Large Monochromatic Components in Two-colored Grids

Jiří Matoušek
Department of Applied Mathematics and
Institute of Theoretical Computer Science (ITI)
Charles University
Malostranské nám. 25, 118 00 Praha 1
Czech Republic
matousek@kam.mff.cuni.cz

Aleš Přívětivý *
Department of Applied Mathematics
Charles University
Malostranské nám. 25, 118 00 Praha 1
Czech Republic
privetivy@kam.mff.cuni.cz

Rev. 3/X/2007 AP

Abstract

Let D_n^d denote the d-dimensional *grid with diagonals*, that is, the graph with vertex set $\{1, 2, \ldots, n\}^d$ and with edges connecting every two vertices that differ by at most 1 in every coordinate. We prove that for an arbitrary coloring of the vertices of D_n^d by two colors there exists a monochromatic connected subgraph with at least $n^{d-1} - d^2 n^{d-2}$ vertices, and thus the "horizontal layer" coloring (by the parity of the first coordinate) is almost optimal.

We also consider a d-dimensional *triangulated grid*; this is the graph (1-skeleton) of a triangulation of the solid cube $[1, n]^d$ that refines the subdivision of $[1, n]^d$ into the grid of unit cubes. Here every 2-coloring has a monochromatic connected subgraph with $\Omega(n^{d-1}/\sqrt{d})$ vertices.

These results are proved by combining combinatorial and topological arguments with suitable isoperimetric inequalities, and they can be viewed as d-dimensional generalizations of the planar HEX lemma.

1 Introduction

The HEX lemma. Let T_n^2 be the graph of an $n \times n$ triangulated square grid in the plane as in Fig. 1. Here n counts the number of rows and columns of vertices; so the picture has $n = 6$. The well-known planar HEX lemma asserts that if the vertices of T_n^2 are colored red and blue in an arbitrary way, then there exists a path in T_n^2 connecting the top and bottom sides and having all vertices red, or a path connecting the left and right sides and having all vertices blue. In particular, under any two-coloring of the vertices of T_n^2 there is a monochromatic connected subgraph with at least n vertices.

Triangulated grids. Our results can be regarded as d-dimensional generalizations of the HEX lemma. First we consider a d-dimensional *triangulated grid* T_n^d, which is defined geometrically as follows. We begin with the d-dimensional solid cube $[1, n]^d$ and we subdivide it into the grid of unit cubes; each unit cube is of the form $[i_1, i_1+1] \times [i_2, i_2+1] \times \cdots \times [i_d, i_d+1]$,

*The second author's research is supported by grant from Czech Science Foundation GD201/05/H014.

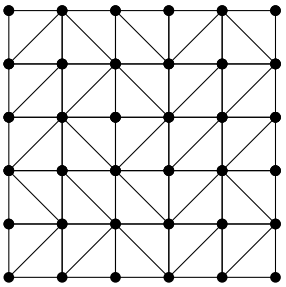

Figure 1: A triangulated square grid.

$1 \le i_1, i_2, \ldots, i_d \le n-1$. Then we triangulate each of the unit cubes in such a way that the simplices of all of these triangulations taken together form a triangulation of $[1,n]^d$ (here a triangulation is a simplicial complex in the sense of algebraic topology, i.e. a face-to-face triangulation with no additional vertices; see Section 2). Then T_n^d is the graph (or 1-skeleton) of such a triangulation; that is, the vertex set is $[n]^d$ (where we use the shorthand $[n]$ for the set $\{1,2,\ldots,n\}$), and the edges of T_n^d are the edges of the triangulation. Thus, for given n and d, T_n^d is not defined uniquely, but rather it stands for an arbitrary graph from a (finite) family.

We consider an arbitrary coloring of $[n]^d$ by two colors and we ask, what can be said about the number of vertices of the largest monochromatic connected subgraph of T_n^d? The obvious "horizontal layer" coloring, by the parity of the first coordinate, shows an upper bound of n^{d-1} for this quantity in the worst case, and one might suspect this coloring to be optimal. However, it turns out that for $d \ge 3$, there are better colorings at least for *some* of the possible triangulated grids. Namely, let us define the *ith diagonal layer* $L_n^d(i)$ by $L_n^d(i) = \{x \in [n]^d : \sum_{j=1}^d x_j = i\}$, and the *diagonal-layer coloring* by coloring all $L_n^d(i)$ with even i red and all $L_n^d(i)$ with odd i blue. As we will verify in Section 7, there are triangulated grids T_n^d with no edges connecting $L_n^d(i)$ and $L_n^d(i')$ with $|i-i'| \ge 2$, and for these, the largest monochromatic connected component is the largest diagonal layer. For $d=3$, for example, $\max_i |L_n^d(i)|$ is approximately $\frac{3}{4}n^2$.

We prove that the diagonal-layer coloring is not far from optimal in the worst case. To state the result precisely, let us define, for $\alpha \in [0,1]$ and given n and d,

$$i_\alpha = \min\left\{ i : \sum_{j \le i} |L_n^d(j)| > \alpha n^d \right\}.$$

It can be checked that $i_{1/2} = \left\lceil \frac{(n+1)d}{2} \right\rceil$ and that $L_n^d(i_{1/2})$ is either the single largest diagonal layer or one of the two largest diagonal layers; see Lemma 5.3. We prove the following.

Theorem 1.1 *For an arbitrary 2-coloring of the vertices of a d-dimensional triangulated grid T_n^d, $n \ge 3$, there exists a monochromatic connected subgraph with at least $|L_n^d(i_{2/3})| - |L_n^{d-1}(i_{1/2})|$ vertices. The last quantity is of order $\frac{n^{d-1}}{\sqrt{d}}$.*

We conjecture that the right answer is $|L_n^d(i_{1/2})|$ instead of $|L_n^d(i_{2/3})| - |L_n^{d-1}(i_{1/2})|$; that is, the diagonal-layer coloring is optimal in the worst case. This would follow from a natural and very plausible-looking conjecture stated in Section 8. The numerical values of $|L_n^d(i_{1/2})|$

and $|L_n^d(i_{2/3})| - |L_n^{d-1}(i_{1/2})|$ are not too far from each other, as is illustrated by the following table (showing approximate values for small fixed d and $n \to \infty$):

	$\lvert L_n^d(i_{1/2})\rvert$	$\lvert L_n^d(i_{2/3})\rvert - \lvert L_n^{d-1}(i_{1/2})\rvert$
$d = 3$	$0.750n^2$	$0.698n^2$
$d = 4$	$0.667n^3$	$0.609n^3$
$d = 5$	$0.599n^4$	$0.549n^4$
$d = 10$	$0.430n^9$	$0.393n^9$
$d = 20$	$0.307n^{19}$	$0.280n^{19}$

Let us remark that the proof of Theorem 1.1 works for graphs G somewhat more general than T_n^d. However, since we don't see any good use of such greater generality at the moment, we stick to the concrete geometric formulation as above. If needed, the properties of the graph actually required can be reconstructed from the proof.

Grid with all diagonals. Next, we consider the d-dimensional grid with *all* diagonals of the unit cubes added. That is, we define the graph D_n^d, the d-dimensional *grid with diagonals*, as the graph with vertex set $[n]^d$ and edge set $\{\{u, v\} : u, v \in [n]^d, \|u - v\|_\infty \leq 1\}$, where $\|u - v\|_\infty = \max_i |u_i - v_i|$. In this case we show that the horizontal-layer coloring is almost optimal (the remaining lower-order term is probably an artifact of the proof method):

Theorem 1.2 *For an arbitrary 2-coloring of the vertices of the grid with diagonals D_n^d, there exists a monochromatic connected subgraph with at least $n^{d-1} - d^2 n^{d-2}$ vertices.*

Related work. There is a different and well-known d-dimensional generalization of the planar HEX lemma appearing e.g. in Gale [G79] (also see Linial and Saks [LS93] for a different proof and an application in computer science). It asserts that if the vertices of T_n^d are colored by d colors, then there is a monochromatic path connecting two opposite facets of the cube $[1, n]^d$. In particular, there is a monochromatic connected component with at least n vertices.

