



TECHNISCHE UNIVERSITÄT
CHEMNITZ

Faculty of Mathematics

Professorship for Algorithmic and
Discrete Mathematics

Spectral Realizations of Symmetric Graphs, Spectral Polytopes and Edge-Transitivity

Dissertation submitted to the Faculty of Mathematics at Technische
Universität Chemnitz in accordance with the requirements for the
degree

doctor rerum naturalium (Dr. rer. nat.)

by **M.Sc. Martin Winter**

Chemnitz, June 26, 2021

Referees: Prof. Dr. Christoph Helmberg (supervisor)
Prof. Dr. Michael Joswig
Prof. Dr. Egon Schulte

submitted April 8, 2021
defended June 18, 2021

<https://nbn-resolving.org/urn:nbn:de:bsz:ch1-qucosa2-752155>

Winter, Martin

Spectral Realizations of Symmetric Graphs, Spectral Polytopes and Edge-Transitivity

Dissertation, Faculty of Mathematics

Technische Universität Chemnitz, 2021

Bibliographical description

Martin Winter

Spectral Realizations of Symmetric Graphs, Spectral Polytopes and Edge-Transitivity

Dissertation, 188 pages, Technische Universität Chemnitz, Faculty of Mathematics, 2021

Abstract

A *spectral graph realization* is an embedding of a finite simple graph into Euclidean space that is constructed from the eigenvalues and eigenvectors of the graph's adjacency matrix. It has previously been observed that some polytopes can be reconstructed from their edge-graphs by taking the convex hull of a spectral realization of this edge-graph. These polytopes, which we shall call *spectral polytopes*, have remarkable rigidity and symmetry properties and are a source for many open questions.

In this thesis we aim to further the understanding of this phenomenon by exploring the geometric and combinatorial properties of spectral polytopes on several levels. One of our central questions is whether already “weak” forms of symmetry can be a sufficient reason for a polytope to be spectral. To answer this, we derive a geometric criterion for the identification of spectral polytopes and apply it to prove that indeed all polytopes of *combined vertex- and edge-transitivity* are spectral, admit a unique reconstruction from the edge-graph and realize all the symmetries of this edge-graph. We explore the same questions for graph realizations and find that realizations of combined vertex- and edge-transitivity are not necessarily spectral. Instead we show that we require a stronger form of symmetry, called *distance-transitivity*.

Motivated by these findings we take a closer look at the class of *edge-transitive polytopes*, for which no classification is known. We state a conjecture for a potential classification and provide complete classifications for several sub-classes, such as distance-transitive polytopes and edge-transitive polytopes that are *not* vertex-transitive. In particular, we show that the latter class contains only polytopes of dimension $d \leq 3$.

As a side result we obtain the complete classification of the *vertex-transitive zonotopes* and a new characterization for root systems.

Keywords

spectral graph realizations, eigenpolytopes, spectral polytopes, edge-transitive polytopes, vertex-transitive zonotopes

Acknowledgments

Foremost, I am extremely grateful to my doctoral advisor Prof. Dr. Christoph Helmberg, who gave me the freedom to follow my interests, supported me in every instance, read all my papers, improved them by his detailed comments and was always available for questions. I also thank Dr. Thomas Jahn and Dr. Frank Göring for many hours of fruitful discussions. As a result of one such discussion Frank came up with an ingenious argument to significantly shorten the final step in one of the classification proofs presented in this thesis. Special thanks go to Prof. Dr. Achill Schürmann and Prof. Dr. Raman Sanyal who provided valuable feedback and put me on the right track with their literature references. Thanks go also to my numerous proofreaders (foremost Christin, Daniela, Marvin, Thomas) who made my work more readable with every hint after I became blind to the shortcomings myself. Finally, I would like to thank Prof. Dr. Michael Joswig and Prof. Dr. Egon Schulte for the time and effort they invested in reviewing this thesis.

The author also gratefully acknowledges the support by the funding of the European Union and the Free State of Saxony who enabled this project in the first place.

Contents

List of Symbols	11
Introduction	13
I Spectrum and Symmetry	25
1 Point Arrangements	27
1.1 The arrangement space	28
1.2 Symmetric arrangements	32
1.3 Separating symmetries	37
1.4 Deformation and rigidity	40
1.5 Transitive arrangements	45
2 Graph Realizations	51
2.1 Spectral and balanced realizations	53
2.2 Edge- and arc-transitive realizations	58
2.3 Full local dimension	62
2.4 Distance-transitive realizations	65
2.5 Cosine vectors and cosine sequences	67
3 Eigenpolytopes and Spectral Polytopes	73
3.1 Eigenpolytopes	74
3.2 Spectral polytopes	76
3.3 The Theorem of Izvestiev	80
II Transitivity in Convex Polytopes	85
4 Edge-Transitive Polytopes	87
4.1 An overview of transitivity in polytopes	88
4.2 Known edge-transitive polytopes	89
4.3 Simultaneously vertex- and edge-transitive polytopes	91
4.4 Arc-transitive polytopes	93
4.5 Half-transitive polytopes	98
4.6 Distance-transitive polytopes	98
5 Vertex-Transitive Zonotopes	101
5.1 Generators, faces, symmetry and permutahedra	102

5.2	The case $d = 2$	106
5.3	The general case	108
5.4	The classification	110
5.5	Related topics	112
6	Purely Edge-Transitive Polytopes	115
6.1	Bipartite polytopes	116
6.2	Some easy cases	119
6.2.1	Bipartite polygons	119
6.2.2	Inscribed bipartite polytopes	120
6.2.3	Strictly bipartite polytopes for $d \geq 4$	121
6.3	Strictly bipartite polyhedra	121
6.4	Adjacent pairs	128
6.4.1	The case $\tau_1 = (4, 6, 10)$	130
6.4.2	The case $\tau_1 = (4, 6, 8)$	131
6.4.3	The case $\tau_1 = (4, 6, 6)$	131
6.4.4	The case $\tau_1 = (4, 4, 4)$	132
6.4.5	The case $\tau_1 = (4, 4, 2k), 2k \geq 6$	133
6.4.6	The case $\tau_1 = (4, 4, 6)$	135
6.5	An almost strictly bipartite polyhedron	135
	Conclusions and Outlook	141
	Theses	145
	Appendix	147
A	Matrix Groups and Representations	149
A.1	Matrix groups	149
A.2	Representations	151
B	Spectral Graph Theory	153
C	Polytope Theory	155
C.1	Zonotopes	156
D	Reflection Groups and Root Systems	161
D.1	Reflection groups	161
D.2	Root systems	162
D.3	Coxeter-Dynkin diagrams	162
E	Wythoffian, Uniform and Regular Polytopes	165
E.1	Orbit polytopes and Wythoffian polytopes	165
E.2	Coxeter-Dynkin diagrams for Wythoffian polytopes	165
E.3	Uniform and regular polytopes	166

F	Mathematica Scripts	169
G	Additional proofs and computations	171
	G.1 Two geometric statements	171
	G.2 A proof for spherical interior angles	172

List of Symbols

\mathbb{R}^d	d -dimensional Euclidean space
dim	dimension of linear or affine space
span	linear span/hull of a set of vectors
aff	affine hull of a set of points
conv	convex hull of a set of points
rank	rank of a matrix or set of vectors, equivalent to $\dim \text{span}$
Id	identity matrix (of appropriate dimension)
$\text{GL}(\mathbb{R}^d)$	group of invertible linear transformations $\mathbb{R}^d \rightarrow \mathbb{R}^d$
$\text{O}(\mathbb{R}^d)$	group of orthogonal transformations $\mathbb{R}^d \rightarrow \mathbb{R}^d$
$\text{SO}(\mathbb{R}^d)$	group of orthogonal transformations $\mathbb{R}^d \rightarrow \mathbb{R}^d$ with determinant one
π_W	orthogonal projection onto the subspace W
$\deg(i)$	vertex degree of the vertex i
$\deg(G)$	(common) vertex degree of the vertices in a regular graph G
$\text{dist}(i, j)$	graph theoretic distance between two vertices $i, j \in V$
$\text{diam}(G)$	diameter of the graph G
$\text{Aut}(\cdot)$	automorphism group of a point arrangement, graph, polytope, etc.
G_P	edge-graph of the polytope P
sk_P	skeleton of P , a graph realization $\text{sk}_P : V(G_P) \rightarrow \mathbb{R}^d$ of the edge-graph
$\mathcal{F}(P)$	face lattice of the polytope P
$\mathcal{F}_\delta(P)$	set of δ -dimensional faces of the polytope P
$\text{Zon}(R)$	zonotope generated by R
$\text{Gen}(Z)$	(standard) generators of the zonotope Z
$\text{Sym}(V)$	symmetric group on V , <i>i.e.</i> , the group of permutations of the set V
$\text{Perm}(\mathbb{R}^n)$	group of permutation matrices on \mathbb{R}^n
Π_σ	permutation matrix associated to the permutation σ
id	identity permutation
$\mathcal{A}_d(\Sigma)$	space of full-dimensional Σ -arrangements in \mathbb{R}^d
$\mathcal{A}_d(G, \Sigma)$	space of full-dimensional Σ -realizations of G in \mathbb{R}^d
$\text{Spec}(G)$	(adjacency) spectrum of the graph G
$\text{Eig}_G(\theta)$	θ -eigenspace of the (adjacency matrix of the) graph G
$P_G(\theta)$	θ -eigenpolytope of the graph G
eig_G^θ	θ -realization used to construct the θ -eigenpolytope
$\ x\ $	Euclidean norm of a vector x
$\langle x, y \rangle$	standard inner product in \mathbb{R}^d
$\sphericalangle(x, y)$	angle between the vectors x and y
S^d	d -dimensional sphere (embedded in \mathbb{R}^{d+1})

List of Symbols

$S_r(c)$ sphere with center c and radius r (of appropriate dimension)
vol (relative) volume of a subset of \mathbb{R}^d

Introduction

The product of mathematics is clarity and understanding. Not theorems, by themselves.


William Thurston

The content of this thesis was developed in an effort to understand a series of curious observations made while working on so-called *spectral graph realizations*. Those are a technique to embed a graph (a purely combinatorial object) into Euclidean space by using the eigenvalues and eigenvectors of its adjacency matrix. Applying this to the edge-graphs of polytopes results in some cases in an embedding that looks just like the skeleton of the original polytope. This phenomenon points to connections between polytope theory and spectral graph theory, between convexity and rigidity, between symmetry in geometry and combinatorics, all to be explored in this thesis.

A curious observation

The best way to motivate the investigation of the described phenomenon is to see it happen.

Consider the 3-dimensional cube and let $G = (V, E)$ be its edge-graph with vertex set $V = \{1, \dots, 8\}$, numbers assigned to the vertices according to the following figure:



To G we can assign a matrix $A \in \{0, 1\}^{8 \times 8}$, the so-called *adjacency matrix*, with $A_{ij} = 1$ if and only if $ij \in E$ is an edge. The field of *spectral graph theory* studies the spectrum, eigenvalues and eigenvectors of A and related matrices. For example, the spectrum of G (by which we mean the set of eigenvalues of A) in the case of the cube-graph is

$$\text{Spec}(G) = \{3^1, 1^3, (-1)^3, (-3)^1\},$$

where the exponents denote multiplicities. The eigenvalues of the adjacency matrix are usually denoted by $\theta_1 > \theta_2 > \dots > \theta_m$ in decreasing order. It is general wisdom in spectral graph theory that one of the more exciting eigenvalues of a graph is the *second-largest* eigenvalue

θ_2 . For the edge-graph of the cube we have $\theta_2 = 1$ of multiplicity *three*. That is, there are three linearly independent eigenvectors to θ_2 . One possible choice of orthogonal bases for the eigenspace is


$$u_1 = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \\ -1 \\ -1 \\ -1 \\ -1 \end{pmatrix}, \quad u_2 = \begin{pmatrix} 1 \\ 1 \\ -1 \\ -1 \\ 1 \\ 1 \\ -1 \\ -1 \end{pmatrix}, \quad u_3 = \begin{pmatrix} 1 \\ -1 \\ 1 \\ -1 \\ 1 \\ -1 \\ 1 \\ -1 \end{pmatrix}.$$

We can already observe the following: we started with the cube, a 3-dimensional polyhedron, and found an eigenspace of dimension *three*. At this point this could well be a coincidence, but it will become clear that there is more to it.

We can write the eigenvectors more compactly in a single matrix $\Phi \in \mathbb{R}^{8 \times 3}$:

$$\Phi = \begin{matrix} & \begin{matrix} u_1 & u_2 & u_3 \end{matrix} \\ \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & -1 \\ 1 & -1 & 1 \\ 1 & -1 & -1 \\ -1 & 1 & 1 \\ -1 & 1 & -1 \\ -1 & -1 & 1 \\ -1 & -1 & -1 \end{pmatrix} & \begin{matrix} \leftarrow v_1 \\ \leftarrow v_2 \\ \leftarrow v_3 \\ \leftarrow v_4 \\ \leftarrow v_5 \\ \leftarrow v_6 \\ \leftarrow v_7 \\ \leftarrow v_8 \end{matrix} \end{matrix}.$$

The *eight* rows of Φ can be naturally assigned to the vertices of G : vertex $i \in V$ corresponds to the i -th row v_i of Φ . In particular, each row can be interpreted as a vector in \mathbb{R}^3 and this defines a map $V \rightarrow \mathbb{R}^3$. This map describes an embedding of the graph into 3-dimensional Euclidean space and is what we call a *spectral graph realization* to the eigenvalue θ_2 . If we do this explicitly for the cube graph we find a structure that looks suspiciously like the skeleton of a cube.



In fact, if we were to take the convex hull of the v_i we would be back at the polyhedron from which we started.


What we just saw happen is that the edge-graph of the cube seems to encode, in its combinatorial structure, information about the dimension and geometry of the polytope it came from, and that we can use this information to reconstruct it.

One can ask what else we could have expected to obtain by this construction? Why have we used θ_2 ? Why these exact eigenvectors? Is this a coincidence or does it happen for other polytopes too? These questions have been the starting point for our studies.

Eigenpolytopes and spectral polytopes

The previously demonstrated phenomenon is not unknown to the literature. In fact, the very polytope we obtain by taking the convex hull of a spectral realization is called *eigenpolytope* of a graph and was introduced by Godsil in 1978 [31]. The eigenpolytope was soon recognized as an interesting link between the algebraic and combinatorial properties of a graph and the geometric properties of a polytope. To this day, their literature is quite scattered [32, 47, 57, 59, 62–64]. A more detailed overview of the previous work will be given in Chapter 3, after the formal introduction of eigenpolytopes.

In many cases the eigenpolytopes of an edge-graph show little resemblance to the original polytope. Inspection of the cube example shows that the dimension of the resulting embedding was not a choice, but was determined by the multiplicity of θ_2 , and mysteriously matched the dimension of the cube. In contrast, the edge-graphs of many prisms have no eigenspace of dimension three, even though prisms are 3-dimensional.



$$\text{Spec}(6\text{-prism}) = \{3^1, 2^2, 1^1, 0^4, (-1)^1, (-2)^2, (-3)^1\}$$

We can then not expect to reconstruct a prism in the same way we reconstructed the cube. We shall see that a polytope being an eigenpolytope of its edge-graph is something very special.

Polytopes with this property will be called *spectral polytopes*. They are a special kind of eigenpolytope. The introductory example demonstrated that the cube is a spectral polytope, or more precisely, a θ_2 -spectral polytope, since we used θ_2 for the reconstruction.

The existence of such polytopes has been noticed before at several instances. Two notable occurrences are in a paper by Licata and Powers [47], and in another paper by Godsil [32]. Licata and Powers observed this “self-reproducing” property for the Platonic solids, and more generally, for all regular polytopes excluding the 4-dimensional exceptions (the 24-cell, 120-cell and 600-cell, for which their work was inconclusive; for a reminder on the regular polytopes consider the classical reference from Coxeter [18], or see Appendix E). Godsil, on the other hand, studied the eigenpolytopes for a class of graphs that is especially accessible by

the techniques of spectral graph theory, namely, the *distance-regular graphs*. He obtained a complete classification of the distance-regular graphs that produce spectral polytopes.

In both papers the authors exclusively studied eigenpolytopes to the eigenvalue θ_2 . In fact, it is not unreasonable to assume that polytopes cannot be spectral for any other eigenvalue, though there is no proof for this claim as of yet. It appears that the second-largest eigenvalue is somehow linked to the convexity of polytopes, a claim that we explore in more detail in this thesis, and to which partial results will be obtained.

All in all there is still little known about eigenpolytopes and spectral polytopes. We ask:

Question 1. Which polytopes are spectral? Can we classify spectral polytopes? What else makes them special besides the property that defines them?

Question 2. Are there spectral polytopes to an eigenvalue other than θ_2 ?

While we will obtain only partial answers, these questions have been the driving force behind many of our investigations.

Spectral methods for convex polytopes

If a polytope P turns out to be spectral, then from the perspective of polytope theory there are at least the following two observations to be made:

O1 P is uniquely determined by its edge-graph.

The construction of the eigenpolytope provides an explicit procedure to obtain P from its edge-graph. The question which polytopes or polytope classes admit a unique reconstruction from the edge-graph (up to combinatorial equivalence) has a long history. Since Steinitz [70] we know that a unique reconstruction is possible in dimension three. It has since become clear that the same fails in higher dimensions. For example, the complete graph $K_n, n \geq 6$ is the edge-graph of many not combinatorially equivalent polytopes, already in dimension four [79, Example 0.6].

Certain sub-classes of polytopes are known to admit a unique reconstruction, e.g. simple polytopes [7, 42] and zonotopes [6, Theorem 6.14]. The spectral polytopes then form another class in this list. In fact, their reconstruction is unlike stronger: the procedure not only determines the combinatorial type, but also favors a particular realization.

