ISRAEL JOURNAL OF MATHEMATICS **TBD** (2018), 1–31
DOI: 10.1007/s11856-018-1803-0

EXTREMAL HYPERCUTS AND SHADOWS OF SIMPLICIAL COMPLEXES

BY

NATI LINIAL*

School of Computer Science and Engineering, The Hebrew University of Jerusalem
Givat Ram, Jerusalem 91904, Israel
e-mail: nati@cs.huji.ac.il

AND

ILAN NEWMAN**

Department of Computer Science, University of Haifa
Mount Carmel, Haifa 31905, Israel
e-mail: ilan@cs.haifa.ac.il

AND

YUVAL PELED†

School of Computer Science and Engineering, The Hebrew University of Jerusalem
Givat Ram, Jerusalem 91904, Israel
e-mail: yuvalp@cs.huji.ac.il

AND

YURI RABINOVICH††

Department of Computer Science, University of Haifa
Mount Carmel, Haifa 31905, Israel
e-mail: yuri@cs.haifa.ac.il

* Research supported by an ERC grant "High-dimensional combinatorics".
** This Research was supported by The Israel Science Foundation (grant number 862/10.)
† Yuval Peled is grateful to the Azrieli Foundation for the award of an Azrieli Fellowship.
†† This Research was supported by The Israel Science Foundation (grant number 862/10.)
Received August 1, 2017 and in revised form March 15, 2018

ABSTRACT

Let F be an n-vertex forest. An edge $e \notin F$ is said to be in F's shadow if $F \cup \{e\}$ contains a cycle. It is easy to see that if F is an "almost tree", i.e., a forest that contains two components, then its shadow contains at least $\lfloor \frac{(n-3)^2}{4} \rfloor$ edges and this is tight. Equivalently, the largest number of edges in an n-vertex cut is $\lfloor \frac{n^2}{4} \rfloor$. These notions have natural analogs in higher d-dimensional simplicial complexes which played a key role in several recent studies of random complexes. The higher-dimensional situation differs remarkably from the one-dimensional graph-theoretic case. In particular, the corresponding bounds depend on the underlying field of coefficients. In dimension $d = 2$ we derive the (tight) analogous theorems. We construct 2-dimensional "$\mathbb{Q}$-almost-hypertrees" (defined below) with an empty shadow. We prove that an "$\mathbb{F}_2$-almost-hypertree" cannot have an empty shadow, and we determine its least possible size. We also construct large hyperforests whose shadow is empty over every field. For $d \geq 4$ even, we construct a d-dimensional $\mathbb{F}_2$-almost-hypertree whose shadow has vanishing density.

Several intriguing open questions are mentioned as well.

1. Introduction

This article is part of an ongoing research in the field of "high-dimensional combinatorics". This research program starts from the observation that a graph can be viewed as a 1-dimensional simplicial complex, and that many basic concepts of graph theory, such as connectivity, forests, cuts, cycles, etc., have natural counterparts in the realm of higher-dimensional simplicial complexes. As may be expected, higher-dimensional objects tend to be more complicated than their 1-dimensional counterparts, and many fascinating phenomena reveal themselves from the present vantage point (see, e.g., [6, 10, 8, 4]). In this paper we study several extremal problems in this domain, and, in particular, we investigate the possible sizes of shadows and, equivalently, hypercuts of simplicial complexes. These concepts are based on homology theory, but they can be understood by means of standard linear algebra (see Section 2).

Let X be an n-vertex d-dimensional simplicial complex with a full skeleton, i.e., every face of dimension smaller than d belongs to X.

Pick a field $\mathbb{F}$ and define $\mathrm{SH}(X;\mathbb{F})$, the $\mathbb{F}$-**shadow** of X, as the set of all d-simplices $\sigma \notin X$ such that $H_d(X;\mathbb{F})$ is a proper subspace of $H_d(X \cup \{\sigma\};\mathbb{F})$. Here $H_d(\cdot)$ denotes the d-dimensional homology, the kernel of the boundary operator $\partial_d(X)$.

Shadows have recently played a significant role in the theory of random simplicial complex, specifically in the Linial–Meshulam binomial model $Y_d(n,p)$ of random simplicial complexes. In light of the classical Erdős–Rényi work, it is very natural to seek a high-dimensional analog of the phase transition and the emergence of the giant component that occurs in $G(n,p)$ at $p = \frac{1}{n}$. A major obstacle is that there is no high-dimensional notion of a connected component, but, as discovered in [8], rather than considering a giant component it is possible to observe the emergence of a giant shadow. In the case of graphs ($d = 1$) the two concepts are equivalent, but whereas there is no good notion of connected components in higher dimensions, shadows make sense in all dimensions. Moreover, it turns out that the density of the $\mathbb{Q}$-shadow of the random simplicial complex undergoes a discontinuous phase transition exactly when the first nontrivial d-cycle appears. Quantitatively, at the phase transition of the Linial–Meshulam complex, the size of the $\mathbb{Q}$-shadow jumps from $\Theta(n)$ to $\Theta(n^{d+1})$ with the addition of only $o(n^d)$ new d-faces. A motivating question for this work is how fast the shadow's size can jump. In particular, can the density of the shadow jump from 0 to 1 with the addition of a single new d-face?.

In addition, $\mathbb{F}$-shadows were crucial in studying the threshold for the integral homological connectivity of a random $Y_d(n,p)$ complex [5, 9]. Recall that the $(d-1)$-homology of a d-complex X with a full skeleton is the quotient $\mathrm{Im}\partial_d/\mathrm{Im}\partial_d(X)$ (see Section 2). The $\mathbb{F}$-shadow of X can be equivalently defined by the set of d-simplices $\sigma \notin X$ such that

$$H_{d-1}(X;\mathbb{F}) = H_{d-1}(X \cup \{\sigma\};\mathbb{F}).$$

The standard way of quantifying the fact that $Y_d(n,p)$ is far from having a trivial $(d-1)$-homology is by studying its Betti numbers. The results in [5, 9] indicate that a better means to this end is to consider the size of its shadow. This is founded on the fact that $H_{d-1}(X;\mathbb{F})$ is trivial if and only if X has a full $\mathbb{F}$-shadow, i.e.,

$$|\mathrm{SH}(X;\mathbb{F})| + |X| = \binom{n}{d+1},$$

where X is an n-vertex d-complex with a full skeleton, and its **size** $|X|$ is the number of d-faces in X.

This naturally suggests the question whether $H_{d-1}(X;\mathbb{F})$ is "small" if and only if $\mathrm{SH}(X;\mathbb{F})$ is "large". One implication does hold, and follows from Björner and Kalai's [1]. Namely, for every integer $k \le n$, if $|\mathrm{SH}(X;\mathbb{F})| + |X| \ge \binom{k}{d+1}$ then

$$\dim(H_{d-1}(X;\mathbb{F})) \le \binom{n-1}{d} - \binom{k-1}{d},$$

and this bound is tight. However, as we show here, the reverse implication fails. We present several constructions of simplicial complexes with a small $(d-1)$-homology, shadow and size.

Let us recall some further basic terminology. A d-complex X is **d-acyclic over** $\mathbb{F}$ if $H_d(X;\mathbb{F})$ is trivial. A $(d-1)$-face τ in a d-complex X is called **exposed** if it is contained in exactly one d-face σ of X. In the **elementary d-collapse** on such τ we remove τ and σ from X. We say that X is **d-collapsible** if it is possible to eliminate all the d-faces of X by a series of elementary d-collapses. It is easy to see that a d-collapsible complex is d-acyclic over every field.

The search for d-complexes with small $(d-1)$-homology, shadow and size can be restricted to d-acyclic complexes. Indeed, every d-complex X can be replaced by an inclusion-maximal d-acyclic subcomplex $X' \subseteq X$. These complexes have the same $(d-1)$-homology and $|X| + |\mathrm{SH}(X;\mathbb{F})| = |X'| + |\mathrm{SH}(X';\mathbb{F})|$.

Kalai [6] introduced the concept of a **d-hypertree** over $\mathbb{F}$. This is a maximal n-vertex d-complex with a full skeleton which is d-acyclic over $\mathbb{F}$. A d-hypertree has exactly $\binom{n-1}{d}$ d-faces, and, in addition, every n-vertex d-acyclic d-complex with this number of d-faces is a d-hypertree. Because d-hypertrees have a trivial $(d-1)$-homology and thus a full shadow, they are not useful for us. However, a slight modification significantly changes the game. A d-hypertree from which one d-face is removed is called a **d-almost-hypertree**. For example, a d-collapsible complex with $\binom{n-1}{d} - 1$ d-faces is an almost-hypertree over every field. Clearly, the $(d-1)$-homology of a d-almost-hypertree is only one-dimensional, but how small can its shadow be?

Let us illustrate with some concrete examples.

Example 1.1 ($d = 1$): In the one-dimensional case, a 1-hypertree is a spanning tree and a 1-almost-hypertree is a forest F with two connected components $A \cup B = V(F)$. Its shadow is comprised of all pairs of vertices within the same connected component, and its size is $\binom{|A|}{2} + \binom{|B|}{2} - (n-2) \ge (\frac{1}{2} - o(1))\binom{n}{2}$, where $n = |V(F)|$.

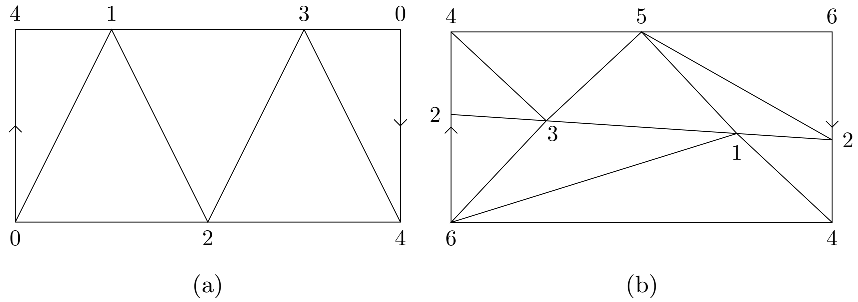

Figure 1. The triangulations M_5 and M_6 of the Möbius band.

