



In Search of Hyperpaths

Amir Dahari¹ · Nathan Linial¹

Received: 20 November 2020 / Revised: 30 August 2021 / Accepted: 13 September 2021

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

Abstract

Hypertrees are high-dimensional counterparts of graph theoretic trees. They have attracted a great deal of attention by various investigators. Here we introduce and study *hyperpaths*—a particular class of hypertrees which are high dimensional analogs of paths in graph theory. A d -dimensional hyperpath is a d -dimensional hypertree in which every $(d - 1)$ -dimensional face is contained in at most $(d + 1)$ faces of dimension d . We introduce a possibly infinite family of hyperpaths for every dimension, and investigate its properties in greater depth for dimension $d = 2$.

Keywords Hypertrees · Simplicial complexes · High dimensional combinatorics · Matrix multiplication · Linear algebra · Finite fields

Mathematics Subject Classification 05E45 · 05A19 · 65F05

1 Introduction

Hypertrees were defined in 1983 by Kalai [7]. An n -vertex d -dimensional hypertree $X = (V, E)$ is a \mathbb{Q} -acyclic d -dimensional simplicial complex with a full $(d - 1)$ -dimensional skeleton and $\binom{n-1}{d}$ faces of dimension d . Note that when $d = 1$ this coincides with the usual notion of a tree in graph theory. Also note that a hypertree is completely specified by its list of d -dimensional faces. There is already a sizable literature, e.g., [11, 12] dealing with hypertrees, but many basic questions in this area

Editor in Charge: János Pach

Supported by BSF US-Israel Grant 2018313 “Between topology and combinatorics”.

Amir Dahari
amir.dahari@mail.huji.ac.il

Nathan Linial
nati@cs.huji.ac.il

¹ School of Computer Science and Engineering, Hebrew University, 91904 Jerusalem, Israel

are still open. Also, in order to develop an intuition for these complexes it is desirable to have a large supply of different constructions. Many investigations in this area are done with an eye to the one-dimensional situation. Arguably the simplest-to-describe (1-dimensional) trees are stars and paths. These two families of trees are also the two extreme examples with respect to various natural graph properties and parameters such as the tree's diameter. Hyperstars are very easy to describe in any dimension d . Namely, we pick a vertex $v \in V$ and put a d -face σ in E iff $v \in \sigma$. On the other hand, it is much less obvious how to define d -dimensional paths. A one-dimensional path is a tree in which every vertex has degree at most 2. Working by analogy we can define a d -dimensional hyperpath as a d -dimensional hypertree in which every $(d - 1)$ -dimensional face is contained in no more than $(d + 1)$ faces of dimension d . We include a summary of the main results presented in this paper:

1. We introduce an infinite family of d -dimensional algebraically-defined simplicial complexes (see Definition 1.2). In dimension $d = 2$ we analyzed fairly large (up to $n \sim 1400$) such complexes most of which turned out to be 2-dimensional hyperpaths. To this end we devised a new fast algorithm that determines whether a matrix with circulant blocks is invertible.
2. Partial characterization: We showed that infinitely many of the 2-dimensional complexes discussed in 1. are *not* \mathbb{Q} -acyclic.
3. We develop several approaches for proving positive results and finding an infinite family of 2-dimensional hyperpaths.

Note 1.1 The necessary background in simplicial combinatorics and in number theory are introduced in Sect. 2.

Definition 1.2 Let \mathbb{F}_n be the field of prime order n . For $c \in \mathbb{F}_n$ and $d \geq 1$ an integer, we define the complex $X = X_{d,n,c}$ on vertex set \mathbb{F}_n . It has a full $(d - 1)$ -dimensional skeleton, and $\{x_0, x_1, \dots, x_d\}$ is a d -face in X iff $cx_d + \sum_{j=0}^{d-1} x_j \equiv 0 \pmod{n}$.¹

Ours is by no means the only sensible definition of a hyperpath. An alternative approach is described in [13]. That paper starts from the observation that a (1-dimensional) path is characterized as a tree that can be made a spanning cycle by adding a single edge. In this view they define a Hamiltonian d -cycle as a simple d -cycle of size $\binom{n-1}{d} + 1$. A Hamiltonian d -dimensional hyperpath is defined as a d -dimensional hypertree which can be made a Hamiltonian d -cycle by adding a single d -face. Other possibilities suggest themselves. For example, when an edge is added to a tree a single cycle is created. One may wonder how the length of this cycle is distributed when the added edge is chosen randomly. A path is characterized as the tree for which the average of this length is maximized. Similar notions clearly make sense also for $d > 1$. These various definitions coincide when $d = 1$ but disagree for $d > 1$. It would be interesting to understand the relations between these different definitions. Note that sum complexes [9] as well as certain hypertrees from [12] are hyperpaths according to our definition.

¹ Throughout this paper, unless stated otherwise, given a prime n , all arithmetic equations are mod n , and we often replace the congruence relation \equiv by an equality sign when no confusion is possible.

Running Example

We repeatedly return throughout the paper to the example corresponding to $d = 2$, $n = 13$, $c = 5$. Some of the 2-faces in $X_{2,13,5}$ are $\{0, 1, 5\}$, $\{2, 3, 12\}$, $\{2, 9, 3\}$ since $0 + 1 + 5 \cdot 5 = 2 + 3 + 5 \cdot 12 = 2 + 9 + 5 \cdot 3 = 0$. As the next claim shows the number of 2-faces in $X_{2,13,5}$ is $\binom{12}{2} = 66$ (out of the total of $\binom{13}{3} = 286$ 2-faces).

Given the vertex set $\{x_0, x_1, \dots, x_d\}$ of a d -face as in Definition 1.2, and if $c \neq 1$, the coordinate that is multiplied by c is uniquely defined. For if $cx_d + \sum_{j=0}^{d-1} x_j = cx_0 + \sum_{j=1}^d x_j = 0$, then $(c - 1) \cdot x_d = (c - 1) \cdot x_0$. This is impossible, since we are assuming that all x_i are distinct, and $c \neq 1$.

Claim 1.3 For an integer $d \geq 1$, a prime n and $c \in \mathbb{F}_n$, if $c \neq -d, 1$ then $X = X_{d,n,c}$ has exactly $\binom{n-1}{d}$ d -faces. If $c = 1$ then $X = X_{d,n,1}$ has exactly $\binom{n-1}{d}/(d+1)$ d -faces.

Proof By induction on d . Let us start with $d = 1$. If $c \neq 0$, then for every $x \neq 0$, there is a unique $y \neq 0$ s.t. $x + c \cdot y = 0$. Also, $x \neq y$, since by assumption $c \neq -1$ ($= -d$). This yields $\binom{n-1}{1} = n - 1$ edges, unless $c = 1$ in which case every edge is counted twice with a total of $\binom{n-1}{1}/2$ different edges. When $c = 0$, the complex has $n - 1$ edges, namely, $\{0, y\}$ for all $y \neq 0$. We proceed to deal with $d > 1$,

- If $c = 0$, then $\{x_0, \dots, x_{d-1}, y\}$ is a d -face iff $\sum_{i=0}^{d-2} x_i + 1 \cdot x_{d-1} = 0$. By the induction hypothesis with $c = 1$ and dimension $d - 1$ there are exactly $\binom{n-1}{d-1}/d$ such different choices of $\{x_0, \dots, x_{d-1}\}$. For every such choice of $\{x_0, \dots, x_{d-1}\}$ there are $n - d$ choices for y , namely, any value not in $\{x_0, \dots, x_{d-1}\}$, yielding a total of

$$\frac{1}{d} \binom{n-1}{d-1} \cdot (n-d) = \binom{n-1}{d}$$

distinct d -faces.

