned on certain events, and in particular on the proportion of the weight of 1-inputs and the weight of 0-inputs when conditioning, this imbalance complicates several computations. The second difficulty comes from the fact that Razborov's entropy "counting" argument no longer works in our case, because in that argument a linear number of terms have their entropy upper-bounded as $H(\frac{1}{2}) = 1$. Since we still have to deal with a linear number of terms, while having much less total entropy, we require a much finer combinatorial counting argument instead.

Theorem 3.2.4. For any sufficiently small $\epsilon > 0$ we have that $D_\epsilon^{\mu_{n,k}}(\text{DISJ}) = \Omega(\sqrt{n(k+1)})$, and hence that $R_\epsilon^{I \leq k}(\text{DISJ}) = \Omega(\sqrt{n(k+1)})$.

Proof. We may assume that $k = o(n)$, since otherwise (if $k = \Omega(n)$), the original proof by Razborov [29] applies directly. Let $l \in \mathbb{N}$ be given and assume that $n = 4l - 1$. Let $\gamma = \log_l(c\sqrt{n(k+1)})$, where $c = (\log e)^{-1}$. Thus $\gamma \in (\frac{1}{2}, 1)$ (for n sufficiently large) and our distribution will pick sets of size $l^\gamma = c\sqrt{n(k+1)}$. Throughout the proof we will treat numbers like l^γ as natural numbers, and avoid using the floor function for

the sake of readability. We will also identify $\mathcal{P}([n])$ with $\{0, 1\}^n$.

We now give an alternative definition for the distribution $\mu = \mu_{n,k}$, as the distribution induced by the following process: First, a triple $T = (T_1, T_2, i)$ is chosen uniformly among all such triples, where $|T_1| = |T_2| = 2l - 1$ and $\{T_1, T_2, \{i\}\}$ form a partition of the set $[n]$. Then, with probability $\frac{1}{2}$ the set x is chosen uniformly among all subsets of $T_1 \cup \{i\}$ with l^γ elements and such that they contain i , and with probability $\frac{1}{2}$ the set x is chosen as a subset of T_1 with l^γ elements, again uniformly among all such subsets of T_1 . Similarly, and independently of the choice of x , y is chosen with probability $\frac{1}{2}$ uniformly as a subset of $T_2 \cup \{i\}$ with l^γ elements and such that it contains i , and with probability $\frac{1}{2}$ uniformly among the subsets of T_2 with l^γ elements (not containing i). Thus non-zero probabilities are assigned only on the set $\{(x, y) \mid x, y \subseteq [n], |x| = |y| = l^\gamma, |x \cap y| \in \{0, 1\}\}$.

Now the statement that $D_\epsilon^{\mu_{n,k}}(\text{DISJ}) = \Omega(\sqrt{n(k+1)})$ for any sufficiently small constant $\epsilon > 0$, follows from Lemma 3.2.5 below by the following argument: Recall that every deterministic protocol induces a partition of $\{0, 1\}^n \times \{0, 1\}^n$ into combinatorial rectangles, such that the protocol returns the same value on all elements of a rectangle. Since we are considering the distributional complexity with respect to μ , the best deterministic protocol may have error (meaning positions where a 1-rectangle of the protocol covers a 0 of the function, or vice-versa) at most ϵ under μ . However, the lemma states that unless the weight of a rectangle under μ is very small (close to $2^{-\Omega(n^\gamma)}$), a constant fraction of this weight will contribute to the error. In other words, either the partition induced by the protocol uses $2^{\Omega(n^\gamma)}$ or more rectangles, *or* its error will be larger than ϵ . $\square$

Lemma 3.2.5. Let γ, μ be defined as in the proof of Theorem 3.2.4. Let $A = \{(x, y) \mid \mu(x, y) > 0 \text{ and } x \cap y = \emptyset\}$ and $B = \{(x, y) \mid \mu(x, y) > 0 \text{ and } x \cap y \neq \emptyset\}$. For any

sufficiently small $\epsilon > 0$ we have for any rectangle $R = C \times D \subseteq \mathcal{P}([n])^2$ that

$$\mu(B \cap R) \geq \Omega(\mu(A \cap R)) - 2^{-\Omega(n^\gamma)}.$$

Proof. We consider $\epsilon > 0$ to be fixed (but will specify its value later). We begin by defining for any triple $T = (T_1, T_2, \{i\})$ as above, the numbers $Row(T) = \mathbf{Pr}[x \in C \mid x \subseteq T_1 \cup \{i\}]$, $Row_0(T) = \mathbf{Pr}[x \in C \mid x \subseteq T_1 \cup \{i\}, i \notin x]$ and $Row_1(T) = \mathbf{Pr}[x \in C \mid x \subseteq T_1 \cup \{i\}, i \in x]$, and similarly $Col(T) = \mathbf{Pr}[y \in D \mid y \subseteq T_2 \cup \{i\}]$, $Col_0(T) = \mathbf{Pr}[y \in D \mid y \subseteq T_2 \cup \{i\}, i \notin y]$ and $Col_1(T) = \mathbf{Pr}[y \in D \mid y \subseteq T_2 \cup \{i\}, i \in y]$. It is important to note that $Row(T) = \frac{1}{2}(Row_0(T) + Row_1(T))$ and $Col(T) = \frac{1}{2}(Col_0(T) + Col_1(T))$, just as in the case of Razborov's original distribution, and for the same reasons.

Next, for a triple $T = (T_1, T_2, \{i\})$ (and under the above distribution μ) we say that T is *x-bad* if $Row_1(T) < \frac{1}{6}Row_0(T) - 2^{-\epsilon n^\gamma}$, and that T is *y-bad* if $Col_0(T) < \frac{1}{6}Col_1(T) - 2^{-\epsilon n^\gamma}$. If T is *x-bad* or *y-bad*, we say that T is *bad*. Let $Bad_x(T)$, $Bad_y(T)$ and $Bad(T)$ be the respective event indicators.

Claim 1. For all $t_2 \subseteq [n]$, with $|t_2| = 2l - 1$, we have that $\mathbf{Pr}[Bad_x(T) = 1 \mid T_2 = t_2] \leq \frac{1}{5}$ and $\mathbf{Pr}[Bad_y(T) = 1 \mid T_2 = t_2] \leq \frac{1}{5}$.

Proof of Claim 1. We prove the first statement, the second one having an almost identical proof.

Let $t_2 \subseteq [n]$, with $|t_2| = 2l - 1$, be fixed. Under our distribution, $Row(T)$ can take different values even when T is restricted to partitions for which $T_2 = t_2$. Thus we first treat the case when $\max\{Row(T) \mid T_2 = t_2\} \leq 2^{-\epsilon n^\gamma}$. If this inequality holds, then for all T with $T_2 = t_2$ we have: $Row(T) \leq 2^{-\epsilon n^\gamma}$, and hence $Row_0(T) \leq 2Row(T) \leq 2 \cdot 2^{-\epsilon n^\gamma}$ so that $\frac{Row_0(T)}{6} - 2^{-\epsilon n^\gamma} < 0 \leq Row_1(T)$ holds trivially (and hence

$\mathbf{Pr}[Bad_x(T) = 1 \mid T_2] = 0$.

Next we treat the case where $\max\{Row(T) \mid T_2 = t_2\} > 2^{-\epsilon n^\gamma}$. Define $S = \{x \in C \mid |x| = l^\gamma, x \subset [n] \setminus t_2\}$. Note that for any T with $T_2 = t_2$, $Row(T)$ measures the conditional probability (conditioned on T) of the same set S , with each $x \in S$ having a different (conditional) probability depending on whether $i \in x$. Specifically, if $i \in x$ then the probability of x being chosen, conditioned on T , is $\frac{1}{2} \binom{2l-1}{l^\gamma-1}^{-1}$, otherwise the probability is $\frac{1}{2} \binom{2l-1}{l^\gamma}^{-1} = \frac{1}{2} \binom{2l-1}{l^\gamma-1}^{-1} \frac{l^\gamma}{2l-l^\gamma} = \frac{1}{2} \binom{2l-1}{l^\gamma-1}^{-1} \frac{1}{2l^{1-\gamma}-1}$. Thus, when T is fixed, the probability of each set x containing i is $2l^{1-\gamma} - 1$ times that of a set which does not contain i .

The proof of this case will proceed as follows: First, we show that under the assumption that a sufficiently large part of the partitions T with $T_2 = t_2$ are x -bad, three quarters of the elements of S (which are subsets of $[n] \setminus T_2$) must have at least $\frac{21}{25}$ of their elements in a subset of $[n] \setminus T_2$ of size $\frac{8l}{5}$. We will then upper-bound the number of subsets of $[n] \setminus T_2$ of size l^γ that have this property (regardless of whether they are in C or not). Next, we will lower-bound $\frac{3}{4}|S|$ in terms of ϵ , and show that for a suitable choice of ϵ , the lower bound for $\frac{3}{4}|S|$ is in fact larger than the upper bound we computed before, which is a contradiction showing that it is not possible for that T with $T_2 = t_2$ to be x -bad for that many choices of i .

Note first that whenever T_2 is fixed (in our case to t_2), the choice of $i \in [n] \setminus T_2$ also fixes T_1 and hence all of T , and that the choice of i determines the proportion of $x \in S$ whose weights are counted in $Row_1(T)$. If for a particular choice of i the resulting T is x -bad, then by definition we have that $Row_1(T) < \frac{1}{6}Row_0(T) - 2^{-\epsilon n^\gamma}$, and in particular that $Row_1(T) < \frac{1}{6}Row_0(T)$. If we let S' be the set of $x \in S$ with $i \in x$, then we may rewrite this inequality as:

$$\frac{|S'|}{\binom{2l-1}{l^\gamma-1}} < \frac{|S| - |S'|}{6\binom{2l-1}{l^\gamma}} \iff |S'| < \frac{|S| - |S'|}{6(2l^{1-\gamma} - 1)}$$

$$\iff |S'| \left(1 + \frac{1}{6(2l^{1-\gamma} - 1)}\right) < \frac{|S|}{6(2l^{1-\gamma} - 1)},$$

and we may conclude that for l sufficiently large, $|S'| < \frac{|S|}{10l^{1-\gamma}}$ (under the assumption that T is x -bad). For the last inequality we have used the fact that $\lim_{n \rightarrow \infty} l^{1-\gamma} = \infty$, which holds by the following computation: $\lim_{n \rightarrow \infty} \log l^{1-\gamma} = \lim_{n \rightarrow \infty} (1 - \gamma) \log l = \lim_{n \rightarrow \infty} (1 - \log_l(c\sqrt{n(k+1)})) \log l \geq \lim_{n \rightarrow \infty} (\log l - \log \sqrt{n(k+1)}) \geq \lim_{n \rightarrow \infty} \log \sqrt{\frac{n+1}{16(k+1)}} = \infty$ (since $k = o(n)$).

Let $B = \{i \in [n] \mid \text{the partition } ([n] \setminus (t_2 \cup \{i\}), t_2, \{i\}) \text{ is } x\text{-bad}\}$, and assume that $|B| \geq \frac{2l}{5}$, that is, assume that for at least one fifth of the possible choices for i the corresponding partition is x -bad. By excluding some elements of B , we may assume that $|B| = \frac{2l}{5}$. Now, if we consider the number of pairs (x, i) with $x \in S$ and $i \in x$, we have by the inequality in the last paragraph that each of the $i \in B$ can be the second element of at most $\frac{|S|}{10l^{1-\gamma}}$ such pairs, and hence B can contribute the second element of at most $\frac{2l}{5} \frac{|S|}{10l^{1-\gamma}} = \frac{1}{25} l^\gamma |S|$ of the total of $l^\gamma |S|$ pairs. Applying the Lemma 3.2.6 below with $X = S$, $Y = [l^\gamma]$, $c(x, i) = 0$ if and only if the i -th smallest element of x is in B (so that $p \geq \frac{24}{25}$) and $r = \frac{21}{25}$, we have that at least three quarters of all $x \in S$ have the property that more than $\frac{21}{25}$ of their elements lie in $G = [n] \setminus (t_2 \cup B)$. Let Q be the set of subsets $x \subseteq B \cup G = [n] \setminus t_2$, with $|x| = l^\gamma$ and the property that $|x \cap G| \geq \frac{21}{25} l^\gamma$. Then we must have that $|Q| \geq \frac{3}{4} |S|$. We will now upper-bound the size of the set Q .

Since every $x \in Q$ can have a proportion of at most $4/25$ of its elements in B , we have that

$$\begin{aligned} \log |Q| &\leq \log \left[\sum_{i=0}^{\frac{4}{25} l^\gamma} \binom{\frac{2l}{5}}{i} \binom{\frac{8l}{5}}{l^\gamma - i} \right] \leq \log \left[\sum_{i=0}^{\frac{4}{25} l^\gamma} \left(\frac{2le}{5i} \right)^i \left(\frac{8le}{5(l^\gamma - i)} \right)^{l^\gamma - i} \right] \\ &\leq \log \left[\frac{4}{25} l^\gamma \left(\frac{2le}{5} \frac{25}{4l^\gamma} \right)^{\frac{4}{25} l^\gamma} \left(\frac{8le}{5} \frac{25}{21l^\gamma} \right)^{\frac{21}{25} l^\gamma} \right] \\ &= \log \left[\frac{4}{25} l^\gamma \left(\frac{5e}{2} l^{1-\gamma} \right)^{\frac{4}{25} l^\gamma} \left(\frac{40e}{21} l^{1-\gamma} \right)^{\frac{21}{25} l^\gamma} \right] \end{aligned}$$

$$\begin{aligned}
&\leq \gamma \log l + \frac{4}{25} l^\gamma \log \left(\frac{5e}{2} l^{1-\gamma} \right) + \frac{21}{25} l^\gamma \log \left(\frac{40e}{21} l^{1-\gamma} \right) + O(1) \\
&= (1 - \gamma) l^\gamma \log l + \left(\frac{4}{25} \log \frac{5e}{2} + \frac{21}{25} \log \frac{40e}{21} \right) l^\gamma + O(\log l) \\
&\leq (1 - \gamma) l^\gamma \log l + 2.43508 \cdot l^\gamma + O(\log l),
\end{aligned}$$

where in the first line we used the inequality $\binom{m}{k} \leq \left(\frac{em}{k}\right)^k$ for each term of the sum.

The inequality sign between the first and second line can be justified as follows: For

$x \in (0, \frac{1}{2})$, consider the expression

$$\begin{aligned}
&\log \left[\left(\frac{\frac{2}{5}el}{x \cdot l^\gamma} \right)^{x \cdot l^\gamma} \left(\frac{\frac{8}{5}el}{(1-x)l^\gamma} \right)^{(1-x)l^\gamma} \right] \\
&= (1 - \gamma) l^\gamma \log l + l^\gamma \left(x \log \frac{2e}{5x} + (1-x) \log \frac{8e}{5(1-x)} \right)
\end{aligned}$$

and set $f(x) = x \log \frac{2e}{5x} + (1-x) \log \frac{8e}{5(1-x)} = x \log \frac{1}{x} + (1-x) \log \frac{1}{1-x} + (1+\log e)x + (3+\log e)(1-x) - \log 5 = 3 + \log e - \log 5 + H(x) - 2x$. Then $f'(x) = H'(x) - 2 = \log \frac{1-x}{x} - 2$.

Note that the function $\frac{1-x}{x}$ is decreasing but positive on $(0, 1)$, and we have that the smallest value $x_0 \in (0, \frac{1}{2})$ for which we can have $f'(x_0) = 0$ is $x_0 = \frac{1}{5}$, which implies that $f(x)$, and hence also the argument of the logarithm in the expression above, is strictly increasing on $(0, \frac{1}{5})$. Thus the terms of the sum

$$\sum_{i=0}^{\frac{4}{25} l^\gamma} \left(\frac{2le}{5i} \right)^i \left(\frac{8le}{5(l^\gamma - i)} \right)^{l^\gamma - i}$$

are increasing, so that each term is upper-bounded by the final term, which justifies the inequality between the the first and second line above.

Next we compute a lower bound for $\frac{3}{4}|S|$. Let T^* be a partition with $T_2^* = t_2$ and $Row(T^*) = \max\{Row(T) \mid T_2 = t_2\}$. Then we have that $\frac{3}{4}|S| = \frac{3}{4}[Row_0(T^*)\binom{2l-1}{l^\gamma} + Row_1(T^*)\binom{2l-1}{l^\gamma-1}] \geq \frac{3}{4}(Row_0(T^*) + Row_1(T^*))\binom{2l-1}{l^\gamma-1} = \frac{6}{4}Row(T^*)\binom{2l-1}{l^\gamma-1} > 2^{-\epsilon n^\gamma} \binom{2l-1}{l^\gamma-1}$.

Finally we have:

$$\log \frac{3}{4}|S| > \log \left[2^{-\epsilon n^\gamma} \binom{2l-1}{l^\gamma-1} \right] \geq l^\gamma \log \left((e - o(1)) \frac{2l-1}{l^\gamma-1} \right) - \epsilon n^\gamma - \Theta(\log l)$$

$$\geq (1 - \gamma)l^\gamma \log l + l^\gamma \log(2(e - o(1))) - \epsilon \cdot (4l - 1)^\gamma - \Theta(\log l)$$

$$(\text{for large } l) \geq (1 - \gamma)l^\gamma \log l + 2.4426 \cdot l^\gamma - \epsilon \cdot (4l)^\gamma - \Theta(\log l).$$

For $\epsilon \leq \frac{1}{1000 \cdot 4^\gamma}$ we get the desired contradiction, that $\frac{3}{4}|S| > |Q|$.

(The lower-bound for $\binom{2l-1}{l^\gamma-1}$ above can be obtained using the Stirling bounds for the factorial, $\sqrt{2\pi n} \left(\frac{n}{e}\right)^n \leq n! \leq e\sqrt{n} \left(\frac{n}{e}\right)^n$, as follows:

$$\begin{aligned} \binom{2l-1}{l^\gamma-1} &\geq \frac{\sqrt{2\pi(2l-1)} \cdot (2l-1)^{2l-1}}{e^2 \sqrt{(l^\gamma-1)(2l-l^\gamma)} \cdot (l^\gamma-1)^{l^\gamma-1} \cdot (2l-l^\gamma)^{2l-l^\gamma}} \\ &= \frac{\sqrt{2\pi(2l-1)}}{e^2 \sqrt{(l^\gamma-1)(2l-l^\gamma)}} \left(\frac{2l-1}{l^\gamma-1}\right)^{l^\gamma-1} \left(\frac{2l-1}{(2l-1)-(l^\gamma-1)}\right)^{2l-l^\gamma} \\ &= \frac{\sqrt{2\pi(2l-1)}}{e^2 \sqrt{(l^\gamma-1)(2l-l^\gamma)}} \left(\frac{2l-1}{l^\gamma-1}\right)^{l^\gamma-1} \left[1 + \frac{l^\gamma-1}{2l-l^\gamma}\right]^{\frac{2l-l^\gamma}{l^\gamma-1} l^\gamma-1} \\ &\geq \frac{\sqrt{2\pi}}{e^2} \frac{1}{\sqrt{(l^\gamma-1)}} \left(\frac{2l-1}{l^\gamma-1}\right)^{l^\gamma-1} (e - o(1))^{l^\gamma-1}. \end{aligned}$$

Claim 2. $\mathbb{E}[\text{Row}_0(T)\text{Col}_0(T)(1 - \text{Bad}(T))] > \frac{1}{5}\mathbb{E}[\text{Row}_0(T)\text{Col}_0(T)]$.

Proof of Claim 2. Since $\text{Bad}(T) \leq \text{Bad}_x(T) + \text{Bad}_y(T)$, it is enough to prove that $\mathbb{E}[\text{Row}_0(T)\text{Col}_0(T)\text{Bad}_x(T)] \leq \frac{2}{5}\mathbb{E}[\text{Row}_0(T)\text{Col}_0(T)]$, with a similar statement being proved in the same fashion for $\text{Bad}_y(T)$. For each $t_2 \subseteq [n]$, with $|t_2| = 2l - 1$, we will show that the desired inequality holds when conditioned on $T_2 = t_2$, which implies that the unconditioned inequality holds. All triples T with $T_2 = t_2$ have the same value for $\text{Col}_0(T)$, so let this value be called c' . Also let $r = \mathbb{E}[\text{Row}(T) \mid T_2 = t_2]$. Now we have:

$$\begin{aligned} &\mathbb{E}[\text{Row}_0(T)\text{Col}_0(T)\text{Bad}_x(T) \mid T_2 = t_2] \\ &\leq c' \mathbb{E}[\text{Row}_0(T)\text{Bad}_x(T) \mid T_2 = t_2] \\ &\leq c' E[2 \cdot \mathbb{E}[\text{Row}(T) \mid T_2 = t_2] \cdot \text{Bad}_x(T) \mid T_2 = t_2] \\ &\leq 2c'r \mathbb{E}[\text{Bad}_x(T) \mid T_2 = t_2] \\ &\leq \frac{2}{5}c'r \text{ (by Claim 1)} \end{aligned}$$

$$\begin{aligned}
&= \frac{2}{5} c' \mathbb{E}[\text{Row}(T) \mid T_2 = t_2] \\
&= \frac{2}{5} c' \mathbb{E}[\text{Row}_0(T) \mid T_2 = t_2] \\
&= \frac{2}{5} \mathbb{E}[\text{Row}_0(T) \text{Col}_0(T) \mid T_2 = t_2]
\end{aligned}$$

The inequality between the second and the third line can be justified as follows: Recall that, as observed in the proof of Claim 1, even when considering only triples $T_2 = t_2$, the value of $\text{Row}(T)$ can differ by a factor of at most $2^{l^{1-\gamma}} - 1$. This is due to the fact that $\text{Row}(T)$ measures the probability (conditioned on T) of the same set $S = \{x \in C \mid |x| = l^\gamma, x \subseteq [n] \setminus t_2\}$, but depending on whether $i \in x$ (for a particular choice of i and hence of T), an $x \in S$ will have (conditional) probability either $\frac{1}{2} \binom{2l-1}{l^\gamma-1}^{-1}$ or $\frac{1}{2} \binom{2l-1}{l^\gamma}^{-1}$. Thus if T^* is a triple with $T_2^* = t_2$ for which $\text{Row}_0(T^*) = \max\{\text{Row}_0(T) \mid T_2 = t_2\}$, then $\text{Row}(T^*)$ must be the minimum among all values of $\text{Row}(T)$ when $T_2 = t_2$, because when $T = T^*$ the largest portion of elements of S have probability $\frac{1}{2} \binom{2l-1}{l^\gamma}^{-1}$ instead of $\frac{1}{2} \binom{2l-1}{l^\gamma-1}^{-1}$. It follows that for all T with $T_2 = t_2$ we have $\text{Row}(T) \geq \text{Row}(T^*) \geq \frac{1}{2} \text{Row}_0(T^*)$, and hence that $\mathbb{E}[2\text{Row}(T) \mid T_2 = t_2] \geq \text{Row}_0(T^*)$. On the other hand we have that for all T with $T_2 = t_2$, $\text{Row}_0(T) \leq \text{Row}_0(T^*)$, so finally we get that $\text{Row}_0(T) \leq 2\mathbb{E}[\text{Row}(T) \mid T_2 = t_2]$ for all T with $T_2 = t_2$. ■

Claim 3. For any rectangle R : $\mu(B \cap R) = \frac{1}{4} \mathbb{E}[\text{Row}_1(T) \text{Col}_1(T)]$ and $\mu(A \cap R) = \frac{3}{4} \mathbb{E}[\text{Row}_0(T) \text{Col}_0(T)]$ (with the expectation taken over all partitions T).

The proof of Claim 3 is identical to the case where μ is the distribution in Razborov's proof (see [26], chap. 4), since the relevant observations also apply to our modified distribution: 1. $\mu(B) = \frac{1}{4}$ (and hence $\mu(A) = \frac{3}{4}$), because for every fixed partition T , $i \in x$ with probability $\frac{1}{2}$ and $i \in y$ with probability $\frac{1}{2}$, indepen-

dently. 2. $i \in x$ and $i \in y$ are independent events (for the same reason). 3. For every (x, y) with $x \cap y = \emptyset$ we have that $\mathbf{Pr}[(x, y) \mid (i \notin x) \wedge (i \notin y)] = \mathbf{Pr}[(x, y) \mid ((i \notin x) \wedge (i \notin y)) \vee ((i \in x) \wedge (i \notin y)) \vee ((i \notin x) \wedge (i \in y))]$, because conditioning on either one of the two events induces the uniform distribution on the set $\{(x, y) \mid x, y \subset [n], x \cap y = \emptyset, |x| = |y| = l^\gamma\}$.