The problem considered in this paper fits in the following general context. For an arbitrary graph G and an integer k, let us define $\mathrm{mcc}_k(G)$ as the smallest m such that there exists a coloring of the vertices of G by k colors with no monochromatic connected subgraph having more than m vertices. (The letters mcc abbreviate "monochromatic connected component". The usual chromatic number $\chi(G)$ equals $\min\{k : \mathrm{mcc}_k(G) = 1\}$.)

The quantity $\mathrm{mcc}_k(G)$ was investigated by Kleinberg, Motwani, Raghavan, and Venkatasubramanian [KMR+97] in the context of data storage management. They call a coloring witnessing $\mathrm{mcc}_k(G) \leq c$ a (k, c)-*fragmented coloring*. They proved, among others, that mcc_3 can be arbitrarily large for planar graphs (we note that $\mathrm{mcc}_4 = 1$ for planar graphs by the Four-Color Theorem), and that for a suitable absolute constant $\varepsilon > 0$ and for every Δ there is an infinite family $\{G_n\}$ of graphs of maximum degree Δ with $\mathrm{mcc}_{\varepsilon\sqrt{\Delta}}(G_n) > \varepsilon|V(G_n)|$.

Apparently independent of the work of Kleinberg et al., $\mathrm{mcc}_k(G)$ has been studied extensively by graph-theorists for graphs of bounded degree, mainly for $k = 2$. It is easy to see that any graph G of maximum degree 3 has $\mathrm{mcc}_2(G) \leq 2$. Alon, Ding, Oporowski, and Vertigan [ADO+03] proved that every graph G of maximum degree 4 satisfies $\mathrm{mcc}_2(G) \leq 57$. Haxell, Szabó, and Tardos [HST03] improved this to $\mathrm{mcc}_2(G) \leq 6$ and proved that $\mathrm{mcc}_2(G) \leq 20000$

for every graph G of maximum degree 5. On the other hand, Alon et al. [ADO+03] constructed 6-regular graphs G with $\mathrm{mcc}_2(G)$ arbitrarily large. For graphs G of maximum degree 3 it was also shown in [BS05] that they admit two-coloring where one color induces an independent set, while the other color induces components of size at most 189. Earlier work on this subject [DOS+96], [JW96] mainly focused on more specific questions concerning line graphs of 3-regular graphs. These investigations culminated in [Tho99] showing that the edges of every 3-regular graph can be 2-colored so that each monochromatic component is a path of length at most 5. Recently Linial, Matoušek, Sheffet, and Tardos [LMST07] investigated the order of magnitude of mcc_k for several graph classes; for example, they showed that $\mathrm{mcc}_2(G) = O(n^{2/3})$ for every planar graph on n vertices, and this bound is tight for some planar graphs.

Outline of the proofs of Theorems 1.1 and 1.2. To prove Theorem 1.1, we consider a red-blue coloring of $[n]^d$. Assuming that no monochromatic connected component in T_n^d is very large, we prove that T_n^d has a connected monochromatic *separator*, where a separator in a graph G is a subset $S \subseteq V(G)$ whose removal disconnects G into components of size at most $\frac{1}{2}|V(G)|$. The only property of T_n^d that we use for this part is that it is the graph of a simply connected simplicial complex. Theorem 1.1 then follows from a known result about the vertex expansion of the ordinary grid graph (not triangulated). Theorem 1.2 follows along similar lines: we verify that D_n^d is also the graph of a simply connected simplicial complex (which is not obvious here, unlike in the case of T_n^d), and then we apply a bound on the vertex expansion of D_n^d. We derive this bound from an edge-isoperimetric inequality for the ordinary grid.

2 Topological preliminaries

Here we present some material from elementary algebraic topology. We briefly recall even very standard things in order to make the paper more accessible; see, e.g., [Hat01] or [Mun84] for more details and background.

Simplicial complexes. A (geometric) *simplex* σ is the convex hull of a finite affinely independent set A in some $\mathbf{R}^d$. The points of A are the *vertices* of σ. The *dimension* of σ is $\dim \sigma := |A| - 1$. Thus a *k-simplex* (k-dimensional simplex) has $k+1$ vertices. The convex hull of an arbitrary subset of vertices of a simplex σ is a *face* of σ (every face is itself a simplex).

A nonempty family K of simplices is a (geometric) *simplicial complex* if the following two conditions hold:

1. Each face of any simplex $\sigma \in K$ is also a simplex of K.
2. The intersection $\sigma_1 \cap \sigma_2$ of any two simplices $\sigma_1, \sigma_2 \in K$ is a face of both σ_1 and σ_2.

In this paper we consider only simplicial complexes with finitely many simplices. A *subcomplex* of a simplicial complex K is a subset L of the simplices of K that constitutes a simplicial complex.

The union of all simplices in a simplicial complex K is the *polyhedron* of K and it is denoted by $\|K\|$. In this situation K is called a *triangulation* of the topological space $\|K\|$.

The *vertex set* of K, denoted by $V(K)$, is the union of the vertex sets of all simplices of K. The *graph* (or 1-*skeleton*) of K, denoted by $G(K)$, has vertex set $V(K)$ and two vertices are connected by an edge if they are contained in a common 1-simplex (edge) of K.

Simplicial maps and simplicial approximation. Let K and L be simplicial complexes. A *simplicial map* of K into L is a continuous map $f\colon \|K\| \to \|L\|$ such that the image of every simplex $\sigma \in K$ is a simplex of L (and in particular, every vertex of K is mapped to a vertex of L), and moreover, f restricted to each simplex $\sigma \in K$ is an affine map.

Let K and $\tilde{K}$ be simplicial complexes. We call $\tilde{K}$ a *refinement* of K if $\|\tilde{K}\| = \|K\|$ and every simplex of $\tilde{K}$ is contained in some simplex of K (and thus for every nonempty simplex σ of K there is a subcomplex of $\tilde{K}$ that is a triangulation of σ).

We will need the following proposition, which is a variation of the standard *simplicial approximation theorem*.

Proposition 2.1 *Let K and L be simplicial complexes and let $f\colon \|K\| \to \|L\|$ be an arbitrary continuous map. Then there is a refinement $\tilde{K}$ of K and a simplicial map $\tilde{f}$ of $\tilde{K}$ into L such that*

(∗) *If $\sigma \in K$ is a simplex such that $f(\sigma)$ is completely contained in some simplex $\tau \in L$, then $\tilde{f}$ maps all vertices of $\tilde{K}$ lying in σ to vertices of τ. In particular, if $f(v) \in V(L)$ for some $v \in V(K)$, then $\tilde{f}(v) = f(v)$.*

The usual simplicial approximation theorem also yields $\tilde{f}$ homotopic to f (but we don't need this part). On the other hand, condition (∗) doesn't appear in the formulations of the simplicial approximation theorem in the literature, but it immediately follows from the standard proof; see, for example, [Hat01] or [Mun84].

A variant of the HEX lemma. Let T be a 2-dimensional simplicial complex with $\|T\|$ homeomorphic to B^2, the unit disk in the plane, and let C_T be the cycle in the graph $G(T)$ corresponding to the boundary of B^2. (So $G(T)$ can be drawn in the plane with the outer face bounded by C_T and with all inner faces being triangles.) We need the following variant of the HEX lemma:

Lemma 2.2 *Let T and C_T be as above, and let the vertices of T be colored red and blue. Let the vertices of C_T be partitioned into four sets, forming consecutive segments along C_T (see Fig. 2), as follows: a set A of red and blue vertices, a single red vertex u, a set B of red and blue vertices, and a single red vertex v. If there is no blue path (path consisting of blue vertices) from A to B, then u and v are connected by a red path. See Fig. 2*

Proof. The case where C_T has four vertices (and thus A and B are singletons) and A and B are blue is Claim 6.1.4 in [MN98]. The case of A, B arbitrary reduces to the previous case by adding two new vertices a and b in the outer face (Fig. 3), coloring them blue, connecting a to all vertices of $A \cup \{u, v\}$ and b to all vertices of $B \cup \{u, v\}$. □

3 Monochromatic connected separators

We recall that a topological space X is *simply connected* if each continuous map $f\colon S^1 \to X$ can be extended to a continuous map $\bar{f}\colon B^2 \to X$ (S^1 denotes the unit circle, the boundary

