Improving Gebauer's Construction of 3-Chromatic Hypergraphs with Few Edges

Jakub Kozik

Theoretical Computer Science Department, Faculty of Mathematics and Computer Science, Jagiellonian University, Kraków, Poland

Abstract

In 1964 Erdős proved, by randomized construction, that the minimum number of edges in a k-graph that is not two colorable is $O(k^2\, 2^k)$. To this day, it is not known whether there exist such k-graphs with smaller number of edges. Known deterministic constructions use much larger number of edges. The most recent one by Gebauer requires $2^{k+\Theta(k^{2/3})}$ edges. Applying a derandomization technique we reduce that number to $2^{k+\widetilde{\Theta}(k^{1/2})}$.

2012 ACM Subject Classification Mathematics of computing → Hypergraphs; Mathematics of computing → Probabilistic algorithms; Theory of computation → Pseudorandomness and derandomization

Keywords and phrases Property B, Hypergraph Coloring, Deterministic Constructions

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.89

Category Track A: Algorithms, Complexity and Games

Related Version `https://arxiv.org/abs/2102.11674`

Funding *Jakub Kozik*: This work was partially supported by Polish National Science Center (2016/21/B/ST6/02165).

1 Introduction

In 1964 Erdős proved in [3] that $(1+o(1))\frac{\mathrm{e}\ln(2)}{4}k^2 2^k$ edges are sufficient to build a k-graph[1] which is not two colorable. To this day that result provides the best known upper bound for the minimum number of edges in such hypergraph. The Erdős' bound results from the fact that random k-graph with that number of edges, built on a set of $k^2/2$ vertices can not be colored properly with two colors with high probability. In some sense this value is not far from being optimal since a straightforward probabilistic argument shows that any k-graph with at most 2^{k-1} edges is two colorable. Using random recoloring method, introduced by Beck in [1], Radhakrishnan and Srinivasan in [8] extended this result to k-graphs with at most $\Theta(\sqrt{k/\log(k)}2^k)$ edges (see also [2] for a simplified proof).

The best known deterministic construction of a k-graph that is not two colorable has been obtained by Gebauer [5]. It requires $2^{k+\Theta(k^{2/3})}$ edges. It is also the first construction in which the number of edges is $2^{k+o(k)}$. The main result of the current paper is an upgrade of this construction that allows to cut down the number of edges to $2^{k+\Theta((k\log(k))^{1/2})}$.

Within the whole paper, $\log(.)$ stands for binary logarithm. We are only concerned with vertex two coloring of hypergraphs. However, just like in the construction from [5], the presented method naturaly generalizes to any fixed number of colors and gives analogous improvement. Vertex coloring is *proper* if no edge is monochromatic. Following common convention we use colors *red* and *blue*.

[1] i.e. k-uniform hypergraph

2 Gebauer's construction

We start with recalling the construction of [5], as we are going to modify it. The whole procedure is parameterized by $t = t(k)$ that takes value roughly k^α for some optimized positive $\alpha < 1$. It it convenient to organize the vertices of the constructed hypergraph into a rectangular matrix $\mathbb{M}$ of $2t-1$ rows and s columns. In particular, the vertex set of the hypergraphs has $(2t-1)\cdot s$ elements. The value of s will be a subject of optimization and in the final construction we are going to chose $s = \lceil \mathrm{e}\,(k/t)\,2^t \rceil$. Slightly abusing the notation, we use $\mathbb{M}$ for both the matrix and the set of vertices. We use the same convention for submatrices of $\mathbb{M}$.

2.1 Preliminary choice of rows

As the vertices of constructed hypergraph are identified with the entries of $\mathbb{M}$, vertex coloring can be seen as assigning colors to the entries of the matrix. A color is *dominating* in a row if at least half of its entries are colored with it (there can be two dominating colors). A matrix for which one of the colors is dominating in all rows will be called *consistently dominated*. The main part of the construction is designed to work with a consistently dominated submatrix of $\mathbb{M}$ with t rows. Wlog we always assume that red is the dominating color in such a matrix.

