How many points can be reconstructed from k projections?

Jiří Matoušek

Department of Applied Mathematics and
Institute of Theoretical Computer Science (ITI)
Charles University, Malostranské nám. 25
118 00 Praha 1, Czech Republic

Aleš Přívětivý

Department of Applied Mathematics
Charles University, Malostranské nám. 25
118 00 Praha 1, Czech Republic

Petr Škovroň

Department of Applied Mathematics
Charles University, Malostranské nám. 25
118 00 Praha 1, Czech Republic

Abstract

Let A be an n-point set in the plane. A *discrete X-ray* of A in direction u gives the number of points of A on each line parallel to u. We define $F(k)$ as the maximum number n such that there exist k directions $u_1, \ldots, u_k$ such that every set of at most n points in the plane can be uniquely reconstructed from its discrete X-rays in these directions. A simple "cube" construction shows $F(k) \leq 2^{k-1}$. We establish the lower bound $F(k) \geq 2^{\Omega(k/\log k)}$ by reducing the problem through linear algebra to a graph-theoretic question, for which we then obtain an almost tight bound. As a part of the proof we establish a result in extremal theory that allows one to conclude that, under certain conditions, a graph has only at most a logarithmic density, which may be of independent interest. We also improve the upper bound to $F(k) \leq O(1.81712^k)$ (or $O(1.79964^k)$ if we allow A to be a multiset).

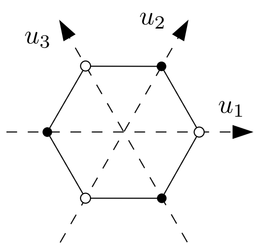

Figure 1: Two 3-point sets with identical X-rays in 3 directions.

1 Introduction

Let A be a finite set of points in the plane. We consider A to be chosen by an adversary and inaccessible by itself, and the information available to us consists only of k discrete X-rays of A. The *discrete X-ray of A in direction u* specifies the number of points of A on every line parallel to u (this convenient terminology is borrowed from the field of geometric tomography). We say that A is *uniquely reconstructible for directions* $u_1, u_2, \dots, u_k$ if there is no $B \neq A$ such that for all $i = 1, 2, \dots, k$, the X-rays of A and B in direction u_i are identical.

It is not hard to see that every set A of $k-1$ or fewer points is uniquely reconstructible for *any* k distinct directions. (To see this, we suppose that some $A \neq B$ have the same X-rays in directions $u_1, \dots, u_k$, we fix a point $a \in A \setminus B$, and we note that each line through a parallel to some u_i has to contain a point of B, forcing $|B| \geq k$.) This has been observed many times; the earliest reference seems to be Rényi [Rén52]. If the directions are chosen by an adversary, then we cannot do any better in general: Fig. 1 shows directions u_1, u_2, u_3 and two distinct point sets A and B (the black points and the white points) that cannot be distinguished. This can be generalized to any number k of equally spaced directions, where A and B are obtained by coloring the vertices of a regular $2k$-gon alternately black and white.

Intuition suggests that the equally spaced directions in this example are "exceptionally bad", and that other sets of directions should allow for unique reconstruction of much larger sets. For given directions $u_1, u_2, \dots, u_k$ let us define

$$f(u_1, \dots, u_k) := \max\{n : \text{every } n\text{-point set is uniq. reconstr. for } u_1, \dots, u_k\}$$

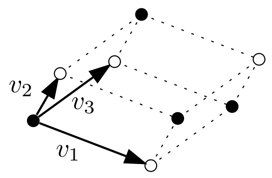

Figure 2: A construction of two sets with identical X-rays in given directions.

and

$$F(k) := \max_{u_1, u_2, \ldots, u_k} f(u_1, \ldots, u_k).$$

We have just noted that $f(u_1, \ldots, u_k) \geq k - 1$ for all k-tuples of distinct $u_1, \ldots, u_k$.

The following simple construction (also rediscovered several times) shows that, perhaps counterintuitively, $F(k)$ is finite for every k; see Fig. 2. Let $u_1, u_2, \ldots, u_k$ be given distinct directions. For suitable nonzero scalars $\alpha_1, \ldots, \alpha_k$ we set $v_i := \alpha_i u_i$ and we define

$$A := \{v(I) : I \subseteq [k], |I| \text{ even}\}, \quad B := \{v(I) : I \subseteq [k], |I| \text{ odd}\}, \tag{1}$$

where we use the notation $[k] := \{1, 2, \ldots, k\}$ and $v(I) := \sum_{i \in I} v_i$. First, it is easy to see that almost all choices of the α_i guarantee that all of the 2^k points $v(I)$, $I \subseteq [k]$, are distinct, and thus $|A| = |B| = 2^{k-1}$. Second, for every $i \in [k]$, each point $v(I) \in A$ can be paired with the point $v(I \triangle \{i\}) = v(I) \pm v_i \in B$ (where $\triangle$ denotes symmetric difference), and this shows that A and B have identical X-rays in direction u_i. Thus $F(k) \leq 2^{k-1} - 1$.

The only published lower bound on $F(k)$ we could find is roughly $k + \Omega(\sqrt{k})$, due to Bianchi and Longinetti [BL90]. After we started working on the problem, we learned from Attila Pór that Tóth [Tót03] announced an $\Omega(k^{3/2})$ lower bound, which has remained unpublished.

We have the following lower bound:

Theorem 1.1 *There are constants $c > 0$ and k_0 such that*

$$F(k) > 2^{ck/\log k}$$

for all $k \geq k_0$. Moreover, for every $k \geq k_0$ there exists a finite set $\mathcal{P}_k$ of nonzero polynomials in $2k$ variables and with integer coefficients such that if

$u_1 = (x_1, y_1)$,..., $u_k = (x_k, y_k)$ are directions with $f(u_1, \ldots, u_k) \le 2^{ck/\log k}$, then $(x_1, x_2, \ldots, x_k, y_1, \ldots, y_k)$ is in the zero set of some polynomial in $\mathcal{P}_k$. Consequently, almost all (in the sense of measure) k-tuples of directions $u_1, \ldots, u_k$ satisfy the stated lower bound.

Thus, the true order of magnitude of $F(k)$ is between $2^{ck/\log k}$ and 2^{k-1}. We don't have any good guess of what it might be, but at least we can improve the upper bound from 2^{k-1} to C^k for a suitable $C < 2$.

To prove an upper bound we need to construct, for any k-tuple $u_1, \ldots, u_k$, two sets A and B with identical X-rays in these directions. It seems technically much easier to construct examples if we allow for *multisets* A and B (formally, an n-point multiset A can be regarded as an arbitrary mapping $[n] \to \mathbb{R}^2$, and in an X-ray of A, the points are counted with the appropriate multiplicities). We thus define $f_{\rm mult}(u_1, \ldots, u_k)$ as the largest n such that every n-point multiset is uniquely reconstructible for $u_1, \ldots, u_k$, and $F_{\rm mult}(k) := \max_{u_1,\ldots,u_k} f_{\rm mult}(u_1, \ldots, u_k)$. Let us remark that the proof of the lower bound in Theorem 1.1 applies equally well to $F_{\rm mult}$.

Theorem 1.2 *We have*

$$F(k) \le O(C^k)$$

for $C = 6^{1/3} \approx 1.81712$ and

$$F_{\rm mult}(k) \le O(C_1^k)$$

for $C_1 = 198^{1/9} \approx 1.79964$.

Higher dimensions. The reconstruction problem can naturally be generalized to point sets in $\mathbb{R}^d$. There are actually several possible generalizations, since we can X-ray the point set with lines as we did in the plane, but also with k-flats for some k between 1 and $d-1$. Here we comment only on the case of line X-rays, for which some simple observations below show that it is not too different from the planar case.

We thus say that $A, B \subset \mathbb{R}^d$ have the same X-rays in direction u if $|A \cap \ell| = |B \cap \ell|$ for every line ℓ parallel to u, and we define the function $F_d(k)$ analogous to $F(k)$.

The simple construction above for the upper bound on $F(k)$ works in any dimension, so $F_d(k) \le 2^{k-1} - 1$ for all $d \ge 2$ and all k. By the method explained in Section 4 below, it can be shown that for every fixed d there exists $\delta_d > 0$ with $F_d(k) = O((2 - \delta_d)^k)$.