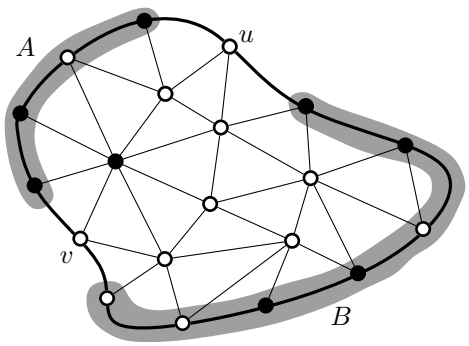

Figure 2: A variant of the HEX lemma

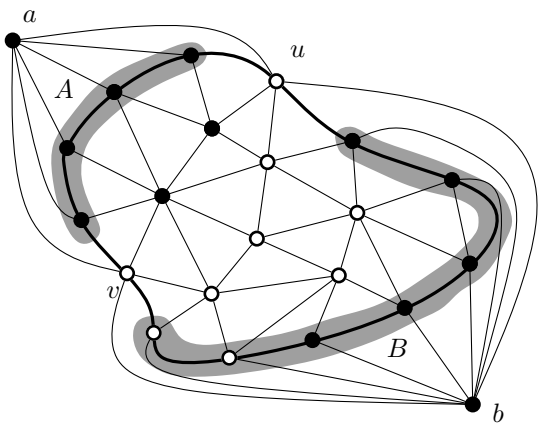

Figure 3: A reduction to the 4-cycle case.

of B^2). A simplicial complex K is simply connected if the topological space $\|K\|$ is simply connected.

Here is the main result of this section.

Proposition 3.1 *Let $G = G(K)$ be the graph of a (finite) simply connected simplicial complex K. For an arbitrary coloring of $V(G)$ by two colors there exists a monochromatic connected separator, that is, a subset $S \subseteq V(G)$ such that all vertices in S have the same color, the subgraph induced by S is connected, and each component of $G \setminus S$ has at most $\frac{1}{2}|V(G)|$ vertices.*

We need the following lemma.

Lemma 3.2 *Let G be a graph as in Proposition 3.1 with $V(G)$ colored red and blue. Let C be a cycle in G, let $u, v \in V(C)$ be two distinct red vertices, and let A and B be the two connected components (arcs) of $C \setminus \{u, v\}$. If there is no blue path from a vertex of A to a vertex of B in G, then u and v are connected by a red path.*

Proof. The goal is to apply Lemma 2.2. Let us take a triangulated disk T with $\|T\|$ homeomorphic to the unit disk and with the boundary cycle C_T such that C_T has the same length as C. Let $f_0 \colon V(C_T) \to V(C)$ be an isomorphism of the two cycles, and let $f \colon \|C_T\| \to \|C\|$ be the (unique) extension of f_0 to a simplicial map. (So here we also regard C_T as a 1-dimensional subcomplex of T and C as a 1-dimensional subcomplex of K.) Since $\|T\|$ is homeomorphic to B^2 and since K is simply connected, f can be extended to a continuous map $\bar{f} \colon \|T\| \to \|K\|$.

Applying our version of the simplicial approximation theorem (Proposition 2.1) to the mapping $\bar{f}$, we obtain a refinement $\tilde{T}$ of T and a simplicial map $\tilde{f}$ of $\tilde{T}$ into K. Let $\tilde{C}_T$ be the subcomplex of $\tilde{T}$ that refines C_T (so $\|\tilde{C}_T\| = \|C_T\|$ is the boundary of the disk $\|T\|$).

According to condition $(*)$ in Proposition 2.1, we have $\tilde{f}(v) = f(v)$ for every vertex v of C_T, and for every edge (1-simplex) e of C_T with vertices u and v, all vertices of $\tilde{C}_T$ lying in e are mapped by $\tilde{f}$ to either $f(u)$ or $f(v)$.

Each vertex of $\tilde{T}$ is mapped to a vertex $\tilde{f}(v) \in V(G)$ since $\tilde{f}$ is a simplicial map, so we can define the color of each $v \in V(\tilde{T})$ as the color of $\tilde{f}(v)$. We note that if $x_1, x_2, \ldots, x_k$ is a monochromatic path in $\tilde{T}$, then $\tilde{f}(x_1), \tilde{f}(x_2), \ldots, \tilde{f}(x_k)$ is a monochromatic walk in G. A monochromatic walk can be shortcut to a monochromatic path.

Now we check that $\tilde{T}$, $\tilde{C}_T$, and the two red vertices $u_T = f^{-1}(u)$ and $v_T = f^{-1}(v)$ satisfy the conditions of Lemma 2.2. Since the two components A and B of $C \setminus \{u, v\}$ are not connected by a blue path in G, it follows that the two components of $\tilde{C}_T \setminus \{u_T, v_T\}$ are not connected by a blue path in $\tilde{T}$ either. Therefore, by Lemma 2.2, u_T and v_T are connected by a red path, and u and v are connected by a red path too, as claimed. □

Proof of Proposition 3.1. If the vertices of G are colored with only one color, we can take $V(G)$ for S and we are done. Otherwise, we choose S as the vertex set of a monochromatic connected subgraph of G that is

- inclusion-maximal (no vertex of the same color can be added), and
- such that the largest component G_1 of $G \setminus S$ has the smallest possible number of vertices.

Without loss of generality, let S be colored blue.

We claim that G_1 has at most $\frac{1}{2}|V(G)|$ vertices. For contradiction, we thus assume $|V(G_1)| > \frac{1}{2}|V(G)|$. We define S' as the union of all inclusion-maximal red connected subgraphs of G_1 that have at least one vertex connected to S (so S' is the "red foam" on the blue surface of S facing G_1). We will show that S' is "better" than S; namely, that S' induces an inclusion-maximal connected red subgraph of G and the largest component of $G \setminus S'$ is smaller than G_1. The set S' is red by definition.

Next, we show that S' induces a connected subgraph (then it will also be clear that it is inclusion-maximal, by the choice of S'). It suffices us to show that any two vertices u and v of S' are connected by a red path in G. Moreover, it suffices us to show this only for u and v that are both connected to S by an edge. Let $u' \in S$ be a blue neighbor of u and let $v' \in S$ be a blue neighbor of v. Since S is a blue connected subgraph, there is a blue path $\mathcal{P}_1$ connecting u' and v'. Since G_1 is connected, u and v are connected by a path $\mathcal{P}_2$ in G_1, which is vertex-disjoint from $\mathcal{P}_1$. The path $\mathcal{P}_2$ may consist of vertices of both colors, but there is no blue path from a vertex of $\mathcal{P}_2$ to $\mathcal{P}_1$ (by the inclusion-maximality of S). The paths $\mathcal{P}_1$ and $\mathcal{P}_2$ together form a cycle C in G. The graph G and the cycle C satisfy all conditions of Lemma 3.2, and so we have a red path connecting u and v. Hence S' indeed induces a connected subgraph.

Each component of $G \setminus S'$ is contained either in $G_1 \setminus S'$, or in $G \setminus G_1$. We have $|V(G_1) \setminus S'| < |V(G_1)|$ since $S' \neq \emptyset$, and $|V(G) \setminus V(G_1)| < \frac{1}{2}|V(G)|$. Thus the largest component of $G \setminus S'$ is strictly smaller than G_1, and this contradicts the choice of S. □

4 D_n^d as the graph of a simply connected complex

In order to apply Proposition 3.1 to some graph G, we need to know that G is the graph of a simply connected simplicial complex. This is obviously true for T_n^d, a graph of a triangulated grid, but for D_n^d, the grid with all diagonals, this has to be established. (One could say that this is a somewhat artificial-looking way of proving a purely combinatorial statement about D_n^d, but it is a convenient one.)

Lemma 4.1 *For all $n, d \geq 2$ there exists a simply connected (2-dimensional) simplicial complex K such that D_n^d is (isomorphic to) the graph of K.*

Let us note that the simplicial complexes K constructed in the proofs below cannot be embedded in $\mathbf{R}^d$ and they are probably hard to visualize.

First proof. (This proof was simplified using a suggestion of a referee.) First let us fix a triangulation T of the grid, i.e. a geometric simplicial complex refining the subdivision of $[1, n]^d$ into unit cubes as in the definition of the graphs T_n^d. Let K_0 be the 2-skeleton of it (consisting of all simplices of this T_n^d of dimension at most 2). This K_0 is simply connected since T is and the fundamental group of a simplicial complex is determined by the 2-skeleton.