We now use claims 2 and 3 to prove the statement of the lemma:

$$\begin{aligned}
\mu(B \cap R) &= \frac{1}{4} \mathbb{E}[\text{Row}_1(T) \text{Col}_1(T)] \\
&\geq \frac{1}{4} \mathbb{E}[\text{Row}_1(T) \text{Col}_1(T) (1 - \text{Bad}(T))] \\
&\geq \frac{1}{4} \mathbb{E} \left[\left(\frac{\text{Row}_0(T)}{6} - 2^{-\epsilon n^\gamma} \right) \left(\frac{\text{Col}_0(T)}{6} - 2^{-\epsilon n^\gamma} \right) (1 - \text{Bad}(T)) \right] \quad (\text{def. of } \text{Bad}) \\
&= \frac{1}{4} \mathbb{E} \left[\left(\frac{\text{Row}_0(T) \text{Col}_0(T)}{36} - \frac{2^{-\epsilon n^\gamma}}{6} (\text{Row}_0(T) + \text{Col}_0(T)) + 2^{-2\epsilon n^\gamma} \right) (1 - \text{Bad}(T)) \right] \\
&\geq \Omega(\mathbb{E}[\text{Row}_0(T) \text{Col}_0(T) (1 - \text{Bad}(T))]) - 2^{-\epsilon n^\gamma} \quad (\text{since } \text{Row}_0(T) + \text{Col}_0(T) \leq 2) \\
&\geq \Omega(\mathbb{E}[\text{Row}_0(T) \text{Col}_0(T)]) - 2^{-\epsilon n^\gamma} \quad (\text{by Claim 2}) \\
&\geq \Omega(\mu(A \cap R)) - 2^{-\epsilon n^\gamma} \quad (\text{by Claim 3})
\end{aligned}$$

Choosing ϵ to be smaller than both the constant in front of $\mu(A \cap R)$ and $\frac{1}{1000 \cdot 4^\gamma}$ completes the proof. $\square$

Lemma 3.2.6. Let X and Y be non-empty finite sets, and let $c : X \times Y \mapsto \{0, 1\}$ be a colouring of $X \times Y$ such that a proportion $p \in (0, 1)$ of the elements of $X \times Y$ are mapped to 1, that is, such that $|c^{-1}(1)|/|X \times Y| = p$. Then for any $r \in (0, p)$ such that $r|Y| \in \mathbb{N}$, we have that for at least $\frac{p-r}{1-r}|X|$ elements $x \in X$, $|(\{x\} \times Y) \cap c^{-1}(1)| > r|Y|$.

Proof. Let us call sets of the form $\{x\} \times Y$ *rows*, and define the number $w(x) = \sum_{y \in Y} c(x, y) = |(\{x\} \times Y) \cap c^{-1}(1)|$ be the *weight* of the row $\{x\} \times Y$, for each $x \in X$. Let c be a colouring of $X \times Y$ as above, but such that the smallest possible proportion of rows have weight $> r|Y|$, and denote this proportion by q . Thus q is such that for

any colouring c' satisfying the conditions of the lemma, at least $q|X|$ elements $x \in X$ satisfy $|(\{x\} \times Y) \cap c^{-1}(1)| > r|Y|$.

We may assume that all rows with weight $\leq r|Y|$ have weight exactly $r|Y|$: If this is not the case, we may repeatedly perform the operation of changing a 0 into 1 on a row with weight $< r|Y|$, and a 1 into 0 on a row with weight $> r|Y|$, until the above statement is true. (It is easy to see that the colouring c must have rows with weight $> r|Y|$, since otherwise the overall proportion of elements mapped to 1 would be $\leq r < p$.) This operation leaves the proportion of elements that are mapped to 1 unchanged, and the minimality of the chosen colouring c guarantees that the number of rows with weight $> r|Y|$ does not decrease (and therefore remains unchanged).

Next, we may assume that all but at most one of the rows with weight $> r|Y|$ have weight exactly $|Y|$: If this is not the case, we may fix one such row, replace all zeroes with ones on all other rows of weight $> r|Y|$ (thus making their weight exactly $|Y|$), and on the fixed row change the same number of ones into zeroes so as to match the changes made on all other rows. Again the overall proportion of elements being mapped to 1 does not change, and the minimality of the colouring c guarantees that the weight of the fixed row stays $> r|Y|$.

Based on the above we now have: $p|X||Y| = q|X||Y| - \alpha|Y| + (1 - q)|X|r|Y|$, where $\alpha \in [0, 1 - r)$ is the proportion of zeroes on the one row that has weight $> r|Y|$ but not necessarily $= |Y|$. Thus we have:

$$p \leq q + (1 - q)r \iff p \leq (1 - r)q + r \iff \frac{p - r}{1 - r} \leq q.$$

□

3.3 A non-deterministic SMP protocol for the Disjointness problem

The *SMP* (*simultaneous message passing*) model of communication was originally introduced by Yao in [40], and can be described as follows: Alice and Bob are once again each given a part of the input $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$ of a Boolean function, but they are not allowed to communicate with each other in any way. Instead, the players have to prepare their messages based on the part of the input they see, and simultaneously send it to a referee, who has to either accept or reject (the referee does not talk back to the players either). This is one of the weakest communication models which has been studied. Evidently this weakness stems from the fact that the players are not permitted to interact with each other, and that neither of them can even choose his message based on the message the other player sends to the referee.

The same variations as in the standard communication model can be introduced for the SMP model. For example, the players may be given access to a source of randomness on which their messages can depend, in which case the referee is required to produce the correct output with probability $\geq \frac{2}{3}$. However, it is important to note that whenever the players are randomised in the SMP model, their sources of randomness are *private* (each is hidden from the other player), because a public coin would allow them to conspire as they choose their messages. In the *non-deterministic SMP* model, the availability of additional advice is modelled through a third player, Merlin, who knows the entire input, and whose goal it is to convince the referee to accept (he is an *untrusted* player, in the sense that he will attempt to give evidence that the referee should accept, even on 0-instances). Thus the goal of Alice and Bob becomes to help the referee catch Merlin if he cheats. In the case of non-deterministic SMP protocols, Alice and Bob are allowed to be randomised, but again their sources

of randomness are hidden from each other, and, crucially, from Merlin as well.

The presence of Merlin and the fact that the players are randomised, makes non-deterministic SMP protocols similar to the *Merlin-Arthur communication protocols* (*MA-protocols*) from the standard communication complexity framework. In an MA-protocol there is once again a Merlin who knows the entire input (x, y) , and who tries to convince Alice and Bob to accept. Merlin speaks first, sending his untrusted advice to both players, after which Alice and Bob execute a randomised protocol (using random bits which are unknown to Merlin at the time when he sends his advice). The cost of such a protocol is the sum of the length of all messages (including Merlin's advice). The complexity class MA , defined as the class of all Boolean functions which have MA-protocols of poly-logarithmic cost, is such that $NP^{cc}(= \Sigma_1^{cc}) \subseteq MA \subseteq \Sigma_2^{cc}$.

In [1], Aaronson and Wigderson give an MA-protocol with cost $O(\sqrt{n} \log n)$ for both Disjointness and Inner Product modulo 2 (their protocol is in fact *one-way*, meaning that both Merlin and Bob send their messages only to Alice, who does not communicate, and either accepts or rejects based on the messages). This upper bound essentially matches the lower bound of $\Omega(\sqrt{n})$, due to Klauck [20], for the complexity of Disjointness under MA-protocols.

In most versions of the SMP model, the Disjointness problem is trivially hard (has complexity $\Omega(n)$). However, below we give a non-deterministic SMP protocol with total cost $O(n^{2/3} \log n)$, which we obtain by adapting the above-mentioned protocol of Aaronson and Wigderson. The increase in the cost is due to the need to avoid introducing shared randomness between the players.

For the lower bound we only have the $\Omega(\sqrt{n})$ bound which holds trivially due to the above-mentioned lower bound of Klauck for the MA-complexity, because every non-deterministic SMP protocol can be converted into an MA-protocol by having Alice and Bob each play the role of the referee for the other.

Theorem 3.3.1. There is a non-deterministic SMP protocol with cost $O(n^{2/3} \log n)$ for DISJ.

Proof. As in the protocol of Aaronson and Wigderson, we divide the set $[n]$ into n^α blocks of size $n^{1-\alpha}$, with $\alpha \geq 1/2$, and write the inputs x and y respectively as functions $a, b : [n^\alpha] \times [n^{1-\alpha}] \rightarrow \{0, 1\}$ such that $x_{n^{1-\alpha}(i-1)+j} = a(i, j)$ and $y_{n^{1-\alpha}(i-1)+j} = b(i, j)$. The problem then is to accept with high probability if

$$\sum_{i \in [n^\alpha]} \sum_{j \in [n^{1-\alpha}]} a(i, j)b(i, j) = 0,$$

and reject with high probability otherwise. Choose a prime $q \in (n, 2n]$ (that a prime always exists in this range is a well-known theorem of Chebyshev). The functions a and b can then be extended uniquely to two-variable polynomials $\tilde{a}, \tilde{b} : \mathbb{F}_q^2 \rightarrow \mathbb{F}_q$ which have total degree $\leq O(n^\alpha)$ (by “extended” we mean that the corresponding polynomial coincides with the function on the function’s domain). We also define the $O(n^\alpha)$ -degree polynomial $s(i) = \sum_{j \in [n^{1-\alpha}]} \tilde{a}(i, j)\tilde{b}(i, j) \pmod q$.

For the protocol, Merlin sends a candidate polynomial s' of degree equal to that of s to the referee. Now if Alice and Bob could each send for some randomly chosen value r the values $\tilde{a}(r, j)$ and $\tilde{b}(r, j)$ (respectively) for all $j \in [n^{1-\alpha}]$, as well as the value r , then the referee could compute $s(r)$ and test it against $s'(r)$: If the size of the set from which r is chosen is sufficiently large, and $s \neq s'$, then with sufficiently high probability $s(r)$ would differ from $s'(r)$, and the referee could reject (otherwise he would assume that $s' = s$ and use s' to decide whether to accept or reject). Unfortunately, the random choice of r , which is necessary in order to keep it hidden from Merlin, also keeps it hidden from the player who is not making the choice. Thus, if Bob chooses r , the only way for Alice to know which block of values $\tilde{a}(i, j)$ to send, is by either communicating or having shared randomness with Bob.

This issue can of course be overcome by having each player send the blocks of $n^{1-\alpha}$

values $\tilde{a}(r, j)$ (resp. $\tilde{b}(r, j)$), together with r , for sufficiently many values of r as to have a collision with high probability. In order to do this efficiently, we need to make the set from which r is chosen as small as possible, while keeping it large enough to be able to distinguish s from s' reliably if they differ. For the Schwartz-Zippel Lemma [31], it suffices for the set S from which r is chosen to be of size, say, $10 \deg(s) = O(n^\alpha)$, while Alice and Bob each choose $100\sqrt{|S|} = O(n^{\alpha/2})$ values of r .

What remains is to analyse the cost of this protocol. Merlin uses $O(n^\alpha \log n)$ bits to send the coefficients of the polynomial s' . Alice and Bob send $O(n^{\alpha/2})$ blocks of $O(n^{1-\alpha})$ values in $[q]$ each, which costs $O(n^{1-\alpha/2} \log n)$ bits of communication. Setting $\alpha = 2/3$, we get the statement of the theorem. $\square$

Chapter 4

The UPP^{cc} vs. $PSPACE^{cc}$ Question

4.1 Overview

In this chapter we present a new approach to the long-standing conjecture of Babai *et al.* [3] regarding the relation between the complexity classes UPP^{cc} and $PSPACE^{cc}$, and the progress we have made in this direction.

In Section 4.2 we start with a brief account of results related to the sign-rank and margin complexity, the significance of the two concepts in Communication Complexity and Computational Learning Theory, and their relationship to the complexity class UPP^{cc} . We then state our first conjecture, 4.2.1, and prove one of our main theorems in this chapter:

Theorem 4.2.2. If Conjecture 4.2.1 holds, then $UPP^{cc} \subset PSPACE^{cc}$.

A second, stronger geometric conjecture, 4.2.3, states that for any arrangement of hyperplanes there is an arrangement with the same dimensions and signature set, every region of which is sufficiently “large” (in the sense that it contains a unit vector which is sufficiently far away from all hyperplanes). The remainder of the chapter will

be devoted to studying this conjecture.

We also prove an important technical result in this section: Lemma 4.2.7 gives a lower bound for the size of the regions of a hyperplane arrangement in terms of the minors of the matrix whose rows are the normal vectors of the arrangement hyperplanes. As mentioned in Section 2.2, this matrix can also be interpreted as the basis of a linear subspace of a Euclidean space, which allows us to view the conjecture in terms of linear subspaces. Thus our goal becomes to show that any given subspace can be “balanced”, in the sense that the minors of a matrix whose columns form an orthonormal basis can all be made sufficiently large in absolute value, without changing the signature set.

Our approach to proving that a balanced version of every subspace exists is to try to show that we can perform a series of improvements on an arbitrary subspace, and obtain a balanced version. Naturally we will want the subspace at every intermediate step to have the same signature set as the original one.

In Section 4.3 we begin our study of operations on subspaces which leave their signature set unchanged. The basic operation we consider consists of replacing a chosen vector v from a d -dimensional subspace with a target vector v' outside of it, leaving $d - 1$ dimensions of the subspace unchanged. We define for every non-zero vector v an arrangement of hyperplanes in $\mathbb{R}^m$, called the *extension arrangement of v* , whose hyperplanes (called *extension hyperplanes*) are each associated with an $(m - d)$ -dimensional canonical subspace of $\mathbb{R}^m$. We show that:

- if v' is in the same region of the extension arrangement as v , the resulting subspace has the same signature set as the original one (Proposition 4.3.5);

- if v' is in one of the extension hyperplanes of v , the resulting subspace is not in general position (Lemma 4.3.3);

- if v' lies on the opposite side of an extension hyperplane of v , the signature set

changes (Remark 4.3.7).

Thus the extension hyperplanes of v delimit the set of vectors with which v can safely be replaced.

In Section 4.4 we make the connection between the distance of vectors from their respective extension hyperplanes and the d -minors of any matrix whose columns form an orthonormal basis of the subspace. In Theorem 4.4.2 we give an exact formula for the distance of a unit vector from the hyperplanes of its extension arrangement. This is then used for the main theorem:

Theorem 4.4.5 Let D be a d -dimensional subspace in general position in $\mathbb{R}^m$, Let $\epsilon > 0$ be given.

If there is an orthonormal basis $\{v_1, \dots, v_d\}$ of D , such that every basis vector has distance at least ϵ from each of its extension hyperplanes, then every d -minor of the matrix $A = (v_1, \dots, v_d)$ is, in absolute value, $\geq \left(\frac{\epsilon}{(2\sqrt{2})^{d-1}} \right)^d$.

The strength of this theorem lies in the fact that it is enough to look at d vectors of the subspace in order to deduce how balanced it is overall. Finally, we show that the conditions of the theorem can be relaxed in the sense that only a strict subset of the extension hyperplanes need to be considered (regardless of the chosen vector) and give bounds on the size of this set in Proposition 4.4.11 and the remark preceding it.

Lacking the full proof that the conditions of Theorem 4.4.5 can be achieved with sufficiently large ϵ for any given subspace by using only safe vector replacement operations, we give numerical evidence for this in Section 4.5, where we describe an algorithm for balancing subspaces, and discuss its performance.

4.2 A new approach to the UPP^{cc} vs. $PSPACE^{cc}$ question

Unbounded-error probabilistic protocols (see Definition 2.1.8) were first defined by Paturi and Simon in [28]. Such protocols are more powerful than almost all other communication protocols that are generally considered, in the sense that the complexity of a Boolean function f under unbounded-error probabilistic protocols, denoted by $U(f)$, is lower than its complexity under most other kinds of protocol. Intuitively, the reason for this fact is that, for every input (x, y) , the protocol is required to predict $f(x, y)$ only slightly better than a fair coin flip.

In the same paper, Paturi and Simon also introduced the notion of *sign-rank* of a matrix with no entries equal to 0 (see Definition 2.1.9) and proved that the logarithm of the sign-rank of $M_f^- \in \mathcal{M}_{2^n}(\{-1, +1\})$ is essentially $U(f)$. Virtually all results regarding the unbounded-error probabilistic communication complexity of functions are based only on the properties of the sign-rank, and do not rely on the details of the protocols.

Much of the previous work on the sign-rank focused on developing methods for proving lower bounds for the sign-rank of matrices. We list some important results here: Early on, Alon, Frankl and Rödl [2] proved that almost all sign matrices (matrices whose entries are -1 or $+1$) of size $m \times m$ have sign-rank $\Omega(m)$. However, as is typical for the probabilistic method, this did not yield any concrete example of such a function. Much later, Forster [12] gave the first lower bound for specific matrices, by showing that Hadamard matrices (and more generally, matrices with small spectral norm) have large sign-rank, thus proving that IP_2 (Inner Product mod 2) is not in UPP^{cc} , the class of functions which have unbounded-error probabilistic protocols of cost $\text{polylog}(n)$. Forster's use of matrix-analytical tools to lower-bound the sign-rank constituted a

major breakthrough in the study of this topic, and his methods were subsequently extended by various authors (see [13], [14] and [21]). Next, Sherstov [34] proved nearly tight bounds on $U(f)$ for functions satisfying $f(x, y) = D(|x \wedge y|)$ for some $D : [n] \rightarrow \{-1, +1\}$. Finally, Razborov and Sherstov [30] proved that the complexity class AC^0 (the class of functions which are computable by constant-depth circuits using only $\wedge$ and $\vee$ gates with unbounded fan-in, and either variables or negations thereof as inputs) is not contained in UPP^{cc} , by exhibiting a specific function in AC^0 for which $U(f) = \Omega(n^{1/3})$. With this they also proved that $\Sigma_2^{cc} \not\subseteq UPP^{cc}$, thereby answering a question posed by Babai *et al.* in [3].

Another important result is related to the class PP^{cc} , a sister-class of UPP^{cc} , first defined in [3]. This is the class of functions which have randomised protocols of cost $\text{polylog}(n)$ such that on each input they correctly compute $f(x, y)$ with probability at least $\frac{1}{2} + 2^{-\text{polylog}(n)}$. In other words, the error of such protocols is moderately bounded away from $\frac{1}{2}$. Both PP^{cc} and UPP^{cc} are Communication Complexity versions of the Turing complexity class PP . Obviously $PP^{cc} \subseteq UPP^{cc}$, but the question of whether the two classes coincide remained open for over two decades, until Buhrman, Vereshchagin and de Wolf [8], and simultaneously Sherstov [33], proved that the inclusion is strict.

Although the unbounded-error probabilistic communication model does not arise from practical considerations, and is not the most reasonable model of probabilistic communication, the notion of sign-rank has also found applications in another field of Computer Science, namely Computational Learning Theory. In the PAC (*probably approximately correct*) learning model defined by Valiant in [37], a learner is tasked with “learning” a *concept class* $\mathcal{C}$ of functions $f : \{-1, +1\}^n \rightarrow \{-1, +1\}$, in the following sense: For any $f \in \mathcal{C}$, given a set of training examples $(z_1, f(z_1)), \dots, (z_m, f(z_m))$, where the values $z_i \in \{-1, +1\}^n$ are chosen according to a distribution μ which

is unknown to the learner, he must compute as efficiently as possible a *hypothesis* $h : \{-1, +1\}^n \rightarrow \{-1, +1\}$ which is close to f under μ . A general method for designing algorithms for PAC learning, which for some concept classes produces very efficient algorithms, works as follows: Let $r \in \mathbb{N}$ be such that there exists a function $\phi : \{-1, +1\}^n \rightarrow \mathbb{R}^r$ such that for every $f \in \mathcal{C}$ there is a vector $a_f \in \mathbb{R}^r$ for which $f(z) = \text{sgn}\langle a, \phi(z) \rangle$. In other words, r is chosen in such a way that there is an embedding ϕ of $\{-1, +1\}^n$ into $\mathbb{R}^r$ such that for every function in $\mathcal{C}$ there is a hyperplane of $\mathbb{R}^r$ which separates the points in $\phi(\{-1, +1\}^n)$ in the same way as the function. (That such an r must exist is clear because there is a trivial choice of ϕ when $r = 2^n$.) With knowledge of the set $\{a_f \mid f \in \mathcal{C}\}$, the learner can choose from this set the vector a which best matches the training examples, and return h defined by $h(z) = \text{sgn}\langle a, \phi(z) \rangle$. Intuitively it is clear that finding the right vector is easier the smaller r is, and indeed there are algorithms for this problem whose running time is polynomial in n and r (see, for example, [5]). The smallest value of r for which the above-mentioned function ϕ exists is called the *dimension complexity* of the concept class $\mathcal{C}$.

If, in the above setting, we define the $(|\mathcal{C}| \times 2^n)$ -matrix M such that $M_{f,z} = f(z)$ for all $f \in \mathcal{C}$ and $z \in \{-1, +1\}^n$, the $(r \times 2^n)$ -matrix B such that $B_{\cdot,z} = \phi(z)$ for all $z \in \{-1, +1\}^n$, and the $(|\mathcal{C}| \times r)$ -matrix A such that $A_f = a_f$ for all $f \in \mathcal{C}$, then it is clear that we have that $\text{sgn}\langle A_f, B_{\cdot,z} \rangle = M_{f,z}$ for all $f \in \mathcal{C}$ and $z \in \{-1, +1\}^n$. It is then also clear that the dimension complexity of $\mathcal{C}$ is precisely the sign-rank of the matrix M . Thus we have that the complexity of certain algorithms for PAC learning is essentially determined by the sign-rank of the matrix associated with the concept class to be learned.

For every condition $\text{sgn}\langle A_f, B_{\cdot,z} \rangle = M_{f,z}$ above, the real value $\langle A_f, B_{\cdot,z} \rangle$ is allowed to be arbitrarily close to 0, as long as it has the right sign. In other words, any point

$\phi(z)$ is allowed to be arbitrarily close to a hyperplane associated with a function, as long as it lies on the right side of the hyperplane. A different, but related, notion of complexity of a sign matrix is obtained if, instead of requiring that the dimension r be as small as possible, one requires that the smallest distance of any point in $\phi(\{-1, +1\}^n)$ from any one of the hyperplanes with normal vectors $\{a_f \mid f \in \mathcal{C}\}$ be as large as possible. In order for this requirement to be meaningful, both the normal vectors of the hyperplanes and the vectors that points in the domain are mapped to must be taken as unit vectors, otherwise any $|\langle a_f, \phi(z) \rangle|$ can be made arbitrarily large, even though the direction line of $\phi(z)$ may be very close to the hyperplane with normal vector a_f . Reindexing the rows of M by $[k]$ and the columns by $[l]$, where $k = |\mathcal{C}|$ and $l = 2^n$, the measure of M that we get is

$$\sup \left\{ \min_{i,j} \frac{|\langle A_i, B_{\cdot,j} \rangle|}{\|A_i\| \|B_{\cdot,j}\|} \mid A, B \text{ are such that } \forall i, j : \text{sgn}(AB)_{i,j} = M_{i,j} \right\}.$$

This is the *margin* of the sign matrix M , denoted $m(M)$ (see Definition 2.1.11), a concept first introduced and studied in the context of Computational Learning Theory ([4]). For convenience, the *margin complexity* of a sign matrix M is defined as $mc(M) = m(M)^{-1}$.