On the other hand, we have $F_d(k) \geq F(k)$ for all d, k, $d \geq 2$. Indeed, let $u_1, \ldots, u_k$ be directions in $\mathbb{R}^2$ with $f(u_1, \ldots, u_k) = F(k)$. We identify $\mathbb{R}^2$ with a (2-dimensional) plane ρ in $\mathbb{R}^d$, so that we can consider the u_i as directions in $\mathbb{R}^d$. Now if $A, B \subset \mathbb{R}^d$ are distinct sets with identical X-rays in directions $u_1, \ldots, u_k$, then for every plane σ parallel to ρ the sets $A \cap \sigma$ and $B \cap \sigma$ have identical X-rays in directions $u_1, \ldots, u_k$ as well, and they have to be distinct for at least one σ. For such a σ we get $|A \cap \sigma| = |B \cap \sigma| > F(k)$ by the choice of $u_1, \ldots, u_k$, and this concludes the proof of $F_d(k) \geq F(k)$.

Related work. Problems similar to those investigated in the present paper have been studied in a lively area called *geometric tomography*; see, e.g., the book by Gardner [Gar06]. The classical tomography problem deals with reconstructing a set, or more generally a density function, from X-rays in all directions. *Discrete tomography* investigates the reconstruction problem for a finite (discrete) set of X-ray directions. Since reconstructing an arbitrary set is generally impossible, most of the work deals with special sets, say convex ones.

For reconstructing finite sets A, most of the results in the literature concern the case where A is a lattice set, $A \subseteq \mathbb{Z}^2$, and the directions of the X-rays are integer vectors. A seminal paper by Gardner and Gritzmann [GG97] thoroughly examines the case where A is guaranteed to be a convex lattice set (that is, the intersection of $\mathbb{Z}^2$ with a convex set). In this case, they show that any 7 lattice directions suffice for unique reconstruction of every convex lattice set, while 6 directions need not suffice. Dulio, Gardner, and Peri [DGP06] studied a variant of the problem ("point X-rays"): instead of k directions, we have k points $p_1, \ldots, p_k$, and we obtain the number of points of A on every line passing through one of the p_i. They show, for example, that the analogue of our function $F(k)$ is unbounded in their setting. Few other papers with somewhat related results are [Hep56], [Gar92], [BDLNP01].

Acknowledgements. We thank Štěpán Holub for raising a problem that brought us to the research reported here, Attila Pór for fruitful discussions concerning the upper bounds, Géza Tóth for the idea of considering the interchange graph, and Noga Alon for pointing out the reference [Alo95].

2 A reduction of the lower bound to a combinatorial problem

This section contains the first part of the proof of the lower bound for $F(k)$. Given two sets A and B with identical X-rays in given k directions, we construct a graph with edges colored by k colors, and we show that for a generic choice of the directions this colored graph satisfies a combinatorial condition; namely, it doesn't contain two color-disjoint spanning trees. The subsequent section then deals with the resulting graph-theoretic problem.

The interchange graph. Let $u_1, \ldots, u_k$ be fixed direction vectors, and let $A \subset \mathbb{R}^2$ be a set with the minimum number of points n that is not uniquely reconstructible for $u_1, \ldots, u_k$. Thus there exists another set $B \subset \mathbb{R}^2$ with the same X-rays in these directions. We note that $A \cap B = \emptyset$ (if not, $A \setminus B$ and $B \setminus A$ would be smaller indistinguishable sets).

For each $i = 1, 2, \ldots, k$, we construct a perfect matching E_i between A and B in such a way that for each edge $\{a, b\} \in E_i$, the points $a \in A$ and $b \in B$ lie on the same line parallel to u_i. Thus, if a line ℓ in direction u_i contains a single point of A, and hence also a single point of B, these two points necessarily form an edge of E_i. If ℓ contains several points of A, these are matched arbitrarily, in a one-to-one fashion, to the points of $\ell \cap B$ (hence E_i need not be determined uniquely, but if there are several choices, we fix one once and for all). We note that $E_i \cap E_j = \emptyset$ for $i \neq j$.

We call the bipartite graph H with vertex set $A \cup B$ and with edge set $E(H) = E_1 \cup \cdots \cup E_k$ an *interchange graph for directions* $u_1, \ldots, u_k$. It has $2n$ vertices and kn edges.

For our exposition it will be convenient to consider the edges of an interchange graph to have colors: We say that an edge $e \in E_i$ has *color* i and we write $c(e) = i$.

2.1 An algebraic necessary condition for interchange graphs

We begin by recalling several notions from algebraic graph theory.

Let $\vec{G} = (V, \vec{E})$ be a directed graph. For an edge $e = (u, v) \in \vec{E}$, we write $u = \mathrm{tail}(e)$ and $v = \mathrm{head}(e)$. The *incidence matrix* $D = D_{\vec{G}}$ is an $V \times \vec{E}$ matrix, with rows indexed by vertices of $\vec{G}$ and columns indexed by edges, given by

$$d_{v,e} := \begin{cases} +1 & \text{if } v = \mathrm{head}(e), \\ -1 & \text{if } v = \mathrm{tail}(e), \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

A *circulation* in $\vec{G}$ is any vector in $\ker D$, that is, a vector $\xi \in \mathbb{R}^{\vec{E}}$ with $D\xi = 0$.

Further we recall that an *orientation* of an undirected graph $G = (V, E)$ is a directed graph $(V, \vec{E})$ obtained by choosing, for every undirected edge $\{u, v\} \in E$, one of its vertices as head and one as tail (more formally, for every $\{u, v\} \in E$ we have exactly one of the directed edges (u, v) and (v, u) in $\vec{E}$). For notational convenience we will use the same letter for an undirected edge $e = \{u, v\}$ and for a directed version (u, v) of it.

We begin with the following necessary condition for interchange graphs:

Lemma 2.1 *Let $H = (V, E)$ be a subgraph of an interchange graph for directions $u_1, u_2, \ldots, u_k$, and let $\vec{H}$ be an arbitrary orientation of H. Then there exist reals α_e, $e \in E$, all of them nonzero, such that for every circulation ξ in $\vec{H}$ we have*

$$\sum_{e \in E} \alpha_e \xi_e x(u_{c(e)}) = 0,$$

and for every circulation ψ in $\vec{H}$ we have

$$\sum_{e \in E} \alpha_e \psi_e y(u_{c(e)}) = 0.$$

Here $c(e)$ denotes the color of the edge e (inherited from the interchange graph), and $x(u)$ and $y(u)$ denote the x-coordinate and y-coordinate, respectively, of a vector $u \in \mathbb{R}^2$.

Proof. First we note that if $G = (V, \vec{E})$ is a directed graph, ξ is a circulation in it, and $z \in \mathbb{R}^V$ is a vector indexed by the vertices of G, we have

$$\sum_{e \in \vec{E}} (z_{\mathrm{head}(e)} - z_{\mathrm{tail}(e)}) \xi_e = 0. \qquad (2)$$

This is immediate using $D\xi = 0$, because the sum equals $z^T D\xi$. (Or, perhaps more intuitively, the equality is obvious if ξ is an elementary circulation along a cycle, and the general case follows since every circulation is a linear combination of elementary circulations.)

Next, let us consider $\vec{H}$ as in the lemma, and let ξ be a circulation in it. The vertices of H are points in the plane, and for every edge $e \in E$ we have $(\mathrm{head}(e) - \mathrm{tail}(e)) = \alpha_e u_{c(e)}$ for some $\alpha_e \neq 0$, since the segment with endpoints head(e) and tail(e) is parallel to $u_{c(e)}$. Then the first equality in the claim of the lemma follows by applying (2) with $z_v := x(v)$ for all vertices v, while the second equality is obtained similarly using the y-coordinates.

□

Let D be the incidence matrix of a directed graph. By D^- we denote D with the last row deleted. We note that since every column of D has one $+1$ and one -1, the sum of all rows of D is the zero vector, and thus each row is a linear combination of the others. Consequently, we have $\ker D^- = \ker D$, and thus it suffices to "test" circulations using D^-.

Let $H = (V, E)$ be a graph with n vertices and m edges and with edges colored by k colors. We define a polynomial matrix $P_H = P_H(x_1, \ldots, x_k, y_1, \ldots, y_k)$ with $m + 2n - 2$ rows and $2m$ columns, where the x_i and the y_i are variables:

$$P_H := \begin{pmatrix} X & Y \\ D^- & 0 \\ 0 & D^- \end{pmatrix},$$

where X is an $E \times E$ diagonal matrix with the entry (e, e) equal to $x_{c(e)}$, Y is defined analogously with $y_{c(e)}$, and D is the incidence matrix of some orientation of H (thus,P_H is not defined uniquely, since it depends on the orientation).

