arXiv:2007.00685v1 [math.CO] 1 Jul 2020

Coloring linear hypergraphs: the Erdős–Faber–Lovász conjecture and the Combinatorial Nullstellensatz

Oliver Janzer*   Zoltán Lóránt Nagy†

Abstract

The long-standing Erdős–Faber–Lovász conjecture states that every n-uniform linear hypergaph with n edges has a proper vertex-coloring using n colors. In this paper we propose an algebraic framework to the problem and formulate a corresponding stronger conjecture. Using the Combinatorial Nullstellensatz, we reduce the Erdős–Faber–Lovász conjecture to the existence of non-zero coefficients in certain polynomials. These coefficients are in turn related to the number of orientations with prescribed in-degree sequences of some auxiliary graphs. We prove the existence of certain orientations, which verifies a necessary condition for our algebraic approach to work.

Keywords: coloring, hypergraphs, Erdős–Faber–Lovász, Combinatorial Nullstellensatz, graph orientations

1 Introduction

A hypergraph $\mathcal{H} = (V, \mathcal{E})$ consists of a nonempty vertex set V, and an edge set $\mathcal{E}$. A hypergraph is called *linear* if the intersection of each pair of edges contains at most one vertex. A *proper vertex coloring* with a color set C of the hypergraph is a function $c : V \to C$ such that each edge consists of vertices of different colors. A well known conjecture of Erdős, Faber and Lovász, dating back to 1972, asserts an upper bound on the minimum number of colors.

Conjecture 1.1 (Erdős–Faber–Lovász)**.** *If a linear hypergraph $\mathcal{H} = (V, \mathcal{E})$ has n edges, each of size at most n, then $\mathcal{H}$ can be colored properly by n colors.*

Note that the statement follows if one considers only linear hypergraphs such that every vertex is incident to at least 2 hyperedges, according to the observation below.

Observation 1.2. *Consider a hypergraph $\mathcal{H} = (V, \mathcal{E})$ in which every edge has size at most n. If one deletes the vertices of degree 1, any proper coloring of the obtained derived hypergraph with at least n colors can be extended to a proper coloring of $\mathcal{H} = (V, \mathcal{E})$ with the same color set.*

*Department of Pure Mathematics and Mathematical Statistics, University of Cambridge, United Kingdom. E-mail: oj224@cam.ac.uk.

†MTA–ELTE Geometric and Algebraic Combinatorics Research Group, Eötvös Loránd University, Budapest, Hungary. The author is supported by the Hungarian Research Grant (NKFI) No. K 120154 and SNN 132625 and by the János Bolyai Scholarship of the Hungarian Academy of Sciences. E-mail: nagyzoli@cs.elte.hu

Erdős himself considered this one of his three favourite combinatorial problems, and offered one of his highest prizes ever for a proof or disproof [5]. By dualizing the problem, its connection to Vizing's theorem becomes clear. To this end, one may assign vertices to the edges of $\mathcal{H}$ and introduce the dual hypergraph $\hat{\mathcal{H}}$ with hyperedges H_v assigned to each vertex $v \in V(\mathcal{H})$ such that H_v consists of the vertices corresponding to hyperedges incident to v in $\mathcal{H}$. If $\mathcal{H}$ is linear with n edges, then the resulting hypergraph $\hat{\mathcal{H}}$ on n vertices is linear as well. Thus this way we get another variant of the conjecture.

Conjecture 1.3 (Erdős–Faber–Lovász, 2nd (dual) variant)**.** *Any linear hypergraph on n vertices has chromatic index at most n.*

The conjecture is confirmed for certain hypergraph families, but the problem is still widely open, even though asymptotic and fractional versions were established by Kahn and Seymour [10, 11, 25]. Some notable hypergraph families for which the conjecture is confirmed are the *dense* derived hypergraphs for which the minimum degree $\delta(\mathcal{H})$ is greater than $\sqrt{n}$ [23], the uniform derived hypergraphs [6] or the cases $n \leq 12$ [8, 22] and some other families [3, 7, 9, 17, 20]. These results mostly apply algorithmic and graph theoretic arguments some with computer-based search. For further results on the topic, we refer to [21].

In this paper we propose an algebraic approach, in connection with the celebrated Combinatorial Nullstellensatz of Alon [1]. We point out that the existence of a suitable proper coloring of a hypergraph $\mathcal{H}$ is strongly connected to the existence of a particular degree-bounded orientation of certain auxiliary graphs obtained from $\mathcal{H}$. In Section 2 we introduce the algebraic tool and present two types of auxiliary graphs assigned to the linear hypergraphs. The application of the algebraic tool will imply that if the total sum of certain signed bounded-degree orientations of the auxiliary graph is nonzero, then there exists a proper coloring with at most n colors. We formulate a conjecture that in fact, this related stronger variant of the Erdős–Faber–Lovász conjecture also holds. In Section 3 we study the strengthened variant of Conjecture 1.1 and confirm it in a weak sense by showing that a special, so-called Vandermonde-type, orientation exists for both families of auxiliary graphs assigned to arbitrary n-uniform hypergraphs $\mathcal{H}$ with n edges. This verifies a necessary condition for our algebraic approach to work. Finally, in Section 4 we give some concluding remarks.

