

Colouring Multijoints

Anthony Carbery · Stefán Ingi Valdimarsson

Received: 6 September 2013 / Revised: 28 May 2014 / Accepted: 5 July 2014 /
Published online: 21 August 2014
© Springer Science+Business Media New York 2014

Abstract Let \mathbb{F} be a field, let L_1, \dots, L_d be pairwise disjoint collections of lines in \mathbb{F}^d , and let $\mathcal{L} = \{L_1, \dots, L_d\}$. We say that a point $x \in \mathbb{F}^d$ is a *multijoint of \mathcal{L}* if x lies on a line from each of the collections in \mathcal{L} , and moreover the directions of these lines span \mathbb{F}^d . We prove that there exists a constant C_d such that if \mathcal{L} is a generic family of collections of lines in \mathbb{F}^d and J is a set of multijoints of \mathcal{L} , then there exists a d -colouring $\kappa : J \rightarrow \{1, 2, \dots, d\}$ such that for each j , for each $l \in L_j$ we have $|\{x \in J \cap l : \kappa(x) = j\}| \leq C_d |J|^{1/d}$.

Keywords Multijoints · Colourings · Multilinear Kakeya

1 Introduction

Let \mathbb{F} be a field, let L_1, \dots, L_d be pairwise disjoint collections of lines in \mathbb{F}^d , and let $\mathcal{L} = \{L_1, \dots, L_d\}$. We say that a point $x \in \mathbb{F}^d$ is a *multijoint of \mathcal{L}* if x lies on a line from each of the collections in \mathcal{L} , and moreover the directions of these lines span \mathbb{F}^d .

Regard each line of L_j as being coloured with colour j . In this note we address the problem of colouring the set J of multijoints of \mathcal{L} with as few colours as possible in such a way that no line of a given colour contains too many points of that same colour. We need to make these notions precise, and do so in the statement of our main result. Further clarification and a discussion of the context of the result follows in the remarks after its statement. The family \mathcal{L} is said to be *generic* if whenever $l_j \in L_j$ meet at x , then the directions of the l_j span \mathbb{F}^d .

A. Carbery
School of Mathematics and Maxwell Institute for Mathematical Sciences, University of Edinburgh,
JCMB, King's Buildings, Mayfield Road, Edinburgh EH9 3JZ, Scotland, UK
e-mail: a.carbery@ed.ac.uk

S. I. Valdimarsson
Science Institute, University of Iceland, Dunhagi 3, 107 Reykjavík, Iceland
e-mail: siv@hi.is

Theorem 1 *Let \mathcal{L} be a generic family of collections of lines in \mathbb{F}^d as above, and let J be a (finite) set of multijoints of \mathcal{L} . Then there exists a constant C_d which depends only on the dimension d and not on \mathcal{L} or J , and a d -colouring $\kappa : J \rightarrow \{1, 2, \dots, d\}$ such that for each j , for each $l \in L_j$,*

$$|\{x \in J \cap l : \kappa(x) = j\}| \leq C_d |J|^{1/d}.$$

Remark 1 We emphasise that while an $x \in J$ can arise as a multijoint in several ways (i.e. there may be many d -tuples of lines $(l_1, \dots, l_d) \in \mathcal{L}_1 \times \dots \times \mathcal{L}_d$ with $x \in l_j$ for all j) we only associate a single colour to it. That is, a multijoint is a point of \mathbb{F}^d , not a d -tuple of lines.

Remark 2 Matters are trivial if we allow more than d colours for J : simply colour every point of J with colour $d+1$. We therefore consider d -colourings $\kappa : J \rightarrow \{1, 2, \dots, d\}$ of J .

Remark 3 We cannot hope for each line of a given colour to contain at most about $|J|^\beta$ points of J of the same colour unless $\beta \geq 1/d$. To see this consider the monkey-bar/jungle-gym example where L_j consists of N^{d-1} lines parallel to the x_j -axis passing through the points $(m_1, \dots, m_{j-1}, 0, m_{j+1}, \dots, m_d)$ for $m_i \in \{1, \dots, N\}$. If each line of L_j contains at most K multijoints of colour j then there are at most $N^{d-1}K$ multijoints of colour j altogether and hence at most $dN^{d-1}K$ multijoints altogether. But there are N^d multijoints in this example, so we must have $N^d \leq dN^{d-1}K$. Hence K must satisfy $K \geq N/d = |J|^{1/d}/d$.

Remark 4 We cannot expect in general to use fewer than d colours. To illustrate this in the case $d = 3$, put $2N$ red lines parallel to e_1 passing through the points $(0, j, 0)$ for $1 \leq j \leq N$ and $(0, 0, j)$ for $1 \leq j \leq N$, and similarly put $2N$ blue lines parallel to e_2 and $2N$ green lines parallel to e_3 in the corresponding places. Then on the plane $x_3 = 0$ we have an $N \times N$ square lattice of N red lines parallel to e_1 and N blue lines parallel to e_2 . Through each lattice point on this plane put a green line to make it a multijoint in J but in such a way that no new multijoints in J are created. Similarly add red lines through lattice points on the plane $x_2 = 0$ and blue lines through lattice points on the plane $x_1 = 0$. Altogether we now have $3N^2$ multijoints, with the colours red, blue and green in symmetry. Can we colour this arrangement of multijoints using only two colours, say red and blue, in such a way that a line of a given colour contains at most $\sim N^{2/3}$ points of that colour? If so, considering the multijoints in the the plane $x_3 = 0$, every red line would have at most $\sim N^{2/3}$ red multijoints, so there would be at most $\sim N \times N^{2/3} = N^{5/3}$ red multijoints on this plane, and similarly at most $\sim N^{5/3}$ blue multijoints. Hence there would be at most $\sim N^{5/3}$ multijoints on this plane, when in fact there are $\sim N^2$. This contradiction shows that we cannot colour this arrangement with fewer than three colours, and similar examples in higher dimensions show that in \mathbb{F}^d we will need d colours in general.