Unlike the sign-rank, margin is fairly well understood as a complexity measure of sign matrices. Ben-David *et al.* [4] showed that almost all sign matrices of size $m \times m$ have margin complexity $\Omega(\sqrt{\frac{m}{\log m}})$, and investigated the relation between this notion and the *VC-dimension*. In [12], Forster proved lower bounds for the margin complexity of certain functions, using the same method as for the sign-rank. Linial and Shraibman investigated the relation between margin complexity and various *factorisation norms* in [22], where they also proved that for any sign matrix M of size $m_1 \times m_2$, $\text{sign-rank}(M) \leq O(mc(M)^2 \log(m_1 + m_2))$. Finally, Linial *et al.* [21] showed that the margin complexity of a sign-matrix is, up to a constant factor, equal to the

discrepancy of the matrix, a well-studied measure which has been used to tightly lower-bound, among others, the randomised complexity of certain functions (see, for example, [29] or [3]).

Although both sign-rank and margin complexity can be defined in terms of arrangements of hyperplanes, for each notion only a single parameter is optimised, while no restrictions are placed on the other. We are interested in restrictions which can be applied *simultaneously* to both the dimension of the space in which the arrangement resides, and the margin. Specifically, we wish to show that if a sign matrix can be realised by an arrangement of hyperplanes (and a set of points) in a space of small dimension, then it can also be realised by an arrangement of hyperplanes (and a set of points) in the same space *and* with a margin that is not too small. We make this precise in the form of a conjecture.

Conjecture 4.2.1. There exists a constant $c \geq 1$ such that for all $m \geq d \geq 2$, we have that if $M \in \mathcal{M}_m(\mathbb{R} \setminus \{0\})$ has sign-rank d , then there exist matrices $A \in \mathcal{M}_{m,d}(\mathbb{R})$ and $B \in \mathcal{M}_{d,m}(\mathbb{R})$, such that the pair (A, B) realises the matrix M with margin $\Omega(m^{-d^c})$.

To check whether the conjecture is plausible, consider the following: In order to prove that $PP^{cc} \neq UPP^{cc}$, both Buhrman *et al.* [8] and Sherstov [33] each exhibit an explicit function f such that M_f^- has sign-rank $d = 2^{O(\log n)}$, but such that for any two matrices $A \in \mathcal{M}_{m,r}(\mathbb{R})$ and $B \in \mathcal{M}_{r,m}(\mathbb{R})$, where $m = 2^n$ and $r = 2^{\text{polylog}(n)} \geq d$, the pair (A, B) can not realise M_f^- with margin $\geq 2^{-(\log n)^{c'}}$, for any constant $c' > 0$. We wish to show that for $r = d$, the “inner” dimension r can be maintained while at the same time achieving a margin $\geq m^{-d^c} = 2^{-2^{O(\log n)} \log m} = 2^{-2^{O(\log n)}}$. Since $2^{O(\log n)} = \omega((\log n)^{c'})$ for any constant $c' > 0$, our conjecture does not come close to contradicting the result separating PP^{cc} from UPP^{cc} .

Despite the modest bound on the margin, the above conjecture would nevertheless be sufficient to solve Babai *et al.*'s long-standing open question regarding the relation between the classes UPP^{cc} and $PSPACE^{cc}$ (see Definition 2.1.3).

Theorem 4.2.2. If Conjecture 4.2.1 holds, then $UPP^{cc} \subset PSPACE^{cc}$.

Proof. Let $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ be a function in UPP^{cc} . Then M_f^- has sign-rank $d = 2^{(\log n)^g}$, for some constant $g \geq 1$. Setting $m = 2^n$, it follows from the statement of Conjecture 4.2.1 that there is a row-normalised $(m \times d)$ -matrix A and a column-normalised $(d \times m)$ -matrix B such that $(AB)_{i,j}(M_f^-)_{i,j} > 0$ for all i and j , and $|(AB)_{i,j}| \geq \Omega(m^{-d^c})$. From this inequality it follows that, for some sufficiently large constant $c' \geq 1$, we may discard all but the first $c' \cdot \log(m^{d^c}) = c'd^cn = 2^{c \log(d) + \log(n+c')} = 2^{c(\log n)^g + \log(n+c')} \leq 2^{O((\log n)^g)} = 2^{\text{polylog}(n)}$ digits of the binary representation of every entry of A and B , with the resulting pair of matrices still satisfying $(AB)_{i,j}(M_f^-)_{i,j} > 0$ for all i and j . Since Alice and Bob's goal is to compute $(M_f^-)_{x,y}$, their problem now becomes to compute the inner product of row x of A and column y of B , and to decide whether the result is greater or less than 0. In other words, they need to compute the sum of $d = 2^{\text{polylog}(n)}$ products of two numbers of $2^{O((\log n)^g)} = 2^{\text{polylog}(n)}$ bits each. As stated in Fact 2.1.14, the summation of k numbers of k bits each can be done by $(\log k)^2$ -depth circuits using only the 'V' and '∧' gates with fan-in 2 and with variables or negations of variables as inputs, and that the multiplication of two numbers of k bits, as well as the comparison of a k -bit number with 0, can be done by $\log(k)$ -depth circuits of the same type. Thus the problem of computing $(M_f^-)_{x,y} = f(x, y)$ can be solved by a circuit of depth $\text{polylog}(2^{\text{polylog}(n)}) = \text{polylog}(n)$ using only the above-mentioned gates, whose inputs are the binary representations of the numbers on row x of A and column y of B . (Note that Alice and Bob can be assumed to know the matrices A and B (with truncated

entries), and it is therefore trivial for Alice to obtain the bits of the numbers on the row x of A , and likewise for Bob with column y of B , by simply looking up the respective vectors. They can then plug the numbers into the circuit as inputs.) By the characterisation of $PSPACE^{cc}$ from Fact 2.1.15, this places f in $PSPACE^{cc}$, and since the function $f \in UPP^{cc}$ was arbitrary, the desired inclusion between the classes follows. $\square$

Recalling Definition 2.2.3, it is clear that if $d = \text{sign-rank}(M)$, $A \in \mathcal{M}_{m,d}(\mathbb{R})$ and $B \in \mathcal{M}_{d,m}(\mathbb{R})$ are such that the pair (A, B) realises $M \in \mathcal{M}_m(\{-1, +1\})$, then the column vectors of M are signatures of the regions of $\mathcal{A}(A)$ which are inhabited by the points whose coordinates are column vectors of B . It follows that for any matrix A whose row vectors define an arrangement whose signature set contains all column vectors of M , we can produce a new matrix B such that (A, B) realises M , and in fact we may assume that each column vector of B is the “best” unit vector in its region. Therefore we may discard the matrix B entirely, and work only with the matrix A , with the restriction that certain signatures of regions of $\mathcal{A}(A)$ must be preserved.

Thus, in order to prove Conjecture 4.2.1, we need to show that for any matrix $M \in \mathcal{M}_m(\{-1, +1\})$ such that there exists a matrix $A \in \mathcal{M}_{m,d}(\mathbb{R})$ for which the signature set of $\mathcal{A}(A)$ contains all column vectors of M , there exists a matrix $A' \in \mathcal{M}_{m,d}(\mathbb{R})$ for which the signature set of $\mathcal{A}(A')$ also contains the column vectors of M , and such that the regions with these signatures are sufficiently “large”. In this case, a region being large means that the radius of the largest open ball centred at a unit vector and contained entirely in the region, is large; evidently this means that the smallest distance from the centre of such a ball to any hyperplane of $\mathcal{A}(A')$ is equal to the radius.

Another important point is that, since we take M as a square matrix, only m

signatures in the signature set of $\mathcal{A}(A)$ are fixed, and only m regions need to be large, whereas from Fact 2.2.7 we know that the number of regions of an arrangement of m hyperplanes in $\mathbb{R}^d$ can be roughly m^{d-1} . On the other hand, the m regions can be arbitrary, and we can make no assumptions about their relative positions, or about the way in which enlarging one might affect the others. We will therefore attempt to show that *all* regions of a hyperplane arrangement can simultaneously be made sufficiently large.

Conjecture 4.2.3. There is a constant $c \geq 1$ such that for all $m \geq d \geq 2$, we have that if $A \in \mathcal{M}_{m,d}(\mathbb{R})$, then there exists a matrix $A' \in \mathcal{M}_{m,d}(\mathbb{R})$ such that the signature set of $\mathcal{A}(A')$ contains that of $\mathcal{A}(A)$, and such that every region of $\mathcal{A}(A')$ contains a unit vector which has Euclidean distance $\geq \Omega(m^{-d^c})$ from every hyperplane of the arrangement.

Using Fact 2.2.11, we will assume from now on that the matrix $A \in \mathcal{M}_{m,d}(\mathbb{R})$ in the statement of the above conjecture is such that $\mathcal{A}(A)$ is in general position.

While the true (tight) lower bound on the margin in the statement of Conjecture 4.2.1 may in fact be much larger (it is not even clear whether the exponent of m must depend on d), for Conjecture 4.2.3 it can be shown that the true lower bound can not be much larger than m^{-d} (from Fact 2.2.7 we have that the number of regions of an arrangement of m hyperplanes in general position in $\mathbb{R}^d$ can be roughly $m^{-(d-1)}$, whereas the surface of the unit sphere in $\mathbb{R}^d$ is only exponential in d (with constant base)).

As indicated above, we consider the “size” of a region as the radius of the largest open ball which is contained entirely in the region and centred at a unit vector. This centre is obviously the unit vector which has the largest distance from any hyper-

plane delimiting the region, and hence from all hyperplanes of the arrangement. It is intuitively clear that in $\mathbb{R}^d$ the open ball will be constricted by at least d of the hyperplanes which delimit the region; this fact is easy to visualise when $m = 3$ and $d = 2$. The following lemma shows that this intuition holds for higher dimensions.

Lemma 4.2.4. Let $d, m \in \mathbb{N}$ be such that $m \geq d \geq 2$, and let $A \in \mathcal{M}_{m,d}(\mathbb{R})$ be such that not all d -minors of A are 0. Let $H_1, \dots, H_m \subset \mathbb{R}^d$ be the hyperplanes of $\mathcal{A} = \mathcal{A}(A)$. Let R be a region of $\mathcal{A}$.

Then there exists a unit vector z in R , which has the same distance from at least d of the hyperplanes of $\mathcal{A}$, such that this distance is the smallest distance from this point to any of the hyperplanes. In other words, $\exists z \in R, |z| = 1 : \exists I \subseteq [m], |I| \geq d$, such that $\text{dist}(z, H_j) = \min\{\text{dist}(z, H_i) \mid i \in [m]\}$ for all $j \in I$.

Proof. Let $n_1, \dots, n_m \in \mathbb{R}^d$ be the normalised row vectors of A .

If $d = 2$ then every region is delimited by two lines, and we may choose $z \in R$ to be on the bisecting line of the angle that is formed by the two lines delimiting R .

For $d \geq 3$, the proof is by induction on d (with $d = 2$ being the base case), and will proceed as follows: First, we prove the existence of a unit vector x_1 which is contained in the intersection of at least $d - 1$ hyperplanes of the arrangement, whose intersection is a line, such that x_1 and R are contained in the same half-space with respect to each of the remaining hyperplanes. Next, we show that a second point $x_2 \in R$ exists, for which the distance from at least $d - 1$ of the hyperplanes in whose intersection x_1 lies is the same (this is the only part where the induction hypothesis is used). The plane (not hyperplane) P which is generated by x_1 and x_2 , will be shown to have the property that every point in P is equidistant from the above-mentioned $d - 1$ (or more) hyperplanes. Finally, using a continuity argument, we prove that there is a point $z \in P \cap R$ with the desired properties.

Claim 1. There exists a point x_1 in the intersection of at least $d - 1$ of the hyperplanes of the arrangement, whose intersection is a line, such that x_1 and R lie in the same half-space with respect to each of the hyperplanes which do not contain x_1 .

Proof of claim. Let y_0 be some point in R , and let k_1 be such that $\text{dist}(y_0, H_{k_1}) = \min\{\text{dist}(y_0, H_k) \mid k \in [m]\}$. Let $y_1 \in H_{k_1}$ be a point for which $\text{dist}(y_0, y_1) = \text{dist}(y_0, H_{k_1})$ (that such a point exists and is unique follows from the completeness of the metric space $\mathbb{R}^d$ and the theory of Hilbert spaces). We can repeat this construction: at step l ($2 \leq l < d$), we have the pairwise distinct indices $k_1, k_2, \dots, k_{l-1}$ and the (not necessarily pairwise distinct) points $y_1, \dots, y_{l-1}$, such that for each $i \in [l-1]$ we have that $y_i \in H_{k_1} \cap \dots \cap H_{k_i}$. To choose k_l and y_l , note first that there must be a $k \neq k_1, \dots, k_{l-1}$ such that $(H_{k_1} \cap \dots \cap H_{k_{l-1}} \cap H_k) \neq (H_{k_1} \cap \dots \cap H_{k_{l-1}})$, otherwise it must hold that $H_{k_1} \cap \dots \cap H_{k_{l-1}} = \bigcap_{j \in [m]} H_j$. We know that the latter set must be $\{0\}$, since d of the hyperplanes in the intersection form an arrangement that is in general position. But $\dim(H_{k_1} \cap \dots \cap H_{k_{l-1}}) \geq d - (l-1) \geq 2$, which is a contradiction. Now let $k_l \in [m]$ be such that $(H_{k_1} \cap \dots \cap H_{k_{l-1}} \cap H_{k_l}) \neq (H_{k_1} \cap \dots \cap H_{k_{l-1}})$ and $\text{dist}(y_{l-1}, H_{k_1} \cap \dots \cap H_{k_{l-1}} \cap H_{k_l}) = \min\{\text{dist}(y_{l-1}, H_{k_1} \cap \dots \cap H_{k_{l-1}} \cap H_k) \mid (H_{k_1} \cap \dots \cap H_{k_{l-1}} \cap H_k) \neq (H_{k_1} \cap \dots \cap H_{k_{l-1}})\}$. Let $y_l \in H_{k_1} \cap \dots \cap H_{k_{l-1}} \cap H_{k_l}$ be the point such that $\text{dist}(y_{l-1}, y_l) = \text{dist}(y_{l-1}, H_{k_1} \cap \dots \cap H_{k_{l-1}} \cap H_{k_l})$.

Next we show that $y_{d-1} \neq 0$. Assume otherwise, and let $i \in [d-1]$ be the smallest value for which $y_i = 0$. Because of the way y_i is chosen, we must have that $\text{dist}(y_{i-1}, 0) = \min\{\text{dist}(y_{i-1}, H_{k_1} \cap \dots \cap H_{k_{i-1}} \cap H_k) \mid k \neq k_1, \dots, k_{i-1}\}$, which implies that y_{i-1} (which is $\in H_{k_1} \cap \dots \cap H_{k_{i-1}}$) is orthogonal to $H_{k_1} \cap \dots \cap H_{k_{i-1}} \cap H_k$ for all $k \neq k_1, \dots, k_{i-1}$. Since $\dim(H_{k_1} \cap \dots \cap H_{k_{i-1}} \cap H_k) \geq \dim(H_{k_1} \cap \dots \cap H_{k_{i-1}}) - 1$, it follows that for all $k \neq k_1, \dots, k_{i-1}$ the subspace $H_{k_1} \cap \dots \cap H_{k_{i-1}} \cap H_k$ is the

same (namely the orthogonal complement of y_{i-1} in $H_{k_1} \cap \dots \cap H_{k_{i-1}}$), and hence that $\dim(\bigcap_{j \in [m]} H_j) = \dim(H_{k_1} \cap \dots \cap H_{k_{i-1}} \cap H_k) = d - i \geq 1$, contradicting the requirement that some d arrangement hyperplanes form an arrangement in general position.

Now y_{d-1} is contained in the intersection of $d - 1$ hyperplanes whose intersection is a line: For each $i \in [d - 1]$, $\dim(H_{k_1} \cap \dots \cap H_{k_i}) = \dim(H_{k_1} \cap \dots \cap H_{k_{i-1}}) - 1$, because k_i was chosen so that the dimensions of the two subspaces are not equal. On the other hand we can not rule out the possibility of this line being contained in other hyperplanes of the arrangement as well, due to the fact that the arrangement is not necessarily in general position. Let K be the set of indices of hyperplanes which contain y_{d-1} . Since the sequence of points (y_i) never leaves a hyperplane it arrives at, and because every point in the sequence is defined by the minimum distance from the previous point to the hyperplanes that were not yet intersected by the sequence (so that it can not “cross” any of them), it follows that if y_{d-1} is in a half-space associated with one of the H_j ’s, then this must necessarily be the same half-space as the one in which y_0 lies. It follows that y_{d-1} is in the same half-space as R with respect to H_j for each $j \notin K$ (and y_{d-1} and R are trivially not in different half-spaces with respect to any H_k when $k \in K$, since then $y_{d-1} \in H_k$). We now set $x_1 = \frac{y_{d-1}}{\|y_{d-1}\|}$ to get the desired unit vector.

Claim 2. There is a point $x_2 \in R$ which is at equal distance from at least $d - 1$ of the hyperplanes in whose intersection x_1 lies, and has at least that distance from the other hyperplanes which contain x_1 .

Proof of claim. As before, let K be the set of indices of hyperplanes in whose intersection x_1 lies (thus $|K| \geq d - 1$), and let $L \subseteq \mathbb{R}^d$ be the subspace spanned by their normal vectors. Since by Claim 1 this intersection is a line, it follows that

$\dim L = d - 1$. Identifying L with $\mathbb{R}^{d-1}$, consider the arrangement of hyperplanes in L whose normal vectors are $\{n_k \mid k \in K\}$ (it is easy to see that these hyperplanes are just the sets $H_k \cap L$ for $k \in K$). This arrangement must have a region $R' \subset L$ whose points lie in the same half-space (of $\mathbb{R}^d$) as R with respect to H_k for each $k \in K$: To see this, consider the orthogonal projection of any vector $x \in R$ onto L ; since this is just $x - \alpha x_1$ for some $\alpha \in \mathbb{R}$, it is clear that $\text{sgn}\langle n_k, x \rangle = \text{sgn}\langle n_k, x - \alpha x_1 \rangle$ for all $k \in K$, because $x_1 \perp n_k$ for all $k \in K$. Thus we may use the induction hypothesis for the above arrangement of $|K|$ hyperplanes in $L (= \mathbb{R}^{d-1})$ and the region R' , to obtain a unit vector $y \in R'$ that is equidistant with respect to at least $d - 1$ of the hyperplanes with indices in K , and has a greater distance with respect to the remaining hyperplanes (of the arrangement in L). Note that since $x_1 \perp L$, we have $\text{dist}(y, H_k) = \text{dist}(y, H_k \cap L)$ for all $k \in K$, since the H_k 's (when $k \in K$) can be obtained from the hyperplanes in L by extending them along the perpendicular vector x_1 .

We have shown that y is in the same half-space as R with respect to H_k for all $k \in K$, but it is not necessarily in the same half-space as R with respect to H_j when $j \notin K$. Consider the function $x(t) : [0, 1] \rightarrow \mathbb{R}^d$ defined by $x(t) = tx_1 + (1 - t)y$, and let $f_1 : [0, 1] \rightarrow \mathbb{R}^m$ be the function defined as $f_1(t) = (\langle x(t), n_1 \rangle, \dots, \langle x(t), n_m \rangle)^T$. Since all functions involved in this definition are continuous, f_1 must be as well (with respect to the standard topologies on $[0, 1]$ and $\mathbb{R}^m$), and we have that $\exists \epsilon > 0$ such that the l -th coordinate of $f_1(1 - \epsilon)$ has the same sign as the l -th coordinate of $f_1(1)$ for all $l \notin K$, and such that for any $k \in K$ the k -th coordinate of $f_1(1 - \epsilon)$ has the same sign as $\langle x, n_k \rangle$ for any point $x \in R$ (because for each $k \in K$ and $t \in [0, 1]$ we have that $\langle x(t), n_k \rangle = t\langle x_1, n_k \rangle + (1 - t)\langle y, n_k \rangle = (1 - t)\langle y, n_k \rangle$, and we have already established that $\text{sgn}\langle y, n_k \rangle = \text{sgn}\langle x, n_k \rangle$ for any $x \in R$). In other words, since $x(1) = x_1$, we have that $x(1 - \epsilon)$ is in the same half-space as x_1 with respect to all H_l for which $l \notin K$, and is in the same half-space as R with respect all H_k with $k \in K$, meaning that

$x(1 - \epsilon)$ is in R . Let $x_2 = x(1 - \epsilon)$.

Claim 3. Every point in $P = \langle \{x_1, x_2\} \rangle$ is at equal distance from the at least $d - 1$ hyperplanes with respect to which x_2 is equidistant.

Proof of claim. Let Q be the set of indices $q \in K$ such that $\text{dist}(x_2, H_q) = \min\{\text{dist}(x_2, H_k) \mid k \in K\}$ (thus $|Q| \geq d - 1$). Let $x = \alpha x_1 + \beta x_2 \in P$, $\alpha, \beta \in \mathbb{R}$, be given. Then, if $p, q \in Q$, $p \neq q$, we have $|\langle x, n_p \rangle| = |\langle \alpha x_1 + \beta x_2, n_p \rangle| = |\beta \langle x_2, n_p \rangle| = |\beta \langle x_2, n_q \rangle| = |\langle x, n_q \rangle|$.

Claim 4. There exists a point $z \in P \cap R$ which is at the same distance from some hyperplane not containing x_1 , as from the at least $d - 1$ hyperplanes in whose intersection x_1 lies and with respect to which x_2 is equidistant.

Proof of claim. Note that we must have that $K \neq [m]$, because of the condition that the given arrangement has a sub-arrangement of d hyperplanes that is in general position.

Let $x_3 \in P$ be the vector such that $x_1 \perp x_3$, $\|x_3\| = 1$ (actually x_3 and y are collinear), and such that $x_2/\|x_2\| = \cos(\alpha)x_1 + \sin(\alpha)x_3$ for some $\alpha \in (0, \pi]$. (Clearly the vector $x_2/\|x_2\|$ can be written as above, with α either positive or negative; we choose x_3 so that it is positive.) For any $t \in \mathbb{R}$, let $u(t)$ be the point $\cos(t) \cdot x_1 + \sin(t) \cdot x_3$. Note that $u(\pi) = -x_1$; since $-x_1$ is in a different half-space than x_1 with respect to every H_j that does not contain x_1 , and since x_1 is in the same half-space as R with respect to all such H_j , it follows that $u(\pi) = -x_1$ is in a different half-space than R with respect to some H_j . It is also clear that $u(t)$, regarded as a function from $\mathbb{R}$ to $\mathbb{R}^d$, is continuous.

Consider the function $f_2 : \mathbb{R} \rightarrow \mathbb{R}$ defined by $f_2(t) = \min\{\text{dist}(u(t), H_i) \mid i \in [m]\}$. Since $u(0) = x_1 \in \bigcap_{k \in K} H_k$, we have that $f_2(0) = 0$, and since $x_1 \notin H_j$ when $j \notin K$

(from the definition of K , see the proof of Claim 1), we have that $\text{dist}(u(0), H_j) > 0$ for all $j \notin K$. Since $u(\pi)$ is in a different half-space than R with respect to some H_j , we must have for some $r \in (0, \pi)$ and some $j \notin K$, that $u(r) \in H_j$, and hence that $f_2(r) = \text{dist}(u(r), H_j) = 0$. To see this, let H_j be one of the hyperplanes not containing x_1 , so that x_1 and $-x_1$ lie in different half-spaces with respect to H_j ; then the continuous function $g : [0, \pi] \rightarrow \mathbb{R}$, $g(t) = \langle u(t), n_j \rangle$, has different signs in 0 and in π , so it follows that for some $r \in (0, \pi)$, $g(r) = 0$, and hence that $0 = \text{dist}(u(r), H_j) \geq f_2(r)$, from which we conclude that $f_2(r) = \text{dist}(u(r), H_j) = 0$.

Based on the above, we can now claim that the set $S = \{t \in [0, \pi) \mid \exists j \notin K : f_2(t) = \text{dist}(u(t), H_j)\}$ is non-empty and does not contain 0. The functions $\text{dist}(\cdot, H_j)$, $u(t)$ and $\min$ are continuous, so f_2 is continuous as well, and we argue that $\inf S \in S$, i.e. that S has a minimum: Let $(t_n) \subseteq S$ be an infinite sequence of points with limit $\inf S$; by definition of S , there corresponds to each term t_n at least one $j \notin K$ for which $f_2(t_n) = \text{dist}(u(t_n), H_j)$. Since there are only finitely many such j , this sequence must contain at least one infinite subsequence (t'_n) such that all terms of the subsequence are associated with the same j , that is: $\forall n \in \mathbb{N} : f_2(t'_n) = \text{dist}(u(t'_n), H_j)$. By the continuity of all functions involved, we have that $f_2(\lim t'_n) = \text{dist}(u(\lim t'_n), H_j)$, proving that $\lim t'_n \in S$.