Let $E = E(D_n^d) \setminus E(K_0)$ be the set of all edges of D_n^d not present in K_0, and let $e_1, e_2, \ldots, e_t$ be an enumeration of E such that shorter edges precede longer ones; that is, $i < j$ implies $\|e_i\|_1 \leq \|e_j\|_1$, where we write $\|e\|_1 = \sum_{k=1}^{d} |u_k - v_k|$ for $e = \{u, v\}$.

We are going to define a sequence $K_0, K_1, \ldots, K_t$ of simplicial complexes inductively, and $K = K_t$ will be as in the lemma. The simplicial complex K_i is obtained from K_{i-1} by adding the edge e_i plus a triangle with e_i as one of the edges, while the remaining two edges of the triangle are already present in K_{i-1}. Obviously K_i is homotopy equivalent to K_{i-1} and hence

by induction we get that each K_i is simply connected. We also have $E(K_t) = E(D_n^d)$. It remains to check that the triangles can indeed be added as described.

Let us write $e_i = \{u, v\}$, and let $a = \min\{k \in [d] : u_k \neq v_k\}$. Let $x \in [n]^d$ be the vertex

$$x = (u_1, \dots, u_a, v_{a+1}, \dots, v_d).$$

Since all edges $e \in E(D_n^d)$ with $\|e\|_1 = 1$ are already present in K_0, we have $\|e_i\|_1 \geq 2$ and $u \neq x \neq v$.

The edges $\{u, x\}$ and $\{x, v\}$ belong to $E(D_n^d)$ and have length strictly smaller than $\|e_i\|$, and hence they are present in K_{i-1}. The triangle added to K_{i-1} to form K_i is $\{u, x, v\}$. □

Second proof (suggested by a referee). For each $u \in [n]^d$ let C_u be the axis-parallel unit cube centered at u, and let K be the nerve of the set system $\{C_u : u \in [n]^d\}$ (that is, a set $I \subseteq [n]^d$ forms a simplex if $\bigcap_{u \in I} C_u \neq \emptyset$). The 1-skeleton of K is exactly D_n^d. By the nerve theorem, K is homotopy equivalent to $\bigcup_{u \in [n]^d} C_u = [\frac{1}{2}, n + \frac{1}{2}]^d$, which is simply connected. □

5 Isoperimetric inequalities

For a graph G and a set $A \subseteq V(G)$, the *vertex boundary* of A in G is the set

$$\text{vert-bd}_G(A) = \{v \in V(G) : v \notin A, \{u, v\} \in E(G) \text{ for some } u \in A\}.$$

An *isoperimetric inequality* for G (or more precisely, a vertex-isoperimetric inequality) bounds from below the quantity $\min\{|\text{vert-bd}_G(A)| : A \subseteq V(G), |A| = s\}$ as a function of s.

The ordinary grid. Let G_n^d be the graph of the ordinary grid (not triangulated). Explicitly, $V(G_n^d) = [n]^d$ and $u, v \in V(G_n^d)$ are connected by an edge if $\|u - v\|_1 = \sum_{i=1}^d |u_i - v_i| = 1$.

A vertex-isoperimetric inequality for G_n^d was established by Bollobás and Leader [BL91a]. For stating it, we define the *simplicial order* on $[n]^d$ by setting $x <_{\text{simpl}} y$ if either $\sum_{i=1}^d x_i < \sum_{i=1}^d y_i$, or if $\sum_{i=1}^d x_i = \sum_{i=1}^d y_i$ and x precedes y lexicographically; that is, for some $i \in [d]$ we have $x_j = y_j$ for $j = 1, 2, \dots, i-1$ and $x_i < y_i$.

Theorem 5.1 ([BL91a]) *For every $A \subseteq [n^d]$ we have $|\text{vert-bd}_{G_n^d}(A)| \geq |\text{vert-bd}_{G_n^d}(I_{\text{simpl}}(|A|))|$, where $I_{\text{simpl}}(s)$ denotes the set of the first s elements of $[n]^d$ in the ordering $<_{\text{simpl}}$.*

To bound from below the size of the vertex boundary of arbitrary set $A \subseteq [n]^d$ we use the following lemma:

Lemma 5.2 *Let $A \subseteq [n]^d$ with $|A| = \beta n^d$. Then we have $|\text{vert-bd}_{G_n^d}(A)| \geq |L_n^d(i_\beta)| - |L_n^{d-1}(i_{1/2})|$.*

Proof. By Theorem 5.1 it suffices to prove the statement only for A of the form $I_{\text{simpl}}(|A|)$. Then A is a subset of the union of the first i_β diagonal layers and the first $i_\beta - 1$ layers are fully contained in A.

Let us write D for the intersection of A with $L_n^d(i_\beta)$. If $D = \emptyset$, then $\text{vert-bd}_{G_n^d}(A) = L_n^d(i_\beta)$ and we are done. If $D \neq \emptyset$, then $\text{vert-bd}_{G_n^d}(A)$ is a disjoint union of two sets B_1 and B_2, where $B_1 \subseteq L_n^d(i_\beta)$ and $B_2 \subseteq L_n^d(i_\beta + 1)$. It is easy to check that $B_1 = L_n^d(i_\beta) \setminus D$ and if we show that $|B_2| \geq |D| - |L_n^{d-1}(i_{1/2})|$, we are done.

Let D' be D shifted by 1 in the first component; that is, $D' = D + (1, 0, \ldots, 0)$. Similarly, let L' be $L_n^d(i_\beta) + (1, 0, \ldots, 0)$. Since $D' \subseteq L'$, we have $D' \setminus [n]^d \subseteq L' \setminus [n]^d$ and from the fact $D' \cap [n]^d \subseteq B_2$ we arrive at $|B_2| \geq |D| - |L' \setminus [n]^d|$.

The rest follows from the fact that $L' \setminus [n]^d$ is exactly a copy of the set $L_n^{d-1}(i_\beta - n)$, which has always smaller size than the middle layer $L_n^{d-1}(i_{1/2})$ as will be shown in Lemma 5.3. □

Let us remark that some error term in the previous estimate of vert-$\mathrm{bd}_{G_n^d}(A)$ is needed. For example, vert-$\mathrm{bd}_{G_n^d}(I_{\mathrm{simpl}}(21)) = 5$ for the case $n = 3$, $d = 3$, but $|L_3^3(i_{21/27})| = 6$.

Lemma 5.3 *For any $n, d \in \mathbf{N}$ the following hold:*

(a) $\left|L_n^d\left(\frac{(n+1)d}{2} - i\right)\right| = \left|L_n^d\left(\frac{(n+1)d}{2} + i\right)\right|$ *for every $i = 0, 1, 2, \ldots,$ and therefore, $i_{1/2} = \left\lceil \frac{(n+1)d}{2} \right\rceil$.*

(b) The function $i \mapsto |L_n^d(i)|$ is non-decreasing on integers $i \leq \left\lceil \frac{(n+1)d}{2} \right\rceil = i_{1/2}$ and non-increasing on integers $i \geq \left\lfloor \frac{(n+1)d}{2} \right\rfloor$.

Proof.

(a) The function $f\big((x_1, \ldots, x_d)\big) = (n + 1 - x_1, \ldots, n + 1 - x_d)$ maps $L_n^d\left(\frac{(n+1)d}{2} - i\right)$ to $L_n^d\left(\frac{(n+1)d}{2} + i\right)$ and it is easy to check from definition that $i_{1/2} = \left\lceil \frac{(n+1)d}{2} \right\rceil$.