Let $\mathcal{M}$ denote the set of submatrices of $\mathbb{M}$ built of every t rows. For every $M \in \mathcal{M}$ we apply the *main construction* described in the next section. The construction outputs hypergraph H_M. The union of the edge sets of these hypergraphs forms the edge set of the resulting hypergraph. For every coloring of $\mathbb{M}$ at least one submatrix $M \in \mathcal{M}$ is consistently dominated. The main construction guarantees that in such a case, H_M contains a monochromatic edge.

2.2 Main construction

Let $M \in \mathcal{M}$, recall that M has t rows. Our goal is to build a hypergraph H_M on the vertex set M such that for every consistently dominated coloring of M, there exists a monochromatic edge in H_M. For $(\sigma_1, \ldots, \sigma_t) \in [s]^t$, we denote by $M(\sigma_1, \ldots, \sigma_t)$ matrix M in which for every $i \in [t]$, the i-th row has been cyclically shifted by σ_i. The construction proceeds as follows.

For every
1. *sequence of shifts* $\sigma \in [s]^t$,
2. and *set of indices* $I \subset [s]$ of size k/t,

add to H_M an edge built from all elements of the columns of $M(\sigma)$ with indices in I.

Note that the edges of H_M are of size k as required.

Let us fix a consistently dominated coloring of M. We assume wlog that red is the dominating color of the rows. When the sequence of shifts is chosen randomly, the probability that some fixed column is red is at least 2^{-t}. As a consequence, for $s \geqslant (k/t)\,2^t$ the expected number of red columns is at least k/t. In particular, for some sequence of shifts, there exists a set of k/t red columns. Hence the edge built for these shifts and columns is monochromatic.

2.3 Counting

We have

$$\binom{2t-1}{t} < 2^{2t}$$

choices for the subset of rows in the preliminary step. Then, in the main construction, every sequence of t elements of $[s]$ and a subset of k/t elements of $[s]$ is used to build an edge. The number of choices is

$$s^t \cdot \binom{s}{k/t} \leqslant s^t \cdot \left(\frac{\mathrm{e}s}{k/t}\right)^{k/t}.$$

For $s = (k/t)\, 2^t$ (we assume for simplicity that it is an integer) we obtain

$$(k/t)^t\, 2^{t^2} \cdot \mathrm{e}^{k/t}\, 2^k = 2^{t\log(k/t)+t^2+k/t\log(\mathrm{e})+k}.$$

The total number of edges is smaller than

$$2^{2t+t\log(k/t)+t^2+k/t\log(\mathrm{e})+k}.$$

Finally we choose t so that the above exponent is minimized. That happens for $t = \Theta(k^{1/3})$. In the end we obtain that the total number of edges is $2^{k+\Theta(k^{2/3})}$.

3 Improved construction

We modify only the main construction. Recall that we work with matrix M with t rows. For a fixed consistently dominating coloring of M, sequence of shifts $\sigma \in [s]^t$ is called *good* if $M(\sigma)$ contains at least $s\, 2^{-t}$ red columns. The set of good sequences for a coloring $\mathcal{C}$ of M is denoted by $\mathcal{G}(\mathcal{C})$.

If we fix a consistently dominating coloring of M and choose the sequence of shifts $\sigma \in [s]^t$ uniformly at random, the expected number of red columns in $M(\sigma)$ is $s\, 2^{-t}$. That observation was used to justify that there exists a good sequence. However, it also suggests that a large number of shift sequences might be good. For the constructed hypergraph not to be two colorable, it is sufficient that for every consistently dominated coloring of M, at least one such sequence is used in the main construction.

We apply derandomization techniques to construct relatively small set of sequences of shifts that can be used in the main construction instead of $[s]^t$. For a family of sets $\mathcal{F}$, a set that intersects every element of that family is called a *hitting set* for $\mathcal{F}$. In these terms we are looking for a small hitting set for family $\mathcal{G}_M = \{\mathcal{G}(\mathcal{C}) : \mathcal{C}$ is a consistently dominating coloring of $M\}$.