- If $c \neq 0$, for each of the $\binom{n}{d}$ $(d - 1)$ -faces $\{x_0, \dots, x_{d-1}\}$ there is a unique y satisfying $\sum_{i=0}^{d-1} x_i + c \cdot y = 0$. This gives a d -face, unless $y \in \{x_0, \dots, x_{d-1}\}$. By reordering the x_i if necessary $y = x_{d-1}$, which yields

$$\sum_{i=0}^{d-2} x_i + (c + 1) \cdot x_{d-1} = 0.$$

Since $c \neq 0, -d$, we know that $c + 1 \neq 1, -(d - 1)$ and we can apply the induction hypothesis for dimension $d - 1$ to obtain $\binom{n-1}{d-1}$ such different choices of $\{x_0, \dots, x_{d-1}\}$. All told there are

$$\binom{n}{d} - \binom{n-1}{d-1} = \binom{n-1}{d}$$

distinct d -faces in $X_{d,n,c}$. If $c = 1$, y has no special role and we over-count by a multiple of $d + 1$, hence we get only $\binom{n-1}{d}/(d + 1)$ distinct d -faces. \square

Henceforth to simplify matters we assume $c \notin \{0, 1, -d, -1\}$. In the 1-dimensional case the resulting graph $G = (V, E)$ has $V = \mathbb{F}_n$ and $E = \{\{x, -x/c\} \mid x \neq 0\}$. Consequently G is the union of $(n - 1)/o(-c^{-1})$ circles of length $o(-c^{-1})$. We will later see that $o(c)$ plays a crucial rule in determining whether $X_{2,n,c}$ is a hypertree.

We note that if $X = X_{d,n,c}$ is a hypertree, then it is a hyperpath, since every $(d - 1)$ -face σ in X is contained in at most $d + 1$ of its d -faces. Indeed, let $y \notin \sigma$ be the vertex that is added to σ to form a d -face. Then either $\sum_{i=0}^{d-1} x_i + c \cdot y = 0$ or there is an index $d - 1 \geq j \geq 0$ such that $cx_j + \sum_{i \neq j} x_i + y = 0$. Let us illustrate this observation with our running example:

Running Example

The edge $\{1, 5\}$ in the complex $X_{2,13,5}$ is included in the faces $\{5, 1, 4\}$, $\{0, 1, 5\}$ and $\{5, 3, 1\}$ with the convention that the last vertex is multiplied by $c = 5$. In contrast $\{1, 4\}$ is included in only two faces, namely $\{1, 4, 12\}$ and $\{5, 1, 4\}$. The equality $4 + 4 + 5 \cdot 1 = 0$ yields the non-face $\{4, 4, 1\}$.

Since we focus mostly on the 2-dimensional case $d = 2$, we use the shorthand $X = X_{2,n,c}$. The boundary operator of X is given by an $\binom{n}{2} \times \binom{n-1}{2}$ matrix which we denote by $A = A_{n,c}$. Clearly X is a hyperpath iff A has a full column rank, and indeed our main technical question is:




Fig. 1 Data on $X = X_{2,n,c}$ for all primes $11 \leq n \leq 59$. A yellow entry means that X is a hypertree. White entries show the (positive) co-dimension of the column space of $A = A_{n,c}$. Red indicates an illegal c , i.e., $c \equiv -2$ or $c \equiv n$

Question 1.4 For which primes n and $c \in \mathbb{F}_n$ is $X = X_{2,n,c}$ a hypertree?

Figure 1 shows the answer for Question 1.4 for all primes $11 \leq n \leq 59$ and each appropriate c . Figure 3 shows the fraction of eligible c for which $X = X_{2,n,c}$ is a hypertree for all primes $11 \leq n \leq 1373$.

The paper is structured as follows: in Sect. 2 we discuss some preliminary facts and outline the necessary background in number theory and simplicial combinatorics. It turns out that the problem whether the complex X is acyclic reduces to the question whether a certain matrix with a special structure is invertible. We study this special structure in Sect. 3, explain the reduction and give a new fast algorithm that determines if a matrix of this kind is invertible. In Sect. 4 we further investigate this reduction. This allows us to exhibit in Sect. 5 an infinite family of non-acyclic 2-dimensional complexes. We conjecture that in our construction, a certain simple criterion asymptotically determines if a complex is acyclic or not. Section 6 is devoted to another approach in search of an infinite family of 2-dimensional hyperpaths.

2 Preliminaries

Many matrices are defined throughout this paper. They are marked throughout by hyperlinks that can return the reader to their definitions.

2.1 Some Relevant Number Theory

For a prime n , we denote by $\mathbb{F}_n = \{0, 1, \dots, n - 1\}$ the field with n elements. Addition and multiplication are done mod n . The multiplicative group of \mathbb{F}_n is comprised of the set $\mathbb{F}_n^* = \{1, \dots, n - 1\}$. It is a cyclic group isomorphic to $\mathbb{Z}/(n - 1)\mathbb{Z}$. The order, $o(x)$ of $x \in \mathbb{F}_n^*$ is the smallest positive integer r s.t. $x^r = 1$. The following easy lemma gives the orders of x 's powers:

Lemma 2.1 *If n is prime and $x \in \mathbb{F}_n^*$, then for every integer j*

$$o(x^j) = \frac{o(x)}{\gcd(j, o(x))}.$$

Recall Euler's totient function φ . Namely, $\varphi(t)$ is the number of integers in $\{1, \dots, t - 1\}$ that are co-prime with t . It is also the order of the multiplicative group mod t .

Clearly x is a generator of \mathbb{F}_n^* iff $o(x) = |\mathbb{F}_n^*| = n - 1$. By the above comments, \mathbb{F}_n^* has exactly $\varphi(n - 1)$ generators. We write logarithms w.r.t. some fixed generator λ of \mathbb{F}_n^* . That is, $\log_\lambda u = \log u$ is the unique $k \in \{0, \dots, n - 2\}$ for which $\lambda^k = u$.

2.2 Background on Simplicial Combinatorics

We follow the setup in [10, Chapter 2]. All simplicial complexes considered here have vertex set $V = \{0, \dots, n - 1\} = \mathbb{F}_n$. 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 $\sigma \in X$. The size $|X|$ of a d -complex X is the number of its d -faces.

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. 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 ∂_d maps an oriented d -simplex $\sigma = (v_0, \dots, v_d)$ to the formal sum

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

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

We linearly extend the boundary operator to free \mathbb{Q} -sums of simplices. 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 $\partial_d z$.

A simple observation shows that the matrix ∂_d has rank $\binom{n-1}{d}$. We denote by A the submatrix of ∂_d restricted to the columns associated with d -faces of a d -complex X . We define $\text{rank}(X)$ to be $\text{rank}(A)$. The rational d -homology of X , denoted by $H_d(X; \mathbb{Q})$, is the right kernel of the matrix A . Elements of $H_d(X; \mathbb{Q})$ are called d -cycles. A d -hypertree X over \mathbb{Q} is a d -complex of size $\binom{n-1}{d}$ with a trivial rational d -dimensional homology. This means that the columns of the matrix A form a basis for the column space of ∂_d .

3 Matrices with Circulant Blocks (MCB)

We remind the reader that $A_{n,c}$ denote the submatrix of the matrix form of the boundary operator ∂_2 restricted to the columns associated with the 2-faces of $X_{2,n,c}$. It turns out that $A_{n,c}$ is closely related to a block matrix whose blocks are circulant matrices. So, we start our work on Question 1.4 by deriving a structure theorem for such matrices. This is what we do in the present section.