Our setting with d families of lines is a variant of the setting of the so-called joints problem. There we have a single collection L of lines in \mathbb{F}^d , and we define a *joint of L*

to be any point which lies at the intersection of d lines from L with the condition that the set of directions of those d lines should span \mathbb{F}^d . In recent years there has been quite a bit of interest in the joints problem and it is now known that if J is the set of joints of L then we have

$$|J| \leq C_d |L|^{d/(d-1)}, \tag{1}$$

where C_d depends only on the dimension d . This was originally proved in the case $\mathbb{F} = \mathbb{R}$ by Guth and Katz in [7] for $d = 3$, then for a general $d \geq 3$ by Quilodrán [10] and independently by Kaplan, Sharir and Shustin [9]. Treatments of the extension to general fields are in [2, 5] and [11].

Let us temporarily pass to the situation $\mathbb{F} = \mathbb{R}$. Let N_1, \dots, N_d be natural numbers and consider the grid \mathcal{G} of points $\{1, 2, \dots, N_1\} \times \dots \times \{1, 2, \dots, N_d\} \subseteq \mathbb{R}^d$. Let us consider the family L_j of $N_1 N_2 \dots N_{j-1} N_{j+1} \dots N_d$ lines parallel to the j th coordinate axes passing through the points of the grid \mathcal{G} . Then it is easy to see that $|J| = N_1 \dots N_d$ while $\prod_{j=1}^d |L_j| = (N_1 \dots N_d)^{d-1}$. This example therefore gives rise to the natural question in our setting as to whether we have, with J now being the set of multijoints of \mathcal{L} ,

$$|J| \leq C_d \prod_{j=1}^d |L_j|^{1/(d-1)}, \tag{2}$$

where C_d depends only on the dimension d . At the moment this question seems difficult (except in two dimensions, in which case it is obvious)¹ and we instead consider the related problem described above.

Let us explain the relevance of our result to inequality (2). In [4] Dvir proved the finite field Kakeya conjecture. Since then, his central idea, dubbed the *polynomial method*, has been used extensively, among other things in the cited work on the joints problem. In another direction, Guth [6] extended the polynomial method to prove the endpoint case of the multilinear Kakeya conjecture in \mathbb{R}^d . This is a continuous version of inequality (2) introduced above. His proof used algebraic topology but see [3] for a treatment which relies only on the Borsuk–Ulam theorem.

Suppose we have d families \mathcal{T}_j of doubly-infinite tubes T_j in \mathbb{R}^d of infinite length and unit cross-section, and suppose that each tube in \mathcal{T}_j points approximately in the direction of the j -th standard basis vector e_j .² The main substance of Guth’s argument consists of proving that for every non-negative function M there exist functions S_j , $j = 1, \dots, d$, such that

$$M(x) \leq \prod_{j=1}^d S_j(x)^{1/d} \tag{3}$$

¹ However, Iliopoulou [8] has recently informed us that she has established (2) when $\mathbb{F} = \mathbb{R}$ for arbitrary d and also for arbitrary \mathbb{F} when $d = 3$. In neither case is the genericity hypothesis needed.

² Guth’s set up is more relaxed than this, see also [1].

and

$$\sum_x S_j(x) \leq C_d \|M\|_d. \tag{4}$$

The domain of the functions M and S_j is the set of unit cubes in \mathbb{R}^d , and inequality (3) is supposed to hold for each cube. The sum in inequality (4) is over cubes x meeting a tube $T_j \in \mathcal{T}_j$, and inequality (4) is supposed to hold for each tube in the collection \mathcal{T}_j , for all $j = 1, \dots, d$. Loosely speaking, Guth’s functions $S_j(x)$ are constructed geometrically as the directional surface areas of $Z_p \cap x$ where Z_p is the zero set of a polynomial p of degree dominated by $\|M\|_d$ which bisects each cube x in the domain of M in a particularly strong way.

In considering inequality (2) one is naturally led to consider inequalities (3) and (4) where the tubes are replaced by lines (of zero width), the unit cubes by points and where we suppose that if $x \in l_j \in L_j$ for $j = 1, \dots, d$ then the directions of the l_j should span \mathbb{F}^d , i.e. that x is a multijoint according to our definition above. Note that in the case of \mathbb{R}^d straightforward limiting arguments applied to the results of [6] and [1] do not yield an answer to the question of the satisfiability of (3) and (4), or of the validity of (2) in this setting. Note also that there is no obvious analogue of the notion of directional surface area in our discrete setting.