3.1 Sequential choice of shifts

We start with estimating the size of the set of good shift sequences. While it is not directly used in our construction, it provides a good opportunity to introduce some tools. It will also allow us to derive a probabilistic argument that small hitting sets actually exist.

The property of being good is generalized to prefixes in the straightforward way – sequence of shifts $(\sigma_1, \ldots, \sigma_i)$ is *good* if the matrix trimmed to the first i rows and shifted according to the sequence, has at least $s\, 2^{-i}$ red columns.

Suppose that $(\sigma_1, \ldots, \sigma_i)$ is good. We want to estimate the number of possible choices of σ_{i+1} for which $(\sigma_1, \ldots, \sigma_i, \sigma_{i+1})$ is good as well. If the coloring of the $(i+1)$-th row was "random", then about half of the choices would be right, and almost all of the choices would be almost right. That property does not hold in the worst case scenario and hence we are going to work with relaxed definitions.

For $\varepsilon > 0$, a sequence of shifts $(\sigma_1, \dots, \sigma_i)$ is *ε-good* if the number of red columns in the shifted matrix trimmed to the first i rows is at least $s\ \left(\frac{1-\varepsilon}{2}\right)^i$. Then, every ε-good sequence of shifts of length t gives a shifted matrix with at least

$$s\frac{(1-\varepsilon)^{t-1}}{2^t}$$

red columns. For $s \geqslant \mathrm{e}\ (k/t)\ 2^t$ and $\varepsilon = 1/t$, the number of red columns is at least k/t as needed. In the modified construction we set s to $\lceil \mathrm{e}\ (k/t)\ 2^t \rceil$.

We also define $\mathcal{G}^\varepsilon(\mathcal{C})$ as the set of ε-good sequences for a coloring $\mathcal{C}$ of M and $\mathcal{G}^\varepsilon_M$ as $\{\mathcal{G}^\varepsilon(\mathcal{C}) : \mathcal{C} \text{ is a consistently dominating coloring of } M\}$.

The following proposition is used to derive a lower bound for the number of ε-good sequences. It is formulated in more general terms that needed here, but we are going to use it again later. For a set $A \subset [s]$ and a number x, set $A + x$ is defined as the set A shifted cyclically within $[s]$ by x, formally $A + x = \{(a - 1 + x)(\mod s) + 1 : a \in A\}$.

▶ **Proposition 1.** *For any positive $\varepsilon < 1$ and sets $A, B \subset [s]$, let $\alpha = |B|/s$, there exist at least*

$$\frac{\varepsilon}{1-(1-\varepsilon)\alpha}\,\alpha s$$

elements $x \in [s]$ for which $|(A + x) \cap B| \geqslant (1 - \varepsilon)\alpha|A|$.

Proof. Let random variable X denote the size of $(A+x)\cap B$, when $x \in [s]$ is chosen uniformly at random. By the fact that $|B| = \alpha s$ and linearity of expectation we obtain

$$\mathbb{E}(X) = \alpha|A|.$$

From the definition of X, we get also

$$X \leqslant |A|.$$

We can observe now that a distribution that minimizes $\Pr[X > (1 - \varepsilon)\alpha|A|]$ and satisfies the above conditions, is supported only by values $(1 - \varepsilon)\alpha|A|$ and $|A|$. There is only one such distribution that satisfies $\mathbb{E}(X) = \alpha|A|$. Straightforward calculations give

$$\Pr[X > (1-\varepsilon)\alpha|A|] \geqslant \frac{\varepsilon\alpha}{1-(1-\varepsilon)\alpha}.$$ ◀

For $|B| \geqslant s/2$ we get that there exist at least

$$\frac{2\varepsilon}{1+\varepsilon}\,s/2$$

elements $x \in [s]$ for which $|(A + x) \cap B| \geqslant (1 - \varepsilon)|A|/2$.

