Cyclic Transversal Polytopes

Dissertation

zur Erlangung des akademischen Grades

doctor rerum naturalium
(Dr. rer. nat.)

von M.Sc. Jonas Frede
geb. am 15. Januar 1994 in Goslar

genehmigt durch die Fakultät für Mathematik
der Otto–von–Guericke–Universität Magdeburg

Gutachter: Prof. Dr. Volker Kaibel
Otto–von–Guericke–Universität Magdeburg

Prof. Dr. Maximilian Merkert
Technische Universität Braunschweig

eingereicht am: 22. November 2022
Verteidigung am: 14. April 2023

Zusammenfassung

In der Literatur wurden klassische kombinatorische Optimierungsprobleme wie das *Max-Cut-Problem* oder das *Stable-Set-Problem*, ersatzweise auch die dazu gehörigen Cut- und Stable-Set-Polytope, bereits vielfach behandelt und diskutiert. Diese Arbeit bietet einen vereinheitlichten Blick auf diese beiden und andere Probleme. Die hier untersuchten Probleme lassen sich mit Hilfe ausgewählter Tupel von Binärvektoren beschreiben, deren Summe eine Paritätsbedingung erfüllen muss. Tupel dieser Art nennen wir *zyklische Transversale.*

Ziel der vorliegenden Arbeit ist es daher, eine Grundlage zur weiteren Erforschung dieses verallgemeinernden Paradigmas der zyklischen Transversale zu schaffen. Davon ausgehend werden zuerst die zentralen Objekte dieser Arbeit, so genannte *zyklische-Transversale-Polytope*, definiert und analysiert. Diese Analyse führt einerseits zu einer Identifikation wesentlicher Meta-Parameter, die bei der Klassifikation dieser Polytope von Bedeutung sind, und andererseits zu einer Untersuchung kombinatorischer Eigenschaften der Polytope, die beispielsweise in einer Charakterisierung der Adjazenz ihrer Ecken mündet.

Nach einem Beweis der anfänglichen Behauptung, dass Cut-Polytope und Stable-Set-Polytope neben anderen in der Literatur bekannten Polytopen zu den zyklische-Transversale-Polytopen gehören, erfolgt der Beweis einer notwendigen Bedingung für zyklische-Transversale-Polytope. Darauf aufbauend ergibt sich, dass nicht alle Kreuzpolytope zyklische-Transversale-Polytope sein können. Daran schließt sich eine Klassifikation der zyklische-Transversale-Polytope bis Dimension drei an.

Durch Betrachtung in gewisser Weise allgemeinster zyklische-Transversale-Polytope wird außerdem ein Einblick in die Struktur gültiger Ungleichungen für zyklische-Transversale-Polytope gewährt. Dadurch kann eine Klasse gültiger Ungleichungen identifiziert werden, die die so genannten Odd-Set-Ungleichungen für Paritätspolytope verallgemeinert. Computergestützte Berechnungen zusätzlicher Ungleichungen und deren schematische Visualisierung ergänzen diese Betrachtung.

Schließlich werden Relaxierungen von zyklische-Transversale-Polytopen und ihre Eigenschaften beschrieben. Die Konstruktion erweiterter Formulierungen von zyklische-Transversale-Polytopen vervollständigen die vorliegende Arbeit.

Abstract

In the literature, classical combinatorial optimization problems such as the *max-cut problem* or the *stable set problem*, alternatively the associated cut and stable set polytopes, have been widely treated and discussed. This work provides a unified view of these two and other problems. The problems studied here can be described in terms of selected tuples of binary vectors whose sum must satisfy a parity condition. We call tuples of this type *cyclic transversals*.

Thus, the aim of the present dissertation is to provide a basis for further research of this generalizing paradigm of cyclic transversals. Starting from this, the central objects of this work, so-called *cyclic transversal polytopes*, are first defined and analyzed. This analysis leads, on the one hand, to an identification of essential meta-parameters relevant in the classification of these polytopes, and, on the other hand, to an investigation of combinatorial properties of the polytopes, resulting, for example, in a characterization of the adjacency of their vertices.

After a proof of the initial assertion that cut polytopes and stable set polytopes belong to the cyclic transversal polytopes among other polytopes known in the literature, a proof of a necessary condition for cyclic transversal polytopes is given. Based on this it follows that not all cross polytopes can be cyclic transversal polytopes. This is followed by a classification of cyclic transversal polytopes up to dimension three.

By considering in some sense most general cyclic transversal polytopes, insight into the structure of valid inequalities for cyclic transversal polytopes is also provided. Thus, a class of valid inequalities can be identified which generalizes the so-called odd-set inequalities for parity polytopes. Computer-aided computations of additional inequalities and their schematic visualizations supplement this review.

Finally, relaxations of cyclic transversal polytopes and their properties are described. The construction of extended formulations of cyclic transversal polytopes complete the present work.

Contents

1 Introduction

The field of combinatorial optimization is concerned with the problem of finding an optimal element out of a finite set of objects with respect to certain properties. Many classical combinatorial optimization problems do not lend themselves to any naive enumeration of possibilities because of the sheer number of objects that need to be enumerated. Therefore, obtaining structural insights into these objects is a necessary and worthwhile endeavor.

Among these ground-breaking insights are the formulation of combinatorial optimization problems as (integer) linear optimization problems over real vector spaces, which enables the use of an expanding portfolio of theoretical and algorithmic results that apply to these problems. Indeed, powerful and sophisticated software such as general-purpose solvers built upon such results allow for solving even some fairly large and practically relevant instances of combinatorial optimization problems in a reasonable amount of time.

Nevertheless, most combinatorial optimization problems remain intractable, which encourages more research to incorporate findings of any valuable sort, as this is one of the primary instincts of mathematicians. Thus, to balance the general purpose approach, other advances are more tailored to specific combinatorial optimization problems, like looking for new methods and patterns to exploit and specializing in solving certain kinds of challenging combinatorial optimization problems.

One such challenging class of combinatorial optimization problems are those whose feasible objects can be characterized as collections of vectors over the binary finite field that fulfill a parity condition. As one will see, although this description seems rather specialized, optimization problems of this kind are numerous and many important standard problems in combinatorial optimization are expressible in this way, such as problems over graph structures like the famous *max-cut problem* or the *stable set problem*. Therefore, as the scope of this work, we concern ourselves with laying the groundwork for a better understanding of this kind of problem class. This is laid out as the *cyclic transversal framework*, since we call feasible objects of the aforementioned kind *cyclic transversals*.

The max-cut and stable set problems are non-trivial and computationally difficult to solve in general, and the viewpoint of cyclic transversals provides a unified framework for one to work with, but as a generalization, optimization problems in the cyclic transversal framework will also be difficult to solve. Hence, aside from examining which problems can be expressed in the cyclic transversal framework, we are concerned with descriptions and valid inequalities of the central geometric objects of the cyclic transversal framework, also called *cyclic transversal polytopes*. To mitigate the difficulty of solving cyclic transversal optimization problems directly, a relaxation hierarchy as well as so-called extended formulations for cyclic transversal polytopes and their relaxations are investigated. The results of this undertaking are laid out in this thesis.

By presenting these results, we hope to convince the reader of the importance and richness of the cyclic transversal framework.

Outline

Although we assume familiarity with mathematical texts and conventions, in the rest of this chapter we give some necessary definitions to delve into the cyclic transversal framework, namely some foundations of linear and Boolean algebra as well as graph theory and polyhedral combinatorics. We advise that this is not a thorough introduction to these topics and ample references will be given, but it should be sufficient to serve as background to readers not too familiar with it.

Starting in Chapter 2, we thoroughly introduce the cyclic transversal framework. Properties of the central objects of this framework called cyclic transversals and cyclic transversal polytopes will be studied. We develop methods to classify and distinguish sets of cyclic transversals based on some essential parameters, and describe an equivalence operation on them. The chapter then finishes with methods called reductions to simplify the defining sets called block configurations, rudimentary bounds on the essential parameters of these configurations and a characterization of the vertex adjacency of cyclic transversal polytopes.

After that in Chapter 3, we reexamine some well-known polytopes introduced at the end of this introduction and describe their relationship to the cyclic transversal framework. Among other results, binary kernel polytopes are identified as a special subclass of cyclic transversal polytopes, and any stable set and set packing polytope is proven to be isomorphic to a cyclic transversal polytope with the essential parameters bounded polynomially by the sizes of the underlying structured sets, i. e., graphs and set collections. We also prove that projections of cyclic transversal polytopes are universal in the following sense: Given a set of 0/1-vectors that are described by a Boolean formula, its convex hull is affinely isomorphic to a projection of a cyclic transversal polytope whose essential parameters are polynomially bounded by the characteristics of this Boolean formula.

Chapter 4 is dedicated to the subfamily of full cyclic transversal polytopes, which feature additional symmetry when compared to general cyclic transversal polytopes. Given the appropriate parameters, every other cyclic transversal polytope is a face of a full cyclic transversal polytope. We find valid inequalities for these polytopes and prove that they are facet-defining. Afterwards we examine full cyclic transversal polytopes with small parameters more closely and show how irreducible complete inequality descriptions for some of these parameters are obtained. Excursions and computations regarding larger parameters complete this chapter.

Relaxations and extensions of cyclic transversal polytopes are the topic of Chapter 5, where we use linear algebra to obtain a hierarchy of polytopes that gradually approximates a given cyclic transversal polytope. For the first steps of this hierarchy, we give an explicit description of the resulting approximating polytopes, using and extending a result from Barahona and Grötschel along the way. Giving extended formulations of cyclic transversal polytopes and their relaxations constitutes the second part of this chapter, where a well-known method of building an extended formulation is generalized and applied to these polytopes. The chapter closes with some suggestions leading into further directions regarding extended formulations.

Finally, we conclude this thesis by summarizing the main results and give an outlook with some leads into further research about cyclic transversal polytopes at the end.

Notation and Background Definitions

Throughout this work, we assume that the reader is well-grounded in mathematics and therefore understands the basic terminology and notational conventions used, particularly in regard to polyhedra in mathematical optimization. Despite that, aside from the chunk of notation that is introduced at the place of its first occurrence, some conventions that are used universally and are needed for this thesis are fixed here. Therefore, we recall some foundations and conventions of linear and Boolean algebra, graph theory and polyhedral combinatorics.

Furthermore, referrals to standard literature for exhaustive background information on the aforementioned fields are given, especially on topics that are beyond the scope or cannot be explained to a satisfactory degree within this section, such as details in complexity theory and with it the concepts of NP, coNP, NP-*complete*, NP-*hard*, for which Arora and Barak [1], Ausiello et al. [3], and Papadimitriou [43] have given adequate textbooks, or the vast spectrum of topics in matroid theory, for which Oxley [41] has written a reference.

Foundations

If some object or expression X is defined by another object or expression Y, we write this as $X \coloneqq Y$ or $Y \eqqcolon X$. Given a set S, our convention is that the statement "T is a subset of S" is written as $T \subseteq S$, while the additional condition that T as a subset of S is not equal to S itself is written as $T \subsetneq S$. If the negation "T is not a subset of S" should be expressed, we write it as $T \nsubseteq S$. The cardinality of a set S is denoted by $|S|$. For a set S, we denote the set of all its subsets by 2^S, and by $\binom{S}{\kappa}$ the set of all subsets of S with cardinality κ. For singleton sets $\{s\}$ we usually omit the brackets in statements. In addition to the usual binary set operations $\cup, \cap$ and $\setminus$ we use $\uplus$ to denote a union of disjoint sets in order to emphasize their disjointedness.

Expressions containing $((...))$ or $(\{...\})$ are typically simplified to $(...)$ or another expression with single brackets, if the meaning is still clear from context.

The set of natural numbers, denoted by $\mathbb{N}$, and the set $\mathbb{N}_0 := \mathbb{N} \cup 0$ are distinct, meaning that the smallest element in $\mathbb{N}$ is 1. As shorthand notation, we define $[n] := \{1, \dots, n\}$ and $[n]_0 := 0 \uplus [n]$ for $n \in \mathbb{N}$.

Given a set S and some other set I of indices, we denote (indexed) sequences of elements $s_i \in S$ by $(s_i \mid i \in I)$, or $(s_1, \dots, s_n)$ for the index set $[n]$. Note that there is a natural bijection between functions $f\colon [n] \to S$ and tuples $(f(1), \dots, f(n))$.

For two maps $\pi\colon A \to B$ and $\tau\colon B \to C$, their composition is denoted by $\tau \circ \pi\colon A \to C$. For maps between structured sets (e. g., vector spaces, groups, polyhedra, lattices and graphs), we follow the usual definitions and conventions: An *isomorphism* is an invertible structure-preserving map. E. g., for vector spaces these are invertible linear maps. An *automorphism* is an isomorphism of an object to itself.

We use the usual Landau notation for asymptotics, i. e., for a given map $g\colon \mathbb{N} \to \mathbb{N}$, we denote by $\mathcal{O}(g)$ and $\Omega(g)$ the sets of all maps that asymptotically grow not faster (not slower, respectively) than g. That means that the set $\mathcal{O}(g)$ (respectively $\Omega(g)$) contains all maps f for which there exist constants $c > 0$ and $n_0 > 0$, depending on f, such that the inequality $f(n) \leq cg(n)$ (respectively $f(n) \geq cg(n)$) holds for all $n \geq n_0$. The set $\Theta(g)$ is defined as the intersection $\mathcal{O}(g) \cap \Omega(g)$, i. e., it contains exactly all the functions that asymptotically grow like g.

Linear Algebra

This short section serves not as a comprehensive review of linear algebra but to reexamine some notation needed in this thesis. Bosch [11], Fischer and Springborn [26], and Lang [36] give thorough introductions into linear algebra.

For every finite set S, we denote by $\mathbb{R}^S$ the real vector space of dimension $|S|$, indexed by the elements in S. If we select an ordering of the elements of S by selecting some map $[|S|] \to S$, we canonically identify $\mathbb{R}^S$ with $\mathbb{R}^{|S|}$. Aside from vector spaces over the real numbers, we will utilize vector spaces over finite fields of order q, denoted by $\mathbb{F}_q$, most prominently over the finite field of order 2, where we will use the same notation $\mathbb{F}_2$.

Given a non-empty set S and a subset $T \subseteq S$, we define its *characteristic* or *incidence vector* as $\chi(T) \in \{0,1\}^S$ with $\chi(T)_i = 1$ if $i \in T$ and 0 otherwise. We also identify zero-one vectors with the subsets of which they are incidence vectors; in that sense we use formulations like a 0/1-vector being a subset of another one.

The letter d will always denote the dimension of some underlying vector space or some dimensional parameter and therefore will usually be a non-negative integer. We identify vector spaces $\mathbb{F}_2^d$ and the sets $\{0,1\}^d \subseteq \mathbb{R}^d$ for $d \in \mathbb{N}_0$ and denote the addition in $\mathbb{F}_2^d$ with $\oplus$ to prevent confusion with the usual addition in $\mathbb{R}^d$.

Letters i, j, k, ℓ are usually indices, which are positive integers unless specified otherwise. The letter n will also denote a positive integer as we have used it before. The letters p and q will be used as well for this purpose. We denote vectors of real vector spaces $\mathbb{R}^d$ with small Latin letters a, b, c or x, y, z, while elements in $\mathbb{F}_2^d$ are denoted by small Greek letters from the rear part of the Greek alphabet, like ω or σ, but θ and ϱ are also used. The letter β in contrast is used as the right-hand side of an inequality $\langle a, x\rangle \leq \beta$ or equation $\langle a, x\rangle = \beta$, and h will be the normal vector of a hyperplane H for any vector space, where $\langle a, x\rangle$ is the standard scalar product of the respective vector space. To denote subsets and subspaces of $\mathbb{R}^d$ or $\mathbb{F}_2^d$, Latin (or Greek, respectively) big letters are used, unless they have been defined for a different purpose. The letter Ξ_i, possibly without an index, will always denote a subset of $\mathbb{F}_2^d$ for some d, and $\xi(i)$ will always be an element of Ξ_i subject to some predetermined condition.

For $\mathbb{R}^d$ and $\mathbb{F}_2^d$, the vector $\mathbb{0}_d$ is the neutral element with respect to vector addition, i. e., the vector where all components are equal to 0, also called zero-vector. Likewise, $\mathbb{1}_d$ is defined as the vector where all components are 1, the usual neutral element of multiplication in the underlying field. It is also called all-ones-vector. If the dimension is clear from context, we usually omit the indices for these vectors. Furthermore, by $\mathbb{e}_i$ we denote the canonical basis vectors for both $\mathbb{R}^d$ and $\mathbb{F}_2^d$, unless the underlying vector space is not clear from context.

The notation $x \leq y$ for two vectors $x, y \in \mathbb{R}^d$ is meant component-wise, i. e., we say that $x \leq y$ if and only if $x_i \leq y_i$ for all $i \in [d]$ is true.

For a given vector space $\mathcal{V}$ (usually $\mathbb{R}^d$ or $\mathbb{F}_2^d$) over a field $\mathbb{K}$ (say, $\mathbb{R}$ or $\mathbb{F}_2$, respectively) and some set $S \subseteq \mathcal{V}$ of vectors, we denote by $\operatorname{span}(S)$ the linear hull of S, i. e., the smallest linear subspace that contains all elements of S. This is the same as the set of all finite linear combinations of elements of S. Likewise, we denote by $\operatorname{aff}(S)$ and $\operatorname{conv}(S)$ the affine and convex hull of S, which are the smallest affine and convex sets that contain S, respectively. Linear and affine subspaces that have dimension $\dim(V) - 1$ are called *hyperplanes*. Equivalently, hyperplanes are exactly those (non-empty) sets of $\mathcal{V}$ that are defined as the points of $\mathcal{V}$ which fulfill a single linear equation $\langle a, x\rangle = \beta$.

Boolean Algebra

In Chapter 3, in particular in Theorem 3.2 as well as the Corollaries 3.6 and 3.12, we make use of *Boolean formulas*, that is, formulas that are built (and well-formed) using $q \geq 1$ input *variables* $z_1, \dots, z_q$, conjunctions $\wedge$, disjunctions $\vee$ and negations $\neg$.

Given a Boolean formula φ, we interpret it as a function $\{0,1\}^q \to \{0,1\}$ by assigning a *truth value* to every variable z_j for $j \in [q]$. Here we use the truth values $\{\perp, \top\}$ representing the values false and true, respectively. We usually identify $\perp$ with 0 and $\top$ with 1 and use them synonymous, except for cases where the usage of these numbers might lead to confusion, like in the proof of Theorem 3.2. Evaluating a Boolean formula in the standard way after applying the assignment $x \in \{0,1\}^q$ gives the result $\varphi(x)$. The set $T_\varphi := \{x \in \{0,1\}^q \mid \varphi(x) = 1\}$ is said to be *defined* by φ.

We say that a Boolean formula φ is in *conjunctive normal form* if φ can be written as

$$\varphi = \bigwedge_{i=1}^{p} \bigvee_{j : A_{i,j} = 1} \lambda_j$$

for some binary matrix $A \in \{0,1\}^{p \times q}$, where λ_j is either z_j or $\neg z_j$. Every λ_j is called a *literal*, and every disjunction of literals is called a *clause*, of which there are $p \geq 1$ many.

A variable of a Boolean formula in conjunctive normal form is called *free* if neither it nor its negation do occur as any literal within the formula. If a Boolean formula in conjunctive normal form has no free variables, we call it *complete*. It is easy to prove that every Boolean formula has an equisatisfiable complete conjunctive normal form. This can be achieved by enumerating all possible allocations of truth values to the input variables, constructing a clause for each allocation that results in the Boolean formula being false. A common construction method involves taking as literals the negation of variables assigned true and the variables themselves if assigned false. That way, every variable or its negation appears exactly once for each clause. The conjunction of these clauses is then in complete conjunctive normal form and has the same truth table, making it equisatisfiable to the original Boolean formula. This particular construction is sometimes also called *canonical conjunctive normal form*.

We call a boolean formula in conjunctive normal form k-SAT if every clause has exactly k literals. This is of course related to the 3-*satisfiability* or 3-SAT-*problem*, which is the problem to decide whether a 3-SAT Boolean formula ψ has a variable assignment so that every clause in ψ is satisfied, or equivalently, whether there is an $x \in \{0,1\}^q$ such that $\psi(x) = 1$. It is well-known that this problem is NP-complete [18, 34].

An encyclopedia on satisfiability and related concepts is available from Biere et al. [9].

Graph Theory

For a comprehensive anthology of graph theory we refer to the book of Diestel [22], but we recapitulate some of its notions here.

By $G = (V, E)$ we denote an undirected (simple) *graph*, consisting of a (usually non-empty) finite set V of *vertices* or *nodes* and a subset $E \subseteq \binom{V}{2}$ of *edges* that are themselves subsets of V of cardinality two. A graph $G = (W, F)$ is sometimes called a graph *on* W. Its nodes W are also denoted as $V(G)$ and its edges F as $E(G)$.

A *subgraph* of a graph $G = (V, E)$ is another graph $G' = (V', E')$ where $V' \subseteq V$ and $E' \subseteq E$. We occasionally denote this by $G' \subseteq G$. Given a graph $G = (V, E)$ and two nodes $v_0, v_k \in V$, a v_0-v_k-*walk* is a non-empty alternating sequence of nodes and edges $v_0, e_0, v_1, e_1, \dots, e_{k-1}, v_k$, such that $e_i = \{v_i, v_{i+1}\}$ for all $i \in [k-1]$. Sometimes we identify a walk with its (unordered) subgraph $H \subseteq G$ with node set $V(H) = \{v_0, \dots, v_k\}$ and edge set $E(H) = \{\{v_0, v_1\}, \dots, \{v_{k-1}, v_k\}\}$. If v_0 and v_k are clear from context, we omit them and just call the subgraph a *walk*. By abuse of notation, since we can recoup the nodes from the edges of a walk, we may also refer to the edge-set itself as a walk. If all nodes of a v_0-v_k-walk are distinct, we call it a v_0-v_k-*path*. Again, if v_0 and v_k are clear from the surrounding context, we omit them and just speak of a *path*. A walk with all nodes distinct except for $v_0 = v_k$ is called a *cycle*. The *length* of a walk (or path, or cycle) H is the cardinality of $E(H)$, i. e., if $|E(H)| = k$ then the walk (or path, or cycle) has length k.

We call two (distinct) nodes *adjacent* when they are connected by an edge, i. e., in a graph $G = (V, E)$, the nodes $v \in V$ and $w \in V$ are adjacent if and only if $\{v, w\} \in E$. Likewise, we call two edges *incident* when they have a vertex in common, meaning that two edges $e, f \in E$ are incident if and only if there exists some $v \in e \cap f$. The relation $v \in e$ for $v \in V$ and $e \in E$ is also called an *incidence*. The set of neighboring edges of a node $v \in V$ is denoted by $\delta(v)$, the set of adjacent nodes or *neighbors* of $v \in V$ is denoted by $\mathrm{N}(v)$.

Sometimes we also need *directed* graphs or *digraphs* $D = (V, A)$, where the set of edges is replaced by a set $A \subseteq V \times V \setminus \{(v, v) \mid v \in V\}$ of directed edges or *arcs*. The first part of any arc is called the *tail*, while the second component is the *head* of an arc. Analogously, one defines *directed* subgraphs, walks, paths, cycles and the other associated graph notions. The set of arcs incident to a node $v \in V$ is partitioned into the in-arcs $\delta^{\text{in}}(v)$ and out-arcs $\delta^{\text{out}}(v)$ depending on whether v is the head or tail of an arc (in this order). Likewise, we do the same with the out-neighbors $\mathrm{N}^{\text{out}}(v)$ and in-neighbors $\mathrm{N}^{\text{in}}(v)$.

For any graph $G = (V, E)$, the vector spaces $\mathbb{F}_2^V$ and $\mathbb{F}_2^E$ are called *vertex-* or *node-space* and *edge-space*, respectively. Addition over these spaces can be interpreted as taking the symmetric difference of the corresponding vertex- and edge-subsets. The characteristic vectors of singleton nodes and edges form bases of these spaces.

The *complete graph* on k nodes is denoted by $K_k = ([k], E_k = \binom{[k]}{2})$ for some $k \in \mathbb{N}$. Further definitions, especially of sub-structures of graphs that are of interest in this work, will be introduced as needed. Among them are stable sets, spanning trees, Hamiltonian cycles and flows.

Polyhedral Combinatorics and Optimization

In this section, we recapitulate the basic notions from polyhedral combinatorics and their relation to optimization. Since we cannot give a comprehensive review of all important definitions from the literature, we will mention some reference works. For a thorough treatise about polytopes and polyhedral combinatorics we suggest Ziegler [53, 54]. Reference texts leaning more towards the theory of linear programming and the interface between combinatorics and optimization are written by Grötschel, Lovasz, and Schrijver [30], for example. Schrijver [48] also published a monumental encyclopedic reference monograph documenting the close relation between integer programming and combinatorial optimization. Moreover, Conforti, Cornuejols, and Zambelli [15] further deal with the theory of integer programming, covering the foundations of relaxations, valid inequalities and hierarchies in particular, which we employ in Chapters 4 and 5. Furthermore, a recent reference book aimed towards the algorithmic treatment of combinatorial optimization problems is given by Korte and Vygen [35].

A *polytope* in $\mathbb{R}^d$ is the convex hull of a finite subset of $\mathbb{R}^d$. It is foundational (but not trivial) that this definition (called the *inner description*) is equivalent to a polytope being the bounded intersection of finitely many *half-spaces* of $\mathbb{R}^d$ (also called the *outer description*). Half-spaces themselves are sets that contain all elements $x \in \mathbb{R}^d$ that satisfy some linear (affine) inequality $\langle a, x \rangle \leq \beta$. Therefore, polytopes can also be described by some system of inequalities, denoted as $Ax \leq b$. Transforming one description into the other is a non-trivial task. A description is said to be *irreducible* or *irredundant* if there is no smaller description (in terms of fewer elements or linear inequalities) that leads to the same polytope (as a subset of $\mathbb{R}^d$).

If an inequality $\langle a, x \rangle \leq \beta$ defines a half-space that contains a given polytope P, we say the inequality is *valid* for P. A *face* of P is any set of the form $F = P \cap \{x \in \mathbb{R}^d \mid \langle a, x \rangle = \beta\}$, where $\langle a, x \rangle \leq \beta$ is a valid inequality for P. The *dimension* of a face is the dimension of its affine hull. Note that since $\langle \mathbb{0}, x \rangle \leq 0$ is a valid inequality for any polytope P, it is itself a face of P. All other faces are called *proper faces*. For the inequality $\langle \mathbb{0}, x \rangle \leq 1$, we get the face $\emptyset$, which is always a face of P. The faces $\emptyset$ and P are called *trivial* faces.

The faces of dimension 0 are called *vertices* of a polytope P, while faces of dimension 1 are called *edges*. The notation is similar to graphs, which is intentional. For example, we say that two vertices of P are *adjacent* if the smallest common face containing both of them is an edge of P. On the other end, faces of dimension $\dim(P) - 1$ are called *facets*. An irreducible outer description of a full-dimensional polytope P consists of exactly one inequality for each facet of P.

The *face lattice* of a polytope P is the set of all faces of P which is partially ordered by inclusion. For lattice theory, we refer to the work of Grätzer [29]. Two polytopes are called *combinatorially isomorphic* if their face lattices are isomorphic as finite lattices. It is the weakest notion of isomorphism between polytopes we discuss in this thesis.

Since we only deal with polytopes whose vertices belong to $\{0,1\}^d$, we also call them 0/1-polytopes. All faces F of a 0/1-polytope are themselves 0/1-polytopes, which can be written as

$$F = \operatorname{conv}\Big(F \cap \{0,1\}^d\Big).$$

The vertices of a 0/1-polytope P are given by $P \cap \{0,1\}^d$. A work that deals specifically with 0/1-polytopes is given by Ziegler [53].

We will make heavy use of the notion of an *affine isomorphism* [54, p. 5]: Two polytopes $P \subsetneq \mathbb{R}^d$ and $Q \subsetneq \mathbb{R}^n$ are *affinely isomorphic*, denoted by $P \cong Q$, if there is an affine map $f\colon \mathbb{R}^d \to \mathbb{R}^n$ that induces a bijection between points of P and points of Q. Note that f need not be injective or surjective on the ambient spaces. In this scenario, the map f is called an affine isomorphism between P and Q. If such a map f is not bijective between points of P and points of Q but satisfies $f(P) = Q$, we call it a *projection* of P onto Q.

It is easy to prove that two 0/1-polytopes that are affinely isomorphic are also combinatorially isomorphic, but not vice-versa [53, Proposition 7]. Ziegler also defines the stronger notion of 0/1-*isomorphic* 0/1-polytopes: Two 0/1-polytopes P and Q are 0/1-isomorphic if P can be transformed into Q by a sequence of coordinate permutations and *switches*, i. e., replacement of some coordinates x_i by $1 - x_i$. This notion of 0/1-isomorphism between 0/1-polytopes implies that they are also affinely isomorphic since both of the allowed operations are affine isomorphisms, but again not vice-versa, as there are affinely isomorphic polytopes which are not 0/1-isomorphic [53, Proposition 7].

For our purposes, the concept of isomorphy we will use most is that of affine isomorphy: Having an affine isomorphism between two polytopes P and Q allows us to solve a linear optimization problem over Q, i. e., maximizing or minimizing a linear function over Q, by solving an auxiliary linear optimization problem over P and mapping the solution to this auxiliary problem back to the original polytope using the affine isomorphism. Therefore, understanding the structure of affinely isomorphic polytopes yields valuable insight into the structure of linear optimization problems over these polytopes as well. Furthermore, if we prove that two polytopes cannot be combinatorially isomorphic, then they cannot be affinely isomorphic either [53, Proposition 7].

Since combinatorial optimization problems naturally give rise to 0/1-polytopes whose vertices represent solutions to the underlying combinatorial problem, it is valuable to also study these polytopes independently. What follows next is a small selection of polytopes that turn out to be significant in the study of cyclic transversal polytopes.

Parity Polytopes

As a simple starting point, consider vectors $x \in \{0,1\}^n$ that have a certain *parity*. The parity of a 0/1-vector is the parity of its sum of entries. Therefore, we call $x \in \{0,1\}^n$ an *even* vector if $\langle \mathbb{1}, x\rangle$ is even. This is expressible over $\mathbb{F}_2^n$ as $\langle \mathbb{1}, x\rangle = 0$, which we will exploit in the following chapters. The *parity polytope* is then the convex hull of even vectors, i. e.,

$$\mathrm{PAR}(n) \coloneqq \mathrm{conv}\big\{x \in \{0,1\}^n \mid \langle \mathbb{1}, x\rangle \text{ is even}\big\}.$$

Jeroslow [31] showed that an irreducible description for $\mathrm{PAR}(n)$ for $n \geq 4$ contains the "trivial" inequalities $\mathbb{0} \leq x \leq \mathbb{1}$ and all so called *odd-set inequalities*

$$\sum_{i \in S} x_i - \sum_{i \in [n]-S} x_i \leq |S| - 1 \text{ for } S \subseteq [n] \text{ where } |S| \text{ is odd.}$$

Note that any odd-set inequality is invalid for exactly one element in $\{0,1\}^n$. There are exactly 2^{n-1} such inequalities, none of which can be omitted. We will see in Chapter 4 how parity polytopes in particular can be generalized as cyclic transversal polytopes, and how a large class of valid inequalities for cyclic transversal polytopes is a generalization of odd-set inequalities, which we will call *odd-hyperplane inequalities*.

Cut Polytopes

To reiterate some definitions, given a graph $G = (V, E)$ and a set $U \subseteq V$, the set of edges with exactly one node in U is called a *cut* of G. The *cut polytope* of a graph $G = (V, E)$ is the convex hull of characteristic vectors of cuts of G. We denote the cut polytope of G as $\mathrm{CUT}(G)$.

Cut polytopes have been of interest for several decades [8, 10]. An extensive treatise on the geometry of cuts is given by Deza and Laurent [21]. Part 5 of their book deals with cut polytopes in particular. Ziegler [53, Chapter 4] also discusses cut polytopes, particularly describing the affine isomorphism between them and correlation polytopes. It is well-known that cut polytopes can be used to model quadratic unconstrained binary problems [20], which further reinforces their importance in combinatorial optimization.

Parity polytopes form a subclass of cut polytopes. In particular, $\mathrm{PAR}(n)$ is affinely isomorphic to $\mathrm{CUT}(G)$ for the case where G is a cycle of length n. In contrast to parity polytopes, facets and valid inequalities for cut polytopes are only known in various special cases [21, Chapters 27-30].

Cut polytopes are a precursor to the class of binary kernel polytopes, at which we will take a closer look.

Binary Kernel Polytopes

While Oxley [41] is a suitable reference to all things connected to matroid theory for this thesis, we recall here that a *matroid* is a pair $\mathcal{M} = (E, \mathcal{I})$ consisting of a finite set E (called the *ground set*) and a non-empty collection $\mathcal{I}$ of subsets of E, which fulfill the conditions

(1) for every $\widetilde{F} \subseteq F$ with $F \in \mathcal{I}$ we also have that $\widetilde{F} \in \mathcal{I}$ holds, and
(2) for every two sets $X, Y \in \mathcal{I}$ with $|X| < |Y|$, there exists an element $y \in Y \setminus X$ such that $X \cup y \in \mathcal{I}$ holds.

In a seminal paper, Barahona and Grötschel [7] introduce the cycle polytope of a binary matroid in order to develop practically efficient cutting plane algorithms for a number of real-world problems such as solving certain quadratic 0/1-problems or determining the ground state of spin glasses at absolute zero temperature (which has applications in physics) [7, p. 41].

Given a matrix $M \in \{0,1\}^{d \times n}$ and a vector $b \in \{0,1\}^d$, Barahona and Grötschel introduce the polytope defined as the convex hull of $\left\{x \in \{0,1\}^n \;\middle|\; Mx \equiv b \pmod 2\right\}$. They then prove that the essential objects for the investigation of these polytopes are the case where $b = \mathbb{0}$. This special case is what we will call *(binary) kernel polytopes*, denoted by $\mathrm{KP}_2(M)$.

The reason for also calling these polytopes binary matroid cycle polytopes is two-fold. First, the column index set of the matrix M can be seen as the ground set E of a matroid, where a set $S \subseteq E$ is independent if and only if the columns of M corresponding to S are linearly independent over $\mathbb{F}_2^d$. This resulting matroid is binary, since it is represented as the linear matroid over the binary matrix M. Second, a set $C \subseteq E$ in a binary matroid is called a *cycle* if either $C = \emptyset$ or C is the disjoint union of circuits, which themselves are minimally dependent sets.

It is a foundational result that characteristic vectors of cuts of a graph G form a subspace of the edge-space $\mathbb{F}_2^E$ of G [22, Prop. 1.9.2]. In this sense, cut polytopes are a subfamily of binary kernel polytopes, as any binary matrix M whose kernel is equal to this particular subspace induces a binary kernel polytope $\mathrm{KP}_2(M)$ which is affinely isomorphic to the cut polytope $\mathrm{CUT}(G)$.

We cite and present appropriate reasons as to why parity polytopes, cut polytopes and binary kernel polytopes are related to each other in Section 3.1.1, where we also show how they relate to the cyclic transversal polytope framework. We will now introduce this framework in the coming chapter.

2 The Cyclic Transversal Framework

To gain a new and unifying structural insight into the polytopes mentioned in the introduction, we familiarize us with the cyclic transversal framework to generalize them and depict possibilities on what can be done within this framework in this chapter.

2.1 Basic Definitions

A common theme in the description of the previously presented polytopes is the concept of "choosing or discarding" from given lists of elements such that the tuple of chosen elements fulfills some collection of predetermined conditions. These lists are to be interpreted as subsets of a vector space, which for our purposes will be $\mathbb{F}_2^d$ with $d \in \mathbb{N}_0$, and the collection of conditions to be fulfilled are modeled as a (homogeneous) linear equation in this vector space. We call this underlying vector space the *venue space* or just *venue*.

With this in mind, let $n \in \mathbb{N}$ and let $\Xi_1, \dots, \Xi_n \subseteq \mathbb{F}_2^d$ be non-empty subsets of $\mathbb{F}_2^d$ which we will call *blocks*. The non-empty set family of blocks will sometimes be called *block family* or *block configuration* and denoted by $\Pi := (\Xi_1, \dots, \Xi_n)$, and we will use both terms interchangeably. While the block configuration is an (ordered) tuple, most subsequent definitions are independent of the ordering of the blocks. Wherever necessary, we express this by stating that something is valid "possibly after renumbering" of the blocks. The (set of) *cyclic transversals* of these blocks are then defined as

$$\mathrm{CT}(\Pi) := \left\{ \xi \colon [n] \to \mathbb{F}_2^d \;\middle|\; \xi(i) \in \Xi_i \text{ for all } i \in [n] \text{ and } \sum_{i \in [n]} \xi(i) = \mathbb{0} \right\}.$$

Instead of $\mathrm{CT}(\Pi)$, we sometimes use $\mathrm{CT}(\Xi_1, \dots, \Xi_n)$ to denote the same set. The images $\xi(i)$ are also called *block elements*. Sometimes it is convenient to denote a cyclic transversal as $(\xi(i) \mid i \in [n])$ or via a tuple $(\xi(1), \dots, \xi(n)) \in \bigtimes_{i \in [n]} \Xi_i$ of block elements. A cyclic transversal $\xi \in \mathrm{CT}(\Pi)$ in this sense resembles a $d \times n$-matrix M_ξ over $\mathbb{F}_2$ whose columns are $\xi(i)$ for $i \in [n]$ and that fulfills the equation $M_\xi \mathbb{1} = \mathbb{0}$.

The name cyclic transversal is borrowed from combinatorics and matroid theory in particular, in which *transversals* (or *cross-sections*) are a set of (usually pair-wise distinct) representatives of a given family of sets, like the independent sets of transversal matroids. While our cyclic transversals taken from a family of blocks do not have to contain pair-wise distinct elements, some authors (in particular, Brualdi [12, Chapter 9] as well as Roberts and Tesman [45, Section 12.2.2]) also address systems of non-distinct representatives in this way, so we chose the name *transversal* also for this case.

The modifier *cyclic* then signifies the condition that the tuple of block elements has to fulfill, namely the homogeneous linear equation.

2.1 Definition (Cyclic map)
Let I be some finite set and V be a vector space. We then call a map $\tau\colon I \to V$ *cyclic* if it fulfills the equation $\sum_{i\in I} \tau(i) = \mathbb{0}$.

Although we defined cyclic maps with respect to vector spaces, Abelian groups would suffice to define them. Cyclic maps form a subgroup of V^I under pointwise addition with respect to V.

The equation $\sum_{i\in[n]} \xi(i) = \mathbb{0}$ for cyclic transversals certainly implies that the tuple of block elements is linearly dependent, which means that it contains a circuit in the language of matroid theory. That in turn is an abstraction of a cycle in a graph.

In the plane over the real numbers, the choice of naming such tuples of vectors *cyclic* is apparent by taking a tuple of finitely many real vectors that sum up to zero and then drawing straight line segments between consecutive partial sums of this tuple, starting (and necessarily ending) in zero. The resulting figure will consist of a collection of closed piecewise linear curves, i. e., a collection of (piecewise linear approximations of) cycles. Equivalently, the chosen tuple forms a polygon that cycles around the zero vector. These visualizations are shown in Figure 2.1.

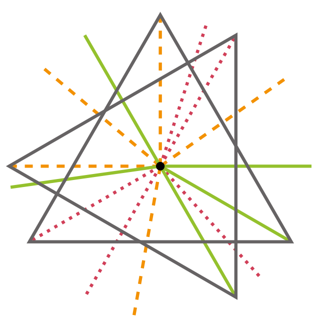

(a) Two cyclic transversals are indicated.

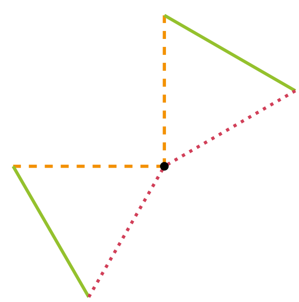

(b) Showing consecutive partial sums of elements.

Figure 2.1: Visualization of cyclic transversals over the real numbers via unit length segments. Dotted, dashed and solid lines indicate the three different blocks of the configuration.

Three important parameters to differentiate block configurations are

(1) the dimension $r := \dim \operatorname{span}\left(\bigcup_{i\in[n]} \Xi_i\right)$ of the smallest subspace of $\mathbb{F}_2^d$ that contains all blocks,
(2) the number n of blocks given, and
(3) the sum $s := \sum_{i\in[n]} |\Xi_i|$ of cardinalities of the blocks.

They are called *rank*, *order* and *size* of the family of blocks and its set of cyclic transversals.

Cyclic transversals naturally form polytopes by taking the incidence vectors of $\xi(i) \in \Xi_i$ for every block, merging them into a tuple of incidence vectors and taking these merged 0/1-vectors as vertices. This construction via incidence vectors results in the *cyclic transversal polytope* of $\Pi = (\Xi_1, \dots, \Xi_n)$, denoted $\mathrm{CTP}(\Pi)$ or $\mathrm{CTP}(\Xi_1, \dots, \Xi_n)$, that can be written as

$$\mathrm{CTP}(\Pi) := \operatorname{conv}\left\{ (\chi(\xi(i)))_{i\in[n]} \in \bigtimes_{i\in[n]} \{0,1\}^{\Xi_i} \;\middle|\; \xi \in \mathrm{CT}(\Pi) \right\}.$$

We address the components of all points $y \in \bigtimes_{i\in[n]} \mathbb{R}^{\Xi_i}$ as $y^i_\omega \in \mathbb{R}$ for all $i \in [n]$ and $\omega \in \Xi_i$. When dealing with cyclic transversal polytopes and coordinates of these vectors, we implicitly assume an ordering on the combinatorial objects at play, like an ordering of the blocks inherited from their indices. The influence of certain coordinate permutations is discussed in Section 2.2.

The notions of *rank*, *order* and *size* of cyclic transversal polytopes are adopted from the respective definitions for cyclic transversals. The size of a cyclic transversal polytope is equal to the dimension of its ambient space $\bigtimes_{i\in[n]} \mathbb{R}^{\Xi_i}$. Their own dimension is at most their size minus their order, as every block Ξ_i implies a valid equation $\sum_{\omega\in\Xi_i} y^i_\omega = 1$ for the cyclic transversal polytope called (i-th) *block equation*, which in turn implies that

$$\dim \mathrm{CTP}(\Xi_1, \dots, \Xi_n) := \dim \operatorname{aff} \mathrm{CTP}(\Xi_1, \dots, \Xi_n) \leq \sum_{i\in[n]} |\Xi_i| - n = s - n \; .$$

Since every cyclic transversal polytope fulfills (at least) these block equations, it is sensible to incorporate them into our framework:

2.2 Definition (Affine transversal space)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a block configuration. The affine space in $\bigtimes_{i\in[n]} \mathbb{R}^{\Xi_i}$ that is defined by the block equations of Π is called *affine transversal space* and is denoted by $\mathbb{A}(\Pi)$, that is,

$$\mathbb{A}(\Pi) := \bigtimes_{i\in[n]} \left\{ y^i \in \mathbb{R}^{\Xi_i} \;\middle|\; \sum_{\omega\in\Xi_i} y^i_\omega = 1 \right\}.$$

The 0/1-vectors in $\mathbb{A}(\Pi)$ are in 1-to-1 correspondence to the elements of $\bigtimes_{i\in[n]} \Xi_i$, a subset of which is isomorphic to the cyclic transversals of Π by the embedding $\xi \mapsto (\xi(1), \dots, \xi(n))$. This is why, unless stated otherwise, we consider $\mathrm{CTP}(\Pi)$ as a subset of $\mathbb{A}(\Pi)$ instead of $\bigtimes_{i\in[n]} \mathbb{R}^{\Xi_i}$. This slight modification of the underlying space will be more important in Chapter 5, where we discuss polytopes which are not necessarily contained in $\mathbb{A}(\Pi)$, but they contain $\mathrm{CTP}(\Pi)$.

Depending on the choice of blocks, various polytopes (especially those defined in the introduction) are manifestations of certain cyclic transversal polytopes, which we see in Chapter 3. There in Section 3.1 we will also see that cyclic transversal polytopes are (in a sense) universal 0/1-polytopes, if we allow projections. In Section 3.2, we establish some necessary criteria for a polytope to be a cyclic transversal polytope itself, without allowing projections, and give examples of polytopes that cannot be represented as cyclic transversal polytopes, giving us a classification of all polytopes up to dimension three that can and cannot be represented as cyclic transversal polytopes. After that, in Chapter 4, the subclass of full cyclic transversal polytopes is examined, culminating in complete descriptions of polytopes in this subclass when the rank or order is small.

But first, we establish some combinatorial properties including relations between the rank, order and size of cyclic transversals and their polytopes. This includes an investigation on possible simplifications of their presentation which helps with their computational enumeration, and a characterization of adjacency of vertices of cyclic transversal polytopes.

2.2 Combinatorial Operations and Properties

In order to better speak about the combinatorics of cyclic transversals and their polytopes, we introduce the notions of *equivalence* and *reduction* of block configurations that translate to similar notions on cyclic transversals and their polytopes. To unify the presentation and describe canonical candidates for the different resulting equivalence classes of cyclic transversal polytopes, we also define *normalization* of cyclic transversal sets and make some assumptions on the cyclic transversals and their polytopes later in this section, which will be used throughout this work and are mostly justified by the existence of equivalent or reduced block configurations that fulfill these assumptions.

2.2.1 Equivalence

A given block configuration Π with its associated set of cyclic transversals and cyclic transversal polytope can be transformed in various ways such that the transformed block configuration generates, in some sense, an isomorphic set of cyclic transversals and an affinely isomorphic cyclic transversal polytope, respectively. We want to call the original and transformed block configurations *equivalent* to each other. Furthermore, these transformations should preserve the rank, order and size of the original block configuration.

2.3 Definition (Equivalent block configurations)
Two block configurations $\Pi = (\Xi_1, \dots, \Xi_n)$ and $\widetilde{\Pi} = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_n)$ are called *equivalent* if and only if there is a bijective linear map

$$\varphi\colon \underbrace{\operatorname{span}\left(\bigcup_{i\in[n]} \Xi_i\right)}_{\subseteq \mathbb{F}_2^d} \to \underbrace{\operatorname{span}\left(\bigcup_{i\in[n]} \widetilde{\Xi}_i\right)}_{\subseteq \mathbb{F}_2^{\tilde{d}}}$$

and a cyclic map $\tau\colon [n] \to \mathbb{F}_2^{\tilde{d}}$, such that, possibly after some renumbering of blocks,

$$\widetilde{\Xi}_i = \{\varphi(\omega) \oplus \tau(i) \mid \omega \in \Xi_i\} \text{ for all } i \in [n].$$

We call the pair (φ, τ) an *equivalence transformation* of block configurations.

Remember that we usually denote the addition of two elements of $\mathbb{F}_2^d$ by $\oplus$ instead of $+$. By insisting on bijective linear maps, i. e., isomorphisms, on the linear spans of the blocks as part of equivalence transformations, one guarantees that while the dimensions of the underlying venue spaces are not fixed, the ranks and sizes of equivalent block configurations are equal, and the order is preserved by mapping each block separately.

2.4 Example
By taking the zero cyclic map or the trivial isomorphism, i. e., the identity map, as part of an equivalence transformation for block configurations, we obtain two special cases, possibly after some renumbering of blocks:

(1) *Block isomorphism*: An isomorphism $\varphi\colon \operatorname{span}\left(\bigcup_{i\in[n]} \Xi_i\right) \to \operatorname{span}\left(\bigcup_{i\in[n]} \widetilde{\Xi}_i\right)$, such that
$$\widetilde{\Xi}_i = \varphi(\Xi_i) = \{\varphi(\omega) \mid \omega \in \Xi_i\} \text{ for all } i \in [n].$$

(2) *Block translation*: A cyclic map $\tau\colon [n] \to \mathbb{F}_2^d$, identified via some tuple $(\tau(1), \dots, \tau(n))$ of vectors in $\mathbb{F}_2^d$ with $\sum_{i\in[n]} \tau(i) = \mathbb{0}$, such that
$$\widetilde{\Xi}_i = \Xi_i \oplus \tau(i) = \{\omega \oplus \tau(i) \mid \omega \in \Xi_i\} \text{ for all } i \in [n].$$

Equivalence transformations of cyclic transversals can be seen as applying some block isomorphism followed by some block translation to the entire block configuration. A block translation is essentially adding a productive zero to the linear equation that defines the set of cyclic transversals, and distributing its parts to the blocks in some way. Additionally, block isomorphisms do not change the validity of the defining linear equation for cyclic maps. This is why we allow these special bijective affine transformations of the venue space as equivalence transformations of block configurations.

The equivalence relation of block configurations as defined in Definition 2.3 is transitive, which is easily verified via the composition rule

$$(\psi, \sigma) \circ (\varphi, \tau) := (\psi \circ \varphi, (\psi \circ \tau) \oplus \sigma)$$

of the corresponding equivalence transformations, where $\psi \circ \varphi$ is an isomorphism, since the composition of bijective linear maps is bijective and linear. Since ψ is a bijective linear map, $(\psi \circ \tau) \oplus \sigma$ is described by $((\psi \circ \tau) \oplus \sigma)(i) = \psi(\tau(i)) \oplus \sigma(i)$ with the property

$$\sum_{i \in [n]} ((\psi \circ \tau) \oplus \sigma)(i) = \psi(\sum_{i \in [n]} \tau(i)) \oplus \sum_{i \in [n]} \sigma(i) = \psi(\mathbb{0}) \oplus \mathbb{0} = \mathbb{0},$$

hence $(\psi \circ \tau) \oplus \sigma$ is a cyclic map.

By this composition rule, it is also easy to establish the symmetry of this equivalence relation. Given an equivalence transformation pair (φ, τ) consisting of a block isomorphism and a block translation, the inverse transformation is $(\varphi^{-1}, \varphi^{-1} \circ \tau)$. Furthermore, any block configuration is equivalent to itself by taking the identity map and the cyclic map $\tau\colon [n] \to \mathbb{0}$ as an equivalence transformation, which means that these maps are also reflexive.

In general, equivalence transformations between two given block configurations are not unique. Indeed, this is already evident by the fact that there are multiple automorphisms on $\mathbb{F}_2^d$.

We now directly transfer the notion of equivalence to cyclic transversals and cyclic transversal polytopes:

2.5 Definition (Equivalent cyclic transversals and cyclic transversal polytopes)
Two sets of cyclic transversals are called *equivalent* if they are generated by two equivalent block configurations Π and $\widetilde{\Pi}$. Likewise, we call two cyclic transversal polytopes *equivalent* if they are derived from equivalent block configurations.

We want to treat equivalent sets of cyclic transversals as essentially the same set, since an equivalence transformation maps cyclic transversals in a straightforward way. More precisely, equivalent sets of cyclic transversals are described as follows:

2.6 Proposition

Let $\Pi = (\Xi_1, \dots, \Xi_n)$ *be a block configuration and let* $\widetilde{\Pi} = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_n)$ *be a block configuration equivalent to* Π*. Then, possibly after some renumbering of blocks,*

$$\mathrm{CT}(\widetilde{\Pi}) = \left\{ (\varphi \circ \xi) \oplus \tau \colon [n] \to \mathbb{F}_2^d \;\middle|\; \xi \in \mathrm{CT}(\Pi) \right\},$$

where (φ, τ) *is a pair of block isomorphism and block translation that describes the equivalence transformation from* Π *to* $\widetilde{\Pi}$*.*

Proof: On the one hand, every map of the form $\varphi(\xi) \oplus \tau$, where ξ is a cyclic transversal in $\mathrm{CT}(\Pi)$, is itself a cyclic transversal of $\widetilde{\Pi}$. Indeed, $\varphi(\xi(i)) \oplus \tau(i) \in \widetilde{\Xi}_i$ is true for all $i \in [n]$ because of the equivalence relation between Π and $\widetilde{\Pi}$. Furthermore,

$$\sum_{i\in[n]} \varphi(\xi(i)) \oplus \tau(i) = \varphi\left(\sum_{i\in[n]} \xi(i)\right) \oplus \sum_{i\in[n]} \tau(i) = \varphi(\mathbb{0}) \oplus \mathbb{0} = \mathbb{0}$$

holds because φ is linear and ξ and τ are cyclic maps.