O2 P realizes all the symmetries of its edge-graph.

This is not quite self-evident, but is more generally a property of spectral graph realizations: if a graph is embedded spectrally, then each of its combinatorial symmetries is expressed in the embedding by a geometric symmetry. This is a well-known property, and we include a proof of this after the formal introduction of spectral graph realizations in [Chapter 2](#).

A polytope is clearly at most as symmetric as its edge-graph. But the converse, realizing *all* symmetries of the edge-graph, is something very special. Trivial examples of this

failing are obtained by deforming a polytope into a less symmetric shape, *e.g.* a square into a rectangle. This “counterexample” is not too convincing, as the rectangle clearly has a realization that is as symmetric as the edge-graph. A version of Steinitz’ theorem ensures that this is true for all polytopes in up to three dimensions [51].

This breaks down in higher dimensions. All *neighborly polytopes* (*i.e.*, polytopes with edge-graph K_n) besides the simplices are less symmetric than K_n . This is another symptom of the general rule that the edge-graph carries very little information about a polytope in higher dimensions. Some 4-polytopes cannot even be realized with all the symmetries of their *face lattice* (which carries much more information than the edge-graph). One such polytope was constructed in [9]. Its polar dual has another special property: its edge-graph is not the edge-graph of any polytope with all the symmetries of the graph.

These properties of spectral polytopes are nice, but also kind of trivial. These observations show their full potential only as we identify classes of polytopes as spectral that are not obviously so.

For example, O2 is especially interesting if P ’s edge-graph has some symmetries to begin with. Wishful thinking leads to the following question:

Question 3. Is every sufficiently symmetric polytope spectral?

This is intentionally vague. What does qualify as “sufficient symmetry”? For example, vertex-transitivity (any two vertices of the polytope can be moved onto each other by a symmetry) does not qualify: there are vertex-transitive polytopes that are not spectral, *e.g.* rectangles, prisms and other more “convincing” counterexamples such as vertex-transitive neighborly polytopes [41].

On the other hand, being a regular polytope (being transitive on faces of *all* dimensions) does qualify. As previously mentioned, this was shown for most cases by Licata and Powers [47]. We shall give several proofs for the general case over the course of this thesis. This statement is not too impressive however, given that the regular polytopes are not a particularly rich class (the classification has been achieved by Schläfli [66] around 1852 and shows that there are only three regular polytopes in any dimension $d \geq 5$).

The “right symmetry” must be somewhere in between. To narrow down where this might be we can take a look at what O1 and O2 would imply for it:

O1’ Every sufficiently symmetric polytope is uniquely determined by its edge-graph up to orientation and scale.

This is related to the concept of a *perfect polytope*. Roughly, a polytope is called perfect if it has a unique realization of highest symmetry (see *e.g.* [28]).

This is also related to ideas from *rigidity theory*: the polytope is *rigid* in the sense that it admits a unique realization that complies with a certain set of constraints.

O2’ Every sufficiently symmetric polytope realizes all the symmetries of its edge-graph.

We think that these statements are quite exciting, if they are true, and if they are formulated with an interesting version of “sufficient symmetry”.

One of our first candidates for sufficient symmetry is *edge-transitivity* (any two edges of the polytope can be moved onto each other by a symmetry). This does not work from the get go: there are several edge-transitive polytopes that are not spectral, e.g. certain polygons and also one 3-dimensional polyhedron (see [Figure 1](#)).




Figure 1. Examples of edge-transitive polytopes that are not spectral. This is evident for the polygons as they not express the vertex-transitivity of their edge-graphs. The polyhedron on the right is the *rhombic triacontahedron*, a Catalan solid, and is the only non-spectral edge-transitive polyhedron.

However, we shall see that this first guess does not miss by much. Consider the following refinement:

Question 4. If P is simultaneously vertex- and edge-transitive, is it spectral?

Providing an answer to [Question 4](#) was a major driving force behind the development of Part I of the thesis. And in fact, one of the main results of this part is to answer this question in the affirmative. As a consequence, polytopes of this symmetry can be uniquely reconstructed from their edge-graphs and realize all their symmetries. This is an example for a statement about convex polytopes whose currently only proof uses spectral graph theory.

Edge-transitive polytopes and more

It have been these findings that sparked our interest in the class of edge-transitive polytopes. Is this class sufficiently rich to make our investigations worthwhile?

The reader might be surprised (we certainly were) to find that despite the classical appeal of this question, there has seemingly not been obtained a complete classification of the edge-transitive convex polytopes.

Question 5. How rich is the class of edge-transitive polytopes?

Is it more like vertex-transitivity (a classification is infeasible), or more like regularity (only a few examples in higher dimensions)?


This interest developed into Part II of the thesis, where we discuss in detail the ramifications and intricacies of [Question 5](#). We do not provide a complete answer, but we certainly make progress. After structuring the class of edge-transitive polytopes by introducing a hierarchical classification scheme, we provide complete classifications for several sub-classes. We derive a classification of distance-transitive polytopes (by means of spectral graph theory), and give a precise conjecture for the classification of arc-transitive polytopes.

Another major result in this regard is the classification of all edge-transitive polytopes that are *not* vertex-transitive. The list turns out to be quite short, and most surprisingly, it contains no polytopes of dimension four or higher. So, in a very precise sense, we show that [Question 4](#) is indeed not far from asking whether all edge-transitive polytopes are spectral.

Part II also discusses a seemingly unrelated question: the classification of all *vertex-transitive zonotopes* (a zonotope is a polytope in which all faces are centrally symmetric, see [Appendix C.1](#)). While initially disconnected to our study of edge-transitivity, the results of this investigation turn out to be crucial ingredients for the classification of edge-transitive polytopes that are not vertex-transitive.

Graph realizations and point arrangements

Even though they are meant to provide the fundamental motivation for this work, polytopes are only one of several objects we study in connection with symmetry and spectral graph theory. Part I of the thesis deals only to a third with the phenomenon of spectral polytopes. Instead it is composed of chapters addressing (in this order) point arrangements, graph realizations, and only then, eigenpolytopes and spectral polytopes. The chapters of Part I represent three levels of increasing structure. We describe briefly the motivation for introducing this hierarchy, as well as some of the directions in which their theory is developed.



First, this hierarchy allows a step-wise development of all relevant concepts, such as symmetry, rigidity and spectral properties. Each concept can be introduced with using the least structure necessary for its definition.

Second, it is not a priori clear that the phenomenon of spectral polytopes can be best explained on the level of polytopes. It might well be that it is just a consequence of a more fundamental result for, say, graph realizations. The development of Part I in stages allows us to address each phenomenon at the earliest possible level. We will also make clear when results do not transfer between levels.

For example, [Question 3](#) makes perfect sense when asked for graph realizations instead of polytopes:

Question 6. Is every sufficiently symmetric graph realization a spectral realization?

A positive answer to [Question 6](#) implies a positive answer to the polytope version. But, it is unclear that “sufficient symmetry” for realizations means the same as for polytopes. We shall see that combined vertex- and edge-transitivity is *not* sufficient here, and must be replaced by a stronger symmetry called *distance-transitivity*.

To figure out what “sufficiently symmetric” might mean, recall the following interpretation of **O1'**: if we want sufficient symmetry to make a polytope/realization spectral, it must come attached with a kind of *rigidity*, that is, uniqueness up to scale and orientation.

Question 7. When is symmetry implying rigidity?

This is a very broad question. For example, whether a framework (in the sense of rigidity theory) with prescribed symmetries is rigid has been investigated in great detail [50, 58, 68]. But **Question 7** makes sense even without any underlying graph structure: all we need is a set of points in Euclidean space. If this “point arrangement” has some non-trivial Euclidean symmetries then we can ask whether it can be “deformed” in a way that does not destroy these symmetries. If this is not possible, then we found a *rigid symmetry*. We later work out the details of how rigidity and spectral properties relate to each other.

Likewise, **O2'** can be considered for realizations too: a sufficiently symmetric realization of a graph G has *all* the symmetries of the graph. This statement has a flavor quite different from its polytope version: suppose that $\Sigma \subset \text{Aut}(G)$ is a “sufficiently large” proper subgroup of the symmetry group $\text{Aut}(G)$ of the graph. **O2'** then says that it is *impossible* to find a graph realization that has all the symmetries in Σ , but not a single additional symmetry. We say that Σ and $\text{Aut}(G)$ cannot be “separated geometrically”.

Question 8. When can symmetry groups, even though they are combinatorially distinct, not be separated geometrically?

Structure and main results

The thesis is divided into two parts.

Part I

Part I was developed in an effort to understand the connections between spectrum and symmetry, and in particular, the phenomenon of spectral polytopes. This part is divided into three chapters that incrementally build a theory of point arrangements, graph realizations and polytopes.

Chapter 1 introduces symmetry and rigidity in the context of point arrangements and defines the central notion of the *arrangement space*, which will permeate all of Part I. It is one of the main goals of this part to characterize various properties of point arrangements, graph realizations and polytopes in terms of their arrangement space.

In **Chapter 2** we introduce spectral and balanced graph realizations. We show that combined vertex- and edge-transitivity is not sufficient to ensure that a realization is spectral, but we show that *distance-transitivity*, an even stronger form of symmetry, is sufficient in this sense (cf. **Question 6**). Already here we can present a procedure by which to verify that certain polytopes are spectral, including the regular polytopes.

Part I culminates in **Chapter 3**, in which we properly define eigenpolytopes and spectral polytopes as well as provide an extensive literature overview. We then present a purely geometric criterion for the identification of θ_2 -spectral polytopes and apply it to polytopes

with combined vertex- and edge-transitivity. This shows that this symmetry class is indeed sufficient in the sense of [Question 3](#).

The following are the **main achievements** of Part I:

- The introduction of the arrangement space and its use in connecting symmetry, rigidity and spectral properties of point arrangements, graph realizations and polytopes.
- The identification of symmetries that are sufficient for a graph realization or polytope to be spectral. For example, we find that distance-transitivity is sufficient for graph realizations.
- The development of a geometric criterion for the identification of spectral polytopes and the application of this criterion to polytopes of combined vertex- and edge-transitivity. We shall find that such polytopes are θ_2 -spectral, and thereby, are uniquely determined by their edge-graphs and realize all the symmetries of their edge-graphs.
- The identification of graph symmetries that cannot be geometrically separated from the other symmetries of the graph.

More generally, we demonstrate that the methods of spectral graph theory add to the toolbox of proof techniques for polytope theory.

Part II

Motivated by the findings that edge-transitivity has a special connection with spectral properties, Part II focuses on the investigation of this symmetry class.

In [Chapter 4](#) we give a general introduction to transitivity phenomena in convex polytopes and conclude that these are still badly understood. We then focus mainly on edge-transitivity. We introduce a hierarchical classification scheme for the systematic investigation of this symmetry class. We take a detailed look at each level of the proposed hierarchy, which includes polytopes that are simultaneously vertex- and edge-transitive, arc-transitive, half-transitive and distance-transitive. We give a complete classification of all distance-transitive polytopes, and we state a conjecture about the classification of arc-transitive polytopes. We explore the potential consequences of this conjecture.

One of the main aspects of this classification scheme had to be postponed until later chapters, namely, the classification of edge-transitive polytopes that are *not* vertex-transitive. The proof of this classification, which is quite technical at some points, will be our main occupation for the rest of this part.

In the first step we have to study a seemingly unrelated problem. In [Chapter 5](#) we prove a classification of zonotopes under additional constraints, *e.g.* vertex-transitive zonotopes and inscribed zonotopes that have all edges of the same length. While not quite evident at this point, results of this chapter will be a major ingredient in [Chapter 6](#), where we finally achieve the complete classification of edge-transitive polytopes that are not vertex-transitive. In particular, we find that all edge-transitive polytopes in dimension $d \geq 4$ are vertex-transitive.

The following are the **main achievements** of Part II:

- The organization of edge-transitive polytopes into the proposed hierarchy.

- The derivation of several deep properties of edge-transitive polytopes that can currently only be achieved with the methods of spectral graph theory.
- The conjecture for the complete classification of arc-transitive polytopes.
- The complete classification of distance-transitive polytopes.
- The complete classification of edge-transitive polytopes that are not vertex-transitive.
- The complete classification of vertex-transitive zonotopes, and the complete classification of inscribed zonotopes in which all edges are of the same length. As a side result we obtain a new characterization for root systems and a classification of hyperplane arrangements whose symmetries act transitively on their chambers.

Appendix

We included an extensive appendix, containing introductions to matrix groups and representations (Appendix A), spectral graph theory (Appendix B), polytope theory (Appendix C), root systems and reflections groups (Appendix D) as well as Wythoffian polytopes, including regular and uniform polytopes (Appendix E).

The appendix further includes a short Mathematica script for the computation of spectral graph realizations and eigenpolytopes (Appendix F), as well as some longer calculations that would have interrupted the flow of the main text (Appendix G).

Some notes on notation

Most of the notation and terminology in this thesis is standard or will be explained at first use. Any unexplained notation and terminology can be found in

- [20] for general graph theory,
- [30] for algebraic graph theory and highly symmetric graphs,
- [15] for spectral graph theory,
- [79] for general polytope theory,
- [18] for the terminology of regular and uniform polytopes.

Still, to be on the same page with the reader, some of the most basic notations we shall introduce already now.

By $G = (V, E)$ we will denote a finite simple *graph* with vertex set $V = \{1, \dots, n\}$ and edge set $E \subseteq \binom{V}{2}$. An edge $\{i, j\} \in E$ will be denoted by $ij \in E$, and the vertices $i, j \in V$ are said to be *adjacent*. The *neighborhood* of a vertex $i \in V$ is defined as

$$N(i) := \{j \in V \mid ij \in E\}.$$

By $\text{Sym}(V)$ we denote the *symmetric group* on the vertex set of G , that is, the group of permutations of V . The (combinatorial) *symmetry group* of G is the set of permutations of V that preserves adjacency, that is,

$$\text{Aut}(G) := \{\sigma \in \text{Sym}(V) \mid ij \in E \Leftrightarrow \sigma(i)\sigma(j) \in E\}.$$

A graph is *vertex-transitive* if $\text{Aut}(G)$ acts transitively on V , that is, if for any two $i, j \in V$ there is a $\sigma \in \text{Aut}(G)$ with $\sigma(i) = j$. Like-wise, a graph is *edge-transitive* if $\text{Aut}(G)$ acts transitively on its edges.

For $d \geq 1$, let $P \subset \mathbb{R}^d$ denote a (*convex*) *polytope*, that is, the convex hull of finitely many points. We assume that the vertices $v_1, \dots, v_n \in P$ are labeled from an index set $V = \{1, \dots, n\}$. By $\mathcal{F}(P)$ we denote the set of faces (or *face lattice*) of P , and by $\mathcal{F}_\delta(P)$ we denote the subset of δ -dimensional faces. By G_P we denote the *edge-graph* of P , which we consider as a purely combinatorial graph with vertex set $V = \{1, \dots, n\}$ as before. Each $i \in V$ is assigned to the vertex $v_i \in \mathcal{F}_0(P)$. By *skeleton* of P we mean the map

$$\text{sk}_P : V \rightarrow \mathbb{R}^d, i \mapsto v_i$$

In particular, $ij \in E$ if and only if $\text{conv}\{v_i, v_j\} \in \mathcal{F}_1(P)$ is an edge of P .

For our purpose, a symmetry of a polytope is an orthogonal transformation that fixes the polytope set-wise. That is, its (orthogonal) symmetry group is

$$\text{Aut}(P) := \{T \in O(\mathbb{R}^d) \mid TP = P\}.$$

Note that these symmetries fix the origin, and so the group is not invariant under translation of P . Vertex- and edge-transitivity for polytopes is defined parallel to graphs.

Publications


This thesis is in parts a compilation of the previous publications [74–78], but contains also large chunks of revised and rewritten material. In particular, the terminology and notation was unified and most proofs have been rewritten.

- **Chapter 1** follows the theory developed in [76], though has some novel features such as the section on separation of symmetries (**Section 1.3**) and its discussion of transitivity and exceptional rigidity (**Section 1.5**).
- **Chapter 2** is based on [78] but was completely rewritten.
- **Chapter 3** is based on [75] but was completely rewritten.
- Parts of **Chapter 4** are also contained in [75] (the properties of simultaneously vertex- and edge-transitive polytopes in **Theorem 4.5**, as well as the classification of distance-transitive polytopes in **Theorem 4.18**), but most of this chapter is unique to this thesis (most notably, the classification scheme for edge-transitive polytopes, as well as the discussion on arc-transitivity and the conjectured classification).
- The classification of vertex-transitive zonotopes in **Chapter 5** can also be found in [74] (which has since been accepted at *Discrete & Computational Geometry*).
- **Chapter 6** can also be found in [77].

Part I

Spectrum and Symmetry

1 Point Arrangements



For this chapter we study the simplest geometric objects for which we can meaningfully talk about notions like symmetry and rigidity – finite sets of points in Euclidean spaces. In many cases one can imagine these points to originate as the vertex set of a convex polytope. The (Euclidean) symmetries of such a “point arrangement” are isometries of the ambient space that permute the points, and aiming to preserving these symmetries is enough to have a notion of rigidity.

The purpose of this chapter is mainly preparatory – to introduce a convenient language to be carried over into the upcoming chapters on graph realizations and spectral polytopes. Our tools are largely linear algebra and real representation theory. The introduced language is not conceptually new in many aspects (it shares parallels to, for example, finite frame theory [13, 72, 73]), but to emphasize the geometric perspective, and in order to be self-contained, it is developed with great care, detailed proofs and many examples and illustrations.

The central object of this chapter are *point arrangements* which are families of finitely many points in Euclidean space. We assume that these points are labeled with natural numbers from an index set $V = \{1, \dots, n\}$ (motivated by later use, V stands for “vertex set”).