Let $s = \min S$, and let $l \in [m] \setminus K$ be such that $f_2(s) = \text{dist}(u(s), H_l)$. Since s is the smallest value in S , we have that for all $t \in [0, s)$, $f_2(t) = \text{dist}(u(t), H_k)$ for all $k \in Q$, and again by the continuity of f_2 we may conclude that $f_2(s) = \text{dist}(u(s), H_l) = \text{dist}(u(s), H_k)$ for all $k \in Q$, which is the desired conclusion for $z = u(s)$ and $I = \{l\} \cup Q$. That the distance from z to any of these hyperplanes is the minimum distance follows from the minimality of s . $\square$

Remark 4.2.5. The centre-like unit vector provided by Lemma 4.2.4 is not, in gen-

eral, unique, unless $d = 2$. In general it can be shown that the d hyperplanes with respect to which the unit vector z is equidistant, are hyperplanes which delimit the region containing z . However, for $d \geq 3$, a region can be delimited by more than d hyperplanes, and it is possible that there are (finitely many) distinct unit vectors with the properties of z in the region, but that the d hyperplanes with respect to which they are equidistant (as well as the distance itself!) are different for each of the vectors. However, for the “best” unit vector in the region, we may simply choose the one with the greatest minimum distance from the arrangement hyperplanes, among the possible choices for z .

Remark 4.2.6. We have proved Lemma 4.2.4 with more restrictive assumptions in order to not overcomplicate the proof. However, there are two useful generalisations which we will need:

1. If we weaken the initial assumptions such that the rank of the matrix of normal vectors can be $< d$, we may still apply Lemma 4.2.4 to $\text{ess}(\mathcal{A})$ in $\mathbb{R}^k$ (see Definition 2.2.2), which may not be in general position, but will have k hyperplanes which form a sub-arrangement that is in general position (in $\mathbb{R}^k$). The distances from the vector z to the hyperplanes of $\mathcal{A}$ are the same as the distances to the hyperplanes of $\text{ess}(\mathcal{A})$, respectively, since the former are obtained by extending the latter by a subspace that is orthogonal with K and hence is orthogonal with both z and the hyperplanes of $\text{ess}(\mathcal{A})$. However, in this case the conclusion of the proposition will be that z has the same distance from k (not d) hyperplanes, and at least that distance from the remaining $m - k$ hyperplanes.

2. Lemma 4.2.4 still holds if instead of the regular Euclidean distance we consider *scaled Euclidean distances*, with a different scalar factor for each of the hyperplanes of the arrangement. In other words, if we allow the normal vectors $n_1, \dots, n_m$ to

be arbitrary non-zero vectors (instead of unit vectors), but still take the distance of any vector v from H_i as being $|\langle v, n_i \rangle|$ (without normalising n_i), then the proof goes through in exactly the same way. However, one side effect is that, unlike the case where the normal vectors are unit vectors, only $d - 1$ of the d hyperplanes with respect to which the vector z is equidistant, are certain to be delimiting the region in which z lies.

The next lemma makes a connection between the d -minors of a matrix $A \in \mathcal{M}_{m,d}(\mathbb{R})$ and the size of the regions of $\mathcal{A}(A)$. The essence of the proof is that, if there is a region that is “small”, in the sense that every unit vector in it is close to at least one hyperplane that delimits the region, then we can identify d hyperplanes among the ones delimiting the region, which are the cause of the region being “small”.

Lemma 4.2.7. Let $A \in \mathcal{M}_{m,d}(\mathbb{R})$ be such that the row vectors of A , $r_1, \dots, r_m$, all have norm 1, and such that $\mathcal{A} = \mathcal{A}(A)$ is in general position. Let $\epsilon \in (0, 1)$ be given.

If the absolute value of every d -minor of A is $> \epsilon$, then we have that for every region R of $\mathcal{A}$, there is a unit vector $v \in R$ such that $\min\{|w_i| \mid w = Av\} \geq \epsilon/\sqrt{2}^d$ (in other words, every region of the arrangement contains a unit vector which is not too close to any of the hyperplanes of the arrangement).

Proof. Let R be some region of the arrangement, and let $v \in R$ be the unit vector whose existence is guaranteed by Lemma 4.2.4; in other words, let v , be a unit vector such that there are d distinct indices $i_1, \dots, i_d$ for which $|\langle v, r_{i_1} \rangle| = \dots = |\langle v, r_{i_d} \rangle|$ and $|\langle v, r_i \rangle| \geq |\langle v, r_{i_1} \rangle|$ for all $i \neq i_1, \dots, i_d$. Let $\gamma = |\langle v, r_{i_1} \rangle|$, and assume that $\gamma < \epsilon/\sqrt{2}^d$. Without loss of generality, we may assume that $v = (1, 0, \dots, 0)$ (since we can certainly apply a rotation to the arrangement so that v corresponds to this vector). For any fixed $j \in [d]$, let w be the orthogonal projection of v onto the hyperplane that is orthogonal

with r_{i_j} . Then $\|v - w\| = \gamma$, and $\langle v - w, w \rangle = 0$. Thus, if $w = (w_1, \dots, w_d)^T$ we have:

$$(w_1 - 1)^2 + w_2^2 + \dots + w_d^2 = \gamma^2, \text{ and}$$

$$(w_1 - 1)w_1 + w_2^2 + \dots + w_d^2 = 0,$$

so that $1 - w_1 = \gamma^2$ and $w_2^2 + \dots + w_d^2 = \gamma^2 - \gamma^4$. Note that since w is the projection of v onto the hyperplane orthogonal with r_{i_j} , we have that $v - w = (1 - w_1, w_2, \dots, w_d)^T$ is collinear with r_{i_j} , so that $\frac{v-w}{\|v-w\|} = \frac{v-w}{\gamma} = (\gamma, \frac{w_2}{\gamma}, \dots, \frac{w_d}{\gamma})^T$ is either r_{i_j} or $-r_{i_j}$. Since j was arbitrary, we may conclude that for each $j \in [d]$ the first coordinate of r_{i_j} is either γ or $-\gamma$. Thus the matrix $B = (r_{i_1}, \dots, r_{i_d})^T$, which is a $(d \times d)$ -submatrix of A , has the value $\pm\gamma$ on the first column, and all rows of norm 1. Then $|\det B|$ equals γ times the absolute value of the determinant of a matrix whose rows all have norm $\leq \sqrt{2}$. Factoring out the norm of each row we get $|\det B| \leq \gamma\sqrt{2}^d < \epsilon$ (since the absolute value of the determinant of a matrix whose rows all have norm 1 is ≤ 1). The contradiction allows us to conclude that $\min\{|u_i| \mid u = Av\} \geq \gamma \geq \epsilon/\sqrt{2}^d$. Since the region R was arbitrary, this statement holds for all regions. $\square$

When a matrix $A \in \mathcal{M}_{m,d}(\mathbb{R})$ is interpreted as the basis of a subspace, it is convenient to assume that its columns form an orthonormal basis. On the other hand, when interpreting it as an arrangement of hyperplanes, it is convenient to assume that its rows are normalised, and this is also a requirement of Lemma 4.2.7. Evidently we can always switch between the two forms without changing the signature set, by either performing Gram-Schmidt orthonormalisation on the columns, or normalising the rows. However, it is worth noting that the minimum absolute value of a d -minor of a matrix whose column vectors form an orthonormal basis, is less than the minimum absolute value of a d -minor after the rows of the same matrix have been normalised. This is due to the fact that each row of a matrix whose columns form an orthonormal basis must have norm ≤ 1 . (For any $v \in \mathbb{R}^d$, $\|v\| = 1$, we have that $|(Av)_i| \leq \|Av\| =$

$\|\sum_{i=1}^d v_i A_{\cdot,i}\| = \sqrt{\sum_{i=1}^d v_i^2 \|A_{\cdot,i}\|^2} = \|v\| = 1$, due to the columns of A being pairwise orthogonal unit vectors. On the other hand we have that $|(Av)_i| = |\langle A_i, v \rangle|$, and by the Cauchy-Schwarz inequality this expression is maximised when $v = \frac{A_i}{\|A_i\|}$, in which case we get $|(Av)_i| = \frac{\langle A_i, A_i \rangle}{\|A_i\|} = \|A_i\|$. Thus, when normalising the rows of A , the absolute value of every d -minor can only increase. For Conjecture 4.2.3 it is therefore sufficient to produce a matrix whose column vectors form an orthonormal basis of a subspace in $\mathbb{R}^m$, and whose d -minors are sufficiently large in absolute value.

4.3 Signature-invariant operations

In this section we introduce a general operation on subspaces and study its properties, particularly the conditions under which it is invariant with respect to the signature set. Throughout this section we assume $m \gg d \geq 2$, unless mentioned otherwise. Recall that in the context of subspaces of a Euclidean space, $H_1, \dots, H_m$ always denote the canonical hyperplanes of $\mathbb{R}^m$.

We now turn to the question regarding operations which leave the signature set of the arrangement associated with a given matrix $A \in \mathcal{M}_{m,d}(\mathbb{R})$ unchanged. Given the two interpretations of such matrices from Section 2.2, we already have several such operations: multiplying rows by positive values, multiplying columns by non-zero values, multiplying A from the right side by a non-singular $(d \times d)$ -matrix (and in particular, performing the Gram-Schmidt orthonormalisation process on the column vectors of A), or any alternation of these. The first operation is stable with respect to the signature set because it is merely a rescaling of the chosen normal vectors of the corresponding hyperplane arrangement, while the second and third operations are stable because they are simply the choosing of new bases of the same subspace. However,

while these operations can be beneficial, they appear to be quite weak with regards to ensuring that *all* d -minors are sufficiently large in absolute value: For example, if A has two rows that are almost, but not exactly, identical, then the determinant of any $(d \times d)$ -submatrix containing the two rows will be very small, but there is no obvious way to pry these vectors apart using only the above operations, without “damaging” the matrix by making most of its rows have norm almost 0. Therefore we require more powerful tools for handling subspaces.

We define a basic operation which, given a d -dimensional subspace D , produces a new subspace which has a $(d - 1)$ -dimensional intersection with D (in other words, the resulting subspace differs from the original one only along one dimension). This has the advantage of making the effect on the signature set easier to study, without sacrificing the strength of the operation, since we can apply a series of such basic operations in order to change a given subspace along all dimensions.

Definition 4.3.1. Given a d -dimensional subspace D of $\mathbb{R}^m$, a non-zero vector $v \in D$, and $v' \in \mathbb{R}^m \setminus C_D^\perp(v)$, let $o(D, v, v')$ denote the subspace $\langle C_D^\perp(v) \cup \{v'\} \rangle$. Thus $o(D, v, v')$ denotes the result of the operation of replacing v with v' while leaving $C_D^\perp(v)$ intact.

Since $v' \notin C_D^\perp(v)$, $o(D, v, v')$ must have dimension d , hence there must be a non-zero vector $v'' \in C_{o(D, v, v')}^\perp(C_D^\perp(v)) \subset o(D, v, v')$. Evidently $o(D, v, v') = \langle C_D^\perp(v) \cup \{v'\} \rangle = \langle C_D^\perp(v) \cup \{v''\} \rangle = o(D, v, v'')$, and therefore we may always assume that $v' \in C^\perp(C_D^\perp(v))$, in other words, that v' is chosen from the $(m - d + 1)$ -dimensional orthogonal complement of $C_D^\perp(v)$ in $\mathbb{R}^m$.

Given any $v \in D \setminus \{0\}$, we are interested in the set of vectors $v' \in \mathbb{R}^m \setminus D$ such that D and $o(D, v, v')$ intersect the same set of orthants. This is the set of vectors with which we may replace v without causing the signature set of the resulting subspace to differ from that of the original subspace. The next lemma provides a hint as to what

the border of this set will be.

Lemma 4.3.2. Let $D \subseteq \mathbb{R}^m$ be a subspace, and let $v \in D \setminus \{0\}$ and $v' \notin D$ be given.

If the signature sets of D and $o(D, v, v')$ are different, then for some $t \in [0, 1]$ and for $v(t) = (1 - t)v + tv'$, the subspace $D(t) = o(D, v, v(t))$ is not in general position.

Proof. We start with some useful observations: First, we may assume that $D = D(0)$ is in general position, since otherwise we are done. Since $v' \notin D$, we have $(1-t)v + tv' \notin D$ for any $t \in (0, 1]$, and hence $D(t) \neq D$ and $\dim D(t) = d$ for all $t \in (0, 1]$. Furthermore there must be an orthant that is intersected by D but not by $o(D, v, v')$, since D is in general position and the signature sets of the two are not equal. Finally, if $\exists t_0 \in (0, 1] : D(t_0) \subseteq H_i$ for some $i \in [m]$, then one of the rows of any orthogonal basis matrix of $D(t_0)$ is zero, and hence $D(t_0)$ is not in general position and we are done. Thus we may assume that $D(t) \not\subseteq H_i$ for all $i \in [m]$ and $t \in [0, 1]$.

We show that for any fixed orthant of $\mathbb{R}^m$ the values of t for which the intersection of $D(t)$ with the orthant is non-empty, must form an open set (in the topology of $[0, 1]$ induced by the standard topology of the real line). First define the function $F : D \times [0, 1] \rightarrow \mathbb{R}^m$ by $F(\alpha v + \beta w, t) = \alpha v(t) + \beta w$ for all $\alpha, \beta \in \mathbb{R}$ and $w \in C_D^\perp(v)$. It is easy to verify that for every fixed $t \in [0, 1]$, $F(\cdot, t)$ is a bijection between D and $D(t)$, and that for every fixed $u \in D$, $F(u, t)$ is continuous as a function of t , because $v(t)$ is continuous. Then for every $u \in D \setminus \{0\}$ and $t \in [0, 1)$, and for every $\epsilon > 0$, there is a $\delta_0 > 0$ such that for all $\delta \in (0, \delta_0)$ we have that $\|F(u, t + \delta) - F(u, t)\| < \epsilon$, and similarly for every $u \in D \setminus \{0\}$ and $t \in (0, 1]$, and for every $\epsilon > 0$, there is a $\delta_0 > 0$ such that for all $\delta \in (0, \delta_0)$: $\|F(u, t - \delta) - F(u, t)\| < \epsilon$. It follows that if for some $u \in D$ and $t \in [0, 1)$ the vector $F(u, t)$ lies in some orthant of $\mathbb{R}^m$, then there exists a $\delta_0 > 0$ such that for all $\delta \in (0, \delta_0)$ we have that $F(u, t + \delta)$ lies in the same orthant. Similarly, if for some $u \in D$ and $t \in (0, 1]$ the vector $F(u, t)$ lies in some orthant of

$\mathbb{R}^m$, then there exists a $\delta_0 > 0$ such that for all $\delta \in (0, \delta_0)$ we have that $F(u, t - \delta)$ lies in the same orthant. Together these two statements imply that the set of values of t for which any one fixed orthant is intersected by $D(t)$ is an open set (in the induced topology of $[0, 1]$).

From the conclusion of the last paragraph it follows that for each orthant intersected by D we may pick some vector u in the intersection of D with the orthant, and there will be a $\delta(u) \in (0, 1)$ such that for all $t \in [0, \delta(u))$ the vector $F(u, t)$ lies in the same orthant, in other words such that $D(t)$ intersects the orthant for all $t \in [0, \delta(u))$. Since there are only finitely many orthants of $\mathbb{R}^m$, we may set δ to be the minimum of the above values $\delta(u)$, and it follows that there is a $\delta \in (0, 1)$ such that for all $t \in [0, \delta)$ (which is an open set in the topology induced in $[0, 1]$ by the standard topology of the real line), $D(t)$ intersects the same set of orthants as D . Let t_1 be the smallest value in $(0, 1]$ for which there is an orthant which is intersected by D but not by $D(t_1)$ (that this minimum exists follows from the fact that it is taken from the complement of a set that is open in $[0, 1]$). Then every orthant intersected by $D(t_1)$ must also be intersected by D : If there is an orthant intersected by $D(t_1)$ but not by D , then it must be intersected for some $t < t_1$ by $D(t)$ as well. But then $D(t)$ can not intersect all of the orthants intersected by D (because D intersects the maximum number of orthants possible), contradicting the minimality of t_1 . We may conclude that the number of orthants intersected by $D(t_1)$ is strictly less than the number of orthants intersected by D , and therefore that $D(t_1)$ can not be in general position. $\square$

From Lemma 4.3.2 we have that if D and $o(D, v, v')$ are both in general position, but have different signature sets, then one of the “intermediate” subspaces is not in general position. This suggests that the vectors v' for which $o(D, v, v')$ is *not* in general

position may be of interest.

Lemma 4.3.3. Let D be a d -dimensional subspace of $\mathbb{R}^m$. If D is in general position, then for any $v \in D \setminus \{0\}$ and $E \in \mathcal{E}_{m,m-d}$ we have that $\dim(\langle C_D^\perp(v) \cup E \rangle) = m - 1$, and that $v \notin \langle C_D^\perp(v) \cup E \rangle$.

If D is not in general position, then there is an $E \in \mathcal{E}_{m,m-d}$ such that for all $v \in D \setminus C_D^\perp(D \cap E)$ we have $v \in \langle C_D^\perp(v) \cup E \rangle$. (If $v \in C_D^\perp(D \cap E) \setminus \{0\}$, then $\dim(\langle C_D^\perp(v) \cup E \rangle) = \dim(\langle D \cup E \rangle) - 1$ and hence $v \notin \langle C_D^\perp(v) \cup E \rangle$.)

Proof. If D is in general position, then the intersection of D with any $(m - d)$ -dimensional canonical subspace E is $\{0\}$ (Fact 2.2.11), and hence for any non-zero vector $v \in D$ we have that $\dim(\langle C_D^\perp(v) \cup E \rangle) = m - 1$ and that $v \notin \langle C_D^\perp(v) \cup E \rangle$ (otherwise $m - 1 = \dim(\langle C_D^\perp(v) \cup E \rangle) = \dim(\langle D \cup E \rangle) = \dim(D) + \dim(E) - \dim(D \cap E) = d + (m - d) - \dim(D \cap E)$, implying that $\dim(D \cap E) = 1$).

For the second statement, we first have that if D is not in general position, then there is an $E \in \mathcal{E}_{m,m-d}$ such that $\dim(D \cap E) \geq 1$. If $v \notin C_D^\perp(D \cap E)$, then $D \cap E \not\subset C_D^\perp(v)$ and hence $\dim(C_D^\perp(v) \cap E) = \dim(D \cap E) - 1$. It follows that $\dim(\langle C_D^\perp(v) \cup E \rangle) = \dim(C_D^\perp(v)) + (m - d) - \dim(C_D^\perp(v) \cap E) = m - 1 - (\dim(D \cap E) - 1) = \dim(\langle D \cup E \rangle)$, and we get that $\langle C_D^\perp(v) \cup E \rangle = \langle D \cup E \rangle$ and that $v \in \langle C_D^\perp(v) \cup E \rangle$ (which is still a hyperplane if $\dim(D \cap E) = 1$).

If $v \in C_D^\perp(D \cap E)$, then $C_D^\perp(v) \cap E = D \cap E$ and hence $\dim(\langle C_D^\perp(v) \cup E \rangle) = m - 1 - \dim(D \cap E) = \dim(\langle D \cup E \rangle) - 1$, in which case $v \notin \langle C_D^\perp(v) \cup E \rangle$. $\square$

The above lemma justifies the following definition.

Definition 4.3.4. For any two subspaces K, L and any non-zero vector $v \in K$, we denote by $H_{K,L}(v)$ the set $\langle C_K^\perp(v) \cup L \rangle$. Given a d -dimensional subspace $D \subset \mathbb{R}^m$, a non-zero vector $v \in D$, and an $(m - d)$ -dimensional canonical subspace $E \subset \mathbb{R}^m$ such that $E \cap C_D^\perp(v) = \{0\}$, we call the set $H_{D,E}(v)$ an *extension hyperplane of v* (with

respect to D and E). The collection of extension hyperplanes $\mathcal{H}_D(v) = \{H_{D,E}(v) \mid E \in \mathcal{E}_{m,m-d}, E \cap C_D^\perp(v) = \{0\}\}$ is called the *extension arrangement of v (with respect to D)*.

As we will prove shortly, the extension arrangement of a vector $v \in D \setminus \{0\}$ delimits the set of vectors with which v can be safely replaced (in the sense that the replacement operation leaves the signature set unchanged). The name “extension arrangement” is appropriate because extending $C_D^\perp(v)$ by any vector from a region of $\mathcal{H}_D(v)$ results in a subspace which intersects the same set of orthants, and thus the regions of this arrangement correspond to the different signature sets which can be obtained by extending $C_D^\perp(v)$ along one dimension. Evidently this means that $\mathcal{H}_D(v)$ should be associated with $C_D^\perp(v)$ alone, rather than with v and D , but the chosen notation is convenient because we will be studying the behaviour of these hyperplane arrangements when v (and $C_D^\perp(v)$ with it) varies over the set of unit vectors in D .

Note that even if $\dim(E \cap C_D^\perp(v)) \geq 1$ it is meaningful to consider the sets $H_{D,E}(v) = \langle C_D^\perp(v) \cup E \rangle$, although they are not hyperplanes but rather subspaces of lower dimension.

From Lemma 4.3.3 we have that if D is in general position, then for all non-zero $v \in D$, v lies in a region R of $\mathcal{H}_D(v)$ (and not in any extension hyperplane), and we will denote this region by $R(v)$. Using Lemma 4.3.2 we can show that the region $R(v)$ of $\mathcal{H}_D(v)$ is contained in the set of vectors with which v can be replaced without changing the signature set.

Proposition 4.3.5. Let D be a d -dimensional subspace in general position in $\mathbb{R}^m$. Let $v \in D \setminus \{0\}$ be given, and denote by $R(v)$ the region of $\mathcal{H}_D(v)$ containing v . Then for any vector $v' \in R(v) \cup (-R(v))$ we have that D and $o(D, v, v')$ have the same signature set.

Proof. We may assume that $v' \perp C_D^\perp(v)$, and therefore that $\mathcal{H}_{D(t)}(v(t)) = \mathcal{H}_D(v)$, where $v(t) = (1-t)v + tv'$ and $D(t) = o(D, v, v(t))$, for all $t \in [0, 1]$. If $v' \in R(v)$, then $[v, v'] \subset R(v)$, and hence this segment does not intersect any hyperplane of the extension arrangement. Each intermediate subspace $D(t)$ intersects the segment in exactly one point, namely $v(t)$. If $D(t)$ were not in general position for some $t \in [0, 1]$, it would follow that either $v(t)$ is contained in an extension hyperplane (which we have just established can not be the case), or that $v(t) \in C_{D(t)}^\perp(D(t) \cap E)$ for some $E \in \mathcal{E}_{m, m-d}$ with $\dim(D(t) \cap E) \geq 1$. But in the latter case we would have $D(t) \cap E \subseteq C_{D(t)}^\perp(v(t)) = C_D^\perp(v)$, implying that D is not in general position. Having ruled out both cases, we may conclude that $D(t)$ must be in general position for all $t \in [0, 1]$, and by Lemma 4.3.2 we have that $o(D, v, v')$ and D have the same orthant intersection.

The same statement holds for vectors in $-R(v)$, because we have that $o(D, v, -v) = D$ and that $o(D, -v, v')$ and D have the same signature set for all $v' \in (-R(v))$ (by the above reasoning, applied to $-v$ and $-R(v)$). However, it can be shown that if $v' \in (-R(v))$ then for some $t \in (0, 1]$ the intermediate subspace is not in general position. $\square$

We leave it as a conjecture that the set of vectors with which a vector v can be safely replaced is exactly $R(v) \cup (-R(v))$. Since we will restrict our applications of the replacement operation to target vectors which lie in the same region of $\mathcal{H}_D(v)$ as the vector v itself (excluding even vectors in $-R(v)$, in order to preserve the fact that all intermediate subspaces also have the same signature set), it is not relevant for our purpose whether an extension arrangement can have more than one pair of opposite regions corresponding to the same signature set of an extended subspace. However, this fact is very relevant for the combinatorial problem of counting the number of distinct signature sets which can be realised by subspaces in general position.

