Three-Coloring Three-Dimensional Uniform Hypergraphs

Ahmad Biniaz* Prosenjit Bose† Jean Cardinal‡ Michael S. Payne§

Abstract

We study the chromatic number of hypergraphs whose vertex-hyperedge incidence poset has dimension at most three. Schnyder (1989) showed that graphs with this property are planar and thus four-colorable. Results of Keszegh and Pálvölgyi (2015) imply that k-uniform hypergraphs with dimension at most three are two-colorable for $k \geq 9$. In this paper we show that k-uniform hypergraphs with dimension at most three are three-colorable for $k \geq 6$. This implies three colorability of k-uniform triangle Delaunay hypergraphs and k-uniform hypergraphs induced by points and octants in 3-space. We also observe that the chromatic number of k-uniform hypergraphs with dimension $d \geq 4$ is not bounded by any function of k and d.

1 Introduction

A *hypergraph* G consists of a set of *vertices* and a collection of non-empty subsets of vertices called *hyperedges*. The *incidence poset* of G is the partially ordered set (poset) describing the vertex-hyperedge containment relationship. The order dimension of a poset $\mathcal{P}$ is defined as the minimum size of a set of total orders on the elements of $\mathcal{P}$ whose intersection is $\mathcal{P}$. The *dimension* of G is the order dimension of its incidence poset. A *k-uniform* hypergraph (or *k-graph*) is a hypergraph in which all hyperedges have cardinality k. A (simple) *graph* is one for which k is 2. A *c-coloring* of G is to color each vertex by one of the colors $\{1, \ldots, c\}$ such that no edge of G has all vertices of the same color. A hypergraph is *c-colorable* if it admits a c-coloring.

In 1989, Schnyder showed that a graph has dimension at most three if and only if it is planar [13]. Therefore, all such graphs are 4-colorable by the Four Colour Theorem. We study the problem of coloring k-graphs of dimension at most three. We will refer to hypergraphs of dimension at most three as *three-dimensional hypergraphs*. It follows from the seminal work of Keszegh and Pálvölgyi [10] that three-dimensional k-graphs are 2-colorable for $k \geq 9$. In this note we adapt their approach and show the following result.

*University of Waterloo, ahmad.biniaz@gmail.com
†Carleton University, jit@scs.carleton.ca
‡Université Libre de Bruxelles (ULB), jcardin@ulb.ac.be
§La Trobe University, m.payne@latrobe.edu.au

Theorem 1 *Every three-dimensional k-uniform hypergraph is 3-colorable, for $k \geq 6$.*

2 Background

The dimension of a hypergraph can be determined by representing the incidence poset as the intersection of a number of total orders on vertices. The following is a well-known characterization of hypergraphs of dimension d that we will rely upon often [13].

Proposition 2 (Schnyder 1989) *A hypergraph H has dimension at most d if and only if there exist d total orders $<_1, \ldots, <_d$ on the vertices of H such that*

- *the intersection of all the orders is empty, and*
- *for each hyperedge e of H and each vertex $z \notin e$ there exists i such that $x <_i z$ for every $x \in e$.*

This characterization implies that any hyperedge e is uniquely determined by its maximum vertices in the d total orders. Every vertex not in e must be above at least one of these maxima. See Figure 6 for an illustration of a hyperedge.

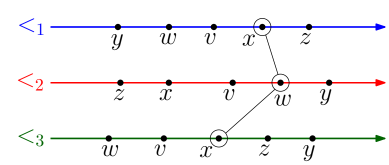

Figure 1: Illustration of a hyperedge $e = \{v, w, x\}$.

Our work is inspired by the work of Keszegh and Pálvölgyi [10] on coloring octant k-graphs (a subclass of k-graphs). Given any finite set P of points in $\mathbb{R}^3$, take as hyperedges every set of k points that is the intersection of P with some axis-parallel octant (which is an open set of the form $(-\infty, x) \times (-\infty, y) \times (-\infty, z)$ with apex point (x, y, z)). They showed that any octant k-graph is 2-colorable for $k \geq 9$, and there are octant 4-graphs that are not 2-colorable. It is implied by Proposition 2 that three-dimensional k-graphs are a subclass of octant k-graphs—just use the three total orders to give coordinates to the vertices. (Octant hypergraphs are more general because the coordinates need not satisfy the first property in Proposition 2.)