Definition 1.1. A d -dimensional (*point*) *arrangement* is a map $v: V \rightarrow \mathbb{R}^d$.




Figure 1.1. In depictions of point arrangements we usually assume that the arrangement is centered at the origin. We use lines to hint to an intended geometry. These lines are for visualization purpose only and are not part of the arrangement. The points are often labeled $1, \dots, n$, even though they are $v_1, \dots, v_n \in \mathbb{R}^d$.

The points in an arrangement v are denoted $v_1, \dots, v_n \in \mathbb{R}^d$. Throughout this thesis, n shall

always denote the number of these points (and later the number of vertices in a graph or polytope) and d shall always denote the dimension of the ambient space.

The points of an arrangement can further be organized as the rows of a matrix

$$\Phi := \begin{pmatrix} - & v_1 & - \\ & \vdots & \\ - & v_n & - \end{pmatrix} \in \mathbb{R}^{n \times d} \quad (1.1)$$

which we shall call the *arrangement matrix* of v . An arrangement is called *full-dimensional* if $\text{span } v := \text{span}\{v_1, \dots, v_n\} = \mathbb{R}^d$, or equivalently, if $\text{rank } \Phi = \text{rank } \Phi^\top = d$.

Chapter overview

Section 1.1 introduces the *arrangement space* of a point arrangement, a tool to discuss symmetry, spectral properties, rigidity etc. under a common terminology. Over the course of Part I we develop a “dictionary”, the *arrangement space dictionary*, to translate between the properties of an arrangement and the properties of its arrangement space.

From **Section 1.2** on we work with *symmetric arrangements*, the points of which can be permuted in prescribed ways by isometries of the ambient space. We characterize them via their arrangement spaces.

In **Section 1.3** we discuss the *geometric separation of symmetries* as asked for in **Question 8**. We elaborate how this question has appeared before in the form of symmetry classification problems. We prove a criterion for geometric separation in terms of arrangement spaces.

In **Section 1.4** we explore the rigidity theory of symmetric arrangements (cf. **Question 7**): we ask under which conditions an arrangement can be deformed without losing its prescribed symmetries. If it cannot be deformed, we consider it a *rigid arrangement*. We prove that rigidity depends on the number and relative placement of invariant subspaces in \mathbb{R}^n .

In **Section 1.5** we address the special case of *transitive arrangements*, in which any two points are identical under symmetry. We show that transitive arrangements provide the natural setting in which to expect rigidity. We formulate simple criteria by which to judge the rigidity of transitive arrangements. We briefly address *exceptional rigidity*.

1.1 The arrangement space

To each arrangement v we can assign a linear subspace $U := \text{span } \Phi \subseteq \mathbb{R}^n$, called its *arrangement space*, defined as the column span of the arrangement matrix Φ .

The arrangement space is a straightforward, yet very useful tool for working with point arrangements, and later, graph realizations and polytopes. It enables us to study those objects up to certain irrelevant transformations and it unifies the language we use to talk about symmetries, spectral properties, rigidity etc.

The arrangement space surely cannot capture all the information of an arrangement and therefore does not allow us to reconstruct it. In fact, for each (non-zero) subspace $U \subseteq \mathbb{R}^n$ there are many arrangements which have U as its arrangement space:

Construction 1.2. Choose a basis $u_1, \dots, u_d \in \mathbb{R}^n$ of U and let $\Phi := (u_1, \dots, u_d) \in \mathbb{R}^{n \times d}$ be the matrix in which the u_i are the columns. Then $U = \text{span } \Phi$. If we now define v_i as the i -th row of Φ , we find that v is an arrangement with arrangement matrix Φ , and thus, arrangement space U .

We see that any choice of a basis of U gives us a distinct arrangement with arrangement space U . Let us call two arrangements *equivalent* if they have the same arrangement space. Whether two arrangements are equivalent is then a matter of basic linear algebra. Consider the following well-known theorem:

Theorem 1.3. Two matrices $\Phi, \Phi' \in \mathbb{R}^{n \times d}$ have the same column span if and only if their rows are related by an invertible linear transformation, i.e., $\Phi^\top = X\Phi'^\top$ for some $X \in \text{GL}(\mathbb{R}^d)$.

This statement is sufficiently standard, so that we did not include a proof. Later we shall give a proof for a more specialized version (see [Lemma 1.10](#)).

Note that for arrangements v and v' with arrangement matrices Φ and Φ' the statement $\Phi^\top = X\Phi'^\top$ resp. $\Phi = \Phi'X^\top$ translates to $v_i = Xv'_i$ for all $i \in V$, or $v = Xv'$ for short. Using [Theorem 1.3](#) this immediately yields

Corollary 1.4. Two arrangements $v, v' : V \rightarrow \mathbb{R}^d$ are equivalent (have the same arrangement space) if and only if they are related by an invertible linear transformation, i.e., $v = Xv'$ for some $X \in \text{GL}(\mathbb{R}^d)$.

The arrangement space determines an arrangement up to invertible linear transformations. Rather than a shortcoming, this is a useful feature in many cases. For example, if one only cares about linear, affine or convex dependencies between points, rather than their exact positioning, then the arrangement space carries all the information one needs.

For our purpose however, ignorance up to invertible linear transformations is not always appropriate. For example, the arrangement space cannot distinguish between (the vertex set of) a rhombus, a rectangle and a square, even though these shapes are very different from a symmetry perspective (e.g. only the rhombus is not vertex-transitive).



If we decide that metric properties like angles and lengths are important for us and should be preserved among all equivalent arrangements, then we can make this happen by working with a smaller class of arrangements.

For a linear subspace $W \subseteq \mathbb{R}^d$, let π_W denote the *orthogonal projection* onto W .

Definition 1.5. Let v be an arrangement with arrangement matrix Φ .

- (i) v is called *spherical*, if $\Phi^\top \Phi = \alpha \pi_W$ for some $W \subseteq \mathbb{R}^d$ and $\alpha > 0$.
- (ii) v is called *normalized*, if $\Phi^\top \Phi = \pi_W$ for some $W \subseteq \mathbb{R}^d$.




Figure 1.2. On the left a non-spherical arrangement, and on the right its spherical counterpart.

Geometrically, the points in a spherical/normalized arrangement are distributed as *isotropically* as possible inside of their linear span, that is, the arrangement appears not stretched or squeezed in any direction. This follows essentially from the fact that $\Phi^\top \Phi$ is (up to some factor) the *covariance matrix* of the point cloud v .

Normalized arrangements are spherical arrangements of a preferred scale (every spherical arrangement can be uniformly scaled into a normalized arrangement). But more generally, one can view normalized arrangements as a normal form of general arrangements: every arrangement is equivalent to a (full-dimensional) normalized arrangement.

Construction 1.6. Choose an *orthonormal basis* $u_1, \dots, u_\delta \in \mathbb{R}^n$ of the arrangement space U of v . Let v' be the δ -dimensional arrangement with arrangement matrix $\Phi' := (u_1, \dots, u_\delta) \in \mathbb{R}^{n \times \delta}$. By construction, v' has arrangement space $U = \text{span } \Phi'$, and so v and v' are equivalent. Since the columns of Φ' are an orthonormal basis, $\Phi'^\top \Phi' = \text{Id}$, which is the special case of $\Phi'^\top \Phi' = \pi_W$ for $W = \mathbb{R}^\delta$. Thus, v' is normalized. Since $\text{rank } \Phi' = \delta$, v' is full-dimensional.

In particular, every subspace $U \subseteq \mathbb{R}^n$ is the arrangement space of a normalized arrangement, and one such can be obtained by using [Construction 1.2](#) with an orthonormal basis.

Example 1.7. A generic rhombus or rectangle (that is, the arrangement of its vertices) is not spherical. Applying [Construction 1.6](#) yields a square in both cases and this square is then spherical. More generally, the vertices of any regular polygon, Platonic solid or regular polytope in a higher dimension yield a spherical arrangement.

If v is full-dimensional, then the definition of “normalized” simplifies to $\Phi^\top \Phi = \text{Id}$. This will be the most common case. The more general definition as given in [Definition 1.5](#) has the advantage that projections of spherical arrangements are spherical again. This can be quite convenient. We prove this property in [Proposition 1.9](#).

The following algebraic properties of normalized/spherical arrangements will be of great use:

Observation 1.8. Let v be normalized (it works analogously for spherical v) with arrangement matrix Φ . The following properties follow via basic linear algebra:

- (i) by definition, $\Phi^\top \Phi = \pi_W$ for some $W \subseteq \mathbb{R}^d$, and one can show that $W = \text{span } v = \text{span } \Phi^\top$ is the smallest subspace of \mathbb{R}^d that contains all the points of the arrangement.
- (ii) $\Phi \Phi^\top = \pi_U$, where $U \subseteq \mathbb{R}^n$ is the arrangement space of v .

Proposition 1.9. *If v is spherical resp. normalized and $W' \subseteq W := \text{span } v$, then the projection $v' := \pi_{W'} v$ is spherical resp. normalized as well.*

Proof. The arrangement $v' := \pi_{W'}v$ has arrangement matrix $\Phi' := \Phi\pi_{W'}$. It follows

$$\Phi'^{\top}\Phi' = \pi_{W'}\Phi^{\top}\Phi\pi_{W'} = \pi_{W'}\alpha\pi_W\pi_{W'} \stackrel{(*)}{=} \alpha\pi_{W'}$$

where in (*) we used that after projecting onto W' , a subsequent projection onto $W \supseteq W'$ does nothing. Likewise a third projection onto W' does nothing. \square

Our main reason for introducing spherical and normalized arrangements was to define a class of arrangements for which the arrangement space determines metric properties. And indeed, such arrangements are determined by their arrangement space up to orientation (if normalized) resp. up to orientation and scale (if spherical):

Lemma 1.10. *Let $v, v' : V \rightarrow \mathbb{R}^d$ be two normalized arrangements, then the following are equivalent:*

- (i) v and v' are equivalent (i.e., they have the same arrangement space), and
- (ii) $Xv = v'$ for some orthogonal transformation $X \in O(\mathbb{R}^d)$.

Proof. Let Φ, Φ' be the arrangement matrices of v resp. v' , and $U, U' \subseteq \mathbb{R}^n$ the corresponding arrangement spaces. Note also that (ii) \implies (i) already follows from [Corollary 1.4](#).

For the other direction set $W := \text{span } v$ and $W' := \text{span } v'$. Then $\dim W = \dim U = \dim W'$, and therefore also $\dim W^{\perp} = \dim W'^{\perp}$.

For the next step we need to choose a normalized arrangement $w : V \rightarrow \mathbb{R}^d$ with $\text{span } w = W^{\perp}$. Such exist: choose any full-dimensional normalized arrangement constructed via [Construction 1.6](#) and project it onto W^{\perp} (and use [Proposition 1.9](#)). Since W^{\perp} and W'^{\perp} are of the same dimension, we can reorient w via an orthogonal transformation to an arrangement w' with $\text{span } w' = W'^{\perp}$. By (ii) \implies (i) w and w' have the same arrangement space $\bar{U} \subseteq \mathbb{R}^n$.

Let Ψ and Ψ' be the arrangement matrices of w and w' respectively and define

$$X := \Phi'^{\top}\Phi + \Psi'^{\top}\Psi.$$

We show that this is the desired orthogonal transformation to establish (ii). We first show that indeed $Xv = v'$:

$$X\Phi^{\top} = \Phi'^{\top}\Phi\Phi^{\top} + \Psi'^{\top}\Psi\Phi^{\top} \stackrel{(*)}{=} \Phi'^{\top}\pi_U + 0 = (\pi_U\Phi')^{\top} \stackrel{(**)}{=} \Phi'^{\top},$$

where in (*) we used $\Phi\Phi^{\top} = \pi_U$ (see [Observation 1.8](#)) and $\Psi\Phi^{\top} = 0$ (the rows of Φ and Ψ are in orthogonal subspaces). In (**) we used that $\pi_U\Phi' = \Phi'$ since $\text{span } \Phi' = U$.

Similar use of the analogous identities for Φ', Ψ and Ψ' shows that X is orthogonal:

$$\begin{aligned} X^{\top}X &= (\Phi'^{\top}\Phi + \Psi'^{\top}\Psi)^{\top}(\Phi'^{\top}\Phi + \Psi'^{\top}\Psi) \\ &= \Phi^{\top}\Phi'\Phi'^{\top}\Phi + \Phi^{\top}\Phi'\Psi'^{\top}\Psi + \Psi^{\top}\Psi'\Phi'^{\top}\Phi + \Psi^{\top}\Psi'\Psi'^{\top}\Psi \\ &= \Phi^{\top}\pi_U\Phi + 0 + 0 + \Psi^{\top}\pi_{\bar{U}}\Psi \\ &= \Phi^{\top}\Phi + \Psi^{\top}\Psi = \pi_W + \pi_{W^{\perp}} = \text{Id} \end{aligned}$$

\square

One goal of the first two chapters is to build a dictionary that translates between properties of arrangements (resp. graph realizations) and properties of their arrangement spaces. This includes properties about symmetry and rigidity, but also metric properties. A trivial example is, if ν is full-dimensional, then its dimension matches $\dim U$. If ν is normalized, then [Lemma 1.10](#) tells us that the arrangement space also determines metric properties. For example, the *radius* $r(\nu)$ of an arrangement is defined as

$$[r(\nu)]^2 := \frac{1}{n} \sum_{i=1}^n \|v_i\|^2 = \frac{1}{n} \operatorname{tr}(\Phi^\top \Phi) = \frac{1}{n} \operatorname{tr}(\Phi \Phi^\top) \stackrel{(*)}{=} \frac{1}{n} \operatorname{tr}(\pi_U) = \frac{\dim U}{n}, \quad (1.2)$$

where $(*)$ applies if ν is normalized. This definition is useful when working with *inscribed* arrangements, *i.e.*, all vertices are on a sphere around the origin. Then $r(\nu)$ is the radius of that sphere.

1.2 Symmetric arrangements

A (*Euclidean*) *symmetry* of a point arrangement is an isometry (that is, a distance preserving transformation) of the ambient space, so that points of the arrangement are again mapped onto points of the arrangement. That is, a symmetry permutes the points.




Figure 1.3. If four points are arranged in the shape of a square like shown in the figure, then certain permutations of these points can be achieved by rigid motions of the square (rotations and reflections in the three left most images), but others cannot (the permutation in the right most image).

In this section we study point arrangements with prescribed symmetries. That means we prescribe permutations (actually, whole groups $\Sigma \subseteq \operatorname{Sym}(V)$ of permutations) that must be realized by an isometry of the ambient space. Slightly stronger, we require these isometries to be orthogonal transformations. Arrangements with prescribed symmetries will be shortly called *symmetric arrangements*.

To give a more formal definition, recall that an (orthogonal) representation $T: \Sigma \rightarrow \operatorname{O}(\mathbb{R}^d)$ of a permutation group $\Sigma \subseteq \operatorname{Sym}(V)$ is a group homomorphism into the orthogonal group $\operatorname{O}(\mathbb{R}^d)$ (for a sufficient introduction to representation theory, see [Appendix A](#)).

Definition 1.11. Let $\Sigma \subseteq \operatorname{Sym}(V)$ be a permutation group on the index set V .

A Σ -arrangement is an arrangement ν for which there exists a (linear orthogonal) representation $T : \Sigma \rightarrow O(\mathbb{R}^d)$ with

$$T_\sigma v_i = v_{\sigma(i)}, \quad \text{for all } i \in V \text{ and } \sigma \in \Sigma. \quad (1.3)$$

T is called a *representation* of ν . A (non-zero) Σ -arrangement is called *irreducible* if the subspace $\text{span } \nu \subseteq \mathbb{R}^d$ is T -irreducible. It is called *reducible* otherwise.

We emphasize that the representation $T : \Sigma \rightarrow O(\mathbb{R}^d)$ of a Σ -arrangement maps to the *orthogonal* matrices. This is purely a design choice: one could develop a completely parallel theory using representations $T : \Sigma \rightarrow GL(\mathbb{R}^d)$ instead. However, orthogonal transformations are closer to how we often think of geometric symmetries as rigid motions. All in all, this choice comes with no loss of generality, it makes some computations and proofs simpler, but requires more work at other places.

While the representation of a Σ -arrangement might not be unique, its action on $\text{span } \nu$ is unique. Whether T acts irreducibly on $\text{span } \nu$ is therefore independent of the choice of the representation, and hence the notion of irreducibility for Σ -arrangements is well-defined.

Equation (1.3) can be compactified using the arrangement matrix Φ of ν : $T_\sigma \Phi^\top = \Phi^\top \Pi_\sigma$ for all $\sigma \in \Sigma$, where $\Pi_\sigma \in \text{Perm}(\mathbb{R}^n)$ denotes the *permutation matrix* corresponding to σ . By transposing both sides of the equations, and using $T_\sigma^\top = T_{\sigma^{-1}}$, $\Pi_\sigma^\top = \Pi_{\sigma^{-1}}$ and $\sigma \in \Sigma \Leftrightarrow \sigma^{-1} \in \Sigma$, we arrive at the following equivalent form:

$$\Phi T_\sigma = \Pi_\sigma \Phi, \quad \text{for all } \sigma \in \Sigma. \quad (1.4)$$

Interpreting (1.4) in the language of representation theory, the arrangement matrix serves as an *equivariant map* $\Phi : \mathbb{R}^d \rightarrow \mathbb{R}^n$ between the representation T and the permutation representation $\sigma \mapsto \Pi_\sigma$ (cf. [Definition A.7](#)).

We can characterize symmetric arrangements via their arrangement space. Recall that a Σ -invariant subspace $U \subseteq \mathbb{R}^n$ satisfies $\Pi_\sigma U \subseteq U$ for all $\sigma \in \Sigma$ (see [Definition A.2](#)).