Here is an algebraic condition for certain subgraphs of interchange graphs:

Lemma 2.2 *Let $H = (V, E)$ be a subgraph of an interchange graph for directions $u_1, \ldots, u_k$. Let us assume that $|E| = 2n - 2$, where $n = |V|$. Then P_H is a square matrix and*

$$\det\Big(P_H(x(u_1), \ldots, x(u_k), y(u_1), \ldots, y(u_k))\Big) = 0$$

(we substitute the coordinates of the directions $u_1, \ldots, u_k$ for the variables in the matrix P_H).

Proof. Let us write $\tilde{P}_H = P_H(x(u_1), \ldots, x(u_k), y(u_1), \ldots, y(u_k))$. Supposing for contradiction that $\det(\tilde{P}_H) \neq 0$, we get that the linear system

$\tilde{P}_H q = b$ has a (unique) solution q for every right-hand side $b \in \mathbb{R}^{2m}$. Let us choose, in particular, $b = \binom{w}{0}$, where $w \in \mathbb{R}^E$ is arbitrary (and 0 stands for an m-component zero vector), and let us also write $q = \binom{\xi}{\psi}$, with $\xi, \psi \in \mathbb{R}^E$. Using the definition of P_H, we get that the system $\tilde{P}_H q = b$ is then equivalent to

$$D^- \xi = 0, \quad D^- \psi = 0, \quad x(u_{c(e)})\xi_e + y(u_{c(e)})\psi_e = w_e \text{ for all } e \in E.$$

The first two equations say that ξ and ψ are circulations.

Lemma 2.1 tells us that there exist nonzero α_e, $e \in E$, such that for every two circulations ξ, ψ we have

$$\sum_{e \in E} \alpha_e \Big(x(u_{c(e)})\xi_e + y(u_{c(e)})\psi_e \Big) = 0.$$

However, in view of the previous paragraph, for every $w \in \mathbb{R}^E$ we can choose circulations ξ, ψ with $x(u_{c(e)})\xi_e + y(u_{c(e)})\psi_e = w_e$ for all $e \in E$, and thus we have $\sum_{e \in E} \alpha_e w_e = 0$ for all w. This forces $\alpha_e = 0$ for all e, though—a contradiction proving the lemma. □

In order that the last lemma provide a nontrivial condition for the directions u_i, we need that the determinant of the matrix $P_H(x_1, \ldots, x_k, y_1, \ldots, y_k)$ is not the zero polynomial. The following section provides a sufficient combinatorial condition for this.

2.2 A combinatorial necessary condition for interchange graphs

Lemma 2.3 *Let $H = (V, E)$ be a subgraph of an interchange graph for directions $u_1, \ldots, u_k$, and suppose that H has two color-disjoint spanning trees T_1 and T_2 (that is, $c(e_1) \neq c(e_2)$ for all $e_1 \in E(T_1)$, $e_2 \in E(T_2)$). Then $(x(u_1), \ldots, x(u_k), y(u_1), \ldots, y(u_k))$ is in the zero set of a nonzero polynomial with integer coefficients (depending on H).*

Proof. Without loss of generality we may assume that $E = E(T_1) \cup E(T_2)$. Writing $n = |V|$, we have $|E| = 2n - 2$, and so Lemma 2.2 tells us that that the vector of the coordinates of the u_i is in the zero set of $\det(P_H)$. It remains to check that $\det(P_H)$ is not identically 0. To this end, we verify

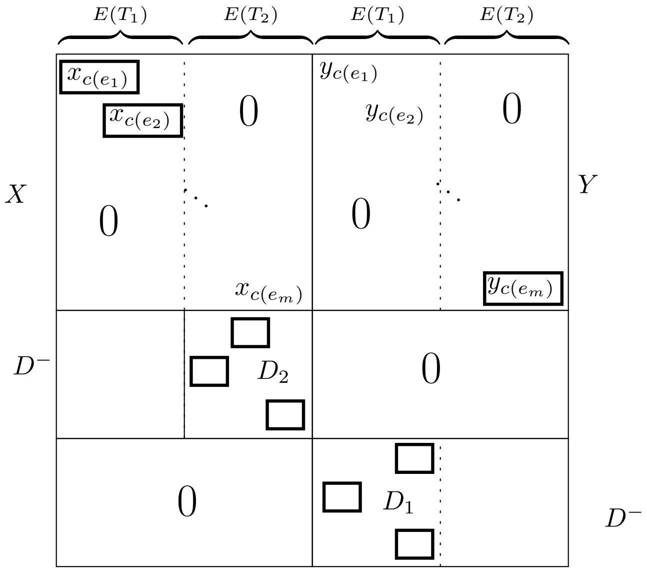

Figure 3: Terms in $\det(\tilde{P}_H)$ contributing to the monomial μ.

that the monomial

$$\mu := \Big(\prod_{e\in E(T_1)} x_{c(e)}\Big)\Big(\prod_{e\in E(T_2)} y_{c(e)}\Big)$$

appears in $\det(P_H)$ with coefficient ± 1. Indeed, let us consider the expansion of $\det(P_H)$ according to the definition of the determinant as a sum over all permutations of products. Since T_1 and T_2 are color-disjoint, if such a product should be a nonzero multiple of the monomial μ, it must use all the entries $x_{c(e)}$ of the diagonal matrix X for $e \in E(T_1)$, and also all the entries $y_{c(e)}$ of the diagonal matrix Y for $e \in E(T_2)$; see Fig. 3. This uniquely determines which entries are used from the top m rows of P_H. In the next $n-1$ rows, the entries must be taken from the submatrix D_2^- of D^- corresponding to $E(T_2)$, and in the last $n-1$ rows the entries must be taken from the submatrix D_1^- of D^- corresponding to $E(T_1)$. Hence the coefficient of μ in $\det(P_H)$ equals $\det(D_1^-)\det(D_2^-)$. By a standard fact (a special case of the matrix-tree theorem; see, e.g., Biggs [Big93]), the incidence matrix of any directed tree has determinant ± 1, and the lemma follows. $\square$

3 Existence of two color-disjoint spanning trees

The last building block in the proof of our lower bound for $F(k)$ is a result stating that every sufficiently dense graph has a subgraph with two color-disjoint spanning trees.

The graphs we will deal with throughout the proof will be disjoint unions of perfect matchings, so it is convenient to introduce the following terminology. A *DUPM-graph* is a graph $G = (V, E)$ together with a fixed partition $E = E_1 \cup \cdots \cup E_k$ of the edge set into perfect matchings on V (we also keep referring to *colors* of edges as in the previous sections). In particular, we note that each vertex is incident to edges of all colors in a DUPM-graph. For an index set $I \subseteq [k]$ and a subset $W \subseteq V$, we write $G[I, W]$ for the subgraph induced by W on edges with colors in I, that is, $G[I, W] = \left(W, \binom{W}{2} \cap \bigcup_{i \in I} E_i\right)$ (instead of $G[I, V]$ we write just $G[I]$). This $G[I, W]$ need not be a DUPM-graph in general, but in our considerations we will choose I and W so that we do get DUPM-graphs.

Before proving the result needed for the proof of Theorem 1.1, we first establish a quantitatively weaker result with a simpler proof (Section 3.1), which can serve as a gradual introduction to the subsequent more involved proof of the stronger result.

3.1 A weaker bound

The weaker result with the simpler proof is the following:

Proposition 3.1 *Let $G = (V, E_1 \cup \cdots \cup E_k)$ be a DUPM-graph on $n \geq 2$ vertices. If $k \geq 5 \log_2^3 n$, then there exist disjoint index sets $I_1, I_2 \subset [k]$ and a subset $W \subseteq V$, $|W| \geq 2$, such that the graphs $G[I_1, W]$ and $G[I_2, W]$ are both connected.*

The construction of the desired two subgraphs proceeds in two stages. In the first stage, captured by the next lemma, we find a DUPM-subgraph of G in which every edge cut has sufficiently many colors.