(b) The proof is by induction on d. For $d = 1, 2$ and any $n \geq 1$ the statement trivially holds. Fix some n and $d > 2$. We can partition the set $L_n^d(i)$ into n sets $L_1, \ldots, L_n$ according to the last component of the vertices; that is

$$L_t = \{x \in L_n^d(i) : x_d = t\} \text{ for all } t \in [n].$$

It is easy to check that each L_t is isomorphic to $L_n^{d-1}(i - t)$ and thus

$$\begin{aligned} \big|L_n^d(i+1)\big| - \big|L_n^d(i)\big| &= \sum_{t=1}^{n} \big|L_n^{d-1}(i+1-t)\big| - \sum_{t=1}^{n} \big|L_n^{d-1}(i-t)\big| \\ &= \big|L_n^{d-1}(i)\big| - \big|L_n^{d-1}(i-n)\big|. \end{aligned}$$

We distinguish four cases:

1. If $i \leq \left\lceil \frac{(n+1)(d-1)}{2} \right\rceil$ then by the inductive hypothesis we have $\big|L_n^{d-1}(i-n)\big| \leq \big|L_n^{d-1}(i)\big|$.
2. If $\left\lceil \frac{(n+1)(d-1)}{2} \right\rceil < i < \left\lceil \frac{(n+1)d}{2} \right\rceil$ then by (a) we have

   $$\big|L_n^{d-1}(i)\big| = \big|L_n^{d-1}\big((n+1)(d-1) - i\big)\big|$$

   and since

   $$(i - n) \leq (n+1)(d-1) - i \leq \left\lceil \frac{(n+1)(d-1)}{2} \right\rceil$$

   we have by the inductive hypothesis

   $$\big|L_n^{d-1}(i-n)\big| \leq \big|L_n^{d-1}\big((n+1)(d-1) - i\big)\big| = \big|L_n^{d-1}(i)\big|.$$

3. If $\lfloor \frac{(n+1)d}{2} \rfloor \le i < \lfloor \frac{(n+1)(d+1)}{2} + n \rfloor$ then by (a) we have

$$\big|L_n^{d-1}(i-n)\big| = \big|L_n^{d-1}((n+1)(d-1)-(i-n))\big|$$

and since

$$i \ge (n+1)(d-1)-(i-n) \ge \Big\lfloor \frac{(n+1)(d-1)}{2} \Big\rfloor$$

we have by the inductive hypothesis

$$\big|L_n^{d-1}(i)\big| \le \big|L_n^{d-1}\big((n+1)(d-1)-(i-n)\big)\big| = \big|L_n^{d-1}(i-n)\big|.$$

4. If $i \ge \lfloor \frac{(n+1)(d+1)}{2} + n \rfloor$ then by the inductive hypothesis we have $\big|L_n^{d-1}(i-n)\big| \ge \big|L_n^{d-1}(i)\big|$.

□

Proposition 5.4 *For every $\alpha \in (0, \frac{1}{2})$ and every $A \subseteq [n]^d$ with $\alpha n^d \le |A| \le (1-\alpha)n^d$ we have* $|\text{vert-bd}_{G_n^d}(A)| \ge |L_n^d(i_{1-\alpha})| - |L_n^{d-1}(i_{1/2})|$.

Proof. Let $|A| = \beta n^d$. By Lemma 5.2 we have $|\text{vert-bd}_{G_n^d}(A)| \ge |L_n^d(i_\beta)| - |L_n^{d-1}(i_{1/2})|$. If $\alpha \le \beta \le \frac{1}{2}$, then also $i_\alpha \le i_\beta \le i_{1/2}$ and by Lemma 5.3 we have $|L_n^d(i_\beta)| \ge |L_n^d(i_\alpha)|$. Similarly, for $\frac{1}{2} < \beta \le (1-\alpha)$ we have $|L_n^d(i_\beta)| \ge |L_n^d(i_{1-\alpha})|$. Since $i_{1/2} \le (n+1)d - i_\alpha \le i_{(1-\alpha)}$, by Lemma 5.3 we have $|L_n^d(i_\alpha)| = |L_n^d((n+1)d - i_\alpha)| \ge |L_n^d(i_{(1-\alpha)})|$, and thus $|\text{vert-bd}_{G_n^d}(A)| \ge |L_n^d(i_{(1-\alpha)})| - |L_n^{d-1}(i_{1/2})|$. □

The grid with diagonals. Here we derive the following vertex-isoperimetric inequality:

Proposition 5.5 *For any set $A \subseteq [n]^d$ with $\frac{1}{4}n^d \le |A| \le \frac{3}{4}n^d$ we have* $|\text{vert-bd}_{D_n^d}(A)| \ge n^{d-1} - d^2 n^{d-2}$.

We are going to derive this result from an edge-isoperimetric inequality for the ordinary grid G_n^d. For $A \subseteq [n]^d$, let

$$\text{edge-bd}_{G_n^d}(A) = \{\{u,v\} \in E(G_n^d) : u \in A, v \notin A\}$$

be the edge boundary of A in G_n^d.

Theorem 5.6 (Bollobás and Leader [BL91b]) *Let $A \subseteq [n]^d$. Then*

$$|\text{edge-bd}_{G_n^d}(A)| \ge \begin{cases} 4|A|/n & \text{for } |A| < n^d/4, \\ n^{d-1} & \text{for } n^d/4 \le |A| \le 3n^d/4, \\ 4(n^d - |A|)/n & \text{for } |A| > 3n^d/4. \end{cases}$$

First we show, following [BL91a] almost verbatim, that it suffices to prove Proposition 5.5 for sets A that are *down-sets*. We begin with the necessary definitions from [BL91a].

A set $A \subseteq [n]^d$ is a *down-set* if $x \in A$ implies $y \in A$ for all $y \in [n]^d$ satisfying $y_i \le x_i$ for all $i \in [d]$. For $A \subseteq [n]^d$ and $1 \le i \le d$, let the *i-section* of A at $x \in [n]^{d-1}$ be defined as

$$A^i(x) = \{t \in [n] : (x_1, \ldots, x_{i-1}, t, x_i, \ldots, x_{d-1}) \in A\}.$$

For $A \subseteq [n]^d$ and $1 \le i \le d$, let the *i-compression* of A be defined by giving its i-sections:

$$(C_i(A))^i(x) = \{1, \dots, |A^i(x)|\} \qquad \text{for all } x \in [n]^{d-1}.$$

In other words, C_i "compresses" each i-section of A downwards. We note that $|C_i(A)| = |A|$, and it is also easy to check that if a set A is i-compressed, then so is $C_j(A)$. So the set $A' = C_n(C_{n-1}(\dots C_1(A)\dots))$ is i-compressed for every $i \in [d]$, and thus it is a down-set.

Lemma 5.7 *Let $A \subseteq [n]^d$. Then there is a down-set $A' \subseteq [n]^d$ with $|A'| = |A|$ and $|\text{vert-bd}_{D_n^d}(A')| \le |\text{vert-bd}_{D_n^d}(A)|$.*

Proof. For $A \subseteq [n]^d$, let $N(A) = A \cup \text{vert-bd}_{D_n^d}(A)$ denote the *closed neighborhood* of A. By the above remarks it suffices to prove that $|N(C_i(A))| \le |N(A)|$ for all $A \subseteq [n]^d$ and all $i \in [d]$.

Let us write B for $C_i(A)$. The case $d = 1$ is easy; we have $|N(A)| \ge |A| + 1 = |N(B)|$ for all $A \subseteq [n]$ apart from $A = \emptyset$ and $A = [n]$. For $A = \emptyset$ and $A = [n]$ we have $|N(A)| = |A| = |B| = |N(B)|$.