Lemma 1.12. *Let ν be an arrangement.*

- (i) *If ν is a Σ -arrangement, then its arrangement space is a Σ -invariant subspace of \mathbb{R}^n .*
- (ii) *If ν is spherical and its arrangement space is Σ -invariant, then ν is a Σ -arrangement.*

Together, (i) and (ii) almost provide a perfect characterization, except that we need ν to be spherical in (ii). For spherical arrangements being a Σ -arrangement and having a Σ -invariant subspace is equivalent, which provides an entry to our “arrangement space dictionary”.

Proof of Lemma 1.12. Suppose that ν is a Σ -arrangement with representation T and arrangement space $U := \text{span } \Phi$. Recall [Theorem 1.3](#), that Φ and ΦT_σ have the same column span because T_σ is invertible. We use this fact in (*):

$$\Pi_\sigma U = \text{span}(\Pi_\sigma \Phi) \stackrel{(1.4)}{=} \text{span}(\Phi T_\sigma) \stackrel{(*)}{=} \text{span } \Phi = U, \quad \text{for all } \sigma \in \Sigma.$$

¹For representations, we write T_σ instead of $T(\sigma)$.

Thus, U is Σ -invariant, which proves (i).

For (ii), assume that the arrangement space $U \subseteq \mathbb{R}^n$ is Σ -invariant and that ν is spherical. W.l.o.g. assume that ν is even normalized, that is, $\Phi^\top \Phi = \pi_W$, where $W := \text{span } \nu$ by [Observation 1.8](#). The crucial step is to define

$$T_\sigma := \Phi^\top \Pi_\sigma \Phi + \pi_{W^\perp}, \quad \text{for each } \sigma \in \Sigma.$$

It remains to prove that this defines an orthogonal representation T that satisfies (1.4). We show (1.4) by expanding the definition of T :

$$\begin{aligned} \Phi T_\sigma &= \Phi \Phi^\top \Pi_\sigma \Phi + \Phi \pi_{W^\perp} \\ &= \Phi \Phi^\top \Pi_\sigma \Phi + 0 && \text{by } \Phi \pi_{W^\perp} = (\pi_{W^\perp} \Phi^\top)^\top = 0 \text{ since } \text{span } \Phi^\top = W \\ &= \pi_U \Pi_\sigma \Phi && \text{by } \Phi \Phi^\top = \pi_U \text{ from } \text{Observation 1.8 (ii)} \\ &= \Pi_\sigma \pi_U \Phi && \text{by } \pi_U \Pi_\sigma = \Pi_\sigma \pi_U \text{ since } U \text{ is } \Sigma\text{-invariant} \\ &= \Pi_\sigma \Phi && \text{by } \pi_U \Phi = \Phi \text{ since } \text{span } \Phi = U. \end{aligned}$$

The same arguments using properties of Φ and Π_σ show that T_σ is orthogonal (which we prove with less explicit detail):

$$\begin{aligned} T_\sigma^\top T_\sigma &= (\Phi^\top \Pi_\sigma \Phi + \pi_{W^\perp})^\top (\Phi^\top \Pi_\sigma \Phi + \pi_{W^\perp}) \\ &= \Phi^\top \Pi_\sigma^\top \Phi \Phi^\top \Pi_\sigma \Phi + \Phi^\top \Pi_\sigma^\top \Phi \pi_{W^\perp} + \pi_{W^\perp}^\top \Phi^\top \Pi_\sigma \Phi + \pi_{W^\perp}^\top \pi_{W^\perp} \\ &= \Phi^\top \Pi_\sigma^\top \pi_U \Pi_\sigma \Phi + 0 + 0 + \pi_{W^\perp} \\ &= \Phi^\top \Pi_\sigma^\top \Pi_\sigma \pi_U \Phi + \pi_{W^\perp} \\ &= \Phi^\top \Phi + \pi_{W^\perp} = \pi_W + \pi_{W^\perp} = \text{Id}. \end{aligned}$$

The proof of $T_\sigma T_\tau = T_{\sigma\tau}$ is analogous. □

The condition ‘‘spherical’’ is necessary in (ii): a (generic) rectangle is just a linear transformation of a square, which means both share the same arrangement space. But clearly they do not share the same symmetries. In a sense, spherical arrangements are the most symmetric arrangements (in the Euclidean sense) among their equivalent counterparts.




Figure 1.4. The permutation $\sigma = (1234)$ can be realized for the square by a rotation, but not for the rectangle, even though Π_σ preserves the arrangement spaces of both arrangements. The reason is that the rectangle is not spherical.

We show that *irreducible* Σ -arrangements are always spherical. Recall that a Σ -invariant subspace is Σ -*irreducible* if it does not contain a non-zero proper Σ -invariant subspace.

Lemma 1.13. *Given an arrangement ν , the following are equivalent:*

- (i) ν is an irreducible Σ -arrangement,
- (ii) ν is spherical and its arrangement space $U \subseteq \mathbb{R}^n$ is Σ -irreducible.

Proof. Suppose (i) that ν is an irreducible Σ -arrangement with representation T . If Φ is its arrangement matrix, then applying (a variant of) (1.4) twice gives

$$T_\sigma \Phi^\top \Phi = \Phi^\top \Pi_\sigma \Phi = \Phi^\top \Phi T_\sigma.$$

We see that $\Phi^\top \Phi$ commutes with T_σ for all $\sigma \in \Sigma$. As a consequence, every eigenspace of $\Phi^\top \Phi$ is T -invariant. One such eigenspace is $\ker \Phi$ to eigenvalue zero. Since $\Phi^\top \Phi$ is symmetric and positive semi-definite, any other eigenspace $\text{Eig}(\alpha)$ corresponds to some positive eigenvalue $\alpha > 0$ and must be contained in $W := (\ker \Phi)^\perp = \text{span } \nu$. Since ν is irreducible, T acts irreducibly on W , and so W cannot have a proper invariant subspace. But $\text{Eig}(\alpha)$ is an invariant subspace of W , and therefore $W = \text{Eig}(\alpha)$. To summarize, $\Phi^\top \Phi$ has exactly two eigenvalues, 0 and $\alpha > 0$, and can therefore be written in the form $\Phi^\top \Phi = \alpha \pi_W + 0 \pi_{W^\perp} = \alpha \pi_W$. Hence, ν is spherical. This shows one part of (ii).

For the other part, consider the arrangement matrix as a linear isomorphism $\Phi: W \rightarrow U$ with $W := \text{span } \nu$ and $U \subseteq \mathbb{R}^n$ the arrangement space. Since W is T -invariant, T restricts to a Σ -representation $T|_W: \Sigma \rightarrow \text{O}(W)$ on W . Likewise, since U is Σ -invariant (by Lemma 1.12 (i)) we obtain the restricted Σ -representation $\sigma \mapsto \Pi_\sigma|_U \in \text{O}(U)$. By (1.4) Φ can be interpreted as an invertible equivariant map, establishing that these restricted representations are isomorphic. In particular, one is irreducible if and only if the other one is.

The claim then follows easily: if ν is irreducible, then W is T -irreducible. By the isomorphism, U is Σ -irreducible. This proves (ii).


Now assume (ii) that ν is spherical and U is Σ -irreducible. In particular, U is Σ -invariant. By Lemma 1.12 (ii), ν is then a Σ -arrangement with some representation T . We have just shown that if U is Σ -irreducible, then $W := \text{span } \nu$ is T -irreducible. This shows (i). \square

The previous results (Lemma 1.12 and Lemma 1.13) in conjunction with Construction 1.2 provide an explicit way to construct Σ -arrangements as soon as we have access to the Σ -invariant subspaces. This is especially useful in the presence of an additional graph structure: as we discuss in Chapter 2, the eigenspaces of a graph G provide easy access to some of the $\text{Aut}(G)$ -invariant subspaces.

Example 1.14. Consider the group $\Sigma := \langle (1234) \rangle \subseteq \text{Sym}(V)$ on $V = \{1, 2, 3, 4\}$ generated by a single cyclic permutation $(1234) \in \text{Sym}(V)$. One can show that \mathbb{R}^4 decomposes uniquely into three Σ -irreducible subspaces $U_1, U_2, U_3 \subseteq \mathbb{R}^4$

$$\begin{aligned} U_1 &= \text{span}\{(1, 1, 1, 1)^\top\} \\ U_2 &= \text{span}\{(1, -1, 1, -1)^\top\} \\ U_3 &= \text{span}\{(1, 1, -1, -1)^\top, (1, -1, -1, 1)^\top\} \end{aligned}$$


Via Construction 1.2 those translate to the following Σ -arrangements:



Note that several points can be positioned on top of each other. The subspace U_1 is invariant for every Σ and corresponds always to the arrangement with all vertices at a single point.

Since U_1, U_2 and U_3 are the only Σ -irreducible subspaces, this is a complete list of irreducible Σ -arrangements up to equivalence (by Lemma 1.13).

The remaining Σ -arrangements (the reducible ones) have arrangement spaces $U_1 \oplus U_2, U_1 \oplus U_3, U_2 \oplus U_3$ and $U_1 \oplus U_2 \oplus U_3 = \mathbb{R}^n$. The figure below depicts an arrangement to $U_2 \oplus U_3$. It projects to the two arrangements it is built from.



Another way to construct Σ -arrangements is via a given Σ -representation T :

Construction 1.15. Suppose that we are given a permutation group $\Sigma \subseteq \text{Sym}(V)$ and a representation $T : \Sigma \rightarrow O(\mathbb{R}^d)$ thereof. Our goal is to construct a Σ -arrangement with representation T .

Consider the decomposition $V = V_1 \cup \dots \cup V_m$ of V into its Σ -orbits and fix an $i_k \in V_k$ for each $k \in \{1, \dots, m\}$. The idea is to define a Σ -arrangement v by choosing the v_{i_k} more or less freely and let the other points be determined by T (see also Figure 1.5).

A completely free choice of the v_{i_k} might not be possible though: let $\Sigma_i := \{\sigma \in \Sigma \mid \sigma(i) = i\}$ denote the stabilizer of Σ at i . Then for any $\sigma \in \Sigma_{i_k}$ there must hold

$$T_\sigma v_{i_k} = v_{\sigma(i_k)} = v_{i_k}.$$

In other words, v_{i_k} must be chosen from the fix space

$$\text{Fix}(T, \Sigma_{i_k}) := \{x \in \mathbb{R}^d \mid T_\sigma x = x \text{ for all } \sigma \in \Sigma_{i_k}\} = \bigcap_{\sigma \in \Sigma_{i_k}} \ker(\text{Id} - T_\sigma),$$

which is itself a linear subspace of \mathbb{R}^d .

Since V_k is an orbit, for any $i \in V_k$ there exists a $\sigma_i \in \Sigma$ with $\sigma_i(i_k) = i$. If we now choose points $x_k \in \text{Fix}(T, \Sigma_{i_k})$ for all $k \in \{1, \dots, m\}$, we can define

$$v_i := T_{\sigma_i} x_k, \quad \text{whenever } i \in V_k.$$




Figure 1.5. Visualization of **Construction 1.15** using $\Sigma := \langle (135)(246), (12)(36)(45) \rangle$. Only the white point was chosen freely. If the $T_\sigma, \sigma \in \Sigma$ include reflections on the three gray lines, then the black points are enforced by the assumption that v is a Σ -arrangement with representation T .

It is then easy to show that v is a Σ -arrangement with representation T and $v_{i_k} = x_k$ for all $k \in \{1, \dots, m\}$. One can show that the definition is independent of the choice of the σ_i .

Note that a necessary criterion for the existence of a non-zero Σ -arrangement with representation T is that $\text{Fix}(T, \Sigma_{i_k}) \neq \{0\}$ for some $k \in \{1, \dots, m\}$.

1.3 Separating symmetries

We have developed enough terminology to discuss **Question 8** raised in the introduction: how to detect whether two permutation groups can be “separated geometrically”. In the language of point arrangements, this reads

Question 1.16. Given a permutation group $\Sigma \subset \text{Sym}(V)$. When does it happen that every Σ -arrangement is also a Σ' -arrangement for a larger group $\Sigma' \supset \Sigma$?

Example 1.17. Each arrangement in **Example 1.14** has a symmetry that permutes its points by $(12)(34) \notin \Sigma$ (which might be the identity transformation). That is, $\Sigma := \langle (1234) \rangle$ cannot be separated geometrically from the larger group $\Sigma' := \langle (1234), (12)(34) \rangle$ (the *dihedral group* on four elements).

The same phenomenon happens for all cyclic groups $\Sigma := \langle (123 \dots n) \rangle, n \geq 3$.

Recall that the initial motivation for **Question 8** resp. **Question 1.16** was the observation that a “sufficiently large” group $\Sigma \subset \text{Aut}(G)$ of symmetries of a graph G cannot be geometrically separated from $\text{Aut}(G)$ itself, where the meaning of “sufficiently large” depends on spectral properties of G . We come back to this in **Chapter 2**. For now we shall take a look at how **Question 1.16** is related to the classification of symmetry groups of geometric objects, and how similar questions have been asked and answered in the literature before.

Group theory is often motivated as the study of symmetries of discrete or geometric objects, such as graphs or polytopes. However, the modern definition of an abstract group makes no use of this connection, and so one might ask whether this motivation is still justified.

Given an abstract group \mathcal{G} , is there a graph/polytope whose symmetry group is isomorphic to \mathcal{G} ?

These questions have since been answered in the affirmative (for graphs [27] and for polytopes [67]), and the focus shifted to modified versions of these questions.

For one possible modification, we restrict to a proper sub-class of polytopes (or graphs). For example, not every finite group appears as the symmetry group of a *vertex-transitive* polytope (a polytope with a single orbit of vertices, see also Definition 2.14). Most cyclic groups are such exceptions.




Figure 1.6. The polygons with the black outline have no mirror symmetry, and $\text{Aut}(P)$ is isomorphic to a cyclic group. These polygons are not vertex-transitive as there are interior angles of different sizes.

A complete classification of these exceptions has been obtained in [2], and in a slightly more general setting (linear symmetries instead of Euclidean symmetries) in [24–26]. Some work has been done for centrally symmetric polytopes in [14].

Another way to modify the polytope version of the question is to replace abstract groups with something more concrete like matrix groups of permutation groups:

- Given a *matrix group* $\Gamma \subseteq O(\mathbb{R}^d)$, is there a polytope $P \subset \mathbb{R}^d$ with $\text{Aut}(P) = \Gamma$ (identical as matrix groups, not just isomorphic)?

The answer to this question is affirmative too. Take a sufficiently large set of generically chosen points $x_1, \dots, x_m \in S^{d-1} \subset \mathbb{R}^d$ on the unit sphere, and define $P := \text{conv}\{Tx_i \mid T \in \Gamma, i \in \{1, \dots, m\}\}$. Then $\text{Aut}(P) = \Gamma$ holds almost surely.

This question has also been asked for vertex-transitive polytopes (see e.g. [24, Chapter 7] for an algebraic approach, or see [23] for unitary symmetry groups).

- Given a *permutation group* $\Sigma \subseteq \text{Sym}(V)$, is there a polytope P whose vertices $v_1, \dots, v_n \in \mathcal{F}_0(P)$ are permuted by $\text{Aut}(P)$ exactly in the ways described by Σ ?

That is, all permutations of Σ are realized by $\text{Aut}(P)$, but also, no other permutations are realized besides these.

Given that a polytope, and its symmetry too, is completely determined by its vertices, all these questions have an equivalent formulation in terms of point arrangements. In particular the permutation group version, expressed for Σ -arrangements, becomes

Given a permutation group $\Sigma \subseteq \text{Sym}(V)$, is there a Σ -arrangement that is not a Σ' -arrangement for any larger group $\Sigma' \supset \Sigma$?

This is just another formulation of Question 1.16. We have then seen that classifying possible symmetries of geometric objects might require an explanation for why certain symmetries cannot be “separated” from larger symmetry groups. For example, Example 1.17 shows that $\Sigma := \langle (1234) \rangle$ cannot be the full (permutation) symmetry group of any arrangement.

To our knowledge “realizable permutation groups” have not been classified, and it could be infeasible to do so.

For later use we prove a criterion for detecting separation of symmetries using the arrangement space (see [Theorem 1.19](#)). We need the following technical proposition:

Proposition 1.18. *Let $v: V \rightarrow \mathbb{R}^d$ be a point arrangement.*

- (i) *If v is a Σ -arrangement with representation T , and $W \subseteq \mathbb{R}^d$ is a T -invariant subspace, then $v' := \pi_W v$ is also a Σ -arrangement. If W is T -irreducible, then v' is an irreducible Σ -arrangement.*
- (ii) *Let $\mathbb{R}^d = W_1 \oplus \cdots \oplus W_m$ be a decomposition of \mathbb{R}^d into pair-wise orthogonal subspaces $W_1, \dots, W_m \subseteq \mathbb{R}^d$. If $v^k := \pi_{W_k} v$ is a Σ -arrangement for all $k \in \{1, \dots, m\}$, then so is v .*

Proof. Let v be a Σ -arrangement with representation T , and $W \subseteq \mathbb{R}^d$ a T -invariant subspace. We show that $v' := \pi_W v$ is a Σ -arrangement with the same representation:

$$T_\sigma v'_i = T_\sigma \pi_W v_i = \pi_W T_\sigma v_i = \pi_W v_{\sigma(i)} = v'_{\sigma(i)},$$

where π_W commutes with T_σ because $W \subseteq \mathbb{R}^d$ is T -invariant.

Suppose now that W is T -irreducible. Since $W' := \text{span } v'$ is T -invariant and $W' \subseteq W$, we either have $W' = \{0\}$ or $W' = W$. In either case T acts irreducibly on W' . This proves (i).