On the other hand, let $\tilde{\xi} \in \mathrm{CT}(\widetilde{\Pi})$ be a cyclic transversal of $\widetilde{\Pi}$. Then, for all $i \in [n]$, the block element $\tilde{\xi}(i)$ is of the form $\varphi(\omega_i) \oplus \tau(i) \in \widetilde{\Xi}_i$ for some $\omega_i \in \Xi_i$, since $\widetilde{\Pi}$ is equivalent to Π. Plugging this into the equation for cyclic transversals we get

$$\mathbb{0} = \sum_{i\in[n]} \tilde{\xi}(i) = \sum_{i\in[n]} \varphi(\omega_i) \oplus \sum_{i\in[n]} \tau(i) = \sum_{i\in[n]} \varphi(\omega_i) = \varphi\left(\sum_{i\in[n]} \omega_i\right),$$

and $\varphi(\sum_{i\in[n]} \omega_i) = \mathbb{0}$ if and only if $\sum_{i\in[n]} \omega_i = \mathbb{0}$, since φ is a bijective linear map. This means that the map $\xi \colon [n] \to \mathbb{F}_2^d$ that is defined via $\xi(i) = \omega_i$ for all $i \in [n]$ is a cyclic transversal of Π. ■

Proposition 2.6 shows that there is an obvious relationship between cyclic transversals of equivalent block configurations. Fundamentally, equivalence transformations of block configurations manifest as certain coordinate permutations of the underlying vector space of the generated cyclic transversal polytopes, which we will now prove:

2.7 Proposition

Let $\Pi = (\Xi_1, \dots, \Xi_n)$ *be a block configuration and* $P = \mathrm{CTP}(\Pi) \subsetneq \mathbb{A}(\Pi)$ *be its cyclic transversal polytope. Then for any equivalent block configuration* $\widetilde{\Pi} = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_n)$*, the cyclic transversal polytope* $Q = \mathrm{CTP}(\widetilde{\Pi}) \subsetneq \mathbb{A}(\widetilde{\Pi})$ *is given by a permutation of the coordinates of* P*.*

Proof: After renumbering the blocks using a permutation $\pi \colon [n] \to [n]$, the equivalence between Π and $\widetilde{\Pi}$ is described by a pair of block isomorphism and block translation. Without loss of generality, we assume $\pi(i) = i$ for all $i \in [n]$, as otherwise we just apply this permutation to the appropriate indices.

Let (φ, τ) be the equivalence transformation to transform Π into $\widetilde{\Pi}$. We construct a coordinate permutation map $p\colon \mathbb{A}(\widetilde{\Pi}) \to \mathbb{A}(\Pi)$ via $p(y)^i_\omega = y^i_{\varphi(\omega)\oplus\tau(i)}$.

This is a coordinate permutation such that the vertices of P and Q are mapped bijectively to one another, i. e., for every vertex v of P with associated cyclic transversal ξ, the map $\tilde{\xi}\colon [n] \to \mathbb{F}_2^d$, defined by $\tilde{\xi}(i) = \varphi(\xi(i)) \oplus \tau(i)$, is a cyclic transversal of $\widetilde{\Pi}$, and its corresponding vertex $\tilde{v}$ of Q is given by $p(\tilde{v}) = v$.

It follows that p maps Q to P and is an isomorphism, as it is a coordinate permutation. Its inverse is the asserted isomorphism that maps P to Q. ■

This notion of equivalence of cyclic transversal polytopes is very restrictive. In contrast to 0/1-isomorphy as defined by Ziegler [53], our notion of equivalence for cyclic transversal polytopes includes only certain coordinate permutations and not switches (which is the replacing of coordinates x_i by $1 - x_i$ as explained in the introduction), as there is no straightforward concept to emulate switching with block configurations without violating the necessary block equations of cyclic transversal polytopes.

Since our notion of equivalence of cyclic transversal polytopes implies that they are also 0/1-isomorphic polytopes, which in turn implies [53, Proposition 7] that they are affinely isomorphic, equivalent cyclic transversal polytopes are affinely isomorphic.

The construction of equivalence classes of block configurations will aid us in figuring out which constructions of cyclic transversal polytopes are fundamentally different from one another. Next, we define reduction as a concept similar to equivalence to relate block configurations to one another that will generate the isomorphic cyclic transversal polytopes but may have different ranks, orders or sizes.

2.2.2 Reduction

Similarly to equivalence transformations, there are operations one may apply to a block configuration to get one that is not equivalent, but still generates an affinely isomorphic cyclic transversal polytope regardless, even if these operations possibly alter the order, rank or size of the block configuration. We call a certain set of operations that reduce at least one of these parameters *reductions*. These operations are later used to shed block configurations of unnecessary parts, and to be able to more clearly argue about block configurations that have no superfluous elements or blocks in Section 3.2.

2.8 Definition (Block element deletion)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a block configuration, $i \in [n]$ fixed and $|\Xi_i| \geq 2$. The *deletion* of $\omega \in \Xi_i$ from Π is the block configuration $\widetilde{\Pi} = (\Xi_1, \dots, \Xi_i \setminus \omega, \dots, \Xi_n)$.

2.9 Proposition
Given a block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$ *with* $i \in [n]$ *fixed and* $|\Xi_i| \geq 2$, *then if* $\widetilde{\Pi} = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_n)$ *denotes the deletion of* $\omega \in \Xi_i$ *from* Π, *the set of cyclic transversals of* $\widetilde{\Pi}$ *is exactly* $\mathrm{CT}(\widetilde{\Pi}) = \{\xi \in \mathrm{CT}(\Pi) \mid \xi(i) \neq \omega\}$.

Proof: This proof is basically by definition, we only write it here for the convenience of the reader.

Let $\tilde{\xi} \in \mathrm{CT}(\widetilde{\Pi})$ be a cyclic transversal of $\widetilde{\Pi}$, then it is also a cyclic transversal of Π, since it is a cyclic map by definition and fulfills $\tilde{\xi}(j) \in \widetilde{\Xi}_j = \Xi_j$ for all $j \in [n] \setminus i$, and $\tilde{\xi}(i) \in \widetilde{\Xi}_i = \Xi_i \setminus \omega \subseteq \Xi_i$. Therefore, also $\tilde{\xi}(i) \neq \omega$ holds.

On the other hand, any cyclic transversal $\xi \in \mathrm{CT}(\Pi)$ with $\xi(i) \neq \omega$ is isomorphic to a cyclic transversal of $\widetilde{\Pi}$, again because it is a cyclic map by definition and fulfills $\xi(j) \in \Xi_j = \widetilde{\Xi}_j$ for all $j \in [n] \setminus i$, with the additional property that $\xi(i) \in \Xi_i \setminus \omega = \widetilde{\Xi}_i$. ■

Since it is clear how the deletion of an element from a block affects the set of cyclic transversals, we emphasize elements whose deletion keeps the set of cyclic transversals unchanged.

2.10 Definition (Trivial block element)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a block configuration, $i \in [n]$ fixed and $|\Xi_i| \geq 2$. We call $\omega \in \Xi_i$ a *trivial block element* of Ξ_i if its deletion results in a block configuration $\widetilde{\Pi}$ with the property that $\mathrm{CT}(\Pi) = \mathrm{CT}(\widetilde{\Pi})$.

The removal of trivial block elements does not have any influence in the set of cyclic transversals by definition. All other elements are *non-trivial*, their removal would shrink the set of cyclic transversals. Non-trivial block elements are identifiable by examining the cyclic transversals of a block configuration in the following way:

2.11 Proposition
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ *be a block configuration. Then for any block* Ξ_i *with* $|\Xi_i| \geq 2$, *the set* $\{\omega \in \Xi_i \mid \forall \xi \in \mathrm{CT}(\Pi) \colon \xi(i) \neq \omega\}$ *contains the trivial elements of* Ξ_i, *and its complement* $\{\omega \in \Xi_i \mid \exists \xi \in \mathrm{CT}(\Pi)$ *with* $\xi(i) = \omega\}$ *contains the non-trivial elements of* Ξ_i.

Proof: Let us define $S_i \coloneqq \{\omega \in \Xi_i \mid \exists \xi \in \mathrm{CT}(\Pi) \text{ with } \xi(i) = \omega\}$. We will show that its complement $\Xi_i \setminus S_i = \{\omega \in \Xi_i \mid \forall \xi \in \mathrm{CT}(\Pi) \colon \xi(i) \neq \omega\} \subseteq \Xi_i$ contains exactly the trivial block elements, which also shows that S_i contains the non-trivial elements.

If $\omega \in \Xi_i$ is a trivial block element, then its deletion does not change the set of cyclic transversals by definition. Therefore, by Proposition 2.9, we know that $\mathrm{CT}(\Pi)$ is equal to $\{\xi \in \mathrm{CT}(\Pi) \mid \xi(i) \neq \omega\}$. This is equivalent to the assertion that for all $\xi \in \mathrm{CT}(\Pi)$, we have $\xi(i) \neq \omega$, so $\omega \notin S_i$.

Conversely, every element $\omega \in \Xi_i \setminus S_i$ fulfills the condition that no $\xi \in \mathrm{CT}(\Pi)$ has the property $\xi(i) = \omega$. Therefore, $\mathrm{CT}(\Pi) = \{\xi \in \mathrm{CT}(\Pi) \mid \xi(i) \neq \omega\}$, which means that ω is trivial. ■

It is a direct consequence of Proposition 2.11 that the deletion of a trivial block element $\omega \in \Xi_i$ reduces the amount of trivial elements of Ξ_i by one, while the number of non-trivial elements of Ξ_i remains the same, i. e., block element deletions do not change non-trivial block elements into trivial ones on Ξ_i. Since this characterization via cyclic transversals is index-wise and the deletion of a trivial block element does not change the set of cyclic transversals of a configuration by Definition 2.10, the deletion of a trivial block element in one block does not change the trivial elements in other blocks, either.

In that way, trivial block elements are inherent to a chosen block configuration and are invariant under equivalence transformations:

2.12 Proposition
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a block configuration and let $\widetilde{\Pi} = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_n)$ be a block configuration equivalent to Π. Then an associated equivalence transformation maps trivial (and non-trivial, respectively) block elements of Ξ_i to trivial (and non-trivial, respectively) block elements of $\widetilde{\Xi}_i$.

Proof: Let (φ, τ) be an equivalence transformation between Π and $\widetilde{\Pi}$, and let $\kappa_i \colon \Xi_i \to \widetilde{\Xi}_i$ with $\kappa_i(\omega) = \varphi(\omega) \oplus \tau(i)$ be the corresponding bijective maps that transform Ξ_i into $\widetilde{\Xi}_i$ under the equivalence transformation, possibly after some renumbering of blocks, for each $i \in [n]$. We have to show that $\kappa_i(\omega)$ is a trivial block element of $\widetilde{\Xi}_i$ if and only if $\omega \in \Xi_i$ is trivial.

First the “only if” part, which we prove by contraposition: If $\omega \in \Xi_i$ is non-trivial, there exists a cyclic transversal $\xi \in \mathrm{CT}(\Pi)$ such that $\xi(i) = \omega$ by Proposition 2.11. We have to show that there also exists some $\tilde{\xi} \in \mathrm{CT}(\widetilde{\Pi})$ such that $\tilde{\xi}(i) = \kappa_i(\omega)$. But this is evident by the fact that transforming the equation $\xi(i) = \omega$ by applying φ on both sides and adding $\tau(i)$ on both sides afterwards yields $\varphi(\xi(i)) \oplus \tau(i) = \kappa_i(\omega)$ by definition of κ_i. This transformation is possible since φ is an isomorphism on the span of blocks. Finally, $\varphi(\xi) \oplus \tau$ is a cyclic transversal of $\widetilde{\Pi}$ by Proposition 2.6, so $\kappa_i(\omega)$ is non-trivial.

If $\omega \in \Xi_i$ is trivial, then all cyclic transversals ξ of Π fulfill $\xi(i) \neq \omega$. We again apply the equivalence transformation to obtain the condition $\varphi(\xi(i)) \oplus \tau(i) \neq \kappa_i(\omega)$ for all $\xi \in \mathrm{CT}(\Pi)$. Since we know that $\mathrm{CT}(\widetilde{\Pi})$ is exactly equal to the set $\{\varphi(\xi) \oplus \tau \colon [n] \to \mathbb{F}_2^{\tilde{d}} \mid \xi \in \mathrm{CT}(\Pi)\}$ by Proposition 2.6, this is the condition that the element $\kappa_i(\omega)$ is trivial. ■

Trivial block elements are elements that are never used in any cyclic transversal of a block configuration. Likewise, we define elements that are always used in a cyclic transversal:

2.13 Definition (Fixed block and block element)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a block configuration and let $i \in [n]$. If every cyclic transversal $\xi \in \mathrm{CT}(\Pi)$ fulfills $\xi(i) = \omega$ for some block element $\omega \in \Xi_i$, then we call ω a *fixed block element* and Ξ_i a *fixed block*. Otherwise, we call the block *unfixed*.

A fixed block Ξ contains exactly one fixed element ω, and all elements $\Xi \setminus \omega$ are trivial block elements, since they do not occur in any cyclic transversal by definition. Therefore, by applying trivial block element deletions, it is possible to transform any fixed block into a block of cardinality 1. Conversely, any block with at least two non-trivial elements is unfixed.

We disclose that, just like trivial block elements in Proposition 2.12, fixed blocks are inherent to a block configuration and equivalence transformations do not change the property that a block is fixed:

2.14 Proposition
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ *be a block configuration and let* $\widetilde{\Pi} = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_n)$ *be a block configuration equivalent to* Π*. Then any corresponding equivalence transformation maps fixed blocks of* Π *to fixed blocks of* $\widetilde{\Pi}$*.*

Proof: Let (φ, τ) be a corresponding equivalence transformation between Π and $\widetilde{\Pi}$. Without loss of generality, we assume that the equivalence transformation transforms Ξ_i into $\widetilde{\Xi}_i$, such that indices are identical. We now need to prove that Ξ_i is fixed if and only if $\widetilde{\Xi}_i$ is fixed as well.

For any fixed block Ξ_i, there is an element $\omega_i \in \Xi_i$ such that for any cyclic transversal $\xi \in \mathrm{CT}(\Pi)$ we have that $\xi(i) = \omega_i$. By applying the equivalence transformation, we see that every cyclic transversal $\tilde{\xi} \in \mathrm{CT}(\widetilde{\Pi})$ satisfies the equation $\tilde{\xi}(i) = \varphi(\omega_i) \oplus \tau(i)$. Therefore, $\widetilde{\Xi}_i$ is a fixed block. ∎

By applying the inverse equivalence transformation $(\varphi^{-1}, \varphi^{-1} \circ \tau)$ to a fixed block $\widetilde{\Xi}_i$, we also get the reverse direction of Proposition 2.14, and since an equivalence transformation does not change the order of a block configuration, Proposition 2.14 remains true if we replace the term *fixed* with *unfixed*. In particular, the number of blocks that are fixed (and consequently, the number of unfixed blocks) is an inherent property of a block configuration that cannot be changed by equivalence transformations.

We now want to extend the analysis of fixed blocks by utilizing the insight after Definition 2.13:

2.15 Definition (Trivial block)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a block configuration. A block Ξ_i with $|\Xi_i| = 1$ is called a *trivial block*.

By applying a suitable block translation, we are able to transform any block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$ with $n \geq 2$ and a trivial block Ξ_i into another block configuration such that $\Xi_i = \{\mathbb{0}\}$ holds.

2.16 Definition (Trivial block deletion)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a block configuration with $n \geq 2$ and Ξ_i be a block of the form $\{\mathbb{0}\}$. The *deletion* of Ξ_i in Π is the block configuration $\widetilde{\Pi} = (\Xi_1, \dots, \Xi_{i-1}, \Xi_{i+1}, \dots, \Xi_n)$.

Since we do not allow empty block configurations, at least one block has to remain after performing a trivial block deletion. With these operations at hand, we define what it means to be a reduction of a block configuration:

2.17 Definition (Reduction of block configurations)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ and $\widetilde{\Pi} = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_{\widetilde{n}})$ be two block configurations. Then $\widetilde{\Pi}$ is called a *reduction* of Π if and only if $\widetilde{\Pi}$ is constructed from Π by a sequence of equivalence transformations as well as deletions of trivial block elements and trivial blocks.

The process of deleting trivial blocks or block elements, possibly after performing an equivalence transformation, is also called *reducing* the block configuration.

Without allowing equivalence transformations in Definition 2.17, there would be block configurations that have no reductions other than themselves, but are equivalent to block configurations that can be further reduced.

2.18 Example
The block configuration $(\mathbb{F}_2^d, \mathbb{F}_2^d, \{\mathbb{1}\})$ does not have trivial block elements, but it has a trivial block, such that it can be transformed into the equivalent block configuration $(\mathbb{F}_2^d, \mathbb{F}_2^d, \{\mathbb{0}\})$ by the block translation $(\mathbb{1}, \mathbb{0}, \mathbb{1})$. This equivalent configuration can then further be reduced to $(\mathbb{F}_2^d, \mathbb{F}_2^d)$. Therefore, $(\mathbb{F}_2^d, \mathbb{F}_2^d)$ is a reduction of $(\mathbb{F}_2^d, \mathbb{F}_2^d, \{\mathbb{1}\})$.

The example shows that an equivalence transformation (φ, τ) of a block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$ is not necessarily an equivalence transformation of a subset of the blocks because $\tau \colon [n] \to \mathbb{F}_2^d$ is not necessarily cyclic on a subset of $[n]$. Therefore, equivalence operations sometimes have to be performed before deletion of trivial block elements or trivial blocks. In fact, we prescribe a partial ordering of reduction steps in Lemma 2.20.

By Proposition 2.14, any two equivalent block configurations have the same number of fixed and unfixed blocks. Since only fixed blocks can be reduced to trivial blocks by deletion of trivial block elements and suitable equivalence transformations, the number of unfixed blocks remains constant even for any reduction of a given block configuration.

Likewise, the number of non-trivial block elements per block remains the same for any reduction of a given block configuration by Proposition 2.12. Since only these non-trivial block elements and unfixed blocks are relevant for the construction of cyclic transversals, we conclude that cyclic transversals of reductions of block configurations can be used to canonically reconstruct the set of cyclic transversals of the original block configuration and therefore show how the set of cyclic transversals behaves when reducing a block configuration. This extends Proposition 2.6, which explains how cyclic transversals behave under equivalence transformations. We call this reconstructive relationship a *canonical bijection* between the cyclic transversals of a block configuration Π and any of its reductions $\widetilde{\Pi}$, which is formalized in the following proposition:

2.19 Proposition
Let $\Pi = \Pi_0, \Pi_1, \dots, \Pi_r = \widetilde{\Pi}$ be a sequence of block configurations for some $r \in \mathbb{N}$, such that Π_i results from Π_{i-1} by performing one of the following three operations: an equivalence transformation, deleting a trivial block element from some block, or deletion of a trivial block for $i \in [r]$.

Then for every $i \in [r]$, there are bijections $\pi_i \colon \mathrm{CT}(\Pi_i) \to \mathrm{CT}(\Pi_{i-1})$ as follows:

(1) If $\Pi_i = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_k)$ is equivalent to $\Pi_{i-1} = (\Xi_1, \dots, \Xi_k)$ with an associated equivalence transformation (φ, τ), then

$$\pi_i(\xi)(j) = \varphi^{-1} \circ (\xi \oplus \tau)(j) \text{ for all } j \in [k].$$

(2) If Π_i is given by deleting a trivial block element from Π_{i-1}, then

$$\pi_i(\xi)(j) = \xi(j) \text{ for all } j \in [k].$$

(3) If $\Pi_i = (\Xi_1, \dots, \Xi_{k-1})$ is given by deleting a trivial block Ξ_k from $\Pi_{i-1} = (\Xi_1, \dots, \Xi_k)$, possibly after some renumbering, then

$$\pi_i(\xi)(j) = \xi(j) \text{ for all } j \neq k \text{ and } \pi_i(\xi)(k) = \mathbb{0}.$$

The canonical bijection $\pi \colon \mathrm{CT}(\widetilde{\Pi}) \to \mathrm{CT}(\Pi)$ then is the concatenation of π_i for $i \in [r]$.

Proof: It suffices to show how the set of cyclic transversals of Π_{i-1} is canonically determined by the cyclic transversals of Π_i for $i \in [r]$.

If $\Pi_i = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_k)$ is equivalent to $\Pi_{i-1} = (\Xi_1, \dots, \Xi_k)$ and (φ, τ) is an associated equivalence transformation from Π_{i-1} to Π_i, then Proposition 2.6 proves that

$$\mathrm{CT}(\Pi_{i-1}) = \{\varphi^{-1} \circ (\xi \oplus \tau) \colon [k] \to \mathbb{F}_2^d \mid \xi \in \mathrm{CT}(\Pi_i)\},$$

where $(\varphi^{-1}, \varphi^{-1} \circ \tau)$ is the inverse equivalence transformation to (φ, τ).

Deletion of a trivial block element does not change the set of cyclic transversals by definition, so in this case, $\mathrm{CT}(\Pi_{i-1}) = \mathrm{CT}(\Pi_i)$ is obvious. Therefore, it remains to show how the set of cyclic transversals of a block configuration behaves after deleting a trivial block, so assume that Π_i emerges from Π_{i-1} by deleting a trivial block. Without loss of generality, let Π_{i-1} be the configuration $(\Xi_1, \dots, \Xi_k)$ on $\mathbb{F}_2^d$, and let Π_i be the configuration $(\Xi_1, \dots, \Xi_{k-1})$, such that the trivial block that is deleted from Π_{i-1} is the last block $\Xi_k = \{\mathbb{0}\}$, possibly after some renumbering of blocks.

The cyclic transversals of Π_{i-1} can then clearly be characterized as

$$\mathrm{CT}(\Pi_{i-1}) = \{\xi\colon [k] \to \mathbb{F}_2^d \mid \xi(k) = \mathbb{0} \text{ and } \exists \tilde{\xi} \in \mathrm{CT}(\Pi_i) \text{ with } \xi(j) = \tilde{\xi}(j)\ \forall j \in [k-1]\},$$

since the validity of the defining equation $\sum_{j\in[k]} \xi(j) = \mathbb{0}$ for cyclic transversals of Π_{i-1} is unaltered after removing a trailing $\mathbb{0}$ from all cyclic transversals to obtain a cyclic transversal of Π_i, and $\sum_{j\in[k-1]} \tilde{\xi}(j) = \mathbb{0}$ for elements $\tilde{\xi} \in \mathrm{CT}(\Pi_i)$ remains valid as well after padding all such cyclic transversals with a trailing $\mathbb{0}$.

This shows that all three reduction operations induce the aforementioned bijections $\pi_i\colon \mathrm{CT}(\Pi_i) \to \mathrm{CT}(\Pi_{i-1})$ for $i \in [r]$. It is clear that concatenation of these bijections for $i \in [r]$ yields the canonical bijection between elements of $\mathrm{CT}(\Pi) = \mathrm{CT}(\Pi_0)$ and elements of $\mathrm{CT}(\widetilde{\Pi}) = \mathrm{CT}(\Pi_r)$. ■

The proof of Proposition 2.19 also shows that the deletion of a trivial block reduces the number of fixed blocks of a block configuration by one and cannot change unfixed blocks into fixed ones. Indeed, since the deletion of a trivial block does not change the number of non-trivial block elements of other blocks because of the canonical bijection between sets of cyclic transversals, the number of unfixed blocks of a block configuration remains invariant under reductions.

Using the knowledge we gathered in the previous results, we may assume a partial ordering in the sequence of steps to perform when reducing a block configuration:

2.20 Lemma
Let Π be a block configuration and $\widetilde{\Pi}$ be any reduction of Π. Then $\widetilde{\Pi}$ can also be obtained from Π by performing a single equivalence transformation first, trivial block element deletions second and trivial block deletions last.

Proof: Let $\Pi = \Pi_0, \Pi_1, \dots, \Pi_r = \widetilde{\Pi}$ for some $r \in \mathbb{N}$ be a sequence of block configurations such that Π_i results from Π_{i-1} by performing an equivalence transformation, deleting a trivial block element from some block, or deleting of a trivial block for $i \in [r]$.

To prove Lemma 2.20, we need to show that whenever an equivalence transformation follows the deletion of a trivial block element or trivial block, we can change the sequence locally such that the equivalence transformation is performed first, while the block configurations before and after both steps are unchanged. Furthermore, we then show that if trivial block element deletions occur after trivial block deletions, the trivial block element deletion can be performed before deleting a trivial block.

Thus, let Π_i be the result of an equivalence transformation (φ, τ) of Π_{i-1}, and Π_{i-1} itself be the block configuration that is obtained after performing a deletion of a trivial block element from Π_{i-2}. Without loss of generality, let $\Xi_1, \dots, \Xi_k$ be the blocks of Π_{i-2}, and possibly after some renumbering, let $\omega \in \Xi_1$ be the trivial block element that is deleted from Ξ_1. Then we know that

$$\Pi_{i-1} = (\Xi_1 \setminus \omega, \Xi_2, \dots, \Xi_k)$$

and

$$\Pi_i = (\varphi(\Xi_1 \setminus \omega) \oplus \tau(1), \varphi(\Xi_2) \oplus \tau(2), \dots, \varphi(\Xi_k) \oplus \tau(k)).$$

Now exchanging the deletion of $\omega \in \Xi_1$ and the equivalence transformation (φ, τ) yields a modified subsequence of block configurations, namely

$$\widetilde{\Pi}_{i-1} = (\varphi(\Xi_1) \oplus \tau(1), \varphi(\Xi_2) \oplus \tau(2), \dots, \varphi(\Xi_k) \oplus \tau(k))$$

and

$$\widetilde{\Pi}_i = ((\varphi(\Xi_1) \oplus \tau(1)) \setminus (\varphi(\omega) \oplus \tau(1)), \varphi(\Xi_2) \oplus \tau(2), \dots, \varphi(\Xi_k) \oplus \tau(k)),$$

where we now deleted the modified trivial block element $(\varphi(\omega) \oplus \tau(1))$. This element is trivial in $\widetilde{\Pi}_{i-1}$ because of Proposition 2.12. We see that $\widetilde{\Pi}_i = \Pi_i$ holds because

$$\varphi(\Xi_1 \setminus \omega) \oplus \tau(1) = (\varphi(\Xi_1) \oplus \tau(1)) \setminus (\varphi(\omega) \oplus \tau(1))$$

is true for any equivalence transformation (φ, τ). Therefore, an equivalence transformation can be performed before a trivial block element deletion while arriving at the same block configuration.

For the second part, take the setup as before, but now let $\Xi_k = \{\mathbb{0}\}$ be the trivial block that is deleted from Π_{i-2} to obtain Π_{i-1}. Then

$$\Pi_{i-1} = (\Xi_1, \dots, \Xi_{k-1}) \quad \text{and} \quad \Pi_i = (\varphi(\Xi_1) \oplus \tau(1), \dots, \varphi(\Xi_{k-1}) \oplus \tau(k-1)).$$

Since $\operatorname{span}\left(\bigcup_{i \in [k]} \Xi_i\right) = \operatorname{span}\left(\bigcup_{i \in [k-1]} \Xi_i\right)$, the pair $(\varphi, \tilde{\tau})$ with $\tilde{\tau}(i) = \tau(i)$ for $i \in [k-1]$ and $\tilde{\tau}(k) = \mathbb{0}$ is an equivalence relation on Π_{i-2}. With this we obtain the modified block configuration

$$\widetilde{\Pi}_{i-1} = (\varphi(\Xi_1) \oplus \tilde{\tau}(1), \varphi(\Xi_2) \oplus \tilde{\tau}(2), \dots, \varphi(\Xi_k) \oplus \tilde{\tau}(k)).$$

Now $\varphi(\Xi_k) \oplus \tilde{\tau}(k) = \varphi(\{\mathbb{0}\}) \oplus \mathbb{0} = \{\mathbb{0}\}$ since φ is a linear map. Therefore, $\varphi(\Xi_k) \oplus \tilde{\tau}(k)$ is still a trivial block that can be deleted to obtain

$$\widetilde{\Pi}_i = (\varphi(\Xi_1) \oplus \tilde{\tau}(1), \ldots, \varphi(\Xi_{k-1}) \oplus \tilde{\tau}(k-1)).$$

Since $\tilde{\tau}(i)$ is equal to $\tau(i)$ for $i \in [k-1]$, we again get the equation $\widetilde{\Pi}_i = \Pi_i$, but now the (modified) equivalence relation comes before the trivial block deletion.

Lastly, if a trivial block deletion step occurs in the sequence of reduction steps directly before the deletion of some trivial block element, it is obvious that the trivial block element does not belong to the (deleted) trivial block. Let Π_{i-2} consist again of the blocks $\Xi_1, \ldots, \Xi_k$. Possibly after some renumbering of blocks, let $\Xi_k = \{\mathbb{0}\}$ be the trivial block that is deleted first to obtain $\Pi_{i-1} = (\Xi_1, \ldots, \Xi_{k-1})$, and let $\omega \in \Xi_1$ be the trivial block element that is deleted to get $\Pi_i = (\Xi_1 \setminus \omega, \Xi_2, \ldots, \Xi_{k-1})$. Then it is clear that $\omega \in \Xi_1$ can be deleted first to get $\widetilde{\Pi}_{i-1} = (\Xi_1 \setminus \omega, \Xi_2, \ldots, \Xi_k)$, then Ξ_k can be deleted afterwards to arrive at $\widetilde{\Pi}_i = (\Xi_1 \setminus \omega, \Xi_2, \ldots, \Xi_{k-1}) = \Pi_i$.

Since the number of reduction operations that are not in the prescribed order is finite and performing these local changes on the sequence of reductions each reduce the number of wrongly ordered steps by one while not changing $\Pi_0 = \Pi$ and $\Pi_r = \widetilde{\Pi}$, we conclude that $\widetilde{\Pi}$ can also be obtained from Π in the desired prescribed order, with the additional final observation that the composition of multiple equivalence transformations after one another can be expressed as a single equivalence transformation. ■

Let Π be a block configuration and $\widetilde{\Pi}$ be any reduction of Π. By Lemma 2.20, there are canonical injections of blocks and block elements of $\widetilde{\Pi}$ to blocks and block elements of Π by tracing these elements along every step of the reduction.

Under Definition 2.17, equivalent block configurations are reductions of one another. Any deletion of trivial block elements or trivial blocks are operations to obtain a reduced configuration that is not equivalent to the original one, but has fewer trivial block elements and fixed blocks as mentioned in the discussion before. By successively performing these deletions and further equivalence transformations, one obtains what we call a *pruned* block configuration:

2.21 Definition (Pruned block configurations)
A block configuration Π with at least two blocks is called *pruned* if it does not contain a fixed block and no block contains a trivial element.

Pruned block configurations contain only blocks of cardinality at least 2, since we can further reduce configurations with trivial blocks, possibly after performing an equivalence transformation. This process either leads to a pruned block configuration or one with less than two blocks. To this end, note that any pruned block configuration Π has at least two unfixed blocks, which implies that $|\mathrm{CT}(\Pi)| \geq 2$ holds. Consequently, any block configuration with fewer than two cyclic transversals cannot be pruned. Nevertheless, it is a straightforward exercise to show that any block configuration Π satisfying $|\mathrm{CT}(\Pi)| = 1$ has a reduction that is equivalent to $(\{\mathbb{0}\})$, and any block configuration with empty set of cyclic transversals has a reduction that is equivalent to $(\{\mathbb{1}\})$, both in the venue space $\mathbb{F}_2$. Furthermore, it is evident that these configurations are the smallest with respect to rank, order and size among all such reductions. We note here that recognizing whether $\mathrm{CT}(\Pi)$ is empty or not is actually a hard problem, as we will show in Lemma 3.3.

Pruned block configurations of a block configuration are *most reduced* among all of its reductions, as they contain no elements or blocks that can be deleted to obtain a block configuration that is even further reduced. The following proposition is then rather straightforward, but we want to utilize our knowledge from the proof of Lemma 2.20 to give a constructive proof here:

2.22 Proposition

For every block configuration Π with at least two unfixed blocks, there exists a reduction of Π that is pruned.

Proof: Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be such that $\Xi_1, \dots, \Xi_k$ are the unfixed blocks and $\omega_{k+1}, \dots, \omega_n$ are the fixed elements of the blocks $\Xi_{k+1}, \dots, \Xi_n$. Then Π is equivalent to

$$\Pi_1 := \left(\Xi_1 \oplus \sum_{i \in [n-k]} \omega_{k+i}, \Xi_2, \dots, \Xi_k, \Xi_{k+1} \oplus \omega_{k+1}, \dots, \Xi_n \oplus \omega_n \right)$$

by using the block translation

$$\left(\sum_{i \in [n-k]} \omega_{k+i}, \underbrace{\mathbb{0}, \dots, \mathbb{0}}_{k-1 \text{ times}}, \omega_{k+1}, \dots, \omega_n \right).$$

By performing this equivalence transformation, the fixed elements of fixed blocks of Π_1 are all equal to $\mathbb{0}$. Furthermore, the fixed blocks of Π_1 are in bijection with the fixed blocks of Π by Proposition 2.14, and there is another bijection between the trivial block elements of every block of Π_1 and those of Π by Proposition 2.12.

Now, deleting all trivial block elements of Π_1 yields the block configuration

$$\Pi_2 := \left(\widetilde{\Xi}_1 \oplus \sum_{i \in [n-k]} \omega_{k+i}, \widetilde{\Xi}_2, \dots, \widetilde{\Xi}_k, \underbrace{\{\mathbb{0}\}, \dots, \{\mathbb{0}\}}_{n-k \text{ times}} \right),$$

where

$$\widetilde{\Xi}_i = \{\omega \in \Xi_i \mid \exists \xi \in \mathrm{CT}(\Pi_1) \text{ with } \xi(i) = \omega\}$$

for $i \in [n]$. We already know that $\widetilde{\Xi}$ equals $\{\mathbb{0}\}$ for any fixed block of Π_1, by the remark after Definition 2.13.

Since the fixed blocks are now transformed into trivial blocks of the form $\{\mathbb{0}\}$, we are also able to remove them and obtain

$$\Pi_3 := \left(\widetilde{\Xi}_1 \oplus \sum_{i\in[n-k]} \omega_{k+i}, \widetilde{\Xi}_2, \dots, \widetilde{\Xi}_k\right),$$

which is a pruned block configuration, as it does not contain any trivial block elements or fixed blocks, which is guaranteed by the aforementioned bijections and the property that deletion of trivial block elements and trivial blocks cannot introduce new fixed blocks or trivial block elements, which is implied by Proposition 2.11 and Proposition 2.19. ■

The proof of Proposition 2.22 offers a procedure to directly obtain a pruned block configuration from any given block configuration that has at least two unfixed blocks, which is the case if there are at least two cyclic transversals. But because we do not want to consider a particular order of reduction operations to obtain a pruned block configuration, it is desirable to show that performing different sequences of reductions lead to equivalent pruned block configurations:

2.23 Theorem
All pruned block configurations that are reductions of a block configuration Π are equivalent.

Proof: For any pruned block configurations that is a reduction of a given block configuration Π, we are allowed to prescribe the order of reduction steps by Lemma 2.20. Therefore, any pruned block configuration is given by performing an equivalence transformation on Π before deleting trivial block elements and trivial blocks.

Let $\widetilde{\Pi}$ and $\widehat{\Pi}$ be two pruned block configurations that are obtained by reducing the block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$, and let $\Xi_{k+1}, \dots, \Xi_n$ be the fixed blocks of Π, with fixed elements ω_{k+i} for $i \in [n-k]$. Then, possibly after some renumbering, let (φ, τ) be an equivalence transformation on Π such that $\widetilde{\Pi} = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_k)$ is obtained from the equivalent block configuration by just deleting trivial block elements and trivial blocks, and let (ψ, σ) be an equivalence transformation on Π such that the same is true for $\widehat{\Pi} = \left(\widehat{\Xi}_1, \dots, \widehat{\Xi}_k\right)$.

Given the composition rule of equivalence transformations, the composition of the inverse equivalence transformation $\left(\varphi^{-1}, \varphi^{-1} \circ \tau\right)$ and (ψ, σ) is an equivalence transformation itself. Let (θ, η) denote this composition. Furthermore, a pruned block configuration does not contain trivial block elements or trivial blocks, and equivalence transformations preserve the property that an element is trivial by Proposition 2.12 and the property that a block is fixed by Proposition 2.14.

By Proposition 2.11, we know that the set $\{\omega \in \Xi_i \mid \exists \xi \in \mathrm{CT}(\Pi) \text{ with } \xi(i) = \omega\}$ contains exactly the non-trivial elements of Ξ_i. Then it is clear that

$$\widetilde{\Xi}_i = \{\varphi(\omega) \oplus \tau(i) \mid \omega \in \Xi_i \text{ and } \exists \xi \in \mathrm{CT}(\Pi) \text{ with } \xi(i) = \omega\}$$

and

$$\widehat{\Xi}_i = \{\psi(\omega) \oplus \sigma(i) \mid \omega \in \Xi_i \text{ and } \exists \xi \in \mathrm{CT}(\Pi) \text{ with } \xi(i) = \omega\}$$

for $i \in [k]$. One easily sees that $\theta(\widetilde{\Xi}_i) \oplus \eta(i) = \widehat{\Xi}_i$ for all $i \in [k]$.

Therefore, what is left to show is that if we restrict (θ, η) to the unfixed blocks, it is an equivalence relation between $\widetilde{\Xi}_1, \ldots, \widetilde{\Xi}_k$ and $\widehat{\Xi}_1, \ldots, \widehat{\Xi}_k$. Since $\theta = \psi \circ \varphi^{-1}$ is clearly still a linear isomorphism when restricted to $\operatorname{span}\left(\bigcup_{i\in[k]} \widetilde{\Xi}_i\right)$, we only have to show that $\eta|_{[k]}$ is a cyclic map, i. e., we will show $\sum_{i\in[k]} \eta(i) = \mathbb{0}$.

Because the pruned block configurations do not contain any fixed blocks, they had to be deleted during the reduction process. To do that, all trivial elements had to have been deleted from them before, such that they are of the form $\{\mathbb{0}\}$ afterwards. This means that the fixed element of all fixed blocks had to be transformed into $\mathbb{0}$ by the respective equivalence transformations. This means that

$$\tau(k+i) = \varphi(\omega_{k+i}) \text{ and } \sigma(k+i) = \psi(\omega_{k+i}) \text{ for all } i \in [n-k]$$

holds. Since $\sum_{i\in[n]} \tau(i) = \sum_{i\in[n]} \sigma(i) = \mathbb{0}$ is true, we have the equations

$$\sum_{i\in[k]} \tau(i) = \sum_{i\in[n-k]} \varphi(\omega_{k+i}),$$

as well as

$$\sum_{i\in[k]} \sigma(i) = \sum_{i\in[n-k]} \psi(\omega_{k+i}).$$

By definition of the composition rule, η is equal to $\psi \circ (\varphi^{-1} \circ \tau) \oplus \sigma$. We evaluate the sum

$$\begin{aligned}
\sum_{i\in[k]} \eta(i) &= \sum_{i\in[k]} \psi(\varphi^{-1}(\tau(i))) \oplus \sigma(i) \\
&= \psi\left(\varphi^{-1}\left(\sum_{i\in[k]} \tau(i)\right)\right) \oplus \sum_{i\in[k]} \sigma(i) \\
&= \psi\left(\varphi^{-1}\left(\sum_{i\in[n-k]} \varphi(\omega_{k+i})\right)\right) \oplus \sum_{i\in[n-k]} \psi(\omega_{k+i}) \\
&= \sum_{i\in[n-k]} \psi(\omega_{k+i}) \oplus \sum_{i\in[n-k]} \psi(\omega_{k+i}) \\
&= \mathbb{0},
\end{aligned}$$

which shows that $\eta|_{[k]}$ is cyclic, since φ and ψ are linear. This means that $\widetilde{\Pi}$ and $\widehat{\Pi}$ are equivalent with the constructed equivalence transformation $\left(\theta, \eta|_{[k]}\right)$. ∎

By Proposition 2.22 and Theorem 2.23, it is sensible to speak of a *pruning* of a given block configuration, i. e., a reduction of a block configuration that is pruned, which is unique up to equivalence.

We now know that all pruned block configurations originating from the same block configuration are equivalent, and that equivalent block configurations result in isomorphic cyclic transversal polytopes. Next, we prove that reducing a block configuration also leads to an isomorphic cyclic transversal polytope, but with an additional desirable property if the block configuration is pruned, which is later shown in Proposition 2.26.

2.24 Proposition
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ *be a block configuration with corresponding cyclic transversal polytope* $P = \mathrm{CTP}(\Pi)$*. Then* $P \cong \mathrm{CTP}(\widetilde{\Pi})$ *holds whenever* $\widetilde{\Pi}$ *is a reduction of* Π*.*

Proof: We will show that for any of the three reduction operations, there is an affine isomorphism between P and $\mathrm{CTP}(\widetilde{\Pi})$ that is easy to describe.

First, if $\widetilde{\Pi}$ is equivalent to Π, then the affine isomorphy follows from Proposition 2.7, since $\mathrm{CTP}(\widetilde{\Pi})$ is given by applying a certain coordinate permutation of P.

If $\widetilde{\Pi}$ is given by the deletion of a trivial block element $\omega \in \Xi_i$ for some $i \in [n]$ from Π, then P is contained in the hyperplane of $\mathbb{A}(\Pi)$ given by $y^i_\omega = 0$. Indeed, since no cyclic transversal $\xi \in \mathrm{CT}(\Pi)$ of Π fulfills $\xi(i) = \omega$, no characteristic vector of the cyclic transversals of Π and therefore no vertex of P fulfill the equation $y^i_\omega = 1$. Since P is the convex hull of its vertices, it lies in the aforementioned hyperplane. The projection onto all other coordinates then is an affine map that is a bijection between P and $\mathrm{CTP}(\widetilde{\Pi})$.

If $\widetilde{\Pi}$ is given by the deletion of a trivial block $\Xi_i = \{\mathbb{0}\}$ for some $i \in [n]$ from Π, then P is contained in the hyperplane given by $y^i_{\mathbb{0}} = 1$. As a matter of fact, P is contained in $\mathbb{A}(\Pi)$ and $\sum_{\omega \in \Xi_i} y^i_\omega = 1$ for all $i \in [n]$ are valid equations for P. If Ξ_i is equal to $\{\mathbb{0}\}$ for some $i \in [n]$, then the equation simplifies to $y^i_{\mathbb{0}} = 1$. Removing this coordinate results in a projection onto $\mathbb{A}(\widetilde{\Pi})$ since all other block equations remain unchanged. This projection again is an affine map and a bijection between P and $\mathrm{CTP}(\widetilde{\Pi})$.

If $\widetilde{\Pi}$ on the other hand is obtained by performing a sequence of multiple such reduction steps on Π, the affine isomorphism between P and $\mathrm{CTP}(\widetilde{\Pi})$ then is the concatenation of the aforementioned simple affine isomorphisms according to the order of a sequence of reduction steps to reduce Π to $\widetilde{\Pi}$. ∎

2.25 Corollary
Let Π *be a block configuration with corresponding cyclic transversal polytope* $P = \mathrm{CTP}(\Pi)$*. Then* $P \cong \mathrm{CTP}(\widetilde{\Pi})$ *holds, where* $\widetilde{\Pi}$ *is any pruning of* Π*.*

Proof: Since a pruning of Π is also a reduction of Π by definition, the result follows immediately from Proposition 2.24. ∎

Now we come to a desirable property of cyclic transversal polytopes that are obtained by pruned block configurations: Given a family of blocks $\Xi_1, \dots, \Xi_n$ and their non-empty cyclic transversal polytope P, being contained in a proper face of $[0,1]^d$ is equivalent to an equation of the form $y^i_\omega = 0$ or $y^i_\omega = 1$ being valid for P, for some block Ξ_i with $i \in [n]$ and some element $\omega \in \Xi_i$.

2.26 Proposition
A non-empty cyclic transversal polytope in $[0,1]^d \subseteq \mathbb{R}^d$ generated by a pruned block configuration is not contained in a proper face of $[0,1]^d$.

Proof: Let $P \subseteq [0,1]^d$ be a cyclic transversal polytope generated by some pruned block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$. Suppose, for the sake of contradiction, that P is contained in a proper face of $[0,1]^d$, i. e., an equation $y^i_\omega = 0$ or $y^i_\omega = 1$ is valid for P, for some block Ξ_i with $i \in [n]$ and some element $\omega \in \Xi_i$.

Case 1: All $y \in P$ fulfill $y^i_\omega = 0$. Then in particular all vertices of P fulfill this equation. This means that the element $\omega \in \Xi_i$ is not part of any cyclic transversal in $\mathrm{CT}(\Pi)$, since any cyclic transversal ξ with $\xi(i) = \omega$ would result in a vertex v of P with $v^i_\omega = 1$.

Furthermore, we know that $|\Xi_i| \geq 2$, and since Ξ_i implies a block equation $\sum_{\sigma \in \Xi_i} y^i_\sigma = 1$, there has to be at least one other element $\tilde{\omega} \in \Xi_i$ such that the block equation is satisfied. Therefore, we may delete $\omega \in \Xi_i$ from Π to obtain another block configuration $\widetilde{\Pi}$, and since ω is not in any cyclic transversal, we have that $\mathrm{CT}(\Pi) = \mathrm{CT}(\widetilde{\Pi})$, i. e., ω is a trivial block element. This is a contradiction to the assertion that Π was pruned.

Case 2: All $y \in P$ fulfill $y^i_\omega = 1$. We then show that Ξ_i is a trivial block, which again is a contradiction to Π being pruned.

First, if $|\Xi_i| \geq 2$, then we switch to the first case, since $y^i_\omega = 1$ together with the implied block equation $\sum_{\sigma \in \Xi_i} y^i_\sigma = 1$ show that $y^i_{\tilde{\omega}} = 0$ for all $\tilde{\omega} \in \Xi_i \setminus \omega$ are also valid equations for P. Therefore, it follows that $|\Xi_i| = 1$, i. e., $\Xi_i = \{\omega\}$, so Ξ_i is trivial, which is the desired contradiction. ∎

While Proposition 2.26 could be extracted as a corollary from the proof of Proposition 2.24, its own proof emphasizes that superfluous coordinates imply redundant elements of corresponding block configurations, while the proof of Proposition 2.24 emphasizes that unnecessary block elements or blocks imply fixed coordinates of cyclic transversal polytopes.

In the same vein as equivalences and reductions, there are operations on block configurations that leave the rank, order and size unchanged, but applying them may generate a set of cyclic transversals that has different properties (like cardinality or fixed block elements) than the one before the operation was applied. An example of one such operation is exchanging a vector in a single block, and another example would be deleting a non-trivial block element and simultaneously enlarging another block by adding to it a new vector it did not previously contain. The behavior of cyclic transversals and their polytopes under these transformations is poorly understood thus far.

Nevertheless, we now have some tools at hand to categorize and distinguish different block configurations and cyclic transversal sets. The focus of this framework will be on the resulting cyclic transversal polytopes in the coming chapter, and we will see some non-equivalent (pruned) block configurations that are not reductions of each other, but still generate affinely isomorphic cyclic transversal polytopes. Therefore, one may also ask the reverse question that underlies this section:

2.27 Question
Can one characterize whether two cyclic transversal polytopes are (affinely) isomorphic by only using block configurations and their properties?

This question cannot yet be answered to a satisfactory degree, and we suggest that further research is required into the structure of cyclic transversals to give a meaningful answer.

To circumvent this shortcoming, we present the definition of a normalized block configuration and possible directions to enumerate cyclic transversal polytopes in the next section. This also leads to so-called signatures of block configurations to see more easily whether two block configurations are not equivalent.

2.2.3 Normalization and Parameter Bounds

If we assume that the set of cyclic transversals of an unspecified block configuration contains at least one element, there may not be a canonical choice of such a cyclic transversal without some kind of normalization of the blocks. This means that the next step towards normalized block configurations is the following definition:

2.28 Definition (Centralization)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a block configuration. If $\mathbb{0} \in \Xi_i$ for all $i \in [n]$, then we call Π a *centralized* block configuration. If Π is centralized, we call its set of cyclic transversals centralized as well.

If we know any cyclic transversal $\xi \in \mathrm{CT}(\Pi)$, the block translation with the map ξ transforms Π into a centralized block configuration. We then call the resulting block configuration *centralized on* ξ. Furthermore, if $\mathbb{0} \in \Xi_i$ for all $i \in [n]$, then the map $\xi\colon [n] \to \mathbb{F}_2^d$ defined by $\xi(i) = \mathbb{0}$ for all $i \in [n]$ will always be a cyclic transversal of Π. This means that block configurations Π such that $\mathrm{CT}(\Pi) \neq \emptyset$ are exactly those that can be transformed into centralized block configurations.

While assuming a block configuration is centralized, one may always assume that a single fixed vertex of a cyclic transversal polytope is the vertex corresponding to the cyclic transversal $(\mathbb{0}, \dots, \mathbb{0})$. We again refer to Lemma 3.3, which shows that checking emptiness of a set of cyclic transversals and cyclic transversal polytopes is not an easy problem. Still, we gather here some (necessary) conditions for $\mathrm{CT}(\Pi) \neq \emptyset$ to hold. Furthermore, we collect some bounds on the parameters of block configurations to bring ourselves closer to the goal of normalizing the presentation of block configurations and cyclic transversals with the future possibility of enumerating affinely non-isomorphic cyclic transversal polytopes.

Usually we even require at least two cyclic transversals to exist, otherwise the resulting polytopes consisting of a single vertex appear to be rather uninteresting from a combinatorial perspective. Vertex adjacency in Section 2.2.4, for example, becomes meaningless. Like before, this also implies that block configurations consist of at least two blocks.

Not every given block configuration is equivalent to or can be reduced to a block configuration that fulfills these assumptions. Regardless, these assumptions are sensible to include to investigate combinatorial properties of cyclic transversals, as empty or particularly small sets of cyclic transversals are not the main focus of this framework.

2.29 Definition (Normalization)
A block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$ on $\mathbb{F}_2^d$ is called *normalized* if it is centralized, pruned and its rank is equal to d.

If a block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$ is normalized, the definition readily implies that $n \geq 2$ and $|\Xi_i| \geq 2$ holds for all $i \in [n]$. We also know that $\mathrm{CT}(\Pi) \geq 2$ is true for any normalized block configuration, since it already has to hold for any pruned block configuration. We purposefully do not prescribe a particular ordering of the blocks in a normalized block configuration, since unless otherwise stated, the results in this thesis are not dependent on any particular ordering of the blocks.

Without loss of generality, we assume that the linear span of $\bigcup_{i\in[n]} \Xi_i$ for any block configuration is the whole venue space $\mathbb{F}_2^d$ by applying an appropriate equivalence transformation, since embedding blocks in a higher-dimensional space without changing the blocks does not alter the cyclic transversals or the cyclic transversal polytope in any meaningful way. Therefore, from this point on, we simply assume that the rank of any block configuration, not only normalized ones, is equal to the dimension d of its underlying venue space.