Lemma 3.2 *Let G and k be as in Proposition 3.1. Then there exist nonempty $J \subseteq [k]$ and $W \subseteq V$, $|W| \geq 2$, such that $H := G[J, W]$ is a DUPM-graph* *in which every edge cut has edges of at least $r := 5\log_2^2 n$ colors. That is, for every partition of W into nonempty sets A and B, the edges in $E(H;A,B) := \{e \in E(H) : e \cap A \neq \emptyset \neq e \cap B\}$ have at least r distinct colors.*

The proof proceeds by a recursive partitioning: As long as we can find an edge cut F with fewer than r colors in the current graph, we discard the *larger* of the two pieces defined by F, together with all edges that do not have both endpoints in the new vertex set or whose color appears in F. We repeat this step as long as we can. Since the size of the vertex set is halved (at least) in each iteration, the number of iterations is at most $\log_2 n$. The key observation is that since we started with $5\log_2^3 n$ colors and we lose fewer than r colors in every iteration, we must finish with a nontrivial graph satisfying the required condition. A more formal proof follows.

Proof. We start with $J_1 = [k]$, $V_1 = V$, and $G_1 = G$, and for $j = 1, 2, \ldots$ we repeat the following step:

> **(Step j)** We have a current DUPM-graph G_j on vertex set V_j with edge colors in J_j. If there is a partition $V_j = A \cup B$ such that the edges in $E(G_j; A, B)$ have fewer than r colors, we construct J_{j+1}, V_{j+1}, G_{j+1} as follows:
>
> - V_{j+1} is the smaller of the sets A and B (a possible tie broken arbitrarily);
> - J_{j+1} is J_j minus all colors appearing on edges in $E(G_j; A, B)$;
> - $G_{j+1} = G_j[V_{j+1}, J_{j+1}]$.
>
> Otherwise, if $E(G_j; A, B)$ has at least r colors for every partition $V_j = A \cup B$, then we set $W = V_j$, $J = J_j$ and finish.

It is clear that the resulting $H = G[J, W]$ is a DUPM-graph, and actually that every vertex of H is incident to $|J|$ edges, and also that every edge cut has at least r colors. It remains to check that $J \neq \emptyset$. As we have observed, the algorithm must finish in at most $\log_2 n$ steps since $|V_{j+1}| \le \lfloor \frac{1}{2}|V_j| \rfloor$ for all j. In each step the number of colors in J_j decreases by less than r, and hence $|J| > k - r\log_2 n \ge 0$. The lemma is proved. $\square$

In the second stage of the proof of Proposition 3.1, having a subgraph where every edge cut has at least r colors, we exhibit two color-disjoint connected subgraphs. To build the first connected subgraph, we start with an empty edge set and add the edge sets E_i one by one, always taking one that decreases the current number of components the most. Since every edge cut contains at least about $1/\log n$ fraction of all colors, it can be shown that adding a random color decreases the current number of component by factor at least roughly $(1 - 1/\log n)$. Consequently, after adding $s = O(\log^2 n)$ colors we obtain a connected graph. Setting the constants so that $s < r$, the remaining colors form a graph where every edge cut still has at least one color, and hence the graph is connected. As was pointed to us by Noga Alon, this stage of our proof resembles his proof in [Alo95].

The next lemma formalizes this second stage.

Lemma 3.3 *Let H be a DUPM-graph on $m \leq n$ vertices and with $k \leq 6\log_2^3 n$ colors in which every edge cut has edges of at least $r := 5\log_2^2 n$ colors. Then there exists a set I of at most $s := 4\log_2^2 n$ colors such that $H[I]$ is connected.*

Proof. We set $I_0 = \emptyset$, and for $j = 0, 1, 2, \ldots$, we do the following: If $H[I_j]$ is connected, we set $I := I_j$ and finish. Otherwise, we let i_j be a color i maximizing $\mathrm{cc}(H[I_j]) - \mathrm{cc}(H[I_j \cup \{i\}])$, where $\mathrm{cc}(G)$ denotes the number of connected components of a graph G. Then we set $I_{j+1} := I_j \cup \{i_j\}$, and we continue with the next step.

We need to show that we obtain a connected graph before exhausting more than the claimed number of colors.

Let $m_j = \mathrm{cc}(H[I_j])$. We have $m_1 = m \leq n$, and we claim that

$$m_{j+1} \leq m_j \left(1 - \frac{r}{2k}\right)$$

whenever $m_j \geq 2$. To see this, let $K_1, K_2, \ldots, K_{m_j}$ be the connected components of $H[I_j]$. Let us say that K_ℓ *gets connected by color i* if there is an edge of color i connecting K_ℓ to the rest of H. Since the edge cut $E(H; V(K_\ell), W \setminus V(K_\ell))$ contains at least r colors (none of them in I_j, since K_ℓ is a connected component of $H[I_j]$), the probability that any particular K_ℓ gets connected by a randomly chosen color i is at least r/k. Hence in expectation at least $q := m_j r/k$ components get connected, and thus there is a particular color i by which at least q components get connected. It is not difficult to check that by adding such a color i to I_j the number of components decreases by at least $q/2$, and thus $m_{j+1} \leq m_j - q/2 = m_j(1 - r/2k)$ as claimed.

Applying this inequality inductively gives $m_j \leq n(1-r/2k)^j \leq ne^{-jr/2k}$. This is at most 1 for $j = 2k(\ln n)/r \leq 2 \cdot 6\log_2^3 n \cdot \log_2 e \cdot (\log_2 n)/(5\log_2^2 n) \leq s = 4\log_2^2 n$, and the lemma is proved. □

Proof of Proposition 3.1. Given a DUPM-graph $G = (V, E_1 \cup \cdots \cup E_k)$ as in the proposition, we first discard colors if needed so that the remaining number of colors is between $5 \log_2^3 n$ and $6 \log_2^3 n$ (we'd actually like to say that we may assume that there are exactly $5 \log_2^3 n$ colors, but that need not be an integer). Then we use Lemma 3.2 and obtain a DUPM-graph $H = G[J, W]$ where every edge cut contains at least r colors (and of course, H has at most n vertices and at most $6 \log_2^3 n$ colors). We apply Lemma 3.3 and we obtain a set I_1 of at most $s = 4 \log_2^2 n < r$ colors with $H[I_1]$ connected. In the DUPM-graph induced in H by colors not belonging to I_1, every edge cut still has at least one color (since we have removed fewer than r colors), and hence this DUPM-graph is connected as well. We can thus set $I_2 := J \setminus I_1$. This proves the proposition. $\Box$

3.2 An improved bound

Here we improve the bound of $\log^3 n$ from Proposition 3.1 to $\log n \log \log n$. The improvement is based on two observations, which are more or less independent (applying each of them alone would save a factor of roughly $\log n$).

The first observation is that in the first stage, i.e., in the proof of Lemma 3.2, we can force unbalanced cuts (with A much smaller than B) to have considerably more colors than balanced ones (with A and B of roughly equal size). Indeed, if we encounter highly unbalanced cuts in most steps of the algorithm, then the algorithm has to finish considerably sooner than after $\log_2 n$ steps, and as a result, fewer colors get discarded. This is expressed quantitatively in Lemma 3.6 below.

In the second stage we build a connected graph by iteratively merging components. Here we note that at the beginning, when the components are small, the property "cuts having sufficiently many colors" is applied with very much unbalanced cuts. Only as the typical size of the current components grows, we have to deal with more balanced cuts. As a result, if we know that unbalanced cuts are richer in colors than balanced ones, the merging procedure progresses faster at the beginning than towards the end. Technically, we are going to divide the merging procedure into roughly $\log \log n$ phases. In each phase, the number of components is typically reduced to approximately its square root (except for the last phase, where the number of components is reduced to 1).

This first improvement alone would bring the original $\log^3 n$ factor down to $\log^2 n \log \log n$, where the $\log \log n$ factors arises because the contribution of each of the $\log \log n$ phases turns out to be roughly the same.

We now turn to the second improvement. The procedure for selecting colors and merging components in the proof of Lemma 3.3 can conceptually be regarded as a randomized algorithm that selects colors one by one at random. (We phrased it deterministically, saying that we always pick a color that reduces the number of components as much as possible, but the reason why this worked was that the expected number of components merged by adding a *random* color is sufficiently large.) Instead of picking the colors one by one, we can think of this as selecting a random subset of s colors, with a suitable s. If we can show that such a random set of colors induces a connected subgraph with probability strictly greater than $\frac{1}{2}$, we can use the following simple observation to assert the existence of *two disjoint* sets of s colors, each inducing a connected subgraph.

Observation 3.4 *Let s, n be natural numbers, $1 \leq s \leq n/2$, and let $\mathcal{P}$ be a property of s-element subsets of $[n]$. If a random s-element subset of $[n]$ (where all s-element subsets are chosen with equal probability) has property $\mathcal{P}$ with probability $p > \frac{1}{2}$, then there exist two disjoint s-element subsets of $[n]$ with property $\mathcal{P}$.*

This observation, as well as the following neat proof, may very well be folklore.