For (ii) let $\mathbb{R}^d = W_1 \oplus \cdots \oplus W_m$ be a decomposition into pair-wise orthogonal subspaces. Let $T^k: \Sigma \rightarrow \text{O}(\mathbb{R}^d)$ be the representation of $v^k := \pi_{W_k} v$. Define

$$T_\sigma := T_\sigma^1 \pi_{W_1} + \cdots + T_\sigma^m \pi_{W_m}, \quad \text{for all } \sigma \in \Sigma.$$

We show that v is a Σ -arrangement with representation T :

$$\begin{aligned} T_\sigma v_i &= T_\sigma^1 \pi_{W_1} v_i + \cdots + T_\sigma^m \pi_{W_m} v_i \\ &= T_\sigma^1 v_i^1 + \cdots + T_\sigma^m v_i^m \\ &= v_{\sigma(i)}^1 + \cdots + v_{\sigma(i)}^m \\ &= \pi_{W_1} v_{\sigma(i)} + \cdots + \pi_{W_m} v_{\sigma(i)} \\ &= \underbrace{(\pi_{W_1} + \cdots + \pi_{W_m})}_{\text{Id}} v_{\sigma(i)} = v_{\sigma(i)}. \end{aligned}$$

One checks that T_σ is indeed orthogonal and defines a representation T . □

Theorem 1.19. *If every Σ -invariant subspace is also Σ' -invariant, then every Σ -arrangement is also a Σ' -arrangement.*

Proof. Suppose that every Σ -invariant subspace of \mathbb{R}^n is also Σ' -invariant, and that v is a Σ -arrangement with representation T .

Then \mathbb{R}^d decomposes as a direct sum $W_1 \oplus \cdots \oplus W_m$ into pair-wise orthogonal T -irreducible subspaces $W_1, \dots, W_m \subseteq \mathbb{R}^d$. Consider the projections $v^k := \pi_{W_k} v$, which are then irreducible Σ -arrangements by [Proposition 1.18](#) (i). In particular, they are spherical by [Lemma 1.13](#).

The arrangement space $U_k \subseteq \mathbb{R}^n$ of ν^k is a Σ -invariant subspace by Lemma 1.12 (i), and by assumption it is also Σ' -invariant. Since ν^k is spherical, it is a Σ' -arrangement by Lemma 1.12 (ii). We then found that each $\pi_{W_k} \nu$ is a Σ' -arrangement for a complete decomposition into pair-wise orthogonal subspaces $W_1 \oplus \dots \oplus W_m$, and hence ν itself is a Σ' -arrangement by Proposition 1.18 (ii). \square

The reason for $\Sigma := \langle(1234)\rangle$ being not geometrically separable from the larger group $\Sigma' := \langle(1234), (12)(34)\rangle$ is, that these have the exact same invariant subspaces.

Other examples of this kind will be discussed in Chapter 2 (see especially Section 2.4 and Corollary 2.35 (ii)).


1.4 Deformation and rigidity

In this section we address Question 7 raised in the introduction: how does symmetry influence rigidity properties of arrangements? For that purpose we consider continuous *deformations* of arrangements that preserve their prescribed symmetries (see Definition 1.21 for formalization of the notions). If no such deformations exist, then the Σ -arrangement can be considered as *rigid*, otherwise as *flexible*.

Example 1.20. Figure 1.5 (on page 37) gives a first impression of how a deformation might look like. Moving the white point changes the shape of the whole arrangement, but keeps its representation, and hence, its symmetry. This arrangement is therefore flexible.

In order to make rigidity a meaningful notion, we have to exclude some undesirable “deformations”: each arrangement can be “deformed” by a continuous change in scale or orientation. Such should not be considered *proper* deformations. We observe that a change in scale or orientation does not change the arrangement space. We can use this to define what makes a deformation proper: it has to change the arrangement space. This definition would also recognize other “trivial deformations” as non-proper, such as non-uniform stretching, shearing, and other more general linear transformations. Note the following subtlety: while it is straightforward to continuously deform an arrangement ν into $X\nu$ for $X \in GL(\mathbb{R}^d)$ with $\det X > 0$, whether such a deformation is possible for $\det X < 0$ is non-trivial. This subtlety was explored in [76] and we shall exclude it from our discussion.

Similarly undesirable, every Σ -arrangement can be shrunk into a single point, which is a Σ -arrangement for all Σ . This can be used to transform any Σ -arrangement into any other Σ -arrangement of the same dimension. For example, the pentagon arrangement can be transformed into the pentagram arrangement in this way:



Both are Σ -arrangements for $\Sigma := \langle(12345)\rangle$. No “non-collapsing” deformation exists between these. Similar problems arise if we allow arrangements to collapse partially in only some directions. To avoid all these problems entirely, and for maximal simplicity, we may assume that the arrangement is full-dimensional during the whole deformation. In particular, we shall investigate deformations for full-dimensional arrangements only.

To formalize deformations, we need to declare a topology on the space of arrangements: the set of d -dimensional arrangements on the index set $V = \{1, \dots, n\}$ can be associated with its set of arrangement matrices, that is $\mathbb{R}^{n \times d}$. The space $\mathbb{R}^{n \times d}$ is naturally equipped with the topology of a vector space, and we can pull back this topology to the set of arrangements.

For $\Sigma \subseteq \text{Sym}(V)$ we then study the following subspace of arrangements:

$$\mathcal{A}_d(\Sigma) := \{v: V \rightarrow \mathbb{R}^d \mid v \text{ is a full-dimensional } \Sigma\text{-arrangement}\}.$$

We consider $\mathcal{A}_d(\Sigma)$ equipped with the respective subspace topology.

Definition 1.21. Given a $\Sigma \subseteq \text{Sym}(V)$.

- (i) A Σ -deformation is a continuous curve $v(\cdot): [0, 1] \rightarrow \mathcal{A}_d(\Sigma)$. We say that the arrangements $v(0)$ and $v(1)$ can be Σ -deformed into each other.
- (ii) A Σ -deformation is *proper* if not all $v(t)$ are equivalent.
- (iii) A Σ -arrangement is said to be Σ -flexible, if it can be Σ -deformed into a non-equivalent arrangement. Otherwise it is called Σ -rigid.

If Σ is clear from the context we just write *deformation*, *rigid* and *flexible*, etc.

If $v(\cdot)$ is a Σ -deformation, then to each arrangement $v(t)$ we can assign the arrangement matrix $\Phi(t)$ and arrangement space $U(t) := \text{span } \Phi(t)$. The map $t \mapsto \Phi(t)$ is continuous by definition of the topology on $\mathcal{A}_d(\Sigma)$. Since $v(t)$ is full-dimensional we also have $\dim U(t) = d$ for all $t \in [0, 1]$.

The goal of this section is to work out how Σ -rigidity can be detected from the arrangement space, contributing to our dictionary. We shall see that this has everything to do with the number and relative positioning of Σ -invariant subspaces in \mathbb{R}^n .

The core intuition is the following: if $v(\cdot)$ is a Σ -deformation, and $U(t) \subseteq \mathbb{R}^n$ is the arrangement space of $v(t)$, then intuitively, a continuous change in $v(t)$ comes with a continuous change in $U(t)$. To make this precise we would need to define a topology on the “space of linear subspaces of \mathbb{R}^n ” (this is possible and the space is called the *Grassmannian*). However, for our purpose it suffices to observe and formalize the following:

- By [Lemma 1.12](#), each $U(t)$ is Σ -invariant. If $v(\cdot)$ is a proper deformation, then not all $U(t)$ are the same, and continuity suggests that $U(t)$ attains infinitely many different values, each of which is a Σ -invariant subspaces of \mathbb{R}^n . In particular, there are infinitely many Σ -invariant subspaces to begin with.
- If $U(0)$ and $U(1)$ are orthogonal, it appears intuitive that as t transitions from 0 to 1, the angle between $U(t)$ and $U(0)$ cannot suddenly jump from 0° to 90° but must attain all intermediate values. We would then expect to find *non-orthogonal* Σ -invariant subspaces.

One might assume that every permutation group $\Sigma \subseteq \text{Sym}(V)$ induces a *unique* decomposition of \mathbb{R}^n into pair-wise *orthogonal* Σ -irreducible subspaces (see also [Observation A.4](#)). We shall see that flexibility in arrangements occurs exactly when this intuition fails².

This is a good moment to recall the following properties of Σ -invariant subspaces (consider [Appendix A](#) for a complete overview):

- Two irreducible subspaces are either identical or have trivial intersection (see [Observation A.3](#)).
- The projection of an invariant subspace onto an invariant subspace is again invariant (see [Observation A.5](#)), and the projection of an *irreducible* subspace onto an invariant subspace is irreducible again (see [Proposition A.6](#)).
- If two irreducible subspaces are non-orthogonal, then they are of the same dimension (see [Lemma A.9 \(i\)](#)).

We first characterize rigidity of a single irreducible arrangement (see [Theorem 1.23](#)). Because its argument is used repeatedly, the following technical proposition is proven separately:

Proposition 1.22. *Let $v(\cdot)$ be a Σ -deformation and $U(t) \subseteq \mathbb{R}^n$ the arrangement space of $v(t)$. If $U \subseteq \mathbb{R}^n$ is a linear subspace, then the sets*

$$I := \{t \in [0, 1] \mid U(t) \subseteq U\} \subseteq [0, 1]$$

$$J := \{t \in [0, 1] \mid U(t) = U\} \subseteq [0, 1]$$

are closed subsets of $[0, 1]$.

Proof. Let $\Phi(t)$ be the arrangement matrix of $v(t)$ and let $u_k(t)$ be its k -th column. All of these are continuous functions in t . The set

$$I_k := \{t \in [0, 1] \mid u_k(t) \in U\} = u_k^{-1}[U], \quad \text{for } k \in \{1, \dots, d\}$$

is the preimage of a closed set U w.r.t. a continuous function $u_k(\cdot)$, and is therefore closed in $[0, 1]$. Since $U(t) = \text{span}\{u_1(t), \dots, u_d(t)\}$ we have $U(t) \subseteq U \Leftrightarrow u_1(t), \dots, u_d(t) \in U$. Hence $I = I_1 \cap \dots \cap I_k$ and I is closed in $[0, 1]$.

Note that if $\dim U \neq d$ then $J = \emptyset$, and if $\dim U = d$ then $I = J$, establishing that J is closed too. \square

Theorem 1.23. *Let $v \in \mathcal{A}_d(\Sigma)$ be irreducible with arrangement space $U \subseteq \mathbb{R}^n$ and representation T . The following are equivalent:*

- (i) v is flexible.
- (ii) there is a Σ -irreducible subspace $U' \neq U$ that is non-orthogonal to U (i.e., $U' \not\subseteq U^\perp$).
- (iii) there is $v' \in \mathcal{A}_d(\Sigma)$ that is non-equivalent to v but has the same representation T .

²In the language of character theory, flexibility occurs whenever the character of the permutation representation is not multiplicity-free. We shall not pursue this perspective but work in terms of invariant subspaces.

Proof. In the following let Φ denote the arrangement matrix of ν . We prove the implications (i) \implies (ii) \implies (iii) \implies (i).

Assume (i) that ν is flexible. Then there is a proper Σ -deformation $\nu(\cdot)$ with $\nu(0) = \nu$. Let $U(t) \subseteq \mathbb{R}^n$ be the arrangement space of $\nu(t)$. The sets

$$\begin{aligned} I &:= \{t \in [0, 1] \mid U(t) \subseteq U^\perp\}, \\ J &:= \{t \in [0, 1] \mid U(t) = U\}, \end{aligned}$$

are closed by [Proposition 1.22](#). Since $U \neq \{0\}$, I and J are clearly disjoint. Furthermore, $I \neq [0, 1]$, and since $\nu(\cdot)$ is proper we also have $J \neq [0, 1]$. But $[0, 1]$ is connected, and thus not the disjoint union of two closed proper subsets. So there must be a $t' \in [0, 1] \setminus (I \cup J)$. Then $U' := U(t')$ satisfies $U' \neq U$ and $U' \not\subseteq U^\perp$. It remains to show that U' is Σ -irreducible. We know that it is Σ -invariant by [Lemma 1.12 \(i\)](#). It then has a Σ -irreducible subspace $U'' \subseteq U'$ which is also non-orthogonal to U . But non-orthogonal irreducible subspaces agree in dimension (see [Lemma A.9 \(i\)](#)), i.e., $\dim U'' = \dim U = d = \dim U'$, and thus $U'' = U'$. Thus, U' is irreducible, which proves (ii).

Next, assume (ii). Let ν' be the arrangement with arrangement matrix $\Phi' := \pi_{U'}\Phi$. Since U and U' are non-orthogonal, $\nu' \neq 0$. The non-zero arrangement space of ν' is $\pi_{U'}U \subseteq U'$. This is a projection of an invariant subspace onto an invariant subspace, hence itself invariant (see [Observation A.5](#)). But since U' is irreducible, it cannot have a proper non-zero invariant subspace, and so we found that ν' must have arrangement space U' . Next, we show that ν' is a Σ -arrangement with representation T by showing that [\(1.4\)](#) holds:

$$\Phi' T_\sigma = \pi_{U'} \Phi T_\sigma \stackrel{(*)}{=} \pi_{U'} \Pi_\sigma \Phi \stackrel{(**)}{=} \Pi_\sigma \pi_{U'} \Phi = \Pi_\sigma \Phi', \quad \text{for all } \sigma \in \Sigma,$$

where in $(*)$ we used [\(1.4\)](#) for ν , and in $(**)$ we used that $\pi_{U'}$ commutes with Π_σ because U' is Σ -invariant. It remains to show that ν' is full-dimensional: U and U' are non-orthogonal irreducible subspace. As such they are of the same dimension (see [Lemma A.9 \(i\)](#)). That is, $\dim U' = \dim U = d$, and ν' is full-dimensional. This proves (iii).

Finally, assume (iii) and consider the following parametrized arrangement:

$$\nu(t) := (1-t)\nu + t\nu', \quad \text{for } t \in [0, 1]. \tag{1.5}$$

We show that this defines a Σ -deformation. By $T_\sigma \nu_i(t) = (1-t)T_\sigma \nu_i + T_\sigma \nu'_i = (1-t)\nu_{\sigma(i)} + \nu'_{\sigma(i)} = \nu_{\sigma(i)}(t)$, $\nu(t)$ is a Σ -arrangement with representation T for each $t \in [0, 1]$. It remains to show that $\nu(t)$ is full-dimensional. Since T is irreducible and $\text{span } \nu(t)$ is T -invariant, $\nu(t)$ must be either full-dimensional or zero. If $\nu(t) = (1-t)\nu + t\nu' = 0$, then $\nu = t/(t-1)\nu'$, in contradiction to the assumption that ν and ν' are non-equivalent. Thus, $\nu(t)$ is always a full-dimensional Σ -arrangement (with representation T) and $\nu(\cdot)$ therefore a proper Σ -deformation from ν into ν' , establishing that ν is flexible. This proves (i). \square

[Theorem 1.23](#) is not addressing *reducible* arrangements. Similar criteria for rigidity could be written down here, but they are more tedious to formulate precisely. We do not follow this path as we later focus on irreducible arrangements anyway. Note however that reducible arrangements can be rigid in a somewhat counter-intuitive way:




Figure 1.7. The depicted Σ -deformation for $\Sigma := \langle (13)(24) \rangle$ is not proper, even though it is intuitively “deforming” the arrangement. In fact, this Σ -arrangement is rigid by our definition. The colors of the points indicate orbits.

The next theorem is similar to [Theorem 1.23](#), but characterizes rigidity for *all* Σ -arrangements simultaneously.

Theorem 1.24. *Given $\Sigma \subseteq \text{Sym}(V)$, then the following are equivalent:*

- (i) *all $v \in \mathcal{A}_d(\Sigma)$ are rigid.*
- (ii) *\mathbb{R}^n can be written as a direct sum of Σ -irreducible subspaces in a unique way.*
- (iii) *there exist only finitely many Σ -irreducible subspaces.*

Proof. We prove (i) \implies (ii),(iii), and then (ii) \implies (iii) \implies (i).

Assume (i) that all (full-dimensional) Σ -arrangements are rigid and let

$$\mathbb{R}^n = U_1 \oplus \cdots \oplus U_m \tag{1.6}$$

be a decomposition into a direct sum of pair-wise orthogonal Σ -irreducible subspaces $U_1, \dots, U_m \subseteq \mathbb{R}^n$. Let $U \subseteq \mathbb{R}^n$ be some Σ -irreducible subspace. According to [Theorem 1.23](#) (ii) \implies (i), if $U \neq U_k$ then U and U_k must be orthogonal (as otherwise there are flexible arrangements). Clearly, U cannot be orthogonal to all U_k (as they span all of \mathbb{R}^n), and thus $U = U_k$ for some $k \in \{1, \dots, m\}$.

From this now follows (iii): every irreducible subspace must be one of the U_k , and so there are only finitely many. To show (ii), note that any other decomposition $\mathbb{R}^n = U'_1 \oplus \cdots \oplus U'_m$ must use some irreducible subspace not among the U_k . Such do not exist. Since every irreducible subspace is contained in a decomposition of the form (1.6) (see also [Observation A.4](#)), we also obtain (ii) \implies (iii).

To establish (iii) \implies (i), assume that there are only finitely many irreducible subspaces. Reducible subspaces are direct sums of irreducible subspaces, and hence there are also only finitely many Σ -invariant subspaces. Let $v(\cdot)$ be a Σ -deformation. We try to show that $v(\cdot)$ cannot be proper, establishing (i), that all (full-dimensional) Σ -arrangements are rigid. Let $U(t) \subseteq \mathbb{R}^n$ be the arrangement space of $v(t)$. Since each $U(t)$ is Σ -invariant (by [Lemma 1.12](#) (i)) and there are only finitely many Σ -invariant subspaces, $U(t)$ can attain only finitely many values $U_1, \dots, U_m \subseteq \mathbb{R}^n$. The disjoint sets

$$I_k := \{t \in [0, 1] \mid U(t) = U_k\}, \quad \text{for } k \in \{1, \dots, m\}$$

are then closed by [Proposition 1.22](#) and satisfy $I_1 \cup \cdots \cup I_m = [0, 1]$. Since $[0, 1]$ is connected, it cannot be the disjoint union of finitely many closed proper subsets. Hence, one of the I_k

is not a proper subset, but instead $I_k = [0, 1]$. Then $U(t) = U_k$ for all $t \in [0, 1]$ and so $\nu(\cdot)$ is not proper. \square

We have reached a notion of rigidity for permutation group: a group $\Sigma \subseteq \text{Sym}(V)$ satisfying the equivalent conditions in [Theorem 1.24](#) is called *rigid*.