2 The algebraic tool and the strengthening of the E–F–L conjecture

Our starting point is Alon's celebrated Combinatorial Nullstellensatz [1], more precisely the Non-vanishing lemma, described below. This tool turned out to be very powerful in several areas of combinatorics; in particular, in graph coloring problems [1, 14, 18, 26]. The connection of graph orientations and this lemma appeared first in the influential paper of Alon and Tarsi [2]. Here we recall the form of the Combinatorial Nullstellensatz that we will apply.

Theorem 2.1 (Combinatorial Nullstellensatz, Non-vanishing lemma [1])**.** *Let $\mathbb{F}$ be an arbitrary field and let $P = P(x_1, \ldots, x_k)$ be a polynomial of k variables over $\mathbb{F}$. Suppose that there exists a monomial $\prod_{i=1}^{k} x_i^{d_i}$, such that the sum $\sum_{i=1}^{k} d_i$ equals the total degree of P, and the coefficient of $\prod_{i=1}^{k} x_i^{d_i}$ in P is nonzero. Then for any set of subsets $A_1, \ldots, A_k$ of $\mathbb{F}$ such that $|A_i| > d_i$, there exists a k-tuple $(s_1, s_2, \ldots, s_k) \in \bigtimes A_i$ for which $P(s_1, s_2, \ldots, s_k) \neq 0$.*

In most applications of the Combinatorial Nullstellensatz, the polynomial can be directly derived from the combinatorial setting, and the choice of the maximal monomial with which Theorem 2.1 is applied is also natural. The main step to make the argument work is to check that the coefficient of this monomial is not zero. In fact, one usually knows or conjectures in advance the extremal structure, which can be helpful in setting up the corresponding polynomials and verifying that the coefficient in view is nonzero. Unlike in those cases, here we have large freedom to consider a suitable polynomial, and we have to pick the polynomial and its maximal monomial carefully so that the coefficient is surely nonzero. This provides a rather novel application of the main tool.

Let us continue by setting the main notations. For a graph or hypergraph $\mathcal{H}$, $d(v)$ denotes the degree of the vertex v. A monomial $\prod_j y_j^{\alpha_j}$ of a multivariate polynomial $Q(\mathbf{y})$ is a *t-bounded degree monomial* if the degree of each variable y_j is bounded from above by t, i.e. $\alpha_j \leq t$. The total degree of a polynomial Q is denoted by $\deg(Q)$.

Our aim is to set up a multivariate polynomial where the variables correspond to vertices of $\mathcal{H}$ and the values taken by the variables correspond to colors. The polynomial encodes the coloring constraints of the hypergraph $\mathcal{H}$. In order to do this, we assign an auxiliary graph $G(\mathcal{H})$ first to the linear hypergraph $\mathcal{H}$. We note that we shall propose two different kinds of polynomial that can be used to encode the colouring constraints.

2.1 Setting up polynomials corresponding to proper colorings

From now on, speaking about a linear hypergraph $\mathcal{H}$ we always assume that it has n hyperedges $\mathcal{E} = \{F_1, F_2, \ldots, F_n\}$ of size n, unless specified otherwise.

We start with introducing *two kinds of auxiliary graph, $G_1(\mathcal{H})$ and $G_2(\mathcal{H})$* assigned to the hypergraph $\mathcal{H}$. They correspond to two separate approaches to Conjecture 1.1. Both graphs consist of n vertex-disjoint cliques of size n, together with a set of so-called *identifier edges* joining vertices from different cliques. We remark in advance that the identifier edges are not uniquely determined by the hypergraph $\mathcal{H}$; we have some freedom how to choose them.

Definition 2.2. The *auxiliary graph of first kind $G_1(\mathcal{H})$* assigned to $\mathcal{H}$ is defined as follows. We take n copies $K_n^{(i)}$ $(i = 1, \ldots, n)$ of the complete graph K_n, where the vertices of $K_n^{(i)}$ are labelled by the vertices of F_i, see Figure 1. We call them the *base cliques*. Here each vertex $v \in V(\mathcal{H})$ appears $d(v)$ times, and for each $v \in V(\mathcal{H})$ we choose an arbitrary spanning tree on the set of those $d(v)$ vertices in $G_1(\mathcal{H})$ which are labelled by v. We call these *identifier spanning trees*. For each edge of these spanning trees, we define an edge of multiplicity $n-1$ in $G_1(\mathcal{H})$, see Figure 2. We call these new edges *identifier edges*. The edge set of $G_1(\mathcal{H})$ consists of the edges of the base cliques and the identifier edges.