Example 1.2 ($d = 2$, $n = 5$, an optimal construction): Consider the 5-vertex triangulation M_5 of the Möbius strip in Figure 1(a) and an arbitrary field $\mathbb{F}$. The 1-skeleton of M_5 is full, and its triangular faces are the 5 arithmetic progressions of the form $\{x, x+1, x+2\}$ in $\mathbb{Z}_5$. It is easy to see that M_5 is 2-collapsible with $\binom{5-1}{2} - 1$ triangular faces, i.e., it is a 2-collapsible almost-hypertree. In addition, for every triangle $\sigma \notin M_5$, the complex $M_5 \cup \{\sigma\}$ is also 2-collapsible, hence $\mathrm{SH}(M_5; \mathbb{F}) = \emptyset$.

Example 1.3 ($d = 2$, $n = 6$, dependence on the field): Consider the 6-vertex triangulation M_6 of the Möbius strip in Figure 1(b). As before, M_6 is a 2-collapsible almost-hypertree. The complex $M_6 \cup \{\sigma\}$, where $\sigma = \{4, 5, 6\}$, is the 6-vertex triangulation of the real projective plane. For all other triangles $\tau \notin M_6$, the complex $M_6 \cup \{\tau\}$ is 2-collapsible. Consequently, $\mathrm{SH}(X; \mathbb{F})$ is not empty if and only if the field $\mathbb{F}$ has characterisitc 2.

Here is a summary of the results and constructions presented in this paper:

(I) We show that for certain integers n, there exists a 2-dimensional almost-hypertree with an empty shadow over $\mathbb{Q}$. Moreover, assuming a well-known conjecture by Artin in number theory, there are infinitely many n.

(II) Over the field $\mathbb{F}_2$, surprisingly, the situation changes. There are no shadowless 2-almost-hypertrees, but the least possible size of the shadow of a 2-almost-hypertree is $\frac{n^2}{4} + \Theta(n)$.

(III) For the same question in dimension $d > 2$ over $\mathbb{F}_2$, the answer depends on d's parity. For odd d every d-almost-hypertree has a shadow of positive density. For even d there are d-almost-hypertrees with a shadow of vanishingly small density.

(IV) For odd n, we construct 2-acyclic simplicial complexes with

$$\binom{n-1}{2} - (n+1)$$

two-faces that are shadowless over every field. Moreover, these complexes are 2-collapsible even after the addition of an arbitrary new two-face. In consequence, for every field $\mathbb{F}$ the density of the $\mathbb{F}$-shadow can jump from 0 to 1 with the addition of $n+1$ 2-faces.

Perfect hypercuts—a different perspective. The theory of simplicial matroids (see, e.g., [3]) suggests a different perspective to the above results which is more useful when working over the field $\mathbb{F}_2$. The elements of the simplicial matroid with parameters n,d integers and a field $\mathbb{F}$ are all the possible d-faces over n vertices. A set S of d-faces is said to be **independent** if the d-complex that consists of S and the full $(d-1)$-skeleton is d-acyclic over $\mathbb{F}$. The bases of the simplicial matroid are d-hypertrees over $\mathbb{F}$, and the shadow of a d-complex is closely related to its closure. For instance, a closed set in the simplicial matroid is a shadowless complex.

The standard concept of a cocircuit in matroid theory leads to the following definiton. A **d-hypercut over** $\mathbb{F}$ is an inclusion-minimal set of d-faces that intersects every d-hypertree. Note that a 1-hypercut is a graphical cut, namely all the edges connecting a subset of the vertices and its complement. Therefore, an almost tree, i.e., a forest F with two trees, uniquely defines the cut consisting of those edges that join vertices of the two connected components of F. As we show in the sequel, the general d-dimensional situation is essentially the same.

Lemma 1.4: *Let $n > d$ be integers and $\mathbb{F}$ a field. A set C of d-faces on n vertices is a d-hypercut over $\mathbb{F}$ if and only if there exists a d-almost-hypertree X over $\mathbb{F}$ such that*

$$C = \binom{[n]}{d+1} \setminus \Big(X \bigcup SH(X;\mathbb{F})\Big).$$

Here $\binom{[n]}{d+1}$ denotes all the d-faces over n vertices. We postpone the proof to Section 2.

It follows that the size of a d-hypercut is at most $\binom{n}{d+1} - (\binom{n-1}{d} - 1)$ and the bound is attained if and only if there exists a shadowless d-almost-hypertree. We say that a d-hypercut of this size is **perfect**.

Note that all 1-dimensional cuts are far from being perfect. The largest size of a cut in an n-vertex graph is $\lfloor \frac{n^2}{4} \rfloor$, whereas the bound given in the lemma is $\binom{n}{2} - n + 2$. The results of this paper, stated in terms of hypercuts, reveal a completely different situation in higher dimensions .

(I′) Perfect 2-hypercuts over $\mathbb{Q}$ exist for certain integers n, and assuming a well-known number-theoretic conjecture of Artin, there are infinitely many such n.

(II′) There are no perfect 2-hypercuts over $\mathbb{F}_2$ for large n. The largest possible 2-hypercut has $\binom{n}{3} - \frac{3}{4}n^2 - \Theta(n)$ faces. All the extremal hypercuts are characterized.

(III′) For $\mathbb{F} = \mathbb{F}_2$ and $d > 2$, the situation depends on d's parity. Namely, for d even, the largest d-hypercuts have $\binom{n}{d+1} \cdot (1 - o_n(1))$ d-faces. When d is odd, the density of all d-hypercuts is bounded away from 1.

Since the notions of shadow and hypercut are formally equivalent (Lemma 1.4), we will opt in each case for the terminology that seems better suited for the context. The rest of the paper is organized as follows. In Section 2 we introduce the necessary notions in the combinatorics of simplicial complexes. Section 3 deals with the problem of a shadowless 2-almost-hypertree over $\mathbb{Q}$. In Section 4 we study the problem of a maximal 2-hypercut over $\mathbb{F}_2$. In Section 5 we construct large d-hypercuts over $\mathbb{F}_2$ for even $d \geq 4$, and show the impossibility of similar constructions for odd d. In Section 6 we construct large acyclic shadowless 2-dimensional complexes over every field. Lastly, in Section 7 we present some of the many open questions in this area.

2. Background on simplicial combinatorics

All simplicial complexes considered here have vertex set $V = [n] = \{1, \ldots, n\}$ or $V = \mathbb{Z}_n$. A simplicial complex X is a collection of subsets of V that is closed under taking subsets. Namely, if $\sigma \in X$ and $\tau \subseteq \sigma$, then $\tau \in X$ as well. Members of X are called **faces** or **simplices**. The **dimension** of the simplex $\sigma \in X$ is defined as $|\sigma| - 1$. A d-dimensional simplex is also called a d-simplex or a d-face for short. The dimension $\dim(X)$ is defined as $\max \dim(\sigma)$ over all faces $\sigma \in X$, and we also refer to a d-dimensional simplicial complex as a d-complex. The **size** $|X|$ of a d-complex X is the number of d-faces in X.

The collection of the faces of dimension $\leq t$ of X, where $t < d$, is called the t-**skeleton** of X. We say that a d-complex X has a **full skeleton** if its $(d-1)$-skeleton contains all the faces of dimensions at most $(d-1)$ spanned by its vertex set. If X has a full skeleton, the **complement** $\bar{X}$ is defined by taking a full $(d-1)$-dimensional skeleton and those d-faces that are not in X.

The permutations on the vertices of a face σ are split in two **orientations** of σ, according to the permutation's sign. The **boundary operator** $\partial = \partial_d$ maps an oriented d-simplex $\sigma = (v_0, \ldots, v_d)$ to the formal sum

$$\partial\sigma = \sum_{i=0}^{d} (-1)^i (\sigma \setminus v_i),$$

where $\sigma \setminus v_i = (v_0, \ldots, v_{i-1}, v_{i+1}, \ldots, v_d)$ is an oriented $(d-1)$-simplex.

We fix some field $\mathbb{F}$ and linearly extend the boundary operator to free $\mathbb{F}$-sums of simplices. Here we mostly consider $\mathbb{F} = \mathbb{Q}$ or $\mathbb{F}_2$. We consider the $\binom{n}{d} \times \binom{n}{d+1}$ matrix form of ∂_d by choosing arbitrary orientations for $(d-1)$-simplices and d-simplices. Note that changing the orientation of a d-simplex (resp. $d-1$-simplex) results in multiplying the corresponding column (resp. row) by -1. Thus the d-boundary of a weighted sum of d simplices, viewed as a vector z (of weights) of dimension $\binom{n}{d+1}$, is just the matrix-vector product $\partial_d z$.

We denote by $\partial_d(X)$ the submatrix of ∂_d restricted to the columns associated with d-faces of a d-complex X. We occasionally use the notation $\partial_d(A)$, where A is a subset of the d-faces rather than $\partial_d(X)$. This creates no problem for X that has a full skeleton. We observe that the matrix ∂_d has rank $\binom{n-1}{d}$, regardless of the choice of field. We define the rank of a d-complex X by $\text{rank}(\partial_d(X))$.

The d-homology of X, denoted by $H_d(X; \mathbb{F})$, is the right kernel of the matrix $\partial_d(X)$. An element in $H_d(X; \mathbb{F})$ is also called a d-**cycle**. The $(d-1)$-th homology of X, denoted by $H_{d-1}(X; \mathbb{F})$, is the quotient $\text{Im}\partial_d \,/\, \text{Im}\partial_d(X)$.

Let us recall the main concepts used in this paper. A d-hypertree X over the field $\mathbb{F}$ is a d-complex of size $\binom{n-1}{d}$ with a trivial d-dimensional homology over $\mathbb{F}$. This means that the columns of the matrix $\partial_d(X)$ form a basis for the column space of ∂_d. It is easily verified that a d-hypertree has a full skeleton. A d-almost-hypertree is a d-complex of size $\binom{n-1}{d} - 1$ with a trivial d-homology over $\mathbb{F}$. The $\mathbb{F}$-shadow of a d-complex Y is comprised of all d-faces $\sigma \notin Y$ such that $\partial\sigma$ is linearly spanned over $\mathbb{F}$ by the columns of $\partial_d(Y)$. Equivalently, $\sigma \in \text{SH}(Y; \mathbb{F})$ if and only if adding σ to Y does not increase its rank. Lastly, a d-hypercut is an inclusion-minimal set of d-faces that intersects every d-hypertree.

We now prove Lemma 1.4, which shows the equivalnce between shadows of d-almost-hypertrees and d-hypercuts.