If in (3) and (4) we replace the general nonnegative function M by a characteristic function χ_J of a set of multijoints J , and the geometric mean by the (larger) arithmetic mean we arrive at the (easier) problem of finding S_j such that

$$\chi_J(x) \leq \frac{1}{d} \sum_{j=1}^d S_j(x) \tag{5}$$

and

$$\sum_x S_j(x) \leq C_d |J|^{1/d}. \tag{6}$$

Theorem 1 is equivalent to this new problem: if we have such S_j , for each x , choose a j with $S_j(x) \geq 1$ and assign colour j to x ; conversely, if we have a colouring satisfying the conclusion of Theorem 1, declare $S_j(x) = d$ if x has colour j and $S_j(x) = 0$ otherwise. (See also Remark 7 at the end of the note.)

We now make some remarks on the relation between our colouring result and the joints theorem (1). In the first place, our result implies a weak form of the joints theorem in so far as if we have a generic family of lines $\{L_1, \dots, L_d\}$ our result enables us to estimate the size of its set of multijoints J as

$$\begin{aligned} |J| &= \sum_{x \in J} 1 \leq C_d \sum_{x \in J} \sum_{j=1}^d S_j(x) = C_d \sum_{j=1}^d \sum_{x \in J} S_j(x) \\ &\leq C_d \sum_{j=1}^d \sum_{l_j \in L_j} \sum_{x \in J \cap l_j} S_j(x) \leq C_d \sum_{j=1}^d |L_j| |J|^{1/d} \end{aligned}$$

from which it follows that $|J| \leq C_d(\sum_{j=1}^d |L_j|)^{d/(d-1)}$. Secondly, our argument uses important aspects of the arguments used to prove the joints theorem. In particular it makes use of a fundamental lemma of Quilodrán from [10] (see Lemma 1 below) and indeed the strategy of his original proof of the joints theorem. Quilodrán’s lemma is proved using the polynomial method and so there is a polynomial underlying our own colouring theorem; however, due to the iterative nature of our argument it is not clear that one can express the quantities S_j directly in terms of the zero set of a suitable polynomial.

In this direction, one strategy for proving our theorem might be as follows: choose a polynomial p of degree dominated by $|J|^{1/d}$ which vanishes at each member of J . (One can do this simply by observing that if $\deg p \gtrsim_d |J|^{1/d}$ one can choose the coefficients of p such that it satisfies the $|J|$ linear equations $p(x) = 0, x \in J$.) Then, if x is a multijoint such that $\nabla p(x) \neq 0$, it must be the case that for some j , for all $l_j \in L_j$, we have $l_j \not\subseteq Z_p$. (Otherwise, for all j there exists an $l_j \in L_j$ such that p vanishes identically on l_j , and thus its directional derivative $\omega_{l_j} \cdot \nabla p$ also vanishes on l_j ; but for $x \in \cap_j l_j$ we have that $\nabla p(x)$ is a linear combination of $\omega_{l_j} \cdot \nabla p(x)$ since the directions ω_{l_j} span \mathbb{F}^d .) If so, choose the colour $\kappa(x)$ of x to be j and observe that if $l_j \in L_j$ then $\{x \in l_j \cap J : \kappa(x) = j \text{ and } \nabla p(x) \neq 0\} \subseteq l_j \cap Z_p$ which has cardinality at most $\deg p \leq C_d |J|^{1/d}$. However this argument does not deal with any critical points of p and for this reason we are forced to adopt an alternative strategy as below.

Finally, we remark that when $d = 2$ there is a simple *ad hoc* argument leading to the conclusion of Theorem 1. Indeed, suppose in \mathbb{F}^2 we have a family of blue lines and a family of red lines (with no line having both colours). If a blue line contains at most $\sqrt{2}|J|^{1/2}$ bijoints, colour all of those bijoints blue. Colour all other bijoints red. Suppose we have a red line with as many as $k = \sqrt{2}|J|^{1/2} + 1$ red bijoints on it. Then each of these bijoints is on a (different) blue line, which must therefore contain more than $\sqrt{2}|J|^{1/2}$ bijoints (as they are not all blue). Hence there are more than

$$\begin{aligned} & \sqrt{2}|J|^{1/2} + (\sqrt{2}|J|^{1/2} - 1) + (\sqrt{2}|J|^{1/2} - 2) + \dots + (\sqrt{2}|J|^{1/2} - (k - 1)) \\ &= \sqrt{2}k|J|^{1/2} - (k - 1)k/2 = \sqrt{2}(\sqrt{2}|J|^{1/2} + 1)|J|^{1/2} - \sqrt{2}|J|^{1/2}(\sqrt{2}|J|^{1/2} + 1)/2 \\ &= |J| + \frac{|J|^{1/2}}{\sqrt{2}} > |J| \end{aligned}$$

distinct bijoints altogether, which is a contradiction. Hence each red line also contains at most $\sqrt{2}|J|^{1/2}$ red points too.

2 Proof of Theorem 1

We let m be a positive integer and J be a set of multijoints of \mathcal{L} . We say that a colouring $\kappa : J \rightarrow \{1, \dots, d\}$ is m -bounded if for any $j = 1, \dots, d$ and any $l_j \in L_j$ we have

$$|\{x \in l_j : \kappa(x) = j\}| \leq m.$$

Otherwise the colouring will be called *m-unbounded*. So in an *m-unbounded* colouring there is some line containing $(m + 1)$ members of J with the same colour as the line. Theorem 1 can be restated as:

Theorem 2 *Let \mathcal{L} be a generic family of collections of lines in \mathbb{F}^d as above and let J be a set of multijoints of \mathcal{L} . Then there exists a constant C_d which depends only on the dimension d and not on \mathcal{L} or J and an integer m with $m \leq C_d |J|^{1/d}$ such that J is colourable with an m -bounded colouring.*

Proof Fix a positive integer m and the set \mathcal{L} . Let J_c be a set of multijoints of \mathcal{L} which is colourable with an m -bounded colouring and let x_0 be a multijoint of \mathcal{L} which does not belong to J_c .