For $d > 1$ it is sufficient to show that for each $x \in [n]^{d-1}$ we have

$$|(N(B))^i(x)| \le |(N(A))^i(x)|.$$

Fixing an arbitrary $x \in [n]^{d-1}$,

$$(N(A))^i(x) = \bigcup_{y \in [n]^{d-1} : \|y-x\|_\infty \le 1} N\left(A^i(y)\right)$$

and

$$(N(B))^i(x) = \bigcup_{y \in [n]^{d-1} : \|y-x\|_\infty \le 1} N\left(B^i(y)\right).$$

The sets $N(B^i(y))$ are initial segments of $[n]$, hence they are nested, and thus

$$|(N(B))^i(x)| \le \max_{y \in [n]^{d-1} : \|y-x\|_\infty \le 1} |N\left(B^i(y)\right)|.$$

From the one-dimensional case we know

$$|N\left(B^i(y)\right)| \le |N\left(A^i(y)\right)|,$$

and therefore,

$$|(N(B))^i(x)| \le |(N(A))^i(x)|.$$

□

Proposition 5.5 now follows from Theorem 5.6, the previous lemma, and the next one:

Lemma 5.8 *Let $A \subseteq [n]^d$ be a down-set. Then $|\text{vert-bd}_{D_n^d}(A)| \ge |\text{edge-bd}_{G_n^d}(A)| - d^2 n^{d-2}$.*

To prove this lemma we need the following claim:

Claim 5.9 *Let $A \subseteq [n]^d$ be a down-set. Then $|\text{edge-bd}_{G_n^d}(A)| \le dn^{d-1}$.*

Proof. Since A is a down-set, for each $k \in [d]$ and $x \in [n]^{d-1}$ there is at most one $t \in [n-1]$ such that $\{u, v\} \in \text{edge-bd}_{G_n^d}(A)$, where $u = (x_1, \dots, x_{k-1}, t, x_{k+1}, \dots, x_d)$ and $v = (x_1, \dots, x_{k-1}, t+1, x_{k+1}, \dots, x_d)$. Hence there can be at most dn^{d-1} edges in $\text{edge-bd}_{G_n^d}(A)$. □

Proof of Lemma 5.8. Let

$$E = \text{edge-bd}_{G_n^d}(A) \setminus \{\{u, v\} \in E(G_n^d) : u_i = v_i = n \text{ for some } i\}.$$

We show that E can be injectively mapped into $\text{vert-bd}_{D_n^d}(A)$. Let us construct a mapping $f\colon E \to [n]^d$ in the following way: For arbitrary edge $\{u, v\} \in E$, let i be the number of the component in which u and v differ by 1, and without loss of generality, let $u_i \le v_i$. Then we put

$$f(\{u, v\}) = u + (\underbrace{1, 1, \dots, 1}_{i \text{ times } 1}, 0, \dots, 0).$$

It is easy to check that $u_j < n$ for $j \le i$, and hence $z = f(\{u, v\}) \in [n]^d$ indeed. We have $u \in A$ and $v \notin A$. If we had $z \in A$, then since A is a down-set and $v_i \le z_i$ for all $i \in [d]$, it would follow that $v \in A$ as well. Thus $z \notin A$, and since $\|u - z\|_\infty = 1$, we arrive at $f(\{u, v\}) \in \text{vert-bd}_{D_n^d}(A)$.

Now we suppose for contradiction that there exist two edges $\{u_1, v_1\}$, $\{u_2, v_2\} \in E$ that are mapped by f to the same vertex z. Let i_1 be the number of the component where u_1 and v_1 differ by 1 and let i_2 be the number of the component where u_2 and v_2 do so. Without loss of generality let $i_1 < i_2$. Then

$$v_2 = (z_1 - 1, \dots, z_{i_2 - 1} - 1, z_{i_2}, z_{i_2+1}, \dots, z_d)$$

is in all components less or equal to

$$u_1 = (z_1 - 1, \dots, z_{i_1 - 1} - 1, z_{i_1} - 1, z_{i_1+1}, \dots, z_d),$$

and since A is a down-set, $u_1 \in A$ implies $v_2 \in A$. This contradiction shows that f is injective as claimed.

It remains to estimate the size of E. For $i \in [d]$ let us put $F_i = \{\{u, v\} \in \text{edge-bd}_{G_n^d}(A) : u_i = v_i = n\}$. It is easy to see that

$$|F_i| = |\text{edge-bd}_{G_n^{d-1}}(A_i)|,$$

where

$$A_i = \{(x_1, \dots, x_{d-1}) \in [n]^{d-1} : (x_1, \dots, x_{i-1}, n, x_i, \dots, x_{d-1}) \in A\}.$$

Since A_i is a down-set, Claim 5.9 gives $|F_i| \le (d-1)n^{d-2}$, and thus

$$|E| \ge |\text{edge-bd}_{G_n^d}(A)| - \sum_{i=1}^{d} |F_i| \ge |\text{edge-bd}_{G_n^d}(A)| - d(d-1)n^{d-2},$$

which is even slightly better than claimed in the lemma. □

6 Proofs of Theorems 1.1 and 1.2

Monochromatic connected separators again. We will need the following standard consequence of Proposition 3.1.

Corollary 6.1 *Let G be the graph of a simply connected simplicial complex, and let $V(G)$ be colored red and blue. Then there exists a partition of $V(G)$ into three disjoint sets A, B and S such that*

- *no edge connects a vertex of A to a vertex of B,*
- *S is a monochromatic connected subgraph of G, and*
- *$|A|, |B| \le \frac{2}{3}|V(G)|$.*

Proof. Let S be as in Proposition 3.1, and let $V_1, V_2, \ldots, V_m$ be the vertex sets of the components of $G \setminus S$ ordered by size in the descending order. Let i be the largest index with $|V_1 \cup V_2 \cup \ldots \cup V_i| \le \frac{2}{3}|V(G)|$. Then $|V_1| + \cdots + |V_i| \le \frac{2}{3}|V(G)|$ and $|V_{i+2}| + \cdots + |V_m| \le \frac{1}{3}|V(G)|$. Let us put $A = V_1 \cup V_2 \cup \cdots \cup V_i$ and $B = V \setminus A \setminus S$. Obviously, $|A| \le \frac{2}{3}|V(G)|$, and since $|V(G)| - |B| \ge |A| \ge |V_{i+1}|$ and $|B| = |V_{i+1}| + |V_{i+2}| + \cdots + |V_m| \le |V_{i+1}| + \frac{1}{3}|V(G)|$, we arrive at $|B| \le \frac{2}{3}|V(G)|$ too. □

Proof of Theorem 1.1. We consider a two-colored graph T_n^d as in the theorem and a partition $V(T_n^d) = A \cup B \cup S$ as in Corollary 6.1 (S is a monochromatic connected separator). We may assume $|S| \le \frac{1}{3}n^d$, for otherwise, we would be done. Then $|A|$ or $|B|$ is between $\frac{1}{3}n^d$ and $\frac{2}{3}n^d$; let us fix the notation A, B so that $\frac{1}{3}n^d \le |A| \le \frac{2}{3}n^d$.

We now consider the grid graph G_n^d as a subgraph of T_n^d. We have $S \supseteq \text{vert-bd}_{T_n^d}(A) \supseteq \text{vert-bd}_{G_n^d}(A)$, and $|\text{vert-bd}_{G_n^d}(A)| \ge |L_n^d(i_{2/3})| - |L_n^{d-1}(i_{1/2})|$ by Proposition 5.4. The required estimates for $|L_n^d(i_{2/3})|$ and $|L_n^{d-1}(i_{1/2})|$ are established in the appendix. Theorem 1.1 is proved. □

Proof of Theorem 1.2. We consider a 2-coloring of the graph D_n^d. Lemma 4.1 tells us that D_n^d is the graph of a simply connected simplicial complex, and so using Corollary 6.1 we again find a partition $V(D_n^d) = A \cup B \cup S$. We may assume $|S| \le \frac{1}{3}n^d$, for otherwise, we would be done. We can again assume $\frac{1}{3}n^d \le |A| \le \frac{2}{3}n^d$, and using Proposition 5.5 we get $|S| \ge |\text{vert-bd}_{D_n^d}(A)| \ge n^{d-1} - d^2 n^{d-2}$. Theorem 1.2 is proved. □

7 An upper bound for certain triangulated grids

In this section we show that there exists a triangulated grid T_n^d in which every edge connects either two vertices from the same diagonal layer $L_n^d(i)$ or vertices from two consecutive diagonal layers $L_n^d(i)$ and $L_n^d(i+1)$. Thus, each monochromatic connected component in the diagonal-layer coloring of such T_n^d is contained in some $L_n^d(i)$.

To see that such triangulation exists, we consider the arrangement $\mathcal{A}$ of the following hyperplanes: $H_j^i = \{x : x_i = j\}$ for $j \in [n]$, $i = 0, \ldots, d$, and $S_j = \{x : \sum x_k = j\}$ for $j = d, \ldots, nd$. The vertices of this arrangement are of two types: the first type are the vertices formed as the intersection of d hyperplanes among H_j^i, and the second type are vertices formed as the intersection of one hyperplane S_j and $d-1$ of the hyperplanes H_j^i.

The vertices of the first type are exactly the grid points of $[n]^d$. An intersection of $d-1$ hyperplanes H_j^i is either empty or a line parallel to one of the coordinate axes, and if we intersect this line by some hyperplane S_j we get a vertex with integral coordinates. Thus, if we restrict ourselves to the hypercube $[1,n]^d$, the only possible vertices of our arrangement are the grid points of $[n]^d$.