Proof of Lemma 1.4. Everything is done over $\mathbb{F}$. We show first that if

$$C = \binom{[n]}{d+1} \setminus \left(X \bigcup \mathrm{SH}(X;\mathbb{F})\right),$$

where X is a d-almost-hypertree, then C is a d-hypercut. This entails two things:

(1) That C meets every d-hypertree, or in other words, that $X \bigcup \mathrm{SH}(X;\mathbb{F})$ contains no d-hypertree, which is clear since its rank is only $\binom{n-1}{d} - 1$.
(2) That for every proper subset $C' \subsetneq C$ there is a d-hypertree that is disjoint from C'. Indeed, if $C' \subseteq C \setminus \sigma$, then $X \cup \{\sigma\}$ is clearly disjoint from C'. Also, $X \cup \{\sigma\}$ is a d-hypertree, since X is an almost d-hypertree and σ is not in its shadow.

Now for the reverse implication: Let C be a d-hypercut and let X be an $\mathbb{F}$-basis for the set of all the d-faces that are not in C. We claim that X is a d-almost-hypertree. Since X is d-acyclic, it remains to show that

$$|X| = \binom{n-1}{d} - 1.$$

Clearly $|X| < \binom{n-1}{d}$, or else X is a d-hypertree that is disjoint from the d-hypercut C. If $|X| \leq \binom{n-1}{d} - 2$, then the collection of all d-faces that are not in C has rank $\leq \binom{n-1}{d} - 2$, which implies that every hypertree has at least two d-faces in common with C. This, however, contradicts the minimality of C. It remains to prove that the $\mathbb{F}$-shadow of X is disoint from C. Indeed, if $\sigma \in C \cap \mathrm{SH}(X;\mathbb{F})$ then $C \setminus \sigma$ intersects every d-hypertree which contradicts the minimality. ∎

There is another useful characterization of d-hypercuts. Recall that the row space of the $\binom{n}{d} \times \binom{n}{d+1}$ matrix ∂_d is the linear space B^d of d-coboundaries. A set C of d-faces intersects every d-hypertree if and only if the rank of its complement $\bar{C}$ is strictly smaller than $\binom{n-1}{d}$. This is equivalent to the existence of a non-zero d-coboundary that is supported on C. Therefore, C is a d-hypercut if and only if it is an inclusion-minimal support of a d-coboundary.

Here are some additional concepts needed for the proofs in the paper. If σ is a face in a complex X, we define its **link** via

$$\mathrm{link}_\sigma(X) = \{\tau \in X : \tau \cap \sigma = \emptyset,\, \tau \cup \sigma \in X\}.$$

This is clearly a simplicial complex. For instance, the link of a vertex v in a graph G is v's neighbour set which we also denote by $N_G(v)$ or $N(v)$.

We occasionally call a set of d-simplices a d-cycle or a d-coboundary if it is the support of a d-cycle or a d-coboundary. Over $\mathbb{F}_2$ this makes no difference, since a vector over $\mathbb{F}_2$ is naturally identified with its support. In this case, a d-coboundary C consists of the d-faces whose boundary has an odd intersection with some set A of $(d-1)$-faces. A convenient way to generate a given d-coboundary C is to let A be the $(d-1)$-faces in $\mathrm{link}_v(C)$ for some vertex v, i.e., it holds that

$$C = \mathrm{link}_v(C) \cdot \partial_d.$$

In words, the characteristic (row) vector of the d-faces of C is equal to the vector-matrix left product of the characteristic (row) vector of the $(d-1)$-faces of $\mathrm{link}_v(C)$ with the boundary matrix ∂_d. Indeed, one needs to verify that a d-face σ belongs to C if and only if the number of the $(d-1)$-faces of σ that are in $\mathrm{link}_v(C)$ is odd. In case $v \in \sigma$, $\mathrm{link}_v(C)$ contains the $(d-1)$-face $\sigma \setminus \{v\}$ if and only if $\sigma \in C$, and does not contain other $(d-1)$-faces of σ. Otherwise, in case $v \notin \sigma$, the number of d-faces in the $(d+1)$-face $\sigma \cup \{v\}$ that belong to C is even, since C is a d-coboundary and

$$\partial_d \partial_{d+1} = 0.$$

Therefore, $\sigma \in C$ if and only if the number of d-faces of $\sigma \cup \{v\}$ that belong to C other than σ is odd. This concludes the argument since the latter equals the number of the $(d-1)$-faces of σ that belong to $\mathrm{link}_v(C)$.

In the 2-dimensional case over the field $\mathbb{F}_2$, there is a combinatorial property that characterizes whether the link graph $G = \mathrm{link}_v(C)$ generates a 2-hypercut C rather than a general 2-coboundary. Two incident edges uv,uw in a graph $G = (V, E)$ are said to be **Λ-adjacent** if $vw \notin E$. We say that G is **Λ-connected** if the transitive closure of the Λ-adjacency relation has exactly one class.

Proposition 2.1 ([12]): *A 2-dimensional coboundary C is a hypercut if and only if the graph $\mathrm{link}_v(C)$ is Λ-connected for every v.*

3. Shadowless almost-hypertrees over $\mathbb{Q}$

The main result of this section is a construction of shadowless 2-dimensional almost-hypertrees over $\mathbb{Q}$. As mentioned above, the complement of such a complex is a perfect 2-hypercut having $\binom{n}{3} - \binom{n-1}{2} + 1$ faces which is the most possible.

THEOREM 3.1: *Let $n \geq 5$ be a prime for which $\mathbb{Z}_n^*$ is generated by $\{-1, 2\}$. Let $X = X_n$ be a 2-dimensional simplicial complex on vertex set $\mathbb{Z}_n$ whose 2-faces are arithmetic progressions of length 3 in $\mathbb{Z}_n$ with difference not in $\{0, \pm 1\}$. Then:*

- *X_n is 2-collapsible, and hence it is an almost-hypertree over every field.*
- *$SH(X_n; \mathbb{Q}) = \emptyset$. Consequently, the complement of X_n is a 2-perfect hypercut over $\mathbb{Q}$.*

The entire construction and much of the discussion of X_n is carried out over $\mathbb{Z}_n$, but the boundary operator $\partial_2(X_n)$ is considered over the rationals.

We start with two simple observations. First, note that X_n has a full 1-skeleton, i.e., every edge is contained in some 2-face of X_n. Also, we note that the choice of omitting the arithmetic triples with difference ± 1 is arbitrary. Namely, for every $a \in \mathbb{Z}_n^*$, the automorphism $r \mapsto ar$ of $\mathbb{Z}_n$ maps X_n to a combinatorially isomorphic complex of arithmetic triples over $\mathbb{Z}_n$, with omitted difference $\pm a$. Consequently, Theorem 3.1 holds equivalently for any difference that we omit. In what follows we indeed assume for convenience that the missing difference is not ± 1, but rather $\pm 2^{-1} \in \mathbb{Z}_n$.

For $d \in Z_n^*$, define

$$E_d = E_{d,n} = ((0, d), (1, d+1), \ldots, (n-1, d+n-1)),$$

where all additions are $\mod n$. This is an ordered subset of directed edges in X_n.

Similarly, we consider the collection of arithmetic triples of difference d,

$$F_d = F_{d,n} = ((0, d, 2d), (1, d+1, 2d+1), \ldots, (n-1, d+n-1, 2d+n-1)).$$

Clearly every directed edge appears in exactly one E_d and then its reversal is in E_{-d}. Likewise for arithmetic triples and the F_d's. Since we assume that Z_n^* is generated by $\{-1, 2\}$, it follows that the powers $\{2^i\} \subset Z_n^*$, $i = 0, \ldots, \frac{n-1}{2} - 1$, are all distinct, and, moreover, no power is an additive inverse of the other. Therefore, the sets $\{E_{2^i}\}$, $i = 0, \ldots, \frac{n-1}{2} - 1$, constitute a partition of the 1-faces of X_n. Similarly, the sets $\{F_{2^j}\}$, $j = 0, \ldots, \frac{n-1}{2} - 2$, constitute a partition of the 2-faces of X_n. The omitted difference is $2^{\frac{n-1}{2}-1} \in \{\pm 2^{-1}\}$, as assumed (the sign is determined according to whether $2^{\frac{n-1}{2}} = 1$ or -1).

LEMMA 3.2: *Ordering the rows of the adjacency matrix M_X by E_{2^i}'s, and ordering the columns by the F_{2^i}'s, the matrix $\partial_2(X_n)$ takes the following form:*

$$\partial_2(X_n) = \begin{pmatrix} I+Q & 0 & 0 & \cdots \\ -I & I+Q^2 & 0 & \cdots \\ 0 & -I & \ddots & \cdots \\ 0 & 0 & \ddots & I+Q^{2^{\frac{n-1}{2}-2}} \\ 0 & 0 & \cdots & -I \end{pmatrix} \tag{1}$$

where each entry is an $n \times n$ matrix (block) indexed by $\mathbb{Z}_n$, and Q is a permutation matrix corresponding to the linear map $b \mapsto b+1$ in $\mathbb{Z}_n$.

Proof. Consider an oriented face $\sigma \in F_{2^i} \subset X_n$. Then,

$$\sigma = (b, b+2^i, b+2^{i+1})$$

for some $b \in \mathbb{Z}_n$ and $0 \leq i \leq \frac{n-1}{2} - 2$, i.e., σ is the b-th element in F_{2^i}. By definition,

$$\partial\sigma = (b, b+2^i) + (b+2^i, b+2^{i+1}) - (b, b+2^{i+1}).$$

The first two terms in $\partial\sigma$ are the b-th and $(b+2^i)$-th elements in E_{2^i} respectively; the third term corresponds to the b-th element in $E_{2^{i+1}}$. Thus, the blocks indexed by $E_{2^i} \times F_{2^i}$ are of the form $I + Q^{2^i}$, the blocks $E_{2^{i+1}} \times E_{2^i}$ are $-I$, and the rest is 0. ∎

We may now establish the main result of this section.

Proof of Theorem 3.1. We start with the first statement of the theorem. Let $m = \frac{n-1}{2}$.

Lemma 3.2 implies that the edges in $E_{2^{m-1}}$ are exposed. Collapsing on these edges leads to elimination of $E_{2^{m-1}}$ and the faces in $F_{2^{m-2}}$. In terms of the matrix $\partial_2(X_n)$, this corresponds to removing the rightmost "supercolumn". Now the edges in $E_{2^{m-2}}$ become exposed, and collapsing them leads to elimination of $E_{2^{m-2}}$ and $F_{2^{m-3}}$. This results in exposure of $E_{2^{m-3}}$, etc. Repeating the argument to the end, all the faces of X_n get eliminated, as claimed.