The next lemma and the remark following it describe in detail how the signature set of the resulting subspace differs from that of the original one when v' lies in a single extension hyperplane of v , or on the opposite side of one. We state them mostly for the sake of completeness, and will only need one fact from the proof of the lemma later on.

Lemma 4.3.6. Given $D \subset \mathbb{R}^m$ in general position, let $v \in D$ be a non-zero vector whose region $R(v)$ of $\mathcal{H}_D(v)$ has a delimiting hyperplane H which is unique in $\mathcal{H}_D(v)$, and let E be the $(m-d)$ -dimensional canonical subspace for which $H = H_{D,E}(v)$. Let v' be such that $(1-t)v + tv' \in R(v)$ for all $t \in [0, 1)$, and such that $v' \in H$, but is not contained in any other extension hyperplane of v . Let $D' = o(D, v, v')$.

Then the signature set of D' is contained in that of D , and there is a single pair of opposite orthants which are intersected by D but not by D' .

Proof. We may assume without loss of generality that $v' \perp C_D^\perp(v)$. In this case it holds for all $(m-d)$ -dimensional canonical subspaces E' that $H_{D,E'}(v) = H_{D',E'}(v')$.

To show that the signature set of D' is contained in that of D , we define for each vector $w \in D$ the continuous mapping $w(t)$ as follows: If $w = \alpha v + P_{C_D^\perp(v)}(w)$, set $w(t) = \alpha((1-t)v + tv') + P_{C_D^\perp(v)}(w)$ for every $t \in [0, 1]$; clearly the map $w \mapsto w(1)$, $w \in D$, is a bijection between D and D' . Now consider any vector $w \in D$ such that $w(1) \in D'$ lies in an orthant (has all coordinates non-zero). By the continuity of $w(t)$ we must have for some $\epsilon > 0$ that $w(1-\epsilon)$ has all coordinates non-zero and of the same sign as those of $w(1)$, implying that the signature of the orthant that $w(1)$ lies in is also an element of the signature set of $D(1-\epsilon)$, and since by Lemma 4.3.2 the signature set of $D(t)$ is the same for all $t \in [0, 1)$, it follows that the same orthant signature is also an element of the signature set of D . Thus the signature set of D' is contained in that of D .

Since $v' \in H_{D,E}(v)$ and $\dim(H_{D,E}(v)) = m - 1$, we must have that $\dim(D' \cap E) = 1$. In fact if $E = \bigcap_{i \in I} H_i$, where $I = \{i_1, \dots, i_d\}$, then for any $i_0 \in I$: $\dim(D' \cap (\bigcap_{i \in I \setminus \{i_0\}} H_i)) = 1$. To see this, assume that for some $i_0 \in I$: $\dim(D' \cap (\bigcap_{i \in I \setminus \{i_0\}} H_i)) \geq 2$. Fix $j \in [m] \setminus I$ and $E' = (\bigcap_{i \in I \setminus \{i_0\}} H_i) \cap H_j$. Then $\dim(D' \cap E') \geq 1$, and hence $v' \in H_{D',E'}(v') = H_{D,E'}(v)$, contradicting the fact that v' lies in only one extension hyperplane. (As in the previous subsection, $\dim(D' \cap E') \geq 1$ implies that $v' \in H_{D',E'}(v')$ because otherwise $H_{D',E'}(v') = H_{D,E'}(v)$ is not a hyperplane, contradicting the fact that D is in general position.)

Consider now the sub-arrangement of the induced arrangement of D' formed by the hyperplanes $D' \cap H_{i_1}, \dots, D' \cap H_{i_d}$. Since the intersection of all d of these hyperplanes, as well as of any $d - 1$ of them, is a line (denote it by L), it follows that the essentialisation of this sub-arrangement is itself an arrangement of d hyperplanes in $\mathbb{R}^{d-1}$, and that it is in general position. It can be checked that such an arrangement has $2^d - 2$ regions, which is then also the number of regions of the sub-arrangement of the d hyperplanes in $\mathbb{R}^d$. Thus exactly two opposite regions of the 2^d regions which can be formed by d hyperplanes in $\mathbb{R}^d$, are not formed by this sub-arrangement of the induced arrangement of D' , whereas in $D(t)$ for any $t \in [0, 1)$, the corresponding d hyperplanes form all possible 2^d regions.

For every $t \in [0, 1)$, $D(t)$ can be regarded as being obtained through a perturbation of D' which leaves all regions of the induced arrangement of D' intact signature-wise, although they can change as sets. The line L is not contained in any other hyperplane of the induced arrangement of D' , hence its direction vectors have non-zero distance from $D' \cap H_i$ for every $i \in [m] \setminus I$. Since for each $i \in [m]$ the hyperplane $D(t) \cap H_i$ changes continuously as a function of t , there must be an $\epsilon > 0$ such that the pair of opposite regions formed by $D(1 - \epsilon) \cap H_{i_1}, \dots, D(1 - \epsilon) \cap H_{i_d}$ but not by $D(1) \cap H_{i_1}, \dots, D(1) \cap H_{i_d}$ are not intersected by any of the hyperplanes

$D(1 - \epsilon) \cap H_i$ for $i \in [m] \setminus I$. Therefore these two regions of the sub-arrangement $\{D(t) \cap H_{i_1}, \dots, D(t) \cap H_{i_d}\}$ are in fact (as sets) also regions of the induced arrangement of $D(t)$ for every $t \in [1 - \epsilon, 1)$, and we conclude that the signatures of these two opposite regions, which are delimited precisely by the hyperplanes $D(t) \cap H_{i_1}, \dots, D(t) \cap H_{i_d}$, are the only signatures in the signature set of D which are not in the signature set of D' . $\square$

Remark 4.3.7. From the proof of Lemma 4.3.6 it can be seen that the two opposite regions which exist in the induced arrangement of D but not in that of D' are delimited by exactly d hyperplanes, and that when t goes continuously from 0 to 1, these d hyperplanes move into a position where they intersect in a line, in such a way that exactly two opposite regions formed by them gradually contract into that line.

From the proof we can also deduce what happens when the vector $(1 - t)v + tv'$ passes the extension hyperplane: The line into which the two opposite regions contract at $t = 1$ expands again into a pair of opposite regions when $t \in (1, 1 + \epsilon]$, for some small $\epsilon > 0$. If ϵ is sufficiently small, the two regions will not be intersected by any other hyperplane $D(t) \cap H_j$ when $j \notin I$, and thus they will each lie in the same halfspace with respect to $D(t) \cap H_j$ for all $j \in [m] \setminus I$. With respect to the hyperplanes $D(t) \cap H_i$, $i \in I$, the signature of each of the two regions will be the opposite of what it is for $t \in (0, 1)$. In conclusion, for any $v \in D$, if v' lies in a region of $\mathcal{H}_D(v)$ that is *adjacent* to $R(v)$ and is separated from $R(v)$ by a single extension hyperplane, then the signature set of $o(D, v, v')$ is the same as that of D , except for a pair of opposite signatures, which have their values on the same d positions flipped (multiplied by -1).

The next lemma states that extension arrangements of vectors from d -dimensional subspaces of $\mathbb{R}^m$ are essentially hyperplane arrangements in $\mathbb{R}^{m-d+1}$. The lemma is needed only for the proof of Proposition 4.5.1.

Lemma 4.3.8. If D is a d -dimensional subspace of $\mathbb{R}^m$, and $v \in D \setminus \{0\}$, then the normal vectors of the extension hyperplanes of v span $C^\perp(C_D^\perp(v))$, and therefore every region of $\text{ess}(\mathcal{H}_D(v))$ is well-formed in the sense that it is delimited by at least $m-d+1$ extension hyperplanes whose normal vectors span $C^\perp(C_D^\perp(v))$.

Proof. Since every extension hyperplane of a non-zero vector $v \in D$ contains the set $C_D^\perp(v)$, it is clear that $\text{ess}(\mathcal{H}_D(v))$ (see Definition 2.2.2) is contained in $C^\perp(C_D^\perp(v))$, which is an $(m-d+1)$ -dimensional subspace. We prove that the normal vectors of the extension hyperplanes of v span precisely this subspace only in the case when $D \not\subseteq H_i$ for any $i \in [m]$, which includes the case when D is in general position: If K , the subspace spanned by the normal vectors of the extension hyperplanes, is a strict subspace of $L = C^\perp(C_D^\perp(v))$, then there must be a unit vector $w \in C_L^\perp(K)$. Taking any unit vector $u \in K$ which lies in a region of the arrangement $\{H \cap K \mid H \in \mathcal{H}_D(v)\}$ (in K , regarded as $\mathbb{R}^{\dim(K)}$), the two unit vectors $v' = \frac{1}{\sqrt{2}}w + \frac{1}{\sqrt{2}}u$ and $v'' = -\frac{1}{\sqrt{2}}w + \frac{1}{\sqrt{2}}u$ are orthogonal and lie in the same region of $\{H \cap L \mid H \in \mathcal{H}_D(v)\}$ (in L , regarded as $\mathbb{R}^{m-d+1}$), and hence in the same region of $\mathcal{H}_D(v)$ (in $\mathbb{R}^m$). Now by Proposition 4.3.5, $D' = o(D, v, v')$ and $D'' = o(D, v, v'')$ have the same signature set. Since $D \not\subseteq H_i$, there exists a unit vector $x \in D'$ which lies in an orthant of $\mathbb{R}^m$, and it follows that for $\alpha > 0$ sufficiently small, the vector $y = v' + \alpha x$ has all coordinates non-zero and is such that y_i and v'_i have the same sign whenever $v'_i \neq 0$. On the other hand, $D'' \subset C^\perp(v')$ (a hyperplane in $\mathbb{R}^m$), because $v' \perp v''$ and $v' \perp C_D^\perp(v)$, which implies that D'' contains no vectors z such that z_i and v'_i have the same sign whenever $v'_i \neq 0$ (since such vectors can not be orthogonal to v'). This means that D'' can not intersect the orthant of y , contradicting the fact that the two subspaces intersect the same set of orthants. $\square$

The main difficulty in working with extension arrangements is that, even though

$\mathcal{H}_D(v) = \mathcal{H}_{o(D,v,v')}(v')$ whenever $v' \perp C_D^\perp(v)$, $v' \neq 0$, the extension arrangement of a vector $w \in C_D^\perp(v)$ can change significantly when v is replaced with v' , in other words $\mathcal{H}_D(w) \neq \mathcal{H}_{o(D,v,v')}(w)$. The next proposition states that if a vector v “is moved further away” from an extension hyperplane $H_{D,E}(v)$ (for some fixed $E \in \mathcal{E}_{m,m-d}$), then for no $w \in C_D^\perp(v)$ can the extension hyperplane corresponding to E “move closer” to w . The Proposition is stated in more general terms, where it is not required that K (corresponding to D) is in general position, and L (corresponding to E) is not necessarily a canonical subspace.

Proposition 4.3.9. Let $K, L \subseteq \mathbb{R}^m$ be subspaces such that $k = \dim(K) \geq 1$, $K \cap L = \{0\}$ and $\langle K \cup L \rangle = \mathbb{R}^m$. Let $v \in K$ be a unit vector and q a normal vector of the hyperplane $H_{K,L}(v)$ such that $\|q\| = 1$ and $\langle q, v \rangle \geq 0$. For some fixed $\alpha > 0$ set $v' := \frac{v + \alpha q}{\|v + \alpha q\|}$ and $K' := o(K, v, v') = \langle C_K^\perp(v) \cup \{v'\} \rangle$.

Then $\text{dist}(v, H_{K,L}(v)) \leq \text{dist}(v', H_{K',L}(v'))$ and $\forall w \in C_K^\perp(v) : \text{dist}(w, H_{K,L}(w)) \leq \text{dist}(w, H_{K',L}(w))$.

Proof. We start with the observation that $\forall u \in K \setminus \{0\} : u \notin H_{K,L}(u)$: If there were such a vector u that is contained in $H_{K,L}(u)$, we would have that $K \subseteq H_{K,L}(u)$, and since $\dim(K) = k$ and $\dim(L) = m - k$, while $\dim(H_{K,L}(u)) = m - 1$, it would follow that $\dim(K \cap L) \geq 1$, which contradicts the hypothesis. In particular, this means that $\langle v, q \rangle > 0$.

Let $w \in C_K^\perp(v)$ be a unit vector. Our first goal is to show that we can write $H_{K,L}(w) = \langle (H_{K,L}(w) \cap H_{K',L}(w)) \cup \{v\} \rangle$ and $H_{K',L}(w) = \langle (H_{K,L}(w) \cap H_{K',L}(w)) \cup \{v'\} \rangle$.

We may assume that $\dim(K \cap K') = k - 1$, since otherwise we have that $K = K'$, in which case the inequality holds trivially with the two distances being equal. Thus we may assume that $K \cap K' = C_K^\perp(v) = C_{K'}^\perp(v')$, and hence we have that

$H_{K,L}(v) = \langle C_K^\perp(v) \cup L \rangle = \langle C_{K'}^\perp(v') \cup L \rangle = H_{K',L}(v')$. Next, we prove that $C_K^\perp(w) = \langle (C_K^\perp(w) \cap C_K^\perp(v)) \cup \{v\} \rangle$ and $C_{K'}^\perp(w) = \langle (C_K^\perp(w) \cap C_K^\perp(v)) \cup \{v'\} \rangle$. To see this, observe first that $C_K^\perp(v) \cap C_K^\perp(w) = C_{K'}^\perp(v') \cap C_{K'}^\perp(w)$: If $u \in C_K^\perp(v) \cap C_K^\perp(w)$ then u is a vector in $C_K^\perp(v) = C_{K'}^\perp(v') \subset K'$ such that u is orthogonal to w , and thus $u \in C_{K'}^\perp(v') \cap C_{K'}^\perp(w)$; the opposite inclusion is proved in the same way. Thus we have a $(k-2)$ -dimensional subspace $C_K^\perp(w) \cap C_K^\perp(v)$ that is contained in both $C_K^\perp(w)$ and $C_{K'}^\perp(w)$, but contains neither v nor v' , allowing us to write $C_K^\perp(w) = \langle (C_K^\perp(w) \cap C_K^\perp(v)) \cup \{v\} \rangle$ and $C_{K'}^\perp(w) = \langle (C_K^\perp(w) \cap C_K^\perp(v)) \cup \{v'\} \rangle$.

Based on the above we have that $M = \langle (C_K^\perp(w) \cap C_K^\perp(v)) \cup L \rangle$ is an $(m-2)$ -dimensional subspace that is contained in both $H_{K,L}(w)$ and $H_{K',L}(w)$, but contains neither v nor v' . To see that $v \notin M$, assume the contrary. Then $M = \langle (C_K^\perp(w) \cap C_K^\perp(v) \cup L \cup \{v\}) \rangle = \langle \langle (C_K^\perp(w) \cap C_K^\perp(v)) \cup \{v\} \rangle \cup L \rangle = C_K^\perp(w) + L = H_{K,L}(w)$ (with a similar statement holding if $v' \in M$); but this is clearly impossible, since the left-most set has dimension $m-2$ while the right-most set has dimension $m-1$. Thus we may conclude that $H_{K,L}(w) = \langle M \cup \{v\} \rangle$ and $H_{K',L}(w) = \langle M \cup \{v'\} \rangle$.

Using these facts, we now compute the two relevant distances as:

$$\begin{aligned} \text{dist}(w, H_{K,L}(w)) &= \sqrt{\|w\|^2 - \|P_{\langle C_K^\perp(w) \cup L \rangle}(w)\|^2} = \\ &= \sqrt{\|w\|^2 - \left(\|P_{\langle (C_K^\perp(v) \cap C_K^\perp(w)) \cup L \rangle}(w)\|^2 + \|P_{\langle \{v - P_M(v)\} \rangle}(w)\|^2 \right)} = \\ &= \sqrt{\|w\|^2 - \|P_M(w)\|^2 - \|P_{\langle \{v - P_M(v)\} \rangle}(w)\|^2} \end{aligned}$$

and

$$\begin{aligned} \text{dist}(w, H_{K',L}(w)) &= \sqrt{\|w\|^2 - \|P_{H_{K',L}(w)}(w)\|^2} = \\ &= \sqrt{\|w\|^2 - \|P_{\langle (C_K^\perp(v) \cap C_K^\perp(w)) \cup L \rangle}(w)\|^2 - \|P_{\langle \{v' - P_M(v')\} \rangle}(w)\|^2} = \\ &= \sqrt{\|w\|^2 - \|P_M(w)\|^2 - \|P_{\langle \{v' - P_M(v')\} \rangle}(w)\|^2}. \end{aligned}$$

Since in the expressions above the first two terms of the sums under the root signs are

respectively equal, the two expressions can only be different if the third terms differ from one another. Focusing on the third term in each of the two expressions, we have:

$$\|P_{\{v-P_M(v)\}}(w)\| = \left| \left\langle w, \frac{v - P_M(v)}{\|v - P_M(v)\|} \right\rangle \right| = \frac{1}{\|v - P_M(v)\|} |\langle w, v - P_M(v) \rangle|$$

and

$$\begin{aligned} \|P_{\{v'-P_M(v')\}}(w)\| &= \left| \left\langle w, \frac{v' - P_M(v')}{\|v' - P_M(v')\|} \right\rangle \right| = \left| \left\langle w, \frac{v + \alpha q - P_M(v + \alpha q)}{\|v + \alpha q - P_M(v + \alpha q)\|} \right\rangle \right| \\ &= \frac{1}{\|v + \alpha q - P_M(v)\|} |\langle w, v - P_M(v) \rangle|, \end{aligned}$$

where for the final line we used the facts that $q \perp M$ and $q \perp w$.

Since $\langle \alpha q, v - P_M(v) \rangle = \alpha \langle q, v \rangle + 0 > 0$, we have that $\|v + \alpha q - P_M(v)\| > \|v - P_M(v)\|$, and hence that $\|P_{\{v'-P_M(v')\}}(w)\| \leq \|P_{\{v-P_M(v)\}}(w)\|$, which completes the proof that $\text{dist}(w, H_{K,L}(w)) \leq \text{dist}(w, H_{K',L}(w))$.

To show that $\text{dist}(v, H_{K,L}(v)) \leq \text{dist}(v', H_{K',L}(v'))$, we only have to prove that $|\langle v, q \rangle| \leq |\langle v', q \rangle|$, since $H_{K,L}(v) = H_{K',L}(v')$ and q is an orthogonal unit vector of this hyperplane. We may drop the modulus signs since $\langle v, q \rangle > 0$ and $\langle v', q \rangle > 0$, and we have:

$$\begin{aligned} \langle v, q \rangle \leq \langle v', q \rangle &\iff \langle v, q \rangle \leq \frac{\langle v + \alpha q, q \rangle}{\|v + \alpha q\|} \iff \langle v, q \rangle \leq \frac{\langle v, q \rangle + \alpha}{\|v + \alpha q\|} \iff \\ &\iff \langle v, q \rangle^2 \leq \frac{(\langle v, q \rangle + \alpha)^2}{\langle v + \alpha q, v + \alpha q \rangle} \iff \langle v, q \rangle^2 \leq \frac{\langle v, q \rangle^2 + 2\alpha \langle v, q \rangle + \alpha^2}{\langle v, v \rangle + 2\alpha \langle v, q \rangle + \alpha^2 \langle q, q \rangle} \iff \\ &\iff \langle v, q \rangle^2 \leq \frac{\langle v, q \rangle^2 + 2\alpha \langle v, q \rangle + \alpha^2}{1 + 2\alpha \langle v, q \rangle + \alpha^2} \iff \\ &\iff \langle v, q \rangle^2 + 2\alpha \langle v, q \rangle^3 + \alpha^2 \langle v, q \rangle^2 \leq \langle v, q \rangle^2 + 2\alpha \langle v, q \rangle + \alpha^2. \end{aligned}$$

The final inequality holds since $1 \geq \langle v, q \rangle \geq 0$, and we conclude that $\text{dist}(v, H_{K,L}(v)) \leq \text{dist}(v', H_{K',L}(v'))$. $\square$

Remark 4.3.10. Considering the fact that we were able to prove the above inequalities for v as well as for all vectors w in $C_K^\perp(v)$, it is reasonable to ask whether a

similar inequality might hold for all vectors in K . Of course, if $v \neq v'$, then vectors in $K \setminus C_K^\perp(v)$ will not be in K' , so we first need to choose a bijective mapping from K to K' that allows us to identify the vectors from K with those from K' . A possible mapping is $u = \beta v + \gamma w \mapsto \beta v' + \gamma w = u'$, where $\beta, \gamma \in \mathbb{R}$ and $w \in C_K^\perp(v)$. This mapping is a natural choice because it extends the mappings used implicitly in Proposition 4.3.9: $v \mapsto v'$ and $\forall w \in C_K^\perp(v) : w \mapsto w$. Now the question is: Under the conditions of Proposition 4.3.9 and using the above mapping $u \mapsto u'$, is it true that $\text{dist}(u, H_{K,L}(u)) \leq \text{dist}(u', H_{K',L}(u'))$ for all $u \in K$? The following counter-example shows that this need not always be the case.

Let $m = 3$, and let K be $\langle\{v, e\}\rangle$, where $e = (0, 1, 0)$ and $v = (\sqrt{1 - \epsilon^2}, 0, \epsilon)$ for $\epsilon \in (0, 1]$ (to be thought of as fixed; will be chosen later). Let $L = \langle\{l\}\rangle$, where $l = (\lambda_1, \lambda_2, 0)$, $\lambda_1, \lambda_2 \in \mathbb{R}$, $\lambda_1 \neq 0$. Thus $H_{K,L}(v) = \langle\{e, l\}\rangle$ is the elementary plane xOy , and $q = (0, 0, 1)$. Now, for $\alpha > 0$ we define $v' = \frac{v + \alpha q}{\|v + \alpha q\|}$. If we set $\delta = \frac{\alpha + \epsilon}{\sqrt{1 + \alpha^2 + 2\alpha\epsilon}}$ then we can simply write $v' = (\sqrt{1 - \delta^2}, 0, \delta)$. It is important to note that for any fixed $\epsilon \in (0, 1]$ we have $\lim_{\alpha \rightarrow \infty} \delta(\alpha) = 1$. As mentioned above, we take $u = \beta v + \gamma e$, where $\beta^2 + \gamma^2 = 1$, and $u' = \beta v' + \gamma e$, and set $u^\perp = -\gamma v + \beta e$ and $u'^\perp = -\gamma v' + \beta e$ (such that $\langle u, u^\perp \rangle = 0$ and $\langle u', u'^\perp \rangle = 0$).