Definition 2.3. The *auxiliary graph of second kind $G_2(\mathcal{H})$* assigned to $\mathcal{H}$ is defined as follows. We take n copies $K_n^{(i)}$ $(i = 1, \ldots, n)$ of the complete graph K_n, where the vertices of $K_n^{(i)}$ are labelled by the vertices of F_i, see Figure 1. We call them the base cliques. Here each vertex $v \in V(\mathcal{H})$ appears $d(v)$ times, and for each $v \in V(\mathcal{H})$ we choose an arbitrary identifier spanning tree on the set of those $d(v)$ vertices in $G_2(\mathcal{H})$ which are labelled by v. For each edge $v_{i,j}v_{k,l}$ of the spanning tree, we either take the set $\{v_{i,j}v_{k,t} : t \neq l\}$ or the set $\{v_{i,t}v_{k,l} : t \neq j\}$ to be edges of $G_2(\mathcal{H})$. We call these new edges *identifier edges* and they have multiplicity 1, see Figure 2. The edge set of the auxiliary graph consists of the edges of the base cliques and the identifier edges.

Remark 2.4. In what follows, when $v_{i,j}v_{k,l}$ is an edge in an identifier spanning tree and $i < k$, then we shall always take the set $\{v_{i,j}v_{k,t} : t \neq l\}$ to be the corresponding edges of $G_2(\mathcal{H})$.

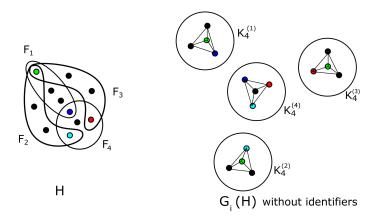
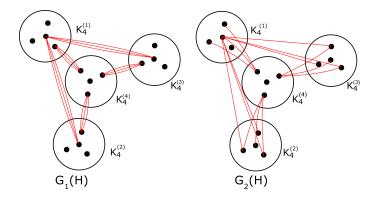Figure 1: $\mathcal{H}$ and the base cliques of $G_i(\mathcal{H})$, $(i \in \{1, 2\})$

Figure 2: Identifier edges in auxiliary graphs $G_1(\mathcal{H})$ and $G_2(\mathcal{H})$

Remark 2.5. Informally, the identifier edges are defined as follows. For a pair of intersecting hyperedges in $\mathcal{H}$, their common vertex v has a copy corresponding to each of the two edges. These two copies are either joined in the corresponding identifier spanning tree or not. If they are, then in $G_1(\mathcal{H})$ we put an identifier edge with multiplicity $n - 1$ between them, while in $G_2(\mathcal{H})$ we put a set of identifier edges forming a star with $n - 1$ leaves whose centre is one of the two copies of v and whose leaves are all vertices in the other base clique, apart from the copy of v.

A suitable coloring for $\mathcal{H}$ with n distinct elements of a field $\mathbb{F}$ possesses the following properties:

(P1) Every clique F_i contains all of the colors (once).

(P2) Vertices from different cliques which correspond to the same vertex in $\mathcal{H}$ are assigned the same color.

4

In terms of the auxiliary graphs, a suitable vertex-colouring of $\mathcal{H}$ corresponds to a vertex-coloring of $G_i(\mathcal{H})$ ($i \in \{1, 2\}$) in which any two vertices in the same base clique have different colors, while any two vertices joined by an edge in an identifier spanning tree have the same color.

In order to determine whether a given coloring is suitable or not, we assign a variable $x_{i,j}$ to each vertex $v_{i,j}$ ($j = 1, \ldots, n$) of the base cliques $K_n^{(i)}$. These variables will take one of n possible values from $\mathbb{F}$ corresponding to the n colors we can use on the vertices. Moreover, we define the following three families of polynomials.

For any $1 \le i \le n$, let

$$Q_i(\mathbf{x}) = \prod_{1 \le j < j' \le n} (x_{i,j} - x_{i,j'}).$$

For any $1 \le i < k \le n$, if there exist some $v_{i,j} \in K_n^{(i)}$ and $v_{k,l} \in K_n^{(k)}$ which form an edge in an identifier spanning tree, then let

$$R_{i,k}(\mathbf{x}) = (x_{i,j} - x_{k,l})^{n-1} - 1.$$

(Note that, as $\mathcal{H}$ is linear, there is at most one such pair of vertices.) Otherwise, let $R_{i,k}(\mathbf{x}) = 1$.