Proof. Let us choose a random permutation of $[n]$. Let S_1 consist of the first s elements in this permutation, and let S_2 consist of the next s elements. We have $S_1 \cap S_2 = \emptyset$ and both S_1 and S_2 can be regarded as random s-element subsets of $[n]$ (not independent, of course). The probability that either of S_1, S_2 fails to have $\mathcal{P}$ is, by the union bound, at most $2(1-p) < 1$. $\square$

Using this observation appropriately will allow us to save another $\log n$ factor in the number of colors.

Here is a formal statement of the improved version of Proposition 3.1.

Proposition 3.5 *There exists a constant C such that the following holds. Let $G = (V, E_1 \cup \cdots \cup E_k)$ be a DUPM-graph on $n \geq 4$ vertices.*[1] *If $k \geq \lfloor C \log_2 n \log_2 \log_2 n \rfloor$, then there exist disjoint index sets $I_1, I_2 \subset [k]$ and a subset $W \subseteq V$, $|W| \geq 2$, such that the graphs $G[I_1, W]$ and $G[I_2, W]$ are both connected.*

[1] We assume $n \geq 4$ so that $\log_2 \log_2 n \geq 1$ and the $C \log_2 n \log_2 \log_2 n$ term can absorb additive constants, for example, if C is chosen sufficiently large.

We begin with the first stage of the proof—a refined version of Lemma 3.2. For a partition of a vertex set into two nonempty subsets A and B, we define the *imbalance* of (A, B) as

$$\mathrm{imb}(A, B) := \frac{|A| + |B|}{\min(|A|, |B|)},$$

and we set

$$\gamma(n, k, A, B) := \frac{k}{2\log_2 n} \cdot \log_2 \mathrm{imb}(A, B);$$

this is going to be the lower bound on the number of colors in every edge cut $E(H; A, B)$.

Lemma 3.6 *Let $G = (V, E_1 \cup \cdots \cup E_k)$ be a DUPM-graph on $n \geq 2$ vertices. Then there exist $J \subseteq [k]$, $|J| \geq k/2$, and $W \subseteq V$, $|W| \geq 2$, such that $H := G[J, W]$ is a DUPM-graph in which every edge cut $E(H; A, B)$ has edges of at least $\gamma(n, k, A, B)$ colors (note that n is the number of vertices of G, not of H).*

Proof. The proof is a simple modification of that for Lemma 3.2 so we proceed quickly. Assuming $k > 0$, we start with $J_1 = [k]$, $V_1 = V$, and $G_1 = G$. If a current DUPM-graph G_j on vertex set V_j with edge colors in J_j has already been constructed and if there is a partition $V_j = A \cup B$ where $E(G_j; A, B)$ has fewer than $\gamma(n, k, A, B)$ colors, we let V_{j+1} be the smaller of the sets A and B, we let J_{j+1} be J_j minus all colors appearing in $E(G_j; A, B)$, and we set $G_{j+1} = G_j[V_{j+1}, J_{j+1}]$. If there is no such partition, we set $W = V_j$, $J = J_j$, $t = j$ and finish.

The only condition whose checking needs some work is $|J| \geq k/2$. In the jth step the number of colors in J_j decreases by less than

$$\gamma(n, k, V_{j+1}, V_j \setminus V_{j+1}) = \frac{k}{2\log_2 n} \log_2 \frac{|V_j|}{|V_{j+1}|},$$

and hence

$$\begin{aligned} |J| &\geq k - \sum_{j=1}^{t-1} \frac{k}{2\log_2 n} (\log_2 |V_j| - \log_2 |V_{j+1}|) \\ &= k - \frac{k}{2\log_2 n} (\log_2 |V_1| - \log_2 |V_t|) \\ &\geq k - \frac{k}{2\log_2 n} \log_2 n = \frac{k}{2}. \end{aligned}$$

The lemma is proved. $\square$

Next, we state a counterpart of Lemma 3.3.

Lemma 3.7 *Let H be a DUPM-graph on $m \le n$ vertices and with k colors, where $k_0/2 \le k \le k_0$ for $k_0 := \lfloor C \log_2 n \log_2 \log_2 n \rfloor$. Let us suppose that every edge cut $E(H; A, B)$ has at least $\gamma(n, k_0, A, B)$ colors. Then a set I of $s := \lfloor \frac{1}{4} k_0 \rfloor$ colors chosen uniformly at random from $[k]$ guarantees $H[I]$ to be connected with probability at least $\frac{3}{4}$.*

Assuming this lemma, the proof of Proposition 3.5 is routine:

Proof of Proposition 3.5. Given a DUPM-graph as in the proposition, we may assume (after deleting colors) that the number of colors is exactly k_0. Lemma 3.6 yields a DUPM-subgraph H with at most n vertices, number of colors between $k_0/2$ and k_0, and every edge cut $E(H; A, B)$ having at least $\gamma(n, k_0, A, B)$ colors. Lemma 3.7 shows that a random subset of s colors induces a connected subgraph in H with probability at least $\frac{3}{4}$, and Observation 3.4 thus yields the existence of two disjoint sets of colors inducing connected subgraphs. $\square$

Now we start working towards the proof of Lemma 3.7. Instead of choosing the random set I of colors all at once, we think of picking the colors one by one. So we set $I_0 = \emptyset$, and for $j = 1, 2, 3, \ldots$, we choose a color $i_j \in [k] \setminus I_{j-1}$ uniformly at random and we set $I_j := I_{j-1} \cup \{i_j\}$.

Moreover, for the purposes of the analysis, we conceptually divide the choice of I into *phases*. We set $b := \lceil 64 \log_2 n \rceil$, and for $i = 0, 1, 2, \ldots$, we let the sets $I_{ib}, I_{ib+1}, \ldots, I_{(i+1)b-1}$ belong to the ith phase.

Let $X_j := \mathrm{cc}(H[I_j])$ be the number of connected components of the graph $H[I_j]$; this is a random variable. The next lemma wants to claim that each phase is reasonably likely to reduce the current number of connected components to no more than its square root. However, if the number of components is already rather small, concretely, below 2^4, we require the phase to finish with a connected graph. So let us call the ith phase *successful* if either $X_{ib} \ge 2^4$ and $X_{(i+1)b} \le \sqrt{X_{ib}}$, or $2 \le X_{ib} < 2^4$ and $X_{(i+1)b} = 1$.

Lemma 3.8 *For every i we can estimate the conditional probability of the ith phase being successful, conditioned on $X_{ib} \ge 2$, as follows:*

$$\mathrm{Prob}\big[i\text{th phase is successful} \,\big|\, X_{ib} \ge 2\big] \ge \frac{1}{2}.$$

Proof. Let us assume that the set I_{ib} is fixed in such a way that $x := X_{ib} \geq 2$. Let $p^* = p^*(I_{ib})$ be the probability that the ith phase is *unsuccessful*, conditioned on I_{ib} fixed as above.[2] It suffices to prove that $p^* \leq \frac{1}{2}$ for every I_{ib} with $X_{ib} \geq 2$.

For contradiction, let us assume $p^* > \frac{1}{2}$. We will show that this implies $\mathbb{E}[\log_2 X_{(i+1)b} \,|\, I_{ib}] < 0$, which is a contradiction, since obviously $X_{(i+1)b} \geq 1$ always. To this end, we define $D_j := \log_2(X_j/X_{j+1})$, and we establish the following:

Claim. *For all $j = ib, ib+1, \ldots, (i+1)b-1$ we have*

$$\mathbb{E}[D_j \,|\, I_{ib}] > \frac{1}{64}\frac{\log_2 x}{\log_2 n}.$$

First, assuming this claim, we can finish the proof of the lemma quickly. We calculate

$$\begin{aligned}
\mathbb{E}[\log_2 X_{(i+1)b} \,|\, I_{ib}] &= \log_2 x - \mathbb{E}[\log_2(X_{ib}/X_{(i+1)b}) \,|\, I_{ib}] \\
&= \log_2 x - \sum_{j=ib}^{(i+1)b-1} \mathbb{E}[D_j \,|\, I_{ib}] \\
&< \log_2 x - \sum_{j=ib}^{(i+1)b-1} \frac{1}{64}\frac{\log_2 x}{\log_2 n} \\
&\leq (\log_2 x)\left(1 - \frac{b}{64\log_2 n}\right) \leq 0
\end{aligned}$$

since $b = \lceil 64 \log_2 n \rceil$. We have shown $\mathbb{E}[\log_2 X_{(i+1)b} \,|\, I_{ib}] < 0$, which is the contradiction announced above.