Our aim is to prove the following claim.

Claim 1 *If $m > C_d |J_c|^{1/d}$ then $\tilde{J} = J_c \cup \{x_0\}$ is colourable with an m -bounded colouring.*

This claim immediately proves the theorem: every singleton subset of J is trivially colourable with an m -bounded colouring and the claim allows us to add points one at a time, preserving the property of being colourable with an m -bounded colouring, until the size of the set reaches $(m/C_d)^d$, which by assumption will not happen before we exhaust J .³

We now turn to the proof of the claim. To simplify notation we use J for what is called J_c in the statement of the claim. Let us label the elements of $J = \{x_1, \dots, x_v\}$ and order the multijoints in \tilde{J} according to the indices. Let K be the set of colourings of J which are m -bounded.

We wish to define a strict partial ordering on K . To do this we construct for every $\kappa \in K$ a *coloured rooted tree* T whose vertices belong to \tilde{J} . The tree will be rooted at x_0 (which is achromatic) and all the other vertices will be members of J and will be assigned the colour given to them by κ . (We shall not colour the edges of the tree.)

2.1 Construction of the Tree

Let us fix a $\kappa \in K$ and describe the construction of the tree T with an iterative process.

At the 0-th step, the tree T_0 has one vertex, x_0 , and no edges. We will maintain an ordering on the vertices, based primarily on the step in which a vertex gets added and secondarily on the ordering inherited from \tilde{J} . In accordance with that we say that x_0 is the first element of the tree and give it the alternative name y_1 .

At the i -th step we consider T_{i-1} , and either construct a T_i , or else stop the procedure and declare $T := T_{i-1}$. We consider the i -th element of the tree T_{i-1} , which we call y_i , and construct T_i by adding one or more children from amongst the members of \tilde{J} not already in T_{i-1} to T_i , and then connect y_i to its children with edges.

If there is no i -th element in the tree T_{i-1} we say that the tree is *fully constructed* and we define $T := T_{i-1}$. Clearly this must happen before or when we reach step $|J|$.

³ As we shall see below in Sect. 2.4, our approach constructs a suitable colouring of \tilde{J} .

Otherwise, we proceed as follows. The children of y_i will be the elements of \tilde{J} which are not already in the tree, and which are the reasons that we may not change the colour of y_i without the colouring becoming m -unbounded or at least in danger of becoming so. Specifically, for each colour j different from $\kappa(y_i)$ ⁴ let $L_j^{(i)}$ be the subset of L_j consisting of the lines l_j going through y_i such that

$$|\{x \in l_j \cap J : \kappa(x) = j\}| + |\{x \in l_j \cap T_{i-1} \cap J : \kappa(x) \neq j\}| \geq m. \quad (7)$$

(Here and later we abuse notation and use T_{i-1} also to denote the vertex set of the tree T_{i-1} .) If the collection $L_j^{(i)}$ is empty for some $j \neq \kappa(y_i)$ we say that the colouring is *advanceable at step i* .⁵ Then we stop the construction, and declare $T := T_{i-1}$.

Otherwise, if the colouring is not advanceable at step i , we define for each $j \neq \kappa(y_i)$ the set $I_j^{(i)}$ which consists of all the points of J of colour j on any line from $L_j^{(i)}$, excluding those points which are already in the tree T_{i-1} . We let the tree T_i be the tree T_{i-1} with the points from $I_j^{(i)}$ for all $j \neq \kappa(y_i)$ added as vertices, specifically as children of y_i . The edges of T_i are those of T_{i-1} together with edges linking y_i to each of its children. Note that no child has the same colour as its parent.⁶

We remark that it is possible that $I_j^{(i)}$ may be empty for some colour j even though $L_j^{(i)}$ is non-empty. This is the case if all the points of colour j in J which lie on any line in $L_j^{(i)}$ are vertices of T_{i-1} . If all of the $I_j^{(i)}$ are empty we let $T_i := T_{i-1}$.

If the colouring is not advanceable at any step then eventually the tree will become fully constructed. We call such a colouring *non-advanceable*.

2.2 The First Stage in the Construction

To fix ideas, let us run through the first stage of the construction. We have $T_0 = \{y_1\} = \{x_0\}$ and so there is a first element of T_0 . For each $j \in \{1, 2, \dots, d\}$ we have that $L_j^{(1)}$ is the subset of L_j consisting of the lines l_j going through y_1 such that

$$|\{x \in l_j \cap J : \kappa(x) = j\}| \geq m \quad (8)$$

since the second term on the left-hand side of (7) is zero. Since κ is m -bounded, (8) means that $|\{x \in l_j \cap J : \kappa(x) = j\}| = m$. Now either $L_j^{(1)}$ is empty for some colour j , or it is non-empty for all j .

⁴ In the case $i = 1$ this simply means all colours j . This understanding applies in several places below.

⁵ See Sect. 2.4 Claim 2(b) for the reason we use this terminology.