Recall that a circulant matrix $C \in M_r(\mathbb{Q})$ has the following form:

$$C = \begin{pmatrix} c_0 & c_{r-1} & & c_1 \\ c_1 & c_0 & & c_2 \\ & & \ddots & \ddots \\ c_{r-1} & & & c_1 & c_0 \end{pmatrix}.$$

Equivalently,

$$C = g(P) = c_0 \cdot I + c_1 \cdot P + c_2 \cdot P^2 + \dots + c_{r-1} \cdot P^{r-1}, \tag{3.1}$$

where $P \in M_r(\mathbb{Q})$ is the cyclic permutation matrix

$$P = \begin{pmatrix} 0 & 0 & & 0 & 1 \\ 1 & 0 & & & 0 \\ 0 & 1 & & & \\ & & \ddots & & \\ 0 & 0 & & 1 & 0 \end{pmatrix}. \tag{3.2}$$

Given positive integers r, t , we denote by $MCB_{r,t}(\mathbb{Q})$ (for Matrices with Circulant Blocks) the set of all matrices of the form

$$E = \begin{pmatrix} C_{0,0} & C_{0,1} & & C_{0,t-1} \\ C_{1,0} & C_{1,1} & & \\ & & \ddots & \\ C_{t-1,0} & & & C_{t-1,t-1} \end{pmatrix}, \tag{3.3}$$

where each $C_{i,j} \in M_r(\mathbb{Q})$ is a circulant matrix. When \mathbb{Q} is omitted, the matrices are over \mathbb{C} . This is not to be confused with the well-studied class of Circulant Block Matrices (CBM) [4]. Such a matrix is circulant as a block matrix, but its blocks need not be circulants. We can clearly express E as follows:

$$E(P) = \begin{pmatrix} g_{0,0}(P) & g_{0,1}(P) & & g_{0,t-1}(P) \\ g_{1,0}(P) & g_{1,1}(P) & & \\ & & \ddots & \\ g_{t-1,0}(P) & & & g_{t-1,t-1}(P) \end{pmatrix}, \tag{3.4}$$

where $g_{i,j}$ are polynomials of degree less than r as in (3.1). Since $P^r = 1$, we can view $g_{i,j}(P)$ as elements of the quotient polynomial ring

$$\mathcal{R} := \mathbb{Q}[P]/(P^r - 1). \tag{3.5}$$

Likewise, we think of $E(P)$ as a member in the matrix ring $M_t(\mathcal{R})$. Associated with every $z \in \mathbb{C}$ is a scalar $t \times t$ complex matrix $\underline{E}(z)$, viz.,

$$\underline{E}(z) = \begin{pmatrix} g_{0,0}(z) & g_{0,1}(z) & & g_{0,t-1}(z) \\ g_{1,0}(z) & g_{1,1}(z) & & \\ & & \ddots & \\ g_{t-1,0}(z) & & & g_{t-1,t-1}(z) \end{pmatrix}. \tag{3.6}$$

Theorem 3.1 A matrix $E \in \text{MCB}_{r,t}(\mathbb{Q})$ is singular iff $\underline{E}(\omega_k)$ is singular for some $k \mid r$. Here $\omega_k = \exp(2\pi i/k)$ is the primitive k -th root of unity.

Proof The proof of Theorem 3.1 uses the next claim:

Claim 3.2 Every $E \in \text{MCB}_{r,t}$ is similar to a block diagonal matrix with $t \times t$ blocks of the form

$$X \cdot E \cdot X^{-1} = \begin{pmatrix} \underline{E}(\omega_r^t) & 0 & & 0 \\ 0 & \underline{E}(\omega_r^{t-1}) & & 0 \\ & & \ddots & \\ 0 & 0 & & \underline{E}(\omega_r^1) \end{pmatrix}$$

for some invertible matrix X .

Proof of the claim We recall the order- r Discrete Fourier Transform (DFT) Matrix \mathcal{F}_r whose entries are $\mathcal{F}_r[k, l] = \exp(-2\pi ikl/r) = \omega_r^{-kl}$, where $\omega_r = \exp(2\pi i/r)$ is a primitive r -th root of unity. It diagonalizes $r \times r$ circulant matrices as follows by the next known lemma:

Lemma 3.3 Let \mathcal{F}_r be the order- r DFT matrix. If C is an $r \times r$ circulant matrix, then

$$\mathcal{F}_r \cdot C \cdot \mathcal{F}_r^{-1} = \Lambda,$$

where Λ is the diagonal matrix whose entries are C 's eigenvalues,

$$\lambda_j = c_0 + c_{r-1}\omega_r^j + c_{r-2}\omega_r^{2j} + \dots + c_1\omega_r^{(r-1)j} \quad \text{for } 0 \leq j \leq r-1.$$

Here $(c_0, c_1, \dots, c_{r-1})$ is C 's first column.

Proof $v_j = (1, \omega_r^j, \omega_r^{2j}, \dots, \omega_r^{(r-1)j})$ is an eigenvector of C with corresponding eigenvalue λ_j as

$$\begin{aligned} (C \cdot v_j)[i] &= c_i + c_{i-1}\omega_r^j + c_{i-2}\omega_r^{2j} + \dots + c_0\omega_r^{ij} + c_{r-1}\omega_r^{(i+1)j} + \dots + c_{i+1}\omega_r^{(r-1)j} \\ &= (c_0 + c_{r-1}\omega_r^j + c_{r-2}\omega_r^{2j} + \dots + c_1\omega_r^{(r-1)j}) \cdot \omega_r^{ij} = \lambda_j \omega_r^{ij}. \end{aligned}$$

The claim follows since $\mathcal{F}_r^{-1} = \mathcal{F}_r^*/r$. □

Set \mathcal{L} as a block diagonal matrix with \mathcal{F}_r on the diagonal. Since $E \in \text{MCB}_{r,t}$, Lemma 3.3 yields

$$\mathcal{L} \cdot E \cdot \mathcal{L}^{-1} = \begin{pmatrix} \Lambda_{0,0} & \Lambda_{0,1} & & \Lambda_{0,t-1} \\ \Lambda_{1,0} & \Lambda_{1,1} & & \\ & & \ddots & \\ \Lambda_{t-1,0} & & & \Lambda_{t-1,t-1} \end{pmatrix},$$

where $\Lambda_{k,l} = \text{diag}(\mathcal{F}_r \cdot c_{k,l})$ is a diagonal $r \times r$ matrix. Here $c_{k,l}$ is the first column vector of the circulant matrix $C_{k,l}$.

This matrix has rt rows and columns which we enumerate from 0 to $rt - 1$. It is made up of $r \times r$ blocks, with t of them at every layer. This suggests that indices in this matrix be written as $\alpha r + \beta$ for some $t > \alpha \geq 0$ and $r > \beta \geq 0$, which we interpret as index β within block number α . We rearrange the matrix to be made up of $t \times t$ blocks, with r of them at every layer, with indices in the form $\gamma t + \delta$ where $r > \gamma \geq 0$ and $t > \delta \geq 0$. Thus the mapping

$$\varphi: \alpha r + \beta \mapsto \beta t + \alpha \tag{3.7}$$

is a permutation which we apply to the rows and columns of $\mathcal{L} \cdot E \cdot \mathcal{L}^{-1}$. Since all blocks in $\mathcal{L} \cdot E \cdot \mathcal{L}^{-1}$ are diagonal, the entry in position $(\alpha r + \beta, \alpha' r + \beta')$ is nonzero only if $\beta = \beta'$. Following the application of φ , the matrix becomes an $r \times r$ diagonal matrix of $t \times t$ blocks,

$$Q \cdot \mathcal{L} \cdot E \cdot \mathcal{L}^{-1} \cdot Q^{-1} = \begin{pmatrix} \Delta_0 & 0 & & 0 \\ 0 & \Delta_1 & & 0 \\ & & \ddots & \\ 0 & 0 & & \Delta_{r-1} \end{pmatrix},$$

where Q is the permutation matrix of φ . The matrix thus becomes an $r \times r$ block matrix, with block size $t \times t$, and

$$\Delta_i[k, l] = \Lambda_{k,l}[i, i]$$

since the mapping (3.7) sends

$$\begin{aligned} \text{row number } k \cdot r + i &\mapsto \text{row number } i \cdot r + k, \\ \text{column number } l \cdot r + i &\mapsto \text{column number } i \cdot r + l. \end{aligned}$$

To complete the proof of Claim 3.2, setting $X = Q \cdot \mathcal{L}$, it only remains to show that $\Delta_i = \underline{E}(\omega_r^{-i})$:

$$\begin{aligned} \Delta_i[k, l] &= \Lambda_{k,l}[i, i] = (\mathcal{F}_r \cdot c_{k,l})[i] = \sum_{j=0}^{r-1} \mathcal{F}_r[i, j] c_{k,l}[j] \\ &= \sum_{j=0}^{r-1} c_{k,l}[j] \cdot \omega_r^{-ij} = \sum_{j=0}^{r-1} c_{k,l}[j] \cdot (\omega_r^{-i})^j = g_{k,l}(\omega_r^{-i}). \quad \square \end{aligned}$$

Claim 3.2 yields one part of Theorem 3.1. Namely, that if $\underline{E}(\omega_k)$ is singular for some $E \in \text{MCB}_{r,t}$ and some $k | r$, then E is singular. In order to prove the other direction of Theorem 3.1, we need the following two lemmas. Recall from (3.5) that for a matrix E in $\text{MCB}_{r,t}$, $E(P)$ is a polynomial matrix in $M_t(\mathcal{R})$ over the quotient

polynomial ring \mathcal{R} . Its determinant $\det(E(P))$ is a polynomial in \mathcal{R} , and we denote by $\det(E(P))(z)$ the evaluation of this polynomial at the complex number $z \in \mathbb{C}$.