Proof of the claim. Let us define y as the desired target value of $X_{(i+1)b}$ making the phase successful; that is, $y := \sqrt{x}$ for $x \geq 2^4$ and $y := 1$ for $2 \leq x < 2^4$. We have

$$\begin{aligned}
\mathbb{E}[D_j \,|\, I_{ib}] &= \mathbb{E}[D_j \,|\, I_{ib} \text{ and } X_j > y] \cdot \mathrm{Prob}[X_j > y \,|\, I_{ib}] \\
&\quad + \mathbb{E}[D_j \,|\, I_{ib} \text{ and } X_j \leq y] \cdot (1 - \mathrm{Prob}[X_j > y \,|\, I_{ib}]) \\
&\geq \mathbb{E}[D_j \,|\, I_{ib} \text{ and } X_j > y\,] \cdot p^*, \qquad (3)
\end{aligned}$$

[2] More precisely, we should say that we first fix an ib-element set $J \subset [k]$ with $\mathrm{cc}(H[J]) \geq 2$, we set $x := \mathrm{cc}(H[J])$, and then condition on $I_{ib} = J$. However, we allow ourselves the luxury of the more concise notation.

since $\mathbb{E}[D_j \mid I_{ib} \text{ and } X_j \le y] \cdot (1 - \text{Prob}[X_j > y \mid I_{ib}]) \ge 0$ and $\text{Prob}[X_j > y \mid I_{ib}] \ge \text{Prob}[X_{(i+1)b} > y \mid I_{ib}] = p^*$.

We now want to bound below $\mathbb{E}[D_j \mid I_{ib} \text{ and } X_j > y]$. Let W denote the vertex set of H and let us recall that $m = |W|$. Let us fix I_j such that $X_j > y$, and for every color $\ell \in [k] \setminus I_j$, let $a_\ell := X_j - \text{cc}(H[I_j \cup \{\ell\}])$ be the decrease in the number of connected components caused by adding color ℓ.

For each component A of $H[I_j]$, the number of colors in the complement of I_j that connect A to the rest of the graph is at least

$$\gamma(n, k_0, A, W \setminus A) = \frac{k_0}{2 \log_2 n} \log_2 \text{imb}(A, W \setminus A),$$

and hence

$$\sum_{\ell \in [k] \setminus I_j} a_\ell \ge \frac{1}{2} \cdot \frac{k_0}{2 \log_2 n} \sum_A \log_2 \text{imb}(A, W \setminus A), \tag{4}$$

where the sum is over all components A of $H[I_j]$.

Now for the case $2 \le x < 2^4$, we simply use $\text{imb}(A, W \setminus A) \ge 2$ for every A, and (4) yields

$$\sum_{\ell \in [k] \setminus I_j} a_\ell \ge \frac{1}{4} X_j \frac{k_0}{\log_2 n}. \tag{5}$$

To derive a similar inequality for $x \ge 2^4$ as well, we note that the number of components of $H[I_j]$ of size exceeding some threshold t is at most m/t, and setting $t = 2m/\sqrt{x}$, we find that if $X_j > \sqrt{x}$, then at least $\frac{1}{2} X_j$ of the components have size at most t. For each such component A, we have $\text{imb}(A, W \setminus A) \ge m/t = \frac{1}{2}\sqrt{x}$, and so

$$\sum_{\ell \in [k] \setminus I_j} a_\ell \ge \frac{k_0}{4 \log_2 n} \cdot \frac{1}{2} X_j \cdot \log_2(\sqrt{x}/2) \ge \frac{1}{32} X_j \frac{k_0 \log_2 x}{\log_2 n}, \tag{6}$$

where we have used $\log_2(\sqrt{x}/2) \ge \frac{1}{4} \log_2 x$, which is valid for all $x \ge 2^4$.

In the case $2 \le x < 2^4$, we have $\log_2 x \le 4$, and so (5) shows that (6) holds in this case too.

Now we can calculate

$$
\begin{aligned}
\mathbb{E}[D_j \mid I_j] &= \mathbb{E}[\log_2(X_j/X_{j+1}) \mid I_j] \\
&= \sum_{\ell \in [k] \setminus I_j} \frac{1}{k-j} \log_2 \frac{X_j}{X_j - a_\ell} \\
&\geq \frac{1}{k_0} \sum_{\ell \in [k] \setminus I_j} \log_2 \frac{1}{1 - a_\ell / X_j} \\
&\geq \frac{1}{k_0} \sum_{\ell \in [k] \setminus I_j} \frac{a_\ell}{X_j},
\end{aligned}
$$

where we have used the inequality $\ln(1+z) \leq z$ with $z = -a_\ell/X_j$. Using (6), we arrive at $\mathbb{E}[D_j \mid I_j] \geq \frac{1}{32} \frac{\log_2 x}{\log_2 n}$, and averaging over all choices of $I_j \supseteq I_{ib}$ and plugging into (3) finally leads to the inequality in the claim. This also concludes the proof of the lemma. ☐

Proof of Lemma 3.7. Let us put $q := \lceil \log_2 \log_2 n \rceil + 1$. First we note that there cannot be more than q successful phases in the algorithm. Indeed, each successful phase starting with at least 2^4 components reduces the current number of components to at most its square root, so there can't be more than $\log_2 \log_2 n$ such phases, and after the number of components drops below 2^4, there can be at most one additional successful phase.

Let Y_i be the indicator variable of the event that the ith phase is successful, and let $p_i := \text{Prob}[X_{ib} \geq 2]$. We have

$$\mathbb{E}[Y_i] = \text{Prob}[Y_i = 1 \mid X_{ib} \geq 2]\, p_i \geq \frac{1}{2} p_i$$

by Lemma 3.8.

Let us consider the first $8q$ phases, and let $p := p_{8q}$. For proving the lemma it suffices to show that $p \leq \frac{1}{4}$. Indeed, then after $8q$ phases we have $X_{8q} = 1$ with probability at least $\frac{3}{4}$.

For contradiction we thus assume $p > \frac{1}{4}$. Since $p_i \geq p > \frac{1}{4}$ for all $i < 8q$, we have $\mathbb{E}[Y_0 + Y_1 + \cdots + Y_{8q-1}] > 8q \cdot \frac{1}{8} = q$. But this contradicts the fact noted above, namely, that at most q of the Y_i can be 1. This finishes the proof. ☐

3.3 Proof of Theorem 1.1

Theorem 1.1 is now a simple consequence of Lemma 2.3 and of Proposition 3.5. Given directions $u_1, \ldots, u_k$, we let $n := 2f(u_1, \ldots, u_k)$. Assuming $k \geq \lfloor C \log_2 n \log_2 \log_2 n \rfloor$, where C is as in Proposition 3.5, we consider an interchange graph G on n vertices for $u_1, \ldots, u_k$. By Proposition 3.5 it has a subgraph H possessing two color-disjoint connected subgraphs, and hence two color-disjoint spanning trees. Then by Lemma 2.3 the vector of the coordinates of the u_i has to be in the zero set of a nonzero polynomial (depending on H). $\square$

3.4 A limitation of the proof method

Our strategy for finding the two color-disjoint spanning trees was to reduce the graph as long as possible by partitioning along "few-colored" cuts, and then to choose two random disjoint sets of colors. Here we present an example showing that our version of the second part of this strategy, namely, Lemma 3.7, is not far from optimal.

Proposition 3.9 *For infinitely many values of n, there exists a DUPM-graph H on n vertices and with k colors, where $k \geq \frac{1}{2} \log n \log \log n / \log^{(3)} n$ (with $\log^{(3)} n = \log \log \log n$), such that every edge cut $E(H; A, B)$ has at least $\gamma(n, k, A, B)$ colors and the probability that $H[I]$ is connected, where $I \subset [k]$ is a random subset of colors obtained by picking each color of $[k]$ independently with probability $\frac{1}{2}$, tends to 0 as $n \to \infty$. (There is nothing special about probability $\frac{1}{2}$; a similar construction with parameters slightly adjusted works for any constant probability $p < 1$.)*

Let us remark that the graph H we construct in the proof of this proposition does have two color-disjoint spanning trees. Hence the proposition shows only that in a setting as in Lemma 3.7 but with considerably fewer than $\log n \log \log n$ colors, we cannot get the desired trees by a purely random choice of colors—it might still be that they could be obtained by some more clever strategy.