Similarly, if there exist some $v_{i,j} \in K_n^{(i)}$ and $v_{k,l} \in K_n^{(k)}$ which form an edge in an identifier spanning tree, then let

$$\Phi_{i,k}(\mathbf{x}) = \frac{\prod_m (x_{i,j} - x_{k,m})}{x_{i,j} - x_{k,l}},$$

and otherwise let $\Phi_{i,k}(\mathbf{x}) = 1$.

Define $P_1(\mathbf{x})$ and $P_2(\mathbf{x})$ as follows:

$$P_1(\mathbf{x}) = \prod_{i=1}^{n} Q_i(\mathbf{x}) \cdot \prod_{1 \le i < k \le n} R_{i,k}(\mathbf{x}),$$

$$P_2(\mathbf{x}) = \prod_{i=1}^{n} Q_i(\mathbf{x}) \cdot \prod_{1 \le i < k \le n} \Phi_{i,k}(\mathbf{x}).$$

Observe that $P_j(\mathbf{x})$ is not uniquely determined yet (it depends on the choice of the identifier spanning trees), but its total degree can be expressed by the degree profile of the hypergraph as

$$\deg P_j(\mathbf{x}) = n\binom{n}{2} + \sum_{v \in V(\mathcal{H})} (d(v) - 1)(n - 1).$$

In order to apply the Non-vanishing lemma (Theorem 2.1), we have to set the field $\mathbb{F}$. When $n = p$ is a prime, let $P_1(\mathbf{x})$ be viewed as a polynomial over $\mathbb{F} = \mathbb{F}_p$. For arbitrary n, let $P_2(\mathbf{x})$ be viewed as a polynomial over $\mathbb{F} = \mathbb{R}$.

Now we can formulate our first contribution which provides an algebraic framework to the main problem.

Conjecture 2.6. *Let $\mathcal{H}$ be an n-uniform linear hypergraph with n edges.*

(a) *One can choose the identifier spanning trees in a way that $P_2(\boldsymbol{x})$ has an $(n-1)$-bounded degree maximal monomial with nonzero coefficient.*

(b) *When n is a prime, one can choose the identifier spanning trees in a way that $P_1(\boldsymbol{x})$ has an $(n-1)$-bounded degree maximal monomial with nonzero coefficient.*

Observe that this conjecture would imply the Conjecture 1.1 of Erdős, Faber and Lovász. Indeed, if one evaluates the polynomial P_1 or P_2 on the Cartesian product $\{0, 1, \ldots, n-1\}^n$, it will vanish except when the values of the variables correspond to a proper coloring, although one has to suppose that n is a prime in the case of P_1. To see this, note that $Q_i(\mathbf{x}) = 0$ holds if and only if there exist two vertices in some hyperedge F_i with the same color. Moreover, when n is a prime, Fermat's little theorem implies that $R_{i,k}(\mathbf{x}) = 0$ for some i, k if and only if not all identified vertices received the same color. Finally, if $Q_i(\mathbf{x}) \neq 0$ for every i, then we have $\Phi_{i,k}(\mathbf{x}) = 0$ for some i, k if and only if not all identified vertices were colored with the same color. Hence, $P_j(\mathbf{x}) \neq 0$ ($j \in \{1, 2\}$) holds if and only if both properties (P1) and (P2) are satisfied.

2.2 Searching for a nonzero coefficient and the connection to orientations

Alon and Tarsi made a connection between a certain coloring problem and the number of Eulerian orientations via the Non-vanishing Lemma [2]. In our case, the situation is somewhat similar.

Define the *sign* of an orientation of the graph $G_a(\mathcal{H})$ ($a \in \{1, 2\}$) to be $(-1)^t$, where t is the number of edges which point from $x_{i,j}$ to $x_{i,j'}$ with $j < j'$ or from $x_{i,j}$ to $x_{k,l}$ with $i < k$. Now, using the correspondence between $G_a(\mathcal{H})$ and $P_a(\mathbf{x})$, it is not hard to see that the coefficient of any maximum-degree monomial $\prod_{i,j} x_{i,j}^{\alpha_{i,j}}$ in $P_a(\mathbf{x})$ is the sum of the signs of those orientations of $G_a(\mathcal{H})$ in which the in-degree of every vertex $v_{i,j}$ is precisely $\alpha_{i,j}$.

To make the Non-vanishing lemma applicable, we clearly need that the exponent of each variable is less than n. We will also rely on the following basic fact.