Another relevant class of geometric hypergraphs are triangle Delaunay k-graphs: Given a finite set P of points in general position in $\mathbb{R}^2$ and a triangle T, a set of k points form a hyperedge if there exists a homothet[1] of T containing just those k points. The classical Delaunay graph of a point set has a similar construction but with respect to a circle instead of a triangle.

With Proposition 2, it can readily be seen that triangle Delaunay hypergraphs have dimension at most three. The three necessary total orders can be obtained by sweeping P with three lines parallel to the three sides of T, as in Figure 2. Combining this with Theorem 1 we get the following corollary.

Corollary 3 *Every k-uniform triangle Delaunay hypergraph is 3-colorable, for $k \geq 6$.*

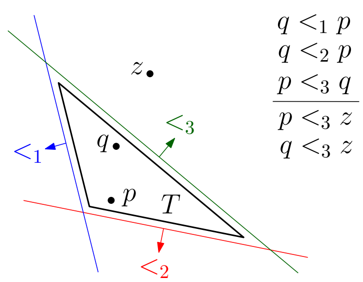

Figure 2: Obtaining three total orders that satisfy the conditions of Proposition 2. The set $\{p, q\}$ is an edge of triangle Delaunay 2-graph.

Returning to graphs, another (perhaps lesser known) result of Schnyder [13, Corollary 5.4] implies that every three-dimensional graph (and thus every planar graph) can be represented as a subgraph of a triangle Delaunay graph. Schnyder called these representations 'barycentric embeddings'. One might wonder whether every three-dimensional k-graph is also a subgraph of a triangle Delaunay k-graph. We note that there exist three-dimensional 10-graphs that are not realizable as a triangle Delaunay 10-graph (Stefan Felsner, private communication). To sum up, we can see that the class of triangle Delaunay hypergraphs is a proper subset of the class of three-dimensional hypergraphs, which is in turn a proper subset of the class of octant hypergraphs.

There exist a large number of fascinating coloring problems for geometric hypergraphs that are closely related to the problem studied here; e.g. [1, 2, 3, 4, 5, 6, 8, 9, 11, 12]. In particular, the question of whether k-uniform Delaunay hypergraphs (induced by points and circles in the plane) are 3-colorable for some k remains open [1], while the analogous questions for homothets of a convex n-gon have relatively loose bounds on k [11]. On the other hand, it is known that there is no finite k such that k-graphs induced by axis parallel rectangles on points in the plane are 2-colorable [6]. For further related problems and results see the discussion in [11].

[1]Homothets of T include translations and scalings but not rotations or reflections.

3 Further results

The result about axis parallel rectangles just mentioned, which is due to Chen, Pach, Szegedy, and Tardos [6], implies that there can be no analogue of Theorem 1 in higher dimensions (see Corollary 5). The statement for $k = 2$ (graphs) was first proved by Ossona de Mendez and Rosenstiehl [12], then rediscovered by Trotter and Wang [14].

Theorem 4 (Chen et al. [6]) *For any positive integers c and k, there is a finite point set in the plane with the property that no matter how we color its elements with c colors, there always exists an axis-parallel rectangle containing at least k points, all of which have the same color.*

Corollary 5 *For any triple of integers $c \geq 1$, $k \geq 1$, and $d \geq 4$, there exists a d-dimensional k-uniform hypergraph that is not c-colorable.*

Proof. Let H_1 be a hypergraph whose vertex set is a finite point set P in the plane that satisfies the conditions of Theorem 4, and whose edge set contains all k-subsets of points that can be obtained by intersecting P with an axis-aligned rectangle. By considering the four total orders obtained by sweeping P horizontally and vertically in both directions, and using Proposition 2, we observe that H_1 is a k-uniform hypergraph with dimension at most 4. Theorem 4 implies that H_1 is not c-colorable.

Let H_2 be any d-dimensional k-uniform hypergraph. Then the disjoint union of H_1 and H_2 is a k-uniform hypergraph that is d-dimensional due to the dimension of H_2, and not c-colorable because H_1 requires more than c colors. □

Finally we note an extension of Theorem 1. Despite the important role of graph planarity in the proof of Theorem 1, the analogous result for octant k-graphs follows as a corollary.