One might also wonder whether the initial reduction, partitioning the graph along few-colored cuts, couldn't be improved. We believe that an improvement cannot be achieved in this way either, but we prefer not to formalize our heuristic reasoning here, since it doesn't seem to bring anything useful towards resolving the problem (while the example in Proposition 3.9 might give some inspiration for future research).

Sketch of proof of Proposition 3.9. We recall that the *Cartesian product* (or *graph product*) $G_1 \Box G_2$ of two graphs has vertex set $V(G_1) \times V(G_2)$ and two vertices (v_1, v_2) and (v_1', v_2') in this product are connected by an edge if $v_1 = v_1'$ and $\{v_2, v_2'\} \in E(G_2)$ or if $v_2 = v_2'$ and $\{v_1, v_1'\} \in E(G_1)$. (So the d-dimensional cube is the dth Cartesian power of an edge.)

We assume n to be sufficiently large and we set (ignoring integrality issues) $b := \frac{1}{2} \log\log n$, assuming b even, and $d := \log n / \log b$. We let H be the d-fold Cartesian product $H_1 \Box H_2 \Box \cdots \Box H_d$, where each H_i is a complete DUPM graph on b vertices such that the sets of colors for H_i and H_j are disjoint for $i \neq j$. Then H is a DUPM-graph on $b^d = 2^{d \log b} = n$ vertices with $k = d(b-1) \in [\frac{1}{2} \log n \log\log n / \log^{(3)} n, \log n \log\log n / \log^{(3)} n]$ colors (each edge in the product "comes from" a particular edge of a particular H_i and it inherits the color of that edge).

If $I \subset [k]$ is a random set of colors as in the proposition, then with probability larger than 2^{-b}, the set I contains no color among those used on H_i, and these events for $i = 1, 2, \ldots, d$ are independent.[3] Hence $H[I]$ is connected with probability at most $(1 - 2^{-b})^d \leq \exp(-2^{-b} d) = o(1)$ as $n \to \infty$.

It remains to verify that all cuts in H have sufficiently many colors. To this end, we show that any subgraph of H on a set $A \subseteq V(H)$ of m vertices has average degree at most $((b-1)/\log b) \log m$, generalizing the well-known isoperimetric inequality for the cube corresponding to the case $b = 2$ (then the average number of edges going from a vertex of A to the complement of A, $|A| \leq \frac{1}{2}|V(H)|$, is at least $d(b-1) - ((b-1)/\log b) \log |A| = ((b-1)/\log b) \log(n/|A|) \geq \gamma(n, k, A, V(H) \setminus A)$). This can be proved by induction; namely, by induction on i we prove that any subgraph on at most m vertices of $H_1 \Box \cdots \Box H_i$ has average degree at most $((b-1)/\log b) \log m$. For $i = 1$ this holds because the average degree is at most $m - 1$ and for $1 \leq m \leq b$ we have $(m-1)/\log m \leq (b-1)/\log b$. For $i > 1$ we consider the partition $A = A_1 \dot\cup \cdots \dot\cup A_b$ of a set $A \subseteq V(H_1 \Box \cdots \Box H_i)$ according to the projection on the last coordinate, we estimate the number of edges on each A_j by induction and we add the edges among distinct A_j, whose number is at most $\sum_{j_1 < j_2} \min(|A_{j_1}|, |A_{j_2}|)$. Elementary calculations, which we omit, finish the inductive step. $\Box$

[3]Here the model of picking the colors independently is convenient; if we fixed the number of colors picked, we would not get independence and the calculations would become more complicated.

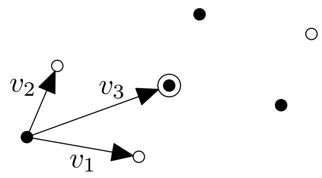

Figure 4: A construction for $k = 3$, using $v_1 + v_2 = v_3$.

4 The upper bound

Here we are going to prove Theorem 1.2. In order to prove an upper bound for $F(k)$, we need to construct, for any given k-tuple of distinct directions $u_1, u_2, \dots, u_k$, two point sets A and B with identical X-rays in all of these directions.

The basic approach. All of the constructions are modification of the basic example mentioned in the introduction, which provides the upper bound $F(k) < 2^{k-1}$. Given $u_1, u_2, \dots, u_k$, we choose suitable nonzero scalars $\alpha_1, \alpha_2, \dots, \alpha_k$ and we define the multisets A', B' as follows:

$$A' := \Big\{ v(I) : I \subseteq [k],\ |I| \text{ even} \Big\}, \quad B' = \Big\{ v(I) : I \subseteq [k],\ |I| \text{ odd} \Big\},$$

where $v(I) = \sum_{i \in I} v_i$ and $v_i = \alpha_i u_i$. For easier exposition we will call the points of A' *black* and the points of B' *white*. We have $|A'| = |B'| = \frac{1}{2} 2^k$ (where each point is counted with multiplicity), and as was discussed in the introduction, A' and B' have identical X-rays in each of the directions u_i.

In order to improve the upper bound of 2^{k-1}, we try to choose the α_i so that A' and B' have many points in common. After removing the common points the resulting multisets still have identical X-rays but they are smaller. More precisely, when a black point coincides with a white point, they cancel each other out (so we have to respect multiplicities; if, say, five black points and three white points all occupy the same location, after cancellation we obtain two black points at that location). The resulting multisets after the cancellation will be denoted by A and B.

An example for $k = 3$. The simplest nontrivial example, with $k = 3$, is depicted in Fig. 4, and it shows that $F(3) < 3$ (the circled point in the middle is in both A' and B' and thus it cancels out).

To make this construction work, we need to choose nonzero $\alpha_1, \alpha_2, \alpha_3$ so that $v_1 + v_2 = v_3$, and we should check that this is possible for any triple of distinct directions u_1, u_2, u_3. This is easy: since the vectors u_1 and u_2 define distinct directions, they are linearly independent, and each vector in the plane is a linear combination of them. In particular, we can write $u_3 = \alpha_1 u_1 + \alpha_2 u_2$, and neither of α_1, α_2 can be 0, for otherwise, u_3 would define the same direction as u_2 or u_1.

We should also note that A' and B' are always sets (no multiple points) in this particular construction, and hence the upper bound holds for $F(3)$ and not only for $F_{\mathrm{mult}}(3)$. Indeed, if two black points were equal or two white points were equal, then there would be two distinct sets $I, J \subseteq [3]$ with $|I| \equiv |J| \pmod 2$ and with $v(I) = v(J)$. This yields a linear combination of the u_j equal to 0 and supported on the symmetric difference $I \triangle J$. But we necessarily have $|I \triangle J| = 2$, which would mean that one of the u_j is a multiple of another, which is excluded.

A product argument. Before considering more involved constructions of the just described type, we present a general "product" trick.

Lemma 4.1 *Let us set $G(k) := 2(F(k) + 1)$. Then we have $G(k + \ell) \leq G(k)G(\ell)$ for any $k, \ell \geq 1$, and in particular, $F(mk) \leq \frac{1}{2}(2F(k) + 2)^m - 1$. An analogous statement holds for F_{mult}.*

Proof. Let $u_1, \ldots, u_k$ and $w_1, \ldots, w_\ell$ be directions, all of them distinct. Let A_u and B_u be sets of $F(k) + 1$ points each, $A_u \cap B_u = \emptyset$, that have identical X-rays in each of the directions u_i, and similarly, A_w and B_w are disjoint $(F(\ell) + 1)$-point sets with identical X-rays in each of the directions w_j.

We choose a sufficiently large real number t and set

$$A := (A_u + tA_w) \dot{\cup} (B_u + tB_w), \quad B := (A_u + tB_w) \dot{\cup} (B_u + tA_w),$$

where "+" represents the Minkowski addition, i.e., $X + Y = \{x + y : x \in X, y \in Y\}$ (and $X + Y$ is understood as a multiset). We have $|A|, |B| = 2(F(k) + 1)(F(\ell) + 1)$, and it is easy to check that A and B have identical X-rays in each of the directions u_i and w_j. Indeed, for the direction u_1, for example, let $\varphi : A_u \to B_u$ be a bijection assigning to each point of A_u a point of B_u lying on the same line parallel to u_1. Each point $a \in A$ can be written as $a_u + ta_w$, $a_u \in A_u$, $a_w \in A_w$, or as $b_u + tb_w$, $b_u \in B_u$, $b_w \in B_w$. Then we define a bijection $\psi : A \to B$ by $\psi(a_u + ta_w) := \varphi(a_u) + ta_w$,

$\psi(b_u + tb_w) := \varphi^{-1}(b_u) + tb_w$, and it is immediate that a and $\psi(a)$ lie on the same line parallel to u_1.