To show that X_n is an almost-hypertree we need to show that the number of its 2-faces is $\binom{n-1}{2} - 1$. Indeed,

$$|X_n| = \sum_{j=0}^{\frac{n-1}{2}-2} |F_{2^j}| = \left(\frac{n-1}{2} - 1\right) \cdot n = \binom{n-1}{2} - 1.$$

We now show the second statement of the theorem, i.e., that $\mathrm{SH}(X_n;\mathbb{Q}) = \emptyset$. Let $u \in \mathbb{Q}^{\binom{n}{2}}$ be a vector indexed by the edges of X_n, where $u_e = 2^i$ when $e \in E_{2^i}$. Here we think of 2^i as an integer (and not an element in $\mathbb{Z}_n$). We claim that for every 2-face $\sigma \in \binom{\mathbb{Z}_n}{3}$,

$$\langle u, \partial\sigma\rangle = 0 \iff \sigma \in X_n.$$

Indeed, for every 2-face σ, exactly three coordinates in the vector $\partial\sigma$ are non-zero, and they are ± 1. Since the entries of u are successive powers of 2, the condition $\langle u, \partial\sigma\rangle = 0$ holds if and only if $\partial\sigma$ (or $-\partial\sigma$) has two 1's in E_{2^i} and one -1 in $E_{2^{i+1}}$ for some $0 \le i \le \frac{n-1}{2} - 1$. This happens if and only if σ is of the form $(b, b+2^i, b+2^{i+1})$, i.e., precisely when $\sigma \in X_n$.

Suppose, in contradiction, that there exists $\sigma' \in \mathrm{SH}(X_n;\mathbb{Q})$, i.e.,

$$\partial\sigma' = \sum_{\sigma \in X_n} \lambda_\sigma \partial\sigma,$$

where the λ_σ's are rational scalars. A contradiction follows by taking the inner product of these two vectors with u. ∎

When the prime n does not satify the assumption of Theorem 3.1 we can still say something about the structure of X_n. Let the group $G_n = \mathbb{Z}_n^*/\{\pm 1\}$, and let H_n be the subgroup of G_n generated by 2. Then:

THEOREM 3.3: *For every prime number n,*

$$\mathrm{rank}_{\mathbb{Q}}(\partial_2(X_n)) = |X_n| - (n-1)\cdot([G_n : H_n] - 1).$$

In particular, X_n is acyclic if and only if $\mathbb{Z}_n^*$ is generated by $\{\pm 1, 2\}$.

We only sketch the proof. Recall the partition of X_n's edges and faces to the sets E_i and F_i. We now consider a coarser partition by joining together all the E_i's and F_i's for which i belongs to some coset of H_n. This induces a block structure on $\partial_2(X_n)$ with $[G_n : H_n]$ blocks. An argument as in the proof of Lemma 3.2 yields the structure of these blocks. Finally, an easy computation shows that one of these blocks is 2-collapsible, and each of the others contribute precisely $n-1$ vectors to the right kernel.

We conclude this section by recalling the following well-known conjecture of Artin which is implied by the generalized Riemann hypothesis [11].

CONJECTURE 3.4 (Artin's Primitive Root Conjecture): *Every integer other than -1 that is not a perfect square is a primitive root modulo infinitely many primes.*

This conjecture clearly yields infinitely many primes n for which $\mathbb{Z}_n^*$ is generated by 2. (It is even conjectured that the set of such primes has positive density.) Clearly this implies that the assumptions of Theorem 3.1 hold for infinitely many primes n.

4. Largest hypercuts over $\mathbb{F}_2$

In this section we discuss questions over the field $\mathbb{F}_2$, where it is more convenient to work with the notion of hypercuts. Recall that a 2-dimensional hypertcut is an inclusion minimal set of 2-faces that meets every 2-hypertree. The main result of this section is:

THEOREM 4.1: *For large enough n, the largest size of an n-vertex 2-dimensional hypercut over $\mathbb{F}_2$ is $\binom{n}{3} - (\frac{3}{4}n^2 - \frac{7}{2}n + 4)$ for even n and $\binom{n}{3} - (\frac{3}{4}n^2 - 4n + \frac{25}{4})$ for odd n.*

Remark 4.2: The proof also provides a characterization of the extremal cases of this theorem.

Since no confusion is possible, in this section we use the shorthand term **cut** for a 2-dimensional hypercut.

The first and more interesting step in proving Theorem 4.1 is the slightly weaker Theorem 4.3. A further refinement that yields the tight upper bound on the size of cuts is given in Appendix B.

Note that since $\binom{[n]}{3}$ is a coboundary, the complement $\bar{C} = \binom{[n]}{3} \setminus C$ of any cut C is a coboundary. Moreover, the complement of the $(n-1)$-vertex graph $\mathrm{link}_v(C)$ is the link $\mathrm{link}_v(\bar{C})$. In what follows, $\mathrm{link}_v(C)$ is always considered as an $(n-1)$-vertex graph with vertex set $[n] \setminus \{v\}$. Occasionally, we will consider the graph $\mathrm{link}_v(C) \cup \{v\}$ which has v as an isolated vertex.

THEOREM 4.3: *The size of every n-vertex cut is at most $\binom{n}{3} - \frac{3}{4} \cdot n^2 + o(n^2)$. In every cut C that attains this bound there is a vertex v for which the graph $G = \mathrm{link}_v(C)$ satisfies either*

(1) *$\bar{G}$ has one vertex of degree $\frac{n}{2} \pm o(n)$ and all other vertices have degree $o(n)$; moreover, $|E(\bar{G})| = n - 1 + o(n)$;*
(2) *$\bar{G}$ has one vertex of degree $n - o(n)$, one vertex of degree $\frac{n}{2} \pm o(n)$, and all other vertices have degree $o(n)$; moreover, $|E(\bar{G})| = 2n \pm o(n)$.*

We need to make some preliminary observations.

Observation 4.4: *Let $G = (V, E)$ be a graph with n vertices, m edges and t triangles. Denote the degree of a vertex $v \in G$ by d_v, and let C be the coboundary generated by G. Then*

$$|C| = nm - \sum_{v \in V} d_v^2 + 4t.$$

Proof. Let $e = (u, v) \in E(G)$. Then $\text{link}_e(C)$ consists of those vertices $x \neq u, v$ that are adjacent to both or none of u, v. Namely,

$$|\text{link}_e(C)| = n - d_u - d_v + 2|N(v) \cap N(u)|.$$

Clearly $|N(v) \cap N(u)|$ is the number of triangles in G that contain e. But $\sum_{e \in E(G)} |\text{link}_e(C)|$ counts every two-face in C three times or once, depending on whether or not it is a triangle in G. Therefore

$$|C| + 2t = \sum_{(u,v) \in E} (n - d_u - d_v + 2|N(v) \cap N(u)|).$$

The claim follows. ∎

Two vertices in a graph are called **clones** if they have the same set of neighbours (in particular they must be nonadjacent).

Observation 4.5: *For every non-empty n-vertex cut C and a vertex $v \in V(C)$, the $(n-1)$-vertex link graph of the complement complex $\bar{G} := \text{link}_v(\bar{C})$ is either (i) a union of two cliques, (ii) a complete graph minus one edge, or (iii) a connected graph with no clones.*

Proof. Directly follows from the fact that

$$G := \text{link}_v(C)$$

is Λ-connected (Proposition 2.1). Indeed, if the edges uv and uw are Λ-adjacenct edges in G, then v and w belong to the same connected component in $\bar{G}$. Therefore, in case $\bar{G}$ is not connected it must be the union of two cliques. In addition, a clone in $\bar{G}$ yields a Λ-isolated edge in G and is permissible only if it is the only edge in G. ∎

The size of an n-vertex cut C for which $\bar{G}$ is of types (i) or (ii) above is of size $O(n^2)$, that is much smaller than the bound in Theorem 4.3. Therefore, we restrict the following discussion to cuts C for which $\bar{G}$ is connected and has no clones.

In the following claims, let C be a cut and $x \in V(C)$. Let

$$G = (V, E) = \text{link}_x(C),$$

where $V = V(C) \setminus \{x\}$, and

$$\bar{G} = (V, \bar{E}) = \text{link}_x(\bar{C}).$$

Note that there are only $n-1$ vertices in $\bar{G}$. In addition, denote by $m = |E(\bar{G})|$ and

$$d = (d_1 \geq d_2 \geq \cdots \geq d_{n-1} \geq 1)$$

the sorted degree sequence of $\bar{G}$. We label the vertices $v_1, \ldots, v_{n-1}$ so that $d(v_i) = d_i$ for all i. Recall that $N_{\bar{G}}(v)$ denotes the set of v's neighbours in $\bar{G}$.

For every $S \subseteq V$, an S-**atom** is a subset $A \subseteq V \setminus S$ which satisfies:

$$(u, v) \in E(\bar{G}) \iff (u', v) \in E(\bar{G})$$

for every $u, u' \in A$ and $v \in S$.

The next claim generalizes Observation 4.5.

Claim 4.6: *Let $S \subseteq V$ and $G' = \bar{G} \setminus S$ be the subgraph induced by $V \setminus S$. Then, for every non-empty S-atom A, at least $|A| - 2$ of the edges in G' meet A.*

Proof. Let H be the subgraph of G' induced by an atom A. If H has at most two connected components, the claim is clear, since a connected graph on r vertices has at least $r-1$ edges. We next consider what happens if H has three or more connected components. We show that every component except possibly one has an edge in $\bar{E}$ that connects it to $V \setminus (S \cup A)$. This clearly proves the claim.

So let C_1, C_2, C_3 be connected components of H, and suppose that neither C_1 nor C_2 is connected in $\bar{G}$ to $V \setminus (S \cup A)$. Let

$$F := \bigcup_{1 \leq i < j \leq 3} C_i \times C_j \subseteq E.$$

Since G is Λ-connected, there must be a Λ-path connecting every edge in $C_1 \times C_2$ to every edge in $C_2 \times C_3$. However, every path that starts in $C_1 \times C_2$ can never leave it. Indeed, let us consider the first time this Λ-path exits $C_1 \times C_2$, say xy that is followed by yw, where $x \in C_1, y \in C_2$, $w \notin C_1 \cup C_2$ and $yw \notin E$. By the atom condition, a vertex in S does not distinguish between vertices $x, y \in A$, whence $w \notin S$. Finally w cannot be in A, for $xw \notin E$ would imply that $w \in C_1$. Hence, C_1 is connected in $\bar{G}$ to $V \setminus (S \cup A)$, a contradiction. ∎

CLAIM 4.7: $d_1 \leq m/2 + 1$.