⁶ Note also that the sequence of points y_1, y_2, y_3, \dots is dependent on the particular colouring κ we are looking at, and two different colourings will give rise in general to two different sequences of points. We should strictly therefore label our sequence as $y_1(\kappa), y_2(\kappa), y_3(\kappa), \dots$ in order to emphasise this point. However we choose not to do this in order that our subsequent notation which involves expressions such as $\kappa(y_i)$ —which strictly would be written $\kappa(y_i(\kappa))$ and means the colour that the colouring κ gives to the i 'th member of the construction of the tree for the colouring κ —should not become too unwieldy.

- (1) In the first case we have that κ is advanceable at step 1, we stop the procedure and declare $T = T_0$. Note that in this case, if for a certain j , $L_j^{(1)} = \emptyset$, then for all $l_j \in L_j$ passing through y_1 we have

$$|\{x \in l_j \cap J : \kappa(x) = j\}| < m.$$

In this case we can simply assign the colour j to x_0 and we are done. (This observation will be important below.)

- (2) Otherwise, when $L_j^{(1)} \neq \emptyset$ for all j , we have

$$I_j^{(1)} = \{x \in J : \kappa(x) = j \text{ and } x \in l_j \text{ for some } l_j \in L_j^{(1)}\}$$

and we note that since every line in each $L_j^{(1)}$ has exactly m members of J of colour j on it, each $I_j^{(1)}$ has at least m members, and so T_1 will be a proper extension of T_0 , and in particular will have a second member ready for the construction of T_2 .

The construction of the tree T depends on the colouring κ ; when we wish to emphasise this we shall use the notation $T(\kappa)$ and likewise $I_j^{(i)}(\kappa)$ to highlight this dependence.

2.3 A Strict Partial Ordering

Now we turn to the definition of the strict partial ordering on K . Take $\kappa_1, \kappa_2 \in K$ and construct the trees $T(\kappa_1)$ and $T(\kappa_2)$. We say that κ_1 is *more advanced than* κ_2 at level i_0 if

- (i) $I_j^{(i)}(\kappa_1) = I_j^{(i)}(\kappa_2)$ for all $j \neq \kappa_1(y_i)$ ⁷ and for all $i < i_0$;
- (ii) $I_j^{(i_0)}(\kappa_1) \subseteq I_j^{(i_0)}(\kappa_2)$ for all $j \neq \kappa_1(y_{i_0})$;
- (iii) at least one of the inclusions in item (ii) is strict.

In other words, κ_1 is more advanced than κ_2 at level i_0 if the coloured trees $T_{i_0-1}(\kappa_1)$ and $T_{i_0-1}(\kappa_2)$ are identical, and $T_{i_0}(\kappa_1)$ is a proper coloured subtree of $T_{i_0}(\kappa_2)$.⁸ This means in particular that $y_i(\kappa_1) = y_i(\kappa_2)$ for all $i \leq i_0$.

We say that κ_1 is *more advanced than* κ_2 if there is a level i_0 such that κ_1 is more advanced than κ_2 at level i_0 . Note that there can be at most one such level because of the requirement of a strict inclusion at level i_0 . It is clear that the notion of being more advanced is a strict partial ordering on K .

⁷ Once again we emphasise that $\kappa_1(y_i)$ means $\kappa_1(y_i(\kappa_1))$.

⁸ Note that this expresses the idea that the construction of the tree for κ_1 as in the previous subsection is closer to termination than that for κ_2 ; hence the terminology “more advanced”.

2.4 A Trichotomy

Now, for a general colouring $\kappa \in K$ there are three possibilities. It may be advanceable at step 1, it may be advanceable at some step $i > 1$ or it may be non-advanceable. We will prove the following claim.

Claim 2

- (a) If κ is advanceable at step 1 then we can extend κ to an m -bounded colouring of \tilde{J} .
- (b) If κ is advanceable at some step $i > 1$ then there is a colouring $\tilde{\kappa} \in K$ which is more advanced than κ .
- (c) If κ is non-advanceable then $|T(\kappa) \cap J| \geq C_d^{-d} m^d$.

We will establish Claim 2 below, but for now we note that Claim 1 follows immediately from it. Indeed, the hypothesis of Claim 1 is that $m > C_d |J|^{1/d}$, so

$$|T(\kappa) \cap J| \leq |J| < C_d^{-d} m^d,$$

meaning that under the hypothesis of Claim 1 the third alternative cannot hold for any $\kappa \in K$. So every $\kappa \in K$ must be advanceable at some step. If κ is advanceable at step 1, part (a) of Claim 2 gives us what we want; if not, κ will be advanceable at some step $i > 1$ and there will be a $\tilde{\kappa} \in K$ which is more advanced than κ . Once again, the third alternative cannot hold for $\tilde{\kappa}$, if the first alternative holds we are happy, and if the second alternative holds we obtain a $\tilde{\tilde{\kappa}}$ which is more advanced than $\tilde{\kappa}$. We iterate this process. Since K is finite, a maximally advanced element of K must exist, meaning that at some point of the iteration the second alternative cannot hold, leaving us with only the first. In summary, if $m > C_d |J|^{1/d}$, for every $\kappa \in K$ there is some $\tilde{\kappa} \in K$ which is more advanced than κ and which is advanceable at step 1. Hence there exists an m -bounded colouring of \tilde{J} as required.

This procedure gives an algorithm for actually constructing an m -bounded colouring, see Algorithm 1.