As we shall discuss in [Chapter 2](#), “sufficient symmetries” in the sense of [Question 6](#) have to be rigid in this sense.

Example 1.25. $\Sigma := \text{Sym}(V)$ is rigid: the only Σ -irreducible subspaces are $\text{span}\{(1, \dots, 1)\}$ and its orthogonal complement, which are finitely many.

Similarly, it can be worked out that cyclic groups $\Sigma := \langle (123 \cdots n) \rangle$ are rigid, and so are the dihedral groups with which they share the same invariant subspaces (cf. [Example 1.14](#)).

1.5 Transitive arrangements

A permutation group $\Sigma \subseteq \text{Sym}(V)$ is said to act *transitively* (or later, *vertex-transitively*) on the underlying set V , if for any two $i, j \in V$ there is a permutation $\sigma \in \Sigma$ with $\sigma(i) = j$. We call a Σ -arrangement *transitive* if Σ is transitive.

In later chapters we mostly work with (vertex-)transitive arrangements, realizations and polytopes. One reason is our interest in rigidity. While rigidity does not imply transitivity, it comes pretty close:

Lemma 1.26. *If $\nu \in \mathcal{A}_d(\Sigma)$ is irreducible and rigid, then it is “essentially transitive”, that is, Σ acts transitively on $V^* := \{i \in V \mid v_i \neq 0\}$.*

The intuition behind this lemma is the following: if there is more than one non-zero orbit, then these orbits can be “deformed” independently, giving rise to a proper deformation and making the arrangement flexible. Irreducibility is necessary: the arrangement in [Figure 1.7](#) (on page [44](#)) is rigid but not “essentially transitive”.

Proof of Lemma 1.26. Consider the decomposition $V = V_1 \cup \cdots \cup V_m$ of the index set into its Σ -orbits and set $\mathcal{V}_k := \{v_i \mid i \in V_k\}$. If T denotes the representation of ν , then $\text{span } \mathcal{V}_k \subseteq \mathbb{R}^d$ is T -invariant. But since ν is irreducible, we must have $\text{span } \mathcal{V}_k = \mathbb{R}^d$ or $\text{span } \mathcal{V}_k = \{0\}$ for all $k \in \{1, \dots, m\}$ (recall that $\nu \in \mathcal{A}_d(\Sigma)$ is full-dimensional).

Suppose then that $\text{span } \mathcal{V}_k = \mathbb{R}^d$ for at least two $k \in \{1, \dots, m\}$, say for $k = 1$ and $k = 2$, and consider the following parametrized arrangement

$$v_i(t) := \begin{cases} (1-t)v_i & \text{if } i \in V_1 \\ v_i & \text{otherwise} \end{cases}.$$

We show that it is a Σ -deformation. One checks easily that $\nu(t)$ is indeed a Σ -arrangement for all $t \in [0, 1]$ with the representation T . Since \mathcal{V}_2 is full-dimensional and $v_i(t) = v_i$ for all $i \in V_2$, $\nu(t)$ is full-dimensional for all $t \in [0, 1]$.

We show that $\nu(\cdot)$ is a *proper* deformation: any potential linear transformation mapping $\nu(0)$ onto $\nu(1)$ must fix the full-dimensional set \mathcal{V}_2 , but must map the other full-dimensional set \mathcal{V}_1 to zero. This cannot be.

We conclude that at most one (actually, exactly one) of the span \mathcal{V}_k is full-dimensional, say \mathcal{V}_1 . In other words, for all $i \in V$ we have $v_i \neq 0 \Leftrightarrow i \in V_1$. Then $V^* = V_1$, and since Σ -acts transitively on V_1 we are done. \square

Lemma 1.26 guides us best in the study of symmetric graph realizations in **Chapter 2**, when rigidity turns out to be important for pinning down “sufficient symmetry”. It then tells us to focus on vertex-transitive realizations.

It is less applicable in the case of polytopes (see **Figure 1.8**), which have additional structure to keep them rigid (recall Cauchy’s rigidity theorem, cf. **Theorem C.2**).




Figure 1.8. The *rhombic dodecahedron* (center) is a polyhedron with two orbits on vertices (emphasized by the point colors). Its skeleton is not “essential transitive” and can therefore be properly deformed while keeping its symmetry. The polyhedron itself however has a unique realization of this symmetry (it is a *perfect polytope*, see [28]), and in fact, is spectral (cf. **Example 3.12**).

Lemma 1.26 does also *not* state that transitive arrangements are rigid. This is not even true for polytopes (e.g. consider the *truncated tetrahedron* in **Example 2.6**).

However, transitivity allows us to give simple criteria for rigidity:

Lemma 1.27. *Given a transitive group $\Sigma \subseteq \text{Sym}(V)$ and a representation T thereof. Recall the fix space from **Construction 1.15**:*

$$\text{Fix}(T, \Sigma_i) := \{x \in \mathbb{R}^d \mid T_\sigma x = x \text{ for all } \sigma \in \Sigma_i\}.$$

If $\dim \text{Fix}(T, \Sigma_i) = 1$ for some (and then all) $i \in V$, then every irreducible $v \in \mathcal{A}_d(\Sigma)$ with representation T is rigid.

Proof. Let v be an irreducible Σ -arrangement with representation T . If v is flexible, then by **Theorem 1.23** there exists a non-equivalent $v' \in \mathcal{A}_d(\Sigma)$ with the same representation T . But by **Construction 1.15** v and v' are already uniquely determined by the positioning of, say, v_i resp. v'_i . Recall that, also by **Construction 1.15**, $v_i, v'_i \in \text{Fix}(T, \Sigma_i)$. By assumption, $\dim \text{Fix}(T, \Sigma_i) = 1$ and thus $v'_i = \alpha v_i$ for some $\alpha \neq 0$. This extends to $v' = \alpha v$ by transitivity, in contradiction to v' and v being non-equivalent. Thus, v must have been rigid. \square

The irreducibility assumption in **Lemma 1.27** is not a real restriction, as one can show that $\dim \text{Fix}(T, \Sigma_i) > 1$ for all reducible arrangements anyway. Also, 1-dimensional arrangements are always irreducible, and from **Lemma 1.27** then follows

Corollary 1.28. *All 1-dimensional transitive arrangements are rigid.*

The converse of [Lemma 1.27](#) is not true, that is, not every transitive arrangement satisfying $\dim \text{Fix}(T, \Sigma_1) \geq 2$ is automatically flexible, not even if it is irreducible. However, examples are surprisingly rare, and we shall call them *exceptionally rigid*.

Digression: exceptional rigidity

The following discussion of exceptional rigidity is not necessary for the development of the thesis, but it shows how exceptional mathematical objects emerge when studying rigidity.

A transitive group $\Sigma \subseteq \text{Sym}(V)$ is *regular*, if $|\Sigma| = |V|$, or equivalently, if for any $i, j \in V$ there is *exactly* one $\sigma \in \Sigma$ with $\sigma(i) = j$. Regular permutation groups have trivial stabilizers, that is, $\Sigma_i = \{\text{id}\}$ for all $i \in V$. For any Σ -representation $T : \Sigma \rightarrow O(\mathbb{R}^d)$ the fix space then is


$$\text{Fix}(T, \Sigma_i) = \mathbb{R}^d,$$

and [Lemma 1.27](#) is no help in concluding rigidity for any Σ -arrangement. Rigidity can still occur though.

Example 1.29. Cyclic groups $\Sigma := \langle (123 \cdots n) \rangle$ are regular. Consider the case $n = 4$ and the 2-dimensional Σ -invariant subspace from [Example 1.14](#):

$$U = \text{span}\{(1, 1, -1, -1), (1, -1, -1, 1)\}.$$

We have seen that it gives rise to a Σ -arrangement in the shape of a square. Its representation T maps (1234) onto the 90° -rotation clock-wise. By [Construction 1.15](#), a Σ -arrangement with representation T is uniquely determined by the placement of v_1 .



However, all these arrangements turn out to be just rotated squares and are thus equivalent.

By [Theorem 1.23](#) flexibility is equivalent to the existence of a non-equivalent arrangement with the same representation. As we have seen, such do not exist here, and the square is rigid despite $\dim \text{Fix}(T, \Sigma_1) \geq 2$.

Equivalently, one finds that the vertices of any regular n -gon form a Σ -arrangements for $\Sigma = \langle (123 \cdots n) \rangle$, which then too is exceptionally rigid.

Besides in dimension two, the only other examples of this kind appear in dimension four.

Theorem 1.30. *If $\Sigma \subseteq \text{Sym}(V)$ acts regularly and $v \in \mathcal{A}_d(\Sigma)$ is irreducible and exceptionally rigid, then $d = 2$ or $d = 4$.*

Proof. Since $v \in \mathcal{A}_d(\Sigma)$ is full-dimensional and irreducible, its representation T is irreducible as well. By [Construction 1.15](#), to every $x \in \text{Fix}(T, \Sigma_1) = \mathbb{R}^d$ exists a unique Σ -arrangement $v(x)$ with representation T and $v_1(x) = x$. Since T is irreducible, $v(x)$ is full-dimensional, in particular, $v(x) \in \mathcal{A}_d(\Sigma)$.

By assumption, ν is rigid, and by [Theorem 1.23](#) (i) \Rightarrow (ii) every Σ -arrangement with representation T must be equivalent to ν . In particular, $\nu(x)$ is equivalent to ν . By [Corollary 1.4](#) there is an $X(x) \in \text{GL}(\mathbb{R}^d)$ with $\nu(x) = X(x)\nu$.

Consider the set $\text{End}(T) := \{X \in \text{GL}(\mathbb{R}^d) \mid XT_\sigma = T_\sigma X\}$ of all transformations that commute with T , also known as the *endomorphism ring* of T (see [Definition A.7](#) (ii)). We show that $x \mapsto X(x)$ is a linear isomorphism between \mathbb{R}^d and $\text{End}(T)$.

Since Σ is regular, for every $i \in V$ there is a unique $\sigma \in \Sigma$ with $\sigma(1) = i$. Then

$$\begin{aligned} X(\alpha a + \beta b)v_i &= v_i(\alpha a + \beta b) = T_{\sigma_i}v_1(\alpha a + \beta b) \\ &= T_{\sigma_i}(\alpha a + \beta b) \\ &= \alpha T_{\sigma_i}a + \beta T_{\sigma_i}b \\ &= \alpha T_{\sigma_i}v_1(a) + \beta T_{\sigma_i}v_1(b) \\ &= \alpha v_i(a) + \beta v_i(b) = \alpha X(a)v_i + \beta X(b)v_i. \end{aligned}$$

for all $i \in V$. Since the v_i contain a basis, this shows $X(\alpha a + \beta b) = X(\alpha a) + X(\beta b)$.

Using the same argument, we show that $X(x)$ commutes with T_σ :

$$X(x)T_\sigma v_i = X(x)v_{\sigma(i)} = v_{\sigma(i)}(x) = T_\sigma v_i(x) = T_\sigma X(x)v_i.$$

for all $i \in V$. This proves $X(x) \in \text{End}(T)$.

It remains to show that there exists an inverse map. This map is $\text{End}(T) \ni X \mapsto x := Xv_1$. We show that if $x := Xv_1$ then $X(x) = X$:

$$X(x)v_i = X(x)T_{\sigma_i}v_1 = T_{\sigma_i}X(x)v_1 = T_{\sigma_i}x = T_{\sigma_i}Xv_1 = XT_{\sigma_i}v_1 = Xv_i.$$

for all $i \in V$.

By Schur's Lemma (see [Theorem A.8](#)) the endomorphism ring $\text{End}(T)$ of an irreducible representation is a division algebra over \mathbb{R} , and thus isomorphic to either \mathbb{R} , \mathbb{C} or \mathbb{H} (the quaternions). On the other hand, we just established that $x \mapsto X(x)$ is a linear isomorphism between \mathbb{R}^d and $\text{End}(T)$, and thus $d \in \{1, 2, 4\}$. But $d \neq 1$ as 1-dimensional arrangements are not *exceptionally* rigid (just ‘‘casually’’ rigid, see [Corollary 1.28](#)). \square

A closer inspection of the representation T in the proof of [Theorem 1.30](#) would yield more information about the exceptionally rigid arrangements. Since this is just a digression, we settle with a short description:

- if $d = 2$, then Σ is a cyclic group and the associated arrangement forms the vertices of a regular polygon (as in [Example 1.29](#)).
- If $d = 4$, then Σ is isomorphic to a finite subgroup of the multiplicative group of unit quaternions $S^1(\mathbb{H})$ (the *binary polyhedral groups*). The corresponding arrangements are the vertices of certain highly-symmetric, but not necessarily regular, 4-polytopes.

The classification of exceptionally rigid arrangements shows many parallels to the classification of symmetry groups of vertex-transitive polytopes (as mentioned in [Section 1.3](#)): the only irreducible matrix groups that do not appear as such Euclidean symmetry groups are in dimension two and four, and are related to the division algebras \mathbb{C} and \mathbb{H} .

Summary

In this chapter we studied point arrangements, the simplest geometric objects for which we can define symmetry and a notion of rigidity. We made an effort to characterize properties of arrangements in terms of their arrangement spaces, building the “arrangement space dictionary”. This will be of great value when we discuss spectral properties in the next chapter.

Let us summarize the various properties characterized by the arrangement space. Let $\nu, \nu' : V \rightarrow \mathbb{R}^d$ be two arrangements with arrangement spaces $U, U' \subseteq \mathbb{R}^n$. Then

$$\begin{aligned} \text{span } \nu \text{ is } \delta\text{-dimensional} &\Leftrightarrow U \text{ is } \delta\text{-dimensional} \\ \nu = X\nu' \text{ for some } X \in \text{GL}(\mathbb{R}^d) &\Leftrightarrow U = U'. \end{aligned}$$

If ν and ν' are spherical, then

$$\begin{aligned} \nu \text{ is a } \Sigma\text{-arrangement} &\Leftrightarrow U \text{ is } \Sigma\text{-invariant} \\ \nu \text{ is an irreducible } \Sigma\text{-arrangement} &\Leftrightarrow U \text{ is } \Sigma\text{-irreducible}. \end{aligned}$$

If ν and ν' are normalized, then

$$\begin{aligned} \nu = X\nu' \text{ for some } X \in \text{O}(\mathbb{R}^d) &\Leftrightarrow U = U' \\ \nu \text{ has radius } r(\nu) &\Leftrightarrow [r(\nu)]^2 = \dim(U)/n. \end{aligned}$$


We furthermore used arrangement spaces to develop a theory on how rigidity of arrangements relates to the number and relative placement of irreducible subspaces. This relation states roughly

$$\begin{aligned} \text{there are flexible arrangements} &\Leftrightarrow \text{there are non-orthogonal irreducible subspaces} \\ &\Leftrightarrow \text{there are infinitely many irreducible subspaces.} \end{aligned}$$

There are many more such connections that are not relevant to this work. For example, arrangement spaces can be used to give a one-line definition of *Gale duality* [79, Section 6.4]: ν and ν' are Gale duals of each other if and only if their arrangement spaces are orthogonal complements of each other. Any property shared between complement subspaces then translates to a property shared between Gale duals, *e.g.* being a Σ -arrangement.

In general, we consider the arrangement space as a very convenient tool through which to study point arrangements and their linear properties. And while the idea is certainly not new, there seems to be no established name for this concept in the literature.

2 Graph Realizations



Graph realizations appear most naturally as the skeleta of convex polytopes, but can be more generally understood as embeddings of graphs into Euclidean spaces, most of which are not in the “convex” configuration of a skeleton. It is the purpose of this chapter to see in how far our questions about spectral polytopes can be already addressed and answered on the level of general graph realizations.




Figure 2.1. A skeleton of a polytope (left), and a realization that is not the skeleton of any polytope (right). We use shading to help the visual understanding. This is not meant to indicate that realizations are “solid” in any way.

Particular focus will be on [Question 6](#) (the realization version of [Question 3](#)), which asks whether every “sufficiently symmetric” graph realization is spectral. The quest for this chapter is then to develop an understanding of the interplay between spectrum and symmetry, and by this, narrow down the notion of “sufficient symmetry”. We shall see that (in contrast to the polytopes case) combined vertex- and edge-transitivity is not quite sufficient, and we instead need so-called *distance-transitivity*.

A classical reference for the spectral properties of highly-symmetric graphs is Lovász [48]. Berkolaiko and Liu [3] studied related questions in the context of quantum graphs and graphs with edge-weights. Most notably, in [21, Section 4] Du and Fan remarked a special property of the Petersen graph, which turned out to match our notion of “sufficient symmetry”. This result will be reproduced and generalized in our study of distance-transitivity.

For this chapter let $G = (V, E)$ denote a finite simple graph with vertex set $V = \{1, \dots, n\}$ and edge set $E \subseteq \{\{i, j\} \mid i, j \in V \text{ and } i \neq j\}$. We write $ij \in E$ instead of $\{i, j\} \in E$.

Definition 2.1. A d -dimensional (graph) realization of G is a map $v: V \rightarrow \mathbb{R}^d$.

We visualize graph realizations as embeddings of graphs into Euclidean spaces, with vertices embedded as points and edges embedded as straight line segments. These points and lines are referred to as vertices and edges of the realization.