Proof. Apply Claim 4.6 with $S = \{v_1\}$ and $A = N_{\bar{G}}(v_1)$. It yields the existence of at least $|A| - 2$ edges in $\bar{G}$ that meet A but not v_1. Since $|A| = d_1$, $m \geq d_1 + (d_1 - 2)$, implying the claim. ∎

CLAIM 4.8: $d_1 + d_2 \leq \frac{m+n+1}{2}$.

Proof. Apply Claim 4.6 with $S = \{v_1, v_2\}$ and $A = N_{\bar{G}}(v_1) \cap N_{\bar{G}}(v_2)$ to conclude that $m \geq d_1 + d_2 + |A| - 3$ (as (v_1, v_2) might be an edge). By inclusion-exclusion, $n - 1 \geq d_1 + d_2 - |A|$. These two inequalities imply the claim. ∎

CLAIM 4.9: *For every $k \geq 2$,*

$$\sum_{i=1}^{k} d_i \leq m - \frac{n}{2} + 2^{k+1}.$$

Proof. Let $S = \{v_1, \ldots, v_k\}$. The sum $\sum_{i=1}^{k} d_i$ equals the number of edges between S and $V \setminus S$ plus twice the number of edges induced by S. Therefore, there are at least $\sum_{i=1}^{k} d_i - \binom{k}{2}$ edges with a vertex in S. In addition, we claim that there are at least $\frac{1}{2}(n - 1 - k - 2 \cdot 2^k)$ edges in $\bar{G} \setminus S$. Indeed, there are 2^k inclusion-maximal atoms of S, and every vertex in $V \setminus S$ belongs to one of them. We apply Claim 4.6 to each such atom A to conclude that there are at least $|A| - 2$ edges having one vertex in A and the other not in S. The sum of $|A| - 2$ over the 2^k inclusion-maximal atoms A is $n - 1 - k - 2 \cdot 2^k$ but every such edge may be counted twice. Consequently,

$$m \geq \sum_{i=1}^{k} d_i - \binom{k}{2} + \frac{1}{2}(n - 1 - k - 2 \cdot 2^k),$$

and the claim follows. ∎

Proof of Theorem 4.3. Let C be a cut, and suppose that $|\bar{C}| = \frac{\gamma}{3} n^2$, for some $\gamma \geq 0$. We assume that $\gamma \leq 9/4$, as otherwise the Theorem follows directly. By averaging, there is a link, say $\bar{G} = (V, \bar{E}) = \mathrm{link}_v(\bar{C})$, of at most

$$m := \frac{3|\bar{C}|}{n} = \gamma n$$

edges, where $V = [n] \setminus \{v\}$. We may also assume that $m \geq n - 2$ since $\bar{G}$ is connected by the discussion following Observation 4.5.

As before, we denote by $d_1 \geq d_2 \geq \cdots \geq d_{n-1} \geq 1$ the sorted degree sequence of $\bar{G}$. Recall that $\bar{C}$ is a 2-coboundary over $\mathbb{F}_2$. Therefore, Observation 4.4 implies that $|\bar{C}| \geq mn - \sum_i d_i^2$. We want to reduce the problem of proving a lower bound on $|\bar{C}|$ to showing a lower bound on $mn - \sum_1^k d_i^2$, where $k = k(n)$ is an appropriately chosen slowly growing function. We note that $d_j \leq \frac{2m}{j}$ for all j whence

$$\sum_{j=k+1}^{n-1} d_j^2 \leq 4m^2 \sum_{j=k+1}^{\infty} \frac{1}{j^2} < \frac{4m^2}{k},$$

i.e,

$$|\bar{C}| \geq mn - \sum_{i=1}^{k} d_i^2 - \frac{21n^2}{k},$$

since $m \leq 9n/4$. In addition,

$$\sum_{i=1}^{k} d_i^2 \leq d_1^2 + \left(\sum_{i=2}^{k} d_i\right) \cdot d_2 \leq d_1^2 + (m - n/2 - d_1) \cdot d_2 + 2^{k+1} n,$$

where the first step follows from the fact that $d_2 \geq d_i$ for $i \geq 2$, and the last step uses Claim 4.9 and $d_2 \leq n$.

We conclude that for $\omega_n(1) \leq k \leq o(\log n)$,

$$|\bar{C}| \geq mn - d_1^2 - (m - n/2 - d_1) \cdot d_2 - o(n^2). \tag{2}$$

We now normalize everything in terms of n, namely, write

$$m = \gamma \cdot n, \quad d_1 = x \cdot n, \quad d_2 = y \cdot n;$$
$$g(\gamma, x, y) := \gamma - x^2 - \gamma \cdot y + \frac{y}{2} + xy.$$

The optimization problem below is a normalized version of minimizing $|\bar{C}|$ subject to our assumptions on γ, $d_2 \leq d_1$, and Claims 4.7, 4.8, and 4.9.

Optimization problem A

Minimize $g(\gamma, x, y)$, subject to:

(1) $1 \leq \gamma \leq \frac{9}{4}$.
(2) $0 \leq y \leq x \leq \min(\frac{\gamma}{2}, 1)$.
(3) $x + y \leq \gamma - \frac{1}{2}$.
(4) $x + y \leq \frac{1+\gamma}{2}$.

This optimization problem is solved in the following Theorem whose proof is in Appendix A.

Theorem 4.10: *The answer to Optmization problem A is* $\min g(\gamma, x, y) = \frac{3}{4}$. *The optimum is attained in exactly two points* $(\gamma = 1, x = \frac{1}{2}, y = 0)$ *and* $(\gamma = 2, x = 1, y = \frac{1}{2})$.

Note that the deviation of the function $g(\gamma, x, y)$ from the lower bound on $|\bar{C}|/n^2$ of (2), as well as the deviations of the constraints from the assumptions and claims, are only of order $o(1)$. Therefore,

$$|\bar{C}| \geq \min g(\gamma, x, y) \cdot n^2 + o(n^2)$$

because g is continuous. In addition, plugging the optimal values on γ, x, y completes the proof of Theorem 4.3. The first case follows immedaitely, and for the second case we apply Claim 4.9 for $k = 3$ and use the optimal values $m = 2n + o(n)$, $d_1 = n + o(n)$, $d_2 = n/2 + o(n)$ to obtain that

$$d_3 \leq m - n/2 - d_1 - d_2 + 16 = o(n). \quad \blacksquare$$

5. Large d-hypercuts over $\mathbb{F}_2$ in even dimensions

In this section we consider large d-dimensional hypercuts over $\mathbb{F}_2$ for $d > 2$. We show that if d is even, the largest d-hypercuts have $\binom{n}{d+1}(1 - o_n(1))$ d-faces. In contrast, for odd d we observe that the density of every d-hypercut is bounded away from 1.

Theorem 5.1: *For every even integer* $d \geq 2$ *there exists an* n*-vertex* d*-hypercut over* $\mathbb{F}_2$ *with* $\binom{n}{d+1}(1 - o_n(1))$ d*-faces.*

Before we prove the theorem, let us explain the difference between odd and even dimensions. Recall that Turan's problem (e.g., [7]) asks for the largest density $ex(n, K_{d+2}^{d+1})$ of a $(d+1)$-uniform hypergraph that does not contain all the $d + 2$ hyperedges on any set of $d + 2$ vertices. This problem is still open for all $d > 1$. For d odd, the support C of a d-dimensional coboundary, and d-hypercuts in particular, has this Turan property. Indeed, since $C \cdot \partial_{d+1} = 0$ (because $\partial_d \partial_{d+1} = 0$), every $d + 2$ vertices span an even number of members from C. A simple double-counting argument shows that the density of C cannot exceed $1 - \frac{1}{d+2}$, and in fact, a better upper bound of $1 - \frac{1}{d+1}$ is known [14]. One of the known constructions for the Turan problem yields d-coboundaries with density $\frac{3}{4} - \frac{1}{2^{d+1}} - o(1)$ for d odd [2]. In particular, for $d = 3$ this gives a lower bound of $\frac{11}{16} = 0.6875$. In [13] an upper bound of 0.6917 was found using flag algebras.

We now turn to prove the theorem for $d \geq 4$. We saw in Section 2 that a d-hypercut C is an inclusion-minimal set of d-faces whose characteristic vector is a coboundary. In addition, every d-coboundary C and every vertex v satisfy $C = \mathrm{link}_v(C) \cdot \partial_d$. Recall that C is a 2-hypercut if and only if $\mathrm{link}_v(C)$ is Λ-connected for some vertex v. In dimension > 2 we do not have such a characterization, but as we show below, an appropriate variant of the sufficient condition for being a hypercut does apply in all dimensions.

Let τ, τ' be two $(d-1)$-faces in a $(d-1)$-complex K. We say that they are Λ-**adjacent** if their union $\sigma = \tau \cup \tau'$ has cardinality $d+1$, and τ, τ' are the only $(d-1)$-dimensional subfaces of σ in K. We say that K is Λ-connected if the transitive closure of the Λ-adjacency relation has exactly one class.

Claim 5.2: *Let C be a d-dimensional coboundary such that the $(d-1)$-complex $K = \mathrm{link}_v(C)$ is Λ-connected for some vertex v. Then C is a d-hypercut.*

Proof. Suppose that $\emptyset \neq C' \subsetneq C$ is a d-coboundary and let $K' = \mathrm{link}_v(C')$. Note that $\emptyset \neq K' \subsetneq K$ and therefore there are $(d-1)$-faces τ, τ' which are Λ-adjacent in K such that $\tau' \in K'$ and $\tau \notin K'$. Consider the d-dimensional simplex $\sigma = \tau \cup \tau'$. On the one hand, since exactly two of the facets of σ are in K, it does not belong to $C = K \cdot \partial_d$. On the other hand, it does belong to C' since exactly one of its facets (τ') is in K'. This contradicts the assumption that $C' \subset C$. ∎

Proof of Theorem 5.1. We start by constructing a random $(n-1)$-vertex $(d-1)$-dimensional complex K that has a full skeleton, where each $(d-1)$-face is placed in K independently with probability $p := 1 - n^{-\frac{1}{3d-3}}$. We show that with probability $1 - o_n(1)$ the complex K is Λ-connected, whence $C := K \cdot \partial_d$ is almost surely a d-hypercut of the desired density.