Algorithm 1 Construct an m -bounded colouring of a set J of multijoints

- (0) We require $m > C_d |\tilde{J}|^{1/d}$.
 - (1) Let J_c be the empty set and κ be a colouring of J_c .
 - (2) For each point x_0 of J do the following:
 - (a) While κ is advanceable at some step $i > 1$ w.r.t. $\tilde{J} = J_c \cup \{x_0\}$ do the following:
 - (i) Let $\tilde{\kappa}$ be a colouring of J_c which is more advanced than κ , constructed as in the proof of Claim 2(b).
 - (ii) Update κ to be $\tilde{\kappa}$.
 - (b) Now κ is advanceable at step 1.
 - (c) Extend κ to \tilde{J} by letting $\kappa(x_0)$ be some colour j for which $L_j^{(1)}$ is empty.
 - (d) Update J_c to be \tilde{J} .
 - (3) Now κ is an m -bounded colouring of J .
-

Note that we have already established case (a) of Claim 2 in the discussion of case (1) in Sect. 2.2. In the remainder of the proof we will verify the remaining two cases of Claim 2.

2.5 Establishing Claim 2(b)

For the second case of Claim 2, let us assume that κ is advanceable at step $i_0 > 1$. That means that there is a colour $j_0 \neq \kappa(y_{i_0})$ such that for all lines $l_{j_0} \in L_{j_0}$ such that $y_{i_0} \in l_{j_0}$ we have

$$|\{x \in l_{j_0} \cap J : \kappa(x) = j_0\}| + |\{x \in l_{j_0} \cap T_{i_0-1} \cap J : \kappa(x) \neq j_0\}| < m. \tag{9}$$

We let $j_1 = \kappa(y_{i_0})$. Let us define a new colouring $\tilde{\kappa}$ which is identical to κ except that $\tilde{\kappa}(y_{i_0}) = j_0$. What we shall show is that there is some $i_1 < i_0$ such that $\tilde{\kappa}$ is more advanced than κ at level i_1 .

We first need to consider the effect of changing the colour of y_{i_0} on the construction of the tree $T(\tilde{\kappa})$, and in particular, we need to bear in mind the possibility that y_{i_0} might occur earlier in the construction of $T(\tilde{\kappa})$ than of $T(\kappa)$. As a rough guide, note that lines of colours other than the old and new colours of y_{i_0} will play exactly the same role in both constructions, as will lines not containing y_{i_0} . We will need to examine vertices of the tree $T(\kappa)$ of colour either the old or new colour of y_{i_0} for possible changes in the construction.

Specifically, we wish to verify that $\tilde{\kappa}$ is m -bounded and that $\tilde{\kappa}$ is more advanced than κ . The m -bounded conditions for $\tilde{\kappa}$ follow immediately from the corresponding conditions for κ except for lines of colour j_0 which go through y_{i_0} . But for those lines we just saw that

$$|\{x \in l_{j_0} \cap J : \kappa(x) = j_0\}| < m.$$

and so

$$|\{x \in l_{j_0} \cap J : \tilde{\kappa}(x) = j_0\}| = |\{x \in l_{j_0} \cap J : \kappa(x) = j_0\}| + 1 \leq m.$$

Thus $\tilde{\kappa}$ belongs to K and it is meaningful to ask whether $\tilde{\kappa}$ is more advanced than κ . Since $y_{i_0} \in T(\kappa)$ we can find an index $i_1 < i_0$ such that y_{i_0} is a vertex of $T_{i_1}(\kappa)$ but not of $T_{i_1-1}(\kappa)$. That means that either $i_1 = 1$ or there exists a colour $j_2 \neq j_1$ such that $\kappa(y_{i_1}) = j_2$ and a line $l_{j_1} \in L_{j_1}$ such that $y_{i_0}, y_{i_1} \in l_{j_1}$ and

$$|\{x \in l_{j_1} \cap J : \kappa(x) = j_1\}| + |\{x \in l_{j_1} \cap T_{i_1-1} \cap J : \kappa(x) \neq j_1\}| \geq m. \tag{10}$$

We want to show that $\tilde{\kappa}$ is more advanced than κ at level i_1 . First let us verify condition (i). If $i_0 = 1$ then condition (i) is vacuous. Otherwise it is clear that there are two types of steps we have to consider, and for other steps before i_1 condition (i) is immediate. The types of steps we have to consider correspond to vertices of the tree $T(\kappa)$ of colour equal to the new colour of y_{i_0} and of colour equal to the old colour of y_{i_0} , and more precisely are:

- steps i_2 such that $i_2 < i_1$ and $\kappa(y_{i_2}) \neq j_0$ but there is a line $l_{j_0} \in L_{j_0}$ such that $y_{i_2}, y_{i_0} \in l_{j_0}$; and
- steps i_3 such that $i_3 < i_1$ and $\kappa(y_{i_3}) \neq j_1$ but there is a line $\tilde{l}_{j_1} \in L_{j_1}$ such that $y_{i_3}, y_{i_0} \in \tilde{l}_{j_1}$.