A graph realization can be considered as a point arrangement whose points are indexed by the vertex set of a graph. We can then reuse the terminology of arrangements and speak of the arrangement matrix and the arrangement space of a realization, of full-dimensional, normalized, spherical, symmetric and irreducible/reducible realizations, as well as Σ -realizations. For the latter it is now the default to assume $\Sigma \subseteq \text{Aut}(G)$, where

$$\text{Aut}(G) := \{\sigma \in \text{Sym}(V) \mid ij \in E \Leftrightarrow \sigma(i)\sigma(j) \in E\}$$

is the (*combinatorial*) *automorphism group* (or *symmetry group*) of the graph G . Moreover

$$\mathcal{A}_d(G, \Sigma) := \{\nu: V \rightarrow \mathbb{R}^d \mid \nu \text{ is a full-dimensional } \Sigma\text{-realization of } G\}$$

is defined in accordance with $\mathcal{A}_d(\Sigma)$ from [Section 1.4](#). This allows us to talk about deformations between Σ -realizations, as well as rigid and flexible realizations.

Chapter overview

We define *spectral graph realizations* in [Section 2.1](#) and show that they are maximally symmetric *w.r.t.* the underlying graph. The definition is in terms of arrangement spaces, which permits us to reformulate the vague notion of “sufficient symmetry” in terms of subspaces of \mathbb{R}^n . It will become clear from this formulation that a “sufficient symmetry” must be rigid (in the sense of [Section 1.4](#)). We introduce *balanced realizations* as a precursor to spectral realizations, and show that every irreducible rigid realization must be balanced.

From [Section 2.2](#) on we focus on graph realizations with many symmetries, in particular, vertex-, edge- and arc-transitive realizations. We show that many metric properties of such realizations, if they are balanced, are already determined by the corresponding eigenvalue. But we shall also see that vertex-, edge- and arc-transitive realizations need not be spectral, or even balanced.

[Section 2.3](#) introduces *full local dimension*, a natural geometric property which in particular is possessed by polytope skeleta. We show that arc-transitive realizations of full local dimension have desirable properties, such as being irreducible, rigid and balanced. They still need not be spectral. We describe a procedure for checking whether a given arc-transitive polytope is spectral, by just comparing metric properties with eigenvalues. We prove that all regular polytopes are θ_2 -spectral (extending a result of Licata and Powers [[47](#)]).

In [Section 2.4](#) we consider a class of realizations of a particularly high degree of symmetry, called *distance-transitivity*. We show that the distance-transitive realizations are “sufficiently symmetric” in the sense of [Question 6](#).

Some results for this symmetry class can be generalized to realizations that are not distance-transitive. We do so in [Section 2.5](#) via the study of *cosine vectors* and *cosine sequences*.

2.1 Spectral and balanced realizations

The idea of assigning geometric information to a graph, constructed from spectral properties of associated matrices, has been around for a long time. To name only a few, applications have been found in data visualization (e.g. graph drawings [45]), semi-definite optimization (e.g. eigenvalue optimization [10, 33]) and geometric combinatorics (e.g. for equiangular lines [46] and balanced point arrangements [16]). Spectral realizations are further related to the Lovász theta function [49] and the Colin de Verdière graph invariant [71].

Recall that the *spectrum*¹ $\text{Spec}(G) = \{\theta_1, \dots, \theta_m\}$ of a graph G denotes the set of eigenvalues of its adjacency matrix $A \in \{0, 1\}^{n \times n}$. It is common to label the distinct eigenvalues $\theta_1 > \theta_2 > \dots > \theta_m$ of A in decreasing order. More generally, by eigenvalues, eigenvectors, eigenspaces, etc. of G , we refer to the respective quantities of its adjacency matrix.

Definition 2.2. A realization of a graph G is called θ -spectral (or just *spectral*) if its arrangement space is the θ -eigenspace $\text{Eig}_G(\theta)$ of G .

This definition fits well with our intention of building the “arrangement space dictionary”, but it is not the standard definition given in the literature. Instead, the literature definition of *spectral realization* often refers to the following concrete construction:

Construction 2.3. Let $\theta \in \text{Spec}(G)$ be an eigenvalue of G of multiplicity d .

We choose an *orthonormal basis* $u_1, \dots, u_d \in \text{Eig}_G(\theta)$ of the corresponding eigenspace and construct the matrix

$$\Phi := \begin{pmatrix} | & & | \\ u_1 & \cdots & u_d \\ | & & | \end{pmatrix} \in \mathbb{R}^{n \times d}$$

in which the u_i are the columns. This matrix has exactly n rows, and we can define a d -dimensional graph realization $v: V \rightarrow \mathbb{R}^d$ for which v_i is the i -th row of Φ .

Φ is the arrangement matrix of v and $U := \text{span } \Phi = \text{Eig}_G(\theta)$ is the corresponding arrangement space. In particular, v is spectral in the sense of [Definition 2.2](#). We shall call this specific spectral realization the θ -realization of G .

Note that $\Phi^\top \Phi = \text{Id}$, and so v is full-dimensional and normalized. In particular, it is unique up to orientation by [Lemma 1.10](#) and it is justified to call it *the* θ -realization of G .


Example 2.4. Let G be the edge-graph of the *regular dodecahedron*. Its spectrum is

$$\text{Spec}(G) = \{3^1, \sqrt{5}^3, 1^5, 0^4, (-2)^4, (-\sqrt{5})^3\},$$

where the exponents denote multiplicities.

Note in particular the two eigenvalues $\pm\sqrt{5}$ of multiplicity three. Using [Construction 2.3](#), those give rise to two 3-dimensional realizations.

¹See [Appendix B](#) for a brief reminder on the most important aspects of spectral graph theory.



Remarkably, the realization to the *second-largest* eigenvalue $\theta_2 = \sqrt{5}$ is precisely the skeleton of the regular dodecahedron.

This is another instance of the “curious observation” from the introduction: the convex hull of the $\sqrt{5}$ -realization is the polyhedron from which we have started. In the language of the upcoming [Chapter 3](#), the regular dodecahedron is a *spectral polytope*.

We also mentioned that the same phenomenon can be observed for *all* regular polytopes (see [Appendix E](#) for a reminder on this polytope class). This was proven by Licata and Powers in [47] for most of them. In [Example 2.29](#) below we describe a procedure by which to verify the remaining cases by a simple comparison of tabulated numerical values.

To explore further examples, in [Appendix F](#) the reader can find a short Mathematica script to compute the spectral realizations of any given graph.

[Example 2.4](#) suggests that the spectral realizations of highly symmetric graphs, such as the edge-graphs of regular polytopes, are themselves highly symmetric. In fact, this is well-known: a θ -realization is always as symmetric as the underlying graph. This is essentially a consequence of the fact that eigenspaces of G are $\text{Aut}(G)$ -invariant:

Theorem 2.5. *If v is spherical and spectral (e.g. if v is a θ -realization), then v realizes all the symmetries of G , that is, v is an $\text{Aut}(G)$ -realization.*

Proof. Recall that the symmetries of G can be defined in terms of its adjacency matrix A : for a permutation $\sigma \in \text{Sym}(V)$ and the corresponding permutation matrix $\Pi_\sigma \in \mathbb{R}^{n \times n}$ holds


$$\sigma \in \text{Aut}(G) \iff A\Pi_\sigma = \Pi_\sigma A. \tag{2.1}$$

Suppose now that v is spherical and θ -spectral, hence has arrangement space $U = \text{Eig}_G(\theta)$. If $\sigma \in \text{Aut}(G)$, then by (2.1) Π_σ commutes with the adjacency matrix A , and Π_σ therefore preserves the eigenspaces of A . In other words, $U = \text{Eig}_G(\theta)$ is $\text{Aut}(G)$ -invariant. Since v is spherical, [Lemma 1.12 \(ii\)](#) implies that it is an $\text{Aut}(G)$ -realization. \square

According to [Theorem 2.5](#), [Construction 2.3](#) provide a fast and robust procedure to construct symmetric realizations in practice. The question remains, *which* symmetric realizations can be obtained by this method.

We give two examples of symmetric realizations that are unobtainable in this way.

Example 2.6. Consider the *truncated tetrahedron*: the polyhedron obtained from the regular tetrahedron by cutting off its vertices. There is a freedom in this construction, from which we obtain an infinite family of distinct realizations:




Each realization has the full symmetry of the regular tetrahedron, and their skeleta are $\text{Aut}(G)$ -realizations of the edge-graph G . The spectrum of G is

$$\text{Spec}(G) = \{3^1, 2^3, 0^2, (-1)^3, (-2)^3\},$$

and we see that there are only three 3-dimensional spectral realizations. But the image shows already *four* of infinitely many distinct stages of an $\text{Aut}(G)$ -deformation. Most of the skeleta are therefore not spectral, even though they are $\text{Aut}(G)$ -realizations.

Example 2.6 provides a first idea for why rigidity is an essential ingredient for “sufficient symmetry”. We come back to this in **Observation 2.9**.

Example 2.7. Consider $G = (V, E)$, the edge-graph of the *hexagonal prism* with $V = \{1, \dots, 12\}$, depicted on the left, and a particular 2-dimensional realization ν of G on the right (note that several edges and vertices are mapped on top of each other):



One can check that ν is an $\text{Aut}(G)$ -realization. The spectrum of G ,

$$\text{Spec}(G) = \{3^1, 2^2, 1^1, 0^4, (-1)^1, (-2)^2, (-3)^1\},$$

has two eigenspaces of dimension two. But ν is a spectral realization to neither: the neighbors of $1 \in V$ (highlighted in the figure) have their “barycenter” at the origin. We shall see in a moment (in **Proposition 2.11**) that then ν can be at most 0-spectral. But the 0-eigenspace has dimension four, and so this can neither be the case.

Note that ν is rigid, and so the problem is not the lack of rigidity for this example.

The two preceding examples point to two distinct obstacles in identifying “sufficient symmetries”. We shall see that both problems are best described in terms of arrangement spaces.

Via the arrangement space, spectral realizations correspond to eigenspaces, and symmetric realizations correspond to invariant subspaces. **Theorem 2.5** was then proven by using that eigenspaces are $\text{Aut}(G)$ -invariant, and **Question 6** can be interpreted as asking for a converse: *when are Σ -invariant subspaces eigenspaces?*

We should not expect this to happen for most *reducible* subspaces. For example, if $\theta, \theta' \in \text{Spec}(G)$ are distinct eigenvalues, and since eigenspaces are Σ -invariant for all $\Sigma \subseteq \text{Aut}(G)$,

$$U := \text{Eig}_G(\theta) \oplus \text{Eig}_G(\theta')$$

is a reducible Σ -invariant subspace that is not an eigenspace.

It is then reasonable to ask [Question 6](#) only for *irreducible* realizations. Let us use this to finally give a formal definition of what we mean by “sufficient symmetry”:

Definition 2.8. A permutation group $\Sigma \subseteq \text{Aut}(G)$ is a *sufficient symmetry*, and irreducible Σ -realizations are called *sufficiently symmetric*, if any (and then both) of the following statements is satisfied:

- (i) all irreducible Σ -realizations are spectral.
- (ii) the Σ -irreducible subspaces of \mathbb{R}^n are exactly the eigenspaces of G .

A realization is spectral if its arrangement space $U \subseteq \mathbb{R}^n$ is an eigenspace. A Σ -realization can then fail to be spectral in essentially two ways:

- U might be a *proper* subspace of an eigenspace. This can happen if an eigenspace is not Σ -irreducible.

This happened in [Example 2.7](#): the 4-dimensional 0-eigenspace of the edge-graph has a 2-dimensional $\text{Aut}(G)$ -irreducible subspace. The given realization has this subspace as its arrangement space.

- U might not be contained in any eigenspace at all.

This happened in [Example 2.6](#): the edge-graph has only three eigenspaces of dimension three, and no larger eigenspaces. So there are only three 3-dimensional subspaces of \mathbb{R}^n that fit in an eigenspace. But there are more than three $\text{Aut}(G)$ -realizations.

It is exactly the second problem that can be avoided by rigidity:

Observation 2.9. In [Definition 2.8](#) we formalized “sufficient symmetry” as the Σ -irreducible subspaces being exactly the eigenspaces. But there are only finitely many eigenspaces. By [Theorem 1.24 \(iii\)](#) \implies (i) “sufficient symmetry” must then come with rigidity.

We can therefore focus on *rigid* permutation groups $\Sigma \subseteq \text{Aut}(G)$.

Of course, rigidity cannot be a sufficient condition for “sufficient symmetry” as we are still left with the first problem (cf. [Example 2.7](#)). We give a name to the problem:

Definition 2.10. A realization with arrangement space $U \subseteq \mathbb{R}^n$ is called θ -*balanced* (or just *balanced*) if $U \subseteq \text{Eig}_G(\theta)$ for some $\theta \in \text{Spec}(G)$.


Spectral realizations are then a special type of balanced realization. The name “balanced” is motivated from the following geometric interpretation:

Proposition 2.11. A realization is θ -balanced if and only if

$$\sum_{j \in N(i)} v_j = \theta v_i, \quad \text{for all } i \in V, \tag{2.2}$$

where $N(i) := \{j \in V \mid ij \in E\}$ denotes the neighborhood of $i \in V$ in G .

Proposition 2.11 states that in a balanced realization every vertex is (up to some factor) in the barycenter of its neighbors (so to say, *balanced* “between” its neighbors). Note that this justifies our conclusion in **Example 2.7**: the realization can be at most 0-spectral because the barycenter of the neighbors is the origin, thus $\theta = 0$.



Proof of Proposition 2.11. Note that (2.2) is equivalent to $A\Phi = \theta\Phi$, where A is the adjacency matrix of G and Φ is the arrangement matrix of v . But this is equivalent to $\theta \in \text{Spec}(G)$ and the columns of Φ being θ -eigenvectors, or $U := \text{span } \Phi \subseteq \text{Eig}_G(\theta)$. \square

It will become apparent that we consider being balanced as a worthwhile middle ground. It is comparatively easy to conclude that realizations of a particular symmetry are balanced (as we shall do in **Theorem 2.27**), and like for spectral realizations, there are strong connections between θ and the metric properties of a θ -balanced realization (see **Lemma 2.20** below). In a sense, spectral realizations are balanced realizations of a maximal dimension, and proving this last step of having the “right dimension” is often the hardest part when working from just symmetry.

Theorem 2.12. *If $v \in \mathcal{A}_d(G, \Sigma)$ is irreducible and rigid, then it is balanced.*

Proof. Let $U \subseteq \mathbb{R}^n$ be the arrangement space of v . Then U is Σ -irreducible by **Lemma 1.13**. Consider the decomposition

$$\mathbb{R}^n = \text{Eig}_G(\theta_1) \oplus \cdots \oplus \text{Eig}_G(\theta_m)$$

of \mathbb{R}^n into the pair-wise orthogonal eigenspaces of G . Since the eigenspaces span all of \mathbb{R}^n , U must be non-orthogonal to some eigenspace, say $\text{Eig}_G(\theta)$. The orthogonal projection $U' \subseteq \text{Eig}_G(\theta)$ of U onto this eigenspace is again Σ -irreducible and non-orthogonal to U (see **Proposition A.6**). Since v is rigid, **Theorem 1.23** $\neg(i) \implies \neg(ii)$ yields $U = U'$. Hence $U \subseteq \text{Eig}_G(\theta)$ and v is θ -balanced. \square

The converse of **Theorem 2.12** is not true: flexible realizations can be balanced. Or even stronger, there are proper deformations $v(\cdot)$ for which *all* $v(t)$, $t \in [0, 1]$ are balanced (see **Example 2.24** in the next section).

We close this section with a theorem about what eigenvalue multiplicities can tell us about the geometry and rigidity of symmetric realizations.

Theorem 2.13. *Let $\mu_1 \geq \mu_2 \geq \cdots$ be the multiplicities of the eigenvalues of G sorted in non-increasing order (note that μ_i not necessarily belongs to θ_i). Let $v \in \mathcal{A}_d(G, \Sigma)$ be irreducible:*

- (i) if $\mu_2 < d$, then ν is θ -balanced and $\theta \in \text{Spec}(G)$ has multiplicity μ_1 .
- (ii) if $\mu_2 < d$ and $\mu_1 < 2d$, then ν is rigid.

Proof. Let $U \subseteq \mathbb{R}^n$ be the arrangement space of ν . For each eigenvalue $\theta \in \text{Spec}(G)$, let $U_\theta \subseteq \text{Eig}_G(\theta)$ denote the orthogonal projection of U onto the θ -eigenspace of G . Since U is Σ -irreducible by [Lemma 1.13](#), the projections U_θ are Σ -irreducible as well (see [Proposition A.6](#)). If U is non-orthogonal to $\text{Eig}_G(\theta)$, then it is non-orthogonal to U_θ too, and non-orthogonal irreducible subspaces have the same dimension (see [Lemma A.9](#)):

$$\dim \text{Eig}_G(\theta) \geq \dim U_\theta = \dim U = d.$$

But by $\mu_2 < d$ there can be only a single eigenspace $\text{Eig}_G(\theta)$ with multiplicity $\mu_1 \geq d$, and U must be orthogonal to all other eigenspaces. Since the eigenspaces induce a decomposition of \mathbb{R}^n into pair-wise orthogonal subspaces, being orthogonal to all but one eigenspaces implies $U \subseteq \text{Eig}_G(\theta)$ for the remaining eigenspace, and ν must be θ -balanced. This shows (i).