We actually show that K satisfies a condition that is stronger than Λ-connectivity. Namely, let $\tau, \tau' \in K$ be two distinct $(d-1)$ faces. We find π, π' where π is Λ-adjacent to τ and π' is Λ-adjacent to τ' and in addition the symmetric differences get smaller, $|\pi \oplus \pi'| < |\tau \oplus \tau'|$.. To this end we pick some vertices $u \in \tau \setminus \tau'$, $u' \in \tau' \setminus \tau$ and aim to show that with high probability there is some $x \notin \tau \cup \tau'$ for which the following event P_x holds:

$$\pi_x := \tau \cup \{x\} \setminus \{u\} \quad \text{and} \quad \pi'_x := \tau' \cup \{x\} \setminus \{u'\} \text{ are in } K, \quad \text{and}$$
$$\tau \text{ is } \Lambda\text{-adjacent to } \pi_x \quad \text{and} \quad \tau' \text{ is } \Lambda\text{-adjacent to } \pi'_x.$$

In other words, it is required that $\pi_x \in K$ and $\tau \cup \{x\} \setminus \{w\} \notin K$ for every $w \in \tau \setminus \{u\}$, and similarly for τ', π'_x. Therefore $\Pr(P_x) = p^2 \cdot (1-p)^{2d-2}$. Moreover, the events $\{P_x \mid x \notin \tau \cup \tau'\}$ are independent. Hence, the claim fails for some τ, τ' with probability at most

$$(1 - p^2 \cdot (1-p)^{2d-2})^{n-2d} = \exp\left[-\Theta(n^{1/3})\right].$$

The proof is concluded by taking the union bound over all pairs τ, τ'. ∎

6. Large shadowless 2-complexes over every field

The main result of this section is the construction of 2-complexes that are shadowless and 2-acyclic over every field $\mathbb{F}$. In comparison to our previous results, recall that assuming Artin's conjecture there are infinitely many shadowless 2-almost hypertrees over $\mathbb{Q}$. We also saw in Section 4 that every 2-almost-hypertree over $\mathbb{F}_2$ has a shadow and there we discussed its minimal cardinality. We now complement this by seeking the largest size of a shadowless acyclic complex. Our construction works at once for all fields since it is based on the combinatorial property of 2-collapsibility.

Theorem 6.1: *For every odd integer n, there exists a 2-collapsible 2-complex $A = A_n$ with $\binom{n-1}{2} - (n+1)$ 2-faces that remains 2-collapsible after the addition of any new 2-face. In particular, this complex is acyclic and shadowless over every field.*

Proof. The vertex set $V = V(A)$ is the additive group $\mathbb{Z}_n$. All additions here are done mod n. Edges in A are denoted $(x, x+a)$ with $a < n/2$, and such an edge is said to have **length** a. Also, for $a > 1$, $b = \lfloor \frac{a}{2} \rfloor$ is uniquely defined subject to $1 \le b < a < n/2$. For every $x \in \mathbb{Z}_n$ and $n/2 > a > 1$ we say that the edge $(x, x+a)$ **yields** the face

$$\rho_{x,a} := \left\{x, x+a, x+\left\lfloor \frac{a}{2} \right\rfloor\right\}$$

of **length** a. These are A's 2-faces:

$$\{\rho_{x,a} \mid n/2 > a \neq 1, 3,\ x \in \mathbb{Z}_n\}$$

It is easy to 2-collapse A by collapsing A's faces in decreasing order of their lengths. In each phase of the collapsing process, the longest edges in the remaining complex are exposed and can be collapsed.

It remains to show that the complex $A \cup \{\sigma\}$ is 2-collapsible for every face $\sigma = \{x, y, z\} \notin A$. To this end, let us carry out as much as we can of the "top-down" collapsing process described above. Clearly some of the steps of this process become impossible due to the addition of σ, and we now describe the complex that remains after all the possible steps of the previous collapsing process are carried out. Subsequently we show how to 2-collapse this remaining complex and conclude that $A \cup \{\sigma\}$ is 2-collapsible, as claimed.

For every $n/2 > a \geq 1$, $x \in \mathbb{Z}_n$ we define a subcomplex $C_{x,x+a} \subset A$. If $a = 1$ or 3, this is just the edge $(x, x + a)$. For all $n/2 > a \neq 1, 3$ it is defined recursively as $C_{x,\ x+\lfloor \frac{a}{2} \rfloor} \cup C_{x+\lfloor \frac{a}{2} \rfloor, x+a} \cup \{\rho_{x,a}\}$.

Note that $C_{x,y}$ is a triangulation of a polygon that is made up of the edge (x, y) and a path of edges of lengths 1 and 3 from x to y (see Figure 2).

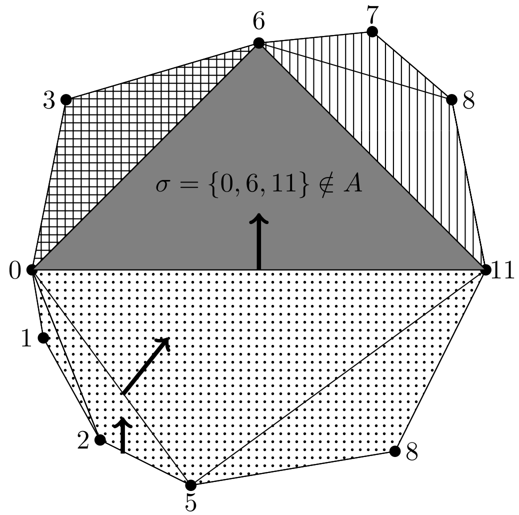

Figure 2. The complexes $C_{0,11}$ (dots), $C_{0,6}$ (grid), $C_{6,11}$ (lines) and $\sigma = \{0, 6, 11\} \notin A$ (filled). The vertex $v = 2$ appears only in $C_{0,11}$ and traversing the complex from the edge $e = \{2, 5\}$ enables a 2-collapse of σ as indicated by the arrows. Note that the edge $\{8, 11\}$ belongs to both $C_{6,11}$ and $C_{0,11}$ and is not exposed.

Our proof will be completed once we (i) observe that this remaining complex is

$$\Delta_\sigma := \{\sigma\} \cup C_{x,y} \cup C_{x,z} \cup C_{y,z},$$

and (ii) show that Δ_σ is 2-collapsible.

Indeed, toward (i), just follow the original collapsing process and notice that Δ_σ is comprised of exactly those faces in A that are affected by the introduction of σ into the complex.

We will show (ii) by proving that the face σ can be collapsed out of Δ_σ. Consequently, Δ_σ is 2-collapsible to a subcomplex of the 2-collapsible complex A, and is therefore 2-collpasible.

As we show below:

Claim 6.2: *There exists a vertex in Δ_σ which belongs to exactly one of the complexes $C_{x,y}, C_{x,z}$ or $C_{y,z}$.*

This allows us to conclude that the face σ can be collapsed out of Δ_σ. Say that the vertex v is in $C_{x,y}$ and only there, and let e be some edge of length 1 or 3 in $C_{x,y}$ that contains v. Follow the recursive construction of $C_{x,y}$ as it leads from (x, y) to e. Every edge that is encountered there appears only in the polygon $C_{x,y}$. By traversing this sequence in reverse, we collapse σ out of Δ_σ (see Figure 2). ∎

Proof of Claim 6.2. By translating mod n if necessary we may assume that $x = 0$ and $0 < y, z - y < \frac{n}{2}$. If $z > \frac{n}{2}$, then

$$V(C_{0,y}) \subseteq \{0, \ldots, y\}, \quad V(C_{y,z}) \subseteq \{y, \ldots, z\} \quad \text{and}$$
$$V(C_{z,0}) \subseteq \{z, \ldots, n-1, 0\},$$

so their vertex sets are nearly disjoint altogether.

We now consider the case $z < \frac{n}{2}$ and assume by contradiction that the claim fails for $\sigma = \{0, y, z\}$. We want to conclude that $\sigma \in A$, and in fact $\sigma \in C_{0,z}$. By the recursive construction of $C_{0,z}$, this, in other words, means that both edges $(0, y)$ and (y, z) are in $C_{0,z}$. We only prove that $(0, y) \in C_{0,z}$, and the claim $(y, z) \in C_{0,z}$ follows by an essentially identical argument.

Fix $0 < y < z < \frac{n}{2}$. We want to show that

$$\text{if } (0, y) \notin C_{0,z} \text{ then } V(C_{0,z}) \cap [0, y] \neq V(C_{0,y}). \tag{3}$$

Consequently, there is a vertex v in $[0, y]$ which belongs to exactly one of the complexes $C_{0,z}$ and $C_{0,y}$. If such a $v < y$ exists, we are done, since $C_{y,z}$ has no vertices in $[0, y-1]$. Otherwise,

$$V(C_{0,z}) \cap [0, y-1] = V(C_{0,y}) \setminus \{y\} \quad \text{and} \quad y \notin C_{0,z}.$$

But the vertices of $C_{0,z}$ form an increasing sequence from 0 to z with differences 1 or 3, so either $(y-2, y+1) \in C_{0,z}$ or $(y-1, y+2) \in C_{0,z}$. In the former case, both $y-2$ and y are vertices in $C_{0,y}$, and therefore $y-1 \in C_{0,y}$ and consequently $y-1 \in C_{0,z}$, contrary to the assumption that the edge $(y-2, y+1)$ is in $C_{0,z}$. In the latter case, $y+2 \in C_{0,z}$ and $y+1 \notin C_{0,z}$. Which vertex succeeds y in $C_{y,z}$? If $(y, y+1) \in C_{y,z}$ then $y+1$ belongs only to $C_{y,z}$. If $(y, y+3) \in C_{y,z}$ then $y+2$ is only in $C_{0,z}$.