Corollary 6 *Every k-uniform octant hypergraph is 3-colorable, for $k \geq 6$.*

Proof. An octant k-graph has vertex set $P \subset \mathbb{R}^3$. We consider the three coordinate directions as three total orders on P. Unlike the case of three-dimensional k-graphs, the intersection of the three orders may not be empty. Consider the poset B on P with the order relation being the intersection of these three orders. Note that if u dominates v in this partial order then every octant containing u also contains v.

Start with the subset $S \subset P$ consisting of all the minimal elements in B. As S is an antichain in B, it induces a three-dimensional k-graph, and so we may 3-color it by Theorem 1.

Now add in a point x that is minimal in $B \setminus S$, and notice that x dominates some point(s) of S. This means that some hyperedges disappear, and some others are created with x and $k-1$ points of S. We need only ensure that these new hyperedges are properly colored, and this can be done by giving x any of the 3 colors that is distinct from the color of a point dominated by x.

By iteratively adding minimal elements from the remaining points, we can build up a 3-coloring for the whole set P. □

4 Proof of Theorem 1

Let H denote the three-dimensional k-graph that we want to color. Following Keszegh and Pálvölgyi, the proof strategy involves constructing a graph F that has an edge in every hyperedge of H. Thus a proper coloring of F is a proper coloring of H. In the proof of Keszegh and Pálvölgyi, F is a forest, and therefore 2-colorable. In our proof, F is a triangle-free three-dimensional graph, and thus planar by Schnyder's theorem and hence 3-colorable by Grötzsch's theorem which says that triangle-free planar graphs are 3-colorable [7].

Let V be the vertex set of H, and let $<_1, <_2, <_3$ be the three total orders on V with empty intersection. For every ordered triple (x, y, z) of (not necessarily distinct) elements of V, we define the *combinatorial triangle* Δxyz as the subset of V determined by three maxima x, y, z:

$$\Delta xyz = \{v \in V \mid v \leq_1 x \ \wedge \ v \leq_2 y \ \wedge \ v \leq_3 z\}.$$

Triangles containing k elements are precisely the hyperedges of H. If two elements of V are reversed by $<_1$ and $<_2$ we say they are *incomparable*, otherwise they are *comparable*. If for two comparable elements x and y we have $x <_1 y$ and $x <_2 y$, then we say y *dominates* x. See Figure 3 for an illustration. Note that if y dominates x then we must have $y <_3 x$ because the intersection of $<_1, <_2, <_3$ is empty.

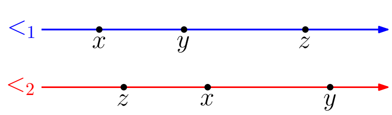

Figure 3: The elements x and z are incomparable, while x and y are comparable and y dominates x.

Without loss of generality we assume that H is edge maximal, that is, H contains all hyperedges (with k vertices) that satisfy the second property in Proposition 2. Let G be the edge-maximal planar graph obtained by the same three orders. We describe an iterative algorithm that processes V in the order of $<_3$ and creates a planar triangle-free subgraph F of G such that every hyperedge of H contains an edge of F. By slightly abusing notation, in the rest of description we refer to F as an edge set (the graph F is induced by this edge set).

We maintain a sequence Y of vertices and a set F of edges that satisfy the following invariants after each iteration of the algorithm:

- Elements of Y are pairwise incomparable and no element in Y dominates a processed element. (By the definition of incomparability, Y is ordered forwards by $<_1$ and backwards by $<_2$.)
- The set F is a subset of G, has no 3-cycles, has no edge between two vertices of Y, and has an edge in every hyperedge formed by processed vertices.

The sequence Y plays the role of "staircase" in [10] that separates the processed vertices from unprocessed vertices. Initially, F is empty and Y contains the least element in $<_3$. The algorithm processes the next vertex m in $<_3$ as follows:

(1) While there exists $v \in Y$ that dominates m, then add vm to F and remove v from Y.

(2) Add m to Y.

(3) While there exist three consecutive vertices $u, v, w \in Y$ such that $u <_2 v <_2 w$ and Δuwm does not contain any vertex outside Y, then add uv and vw to F and remove v from Y.