Applying the proposition iteratively, we obtain that the number of ε-good sequences of length j is at least

$$\left(\frac{\varepsilon}{1+\varepsilon}\,s\right)^j.$$

(For a fixed j, and some ε-good sequence σ of length $j - 1$, let A be the set of indices of the red columns in the matrix trimmed to the first $j - 1$ rows and shifted according to σ, and B be the set of indices of red entries of the j-th row.)

For $j = t$ we get a lower bound for the number of ε-good sequences. Once we have that bound, a typical application of the probabilistic method (along the lines of the proof from [3]) allows to proof that there exists a hitting set for $\mathcal{G}_M^\varepsilon$ of size $2^{O(t\log(t))}$. That argument is briefly described below.

Recall that, for a fixed consistently dominated coloring of matrix M with t rows, the volume of ε-good sequences is at least

$$p = \left(\frac{\varepsilon}{1+\varepsilon}\right)^t.$$

The volume is exactly the probability that uniformly random sequence is ε-good. Let S be a set built from m uniformly and independently sampled random sequences from $[s]^t$. (Since, the sequences are sampled with repetitions, it may happen that $|S| < m$.) The following formula upperbounds the expected number of consistently dominated colorings of M, for which the set of ε-good sequences is not hit by S

$$2^{st} \cdot (1-p)^m < \exp(st\ln(2) - mp).$$

Therefore, whenever $st\ln(2) - mp \leqslant 0$, some set of m sequences hits all the sets of ε-good sequences for consistently dominating colorings. For $s = \lceil \mathrm{e}\,(k/t)\,2^t \rceil$ and $\varepsilon = 1/t$ it is sufficient to take m of the order $2^{O(t\log(t))}$ to satisfy the inequality. As a consequence there exists a hitting set for $\mathcal{G}_M^\varepsilon$ of size $2^{O(t\log(t))}$.

3.2 Expanders for hitting sets

Linial, Luby, Saks and Zuckerman [6] worked on deterministic constructions of small hitting sets for combinatorial rectangles. We summarize in this section, their results that are relevant for our developments. We follow closely their definitions.

Graph $G = (V, E)$ is an (m, Δ, α)-*expander* if it has m vertices, maximum degree Δ and for any $A \subset V$, the fraction of vertices in $V - A$ that have a neighbor in A is at least $\alpha|A|/m$. For a fixed graph G let W_r denote the set of walks in G of length r. Let $W_{r,d}$ be the set of subsequences of elements of W_r of length d (not necessarily subsequences of consecutive elements). Set $R \subset [m]^d$ is a *combinatorial rectangle* if it is of a form $R_1 \times \ldots \times R_d$ for some $R_1, \ldots, R_d \subset [m]$. The volume of rectangle R, denoted as $vol(R)$, is defined as $|R|/m^d$.

▶ **Lemma 2** **([6]).** *Let m, d be positive integers and R be a rectangle in $[m]^d$. Suppose G is an (m, Δ, α)-expander with $1/2 > \alpha > 0$. If $r = 1 + (4/\alpha)(d + \log(1/vol(R)))$, then $W_{r,d}$ contains a point from R.*

The above lemma implies that a specific set of sequences $W_{r,d}$ hits every combinatorial rectangle in $[m]^d$ of sufficiently large volume.

The following rough estimations for the size of $W_{r,t}$ will be sufficient for our needs. We have

$$|W_r| \leqslant m(\Delta+1)^r$$

and

$$|W_{r,d}| < 2^r |W_r| \leqslant m\,(2(\Delta+1))^r.$$

Lemma 2 leaves some space for the choice of expander graph. Authors of [6] used the construction of Margulis [7] (see also [4]) which allows to build an expander with $\Delta = 8$ and $\alpha = (2-\sqrt{3})/4$. A minor inconvenience is that the construction requires the number of vertices to be a perfect square. However, as observed already in [6], we can consider the rectangles of our interest as subsets of a larger space $[m']^d$, and apply the lemma in that space. For every m we can choose number m' that is a perfect square and satisfies $m \leqslant m' \leqslant 2m$. While that change affects the volumes of rectangles, they get smaller at most by a factor of 2^{-d}. For our purposes this cost is negligible.