Lemma 3.4 $\det(E(P))(\omega_r^j) = \det(\underline{E}(\omega_r^j))$ for every $E \in \text{MCB}_{r,t}$.

Proof

$$\begin{aligned} \det(E(P))(\omega_r^j) &= \left(\sum_{\sigma \in S_t} \left(\prod_{i=0}^t g_{i,\sigma(i)}(P) \right) \right) (\omega_r^j) \\ &= \sum_{\sigma \in S_t} \left(\prod_{i=0}^t g_{i,\sigma(i)}(P) \right) (\omega_r^j) \stackrel{(a)}{=} \sum_{\sigma \in S_t} \prod_{i=0}^t g_{i,\sigma(i)}(P)(\omega_r^j) \\ &= \sum_{\sigma \in S_t} \prod_{i=0}^t g_{i,\sigma(i)}(\omega_r^j) = \det(\underline{E}(\omega_r^j)). \end{aligned}$$

Equality (a) holds, because $P^r = (\omega_r^j)^r = 1$. □

The next lemma appears without proof in [16]. We provide a proof, since we could not find it in the literature:

Lemma 3.5 *The non-singular matrices in $\text{MCB}_{r,t}$ form a group w.r.t. matrix multiplication.*

Proof Clearly $\text{MCB}_{r,t}$ is closed under product, since the product of two circulant matrices is circulant, and matrix multiplication respects block product. We only need to show closure under inverse for invertible matrices in $\text{MCB}_{r,t}$.

As mentioned above, $M_t(\mathcal{R})$ and $\text{MCB}_{r,t}$ are in one-to-one correspondence. An inverse of $E(P)$ as in (3.4) in $M_t(\mathcal{R})$ is an inverse under this bijection of E in $\text{MCB}_{r,t}$, so it remains to prove that if E is invertible, then $E(P)$ has an inverse in $M_t(\mathcal{R})$.

The determinant of a matrix over a commutative ring is defined as usual as the alternating sum of products over permutations. Such a matrix has an inverse iff its determinant is invertible, as an element of the underlying ring. The proof of this fact (see, e.g., [14, §1.D]) goes by establishing the Cauchy–Binet formula for matrices over commutative rings. We apply this to the commutative polynomial ring \mathcal{R} , and conclude that $E(P)$ has an inverse in $M_t(\mathcal{R})$ iff its determinant (which is also a polynomial in \mathcal{R}) is invertible.

To prove the lemma, let $E \in \text{MCB}_{r,t}$ be invertible. Recall the definitions of $E(P)$ in (3.4) and $\underline{E}(\omega_r^j)$ in (3.6). By Claim 3.2 and since E is invertible, for all j we obtain

$$\det(\underline{E}(\omega_r^j)) \neq 0.$$

Using Lemma 3.4 this translates into

$$\det(E(P))(\omega_r^j) \neq 0$$

for all j , thus $\det(E(P))$ and $P^r - 1$ do not share any root and

$$\gcd(\det(E(P)), P^r - 1) = 1,$$

so the determinant is invertible, finishing the proof of the lemma. □

With Lemmas 3.4 and 3.5 we can complete the proof of Theorem 3.1. Let $E \in \text{MCB}_{r,t}(\mathbb{Q})$ be singular. By Lemma 3.5, E has no inverse in $\text{MCB}_{r,t}(\mathbb{Q})$, implying that $E(P)$ has no inverse in $M_t(\mathcal{R})$. Consequently, $\det(E(P))$ is not invertible and thus $\det(E(P))$ and $P^r - 1$ have a non-trivial common divisor. But

$$P^r - 1 = \prod_{k|r} \Psi_k(P)$$

where

$$\Psi_k(P) = \prod_{\substack{1 \leq l \leq k \\ \gcd(l,k)=1}} (P - \omega_k^l)$$

is the k -th cyclotomic polynomial. It is a well-known fact that Ψ_k is irreducible over \mathbb{Q} (e.g., [5]). Therefore, $\det(E(P))$ and $P^r - 1$ have a non-trivial common divisor iff one of the cyclotomic polynomials divides the determinant, i.e., there exists a divisor $k | r$ s.t.

$$\Psi_k(P) | \det(E(P)).$$

Since $\Psi_k(\omega_k) = 0$, there exists a divisor k of r s.t. $\det(E(P))(\omega_k) = 0$. By Lemma 3.4 this implies

$$\det(\underline{E}(\omega_k)) = 0,$$

completing the proof of Theorem 3.1. □

3.1 Computational Aspects of MCB

Theorem 3.1 has interesting computational aspects. In order to present them, we need some preparations. We recall that $d(m)$ denotes the number of distinct divisors of the integer m , and that for every $\epsilon > 0$, $d(m) = O(m^\epsilon)$. More precisely (see [6]),

$$\limsup_{m \rightarrow \infty} \frac{\log d(m)}{\log m / \log \log m} = \log 2.$$

Note 3.6 It is a classical fact (e.g., [15]) that matrix multiplication and matrix inversion have essentially the same computational complexity. It is, however, still unknown if

these problems are also equivalent to the decision problem whether a given matrix is invertible (e.g., [2, pp. 35–36]).

The smallest exponent for matrix multiplication is commonly denoted by ω . This is the least real number such that two $n \times n$ matrices can be multiplied using $O(n^{\omega+\epsilon})$ arithmetic operations for every $\epsilon > 0$. Presently the best known bounds [8] are $2 \leq \omega \leq 2.373$.

Proposition 3.7 *For every $\epsilon > 0$, it is possible to determine in time*

$$O(r^{1+\epsilon} \cdot t^2 + r^\epsilon \cdot t^\omega)$$

whether a matrix in $MCB_{r,t}(\mathbb{Q})$ is invertible.

Note 3.8 It follows that when $r \rightarrow \infty$, it is easier to decide the invertibility of matrices in $MCB_{r,t}(\mathbb{Q})$ than general $rt \times rt$ matrices, because

$$(rt)^\omega \gg r^{1+\epsilon} \cdot t^2 + r^\epsilon \cdot t^\omega.$$

An algorithm that decides invertibility for matrices in $MCB_{r,t}(\mathbb{Q})$ clearly needs the matrix to be fully specified, yielding a lower bound of $\Omega(r \cdot t^2)$.

Proof of Proposition 3.7 The proof of Theorem 3.1 yields an algorithm to decide if $E \in MCB_{r,t}(\mathbb{Q})$ is invertible:

1. Produce the matrix $E(P)$ as in (3.4).
2. For each divisor k of r :
 - a. Calculate the matrix $\underline{E}(\omega_k)$ as in (3.6) by evaluating the polynomial matrix $E(P)$ with $\omega_k = \exp(2\pi i/k)$.
 - b. Determine if the $t \times t$ matrix $\underline{E}(\omega_k)$ is invertible. If it is singular, return ‘ E is singular’.
3. If $\underline{E}(\omega_k)$ has full rank for every divisor $k \mid r$, return ‘ E is invertible’.