For the former case we note that $y_{i_0} \notin T_{i_2}(\kappa)$ since this vertex set is a subset of $T_{i_1-1}(\kappa)$ which by assumption y_{i_0} does not belong to. Therefore we see that as $i_0 > i_2$

$$|\{x \in l_{j_0} \cap T_{i_0-1} \cap J : \kappa(x) \neq j_0\}| \geq |\{x \in l_{j_0} \cap T_{i_2-1} \cap J : \kappa(x) \neq j_0\}| + 1$$

since y_{i_0} is a member of the former set but not the latter. So (9) shows that

$$|\{x \in l_{j_0} \cap J : \kappa(x) = j_0\}| + |\{x \in l_{j_0} \cap T_{i_2-1} \cap J : \kappa(x) \neq j_0\}| < m - 1.$$

Now note that

$$|\{x \in l_{j_0} \cap J : \tilde{\kappa}(x) = j_0\}| = |\{x \in l_{j_0} \cap J : \kappa(x) = j_0\}| + 1$$

and

$$|\{x \in l_{j_0} \cap T_{i_2-1} \cap J : \tilde{\kappa}(x) \neq j_0\}| = |\{x \in l_{j_0} \cap T_{i_2-1} \cap J : \kappa(x) \neq j_0\}|,$$

where the second equality follows since $\tilde{\kappa}$ and κ are identical on T_{i_2-1} . So we obtain

$$|\{x \in l_{j_0} \cap J : \tilde{\kappa}(x) = j_0\}| + |\{x \in l_{j_0} \cap T_{i_2-1} \cap J : \tilde{\kappa}(x) \neq j_0\}| < m$$

and this shows that $l_{j_0} \notin L_{j_0}^{(i_2)}(\tilde{\kappa})$. Moreover, $l_{j_0} \notin L_{j_0}^{(i_2)}(\kappa)$ since if it were in this set then we would have $y_{i_0} \in T_{i_2}(\kappa)$ which is not possible as we saw above. So y_{i_0} is not added to the vertices of $T(\tilde{\kappa})$ at step i_2 , and we deduce that the iteration in the definition of the trees proceeds identically at this step for κ and $\tilde{\kappa}$. That is, the coloured trees $T_{i_2}(\tilde{\kappa})$ and $T_{i_2}(\kappa)$ are identical.

For the latter case we note that $y_{i_0} \notin T_{i_3}(\kappa)$ since this vertex set is a subset of $T_{i_1-1}(\kappa)$ which by assumption y_{i_0} does not belong to. That means that

$$|\{x \in \tilde{l}_{j_1} \cap J : \kappa(x) = j_1\}| + |\{x \in \tilde{l}_{j_1} \cap T_{i_3-1} \cap J : \kappa(x) \neq j_1\}| < m.$$

Now note that

$$|\{x \in \tilde{l}_{j_1} \cap J : \tilde{\kappa}(x) = j_1\}| = |\{x \in \tilde{l}_{j_1} \cap J : \kappa(x) = j_1\}| - 1$$

and

$$|\{x \in \tilde{l}_{j_1} \cap T_{i_3-1} \cap J : \tilde{\kappa}(x) \neq j_1\}| = |\{x \in \tilde{l}_{j_1} \cap T_{i_3-1} \cap J : \kappa(x) \neq j_1\}|,$$

where the second equality follows since $\tilde{\kappa}$ and κ are identical on T_{i_3-1} . Therefore

$$|\{x \in \tilde{l}_{j_1} \cap J : \tilde{\kappa}(x) = j_1\}| + |\{x \in \tilde{l}_{j_1} \cap T_{i_3-1} \cap J : \tilde{\kappa}(x) \neq j_1\}| < m - 1,$$

and this shows that $\tilde{l}_{j_1} \notin L_{j_1}^{(i_3)}(\tilde{\kappa})$. Moreover, $\tilde{l}_{j_1} \notin L_{j_1}^{(i_3)}(\kappa)$ since if it were in this set then we would have $y_{i_1} \in T_{i_3}(\kappa)$ which is not possible as we saw above. So y_{i_0} is not added to the vertices of $T(\tilde{\kappa})$ at step i_3 , and we deduce that the iteration in the definition of the trees proceeds identically at this step for κ and $\tilde{\kappa}$. That is, the coloured trees $T_{i_3}(\tilde{\kappa})$ and $T_{i_3}(\kappa)$ are identical.

Hence we conclude that the coloured trees $T_{i_1-1}(\tilde{\kappa})$ and $T_{i_1-1}(\kappa)$ are identical.

Now we verify conditions (ii) and (iii). We note that the only possible difference between the sets $L_j^{(i_1)}(\kappa)$ and $L_j^{(i_1)}(\tilde{\kappa})$ for some colour j is that a line containing both y_{i_0} and y_{i_1} could be in one of these sets and not the other. We already know that the line joining these points is of colour j_1 so for other colours we have that $L_j^{(i_1)}(\kappa)$ and $L_j^{(i_1)}(\tilde{\kappa})$ are identical and so $I_j^{(i_1)}(\kappa)$ and $I_j^{(i_1)}(\tilde{\kappa})$ are identical too. For colour j_1 we have that $\kappa(y_{i_0}) = j_1 \neq \tilde{\kappa}(y_{i_0})$. This shows that y_{i_0} is an element of $I_{j_1}^{(i_1)}(\kappa)$ but not of $I_{j_1}^{(i_1)}(\tilde{\kappa})$. Hence we conclude that the coloured tree $T_{i_1}(\tilde{\kappa})$ is a proper coloured subtree of $T_{i_1}(\kappa)$ and conditions (ii) and (iii) are verified.

This establishes the second case of Claim 2.