Assuming that the rank of a pruned block configuration is equal to d in turn means that

$$d \leq \sum_{i\in[n]} |\Xi_i| - \max\{|\Xi_1|, \dots, |\Xi_n|\} = s - \max\{|\Xi_1|, \dots, |\Xi_n|\},$$

since we can project the vectors of the blocks to an appropriate venue space of dimension $s - \max\{|\Xi_1|, \dots, |\Xi_n|\}$ via a linear map. This is because for pruned block configurations, we have

$$\Xi_k \subseteq \operatorname{span}\left(\bigcup_{i\in[n]\setminus k} \Xi_i\right) \text{ for all } k \in [n],$$

since any block elements of $\Xi_k \setminus \operatorname{span}\left(\bigcup_{i\in[n]\setminus k} \Xi_i\right)$ necessarily would be trivial elements. For general block configurations that are not necessarily pruned, the weaker condition

$$\Xi_k \cap \operatorname{span}\left(\bigcup_{i\in[n]\setminus k} \Xi_i\right) \neq \emptyset \text{ for all } k \in [n]$$

is necessary for $\mathrm{CT}(\Xi_1, \dots, \Xi_n) \neq \emptyset$.

Pruning a block configuration essentially does not change the set of cyclic transversals itself as shown in Proposition 2.19, but shrinks the size parameter of the block configuration. We therefore may assume that given a set of cyclic transversals, the size parameter of its block configuration is minimal, so that every element of every block is part of some cyclic transversal. Minimizing the size of block configurations leads to an interesting research question regarding the resulting dimension of cyclic transversal polytopes, which is beyond the scope of this work:

2.30 Question
Let $s(\delta)$ be the smallest number for which the following holds: For every cyclic transversal polytope of dimension less or equal than δ, there is an affinely isomorphic cyclic transversal polytope of size s less or equal than $s(\delta)$. What are non-trivial bounds on $s(\delta)$?

It is clear that $s(\delta)$ is finite for all $\delta \in \mathbb{N}$, since there are only finitely many 0/1-polytopes of dimension at most δ. Among these, the finitely many cyclic transversal polytopes (or rather, their finitely many equivalence classes with respect to affine isomorphisms) each have a construction using some block configuration that is the smallest possible with respect to size, and the maximum of these finitely many minimal sizes is exactly $s(\delta)$ by definition. Once we know at least one construction for every cyclic transversal polytope of dimension at most δ, the maximum of the sizes of the involved block configurations is an upper bound on $s(\delta)$. For example, we will see that $s(3) \leq 12$, and this upper bound is attained by the construction of the 3-cube $[0,1]^3$ as a cyclic transversal polytope in Corollary 3.19. For further discussion on dimension 3, see also Section 3.2.2.

Structurally, we can and will now refine the size and rank parameters of block configurations using so-called *signatures*. These are helpful for the enumeration of all possible block configurations.

2.31 Definition (Size Signature)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a block configuration, then the sequence $(s_1, \dots, s_n)$ with $s_i = |\Xi_i|$ for $i \in [n]$ is called *size signature* of Π.

The sum of the terms of the size signature is the size of the block configuration by definition. It is rather evident that reduction operations of block configurations either decrease one term of the size signature by one or remove an element $s_i = 1$ from the sequence. Furthermore, since equivalence relations do not change the size signature, non-identical size signatures provide a certificate for two block configurations to be not equivalent.

If Π is a pruned block configuration on $\mathbb{F}_2^d$, the bounds $2 \leq s_i \leq 2^d = \left|\mathbb{F}_2^d\right|$ are trivially true by observing that unfixed blocks contain at least two elements and Π consists only of unfixed blocks that are subsets of $\mathbb{F}_2^d$. Then, the inequalities

$$\frac{s}{2^d} = \sum_{i \in [n]} \frac{s_i}{2^d} \leq n \leq \sum_{i \in [n]} \frac{s_i}{2} = \frac{s}{2}$$

obviously hold.

With size signatures, we can easily prove some existence and uniqueness results, like in the following example:

2.32 Example
Up to equivalence, there is exactly one pruned block configuration $\Pi = (\Xi_1, \dots, \Xi_3)$ with size signature $(4, 2, 2)$. Indeed, let $\Xi_2 = \{\omega_1, \omega_2\}$ and $\Xi_3 = \{\sigma_1, \sigma_2\}$ be the two blocks of cardinality 2, then the set $\{\omega_1 \oplus \sigma_1, \omega_1 \oplus \sigma_2, \omega_2 \oplus \sigma_1, \omega_2 \oplus \sigma_2\}$ clearly contains all elements of Ξ_1. Since the assumed size signature implies $|\Xi_1| = 4$, we know that these four sums have to be distinct, i. e.,

$$\Xi_1 = \{\omega_1 \oplus \sigma_1, \omega_1 \oplus \sigma_2, \omega_2 \oplus \sigma_1, \omega_2 \oplus \sigma_2\},$$

as any other possibility would be a contradiction to the assumption that $\Pi = (\Xi_1, \dots, \Xi_3)$ is pruned.

Without loss of generality, we can apply a block translation using the cyclic map identified via the tuple $(\omega_2 \oplus \sigma_2, \omega_2, \sigma_2)$ to transform these blocks into

$$\Xi_1 = \{\mathbb{0}, \omega, \sigma, \omega \oplus \sigma\}, \quad \Xi_2 = \{\mathbb{0}, \omega\} \quad \text{and} \quad \Xi_3 = \{\mathbb{0}, \sigma\},$$

where $\omega = \omega_1$ and $\sigma = \sigma_1$. Furthermore, $|\Xi_1| = 4$ then implies that ω and σ are linearly independent, again because the block configuration is pruned. This means that $\dim\operatorname{span}(\Xi_2 \cup \Xi_3) = \dim\operatorname{span}(\Xi_1) = 2$ with $\{\omega, \sigma\}$ as a common basis.

Now, any pruned block configuration $\widetilde{\Pi} = (\widetilde{\Xi}_1, \dots, \widetilde{\Xi}_3)$ with size signature $(4, 2, 2)$ can clearly be transformed analogously, such that $\widetilde{\Xi}_1 = \{\mathbb{0}, \tilde{\omega}, \tilde{\sigma}, \tilde{\omega} \oplus \tilde{\sigma}\}$, $\widetilde{\Xi}_2 = \{\mathbb{0}, \tilde{\omega}\}$ and $\widetilde{\Xi}_3 = \{\mathbb{0}, \tilde{\sigma}\}$. Then, since $\tilde{\omega}$ and $\tilde{\sigma}$ are also linearly independent and form a basis of $\operatorname{span}\left(\bigcup_{i\in[3]} \widetilde{\Xi}_i\right)$, there is a unique block isomorphism φ defined by $\varphi(\omega) = \tilde{\omega}$ and $\varphi(\sigma) = \tilde{\sigma}$. This implies that Π and $\widetilde{\Pi}$ are equivalent, finishing this example.

Additional information can be gathered from another type of signature:

2.33 Definition (Dimensional Signature)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a block configuration, then the sequence $(d_1, \dots, d_n)$ with

$$d_1 = \dim\operatorname{span}(\Xi_1) \text{ and } d_k = \dim\operatorname{span}\left(\bigcup_{i\in[k]} \Xi_i\right) - \dim\operatorname{span}\left(\bigcup_{i\in[k-1]} \Xi_i\right) \text{ for } k \in [n] \setminus 1$$

is called *dimensional signature* of Π.

It is clear that the dimensional signature $(d_1, \dots, d_n)$ of a block configuration Π fulfills

$$\sum_{i\in[n]} d_i = \dim\operatorname{span}\left(\bigcup_{i\in[n]} \Xi_i\right)$$

by evaluation of its telescope sum property, and is dependent on the ordering of blocks.

We can categorize block configurations by their dimensional signature, which describes how many new dimensions each block adds to the union of the previous blocks. The notion of dimensional and size signatures will help with computationally enumerating essentially different sets of cyclic transversals by listing their block configurations in a structured fashion less prone to redundancy than simply generating all families of venue space subsets.

2.34 Example
The configuration $(\{\mathbb{0}, \mathbb{e}_1\}, \{\mathbb{0}, \mathbb{e}_2\}, \{\mathbb{0}, \mathbb{1}\})$ over $\mathbb{F}_2^2$ has dimensional signature $(1, 1, 0)$, and the block configurations from Example 2.32 have dimensional signature $(2, 0, 0)$, using the prescribed ordering of the blocks.

If Π is pruned, then this already implies $d_n = 0$ regardless of the ordering of blocks, but $d_n = 0$ on its own is not a sufficient condition for $\mathrm{CT}(\Pi) \neq \emptyset$. Indeed for general block configurations, $d_n = 0$ is equivalent to the last block being contained in the linear span of the other blocks in the particular ordering given by the block indices. Note that this does not have to hold for all orderings of the same blocks. Therefore, if Π is not pruned, there is no relation between d_n and whether $\mathrm{CT}(\Pi)$ is empty or not, i. e., $d_n = 0$ may hold even when $\mathrm{CT}(\Pi)$ is empty, for example for $\Pi = (\Xi_1, \Xi_2)$ with $\Xi_1 = \{\mathbb{1}\}$ and $\Xi_2 = \{\mathbb{0}\}$, and $d_n \neq 0$ may be true for general (non-pruned) block configurations that have cyclic transversals, when the blocks are ordered in a particular way. The block configuration $\Pi = (\Xi_1, \Xi_2)$ with $\Xi_1 = \{\mathbb{e}_1\}$ and $\Xi_2 = \{\mathbb{e}_1, \mathbb{e}_2\}$ is such an example.

Using the signatures, we will now prove a bound on the parameter space of pruned block configurations with non-empty set of cyclic transversals:

2.35 Proposition
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a normalized block configuration. Then for its rank d, order n and size s, we have the inequality

$$d \leq s - n - 1,$$

and this bound is sharp, i. e., there are block configurations that attain this relationship of parameters with equality.

Proof: Let $(s_1, \dots, s_n)$ and $(d_1, \dots, d_n)$ be the size and dimensional signature of Π, respectively. Since Π is centralized, every block is of the form $\Xi_i = U_i \cup \mathbb{0}$ for some set $U_i \subseteq \mathbb{F}_2^d \setminus \mathbb{0}$. We first show that $d_i \leq s_i - 1$ for all $i \in [n]$ holds, which implies the weaker inequality $d \leq s - n$ by summing over $i \in [n]$ on both sides.

Since $d_i \leq \dim \operatorname{span}(\Xi_i) = \dim \operatorname{span}(U_i)$ and $s_i = |U_i| + 1$ are true by definition of the signatures, this translates to the inequality

$$\dim \operatorname{span}(U_i) \leq |U_i|$$

for all $i \in [n]$. This is a well-known result from linear algebra: If U_i is linearly dependent, it contains a linearly independent subset that generates the same subspace as U_i. Then if U_i is a linearly independent set, the inequality is attained with equality by the definition of dimension.

Since Π is normalized, the equation $d_n = 0 \leq s_n - 2$ is true, which diminishes the right-hand side of $d \leq s - n$ by one and therefore proves the initial inequality.

To prove that the inequality is sharp, observe the block configuration $\Xi_0 = \{\mathbb{0}, \mathbb{1}\}$ and $\Xi_i = \{\mathbb{0}, \mathbb{e}_i\}$ for $i \in [d]$ on $\mathbb{F}_2^d$, which obviously has rank d. This configuration is shown to be centralized and pruned by discovering that its only two cyclic transversals are $(\mathbb{0}, \dots, \mathbb{0})$ and $(\mathbb{1}, \mathbb{e}_1, \dots, \mathbb{e}_d)$. It is therefore normalized and consists of $n = d + 1$ blocks and has size $s = 2n$. Therefore, the right-hand side of the inequality $d \leq s - n - 1$ evaluates to

$$s - n - 1 = 2n - n - 1 = n - 1 = d + 1 - 1 = d.$$

■

The bound from Proposition 2.35 proves to be useful in finding small examples of non-trivial cyclic transversal polytopes and enumerating over the appropriate parameter space. To save time in finding these examples, one could sort the blocks in descending order of dimensional signature. As a rough overview, this bound may sometimes be stronger than the other bound

$$d \leq s - \max\{|\Xi_1|, \dots, |\Xi_n|\}$$

in cases where there are many small blocks. The latter bound may be stronger in cases where there is at least one big block (compared to the number of blocks).

For any block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$, we remind the reader that

$$\dim \mathrm{CTP}(\Xi_1, \dots, \Xi_n) \leq s - n$$

is true, but we are not aware of any immediate relation between the dimension of the associated cyclic transversal polytope and the rank d of the block configuration.

We close this section with an open-ended question related to Question 2.27, for which we have displayed the groundwork and given partial answers here:

2.36 Question
How many non-equivalent sets of cyclic transversals and non-isomorphic cyclic transversal polytopes are there, given a set of prescribed parameters of block configurations?

2.2.4 Vertex adjacency

Although we do not know much about the combinatorics of cyclic transversal polytopes so far, we are able to characterize their vertex adjacency. For that let $\Xi_1, \dots, \Xi_n \subseteq \mathbb{F}_2^d$ be a family of blocks and $\mathrm{CTP}(\Xi_1, \dots, \Xi_n)$ its cyclic transversal polytope. For two cyclic transversals $\xi_1 := (\xi_1(1), \dots, \xi_1(n))$ and $\xi_2 := (\xi_2(1), \dots, \xi_2(n))$, we define the set

$$B(\xi_1, \xi_2) := \{i \in [n] \mid \xi_1(i) \neq \xi_2(i)\},$$

which is the set of block indices where the two cyclic transversals differ. With that, it is easy to state when two vertices of $\mathrm{CTP}(\Xi_1, \dots, \Xi_n)$ are adjacent:

2.37 Lemma (Adjacency of vertices of cyclic transversal polytopes)
Let $\Xi_1, \dots, \Xi_n \subseteq \mathbb{F}_2^d$ be a family of blocks and let $v, \hat{v}$ be two vertices of $\mathrm{CTP}(\Xi_1, \dots, \Xi_n)$ with associated cyclic transversals $\xi := (\xi(1), \dots, \xi(n))$ and $\hat{\xi} := (\hat{\xi}(1), \dots, \hat{\xi}(n))$. Then v is adjacent to $\hat{v}$ if and only if $B := B\left(\xi, \hat{\xi}\right)$ does not contain a proper subset $T \subsetneq B$ with which the equation

$$\sum_{i \in T} \xi(i) = \sum_{i \in T} \hat{\xi}(i) \tag{X}$$

holds.

Proof: We will show that the line segment that connects v and $\hat{v}$ is a face of $\mathrm{CTP}(\Xi_1, \dots, \Xi_n)$ if and only if there is no subset $T \subsetneq B$ as described. To do that, let us define a valid inequality for $\mathrm{CTP}(\Xi_1, \dots, \Xi_n)$:

$$\sum_{i \notin B} y^i_{\xi(i)} + \sum_{i \in B} \left(y^i_{\xi(i)} + y^i_{\hat{\xi}(i)} \right) \leq n \ . \tag{L}$$

Inequality (L) is valid since it is the sum of the inequalities $y^i_{\xi(i)} \leq 1$ for $i \notin B$ and $y^i_{\xi(i)} + y^i_{\hat{\xi}(i)} \leq 1$ for $i \in B$ respectively, which are themselves relaxations of the block equations by adding the non-negativity conditions $-y^i_\omega \leq 0$ for all $\omega \in \Xi_i \setminus \left\{ \xi(i), \hat{\xi}(i) \right\}$ to them. Note that v and $\hat{v}$ clearly attain equality for Inequality (L).

Now if there is no subset $T \subsetneq B$ as described, then v and $\hat{v}$ are the only vertices that attain equality for Inequality (L): Suppose that some vertex w attains equality for Inequality (L) and let ζ be its associated cyclic transversal. Because w is integral and

$$\sum_{i \notin B} w^i_{\xi(i)} + \sum_{i \in B} w^i_{\xi(i)} + \sum_{i \in B} w^i_{\hat{\xi}(i)} = n$$

is true, ζ satisfies $\zeta(i) \in \left\{ \xi(i), \hat{\xi}(i) \right\}$ for all $i \in B$ and $\zeta(i) = \xi(i)$ for all $i \notin B$ because of the block equations. For the indices in B, we observe that $B\left(\zeta, \hat{\xi}\right)$ is a subset of $B = B\left(\xi, \hat{\xi}\right)$ by the above equation. By the fact that both ζ and $\hat{\xi}$ are cyclic transversals, we obtain

$$\sum_{i \in B\left(\zeta, \hat{\xi}\right)} \zeta(i) = \sum_{i \in B\left(\zeta, \hat{\xi}\right)} \hat{\xi}(i) \ .$$

Because of that, $B\left(\zeta, \hat{\xi}\right)$ is not a proper subset of B, as otherwise it would be a suitable proper subset $T \subsetneq B$ which we require not to exist. The only possibilities are $B\left(\zeta, \hat{\xi}\right) = \emptyset$ and $B\left(\zeta, \hat{\xi}\right) = B$. In the first case, we get that $\zeta = \hat{\xi}$ and $w = \hat{v}$ holds, while in the second case, we get that $\zeta = \xi$ and $w = v$ is true. Therefore, any vertex that attains equality for Inequality (L) is either v or $\hat{v}$, given that no subset $T \subsetneq B$ as described exists.

On the other hand, if there is a proper subset $T \subsetneq B$ such that Equation (X) holds, then we construct two cyclic transversals that are distinct from ξ and $\hat{\xi}$ such that their associated vertices have the same midpoint as v and $\hat{v}$, which implies that all four vertices must lie in the same face (which is defined by Inequality (L)) and v is not adjacent to $\hat{v}$. The midpoint of v and $\hat{v}$ is described by $y^i_{\xi(i)} = y^i_{\hat{\xi}(i)} = 1/2$ for $i \in [n]$ and $y^i_\omega = 0$ otherwise.

Let $\xi_1 := (\xi_1(1), \dots, \xi_1(n))$ and $\xi_2 := (\xi_2(1), \dots, \xi_2(n))$ be defined via

$$\xi_1(i) := \begin{cases} \xi(i) & \text{, if } i \in T, \\ \hat{\xi}(i) & \text{, otherwise,} \end{cases} \qquad \text{and} \qquad \xi_2(i) := \begin{cases} \hat{\xi}(i) & \text{, if } i \in T, \\ \xi(i) & \text{, otherwise.} \end{cases}$$

Then obviously ξ_1 and ξ_2 are cyclic transversals because of Equation (X), but both ξ_1 and ξ_2 are distinct from ξ and $\hat{\xi}$. Since their associated vertices also attain Inequality (L) with equality and their midpoint is also described by $y^i_{\xi(i)} = y^i_{\hat{\xi}(i)} = 1/2$ for $i \in [n]$ and $y^i_\omega = 0$ otherwise, v is not adjacent to $\hat{v}$. ■

After having characterized the vertex adjacency of cyclic transversal polytopes, we can ask the following questions:

2.38 Question
Are there any non-trivial bounds on vertex degrees of a cyclic transversal polytope, given the parameters of a corresponding (possibly pruned) block configuration?

2.39 Question
For any block configuration Π, let *elementary* cyclic transversals be those elements ξ of $\mathrm{CT}(\Pi)$ such that there is no proper subset $I \subsetneq [n]$ such that $\sum_{i \in I} \xi(i) = \mathbb{0}$. What is the number of elementary cyclic transversals of a given block configuration?

We will utilize the methods developed in this chapter to obtain some properties in Section 3.2 that will certify that not all polytopes are cyclic transversal polytopes. In fact, we will be able to give simple explicit examples of polytopes, namely, cross polytopes, that are not even combinatorially isomorphic to any cyclic transversal polytope.

Before that, we focus on constructive results that certify the representability of a wide range of polytopes as special subclasses of cyclic transversal polytopes or its projections in the following chapter.

3 Cyclic Transversal Polytopes

In the last chapter, we made ourselves familiar with the concept of cyclic transversals and their polytopes, but we have not yet seen their connections to other well-known families of polytopes. This is the purpose of this chapter: We will present universality results to represent any 0/1-polytope as a projection of a cyclic transversal polytope in Section 3.1, and in the subsections we will be able to show that binary kernel and stable set polytopes as well as other 0/1-polytopes are affinely isomorphic to certain cyclic transversal polytopes.

Section 3.2 then will go in the other direction: we will prove an important necessary combinatorial condition that cyclic transversal polytopes satisfy, which allows us to prove that except for certain special cases, the family of cross polytopes is not contained in the family of cyclic transversal polytopes. Some discussion on representability of other families of polytopes, like spanning tree and traveling salesman polytopes, finish this chapter.

3.1 Constructions for CTPs

The general setup of cyclic transversal polytopes allows us to describe various classes of well-known families of polytopes in the cyclic transversal framework. One distinguishing element of these descriptions is whether a family of polytopes is a special case of cyclic transversal polytopes themselves or whether each polytope is a projection of a cyclic transversal polytope.

First, let us state that a naive construction to represent $\operatorname{conv} X$ for any $X \subseteq \{0,1\}^d$ as a projection of a cyclic transversal polytope is possible:

3.1 Proposition (Naive construction)
Let $X \subseteq \{0,1\}^d$ be a set of vectors. Then the polytope $\operatorname{conv} X$ is a projection of a cyclic transversal polytope of rank d, order $d+1$ and size $2d+|X|$.

Proof: We begin by defining blocks $\Xi_i := \{\mathbb{0}, \mathbb{e}_i\} \subseteq \mathbb{F}_2^d$ for $i \in [d]$ and another block $\Xi_0 = X$ by embedding X into $\mathbb{F}_2^d$ via the canonical identification.

Further, let $\pi\colon \operatorname{CTP}(\Xi_0, \Xi_1, \dots, \Xi_d) \to [0,1]^d$ be defined via $\pi(y)_i = y^i_{\mathbb{e}_i}$. As this is obviously a coordinate projection, it just remains to show that its image is $\operatorname{conv} X$. The set of vertices of $\operatorname{CTP}(\Xi_0, \Xi_1, \dots, \Xi_d)$ (i. e., the set of cyclic transversals of the blocks) is in bijection with X using π. Therefore, there is exactly one cyclic transversal for every choice of block element in Ξ_0 because every element of X has a unique coordinate decomposition. This means that there is exactly one way to choose elements from the other blocks $\Xi_1, \dots, \Xi_d$ to obtain a cyclic transversal. Vice versa, from a cyclic transversal $(\xi(0), \xi(1), \dots, \xi(d))$ one can just read off the corresponding element in X by checking $\xi(0)$. Therefore, $\pi(\operatorname{CTP}(\Xi_0, \Xi_1, \dots, \Xi_d))$ is equal to $\operatorname{conv} X$. ■

The construction from Proposition 3.1 can be seen as mapping X to the vertices of an appropriate simplex $\operatorname{conv}\left\{\mathbb{0}, \mathbb{e}_1, \dots, \mathbb{e}_{|X|}\right\} \subseteq \mathbb{R}^{|X|}$ and then taking the coordinates of a point in this simplex as coefficients for a convex combination in $\operatorname{conv} X$. The possibility of this construction of 0/1-polytopes as projections of simplices is a folklore result about 0/1-polytopes [53, Example/Exercise 4] and is indeed even possible for all polytopes as they by definition are the convex hulls of their vertices. Since this construction alone does not provide much insight into the possibilities of cyclic transversal polytopes, we also prove the following alternative universality theorem using Boolean formulas and projections:

3.2 Theorem (0/1-Polytopes as Projections of CTPs)
For every set $T_\varphi := \left\{x \in \{0,1\}^q \mid \varphi(x) = 1\right\}$ defined by a complete k-SAT Boolean formula φ with p clauses and q variables, the convex hull of T_φ can be represented as a projection of a cyclic transversal polytope of rank kp, order $p+q$ and size $2q + p(2^k - 1)$.

Proof: Our strategy is to establish a bijection between T_φ that is defined by φ and the cyclic transversals $\operatorname{CT}(\varphi)$, which are to be defined. The bijection will have the property that it is induced by a linear map between $\operatorname{CTP}(\varphi)$ and $[0,1]^q$. This linear map then is the required projection of $\operatorname{CTP}(\varphi)$ onto $\operatorname{conv} T_\varphi$.

Let $z_1, \dots, z_q$ be the variables of φ, and let $C_1, \dots, C_p$ be its clauses. For every clause C_i we set

$$\Xi_{C_i} := \mathbb{0} \times \dots \times \mathbb{0} \times \underbrace{\left(\mathbb{F}_2^k \setminus \mathbb{0}\right)}_{i\text{th factor}} \times \mathbb{0} \times \dots \times \mathbb{0} \subseteq \left(\mathbb{F}_2^k\right)^p = \mathbb{F}_2^{kp},$$

so that every block Ξ_{C_i} is a Cartesian product with p factors, which we will refer to as *sections*. Given a vector $\omega \in \mathbb{F}_2^{kp}$, let $\omega|_i \in \mathbb{F}_2^k$ denote the vector of its elements in the i-th section or just its i-th section for $i \in [p]$.

For every variable z_j, the block $\Xi_{z_j} \subseteq \mathbb{F}_2^{kp}$ consists of two vectors $\omega_{j,\top}$ and $\omega_{j,\bot}$, representing the choice of true or false for the variable z_j. These two vectors depend on whether z_j, $\neg z_j$ or neither occur in clause C_i, and also which other variables occur in C_i. To represent these clause-literal incidences, the vectors will also be partitioned into p sections of length k each. These sections will now be defined explicitly and simultaneously.

Let $z_{j_1}, \dots, z_{j_k}$ be the (ordered) variables and $\lambda_{j_1}, \dots, \lambda_{j_k}$ be the respective literals in clause C_i. Then the i-th sections of the vectors in block Ξ_{z_j} are defined according to the following cases: If z_j does not occur in C_i, then $\omega_{j,\bot}|_i = \omega_{j,\top}|_i := \mathbb{0}$. Otherwise, z_j occurs as some z_{j_ℓ} in C_i. If $\lambda_{j_\ell} = z_{j_\ell}$, we define $\omega_{j,\top}|_i := \mathbb{e}_\ell$ and $\omega_{j,\bot}|_i := \mathbb{0}$. In the third case we have $\lambda_{j_\ell} = \neg z_{j_\ell}$, in which case $\omega_{j,\bot}|_i$ is $\mathbb{e}_\ell \in \mathbb{F}_2^k$ and $\omega_{j,\top}|_i$ is defined to be $\mathbb{0}$. Since φ is complete, every variable occurs in some clause, so every variable block will have 2 distinct vectors.

This definition implies that for every clause C_i in which an assignment of some variable z_j satisfies the clause, the sum of chosen vectors $\xi(z_j) \in \Xi_{z_j}$ for $j \in [q]$ are not equal to $\mathbb{0}$ in the respective i-th section where there are parity-checking entries in the vectors of Ξ_{C_i}, meaning that this establishes the equivalence between a clause C_i being satisfied and $\sum_{j\in[q]} \xi(z_j)|_i$ being not equal to $\mathbb{0}$.

It is now straightforward to define a bijection between elements of T_φ and cyclic transversals in $\mathrm{CT}(\varphi) := \mathrm{CT}\big(\Xi_{z_1}, \dots, \Xi_{z_q}, \Xi_{C_1}, \dots, \Xi_{C_p}\big)$, which represent the vertices of

$$\mathrm{CTP}(\varphi) := \mathrm{CTP}\big(\Xi_{z_1}, \dots, \Xi_{z_q}, \Xi_{C_1}, \dots, \Xi_{C_p}\big),$$

such that this bijection is induced by a linear projection (i. e., a surjective map) of $\mathrm{CTP}(\varphi)$ onto $\operatorname{conv} T_\varphi$.

For $x \in T_\varphi$, set $\xi(z_j)$ to be the vector that represents the Boolean variable assignment for $z_j \in \{\bot, \top\}$ expressed by $x_j \in \{0, 1\}$, i. e., $\xi(z_j) \in \Xi_{z_j}$ is equal to $\omega_{j,\bot}$ if $x_j = 0$ and $\omega_{j,\top}$ otherwise. After fixing these vectors, the remaining $\xi(C_i)$ serve as parity-checks as mentioned, namely $\xi(C_i)|_\ell = \mathbb{0}$ for $\ell \neq i$ and $\xi(C_i)|_i = \sum_{j\in[q]} \xi(z_j)|_i$. This sum is different from $\mathbb{0}$ since x represents a variable assignment such that φ is satisfied ($\varphi(x) = 1$), meaning that there exists at least one variable z_j for every clause C_i such that C_i is satisfied with the given assignment. With these fixed vectors, the condition $\sum \xi(z_j) \oplus \sum \xi(C_i) = \mathbb{0}$ is fulfilled and $\big(\xi(z_1), \dots, \xi(z_q), \xi(C_1), \dots, \xi(C_p)\big)$ actually represents a cyclic transversal in $\mathrm{CT}(\varphi)$. Conversely, from a cyclic transversal we can uniquely recover $x \in T_\varphi$ by inspection of the first q vectors.

To extend this bijection of vertices to a projection of polytopes, let

$$\pi\colon\ \mathrm{CTP}(\varphi) \to [0,1]^q$$

be defined via $\pi(y)_j := y^{z_j}_{\omega_{j,\top}} = 1 - y^{z_j}_{\omega_{j,\bot}}$. Since this projection bijectively maps vertices of $\mathrm{CTP}(\varphi)$ to vertices of $\operatorname{conv} T_\varphi$ and is linear, the initial claim follows. ■

Note that the property that π is a bijection between the vertices of $\mathrm{CTP}(\varphi)$ and T_φ does not imply that π is an affine isomorphism, since it is not necessarily a bijection between $\operatorname{conv} T_\varphi$ and $\mathrm{CTP}(\varphi)$.

We mention here that for every set $T_\varphi := \{x \in \{0,1\}^q \mid \varphi(x) = 1\}$ which is defined by a k-SAT Boolean formula φ form with p clauses and q variables, there is also a 3-SAT formula ψ such that the number of its clauses is polynomially bounded in q and there are q variables in ψ such that the projection of the solution set of ψ onto these variables is T_φ, while having only polynomially many auxiliary variables [34, Problem 11]. Therefore, contrary to Proposition 3.1, Theorem 3.2 shows that we can represent polytopes whose membership problem is in NP, even those with an exponential number of vertices, as projections of cyclic transversal polytopes whose parameters, especially its size, are polynomial in the size of the input formula.

With Theorem 3.2 we now show that the algorithmic problem of recognizing, given a family of blocks, whether its set of cyclic transversals, or equivalently its cyclic transversal polytope, is empty, is NP-complete:

3.3 Lemma
Deciding whether $\mathrm{CT}(\Pi)$ *is empty for a given block configuration* $\Pi = (\Xi_1, \dots, \Xi_n)$ *is* NP*-complete, and thus the same holds for* $\mathrm{CTP}(\Pi)$*.*

Proof: From the construction in Theorem 3.2 we easily reduce the problem of finding a satisfying assignment of a 3-SAT Boolean formula with p clauses and q variables to the problem of checking whether a certain cyclic transversal polytope of rank $3p$, order $p + q$ and size $2q + 7p$ is empty. Since free variables do not influence the satisfiability of a Boolean formula, we can simply ignore them, so the formula can be assumed to be complete. As referenced in the literature, 3-SAT is known to be NP-complete, which has been shown by Karp [34, Problem 11].

Furthermore, checking whether a set of cyclic transversals is empty is obviously a problem in NP: Indeed, given a collection of n vectors in $\mathbb{F}_2^d$, checking whether they form a transversal of Π and checking whether it is cyclic can both be done in polynomial time in the input length, the input being the explicit block configuration itself. Therefore, the recognition problem for empty cyclic transversals and cyclic transversal polytopes is NP-complete. ■

To illustrate the proof of Theorem 3.2, let us give a concrete example:

3.4 Example (2-SAT Boolean formula)
Let z_1, z_2, z_3, z_4 be Boolean variables and let $\varphi = (z_1 \vee \neg z_2) \wedge (z_2 \vee z_3) \wedge (\neg z_1 \vee \neg z_4)$, i. e., we have three clauses $C_1 = z_1 \vee \neg z_2$, $C_1 = z_2 \vee z_3$ and $C_1 = \neg z_1 \vee \neg z_4$. Then, by evaluation of all 16 possibilities, we deduce that the set of valid assignments is

$$T_\varphi = \{(0,0,1,0)^\top, (0,0,1,1)^\top, (1,0,1,0)^\top, (1,1,0,0)^\top, (1,1,1,0)^\top\}.$$

One readily checks that these five vectors are affinely independent, and thus $\operatorname{conv} T_\varphi$ is affinely isomorphic to a 4-dimensional simplex such as $\operatorname{conv}\{\mathbb{0}, \mathbb{e}_1, \dots, \mathbb{e}_4\} \subseteq \mathbb{R}^4$. The blocks consist then of the columns shown in Table 3.1.

Block index	z_1 $\perp$	z_1 $\top$	z_2 $\perp$	z_2 $\top$	z_3 $\perp$	z_3 $\top$	z_4 $\perp$	z_4 $\top$	C_1			C_2			C_3		
$\mathbb{F}_2$-vector	0	1	0	0	0	0	0	0	1	0	1	0	0	0	0	0	0
	0	0	1	0	0	0	0	0	0	1	1	0	0	0	0	0	0
	0	0	0	1	0	0	0	0	0	0	0	1	0	1	0	0	0
	0	0	0	0	0	1	0	0	0	0	0	0	1	1	0	0	0
	1	0	0	0	0	0	0	0	0	0	0	0	0	0	1	0	1
	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	1	1

Table 3.1: Blocks for CTP representation of $\operatorname{conv} T_\varphi$ in Example 3.4

Since numerous families of polytopes can already be described by k-SAT Boolean formulas of polynomial length and fixed k, the constructions that will follow from now on reduce to the distinction of which polytopes are affinely isomorphic to cyclic transversal polytopes themselves, without allowing projections. Another example of this construction occurs in Corollary 3.12, by introducing a scheme to build 2-SAT Boolean formulas for the stable set polytope.

The idea of *sections* that emerged in the proof of Theorem 3.2 will continue to help with constructing further examples of cyclic transversal polytopes. If $[d]$ is partitioned into contiguous sets of consecutive numbers, we refer to the restriction of $\mathbb{F}_2^d$ to one of these sets of coordinates as a *section*. The venue space $\mathbb{F}_2^d$ is then described as the Cartesian product of its sections. Similarly, if a block $\Xi \subseteq \mathbb{F}_2^d$ is described as a Cartesian product $F_1 \times \ldots \times F_p$, we refer to its p factors as *sections*. Analogously, the notion of *sections* of a vector $\omega \in \mathbb{F}_2^d$ is inherited from the sections of $\mathbb{F}_2^d$. Here $\omega|_i \in \mathbb{F}_2^k$ denotes the vector of its elements in the i-th section or just called its i-th section, where $k \in [d]$ is the length of the section. In simple terms, there are now two hyperparameters to construct CTPs: The sections of the venue space, and the blocks to choose.

Using the language of sections, we prove an immediate corollary:

3.5 Corollary (Cartesian Products)
The Cartesian product of two cyclic transversal polytopes is also a cyclic transversal polytope. The rank, order, and size of the cyclic transversal polytope representing the Cartesian product can be chosen to be at most the sum of ranks, orders and sizes of the respective cyclic transversal polytope factors.

Proof: Let $\Xi_1^1, \ldots, \Xi_n^1 \subseteq \mathbb{F}_2^{d_1}$ and $\Xi_1^2, \ldots, \Xi_{n'}^2 \subseteq \mathbb{F}_2^{d_2}$ be two families of blocks and let P and Q be their associated cyclic transversal polytopes. For the Cartesian product, we define new blocks $\Xi_1, \ldots, \Xi_{n+n'} \subseteq \mathbb{F}_2^{d_1+d_2}$ as follows:

$$\Xi_i = \Xi_i^1 \times \mathbb{0} \subseteq \mathbb{F}_2^{d_1} \times \mathbb{F}_2^{d_2} \text{ for } i \in [n], \text{ and}$$
$$\Xi_{n+j} = \mathbb{0} \times \Xi_j^2 \subseteq \mathbb{F}_2^{d_1} \times \mathbb{F}_2^{d_2} \text{ for } j \in [n'].$$

Then the Cartesian product $P \times Q$ is the cyclic transversal polytope of these newly defined blocks, as there is a bijection between pairs of cyclic transversals $((\xi_1(1), \dots, \xi_1(n)), (\xi_2(1), \dots, \xi_2(n')))$ in the two block families that define P and Q, and a cyclic transversal $(\xi(1), \dots, \xi(n+n'))$ in the new blocks, namely by embedding the pair into the new blocks by mapping vectors to their counterpart in the respective section. This bijection linearly extends to a map between pairs of convex combinations of characteristic vectors of cyclic transversals for P and Q and convex combinations of characteristic vectors of cyclic transversals over $\Xi_1, \dots, \Xi_{n+n'}$. Checking the properties is straightforward. ■

Examining the proof of Theorem 3.2 also brings to light the following improvement:

3.6 Corollary (Improvement of Theorem 3.2 for $k = 2$)
Let φ be a complete 2-SAT Boolean formula with p clauses and q variables. Then the convex hull of $T_\varphi := \{x \in \{0,1\}^q \mid \varphi(x) = 1\}$ is affinely isomorphic to a cyclic transversal polytope of rank $2p$, order $p + q$ and size $2q + 3p$.

Proof: For $k = 2$ (such that φ represents a complete 2-SAT Boolean formula) the constructed cyclic transversal polytope in the proof of Theorem 3.2 is not only an extension, but already affinely isomorphic to $\operatorname{conv} T_\varphi$, as the values on variable blocks in this case uniquely determine the corresponding value on clause blocks, meaning that the constructed map is also injective. In particular, given a clause C with 2 literals $(\lambda_1 \vee \lambda_2)$, let x be a satisfying assignment. Then the vertex of $\mathrm{CTP}(\varphi)$ representing this assignment fulfills the linear equations

$$\begin{array}{ccccccc} y_{10} & & & + & y_{11} & = & x_1 \\ & & y_{01} & + & y_{11} & = & x_2 \\ y_{10} & + & y_{01} & + & y_{11} & = & 1, \end{array}$$

where (y_{01}, y_{10}, y_{11}) is the section of $y \in \mathrm{CTP}(\varphi)$ representing the selection of an element in the block associated to the clause C, namely, Ξ_C. The first two equations are given by the relation between the vectors of the blocks associated with λ_1 and λ_2, and the last equation is the block equation for Ξ_C.

Since this system of equations has full rank and the systems of equations for different clause-blocks are mutually exclusive by the sectioned construction, we can uniquely recover the chosen vectors of the clause-blocks given the choices on the blocks associated to the variables. Therefore, it is straightforward to see that the projection map is also injective and therefore constitutes an affine isomorphism. ■

For higher values of k, i. e., $k \geq 3$, this phenomenon cannot be reproduced anymore, since in general the number of equations does not match the number of additional variables: One obtains at least $2^k - 1 \geq 2^3 - 1 = 7$ variables for every clause block (one variable for every vector in $\mathbb{F}_2^k \setminus \mathbb{0}$), but the number of implied equations only grows linearly in k. It is in fact equal to k because every section has length k.

We now demonstrate various block configuration constructions for important families of polytopes that have been defined in the introduction, such as parity, cut and binary kernel polytopes, as well as some more families present in the literature, such as stable set and set packing polytopes. It is well-known that parity polytopes are a special subclass of cut polytopes, and that for every stable set polytope, there is an affinely isomorphic set packing polytope and vice versa, but different possible non-equivalent block configuration constructions emerge when we separately present them as cyclic transversal polytopes.

3.1.1 Parity, Cut and Binary Kernel Polytopes

With Theorems 3.7 and 3.8 we will now show that binary kernel polytopes are isomorphic to a particular subclass of cyclic transversal polytopes:

3.7 Theorem (Binary Kernel Polytopes are CTPs)
Let $M \in \mathbb{F}_2^{d \times n}$ be a binary matrix. Then the binary kernel polytope $\mathrm{KP}_2(M)$ is affinely isomorphic to a cyclic transversal polytope whose rank is equal to the rank of M, its order is n and its size is $2n$.

Proof: Without loss of generality, we may assume that M has full rank. Otherwise, we transform it into a (smaller) full-rank matrix and proceed with the modified matrix.

Now, let the columns of M be denoted by $M_{*,i} = \omega_i$ for $i \in [n]$. The block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$, given by $\Xi_i = \{\mathbb{0}, \omega_i\}$, generates the required cyclic transversal polytope with the appropriate parameters.

To show why $\mathrm{CTP}(\Pi) \cong \mathrm{KP}_2(M)$ is true, observe the bijective linear map

$$\pi\colon\ \mathrm{CTP}(\Pi) \to \mathrm{KP}_2(M),$$

which is defined via $\pi(y)_i = y^i_{\omega_i}$. Since a point in $\mathrm{CTP}(\Pi)$ is uniquely determined by $\left(y^1_{\omega_1}, \dots, y^n_{\omega_n}\right)$ because of the block equations, and π linearly extends the obvious bijection between vertices of both polytopes, π itself is bijective (and obviously linear), which makes it an affine isomorphism between $\mathrm{CTP}(\Pi)$ and $\mathrm{KP}_2(M)$. ■

Barahona and Grötschel [7] provided a description of $\mathrm{KP}_2(M)$ via inequalities, provided the associated matroid does not contain certain minors. For further information on the terms and concepts of matroid theory, we again refer to Oxley [41]. We will use and reformulate this result later in Theorem 5.11.

The proof of Theorem 3.7 can simply be reversed to prove Theorem 3.8, given some prerequisites on the block configuration:

3.8 Theorem (CTPs with blocks of cardinality 2 are Binary Kernel Polytopes)
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ *be a normalized block configuration over* $\mathbb{F}_2^d$*. Then, if* $|\Xi_i| = 2$ *is true for all* $i \in [n]$*, the polytope* $\mathrm{CTP}(\Pi)$ *is affinely isomorphic to a binary kernel polytope* $\mathrm{KP}_2(M)$ *for some matrix* $M \in \mathbb{F}_2^{d\times n}$*.*

Proof: Since $\Pi = (\Xi_1, \dots, \Xi_n)$ is normalized, all blocks are of the form $\Xi_i = \{\mathbb{0}, \omega_i\}$ for some $\omega_i \in \mathbb{F}_2^d \setminus \mathbb{0}$. Taking all ω_i as columns, we obtain a matrix $M \in \mathbb{F}_2^{d\times n}$. The affine map $\mu\colon \mathrm{KP}_2(M) \to \mathrm{CTP}(\Pi)$, defined by $\mu(x)^i_{\omega_i} = x_i$ and $\mu(x)^i_{\mathbb{0}} = 1 - x_i$ is then obviously an affine isomorphism. ■

It is easy to see that π from the proof of Theorem 3.7 is the inverse map to μ from the proof of Theorem 3.8. Given Theorems 3.7 and 3.8, we immediately see that binary kernel polytopes are characterized as the special case of cyclic transversal polytopes corresponding to (normalized) block configurations whose blocks have cardinality 2. We will apply this knowledge immediately after proving the next corollary, since any result about binary kernel polytopes applies to this subfamily of cyclic transversal polytopes.

3.9 Corollary
Any binary kernel polytope $\mathrm{KP}_2(M)$ *for* $M \in \mathbb{F}_2^{d\times n}$ *has a number of vertices equal to* 2^k *for some* $k \in \mathbb{N}$*.*

Proof: Since the kernel of a matrix $M \in \mathbb{F}_2^{d\times n}$ is a linear subspace of $\mathbb{F}_2^n$ and any subspace of $\mathbb{F}_2^n$ has 2^k elements for some $k \in \mathbb{N}$, the claim follows, since the set of vertices of $\mathrm{KP}_2(M)$ is in bijection with the kernel of M. ■

We now turn our attention to binary kernel polytopes which have come up in other contexts. Among the most prominent examples are cut polytopes of graphs. We remind the reader that a thorough discussion about cuts and their geometry in general is given by Deza and Laurent [21].

Since we know that the characteristic vectors of cuts of a graph $G = (V, E)$ form a subspace of the edge-space $\mathbb{F}_2^E$ of G, any matrix M that has this subspace as its kernel is suitable to represent $\mathrm{CUT}(G)$ as a binary kernel polytope $\mathrm{KP}_2(M)$. Corollary 3.10 then follows immediately:

3.10 Corollary (Cut Polytopes are CTPs)
For any graph $G = (V, E)$*, its cut polytope* $\mathrm{CUT}(G)$ *is affinely isomorphic to a cyclic transversal polytope.*

The subspace of cuts and the subspace of cycles of a graph $G = (V, E)$ are orthogonal subspaces to each other (see, e. g., the textbook by Diestel [22, Theorem 1.9.4]). Therefore, a suitable matrix $M \in \mathbb{F}_2^{d \times |E|}$ to represent $\mathrm{CUT}(G)$ consists of rows that are characteristic vectors of cycles of G. In particular, bases of the cycle subspace generate suitable matrices (see [22, Section 1.9] for more information). This means that for a connected graph, the cyclic transversal polytope from Corollary 3.10 can be given by a block configuration that has rank $d = |E| - |V| + 1$, order $n = |E|$ and size $s = 2|E|$.

As seen in the introduction, cut polytopes of graphs that are themselves cycles of length n are also known as parity polytopes $\mathrm{PAR}(n)$, since they are exactly the polytopes that are the convex hulls of all 0/1-vectors of length n with an even number of ones. Equivalently, they are the binary kernel polytopes for matrices $\mathbb{1}_{1,n}$, where $\mathbb{1}_{d,n}$ is the matrix of shape $d \times n$ with every entry equal to one. Therefore, we state Corollary 3.11 without proof:

3.11 Corollary (Parity Polytopes are CTPs)
For any $n \in \mathbb{N}$, the polytope $\mathrm{PAR}(n)$ is affinely isomorphic to $\mathrm{CTP}(\Xi_1, \ldots, \Xi_n)$ with Ξ_i equal to $\mathbb{F}_2$ for all $i \in [n]$.

It is clear that the cyclic transversal polytope from Corollary 3.11 has rank 1. Since the automorphism group of $\mathbb{F}_2$ is trivial, it is generated by the unique normalized block configuration of rank 1, order n and size $2n$.

After having proven that binary kernel polytopes, together with some important exemplary subfamilies of them, are themselves a subfamily of cyclic transversal polytopes, we continue with more elaborate constructions where larger blocks will be involved.

3.1.2 Stable Set and Set Packing Polytopes

A *stable* or *independent set* is a set of nodes of a graph $G = (V, E)$ such that no two nodes of this set are adjacent. Naturally, taking the convex hull of the characteristic vectors of these sets, we define a polytope, called the *stable set polytope* $\mathrm{STAB}(G)$. This polytope and the associated combinatorial optimization problem has been extensively investigated in the literature, e. g., in early contributions by Chvátal [14], Nemhauser and Trotter [40], and Padberg [42], as well as newer works, e. g., by Conforti et al. [17] and Lipták and Lovász [37] and the references therein.

Formally, the stable set polytope of a graph $G = (V, E)$ is also defined as

$$\mathrm{STAB}(G) := \mathrm{conv}\Big\{x \in \{0,1\}^V \;\Big|\; x_u + x_v \leq 1 \text{ for all } \{u, v\} \in E\Big\}.$$

Using our knowledge about cyclic transversal polytopes of complete 2-SAT formulas, we now prove the following corollary:

3.12 Corollary (Stable Set Polytopes)
Let $G = (V, E)$ be a graph. Then if every node of G has at least one edge incident to it, there exists a cyclic transversal polytope of rank $2|E|$, order $|V| + |E|$ and size $2|V| + 3|E|$ which is affinely isomorphic to $\mathrm{STAB}(G)$. Otherwise, there exists a cyclic transversal polytope of rank $2(|E| + |S|)$, order $|V| + |E| + |S| + 1$ and size $1 + 2|V| + 3(|E| + |S|)$ which is affinely isomorphic to $\mathrm{STAB}(G)$, where $S \subseteq V$ is the set of nodes without any incident edges in G.

Proof: If every node in $G = (V, E)$ has at least one edge incident to it, we construct a monotone complete Boolean formula in $|V|$ variables with $|E|$ clauses and $k = 2$ literals for every clause, such that any satisfying assignment of truth values to the variables represents a stable set and vice versa. Then we use this constructed 2-SAT formula and a slightly different projection map with Theorem 3.2 to obtain an appropriate cyclic transversal polytope which, by an augmentation of the previous argument in Corollary 3.6, is affinely isomorphic to $\mathrm{STAB}(G)$.

First, let $\varphi_{\mathrm{STAB}(G)}$ be defined as

$$\varphi_{\mathrm{STAB}(G)} \coloneqq \bigwedge_{e=\{u,v\}\in E} (z_u \vee z_v) \eqqcolon \bigwedge_{e=\{u,v\}\in E} C_e.$$

Here z_v is a Boolean variable that is true if $v \in V$ is not in a stable set, i. e., if T is a stable set of G, we set $z_v = \top$ if and only if $v \notin T$ and $z_v = \perp$ otherwise. This way, the truth value assignments that satisfy $\varphi_{\mathrm{STAB}(G)}$ are in bijection with the stable sets of G, since a clause C_e is unsatisfied if and only if adjacent nodes are picked to both be in a stable set, but this directly contradicts the definition of a stable set. This argument works both ways, so any satisfying assignment of $\varphi_{\mathrm{STAB}(G)}$ translates into a unique stable set of G.

Therefore, by Theorem 3.2 and after ordering the edges in some fashion, we obtain the blocks

$$\Xi_e \coloneqq \Xi_{C_e} \coloneqq \mathbb{0} \times \ldots \times \mathbb{0} \times \underbrace{(\mathbb{F}_2^2 \setminus \mathbb{0})}_{\text{edge } e} \times \mathbb{0} \times \ldots \times \mathbb{0} \subseteq \left(\mathbb{F}_2^2\right)^{|E|} = \mathbb{F}_2^{2|E|},$$

as well as $\Xi_v \coloneqq \Xi_{z_v} \coloneqq \{\mathbb{0}, \omega_v\} \subseteq \mathbb{F}_2^{2|E|}$ for $v \neq v_0$. Here ω_v for $v \in V$ are vectors that are decomposable into sections, indexed by $e \in E$. The section $\omega_v|_e \in \mathbb{F}_2^2 \setminus \mathbb{1}$ then is the characteristic vector of the incidence $v \in e$ for $v \in V$ and $e \in E$.

In contrast to the projection map from Theorem 3.2, we define

$$\pi\colon\ \mathrm{CTP}\left(\varphi_{\mathrm{STAB}(G)}\right) \to [0,1]^V$$

via $\pi(y)_v := y^v_{\mathbb{0}}$ for $v \in V$, since the variables z_j represent the negation of $v \in V$ being in a stable set. Since every variable block only contains two elements, the representing coordinates of which are linked by a block equation, this does not pose a problem to π being an affine isomorphism. Indeed, π is an affine isomorphism as it is the identity on $y^v_{\mathbb{0}}$ for $v \in V$, and these coordinates alone already uniquely determine a point in $\mathrm{CTP}\big(\varphi_{\mathrm{STAB}(G)}\big)$ because of the block equations for Ξ_v with $v \in V$ and additional equations present in the 2-SAT construction from Corollary 3.6, namely

$$y^v_{\omega_v} = 1 - y^v_{\mathbb{0}} \quad \text{for all } v \in V$$

and

$$\begin{array}{ccccccc} y^e_{\mathbb{e}_1} & & & + & y^e_{\mathbb{1}} & = & 1 - y^u_{\mathbb{0}} \\ & & y^e_{\mathbb{e}_2} & + & y^e_{\mathbb{1}} & = & 1 - y^v_{\mathbb{0}} \\ y^e_{\mathbb{e}_1} & + & y^e_{\mathbb{e}_2} & + & y^e_{\mathbb{1}} & = & 1 \end{array}$$

for all $e = \{u, v\} \in E$ where, by abuse of notation, $\big(y^e_{\mathbb{e}_1}, y^e_{\mathbb{e}_2}, y^e_{\mathbb{1}}\big)$ means the section of $y \in \mathrm{CTP}(\varphi)$ representing the selection of an element in the block Ξ_e for any $e \in E$.