A circulant block is clearly completely defined by its first row. Therefore the matrix $E(P)$ can be found in time $O(r \cdot t^2)$. To find all the divisors of r we can even factor r using Erasthones’ sieve with no harm to the complexity. Step 2 is repeated $d(r) = r^{o(1)}$ times. Horner’s Rule allows to evaluate a degree r polynomial with r additions and r multiplications, so step 2a takes time $O(r \cdot t^2)$. The running time of step 2b is at most $O(t^\omega)$. All told the combined running time is

$$O(r \cdot t^2 + r + d(r) \cdot (r \cdot t^2 + t^\omega)) = O(r^{1+\epsilon} \cdot t^2 + r^\epsilon \cdot t^\omega)$$

for every $\epsilon > 0$. □

In step 2 we need to check whether $\underline{E}(\omega_k)$ is invertible for different k . These calculations can clearly be done in parallel. In [17] an iterative algorithm is presented to

invert a matrix in $MCB_{r,t}$, with run time

$$\begin{cases} O(2^{3l} \cdot r + t^2 \cdot r^2) & r \text{ is not a power of } 2, \\ O(2^{3l} \cdot r + t^2 \cdot r \log r) & r \text{ is a power of } 2, \end{cases} \tag{3.8}$$

where $l := \lceil \log_2 t \rceil$. If we only need to decide whether the matrix is invertible, then the algorithm in Proposition 3.7 is faster. We still do not know whether Theorem 3.1 yields an algorithm to invert a matrix in $MCB_{r,t}(\mathbb{Q})$ that is faster than the algorithm from [17].

3.2 From A to MCB

It turns out that there is a rank-preserving transformation of the boundary operator $A = A_{n,c}$ as defined in Sect. 2.2 into a matrix in $MCB_{n-1,(n-3)/2}$ as defined in (3.3). The transformation is fairly simple and only involves reordering of the rows and columns plus removal of $n - 1$ rows that are linearly dependent on the other rows and Gaussian elimination of $(n - 1)/2$ rows and columns.

As mentioned in Sect. 2.2, and maintaining the same terminology, the boundary operator $A = A_{n,c}$ of $X = X_{2,n,c}$ is given by an $\binom{n}{2} \times \binom{n-1}{2}$ matrix. We next find a square $\binom{n-1}{2} \times \binom{n-1}{2}$ submatrix of A of rank $\text{rank}(A)$. To this end, we remove $n - 1$ rows of A which are linearly spanned by the other $\binom{n-1}{2}$ rows. Rows in A are indexed by edges (1-dimensional faces). It is well known and easy to prove that this is the case with any $n - 1$ rows that represent the edge set of a spanning tree of the complete graph K_n . We apply this with the star rooted at vertex 0. In other words, we remove the rows corresponding to pairs $\{0, j\} \mid j \in \mathbb{F}_n^*$.

Nonzero elements of \mathbb{F}_n act on subsets of \mathbb{F}_n . Namely, if $u \in \mathbb{F}_n^*$, and $\sigma = \{x_0, \dots, x_k\} \subseteq \mathbb{F}_n$ we denote

$$u \cdot \sigma = u \cdot \{x_0, \dots, x_k\} := \{u \cdot x_0, \dots, u \cdot x_k\}.$$

We note that X is closed under such action, and linearly extend this definition as $u \cdot (\sigma_1 + \sigma_2) = u \cdot \sigma_1 + u \cdot \sigma_2$.

We organize A 's rows and columns by blocks corresponding to the orbits under the action of \mathbb{F}_n^* . A simple calculation shows that X 's 2-faces form $(n - 3)/2$ orbits of size $n - 1$, and one orbit of size $(n - 1)/2$. The latter is comprised of all 2-faces of the form $u \cdot \{1, -1, 0\}$. The 1-faces (edges) also form $(n - 3)/2$ orbits of size $n - 1$, and one orbit of size $(n - 1)/2$ that includes the edges $u \cdot \{1, -1\}$.

For $x \in \mathbb{F}_n^*$, and $y \in \mathbb{F}_n$, the block called $B_{[x,y]}$ is characterized by having the edge $\{1, x\}$ as a row and the 2-face $\{1, y, -(1 + y)/c\}$ as a column. We refer to x and y as the row and column *leaders* of B . This creates some ambiguity, since $\{1, x\}$ and $\{1, x^{-1}\}$ belong to the same block. Between x and x^{-1} the leader is the one with the smaller logarithm (as defined in Sect. 2). The same ambiguity and the way around it apply as well to y and y^{-1} . Since λ is a generator of the field, we order the rows of

the block indexed by x as follows:

$$[(1, x), \lambda \cdot (1, x), \lambda^2 \cdot (1, x), \dots, \lambda^{n-2} \cdot (1, x)]. \tag{3.9}$$

Likewise, the columns in a block whose column index is y are ordered as follows:

$$[(1, y, z), \lambda \cdot (1, y, z), \lambda^2 \cdot (1, y, z), \dots, \lambda^{n-2} \cdot (1, y, z)], \tag{3.10}$$

where $z = -(1 + y)/c$. Equations (3.9) and (3.10) also represent the orientation that we use for the edges and 2-faces, indicated by the use of tuples over sets.

Note that the column of a 2-face $\{u, -u, 0\}$ has a single non-zero entry in row $\{u, -u\}$, since we have removed the rows that correspond to the star centered at vertex 0. We eliminate these $(n - 1)/2$ rows and columns by Gaussian elimination as in [1] and arrive at

$$S = S_{n,c}, \tag{3.11}$$

an $((n - 3)(n - 1)/2) \times ((n - 3)(n - 1)/2)$ submatrix of A . To recap, X is a hypertree iff S is non-singular.

Claim 3.9 *Every block $B_{[x,y]}$ of S is circulant, i.e., $S \in \text{MCB}_{n-1,(n-3)/2}(\mathbb{Q})$.*

Proof The boundary operator is a signed inclusion matrix where the column that corresponds to the oriented face (u, v, w) is

$$e_{(u,v)} - e_{(u,w)} + e_{(v,w)},$$

where for an oriented edge (u, v) we define $e_{(u,v)}$ to be the 0/1 column vector with a single 1 in position (u, v) . If the direction is opposite, i.e., the edge (v, u) is present and not (u, v) , then $e_{(v,u)} = -e_{(u,v)}$. Note that the column that corresponds to the oriented face $(1, y, z)$ is

$$S[:, (1, y, z)] = e_{(1,y)} - e_{(1,z)} + e_{(y,z)} \tag{3.12}$$

while the column that corresponds to the oriented face in the same block $\lambda^k \cdot (1, y, z)$ is

$$S[:, \lambda^k \cdot (1, y, z)] = e_{\lambda^k \cdot (1,y)} - e_{\lambda^k \cdot (1,z)} + e_{\lambda^k \cdot (y,z)}. \tag{3.13}$$

To show the blocks in S are circulant, we need to show that

$$B_{[x,y]}[i, 0] = B_{[x,y]}[i + k, k],$$

where addition is done modulo $n - 1$. Before we give this proof, let us illustrate it with our running example:

Running Example

Our chosen generator of \mathbb{F}_{13}^* is $\lambda = 2$. The next matrix $B_{[3,4]}$ is an example of a block matrix in $S = S_{13,5}$ whose row leader is 3 and column leader 4:

Edge	2-Face						
	(1, 4, 12)	$2 \cdot (1, 4, 12)$ = (2, 8, 11)	$2^2 \cdot (1, 4, 12)$ = (4, 3, 9)	...	$2^9 \cdot (1, 4, 12)$ = (5, 7, 8)	$2^{10} \cdot (1, 4, 12)$ = (10, 1, 3)	$2^{11} \cdot (1, 4, 12)$ = (7, 2, 6)
(1,3)	0	0	0	...	0	1	0
$2 \cdot (1, 3) = (2, 6)$	0	0	0		0	0	1
$2^2 \cdot (1, 3) = (4, 12)$	1	0	0		0	0	0
$2^3 \cdot (1, 3) = (8, 11)$	0	1	0		0	0	0
$2^4 \cdot (1, 3) = (3, 9)$	0	0	1		0	0	0
⋮	⋮			⋮			
$2^{11} \cdot (1, 3) = (7, 8)$	0	0	0	...	1	0	0

For example, the leader column of (1, 4, 12) has non-zero entries in rows (1, 4), (1, 12), and (4, 12). In $B_{[3,4]}$ only row (4, 12) is present. Indeed, $B_{[3,4]} = P^2$ is a circulant matrix

$$\begin{aligned}
 B_{[x,y]}[i, 0] &= S[\lambda^i \cdot (1, x), (1, y, z)] = S[\lambda^{i+k} \cdot (1, x), \lambda^k(1, y, z)] \\
 &= B_{[x,y]}[i + k, k].
 \end{aligned}$$

The second equality stems from the indexation of the rows and 2-faces in the block as in (3.9), (3.10), and the definition of the boundary matrix for the appropriate columns in (3.12) and (3.13).