Fact 2.7. $(-1)^{\binom{n}{2}} \cdot Q_i(\boldsymbol{x}) = \prod_{j<j'}(x_{i,j'} - x_{i,j})$ *equals the Vandermonde determinant*

$$
\begin{bmatrix}
1 & x_{i,1} & x_{i,1}^2 & \cdots & x_{i,1}^{n-1} \\
1 & x_{i,2} & x_{i,2}^2 & \cdots & x_{i,2}^{n-1} \\
1 & x_{i,3} & x_{i,3}^2 & \cdots & x_{i,3}^{n-1} \\
\vdots & \vdots & \vdots & \ddots & \vdots \\
1 & x_{i,n} & x_{i,n}^2 & \cdots & x_{i,n}^{n-1}
\end{bmatrix}.
$$

Using Fact 2.7, the product $\prod_i Q_i(\mathbf{x}) \cdot \prod_{i,j} x_{i,j}^{\beta_{i,j}}$ is a linear combination of monomials of the form $\prod_{i,j} x_{i,j}^{\beta_{i,j}+\sigma_i(j)}$, where $\{\sigma_i(j) : 1 \leq j \leq n\} = \{0, 1, \ldots, n-1\}$ for every i. Thus, to determine the coefficients of the $(n-1)$-degree bounded monomials with maximal degree in $P_1(\mathbf{x})$, it suffices

to consider those monomials $\prod_{i,j} x_{i,j}^{\beta_{i,j}}$ in $\prod_{i<k} R_{i,k}(\mathbf{x})$ which have maximal degree and for which there exist functions σ_i as above such that $\beta_{i,j} + \sigma_i(j) \leq n-1$ for every i,j. Call such monomials *Vandermonde-completable*. Similarly, to compute the coefficients of the $(n-1)$-bounded degree polynomials in $P_2(\mathbf{x})$ with maximal degree, we only need to consider Vandermonde-completable monomials in $\prod_{i<k} \Phi_{i,k}(\mathbf{x})$.

We call an orientation of the identifier edge set (in $G_1(\mathcal{H})$ or $G_2(\mathcal{H})$) *Vandermonde-completable* if one can orient the edges in the base cliques such that each clique spans a transitive tournament and every in-degree in the whole graph is bounded by $n-1$ from above. Under the correspondence between the orientations of the identifier edges in $G_1(\mathcal{H})$ and the monomials in $\prod_{i<k} R_{i,k}(\mathbf{x})$, Vandermonde-completable orientations correspond to Vandermonde-completable monomials and vice versa. The same holds in the case of $G_2(\mathcal{H})$ and $\prod_{i<k} \Phi_{i,k}(\mathbf{x})$.

2.3   Main results

Our main result states that Vandermonde-completable orientations exist.

Theorem 2.8. *Let $\mathcal{H}$ be an n-uniform linear hypergraph with n hyperedges. Then, for any choice of the identifier spanning trees, there is a Vandermonde-completable orientation of the identifier edges in $G_1(\mathcal{H})$.*

When n is a prime, the number of Vandermonde-completable orientations corresponding to a given Vandermonde-completable monomial in $\prod_{i<k} R_{i,k}(\mathbf{x})$ is not divisible by n. Moreover these orientations all have the same sign, so we obtain the following corollary.

Corollary 2.9. *Let n be a prime and let $\mathcal{H}$ be an n-uniform linear hypergraph with n hyperedges. Then, for any choice of the identifier spanning trees, there is a Vandermonde-completable monomial (with non-zero coefficient) in $\prod_{i<k} R_{i,k}(\boldsymbol{x})$.*

We also prove the analogue of Theorem 2.8 in the case of $G_2(\mathcal{H})$, although in this case we make a specific choice for the identifier spanning trees.

Theorem 2.10. *Let $\mathcal{H}$ be an n-uniform linear hypergraph with n hyperedges. Then one can choose the identifier spanning trees in a way that there is a Vandermonde-completable orientation of the identifier edges in $G_2(\mathcal{H})$.*

Unfortunately, we cannot prove an analogue of Corollary 2.9 because we cannot compute the sum of the signs of the orientations that yield the same monomial in $\prod_{i<k} \Phi_{i,n}(\mathbf{x})$. We leave it as a conjecture.

Conjecture 2.11. *Let $\mathcal{H}$ be an n-uniform linear hypergraph with n hyperedges. Then one can choose the identifier spanning trees in a way that there is a Vandermonde-completable monomial (with non-zero coefficient) in $\prod_{i<k} \Phi_{i,k}(\boldsymbol{x})$.*

3 The proofs of the main results

3.1 The case of auxiliary graph $G_1(\mathcal{H})$

Proof of Theorem 2.8. We will use the following claim.