Now in case $G = (V, E)$ has any isolated nodes, i. e., nodes without any edges incident to them, we modify G as follows: We introduce a fictitious node v_0 and add a new edge $\{v, v_0\}$ for every isolated node $v \in S$. In this augmented graph $G' = (V', E')$, there are no isolated nodes. We observe that $|V'| = |V| + 1$ and $|E'| = |E| + |S|$ holds. We can then apply the previously described construction to G', noting that the stable sets of G are in bijection with the stable sets of G' that do not contain v_0. Consequently, we set $z_{v_0} = \top$ in this case and modify the construction of block Ξ_{v_0} to be the trivial block $\big\{\omega_{v_0}\big\}$, reducing the size parameter of the resulting cyclic transversal polytope by one. This yields the correct parameters. By defining the affine map π analogously to before, we establish the assertion that the resulting cyclic transversal polytope and $\mathrm{STAB}(G)$ are affinely isomorphic also in this case. ■

We note that in general, the parameters in Corollary 3.12 are not optimal and can be reduced further, for example by employing the techniques from Section 2.2.2, but the clarity of the construction would suffer in doing so within the proof above.

Interestingly, since the coordinates of $\mathbb{F}_2^{2|E|}$ in the preceding proof are basically indexed by the pairs $(v, e) \in V \times E$ with $v \in e$, there are two ways to partition the venue space for the stable set polytope construction into sections: One where every section has length 2 and is indexed by the edges $e \in E$ as done in the proof, and another one where the sections are of (varying) length $|\delta(v)|$ for $v \in V$, where $\delta(v)$ are the edges incident with v. By reordering of the venue space coordinates, one may therefore write the node blocks as

$$\Xi_v = \mathbb{0} \times \ldots \times \mathbb{0} \times \underbrace{(\{\mathbb{0}, \mathbb{1}\})}_{\text{length } |\delta(v)|} \times \mathbb{0} \times \ldots \times \mathbb{0} \subseteq \mathbb{F}_2^{2|E|},$$

which leads to an equivalent cyclic transversal polytope construction for stable set polytopes, where the appropriate equivalence transformation of the block configuration is simply the block isomorphism that permutes the venue space coordinates.

Since the empty set $\emptyset \subsetneq V$ is always a valid stable set for $G = (V, E)$, its associated cyclic transversal may be used to centralize the participating blocks, since the construction in the proof is obviously not centralized.

Note that the classical stable set (or *vertex packing*) decision problem, which is known to be NP-complete [34, Problem 4], asks whether a given graph G contains a stable set of size ℓ, for some parameter ℓ which is part of the input. Therefore, constructing a 2-SAT Boolean formula $\varphi_{\mathrm{STAB}(G)}$ that has exactly the characteristic vectors of stable sets as satisfying assignments does not contradict the hardness of the former problem or imply that it is actually easy to solve, but finding a satisfying assignment for such formulas is already a problem in P. The classical stable set decision problem translates to whether there is a satisfying assignment of $\varphi_{\mathrm{STAB}(G)}$ which has (at least) ℓ of the variables set to $\perp$ for false, or its corresponding cyclic transversal being $\mathbb{0}$ in at least ℓ elements.

As the construction for stable set polytopes is rather involved to show that they belong to the family of cyclic transversal polytopes, one could be inclined to ask about the reverse direction: Whether any given cyclic transversal polytope is affinely isomorphic to a stable set polytope of some graph G. This could be suspected especially since stable set polytopes are not necessarily representable as kernel polytopes, since the number of stable sets of a given graph in general is not a power of 2 by necessity. The following example using Corollary 3.11 shows us that this is not the case, and cyclic transversal polytopes really are a more general class of polytopes that contain both stable set and cut polytopes:

3.13 Example (PAR(4) is not a stable set polytope)
Since every singleton set of vertices is a stable set of a graph $G = (V, E)$, we know that $\dim \mathrm{STAB}(G) = |V|$. Since the parity polytope PAR(4) is of dimension 4 and has 8 vertices, we can enumerate the 11 non-isomorphic graphs on 4 vertices to find that only two of them have 8 stable sets. One of them is the disjoint union of an isolated node and a complete graph on 3 vertices, giving rise to a stable set polytope which is affinely isomorphic to the Cartesian product

$$[0,1] \times \left\{ x \in \mathbb{R}^3 \;\middle|\; x \geq \mathbb{0}, \sum_{i \in [3]} x_i \leq 1 \right\}.$$

This polytope is not even combinatorially isomorphic to PAR(4), therefore it is also not affinely isomorphic to it. The other possible graph is the path on 4 vertices whose stable set polytope only has 7 facets, while PAR(4) has 16 facets, again showing that these polytopes cannot be combinatorially isomorphic.

We do note that this example is actually independent of cyclic transversal polytopes and shows that the families of cut and stable set polytopes are incomparable, i. e., neither family is contained in the other. The direction not shown in the example can easily be seen by observing that the number of (possibly empty) cuts in a graph are necessarily even (they are always a power of 2) and therefore not even all simplices (see Proposition 3.18) are cut polytopes, while the number of stable sets of a graph may be odd.

Another construction related to stable set polytopes concerns set packing polytopes, which are also well-known [42], including their relation to stable set polytopes [40]. Given a universe $\mathcal{U}$ and a family $\mathcal{S}$ of subsets of $\mathcal{U}$, a *packing* is a collection of subsets $\mathcal{C} \subseteq \mathcal{S}$ such that any two sets in $\mathcal{C}$ are disjoint. The convex hull of characteristic vectors of such packings gives rise to the set packing polytope, which is also defined as

$$\mathrm{PACK}(\mathcal{U}, \mathcal{S}) := \operatorname{conv}\left\{ x \in \{0,1\}^{\mathcal{S}} \;\middle|\; \sum_{T \in \mathcal{S} \colon e \in T} x_s \leq 1 \text{ for all } e \in \mathcal{U} \right\}.$$

The optimization problems of finding a stable set of maximum cardinality in a given graph G, also called the *maximum stable set problem*, and finding a maximum number of pairwise disjoint sets in a set family, i. e., the *maximum set packing problem*, are in fact reducible to one another in a structure-preserving manner [2, Theorem 4]: Taking the graph

$$G(\mathcal{S}) := (\mathcal{S}, E_{\mathcal{S}}) \text{ with } E_{\mathcal{S}} := \{\{T_1, T_2\} \mid T_1 \cap T_2 \neq \emptyset\}$$

as an instance of the maximum stable set problem reduces the maximum set packing problem to it. Here every chosen node in a maximum stable set corresponds to a chosen subset in a valid collection of disjoint subsets, which necessarily is also maximum. Reducing the stable set problem to the set packing problem on the other hand is done by choosing $\mathcal{U} = E$ and $\mathcal{S} = \{\delta(v) \mid v \in V\}$ as an instance for the maximum set packing problem, where $\delta(v)$ is the set of incident edges for the node $v \in V$. The solution for the former problem can be recovered by identifying every node $v \in V$ with its set $\delta(v)$.

Note that this reduction is formally only correct if the sets of neighboring edges are pairwise non-identical, which can be mitigated by allowing identical edge subsets in $\mathcal{S}$. Identical edge sets $\delta(u)$ and $\delta(v)$ only occur when $\{u, v\}$ is an isolated edge or both u and v are isolated nodes, on which the stable set problem is trivial to solve. It goes without saying that splitting a graph instance into its connected components and solving the maximum stable set problem on each component separately yields an optimal solution to the maximum stable set problem on the whole graph.

Although these reductions are straightforward, we now surprisingly construct a family of blocks for the set packing polytope such that the resulting cyclic transversal polytope is affinely isomorphic to the appropriate stable set polytope and its linked cyclic transversal polytope construction from Corollary 3.12, but the two block configurations for the set packing and stable set polytopes are not equivalent under Definition 2.3.

3.14 Theorem (Set Packing Polytopes)
Given a universe $\mathcal{U}$ and a family $\mathcal{S}$ of its subsets, there is a cyclic transversal polytope of rank $\sum_{T\in\mathcal{S}}(|T|-1)$, order $|\mathcal{U}|$ and size $|\mathcal{U}|+\sum_{T\in\mathcal{S}}|T|$ that is isomorphic to $\mathrm{PACK}(\mathcal{U},\mathcal{S})$.

Proof: We will construct one block for every $e \in \mathcal{U}$ that contains canonical basis vectors depending on which sets $T \in \mathcal{S}$ the element e is contained in. To make identification of coordinates of the venue space easier we order the subsets in $\mathcal{S}$ such that $\mathcal{S} = \{T_1, \dots, T_k\}$ and define

$$\mathcal{D} := \dot{\bigcup_{T_i\in\mathcal{S}}} \{(i,j) \mid j \in [|T_i|-1]\},$$

which will serve as the set of coordinates for this construction. Note that its size is $\sum_{T_i\in\mathcal{S}}(|T_i|-1)$ as claimed. The elements $e \in \mathcal{U}$ are also assumed to be ordered as $\mathcal{U} = \{e_1, \dots, e_m\}$, and every $T_i \in \mathcal{S}$ inherits this order so that $T_i = \left\{e_{(i,1)}, \dots, e_{(i,|T_i|)}\right\}$ corresponds to some restriction of $\mathcal{U}$ with the same order. As a shorthand we define $\mathbb{e}_{(i,0)}$ and $\mathbb{e}_{(i,|T_i|)}$ to be $\mathbb{0}$. Now every block Ξ_e contains $\mathbb{0}$ and some (sums of) canonical basis vectors from the venue space $\mathbb{F}_2^{\mathcal{D}}$, namely

$$\Xi_e := \mathbb{0} \cup \left(\dot{\bigcup_{T_i\in\mathcal{S}}} \left\{ \mathbb{e}_{(i,j-1)} \oplus \mathbb{e}_{(i,j)} \;\middle|\; e = e_{(i,j)} \in T_i \right\} \right).$$

The cyclic transversal polytope $\mathrm{CTP}\left(\Xi_{e_1}, \dots, \Xi_{e_m}\right)$ using these blocks is then claimed to be isomorphic to $\mathrm{PACK}(\mathcal{U},\mathcal{S})$.

The sum of two consecutive canonical basis vectors in every block, or exactly one canonical basis vector in the two edge cases, ensures that if we choose one of the canonical basis vectors $\mathbb{e}_{(i,j)}$ for any set T_i and $j \in [|T_i|-1]$, there is exactly one other block that also uses this canonical basis vector in a sum. To fulfill the cyclic transversal condition that the sum of all chosen vectors equals $\mathbb{0}$, the only possibility is to choose all or none of the $\mathbb{e}_{(i,j)}$ for any T_i, because $2\sum_{j\in[|T_i|]} \mathbb{e}_{(i,j)} = \mathbb{0}$ is trivially true, whereas omission of one of these vectors or two distinct ones results in a sum different from $\mathbb{0}$. This condition translates to an equation of the corresponding coordinates for the cyclic transversal polytope, i. e., all values $y^{e_{(i,j)}}_{\mathbb{e}_{(i,j-1)}\oplus\mathbb{e}_{(i,j)}}$ are equal for fixed $i \in [k]$. For 0/1-points this means that either the whole set $T_i \in \mathcal{S}$ is part of the collection, or the whole set is not in the collection represented by the point.

The distribution of the sums to the blocks in turn ensures that for every element $e \in \mathcal{U}$, at most one of the sets T_i that contain it is chosen to be in a collection, which is exactly the set packing condition. That means the vertices of $\mathrm{CTP}\left(\Xi_{e_1}, \dots, \Xi_{e_m}\right)$ are in bijection to the vertices of $\mathrm{PACK}(\mathcal{U}, \mathcal{S})$.

To prove the claim of isomorphy of both polytopes, we finally argue that the linear map $\pi\colon \mathrm{CTP}\left(\Xi_{e_1}, \dots, \Xi_{e_m}\right) \to \mathrm{PACK}(\mathcal{U}, \mathcal{S})$, defined by $\pi(y)_{T_i} = y^{e_{(i,1)}}_{\mathbb{e}_{(i,1)}}$, is bijective, since the aforementioned equations imply that a point in $\mathrm{CTP}\left(\Xi_{e_1}, \dots, \Xi_{e_m}\right)$ is uniquely determined by the values on $\left(y^{e_{(1,1)}}_{\mathbb{e}_{(1,1)}}, \dots, y^{e_{(k,1)}}_{\mathbb{e}_{(k,1)}}\right)$ alone. ■

To make this proof easier to digest, we provide an example:

3.15 Example (Set Packing)
Let $\mathcal{U} = [6]$ be the elements with their canonical order and $\mathcal{S} = \{T_1, T_2, T_3\}$ be a family of subsets with

$$T_1 = \{1, 2, 5\}, T_2 = \{2, 3, 4\} \text{ and } T_3 = \{3, 4, 5, 6\}.$$

By examination, it is clear that at most one of the sets is chosen in a set packing collection, as they have pairwise common elements. Nevertheless, we want to understand the cyclic transversal construction. According to the proof, we have the coordinates

$$\mathcal{D} = \{(1,1), (1,2), (2,1), (2,2), (3,1), (3,2), (3,3)\}.$$

Then we obtain one block for every element of $\mathcal{U}$, which are broken down in Table 3.2.

The choice of sums of basis vectors of $\mathbb{F}_2^{\mathcal{D}}$ and their distribution among the blocks ensures that any element y of $\mathrm{CTP}(\Xi_1, \dots, \Xi_6)$ fulfills the block equations, as well as the additional equations

$$\begin{aligned}
y^1_{\mathbb{e}_{(1,1)}} &= y^2_{\mathbb{e}_{(1,1)} \oplus \mathbb{e}_{(1,2)}} = y^5_{\mathbb{e}_{(1,2)}}, \\
y^2_{\mathbb{e}_{(2,1)}} &= y^3_{\mathbb{e}_{(2,1)} \oplus \mathbb{e}_{(2,2)}} = y^4_{\mathbb{e}_{(2,2)}}, \\
y^3_{\mathbb{e}_{(3,1)}} &= y^4_{\mathbb{e}_{(3,1)} \oplus \mathbb{e}_{(3,2)}} = y^5_{\mathbb{e}_{(3,2)} \oplus \mathbb{e}_{(3,3)}} = y^6_{\mathbb{e}_{(3,3)}},
\end{aligned}$$

which means that the values are row-wise constant when viewed as in Table 3.2. Obviously, an element $y \in \mathrm{CTP}(\Xi_1, \dots, \Xi_6)$ is then uniquely determined by $\left(y^1_{\mathbb{e}_{(1,1)}}, y^2_{\mathbb{e}_{(2,1)}}, y^3_{\mathbb{e}_{(3,1)}}\right)$. Since the equations above together with the block equations imply all set packing conditions on these coordinates and $\mathrm{CTP}(\Xi_1, \dots, \Xi_6)$ is integral, $\mathrm{CTP}(\Xi_1, \dots, \Xi_6) \cong \mathrm{PACK}(\mathcal{U}, \mathcal{S})$ holds.

	Ξ_1	Ξ_2	Ξ_3	Ξ_4	Ξ_5	Ξ_6
T_1	$\mathbb{e}_{(1,1)}$	$\mathbb{e}_{(1,1)} \oplus \mathbb{e}_{(1,2)}$			$\mathbb{e}_{(1,2)}$	
T_2		$\mathbb{e}_{(2,1)}$	$\mathbb{e}_{(2,1)} \oplus \mathbb{e}_{(2,2)}$	$\mathbb{e}_{(2,2)}$		
T_3			$\mathbb{e}_{(3,1)}$	$\mathbb{e}_{(3,1)} \oplus \mathbb{e}_{(3,2)}$	$\mathbb{e}_{(3,2)} \oplus \mathbb{e}_{(3,3)}$	$\mathbb{e}_{(3,3)}$
	$\mathbb{0}$	$\mathbb{0}$	$\mathbb{0}$	$\mathbb{0}$	$\mathbb{0}$	$\mathbb{0}$

Table 3.2: Blocks of the set packing construction for Example 3.15

We can represent the stable set polytope $\mathrm{STAB}(G)$ of a graph $G = (V, E)$ that is connected and has at least three nodes as a set packing polytope $\mathrm{PACK}(\mathcal{U}, \mathcal{S})$ by choosing $\mathcal{U} = E$ and $\mathcal{S} = \{\delta(v) \mid v \in V\}$, just as in the reduction from the maximum stable set problem to the maximum set packing problem.

The block family construction for the set packing polytope from Theorem 3.14 can then be applied to these sets $\mathcal{U}$ and $\mathcal{S}$. Therefore, we get a second (affinely isomorphic) representation of stable set polytopes as cyclic transversal polytopes for certain graphs, but this second construction is not equivalent to the construction in Corollary 3.12 under Definition 2.3. Indeed, that the block configurations are not equivalent is apparent in their parameters: the set packing construction has smaller parameters than the stable set construction. The rank reduces from $2|E|$ to $2|E| - |V|$, the order reduces from $|E| + |V|$ to $|E|$, and the size reduces from $2|V| + 3|E|$ to $3|E|$. Yet, both constructions are pruned: The empty solution and all singleton solutions to both problems generate a subset of the cyclic transversals which utilize all block elements.

Example 3.15 and the equations confirm how the coordinates $y^e_{\mathbb{0}}$ for every element $e \in \mathcal{U}$ act as slack variables for the set packing polytope, signifying how much every element e is not contained in any of the sets of a chosen collection of a set packing.

While meaningful sections of the venue space, given by the subsets in $\mathcal{S}$, are possible to introduce and investigate, there are very few blocks only containing vectors that are non-zero in exactly one section. Namely, these are blocks whose element is only contained in one subset. However, the number of such sections that a block contains non-zero elements from is clearly equal to the number of subsets the element is contained in.

For an emergent corollary, a *matching* of a graph $G = (V, E)$ is a subset of the edges such that no two edges are incident. Based on this definition, one readily sees that the matching polytope of G, denoted as $\mathrm{MP}(G)$ and defined as the convex hull of characteristic vectors of matchings of G, is a special case of a set packing polytope, with $\mathcal{U} = V$ and $\mathcal{S} = E$.

3.16 Corollary (Matching Polytopes)
Let $G = (V, E)$ be a graph. Then $\mathrm{MP}(G)$ is isomorphic to a cyclic transversal polytope of rank $|E|$, order $|V|$ and size $|V| + 2|E|$.

The proof of Corollary 3.16 follows directly from the proof of Theorem 3.14 for set packing polytopes. Since every edge has cardinality 2, the proof could be simplified further as there are no sums of two canonical basis vectors. We show this here using an example.

3.17 Example (The square pyramid in three dimensions)
As the matching polytope of the path with 3 edges (on 4 vertices), a square pyramid in three dimensions is a cyclic transversal polytope. According to the set packing construction applied to $\mathcal{U} = [4]$ and $\mathcal{S} = \{\{1, 2\}, \{2, 3\}, \{3, 4\}\}$, the configuration $\Pi = (\Xi_1, \dots, \Xi_4)$

turns out to be equivalent to

$$\Xi_1 := \{\mathbb{0}, \mathbb{e}_1\}, \Xi_2 := \{\mathbb{0}, \mathbb{e}_1, \mathbb{e}_2\}, \Xi_3 := \{\mathbb{0}, \mathbb{e}_2, \mathbb{e}_3\}, \text{ and } \Xi_4 := \{\mathbb{0}, \mathbb{e}_3\}.$$

Then $\operatorname{CTP}(\Xi_1, \dots, \Xi_4)$ is (affinely isomorphic to) a three-dimensional pyramid over the square. Its vertices can be read off of Table 3.3.

	Ξ_1		Ξ_2			Ξ_3			Ξ_4	
	$\mathbb{0}$	$\mathbb{e}_1$	$\mathbb{0}$	$\mathbb{e}_1$	$\mathbb{e}_2$	$\mathbb{0}$	$\mathbb{e}_2$	$\mathbb{e}_3$	$\mathbb{0}$	$\mathbb{e}_3$
v_1	1	0	1	0	0	1	0	0	1	0
v_2	0	1	0	1	0	1	0	0	1	0
v_3	1	0	0	0	1	0	1	0	1	0
v_4	1	0	1	0	0	0	0	1	0	1
v_5	0	1	0	1	0	0	0	1	0	1

Table 3.3: Vertices of $\operatorname{CTP}(\Xi_1, \dots, \Xi_4)$ and their cyclic transversals from Example 3.17, with columns marked according to the affine isomorphism

The matching polytope $\operatorname{MP}(G)$ is also the stable set polytope of the *line graph* of G. The line graph of a graph $G = (V, E)$ is $L(G) := (E, F)$, where $F := \{\{e, f\} \mid e \cap f \neq \emptyset\}$ is the set of unordered pairs of intersecting edges of G. This connection between the matching polytope of G and the stable set polytope of $L(G)$ is well-known and readily discernible. Consequently, the cyclic transversal polytope construction described in Corollary 3.12 applies to matching polytopes of graphs, but it works differently than the set packing construction. In particular, for graphs G without isolated edges, the parameters for this construction are $2|F|$ for the rank, $|E| + |F|$ for the order and $2|E| + 3|F|$ for the size, and it is worth noting that these parameters generally differ from those in Corollary 3.16.

3.1.3 Further Constructions

Some additional and more straightforward constructions for more foundational polytopes are collected in this section.

An immediate observation concerns the *standard simplex*

$$\Delta_k := \operatorname{conv}\{\mathbb{0}, \mathbb{e}_1, \dots, \mathbb{e}_k\} \subseteq \mathbb{R}^k.$$

Because Δ_k is the stable set polytope of the complete graph $K_k = ([k], E_k = \binom{[k]}{2})$ and all stable set polytopes are cyclic transversal polytopes by Corollary 3.12, it is an immediate corollary that Δ_k can be described as a cyclic transversal polytope as in the proof of Corollary 3.12. But this is not the only construction to obtain a cyclic transversal polytope that is affinely isomorphic to the standard simplex:

3.18 Proposition (Simplex)
Let $\Xi \subseteq \mathbb{F}_2^d$ *be of cardinality* $|\Xi| = k+1$*. Then the standard simplex* Δ_k *is affinely isomorphic to the cyclic transversal polytope* $\mathrm{CTP}(\Xi, \Xi)$*.*

Proof: Let the elements of $\Xi \subseteq \mathbb{F}_2^d$ be enumerated as $\xi_0, \dots, \xi_k$. We have to show that the linear map $\pi\colon \mathbb{R}^k \to \mathbb{R}^\Xi \times \mathbb{R}^\Xi$, defined by $\pi(x)^1_{\xi_i} = \pi(x)^2_{\xi_i} = x_i$ for $i \in [k]$ as well as $\pi(x)^1_{\xi_0} = \pi(x)^2_{\xi_0} = 1 - \sum_{i\in[k]} x_i$ maps the points of Δ_k to $\mathrm{CTP}(\Xi, \Xi)$ in a bijective manner.

Note that aside from the non-negativity constraints and block equations, the equations $y^1_{\xi_i} = y^2_{\xi_i}$ for all $i \in [k]_0$ are valid for $\mathrm{CTP}(\Xi, \Xi)$, since the cyclic transversal condition $\xi(1) \oplus \xi(2) = \mathbb{0}$ implies $\xi(1) = \xi(2)$. Further, the block equations imply that

$$y^\ell_{\xi_0} = 1 - \sum_{i\in[k]} y^\ell_{\xi_i} \text{ and } \sum_{i\in[k]} y^\ell_{\xi_i} \leq 1 \text{ for all } \ell \in [2].$$

Bijectivity of π then follows from the simple fact that π is the identity when restricted to the coordinates $(y^1_{\xi_1}, \dots, y^1_{\xi_k}) \in \Delta_k$, and these coordinates already uniquely define a point in $\mathrm{CTP}(\Xi, \Xi)$ as seen by the equations and inequalities. ∎

The construction in the proof of Corollary 3.12 applied to K_k yields a cyclic transversal polytope of rank $2\binom{k}{2} = k(k-1)$, order $k + \binom{k}{2}$ and size $2k + 3\binom{k}{2}$, while Proposition 3.18 yields an affinely isomorphic cyclic transversal polytope with rank in $\mathcal{O}(\log k)$, order 2 and size $2(k+1)$.

It follows directly from Corollary 3.5, together with Proposition 3.18, that Cartesian products of any family of simplices are cyclic transversal polytopes. In fact, that means that every simple 0/1-polytope is affinely isomorphic to a cyclic transversal polytope. This follows from the fact that every simple 0/1-polytope is equal to a Cartesian product of 0/1-simplices, which was proven by Kaibel and Wolff [33, Theorem 1].

Nevertheless, one such simple 0/1-polytope shall be emphasized:

3.19 Corollary (Cube)
Let $d \in \mathbb{N}$ *be fixed, then the* d*-cube* $[0,1]^d$ *is isomorphic to a cyclic transversal polytope. The rank, order and size of this cyclic transversal polytope can be chosen to be at most* d*,* $2d$ *and* $4d$*, respectively.*

Proof: Since $[0,1]^d$ is the Cartesian product of d simplices of dimension 1, the result follows directly from Corollary 3.5 and Proposition 3.18. The statement about rank, order and size is seen by observing that every coordinate is its own section. The $2d$ blocks are then given by taking twice every subspace of the form $\{\mathbb{0}, \mathbb{e}_i\} \subseteq \mathbb{F}_2^d$ with $i \in [d]$. ∎

With these finishing constructive results regarding cyclic transversal polytopes, we now switch to properties of polytopes which prohibit their representability as cyclic transversal polytopes, called *obstructions*.

3.2 Obstructions for CTPs

Being a 0/1-polytope is a necessary condition for cyclic transversal polytopes by definition. In this section, we explain why not all 0/1-polytopes are affinely isomorphic to cyclic transversal polytopes and give combinatorial necessary conditions for cyclic transversal polytopes. With that, we are able to produce explicit examples and a known class of 0/1-polytopes that are not combinatorially isomorphic to any cyclic transversal polytope, leading us to describe all cyclic transversal polytopes up to dimension 3 with respect to combinatorial isomorphism. The combinatorial necessary conditions are as follows:

3.20 Theorem (Necessary conditions for cyclic transversal polytopes)
Let P be combinatorially isomorphic to a cyclic transversal polytope. Then P is already combinatorially isomorphic to a kernel polytope, or every pair of vertices of P lies in a common proper face of P.

Proof: We have already seen in Theorems 3.7 and 3.8 that kernel polytopes are exactly those cyclic transversal polytopes that have block configurations which consist only of blocks of cardinality 2. Hence, what is left to show is that if P is not combinatorially isomorphic to a kernel polytope, every pair of vertices of P lies in a common proper face of P.

To prove that, we will show the contrapositive: If P has a pair of vertices that does not lie in a common proper face, then P is combinatorially isomorphic to a kernel polytope.

Let $\Pi = (\Xi_i \mid i \in [n])$ be a block configuration over $\mathbb{F}_2^d$ such that P is combinatorially isomorphic to $\mathrm{CTP}(\Pi)$. For ease of the argument, we assume that P is already equal to $\mathrm{CTP}(\Pi)$. Then let v and $\tilde{v}$ be two vertices of P that do not lie in a common proper face, that means the line segment connecting v and $\tilde{v}$ lies in the interior of P.

Let ξ and $\tilde{\xi}$ be the respective cyclic transversals of these vertices. We have to show that there is a block configuration $\widetilde{\Pi}$ that consists only of blocks of the form $\widetilde{\Xi}_i = \{\mathbb{0}, \omega_i\}$ for some $\omega_i \in \mathbb{F}_2^d$, such that the kernel polytope $\mathrm{CTP}(\widetilde{\Pi})$ is also combinatorially isomorphic to P.

Without loss of generality, we assume that Π is pruned (as defined in Definition 2.21 and the subsequent results), since we can readily replace Π with its pruning. This assurance arises from the fact that Π can be pruned, a conclusion derived from the fact that $|\mathrm{CT}(\Pi)|$ is equal to the number of vertices of P, which implies that $|\mathrm{CT}(\Pi)| \geq 2$ holds, making Proposition 2.22 applicable to Π, as otherwise there would be nothing to show for P.

From the fact that v and $\tilde{v}$ are not contained in a common proper face, we deduce that the cyclic transversals ξ and $\tilde{\xi}$ have no common block element $\omega \in \Xi_i$, otherwise the equation $y^i_\omega = 1$ would be valid for both vertices and therefore for P itself, which is a contradiction to Π being pruned. Therefore, every block $\Xi_i \in \Pi$ contains at least the two elements $\xi(i)$ and $\tilde{\xi}(i)$.

Likewise, an equation $y^i_\omega = 0$ for any $i \in [n]$ cannot be valid for both v and $\tilde{v}$, since this would force the equation to be valid for all of P. This is again a contradiction to Π being pruned. It follows that every block of Π is of the form $\Xi_i = \{\xi(i), \tilde{\xi}(i)\}$.

By applying the block translation with $(\xi(1), \dots, \xi(n))$, we equivalently transform this block configuration into $\widetilde{\Pi} = (\widetilde{\Xi}_i \mid i \in [n])$ with $\widetilde{\Xi}_i = \{\mathbb{0}, \xi(i) \oplus \tilde{\xi}(i)\}$, which is of the desired normalized form. It follows that if P is not combinatorially isomorphic to a kernel polytope, then it cannot have a pair of vertices of this kind, i. e., all pairs of vertices need to lie in a common proper face. ■

The two necessary conditions in Theorem 3.20 are not exclusive of each other since the 3-dimensional parity polytope is a kernel polytope as well as a simplex. As such, it satisfies both conditions simultaneously.

Having a proper common face for every pair of vertices is an interesting condition for polytopes in and of itself. Regardless, we think that the answer to the following question is negative:

3.21 Question
Are the conditions in Theorem 3.20 also sufficient? In particular, is every polytope where all pairs of vertices lie in a common proper face combinatorially isomorphic to a cyclic transversal polytope?

Even a general proof of the negative answer will be insightful. On the contrary, the answer is ‘yes’ for polytopes of dimension three, as we will see in Section 3.2.2. But first, with Theorem 3.20 at hand, we turn our attention to polytopes that cannot necessarily be represented as cyclic transversal polytopes.

3.2.1 Cross Polytopes

Usually, the d-dimensional cross polytope is defined as the convex hull of all $\pm\mathbb{e}_i$ for $i \in [d]$. But this does not give us a 0/1-polytope. Instead for $d \geq 3$, we call the polytope $\mathrm{CP}(d)$, defined by

$$\mathrm{CP}(d) := \mathrm{conv}\{\mathbb{e}_1, \dots, \mathbb{e}_d, \mathbb{1} - \mathbb{e}_1, \dots, \mathbb{1} - \mathbb{e}_d\},$$

the d-dimensional *cross polytope*. It has $2d$ vertices for $d \geq 3$ and is centrally symmetric with respect to $\frac{1}{2}\mathbb{1}$. Hence, it is affinely isomorphic to the canonical d-dimensional cross polytope $\operatorname{conv}\{\pm\mathbb{e}_i \mid i \in [d]\}$ for $d \geq 3$, according to Ziegler [53, Section 2.1]. We will prove this small fact explicitly:

3.22 Lemma (Equivalence of Cross Polytope Definitions [cf. 53, Section 2.1 and p. 21]) *Let $\{v_1, \dots, v_d\} \subsetneq \{0,1\}^d$ be $d \geq 3$ affinely independent vectors in a common coordinate hyperplane defined by $\langle \mathbb{e}_i, x\rangle = 0$ and let $\{\mathbb{1} - v_1, \dots, \mathbb{1} - v_d\} \subsetneq \{0,1\}^d$ lie in the parallel hyperplane $\langle \mathbb{e}_i, x\rangle = 1$, both for the same fixed $i \in [d]$. Then the canonical cross polytope $P = \operatorname{conv}\big\{\pm\mathbb{e}_j \bigm| j \in [d]\big\}$ is affinely isomorphic to $Q = \operatorname{conv}\{v_1, \dots, v_d, \mathbb{1} - v_1, \dots, \mathbb{1} - v_d\}$.*

Proof: We give the affine isomorphism $\pi\colon \mathbb{R}^d \to \mathbb{R}^d$ directly. Notice that the center of Q, i. e., the convex sum of all vertices $\frac{1}{2d}\sum_{j\in[d]} v_j + (\mathbb{1} - v_j)$, is the point $\frac{1}{2}\mathbb{1}$, while the center is $\mathbb{0}$ for the canonical cross polytope P. Therefore, the condition $\pi(\mathbb{0}) = \frac{1}{2}\mathbb{1}$ must hold. Furthermore, since an affine map is uniquely determined by its value at $\mathbb{0}$ and by an associated linear map φ via the equation $\pi(x) = \varphi(x) + \pi(\mathbb{0})$, the conditions $\pi(\mathbb{e}_j) = v_j$ determine π uniquely.

The equation $\pi(x) = \varphi(x) + \pi(\mathbb{0})$ implies $\varphi(\mathbb{e}_j) = v_j - \frac{1}{2}\mathbb{1}$ by plugging in the required values for $\mathbb{e}_j$. This in turn implies that φ is invertible: Since $\{v_1, \dots, v_d\}$ are affinely independent and their i-th coordinate is zero, the set $\{v_1, \dots, v_d\} \cup \frac{1}{2}\mathbb{1}$ is also affinely independent. Therefore, the d vertices $\big\{v_j - \frac{1}{2}\mathbb{1} \bigm| j \in [d]\big\}$ form a basis. Since φ is then an automorphism over $\mathbb{R}^d$, the affine map π is also invertible, and $\pi^{-1}(v_j) = \mathbb{e}_j$.

What is left to show is that the invertible affine map π actually is the required isomorphism. It suffices to show that the remaining vertices $-\mathbb{e}_j$ map to $\mathbb{1} - v_j$ for $j \in [d]$, since every convex combination of the vertices of P is a convex combination of the corresponding vertices in Q under π. Thus, the simple calculation

$$\pi(-\mathbb{e}_j) = \varphi(-\mathbb{e}_j) + \pi(\mathbb{0}) = -\varphi(\mathbb{e}_j) + \pi(\mathbb{0}) = -(v_j - \frac{1}{2}\mathbb{1}) + \frac{1}{2}\mathbb{1} = \mathbb{1} - v_j$$

finishes the proof. ■

Note that for $\mathrm{CP}(d)$, the vertices in $\{\mathbb{e}_1, \dots, \mathbb{e}_d\}$ for $d \geq 3$ do not lie in a common coordinate hyperplane as required in Lemma 3.22. Instead, the vectors $\big\{\mathbb{e}_j \bigm| j \in [d] \setminus i\big\} \cup \{\mathbb{1} - \mathbb{e}_i\}$ with fixed $i \in [d]$ lie in the common coordinate hyperplane that is defined by $\langle \mathbb{e}_i, x\rangle = 0$, so Lemma 3.22 still applies to $\mathrm{CP}(d)$, but not in the obvious way.

Ziegler [53, 54, pp. 8–23] discusses other properties of cross polytopes, such as central symmetry and general facial structure. In fact, exercise 0.2 in [54] is the task to prove that every centrally symmetric polytope is the projection of a cross polytope.

Cross polytopes not only contain vertices that do not lie in a common proper face, in fact, their set of vertices can be partitioned into such pairs. Furthermore, these pairs are unique, since every vertex is adjacent to all but one other vertex. For $\mathrm{CP}(d)$ these are the pairs $\{\mathbb{e}_i, \mathbb{1} - \mathbb{e}_i\}$ for $i \in [d]$. The polytopes $\mathrm{CP}(d)$ are therefore an interesting class to obtain results on non-representability as cyclic transversal polytopes.

3.23 Proposition
Let $d \geq 3$ be different from a power of two. Then the d-dimensional cross polytope $\mathrm{CP}(d)$ cannot be combinatorially isomorphic to a cyclic transversal polytope.

Proof: The number of vertices of $\mathrm{CP}(d)$ is $2d$ for $d \geq 3$. Thus, when d is not a power of two, the polytope $\mathrm{CP}(d)$ cannot be represented as a kernel polytope, since the number of vertices of any kernel polytope needs to be of the form 2^k for some $k \in \mathbb{N}$, as we recall from Corollary 3.9. Since $\mathrm{CP}(d)$ also contains pairs of vertices that do not lie in a common proper face such as $\mathbb{e}_i$ and $\mathbb{1} - \mathbb{e}_i$ for any $i \in [d]$, whose midpoint $\frac{1}{2}\mathbb{1}$ lies in the interior of $\mathrm{CP}(d)$, Theorem 3.20 shows that it cannot be combinatorially isomorphic to a cyclic transversal polytope. ■

So far, this leaves open the representability of all polytopes $\mathrm{CP}(2^m)$ for $m \in \mathbb{N}_0$. For these, we will actually give a construction as cyclic transversal polytopes. Therefore, the main result of this section is the following:

3.24 Theorem
The polytope $\mathrm{CP}(d)$ is affinely isomorphic to a cyclic transversal polytope if and only if d is a power of two.

The necessity of d being a power of two is clear from Proposition 3.23, while the rest of this section is devoted to giving a construction to represent $\mathrm{CP}(2^m)$ as a cyclic transversal polytope for all $m \in \mathbb{N}$.

The proof of the cyclic transversal polytope construction for $\mathrm{CP}(2^m)$ will work by induction on $m \in \mathbb{N}_0$. Since $\mathrm{CP}(d)$ is a d-simplex for $d \leq 2$, the cases $m \in \{0, 1\}$ are already proven by Proposition 3.18. For the canonically defined cross polytope $\mathrm{conv}\{\pm\mathbb{e}_i \mid i \in [d]\}$, the cases $d \leq 2$ are different from $\mathrm{CP}(d)$ and reduce to being affinely isomorphic to the cube in $d \leq 2$ dimensions, which we also have proven to be cyclic transversal polytopes in Corollary 3.19. The next case $m = 2$ is the start of our induction and the smallest case where $\mathrm{CP}(2^m)$ is affinely isomorphic to a canonical cross polytope. Throughout this section, we write $\mathbb{1}_{n,m}$ for the matrix of shape $n \times m$ where every entry is equal to one.

3.25 Lemma (Representation of the 4-dimensional cross polytope)
The cross polytope $\mathrm{CP}(4)$ is affinely isomorphic to the parity polytope $\mathrm{PAR}(4)$.

Proof: We will prove that the vertices of PAR(4) can be partitioned into two sets such that their respective convex hulls form 3-simplices in two parallel hyperplanes $x_1 = 0$ and $x_1 = 1$. Since the convex hull of two 3-simplices in parallel hyperplanes with vertices v and $\mathbb{1} - v$ respectively is affinely isomorphic to a 4-dimensional cross polytope by Lemma 3.22, the statement follows.

We define the matrix

$$A_4 := \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 \end{pmatrix},$$

then the columns of the matrix

$$B_4 := \left(A_4 \;\middle|\; \mathbb{1}_{4,4} - A_4 \right)$$

consist of the eight vertices of PAR(4), i. e., all vectors of length four with an even number of entries equal to one. These vertices form a linear subspace over $\mathbb{F}_2$, given as the kernel of the matrix

$$M_4 = \begin{pmatrix} 1 & 1 & 1 & 1 \end{pmatrix}.$$

Furthermore, the columns of A_4 lie in the hyperplane defined by $x_1 = 1$ and are affinely independent: A small calculation shows that the determinant of A_4 is -2, so the columns are linearly independent and lie in a common hyperplane that does not go through the origin. This means that the convex hull of the columns of A_4 form a simplex in the hyperplane $x_1 = 1$.

Consequently, the columns of $\mathbb{1}_{4,4} - A_4$ also lie in a common hyperplane $x_1 = 0$, and since the bijective affine transformation $v \mapsto \mathbb{1} - v$ does not change the affine independence of this set of vectors, these columns also form a simplex in the hyperplane $x_1 = 0$.

Therefore, the convex hull of all eight columns, which is the parity polytope PAR(4), is affinely isomorphic to the cross polytope CP(4), by using Lemma 3.22 twice, together with the fact that affine isomorphy is transitive. ■

We will use the matrices A_4 and $\mathbb{1}_{4,4} - A_4$ to construct vertices of higher-dimensional polytopes which will be affinely isomorphic to $\mathrm{CP}(2^m)$ for $m > 2$. To do this, we need to show that these constructed vertices form two simplices in parallel hyperplanes, so the columns of certain submatrices need to be affinely independent. To show this, we quickly prove the following lemma:

3.26 Lemma (Regularity of $\mathbb{1} - 2A$)
If $A \in \{0,1\}^{k \times k}$ is a regular matrix where one row is equal to $\mathbb{1}_k^\top$, then $\mathbb{1}_{k,k} - 2A$ is also regular.

Proof: Given that A is regular, we want to show that the only solution to the equation $(\mathbb{1}_{k,k} - 2A)c = 0$ with $c \in \mathbb{R}^k$ is $c = \mathbb{0}$.

From $(\mathbb{1}_{k,k} - 2A)c = 0$ follows that $Ac = -\frac{1}{2}\mathbb{1}_{k,k}c$. Since one row of A is equal to $\mathbb{1}_k^\top$, we have the equation $\sum_{i\in[k]} c_i = \sum_{i\in[k]} -\frac{1}{2}c_i$, or equivalently, $\frac{3}{2}\sum_{i\in[k]} c_i = 0$, from which we deduce $\sum_{i\in[k]} c_i = 0$. Then the right-hand side of $Ac = -\frac{1}{2}\mathbb{1}_{k,k}c$ evaluates to $\mathbb{0}$, leaving $Ac = \mathbb{0}$. Because A is regular, the only solution to this equation system is $c = \mathbb{0}$. Therefore, $\mathbb{1}_{k,k} - 2A$ is regular. ■

We also need to ensure that the vertices we construct are the vectors of a linear subspace over $\mathbb{F}_2$, so that the resulting polytope is a kernel polytope.

3.27 Lemma (Linear subspace of complementary vectors)
Let $A \in \{0,1\}^{k\times k}$. If the columns of

$$\widetilde{B} = \left(A \;\middle|\; \mathbb{1}_{k,k} - A\right)$$

are the vectors of a linear subspace in $\mathbb{F}_2^k$, then the columns of

$$B = \left(\begin{array}{c|c|c|c} A & A & \mathbb{1}_{k,k} - A & \mathbb{1}_{k,k} - A \\ \hline \mathbb{1}_{k,k} - A & A & A & \mathbb{1}_{k,k} - A \end{array}\right)$$

are the vectors of a linear subspace in $\mathbb{F}_2^{2k}$.

Proof: We only need to show that the $\mathbb{F}_2$-sum of two columns of the larger matrix B is again a column of B, since scalar multiplication over $\mathbb{F}_2$ is trivial, and because some column of $\widetilde{B}$ is equal to $\mathbb{0}_k$, the vector $\mathbb{0}_{2k}$ will show up as some column in B as well, either in the second or the fourth column block.

Given a column (a_1, a_2) of B, where $a_1, a_2 \in \mathbb{F}_2^k$, we know that $a_2 \in \{a_1, \mathbb{1}_k \oplus a_1\}$. For two such columns (x_1, x_2) and (y_1, y_2), their sum is $(x_1 \oplus y_1, x_2 \oplus y_2)$. Since x_1 and y_1 are columns of $\widetilde{B}$ and the columns of $\widetilde{B}$ are assumed to form a subspace, $x_1 \oplus y_1$ is either a column of A or a column of $\mathbb{1}_{k,k} - A$. Thus, we only need to show that $x_2 \oplus y_2$ is one of $x_1 \oplus y_1$ or $\mathbb{1}_k \oplus x_1 \oplus y_1$.

Now we proceed on a case-by-case basis: If $x_2 = x_1$ and $y_2 = y_1$, then of course $x_2 \oplus y_2$ is equal to $x_1 \oplus y_1$. If either $x_2 = \mathbb{1}_k \oplus x_1$ or $y_2 = \mathbb{1}_k \oplus y_1$, then we get $x_2 \oplus y_2 = \mathbb{1}_k \oplus x_1 \oplus y_1$, and lastly, if both $x_2 = \mathbb{1}_k \oplus x_1$ and $y_2 = \mathbb{1}_k \oplus y_1$, we see that $x_2 \oplus y_2 = \mathbb{1}_k \oplus x_1 \oplus \mathbb{1}_k \oplus y_1$, but since $\mathbb{1}_k \oplus \mathbb{1}_k = \mathbb{0}_k$ over $\mathbb{F}_2$, we again obtain $x_2 \oplus y_2 = x_1 \oplus y_1$.

Therefore, the columns of B are the vectors of a linear subspace in $\mathbb{F}_2^{2k}$. ■

With these ingredients, a proof of Theorem 3.24 comes within our grasp.

Proof of Theorem 3.24: The base case $m = 2$ for our induction is proven by Lemma 3.25, so what is left is the induction step, using Lemmas 3.26 and 3.27.

Let $k = 2^m$, and let

$$B_k := \left(A_k \;\middle|\; \mathbb{1}_{k,k} - A_k \right)$$

be the matrix whose columns are the vertices of a polytope that is affinely isomorphic to the cross polytope $\mathrm{CP}(2^m) = \mathrm{CP}(k)$, such that $A_k \in \{0,1\}^{k \times k}$ is the regular matrix whose columns are those vertices that lie in the hyperplane $x_1 = 1$, meaning that the first row of A_k is equal to $\mathbb{1}_k^\top$.

For the step $m+1$, we define the block matrix

$$A_{2k} := \begin{pmatrix} A_k & A_k \\ A_k & \mathbb{1}_{k,k} - A_k \end{pmatrix},$$

whose first row is equal to $\mathbb{1}_{2k}^\top$, so its columns lie in the hyperplane $x_1 = 1$ as well. By column operations, we transform A_{2k} into the form

$$\begin{pmatrix} A_k & \mathbb{0}_{k,k} \\ A_k & \mathbb{1}_{k,k} - 2A_k \end{pmatrix},$$

where $\mathbb{0}_{k,k}$ is the $k \times k$-matrix with all entries equal to zero. The determinant of this matrix is equal to $\det(A_k) \cdot \det(\mathbb{1}_{k,k} - 2A_k)$. Since A_k is assumed to be regular, $\mathbb{1}_{k,k} - 2A_k$ is regular as well by Lemma 3.26. Therefore, A_{2k} has a determinant not equal to zero, so it is regular as well, and this shows that its columns are affinely independent in their common hyperplane. We conclude that the convex hull of the columns of A_{2k} form a simplex in the hyperplane $x_1 = 1$.

Similarly, the columns of $\mathbb{1}_{2k,2k} - A_{2k}$ are affinely independent and form a simplex in the hyperplane $x_1 = 0$, so the convex hull of the columns of

$$B_{2k} = \left(A_{2k} \;\middle|\; \mathbb{1}_{2k,2k} - A_{2k} \right)$$

are the vertices of a 0/1-polytope which is affinely isomorphic to $\mathrm{CP}(2k) = \mathrm{CP}(2^{m+1})$.

Finally, the polytope constructed this way is a kernel polytope because the columns of B_{2k} are elements of a linear subspace in $\mathbb{F}_2^{2k}$ by Lemma 3.27, with $B = B_{2k}$ and $\widetilde{B} = B_k$. ■

Regardless of the negative result from before, the 3-dimensional cross polytope CP(3), also called the octahedron, can be described in a rather compact way as the projection of a 5-dimensional cyclic transversal polytope, similar to Proposition 3.1:

3.28 Example

Let $\Pi_O = (\Xi_1, \dots, \Xi_4)$ be a block configuration over the venue space $\mathbb{F}_2^2$, and let the blocks be defined as

$$\Xi_0 := \mathbb{F}_2^2 \setminus \mathbb{0}, \quad \Xi_1 := \{\mathbb{0}, \mathbb{e}_1\}, \quad \Xi_2 := \{\mathbb{0}, \mathbb{e}_2\}, \quad \Xi_3 := \{\mathbb{0}, \mathbb{1}\}.$$

Note that this block configuration is pruned, since every block element occurs in some cyclic transversal. The resulting polytope $\mathrm{CTP}(\Pi_O)$ is affinely isomorphic to a 5-dimensional simplex, since its size is 9 and its order is 4, so its dimension is at most 5. Since $\mathrm{CTP}(\Pi_O)$ contains 6 vertices that are determined by the choice of the last three block elements and do not all lie on the same hyperplane, it must be a simplex.

	Ξ_0			Ξ_1		Ξ_2		Ξ_3	
	$\mathbb{e}_1$	$\mathbb{e}_2$	$\mathbb{1}$	$\mathbb{0}$	$\mathbb{e}_1$	$\mathbb{0}$	$\mathbb{e}_2$	$\mathbb{0}$	$\mathbb{1}$
v_1	1	0	0	0	1	1	0	1	0
v_2	1	0	0	1	0	0	1	0	1
v_3	0	1	0	0	1	1	0	0	1
v_4	0	1	0	1	0	0	1	1	0
v_5	0	0	1	0	1	0	1	1	0
v_6	0	0	1	1	0	1	0	0	1

Table 3.4: Vertices of $\mathrm{CTP}(\Pi_O)$ and their cyclic transversals from Example 3.28, with the coordinate projection to the vertices of the octahedron marked

The block Ξ_0 ensures that the block elements $\xi(1), \xi(2), \xi(3)$ cannot all be equal to $\mathbb{0}$ or all simultaneously be different from $\mathbb{0}$. Therefore, the elements $y^j_{\mathbb{0}}$ are not all equal for $j \in [3]$. The affine map $\pi\colon \mathrm{CTP}(\Pi_O) \to [0,1]^3$, defined via $\pi(y)_j = 1 - y^j_{\mathbb{0}}$ for $j \in [3]$, is then a projection onto the octahedron.

Using the naive construction from Proposition 3.1 for the octahedron results in the block configuration defined by

$$\Xi_0 = \mathbb{F}_2^3 \setminus \{\mathbb{0}, \mathbb{1}\}, \quad \Xi_1 := \{\mathbb{0}, \mathbb{e}_1\}, \quad \Xi_2 := \{\mathbb{0}, \mathbb{e}_2\}, \quad \Xi_3 := \{\mathbb{0}, \mathbb{e}_3\}.$$

The image of these blocks under any map $f\colon \mathbb{F}_2^3 \to \mathbb{F}_2^2$ that fulfills $\ker(f) = \{\mathbb{0}, \mathbb{1}\}$ yields a configuration that is equivalent to the one outlined in Example 3.28. One such map is defined by $f(\mathbb{e}_1) = \mathbb{e}_1$, $f(\mathbb{e}_2) = \mathbb{e}_2$ and $f(\mathbb{e}_3) = \mathbb{1}_2$. Note that under our notion of equivalence from Definition 2.3, the naive block configuration and the one in Example 3.28 are neither equivalent since no such f is bijective, nor is the construction in Example 3.28 a reduction of this one, since there is no deletion of fixed blocks or trivial block elements along with equivalence transformations, as only the size of Ξ_0 changes under f and this configuration is already pruned without applying f. The fact that the set of non-vertices of $\mathrm{CP}(3)$ is isomorphic to a linear subspace over $\mathbb{F}_2^3$ allows this construction to be further improved by f while retaining the projection property onto the octahedron.