When we are interested in rectangles of volume at least $\mathcal{V}$, Lemma 2 instructs to take

$$r = r(d, \mathcal{V}) = 1 + (4/\alpha)(d + \log(2^d/\mathcal{V})).$$

For some specific constant $\hat{C}$ and for all positive d and $\mathcal{V}$ we have

$$r(d, \mathcal{V}) \leqslant \hat{C}(d + \log(1/\mathcal{V})).$$

▶ **Corollary 3.** *There exists a constant $C > 0$ such that, for every integers m, d, and $\mathcal{V} > 0$ there exists an efficiently constructible subset of $[m]^d$ of size at most*

$$m \cdot 2^{C(d+\log(1/\mathcal{V}))},$$

that intersects every combinatorial rectangle in $[m]^d$ of volume at least $\mathcal{V}$.

We apply that result, to construct a small hitting set for $\mathcal{G}^{\varepsilon}_{M}$. That set is then used in the modified main construction instead of the set of all shift sequences.

3.3 Under false assumption

Unfortunately, for a fixed consistently dominating coloring of M, the set of good or ε-good shift sequences does not need to form a combinatorial rectangle. It is instructive to pretend for a moment that it does. We assume (falsely) in this subsection that $\mathcal{G}^{\varepsilon}_{M}$ contains only combinatorial rectangles.

By the discussion that follows Proposition 1, for every consistently dominating coloring of M, the set of ε-good shift sequences has volume at least

$$\nu = \left(\frac{\varepsilon}{2(1+\varepsilon)}\right)^t.$$

By Corollary 3 there exists a hitting set *HS* for all rectangles of volume ν of size $s \cdot 2^{C(t+\log(1/\nu))}$. For $\varepsilon = 1/t$ and $s = \lceil \mathrm{e}\ (k/t)\ 2^t \rceil$, the size of *HS* is at most $2^{2Ct\log(t)}$ (assuming that t is sufficiently large). Note that in the original construction all possible shift sequences were used. Using set *HS* instead of $[s]^t$ and choosing $t = (k\log(k))^{1/2}$, the total number of edges becomes

$$2^{k+O((k\log(k))^{1/2})}.$$

3.4 Decomposing good shift sequences

We showed in Section 3.1 that, for every consistently dominating coloring of M, the set of ε-good shift sequences is large. While, in general, it does not have a structure of combinatorial rectangle, in some sense it can be decomposed into a small number of such. We start by altering the way that the sequences of shifts are represented. For the clarity of the exposition we assume that t is a power of 2.

Let T be a rooted plane complete binary tree with t leaves[2]. A subtree rooted at some internal node of T consists of that node and all its descendants. A node of T is *at level* j if its distance to the set of leaves is j. Let S_j be the set of inner nodes at level j. Note that $|S_j| = t\,2^{-j}$, we denote that value by d_j. For $h = \log(t)$, the tree has $h+1$ levels with all the leaves on level 0.

We associate leaves of T with rows of M in such a way that the i-th leaf from the left, corresponds to the i-th row. Inner nodes of the tree are going to be labeled by elements of $[s]$. These labels represent the relative shifts between neighboring rows of M. For an inner node v, if l is the rightmost leaf of the left subtree of v and r is the leftmost leaf of the right subtree of v, then the label of v describes how row r is shifted wrt l.

Labeling of a subtree rooted at node v is ε-good, if for r being the number of descendant leaves of v, the submatrix of the rows that correspond to these leaves, shifted according to the labels of the inner nodes of the subtree, has at least $s\,((1-\varepsilon)/2)^r$ red columns. Note that ε-good labellings of the whole tree correspond to ε-good sequences (up to a cyclic shift of the whole matrix, which is clearly redundant in the original construction).