Claim 3.1. *Let T be a tree with k vertices. Write T' for the multigraph on the same vertex set, obtained by taking $n-1$ copies of each edge of T. Suppose that for every $v \in V(T)$ there is an integer α_v such that $\alpha_v \leq n-1$ and $\sum_{v \in V(T)} \alpha_v = (k-1)(n-1)$. Then there exists an orientation of the edges of T' in which every v has in-degree α_v.*

Proof. We use induction on k. The statement is clear for $k = 0, 1$. Assume that $k > 1$. Let w be a leaf of T. Let uw be the unique edge of T containing w. Direct α_w of the $(n-1)$ edges of T' corresponding to uw towards w and direct the rest towards u. (Note that the conditions $\alpha_v \leq n-1$ and $\sum_{v \in V(T)} \alpha_v = (k-1)(n-1)$ ensure that $\alpha_w \geq 0$.) Let S be the tree obtained from T by deleting w. For every $v \in V(S) \setminus \{u\}$, let $\beta_v = \alpha_v$, while let $\beta_u = \alpha_u - (n-1-\alpha_w)$. Clearly, $\sum_{v \in V(S)} \beta_v = (k-2)(n-1)$. Hence, by the induction hypothesis, we can orient the edges of S' (which is the multigraph obtained by taking $n-1$ copies of each edge of S) in a way that the in-degree of every vertex $v \in V(S)$ is β_v. Together with the orientation of the edges corresponding to uw, we get a suitable orientation of T', completing the induction step. $\square$

For any $1 \leq a, b \leq n$ with $a \neq b$, let $s(a,b) = \begin{cases} a - b & \text{if } a > b \\ n + a - b & \text{if } a < b \end{cases}$

Let $v \in V(\mathcal{H})$. Let the edges of $\mathcal{H}$ which contain v be $F_{i_1}, \ldots, F_{i_k}$, where $i_1 < \cdots < i_k$. Then each base clique $K_n^{(i_j)}$ $(1 \leq j \leq k)$ contains a vertex labelled by v, call it w_j. Let T be the identifier spanning tree on the vertex set $\{w_1, \ldots, w_k\}$. Define i_0 to be i_k. Note that $s(i_{j-1}, i_j) \leq n-1$ for every $1 \leq j \leq k$ and $\sum_{1 \leq j \leq k} s(i_{j-1}, i_j) = (k-1)n \geq (k-1)(n-1)$. Thus, by Claim 3.1, we can orient the edges of $G_1(\mathcal{H})$ corresponding to the spanning tree T in a way that every vertex w_j gets in-degree at most $s(i_{j-1}, i_j)$. Hence, the in-degree at every w_j is at most $\max(s(i_q, i_j) : q \neq j)$.

Performing this for every $v \in V(\mathcal{H})$, we obtain an orientation of the identifier edges in $G_1(\mathcal{H})$. We claim that it is Vandermonde-completable. Indeed, for every $1 \leq i \leq n$, using the fact that $\{s(q,i) : q \neq i\} = \{1, \ldots, n-1\}$ and that $\mathcal{H}$ is linear, the number of vertices of $K_n^{(i)}$ with in-degree at least $n-j$ is at most j for every $1 \leq j \leq n$. $\square$

3.2 The case of auxiliary graph $G_2(\mathcal{H})$

Proof of Theorem 2.10. For any $x \in V(\mathcal{H})$, let the identifier spanning tree corresponding to x be the path $v_{i_1, j_1} v_{i_2, j_2} \ldots v_{i_t, j_t}$, where $i_1 < i_2 < \cdots < i_t$ (so the edges of $\mathcal{H}$ containing x are $F_{i_1}, \ldots, F_{i_t}$). We call the auxiliary graphs $G_2(\mathcal{H})$ obtained this way *path-like*. Note that if $v_{i,j}$ and $v_{k,l}$ $(i < k)$ are identified, then the edges between the ith and kth base clique are joining each vertex of the kth base clique to $v_{i,j}$ except for $v_{k,l}$.

We claim that for this choice, there is a Vandermonde-completable orientation of the identifier edges in $G_2(\mathcal{H})$. The proof is by induction on n. The case $n = 1$ is trivial. For the inductive step, suppose that we have proved the claim for n and let us consider a path-like auxiliary graph

$G = G_2(\mathcal{H})$ of a hypergraph $\mathcal{H}$ with $n+1$ edges $F_0, F_1, \ldots, F_n$, each of size $n+1$. We call a vertex $v_{i,j}$ a *source* if it forms an edge with $v_{k,l}$ for some $i < k$ in an identifier spanning path.
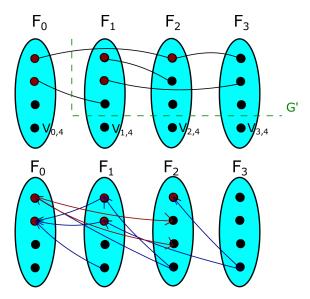

Figure 3: Spanning paths on identified vertices in $G_2(\mathcal{H})$ (top) and the orientation of the identifier edges adjacent to F_0 or S.