The arrangement $\mathcal{A}$ is a polyhedral complex. Let $\mathcal{C}$ be the subcomplex of $\mathcal{A}$ consisting of all faces contained in $[1,n]^d$. The union of all faces of $\mathcal{C}$ is exactly $[1,n]^d$, $\mathcal{C}$ refines the subdivision of $[1,n]^d$ into unit cubes, and it remains to refine $\mathcal{C}$ to a triangulation. This can be done, for example, using the *bottom-vertex triangulation* (or *canonical triangulation*) of Clarkson [Cla88].

For each face F of $\mathcal{C}$, which is a convex polytope, we define the *bottom vertex* v as the vertex of F with the lexicographically smallest coordinate vector.

The bottom-vertex triangulation of $\mathcal{C}$ is defined inductively. For $k \geq 1$, we assume that all faces of $\mathcal{C}$ dimension at most k have already been triangulated (for $k=1$ this is satisfied automatically, since the 1-dimensional faces are segments). Given a $(k+1)$-dimensional face F, we let v be the bottom vertex of F, and we triangulate F as follows. For each simplex σ from the union of the triangulations of all proper faces of F, we consider the cone over σ with apex v, that is, the simplex $conv(\sigma \cup \{v\})$. Taking all of these simplices plus all simplices in the triangulation of all proper faces of F yields the bottom-vertex triangulation of F.

It is not hard to verify that this indeed defines a simplicial complex refining $\mathcal{C}$, and thus the desired triangulated grid T_n^d.

8 Open problems

We have considered colorings of the vertices of $[n]^d$ by two colors. The case of d colors is essentially solved by the d-dimensional HEX lemma. A natural and, in our opinion, very interesting question is, what happens for colorings by k colors for $3 \leq k \leq d-1$? An obvious conjecture is that the minimum possible size of the largest monochromatic connected subgraph should be about n^{d+1-k} in this case (for d and k fixed and $n \to \infty$), but at present we have nothing nontrivial to say about this question.

A natural bound one would expect in Theorem 1.1 (for a triangulated grid) would be the size of the "middle diagonal layer" $L_n^d(i_{1/2})$. However, our method yields only $|L_n^d(i_{2/3})|$. The improved bound of $|L_n^d(i_{1/2})|$ would follow immediately by our method from the following conjecture: *If $A \subseteq [n]^d$ is such that removing it from the grid G_n^d leaves no connected component larger than $\frac{1}{2}n^d$, then $|A| \geq |L_n^d(i_{1/2})|$.* This doesn't seem to follow from the vertex-isoperimetric inequality for G_n^d, since in principle, for example, we might have a situation where $G_n^d \setminus A$ has three connected components of size about $\frac{1}{3}n^d$ each, and A is a common boundary of all three. While this or similar situations seem unlikely, we currently don't have a proof of the conjecture. We note that Kleitman [Kle86] proved a similar conjecture for the hypercube (that is, for the grid G_2^d), but adapting his method to larger grids seems nontrivial.

In the appendix below, we derive asymptotic bounds for the size of the diagonal layers $|L_n^d(i_{1/2})|$ and $|L_n^d(i_{2/3})|$, using a quantitative version of the central limit theorem. However, the estimated quantity is of a quite basic nature and it might well be that considerably more precise results are known.

Acknowledgments

We would like to thank Nati Linial for discussions that were the starting point of our research reported in this paper, and Imre Leader for useful advice concerning isoperimetric inequalities. We also thank two anonymous referees for many useful comments and suggestions.

References

[ADO+03] N. Alon, G. Ding, B. Oporowski, D. Vertigan. Partitioning into graphs with only small components. *Journal of Combinatorial Theory* (Series B), 87:231–243, 2003.

[BL91a] B. Bollobás and I. Leader. Compressions and Isoperimetric Inequalities. *Journal of Combinatorial Theory* (Series A) 56:47-62, 1991.

[BL91b] B. Bollobás and I. Leader. Edge-isoperimetric inequalities in the grid. *Combinatorica* 11(4):299–314, 1991.

[BS05] R. Berke and T. Szabó. Relaxed two-coloring of cubic graphs. *DMTCS Proceedings of 2005 European Conference on Combinatorics, Graph Theory and Applications*, 2005.

[Cla88] K. L. Clarkson. A randomized algorithm for closest-point queries. *SIAM J. Comput.*, 17:830-847, 1998.

[DOS+96] G. Ding, B. Oporowski, D. Sanders, D. Vertigan, preprint, Louisiana State University, Baton Rouge, 1996.

[Fel43] W. Feller. Generalization of a Probability Limit Theorem of Cramer. *Transactions of the American Mathematical Society* Vol. 54, No. 3, pp. 361-372, 1943.

[G79] D. Gale. The game of hex and the Brouwer fixed-point theorem. *Amer. Math. Monthly* 86:818–827, 1979.

[Hat01] A. Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge 2001. Electronic version available at `http://math.cornell.edu/~hatcher#AT1`.

[HST03] P. Haxell, T. Szabó, G. Tardos. Bounded size components—partitions and transversals. *Journal of Combinatorial Theory* (Series B) 88:281–297, 2003.

[JW96] B. Jackson and N. Wormald. On the linear k-arboricity of cubic graphs. *Discrete Math* 162:293–297, 1996.

[Kle86] D. Kleitman. On a problem of Yuzvinsky on separating the n-cube. *Discrete Math.* 60:207–213, 1986.

[KMR+97] J. Kleinberg, R. Motwani, P. Raghavan, S. Venkatasubramanian. Storage management for evolving databases. *Proc. 38th IEEE Symposium on Foundations of Computer Science*, 1997.

[LMST07] N. Linial, J. Matoušek, O. Sheffet, and G. Tardos. Graph coloring with no large monochromatic components. Submitted, 2007. Extended abstract to appear in *Proc. EUROCOMB 2007*.

[LS93] N. Linial and M. Saks. Low diameter graph decompositions. *Combinatorica* 13:441–454, 1993.

[MN98] J. Matoušek and J. Nešetřil. *Invitation to Discrete Mathematics*. Oxford University Press, Oxford 1998.

[Mun84] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Reading, MA, 1984.

[Tho99] C. Thomassen. Two-colouring the edges of a cubic graph such that each monochromatic component is a path of length at most 5. *Journal of Combinatorial Theory* (Series B) 75:100–109, 1999.

A Appendix: Bounding the size of diagonal layers

Let us recall that $L_n^d(i) = \{x \in [n]^d : \sum_{j=1}^d x_j = i\}$ and $i_\alpha = \min\{i : \sum_{j\le i} |L_n^d(j)| > \alpha n^d\}$. Here we prove the following estimates:

Proposition A.1 *There exist constants $C \ge c > 0$ such that for all $n > 1$ and all d we have*

$$\frac{c}{\sqrt{d}} n^{d-1} \le |L_n^d(i_{2/3})| \le |L_n^d(i_{1/2})| \le \frac{C}{\sqrt{d}} n^{d-1}.$$

We use the following quantitative version of the central limit theorem by Feller [Fel43].

Theorem A.2 *Let $X_1, \dots, X_d$ be independent random variables and suppose that their second moments $\sigma_1^2, \dots, \sigma_d^2$ exist. Let $S_i = X_1 + \dots + X_i$ and $s_i^2 = \sigma_1^2 + \dots + \sigma_i^2$ for $i = 1, \dots, d$, let $F_d(t) = Prob[S_d \le t]$ be the distribution function of S_d and let $\Phi(u)$ be the distribution function of the standard normal distribution. Let there exist a non-increasing sequence $\lambda_1, \dots, \lambda_d$, such that $|X_k| \le \lambda_k s_k$ for $k = 1, \dots, d$.*

If $0 < \lambda_d u < (3 - \sqrt{5})/4$, then

$$1 - F_d(us_d) = e^{-(1/2)u^2 Q_d(u)} \left((1 - \Phi(u)) + \theta \lambda_d e^{-(1/2)u^2} \right).$$

Here θ has an upper bound independent of the sequence $\{X_k\}$ and of d; furthermore

$$Q_d(u) = \sum_{\nu=1}^{\infty} q_{d,\nu} u^\nu$$

where coefficient $q_{d,\nu}$ depends only on the ν first moments of $X_1, \dots, X_d$.