We start by noting that $C_{K'}^\perp(u') = \langle\{u'^\perp\}\rangle$, and thus $H_{K',L}(u') = \langle\{u'^\perp, l\}\rangle$. Using the cross product of $u'^\perp = (-\gamma\sqrt{1 - \delta^2}, \beta, -\gamma\delta)$ and $l = (\lambda_1, \lambda_2, 0)$ we find a normal vector for this plane:

$$\begin{aligned} N = u'^\perp \times l &= \left(\begin{vmatrix} \lambda_2 & 0 \\ \beta & -\gamma\delta \end{vmatrix}, - \begin{vmatrix} \lambda_1 & 0 \\ -\gamma\sqrt{1 - \delta^2} & -\gamma\delta \end{vmatrix}, \begin{vmatrix} \lambda_1 & \lambda_2 \\ -\gamma\sqrt{1 - \delta^2} & \beta \end{vmatrix} \right) = \\ &= (-\lambda_2\gamma\delta, \lambda_1\gamma\delta, \lambda_1\beta + \lambda_2\gamma\sqrt{1 - \delta^2}). \end{aligned}$$

Recalling that $u' = (\beta\sqrt{1 - \delta^2}, \gamma, \beta\delta)$, our distance function becomes:

$$\frac{|\langle u', N \rangle|}{\|N\|} = \frac{|-\lambda_2\beta\gamma\delta\sqrt{1 - \delta^2} + \lambda_1\gamma^2\delta + \lambda_1\beta^2\delta + \lambda_2\beta\gamma\delta\sqrt{1 - \delta^2}|}{\sqrt{(\lambda_2\gamma\delta)^2 + (\lambda_1\gamma\delta)^2 + (\lambda_1\beta + \lambda_2\gamma\sqrt{1 - \delta^2})^2}} =$$

$$\begin{aligned}
&= \frac{|\lambda_1 \delta|}{\sqrt{(\lambda_2 \gamma \delta)^2 + (\lambda_1 \gamma \delta)^2 + (\lambda_1 \beta + \lambda_2 \gamma \sqrt{1 - \delta^2})^2}} = \\
&= \frac{|\lambda_1|}{\sqrt{(\lambda_2 \gamma)^2 + (\lambda_1 \gamma)^2 + \left(\frac{\lambda_1 \beta + \lambda_2 \gamma \sqrt{1 - \delta^2}}{\delta}\right)^2}}.
\end{aligned}$$

Similarly we have that $C_K^\perp(u) = \langle \{u^\perp\} \rangle$, and therefore $H_{K,L}(u) = \langle \{u^\perp, l\} \rangle$. Because u has the same form as u' , except with ϵ instead of δ , it is easy to see that the above computations are virtually identical for u , and we obtain that

$$\text{dist}(u, H_{K,L}(u)) = \frac{|\lambda_1|}{\sqrt{(\lambda_2 \gamma)^2 + (\lambda_1 \gamma)^2 + \left(\frac{\lambda_1 \beta + \lambda_2 \gamma \sqrt{1 - \epsilon^2}}{\epsilon}\right)^2}}.$$

In the expression above all terms are independent of ϵ , except for the term $x(\epsilon) := \left(\frac{\lambda_1 \beta + \lambda_2 \gamma \sqrt{1 - \epsilon^2}}{\epsilon}\right)^2$, which is non-negative and differentiable on $(0, 1)$. Now if $\beta, \gamma, \lambda_1, \lambda_2$ are such that $\lambda_2 \gamma < 0 < \lambda_1 \beta$ and $|\lambda_2 \gamma| > |\lambda_1 \beta|$, then for ϵ sufficiently close to 0 the numerator will be negative, as $\sqrt{1 - \epsilon^2}$ will be close to 1. However, for ϵ close to 1, the term $\sqrt{1 - \epsilon^2}$ approaches 0, and thus for ϵ sufficiently close to 1 the numerator will be positive. It follows that for some $\epsilon \in (0, 1)$ the numerator (and therefore $x(\epsilon)$) is 0. Let us choose this value for ϵ . Independently of this value, we are still free to choose any $\alpha > 0$ (thereby choosing $\delta \in (\epsilon, 1)$), and from the above analysis of the monotonicity of $x(\epsilon)$ (which has the same monotonicity as $x(\delta)$), it follows that for some α close to 0 we must have that the term $\left(\frac{\lambda_1 \beta + \lambda_2 \gamma \sqrt{1 - \delta^2}}{\delta}\right)^2$ is positive. But then $\text{dist}(u', H_{K',L}(u')) < \text{dist}(u, H_{K,L}(u))$, which concludes the construction of our counter-example.

4.4 The geometry of extension arrangements

In this section we study the relationship between the extension arrangements of vectors of a subspace, and the position of that subspace in $\mathbb{R}^m$. As in the previous section, we assume that $m \gg d \geq 2$, unless stated otherwise, and denote by H_i , $i \in [m]$, the

m canonical hyperplanes of $\mathbb{R}^m$.

We begin by discussing the significance of the distance of a unit vector from one of its extension hyperplanes. Assume once again that D is a d -dimensional subspace in general position in $\mathbb{R}^m$. If $E \in \mathcal{E}_{m,m-d}$, say $E = \bigcap_{i \in [d]} H_i$, and $v \in D$ is a unit vector, then $\text{dist}(v, E) = \|v - P_E(v)\| = \sqrt{v_1^2 + \dots + v_d^2}$ is the weight of the vector v on coordinates $1, \dots, d$. Now if A is a matrix such that its column vectors form an orthonormal basis of D , then $|\det(A[[d], [d]])|$ is upper-bounded by the product of the distances from E of the basis vectors. Thus if D has a unit vector which is very close to E , then the corresponding d -minor will be very small in absolute value. Unfortunately, it may be happen that the column vectors of A all have large distance from E , even though there is a unit vector in D which is very close to E , far away from them. It is therefore not possible to say anything about the position of a subspace relative to an $(m-d)$ -dimensional canonical subspace (or about the corresponding d -minor), just by looking at one basis of the subspace. On the other hand, we have already seen in Lemma 4.3.3 that if for some $E \in \mathcal{E}_{m,m-d}$ there is a unit vector $v \in D \cap E$, meaning that $\text{dist}(v, E)$ and the corresponding d -minor are 0, then *almost all* unit vectors $w \in D$ are such that $\text{dist}(w, H_{D,E}(w)) = 0$. This suggests that the distance of almost any unit vector w from its extension hyperplane $H_{D,E}(w)$ contains *global* information about the position of D relative to E , and we will prove this rigorously in Theorem 4.4.2.

First, we show that if for some $E \in \mathcal{E}_{m,m-d}$ every unit vector in an orthonormal basis of D is far from its extension hyperplane with respect to E , then *every* unit vector $w \in D$ is far from $H_{D,E}(w)$ (in contrast with the observation above that the same does not hold if we take the distance from E for every vector). The lemma is again stated in more general terms, where it is not required that K (corresponding to D) is in general position, and L (corresponding to E) is not necessarily a canonical

subspace.

Lemma 4.4.1. Let K, L be two subspaces of $\mathbb{R}^m$ such that $k = \dim(K) \geq 1$, $K \cap L = \{0\}$ and $\langle K \cup L \rangle = \mathbb{R}^m$. Let $\{v_1, \dots, v_k\}$ be an orthonormal basis of K . Let $\delta = \min\{\text{dist}(v_i, H_{K,L}(v_i)) \mid i \in [k]\}$.

Then for any $u \in K$ with $\|u\| = 1$ we have that $\text{dist}(u, H_{K,L}(u)) \geq \frac{\delta}{(2\sqrt{2})^{k-1}}$.

Proof. The proof is by induction on k . For $k = 1$ the conclusion holds trivially. The first non-trivial case, $k = 2$, will be used in the induction step, so we take it as the base case.

Base case: $k = 2$. For any $u \in K$ such that $\|u\| = 1$ we may write $u = \alpha_1 v_1 + \alpha_2 v_2$ for some $\alpha_1^2 + \alpha_2^2 = 1$. We may assume that both α_1 and α_2 are non-negative, since we may replace v_i with $-v_i$ if necessary, without changing any of the distances between basis vectors and hyperplanes. Denote by $u^\perp$ the vector $-\alpha_2 v_1 + \alpha_1 v_2$. Clearly we have $u^\perp \in K$, $\|u^\perp\| = 1$ and $u \perp u^\perp$. Thus in this case we may write $H_{K,L}(u) = \langle L \cup \{u^\perp\} \rangle$.

We start by taking a closer look at the distances of the chosen basis vectors from their respective hyperplanes.

$$\begin{aligned} \text{dist}(v_1, H_{K,L}(v_1)) &= \|v_1 - P_{H_{K,L}(v_1)}(v_1)\| = \|v_1 - P_{\langle L \cup \{v_2\} \rangle}(v_1)\| = \\ &= \|v_1 - P_{\langle L \cup \{v_2 - P_L(v_2)\} \rangle}(v_1)\| = \|v_1 - (P_L(v_1) + P_{\{v_2 - P_L(v_2)\}}(v_1))\| = \\ &= \left\| v_1 - P_L(v_1) - \left\langle v_1, \frac{v_2 - P_L(v_2)}{\|v_2 - P_L(v_2)\|} \right\rangle \frac{v_2 - P_L(v_2)}{\|v_2 - P_L(v_2)\|} \right\| = \\ &= \left\| v_1 - P_L(v_1) - \frac{\langle v_1, v_2 - P_L(v_2) \rangle}{\|v_2 - P_L(v_2)\|^2} (v_2 - P_L(v_2)) \right\| \\ &= \|v_1 - P_L(v_1) - \eta_1(v_2 - P_L(v_2))\|, \end{aligned}$$

where $\eta_1 = \frac{\langle v_1, v_2 - P_L(v_2) \rangle}{\|v_2 - P_L(v_2)\|^2}$. Now, if we set $a_1 := v_1 - P_L(v_1) - \eta_1(v_2 - P_L(v_2))$, $b_1 := P_L(v_1)$ and $c_1 := \eta_1(v_2 - P_L(v_2))$, we can write $v_1 = a_1 + b_1 + c_1$, and have that $a_1 \perp b_1 \perp c_1 \perp a_1$: To see that $a_1 \perp b_1$ and $a_1 \perp c_1$, note that a_1 is orthogonal to

$H_{K,L}(v_1)$, as we obtain a_1 by subtracting from v_1 its orthogonal projection onto that hyperplane, so that what is left must be orthogonal to it. Now b_1 and c_1 are both in $H_{K,L}(v_1)$ so it follows that a_1 must be orthogonal to both. That $b_1 \perp c_1$ follows from the fact that $H_{K,L}(v_1)$ can further be decomposed as the sum of two orthogonal subspaces, L and $\langle\{v_2 - P_L(v_2)\}\rangle$, where b_1 resides in the former and c_1 in the latter.

Similarly, we can define $\eta_2 := \frac{\langle v_2, v_1 - P_L(v_1) \rangle}{\|v_1 - P_L(v_1)\|^2}$, $a_2 := v_2 - P_L(v_2) - \eta_2(v_1 - P_L(v_1))$, $b_2 := P_L(v_2)$ and $c_2 := \eta_2(v_1 - P_L(v_1))$, and by almost identical reasoning find that $v_2 = a_2 + b_2 + c_2$ and $a_1 \perp b_2 \perp c_2 \perp a_2$. Note that the bounds in the hypothesis can now be written as: $\|a_1\| \geq \delta$ and $\|a_2\| \geq \delta$. On the other hand it is at least in principle possible to have $b_i = 0$ or $c_i = 0$ ($i = 1$ or 2).

The crucial observation based on these computations is that $a_1 \perp (a_2 + c_2)$ and $a_2 \perp (a_1 + c_1)$. To show that the first relation holds, consider first the case when $c_1 \neq 0$: From the equation $c_1 = \eta_1(v_2 - P_L(v_2)) = \eta_1(v_2 - b_2) = \eta_1(a_2 + c_2)$ we get that c_1 is collinear with the non-zero vector $a_2 + c_2$, and since $a_1 \perp c_1$, the desired conclusion follows. On the other hand if $c_1 = 0$, we must have that $\eta_1 = 0$ (since $\|a_2 + c_2\| \geq \|a_2\| > 0$), which is equivalent to the condition that $\langle v_1, v_2 - b_2 \rangle = 0$. Using the fact that $\langle a_1, b_2 \rangle = 0$ (because $b_2 \in L$ and $a_1 \perp (\langle L \cup \{v_2\} \rangle) \supset L$), and that $\langle v_1, v_2 \rangle = 0$, we then have $0 = \langle v_1, v_2 - b_2 \rangle = \langle v_1, -b_2 \rangle = \langle a_1 + b_1, -b_2 \rangle = \langle b_1, -b_2 \rangle$. Thus we have that $\langle b_1, b_2 \rangle = 0$, and since $\langle a_2, b_1 \rangle = 0 = \langle c_2, b_1 \rangle$ (as both a_2 and c_2 are orthogonal to L), we get that $\langle b_1, v_2 \rangle = \langle b_1, a_2 + b_2 + c_2 \rangle = 0$. Thus we finally get $0 = \langle v_1, v_2 \rangle = \langle a_1 + b_1, a_2 + b_2 + c_2 \rangle = \langle a_1, a_2 + b_2 + c_2 \rangle = \langle a_1, a_2 + c_2 \rangle$, which is the desired conclusion. By similar reasoning we obtain $a_2 \perp (a_1 + c_1)$.

We are now in a position to lower-bound $\text{dist}(u, H_{K,L}(u))$. In order to avoid cluttering, we will denote the expression $\frac{\langle u, u^\perp - P_L(u^\perp) \rangle}{\|u^\perp - P_L(u^\perp)\|^2}$ by η in the third row of the com-

putation. We have:

$$\begin{aligned}
dist(u, H_{K,L}(u)) &= \|u - P_{H_{K,L}(u)}(u)\| = \|u - P_{\langle L \cup \{u^\perp\} \rangle}(u)\| \\
&= \|u - P_{\langle L \cup \{u^\perp - P_L(u^\perp)\} \rangle}(u)\| \\
&= \left\| u - P_L(u) - \left\langle u, \frac{u^\perp - P_L(u^\perp)}{\|u^\perp - P_L(u^\perp)\|} \right\rangle \frac{u^\perp - P_L(u^\perp)}{\|u^\perp - P_L(u^\perp)\|} \right\| \\
&= \left\| u - P_L(u) - \frac{\langle u, u^\perp - P_L(u^\perp) \rangle}{\|u^\perp - P_L(u^\perp)\|^2} (u^\perp - P_L(u^\perp)) \right\| \\
&= \|u - P_L(u) - \eta(u^\perp - P_L(u^\perp))\| \\
&= \|\alpha_1 v_1 + \alpha_2 v_2 - P_L(\alpha_1 v_1 + \alpha_2 v_2) - \eta(-\alpha_2 v_1 + \alpha_1 v_2 - P_L(-\alpha_2 v_1 + \alpha_1 v_2))\| \\
&= \|\alpha_1 v_1 + \alpha_2 v_2 - \alpha_1 P_L(v_1) - \alpha_2 P_L(v_2) + \eta(\alpha_2(v_1 - P_L(v_1)) - \alpha_1(v_2 - P_L(v_2)))\| \\
&= \|\alpha_1(v_1 - b_1) + \alpha_2(v_2 - b_2) + \eta(\alpha_2(v_1 - b_1) - \alpha_1(v_2 - b_2))\| \\
&= \|\alpha_1(a_1 + c_1) + \alpha_2(a_2 + c_2) + \eta(\alpha_2(a_1 + c_1) - \alpha_1(a_2 + c_2))\|.
\end{aligned}$$

Now, based on the above we have on the one hand that $dist(u, H_{K,L}(u))$ is the norm of a vector that is the sum of a scalar multiple of a_1 and some vectors that are orthogonal to it (namely scalar multiples of c_1 and $a_2 + c_2$), which allows us write $dist(u, H_{K,L}(u)) \geq \|\alpha_1 a_1 + \eta \alpha_2 a_1\| \geq |\alpha_1 + \eta \alpha_2| \delta$. On the other hand we may regard $dist(u, H_{K,L}(u))$ as the norm of a vector that is the sum of a scalar multiple of a_2 and some vectors that are orthogonal to it, so that $dist(u, H_{K,L}(u)) \geq |\alpha_2 - \eta \alpha_1| \delta$.

All that is left to show now is that $M = \max\{|\alpha_1 + \eta \alpha_2|, |\alpha_2 - \eta \alpha_1|\} \geq \frac{1}{2\sqrt{2}}$. Recall that we are working under the assumption that $\alpha_1, \alpha_2 > 0$. Consider first the case where $\alpha_1 \geq \alpha_2$. If $\eta \geq 0$ then $M \geq \alpha_1$. If $\eta < 0$ we consider two sub-cases: If $-\eta \alpha_2 \geq \alpha_1/2$, meaning $-\eta \geq \frac{\alpha_1}{2\alpha_2}$, then $|\alpha_2 - \eta \alpha_1| \geq |\alpha_2 + \frac{\alpha_1}{2\alpha_2} \alpha_1| \geq |\alpha_2 + \alpha_1/2| \geq \alpha_1/2$; on the other hand if $-\eta \alpha_2 \in (0, \alpha_1/2)$, then $\eta \alpha_2 > -\alpha_1/2$ and hence $|\alpha_1 + \eta \alpha_2| > |\alpha_1 - \alpha_1/2| = \alpha_1/2$. Thus if $\alpha_1 \geq \alpha_2$ then $M \geq \alpha_1/2$. Now consider the case when $\alpha_2 > \alpha_1$. If $\eta \leq 0$ then $|\alpha_2 - \eta \alpha_1| \geq \alpha_2$. If $\eta > 0$, consider first the sub-case when $\eta \geq \frac{1}{2}$: Then $|\alpha_1 + \eta \alpha_2| \geq \alpha_2/2$. Finally, if $\eta \in (0, \frac{1}{2})$, we have

$|\alpha_2 - \eta\alpha_1| \geq |\alpha_2 - \frac{1}{2}\alpha_1| \geq |\alpha_2 - \alpha_2/2| = \alpha_2/2$. Combining the conclusions of all cases, we have that $M \geq \max\{\alpha_1, \alpha_2\}/2$. Since $\alpha_1^2 + \alpha_2^2 = 1$, at least one of the α_i 's must be $\geq \frac{1}{\sqrt{2}}$, so we may finally conclude: $\text{dist}(u, H_{K,L}(u)) \geq M\delta \geq \frac{\delta}{2\sqrt{2}}$, which completes the proof of the base case.

Induction step. Assume that for some $k_0 \geq 2$ the lemma holds for all $k = 1, \dots, k_0$. Let $k = k_0 + 1$, and let K be a k -dimensional vector space with orthonormal basis $\{v_1, \dots, v_k\}$. Let L be a vector space of dimension $m - k$ such that $L \cap K = \{0\}$, and assume that $\text{dist}(v_i, H_{K,L}(v_i)) \geq \delta$ for all $i \in [k]$, for some fixed $\delta \in (0, 1]$.

Let $u \in K$ be a unit vector, and write $u = \alpha_1 v_1 + \dots + \alpha_{k_0+1} v_{k_0+1}$, where $\sum_{i=1}^k \alpha_i^2 = 1$. We may assume that $u \neq \pm v_i$ for all i , since the desired inequality holds trivially for the basis vectors, and hence that $|\alpha_i| < 1$ for all i . Then $u' = \frac{1}{\sqrt{1-\alpha_1^2}}(\alpha_2 v_2 + \dots + \alpha_{k_0+1} v_{k_0+1})$ is a unit vector in $K' = \langle \{v_2, \dots, v_{k_0+1}\} \rangle$, and together with $\langle L \cup \{v_1\} \rangle$, this pair of subspaces satisfies the conditions of the lemma with $\dim(K') = k_0$ and $\text{dist}(v_i, H_{K', \langle L \cup \{v_1\} \rangle}(v_i)) \geq \delta$, for $i = 2, \dots, k$ (to see this, note that $H_{K', \langle L \cup \{v_1\} \rangle}(v_i) = \langle C_{K'}^\perp(v_i) \cup \{v_1\} \cup L \rangle = \langle C_K^\perp(v_i) \cup L \rangle = H_{K,L}(v_i)$). Thus, by the induction hypothesis, we have that $\text{dist}(u', H_{K', \langle L \cup \{v_1\} \rangle}(u')) \geq \frac{\delta}{(2\sqrt{2})^{k_0-1}}$.

Now let $K'' = \langle \{v_1, u'\} \rangle$ and $L'' = \langle L \cup C_{K'}^\perp(u') \rangle$. Clearly $\{v_1, u'\}$ is an orthonormal basis of K'' and $u = \alpha_1 v_1 + \sqrt{1-\alpha_1^2} u'$ is a unit vector in K'' . Also, $\dim(L'') = m - 2$, because $\dim(\langle K'' \cup L'' \rangle) = \dim(\langle K \cup L \rangle) = m$ so that $\dim(L'') \geq m - 2$, and because $\dim(C_{K'}^\perp(u')) = k - 2$ so that $\dim(L'') \leq \dim(L) + k - 2 = (m - k) + k - 2$. Now $H_{K'', L''}(u') = \langle C_{K''}^\perp(u') \cup L'' \rangle = \langle \{v_1\} \cup (C_{K'}^\perp(u') \cup L) \rangle = \langle C_{K'}^\perp(u') \cup (\{v_1\} \cup L) \rangle = H_{K', \langle L \cup \{v_1\} \rangle}(u')$, and $H_{K'', L''}(v_1) = \langle \{u'\} \cup (C_{K'}^\perp(u') \cup L) \rangle = \langle K' \cup L \rangle = \langle C_K^\perp(v_1) \cup L \rangle = H_{K,L}(v_1)$. Thus we get the following conditions on the distances of the points u' and v_1 from their assigned hyperplanes: $\text{dist}(u', H_{K'', L''}(u')) = \text{dist}(u', H_{K', \langle L \cup \{v_1\} \rangle}(u')) \geq$

$\frac{\delta}{(2\sqrt{2})^{k_0-1}}$ and $\text{dist}(v_1, H_{K'',L''}(v_1)) = \text{dist}(v_1, H_{K,L}(v_1)) \geq \delta \geq \frac{\delta}{(2\sqrt{2})^{k_0-1}}$. Finally, applying our base case reasoning to the pair of subspaces (K'', L'') with the above lower bound, we may conclude that $\text{dist}(u, H_{K,L}(u)) \geq \frac{1}{2\sqrt{2}} \frac{\delta}{(2\sqrt{2})^{k_0-1}} = \frac{\delta}{(2\sqrt{2})^{k_0}}$, as required. $\square$

The generality of Lemma 4.4.1 (the fact that it holds for any orthonormal basis of the subspace K) is paid for with a factor of $(2\sqrt{2})^{-(k-1)}$ in the inequality. In relation to Conjecture 4.2.3, where $k = d = \dim(D)$, this factor is not significant, considering the lower bound we are aiming for. However, for a well-chosen basis of K , the distance of u from $H_{K,L}(u)$ can be computed exactly. In order to do this, we first review the notion of *principal angles between subspaces*, first introduced by Jordan (see also [35]).

Given two subspaces K and L of $\mathbb{R}^m$, with $k = \dim(K) \leq \dim(L)$, let the pair of unit vectors (v_1, w_1) , with $v_1 \in K$ and $w_1 \in L$, be such that $\langle v_1, w_1 \rangle = \max\{\langle v, w \rangle \mid v \in K, w \in L, \|v\| = \|w\| = 1\}$. For $i = 2, \dots, k$, let the pair (v_i, w_i) of unit vectors with $v_i \in C_K^\perp(\{v_1, \dots, v_{i-1}\})$ and $w_i \in C_L^\perp(\{w_1, \dots, w_{i-1}\})$ be such that

$$\langle v_i, w_i \rangle = \max \left\{ \frac{\langle v, w \rangle}{\|v\| \|w\|} \mid v \in C_K^\perp(\{v_1, \dots, v_{i-1}\}), w \in C_L^\perp(\{w_1, \dots, w_{i-1}\}) \right\}.$$

The angles formed by the pairs of vectors (v_i, w_i) are called the *principal angles between K and L* . Note that while the vector pairs chosen in this way are generally not unique, the sequence $\langle v_1, w_1 \rangle, \dots, \langle v_k, w_k \rangle$ is the same regardless of the choice of vectors. The set of vectors $\{v_1, \dots, v_k\}$ form an orthonormal basis of K , and we will call any orthonormal basis of K whose vectors can be chosen in this way a *principal basis of K with respect to L* . Based on the construction of the two bases one can check that for all $i \in [k]$ we have $P_L(v_i) = \langle v_i, w_i \rangle w_i$, and from this it follows that $\forall i, j \in [k], i \neq j$: $(v_i - P_L(v_i)) \perp v_j$ and $(v_i - P_L(v_i)) \perp (v_j - P_L(v_j))$. Since $H_{K,L}(v_1) = \langle C_K^\perp(v_1) \cup L \rangle = \langle \{v_2 - P_L(v_2), \dots, v_k - P_L(v_k)\} \cup L \rangle$, and $v_1 \perp \langle \{v_2 - P_L(v_2), \dots, v_k - P_L(v_k)\} \rangle \perp L$, we have that $\text{dist}(v_1, H_{K,L}(v_1)) = \|v_1 - P_{H_{K,L}(v_1)}(v_1)\| = \|v_1 - P_L(v_1)\| = \text{dist}(v_1, L)$,

and it is clear that such an equality holds for the other v_i 's as well. From these facts it is also not difficult to see that the vectors of a principal basis of K with respect to L satisfy the properties that $\text{dist}(v_1, L) = \min\{\text{dist}(v, L) \mid v \in K, \|v\| = 1\}$ and that $\text{dist}(v_i, L) = \min\{\text{dist}(v, L) \mid v \in K, v \perp v_j \text{ for } j < i, \|v\| = 1\}$ for $i > 1$ (in fact these properties provide an alternative definition for the principal angles).