2.6 Establishing Claim 2(c)

For the last case of Claim 2, let us recall the statement of Quilodrán’s lemma, which in [10] is proved for the case $\mathbb{F} = \mathbb{R}$. For the general case, see [2,5] or [11]. See also [9].

Lemma 1 *Let L be a collection of lines in \mathbb{F}^d and let J be a subset of the set of joints of L . Suppose that J has the property that for every line $l \in L$ the cardinality of $l \cap J$ is either 0 or at least m . Then $|J| \geq C_d m^d$.*

Let us then assume that κ is non-advanceable. We let \bar{J} be the set of points in J which are vertices of the tree $T(\kappa)$. For a colour j we let \bar{L}_j be the union of the sets $L_j^{(i)}$ for those indices i such that $\kappa(y_i) \neq j$. Then we let $\bar{L} = \bar{L}_1 \cup \dots \cup \bar{L}_d$.

We need to verify that the hypotheses of the lemma are satisfied for \bar{J} and \bar{L} . First we note that the elements of \bar{J} are in fact joints of \bar{L} . To see this, take $y_i \in \bar{J}$ with $i > 1$ and assume that $\kappa(y_i) = j$. Then there is an $\tilde{i} < i$ and a line $l_j \in L_j^{(\tilde{i})}$ such that $y_i \in l_j$ and $\kappa(y_{\tilde{i}}) \neq j$ (since children always have a different colour from their parents). So $y_i \in \bar{L}_j$. Furthermore, for all colours $\tilde{j} \neq j$ we have by non-advanceability that the set $L_{\tilde{j}}^{(i)}$ is non-empty and the lines in these sets all go through y_i . So $y_i \in \bar{L}_{\tilde{j}}$. Thus for each $j_* \in \{1, \dots, d\}$, we have $y_i \in \bar{L}_{j_*} \subseteq L_{j_*}$. Since by hypothesis the collection \mathcal{L} is generic, we conclude that y_i is a joint of \bar{L} .

Now consider a line in \bar{L} , say $l_j \in L_j^{(i)}$. Then by definition of $L_j^{(i)}$ we have

$$|\{x \in l_j \cap J : \kappa(x) = j\}| + |\{x \in l_j \cap T_{i-1} \cap J : \kappa(x) \neq j\}| \geq m.$$

Note that all the points which are elements of the first of these sets will be vertices of T_i . Therefore the two sets occurring in this expression are subsets of $l_j \cap \bar{J}$ which

are disjoint. Hence we have $|l_j \cap \bar{J}| \geq m$. This shows that all the hypotheses of the lemma are satisfied and so we deduce that $|\bar{J}| \geq C_d^{-d} m^d$.

Remark 5 The reader will observe that we use the hypothesis of genericity only in establishing Claim 2(c). We conjecture that the main result holds without this hypothesis.

Remark 6 We do not claim any particular efficiency for Algorithm 1. We simply wish to point out that our methods are constructive rather than existential.

Remark 7 It is natural to ask whether the method we have presented might be pushed in order to obtain the full multijoints bound (2). As noted above, the main difficulty is to replace the arithmetic mean bound $\sum_j S_j(x) \geq C_d$ for all $x \in J$ by the corresponding geometric mean bound $\prod_j S_j(x)^{1/d} \geq C_d$. In fact it would suffice to obtain weighted arithmetic bounds $\sum_j \beta_j S_j(x) \geq C_d$ for all $x \in J$ for suitably chosen parameters β_j . Our methods, in conjunction with a suitable multiparameter version of Quilodrán's lemma, would seem to be applicable to this problem. We plan to return to this matter at a later date.

References

1. Bourgain, J., Guth, L.: Bounds on oscillatory integral operators based on multilinear estimates. *Geom. Funct. Anal.* **21**(6), 1239–1295 (2011)
2. Carbery, A., Iliopoulou, M.: Counting joints in vector spaces over arbitrary fields (2014). [arXiv:1403.6438](https://arxiv.org/abs/1403.6438)
3. Carbery, A., Valdimarsson, S.I.: The endpoint multilinear Kakeya theorem via the Borsuk–Ulam theorem. *J. Funct. Anal.* **264**(7), 1643–1663 (2013)
4. Dvir, Z.: On the size of Kakeya sets in finite fields. *J. Am. Math. Soc.* **22**(4), 1093–1097 (2009)
5. Dvir, Z.: Incidence theorems and their applications (2012). [arXiv:1208.5073](https://arxiv.org/abs/1208.5073)
6. Guth, L.: The endpoint case of the Bennett–Carbery–Tao multilinear Kakeya conjecture. *Acta Math.* **205**(2), 263–286 (2010)
7. Guth, L., Katz, N.H.: Algebraic methods in discrete analogs of the Kakeya problem. *Adv. Math.* **225**(5), 2828–2839 (2010)
8. Iliopoulou, M.: Personal Communication. IPAM (2014)
9. Kaplan, H., Sharir, M., Shustin, E.: On lines and joints. *Discrete Comput. Geom.* **44**(4), 838–843 (2010)
10. Quilodrán, R.: The joints problem in \mathbb{R}^n . *SIAM J. Discrete Math.* **23**(4), 2211–2213 (2009/2010)
11. Tao, T.: Algebraic combinatorial geometry: the polynomial method in arithmetic combinatorics, incidence combinatorics, and number theory. *EMS Surv. Math. Sci.* **1**(1), 1–46 (2014)