Since we apply step (3) greedily, any triple considered in this step contains m, that is $m \in \{u, v, w\}$. Moreover, since $u <_2 v <_2 w$ and the elements of Y are pairwise incomparable, we have $w <_1 v <_1 u$. This in turn implies that Δuwm contains u, v, and w.

It remains to show that each of the claimed properties for Y and F holds at the end of every iteration. In the proof of these properties we use the fact that m is the maximum element in $<_3$ that is processed so far, without further mentioning. Let X denote the set of processed elements that are not in Y.

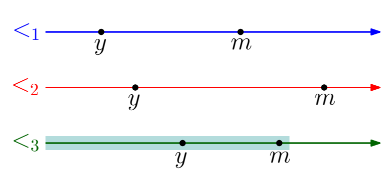

Figure 4: The processed elements (in the order $<_3$) are shaded; m is processed in the current iteration.

– *Elements of Y are pairwise incomparable:* Before we add a new vertex m to Y in step (2), we remove all vertices that dominate m in step (1). The vertex m does not dominate any vertex $y \in Y$ because otherwise we would have $y <_1 m$, $y <_2 m$, and $y <_3 m$ which contradicts the intersection of total orders being empty; see Figure 4 for an illustration. Therefore, the elements of Y are pairwise incomparable after each iteration.

– *Elements of Y do not dominate elements of X:* We need to ensure this only when we add the current element m to Y. If m dominates an element of X, then similar to the previous claim (as in Figure 4) we get a contradiction to the emptiness of the intersection of total orders.

– *F has no edge between two vertices of Y:* In step (1) after adding the edge vm, we remove v from Y. In step (3) after adding the edges uv and vw, we remove v from Y. Therefore, this claim follows.

– *F is a subset of G:* Consider an edge vm added to F in step (1). We show that the triangle Δvvm contains only v and m; this implies that vm is an edge of G. This triangle contains m because $m <_1 v$ and $m <_2 v$ (as v dominates m) and contains v because $v <_3 m$ (as m is the largest element of $<_3$ processed so far). Now we show, by contradiction, that Δvvm does not contain any other point. Recall that before adding vm, the vertex v belongs to Y. If Δvvm contains another element $y \in Y$ (as in Figure 5) then v dominates y; this contradicts the fact that elements of Y are pairwise incomparable. If Δvvm contains an element $x \in X$ (as in Figure 5) then v dominates x; this contradicts the fact that elements of Y do not dominate elements of X.

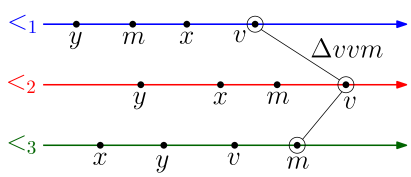

Figure 5: The triangle Δvvm contains m, v, x, and y.

Now consider edges uv and vw added to F in step (3), and recall that due to greedily application of this step we have $m \in \{u, v, w\}$. Our choices of u, v, w (as three consecutive elements of Y) and Δuwm (as having no element of X) ensures that Δuwm contains only u, v, w. In this setting the triangle $\Delta uv*$ contains only u, v, and the triangle $\Delta vw*$ contains only v, w ($*$ represents the maximum of the first two elements with respect to $<_3$). Thus, uv and vw are edges of G. See Figure 6 for an illustration.

– *F has no 3-cycle:* Since there are no edges between elements of Y, the edges added in step (1)—between m and elements of Y—do not create any 3-cycle. Consider edges uv and vw added in step (3). Since there was no edge between u and w which belong to Y, the three vertices u, v, w cannot form a 3-cycle. If uv creates a 3-cycle then u and v were joined by a path of length two

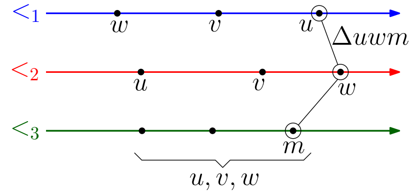

Figure 6: The triangle Δuwm contains only u, v, w.

through x say which now belongs to X. For x to have two neighbors in Y, the edges ux and xv must come from a prior application of step (3), and thus we must have $u <_2 x <_2 v$ and $v <_1 x <_1 u$, as in Figure 7. In this setting the triangle Δuwm contains x, which contradicts the current application of step (3). Thus uv does not create a 3-cycle. A similar argument applies for vw.