If, more particularly,

$$0 < \lambda_d u < 1/12$$

then

$$|\theta| < 9$$

and

$$|q_{d,\nu}| < (1/7)(12\lambda_d)^\nu.$$

Moreover, for any $0 < i < j \le d$ we have

$$|Q_j(u) - Q_i(u)| \le (1/2)(s_j^2 - s_i^2)/s_n^2.$$

Let us assume that n is odd, $n = 2k+1$; the case of even n is similar. Let us consider the sequence of independent random variables $X_1, \ldots, X_d$, where each X_i is uniformly distributed on the set of integers $\{-k, -k+1, \ldots, k\}$. For this sequence $X_1, \ldots, X_d$ we define S_d, F_d, s_d as in Feller's theorem. It is easy to see that $n^d F_d\left(t - \frac{nd+d}{2}\right) = \sum_{i=0}^{t} |L_n^d(i)|$ for all $t \in \mathbf{N}$.

In our case $\frac{\sqrt{d}k}{3} \le s_d \le \sqrt{d}k$, so we can put $\lambda_d = \frac{3}{\sqrt{d}}$. For $0 \le u \le 1$ and $d > 10^6$ it is seen that $\lambda_d u < \frac{1}{300}$ and thus

$$\begin{aligned} |Q_d(u)| &\le \sum_{\nu=1}^{\infty} |q_{d,\nu} u^\nu| \le \sum_{\nu=1}^{\infty} |q_{d,\nu}| \le \sum_{\nu=1}^{\infty} \frac{1}{7}(12\lambda_d)^\nu \\ &\le \frac{12\lambda_d}{7} \sum_{\nu=0}^{\infty} (12\lambda_d)^\nu \le \frac{36}{7\sqrt{d}} \sum_{\nu=0}^{\infty} \left(\frac{12}{300}\right)^\nu \le \frac{6}{\sqrt{d}}. \end{aligned}$$

So we arrive at

$$\begin{aligned} e^{-3d^{-1/2}} &\le e^{-\frac{1}{2}u^2 Q_d(u)} \le e^{3d^{-1/2}} \\ 1 - 3d^{-1/2} &\le e^{-\frac{1}{2}u^2 Q_d(u)} \le 1 + 4d^{-1/2}. \end{aligned}$$

From

$$1 - F_d(us_d) = e^{-(1/2)u^2 Q_d(u)} \left((1 - \Phi(u)) + \theta \lambda_d e^{-(1/2)u^2} \right)$$

we derive

$$\begin{aligned} 1 - F_d(us_d) &\le \left(1 + 4d^{-1/2}\right)\left(1 - \Phi(u) + 27d^{-1/2}\right) \\ F_d(us_d) &\ge \Phi(u) - 32d^{-1/2} \end{aligned}$$

and also

$$\begin{aligned} 1 - F_d(us_d) &\ge \left(1 - 3d^{-1/2}\right)\left(1 - \Phi(u) - 27d^{-1/2}\right) \\ F_d(us_d) &\le \Phi(u) + 32d^{-1/2}. \end{aligned}$$

Let $q_1 = \Phi^{-1}(0.7)$ and $q_2 = \Phi^{-1}(0.8)$; both of them satisfy $0 < q_1, q_2 < 1$. Then the i-layers for $i \in \left[\lfloor q_1 s_d + \frac{nd+d}{2} \rfloor, \lceil q_2 s_d + \frac{nd+d}{2} \rceil\right]$ contain together at least $\frac{1}{30} n^d$ vertices, since

$$\begin{aligned} F_d(q_2 s_d) - F_d(q_1 s_d) &\ge \Phi(q_2) - 32d^{-1/2} - \Phi(q_1) - 32d^{-1/2} \\ &= 1/10 - 64d^{-1/2} \ge 1/30. \end{aligned}$$

There are at most $(q_2 s_d - q_1 s_d) \le k d^{1/2}$ such layers, so the largest of them, $L_n^d(q_1 s_d + \frac{nd+d}{2})$, contains at least $\frac{1}{15d^{1/2}} n^{d-1}$ vertices. Now the ith layer for $i = q_1 s_d + \frac{nd+d}{2}$ cuts off a set of size at least $\frac{2}{3} n^d$ from the grid, since

$$\begin{aligned} F_d(q_1 s_d) &\ge \Phi(q_1) - 32d^{-1/2} \\ &= 7/10 - 32d^{-1/2} \ge 0.668 > 2/3. \end{aligned}$$

Thus, $|L_n^d(i_{2/3})| > |L_n^d(q_1 s_d + \frac{nd+d}{2})|$ and therefore, $|L_n^d(i_{2/3})| \ge \frac{1}{15d^{1/2}} n^{d-1}$. Let us remind that during all the previous steps we still assume that $d > 10^6$.

In the following proposition we show the upper bound on the size of the middle layer. To prove this, we need to show that two consecutive layers cannot differ too much in size.

Lemma A.3 *For any $n, d, i \in \mathbf{N}$ the following hold:*

(a) $|L_n^d(i_{1/2})| \leq 1000\frac{n^{d-1}}{\sqrt{d}}$,

(b) $|L_n^d(i)| - |L_n^d(i+1)| \leq |L_n^{d-1}(i-n)| \leq 2000\frac{n^{d-2}}{\sqrt{d}}$.

Proof. We prove both parts together by induction on d. First we show that $|L_n^d(i)| - |L_n^d(i+1)| \leq |L_n^{d-1}(i-n)| \leq |L_n^{d-1}(i_{1/2})|$. Let L' be $L_n^d(i)$ shifted by 1 in the first component, i.e. $L' = L_n^d(i) + (1, 0, \ldots, 0)$. Since $(L' \cap [n]^d) \subseteq L_n^d(i+1)$, we arrive at $|L_n^d(i)| - |L_n^d(i+1)| \leq |L' \setminus [n]^d|$. It is easy to check that $L' \setminus [n]^d$ is exactly a copy of the set $L_n^{d-1}(i-n)$, which has always smaller size than the middle layer $L_n^{d-1}(i_{1/2})$. The rest of (b) follows from (a).

To prove (a) for $d \leq 10^6$, we note that size of any layer $L_n^d(i)$ is bounded above by n^{d-1} and for $d \leq 10^6$ we have $n^{d-1} \leq 1000\frac{n^{d-1}}{\sqrt{d}}$. To prove (a) for $d > 10^6$, we estimate the size of the $\frac{n}{6}$ consecutive layers $L_n^d(i)$ for $i \in \left[i_{1/2}, i_{1/2} + \frac{n}{6}\right]$:

$$\begin{aligned} F_n^d\left(i_{1/2} + \frac{n}{6}\right) - F_n^d\left(i_{1/2}\right) &\leq \left(\Phi\left(\frac{n}{6s_d}\right) + \frac{32}{\sqrt{d}}\right) - \left(\Phi\left(0\right) - \frac{32}{\sqrt{d}}\right) \\ &\leq \Phi\left(\frac{1}{\sqrt{d}}\right) - \Phi(0) + \frac{64}{\sqrt{d}} \\ &\leq \frac{1}{\sqrt{2\pi}}\left(\int_0^{d^{-1/2}} e^{-(1/2)y^2} dy\right) + \frac{64}{\sqrt{d}} \\ &\leq 65d^{-1/2}. \end{aligned}$$

Among the $n/6$ layers the size of the smallest layer $L_n^d\left(i_{1/2} + \frac{n}{6}\right)$ is at most

$$\left|L_n^d\left(i_{1/2} + \frac{n}{6}\right)\right| \leq \frac{65d^{-1/2}n^d}{n/6} \leq 400\frac{n^{d-1}}{\sqrt{d}}.$$

Using (b) we estimate the size of the middle layer as

$$\begin{aligned} |L_n^d(i_{1/2})| &\leq \left|L_n^d\left(i_{1/2} + \frac{n}{6}\right)\right| + \frac{n}{6}\left(2000\frac{n^{d-2}}{\sqrt{d}}\right) \\ &\leq 1000\frac{n^{d-1}}{\sqrt{d}}. \end{aligned}$$

□