We prove the implication (3) by induction on y. The base cases where $y = 1$ or $y = 3$ are straightforward. If $(0, \lfloor \frac{y}{2} \rfloor) \notin C_{0,z}$ then by induction $V(C_{0,z}) \cap [0, \lfloor \frac{y}{2} \rfloor] \neq V(C_{0,\lfloor \frac{y}{2} \rfloor})$. But $V(C_{0,y}) \cap [0, \lfloor \frac{y}{2} \rfloor] = V(C_{0,\lfloor \frac{y}{2} \rfloor})$ and the conclusion that

$$V(C_{0,z}) \cap [0, y] \neq V(C_{0,y})$$

follows. We now consider what happens if $(0, \lfloor \frac{y}{2} \rfloor) \in C_{0,z}$. Which edge has yielded the 2-face of $C_{0,z}$ that contains the edge $(0, \lfloor \frac{y}{2} \rfloor)$? It can be either $(0, 2 \cdot \lfloor \frac{y}{2} \rfloor)$ or $(0, 2 \cdot \lfloor \frac{y}{2} \rfloor + 1)$. But one of these two edges is $(0, y)$ which, by assumption, is not in $C_{0,z}$, so it must be the other one. Namely, either $y = 2r$ and $(0, 2r+1) \in C_{0,z}$ or $y = 2r+1$ and $(0, 2r) \in C_{0,z}$.

Let us deal first with the case $y = 2r$. Assume, in contradiction to (3), that

$$V(C_{0,z}) \cap [0, 2r] = V(C_{0,2r}).$$

In particular $V(C_{0,z}) \cap [r, 2r] = V(C_{0,2r}) \cap [r, 2r]$. But since $(0, 2r+1)$ is an edge of $C_{0,z}$ it also follows that $V(C_{0,2r+1}) \cap [r, 2r] = V(C_{0,z}) \cap [r, 2r]$. Therefore,

$$V(C_{0,2r+1}) \cap [r, 2r] = V(C_{0,2r}) \cap [r, 2r].$$

By the recursive construction of $C_{0,2r+1}$ and $C_{0,2r}$ we obtain that

$$V(C_{r,2r+1}) \cap [r, 2r] = V(C_{r,2r}).$$

By using the rotational symmetry of A we can translate this equation by r to conclude that $V(C_{0,r+1}) \cap [0, r] = V(C_{0,r})$. By induction, using the contrapositive of Equation (3) this implies that $(0, r) \in C_{0,r+1}$, hence $r = 1$. However, $C_{0,z}$ cannot contain both $(0, 3)$ and $(0, 1)$ so we are done.

The argument for $y = 2r+1$ is essentially the same and is omitted. ∎

7. Open problems

- There are several problems that we solved here for 2-dimensional complexes. It is clear that some completely new ideas will be required in order to answer these questions in higher dimensions. In particular, it would be interesting to extend the construction based on arithmetic triples for $d > 2$.
- An interesting aspect of the present work is that the behaviour over $\mathbb{F}_2$ and $\mathbb{Q}$ differ, sometimes in a substantial way. It would be of interest to investigate the situation over other coefficient rings.
- How large can an acyclic shadowless 2-complex over $\mathbb{F}_2$ be? Theorem 6.1 gives a bound, but we do not know the exact answer yet.
- Many basic (approximate) enumeration problems remain wide open. How many n-vertex d-hypertrees are there? What about d-collapsible complexes? A fundamental work of Kalai [6] provides some estimates for the former problem, but these bounds are not sharp. In one dimension there are exactly $\frac{(n-1)!}{2}$ inclusion-minimal n-vertex cycles. We know very little about the higher-dimensional counterparts of this fact.

Appendix A. Proof of Theorem 4.10

Let $h(\gamma, x, y) = g(\gamma, x, y) - \frac{3}{4} = \gamma - x^2 - \frac{3}{4} - y(\gamma - \frac{1}{2} - x)$. We need to show that $h \geq 0$ under the conditions of the Optimization problem. This involves some case analysis.

First note $\gamma - \frac{1}{2} - x \geq 0$ by condition 3, so that for fixed γ, x we have that h is a decreasing function of y. Thus, to minimize h, we need to determine the largest possible value of y.

(1) We first consider the range $\gamma \leq 2$. Here condition 4 is redundant, and $y \leq \min\{x, \gamma - \frac{1}{2} - x\}$.

 (a) We further restrict to the range $x \leq \frac{\gamma}{2} - \frac{1}{4}$, where $x \leq \gamma - \frac{1}{2} - x$, so the largest feasible value of y is $y = x$. Note that

$$h|_{y=x} = \gamma - \frac{3}{4} - x\Big(\gamma - \frac{1}{2}\Big).$$

 But $\gamma - \frac{1}{2} \geq 0$ by condition 1, so h is minimized by maximizing x, namely taking $x = \frac{\gamma}{2} - \frac{1}{4}$. This yields $h = \frac{1}{8} - \frac{1}{2}(\gamma - 1)(\gamma - 2)$ which is positive in the relevant range $2 \geq \gamma \geq 1$.

(b) In the complementary range $\frac{\gamma}{2} - \frac{1}{4} \le x$ the largest value for y is $y = \gamma - \frac{1}{2} - x$ which yields

$$\begin{aligned} h =& \gamma - x^2 - \frac{3}{4} - \Big(\gamma - \frac{1}{2} - x\Big)^2 \\ =& -2\Big(x - \frac{\gamma}{2}\Big)\Big(x - \frac{\gamma - 1}{2}\Big) - \frac{(\gamma-1)(\gamma-2)}{2}. \end{aligned}$$

It suffices to check that $h \ge 0$ at both extreme values of x, namely $\frac{\gamma}{2} - \frac{1}{4}$ and $\gamma/2$. Also $h = 0$ only at $x = \gamma/2$ with $\gamma = 1$ or 2.

(2) In the complementary range $\gamma \ge 2$, condition 3 is redundant and condition 4 takes over.

(a) Assume first that $x \le \frac{1+\gamma}{4}$, then $x \le \frac{1+\gamma}{2} - x$ and the extreme value for y is $y = x$. Again $h|_{y=x} = \gamma - \frac{3}{4} - x(\gamma - \frac{1}{2})$ and now the largest possible value of x is $x = \frac{1+\gamma}{4}$ which yields $h = \frac{(5-2\gamma)(\gamma-1)}{8}$. This is positive at the range $\frac{9}{4} \ge \gamma \ge 2$.

(b) When $x \ge \frac{1+\gamma}{4}$ the minimum h is attained at $y = \frac{1+\gamma}{2} - x$, so that

$$\begin{aligned} h =& \gamma - x^2 - \frac{3}{4} - \Big(\frac{1+\gamma}{2} - x\Big) \cdot \Big(\gamma - \frac{1}{2} - x\Big) \\ =& -2(x-1)\Big(x + 1 - \frac{3\gamma}{4}\Big) - \frac{1}{2}(\gamma - 2)\Big(\gamma - \frac{5}{2}\Big). \end{aligned}$$

For fixed γ it suffices to check that $h \ge 0$ at the two ends of the range $1 \ge x \ge \frac{1+\gamma}{4}$. At $x = 1$ we get $h = -\frac{1}{2}(\gamma-2)(\gamma - \frac{5}{2})$ which is nonnegative when $\frac{9}{4} \ge \gamma \ge 2$ with $h = 0$ only at $\gamma = 2$. When $x = \frac{1+\gamma}{4}$, we get $h = \frac{(5-2\gamma)(\gamma-1)}{8}$ which is positive for $\frac{9}{4} \ge \gamma \ge 2$.

To sum up, $h \ge 0$ throughout the relevant range with two points where $h = 0$, namely $\gamma = 2$, $x = 1$, $y = \frac{1}{2}$ and $\gamma = 1$, $x = \frac{1}{2}$, $y = 0$.

Appendix B. Proof of Theorem 4.1

Let us recall some of the facts proved in Section 4 concerning the largest n-vertex 2-hypercut C. Pick an arbitrary vertex v. Since C is a coboundary, it can be generated by an n-vertex graph which consists of the isolated vertex v, and $G = \mathrm{link}_v(C)$, an $(n-1)$-vertex Λ-connected graph. Similarly, $\bar{C}$ can be generated by the disjoint union of v and $\bar{G}$. As we saw, there exists some v for which the corresponding $\bar{G}$ satisfies either

Case (I): $m = n - 1 + o(n)$, $d_1 = \frac{n}{2} \pm o(n)$ and $d_2 = o(n)$, or

Case (II): $m = 2n \pm o(n)$, $d_1 = n - o(n)$, $d_2 = \frac{n}{2} \pm o(n)$ and $d_3 = o(n)$, where, as before, $m = |E(\bar{G})|$, $d_1 \geq d_2 \geq \cdots \geq d_{n-1}$ is the degree sequence of $\bar{G}$, with $d_i = d(v_i)$. We denote by t the number of triangles in $\bar{G}$. Since C is the largest cut, the graph $\bar{G}$ attains the minimum of $f(\bar{G}) = nm - \sum d_i^2 + 4t$ among all graphs whose complement is Λ-connected.

We now turn to further analyse the structure of $\bar{G}$, in Case (I).

Lemma B.1: *Suppose that $\bar{G}$ satisfies Case (I) and let $H = \bar{G} \setminus v_1$. Then H is either (i) a perfect matching, or (ii) a perfect matching plus an isolated vertex, or (iii) a perfect matching plus an isolated vertex and a 3-vertex path.*

Proof. The proof proceeds as follows: for every H other than the above, we find a local variant $\bar{G}_1$ of $\bar{G}$ with $f(\bar{G}_1) < f(\bar{G})$. We then likewise modify G_1 to G_2 etc., until for some $k \geq 1$ the graph G_k is Λ-connected. The process proceeds as follows.

For every connected component U of H of even size $|U| \geq 4$, we replace $H|_U$ with a perfect matching on U, and connect v_1 to one vertex in each of these $\frac{|U|}{2}$ edges. Now all connected components of H are either an edge or have an odd size.

Consider now odd-size components. Note that H can have at most one isolated vertex. Otherwise $\bar{G}$ is disconnected or it has clones, so that G is not Λ-connected. As long as H has two odd connected components which together have 6 vertices or more, we replace this subgraph with a perfect matching on the same vertex set, and connect v_1 to one vertex in each of these edges. If the remaining odd connected components are a triangle and an isolated vertex, remove one edge from the triangle, and connect v_1 only to one endpoint of the obtained 3-vertex path. In the last remaining case H has at most one odd connected component U.

If no odd connected components remain or if $|U| = 1$, we are done.

In the last remaining case H has a single odd connected component of order $|U| \geq 3$. We replace $H|_U$ with a matching of $(|U|-1)/2$ edges, connect v_1 to one vertex in each edge of the matching and to the isolated vertex. If, in addition, there is a connected component of order 2 with both vertices adjacent to v_1 (note that by the proof of Claim 4.6 there is at most one such component), we remove as well one edge between v_1 and this component.

All these steps strictly decrease f. We show this for the first kind of steps. The other cases are nearly identical.