Returning to the more general context of the subspaces K and L , given a principal basis $\{v_1, \dots, v_k\}$ of K with respect to L , the next theorem provides a formula for $\text{dist}(u, H_{K,L}(u))$ for any unit vector $u \in K$, expressed in terms of the values $d_i = \text{dist}(v_i, L) = \text{dist}(v_i, H_{K,L}(v_i))$.

Theorem 4.4.2. Let K, L be two subspaces of $\mathbb{R}^m$ such that $\dim(L) \geq k = \dim(K) \geq 2$, $K \cap L = \{0\}$ and $\langle K \cup L \rangle = \mathbb{R}^m$.

Then there is a basis $\{v_1, \dots, v_k\}$ of K such that if $d_i = \text{dist}(v_i, L) = \|v_i - P_L(v_i)\|$ for all $i \in [k]$, then for any unit vector $u = \alpha_1 v_1 + \dots + \alpha_k v_k$ we have:

$$\text{dist}(u, H_{K,L}(u)) = \frac{1}{\sqrt{\left(\frac{\alpha_1}{d_1}\right)^2 + \dots + \left(\frac{\alpha_k}{d_k}\right)^2}}.$$

Proof. Let $\{v_1, \dots, v_k\}$ be a principal basis of K with respect to L . We make use of the facts that with the vectors v_i defined as in the hypothesis, $\forall i, j$ with $i \neq j$ we have $\langle v_i - P_L(v_i), v_j - P_L(v_j) \rangle = 0$ and $\langle v_i - P_L(v_i), P_L(v_j) \rangle = 0$, and hence for all $i \in [k]$,

$$\text{dist}(v_i, H_{K,L}(v_i)) = \text{dist}(v_i, \langle L \cup \{v_j - P_L(v_j) \mid j \neq i\} \rangle) = \text{dist}(v_i, L) = d_i.$$

We first treat the case where $k = 2$. We may write $u = \alpha_1 v_1 + \alpha_2 v_2$, with $\alpha_1^2 + \alpha_2^2 = 1$, and take $u^\perp = -\alpha_2 v_1 + \alpha_1 v_2$. We are interested in $\text{dist}(u, H_{K,L}(u)) = \text{dist}(u, \langle L \cup \{u^\perp\} \rangle) = \text{dist}(u, \langle L \cup \{u^\perp - P_L(u^\perp)\} \rangle)$. Note however that $u^\perp - P_L(u^\perp) = -\alpha_2(v_1 -$

$P_L(v_1)) + \alpha_1(v_2 - P_L(v_2))$ and hence

$$\begin{aligned} \text{dist}(u, H_{K,L}(u)) &= \text{dist}((u - P_L(u)) + P_L(u), \langle L \cup \{u^\perp - P_L(u^\perp)\} \rangle) \\ &= \text{dist}(u - P_L(u), \langle L \cup \{u^\perp - P_L(u^\perp)\} \rangle) \\ &= \text{dist}(u - P_L(u), \langle \{u^\perp - P_L(u^\perp)\} \rangle). \end{aligned}$$

For any vector $w \in K$, denote $\tilde{w} := w - P_L(w)$ (but note that by $\tilde{u}^\perp$ we mean $u^\perp - P_L(u^\perp)$). To simplify the next computation, we first calculate:

$$|\langle \tilde{u}, \tilde{u}^\perp \rangle| = |\langle \alpha_1 \tilde{v}_1 + \alpha_1 \tilde{v}_2, -\alpha_2 \tilde{v}_1 + \alpha_1 \tilde{v}_2 \rangle| = \alpha_1 \alpha_2 (-d_1^2 + d_2^2),$$

where we have used the facts that $\langle \tilde{v}_1, \tilde{v}_2 \rangle = 0$ and that $d_1 \leq d_2$ (and hence that $d_2^2 - d_1^2 \geq 0$). Now we have:

$$\begin{aligned} \text{dist}(\tilde{u}, \langle \{\tilde{u}^\perp\} \rangle) &= \|\tilde{u} - P_{\langle \{\tilde{u}^\perp\} \rangle}(\tilde{u})\| \\ &= \sqrt{\|\tilde{u}\|^2 - \left(\left\langle \tilde{u}, \frac{\tilde{u}^\perp}{\|\tilde{u}^\perp\|} \right\rangle \right)^2} = \frac{\sqrt{\|\tilde{u}\|^2 \|\tilde{u}^\perp\|^2 - \langle \tilde{u}, \tilde{u}^\perp \rangle^2}}{\|\tilde{u}^\perp\|} \\ &= \frac{\sqrt{(\alpha_1^2 d_1^2 + \alpha_2^2 d_2^2)(\alpha_2^2 d_1^2 + \alpha_1^2 d_2^2) - \alpha_1^2 \alpha_2^2 (d_1^4 + d_2^4 - 2d_1^2 d_2^2)}}{\sqrt{\alpha_2^2 d_1^2 + \alpha_1^2 d_2^2}} \\ &= \frac{\sqrt{d_1^2 d_2^2 (\alpha_1^4 + 2\alpha_1^2 \alpha_2^2 + \alpha_2^4)}}{\sqrt{\alpha_2^2 d_1^2 + \alpha_1^2 d_2^2}} = \frac{\sqrt{d_1^2 d_2^2}}{\sqrt{\alpha_2^2 d_1^2 + \alpha_1^2 d_2^2}} = \frac{1}{\sqrt{\left(\frac{\alpha_1}{d_1}\right)^2 + \left(\frac{\alpha_2}{d_2}\right)^2}}. \end{aligned}$$

For the general case we show that if the distance formula holds when $k = k_0$, for some $k_0 \geq 2$, then it also holds for $k = k_0 + 1$. In this case we have $\dim(K) = k = k_0 + 1$, $\dim(L) = m - k_0 - 1$, and the basis $\{v_1, \dots, v_k\}$ of K defined as above. Let $u = \alpha_1 v_1 + \dots + \alpha_k v_k$ be a unit vector.

The crucial fact now is that $v_i \perp (v_j - P_L(v_j))$ for all $i, j \in [k]$ with $i \neq j$. Since $\langle L \cup \{v_k\} \rangle = \langle L \cup \{\tilde{v}_k\} \rangle$, it then follows that for all $i, j \in [k]$, $i \neq j$, that $\tilde{v}_i = v_i - P_L(v_i) = v_i - P_{\langle L \cup \{v_j\} \rangle}(v_i)$. Define $K' = \langle \{v_1, \dots, v_{k_0}\} \rangle$ and $L' = \langle L \cup \{v_k\} \rangle$. We may assume that $|\alpha_k| < 1$ (otherwise rearrange the v_i 's so that this holds), and we define $u' = \frac{1}{\sqrt{1-\alpha_k^2}}(\alpha_1 v_1 + \dots + \alpha_{k_0} v_{k_0})$. By the induction hypothesis (applied to

K' , L' and the basis $\{v_1, \dots, v_{k_0}\}$ of K), we have that

$$\text{dist}(u', H_{K', L'}(u')) = \frac{1}{\sqrt{\frac{1}{(1-\alpha_k^2)} \left(\left(\frac{\alpha_1}{d_1} \right)^2 + \dots + \left(\frac{\alpha_{k_0}}{d_{k_0}} \right)^2 \right)}}.$$

Now define $L'' = \langle L \cup C_{K'}^\perp(u') \rangle$ and $K'' = \langle \{v_k, u'\} \rangle$. Since $v_k \perp \tilde{v}_i$ for all $i \in [k_0]$, it follows that $v_k \perp P_L(K')$ and that $\text{dist}(v_k, L'') = \|v_k - P_{L''}(v_k)\| = \|v_k - P_{\langle L \cup K' \rangle}(v_k)\| = \|v_k - P_{\langle L \cup \{\tilde{v}_1, \dots, \tilde{v}_{k_0}\} \rangle}(v_k)\| = \|v_k - P_L(v_k)\| = d_k$. Then

$$\begin{aligned} \text{dist}(u', H_{K'', L''}(u')) &= \frac{1}{\sqrt{\frac{1}{(1-\alpha_k^2)} \left(\left(\frac{\alpha_1}{d_1} \right)^2 + \dots + \left(\frac{\alpha_{k_0}}{d_{k_0}} \right)^2 \right)}} \\ &\leq \frac{1}{\sqrt{\frac{1}{(1-\alpha_k^2)} \left(\left(\frac{\alpha_1}{d_k} \right)^2 + \dots + \left(\frac{\alpha_{k_0}}{d_k} \right)^2 \right)}} = \frac{1}{\sqrt{\frac{\alpha_1^2 + \dots + \alpha_{k_0}^2}{1-\alpha_k^2} \left(\frac{1}{d_k} \right)^2}} = d_k, \end{aligned}$$

and since $v_k \perp u'$, we now have:

$$\begin{aligned} \min\{\text{dist}(v, L'') \mid v \in K'', \|v\| = 1\} &= \min\{\|v - P_{L''}\| \mid v \in K'', \|v\| = 1\} \\ &= \min\{\sqrt{\beta^2 \text{dist}(u', L'')^2 + (1 - \beta^2) \text{dist}(v_k, L'')^2} \mid \beta \in [0, 1], \} \\ &= \min\{\sqrt{\beta^2 \text{dist}(u', L'')^2 + (1 - \beta^2) d_k^2} \mid \beta \in [0, 1], \} \\ &\geq \text{dist}(u', L''). \end{aligned}$$

Since the last value in the above sequence of inequalities is actually attained when $\beta = 1$, the desired conclusion that $\text{dist}(u', L'') = \min\{\text{dist}(v, L'') \mid v \in K'', \|v\| = 1\}$ follows.

Recalling that $u = \alpha_k v_k + \sqrt{1 - \alpha_{k+1}^2} u'$, we finally apply the argument for the case when $k = 2$ to the subspaces K'' and L'' with the basis $\{u', v_k\}$ for K'' :

$$\begin{aligned} \text{dist}(u, H_{K, L}(u)) &= \text{dist}(u, H_{K'', L''}(u)) \\ &= \frac{1}{\sqrt{\left(\frac{\sqrt{1-\alpha^2}}{\text{dist}(u', H_{K'', L''}(u'))} \right)^2 + \left(\frac{\alpha_k}{\text{dist}(v_k, H_{K'', L''}(v_k))} \right)^2}} \\ &= \frac{1}{\sqrt{\left(\frac{\alpha_1}{d_1} \right)^2 + \dots + \left(\frac{\alpha_{k_0}}{d_{k_0}} \right)^2 + \left(\frac{\alpha_k}{d_k} \right)^2}}, \end{aligned}$$

which is the desired conclusion. $\square$

Remark 4.4.3. From the formula for $\text{dist}(u, H_{K,L}(u))$, and because $d_1 \leq \dots \leq d_d$ by definition, it is clear that d_1 is the smallest value that this expression can attain, and that it is attained when $\alpha_1 = 1$ and $\alpha_i = 0$ for all $i \geq 2$.

Remark 4.4.4. If D is a d -dimensional subspace in general position (recall that $d \ll m$ so that $d \ll m - d$), and $E = \bigcap_{i \in I} H_i$ an $(m - d)$ -dimensional canonical subspace ($I = \{i_1, \dots, i_d\}$, $|I| = d$), then the principal basis of D with respect to E is closely related to the d -minor on the set of rows I of the matrix A (whose columns form an orthonormal basis of D), as follows: Let $\{v_1, \dots, v_d\}$ be the principal basis of D with respect to E . We choose A to be the $(m \times d)$ -matrix $(v_1, \dots, v_d)$, and let $B = A[I, [d]]$. Because E is a canonical subspace, we have for each $i \in [d]$ that $P_E(v_i)$ is the same as the vector v_i , except with 0 on the positions $i_1, \dots, i_d$. Therefore for each $i \in [d]$ the i -th column of the matrix B is the non-zero part of $v_i - P_E(v_i)$, and hence $\|B_{\cdot, i}\| = \|v_i - P_E(v_i)\| = \text{dist}(v_i, E)$. Now by the properties mentioned in the last paragraph, the vectors $v_i - P_E(v_i)$, $i \in [d]$, are pairwise orthogonal, and hence so are the columns of B . Setting $d_i = \text{dist}(v_i, E)$ for all i , we then have $\prod_{i \in [d]} d_i = |\det B| = |\det(A[I, [d]])|$. The last equation also provides a lower bound for the absolute value of a d -minor in terms of the distance of a unit vector from a corresponding extension hyperplane: $(d_1)^d \leq |\det(A[I, [d]])|$.

We can now prove the main result in this section. The next theorem states that if *any* orthonormal basis of a subspace is such that each basis vector is far from all of its extension hyperplanes, then the d -minors of any matrix whose column vectors form an orthonormal basis of the subspace are lower-bounded in absolute value.

Theorem 4.4.5. Let D be a d -dimensional subspace in general position in $\mathbb{R}^m$. Let $\epsilon > 0$ be given.

If there is an orthonormal basis $\{v_1, \dots, v_d\}$ of D such that for each $i \in [d]$ and all $E \in \mathcal{E}_{m,m-d}$ we have that $\text{dist}(v_i, H_{D,E}(v_i)) \geq \epsilon$, then every d -minor of the matrix $A = (v_1, \dots, v_d)$ is, in absolute value, $\geq \left(\frac{\epsilon}{(2\sqrt{2})^{d-1}} \right)^d$.

Proof. Let $E = \bigcap_{i \in I} H_i \in \mathcal{E}_{m,m-d}$ be fixed. By Lemma 4.4.1 we have for all unit vectors $u \in D$ that $\text{dist}(u, H_{D,E}(u)) \geq \frac{\epsilon}{(2\sqrt{2})^{d-1}}$, and in particular, the vector with the shortest distance from E among all the unit vectors in D has at least this distance from E . But we know from Remark 4.4.4 that $|\det(A[I, [d]])|$ is greater than the d -th power of the smallest distance of a unit vector from E . Thus the conclusion follows. $\square$

To summarise, a subspace D is close to some $E \in \mathcal{E}_{m,m-d}$ if and only if for some basis of D at least one basis vector v is close to $H_{D,E}(v)$ (Lemma 4.4.1), and moving such a vector v away from $H_{D,E}(v)$ (in other words, replacing it with a vector that has greater distance from the same extension hyperplane) can not cause any other basis vector v' to get closer to $H_{D,E}(v')$ (Proposition 4.3.9). Therefore, if every basis vector is situated in a sufficiently “large” region of its own extension arrangement, then the whole subspace can be balanced without changing its signature set. In fact, it would suffice to show that: for every $E \in \mathcal{E}_{m,m-d}$, if D is close to E , then there is a unit vector $v \in D$ which can be moved away from $H_{D,E}(v)$ without having it cross any of its extension hyperplanes. This would show that the conditions of Theorem 4.4.5 can be satisfied with sufficiently large ϵ , for any given starting subspace.

In the remainder of this section we improve the previous theorem slightly, by showing that it is not necessary to consider *all* extension hyperplanes. We show that there is a strict subset $S \subset \mathcal{E}_{m,m-d}$ such that for every unit vector $v \in D$, the region $R(v)$ of $\mathcal{H}_D(v)$ is formed by extension hyperplanes associated with canonical subspaces from S . This is less of theoretical interest, and mostly useful when computing with exten-

sion arrangements, as it can help reduce the number of hyperplanes that need to be considered.

First we need a lemma which allows us to assume in most cases that no two hyperplanes of an extension arrangement coincide as sets.

Lemma 4.4.6. If $D \subset \mathbb{R}^m$ is in general position, with $d = \dim(D)$ and $d \geq 2$, then for almost all $v \in D$, the arrangement $\mathcal{H}_D(v)$ contains $\binom{m}{d}$ distinct hyperplanes. In other words, the extension hyperplanes of v are pairwise distinct for almost all vectors v in D .

Proof. We show first that $\mathcal{H}_D(v)$ contains coinciding hyperplanes if and only if $v \in C_D^\perp(D \cap E)$ for some canonical subspace E of dimension $> m - d$.

“ $\Leftarrow$ ”: If $k \in [d - 1]$ and E is an canonical subspace of dimension $m - d + k$, then since D is in general position, we have that $K = D \cap E$ has dimension k . It follows then that for any $v \in C_D^\perp(K)$ (which is a subspace of dimension $d - k$), we have that $K \subseteq C_D^\perp(v)$, and hence for any $(m - d)$ -dimensional canonical subspace $E' \subset E$ it holds that $H_{D,E'}(v) = \langle C_D^\perp(v) \cup E' \rangle = \langle C_D^\perp(v) \cup K \cup E' \rangle = \langle C_D^\perp(v) \cup E \rangle$, where the equality $\langle K \cup E' \rangle = E$ holds because $C_D^\perp(v) \cap E' = \{0\}$ and $\dim(E) = \dim(K) + \dim(E') - \dim(K \cap E')$. In other words for every E' among the $\binom{m-d+k}{m-d}$ canonical subspaces of E that have dimension $m - d$, $H_{D,E'}(v)$ is the same set.

“ $\Rightarrow$ ”: Conversely, if $v \in D$, and E', E'' are two canonical subspaces of dimension $m - d$ such that $H_{D,E'}(v) = H_{D,E''}(v)$, then we have that $\langle C_D^\perp(v) \cup E' \rangle = \langle C_D^\perp(v) \cup \langle E' \cup E'' \rangle \rangle$, and hence that $\dim(C_D^\perp(v)) + \dim(\langle E' \cup E'' \rangle) - \dim(C_D^\perp(v) \cap \langle E' \cup E'' \rangle) = m - 1 \Leftrightarrow d + \dim(\langle E' \cup E'' \rangle) - \dim(C_D^\perp(v) \cap \langle E' \cup E'' \rangle) = m$. On the other hand it is also true, due to D being in general position, that $d + \dim(\langle E' \cup E'' \rangle) - \dim(D \cap \langle E' \cup E'' \rangle) = m$, and we may conclude that $D \cap \langle E' \cup E'' \rangle = C_D^\perp(v) \cap \langle E' \cup E'' \rangle$, which in turn implies that for $K = D \cap \langle E' \cup E'' \rangle$ we have $K \subseteq C_D^\perp(v)$ so that $v \in C_D^\perp(K)$.

Now let S be the set of vectors in D which do *not* lie in $C_D^\perp(D \cap E)$ for any canonical subspace E of dimension $> m - d$. The sets $C_D^\perp(D \cap E)$ have dimension at most $d - 1$, thus S is obtained by removing a finite number of at most $(d - 1)$ -dimensional subspaces from the d -dimensional subspace D . By the converse implication above the extension hyperplanes are pairwise distinct for all $v \in S$. $\square$

Remark 4.4.7. From the proof of Lemma 4.4.6 it can be seen that the vectors $v \in D$ for which $\mathcal{H}_D(v)$ has $\binom{m}{d}$ pairwise distinct hyperplanes are precisely the non-zero vectors for which $C_D^\perp(v)$ is in general position in $\mathbb{R}^m$.

As we have seen in the previous subsection, the shortest distance of any unit vector in D from any $E \in \mathcal{E}_{m,m-d}$, which is also the shortest distance of any unit vector from any of its extension hyperplanes, is an accurate measure of the quality of the subspace in the sense of Problem 2, because in terms of this value the smallest (in absolute value) d -minor of an orthonormal basis matrix of D can be both upper- and lower-bounded. However, any single unit vector $v \in D$ will be closest, among all of its extension hyperplanes, to one that delimits the region $R(v)$ of $\mathcal{H}_D(v)$. Furthermore, since by Lemma 4.4.6 the set of unit vectors whose extension hyperplanes are pairwise distinct is dense in the unit sphere, it is easy to show via a continuity argument that if a delimiting hyperplane of $R(v)$ is not unique (if it coincides as a set with other extension hyperplanes), then among the coinciding extension hyperplanes which delimit it there must be at least one, say $H_{D,E'}(v)$, such that $H_{D,E'}(w)$ is a unique delimiting hyperplane of $R(w)$ for some $w \in D$. It is therefore of interest to study the following set:

Definition 4.4.8. Let D be a d -dimensional subspace in general position in $\mathbb{R}^m$. We define $del(D) := \{E \in \mathcal{E}_{m,m-d} \mid \exists v \in D \setminus \{0\} : H_{D,E}(v) \text{ is a unique delimiter of } R(v)\}$.

We have:

$$\begin{aligned}
& \min\{dist(v, E) \mid v \in D, \|v\| = 1, E \in \mathcal{E}_{m, m-d}\} \\
&= \min\{dist(v, H_{D,E}(v)) \mid v \in D, \|v\| = 1, E \in \mathcal{E}_{m, m-d}\} \\
&= \min\{\min\{dist(v, H_{D,E}(v)) \mid E \in \mathcal{E}_{m, m-d}\} \mid v \in D, \|v\| = 1\} \\
&= \min\{\min\{dist(v, H_{D,E}(v)) \mid E \in del(D)\} \mid v \in D, \|v\| = 1\} \\
&= \min\{dist(v, H_{D,E}(v)) \mid v \in D, \|v\| = 1, E \in del(D)\} \\
&= \min\{dist(v, E) \mid v \in D, \|v\| = 1, E \in del(D)\}.
\end{aligned}$$

This leads directly to the following

Theorem 4.4.9. Let D be a d -dimensional subspace in general position in $\mathbb{R}^m$, Let $\epsilon > 0$ be given.

If there is an orthonormal basis $\{v_1, \dots, v_d\}$ of D such that for each $i \in [d]$ and all $E \in del(D)$ we have that $dist(v_i, H_{D,E}(v_i)) \geq \epsilon$, then every d -minor of the matrix $A = (v_1, \dots, v_d)$ is, in absolute value, $\geq \left(\frac{\epsilon}{(2\sqrt{2})^{d-1}}\right)^d$.

Proof. The theorem follows directly from the preceding remark and Theorem 4.4.5. □

Up to this point we only know that $m - d + 1 \leq |del(D)| \leq \binom{m}{d}$, where the first inequality follows from Lemma 4.3.8 and the second one is trivial. In what follows we will only improve these inequalities to $m \leq |del(D)| \leq O(m^{d-1})$, but we will show that $del(D)$ is a subset of a set for which deciding membership is computationally (comparatively) efficient.

If $m > d$ and an arrangement of m hyperplanes in general position in $\mathbb{R}^d$ has a region that is delimited by exactly d hyperplanes, then this region and its opposite are the only regions of the arrangement which are delimited by precisely these d hyperplanes. To see this, it is enough to consider the case when $m = d + 1$, and the

first d of the arrangement hyperplanes are the canonical hyperplanes of $\mathbb{R}^d$, while the final arrangement hyperplane is some hyperplane that is in general position in $\mathbb{R}^d$. In this case the $(d+1)$ -st arrangement hyperplane intersects all orthants (which are also the regions of the sub-arrangement formed by the first d arrangement hyperplanes), except the two which contain its normal vectors. Therefore the two opposite orthants containing the normal vectors of the last arrangement hyperplane are the only regions of the arrangement which are delimited only by the first d hyperplanes. Thus to any choice of d hyperplanes of an arrangement (of more than d hyperplanes) in general position, there corresponds at most one region of the arrangement which is delimited by exactly these d hyperplanes. In terms of subspaces and induced arrangements we would say that to any $E = \bigcap_{i \in I} H_i \in \mathcal{E}_{m, m-d}$ there corresponds at most one pair of regions of the induced arrangement which are delimited precisely by the hyperplanes $D \cap H_i, i \in I$.

Definition 4.4.10. If $\mathcal{A}$ is an arrangement of m hyperplanes in $\mathbb{R}^d$, we say that a region R of $\mathcal{A}$ is a *simple* region if it is delimited by exactly d hyperplanes of $\mathcal{A}$. If $D \subset \mathbb{R}^m$ is a d -dimensional subspace, let $\text{sim}(D)$ be the set of $(m-d)$ -dimensional canonical subspaces $E = \bigcap_{i \in I} H_i$, such the induced arrangement of D has a simple region delimited by the hyperplanes $D \cap H_i, i \in I$. If A is a matrix whose column vectors form an orthonormal basis of D , we call a d -minor $\det(A[I, [d]])$ *simple* if $\bigcap_{i \in I} H_i \in \text{sim}(D)$.