Incidentally, both the polytopes in Example 3.28 and the preceding construction are simplices. In general, transforming cyclic transversals using such functions f that map blocks to other blocks in a venue space of smaller dimension will result in a *relaxation* of the corresponding cyclic transversal polytope, when canonically lifting the cyclic transversal polytope induced by the transformed blocks back into the affine hull of the original polytope. We refine these definitions and investigate relaxations of cyclic transversal polytopes in Chapter 5.

A 3-SAT Boolean formula that characterizes the octahedron is

$$(x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee \neg x_3),$$

which leads to a cyclic transversal polytope of rank 6, order 5 and size 20 when utilizing the construction from Theorem 3.2.

With that, we conclude our discussion on cross polytopes and their relation to cyclic transversal polytopes by proving several corollaries about cyclic transversal polytopes.

For the first of these corollaries, we define a *vertex figure* ([54, p. 54]): Given a polytope P and a vertex v, let $\langle c, x\rangle \leq c_0$ be a valid inequality for P for which

$$\{v\} = P \cap \{x \mid \langle c, x\rangle = c_0\}$$

holds. Furthermore, choose some $\gamma < c_0$ such that $\langle c, \tilde{v}\rangle < \gamma$ for all vertices $\tilde{v}$ of P other than v. Then, a *vertex figure* of P at v is the polytope $P \cap \{x \mid \langle c, x\rangle = \gamma\}$. Although this construction depends on the choice of γ and of the inequality $\langle c, x\rangle \leq c_0$, the combinatorial type of the vertex figure is independent of these choices [54, Proposition 2.4], i. e., all vertex figures at a fixed vertex are combinatorially isomorphic.

3.29 Corollary
Kernel polytopes and cyclic transversal polytopes are not closed under taking vertex figures.

Proof: Any vertex figure of a d-dimensional cross polytope is (combinatorially isomorphic to) a $(d-1)$-dimensional cross polytope, since a d-dimensional cross polytope is the bipyramid over a $(d-1)$-dimensional cross polytope [54, p. 9]. Therefore, they provide counterexamples. In particular, since the 4-dimensional cross polytope is a kernel polytope (and therefore also a cyclic transversal polytope) and its vertex figures are 3-dimensional cross polytopes, i. e., octahedra, that are neither cyclic transversal nor kernel polytopes, the corollary follows. ■

For the second corollary, the *free sum* (or direct sum) of two polytopes P and Q is as follows ([53, Proposition 10]): Embed P and Q in subspaces such that $\mathbb{0}$ is the only element in $\operatorname{relint}(P) \cap \operatorname{relint}(Q)$, then the *free sum* of P and Q is defined as $\operatorname{conv}(P \cup Q)$. Here $\operatorname{relint}(P)$ is the *relative interior* of P, meaning the topological interior of P within the affine hull of P. As polytopes are convex sets, the relative interior is also formally defined as

$$\mathrm{relint}(P) := \{x \in P \mid \forall y \in P, \text{ there exists some } \lambda > 1 \text{ such that } \lambda x + (1-\lambda)y \in P\}.$$

3.30 Corollary
Kernel polytopes and cyclic transversal polytopes are not closed under taking free sums.

Proof: The free sum of a square and a 1-simplex (both of which are kernel polytopes) is an octahedron [53, p. 16], which is neither a kernel polytope nor a cyclic transversal polytope. This proves the corollary. ■

Having stated these two corollaries, it is simple to ask the following question:

3.31 Question
What is the smallest family of polytopes that contains the cyclic transversal polytopes and is closed under construction of free sums of polytopes or vertex figures, respectively?

While vertex figures and free sums are interesting, but comparatively involved constructions, another natural question is asking the same about the faces or duals of cyclic transversal polytopes. In the case of dual (or polar [54, Section 2.3]) polytopes, there is an immediate answer:

3.32 Corollary
Kernel polytopes and cyclic transversal polytopes are not closed under taking duals of polytopes.

Proof: The dual of a 3-cube (which is a kernel polytope) is an octahedron, which is neither a kernel polytope nor a cyclic transversal polytope. This proves the corollary. ■

For the purpose of understanding the duals of cyclic transversal polytopes, we ask analogously to before:

3.33 Question
What is the smallest family of polytopes that contains the cyclic transversal polytopes and is closed under duality of polytopes?

The analogous question about faces of cyclic transversal polytopes is equally open. With these questions asked, we direct our attention to further families of polytopes which contain members that cannot be combinatorially isomorphic to any cyclic transversal polytope.

3.2.2 Other families of polytopes

In Theorem 3.20 we have seen a necessary condition for a polytope to be combinatorially isomorphic to some cyclic transversal polytope. Based on this condition, we have also proven that there are 3-dimensional polytopes which are not combinatorially isomorphic to any cyclic transversal polytope, e. g., the octahedron. This and previous results actually lead us to a characterization of combinatorial types of cyclic transversal polytopes up to dimension 3. We note that all at most 2-dimensional 0/1-polytopes are combinatorially (and even affinely) isomorphic to a cyclic transversal polytope by Proposition 3.18 and Corollary 3.19. These are the simplices $\Delta_0, \Delta_1, \Delta_2$ and the product $\Delta_1 \times \Delta_1$, which is the square or 2-dimensional cube. In dimension 3, there are 8 different 0/1-polytopes up to combinatorial isomorphism. These are visualized in Figure 3.1.

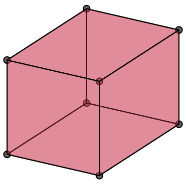

(a) The cube $[0,1]^3$

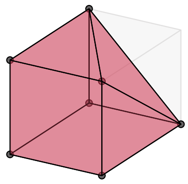

(b) The chipped cube $\operatorname{conv}\{0,1\}^3 \setminus \mathbb{1}$

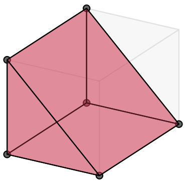

(c) The prism $\Delta_2 \times \Delta_1$

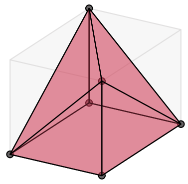

(d) The nameless polytope

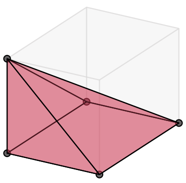

(e) The pyramid from Example 3.17

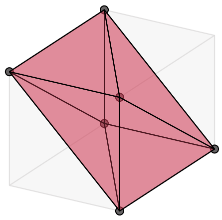

(f) The octahedron CP(3)

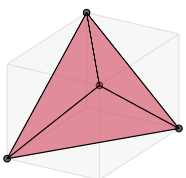

(g) The simplex PAR(3)

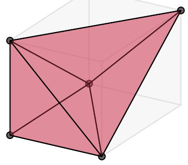

(h) The triangular bipyramid

Figure 3.1: Visualization of 3-dimensional 0/1-polytopes. Polytopes on the left are affinely isomorphic to cyclic transversal polytopes, those on the right are not.

We have not yet proven why the chipped cube, the nameless polytope and the triangular bipyramid are not affinely (or even combinatorially) isomorphic to any cyclic transversal polytope, but the reason again stems from Theorem 3.20: It is easy to see that all of these polytopes have at least one pair of vertices that do not have a common proper face, and all of these polytopes have a number of vertices which is not a power of two. Therefore, they cannot be combinatorially isomorphic to a cyclic transversal polytope. Since any two affinely isomorphic polytopes are also combinatorially isomorphic, they are also not affinely isomorphic to any cyclic transversal polytope.

We note that the octahedron CP(3) is affinely isomorphic to a *hypersimplex*. The hypersimplex $\Delta_{d,k}$ is the polytope given as the convex hull of all 0/1-vectors in $\mathbb{R}^d$ with exactly k entries set to 1. The octahedron is then affinely isomorphic to $\Delta_{4,2}$, which is well-known and can easily be seen from the fact that $\Delta_{4,2}$ is a 3-dimensional polytope with 6 vertices and uniform facets. Therefore, not all hypersimplices are cyclic transversal polytopes.

Furthermore, the octahedron is the *matroid basis polytope* of the uniform matroid with rank 2 on 4 elements. A matroid basis polytope is the convex hull of characteristic vectors of all bases of a matroid, which leads to the same characterization of the octahedron as the hypersimplex $\Delta_{4,2}$.

Matroid basis polytopes are faces of *matroid independence polytopes*, which itself are a special case of polymatroids [23]. A matroid independence polytope is the convex hull of characteristic vectors of independent sets of a matroid. We remark that matroid independence polytopes themselves are not contained in the family of cyclic transversal polytopes. This is evident by the fact that the chipped cube in Figure 3.1(b) is a matroid independence polytope, namely the independence polytope of the uniform matroid U_3^2, the matroid on 3 elements with rank 2.

Since we know of necessary conditions for cyclic transversal polytopes from Theorem 3.20, have discussed cross polytopes as a concrete family of polytopes which cannot be combinatorially isomorphic to cyclic transversal polytopes in general, and have seen some other families of polytopes which are not (combinatorially isomorphic to) cyclic transversal polytopes, we can ask the following rather open-ended question:

3.34 Question
What are other examples of (families of) polytopes that are and that are not representable as (i. e., affinely isomorphic to) cyclic transversal polytopes?

To finish with two concrete example families that are of interest and for which we can give a partial answer to Question 3.34, we revisit some graph theory [cf. 22, pp. 4–14]: Remember that a *spanning tree* of a graph $G = (V, E)$ is a subgraph $H = (W, F)$ such that $W = V$, which explains the *spanning* part, and H is a tree, which means that H is minimally connected, i. e., removal of any edge $e \in F$ would make H disconnected. This especially implies the fact that H does not contain any cycles.

The *spanning tree polytope* $\mathrm{STP}(G)$ of a graph $G = (V, E)$ is the convex hull of the incidence vectors of spanning trees of G. An outer description of the spanning tree polytope of a graph $G = (V, E)$ is given by the inequality system

$$\begin{aligned} &\sum_{e\in E} x_e = n-1, \\ &\sum_{e\in E(X)} x_e \le |X| - 1 \quad \text{for all } \emptyset \ne X \subsetneq V, \text{ and} \\ &\mathbb{0} \le x \le \mathbb{1}, \end{aligned}$$

where $E(X)$ is the set of edges with both endpoints in X. The description by these inequalities and as a convex hull of spanning tree incidence vectors are equivalent [35, Theorem 6.13], which was originally shown in a more general result by Edmonds [23].

Now as a special case, $\mathrm{STP}(n)$ is the spanning tree polytope of $K_n = ([n], E_n = \binom{[n]}{2})$. Note that $\mathrm{STP}(n)$ is trivial for $n \le 2$.

3.35 Example

The graph K_4 can be decomposed into two edge-disjoint trees, as shown by the two trees marked in Figure 3.2. This means that the spanning tree polytope $\mathrm{STP}(4)$ contains vertices that do not have a common face, namely the corresponding vertices of these spanning trees: They do not lie in a common face, since they do not fulfill any of the equations of the form $x_e = 0$ or $x_e = 1$ for any edge $e \in E_4$ as they cover all edges of K_4 and are disjoint to one another, nor do they fulfill any inequality $\sum_{e\in E(X)} x_e \le |X| - 1$ with equality simultaneously, since every subset $X \subsetneq V$ of size 2 only has one edge which is only used by one of the trees, so the left-hand side of the inequality equates to 1 for this tree and 0 for the other, and every such subset X of size 3 induces a subgraph that has exactly two edges in one tree and one in the other, again resulting in different values for the left-hand side of the inequality. Another way to see this is that the midpoint of both corresponding vertices is $\frac{1}{2}\mathbb{1}$, which lies in the (relative) interior of $\mathrm{STP}(4)$. This implies that if $\mathrm{STP}(4)$ is combinatorially isomorphic to a cyclic transversal polytope, it necessarily has to be a kernel polytope, by Theorem 3.20.

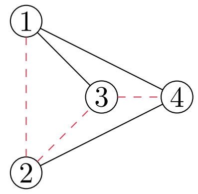

Figure 3.2: The graph K_4, partitioned into two marked edge-disjoint trees

Note that for $n \geq 5$, two edge-disjoint spanning trees cannot cover all edges of K_n, since the number of edges $\binom{n}{2} = n(n-1)/2$ grows faster than the number of edges in two edge-disjoint spanning trees, which is $2(n-1)$. Therefore, at least one edge $e \in E_n$ is unused by both spanning trees, and the equation $x_e = 0$ generates a common face of the vertices in $\mathrm{STP}(n)$ corresponding to these spanning trees.

The fact that K_4 is the only complete graph that contains two edge-disjoint trees that cover all its edges is a consequence of the fact that the path on 4 vertices is the only (non-trivial) tree whose complement is also a tree: For the complement of a tree on n vertices to be a tree, we need to fulfill the equation $\binom{n}{2} = 2(n-1)$ since both the tree and its complement contain exactly $n-1$ edges and both trees together cover the complete graph. This gives us $4(n-1) = n(n-1)$, whose only solutions are $n = 1$, resulting in the trivial tree with only one node, and the other solution being $n = 4$. Exhaustion of all possible trees on 4 nodes gives the desired result: A node of degree 3 would result in a triangle in the complement, which is forbidden in a tree, so a path on 4 nodes is the only possible (non-trivial) option.

Note that it is a well-known result that the number of vertices of $\mathrm{STP}(n)$ is n^{n-2}, since it is equal to the number of distinct labelled trees on n nodes [cf. 50, Proposition 5.3.2]. This number is a power of two if and only if n itself is. Therefore, $\mathrm{STP}(n)$ cannot be a kernel polytope except for the cases where n is a power of two. This also means that $\mathrm{STP}(4)$ could still be combinatorially isomorphic to a kernel polytope since it has 16 vertices, but if it is not, then together with Example 3.35 we know that it is not even a cyclic transversal polytope.

This discussion prompts the following questions, as more explicit open problems than Question 3.34:

3.36 Question

Is $\mathrm{STP}(4)$ a kernel polytope? More general, are $\mathrm{STP}(n)$ representable as cyclic transversal polytopes, or at least when n is a power of two? What about other spanning tree polytopes $\mathrm{STP}(G)$, for general graphs G?

Lastly, another well-investigated family of polytopes are the so-called traveling salesman (or salesperson) polytopes [cf. 15, section 7.4]. For a given graph $G = (V, E)$, a cycle that goes through every node exactly once is called *Hamiltonian cycle*. The (symmetric) *traveling salesman polytope* $\mathrm{TSP}(G)$ then is the convex hull of all such Hamiltonian cycles. In the case of directed graphs, we use directed cycles and the problem and polytope are called asymmetric. Analogously to before, let $\mathrm{TSP}(n)$ be the traveling salesman polytope of the complete graph K_n.

An explicit description of $\mathrm{TSP}(n)$ by linear inequalities is not known (which is expected [cf. 48, Corollary 5.16a]), but several integer formulations have been introduced in the literature [19, 39]. The formulation by Dantzig, Fulkerson, and Johnson [19] for graphs $G = (V, E)$ is

$$\mathrm{TSP}(G) := \operatorname{conv}\left\{ x \in \{0,1\}^E \;\middle|\; \begin{array}{l} \sum_{e \in \delta(v)} x_e = 2 \text{ for all } v \in V, \\ \sum_{e \in \delta(S)} x_e \geq 2 \text{ for all } \emptyset \subsetneq S \subsetneq V \end{array} \right\},$$

where $\delta(S) \subseteq E$ are the edges that have exactly one vertex in $S \subseteq V$.

3.37 Example
The polytope $\mathrm{TSP}(5)$ is equal to

$$\sum_{e \in \delta(v)} x_e = 2 \text{ for all } v \in [5] \text{ and } \mathbb{0} \leq x \leq \mathbb{1},$$

since this description yields an integral polytope and the degree inequalities $\sum_{e \in \delta(S)} x_e \geq 2$ for $S \subsetneq V$ with $|S| = 2$ are directly implied by the degree constraints for singleton vertices: There is at most one edge between these two vertices in the Hamiltonian cycle, so at least two edges have to connect between S and $V \setminus S$ to fulfill the degree condition. For $|S| \geq 3$, one replaces S with $V \setminus S$ to infer the remaining constraints.

By enumeration, $\mathrm{TSP}(5)$ has 12 vertices, which is not a power of two, so $\mathrm{TSP}(5)$ cannot be combinatorially isomorphic to a kernel polytope. But even more is true: Figure 3.3 shows an example of how all edges of K_5 are covered by two edge-disjoint Hamiltonian cycles. Therefore, the two vertices corresponding to these cycles are not contained in a common face of $\mathrm{TSP}(5)$. This can also be proven by observing that the midpoint of both vertices is $\frac{1}{2}\mathbb{1}$, which lies in the (relative) interior of $\mathrm{TSP}(5)$, or by investigating all inequalities of the above description, like with the spanning tree polytopes from before. It follows that $\mathrm{TSP}(5)$ is not combinatorially isomorphic to any cyclic transversal polytope, by Theorem 3.20.

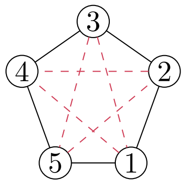

Figure 3.3: The graph K_5, partitioned into two marked edge-disjoint Hamiltonian cycles that cover all edges

The number of vertices of TSP(n) is $(n-1)!/2$ by counting the number of Hamiltonian cycles of K_n as the number $n!$ of permutations of n nodes, selecting one of the n nodes as a starting point and one of two directions for the cycle: $n!/(2n) = (n-1)!/2$. This number is not a power of two except for $n \leq 3$.

We remark that parameters $n \leq 4$ are insignificant because TSP(4) only has 3 vertices. Ergo, it necessarily is affinely isomorphic to a simplex and as such is isomorphic to a cyclic transversal polytope. Trivially, as is TSP(3) with only 1 vertex, and smaller traveling salesman polytopes are empty because of degree constraints.

Consequently, the number of vertices forbids the construction of TSP(n) as a kernel polytope for $n \geq 4$, but since $\binom{n}{2}$ grows faster than $2n$, the union of two Hamiltonian cycles with n edges leaves at least one edge e of K_n uncovered for $n \geq 6$. Therefore, the equation $x_e = 0$ generates a common face of TSP(n) which contains both vertices corresponding to these cycles, similar to the spanning tree polytope from before. Hence, a construction of TSP(n) as a cyclic transversal polytope cannot be ruled out yet.

Nevertheless, we suspect the answer for the following closing question to be negative:

3.38 Question
Can any TSP(n) be represented as a cyclic transversal polytope for $n \geq 6$?

4 Full Cyclic Transversal Polytopes

By choosing the blocks of a cyclic transversal polytope to be equal to the venue space $\mathbb{F}_2^d$, one obtains a subclass of cyclic transversal polytopes which we call *full* or *full-block* cyclic transversal polytopes, just like the associated block configuration. Since the choice of blocks is prescribed, full cyclic transversal polytopes are governed only by the rank d and the order n, that is

$$\mathrm{CTP}(d,n) := \mathrm{CTP}(\underbrace{\mathbb{F}_2^d, \dots, \mathbb{F}_2^d}_{n \text{ times}}) \subseteq \left(\mathbb{R}^{\mathbb{F}_2^d}\right)^n.$$

We define the set of *full* cyclic transversals $\mathrm{CT}(d,n)$ analogously, and overload the notation $\mathbb{A}(d,n)$ accordingly for the affine space of corresponding block equations. One readily sees that every cyclic transversal polytope is a face of a full cyclic transversal polytope via restriction of its incidence vector coordinates, making full cyclic transversal polytopes and its faces useful instruments to study smaller cyclic transversal polytopes. One of the advantages of full cyclic transversal polytopes over other cyclic transversal polytopes is their symmetry, which will be exploited to find large classes of valid inequalities and even some complete descriptions in Section 4.1.

Note that the full cyclic transversal polytopes $\mathrm{CTP}(0,n) = \mathbb{1}_n \in \mathbb{R}^n$ are singletons for all $n \in \mathbb{N}$, which is why we exclude the case $d = 0$ in this chapter, although many of the results presented in this chapter can be generalized to this edge case.

The number of vertices of full cyclic transversal polytopes is derived by counting:

4.1 Proposition
For all parameters $d, n \in \mathbb{N}$, the number of vertices of $\mathrm{CTP}(d,n)$ is $2^{d(n-1)} = \left(2^d\right)^{n-1}$.

Proof: The vertices of $\mathrm{CTP}(d,n)$ are uniquely determined by choosing n vectors $\xi(1), \dots, \xi(n)$ in $\mathbb{F}_2^d$ that sum to $\mathbb{0}$. Since choosing $\xi(1), \dots, \xi(n-1) \in \mathbb{F}_2^d$ arbitrarily and then requiring $\xi(n) = \sum_{i \in [n-1]} \xi(i)$ is equivalent, the number of vertices follows. ■

The size of full cyclic transversal polytopes is $\sum_{i \in [n]} \left|\mathbb{F}_2^d\right| = n2^d$. Aside from the equations implied by the n blocks and the degenerated cases $n \leq 2$, there are no other valid equations for full cyclic transversal polytopes:

4.2 Proposition
For all parameters $d, n \in \mathbb{N}$ with $n \geq 3$,

$$\operatorname{aff}\mathrm{CTP}(d,n) = \left\{ y \in \left(\mathbb{R}^{\mathbb{F}_2^d}\right)^n \;\middle|\; \sum_{\omega \in \mathbb{F}_2^d} y^i_\omega = 1 \textit{ for all } i \in [n] \right\} = \mathbb{A}(d,n),$$

and therefore

$$\dim \mathrm{CTP}(d,n) = \dim \operatorname{aff}\mathrm{CTP}(d,n) = n(2^d - 1).$$

Proof by blockwise constant valid equations: Let $\langle a, y\rangle = \beta$ be a valid equation for $\mathrm{CTP}(d,n)$. By adding appropriate multiples of equations from the equation system given in the proposition statement, we assume that $a^i_{\mathbb{0}} = 0$ for all $i \in [n]$ and only need to show $a = \mathbb{0}$.

For this reason, consider two vertices $v_0, \hat{v} \in \mathrm{CTP}(d,n)$, where v_0 is the vertex belonging to the cyclic transversal $(\mathbb{0}, \dots, \mathbb{0})$ and $\hat{v}$ is a vertex for which there exist exactly two blocks $i_1, i_2 \in [n]$ and a vector $\omega \in \mathbb{F}_2^d \setminus \mathbb{0}$ such that it corresponds to the cyclic transversal $(\xi(1), \dots, \xi(n))$ with $\xi(i_1) = \xi(i_2) = \omega$ and $\xi(k) = \mathbb{0}$ otherwise.

Since the equation $\langle a, y\rangle = \beta$ is valid for both vertices, it follows that

$$\langle a, v_0\rangle - \langle a, \hat{v}\rangle = \beta - \beta = 0,$$

which, together with our initial assumption $a^i_{\mathbb{0}} = 0$ for all $i \in [n]$, implies that $a^{i_1}_\omega + a^{i_2}_\omega = 0$.

Now for three distinct indices $i_1, i_2, i_3 \in [n]$ (since $n \geq 3$) and arbitrary $\omega \in \mathbb{F}_2^d \setminus \mathbb{0}$, we construct the following system of equations with the same procedure:

$$\begin{array}{ccccccc} a^{i_1}_\omega & + & a^{i_2}_\omega & & & = 0 \\ & & a^{i_2}_\omega & + & a^{i_3}_\omega & = 0 \\ a^{i_1}_\omega & & & + & a^{i_3}_\omega & = 0 \end{array}$$

The only solution in a for this system of equations is $a^{i_1}_\omega = a^{i_2}_\omega = a^{i_3}_\omega = 0$, since the matrix of coefficients has full rank. ■

It is not hard to prove that for $n \geq 3$, the construction that is outlined in the proof above yields a suitable family of affinely independent vertices to certify that $\mathrm{CTP}(d,n)$ is of dimension $n(2^d - 1)$: For $\omega \in \mathbb{F}_2^d$ and $i, j \in [n]$ with $i \neq j$ let $v_{i,j}(\omega)$ be the vertex corresponding to the cyclic transversal $(\xi(1), \dots, \xi(n))$ with $\xi(i) = \xi(j) = \omega$ and $\xi(k) = \mathbb{0}$ otherwise. Then the set

$$v_{1,2}(\mathbb{0}) \cup \left\{ v_{1,j}(\omega) \;\middle|\; \omega \in \mathbb{F}_2^d \setminus \mathbb{0}, j \in [n] \setminus 1 \right\} \cup \left\{ v_{2,3}(\omega) \;\middle|\; \omega \in \mathbb{F}_2^d \setminus \mathbb{0} \right\}$$

is an affinely independent subset of vertices of size $n(2^d - 1) + 1$. Therefore, the affine space containing them has dimension at least $n(2^d - 1)$, while the upper bound is given by the block equations.

After characterizing the valid equations for $\mathrm{CTP}(d, n)$, we now turn to valid inequalities $\langle a, y \rangle \geq \beta$, where we assume $a \in \left(\mathbb{Z}^{\mathbb{F}_2^d}\right)^n$ and $\beta \in \mathbb{Z}$, since $\mathrm{CTP}(d, n)$ is a 0/1-polytope [cf. 53, Theorem 5]. We restrict the inequalities further to coprime entries in (a, β) by dividing by their greatest common divisor. Without loss of generality, by adding appropriate multiples of the equations in Proposition 4.2 to a valid inequality $\langle a, y \rangle \geq \beta$, we also assume that $a \geq \mathbb{0}$ and that for every $i \in [n]$, there is at least one $\omega \in \mathbb{F}_2^d$ for which $a^i_\omega = 0$, that is,

$$\min\{a^i_\omega \mid \omega \in \mathbb{F}_2^d\} = 0 \text{ for all } i \in [n].$$

Vectors and inequalities of this form are called *normalized*.

An additional important property of valid inequalities becomes apparent when looking at the symmetries of $\mathrm{CTP}(d, n)$: If one has a valid inequality for $\mathrm{CTP}(d, n)$, the coefficients can by definition be partitioned into n blocks of size 2^d, each block corresponding to the elements in $\mathbb{F}_2^d$. Another inequality is then potentially obtained by permuting block indices and/or applying an equivalence transformation (as given in Definition 2.3) on the block elements, and this other inequality necessarily has to be valid as well: Since full cyclic transversal polytopes are symmetric with respect to the resulting coordinate permutations obtained by utilizing Proposition 2.7, the set of all valid inequalities needs to be invariant with respect to these operations. More precise, the coordinate permutations induced by these operations form a subgroup of the symmetries of $\mathrm{CTP}(d, n)$, and therefore these specific coordinate permutations are automorphisms of $\mathrm{CTP}(d, n)$. Let us denote by $\Gamma(d, n)$ the group of automorphisms of $\mathrm{CTP}(d, n)$ and let us call the specific subgroup of automorphisms $\Upsilon(d, n)$, which is also obviously a subgroup of all $n2^d!$ coordinate permutations of the ambient space.

Both groups of automorphisms *act* on the set of all valid inequalities and partition this set into *orbits*. An orbit of an inequality is the set of all inequalities that are obtained from each other by utilizing the aforementioned operations, just like groups acting on other sets form orbits. In order to obtain well-defined orbits, we assume all inequalities to be normalized. A classification of these normalized inequality orbits with respect to $\Upsilon(d, n)$ for small parameters $d \leq 3$ is found in Section 4.1, where we will present useful visualizations of valid inequality orbits.

Notably, the same orbit classification can be done to the vertices of $\mathrm{CTP}(d,n)$, but it is rather straightforward: All vertices lie in the same orbit with respect to the group of automorphisms generated by equivalence transformations and permuting block indices. In other words, for any two vertices $u, v \in \mathrm{CTP}(d,n)$, there is an automorphism $f_{u,v}\colon \mathrm{CTP}(d,n) \to \mathrm{CTP}(d,n)$ that maps u to v. If that is the case for any polytope P, we say that P is *vertex-transitive*.

4.3 Proposition
The polytope $\mathrm{CTP}(d,n)$ *is vertex-transitive.*

Proof: Let $u, v \in \mathrm{CTP}(d,n)$ be two vertices and $\xi, \tilde{\xi}$ be their associated cyclic transversals. Then the coordinate permutation induced by the block translation with the cyclic transversal $\xi \oplus \tilde{\xi}$ is an automorphism of $\mathrm{CTP}(d,n)$ by Proposition 2.7 and maps u to v as well as v to u. ■

Note that we have not proven that the coordinate permutations given by permuting block indices and/or applying an equivalence transformation on the block elements form the full group of automorphisms of $\mathrm{CTP}(d,n)$, yet coordinate permutations from block translations already show that $\mathrm{CTP}(d,n)$ is vertex-transitive.

A first orbit of inequalities to study for $\mathrm{CTP}(d,n)$ are the obviously valid non-negativity constraints $y^i_\omega \geq 0$ for any $\omega \in \mathbb{F}_2^d$ and $i \in [n]$. All inequalities of this form make up a single orbit under $\Upsilon(d,n)$, as ω can be transformed into any other element of $\mathbb{F}_2^d$ by an equivalence transformation, and any block permutation involving i changes the inequality to another block index. Examining the proof of Proposition 4.2 more closely, we give a first result for this inequality orbit given by the non-negativity conditions:

4.4 Corollary (Non-negativity facets of $\mathrm{CTP}(d,n)$)
For all parameters $d, n \in \mathbb{N}$ *with* $n \geq 3$*, except for* $(d,n) = (1,3)$*, the set*

$$\mathrm{F}_{(d,n)}(\omega, i) := \mathrm{CTP}(d,n) \cap \{y \in \mathbb{A}(d,n) \mid y^i_\omega = 0\}$$

is a facet of $\mathrm{CTP}(d,n)$*, induced by the non-negativity constraint described by* ω *and* i*.*

Proof: Without loss of generality, we assume that the facet is given by the non-negativity constraint with $\omega = \mathbb{1}$ and $i = 1$ by applying an appropriate automorphism on $\mathrm{CTP}(d,n)$. Then, following the structure of the proof of Proposition 4.2, we also assume that any equation $\langle a, y \rangle = \beta$ that is valid for $\mathrm{F}_{(d,n)}(\mathbb{1}, 1)$ already fulfills $a^i_{\mathbb{0}} = 0$ for all $i \in [n]$ by adding an appropriate multiple of the block equations. Then we follow the rest of the proof of Proposition 4.2 for any $\omega \in \mathbb{F}_2^d \setminus \{\mathbb{0}, \mathbb{1}\}$ and any $i \in [n]$ to obtain the equations $a^i_\omega = 0$ as well, since all vertices involved in these cases are necessarily also vertices of

$\mathrm{F}_{(d,n)}(\mathbb{1},1)$. What is left to show is that the coefficient vector a of any such equation also fulfills $a^i_{\mathbb{1}} = 0$ for all $i \in [n] \setminus 1$, i. e., the valid equation $\langle a, y\rangle = \beta$ for $\mathrm{F}_{(d,n)}(\mathbb{1},1)$ then has to be a multiple of $y^i_\omega = 0$, since evaluating $\langle a, y\rangle = \beta$ at any point of $\mathrm{F}_{(d,n)}(\mathbb{1},1)$ then implies $\beta = 0$.

If $d = 1$ then we assume $n \geq 4$, which means that there are at least 3 full blocks with indices $[n] \setminus 1$ for which we repeat the steps in the proof of Proposition 4.2 to show $a^i_{\mathbb{1}} = 0$ for all $i \in [n] \setminus 1$.

For $d \geq 2$, the condition $a^i_{\mathbb{1}} = 0$ for all $i \in [n] \setminus 1$ is achieved by observing that for any $\omega_0 \in \mathbb{F}_2^d$, there are $\omega_1, \omega_2 \in \mathbb{F}_2^d \setminus \omega_0$ such that $\omega_0 = \omega_1 + \omega_2$. Taking $\omega_0 = \mathbb{1}$, the set $\mathrm{F}_{(d,n)}(\mathbb{1},1)$ contains vertices $v_{i,j}$ for $i, j \in [n] \setminus 1$ that correspond to cyclic transversals $(\xi(1), \dots, \xi(n))$ with $\xi(1) = \omega_1$, $\xi(i) = \mathbb{1}$, $\xi(j) = \omega_2$ and $\xi(k) = \mathbb{0}$ otherwise. With v_0 being the vertex belonging to the cyclic transversal $(\mathbb{0}, \dots, \mathbb{0})$, we calculate the differences

$$0 = \left\langle a, v_{i,j} \right\rangle - \langle a, v_0\rangle,$$

which evaluate to

$$a^1_{\omega_1} + a^i_{\mathbb{1}} + a^j_{\omega_2} = 0,$$

as $a^k_{\mathbb{0}} = 0$ is fulfilled for all $k \in [n]$. Since also $a^1_{\omega_1} = a^i_{\omega_2} = 0$ for any $\omega_1, \omega_2 \in \mathbb{F}_2^d \setminus \mathbb{1}$ has been shown in the previous step, this implies the equation $a^i_{\mathbb{1}} = 0$ for all $i \in [n] \setminus 1$. ■

The fact that the non-negativity constraints do not generate facets for smaller parameters will become clear in Section 4.1. We now describe another important orbit of normalized inequalities for $\mathrm{CTP}(d, n)$ and provide a proof of their validity:

4.5 Lemma (Odd-Hyperplane inequalities)
For all parameters $d, n \in \mathbb{N}$, any subspace $H \subseteq \mathbb{F}_2^d$ with $\dim H = d-1$ and any set $I \subseteq [n]$ with $|I|$ odd, the inequality

$$\sum_{i\in I} \sum_{\omega\in H} y^i_\omega + \sum_{i\notin I} \sum_{\omega\notin H} y^i_\omega \geq 1 \qquad [H, I]$$

is valid for $\mathrm{CTP}(d, n)$.

Proof: Let $h \in \mathbb{F}_2^d$ be a vector satisfying $H = \left\{\omega \in \mathbb{F}_2^d \;\middle|\; \langle h, \omega\rangle = 0\right\}$. Furthermore, let $(\xi(1), \dots, \xi(n)) \in \mathrm{CT}(d, n)$ be a cyclic transversal and $v \in \mathrm{CTP}(d, n)$ be its associated vertex. Assume for the sake of contradiction that v violates the inequality $[H, I]$, namely,

$$\sum_{i\in I} \sum_{\omega\in H} v^i_\omega + \sum_{i\notin I} \sum_{\omega\notin H} v^i_\omega = 0$$

holds. Thus, we have

$$\xi(i) \notin H \text{ for all } i \in I \text{ and}$$
$$\xi(i) \in H \text{ for all } i \notin I,$$

hence

$$\begin{aligned} 0 = \langle h, \mathbb{0} \rangle = \left\langle h, \sum_i \xi(i) \right\rangle = \sum_i \langle h, \xi(i) \rangle \\ = \sum_{i \in I} 1 \oplus \sum_{i \notin I} 0 \underset{|I| \text{ odd}}{=} 1, \end{aligned}$$

which is an obvious contradiction. ■

It is easy to see that the odd-hyperplane inequalities form a single orbit under the aforementioned automorphisms of $\mathrm{CTP}(d, n)$: For any pair of hyperplanes, there is an automorphism that maps one hyperplane to the other, for two index sets of the same odd cardinality, there is a permutation of block indices that transforms one index set to the other, and there are appropriate cyclic transversals to apply as block translations that increase or decrease the cardinality of an index set I by an even number.

If $\Pi = (\Xi_1, \dots, \Xi_n)$ is a block configuration where blocks are not necessarily the whole venue space, the odd-hyperplane inequalities for $\mathrm{CTP}(\Pi)$ are of the form

$$\sum_{i \in I} \sum_{\omega \in H \cap \Xi_i} y^i_\omega + \sum_{i \notin I} \sum_{\omega \in \Xi_i \setminus H} y^i_\omega \geq 1,$$

and they are also valid for $\mathrm{CTP}(\Pi)$, which is shown by means of substitution: Since $\mathrm{CTP}(\Pi)$ is isomorphic to a face of some full cyclic transversal polytope $\mathrm{CTP}(d, n)$, obtained by adding equations of the form $y^i_\omega = 0$ for all $i \in [n]$ and $\omega \in \mathbb{F}_2^d \setminus \Xi_i$, the modified odd-hyperplane inequalities are given by replacing these variables y^i_ω with 0.

Colloquially, these inequalities describe that over $\mathbb{F}_2^d$, one cannot sum an odd number of vectors in an affinely translated hyperplane and get the zero vector as a result, since at least the non-zero translation vector will be summed an odd number of times.

For small parameters $d \leq 2$, the odd-hyperplane inequalities already give enough information to describe full cyclic transversal polytopes completely, which we will prove in Section 4.1. Since other cyclic transversal polytopes are faces of full cyclic transversal polytopes where some coordinates are set to zero, these inequalities suffice to describe them as well, if the venue space has at most dimension 2.

Before that, we turn our attention to a closer investigation of odd-hyperplane inequalities for all full cyclic transversal polytopes. First, we show that their inclusion is necessary for the description of full cyclic transversal polytopes with large enough parameters, as they actually generate facets:

4.6 Theorem (Odd-Hyperplane Facets of $\mathrm{CTP}(d,n)$)
For all parameters $d,n \in \mathbb{N}$ with $n \geq 3$, every hyperplane $H \subseteq \mathbb{F}_2^d$ and any set $I \subseteq [n]$ with $|I|$ odd, the set

$$\mathrm{F}_{(d,n)}(H,I) \coloneqq \mathrm{CTP}(d,n) \cap \left\{ y \in \mathbb{A}(d,n) \;\middle|\; \sum_{i \in I} \sum_{\omega \in H} y_\omega^i + \sum_{i \notin I} \sum_{\omega \notin H} y_\omega^i = 1 \right\}$$

is a facet of $\mathrm{CTP}(d,n)$*, induced by the odd-hyperplane inequality described by H and I.*

Proof: Let $H \subseteq \mathbb{F}_2^d$ and $I \subseteq [n]$ with $|I|$ odd be given. To show that $\mathrm{F}_{(d,n)}(H,I)$ is a facet, suppose that some equation $\langle a, y \rangle = \beta$ defines a hyperplane of the affine hull $\mathbb{A}(d,n)$ of $\mathrm{CTP}(d,n)$ and contains $\mathrm{F}_{(d,n)}(H,I)$, i.e., the equation is valid for all vertices of $\mathrm{F}_{(d,n)}(H,I)$, but not for all of $\mathbb{A}(d,n)$. We will show that up to a linear combination of the block equations and scaling, the vector a of the equation $\langle a, y \rangle = \beta$ is already equal to the coefficient vector of the odd-hyperplane constraint

$$\sum_{i \in I} \sum_{\omega \in H} y_\omega^i + \sum_{i \notin I} \sum_{\omega \notin H} y_\omega^i = 1$$

defining $\mathrm{F}_{(d,n)}(H,I)$, which then also implies $\beta = 1$.

We will first define some vertices of $\mathrm{F}_{(d,n)}(H,I)$ that will be used throughout this proof. For this purpose, we introduce a notation to represent a cyclic transversal associated with a vertex in a compact way: Given some indices $i_1, \ldots, i_s \in I$ and $k_1, \ldots, k_t \in [n] \setminus I$ as well as some vectors $\theta_0, \theta_1, \ldots, \theta_s, \varrho_0, \varrho_1, \ldots, \varrho_t \in \mathbb{F}_2^d$, the transversal $\xi = (\xi(1), \ldots, \xi(n))$ over $\mathbb{F}_2^d$ with $\xi(i_r) = \theta_r$ for $r \in [s]$ and $\xi(i) = \theta_0$ for $i \in I \setminus \{i_1, \ldots, i_s\}$, as well as $\xi(k_r) = \varrho_r$ for $r \in [t]$ and $\xi(k) = \varrho_0$ for $k \in [n] \setminus (I \cup \{k_1, \ldots, k_t\})$ is represented in a tabular format as

I				$[n] \setminus I$			
$*$	i_1	$\ldots$	i_s	$*$	k_1	$\ldots$	k_t
θ_0	θ_1	$\ldots$	θ_s	ϱ_0	ϱ_1	$\ldots$	ϱ_t

,

where we note that $[n] \setminus I$ could be empty. Using this notation, it is straightforward to check whether a transversal is cyclic, as one can read off the block elements accordingly to sum them up, using the fact that an even number of occurrences evaluates to $\mathbb{0} \in \mathbb{F}_2^d$. In order for the vertex of $\mathrm{CTP}(d,n)$ that is associated to some cyclic transversal to belong to $\mathrm{F}_{(d,n)}(H,I)$, it needs to fulfill the equation

$$|\{i \in I \mid \xi(i) \in H\}| + \left|\{k \in [n] \setminus I \;\middle|\; \xi(k) \in \mathbb{F}_2^d \setminus H\}\right| = 1,$$

that is, there either is exactly one $q \in I$ such that the element $\xi(q) \in \mathbb{F}_2^d$ belongs to H, or exactly one $q \in [n] \setminus I$ such that $\xi(q) \in \mathbb{F}_2^d \setminus H$. When representing a cyclic transversal in the table notation, we mark this respective unique index with a circle around it.

In the following, we use $\omega \in \mathbb{F}_2^d$ to denote an arbitrary element of H and $\sigma \in \mathbb{F}_2^d$ to represent an arbitrary element in the complement $\mathbb{F}_2^d \setminus H$. Likewise, $i, j \in [n]$ denote arbitrary indices in I, while $k, \ell \in [n]$ denote arbitrary indices in its complement $[n] \setminus I$.

Throughout the proof, we will use six different types of vertices of $\mathrm{F}_{(d,n)}(H, I)$, two of which are derived from the other four. We mention here that we do not claim these vertex types to be distinct, since for particular choices of indices and block elements they may coincide, and there may also be vertices contained in $\mathrm{F}_{(d,n)}(H, I)$ which do not fit any of these types. One of the vertex types we introduce will only be used for the first step of the proof, which we call *relative complementarity*. The three types of vertices that will also be relevant for other parts of the proof are

$$u(k,\ell,\omega,\sigma) \quad := \quad \begin{array}{c|ccc} I & & [n]\setminus I & \\ \hline * & k & \ell & * \\ \hline \sigma\oplus\omega & \sigma & \omega & \mathbb{0} \end{array} \quad \text{for } k \neq \ell,$$

$$v(i,k,\ell,\omega,\sigma) \quad := \quad \begin{array}{cc|ccc} I & & & [n]\setminus I & \\ \hline i & * & k & \ell & * \\ \hline \mathbb{0} & \sigma & \omega & \omega & \mathbb{0} \end{array} \quad \text{for } k \neq \ell, \text{ and}$$

$$w(i,j,\omega,\sigma) \quad := \quad \begin{array}{ccc|c} & I & & [n]\setminus I \\ \hline i & j & * & * \\ \hline \omega & \sigma\oplus\omega & \sigma & \mathbb{0} \end{array} \quad \text{for } i \neq j.$$

Note that vertices of type $w(i, j, \omega, \sigma)$ only exist when $|I| \geq 2$. Since $|I|$ is always odd, this can be strengthened to $|I| \geq 3$. The other two types of vertices, $u(k, \ell, \omega, \sigma)$ and $v(i, k, \ell, \omega, \sigma)$, exist when $|[n] \setminus I| \geq 2$. This will lead to a case distinction in parts of the remaining proof because at least one of these cases is always fulfilled with $|I|$ odd and $n \geq 3$, but not necessarily both.

By observation of these tables we convince ourselves that vertices of these types really do belong to $\mathrm{F}_{(d,n)}(H, I)$ since their associated transversals are cyclic and therefore also satisfy the constraint $\langle a, y\rangle = \beta$. As an example, the elements of the transversals associated to vertices of type $w(i, j, \omega, \sigma)$ sum to

$$\omega \oplus (\sigma \oplus \omega) \oplus \sum_{I\setminus\{i,j\}} \sigma = \underbrace{\omega \oplus \omega}_{=\,\mathbb{0}} \oplus \underbrace{\sum_{|I|-1}}_{=\,\mathbb{0}} \sigma = \mathbb{0},$$

because $|I| - 1$ is even, as I has odd cardinality by assumption. The other vertex types can be checked analogously.

The one vertex type we only use for relative complementarity that does not fit the previous types is

$$z(i,k,\omega,\sigma) \quad := \quad \begin{array}{cc|cc} \multicolumn{2}{c|}{I} & \multicolumn{2}{c}{[n]\setminus I} \\ \hline i & * & k & * \\ \hline \omega & \sigma & \omega & \mathbb{0} \end{array} .$$

Vertices of this type exist as long as $I \neq [n]$, since we already require I to be non-empty.

Additionally, the vertices of type $u(k,\ell,\mathbb{0},\sigma)$ do not depend on the choice of $\ell \in [n]\setminus I$, which is why we can overload the notation and define a derived vertex type

$$u(k,\sigma) \quad := \quad \begin{array}{c|cc} I & \multicolumn{2}{c}{[n]\setminus I} \\ \hline * & k & * \\ \hline \sigma & \sigma & \mathbb{0} \end{array} ,$$

which also always exists if $I \neq [n]$.

The last definition for vertex types stems from the observation that

$$v(i,k,\ell,\mathbb{0},\sigma) = w(i,j,\mathbb{0},\sigma) = z(i,k,\mathbb{0},\sigma),$$

regardless of the choices of $j \in I$ and $k,\ell \in [n]\setminus I$, provided that all these vertex types exist. Vertices of this derived type can be defined more generally, also by overloading the notation, as

$$v(i,\sigma) \quad := \quad \begin{array}{cc|c} \multicolumn{2}{c|}{I} & [n]\setminus I \\ \hline i & * & * \\ \hline \mathbb{0} & \sigma & \mathbb{0} \end{array} ,$$

even if the other three types of vertices do not necessarily exist.

We next prove the following two so-called relative complementarity conditions, one for arbitrary $\omega \in H$ and one for arbitrary $\sigma \in \mathbb{F}_2^d \setminus H$, namely

$$a^k_\sigma - a^k_{\mathbb{0}} = a^i_{\mathbb{0}} - a^i_\sigma \quad \text{as well as} \quad a^k_\omega - a^k_{\mathbb{0}} = a^i_{\mathbb{0}} - a^i_\omega \tag{A}$$

for any $i \in I$ and $k \in [n]\setminus I$.

If $|[n]\setminus I| \geq 2$, we prove two more complementarity conditions, namely

$$a^k_\sigma - a^k_{\mathbb{0}} = a^i_\omega - a^i_{\sigma\oplus\omega} \quad \text{as well as} \quad a^k_\omega - a^k_{\mathbb{0}} = a^i_\omega - a^i_{\mathbb{0}} \tag{B}$$

for any $i \in I$ and $k \in [n]\setminus I$. Note that $|[n]\setminus I| \geq 2$ is implied by $|I| = 1$, because $n \geq 3$ is a prerequisite.

Note that both (A) and (B) require that the set $[n] \setminus I$ is non-empty, i. e., the vertex types $u(k,\sigma)$ and $z(i,k,\omega,\sigma)$ actually exist.

We see that for any $i \in I$ and any $k \in [n] \setminus I$,

$$\begin{aligned} 0 &= \langle a, v(i,\sigma)\rangle - \langle a, u(k,\sigma)\rangle \\ &= a^i_{\mathbb{0}} + \sum_{j\in I\setminus i} a^j_\sigma + \sum_{\ell\in[n]\setminus I} a^\ell_{\mathbb{0}} - a^k_\sigma - \sum_{j\in I} a^j_\sigma - \sum_{\ell\in[n]\setminus(I\cup k)} a^\ell_{\mathbb{0}} \\ &= a^i_{\mathbb{0}} + a^k_{\mathbb{0}} - a^k_\sigma - a^i_\sigma. \end{aligned}$$

Therefore, the relative complementarity condition

$$a^k_\sigma - a^k_{\mathbb{0}} = a^i_{\mathbb{0}} - a^i_\sigma \quad \text{for all } i \in I, k \in [n] \setminus I \tag{A1}$$

holds for the vector a of the equation $\langle a, y\rangle = \beta$.

The second complementarity condition of type (A) is also shown via the difference of the appropriate scalar products that evaluates to

$$\begin{aligned} 0 &= \langle a, v(i,\sigma)\rangle - \langle a, z(i,k,\omega,\sigma)\rangle \\ &= a^i_{\mathbb{0}} + \sum_{j\in I\setminus i} a^j_\sigma + \sum_{\ell\in[n]\setminus I} a^\ell_{\mathbb{0}} - a^i_\omega - a^k_\omega - \sum_{j\in I\setminus i} a^j_\sigma - \sum_{\ell\in[n]\setminus(I\cup k)} a^\ell_{\mathbb{0}} \\ &= a^i_{\mathbb{0}} + a^k_{\mathbb{0}} - a^i_\omega - a^k_\omega, \end{aligned}$$

which therefore means that

$$a^k_\omega - a^k_{\mathbb{0}} = a^i_{\mathbb{0}} - a^i_\omega \quad \text{for all } i \in I, k \in [n] \setminus I. \tag{A2}$$

The more general first relative complementarity condition of type (B) is established by evaluating the equation

$$\begin{aligned} 0 &= \langle a, z(i,\ell,\omega,\sigma\oplus\omega)\rangle - \langle a, u(k,\ell,\omega,\sigma)\rangle \\ &= a^i_\omega + a^\ell_\omega + \sum_{j\in I\setminus i} a^j_{\sigma\oplus\omega} + \sum_{r\in[n]\setminus(I\cup\ell)} a^r_{\mathbb{0}} - a^k_\sigma - a^\ell_\omega - \sum_{j\in I} a^j_{\sigma\oplus\omega} - \sum_{r\in[n]\setminus(I\cup\{k,\ell\})} a^r_{\mathbb{0}} \\ &= a^i_\omega + a^k_{\mathbb{0}} - a^k_\sigma - a^i_{\sigma\oplus\omega}, \end{aligned}$$

for some $\ell \in [n] \setminus I$ distinct from $k \in [n] \setminus I$, which clearly exists since for type (B) we assume $|[n] \setminus I| \geq 2$. From this we obtain

$$a^i_\omega - a^i_{\sigma\oplus\omega} = a^k_\sigma - a^k_{\mathbb{0}} \quad \text{for all } i \in I, k \in [n] \setminus I \tag{B1}$$

by reorganizing the terms.

The second condition of type (B) with the assumption $|[n] \setminus I| \geq 2$ is established by similarly evaluating, again for some $\ell \in [n] \setminus I$ distinct from $k \in [n] \setminus I$, the equation

$$\begin{aligned}
0 &= \langle a, z(i,\ell,\omega,\sigma)\rangle - \langle a, v(i,k,\ell,\omega,\sigma)\rangle \\
&= a^i_\omega + a^\ell_\omega + \sum_{j\in I\setminus i} a^j_\sigma + \sum_{r\in[n]\setminus(I\cup\ell)} a^r_{\mathbb{0}} - a^i_{\mathbb{0}} - a^k_\omega - a^\ell_\omega - \sum_{j\in I\setminus i} a^j_\sigma - \sum_{r\in[n]\setminus(I\cup\{k,\ell\})} a^r_{\mathbb{0}} \\
&= a^i_\omega + a^k_{\mathbb{0}} - a^i_{\mathbb{0}} - a^k_\omega
\end{aligned}$$

which leads us to

$$a^i_\omega - a^i_{\mathbb{0}} = a^k_\omega - a^k_{\mathbb{0}} \quad \text{for all } i \in I, k \in [n] \setminus I. \tag{B2}$$

Therefore, this relative complementarity condition holds as well if $|[n] \setminus I| \geq 2$.