It remains to verify that, with appropriate choice of t, A and B are sets and $A \cap B = \emptyset$. We choose t so large that the diameter of $A_u \cup B_u$ is smaller than the distance of any two points in $t(A_w \cup B_w)$. Then if two points $a, a' \in A \cup B$ coincide, we can write $a = x + ty$, $a' = x' + ty'$, $x, x' \in A_u \cup B_u$, $y, y' \in A_w \cup B_w$, and by the choice of t we obtain $y = y'$ and $x = x'$. □

Proof of Theorem 1.2 for sets. By combining Lemma 4.1 with the example above showing $F(3) \leq 2$, we have $F(3m) \leq \frac{1}{2}6^m - 1$, and by monotonicity of F we then obtain $F(k) = O((6^{1/3})^k)$ for all k. This proves the claim in Theorem 1.2 concerning sets.

A better construction for multisets. Here we return to the basic construction, depending on a good choice of the α_i. We consider the case $k = 9$. We set $\alpha_9 = 1$, say, and we choose nonzero α_3 and α_6 so that $v_9 = v_3 + v_6$ (by the same argument as in the example for $k = 3$). Then we set the remaining α_i so that $v_1 + v_2 = v_3$, $v_4 + v_5 = v_6$, and $v_7 + v_8 = v_9$.

As we have verified by a computer, in this way 157 pairs of points always cancel out and we are left with $|A| = |B| \leq 99$ (see Fig. 5). Hence $F_{\text{mult}}(9) < 99$, and by Lemma 4.1 we then have $F_{\text{mult}}(k) = O((198^{1/9})^k)$, which verifies the part of Theorem 1.2 concerning F_{mult}.

An example where the construction gives multisets. It can easily be shown that for a generic choice of $u_1, \dots, u_9$, the construction just given produces *sets* A' and B' (no two black points and no two white points coincide). However, it turns out that if we insist on the particular dependences among the v_i used in the construction, i.e. on the relations $v_9 = v_3 + v_6$, $v_3 = v_1 + v_2$, $v_4 + v_5 = v_6$, and $v_7 + v_8 = v_9$, then for some choices of the u_i multiple points in A or B cannot be avoided. Indeed, let us consider the following 9 distinct directions:

$$\begin{array}{lllll} u_1 = (-1,4), & u_2 = (2,3), & u_3 = (1,7), & u_4 = (-1,5), & u_5 = (2,1), \\ u_6 = (1,6), & u_7 = (1,1), & u_8 = (1,12), & u_9 = (2,13). & \end{array}$$

Up to scaling, the only choice of the α_i providing the desired relations among the v_i is $\alpha_1 = \alpha_2 = \cdots = \alpha_9 = 1$. But then we also have $u_1 + u_3 = (0,13) = u_4 + u_6$, i.e., two points in A' coincide. We have also checked that these two black points don't cancel out with any white points, and so A also has multiple points

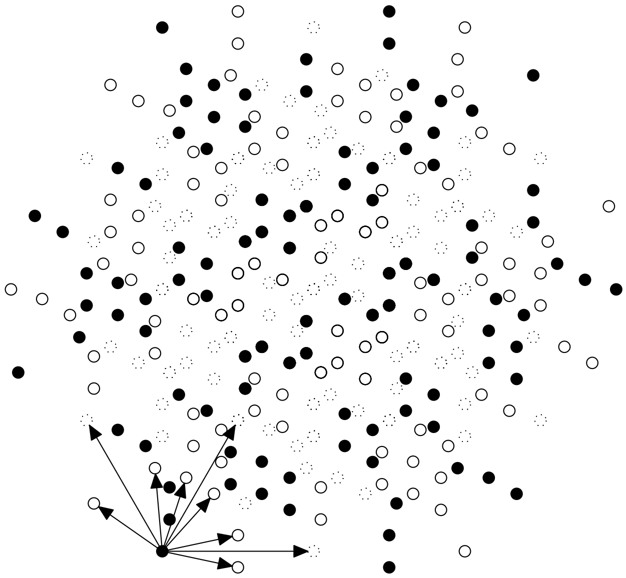

Figure 5: Sets A, B with the same X-rays in 9 directions, with $|A| = |B| = 99$. Dotted circles denote points that have been cancelled out.

It seems that for the choices of u_i that lead to multiple points in this particular construction, we can use alternative constructions, obtained from the present construction by a suitable renumbering of the u_i. But so far we have no systematic approach to such problems.

Concluding remarks. It is clear that the upper bound for F_{mult} can be improved further, by using a larger number of directions in the basic construction. For example, we can take $k = 27$, divide the 27 vectors u_i into 3 groups by 9, apply the choice of the α_i as in the construction for $k = 9$ for each of these groups separately, and then force the additional relation $v_{27} = v_9 + v_{18}$, say. We can then continue and get better and better constructions for k's of the form 3^m.

But analyzing these constructions seems demanding; already the case $k = 27$ looks challenging even with the help of a computer. Moreover, it appears that the improvements get very small as m grows. There might be a nice theory hiding behind such constructions, but since this is only one particular class of examples, it is not clear whether understanding them better would be of much help for estimating the functions F and F_{mult}, which was our main objective in this paper.

5 Open problems

Our work was motivated by a question of Holub [Hol03], which in our language can be re-formulated as estimating $f(u_1, \ldots, u_k)$ for $u_i = (1, i)$ (i.e., u_i is the direction with slope i). The problem emerged in study of certain systems of equations over free semigroups; originally it deals with a set $F = \{x \mapsto a_i x + b_i : i = 1, 2, \ldots, n\}$ of n linear functions, and it asks for what k we can be sure that every F is uniquely determined by the multiset of its values for $x = 1, 2, \ldots, k$. An almost identical question appeared few years later as a problem by Hillar and Levine in the American Mathematical Monthly [HL06]. In our geometric context, these special directions $(1, i)$ do not seem to have any special significance; yet, even disregarding the algebraic motivation, it would be interesting to have a simple explicit set of directions that allows unique reconstruction of large sets.[4] Unfortunately, our current methods are unlikely to yield any nontrivial lower bound in this case.

[4]We can get an "explicit" set of directions from Theorem 1.1, namely, k directions whose coordinates are algebraically independent real numbers, but this is perhaps not what one might call "simple".

References

[Alo95] N. Alon. A note on network reliability. In *Discrete Probability and Algorithms (D. Aldous, P. Diaconis, J. Spencer and J. M. Steele eds.), IMA Volumes in Mathematics and its applications, Vol. 72*, pages 11–14. Springer Verlag, Berlin etc., 1995.

[BDLNP01] E. Barcucci, A. Del Lungo, M. Nivat, and R. Pinzani. X-rays characterizing some classes of discrete sets. *Linear algebra and its applications*, 339:3–21, 2001.

[Big93] N. Biggs. *Algebraic Graph Theory*. Cambridge Univ. Press, Cambridge, 1993. 2nd edition.

[BL90] G. Bianchi and M. Longinetti. Reconstructing plane sets from projections. *Discrete and Computational Geometry*, 5:223–242, 1990.

[DGP06] P. Dulio, R. J. Gardner, and C. Peri. Discrete point X-rays. *SIAM J. Discrete Math*, 20(1):171–188, 2006.

[Gar92] R. J. Gardner. X-rays of polygons. *Discrete and Computational Geometry*, 7(1):281–293, 1992.

[Gar06] R. J. Gardner. *Geometric Tomography*. Western Washington University, 2nd edition, 2006.

[GG97] R. J. Gardner and P. Gritzmann. Discrete tomography: determination of finite sets by X-rays. *Trans. Amer. Math. Soc.*, 349:2271–2295, 1997.

[Hep56] A. Heppes. On the determination of probability distributions of more dimensions by thier projections. *Acta Math. Acad. Sci. Hungar.*, 7:403–410, 1956.

[HL06] C. Hillar and L. Levine. Problem 11223. *Amer. Math. Monthly*, 5:459, 2006.

[Hol03] Š. Holub. Private communication (e-mail), 2003.

[Rén52] A. Rényi. On projections of probability distributions. *Acta Math. Acad. Sci. Hungar.*, 3:131–142, 1952.

[Tót03] G. Tóth. Private communication (through Attila Pór), 2003.