We order the nodes of S_j from left to right and represent labellings of the nodes of S_j as elements of $[s]^{d_j}$. We are going to work bottom up and label inner nodes in groups consisting of the nodes of the same level. A labeling of T is *ε-good up to level j* if all the subtrees rooted at level at most j are ε-good. In all the places where we use this definition, it can be assumed that the labeling is undefined for the nodes of higher levels. Suppose that τ is a labeling of T that is ε-good up to level $j-1$. Then, a sequence of labels $\sigma \in [s]^{d_j}$ is called *an ε-good level j extension (of τ)* if the labeling τ in which the labels of the nodes of level j has been set to σ is ε-good up to level j.

▶ **Proposition 4.** *Suppose, that a labelling of T is ε-good up to level $j-1$. Then, the set of its ε-good level j extensions forms a combinatorial rectangle of volume at least*

$$\nu_j = \left(\varepsilon\,((1-\varepsilon)/2)^{-2^{j-1}}\right)^{d_j}.$$

Proof. Fix j and suppose that labeling τ is ε-good up to level $j-1$. We want to assign labels to the nodes of S_j in such a way that all the subtrees rooted at depth j are ε-good shift trees as well. Note that for any pair of distinct nodes of level j, the property of the corresponding subtrees of being ε-good shift trees is determined by disjoint sets of rows of the underlying matrix. That justify that the set of ε-good level j extensions forms a combinatorial rectangle.

Let v be a node of S_j and let A and B be the sets of indices of red columns respectively in the shifted submatrices corresponding to the left and right subtrees of v. By the assumptions we know that both these sets have cardinality at least

$$s\,((1-\varepsilon)/2)^{-2^{j-1}}.$$

We need to estimate the number of $x \in [s]$ for which the set $A \cap (B+x)$ has cardinality at least

$$s\,((1-\varepsilon)/2)^{-2^{j}}.$$

Proposition 1 gives that there exist at least

$$\varepsilon\,((1-\varepsilon)/2)^{-2^{j-1}}\,s$$

[2] i.e. all the internal nodes of T have two children (left and right) and all the leaves are of the same distance from the root

such values. We obtain that the volume of combinatorial rectangle of ε-good level j extensions is at least

$$\left(\varepsilon\left((1-\varepsilon)/2\right)^{-2^{j-1}}\right)^{d_j}$$ ◀

By Corollary 3, there exists a set HS_j of cardinality

$$s \cdot 2^{C(d_j+\log(1/\nu_j))},$$

that is a hitting set for the family of ε-good level j extensions for labellings that are ε-good up to level $j-1$. That implies the following proposition.

▶ **Proposition 5.** *Set $HS = HS_1 \times \ldots \times HS_h$ is a hitting set for the family of sets of ε-good labellings of T.*

It remains to estimate the size of HS. We have

$$\begin{aligned}|HS| &\leqslant \prod_{j=1\ldots h} s \cdot 2^{C(d_j+\log(1/\nu_j))}\\ &= s^{\log(t)} \cdot 2^{C\sum_{j=1\ldots h}(d_j+\log(1/\nu_j))}\\ &< s^{\log(t)} \cdot 2^{Ct} \cdot 2^{C\sum_{j=1\ldots h}\log(1/\nu_j)},\end{aligned}$$

and

$$\begin{aligned}\sum_{j=1\ldots h}\log(1/\nu_j) &= \sum_{j=1\ldots h} d_j(\log(1/\varepsilon)+2^{j-1}\log(2/(1-\varepsilon)))\\ &< t\log(1/\varepsilon) + t\sum_{j=1\ldots h}\log(4) \quad (\text{for } \varepsilon<1/2)\\ &= t\cdot\log(1/\varepsilon) + 2t\cdot\log(t)\end{aligned}$$