We now observe that there must be some $q_0 \in [n]$ and $\xi_1, \xi_2 \in \mathbb{F}_2^d$ with

$$a^{q_0}_{\xi_1} \neq a^{q_0}_{\xi_2}.$$

By transitivity of equality, we can even assume that $\xi_1 \in H$ as well as $\xi_2 \in \mathbb{F}_2^d \setminus H$, and rename them according to our scheme as $\omega_0 = \xi_1$ as well as $\sigma_0 = \xi_2$, such that a satisfies the condition

$$a^{q_0}_{\omega_0} \neq a^{q_0}_{\sigma_0}.$$

Indeed, if for every $i \in [n]$ and every pair $(\omega, \sigma) \in H \times (\mathbb{F}_2^d \setminus H)$ we would have $a^i_\omega = a^i_\sigma$, then the vector a is constant on every block $i \in [n]$, contradicting the fact that $\langle a, y\rangle = \beta$ defines a hyperplane of the affine transversal space $\mathbb{A}(d,n)$, as the vector a could then be combined by appropriately scaled coefficient vectors of block equations, so the plane defined by the equation $\langle a, y\rangle = \beta$ would contain all of $\mathbb{A}(d,n)$. We use here the fact that the block equations are valid for the affine hull $\mathbb{A}(d,n)$ of $\mathrm{CTP}(d,n)$.

Without loss of generality, we assume that there even exists an index $i_0 \in I$ with $a^{i_0}_{\omega_0} \neq a^{i_0}_{\sigma_0}$, using relative complementarity: If $q_0 \in I$, then choose $i_0 = q_0$. Otherwise, $q_0 \in [n] \setminus I \neq \emptyset$. By the relative complementarity conditions (A), this implies $a^{q_0}_{\sigma_0} - a^{q_0}_{\mathbb{0}} = a^i_{\mathbb{0}} - a^i_{\sigma_0}$ and $a^{q_0}_{\omega_0} - a^{q_0}_{\mathbb{0}} = a^i_{\mathbb{0}} - a^i_{\omega_0}$ for any $i \in I$. By rearranging the terms, we get the two equations

$$a^i_{\sigma_0} = a^i_{\mathbb{0}} + a^{q_0}_{\mathbb{0}} - a^{q_0}_{\sigma_0} \quad \text{and} \quad a^i_{\omega_0} = a^i_{\mathbb{0}} + a^{q_0}_{\mathbb{0}} - a^{q_0}_{\omega_0}.$$

This implies that the left-hand side coefficients cannot be equal since the last terms on the right-hand side are not equal by assumption. Therefore, we can choose any $i \in I$ as i_0.

By adding multiples of the aforementioned block equations $\sum_{\xi\in\mathbb{F}_2^d} y^q_\xi = 1$ that define $\mathbb{A}(d,n)$ and scaling, we do not change the set of solutions to $\langle a, y\rangle = \beta$ that lie in $\mathbb{A}(d,n)$. Because we know that for the block i_0, there exist ω_0 and σ_0 such that $a^{i_0}_{\omega_0} \neq a^{i_0}_{\sigma_0}$, we normalize this block and assume

$$a^{i_0}_{\omega_0} = 1 \quad \text{and} \quad a^{i_0}_{\sigma_0} = 0, \tag{N1}$$

which is achieved by an addition of appropriately scaled block equations and a rescaling of $\langle a, y\rangle = \beta$. For every other block $q \in [n] \setminus i_0$, we assume that one coefficient $a^q_{\xi_q}$ equals a specific value by also adding appropriate multiples of block equations to $\langle a, y\rangle = \beta$. Since the relative complementarity conditions relate the entries of a to $a^q_{\mathbb{0}}$ for every $q \in [n]$, this element presents itself as a canonical choice. Therefore, we assume

$$a^i_{\mathbb{0}} = 1 \quad \text{for } i \in I \setminus i_0 \quad \text{and} \quad a^k_{\mathbb{0}} = 0 \quad \text{for } k \in [n] \setminus I. \tag{N2}$$

In fact, neither involvement of the block equation for i_0 nor additional rescaling are necessary to achieve (N2), so (N1) still holds after this step.

The normalization conditions (N1) and (N2) of $\langle a, y\rangle = \beta$ imply that the relative complementarity conditions can be written as absolute complementarity conditions of the form

$$a^k_\omega = 1 - a^i_\omega \quad \text{and} \quad a^k_\sigma = 1 - a^i_\sigma \quad \text{for all } i \in I \setminus i_0 \text{ and } k \in [n] \setminus I. \tag{C}$$

For the sake of clearness of the remaining exposition, especially since the distinguished element $i_0 \in I$ is handled differently, we will distinguish two cases, depending on the cardinality of I:

Case A: $|I| \geq 3$, and
Case B: $|I| = 1$, which in particular implies $|[n] \setminus I| \geq 2$.

Case A: The rest of the proof for the case $|I| \geq 3$ will work in multiple steps, which we establish as follows:

(A.1) Constant I-blocks for non-H-elements: $a^{i_0}_\sigma = 0$ and $a^i_\sigma = a^i_{\sigma\oplus\omega}$ for all $i \in I$ and all $\sigma \in \mathbb{F}_2^d \setminus H$ as well as all $\omega \in H$,
(A.2) Coefficients for H-elements: $a^i_\omega = 1$ for all $(i, \omega) \in I \times H$,
(A.3) Coefficients for non-H-elements: $a^i_\sigma = 0$ for all $i \in I$ and all $\sigma \in \mathbb{F}_2^d \setminus H$.

Due to complementarity (C), this also fixes, for all $k \in [n] \setminus I$, the coefficients $a^k_\sigma = 1$ for arbitrary $\sigma \in \mathbb{F}_2^d \setminus H$ and $a^k_\omega = 0$ for arbitrary $\omega \in H$.

(A.1): For two vertices $v(i, \sigma)$ and $v(i, \sigma \oplus \omega)$, we get the following equation:

$$\begin{aligned} 0 &= \langle a, v(i, \sigma)\rangle - \langle a, v(i, \sigma \oplus \omega)\rangle \\ &= a^i_{\mathbb{0}} + \sum_{j \in I \setminus i} a^j_\sigma + \sum_{k \in [n] \setminus I} a^k_{\mathbb{0}} - a^i_{\mathbb{0}} - \sum_{j \in I \setminus i} a^j_{\sigma\oplus\omega} - \sum_{k \in [n] \setminus I} a^k_{\mathbb{0}}. \end{aligned}$$

This can be simplified to the equation

$$\sum_{j \in I \setminus i} a^j_\sigma = \sum_{j \in I \setminus i} a^j_{\sigma\oplus\omega},$$

which implies that $\sum_{j\in I\setminus i}(a^j_\sigma - a^j_{\sigma\oplus\omega}) = 0$ for all $i \in I$. Introducing auxiliary variables $\hat{a}_j = (a^j_\sigma - a^j_{\sigma\oplus\omega})$ for the differences, the resulting system of equations for $i \in I$ is regular for $|I| \geq 3$, since the coefficient matrix can be written as an all-ones matrix with the diagonal set to zero, which is a regular matrix if it has at least 2 rows and columns. It follows that

$$a^j_\sigma = a^j_{\sigma\oplus\omega} \quad \text{for all } j \in I \text{ and } (\omega,\sigma) \in H \times \mathbb{F}_2^d \setminus H.$$

In particular, since $a^{i_0}_{\sigma_0} = 0$, we get

$$a^{i_0}_\sigma = a^{i_0}_{\sigma_0} = 0 \quad \text{for all } \sigma \in \mathbb{F}_2^d \setminus H.$$

(A.2): For $i \in I \setminus i_0$, it follows that

$$\begin{aligned} 0 &= \langle a, v(i,\sigma)\rangle - \langle a, w(i, i_0, \omega, \sigma)\rangle \\ &= \underbrace{a^i_{\mathbb{0}}}_{=1} + a^{i_0}_\sigma - a^i_\omega - a^{i_0}_{\omega\oplus\sigma}. \end{aligned}$$

From this and **(A.1)** we know that

$$a^i_\omega = 1 + \underbrace{a^{i_0}_\sigma - a^{i_0}_{\omega\oplus\sigma}}_{=0} = 1 \quad \text{for all } i \in I \setminus i_0.$$

Using the equation

$$\begin{aligned} 0 &= \langle a, w(i_0, j, \omega_0, \sigma)\rangle - \langle a, w(i_0, j, \omega_0 \oplus \omega, \sigma)\rangle \\ &= \underbrace{a^{i_0}_{\omega_0}}_{=1} + a^j_{\sigma\oplus\omega_0} - a^{i_0}_{\omega_0\oplus\omega} - a^j_{\omega_0\oplus\omega\oplus\sigma} \end{aligned}$$

for any $j \in I \setminus i_0$, we also get from **(A.1)** that

$$a^{i_0}_{\omega_0\oplus\omega} = 1 + a^j_{\sigma\oplus\omega_0} - a^j_{\omega_0\oplus\omega\oplus\sigma} = 1$$

holds, and therefore, by freely being able to choose $\omega \in H$ such that $\omega_0 \oplus \omega$ can attain any value in H,

$$a^i_\omega = 1 \quad \text{for all } (i,\omega) \in I \times H.$$

(A.3): For any $i \in I \setminus i_0$ and some $j \in I \setminus \{i, i_0\}$ (which exists because in this case we assume $|I| \geq 3$), we obtain that

$$\begin{aligned} 0 &= \langle a, w(i,j,\omega,\sigma)\rangle - \langle a, v(i_0, \sigma\oplus\omega)\rangle \\ &= \underbrace{a^i_\omega}_{=1} + \underbrace{a^j_{\omega\oplus\sigma}}_{=a^j_\sigma} + \sum_{h\in I\setminus\{i,j\}} a^h_\sigma + \sum_{\ell\in[n]\setminus I} a^\ell_{\mathbb{0}} - \underbrace{a^{i_0}_{\mathbb{0}}}_{=1} - \sum_{h\in I\setminus i_0} \underbrace{a^h_{\sigma\oplus\omega}}_{a^h_\sigma} - \sum_{\ell\in[n]\setminus I} a^\ell_{\mathbb{0}} \\ &= a^{i_0}_\sigma - a^i_\sigma, \end{aligned}$$

which shows that

$$a^i_\sigma = a^{i_0}_\sigma = 0 \quad \text{for all } i \in I \setminus i_0.$$

Case B: For the case $|I| = 1$, we also elaborate on the steps we take during the rest of the proof. Note that since here $i_0 \in I$ is the only element in I, we cannot use (C), as it constitutes a vacuous truth over the empty set $I \setminus i_0$. Nevertheless, we can use the two additional relative complementarity conditions (B1) and (B2) because $|[n] \setminus I| \geq 2$ is implied by $n \geq 3$. The remaining steps are as follows:

(B.1) Constant differences for different blocks: $a^{i_0}_\sigma - a^{i_0}_{\mathbb{0}} = a^k_\omega - a^k_{\sigma\oplus\omega}$ for any $k \in [n] \setminus I$ and arbitrary $\omega \in H$, $\sigma \in \mathbb{F}_2^d \setminus H$,
(B.2) Coefficients for H-elements: $a^{i_0}_\omega = 1$ and $a^k_\omega = 0$ for all $k \in [n] \setminus I$ and all $\omega \in H$,
(B.3) Coefficients for non-H-elements: $a^{i_0}_\sigma = 0$ and $a^k_\sigma = 1$ for all $k \in [n] \setminus I$ and all $\sigma \in \mathbb{F}_2^d \setminus H$.

(B.1): Since $|I| = 1$ and $n \geq 3$ imply that there exist at least two elements $k, \ell \in [n] \setminus I$ with $k \neq \ell$, another appropriate difference of scalar products is used to arrive at the first conclusion

$$\begin{aligned}
0 &= \langle a, u(k, \ell, \omega, \sigma \oplus \omega)\rangle - \langle a, v(i_0, k, \ell, \omega, \sigma)\rangle \\
&= a^{i_0}_\sigma + a^k_{\omega\oplus\sigma} + a^\ell_\omega + \sum_{r \in [n]\setminus\{i_0,k,\ell\}} a^r_{\mathbb{0}} - a^k_\omega - a^\ell_\omega - \sum_{r \in [n]\setminus\{k,\ell\}} a^r_{\mathbb{0}} \\
&= a^{i_0}_\sigma + a^k_{\omega\oplus\sigma} - a^{i_0}_{\mathbb{0}} - a^k_\omega,
\end{aligned}$$

which leads us to

$$a^{i_0}_\sigma - a^{i_0}_{\mathbb{0}} = a^k_\omega - a^k_{\sigma\oplus\omega} \quad \text{for all } k \in [n] \setminus I.$$

Note that this equation looks similar to (B1), but the block elements are switched between $i_0 \in I$ and $k \in [n] \setminus I$.

(B.2): Since both the relative complementarity conditions (A2) and (B2) hold, we equate them for some arbitrary $k \in [n] \setminus I \neq \emptyset$ to obtain

$$a^{i_0}_{\mathbb{0}} - a^{i_0}_\omega = a^{i_0}_\omega - a^{i_0}_{\mathbb{0}},$$

which we further simplify to

$$2a^{i_0}_{\mathbb{0}} = 2a^{i_0}_\omega.$$

By plugging in $\omega = \omega_0$, this leads to

$$a^{i_0}_{\mathbb{0}} = 1,$$

and therefore,

$$a^{i_0}_\omega = 1.$$

Using relative complementarity (A2) again, the condition

$$a_\omega^k = 0 \quad \text{for all } k \in [n] \setminus I$$

also follows.

(B.3): Since **(B.2)** shows $a_\omega^{i_0} = 1$ and **(B.1)** proves $a_\sigma^{i_0} - a_{\mathbb{0}}^{i_0} = a_\omega^k - a_{\sigma\oplus\omega}^k$ for all $k \in [n]\setminus I$, the equation

$$a_\sigma^{i_0} - 1 = 0 - a_\sigma^k \quad \text{for all } k \in [n] \setminus I$$

holds, by using $\omega = \mathbb{0}$ and the normalizations (N1) and (N2). Therefore, in the case $\sigma = \sigma_0$, we get

$$a_{\sigma_0}^k = 1 \quad \text{for all } k \in [n] \setminus I.$$

From this, (B1) and the normalizations we confirm that

$$a_\omega^{i_0} - a_{\sigma_0\oplus\omega}^{i_0} = \underbrace{a_{\sigma_0}^k}_{=1} - \underbrace{a_{\mathbb{0}}^k}_{=0} = 1.$$

Therefore, since $a_\omega^{i_0} = 1$ holds by **(B.2)**, we find $a_{\sigma_0\oplus\omega}^{i_0} = 0$, hence, with $\omega := \sigma_0 \oplus \sigma$,

$$a_\sigma^{i_0} = 0$$

holds as well. Now since all elements except for a_σ^k are fixed in the relative complementarity condition (B1), the final result

$$a_\sigma^k = 1 \quad \text{for all } k \in [n] \setminus I$$

follows directly. This finishes our proof for the second and final case $|I| = 1$. ■

The second result regarding odd-hyperplane inequalities concerns their sufficiency to provide an integral description of not only full but all cyclic transversal polytopes, as they cut off all integer vectors that do not correspond to cyclic transversals:

4.7 Theorem (Integer Hull of Odd-Hyperplane Inequalities)
For all block configurations $\Pi = (\Xi_1, \ldots, \Xi_n)$, the integer solutions of the system defined by

$$y \geq \mathbb{0},$$

$$\sum_{\omega\in\Xi_i} y_\omega^i = 1 \textit{ for all } i \in [n] \textit{ and}$$

$$\sum_{i\in I}\sum_{\omega\in H\cap\Xi_i} y_\omega^i + \sum_{i\notin I}\sum_{\omega\in\Xi_i\setminus H} y_\omega^i \geq 1 \textit{ for all hyperplanes } H \subsetneq \mathbb{F}_2^d \textit{ and } I \subseteq [n] \textit{ with } |I| \textit{ odd}$$

are exactly the vertices of $\mathrm{CTP}(\Pi)$.

Proof: Let $P \subseteq \mathbb{A}(\Pi)$ be the polytope described by the inequality system given in the theorem statement.

First, it is clear that every vertex of $\mathrm{CTP}(\Pi)$ lies in P, since all odd-hyperplane inequalities as well as the block equations and non-negativity constraints are valid for $\mathrm{CTP}(\Pi)$ by definition or by Lemma 4.5. It is also clear that all integral points of P must be 0/1-valued because of non-negativity constraints and block equations. This means we only have to show that P contains no 0/1-points other than the vertices of $\mathrm{CTP}(\Pi)$.

Let $v \in \mathbb{A}(\Pi) \setminus \mathrm{CTP}(\Pi)$ be an arbitrary 0/1-vector. Note that without loss of generality we require that v belongs to $\mathbb{A}(\Pi)$, so it corresponds to a selection of block elements that is not sum-zero, i. e., not a cyclic transversal. Otherwise, v would violate at least one of the block equations and so could not be an element of P.

Now let $\xi = (\xi(1), \dots, \xi(n))$ be the selection that corresponds to $v \in \mathbb{A}(\Pi)$, with $\xi(i) \in \Xi_i$. Since $v \notin \mathrm{CTP}(\Pi)$, we know that $\xi \notin \mathrm{CT}(\Xi_1, \dots, \Xi_n)$, so $\sum_{i\in[n]} \xi(i) \neq \mathbb{0}$. Therefore, there is already some coordinate $j \in [d]$ such that $\sum_{i\in[n]} \xi(i)_j \neq 0 \in \mathbb{F}_2$.

Let H be the hyperplane $\left\{\omega \in \mathbb{F}_2^d \mid \omega_j = 0\right\}$. If we show that v violates any odd-hyperplane inequality for this H and some $I \subseteq [n]$ of odd cardinality, we are done, so let I be the set $\{i \in [n] \mid \xi(i) \notin H\}$. Since this is the set of indices for which $\xi(i)_j = 1$ and the jth coordinate of the sum of all such vectors is 1, the set I is of odd cardinality.

We now evaluate the left-hand side of the odd-hyperplane inequality for v with the given H and I. The first double sum of this inequality

$$\sum_{i\in I} \sum_{\omega\in H\cap\Xi_i} v^i_\omega$$

has to be 0 because I is defined to contain the indices of block elements that are not in H and v^i_ω is 0 for all $\omega \neq \xi(i)$, and the second double sum

$$\sum_{i\notin I} \sum_{\omega\in \Xi_i\setminus H} v^i_\omega$$

is 0 as well for the same reason.

Therefore,

$$\sum_{i\in I} \sum_{\omega\in H\cap\Xi_i} v^i_\omega + \sum_{i\notin I} \sum_{\omega\in \Xi_i\setminus H} v^i_\omega = 0,$$

and so v violates this odd-hyperplane inequality, and thus does not belong to P. ■

This proof shows that the integral points of $\mathrm{CTP}(\Pi)$ are already cut out by the odd-hyperplane inequalities for the axis-parallel hyperplanes of $\mathbb{F}_2^d$ that are generated by the canonical basis vectors, of which there are only d many, in contrast to the $2^d - 1$ general hyperplanes of $\mathbb{F}_2^d$, one for every non-zero vector. Note that the number of inequalities is then still exponential in n, since there is one inequality for every odd subset of $[n]$.

With the versatile class of odd-hyperplane inequalities we can start to understand the intricacies of full cyclic transversal polytopes, at least for small parameters.

4.1 Descriptions for small parameters

In this section, we fix one of the parameters $d, n \in \mathbb{N}$ and attempt to describe the one-parameter family of full cyclic transversal polytopes that arises when varying the other non-fixed parameter. We give complete results for $d \leq 2$ or $n \leq 2$, and also explore partial results for $d = 3$ by means of visualizing the inequality orbits of $\mathrm{CTP}(3, n)$ for $n \leq 4$.

First, note that $\mathrm{CTP}(1, n)$ is (affinely) isomorphic to the known parity polytope

$$\mathrm{PAR}(n) = \mathrm{conv}\big\{x \in \{0,1\}^n \mid \langle \mathbb{1}, x\rangle \text{ is even}\big\}$$

via the affine map

$$y_1^i = x_i \quad \text{and} \quad y_0^i = 1 - x_i \quad \text{for all } i \in [n],$$

which we have seen in Corollary 3.11. By that isomorphism, inequalities for $\mathrm{CTP}(1, n)$ for $n \geq 4$ are readily available from the box constraints $\mathbb{0} \leq x \leq \mathbb{1}$ and the odd-set inequalities given by Jeroslow [31] which define the respective parity polytope $\mathrm{PAR}(n)$:

$$\sum_{i \in S} x_i - \sum_{i \in [n]-S} x_i \leq |S| - 1 \text{ for } S \subseteq [n] \text{ where } |S| \text{ is odd.}$$

Note that the odd-hyperplane inequalities defined in Lemma 4.5 reduce to an equivalent formulation of these odd-set inequalities for $d = 1$, and so odd-hyperplane inequalities constitute a generalization of these odd-set inequalities. This generalization will also emerge again in Example 5.5, where we will explicitly show the equivalence between both inequality formulations.

Moving to $d = 2$, we show that $\mathrm{CTP}(2, n)$ is completely described by non-negativity constraints, block equations and odd-hyperplane inequalities:

4.8 Theorem

The full cyclic transversal polytopes $\mathrm{CTP}(2, n)$ *are described by*

$$\begin{gathered} y \geq \mathbb{0}, \\ \sum_{\omega \in \mathbb{F}_2^2} y_\omega^i = 1 \textit{ for all } i \in [n] \textit{ and} \\ \sum_{i \in I} \sum_{\omega \in H} y_\omega^i + \sum_{i \notin I} \sum_{\omega \notin H} y_\omega^i \geq 1 \textit{ for all hyperplanes } H \subsetneq \mathbb{F}_2^d \textit{ and } I \subseteq [n] \textit{ with } |I| \textit{ odd.} \end{gathered}$$

We do not claim that the system given in Theorem 4.8 is irreducible, and in fact it is not in general: When $n = 1$, this is shown by means of the easy observation that $\mathrm{CTP}(2,1)$ consists of only one point, namely, the characteristic vector of $\mathbb{0} \in \mathbb{F}_2^2$. The system is also reducible for $n = 2$ by Proposition 3.18, because $\mathrm{CTP}(2,2)$ fulfills the additional equations $y^1_\omega = y^2_\omega$ for all $\omega \in \mathbb{F}_2^2$ that are not present in the description of $\mathrm{CTP}(2,n)$ above. In the remaining cases for $n \geq 3$ however, all odd-hyperplane inequalities generate facets for $\mathrm{CTP}(2,n)$ as shown in Theorem 4.6, and the non-negativity constraints also induce facets of these polytopes by Corollary 4.4. Therefore, we can enhance Theorem 4.8 by asserting that the given description is irreducible for $n \geq 3$.

Proof of Theorem 4.8: This proof will be done via strong duality certificates for every given objective function. Therefore, it suffices to construct, for every $c \in \left(\mathbb{R}^{\mathbb{F}_2^2}\right)^n$:

(1) some $\tilde{y} = (\chi(\tilde{\xi}(i)))_{i\in[n]}$ with $(\tilde{\xi}(i))_{i\in[n]} \in \mathrm{CT}(2,n)$, and

(2) some multipliers for the constraints listed in the theorem (non-negative ones for the inequalities) that are satisfied with equality by $\tilde{y}$, such that the corresponding linear combination $a \in \left(\mathbb{R}^{\mathbb{F}_2^2}\right)^n$ of the coefficient vectors satisfies $a = c$.

In fact, due to the equations in the system, it is enough to consider the cases with *normalized* c, i.e.,

$$\min\{c^i_\omega \mid \omega \in \mathbb{F}_2^2\} = 0 \text{ for all } i \in [n].$$

In particular, we then have $c \geq \mathbb{0}$ and, due to the non-negativity constraints in the description, only have to ensure

$$a \leq c \quad \text{ and } \quad a^i_{\tilde{\xi}(i)} = c^i_{\tilde{\xi}(i)} \text{ for all } i \in [n]. \tag{0}$$

(1) Construction of $\tilde{y}$

For each $i \in [n]$, choose one $\xi_0(i) \in \mathbb{F}_2^2$ with $c^i_{\xi_0(i)} = 0$ and define

$$\omega_0 := \sum_{i\in[n]} \xi_0(i).$$

If $\omega_0 = \mathbb{0}$, then we set $(\tilde{\xi}(i))_{i\in[n]} := (\xi_0(i))_{i\in[n]} \in \mathrm{CT}(2,n)$. Otherwise, we start by choosing $\ell_0 \in [n]$ with

$$c^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0} = \min\left\{c^i_{\xi_0(i)\oplus\omega_0} \,\middle|\, i \in [n]\right\}. \tag{1}$$

Then, we choose ω_1 and ω_2 such that $\mathbb{F}_2^2 = \{\mathbb{0}, \omega_0, \omega_1, \omega_2\}$ and $\ell_1, \ell_2 \in [n]$ with $\ell_1 \neq \ell_2$ such that

$$c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} = \min\left\{\sum_{k\in[2]} c^{i_k}_{\xi_0(i_k)\oplus\omega_k} \,\middle|\, \{i_1,i_2\} \in \binom{[n]}{2}\right\}. \tag{2}$$

This equation implies

$$c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} \le c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_1} + c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2},$$

thus

$$\left(c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} - c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_1}\right) + \left(c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} - c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2}\right) \le 0,$$

and hence we have

$$c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} \le c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_1} \quad \text{or} \quad c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} \le c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2}.$$

Note that by exchanging ω_1 and ω_2 as well as ℓ_1 and ℓ_2, that equation remains valid. Therefore, we choose the numbering such that we have

$$c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} \le c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_1}. \tag{3}$$

We finally distinguish two cases in the definition of $\tilde{\xi}$:

Case I: $c^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0} \le c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2}$, where we set

$$\begin{aligned} \tilde{\xi}(i) &:= \xi_0(i) \quad \text{for all } i \neq \ell_0 \\ \tilde{\xi}(\ell_0) &:= \xi_0(\ell_0) \oplus \omega_0. \end{aligned}$$

We have

$$\sum_{i\in[n]} \tilde{\xi}(i) = \sum_{i\in[n]} \xi_0(i) \oplus \omega_0 = \omega_0 \oplus \omega_0 = \mathbb{0},$$

which implies that $(\tilde{\xi}(i))_{i\in[n]} \in \mathrm{CT}(2,n)$ as required.

Case II: $c^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0} > c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2}$, in which case

$$\begin{aligned} \tilde{\xi}(i) &:= \xi_0(i) \quad \text{for all } i \neq \ell_1, \ell_2 \\ \tilde{\xi}(\ell_1) &:= \xi_0(\ell_1) \oplus \omega_1 \\ \tilde{\xi}(\ell_2) &:= \xi_0(\ell_2) \oplus \omega_2. \end{aligned}$$

Here we have

$$\sum_{i\in[n]} \tilde{\xi}(i) = \sum_{i\in[n]} \xi_0(i) \oplus \omega_1 \oplus \omega_2 = \omega_0 \oplus \omega_1 \oplus \omega_2 = \sum_{\omega\in\mathbb{F}_2^2} \omega = \mathbb{0},$$

which again implies that $(\tilde{\xi}(i))_{i\in[n]} \in \mathrm{CT}(2,n)$ as required.

Therefore, for any (normalized) objective function c, we constructed a cyclic transversal $(\tilde{\xi}(i))_{i\in[n]}$ and obtain an associated vertex $\tilde{y}$. To prove that this vertex minimizes $\langle c, y\rangle$ over $\mathrm{CTP}(2,n)$, we next construct a duality certificate, by defining multipliers for the constraints given in the theorem.

(2) Construction of constraint multipliers
If $\omega_0 = \mathbb{0}$, then the empty combination with $a = \mathbb{0}$ satisfies the conditions in (0). Otherwise, we construct a combination from the following inequalities. For $\omega \in \mathbb{F}_2^2 \setminus \mathbb{0}$, we remind the reader that we denote by $\operatorname{span}(\omega)$ the one-dimensional subspace of $\mathbb{F}_2^2$ generated by ω, i. e., the linear hull of ω.

For $\sigma_1, \sigma_2 \in \mathbb{F}_2^2$, we define

$$\sigma_1 \equiv_\omega \sigma_2 :\Longleftrightarrow \sigma_1 \oplus \sigma_2 \in \operatorname{span}(\omega).$$

Hence, for $\sigma \in \mathbb{F}_2^2$ we have $\sigma \in \operatorname{span}(\omega)$ if and only if $\sigma \equiv_\omega \mathbb{0}$. For $\omega \in \{\omega_1, \omega_2\}$ and

$$I(\omega) := \{i \in [n] \mid \xi_0(i) \notin \operatorname{span}(\omega)\}$$

it is true that $|I(\omega)| \in 2\mathbb{Z} + 1$, i. e., $|I(\omega)|$ is odd.

Proof of $|I(\omega)| \in 2\mathbb{Z}+1$: With $\operatorname{span}(\omega) = \{\sigma \in \mathbb{F}_2^2 \mid \langle h, \sigma\rangle = 0\}$, we have

$$1 \underset{\omega_0 \notin \operatorname{span}(\omega)}{=} \langle h, \omega_0\rangle = \sum_{i\in[n]} \langle h, \xi_0(i)\rangle = \sum_{i\in I(\omega)} \underbrace{\langle h, \xi_0(i)\rangle}_{=\,1} \oplus \sum_{i\notin I(\omega)} \underbrace{\langle h, \xi_0(i)\rangle}_{=\,0}.$$

The $[H, I]$-inequality $[\operatorname{span}(\omega), I(\omega)]$ for $\omega \in \{\omega_1, \omega_2\}$ is

$$\sum_{i\in I(\omega)} \sum_{\sigma \equiv_\omega \mathbb{0}} y_\sigma^i + \sum_{i\notin I(\omega)} \sum_{\sigma \not\equiv_\omega \mathbb{0}} y_\sigma^i \geq 1.$$

Since $i \in I(\omega)$ means that $\xi_0(i) \not\equiv_\omega \mathbb{0}$, and $i \notin I(\omega)$ respectively means that $\xi_0(i) \equiv_\omega \mathbb{0}$, we know that $\sigma \not\equiv_\omega \xi_0(i)$ for the indices of both inner sums. In fact, this means that $\sigma \not\equiv_\omega \xi_0(i)$ is an equivalent condition for the concatenation of the double sums, and so we obtain the simpler looking inequalities

$$\sum_{i\in[n]} \sum_{\sigma \not\equiv_{\omega_1} \xi_0(i)} y_\sigma^i \geq 1 \tag{ω_1}$$

and

$$\sum_{i\in[n]} \sum_{\sigma \not\equiv_{\omega_2} \xi_0(i)} y_\sigma^i \geq 1. \tag{ω_2}$$

For every $j \in [n]$ we also have the $[H, I]$-inequality $[\operatorname{span}(\omega_0), j]$:

$$\sum_{\sigma \equiv_{\omega_0} \mathbb{0}} y_\sigma^j + \sum_{\substack{i\in[n] \\ i\neq j}} \sum_{\sigma \not\equiv_{\omega_0} \mathbb{0}} y_\sigma^i \geq 1,$$

from which we subtract the block equation $\sum_{\sigma\in\mathbb{F}_2^2} y_\sigma^j = 1$ in order to obtain

$$-y^j_{\xi_0(j)\oplus\omega_1} - y^j_{\xi_0(j)\oplus\omega_2} + \sum_{\substack{i\in[n]\\ i\neq j}} \left(y^i_{\xi_0(i)\oplus\omega_1} + y^i_{\xi_0(i)\oplus\omega_2} \right) \geq 0. \tag{j}$$

We first show that the Inequalities (ω_1), (ω_2) in both cases are binding for $\tilde{y}$. Additionally, we show that in Case I the Inequality (j) for all $j \in [n]$ is binding, as well as in Case II that (ℓ_1) is binding for $\tilde{y}$, which is just (j) for $j = \ell_1$. The inequality (ℓ_2) is binding in Case II as well, but this is redundant for the construction of a.

Case I: For the first two inequalities, let $\omega \in \{\omega_1, \omega_2\}$. Then we have

$$\sum_{i\in[n]} \sum_{\sigma \not\equiv_\omega \xi_0(i)} \tilde{y}^i_\sigma = \sum_{\sigma \not\equiv_\omega \xi_0(\ell_0)} \tilde{y}^{\ell_0}_\sigma + \sum_{\substack{i\in[n]\\ i\neq \ell_0}} \sum_{\sigma \not\equiv_\omega \xi_0(i)} \tilde{y}^i_\sigma = 1,$$

since we know that $\xi_0(\ell_0) \oplus \omega_0 \not\equiv_\omega \xi_0(\ell_0)$ because $\omega_0 \not\equiv_\omega \mathbb{0}$. This gives us 1 in the first sum, and we deduce that all terms of the second sum are 0 since $\sigma \not\equiv_\omega \xi_0(i)$ implies $\sigma \neq \xi_0(i)$. For the second set of inequalities, for all $i \in [n]$ we know

$$\{\xi_0(i), \xi_0(i) \oplus \omega_0\} \cap \{\xi_0(i) \oplus \omega_1, \xi_0(i) \oplus \omega_2\} = \emptyset$$

which for $j \in [n]$ gives us

$$-\underbrace{\tilde{y}^j_{\xi_0(j)\oplus\omega_1}}_{=\,0} - \underbrace{\tilde{y}^j_{\xi_0(j)\oplus\omega_2}}_{=\,0} + \sum_{\substack{i\in[n]\\ i\neq j}} \Big(\underbrace{\tilde{y}^i_{\xi_0(i)\oplus\omega_1}}_{=\,0} + \underbrace{\tilde{y}^i_{\xi_0(i)\oplus\omega_2}}_{=\,0} \Big) = 0.$$

Case II: We first show that inequality (ω_1) is binding for $\tilde{y}$. The proof for (ω_2) is similar. Therefore,

$$\sum_{i\in[n]} \sum_{\sigma \not\equiv_{\omega_1} \xi_0(i)} \tilde{y}^i_\sigma = \sum_{\sigma \not\equiv_{\omega_1} \xi_0(\ell_2)} \tilde{y}^{\ell_2}_\sigma + \sum_{\substack{i\in[n]\\ i\neq \ell_2}} \sum_{\sigma \not\equiv_{\omega_1} \xi_0(i)} \tilde{y}^i_\sigma = 1,$$

because the first sum is 1 as a result of $\xi_0(\ell_2) \oplus \omega_2 \not\equiv_{\omega_1} \xi_0(\ell_2)$ (to recognize that, observe $\omega_2 \not\equiv_{\omega_1} \mathbb{0}$ analogous to Case I), and all terms $\tilde{y}^i_\sigma$ for $\sigma \not\equiv_{\omega_1} \xi_0(i)$ in the second sum again come out as 0, since the condition $\sigma \not\equiv_{\omega_1} \xi_0(i)$ means that $\tilde{y}^i_\sigma = 0$ for $i \neq \ell_1$ due to $\sigma \neq \xi_0(i)$, as well as $\tilde{y}^{\ell_1}_\sigma = 0$ due to $\sigma \neq \xi_0(i) \oplus \omega_1$. Furthermore, to show that (ℓ_1) is binding,

$$-\underbrace{\tilde{y}^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1}}_{=\,1} - \underbrace{\tilde{y}^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2}}_{=\,0} + \sum_{\substack{i\in[n]\\ i\neq \ell_1}} \Big(\underbrace{\tilde{y}^i_{\xi_0(i)\oplus\omega_1}}_{=\,0} + \underbrace{\tilde{y}^i_{\xi_0(i)\oplus\omega_2}}_{(*)} \Big) = 0,$$

because $(*)$ is 0 if $i \neq \ell_2$ and 1 for $i = \ell_2$.

We now define nonnegative multipliers λ_1 for Inequality (ω_2) and λ_2 for Inequality (ω_1) respectively, as well as in some cases some $j \in [n]$ and a nonnegative multiplier μ for Inequality (j) that yield some combination $a \in \left(\mathbb{R}^{\mathbb{F}_2^2}\right)^n$ as required in (0) above. The numbering for λ_1 and λ_2 will become clear when evaluating coefficients.

The coefficients of a for $i \in [n]$ are:

$$\begin{array}{lcl} a^i_{\xi_0(i)} & = & 0, \\ a^i_{\xi_0(i)\oplus\omega_0} & = & \lambda_1 + \lambda_2, \\ a^i_{\xi_0(i)\oplus\omega_1} & = & \lambda_1 \quad \begin{cases} -\mu & \text{, if } i = j \\ +\mu & \text{, if } i \neq j \end{cases}, \\ a^i_{\xi_0(i)\oplus\omega_2} & = & \lambda_2 \quad \begin{cases} -\mu & \text{, if } i = j \\ +\mu & \text{, if } i \neq j. \end{cases} \end{array}$$

What is left to determine are possible values for λ_1, λ_2 and μ to satisfy (0).

Case I: We need to satisfy the following system:

$$\begin{array}{lcll} a^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0} & = & c^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0}, & \\ a^i_{\xi_0(i)\oplus\omega_0} & \leq & c^i_{\xi_0(i)\oplus\omega_0} & \text{for all } i \in [n], i \neq \ell_0, \\ a^i_{\xi_0(i)\oplus\omega_1} & \leq & c^i_{\xi_0(i)\oplus\omega_1} & \text{for all } i \in [n], \\ a^i_{\xi_0(i)\oplus\omega_2} & \leq & c^i_{\xi_0(i)\oplus\omega_2} & \text{for all } i \in [n]. \end{array}$$

We choose $\lambda_1^*, i_1, \lambda_2^*$ and i_2 such that

$$\lambda_1^* = \min\Big\{ c^i_{\xi_0(i)\oplus\omega_1} \;\Big|\; i \in [n] \Big\} = c^{i_1}_{\xi_0(i_1)\oplus\omega_1}$$

as well as

$$\lambda_2^* = \min\Big\{ c^i_{\xi_0(i)\oplus\omega_2} \;\Big|\; i \in [n] \Big\} = c^{i_2}_{\xi_0(i_2)\oplus\omega_2},$$

and furthermore set $\delta^* := \lambda_1^* + \lambda_2^* - c^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0}$.

If we have $\delta^* \geq 0$, then we choose $\mu := 0$ and (λ_1, λ_2) as some solution to the following system:

$$\begin{array}{ccccc} \lambda_1 & + & \lambda_2 & = & c^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0}, \\ 0 & \leq & \lambda_1 & \leq & \lambda_1^*, \\ 0 & \leq & \lambda_2 & \leq & \lambda_2^*. \end{array}$$

We then have

$$a^i_{\xi_0(i)\oplus\omega_0} \;=\; \lambda_1 + \lambda_2 \;=\; c^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0} \underset{\text{because of (1)}}{\leq} c^i_{\xi_0(i)\oplus\omega_0},$$

with equality for $i = \ell_0$, and for the other inequalities we get

$$a^i_{\xi_0(i)\oplus\omega_1} \;=\; \lambda_1 \;\leq\; \lambda_1^* \;\leq\; c^i_{\xi_0(i)\oplus\omega_1}$$

and

$$a^i_{\xi_0(i)\oplus\omega_2} \;=\; \lambda_2 \;\leq\; \lambda_2^* \;\leq\; c^i_{\xi_0(i)\oplus\omega_2}.$$

Otherwise, we have $\delta^* < 0$, hence due to (2) and since we are in Case I, we know that i_1 and i_2 have to be identical, else the inequality for Case I would be violated. Now we use an Inequality (j) by fixing $i_1 = i_2 =: j$. We then choose the coefficients as follows:

$$\begin{array}{rcl} \lambda_1 & := & \frac{|\delta^*|}{2} \;+\; \lambda_1^*, \\ \lambda_2 & := & \frac{|\delta^*|}{2} \;+\; \lambda_2^*, \\ \mu & := & \frac{|\delta^*|}{2}, \end{array}$$

which yields

$$\begin{array}{rcccl} a^i_{\xi_0(i)\oplus\omega_0} & = & \lambda_1 + \lambda_2 & = & \lambda_1^* + \lambda_2^* + |\delta^*| \\ & & & = & \lambda_1^* + \lambda_2^* - \delta^* \\ & & & = & c^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0} \underset{(1)}{\leq} c^i_{\xi_0(i)\oplus\omega_0}, \end{array}$$

again with equality for $i = \ell_0$. Now for the other two inequalities we distinguish the cases $i = j$ and $i \neq j$, meaning that first we obtain

$$a^j_{\xi_0(j)\oplus\omega_1} \;=\; \lambda_1 - \mu \;=\; \lambda_1^* \;\leq\; c^j_{\xi_0(j)\oplus\omega_1}$$

and

$$a^j_{\xi_0(j)\oplus\omega_2} \;=\; \lambda_2 - \mu \;=\; \lambda_2^* \;\leq\; c^j_{\xi_0(j)\oplus\omega_2},$$

with equality because of the definition of λ_1^* and λ_2^*. For $i \neq j$ we firstly get

$$\begin{array}{rcccl} a^i_{\xi_0(i)\oplus\omega_1} & = & \lambda_1 + \mu & = & \lambda_1^* - \delta^* \\ & = & c^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0} - \lambda_2^* & = & c^{\ell_0}_{\xi_0(\ell_0)\oplus\omega_0} - c^j_{\xi_0(j)\oplus\omega_2} \\ & & & \leq & c^i_{\xi_0(i)\oplus\omega_1}, \end{array}$$

where the last relation stems from the fact that $i \neq j$ and we are in Case I. Secondly, the relation $a^i_{\xi_0(i)\oplus\omega_2} = ... \leq c^i_{\xi_0(i)\oplus\omega_2}$ is deduced similarly. This finishes the construction of a certificate in this case.

Case II: We need to satisfy the following system:

$$\begin{array}{rcll} a^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} & = & c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1}, & \\ a^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} & = & c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2}, & \\ a^i_{\xi_0(i)\oplus\omega_0} & \leq & c^i_{\xi_0(i)\oplus\omega_0} & \text{for all } i \in [n], \\ a^i_{\xi_0(i)\oplus\omega_1} & \leq & c^i_{\xi_0(i)\oplus\omega_1} & \text{for all } i \in [n], i \neq \ell_1 \\ a^i_{\xi_0(i)\oplus\omega_2} & \leq & c^i_{\xi_0(i)\oplus\omega_2} & \text{for all } i \in [n], i \neq \ell_2. \end{array}$$

To define λ_1, λ_2 and μ let

$$\gamma^* := \max\Big\{0, c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} - c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2}\Big\}.$$

We choose $j := \ell_1$ and define

$$\begin{aligned} \lambda_1 &:= \quad \frac{\gamma^*}{2} + c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1}, \\ \lambda_2 &:= -\frac{\gamma^*}{2} + c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2}, \\ \mu &:= \quad \frac{\gamma^*}{2}. \end{aligned}$$

Note that $\lambda_2 \geq 0$ due to $c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2} \geq 0$. For the first set of inequalities we then have

$$a^i_{\xi_0(i)\oplus\omega_0} = \lambda_1 + \lambda_2 = c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} \underset{\text{Case II}}{<} c^i_{\xi_0(i)\oplus\omega_0}.$$

To settle the other two sets of inequalities and the equations we calculate first, for $i \neq \ell_1$,

$$a^i_{\xi_0(i)\oplus\omega_1} = \lambda_1 + \mu = c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + \gamma^* \leq c^i_{\xi_0(i)\oplus\omega_1}.$$

The last inequality is clearly true if $\gamma^* = 0$, on the one hand because of (3) (if $i = \ell_2$) and on the other hand because of the choice of ℓ_1 in (2) (if $i \in [n] \setminus \{\ell_1, \ell_2\}$).
Otherwise, we know that

$$\gamma^* = c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} - c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2} > 0,$$

which results in

$$c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + \gamma^* = c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} - c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2} \underset{\text{because of (2)}}{\leq} c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_1},$$

so the case $i = \ell_2$ is also solved in this case. What remains to be shown is that the inequality

$$c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + \gamma^* = c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} - c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2} \leq c^i_{\xi_0(i)\oplus\omega_1}$$

holds when $i \neq \ell_1, \ell_2$ and we have $c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} > c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2}$. But this is true since (ℓ_1, ℓ_2) is one pair that minimizes the sum in (2), therefore any pair (i, ℓ_1) with $i \neq \ell_1, \ell_2$ results in

$$c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} + c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} \leq c^i_{\xi_0(i)\oplus\omega_1} + c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2},$$

and subtraction of $c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2}$ on both sides gives us the required inequality.
We also have for $i \neq \ell_1$:

$$a^i_{\xi_0(i)\oplus\omega_2} = \lambda_2 + \mu = c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} \underset{i\neq\ell_1}{\leq} c^i_{\xi_0(i)\oplus\omega_2},$$

where the last inequality comes from the fact that any $i \neq \ell_1$ that would give a coefficient $c^i_{\xi_0(i)\oplus\omega_2} < c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2}$ would have been chosen as ℓ_2 in the beginning of the proof, as seen in (2). This also resolves the second equation because $j \neq \ell_2$ and so

$$a^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} = \lambda_2 + \mu = c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2}$$

holds trivially. On the other hand, if $i = \ell_1$, the coefficient is

$$a^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2} = \lambda_2 - \mu = c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2} - \gamma^*,$$

which for $\gamma^* > 0$ is equal to $c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2}$, so the inequality is fulfilled with equality. Otherwise, for $\gamma^* = 0$ it is equal to $c^{\ell_2}_{\xi_0(\ell_2)\oplus\omega_2}$, which is then smaller than $c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2}$, so the inequality $a^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2} \leq c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_2}$ still holds. Finally, for the first equation where $i = \ell_1$, we get

$$a^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1} = \lambda_1 - \mu = c^{\ell_1}_{\xi_0(\ell_1)\oplus\omega_1}$$

by definition of λ_1 and μ.

This finishes the dual certificate construction, and so we have proven that $\mathrm{CTP}(2, n)$ is described as claimed. ■

This concludes our discussion regarding $\mathrm{CTP}(2, n)$. For $\mathrm{CTP}(3, n)$, of which we have calculated irreducible systems of inequalities for $n \leq 4$, there are some patterns arising, which will need to be investigated further in future works. SageMath [47] was used to generate (descriptions of) these cyclic transversal polytopes and their automorphism groups.

We remind the reader that we assume all inequalities to be normalized in the way described in the beginning of this chapter, and we partition valid inequalities of $\mathrm{CTP}(d, n)$ into orbits using the automorphism (sub)group $\Upsilon(d, n)$ generated by the coordinate permutations that arise from block permutations and equivalence transformations of the full blocks.

Calculation of the subgroups $\Upsilon(3, 3)$ and $\Upsilon(3, 4)$ reveals that they are already equal to the full automorphism groups $\Gamma(3, 3)$ and $\Gamma(3, 4)$, respectively. The group $\Upsilon(3, 3)$ consists of 64 512 elements, while $\Upsilon(3, 4)$ contains $2\,064\,384 = 2^5|\Upsilon(3, 3)|$ automorphisms. This leads to an obvious question:

4.9 Question
Is $\Upsilon(d, n)$ always equal to $\Gamma(d, n)$, or is there a symmetry of $\mathrm{CTP}(d, n)$ for some parameters $d \geq 3$ and $n \geq 5$ that cannot be described by coordinate permutations obtained from block permutations and equivalence transformations?

To explain our method of visualizing the inequalities for $\mathrm{CTP}(3, n)$, we assume the reader to be familiar with the *Fano plane*, which is shown in Figure 4.1: Each of the points of this plane represents a non-zero element of $\mathbb{F}_2^3$, and the lines (including the circular line) connecting three points represent two-dimensional subspaces within this venue space, excluding $\mathbb{0}$. To represent an inequality of $\mathrm{CTP}(3, n)$ using the Fano plane, observe that each block can be represented by a Fano plane together with another unconnected point that represents $\mathbb{0}_3$. Afterwards, we mark the points by different shapes corresponding to the coefficient they have in an appropriate representative of an orbit of normalized inequalities. A legend for these markings is given in Figure 4.1 as well.

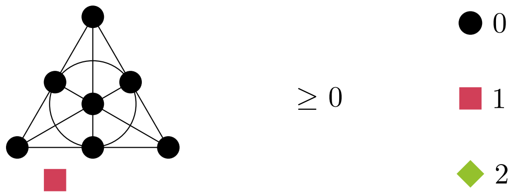

Figure 4.1: The Fano plane and a square representing $\mathbb{0}_3$ visualize a single block $\mathbb{F}_2^3$, and a right-hand side shows the rest of an inequality. On the right, there is a legend for the coefficients. This picture represents a non-negativity constraint $y_{\mathbb{0}} \geq 0$.

For the visualizations that follow, we chose a representative of an inequality orbit in such a way that the vertex representing the cyclic transversal $\xi\colon [n] \to \mathbb{0}$ lies in the facet defined by this representative. Note that even with this choice, the representative is not unique, as an automorphism on the blocks maps $\mathbb{0}$ to itself. With that in mind, we give priority to inequalities where the first block has a non-zero coefficient for $\mathbb{0}$ and identical block coefficients are grouped in the last blocks.

CTP$(3,3)$ is minimally described by a system of 2 740 inequalities which are partitioned into 4 orbits by actions of $\Upsilon(3,3)$ that are seen in Figure 4.2 together with their sizes, while an irreducible description of CTP$(3,4)$ contains 35 928 inequalities that are partitioned via actions of $\Upsilon(3,4)$ into 4 orbits as well, which in the same way are shown in Figure 4.3. One of these orbits in each case belongs to the non-negativity constraints, and another one belongs to the odd-hyperplane inequalities, which leaves two additional orbits that do not occur in inequality descriptions of CTP$(2,n)$ for any $n \in \mathbb{N}$. We can pair these orbits and see that they exhibit the same structure in both polytopes, but the sizes of their respective orbits grow differently.

If we fix the order n of full cyclic transversal polytopes and vary their rank d, we first observe that CTP$(d,1)$ only consists of one point, namely the characteristic vector of $\mathbb{0}$ in $\mathbb{F}_2^d$. We also obtain the result that CTP$(d,2)$ is isomorphic to the simplex Δ_{2^d-1}, as already seen in Proposition 3.18. While an explicit description of CTP$(3,3)$ and CTP$(3,4)$ is still feasible to calculate using conventional methods in SageMath [47], general conclusions on the structure of CTP$(d,3)$ are beyond the scope of this thesis. Curiously, although CTP$(4,3)$ has only $2^{4(3-1)} = 256$ vertices while CTP$(3,4)$ has $2^{3(4-1)} = 512$ vertices, a full description of CTP$(4,3)$ or full cyclic transversal polytopes with higher parameters is computationally infeasible at this point. Since almost nothing with a sufficient level of generality is known about CTP(d,n) for $d \geq 3$ and $n \geq 3$, we conclude this part with the obvious question:

4.10 Question

How can the (orbits of) facet-defining inequalities for $\mathrm{CTP}(d, n)$ with parameters $d \geq 3$ and $n \geq 3$ be characterized?