The proof of Lemma 4.3.6 shows that $\text{del}(D) \subseteq \text{sim}(D)$, which implies that $\text{sim}(D) \neq \emptyset$, in other words that every arrangement in general position has simple regions (a fact which is not obvious when $m \gg d$). Since $|\text{sim}(D)|$ is at most $\frac{1}{2}f(m, d)$, meaning half of the number of regions of an arrangement of m hyperplanes in general position in $\mathbb{R}^d$, this implies an upper bound of $O(m^{d-1})$ on $|\text{del}(D)|$. The

next proposition improves the lower bound on $|del(D)|$ from $m - d + 1$ to m .

Proposition 4.4.11. If D is a d -dimensional subspace in general position in $\mathbb{R}^m$, with $m > d \geq 2$, then $|del(D)| \geq m$.

Proof. It is easy to show via a continuity argument that if a delimiting hyperplane of $R(v)$ is not unique, then among the coinciding extension hyperplanes which delimit it there must be at least one, say $H_{D,E'}(v)$, such that to $H_{D,E'}(w)$ is a unique delimiting hyperplane of $R(w)$ for some $w \in D \setminus \{0\}$. Therefore it holds for all $v \in D \setminus \{0\}$ that among the extension hyperplanes delimiting $R(v)$ there are $m - d + 1$, whose defining canonical subspaces are in $del(D)$. For each $i \in [m]$ let $h_i \in D$ be one of the two unit vectors such that $C_D^\perp(h_i) = D \cap H_i$. Since D is in general position in $\mathbb{R}^m$, $D \cap H_i$ must be in general position in H_i (regarded as $\mathbb{R}^{m-1}$), and it follows that for any $E \in \mathcal{E}_{m,m-d}$ such that $E \subseteq \mathbb{R}^{m-1}$ we have $H_{D,E}(h_i) = \langle C_D^\perp(h_i) \cup E \rangle = H_i$, since $C_D^\perp(h_i) \cap E = \{0\}$. Thus if $E \in \mathcal{E}_{m,m-d}$, and $e_i \notin E$, then $H_{D,E}(h_i) = H_i$.

Let $E_1, \dots, E_k$ be the pairwise distinct $(m-d)$ -dimensional canonical subspaces in $del(D)$. On the one hand, since every E_i is spanned by $m-d$ vectors of the canonical basis of $\mathbb{R}^m$, there must be $k(m-d)$ pairs (e_j, E_l) with $e_j \in E_l$. On the other hand, since for each $i \in [m]$ all E_l not containing e_i are such that $H_{D,E_l}(h_i) = H_i$, and the region $R(h_i)$ of $\mathcal{H}_D(h_i)$ must have at least $m-d$ other delimiting hyperplanes (whose defining canonical subspaces must contain e_i), we have that there must be at least $m(m-d)$ pairs (e_j, E_l) with $e_j \in E_l$. Thus $k \geq m$. $\square$

Corollary 4.4.12. If $\mathcal{A}$ is an arrangement of m hyperplanes in general position in $\mathbb{R}^d$, with $m > d \geq 2$, then $\mathcal{A}$ has at least m distinct pairs of opposite simple regions.

Finally, we mention briefly that parameters of arrangements of hyperplanes such as $|sim(D)|$ have been studied in Combinatorial Geometry. Given any arrangement $\mathcal{A}$ of m hyperplanes in $\mathbb{R}^d$, one can define the numbers p_k , $k \in \{d, \dots, m\}$, such that p_k is the

number of pairs of opposite regions of $\mathcal{A}$ which are delimited by exactly k hyperplanes of $\mathcal{A}$. We have $|\text{sim}(D)| = p_d$, where the arrangement is the induced arrangement of D . Fixing only m and d , one can also ask what the minimum and maximum values of p_k are, taken over all arrangements of m hyperplanes in $\mathbb{R}^d$ (possibly restricted to arrangements in general position). These numbers have been studied at least when $d = 3$, and much is known about this case, although even here there are important open questions. For $d = 3$ the statement of Corollary 4.4.12 appeared as far back as 1890 in [10]; see [15], chap. 18, for an overview of known facts in this case, including an upper bound of $\frac{1}{3}m(m-1)$ for p_3 when $m \geq 4$, as well as a table of known maximal values of this number (for different values of m). In general (when $d \geq 4$), upper bounds on p_d for arrangements in general position are of interest to us because they are also upper bounds on $|\text{del}(D)|$. On the other hand the study of $\text{del}(D)$ could lead to (not necessarily tight) lower bounds for p_d .

4.5 An algorithm for balancing subspaces

In this section we use the results from the previous sections to devise an algorithm for balancing any given subspace without changing its signature set.

Recall from Remark 4.2.6 that the “centre” of a region of an arrangement of hyperplanes exists even if we consider scaled Euclidean distances with a different scalar factor for every hyperplane. Since the set of vectors with which we can safely replace a given vector of a subspace is itself a region of an arrangement of hyperplanes, we may apply the generalisations of Lemma 4.2.4 from Remark 4.2.6 to this region, in order to find its centre. The next proposition shows that if we pick the scalar factors for each distance from a hyperplane in a clever way, we can get that the scaled distances are precisely the d -minors of an orthonormal basis matrix.

Proposition 4.5.1. Let D be a d -dimensional subspace in general position in $\mathbb{R}^m$, and let $v \in D \setminus \{0\}$ be given. Let $R(v)$ be the region of $\mathcal{H}_D(v)$ containing v .

There exists a vector $v' \in R(v)$ such that for any $v'' \in R(v)$, if A' and A'' are matrices whose columns form orthonormal bases of $o(D, v, v')$ and $o(D, v, v'')$, respectively, then

$$\min_{I \subset [m], |I|=d} |\det(A'[I, [d]])| \geq \min_{I \subset [m], |I|=d} |\det(A''[I, [d]])|,$$

and furthermore, at least $m - d + 1$ of the d -minors of A' are equal in absolute value with $\min_{I \subset [m], |I|=d} |\det(A'[I, [d]])|$.

Proof. Let $A \in \mathcal{M}_{m,d}(\mathbb{R})$ be such that the first column is the vector v , and the remaining columns form an orthonormal basis of $C_D^\perp(v)$.

First we compute an orthogonal vector, with a well chosen norm, for each of the extension hyperplanes of v . Let $E = \bigcap_{i \in I} H_i \in \mathcal{E}_{m,m-d}$, where $I = \{i_1, \dots, i_d\} \subseteq [m]$ and $|I| = d$. Any normal vector of $H_{D,E}(v)$ must have 0 on the positions in $[m] \setminus I$. Let $B = A[I, [d]]$. Then by the Laplace expansion along the first column we have that $\det B = \langle B_{\cdot,1}, h' \rangle$, where h' is the vector of co-factors $((-1)^{1+j} \det(B[[d] \setminus \{j\}, [d] \setminus \{1\}]))_{j \in [d]}$. If $h_{i_j}^E = h'_j$ for $j \in [d]$, and $h_k^E = 0$ for all $k \notin I$, we have that $h^E \perp H_{D,E}(v)$ and $\langle v, h^E \rangle = \det B$. Furthermore, with h^E defined in this way, we also have for every vector $w \in C^\perp(C_D^\perp(v))$ and matrix $A(w)$ obtained by replacing the first column of A with w , that $|\det(A(w)[I, [d]])| = |\langle w, h^E \rangle| = \|h^E\| \cdot \text{dist}(w, H_{D,E}(v))$.

We now apply the generalisations of Lemma 4.2.4 from Remark 4.2.6 to the region of $\text{ess}(\mathcal{H}_D(v))$ which contains v , but for each hyperplane $H_{D,E}(v)$ we use the scaled Euclidean distance with scalar factor $\|h^E\|$, and we obtain a unit vector $v' \in C^\perp(C_D^\perp(v))$ in the same region. By Remark 4.2.5, we may assume that the vector v' has the largest minimum scaled distance among all the unit vectors with the properties of the vector whose existence is proved in Lemma 4.2.4. Denote by A' the matrix obtained by

replacing the first column of A with the vector v' . Clearly the columns of A' form an orthonormal basis of $o(D, v, v')$. By Proposition 4.3.5, $o(D, v, v')$ and D have the same signature set. By Lemma 4.3.8 and the above relationship between the scaled distances and the absolute value of d -minors, we have that $m - d + 1$ of the d -minors of A' have the same absolute value, and that all other d -minors of A' are (in absolute value) larger than or equal to this minimum.

Finally, let v'' be in $R(v)$, let A'' be a matrix whose column vectors form an orthonormal basis of $o(D, v, v'')$, and assume that

$$\min_{I \subseteq [m], |I|=d} |\det(A'[I, [d]])| < \min_{I \subseteq [m], |I|=d} |\det(A''[I, [d]])|.$$

Let $I_1, \dots, I_l \subseteq [m]$ be all of the pairwise distinct sets with d elements and with the property that $|\det(A'[I_i, [d]])| = \min_{I \subseteq [m], |I|=d} |\det(A'[I, [d]])|$ for all $i \in [l]$ (thus $l \geq m - d + 1$). For each $i \in [l]$, let $E_i = \bigcap_{j \in I_i} H_j$. Then for all $i \in [l]$ we have $|\det(A'[I_i, [d]])| < |\det(A[I_i, [d]])|$. This implies that $|\langle h^{E_i}, v \rangle| > |\langle h^{E_i}, v' \rangle|$ for all $i \in [l]$. But from the way v' was chosen, it is clear that every other unit vector in the region of v must be closer to at least one of the hyperplanes whose normal vectors are h^{E_i} , with $i \in [l]$. Thus we may conclude that the inequality must hold. $\square$

The core of our algorithm is a procedure which, given the matrix A (whose columns span the subspace D) and a vector $v \in D$, computes the vector v' from Proposition 4.5.1 above. Then the properties of the subspace $o(D, v, v')$ and the matrix A' are clear from the proposition. We give a high-level description:

Description of sub-routine IMPROVE-SUBSPACE-ALONG-VECTOR(v, A):

- *Step 1.* Compute and store in the $(m \times d)$ -matrix C an orthonormal basis of D such that the first $d - 1$ columns of C form an orthonormal basis of $C_D^\perp(v)$, and the last column is v .

- *Step 2.* For each $I = \{i_1, \dots, i_d\} \subset [m]$ with $|I| = d$, compute a unit vector h^I such that $h_i^I = 0$ for all $i \in [m] \setminus I$, and for each $j \in [d]$, $h_{i_j}^I = (-1)^{j+1} \det C[I \setminus \{i_j\}, [d-1]]$.
- *Step 3.* Compute the unit vector $v' \in C^\perp(C_D^\perp(v))$ such that $\langle v, h^I \rangle \langle v', h^I \rangle > 0$ for all $I \subset [m]$ with $|I| = d$, and such that $\min_I |\langle v', h^I \rangle|$ is maximal (over all vectors w with $w \perp C_D^\perp(v)$ and $\langle v, h^I \rangle \langle w, h^I \rangle > 0$ for all $I \subset [m]$ with $|I| = d$).

Our algorithm for improving the quality of subspaces works as follows: A unit vector $v \in D$ is chosen at random and `IMPROVE-SUBSPACE-ALONG-VECTOR`(v, A) is called, after which A is set to a matrix whose columns form an orthonormal basis of $o(D, v, v')$ (and thus D is replaced with $o(D, v, v')$); this is repeated for some fixed number of iterations. The output is the final subspace D .

Behaviour and performance of the algorithm.

In practice the algorithm converges very rapidly, in the sense that the subspace, and hence the d -minors of an orthonormal basis matrix of the subspace, stabilise after relatively few iterations (usually < 50). It is not difficult to prove this convergence rigorously, using the fact that at every iteration the smallest (in absolute value) d -minor of A can only increase, and there are only finitely many d -minors, whose absolute values are upper-bounded by 1.

It is worth noting that the limit of the absolute value of the smallest d -minor does depend on the sequence of choices of vectors along which to improve each intermediate subspace. However, this is not an issue, since we are only interested in the best possible improvement of the subspace, and in practice can simply run the algorithm with the same starting subspace several times, and consider the best result as the final output. Furthermore, even for the worst sequence of choices of vectors along which to

improve the subspace, for the final subspace the smallest absolute value of a d -minor is comparable to the that of the best case, and is never arbitrarily close to 0.

To measure the performance of the algorithm, we performed the following numerical experiment for different values of the parameters: For each pair of parameter values, we chose a number of subspaces of corresponding dimensions at random, ran the algorithm on each subspace, and recorded the worst output (meaning the smallest value, over all tested subspaces, of the minimum absolute value of a d -minor after the algorithm finished improving the subspace). The last column of the table below lists an easy to prove upper bound on the minimum absolute value of a d -minor, for comparison.

	nr. of tests	smallest min-minor after improvement	$\binom{m}{d}^{-1/2}$
$d = 3, m = 5$	200	0.2221	0.3162
$d = 3, m = 6$	272	0.09332	0.2236
$d = 3, m = 7$	211	0.0397	0.1690
$d = 4, m = 7$	261	0.03981	0.1690

Table 4.1: Results of the experiment measuring the performance of the algorithm.

Guessing the asymptotics of the performance of the algorithm is difficult, especially since the exponential running time of the algorithm (with $\binom{m}{d}$ in the exponent) prevents us from solving instances for large values of the parameters. Furthermore, proving that the algorithm performs well when the input subspace is randomly generated does not exclude the possibility that there might be subspaces which can not be improved significantly, but which are highly unlikely to be found through random choice. However, for $d \geq 3$ and $m \geq 7$, the problem of improving a given subspace already becomes highly complex, and it is not at all obvious that good solutions must always exist. Thus the fact that the algorithm consistently performs well in this

situation seems significant.

It is worth noting that in all observed cases, our algorithm outputs a matrix where at least m of the d -minors have the minimum absolute value, and that at least m of them correspond to simple regions (of the corresponding hyperplane arrangement). This fact is noteworthy because the subroutine `IMPROVE-SUBSPACE-ALONG-VECTOR( $v, A$ )` treats all d -minors equally, regardless of whether they correspond to simple regions or not. Therefore, this observed phenomenon confirms, as we have seen in Section 4.4, that the extension hyperplanes corresponding to simple regions play a special role when balancing the subspace.

Two variations of the algorithm.

Instead of choosing a new vector randomly at each iteration, we can also fix an arbitrary basis once at the beginning and then repeatedly cycle through the basis vectors and apply the improvement operation. Since the vector v' which replaces a basis vector v is already orthogonal to the remaining basis vectors, the new set of vectors (with v' instead of v) forms an orthonormal basis of the subspace for the next iteration. Convergence of the smallest absolute value of a d -minor follows as before, and from Lemma 4.4.1 we have that if the basis vectors of the limit subspace are sufficiently far from all of their respective extension hyperplanes, then the same holds for every unit vector.

Another variation is based on the observation that, while our algorithm evens out the d -minors of a matrix with orthonormal columns, we are actually interested in the size of the d -minors after the rows of the matrix are normalised, since this is the matrix to which we would apply Lemma 4.2.7 in order to argue that every region of the corresponding arrangement has a vector which is sufficiently far from every arrangement hyperplane. However, it is entirely possible for a d -minor to be small

only because the norms of the corresponding rows are small, while the hyperplanes whose normals are the rows of the d -minor form large regions together. Thus it would make sense to modify the algorithm in such a way that before every call to `IMPROVE-SUBSPACE-ALONG-VECTOR( $v, A$ )`, the rows of (the current) A are normalised and the Gram-Schmidt orthonormalisation process is applied to the columns. Note that after the row normalisation, the d -minors are the ones we are actually interested in making as large as possible. Also, recall that G-S orthogonalisation means that at step k , the projection of the k -th vector onto the subspace spanned by the previous vectors is subtracted from it, after which it is normalised. Since subtracting columns from one another does not change determinants, and since multiplying a column by a positive constant also means multiplying every d -minor of the matrix by the same value, it follows that after the row normalisation and G-S orthonormalisation, the d -minors of the resulting matrix are, respectively, the d -minors of the matrix whose rows are normalised, up to a factor which is the same for all d -minors.

Note, however, that our improvement operation still evens out the d -minors of a matrix with orthonormal columns, and not those of the row-normalised matrix. While it is entirely possible to define the centre of a region of an extension arrangement in such a way that replacing a vector with the centre of its region in its extension hyperplane equalises a set of d -minors of the *row-normalised matrix*, determining such a centre is not a linear problem, making both the computation and the theoretical treatment much more difficult. In practice, the above modification to the algorithm improves the resulting subspaces significantly, but due to the discrepancy between the goals of the alternating operations, the subspace does not always converge, and can oscillate between different (albeit “good”) subspaces.

Conclusion

The investigation of the relation between the amount of correlation between the marginals of a distribution, and the difficulty of the distribution with respect to the distributional complexity, has yielded a better understanding of the randomised communication complexity of the Disjointness problem. We showed that increasing the allowed correlation by any amount, results in a quantifiable increase in the bounded-information distributional complexity of Disjointness.

We have also shown that Disjointness is not entirely uninteresting to consider in the context of simultaneous message passing, despite the weakness of the setting in general, by giving a non-trivial upper bound for the non-deterministic SMP complexity of this problem. The question of the complexity of Disjointness with respect to such protocols deserves further study, as the known upper and lower bounds for it do not yet match.

For our second, more ambitious topic, we have presented a new approach to a long-standing open problem related to the largest and least well understood probabilistic communication complexity class, UPP^{cc} . We have reduced the open question of whether UPP^{cc} is contained in $PSPACE^{cc}$ to a conjecture in Combinatorial Geometry, and, in order to handle this type of statement, have laid the foundations of a theory of operations on hyperplane arrangements which are invariant with respect to the structure of the arrangements. Finally, we argued for the plausibility of our

conjecture by providing numerical evidence to support it. We are confident that further developing our theory of signature-invariant operations will lead to a proof that UPP^{cc} is indeed a subclass of $PSPACE^{cc}$.

Bibliography

- [1] S. Aaronson, A. Wigderson. Algebrization: a new barrier in complexity theory. *ACM Transactions on Theory of Computing*. 1(1), 2009.
- [2] N. Alon, P. Frankl, and V. Rödl. Geometrical realization of set systems and probabilistic communication complexity. In *Proceedings of the 26th Symposium on Foundations of Computer Science (FOCS)*, pages 277-280, 1985.
- [3] L. Babai, P. Frankl, and J. Simon. Complexity classes in communication complexity theory. In *Proceedings of the 27th IEEE FOCS*, pages 337-347, 1986.
- [4] S. Ben-David, N. Eiron, and H.U. Simon. Limitations of learning via embeddings in Euclidean half-spaces. *Journal of Machine Learning Research*, 3:441-461, 2002.
- [5] A. Blum, A.M. Frieze, R. Kannan, and S. Vempala. A polynomial-time algorithm for learning noisy linear threshold functions. *Algorithmica*, 22(1/2):3552, 1998.
- [6] R.C. Bottesch, D. Gavinsky, and H. Klauck. Correlation in Hard Distributions in Communication Complexity. To appear. (To be presented at *RANDOM 2015*.)
- [7] R.C. Bottesch, D. Gavinsky, and H. Klauck. Equality, Revisited. To appear. (To be presented at *MFCS 2015*.)
- [8] H. Buhrman, N. Vereshchagin, and R. de Wolf. On computation and communication with small bias. In *Proc. IEEE Complexity*, pages 24-32, 2007.

- [9] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, 1991.
- [10] V. Eberhard. Eine Classification der allgemeinen Ebenensysteme. *J. reine angew. Math.*, 106, 1890.
- [11] A. Chattopadhyay, J. Edmonds, F. Ellen, and T. Pitassi. A little advice can be very helpful. In *Proc. 23rd ACM/SIAM Symposium on Discrete Algorithms*, pages 615-625, 2012
- [12] J. Forster. A linear lower bound on the unbounded error probabilistic communication complexity. In *Proceedings of the 16th Annual Conference on Computational Complexity*, pages 100-106, 2001.
- [13] J. Forster, M. Krause, S.V. Lokam, R. Mubarakzjanov, N. Schmitt, and H.U. Simon. Relations between communication complexity, linear arrangements, and computational complexity. In *Proc. of the 21st Conf. on Foundations of Software and Theoretical Computer Science (FST TCS)*, pages 171-182, 2001.
- [14] J. Forster and H.U. Simon. On the smallest possible dimension and the largest possible margin of linear arrangements representing given concept classes. *Theor. Comput. Sci.*, 350(1):40-48, 2006.
- [15] B. Grünbaum. *Convex Polytopes*. Springer, 2003.
- [16] P. Harsha, R. Jain, D.A. McAllester, and J. Radhakrishnan. The communication complexity of correlation. *IEEE Transactions on Information Theory*, 56(1):438-449, 2010.
- [17] J. Håstad and A. Wigderson. The randomized communication complexity of set disjointness. *Theory of Computing*, 3(1):211-219, 2007.

- [18] R. Jain and S. Zhang. New bounds on classical and quantum one-way communication complexity. *Theoretical Computer Science*, 410(26):2463-2477, 2009.
- [19] B. Kalyanasundaram and G. Schnitger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545-557, 1992. Earlier version in Structures '87.
- [20] H. Klauck. Rectangle size bounds and threshold covers in communication complexity. In *Proc. IEEE Conference on Computational Complexity*, pages 118-134, 2003.
- [21] N. Linial, S. Mendelson, G. Schechtman, and A. Shraibman. Complexity measures of sign matrices. *Combinatorica*, 27(4):439-463, 2007.
- [22] N. Linial and A. Shraibman. Learning Complexity vs. Communication Complexity.
- [23] S.V. Lokam. Spectral methods for matrix rigidity with applications to size-depth tradeoffs and communication complexity. In *IEEE Symposium on Foundations of Computer Science*, pages 615, 1995.
- [24] L. Lovasz. Communication complexity: a survey. In *Paths, Flows, and VLSI Layout*, edited by B.H. Korte, Springer Verlag, Berlin New York, 1990.
- [25] L. Lovasz and M. Saks. Lattices, Möbius functions and communication complexity. *Annual Symposium on Foundations of Computer Science*, pages 81-90, 1988.
- [26] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [27] P. Orlik and H. Terao. *Arrangements of Hyperplanes*. Springer-Verlag, 2010.

- [28] R. Paturi and J. Simon. Probabilistic communication complexity. *J. Comput. Syst. Sci.*, 33(1):106-123, 1986.
- [29] A.A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385-390, 1992.
- [30] A.A. Razborov and A.A. Sherstov. The Sign-Rank of AC^0 . In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 57-66, 2008.
- [31] R. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27:701-717.
- [32] A.A. Sherstov. Communication complexity under product and nonproduct distributions. In *Proceedings of the 23rd Annual IEEE Conference on Computational Complexity, CCC*, pages 64-70, 2008.
- [33] A.A. Sherstov. Halfspace matrices. In *Proceedings of 22nd Conference on Computational Complexity*, 2007.
- [34] A.A. Sherstov. Unbounded-error communication complexity of symmetric functions. Technical Report TR-07-53, The Univ. of Texas at Austin, Dept. of Computer Sciences, 2007.
- [35] C. Shonkwiler. Principal angles in terms of inner products. Technical Report, University of Pennsylvania, 2009.
- [36] R.P. Stanley. An introduction to hyperplane arrangements. In *Geometric Combinatorics*, IAS/Park City Mathematics Series, Volume 14, 2004.
- [37] L.G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134-1142, 1984.

- [38] I. Wegener. *The Complexity of Boolean Functions*. Wiley, 1987.
- [39] A. C.-C. Yao. Lower bounds by probabilistic arguments. In *Proc. of the 24th Symposium on Foundations of Computer Science (FOCS)*, pages 420-428, 1983.
- [40] A. C.-C. Yao. Some complexity questions related to distributed computing. In *Proc 11th Symp. on Theory of Computing*, pages 209-123, 1979.
- [41] T. Zaslavsky. Facing up to arrangements: Face-count formulas for partitions of space by hyperplanes. *Memoirs Amer. Math. Soc.* 154, 1975.