Note that each edge F_i contains a vertex which is not contained in any other edge F_j. That is, each base clique contains at least one vertex with no identification to other vertices. Without loss of generality, $v_{1,n+1} \in F_1, \ldots, v_{n,n+1} \in F_n$ are such vertices. Let $S = \{v_{1,n+1}, \ldots v_{n,n+1}\}$ and let $\mathcal{H}'$ be the n-uniform linear hypergraph whose edges are $F_1' = F_1 \setminus \{v_{1,n+1}\}, \ldots, F_n' = F_n \setminus \{v_{n,n+1}\}$. Note that the identifier edges of $G_2(\mathcal{H})$ induced by the set $F_1' \cup \cdots \cup F_n'$ are the identifier edges of $G' = G_2(\mathcal{H}')$. Hence, by the induction hypothesis, we may orient them in a Vandermonde-completable way. We extend this to a Vandermonde-completable orientation of the identifier edges of $G_2(\mathcal{H})$ as follows.

- Orient each identifier edge with one endpoint in S towards the other endpoint of the edge (which is necessarily the source). Clearly, the path-like property of $G_2(\mathcal{H})$ implies that the in-degree of any source vertex is increased by at most one, while the in-degree of any non-source vertex is unchanged.

- Consider all the identifier edges adjacent to the vertices of F_0. Each pair of base cliques F_0 and F_k spans either zero or n identifier edges, all of them incident to a single vertex $v_{0,i}$ of F_0. Orient these edges $v_{0,i} v_{k,j}$ towards $v_{0,i}$ if and only if $v_{k,j}$ is a source or $j = n + 1$. This does not change the in-degree of the source vertices in $F_1' \cup \cdots \cup F_n'$, and it increases the in-degree of any non-source vertex in $F_1' \cup \cdots \cup F_n'$ by at most one.

It is straightforward to see that the resulting orientation is Vandermonde-completable in the base cliques F_i, $i > 0$, since each in-degree is increased by at most one and we added a new vertex of in-degree zero to each F_i. Thus, we only have to confirm that F_0 is also Vandermonde-completable.

9

Since F_i has at most $n - i$ source vertices, any vertex in F_0 which is joined to an element of F_i in an identifier spanning tree has in-degree at most $n - i + 1$. Moreover, any non-source vertex in F_0 has in-degree 0. Hence, for each $i \geq 1$, the number of vertices in F_0 with in-degree at least $n - i + 1$ is at most i, as required. □

4   Concluding remarks

Theorems 2.8 and 2.10 suggest that Conjecture 2.6 (and, consequently, the Erdős–Faber–Lovász conjecture) is likely to hold.

For some linear hypergraph families, we can find $(n - 1)$-bounded degree monomials which correspond to an (almost) unique orientation of the auxiliary graph, hence their coefficient is non-zero. For example, this is the case when the degree of every vertex in $\mathcal{H}$ is either 1 or at least $\sqrt{n}$, or when we have a decent proportion of pairs of hyperedges that do not intersect each other. However, we were unable to extend this approach to be applicable to all linear hypergraphs.

Since the polynomials P_1 and P_2 have a rather difficult structure and involve a huge sum of monomials, approaches which provide a simplified sum or an exact formula for the coefficients seem essential to resolve the problem. As we have seen, Fact 2.7 already provides some simplification. We mention that suitable simplified formulae have been obtained in other settings of the Combinatorial Nullstellensatz, see [12, 13, 15, 24]. In these results, the key ingredient was the following coefficient formula, also mentioned as Quantitative Nullstellensatz.

Lemma 4.1 (Coefficient formula). *Let $\mathbb{F}$ be an arbitrary field and $P \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ a polynomial of degree $\deg(P) \leq d_1 + d_2 + \cdots + d_n$. For arbitrary subsets $C_1, C_2, \ldots, C_n$ of $\mathbb{F}$ with $|C_i| = d_i + 1$, the coefficient of $\prod x_i^{d_i}$ in P is*

$$\sum_{c_1 \in C_1} \sum_{c_2 \in C_2} \cdots \sum_{c_n \in C_n} \frac{P(c_1, c_2, \ldots, c_n)}{\phi_1'(c_1)\phi_2'(c_2)\ldots\phi_n'(c_n)},$$

where $\phi_i(z) = \prod_{c \in C_i}(z - c)$.