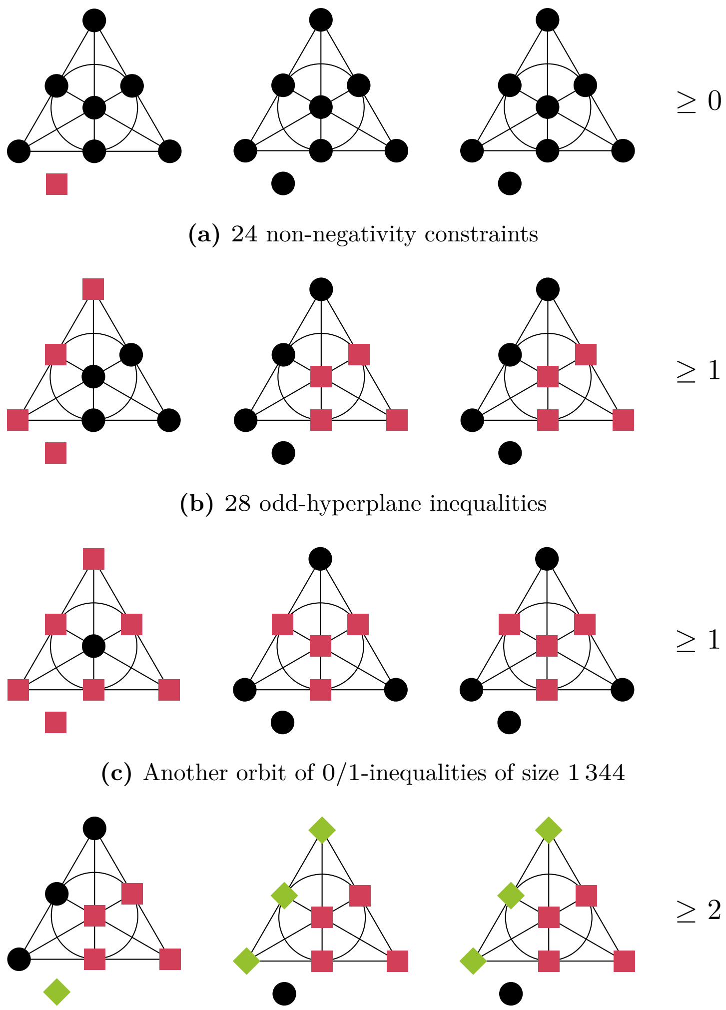

(a) 24 non-negativity constraints

(b) 28 odd-hyperplane inequalities

(c) Another orbit of 0/1-inequalities of size 1 344

(d) An orbit of 1 344 facet-defining inequalities with coefficients in $[2]_0$

Figure 4.2: Visualizations of inequality orbits of $\mathrm{CTP}(3, 3)$

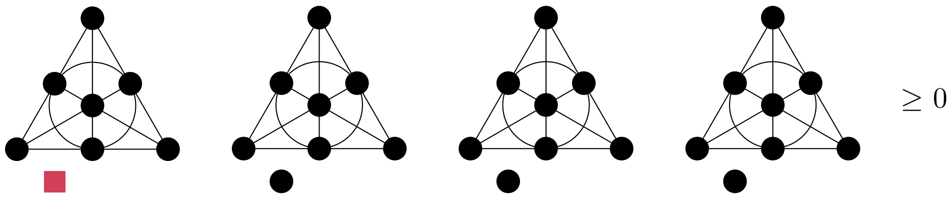

(a) 32 non-negativity constraints

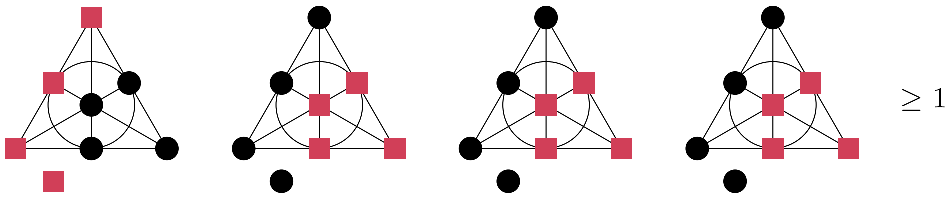

(b) 56 odd-hyperplane inequalities

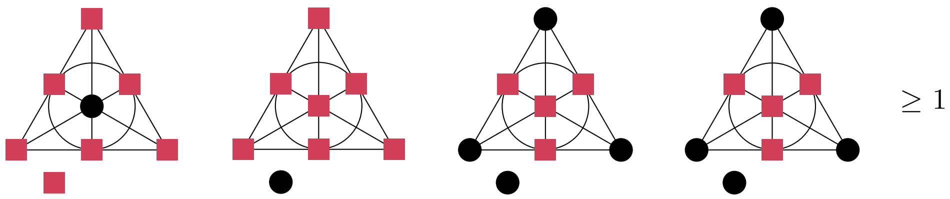

(c) Another orbit of 0/1-inequalities of size 21 504

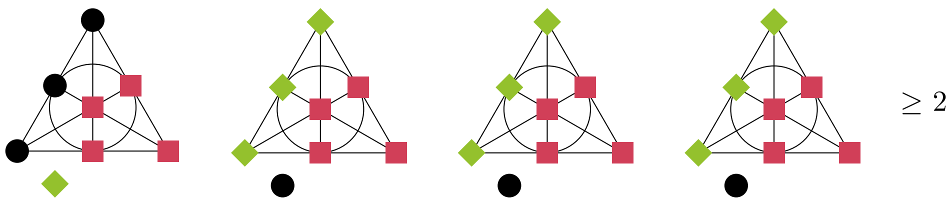

(d) An orbit of 14 336 facet-defining inequalities with coefficients in $[2]_0$

Figure 4.3: Visualizations of inequality orbits of $\mathrm{CTP}(3,4)$

5 Relaxations and Extended Formulations of CTPs

During Chapter 2, we familiarized ourselves with the notion of cyclic transversals and cyclic transversal polytopes as well as special subclasses and conditions that these polytopes have to fulfill. In this chapter, we want to find relaxations of these polytopes that enrich our theory by using projections (and lifts) mainly in two ways: First, by mapping the venue space of a given block configuration underlying a cyclic transversal polytope to another, smaller, venue space, relaxing a cyclic transversal polytope this way, and second, by identifying cyclic transversal polytopes themselves as projections of certain graph polytopes. This generalizes a well-known construction in combinatorial optimization.

5.1 Projected Sub-Transversals

For this section, we are given a block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$ on the venue space $\mathbb{F}_2^d$, which is of rank d. Unless specified otherwise, we assume that Π is normalized. Its order n and size s are assumed to be fixed. This block configuration Π induces a set of cyclic transversals $\mathrm{CT}(\Pi)$ and a cyclic transversal polytope $\mathrm{CTP}(\Pi)$ as before. Assume that we assign some weight $w^i_\omega \in \mathbb{R}$ to every vector $\omega \in \Xi_i$, and we want to find the cyclic transversal with the minimal joint weight, subject to this assignment. This is denoted as a (combinatorial) optimization problem (OP):

$$\begin{aligned} &\min \sum w^i_{\xi(i)} \\ &\text{s.t. } \xi \in \mathrm{CT}(\Pi). \end{aligned} \tag{OP}$$

The weights in this case can be seen as a vector $w \in \bigtimes_{i\in[n]} \mathbb{R}^{\Xi_i}$, and (OP) can then be reformulated using the cyclic transversal polytope $\mathrm{CTP}(\Pi)$, whose vertices correspond to cyclic transversals of Π. Then (OP) is equivalent to the (integer) linear optimization problem or (integer) linear program (LP):

$$\begin{aligned} &\min \langle w, y\rangle \\ &\text{s.t. } y \in \mathrm{CTP}(\Pi). \end{aligned} \tag{LP}$$

Since CTP(Π) is an integral polytope, we do not need to explicitly enforce integrality constraints on y. But to practically solve this linear program exactly with black-box methods, we effectively need a description of CTP(Π) by a finite set of linear inequalities. These descriptions are only available in very special cases so far, and even then, the number of necessary inequalities may lead to exact solution methods being infeasible due to time constraints.

Therefore, one of the main motivations to find relaxations of CTP(Π) is the ability to find approximations of solutions of (LP) with the same black-box algorithms, in general without knowing a complete description of CTP(Π) by inequalities.

Since the cyclic transversal polytopes encompass many well-known combinatorial polytopes like cut or stable set polytopes that represent important combinatorial optimization problems, optimization over the cyclic transversal polytopes generalizes these optimization problems as well. Due to this inherited importance to be able to optimize over the cyclic transversal polytopes, it is valuable to set up possible directions for further research into cyclic transversal polytope relaxations.

Because the set of coordinates for cyclic transversal polytopes are just tuples of subsets of the venue space $\mathbb{F}_2^d$, potentially with some additional structure, it is natural to consider linear maps of this venue space to help us obtain relaxations. A straight-forward choice for such maps are (orthogonal) projections onto subspaces of the same venue space. We will describe the relaxation procedure more generally for all linear maps whose codomain has dimension at most that of the domain. This will simplify descriptions of CTP(Π) by approximating this polytope, in a way, with lifted cyclic transversal polytopes given by block configurations over venue spaces of smaller (or equal) dimension.

That is, let $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^{\tilde{d}}$ be a linear map with $\tilde{d} \leq d$. The *mapping of a block configuration* $\Pi = (\Xi_1, \dots, \Xi_n)$ with respect to f, denoted by $f(\Pi)$, is defined to be the image of its blocks with respect to f, i. e., $f(\Pi) := (f(\Xi_1), \dots, f(\Xi_n))$, where $f(S) = \{f(\omega) \mid \omega \in S\}$ is the image of S with respect to f for $S \subseteq \mathbb{F}_2^d$ as usual. This way, we will examine the mapped set of cyclic transversals $\mathrm{CT}(f(\Pi))$ with respect to f and the mapped cyclic transversal polytope $\mathrm{CTP}(f(\Pi))$ with respect to f.

It is possible to define the central concepts of this chapter with an abstract $\mathbb{F}_2$-vector space V, a subspace $W \subseteq V$, and the linear map $f\colon V \to V/W$ that maps every $v \in V$ to its equivalence class in the quotient space V/W. This is where the *projection relaxation* notion in Definition 5.4 will come from, as these maps are usually called projections onto the quotient space. But when we are dealing with coordinates, for example in Example 5.3 and Corollary 5.10, we choose a basis for V, so we treat it as $\mathbb{F}_2^d$ from the beginning. In fact, specialization to vector spaces over $\mathbb{F}_2$ is also largely unnecessary, but we stay in this setting in the context of this work.

5.1 Proposition
Let $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^{\tilde{d}}$ *be a fixed linear map. Then the mapping of any block configuration* $\Pi = (\Xi_1, \dots, \Xi_n)$ *over* $\mathbb{F}_2^d$ *with respect to* f *induces a linear map*

$$\kappa_f\colon \mathbb{A}(\Pi) \to \mathbb{A}(f(\Pi)),$$

via

$$\kappa_f(y)^i_\sigma = \sum_{\substack{\omega\in\Xi_i \\ f(\omega)=\sigma}} y^i_\omega$$

for $i \in [d]$ *and* $\sigma \in f(\Xi_i)$*. This map has the property that* $\kappa_f(\mathrm{CTP}(\Pi)) \subseteq \mathrm{CTP}(f(\Pi))$*.*

Remember that $\mathbb{A}(\Pi)$ is the affine space of block equations for Π, defined in Definition 2.2, also called the affine transversal space.

Proof: Since κ_f is linear, we only have to prove that the vertices of $\mathrm{CTP}(\Pi)$ are mapped to some vertices of $\mathrm{CTP}(f(\Pi))$, which then implies the containment relation.

Let $\xi \in \mathrm{CT}(\Pi)$ be a cyclic transversal, and $v \in \mathrm{CTP}(\Pi)$ be the corresponding vertex of ξ. Then $\tilde{\xi} := (f(\xi(1)), \dots, f(\xi(n)))$ is the cyclic transversal that ξ is mapped to under f, i. e., it is clear that $\tilde{\xi}(i) = f(\xi(i)) \in f(\Xi_i)$ for all $i \in [n]$, and

$$\sum_{i\in[n]} \tilde{\xi}(i) = \sum_{i\in[n]} f(\xi(i)) = f\left(\sum_{i\in[n]} \xi(i)\right) = f(\mathbb{0}) = \mathbb{0}$$

holds because f is linear. So $\tilde{\xi} \in \mathrm{CT}(f(\Pi))$ holds. Now since $\kappa_f(v)^i_\sigma = 1$ is true if and only if $\sigma = f(\xi(i))$ holds and $\kappa_f(v)^i_\sigma = 0$ holds otherwise, it follows that $\kappa_f(v)$ is the vertex of $\mathrm{CTP}(f(\Pi))$ that corresponds to $\tilde{\xi}$. ■

For the sake of completeness, we now prove a small result about these induced maps κ_f which we will use in Lemma 5.6:

5.2 Proposition
For any two maps $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^{\tilde{d}}$ *and* $g\colon \mathbb{F}_2^{\tilde{d}} \to \mathbb{F}_2^{\hat{d}}$*, their composition* $g \circ f$ *induces a map* $\kappa_{g\circ f}$ *with the property that* $\kappa_{g\circ f} = \kappa_g \circ \kappa_f$*.*

Proof: We compare coefficients in the sums for both induced maps:

$$\kappa_{g\circ f}(y)^i_\sigma = \sum_{\substack{\omega\in\Xi_i \\ g\circ f(\omega)=\sigma}} y^i_\omega = \sum_{\substack{\omega\in\Xi_i \\ g(f(\omega))=\sigma}} y^i_\omega = \sum_{\substack{\theta\in f(\Xi_i) \\ g(\theta)=\sigma}} \sum_{\substack{\omega\in\Xi_i \\ f(\omega)=\theta}} y^i_\omega = \sum_{\substack{\theta\in f(\Xi_i) \\ g(\theta)=\sigma}} \kappa_f(y)^i_\theta = \kappa_g \circ \kappa_f(y)^i_\sigma.$$

Therefore, both maps are equal. ■

There are block configurations such that the containment in Proposition 5.1 is strict, even when the original cyclic transversal polytope is empty, so equality in general is impossible:

5.3 Example
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be the block configuration consisting of blocks $\Xi_i := \{\mathbb{e}_{2i}, \mathbb{e}_{2i+1}\} \subsetneq \mathbb{F}_2^d$ for $i \in [n]$ and $d \geq 2n$. Notice that $\mathbb{e}_1$ is not in any of the blocks, and no block element can be represented as the sum of other block elements, so $\mathrm{CT}(\Pi) = \emptyset$ and consequently $\mathrm{CTP}(\Pi) = \emptyset$. Now let $f\colon \mathbb{F}_2^d \to \mathbb{F}_2$ be the projection onto the first coordinate, $f(\omega) = \omega_1$. Then we see that the mapping of every block is just $f(\Xi_i) = \{0\}$, and so $\xi = (0, \dots, 0)$ is a valid cyclic transversal of $f(\Pi)$. This in turn means that $\mathrm{CTP}(f(\Pi)) \neq \emptyset$, so the inclusion $\emptyset = \kappa_f(\mathrm{CTP}(\Pi)) \subsetneq \mathrm{CTP}(f(\Pi))$ is strict.

Nevertheless, the preimages of $\mathrm{CTP}(f(\Pi))$ under these induced maps κ_f will be our first relaxation candidates:

5.4 Definition (Projection Relaxation with respect to a map)
For any block configuration Π with venue space $\mathbb{F}_2^d$ and any linear map $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^{\tilde{d}}$ we define

$$\mathrm{R}_f(\Pi) := \kappa_f^{-1}(\mathrm{CTP}(f(\Pi))) \cap \{y \in \mathbb{A}(\Pi) \mid y \geq \mathbb{0}\} \subseteq \mathbb{A}(\Pi)$$

and call it the *projection relaxation with respect to* f.

For full cyclic transversal polytopes $\mathrm{CTP}(d, n)$, we denote the projection relaxation with respect to f with $\mathrm{R}_f(d, n)$.

We reiterate that $\mathbb{A}(\Pi)$ is the affine transversal space defined by the block equations for Π, so $\mathrm{R}_f(\Pi)$ always fulfills these equations. All 0/1-vectors in this affine transversal space correspond to transversals of Π, hence the name, though they might not be cyclic.

The intersection with the set

$$\mathrm{C}(\Pi) := \{y \in \mathbb{A}(\Pi) \mid y \geq \mathbb{0}\} = \left\{y \in \mathbb{A}(\Pi) \;\middle|\; y^i_\omega \geq 0 \text{ for all } i \in [n] \text{ and } \omega \in \Xi_i\right\}$$

of non-negativity constraints will sometimes be denoted by $S_+ := S \cap \mathrm{C}(\Pi)$ for any set $S \subseteq \mathbb{A}(\Pi)$ in this chapter. Note that if we have non-negativity constraints, then $y^i_\omega \leq 1$ is also implied via the block equations, so this set is a subset of the 0/1-cube in $\mathbb{A}(\Pi)$.

For any $x \in \mathbb{A}(f(\Pi))$, we know that the preimage of x under κ_f is

$$\kappa_f^{-1}(x) = \left\{ y \in \mathbb{A}(\Pi) \;\middle|\; \sum_{\substack{\omega \in \Xi_i \\ f(\omega) = \sigma}} y^i_\omega = x^i_\sigma \text{ for all } i \in [n] \right\}.$$

Therefore, $\kappa_f^{-1}(x)$ itself is not necessarily contained in the 0/1-cube, so adding the non-negativity inequalities y^i_ω for $i \in [n]$ and $\omega \in \Xi_i$ is necessary for $\mathrm{R}_f(\Pi)$ to be a subset of the 0/1-cube in $\mathbb{A}(\Pi)$.

By Proposition 5.1, we know that $\kappa_f(\mathrm{CTP}(\Pi)) \subseteq \mathrm{CTP}(f(\Pi))$ holds. On the one hand, it is clear that $\mathrm{CTP}(\Pi) \subseteq \kappa_f^{-1}\big(\kappa_f(\mathrm{CTP}(\Pi))\big)$ is also true, by properties of the preimage. On the other hand, taking the preimage of $\mathrm{CTP}(f(\Pi))$ under κ_f will result in a set that is potentially larger than $\mathrm{R}_f(\Pi)$. But since $\mathrm{CTP}(\Pi)$ fully lies inside the set $\mathrm{C}(\Pi) = \{y \in \mathbb{A}(\Pi) \mid y \geq \mathbb{0}\}$, we have that

$$\mathrm{CTP}(\Pi) \subseteq \kappa_f^{-1}\big(\kappa_f(\mathrm{CTP}(\Pi))\big) \cap \mathrm{C}(\Pi) \subseteq \kappa_f^{-1}(\mathrm{CTP}(f(\Pi))) \cap \mathrm{C}(\Pi) = \mathrm{R}_f(\Pi),$$

which is why we call $\mathrm{R}_f(\Pi)$ a relaxation of $\mathrm{CTP}(\Pi)$.

Enforcing that linear maps f have to be surjective does not have an influence on the relaxation $\mathrm{R}_f(\Pi)$, as all linear maps can be written as the composition of a surjective linear map $g\colon \mathbb{F}_2^d \to \mathbb{F}_2^k$ and an embedding (an injective linear map) of $\mathbb{F}_2^k$ into $\mathbb{F}_2^{\tilde{d}}$. In fact, by definition of block isomorphisms, applying this injective linear map to a family of blocks is equivalent to applying a block isomorphism that maps the same family of blocks to a span of block elements over a (potentially higher-dimensional) venue space $\mathbb{F}_2^{\tilde{d}}$. Therefore, we may disregard this artificial block isomorphism and restrict ourselves to work with maps f where $\dim \mathrm{im}(f) = \tilde{d}$.

Let us look at a first example of this newly defined relaxation:

5.5 Example
Let Π be an n-tuple consisting of blocks $\Xi_i = \mathbb{F}_2^2$ for $i \in [n]$ and let $f\colon \mathbb{F}_2^2 \to \mathbb{F}_2$ be the orthogonal projection onto the first coordinate, then $\ker(f) = \mathrm{span}(\mathbb{e}_2)$. The cyclic transversal polytope $\mathrm{CTP}(\Pi)$ is the full cyclic transversal polytope $\mathrm{CTP}(2,n)$, and its projected cyclic transversal polytope $\mathrm{CTP}(f(\Pi))$ with respect to f is isomorphic to $\mathrm{CTP}(1,n)$. This polytope itself is isomorphic to the parity polytope $\mathrm{PAR}(n)$ as seen in Corollary 3.11 and Section 4.1.

The relaxation $\mathrm{R}_f(\Pi)$ is then described as follows:

$$\begin{aligned}
\mathrm{R}_f(\Pi) &= \kappa_f^{-1}(\mathrm{CTP}(f(\Pi)))_+ = \kappa_f^{-1}(\mathrm{CTP}(1,n))_+ \\
&= \kappa_f^{-1}\left\{ x \in \mathbb{A}(1,n) \;\middle|\; \sum_{i \in I} x_1^i - \sum_{i \notin I} x_1^i \leq |I| - 1 \text{ for all } I \subseteq [n] \text{ with } |I| \text{ odd} \right\}_+ \\
&= \mathrm{C}(2,n) \cap \left\{ y \in \mathbb{A}(2,n) \;\middle|\; \begin{array}{l} x_0^i + x_1^i = 1 \text{ for all } i \in [n], \\ \displaystyle\sum_{\omega \in \ker(f)} y^i_\omega = x_0^i \text{ and } \sum_{\omega \notin \ker(f)} y^i_\omega = x_1^i \text{ for all } i \in [n], \\ \displaystyle\sum_{i \in I} x_1^i - \sum_{i \notin I} x_1^i \leq |I| - 1 \text{ for all } I \subseteq [n] \text{ with } |I| \text{ odd} \end{array} \right\}
\end{aligned}$$

We undertake some straightforward transformations to eliminate the x-variables using the given equations to reveal a concrete description of $\mathrm{R}_f(\Pi)$, so for any $I \subseteq [n]$ with $|I|$ odd:

$$\begin{aligned}
& \sum_{i\in I} x_1^i - \sum_{i\notin I} x_1^i \leq |I| - 1 \\
\underset{\cdot(-1)}{\Longleftrightarrow}\quad & \sum_{i\notin I} x_1^i - \sum_{i\in I} x_1^i \geq 1 - |I| \\
\underset{+|I|}{\Longleftrightarrow}\quad & \sum_{i\notin I} x_1^i + \sum_{i\in I} (1 - x_1^i) \geq 1 \\
\underset{x_0^i + x_1^i = 1}{\Longleftrightarrow}\quad & \sum_{i\notin I} x_1^i + \sum_{i\in I} x_0^i \geq 1 \\
\underset{\text{express with } y}{\Longleftrightarrow}\quad & \sum_{i\notin I} \sum_{\omega\notin \ker(f)} y_\omega^i + \sum_{i\in I} \sum_{\omega\in \ker(f)} y_\omega^i \geq 1
\end{aligned}$$

These inequalities are the odd-hyperplane inequalities described in Lemma 4.5, with the hyperplane $H = \ker(f) = \mathrm{span}(\mathbb{e}_2)$ and I as presented. The projection relaxation $\mathrm{R}_f(\Pi)$ is described by these inequalities and the block equations.

We note that the example works analogously for any other linear map $f: \mathbb{F}_2^2 \to \mathbb{F}_2$. More precisely, the same procedure using, for example, the projection onto the second coordinate would give us another subset of the odd-hyperplane inequalities that describe $\mathrm{CTP}(2, n)$, since the kernel then is equal to $H = \mathrm{span}(\mathbb{e}_1)$. The argument would be the exact same except for switching the venue space coordinates. Since both of the resulting preimages of the induced maps are understood quite easily on their own with the given inequality description, it might be convenient to use various relaxations with different kernels at the same time to get a better relaxation. Therefore, instead of focusing on a single linear map, we will focus on multiple such maps and intersect all resulting relaxations.

There are numerous possibilities on which linear maps to use simultaneously, even if we would restrict ourselves to (orthogonal) projections. One may think about a finite sequence $(f_1, \dots, f_t)$ of maps where $\ker(f_j) \subsetneq \ker(f_i)$ whenever $i < j$. The corresponding sequence of images are also called a (partial) flag, and they form a *chain* in the set of all subspaces of our venue space when ordered by inclusion. This choice is not an improvement over choosing the map with the smallest kernel, as the following lemma shows:

5.6 Lemma
Let $f, g: \mathbb{F}_2^d \to \mathbb{F}_2^{\tilde{d}}$ be two linear maps such that $\ker(f) \subseteq \ker(g)$. *Then* $\mathrm{R}_f(\Pi) \subseteq \mathrm{R}_g(\Pi)$.

Proof: Since $\ker(f) \subseteq \ker(g)$, there is a linear map h such that $g = h \circ f$, as follows:

We define $h\colon \mathrm{im}(f) \to \mathbb{F}_2^{\tilde{d}}$ with $f(\omega) \mapsto g(\omega)$. This map is well-defined: if $f(\omega_1) = f(\omega_2)$, we have $\omega_1 - \omega_2 \in \ker(f) \subseteq \ker(g)$, so it follows that $g(\omega_1) = g(\omega_2)$. This is the map we are looking for, since by definition $h(f(\omega)) = g(\omega)$ for all $\omega \in \mathbb{F}_2^d$.

Now to show $\mathrm{R}_f(\Pi) = \kappa_f^{-1}(\mathrm{CTP}(f(\Pi))) \subseteq \kappa_g^{-1}(\mathrm{CTP}(g(\Pi))) = \mathrm{R}_g(\Pi)$, we use this composition representation of g. It is clear that $\mathrm{CTP}(g(\Pi)) = \mathrm{CTP}(h(f(\Pi)))$, and by Proposition 5.2, we have $\kappa_g = \kappa_h \circ \kappa_f$.

Since $\mathrm{R}_g(\Pi) = \kappa_g^{-1}(\mathrm{CTP}(g(\Pi)))$ is therefore equal to $\kappa_f^{-1}(\kappa_h^{-1}(\mathrm{CTP}(h(f(\Pi)))))$, it suffices to show that $\mathrm{CTP}(f(\Pi))$ is contained in $\kappa_h^{-1}(\mathrm{CTP}(h(f(\Pi)))) = \mathrm{R}_h(f(\Pi))$, but this is clear from Proposition 5.1 using the definition of the projection relaxation with respect to h, with Π replaced by $f(\Pi)$. Applying the preimage with respect to κ_f on both $\mathrm{CTP}(f(\Pi))$ and $\mathrm{R}_h(f(\Pi))$ then yields the desired relationship between the relaxations. ■

Since Lemma 5.6 shows that any inclusion relation of kernels does not lead to an improvement in the relaxation, the only other apparent possibility would be to use an *anti-chain* in the language of partially ordered sets, which in our case is a tuple of maps whose kernels (and images) are pairwise incomparable with respect to inclusion. Simple anti-chains in this setting are maps that have images (and kernels) of fixed dimension, similar to what we theorized after Example 5.5:

5.7 Definition (Projection Relaxation of dimension k)
We define $\mathrm{R}_k(\Pi)$ to be the intersection of all $\mathrm{R}_f(\Pi)$ with $\dim \mathrm{im}(f) = k$ and call it the *projection relaxation of dimension k*.

Analogous to before, we denote the projection relaxation of dimension k for the full cyclic transversal polytope $\mathrm{CTP}(d, n)$ with $\mathrm{R}_k(d, n)$.

Note that all definitions regarding relaxations using kernels or images of linear maps in this chapter can be defined just as well with the respective counterpart. This is due to the fundamental homomorphism theorem from linear algebra [11, Satz 8], as there is a well-defined isomorphism η between the quotient space $V/\ker(f)$ and $\mathrm{im}(f)$ via $\eta(v + \ker(f)) = f(v)$, for any linear map $f\colon V \to W$ with vector spaces V and W.

It is clear that $\mathrm{R}_d(\Pi)$ is equivalent to $\mathrm{CTP}(\Pi)$, since an image of dimension d implies that the linear maps at play are isomorphisms of $\mathbb{F}_2^d$, which identifies the maps f as block isomorphisms in the language of cyclic transversal equivalence transformations.

We call $\mathrm{R}_1(\Pi)$ a *linear form relaxation*, the simplest step in our relaxation hierarchy. Increasing the dimension parameter k in the projection relaxation, the polytopes more accurately approximate the original cyclic transversal polytope of Π:

5.8 Lemma (Projection Relaxation Hierarchy)
For any two parameters $k_1 \leq k_2 \leq d$, we have $\mathrm{R}_{k_2}(\Pi) \subseteq \mathrm{R}_{k_1}(\Pi)$.

Proof: This follows from Lemma 5.6, since for any linear map g with $\dim \operatorname{im}(g) = k_1$ there is a linear map f with $\dim \operatorname{im}(f) = k_2$ and $\operatorname{im}(g) \subseteq \operatorname{im}(f)$, possibly by embedding $\operatorname{im}(g)$ into a larger venue space. Therefore, $\mathrm{R}_f(\Pi) \subseteq \mathrm{R}_g(\Pi)$ for these maps, so the intersection of all relaxations from the left-hand side is contained in the intersection of all relaxations from the right-hand side. ■

Lemma 5.8 shows that the intersection $\mathrm{R}_k(\Pi)$ does not change if we relax the condition the maps have to fulfill from $\dim \operatorname{im}(f) = k$ to $\dim \operatorname{im}(f) \leq k$.

We invoke Example 5.5 once more to make a statement about the linear form relaxation $\mathrm{R}_1(\Pi)$ in general, since we have not really used special properties of the full cyclic transversal polytope $\mathrm{CTP}(2, n)$ there:

5.9 Lemma (Description of $\mathrm{R}_1(\Pi)$ by inequalities)
For any block configuration Π, the linear form relaxation $\mathrm{R}_1(\Pi)$ is equal to the polytope defined by all block equations for Π, all non-negativity constraints and the odd-hyperplane inequalities that are given in Lemma 4.5. Furthermore, if Π is a full block configuration of rank d and order n, then in the intersection of all $\mathrm{R}_f(\Pi)$ over all linear forms f which defines $\mathrm{R}_1(\Pi)$, every linear form contributes to the facets of $\mathrm{R}_1(\Pi)$ meaning no such linear form is redundant.

Proof: We differentiate two cases: First, we prove that the linear form relaxation $\mathrm{R}_1(d, n)$ of the full cyclic transversal polytope $\mathrm{CTP}(d, n)$ is equal to the polytope given by all block equations and all odd-hyperplane inequalities, and second, when blocks are allowed to be proper subsets of their respective venue space, the statement follows by substitution of some variables.

We know that $\mathrm{R}_1(d, n)$ is the intersection of all $\mathrm{R}_f(d, n)$ with linear maps $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^{\tilde{d}}$ with $\dim \operatorname{im}(f) = 1$. Without loss of generality, we assume that all linear maps $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^{\tilde{d}}$ necessary in the linear form relaxation are surjective, so in this proof, we only need to consider linear maps $f\colon \mathbb{F}_2^d \to \mathbb{F}_2$ with non-trivial image.

Now for any hyperplane $H = \left\{\omega \in \mathbb{F}_2^d \;\middle|\; \langle h, \omega\rangle = 0\right\}$ where $h \neq \mathbb{0}$, there is, up to scaling, exactly one such map $f\colon \mathbb{F}_2^d \to \mathbb{F}_2$ such that $\ker(f) = H$, namely the linear form $\langle h, \cdot\rangle$. Since 1 is the only scalar in $\mathbb{F}_2$, this linear form is actually unique, and we call it f_h.

The inequalities defining $\mathrm{R}_{f_h}(d, n)$ are determined analogously to Example 5.5. We know that $\mathrm{R}_{f_h}(d, n)$ is equal to

$$\mathrm{C}(d, n) \cap \left\{ y \in \mathbb{A}(d, n) \;\middle|\; \begin{array}{l} x_0^i + x_1^i = 1 \text{ for all } i \in [n], \\ \displaystyle\sum_{\omega \in \ker(f_h)} y_\omega^i = x_0^i \text{ and } \sum_{\omega \notin \ker(f_h)} y_\omega^i = x_1^i \text{ for all } i \in [n], \\ \displaystyle\sum_{i \in I} x_1^i - \sum_{i \notin I} x_1^i \leq |I| - 1 \text{ for all } I \subseteq [n] \text{ with } |I| \text{ odd} \end{array} \right\}.$$

Then we substitute the x-variables in the odd-set inequalities using the given equations and employ some arithmetic to obtain the inequalities

$$\sum_{i\notin I}\sum_{\omega\notin\ker(f_h)} y^i_\omega + \sum_{i\in I}\sum_{\omega\in\ker(f_h)} y^i_\omega \geq 1,$$

which are exactly the odd-hyperplane inequalities for all $I \subseteq [n]$ of odd cardinality and the chosen hyperplane $H = \{\omega \in \mathbb{F}_2^d \mid \langle h,\omega\rangle = 0\} = \ker(f_h)$.

Since all surjective linear maps $f\colon \mathbb{F}_2^d \to \mathbb{F}_2$ with non-trivial image are already of the form $\langle h,\cdot\rangle$ for some $h \in \mathbb{F}_2^d \setminus \mathbb{0}$ and all such maps are distinct, no other inequalities are generated by the intersection of all projection relaxations $\mathrm{R}_f(d,n)$ with respect to linear forms f and every such relaxation generates a distinct set of odd-hyperplane inequalities. By Theorem 4.6, all of the odd-hyperplane inequalities induce facets of $\mathrm{CTP}(d,n)$. Since these inequalities are also valid for its linear form relaxation $\mathrm{R}_1(d,n)$ (which by Proposition 4.2 has the same dimension as $\mathrm{CTP}(d,n)$), they necessarily also have to be facets of $\mathrm{R}_1(d,n)$, which completes the proof of this case.

If $\Pi = (\Xi_1,\dots,\Xi_n)$ contains blocks $\Xi_i \subsetneq \mathbb{F}_2^d$, the cyclic transversal polytope $\mathrm{CTP}(\Pi)$ is isomorphic to the face of $\mathrm{CTP}(d,n)$ generated by intersecting it with the hyperplanes $y^i_\omega = 0$ for all $i \in [n]$ and $\omega \in \mathbb{F}_2^d \setminus \Xi_i$. In fact, replacing these variables y^i_ω with 0 in the first case of this proof, the rest of it remains analogous, so that the projection relaxation $\mathrm{R}_{f_h}(\Pi)$ with respect to the linear form f_h is equal to

$$\mathrm{C}(\Pi) \cap \left\{ y \in \mathbb{A}(\Pi) \;\middle|\; \begin{array}{l} x^i_0 + x^i_1 = 1 \text{ for all } i \in [n], \\ \displaystyle\sum_{\omega\in\Xi_i\cap\ker(f_h)} y^i_\omega = x^i_0 \text{ and } \sum_{\omega\in\Xi_i\setminus\ker(f_h)} y^i_\omega = x^i_1 \text{ for all } i \in [n], \\ \displaystyle\sum_{i\in I} x^i_1 - \sum_{i\notin I} x^i_1 \leq |I| - 1 \text{ for all } I \subseteq [n] \text{ with } |I| \text{ odd} \end{array} \right\},$$

from which we obtain again the odd-hyperplane inequalities

$$\sum_{i\notin I}\sum_{\omega\in\Xi_i\setminus\ker(f_h)} y^i_\omega + \sum_{i\in I}\sum_{\omega\in\Xi_i\cap\ker(f_h)} y^i_\omega \geq 1.$$

These are odd-hyperplane inequalities for $\mathrm{CTP}(\Pi)$. ■

The linear form relaxation has some other desirable properties:

5.10 Corollary
The integer points of the linear form relaxation $\mathrm{R}_1(\Pi)$ *coincide with the vertices of the cyclic transversal polytope* $\mathrm{CTP}(\Pi)$.

Proof: This lemma follows directly from Lemma 5.9 together with Theorem 4.7. If Π is not full, identification of $\mathrm{CTP}(\Pi)$ as a face of $\mathrm{CTP}(d,n)$ by means of additional equations of the form $y^i_\omega = 0$ for $i \in [n]$ and $\omega \in \mathbb{F}_2^d \setminus \Xi_i$ make the statement valid for these block configurations as well. ■

From the proof of Theorem 4.7 we gather the even stronger statement that the cyclic transversal polytope $\mathrm{CTP}(\Pi)$ is already equal to the integer hull of the intersection of all projection relaxations with respect to coordinate projections f_j for $j \in [d]$, defined via $f_j(\omega) = \omega_j$. These are a subset of all linear forms of cardinality d, whereas there are $2^d - 1$ many linear forms over $\mathbb{F}_2^d$ in general, namely one for every non-zero vector.

Of course, the linear form relaxation will in general still contain fractional vertices, but it is valuable and helpful for solving (LP) to characterize when a projection relaxation, especially the linear form relaxation, is already equal to the cyclic transversal polytope, or more general, which relaxations generate facets for the cyclic transversal polytope.

An important paper by Barahona and Grötschel [7] characterizes, among other results, the case $\mathrm{R}_1(\Pi) = \mathrm{CTP}(\Pi)$ given that $\mathrm{CTP}(\Pi)$ is (affinely isomorphic to) a binary kernel polytope, using matroid theory. To meaningfully introduce the next theorem which builds on this paper, we need the notion of a minor of a matroid (see [41, Chapter 3]). Three minors are of particular interest: F_7^*, R_{10}, and $\mathcal{M}(K_5)^*$. The matroid F_7^* is called the *dual Fano matroid*. The matroid R_{10} is the binary matroid associated with the 5×10 matrix whose columns are the ten 0/1-vectors with 3 ones and 2 zeros, and $\mathcal{M}(K_5)^*$ is the cographic matroid of the complete graph K_5. For an extensive discussion on matroid theory and its terminology, we again refer to Oxley [41].

5.11 Theorem
Let $\Pi = (\Xi_1, \dots, \Xi_n)$ be a normalized block configuration with $|\Xi_i| = 2$ for all $i \in [n]$. Then, if $\omega_i \in \Xi_i$ denotes the non-zero element in Ξ_i, let M denote the matrix that is defined via the columns $M_{*,i} = \omega_i$. The linear form relaxation $\mathrm{R}_1(\Pi)$ is equal to the (binary) kernel polytope $\mathrm{CTP}(\Pi) \cong \mathrm{KP}_2(M)$ if and only if the matroid induced by (the columns of) M does not contain F_7^*, R_{10}, or $\mathcal{M}(K_5)^*$ as a minor.

The proof of Theorem 5.11 relies on a result of Barahona and Grötschel [7, Theorem 3.5]. In their notation, Barahona and Grötschel argued that $P(\mathcal{M})$ (the cycle polytope of a binary matroid $\mathcal{M} = (E, \mathcal{I})$) is equal to the polytope

$$\left\{x \in [0,1]^E \;\middle|\; x(F) - x(C \setminus F) \le |F| - 1 \;\forall \text{ cocircuits } C \subseteq E \text{ and all } F \subseteq C, |F| \text{ odd}\right\}$$

if and only if $\mathcal{M}$ fulfills the *sums of circuits* property investigated by Seymour [49, chapter 16], which Seymour in turn proved to be equivalent to $\mathcal{M}$ not containing the three minors F_7^*, R_{10}, or $\mathcal{M}(K_5)^*$. We call the non-trivial inequalities of $P(\mathcal{M})$ the *cocircuit inequalities*. The sums of circuits property has also been discussed by Goddyn [27] and some of the references therein.

Now let $\mathcal{M}$ be the matroid induced by the columns of M from Theorem 5.11, then the description by cocircuit inequalities can easily be seen to be isomorphic to

$$Q(\Pi) := \left\{ y \in \mathbb{A}(\Pi) \;\middle|\; y \geq \mathbb{0} \text{ and } \sum_{i \in F} y^i_{\omega_i} - \sum_{i \in C \setminus F} y^i_{\omega_i} \leq |F| - 1 \quad \begin{matrix} \forall C \subseteq [n] \text{ cocircuit}, \\ F \subseteq C, |F| \text{ odd} \end{matrix} \right\},$$

where $\omega_i \in \Xi_i$ are the same as in the theorem statement.

Since by Lemma 5.9 the linear form relaxation $\mathrm{R}_1(\Pi)$ is equal to the polytope described by the odd-hyperplane inequalities for $\mathrm{CTP}(\Pi)$, and $\mathrm{CTP}(\Pi)$ is isomorphic to $P(\mathcal{M})$ by Theorem 3.7, we only need to show that $Q(\Pi)$ is the set of all points that satisfy the odd-hyperplane inequalities, i. e., $Q(\Pi) = \mathrm{R}_1(\Pi)$.

Proof of Theorem 5.11: To show that $Q(\Pi) \supseteq \mathrm{R}_1(\Pi)$, let $C \subseteq [n]$ be a cocircuit of the matroid $\mathcal{M}$ and $F \subseteq C$ with $|F|$ odd. We have to show that the inequalities

$$\sum_{i \in F} y^i_{\omega_i} - \sum_{i \in C \setminus F} y^i_{\omega_i} \leq |F| - 1$$

for all such C and F are valid for $\mathrm{R}_1(\Pi)$.

Note that $C \subseteq [n]$ being a cocircuit of $\mathcal{M}$ implies that there exists some multiplier $\lambda \in \mathbb{F}_2^d$ such that $\lambda^\top M = \chi(C)$, where M is the aforementioned matrix and $\chi(C)$ is the characteristic vector of C.

With the associated linear form $f_\lambda(\omega) = \langle \lambda, \omega \rangle$, we know that the odd-hyperplane inequality

$$\sum_{i \notin F} \sum_{\omega \in \Xi_i \setminus \ker(f_\lambda)} y^i_\omega + \sum_{i \in F} \sum_{\omega \in \Xi_i \cap \ker(f_\lambda)} y^i_\omega \geq 1$$

is valid for $\mathrm{R}_1(\Pi)$ by Lemma 5.9.

Now since $\Xi_i = \{\mathbb{0}, \omega_i\}$, and ω_i is an element of $\ker(f_\lambda)$ if and only if $\langle \lambda, \omega_i \rangle = 0$, which is true if and only if $i \notin C$ (remember that $\omega_i \neq \mathbb{0}$ are the columns of M), the first double-sum of the left-hand side of this inequality can be written as $\sum_{\substack{i \notin F \\ i \in C}} y^i_{\omega_i}$, while the second double sum can be decomposed into $\sum_{i \in F} y^i_{\mathbb{0}} + \sum_{\substack{i \in F \\ i \notin C}} y^i_{\omega_i}$. The index set of the last sum is actually empty because $F \subseteq C$, so we transform the sum further. This leads to the following complete chain of equivalences:

$$\begin{aligned}
& \sum_{i \notin F} \sum_{\omega \in \Xi_i \setminus \ker(f_\lambda)} y^i_\omega + \sum_{i \in F} \sum_{\omega \in \Xi_i \cap \ker(f_\lambda)} y^i_\omega && \geq 1 \\
\underset{\text{as described}}{\iff} \quad & \sum_{\substack{i \notin F \\ i \in C}} y^i_{\omega_i} + \sum_{i \in F} y^i_{\mathbb{0}} + \sum_{\substack{i \in F \\ i \notin C}} y^i_{\omega_i} && \geq 1 \\
\underset{F \subseteq C}{\iff} \quad & \sum_{i \in C \setminus F} y^i_{\omega_i} + \sum_{i \in F} y^i_{\mathbb{0}} && \geq 1 \\
\underset{y^i_{\mathbb{0}} = (1 - y^i_{\omega_i})}{\iff} \quad & \sum_{i \in C \setminus F} y^i_{\omega_i} + \sum_{i \in F} (1 - y^i_{\omega_i}) && \geq 1 \\
\underset{\text{rearrangement of terms}}{\iff} \quad & \sum_{i \in F} y^i_{\omega_i} - \sum_{i \in C \setminus F} y^i_{\omega_i} && \leq |F| - 1.
\end{aligned}$$

This shows that $\sum_{i \in F} y^i_{\omega_i} - \sum_{i \in C \setminus F} y^i_{\omega_i} \leq |F| - 1$ is valid for $\mathrm{R}_1(\Pi)$, which completes this direction of the proof.

For the other direction, namely $Q(\Pi) \subseteq \mathrm{R}_1(\Pi)$, let $I \subseteq [n]$ with $|I|$ odd be an index set and $H \subseteq \mathbb{F}_2^d$ be a hyperplane. We need to show that the odd-hyperplane inequality

$$\sum_{i \notin I} \sum_{\omega \Xi_i \setminus H} y^i_\omega + \sum_{i \in I} \sum_{\omega \in \Xi_i \cap H} y^i_\omega \geq 1$$

is valid for $Q(\Pi)$. To do that, let $h \in \mathbb{F}_2^d$ be such that $H = \{\omega \in \mathbb{F}_2^d \mid \langle h, \omega \rangle = 0\}$, and let $C \subseteq [n]$ be the cocircuit with $\chi(C) = h^\top M$, that is, $i \in C$ if and only if $\langle h, \omega_i \rangle = 1$, which is true if and only if $\omega_i \notin H$.

Now the odd-hyperplane inequality is written as

$$\sum_{\substack{i \notin I \\ \omega_i \notin H}} y^i_{\omega_i} + \sum_{i \in I} y^i_{\mathbb{0}} + \sum_{\substack{i \in I \\ \omega_i \in H}} y^i_{\omega_i} \geq 1,$$

again because $\Xi_i = \{\mathbb{0}, \omega_i\}$, similar to the first proof direction. This inequality is further reformulated into

$$\sum_{i \in C \setminus I} y^i_{\omega_i} + \underbrace{\sum_{i \in I \setminus C} \underbrace{(y^i_{\mathbb{0}} + y^i_{\omega_i})}_{=1}}_{=|I \setminus C|} + \underbrace{\sum_{i \in I \cap C} y^i_{\mathbb{0}}}_{|I \cap C| - \sum_{i \in I \cap C} y^i_{\omega_i}} \geq 1,$$

and plugging in the expressions as well as rearranging the terms results in the inequality

$$\sum_{i \in C \cap I} y^i_{\omega_i} - \sum_{i \in C \setminus I} y^i_{\omega_i} \leq |I| - 1.$$

This final inequality holds for $Q(\Pi)$ if $I \subseteq C$, since then I is a subset of a cocircuit with odd cardinality. Otherwise, $I \nsubseteq C$ implies that $|I| > |C \cap I|$, and the inequality is therefore already implied by $\mathbb{0} \leq y \leq \mathbb{1}$, since the left-hand side cannot get larger than $|C \cap I|$. ■

This concludes our peek into the linear form relaxation. For the second projection relaxation, we are in for a little surprise:

5.12 Theorem
For any block configuration Π*, the second projection relaxation* $\mathrm{R}_2(\Pi)$ *is equal to the linear form relaxation:* $\mathrm{R}_2(\Pi) = \mathrm{R}_1(\Pi)$.

Proof: Lemma 5.8 already shows that $\mathrm{R}_2(\Pi) \subseteq \mathrm{R}_1(\Pi)$ holds. For the other direction, we prove that every inequality that defines $\mathrm{R}_2(\Pi)$ is already valid for $\mathrm{R}_1(\Pi)$.

Without loss of generality, we assume that the second projection relaxation $\mathrm{R}_2(\Pi)$ is the intersection of the projection relaxations $\mathrm{R}_f(\Pi)$ with respect only to surjective linear maps $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^2$ instead of all linear maps $\tilde{f}\colon \mathbb{F}_2^d \to \mathbb{F}_2^{\tilde{d}}$ with $\dim \operatorname{im}(\tilde{f}) = 2$ by concatenation of $\tilde{f}$ with an isomorphism from $\operatorname{im}(\tilde{f})$ to $\mathbb{F}_2^2$. Let $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^2$ be any such surjective linear map. The projection relaxation with respect to f is described as

$$\mathrm{R}_f(\Pi) = \kappa_f^{-1}(\mathrm{CTP}(f(\Pi)))_+ = \mathrm{C}(\Pi) \cap \left\{ y \in \mathbb{A}(\Pi) \;\middle|\; \begin{array}{l} \displaystyle\sum_{\substack{\omega\in\Xi_i\\ f(\omega)=\sigma}} y_\omega^i = x_\sigma^i \text{ for all } i \in [n], \\ x \in \mathrm{CTP}(f(\Pi)) \end{array} \right\},$$

where the projected cyclic transversal polytope $\mathrm{CTP}(f(\Pi))$ is isomorphic to a face of $\mathrm{CTP}(2,n)$ and therefore a description of $\mathrm{CTP}(f(\Pi))$ by inequalities is obtained in the following way: Since $\mathrm{CTP}(2,n)$ is described by block equations, non-negativity and odd-hyperplane inequalities (as seen in Theorem 4.8), we know that a description of $\mathrm{CTP}(f(\Pi))$ as its face consists of these inequalities and additional equations of the form $x_\sigma^i = 0$ for all $i \in [n]$ and $\sigma \in \mathbb{F}_2^2 \setminus \Xi_i$.

For any hyperplane $H \subsetneq \mathbb{F}_2^d$ and any $I \subseteq [n]$ of odd cardinality, plugging the equation defining x_σ^i into the odd-hyperplane inequality $[H, I]$ gives us

$$\sum_{i\in I}\sum_{\sigma\in H}\sum_{\substack{\omega\in\Xi_i\\ f(\omega)=\sigma}} y_\omega^i + \sum_{i\notin I}\sum_{\sigma\notin H}\sum_{\substack{\omega\in\Xi_i\\ f(\omega)=\sigma}} y_\omega^i \geq 1,$$

which by observation is simplified to

$$\sum_{i\in I}\sum_{\omega\in f^{-1}(H)\cap\Xi_i} y_\omega^i + \sum_{i\notin I}\sum_{\omega\in f^{-1}(\mathbb{F}_2^d\setminus H)\cap\Xi_i} y_\omega^i \geq 1.$$

Since $f^{-1}(\mathbb{F}_2^d \setminus H) = \mathbb{F}_2^d \setminus f^{-1}(H)$, this inequality has the same format as the odd-hyperplane inequalities, as the second sum is reformulated to sum over $\omega \in \Xi_i \setminus f^{-1}(H)$. What remains to be shown is that

$$f^{-1}(H) = \left\{\omega \in \mathbb{F}_2^d \;\middle|\; f(\omega) \in H\right\}$$

is actually a hyperplane of $\mathbb{F}_2^d$.

It is clear that $f^{-1}(H)$ is some linear subspace because H is a linear subspace and f is linear. We know that $f^{-1}(H)$ is not equal to $\mathbb{F}_2^d$, since f is assumed to be surjective and H is not equal to $\mathbb{F}_2^2$. We further know that $\ker(f) \subseteq f^{-1}(H)$, since $\mathbb{0}_2 \in H$ holds. Now since H is a hyperplane in $\mathbb{F}_2^2$, i. e., a line, there is exactly one non-zero vector in H, so $f^{-1}(H)$ is also not equal to $\ker(f)$. But since it is well known from linear algebra that the dimension formula

$$d = \dim \operatorname{im}(f) + \dim \ker(f)$$

holds and $\dim \operatorname{im}(f) = 2$ is true by our initial assumption that f is surjective, we deduce that the dimension of $f^{-1}(H)$ is $d-1$, since it can neither be $d-2$ (which is the dimension of the kernel) nor can it be d. Therefore, $f^{-1}(H)$ is itself a hyperplane, and this shows that all non-trivial inequalities for $\mathrm{R}_2(\Pi)$ are odd-hyperplane inequalities, which are already valid for $\mathrm{R}_1(\Pi)$ by Lemma 5.9. ■

In contrast to the first two levels of our hierarchy, we know that $\mathrm{R}_k(\Pi)$ for $k \geq 3$ is not equal to the linear form relaxation in general. This is already evident by the fact that $\mathrm{R}_3(3,n) = \mathrm{CTP}(3,n)$ for all $n \in \mathbb{N}$, but the odd-hyperplane inequalities do not suffice to describe $\mathrm{CTP}(3,n)$ completely for $n \geq 3$, as seen in Chapter 4.

Theorem 5.12 does not say which projection relaxations are actually necessary to fully describe $\mathrm{R}_2(\Pi)$, but since there are in general more surjective linear maps of the form $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^2$ than of the form $f\colon \mathbb{F}_2^d \to \mathbb{F}_2$, it is interesting to ask the following questions:

5.13 Question
How many projection maps are necessary to describe $\mathrm{R}_k(\Pi)$ as an intersection of projection relaxations? How much weaker is the intersection of all $\mathcal{O}(d^k)$ many projection relaxations with respect to k-coordinate projections, compared to $\mathrm{R}_k(\Pi)$?