The fact that the blocks $B_{[x,y]}$ are circulant, follows also from the invariance of the boundary operator under the action of \mathbb{F}_n^* ,

$$\lambda \cdot \partial_2(\sigma) = \partial_2(\lambda \cdot \sigma)$$

for

$$\begin{aligned}
 \lambda \cdot \partial_2(\sigma) &= \lambda \cdot ((u_0, u_1) - (u_0, u_2) + (u_1, u_2)) \\
 &= (\lambda \cdot u_0, \lambda \cdot u_1) - (\lambda \cdot u_0, \lambda \cdot u_2) + (\lambda \cdot u_1, \lambda \cdot u_2) \\
 &= \partial_2(\lambda \cdot u_0, \lambda \cdot u_1, \lambda \cdot u_2) = \partial_2(\lambda \cdot \sigma).
 \end{aligned}$$

As S is a submatrix of the matrix form of ∂_d , the claim follows from the indexation of the rows and columns. □

4 S as a Polynomial Matrix

This section we find out how to express the matrix $S = S_{n,c}$ as defined in (3.11) as a polynomial matrix. This matrix is a sum of three sparse matrices, each having at most one nonzero block per row and column. All these blocks have the form $\pm P^j$ for some

$n - 2 \geq j \geq 0$, because the boundary operator is a signed inclusion matrix. Thus the column of the oriented triple $(1, y, z)$ may only contain terms corresponding to $(1, y)$, to $(1, z)$ and to (y, z) . But any particular row may be labeled by either x or x^{-1} , so there are (at most) three possible cases where the block $B_{[x,y]}$ is nonzero indexed as follows:

$$\begin{aligned} i = 1: & y \in \{x, x^{-1}\}, \\ i = 2: & z \in \{x, x^{-1}\}, \\ i = 3: & zy^{-1} \in \{x, x^{-1}\}. \end{aligned}$$

In case i we refer to the relevant x as x_i . To sum up,

$$B_{[x,y]} = T_1 + T_2 + T_3$$

where

$$\begin{aligned} T_1 &= \begin{cases} I & y = x, \\ -P^{-\log x} & y = x^{-1}, \end{cases} & T_2 &= \begin{cases} -I & z = x, \\ P^{-\log x} & z = x^{-1}, \end{cases} \\ T_3 &= \begin{cases} P^{\log y} & z \cdot y^{-1} = x, \\ -P^{\log z} & z \cdot y^{-1} = x^{-1}. \end{cases} \end{aligned}$$

Most columns indeed have exactly three nonzero terms, although two of them possibly reside in the same block. We return again to our running example for illustration.

Running Example

The next table shows the logarithm to base $\lambda = 2$ of every $x \in \mathbb{F}_{13}^*$, and its order $o(x)$:

x	1	2	4	8	3	6	12	11	9	5	10	7
$\log x$	0	1	2	3	4	5	6	7	8	9	10	11
$o(x)$	1	12	6	4	3	12	2	12	3	4	6	12

We next present the matrix $S = S_{13,5}$ in block form. To the left and above the matrix appear the blocks' edge and 2-face leaders. For example, $(1, 4, 12) \in X$, since $1 + 4 + 5 \cdot 12 \equiv 0 \pmod{13}$ and it is the 2-face leader of the first column. No leader has the form $(1, 2, z)$, since $1 + 2 + 5 \cdot 2 \equiv 0$, i.e., $z = y = 2$.

Edge	2-Face				
	(1, 4, 12)	(1, 8, 6)	(1, 3, 7)	(1, 6, 9)	(1, 0, 5)
(1,2)	0	0	P^{11}	0	0
(1,4)	I	P^3	0	0	0
(1,8)	0	I	0	P^5	P^9
(1,3)	P^2	0	I	P^8	0
(1,6)	0	$-I$	$-P^{11}$	I	0

The block $B_{[3,4]}$ from our previous discussion of the running example resides in the first column and fourth row. The matrix S is 60×60 , being a 5×5 matrix of 12×12 blocks, whence $\underline{S}(z)$ is a 5×5 matrix. For example, $\underline{S}(\omega_3)$ is

Edge	2-Face				
	(1, 4, 12)	(1, 8, 6)	(1, 3, 7)	(1, 6, 9)	(1, 0, 5)
(1,2)	0	0	ω_3^2	0	0
(1,4)	1	1	0	0	0
(1,8)	0	1	0	ω_3^2	1
(1,3)	ω_3^2	0	1	ω_3^2	0
(1,6)	0	-1	$-\omega_3^2$	1	0

5 Non-acyclic Complexes

In this section we present an infinite family of non-acyclic 2-dimensional complexes. As subsequently discussed in the section, we suspect that asymptotically almost every non-acyclic complex is in this family and asymptotically almost every complex that is not in this family is acyclic.

Before we move on to the main subject of this section, we prove the following curious connection between $o(c)$ and $\log c$, which we could not find in the literature:

Lemma 5.1 For n prime and $c \in \mathbb{F}_n^*$,

$$\gcd\left(\log c, \frac{n-1}{2}\right) = \gcd\left(\frac{n-1}{o(c)}, \frac{n-1}{2}\right).$$

Proof Let $k_1 = \gcd(\log c, (n-1)/2)$ and $k_2 = \gcd((n-1)/o(c), (n-1)/2)$.

$k_2 \mid k_1$: $\lambda^{\log c} = c$ so $\lambda^{\log c \cdot o(c)} = 1$. But λ is a generator whence $(n-1) \mid \log c \cdot o(c)$, thus

$$\frac{n-1}{o(c)} \mid \log c. \tag{5.1}$$

Since $k_2 \mid (n-1)/o(c)$ and $k_2 \mid (n-1)/2$, from (5.1) we conclude that

$$k_2 \mid \gcd\left(\log c, \frac{n-1}{2}\right) = k_1.$$

$k_1 \mid k_2$: It is sufficient to prove that $k_1 \mid (n-1)/o(c)$. We apply Lemma 2.1 with $x = \lambda^2$ and $j = \log c$ and conclude that

$$o(\lambda^{2 \log c}) = \frac{o(\lambda^2)}{\gcd(\log c, o(\lambda^2))}. \tag{5.2}$$

But $o(\lambda^2) = (n - 1)/2$ and $o(\lambda^{2 \log c}) = o(c^2)$. Rearranging (5.2) gives

$$k_1 = \gcd\left(\log c, \frac{n - 1}{2}\right) = \frac{n - 1}{2 \cdot o(c^2)}.$$

Lemma 2.1 also tells us that if $o(c)$ is even, $o(c^2) = o(c)/2$ and if $o(c)$ is odd, $o(c^2) = o(c)$. In either case $k_1 \mid (n - 1)/o(c)$, as required. \square

Corollary 5.2 For n prime and $c \in \mathbb{F}_n^*$,

$$\gcd\left(\log(-c), \frac{n - 1}{2}\right) = \gcd\left(\frac{n - 1}{o(c)}, \frac{n - 1}{2}\right).$$

Proof It is sufficient to prove that $\gcd(\log(-c), (n - 1)/2) = \gcd(\log c, (n - 1)/2)$. Since $\log(-1) = \pm(n - 1)/2$ it follows that $\log(-c) = \log c \pm (n - 1)/2$. This yields

$$\gcd\left(\log(-c), \frac{n - 1}{2}\right) = \gcd\left(\log c \pm \frac{n - 1}{2}, \frac{n - 1}{2}\right) = \gcd\left(\log c, \frac{n - 1}{2}\right),$$

as claimed. \square

Theorem 5.3 The complex $X_{2,n,c}$ is non-acyclic in the following cases:

- $n \equiv 1 \pmod 4$ and c is not a primitive element of \mathbb{F}_n ;
- $n \equiv 3 \pmod 4$ and c is neither primitive nor of order $(n - 1)/2$.