Once we have this coefficient formula, we may seek for a suitable Cartesian product set (or grid) $C_1 \times C_2 \times \ldots \times C_n$ on which polynomial P vanishes in most cases. This approach proved to be successful in several different combinatorial problems, see [13]. To reach an analogous goal, one may take a suitable describing polynomial and suitable monomial (rather than a different grid) in order to guarantee vanishing terms in the sum for the coefficient, via Vandermonde-completability.

References

[1] Alon, N., Combinatorial Nullstellensatz. Combinatorics, Probability and Computing, 8(1-2) (1999). 7–29.

[2] Alon, N., Tarsi, M., Colorings and orientations of graphs. Combinatorica, 12(2) (1992) 125–134.

[3] Araujo-Pardo, G., Vázquez-Ávila, A., (2016). A note on Erdös-Faber-Lovász Conjecture and edge coloring of complete graphs. arXiv preprint arXiv:1605.03374.

[4] Erdős, P. (1981). On the combinatorial problems which I would most like to see solved. Combinatorica, 1(1), 25–42.

[5] Erdős, P. (1988). Problems and results in combinatorial analysis and graph theory. In Annals of Discrete Mathematics (Vol. 38, pp. 81–92). Elsevier.

[6] Faber, V. (2010). The Erdös-Faber-Lovász conjecture–the uniform regular case. Journal of Combinatorics, 1(2), 113–120.

[7] Faber, V., Harris, D. G. (2019) Edge-coloring linear hypergraphs with medium-sized edges. Random Structures and Algorithms, 55(1), 153–159.

[8] Hindman, N. (1981) On a conjecture of Erdös, Faber and Lovász about n-colourings, Canadian J. Math. 33., 545–549.

[9] Jackson, B., Sethuraman, G., Whitehead, C. (2007). A note on the Erdős–Farber–Lovász conjecture. Discrete Mathematics, 307(7-8), 911–915.

[10] Kahn, J. (1992). Coloring nearly-disjoint hypergraphs with n+ o (n) colors. Journal of combinatorial theory, Series A, 59(1), 31–39.

[11] Kahn, J., Seymour, P. D. (1992). A fractional version of the Erdős–Faber–Lovász conjecture. Combinatorica, 12(2), 155–160.

[12] Karasev, R.N., Petrov, F.V. (2012). Partitions of nonzero elements of a finite field into pairs. Israel J. Math. 192, 143–156.

[13] Károlyi, G., Nagy, Z. L., Petrov, F. V., Volkov, V. (2015). A new approach to constant term identities and Selberg-type integrals. Advances in Mathematics, 277, 252–282.

[14] Kaul, H., Mudrock, J. A. (2020). Combinatorial Nullstellensatz and DP-coloring of Graphs. arXiv preprint arXiv:2003.01112.

[15] Lasoń, M. (2010). A generalization of Combinatorial Nullstellensatz. the electronic journal of combinatorics, N32-N32.

[16] Lovász, L. (2007). Combinatorial problems and exercises (Vol. 361). American Mathematical Soc.

[17] Mitchem, J., Schmidt, R. L. (2010). On the Erdos-Faber-Lovasz Conjecture. Ars Combinatoria, 97, 497–505.

[18] Murthy, T. S. (2020). A proof of the Total Coloring Conjecture. arXiv preprint arXiv:2003.09658.

[19] Paul, V., Germina, K. A. (2012). On edge coloring of hypergraphs and Erdős–Faber–Lovász conjecture, Discrete Mathematics, Algorithms and Applications, 4(01), 1250003.

[20] Romero, D., Sanchez-Arroyo, A. (2007). Adding evidence to the Erdos-Faber-Lovasz conjecture. Ars Combinatoria, 85, 71–84.

[21] Romero, D., Sanchez-Arroyo, A. (2007). Advances on the Erdos-Faber-Lovasz conjecture. Oxford Lecture Series in Mathematics and its Applications, Oxford University Press **34** 285–298.

[22] Romero, D., Alonso-Pecina, F., The Erdős–Faber–Lovász conjecture is true for $n \leq 12$. Discrete Mathematics, Algorithms and Applications, 6(03) (2014) 1450039.

[23] Sánchez-Arroyo, A. (2008). The Erdős–Faber–Lovász conjecture for dense hypergraphs. Discrete Mathematics, 308(5-6), 991–992.

[24] Schauz, U. (2008). Algebraically Solvable Problems: Describing Polynomials as Equivalent to Explicit Solutions. the electronic journal of combinatorics, 15(R10), 1.

[25] Seymour, P. D. (1982). Packing nearly-disjoint sets. Combinatorica, 2(1), 91–97.

[26] Zhu, X. (2019). The Alon–Tarsi number of planar graphs. Journal of Combinatorial Theory, Series B, 134, 354–358.