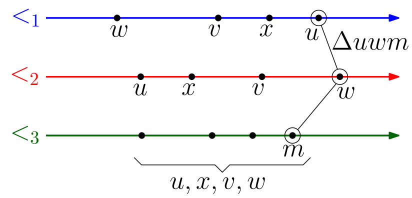

Figure 7: The triangle Δuwm contains u, v, w and x.

– *Every hyperedge formed by processed vertices contains an edge in F:* It suffices to show this only for hyperedges containing m. Consider any such hyperedge h, and recall that $|h| = k \geq 6$. Consider the state of Y and F at the end of current iteration, and set $i := |h \cap X|$. Depending on i, we consider three cases.

If $i = 0$, then all elements of h belong to Y and an edge would be added inside h in step (3).

Now suppose that $i = 1$, and let v be the only element in $h \cap X$ (here is the place where we use $k \geq 6$). Then $|h \cap Y| \geq 5$. Let $a <_1 b <_1 c <_1 d <_1 e$ be five consecutive elements of $h \cap Y$, and note that $e <_2 d <_2 c <_2 b <_2 a$. Thus $\Delta eam \subseteq h$, as in Figure 8. Let m_1 be the greatest of a, b, c and let m_2 be the greatest of c, d, e both with respect to $<_3$. In this setting either Δecm_2 or Δcam_1 does not contain v because otherwise $v <_1 c$, $v <_2 c$, and $v <_3 c$ (as in Figure 8) which contradicts the emptiness of the intersection of total orders. Therefore, in step (3) we get edges ed, dc or cb, ba, and thus h contains an edge of F.

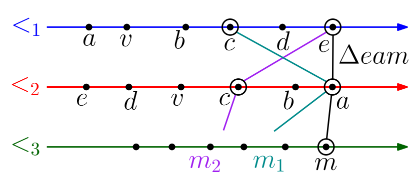

Figure 8: The triangle Δeam is a subset of h.

Now suppose that $i \geq 2$. Then $h \cap X$ contains two elements that are either comparable or incomparable.

First suppose that $h \cap X$ contains two comparable elements v and t. We may assume that t dominates v (this implies that $t <_3 v$). Moreover, we may assume that v was added to X in step (3) because otherwise by step (1) we get an edge in Δttm which is a subset of h. Since step (3) has been applied, as its prerequisites v has two incomparable neighbors u, w such that $u <_2 v <_2 w$ and Δuwm_1 does not contain t, where $m_1 \leq_3 m$ is the maximum of u, v, w with respect to $<_3$; see Figure 9. In step (3) the edges uv and vm were added to F. Recall that $t <_3 v$, and thus $t <_3 m_1$. In this setting either $u, v <_1 t$ (as in Figure 9) or $v, w <_2 t$. In the first case u and v are contained in Δttm_1 while in the second case v and w are contained in Δttm_1. Since Δttm_1 is a subset of Δttm which is in turn a subset of h, we get an edge of F in h.

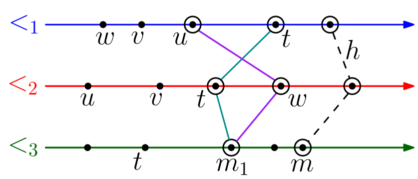

Figure 9: The triangle Δuwm_1 does not contain t.

Now suppose that $h \cap X$ contains two incomparable elements v and t. We may assume that both were added to X in step (3) because otherwise by step (1) we get an edge in Δvvm or in Δttm which are subsets of h. Without loss of generality assume that $t <_2 v$ and that v was added to X after t. As prerequisite of step (3) the vertex v has two incomparable neighbors $u <_2 v <_2 w$ in the triangle Δuwm_1 that does not contain t, where $m_1 \leq_3 m$ is the maximum of u, v, w with respect to $<_3$. In step (3) the edges uv and vm were added to F. We show (by contradiction) that Δtvm contains u and v, or v and w; this would imply our claim because Δtvm is a subset of h.