Proof Note that in both cases $\gcd((n - 1)/o(c), (n - 1)/2) > 1$. When $n \equiv 1 \pmod 4$, the integer $(n - 1)/2$ has all the prime factors of $n - 1$, so $(n - 1)/o(c)$ and $(n - 1)/2$ are relatively prime only if $o(c) = n - 1$, i.e., c is primitive. When $n \equiv 3 \pmod 4$, the only divisor of $n - 1$ that fails to divide $(n - 1)/2$ is 2. Therefore $(n - 1)/o(c)$ and $(n - 1)/2$ are relatively prime only if c is primitive or $o(c) = (n - 1)/2$.

So let $k \geq 2$ be a common divisor of $(n - 1)/o(c)$ and $(n - 1)/2$. We show below that the vector $v_{n,k}$ whose x entry is

$$v_{n,k}[x] = 1 - \omega_k^{\log x}$$

is in the left kernel of $\underline{S}(\omega_k)$. We first illustrate this in our running example:

Running Example

Theorem 5.3 applies to our running example, since $o(5) = 4$, and $\gcd(12/4, 12/2) = 3 > 1$. It is easily verified that the following vector is in the left kernel of $\underline{S}(\omega_3)$:

$$v_{13,3} = (1 - \omega_3^1 \ 1 - \omega_3^2 \ 0 \ 1 - \omega_3^1 \ 1 - \omega_3^2),$$

$$(1 - \omega_3^1 \ 1 - \omega_3^2 \ 0 \ 1 - \omega_3^1 \ 1 - \omega_3^2) \cdot \begin{pmatrix} 0 & 0 & \omega_3^2 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & \omega_3^2 & -1 \\ \omega_3^2 & 0 & 1 & \omega_3^2 & 0 \\ 0 & -1 & -\omega_3^2 & 1 & 0 \end{pmatrix} = \vec{0}.$$

For $y \in \mathbb{F}_n^*$ we compute the y -th entry of $v_{n,k} \cdot \underline{S}(\omega_k)$, i.e., the coordinate that corresponds to the 2-face $(1, y, z)$,

$$(v_{n,k} \cdot \underline{S}(\omega_k))[y] = \sum_{i=1,2,3} (1 - \omega_k^{\log x_i}) \cdot T_i(\omega_k). \tag{5.3}$$

We define $\Theta_{k,i}$ as the scalar term that is obtained upon evaluating T_i at ω_k .

Lemma 5.4 *Let $t_{k,i} := (1 - \omega_k^{\log x_i}) \cdot \Theta_{k,i}$ be the i -th term in (5.3). Then*

- (i) $t_{k,1} = 1 - \omega_k^{\log y}$,
- (ii) $t_{k,2} = \omega_k^{\log z} - 1$,
- (iii) $t_{k,3} = \omega_k^{\log y} - \omega_k^{\log z}$.

Proof Let us go through the cases:

- (i) If $x_1 = y$, then $\Theta_{k,1} = 1$. If $x_1 = y^{-1}$, then $\Theta_{k,1} = -\omega_k^{-\log x_1}$.
- (ii) If $x_2 = z$, then $\Theta_{k,2} = -1$. If $x_2 = z^{-1}$, then $\Theta_{k,2} = \omega_k^{-\log x_2}$.
- (iii) If $x_3 = z \cdot y^{-1}$, then $\Theta_{k,3} = \omega_k^{\log y}$. If $x_3 = z^{-1} \cdot y$, then $\Theta_{k,3} = -\omega_k^{\log z}$.

All cases are readily verifiable. □

It is easy to check that $t_{k,1} + t_{k,2} + t_{k,3} = 0$ for those columns y where T_1, T_2, T_3 are well defined. We turn to deal with the exceptional cases where some of the definitions fail.

The matrix T_2 is undefined for the column of $y = c - 1, z = -1$, since row $(1, -1)$ is absent. But then

$$t_{k,1} + t_{k,3} = 1 - \omega_k^{\log y} + \omega_k^{\log y} - \omega_k^{\log z} = 1 - \omega_k^{\log(-1)} = 1 - \omega_k^{(n-1)/2} = 0,$$

since $k \mid (n - 1)/2$. Also, T_1, T_3 are undefined when $y = 0, z = -c^{-1}$. However, in this case $t_{k,2} = \omega_k^{-\log(-c)} - 1 = 0$, a deduction from Corollary 5.2 since $k \mid \log(-c)$. □

5.1 A Conjecture About a Single Criterion

Again let n be a prime, and $c \neq 0, \pm 1, -2$. As illustrated in Fig. 2, our computer simulations suggest that for large n , Theorem 5.3 captures asymptotically almost all cases in which $X_{2,n,c}$ is non-acyclic. For the acyclic cases, asymptotically almost all other cases that Theorem 5.3 does not capture are acyclic.

This suggests that the following simple criterion asymptotically determines whether or not a complex is acyclic. The asymptotics is w.r.t. $n \rightarrow \infty$.

Conjecture 5.5 *Let the complex $X_{2,n,c}$ with n prime and $c \in \mathbb{F}_n^*$ be as above.*

- *When $n \equiv 1 \pmod 4$:*
 - *Recall that $X_{2,n,c}$ is non-acyclic if $c \in \mathbb{F}_n^*$ is non-primitive.*
 - *We conjecture that $X_{2,n,c}$ is acyclic for asymptotically almost every primitive $c \in \mathbb{F}_n^*$.*
- *When $n \equiv 3 \pmod 4$:*
 - *Recall that $X_{2,n,c}$ is non-acyclic if $c \in \mathbb{F}_n^*$ is neither primitive, nor of order $(n - 1)/2$.*
 - *We conjecture that $X_{2,n,c}$ is acyclic for asymptotically almost every $c \in \mathbb{F}_n^*$ that is either primitive or of order $(n - 1)/2$.*

In light of Fig. 2, we can state Conjecture 5.5 in terms of acyclic ratios:

Claim 5.6 *By Theorem 5.3, for every prime n the acyclic ratio (i.e., the fraction of complexes $X_{2,n,c}$ that are acyclic) is at most*

$$\frac{\varphi((n - 1)/2)}{(n - 1)/2},$$

where φ is Euler’s function. Conjecture 5.5 posits that this bound is asymptotically tight.

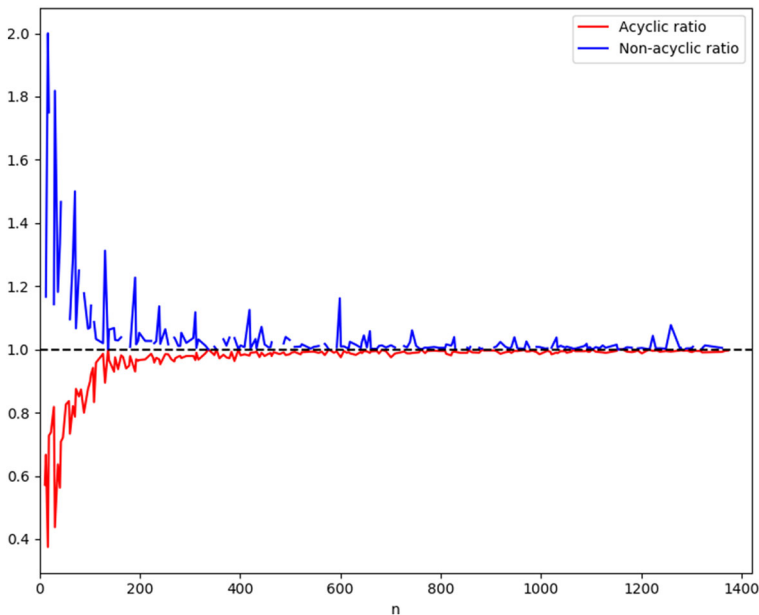


Fig. 2 Acyclic and Non-Acyclic ratios for all primes $11 \leq n \leq 1373$. The acyclic ratio is $A_n/(n - 4 - E_n)$ and the non-acyclic ratio is N_n/E_n . Here A_n and N_n are the number of $c \in \mathbb{F}_n^* \setminus \{\pm 1, -2\}$ for which $X_{2,n,c}$ is acyclic, resp. non-acyclic. E_n is the number of non-acyclic complexes explained by Theorem 5.3