Recall that $|E(H)| = \frac{n}{2} \pm o(n)$ and that H has at most one isolated vertex. Therefore every connected component in H has only $o(n)$ vertices. Let U be a connected component with $2u \geq 4$ vertices of which $0 < r \leq 2u$ are neighbours of v_1, and let $\beta = |E(H|_U)| - (2u-1) \geq 0$. Let $\bar{G}'$ be the graph after the aforementioned modification w.r.t. U. We denote its number of edges and triangles by m' and t' resp., and its degree sequence by d_i'. Then,

$$\begin{aligned} f(\bar{G}) - f(\bar{G}') =& n(m-m') - \sum_i (d_i^2 - d_i'^2) + 4(t-t') \\ \geq& n(\beta + r - 1) - (d_1^2 - (d_1 - r + u)^2) - \sum_{i \in U} d_i^2 \\ \geq& n(\beta + r - 1) + (u-r)(2d_1 + u - r) - 2u(4u - 2 + 2\beta + r). \end{aligned}$$

In the second row we use $t \geq t'$, which is true since the modification on U creates no new triangles. In the third row we use $\sum_{i\in U} d_i^2 \leq (\max_{i \in U} d_i)(\sum_{i \in U} d_i)$.

Let us express $d_1 = \frac{n-w}{2}$ where $w = o(n)$. What remains to prove is that

$$n(\beta + u - 1) + (u-r)(u-r-w) \geq 2u(4u - 2 + 2\beta + r).$$

Or, after some simple manipulation, and using the fact $r \leq 2u$, that

$$\beta n + (u-1)n \geq 4\beta u + u(7u + w + 3r - 4).$$

This is indeed so since $u = o(n)$ implies that $\beta n \gg \beta u$ and $2 \leq u \leq o(n)$ implies $(u-1)n \gg u(7u + w + 4r - 4)$.

The other cases are treated very similarly, with only minor changes in the parameters. In the case of two odd connected components which together have $2u \geq 6$ vertices, in the final step the main term is $(\beta + u - 2)n \geq n + \beta u$ since $u > 2$. In the case of changing a triangle to a 3-vertex path the main term in the final inequality is $(\beta + u - 1)n = n$. ∎

The structure of $\bar{G}$ for Case (I) is almost completely determined by Lemma B.1. Since G is Λ-connected, in $\bar{G}$ v_1 must have a neighbour in each component of H, and can be fully connected to at most one component. In addition, if P is a 3-vertex path in H, then v_1 has exactly one neighbour in P which is an endpoint. Otherwise we get clones. Therefore the only possible graphs are those that appear in Figure 3. The first row of the figure applies to odd n, where the optimal $\bar{G}$ satisfies $f(\bar{G}) = \frac{3}{4}n^2 - 4n + \frac{25}{4}$. The other rows correspond to n even, with four optimal graphs that satisfy $f(\bar{G}) = \frac{3}{4}n^2 - \frac{7}{2}n + 4$.

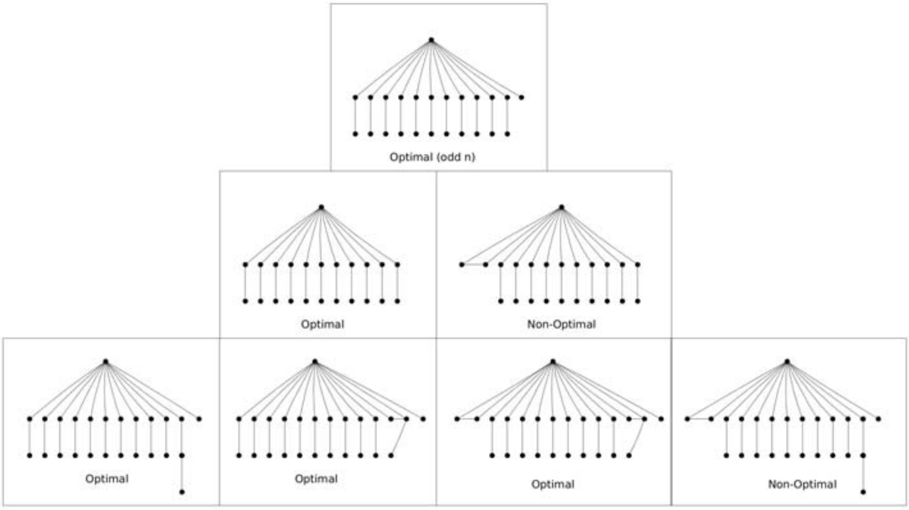

Figure 3. The graphs $\bar{G}$ that are considered in the final stage of the proof of Case (I). The first row refers to the only possibility for odd n. The second row to even n, where H is a perfect matching. The third row refers to even n, where H is a disjoint union of an isolated vertex, 3-path and a matching.

This concludes Case (I), and we now turn to Case (II). Our goal here is to reduce this back to Case (I), and this is done as follows.

Claim B.2: *Let $\bar{G} = \mathrm{link}_v(\bar{C})$ be a graph on $n-1$ vertices with parameters as in Case (II). If $H = \bar{G} \setminus \{v_1, v_2\}$ has an isolated vertex z that is adjacent in $\bar{G}$ to v_1, then $f(\bar{G})$ is bounded by the extremal examples found in Case (I).*

Proof. Let S be the star graph on vertex set $V \cup \{v\}$ with vertex v_1 in the center and $n-1$ leaves. Consider the graph

$$F := \bar{G} \oplus S$$

on the same vertex set, whose edge set is the symmetric difference of $E(S)$ and $E(\bar{G})$. Since every triplet meets S in an even number of edges, the coboundary that F generates equals the coboudnary that $\bar{G}$ generates, which is $\bar{C}$.

In addition, z is an isolated vertex in F since its only neighbour in $\bar{G}$ is v_1. Consequently, $F = \mathrm{link}_{\mathrm{z}}(\bar{C})$, and the claim will follow by showing that F agrees with the conditions of Case (I). Indeed, $\deg_F(v_1) = n-1-\deg_{\bar{G}}(v_1) = o(n)$, and $|\deg_F(u) - \deg_{\bar{G}}(u)| \leq 1$ for every other vertex u. Hence, $\deg_F(v_2) = \frac{n}{2} \pm o(n)$, and $\deg_F(u) = o(n)$ for every other vertex u. ∎

If H has no isolated vertex that is adjacent in $\bar{G}$ to v_1, we show how to modify $\bar{G}$ to a graph $\bar{G}_1$ such that (i) G_1 is Λ-connected, (ii) $\bar{G}_1 \setminus \{v_1, v_2\}$ has an isolated vertex which is adjacent to v_1 in $\bar{G}_1$, and (iii) $f(\bar{G}_1) < f(\bar{G})$.

Since G is Λ-connected and using the proof of Claim 4.6, H has at most one connected component U_1 in H where all vertices are adjacent to v_1 and not to v_2 in $\bar{G}$. Similarly, it has at most one connected component U_2 where all vertices are adjacent to both v_1 and v_2. Also, since $d_1 = n - o(n)$, $d_2 = \frac{n}{2} \pm o(n)$ and H has at most 3 isolated vertices, there exists an edge $xy \in E(H)$ such that $xv_1, xv_2, yv_1 \in E(\bar{G})$, but $yv_2 \notin E(\bar{G})$.

G_1 is constructed as follows:

(1) If neither component U_1 nor U_2 exists, remove the edge xy and the edge v_1v_2, if it exists. Otherwise, let $r := |U_1 \cup U_2|$.

(2) If r is even, replace it in H with a perfect matching on $u - 2$ vertices and two isolated vertices. Connect v_1 to every vertex in $U_1 \cup U_2$. Make v_2 a neighbour of one of the isolated vertices, and one vertex in each of the edges of the matching. Additionally, remove the edge v_1v_2 if it exists.

(3) If u is odd, replace it in H by a perfect matching on $u - 1$ vertices and one isolated vertex. Connect v_1 to every vertex in $U_1 \cup U_2$, and v_2 to one vertex in each edge of the matching.

The fact that the value of f decreased is shown similarly to the calculation in Case (I).

References

[1] A. Björner and G. Kalai, *An extended euler-poincaré theorem*, Acta Mathematica **161** (1988), 279–303.

[2] D. de Caen, D. L. Kreher and J. Wiseman, *On constructive upper bounds for the Turán numbers $T(n, 2r+1, r)$*, Congressus Numerantium **65** (1988), 277–280.

[3] R. Cordovil and B. Lindström, *Simplicial matroids*, in *Combinatorial Geometries*, Encyclopedia of Mathematics and its Applications, Vol. 29, Cambridge University Press, Cambridge, 1987, pp. 98–113.

[4] D. Dotterrer, L. Guth and M. Kahle, *2-Complexes with Large 2-Girth*, Discrete & Computational Geometry **59** (2018), 383–412.

[5] C. Hoffman, M. Kahle and E. Paquette, *The threshold for integer homology in random d-complexes*, Discrete & Computational Geometry **57** (2017), 810–823.

[6] G. Kalai, *Enumeration of $\mathbb{Q}$ -acyclic simplicial complexes*, Israel Journal of Mathematics **45** (1983), 337–351.

[7] P. Keevash, *Hypergraph Turán problems*, in *Surveys in Combinatorics 2011*, London Mathematical Society Lecture Note Series, Vol. 392, Cambridge University Press, Cambridge, 2011, pp. 83–140.

[8] N. Linial and Y. Peled, *On the phase transition in random simplicial complexes*, Annals of Mathematics **184** (2016), 745–773.

[9] T. Łuczak and Y. Peled, *Integral homology of random simplicial complexes*, Discrete & Computational Geometry **59** (2018), 131–142.

[10] R. Mathew, I. Newman, Y. Rabinovich and D. Rajendraprasad, *Boundaries of hypertrees, and Hamiltonian cycles in simplicial complexes*, preprint, arXiv:1507.04471.

[11] P. Moree, *Artin's primitive root conjecture—a survey*, Integers **12** (2012), 1305–1416.

[12] I. Newman and Y. Rabinovich, *On multiplicative λ-approximations and some geometric applications*, SIAM Journal on Computing **42** (2013), 855–883.

[13] Y. Peled, *Combinatorics of simplicial cocycles and local distributions in graphs*, Master's thesis, The Hebrew University of Jerusalem, 2012.

[14] A. Sidorenko, *The method of quadratic forms and Turán's combinatorial problem*, Moscow University Mathematics Bulletin **37** (1982), 1–5.