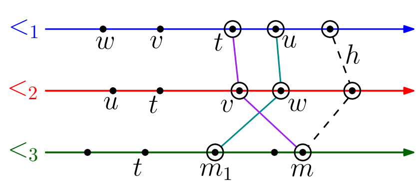

Figure 10: The triangle Δuwm_1 does not contain t.

Observe that Δtvm contains v. For the sake of contradiction assume that Δtvm does not contain any of u and w, and thus $t <_1 u$ and $v <_2 w$, as in Figure 10. Recall that v was added to X when we were processing the greatest element of $\{u, v, w\}$ in $<_3$, which is m_1. At that time, t was already in X which means that t was processed before m_1, i.e., $t <_3 m_1$, as in Figure 10. In this setting Δuwm_1 contains t, which contradicts the application of step (3) on u, v, w. This completes the proof.

Acknowledgement. This work initiated at the *Fifth Annual Workshop on Geometry and Graphs*, March 5-10, 2017, at the Bellairs Research Institute of McGill University, Barbados. We are grateful to the organizers and to the participants for a wonderful workshop.

Ahmad Biniaz is upported by NSERC Postdoctoral Fellowship. Prosenjit Bose is supported by NSERC. Jean Cardinal is supported by the F.R.S.-FNRS under grant CDR J.0146.18. Michael S. Payne is supported by a Discovery Early Career Researcher Award funded by the Australian Government.

References

[1] E. Ackerman, B. Keszegh, and D. Pálvölgyi. Coloring hypergraphs defined by stabbed pseudo-disks and ababfree hypergraphs. 2019. `http://arxiv.org/abs/1902.08468`.

[2] A. Asinowski, J. Cardinal, N. Cohen, S. Collette, T. Hackl, M. Hoffmann, K. B. Knauer, S. Langerman, M. Lason, P. Micek, G. Rote, and T. Ueckerdt. Coloring hypergraphs induced by dynamic point sets and bottomless rectangles. In *Proceedings of the 13th International Symposium on Algorithms and Data Structures* (*WADS*), pages 73–84, 2013.

[3] J. Cardinal, K. B. Knauer, P. Micek, and T. Ueckerdt. Making triangles colorful. *Journal of Computational Geometry*, 4(1):240–246, 2013.

[4] J. Cardinal, K. B. Knauer, P. Micek, and T. Ueckerdt. Making octants colorful and related covering decomposition problems. *SIAM Journal on Discrete Mathematics*, 28(4):1948–1959, 2014.

[5] J. Cardinal and M. Korman. Coloring planar homothets and three-dimensional hypergraphs. *Computational Geometry: Theory and Applications*, 46(9):1027–1035, 2013.

[6] X. Chen, J. Pach, M. Szegedy, and G. Tardos. Delaunay graphs of point sets in the plane with respect to axis-parallel rectangles. *Random Structures Algorithms*, 34(1):11–23, 2009.

[7] H. Grötzsch. Zur Theorie der diskreten Gebilde, VII: Ein Dreifarbensatz für dreikreisfreie Netze auf der Kugel. *Wiss. Z. Martin-Luther-Univ. Halle-Wittenberg. Math.-Nat. Reihe*, 8:109–120, 1959.

[8] B. Keszegh and D. Pálvölgyi. Octants are cover-decomposable. *Discrete & Computational Geometry*, 47(3):598–609, 2012.

[9] B. Keszegh and D. Pálvölgyi. Octants are cover-decomposable into many coverings. *Computational Geometry: Theory and Applications*, 47(5):585–588, 2014.

[10] B. Keszegh and D. Pálvölgyi. More on decomposing coverings by octants. *Journal of Computational Geometry*, 6(1):300–315, 2015.

[11] B. Keszegh and D. Pálvölgyi. Proper coloring of geometric hypergraphs. In *Proceedings of the 33rd International Symposium on Computational Geometry* (*SoCG*), pages 47:1–47:15. 2017.

[12] P. Ossona de Mendez and P. Rosenstiehl. Homomorphism and dimension. *Combinatorics, Probability and Computing*, 14(5-6):861–872, 2005.

[13] W. Schnyder. Planar graphs and poset dimension. *Order*, 5(4):323–343, 1989.

[14] W. T. Trotter and R. Wang. Incidence posets and cover graphs. *Order*, 31(2):279–287, 2014.