5.2 An Arc-Based Extended Formulation for CTPs

As seen in Chapter 3, cyclic transversal polytopes include, among others, cut polytopes and stable set polytopes. These polytopes have exponentially many facets compared to their dimension, making it generally impossible to describe them in their ambient space using a small amount of linear inequalities, i. e., a number of inequalities polynomial in their dimension. Then again, there are many examples of polytopes in combinatorial optimization that can be identified as projections of other polytopes which in turn may have much fewer facets. This has been done for polygons [25], spanning tree polytopes [38] and parity polytopes [13] among many others [16] and leads us exactly to the concept of an extended formulation of a polytope.

Based on the work of Carr and Konjevod [13] for parity polytopes, we develop a theory of extended formulations for cyclic transversal polytopes based on directed graphs in this section. To do that, we will define what an extended formulation is and then reiterate the extended formulation construction for parity polytopes, where we pick up the needed tools to generalize the idea behind this construction afterwards.

5.2.1 The Carr-Konjevod Extended Formulation for Parity Polytopes

Generally, given a polyhedron $P \subseteq \mathbb{R}^n$, if there exists another polyhedron $Q \in \mathbb{R}^m$ and a projection $\pi\colon \mathbb{R}^m \to \mathbb{R}^n$ such that $\pi(Q) = P$, we call the pair (Q, π) and *extension* of P. If Q in this context is given via an outer description of linear inequalities and equations, we call this description an *extended formulation* of P. The number of facets of the polytope Q is called the *size* of an extension and, respectively, an extended formulation. Since every extended formulation needs at least as many inequalities as its size, it is crucial to consider this quantity of an extended formulation and control it. We call an extension of P with size that is polynomial in the dimension of P a *small* extension and the corresponding description a *small* extended formulation, and will use words like *big* otherwise. In contrast to the results mentioned before considering polytopes possessing such small extensions, lower bounds on the sizes of extensions of well-known families of polytopes have been considered by Fiorini et al. [24], Rothvoss [46], Kaibel and Weltge [32], and Yannakakis [52], just to name a few authors and works. They show that various other families of polytopes do not have small extended formulations.

There are many techniques to obtain extended formulations for certain interesting families of polytopes, like disjunctive programming [4, 5, 6] or the aggregation of higher-order relations between variables into new auxiliary variables, for example for the spanning tree polytope of a graph [38]. The extended formulation for the parity polytope by Carr and Konjevod [13] is motivated by dynamic programming, similar to a well-known extension for the knapsack polytope, where we record intermediate results like partial sums in auxiliary variables. This is aptly represented by a directed graph structure and arc variables, as we will explain in the rest of this section.

The number of facets of the parity polytope $\mathrm{PAR}(n)$ is exponential in n, and an irreducible description using linear inequalities was given by Jeroslow [31] using the odd-set inequalities defined in the introduction, namely

$$\sum_{i \in S} x_i - \sum_{i \in [n]-S} x_i \leq |S| - 1 \text{ for } S \subseteq [n] \text{ where } |S| \text{ is odd},$$

and the "trivial" inequalities $\mathbb{0} \leq x \leq \mathbb{1}$.

But one can express the parity polytope as the projection of a polytope with a much smaller number of facets, by relating the vertices of $\mathrm{PAR}(n)$ and the directed s-t-path polytope of an appropriately defined directed acyclic graph whose nodes express the parity of initial partial sums of $\langle \mathbb{1}, x\rangle$. This gives rise to an extension of $\mathrm{PAR}(n)$, and consequently also an extended formulation by describing the directed s-t-path polytope with inequalities.

5.14 Definition (Directed s-t-path polytope [cf. 15, Section 4.3.2])
Let $D = (V, A)$ be a directed graph with a source s and a sink t. The (*directed*) *s-t-path polytope* $P(D, s, t)$ is defined as the convex hull of characteristic vectors of directed s-t-paths of D. It is described by

$$P(D,s,t) := \left\{ y \in [0,1]^A \;\middle|\; \begin{aligned} 1 &= \sum_{a \in \delta^{\mathrm{out}}(s)} y_a - \sum_{a \in \delta^{\mathrm{in}}(s)} y_a \\ -1 &= \sum_{a \in \delta^{\mathrm{out}}(t)} y_a - \sum_{a \in \delta^{\mathrm{in}}(t)} y_a \\ \sum_{a \in \delta^{\mathrm{out}}(v)} y_a &= \sum_{a \in \delta^{\mathrm{in}}(v)} y_a \;\forall v \in V \setminus \{s,t\} \end{aligned} \right\}.$$

If s and t are clear from context, we omit them and write $P(D)$.

In this section we only write "s-t-path" for "directed s-t-path", since we do not consider undirected paths and their polytopes. The equation system in $P(D, s, t)$ can also be written as $A_D y = b$, where A_D is the incidence matrix of D, and b is the vector satisfying $b_s = 1$, $b_t = -1$ and $b_v = 0$ for all $v \in V \setminus \{s, t\}$, and the only inequalities are the box constraints $\mathbb{0} \leq y \leq \mathbb{1}$, meaning that $P(D)$ is described by $2|A|$ inequalities in general. The equations are called *source* constraint, *sink* constraint and *flow conservation* constraints, respectively, as the convex combination of characteristic vectors of s-t-paths of a digraph is also called a unit-*flow* on this digraph.

5.15 Example (Extended formulation for $\mathrm{PAR}(n)$ [13])
Consider the directed acyclic graph $D(1, n) = (V(1, n), A(1, n))$ with

$$V(1,n) := \{0,1\} \times [n]_0$$

and

$$A(1,n) := \{((p, j-1), (q, j)) \mid p, q \in \{0,1\}, j \in [n]\}.$$

Then, relabeling $s = (0, 0)$ and $t = (0, n)$, the convex hull of the characteristic vectors of s-t-paths in $D(1, n)$, which is equal to $P(D(1, n))$, together with the affine projection map $\pi \colon \mathbb{R}^{A(1,n)} \to \mathbb{R}^n$,

$$\pi(y)_j := y_{((0,j-1),(1,j))} + y_{((1,j-1),(0,j))} = 1 - y_{((0,j-1),(0,j))} - y_{((1,j-1),(1,j))},$$

defines an extension of the parity polytope $\mathrm{PAR}(n)$, as shown by Carr and Konjevod [13].

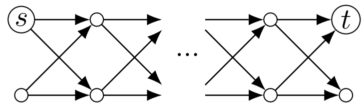

Figure 5.1: Schematic for the acyclic digraph $D(1, n)$

It is no problem to remove the nodes $(1, 0)$ and $(1, n)$ and any arcs incident to them from $D(1, n)$ to make the source and sink unique. By counting the inequalities needed for the s-t-path polytope of $D(1, n)$, we also obtain a proof for the size of this extension:

5.16 Proposition (Size of the extension of $\mathrm{PAR}(n)$ [13])
There exists an extension of size $\mathcal{O}(n)$ for the parity polytope $\mathrm{PAR}(n)$.

Proof: Since the s-t-path polytope of a digraph is defined as the intersection of an affine subspace with the 0/1-cube, the only inequalities needed are the ones defining the cube. As the number of these inequalities is purely dependent on the number of arcs in the given digraph, it suffices to count that $D(1, n) = (V(1, n), A(1, n))$ contains $\mathcal{O}(n)$ arcs. There are 4 arcs per layer $j \in [n]$, which means that we need $8n \in \mathcal{O}(n)$ inequalities to describe $P(D(1, n))$ in its affine space, which completes the proof. ■

5.2.2 Generalization of the Carr-Konjevod EF to CTPs

Cyclic transversal polytopes generalize parity polytopes. Thus, we want to find an extended formulation for cyclic transversal polytopes that can be derived in similar fashion to the one obtained for parity polytopes. Nevertheless, there are some pitfalls when adopting the dynamic programming idea in a straightforward way.

These concern firstly the possible asymmetry of the blocks, which is not existent in the parity polytope, since the parity polytope $\mathrm{PAR}(n)$ is isomorphic to the full cyclic transversal polytope $\mathrm{CTP}(1, n)$. Therefore, we will initially concentrate on the symmetric case of full cyclic transversal polytopes and deal with other cyclic transversal polytopes later, as their extensions can be deduced via restriction to subspaces of the ambient space of an extension of a suitable full cyclic transversal polytope.

Secondly, another concern is the resulting size of such an extension. A large size is expected because vertices of cyclic transversal polytopes represent solution sets of hard problems in general, as we have seen in Chapter 3. But the existence of a small and exact extended formulation for cyclic transversal polytopes would more directly contradict other known results in complexity theory: By Corollary 3.16, matching polytopes are a subclass of cyclic transversal polytopes, but Rothvoss [46] showed that these polytopes cannot have a polynomial-size extended formulation in general. Therefore, we can certainly not expect for cyclic transversal polytopes to have a small and exact extended formulation in general. To counterbalance that, we ignore one of these properties here and construct an extended formulation for cyclic transversal polytopes that has exponential size. Then we deal with extensions of relaxations of cyclic transversal polytopes that we defined, and suggest possible ways to strengthen these constructions for future work.

In the definition of $D(1,n) = (V(1,n), A(1,n))$ for the extension of the parity polytope $\mathrm{PAR}(n)$, we chose the node set to be the Cartesian product of the possible (parities of) partial sums $\{0,1\} = \mathbb{F}_2^1$ together with an index of which blocks have contributed to this partial sum. As a generalization, we define $D(d,n) = (V(d,n), A(d,n))$ similarly:

5.17 Definition ((Directed) Transversal network)
For any $d, n \in \mathbb{N}$, the *(directed) transversal network* $D(d,n) = (V(d,n), A(d,n))$ consists of the vertex set

$$V(d,n) := \mathbb{F}_2^d \times [n]_0$$

and the arc set

$$A(d,n) := \big\{((\sigma, i-1), (\theta, i)) \;\big|\; \sigma, \theta \in \mathbb{F}_2^d, i \in [n]\big\}.$$

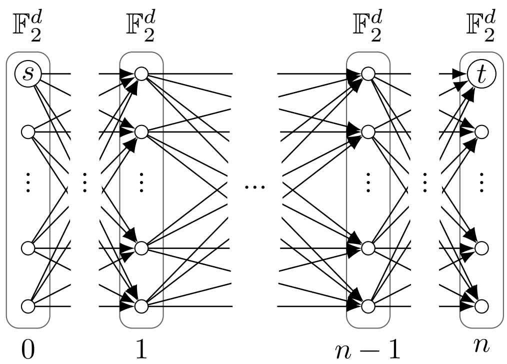

Figure 5.2: Schematic of the transversal network $D(d,n)$

A visualization of $D(d,n)$ is shown in Figure 5.2. It is straightforward to see that the s-t-path polytope $P(D(d,n), (\mathbb{0}, 0), (\mathbb{0}, n))$ of this transversal network, with $s = (\mathbb{0}, 0)$ and $t = (\mathbb{0}, n)$, is a candidate for an extension of $\mathrm{CTP}(d,n)$, which we actually prove in Theorem 5.19. Compared to the special case in Example 5.15, the projection will be a bit different though, to account for the differences between full cyclic transversal polytopes and parity polytopes:

5.18 Definition (Transversal network projection)
The transversal network projection $\pi_{(d,n)}$ for parameters $d, n \in \mathbb{N}$ is defined as the affine map $\pi_{(d,n)} \colon P(D(d,n)) \to \mathbb{A}(d,n)$ with

$$\pi_{(d,n)}(y)^i_\omega := \sum_{\substack{\sigma,\theta \in \mathbb{F}_2^d \\ \omega = \sigma \oplus \theta}} y_{((\sigma, i-1),(\theta, i))}.$$

By means of the flow conservation constraints, is easy to check that for all $y \in P(D(d,n))$, we actually have that $\pi_{(d,n)}(y)$ lies in $\mathbb{A}(d,n)$: The source constraint for $D(d,n)$ and the source $s = (\mathbb{0}, 0)$ results in the equation

$$1 = \sum_{\theta \in \mathbb{F}_2^d} y_{((\mathbb{0},0),(\theta,1))},$$

and since the vertices $(\sigma, 0)$ with $\sigma \in \mathbb{F}_2^d \setminus \mathbb{0}$ do not have incoming arcs, their flow conservation constraints evaluate to $0 = \sum_{\theta \in \mathbb{F}_2^d} y_{((\sigma,0),(\theta,1))}$. This means we sum over all vertices of the form $(\theta, 1)$ with $\theta \in \mathbb{F}_2^d$ and repeatedly apply the flow conservation constraints to see that

$$1 = \sum_{\sigma,\theta \in \mathbb{F}_2^d} y_{((\sigma,0),(\theta,1))} = \sum_{\sigma,\theta \in \mathbb{F}_2^d} y_{((\sigma,i-1),(\theta,i))} \qquad \text{for any } i \in [n]$$

holds, which shows that the block equation $\sum_{\omega \in \mathbb{F}_2^d} \pi_{(d,n)}(y)^i_\omega = 1$ for $i \in [n]$ is satisfied.

5.19 Theorem (Extension of $\mathrm{CTP}(d,n)$)
The s-t-path polytope $P(D(d,n))$ of the directed transversal network $D(d,n)$ with $s = (\mathbb{0}, 0)$ and $t = (\mathbb{0}, n)$, together with the transversal network projection $\pi_{(d,n)}$, is an extension of size $\mathcal{O}(n \cdot 4^d)$ of the full cyclic transversal polytope $\mathrm{CTP}(d,n)$.

Proof: We have to show that $\pi_{(d,n)}(P(D(d,n)))$ is equal to $\mathrm{CTP}(d,n)$. Since $\pi_{(d,n)}$ is an affine map, it suffices to show that $\pi_{(d,n)}$ maps the vertices of $P(D(d,n))$ to those of $\mathrm{CTP}(d,n)$ and every vertex of $\mathrm{CTP}(d,n)$ has a preimage in $P(D(d,n))$, since convex sums of these vertices distribute under the affine map.

For the first inclusion showing that $\pi_{(d,n)}(P(D(d,n)))$ is contained in $\mathrm{CTP}(d,n)$, let $v \in P(D(d,n))$ be a vertex of the s-t-path polytope of $D(d,n)$ with $s = (\mathbb{0}, 0)$ and $t = (\mathbb{0}, n)$. Then v corresponds to some s-t-path of $D(d,n)$ by definition. Let $\{(\sigma_0, 0), (\sigma_1, 1), (\sigma_2, 2), \dots, (\sigma_{n-1}, n-1), (\sigma_n, n)\}$ be the nodes of this directed path with $\sigma_0 = \sigma_n = \mathbb{0}$. By definition of this vertex we know that $v_{((\sigma_{i-1},i-1),(\sigma_i,i))} = 1$ for $i \in [n]$ and $v_a = 0$ for all other arcs $a \in A(d,n) \setminus \{((\sigma_{i-1}, i-1), (\sigma_i, i)) \mid i \in [n]\}$, and therefore the network projection $\pi_{(d,n)}$ maps v to a vertex of $\mathrm{CTP}(d,n)$ as follows: Defining, for all $i \in [n]$, the element $\xi(i) = \sigma_{i-1} \oplus \sigma_i \in \mathbb{F}_2^d$, we have that $\pi_{(d,n)}(v)^i_\omega = 0$ for all $\omega \in \mathbb{F}_2^d \setminus \xi(i)$, and $\pi_{(d,n)}(v)^i_{\xi(i)} = v_{((\sigma_{i-1},i-1),(\sigma_i,i))} = 1$ for all $i \in [n]$, i. e., the image of v under $\pi_{(d,n)}$ is the vector in $\mathbb{A}(d,n)$ belonging to the transversal $(\xi(1), \dots, \xi(n))$, which is cyclic because $\sum_{i \in [n]} \xi(i) = \sum_{i \in [n]} \sigma_{i-1} \oplus \sigma_i = \sigma_0 \oplus \sigma_n = \mathbb{0}$, which by definition is a vertex of $\mathrm{CTP}(d,n)$.

For the other inclusion, we take a vertex y of $\mathrm{CTP}(d,n)$ and its associated cyclic transversal $(\xi(1), \dots, \xi(n))$. We construct an s-t-path in $D(d,n)$ using this transversal by running through the nodes

$$\left\{(\mathbb{0}, 0), (\xi(1), 1), (\xi(1) \oplus \xi(2), 2), \dots, (\sum_{i \in [n-1]} \xi(i), n-1), (\mathbb{0}, n)\right\},$$

where the first part of the i-th node is the partial sum of the first i block elements of the given cyclic transversal. This path clearly exists in $D(d,n)$, and thus it corresponds to a vertex v of $P(D(d,n))$. The previous part of the proof then readily shows that $\pi_{(d,n)}(v) = y$ by observing that $\xi(k) = \sum_{i\in[k-1]} \xi(i) \oplus \sum_{j\in[k]} \xi(j)$. Thus, the pair $\left(P(D(d,n)), \pi_{(d,n)}\right)$ is an extension of $\mathrm{CTP}(d,n)$.

To prove that this extension is of the claimed size, we count the number of inequalities of $P(D(d,n))$, which only depends on $|A(d,n)|$ and is rather straightforward: For every layer $i \in [n]$, we have $2^{2d} = 4^d$ arcs, one for every pair $(\sigma,\theta) \in \mathbb{F}_2^d \times \mathbb{F}_2^d$. Hence, the extension is of size $2n \cdot 4^d \in \mathcal{O}(n \cdot 4^d)$ by counting two inequalities for every arc. ■

Since an outer description by inequalities is readily available for these s-t-path polytopes of $D(d,n)$, they constitute extended formulations for $\mathrm{CTP}(d,n)$, which are small if d is constant.

We can theoretically improve the number of necessary inequalities from $2|A(d,n)|$ to $|A(d,n)|$ in an irreducible description of $P(D(d,n))$, since the constraint $y_a \leq 1$ for any arc $((\omega_{k-1}, k-1), (\omega_k, k)) \in A(d,n)$ is recovered by subtraction of the appropriate non-negativity constraints $y_a \geq 0$ with $a \in \left\{((\sigma, k-1), (\theta, k)) \;\middle|\; \sigma, \theta \in \mathbb{F}_2^d\right\}$ and $a \neq ((\omega_{k-1}, k-1), (\omega_k, k))$ from the equation $1 = \sum_{\sigma,\theta\in\mathbb{F}_2^d} y_{((\sigma,k-1),(\theta,k))}$, which was obtained in the calculation before Theorem 5.19.

The definition of the transversal network projection $\pi_{(d,n)}$ as a sum of certain coefficients and its use in the proof of Theorem 5.19 indicates a natural perspective from which to view this extension: We can imagine the arcs of $D(d,n)$ to be labelled by elements $\omega \in \mathbb{F}_2^d$, depending on the difference of the first element of their head and tail nodes, i. e., the arc $((\sigma, i-1)(\theta, i))$ receives the label $\omega = \sigma \oplus \theta$, since the sum and difference are identical over $\mathbb{F}_2^d$, as shown in Figure 5.3. Then identifying all arcs with the same such label and summing their coefficients leads to this particular transversal network projection $\pi_{(d,n)}$.

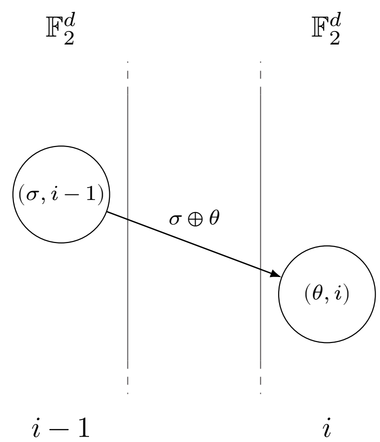

Figure 5.3: A labelled arc of $D(d,n)$

Now, if we are given any block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$ not necessarily consisting of full blocks, we can imagine removing certain arcs from the directed transversal networks, namely those arcs $((\sigma, i-1)(\theta, i))$ whose label $\omega = \sigma \oplus \theta$ fulfills $\omega \in \mathbb{F}_2^d \setminus \Xi_i$. The resulting s-t-path polytope of the restricted transversal network, denoted by $P(D(\Pi))$, and the corresponding restricted transversal network projection, denoted by π_Π, are then also an extension of $\mathrm{CTP}(\Pi)$, and the proof is analogous to Theorem 5.19.

5.2.3 An Extended Formulation for Relaxations of CTPs

In this section, we combine the concept of an extended formulation for a cyclic transversal polytope with the concept of a relaxation of a (full) cyclic transversal polytope from the beginning of this chapter.

Of course, for any linear map $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^{\tilde{d}}$ with $\tilde{d} \le d$, the existence of a network extension for the mapped cyclic transversal polytope $\mathrm{CTP}(f(\Pi))$ for any block configuration Π is clear from Theorem 5.19 and the proximate discussion on arc labels. Since f induces a map $\kappa_f\colon \mathbb{A}(\Pi) \to \mathbb{A}(f(\Pi))$ as shown in Proposition 5.1, we now combine the transversal network projection $\pi_{f(\Pi)}$ with the projection map κ_f and conclude that for all $y \in \mathrm{CTP}(\Pi)$, there exists a point $z \in P(D(f(\Pi)))$ such that the condition

$$\kappa_f(y) = \pi_{f(\Pi)}(z)$$

holds. Conditions of this form are called *consistency conditions*. Note that both $\kappa_f(y)$ for any $y \in \mathrm{CTP}(\Pi)$ and $\pi_{f(\Pi)}(z)$ for any $z \in P(D(f(\Pi)))$ are elements of $\mathrm{CTP}(f(\Pi))$, but the equation does not reference any point in $\mathrm{CTP}(f(\Pi))$ explicitly. Rather, we obtain a condition between the network extension of $\mathrm{CTP}(f(\Pi))$ and the original cyclic transversal polytope $\mathrm{CTP}(\Pi)$.

Taking multiple such linear maps f_i for $i \in [r]$, we get multiple conditions of the above form, namely that for all $y \in \mathrm{CTP}(\Pi)$, there exist $z_i \in P(D(f_i(\Pi)))$, such that

$$\kappa_{f_i}(y) = \pi_{f_i(\Pi)}(z_i) \quad \text{for all } i \in [r]$$

holds simultaneously. This is akin to intersecting the projection relaxations with respect to f_i to obtain a better approximation for $\mathrm{CTP}(\Pi)$, compared to only using a single map f. Now, we already know that the image of $P(D(f_i(\Pi)))$ under $\pi_{f_i(\Pi)}$ is equal to $\mathrm{CTP}(f_i(\Pi))$, while the image of $\mathrm{CTP}(\Pi)$ under κ_{f_i} is only a certain subset of $\mathrm{CTP}(f_i(\Pi))$ in general. In fact, by applying the preimage $\kappa_{f_i}^{-1}$ as in Definition 5.4, we have proven the following:

5.20 Proposition
Let $\mathrm{Epi}(\mathbb{F}_2^d, \mathbb{F}_2^k)$ *be the set of surjective linear maps* $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^k$. *Then the projection relaxation* $\mathrm{R}_k(\Pi)$ *of dimension* k *of any block configuration* Π *is described as*

$$\Big\{y \in \mathrm{C}(\Pi) \;\Big|\; \textit{For all } f \in \mathrm{Epi}\big(\mathbb{F}_2^d, \mathbb{F}_2^k\big),\ \textit{there is } z_f \in P(D(f(\Pi)))\ \textit{with } \kappa_f(y) = \pi_{f(\Pi)}\big(z_f\big)\Big\}.$$

As a reminder, the set $\mathrm{C}(\Pi)$ is defined solely by the block equations and non-negativity constraints. This yields a description of $\mathrm{R}_k(\Pi)$ by linear inequalities and equations, since κ_f and $\pi_{f(\Pi)}$ are affine linear maps, and a description of all s-t-path polytopes is known by Definition 5.14. Taking the collection of s-t-path polytopes and network transversal maps, this satisfies the definition of an extended formulation, although in an unfamiliar format.

The above description provides an upper bound on the size of an extended formulation of $\mathrm{R}_k(\Pi)$ that depends on the cardinality of $\mathrm{Epi}\big(\mathbb{F}_2^d, \mathbb{F}_2^k\big)$, but we do not claim that every s-t-path polytope, enumerated over the surjective linear maps $f\colon \mathbb{F}_2^d \to \mathbb{F}_2^k$, is necessary in an irreducible description. In fact, we know by Theorem 5.12 that for $k = 2$, some maps are unnecessary to completely describe $\mathrm{R}_2(\Pi)$. Therefore, especially in light of Question 5.13, we cannot give a definitive answer on the size of this extension.

Our aim instead is that the above description with the additional conditions may lead to the discovery of further relations between the integer points of relaxations, such that the addition of these relations to the given description improves upon the relaxation:

5.21 Question
Which additional relations can be identified and described between the integer points of the different s-t-path polytopes of the respective transversal networks?

5.3 A Disjunctive Programming Formulation for CTPs

There are many more possibilities to develop the theory of cyclic transversals and their polytopes, in particular with regard to relaxations and extensions. When compared to the construction in Section 5.2.2, we will outline the starting point of a very different approach to obtain an extended formulation here, using the idea of disjunctive programming which was mentioned in the introduction of Section 5.2.1: If we describe a polytope P as the convex hull of the union of smaller polytopes $P_1, \dots, P_r$ for some $r \in \mathbb{N}$ of which we know an outer description, then there exists an extended formulation for $P = \mathrm{conv}\Big(\bigcup_{i\in[r]} P_i\Big)$ that involves the descriptions of P_i for $i \in [r]$. If we extend this idea by replacing the outer description of P_i for $i \in [r]$ with an extended formulation, we know that the size of an extended formulation for P is bounded above roughly by the number r of smaller polytopes, plus the sum of sizes of extended formulations of P_i for $i \in [r]$. In particular, the extended formulation for P will have small size, given that the extended formulations of P_i for $i \in [r]$ are small. For a proof of an explicit construction of this description for P that was developed by Balas [5], see also the thesis of Pashkovich [44, Theorem 2.1], and for an improvement on the bound of the size of this formulation, see the thesis of Weltge [51, Proposition 3.1.1].

While there are also many possibilities to describe a given polytope as the convex hull of the union of other polytopes (e. g., by choosing the singleton sets of vertices as the other polytopes), we outline here one specific description for cyclic transversal polytopes. Let X be the set of integer vectors $x \in [n]_0^d$ having only even entries. This set X is obviously finite with $|X| = \mathcal{O}\Big({(n/2)}^d\Big)$. Then, given a block configuration $\Pi = (\Xi_1, \dots, \Xi_n)$ over $\mathbb{F}_2^d$, we partition its set of cyclic transversals into sets indexed by elements $x \in X$ as follows: Let every $\xi \in \mathrm{CT}(\Pi)$ be identified with the corresponding tuple $\tilde{\xi}$ of real vectors, i. e., let $\xi(i)$ for $i \in [n]$ be identified with its corresponding 0/1-vector $\tilde{\xi}(i)$ over the real numbers. Then we define $\mathrm{CT}_x(\Pi)$ to be the set of all cyclic transversals $\xi \in \mathrm{CT}(\Pi)$ such that

$$\sum_{i \in [n]} \tilde{\xi}(i) = x$$

is fulfilled.

Clearly, every cyclic transversal $\xi \in \mathrm{CT}(\Pi)$ is an element of $\mathrm{CT}_x(\Pi)$ for some $x \in X$ since the sum $\sum_{i \in [n]} \tilde{\xi}(i)$ has to have even entries, as otherwise $\sum_{i \in [n]} \xi(i)$ would not sum to $\mathbb{0}$ over $\mathbb{F}_2^d$, which is a contradiction to ξ being a cyclic transversal. Furthermore, since the vertices of the cyclic transversal polytope $\mathrm{CTP}(\Pi)$ are identified with the cyclic transversals $\xi \in \mathrm{CT}(\Pi)$, this partitioning by elements $x \in X$ applies to the vertices of $\mathrm{CTP}(\Pi)$ as well. We define $\mathrm{CTP}_x(\Pi)$ to be the convex hull of the vertices that correspond exactly to the cyclic transversals in $\mathrm{CT}_x(\Pi)$ and call such a polytope a *slice* of a cyclic transversal polytope, parametrized by $x \in X$.

It is evident that

$$\mathrm{CTP}(\Pi) = \mathrm{conv}\left(\bigcup_{x \in X} \mathrm{CTP}_x(\Pi)\right)$$

holds, which allows us to apply the idea of disjunctive programming to cyclic transversal polytopes in this specific manner.

5.22 Example (Size of the Disjunctive Programming Formulation for Parity Polytopes) Since parity polytopes $\mathrm{PAR}(n)$ are affinely isomorphic to full cyclic transversal polytopes $\mathrm{CTP}(1, n)$, we will explain how the suggested approach applies to parity polytopes and results in an extended formulation for $\mathrm{PAR}(n)$ that has a size of $\mathcal{O}(n^2)$ at most.

For polytopes $\mathrm{CTP}(1,n)$, whose rank is $d=1$, the set X is equal to the even elements of $[n]_0$, and an element $\xi \in \mathrm{CT}(1,n)$ belongs to $\mathrm{CT}_x(1,n)$ if and only if there are exactly x many indices $i \in [n]$ such that $\xi(i)=1$. That means the slices $\mathrm{CTP}_x(1,n)$ for $x \in X$ are exactly the 0/1-polytopes whose vertices have x many coordinates y_1^i equal to 1. By the affine isomorphism between $\mathrm{CTP}(1,n)$ and $\mathrm{PAR}(n)$, it is easy to see that these slices are (affinely isomorphic to) hypersimplices $\Delta_{n,x}$ that are the convex hull of all 0/1-vectors in $\mathbb{R}^n$ with exactly x entries set to 1, that have been part of the discussion in Section 3.2.2.

An article from Grande, Padrol, and Sanyal [28] shows that the size of an extended formulation of any hypersimplex $\Delta_{n,x}$ is at most $2n$, with a better constant for certain exceptions. Since the number of such slices is equal to $|X|$, which is at most $(n+1)/2$, and each extended formulation of a slice has size at most $2n$, we can estimate the size of the resulting extended formulation for $\mathrm{CTP}(1,n)$ to be at most of the size $\mathcal{O}(n^2)$.

The fact that $\mathrm{CTP}_2(1,4)$ is affinely isomorphic to $\Delta_{4,2}$, i. e., an octahedron, shows that the resulting slices of cyclic transversal polytopes are not necessarily cyclic transversal polytopes themselves.

We suggest that the structure of the slices $\mathrm{CTP}_x(\Pi)$ for $x \in X$ might also be of independent interest, lending themselves to explorative results similar to those about cyclic transversal polytopes in the rest of this thesis. Among other questions, the structural insights about this other extended formulation are condensed into the closing, but open-ended, question, which we will refer to again in the conclusion:

5.23 Question

Aside from a different upper bound on the size of the described disjunctive programming formulation, what are meaningful differences and similarities between this formulation and the network formulation for cyclic transversal polytopes, especially in light of Question 5.21? How do the extended formulations relate to one another?

6 Conclusion

At this point, we have presented several aspects of the cyclic transversal framework and have demonstrated its usefulness and value for investigating and generalizing other well-known families of polytopes. Furthermore, we have made several advancements in understanding cyclic transversal polytopes as well as their relaxations and extended formulations. Therefore, it is time to look back and summarize our results, as well as to look forward and show how cyclic transversals and their polytopes might be valuable for and investigated in future research.

Looking back

Our novel approach of using the linear algebra of $\mathbb{F}_2^d$ to dictate conditions on the coordinates of points, which we explored in Chapter 2, gives ample structure to the investigated cyclic transversal polytopes. We were able to exploit this structure in positive and negative ways, as we have shown how several other classes of polytopes relate to cyclic transversal polytopes in Chapter 3. This is summarized in Figure 6.1 on the next page.

Furthermore, in Chapter 4, we used the symmetries of full cyclic transversal polytopes to prove the validity of the class of odd-hyperplane inequalities, which lead to a complete outer description of full cyclic transversal polytopes of the form $\mathrm{CTP}(2, n)$ with $n \in \mathbb{N}$. We do note that due to our ability to generate some higher-parameter full cyclic transversal polytopes and their inequality orbits using code built upon SageMath [47], we also found two additional inequality orbits for $\mathrm{CTP}(3, 3)$ and $\mathrm{CTP}(3, 4)$ and visualized them in Figures 4.2 and 4.3. While a full characterization of these additional orbits has yet to be determined, we already suspect that their number grows faster than the number of odd-hyperplane inequalities when increasing the order of the cyclic transversals, which strongly emphasizes their importance in understanding cyclic transversal polytopes in general.

Finally, the purpose of Chapter 5 was two-fold: In the first part, we showed how linear maps from $\mathbb{F}_2^d$ to its subspaces gave rise to a relaxation hierarchy of cyclic transversal polytopes, and proved some of the properties that these relaxations possess. This resulted in the complete description of the first two levels of the hierarchy, and an inclusion of a result by Barahona and Grötschel [7] into the cyclic transversal framework. Their result characterizes when the first level of the relaxation hierarchy is equal to the original cyclic transversal polytope, given that the cyclic transversal polytope is already isomorphic to

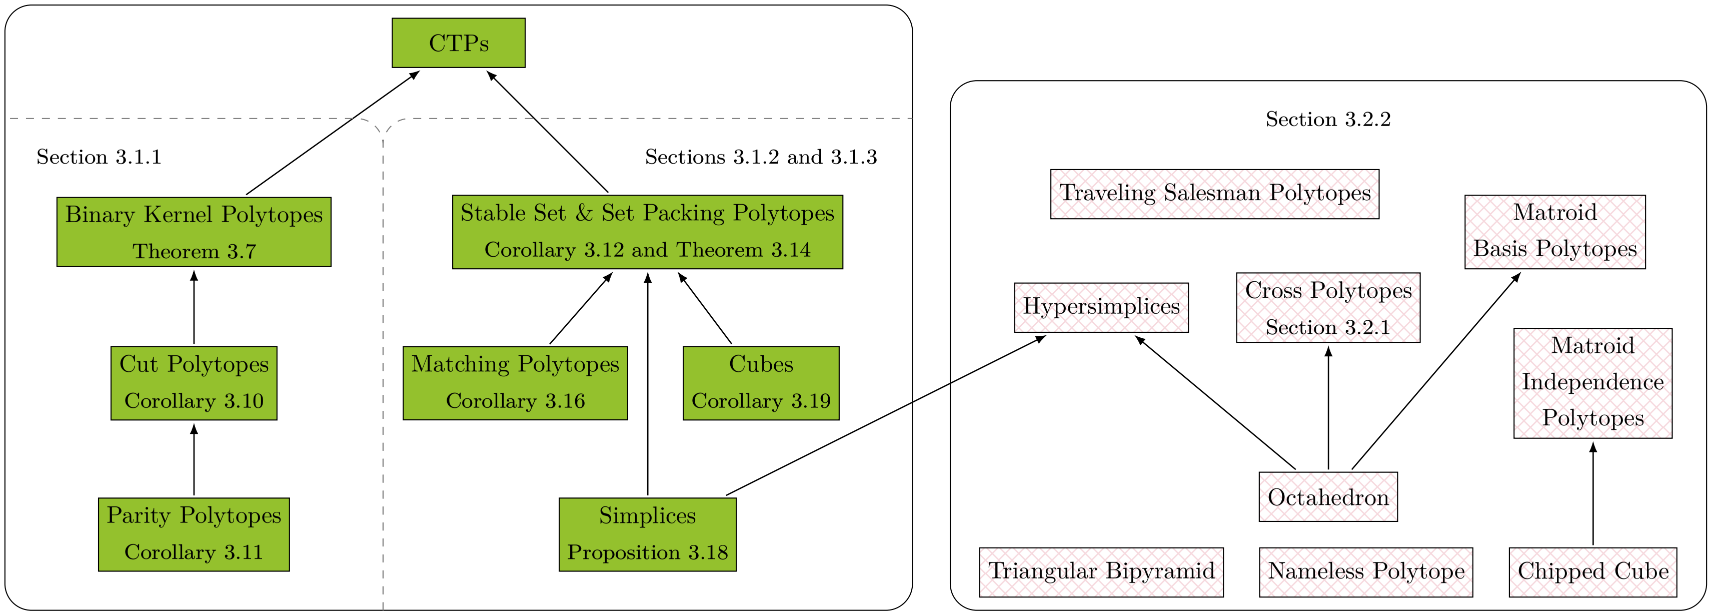

Figure 6.1: Schematic for the containment relations between cyclic transversal polytopes and other polytope families. Solid blocks represent subfamilies of cyclic transversal polytopes, while patterned blocks represent families that contain at least one polytope which is not a cyclic transversal polytope. Additional relations between these families are not depicted, except in cases where they have been discussed in this work.

a binary kernel polytope. In the second part of Chapter 5, we generalized an extended formulation from Carr and Konjevod [13] for parity polytopes to all cyclic transversal polytopes and their relaxations. This extended formulation relies on directed acyclic networks. While it has small size for parity polytopes, it grows exponentially in the dimension of the underlying venue space $\mathbb{F}_2^d$. Using this generalized extended formulation, we could also formulate an extended formulation of the relaxations of cyclic transversal polytopes. To balance this one method of obtaining an extended formulation, we also outlined how the very different disjunctive programming approach leads to an extension of cyclic transversal polytopes that is unlike the one based on network polytopes.

Looking ahead

The ideas in this thesis originated in the investigation of the results from Barahona and Grötschel [7] and their polytopes, together with the extended formulation based on the work of Carr and Konjevod [13]. We wanted to understand how this extended formulation could be applied to the binary kernel polytopes and relaxations we constructed at that time, and how the extended formulation could lead to the discovery of additional constraints that improve the existing relaxations. A reminder of this original goal can still be found in Question 5.21. The pursuit of this goal then uncovered the cyclic transversal framework that is shown in this work, which evidently grew to a stage that paints a much bigger picture. Therefore, it is unsurprising to leave some questions behind, which we have asked at the appropriate places in this thesis. We believe that an answer to any of these questions will lead to a better understanding of the cyclic transversal framework. That said, at one point in every chapter, a leading question is posed that we deem more important and overarching than some of the others.

For Chapter 2, this leading question is Question 2.27, asking about a possible characterization of cyclic transversal polytopes according to the underlying block configurations. Results in this direction are Theorems 3.7 and 3.8, for example. They together state that the cyclic transversal polytopes which are given by a normalized block configuration that consists of blocks with cardinality 2 are exactly the binary kernel polytopes that have been investigated independently before [7]. Another result in this regard is Proposition 3.18, from which it is easy to see that every cyclic transversal polytope $\mathrm{CTP}(\Xi_1, \Xi_2)$ constructed from a block configuration consisting of only two blocks is affinely isomorphic to a simplex of some dimension, since it is a face of the full cyclic transversal polytope $\mathrm{CTP}(d, 2)$ for some rank $d \in \mathbb{N}$, which itself is a simplex. We are certain that more results in this direction lead to a better understanding of the interplay between block configurations and the resulting cyclic transversal polytopes. Another question in this direction is Question 2.30, which asks about a relationship between the dimension of a cyclic transversal polytope and the size of the underlying block configuration.

In Chapter 3, Question 3.34 is asking about further families of polytopes which may be identified as cyclic transversal polytopes. Also, the more specific questions about spanning tree and traveling salesman polytopes that finish the chapter can be viewed as an open-ended endeavor to precisely describe what can be represented in the cyclic transversal framework. Possibly, identifying additional families of polytopes that belong to the cyclic transversal polytopes can shed more light on the existing results as well.

The majority of Chapter 4 is concerned with the description of full cyclic transversal polytopes. Since limited computational resources allow only for small examples of inequality descriptions to be calculated, Question 4.10 has been stated to ask about further theoretical results on the inequalities of these polytopes, since any knowledge about facet-defining inequalities for cyclic transversal polytopes with larger parameters may also lead to practical improvements for solving optimization problems over these polytopes.

In Chapter 5, we already mentioned the leading question and the initial scope of this thesis, which is represented in Question 5.21. To this end, we also asked Question 5.23. Surely, more methods to derive extended formulations and find additional relations between extension variables will lead to further insights on how to properly handle cyclic transversal polytopes themselves, be it algorithmically or theoretically.

In addition to the 16 questions that we asked during this thesis, including the ones mentioned before, we direct our attention to the vast amount of existing literature regarding well-established polytopes that have been identified as cyclic transversal polytopes. There are many families of inequalities that are valid for cut polytopes [8, 10, 20, 21] as well as stable set or set packing polytopes [14, 17, 37, 40, 42], respectively, such as clique inequalities for stable set polytopes or cycle inequalities for cut polytopes. It is natural to try and build upon this heap of knowledge and translate inequalities such as these into the cyclic transversal framework context, in the same way we have generalized the odd-set inequalities of parity polytopes to obtain odd-hyperplane inequalities.

List of Figures

List of Tables

Bibliography

[1] Sanjeev Arora and Boaz Barak. *Computational complexity.* Cambridge University Press, 2009 (cited on p. 3)

[2] Giorgio Ausiello, Alessandro D. D'Atri, and Marco Protasi. "Structure preserving reductions among convex optimization problems." *Journal of Computer and System Sciences* 21.1 (1980), pp. 136–153 (cited on p. 55)

[3] Giorgio Ausiello, Alberto Marchetti-Spaccamela, Pierluigi Crescenzi, Giorgio Gambosi, Marco Protasi, and Viggo Kann. *Complexity and Approximation.* Springer, 1999 (cited on p. 3)

[4] Egon Balas. "Disjunctive Programming." *Discrete Optimization II.* Ed. by P.L. Hammer, E.L. Johnson, and B.H. Korte. Vol. 5. Annals of Discrete Mathematics. Elsevier, 1979, pp. 3–51 (cited on p. 119)

[5] Egon Balas. "Disjunctive programming: Properties of the convex hull of feasible points." *Discrete Applied Mathematics* 89.1 (1998), pp. 3–44 (cited on pp. 119, 127)

[6] Egon Balas. *Disjunctive Programming.* Springer, 2018 (cited on p. 119)

[7] Francisco Barahona and Martin Grötschel. "On the cycle polytope of a binary matroid." *Journal of Combinatorial Theory, Series B* 40.1 (1986), pp. 40–62 (cited on pp. 3, 11, 49, 114, 129, 131)

[8] Francisco Barahona and Ali Ridha Mahjoub. "On the cut polytope." *Mathematical Programming* 36.2 (1986), pp. 157–173 (cited on pp. 11, 132)

[9] Armin Biere, Marijn Heule, Hans van Maaren, and Toby Walsh. *Handbook of Satisfiability.* 2nd ed. Vol. 336. Frontiers in Artificial Intelligence and Applications. IOS Press, 2021 (cited on p. 6)

[10] Thorsten Bonato, Michael Jünger, Gerhard Reinelt, and Giovanni Rinaldi. "Lifting and separation procedures for the cut polytope." *Mathematical Programming* 146.1-2 (2013), pp. 351–378 (cited on pp. 11, 132)

[11] Siegfried Bosch. *Lineare Algebra.* 5th ed. Springer-Lehrbuch. Springer, 2014 (cited on pp. 4, 111)

[12] Richard A. Brualdi. *Introductory Combinatorics.* 5th ed. Pearson Education, 2009 (cited on p. 14)

[13] Robert D. Carr and Goran Konjevod. "Polyhedral Combinatorics." *Tutorials on Emerging Methodologies and Applications in Operations Research*. Ed. by Harvey J. Greenberg. Vol. 76. International Series in Operations Research & Management Science. Springer, 2005. Chap. 2 (cited on pp. 118–121, 131)

[14] Vašek Chvátal. "On certain polytopes associated with graphs." *Journal of Combinatorial Theory, Series B* 18.2 (1975), pp. 138–154 (cited on pp. 51, 132)

[15] Michele Conforti, Gerard Cornuejols, and Giacomo Zambelli. *Integer Programming*. Graduate Texts in Mathematics 271. Springer, 2014 (cited on pp. 8, 75, 120)

[16] Michele Conforti, Gérard Cornuéjols, and Giacomo Zambelli. "Extended formulations in combinatorial optimization." *4OR* 8.1 (2010), pp. 1–48 (cited on p. 118)

[17] Michele Conforti, Samuel Fiorini, Tony Huynh, and Stefan Weltge. "Extended formulations for stable set polytopes of graphs without two disjoint odd cycles." *Mathematical Programming* 192 (2022), pp. 547–566 (cited on pp. 51, 132)

[18] Stephen A. Cook. "The complexity of theorem-proving procedures." *Proceedings of the third annual ACM symposium on Theory of computing - STOC '71*. ACM Press, 1971 (cited on p. 6)

[19] George B. Dantzig, Delbert R. Fulkerson, and Selmer M. Johnson. "Solution of a Large-Scale Traveling-Salesman Problem." *Journal of the Operations Research Society of America* 2.4 (1954), pp. 393–410 (cited on p. 75)

[20] Caterina De Simone. "The cut polytope and the Boolean quadric polytope." *Discrete Mathematics* 79.1 (1990), pp. 71–75 (cited on pp. 11, 132)

[21] Michel Marie Deza and Monique Laurent. *Geometry of Cuts and Metrics*. Algorithms and Combinatorics 15. Springer, 1997 (cited on pp. 11, 50, 132)

[22] Reinhard Diestel. *Graph Theory*. 5th ed. Graduate Texts in Mathematics 173. Springer, 2017 (cited on pp. 7, 12, 51, 73)

[23] Jack Edmonds. "Submodular functions, matroids and certain polyhedra." *Combinatorial Structures and Their Applications; Proceedings of the Calgary International Conference on Combinatorial Structures and Their Applications 1969*. Ed. by Richard K. Guy, Haim Hanani, Norbert Sauer, and Johanan Schönheim. Gordon and Breach, 1970, pp. 69–87 (cited on pp. 72, 73)

[24] Samuel Fiorini, Serge Massar, Sebastian Pokutta, Hans Raj Tiwary, and Ronald de Wolf. "Exponential Lower Bounds for Polytopes in Combinatorial Optimization." *Journal of the ACM* 62.2 (2015), pp. 1–23 (cited on p. 119)

[25] Samuel Fiorini, Thomas Rothvoß, and Hans Raj Tiwary. "Extended Formulations for Polygons." *Discrete & Computational Geometry* 48.3 (2012), pp. 658–668 (cited on p. 118)

[26] Gerd Fischer and Boris Springborn. *Lineare Algebra*. 19th ed. Grundkurs Mathematik. Springer, 2020 (cited on p. 4)

[27] Luis A. Goddyn. “Cones, Lattices and Hilbert Bases of Circuits and Perfect Matchings.” *Graph Structure Theory.* Contemporary Mathematics 147 (1993). Ed. by Neil Robertson and Paul Seymour, pp. 419–439 (cited on p. 115)

[28] Francesco Grande, Arnau Padrol, and Raman Sanyal. “Extension Complexity and Realization Spaces of Hypersimplices.” *Discrete & Computational Geometry* 59.3 (2017), pp. 621–642 (cited on p. 128)

[29] George Grätzer. *Lattice Theory: Foundation.* Springer, 2011 (cited on p. 9)

[30] Martin Grötschel, Laszlo Lovasz, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization.* 2nd ed. Algorithms and Combinatorics 2. Springer, 1993 (cited on p. 8)

[31] Robert G. Jeroslow. “On defining sets of vertices of the hypercube by linear inequalities.” *Discrete Mathematics* 11.2 (1975), pp. 119–124 (cited on pp. 10, 93, 119)

[32] Volker Kaibel and Stefan Weltge. “A Short Proof that the Extension Complexity of the Correlation Polytope Grows Exponentially.” *Discrete & Computational Geometry* 53.2 (2015), pp. 397–401 (cited on p. 119)

[33] Volker Kaibel and Martin Wolff. “Simple 0/1-Polytopes.” *European Journal of Combinatorics* 21.1 (2000), pp. 139–144 (cited on p. 60)

[34] Richard M. Karp. “Reducibility among Combinatorial Problems.” *Complexity of Computer Computations.* Ed. by Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger. Plenum Press, 1972, pp. 85–103 (cited on pp. 6, 46, 54)

[35] Bernhard Korte and Jens Vygen. *Combinatorial Optimization. Theory and Algorithms.* 6th ed. Algorithms and Combinatorics 21. Springer, 2018 (cited on pp. 8, 73)

[36] Serge Lang. *Linear Algebra.* 3rd ed. Undergraduate Texts in Mathematics. Springer, 1987 (cited on p. 4)

[37] László Lipták and László Lovász. “Critical Facets of the Stable Set Polytope.” *Combinatorica* 21 (2001), pp. 61–88 (cited on pp. 51, 132)

[38] R. Kipp Martin. “Using separation algorithms to generate mixed integer model reformulations.” *Operations Research Letters* 10.3 (1991), pp. 119–128 (cited on pp. 118, 119)

[39] C. E. Miller, A. W. Tucker, and R. A. Zemlin. “Integer Programming Formulation of Traveling Salesman Problems.” *Journal of the ACM* 7.4 (1960), pp. 326–329 (cited on p. 75)

[40] George L. Nemhauser and Leslie E. Trotter. “Vertex packings: Structural properties and algorithms.” *Mathematical Programming* 8 (1975), pp. 232–248 (cited on pp. 51, 55, 132)

[41] James Oxley. *Matroid Theory.* Oxford Graduate Texts in Mathematics 21. Oxford University Press, 2011 (cited on pp. 3, 11, 49, 114)

[42] Manfred W. Padberg. "On the facial structure of set packing polyhedra." *Mathematical Programming* 5 (1973), pp. 199–215 (cited on pp. 51, 55, 132)

[43] Christos Harilaos Papadimitriou. *Computational Complexity.* Pearson, 1993 (cited on p. 3)

[44] Kanstantsin Pashkovich. "Extended Formulations for Combinatorial Polytopes." PhD thesis. Otto-von-Guericke-Universität Magdeburg, 2012 (cited on p. 127)

[45] Fred Roberts and Barry Tesman. *Applied Combinatorics.* 2nd ed. CRC Press, 2009 (cited on p. 14)

[46] Thomas Rothvoss. "The Matching Polytope has Exponential Extension Complexity." *Journal of the ACM* 64.6 (2017), pp. 1–19 (cited on pp. 119, 121)

[47] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 9.6).* 2022 (cited on pp. 101, 102, 129)

[48] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency.* Algorithms and Combinatorics 24. Springer, 2003 (cited on pp. 8, 75)

[49] Paul D. Seymour. "Matroids and Multicommodity Flows." *European Journal of Combinatorics* 2.3 (1981), pp. 257–290 (cited on p. 115)

[50] Richard P. Stanley. *Enumerative Combinatorics.* Vol. 2. Cambridge Studies in Advanced Mathematics 62. Cambridge University Press, 1999 (cited on p. 74)

[51] Stefan Weltge. "Sizes of linear descriptions in combinatorial optimization." PhD thesis. Otto-von-Guericke-Universität Magdeburg, 2016 (cited on p. 127)

[52] Mihalis Yannakakis. "Expressing combinatorial optimization problems by Linear Programs." *Journal of Computer and System Sciences* 43.3 (1991), pp. 441–466 (cited on p. 119)

[53] Günter M. Ziegler. "Lectures on 0/1-Polytopes." *Polytopes — Combinatorics and Computation.* DMV Seminar. Birkhäuser Basel, 2000, pp. 1–41. arXiv: `math/9909177v1 [math.CO]` (cited on pp. 8–11, 20, 44, 63, 69, 70, 79)

[54] Günter M. Ziegler. *Lectures on Polytopes.* Graduate Texts in Mathematics 152. Springer, 2007 (cited on pp. 8, 9, 63, 69, 70)