Therefore, for our parametrization (i.e. $s = \lceil \mathrm{e}\,(k/t)\,2^t\rceil$ and $\varepsilon = 1/t$), and for all sufficiently large t we get

$$|HS| \leqslant 2^{4t\log(t)}.$$

3.5 Modified main construction

Let HS be the set from Proposition 5. As we already observed labellings of T correspond to shift sequences up to a cyclic shift of the whole matrix. For a labeling τ let $\sigma(\tau)$ be a shift sequence that is compatible with τ. Observe, that if τ is an ε-good labeling, then $\sigma(\tau)$ is ε-good shift sequence. Recall that we chose $s = \mathrm{e}\,(k/t)\,2^t$ so that if σ is an ε-good sequence for some consistently dominated coloring of M, then $M(\sigma)$ has at least k/t red columns. The modified main construction proceeds as follows.

`For every`
1. `labeling of the tree` $\tau \in HS$`,`
2. `and set of indices` $I \subset [s]$ `of size` k/t`,`

`add to` H_M `an edge build from all elements of the columns of` $M(\sigma(\tau))$ `with indices in` I`.`

By Proposition 5 for every consistently dominated coloring of M, at least one ε-good labeling τ is used in the construction. Then, for every such coloring, matrix M shifted according to $\sigma(\tau)$ has at least k/t red columns. As a consequence at least one of the edges of H_M is monochromatic.

Counting

Just like in the original construction, we have less than 2^{2t} choices for the subset of rows in the preliminary step. Then, in the modified main construction, we use every sequence of HS with every subset of k/t elements of $[s]$ to build an edge. The number of choices is smaller than

$$2^{4t\log(t)} \cdot \binom{s}{k/t} < 2^{4t\log(t)} \cdot \left(\frac{\mathrm{e}s}{k/t}\right)^{k/t}.$$

Substituting the value of s we obtain a value that is smaller than

$$2 \cdot 2^{4t\log(t)} \cdot \mathrm{e}^{2k/t} 2^k = 2^{1+4t\log(t)+(2k/t)\log(\mathrm{e})+k}.$$

The bound is multiplied by 2 to compensate for the ceiling in the definition of s. Taking into account preliminary choices of rows, the total number of edges is smaller than

$$2^{2t+1+4t\log(t)+(2k/t)\log(\mathrm{e})+k}.$$

For $t = (k/\log(k))^{1/2}$, the total number of edges becomes $2^{k+\Theta((k\log(k))^{1/2})}$.

References

1 József Beck. On 3-chromatic hypergraphs. *Discrete Mathematics*, 24(2):127–137, 1978. doi:10.1016/0012-365X(78)90191-7.

2 Danila D. Cherkashin and Jakub Kozik. A note on random greedy coloring of uniform hypergraphs. *Random Structures & Algorithms*, 47(3):407–413, 2015. doi:10.1002/rsa.20556.

3 Paul Erdős. On a combinatorial problem. II. *Acta Mathematica Academiae Scientiarum Hungaricae*, 15:445–447, 1964.

4 Ofer Gabber and Zvi Galil. Explicit constructions of linear-sized superconcentrators. *J. Comput. System Sci.*, 22(3):407–420, 1981. Special issued dedicated to Michael Machtey. doi:10.1016/0022-0000(81)90040-4.

5 Heidi Gebauer. On the construction of 3-chromatic hypergraphs with few edges. *Journal of Combinatorial Theory. Series A*, 120(7):1483–1490, 2013. doi:10.1016/j.jcta.2013.04.007.

6 Nathan Linial, Michael Luby, Michael Saks, and David Zuckerman. Efficient construction of a small hitting set for combinatorial rectangles in high dimension. *Combinatorica*, 17(2):215–234, 1997. doi:10.1007/BF01200907.

7 G. A. Margulis. Explicit constructions of expanders. *Problemy Peredači Informacii*, 9(4):71–80, 1973.

8 Jaikumar Radhakrishnan and Aravind Srinivasan. Improved bounds and algorithms for hypergraph 2-coloring. *Random Structures & Algorithms*, 16(1):4–32, 2000. doi:10.1002/(SICI)1098-2418(200001)16:1<4::AID-RSA2>3.3.CO;2-U.