Fig. 3 For each prime n , $11 \leq n \leq 1373$, this is the probability that $X_{2,n,c}$ is acyclic over all eligible $c \in \mathbb{F}_n^* \setminus \{\pm 1, -2\}$

Proof Recall that the number of elements of order d in a cyclic group of order m , is $\varphi(d)$ for every divisor d of m . Since the cyclic group \mathbb{F}_n^* contains $n - 1$ elements, it is left to prove that

- When $n \equiv 1 \pmod 4$:

$$\frac{\varphi((n - 1)/2)}{(n - 1)/2} = \frac{\varphi(n - 1)}{n - 1}.$$

- When $n \equiv 3 \pmod 4$:

$$\frac{\varphi((n - 1)/2)}{(n - 1)/2} = \frac{\varphi(n - 1) + \varphi((n - 1)/2)}{n - 1}.$$

These claims follow from the multiplicative properties of φ . Namely,

$$\varphi(m \cdot k) = \frac{d \cdot \varphi(m) \cdot \varphi(k)}{\varphi(d)} \tag{5.4}$$

for two positive integers m, k , where $d = \gcd(m, k)$. The claim follows by applying (5.4) with $m = 2, k = (n - 1)/2$. If $n \equiv 1 \pmod 4$, then $\varphi(n - 1) = 2 \cdot \varphi((n - 1)/2)$ and when $n \equiv 3 \pmod 4$, $\varphi(n - 1) = \varphi((n - 1)/2)$. \square

In Fig. 3 observe the good agreement with the function presented in Claim 5.6. For example, the acyclic percentage is close to 1 for primes n of the form $(n - 1)/2 = p$,

where p is prime. The lowest acyclic percentage is attained when $(n - 1)/2 = 2 \cdot 3 \cdot 5 \cdot \dots$. All of this is in agreement with our Conjecture 5.5.

5.2 More Non-acyclic Cases

As the two last figures illustrate, Theorem 5.3 does not capture all the cases in which $X_{n,c}$ is non-acyclic. Indeed, the matrices $\underline{S}(1)$ and $\underline{S}(-1)$ may be singular even when c is a generator of \mathbb{F}_n^* or has order $(n - 1)/2$ when $n \equiv 3 \pmod 4$. For example, and this is not captured by Theorem 5.3, $X_{n,c}$ is non-acyclic when

$$c^2 + c - 1 \equiv 0 \pmod n.$$

To see why, let $k = (n - 1)/2$, and define $\tilde{v}_{n,k}$ to be $v_{n,k}$ as in the proof of Theorem 5.3, with the single change that $\tilde{v}_{n,k}[-c - 1] = 0$. A small variation on the same proof shows that $\tilde{v}_{n,k} \cdot \underline{S}(\omega_k) = \vec{0}$, since $-c - 1 \equiv -1/c \pmod n$.

6 Full Matrices

In this section we sketch another possible approach to the construction of an infinite family of 2-dimensional hyperpaths. A certain extension of the boundary operator matrix plays a key role in these developments.

Definition 6.1 For a prime n and $c \in \{2, \dots, n - 3\}$ the matrix $F = F_{n,c}$ is an $(n^2 - 1) \times (n^2 - 1)$ binary matrix. A row ρ_{xy} of F is indexed by an ordered pair $(x, y) \in \mathbb{F}_n^2 \setminus (0, 0)$. A column $f_{(xyz)}$ is indexed by an ordered triple (x, y, z) , where $(x, y) \in \mathbb{F}_n^2 \setminus (0, 0)$ and $z = -(x + y)/c$, i.e., $x + y + c \cdot z = 0$. The $[(u, v), (x, y, z)]$ entry of F is 1 iff

- $(u, v) = (x, y)$, this is called an xy -entry;
- $(u, v) = (y, z)$, this is called an yz -entry;
- $(u, v) = (z, x)$, this is called an zx -entry.

Note that the xy entries in F form a permutation matrix which we call P_{xy} . Likewise for P_{yz} and P_{zx} . Also $P_{xy} = I$. Consequently:

$$F = I + P_{yz} + P_{zx}, \tag{6.1}$$

where P_{yz}, P_{zx} are permutation matrices. Note that $\text{rank}(F_{n,c}) \leq n^2 - n$ since there are $n - 1$ pairs of identical columns in $F_{n,c}$, namely, $f_{(a, -(c+1)a, a)} = f_{(-(c+1)a, a, a)}$ for every $a \in \mathbb{F}_n^*$.

Claim 6.2 *If $\text{rank}(F_{n,c}) = n^2 - n$, then $\text{rank}(A_{n,c}) = \binom{n-1}{2}$, in which case $X_{n,c}$ is a hypertree.*

Proof The proof is omitted, but can be found in the extended version of this work [3]. □

The reverse implication in Claim 6.2 does not hold. This suggests the study of $F_{n,c}$ as an approach to the construction of infinitely many 2-dimensional hyperpaths.

Supporting Data

Data used to generate Fig. 3 is available from the authors.

Acknowledgements We thank Roy Meshulam for insightful comments on this manuscript.

References

1. Aronshtam, L., Linial, N., Łuczak, T., Meshulam, R.: Collapsibility and vanishing of top homology in random simplicial complexes. *Discrete Comput. Geom.* **49**(2), 317–334 (2013)
2. Bläser, M.: Fast matrix multiplication and related problems. <http://www.ens-lyon.fr/LIP/MC2/data/uploads/slidesmarkusblaser.pdf>
3. Dahari, A., Linial, N.: Hyperpaths (2020). [arXiv:2011.09936](https://arxiv.org/abs/2011.09936)
4. Davis, P.J.: *Circulant Matrices*. A Wiley-Interscience Publication. Pure and Applied Mathematics. Wiley, New York (1979)
5. Gauss, C.F.: *Untersuchungen über höhere Arithmetik*. American Mathematical Society, Providence (2006)
6. Hardy, G.H., Wright, E.M.: *An Introduction to the Theory of Numbers*. Oxford University Press, New York (1979)
7. Kalai, G.: Enumeration of \mathbf{Q} -acyclic simplicial complexes. *Israel J. Math.* **45**(4), 337–351 (1983)
8. Le Gall, F.: Powers of tensors and fast matrix multiplication. In: 39th International Symposium on Symbolic and Algebraic Computation (Kobe 2014), pp. 296–303. ACM, New York (2014)
9. Linial, N., Meshulam, R., Rosenthal, M.: Sum complexes—a new family of hypertrees. *Discrete Comput. Geom.* **44**(3), 622–636 (2010)
10. Linial, N., Newman, I., Peled, Y., Rabinovich, Y.: Extremal hypercuts and shadows of simplicial complexes. *Israel J. Math.* **229**(1), 133–163 (2019)
11. Linial, N., Peled, Y.: On the phase transition in random simplicial complexes. *Ann. Math.* **184**(3), 745–773 (2016)
12. Linial, N., Peled, Y.: Enumeration and randomized constructions of hypertrees. *Random Struct. Algorithms* **55**(3), 677–695 (2019)
13. Mathew, R., Newman, I., Rabinovich, Y., Rajendraprasad, D.: Hamiltonian and pseudo-Hamiltonian cycles and fillings in simplicial complexes. *J. Combin. Theory Ser. B* **150**, 119–143 (2021)
14. McDonald, B.R.: *Linear Algebra Over Commutative Rings*. Monographs and Textbooks in Pure and Applied Mathematics, vol. 87. Marcel Dekker, New York (1984)
15. Petković, M.D., Stanimirović, P.S.: Generalized matrix inversion is not harder than matrix multiplication. *J. Comput. Appl. Math.* **230**(1), 270–282 (2009)
16. Rjasanow, S.: Effective algorithms with circulant-block matrices. *Linear Algebra Appl.* **202**, 55–69 (1994)
17. Tsitsas, N.L., Alivizatos, E.G., Kalogeropoulos, G.H.: A recursive algorithm for the inversion of matrices with circulant blocks. *Appl. Math. Comput.* **188**(1), 877–894 (2007)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.