To show (ii), suppose that ν is flexible. By [Theorem 1.23](#) (i) \implies (ii) there is an irreducible realization $\nu' \in \mathcal{A}_d(G, \Sigma)$ with Σ -irreducible arrangement space $U' \neq U$. By (i), ν and ν' must both be θ -balanced, that is, $U, U' \subseteq \text{Eig}_G(\theta)$. Since U and U' are distinct Σ -irreducible subspaces, they have trivial intersection (see [Observation A.3](#)), and

$$\mu_1 = \dim \text{Eig}_G(\theta) \geq \dim(U + U') = \dim(U) + \dim(U') - \underbrace{\dim(U \cap U')}_{=0} = 2d.$$

Thus, if $\mu_1 < 2d$, then ν must have been rigid, proving (ii). □

2.2 Edge- and arc-transitive realizations

In our search for instances of “sufficient symmetry” we turn to the investigation of first concrete symmetry classes. Recall the definition of vertex-transitivity (cf. [Section 1.5](#)):

Definition 2.14.

- (i) A group $\Sigma \subseteq \text{Aut}(G)$ is said to act *vertex-transitively* on G , if for any two vertices $i, j \in V$ there is a symmetry $\sigma \in \Sigma$ with $\sigma(i) = j$. We also say that Σ is *vertex-transitive*.
- (ii) G is said to be *vertex-transitive* if $\text{Aut}(G)$ is vertex-transitive.
- (iii) A realization of G is *vertex-transitive* if it is a Σ -realization for some vertex-transitive group $\Sigma \subseteq \text{Aut}(G)$.

[Observation 2.9](#) shows that a “sufficient symmetry” must be rigid, and [Lemma 1.26](#) shows that an (irreducible) rigid realization is “essentially vertex-transitive”. It is therefore reasonable for us to focus on vertex-transitive realizations.

However, we have also seen (e.g. in [Examples 2.6](#) and [2.7](#)) that vertex-transitivity alone is not sufficient and should be supplemented by further symmetries. In this section we consider edge- and arc-transitivity. Recall that an *arc* in a graph is an incident vertex-edge pair.

Definition 2.15. Let $\Sigma \subseteq \text{Aut}(G)$ be a group of symmetries of G .

- (i) Σ is said to act *edge-transitively* on G , if for any two edges $ij, \hat{i}\hat{j} \in E$ there is a symmetry $\sigma \in \Sigma$ with $\{\sigma(i), \sigma(j)\} = \{\hat{i}, \hat{j}\}$.
- (ii) Σ is said to act *arc-transitively* on G , if for any two edges $ij, \hat{i}\hat{j} \in E$ there is a symmetry $\sigma \in \Sigma$ with $\sigma(i) = \hat{i}$ and $\sigma(j) = \hat{j}$.

Edge- and arc-transitive graphs and realizations are then defined parallel to [Definition 2.14](#) (ii) and (iii).

Remark 2.16. Arc-transitive graphs are both vertex- and edge-transitive, but the converse is not true. Graphs that are both vertex- and edge-transitive, but not arc-transitive, are called *half-transitive* and are comparatively rare (the smallest such graph is known as the *Holt graph* and has 27 vertices [[37](#)]).

We shall ignore this intermediate class for now, and often default to arc-transitive graphs. The convenience of arc-transitive groups $\Sigma \subseteq \text{Aut}(G)$ comes from the observation that even when fixing a vertex $i \in V$, the group Σ (or better, its stabilizer Σ_i) still acts transitively on the neighborhood $N(i)$.

We shall come back to half-transitivity in the case of polytopes in [Section 4.5](#).

Remark 2.17. Combined vertex- and edge-transitivity might appear as an arbitrary choice of symmetry. We argue that it is a natural choice if one tries to arrive at an especially strong connection between the combinatorics of the graph and the geometry of its spectral realizations.

The reasoning is as follows: the results of spectral graph theory are quite sensitive to the matrix chosen to represent the graph. The choice of the adjacency matrix might seem natural, but is far from the only choice. Other matrices can be associated to the graph, such as the Laplace matrix (see [Appendix B](#)) or any other matrix $M \in \mathbb{R}^{n \times n}$ with $M_{ij} = 0$ for $ij \notin E$ (a precursor to so-called *discrete Schrödinger operators*).

For highly-symmetric graphs, a reasonable choice of a matrix representation should reflect the symmetries of the graph. That is, if G is vertex-transitive, then M_{ii} should be independent of $i \in V$. Likewise, if the graph is edge-transitive, then M_{ij} should be independent of $ij \in E$. Taken together, for combined vertex- and edge-transitivity, there are $\alpha, \beta \in \mathbb{R}$ with

$$M = \alpha A + \beta \text{Id}, \tag{2.3}$$

where A denotes the adjacency matrix. Spectral realizations are build from the eigenspaces of the graph, and in fact, all matrices of the form [\(2.3\)](#) have the same eigenspaces (*i.e.*, the same decomposition of \mathbb{R}^n into subspaces, the corresponding eigenvalues will differ).

In this sense, combined vertex- and edge-transitivity is a setting in which the choice of the adjacency matrix is canonical, and we can expect the spectral properties of the adjacency matrix to be especially expressive.

Observation 2.18. If ν is vertex-transitive, then $\|\nu_i\|$ is independent of $i \in V$. In particular, the radius (as defined in [\(1.2\)](#)) satisfies $r(\nu) = \|\nu_i\|$. If ν is normalized and full-dimensional then (also by [\(1.2\)](#))

$$\|\nu_i\|^2 = [r(\nu)]^2 = \frac{d}{n}. \tag{2.4}$$

Observation 2.19. If v is edge-transitive, then all edges have the same length, and their end vertices have the same inner product. The following notions are then well-defined:

$$\omega(v) := \langle v_i, v_j \rangle, \quad \ell(v) := \|v_i - v_j\|$$

whenever $ij \in E$. The latter is called the *edge length* of v .

In a vertex-transitive graphs all vertices have the same vertex-degree, which we abbreviate by $\deg(G)$. The quantities from [Observation 2.19](#) can then be computed explicitly:

Lemma 2.20. *Let v be full-dimensional, vertex-transitive, edge-transitive, θ -balanced and normalized. Let $\lambda := \deg(G) - \theta$ be the corresponding Laplacian eigenvalue (cf. [Appendix B](#)). Then*

$$\omega(v) = \frac{\theta d}{2|E|}, \quad [\ell(v)]^2 = \frac{\lambda d}{|E|}. \quad (2.5)$$

Proof. Using the radius equation [\(2.4\)](#) as well as the geometric interpretation [\(2.2\)](#) for balanced realizations, for each $ij \in E$ there holds

$$\deg(G) \cdot \omega(v) = \sum_{j \in N(i)} \langle v_i, v_j \rangle = \left\langle v_i, \sum_{j \in N(i)} v_j \right\rangle \stackrel{(2.2)}{=} \langle v_i, \theta v_i \rangle = \theta \cdot \|v_i\|^2 \stackrel{(2.4)}{=} \frac{\theta d}{n}.$$

Expression [\(2.5\)](#) for $\omega(v)$ then follows from the identity $\deg(G) \cdot n = 2|E|$.

The expression for $\ell(v)$ follows via

$$\begin{aligned} [\ell(v)]^2 &= \|v_i - v_j\|^2 \\ &= \|v_i\|^2 + \|v_j\|^2 - 2\langle v_i, v_j \rangle \\ &= [r(v)]^2 + [r(v)]^2 - 2\omega(v) \\ &= 2([r(v)]^2 - \omega(v)) \\ &= 2\left(\frac{d}{n} - \frac{d\theta}{2|E|}\right) = 2d\left(\frac{\deg(G)}{2|E|} - \frac{\theta}{2|E|}\right) = \frac{2d(\deg(G) - \theta)}{2|E|} = \frac{\lambda d}{|E|}. \end{aligned}$$

□

If the realization is just spherical instead of normalized, we can still compute some relative quantities (the computations are straightforward):

Corollary 2.21. *Let v be vertex-transitive, edge-transitive, θ -balanced, and spherical. Then*

$$\frac{\omega(v)}{[r(v)]^2} = \frac{\theta}{\deg(G)}, \quad \left[\frac{\ell(v)}{r(v)}\right]^2 = \frac{2\lambda}{\deg(G)} = 2\left(1 - \frac{\theta}{\deg(G)}\right). \quad (2.6)$$

Note that we dropped the condition that v is full-dimensional, which was required for substituting $\dim U$ by d in [\(1.2\)](#) in the proof of [Lemma 2.20](#). But [\(2.6\)](#) is independent of d , and full dimension is no longer required.

The quantity ω/r^2 is called *cosine* of v since for $ij \in E$ holds

$$\cos \sphericalangle(v_i, v_j) = \frac{\langle v_i, v_j \rangle}{\|v_i\| \|v_j\|} = \frac{\omega(v)}{[r(v)]^2} \stackrel{(2.6)}{=} \frac{\theta}{\deg(G)}. \quad (2.7)$$

Example 2.22. In [Example 2.4](#) we have seen that the skeleton of the regular dodecahedron can be obtained as the θ_2 -realization of its edge graph (where $\theta_2 = \sqrt{5}$). We can then use [Corollary 2.21](#) to compute the circumradius of the dodecahedron with edge length $\ell = 1$:


$$r(v) \stackrel{(2.6)}{=} \sqrt{\frac{\deg(G)}{2\lambda_2}} = \sqrt{\frac{\deg(G)}{2(\deg(G) - \theta_2)}} = \sqrt{\frac{3}{2(3 - \sqrt{5})}} \approx 1.401258.$$

Obtaining this value by elementary geometric techniques is rather laborious.

Example 2.23. The *dihedral angle* of a d -dimensional polytope $P \subset \mathbb{R}^d$ at a $(d-2)$ -dimensional face is the angle between the two facets that meet at that face. In the 3-dimensional case, dihedral angles are measured at edges.

For the *regular icosahedron* (the dual of the regular dodecahedron) this angle α is the same for all edges, and is exactly π minus the angle between adjacent vertices in the dodecahedron. By [\(2.7\)](#) is yields

$$\alpha = \pi - \arccos\left(\frac{\theta_2}{\deg(G)}\right) = \pi - \arccos\left(\frac{\sqrt{5}}{3}\right) = \arccos\left(-\frac{\sqrt{5}}{3}\right) \approx 138.1896^\circ.$$



The computations in [Example 2.22](#) and [Example 2.23](#) work equivalently for all the other regular polytopes. Note that this not yet requires that the regular polytopes are spectral, but just that their skeleta are balanced (which we prove in the next section, see [Theorem 2.27](#)).

We close this section with an example of an arc-transitive realization that is neither spectral nor rigid.

Example 2.24. Consider the complete bipartite graph $K_{n,n}$ on $2n$ vertices. This graph is arc-transitive and has spectrum

$$\text{Spec}(K_{n,n}) = \{n^1, 0^{2(n-1)}, (-n)^1\}.$$

One can check that the corresponding eigenspaces are irreducible *w.r.t.* $\text{Aut}(K_{n,n})$, and thus all $\text{Aut}(K_{n,n})$ -realizations are spectral (by the equivalent statements in [Definition 2.8](#)). In order to construct arc-transitive realizations that are *not* spectral and *not* rigid, we therefore cannot use the full symmetry group, but have to use some proper subgroup $\Sigma \subset \text{Aut}(K_{n,n})$ that is still arc-transitive, but with respect to which some eigenspace (apparently the 0-eigenspace) becomes reducible.

To obtain such a group, start with a set $V := \{1, \dots, n\}$ and choose a proper $\bar{\Sigma}$ -deformation (of arrangements) $\bar{v}(\cdot): [0, 1] \rightarrow \mathcal{A}_d(\bar{\Sigma})$ for some transitive group $\bar{\Sigma} \subseteq \text{Sym}(V)$ and some $d < n - 1$. The $\bar{v}(t)$ are then flexible, and we have $d \geq 2$ by [Corollary 1.28](#).

Consider $K_{n,n}$ with vertex set $V_1 \cup V_2$, where the partition classes V_i are disjoint copies of V . Let $\Sigma_i \subseteq \text{Aut}(K_{n,n})$ be a copy of $\bar{\Sigma}$ acting on V_i instead of V . Let further $\tau \in \text{Aut}(K_{n,n})$ be the involution that exchanges $V_1, V_2 \subseteq V(K_{n,n})$ in the obvious way. One can check that the group $\Sigma := \langle \Sigma_1, \Sigma_2, \tau \rangle \subseteq \text{Aut}(K_{n,n})$ acts arc-transitively on $K_{n,n}$.

It remains to show that the 0-eigenspace is Σ -reducible. We do so by constructing a flexible Σ -realization of $K_{n,n}$: consider the following Σ -deformation $v(\cdot): [0, 1] \rightarrow \mathcal{A}_{2d}(K_{n,n}, \Sigma)$:

$$v_i(t) := \begin{cases} (\bar{v}_i(t), 0) & \text{if } i \in V_1 \\ (0, \bar{v}_i(t)) & \text{if } i \in V_2 \end{cases}.$$

Indeed, each $v(t)$ is an irreducible Σ -realization. Since $\bar{v}(\cdot)$ is proper, so is $v(\cdot)$, and all $v(t)$ are flexible. Also, the $v(t)$ are of dimension $2d \geq 4$. Since only the 0-eigenspace of $K_{n,n}$ has a dimension of at least four, it follows from [Theorem 2.13](#) (i) that the $v(t)$ must be 0-balanced. Since $d < n - 1$, the $v(t)$ are of dimension $2d < 2n - 2$ and none of them can be 0-spectral.

Note the following features of [Example 2.24](#):

- The constructed realizations $v(t)$ are flexible, but they are also all balanced to the same eigenvalue (i.e., the converse of [Theorem 2.12](#) is false).

It is unclear whether irreducible arc-transitive realizations are always balanced.

- The Σ -realization in [Example 2.24](#) uses a proper subgroup $\Sigma \subset \text{Aut}(G)$.

It is unclear whether an arc-transitive $\text{Aut}(G)$ -realization must be rigid. We know that they are not always spectral (see [Example 2.30](#) below).

Similar questions can be asked for graphs and groups of combined vertex- and edge-transitivity or half-transitivity.

2.3 Full local dimension

In this section we consider an additional geometric constraint which is naturally satisfied by the skeleta of (full-dimensional) polytopes.

Definition 2.25. A realization v is of *full local dimension* if

$$\text{span}\{v_j - v_i \mid j \in N(i)\} = \mathbb{R}^d, \quad \text{for all } i \in V.$$

That is, the edge-directions at every vertex contain a basis of \mathbb{R}^d .

Full local dimension implies full dimension. The converse is not true, not even for irreducible arc-transitive realizations:

Example 2.26. As seen in [Example 2.4](#), the edge-graph of the dodecahedron has eigenspaces of dimension four and five. We shall see in [Section 2.4](#) that these eigenspaces are irreducible *w.r.t.* the symmetry group of the edge-graph.

Those eigenspaces give rise to 4- and 5-dimensional irreducible arc-transitive realizations of full dimension (via [Construction 2.3](#)). However, since the edge-graph has degree three, those realizations cannot be of full *local* dimension.

General realizations of full local dimension are not necessarily rigid, balanced or irreducible (see [Examples 2.6](#) and [2.7](#)). This changes when we require arc-transitivity:

Theorem 2.27. *Let $v \in \mathcal{A}_d(G, \Sigma)$ be arc-transitive and of full local dimension. Then*

- (i) v is Σ -rigid.
- (ii) v is irreducible.
- (iii) v is balanced.

Proof. Let $T : \Sigma \rightarrow \text{O}(\mathbb{R}^d)$ denote the representation of v . Let further $\Sigma_i \subseteq \Sigma$ be the stabilizer of Σ at $i \in V$, and consider the restriction $T^i : \Sigma_i \rightarrow \text{O}(\mathbb{R}^d)$ of T onto Σ_i . Clearly, $\text{span}\{v_i\}$ is T^i -invariant and is acted on by identity. That is, $\text{span}\{v_i\} \subseteq \text{Fix}(T, \Sigma_i)$ (where $\text{Fix}(T, \Sigma_i)$ is the fix space as defined in [Construction 1.15](#)). We prove $\text{span}\{v_i\} = \text{Fix}(T, \Sigma_i)$.

By arc-transitivity, T^i acts transitively on the set $\mathcal{N}_i := \{v_j \mid j \in N(i)\}$. That is, for any two $w_1, w_2 \in \mathcal{N}_i$ there exists a $\sigma \in \Sigma_i$ with $T_\sigma w_1 = w_2$. For any $x \in \text{Fix}(T, \Sigma_i)$ (and using that T_σ is orthogonal) we then find

$$\langle x, w_1 \rangle = \langle T_\sigma x, T_\sigma w_1 \rangle = \langle x, w_2 \rangle \implies \langle x, w_1 - w_2 \rangle = 0.$$

This is independent of $x \in \text{Fix}(T, \Sigma_i)$ and the pair $w_1, w_2 \in \mathcal{N}_i$, and so $\text{Fix}(T, \Sigma_i) \subseteq \text{span}\{w_1 - w_2 \mid w_1, w_2 \in \mathcal{N}_i\}^\perp = \text{aff}(\mathcal{N}_i)^\perp$. We therefore established

$$\text{span}\{v_i\} \subseteq \text{Fix}(T, \Sigma_i) \subseteq \text{aff}(\mathcal{N}_i)^\perp. \quad (2.8)$$

From full local dimension follows $\dim \text{aff}(\mathcal{N}_i) \geq d - 1 \implies \dim \text{aff}(\mathcal{N}_i)^\perp \leq 1$. By full dimension and vertex-transitivity, we further have $\dim \text{span}\{v_i\} = 1$. The dimensions of the spaces in [\(2.8\)](#) must then agree, and moreover $\text{span}\{v_i\} = \text{Fix}(T, \Sigma_i) = \text{aff}(\mathcal{N}_i)^\perp$. In particular, we have $\dim \text{Fix}(T, \Sigma_i) = 1$ for all $i \in V$, and v is rigid by [Lemma 1.27](#). This proves (i).

We further show that all T^i -irreducible subspaces, besides $\text{span}\{v_i\}$, are contained in the orthogonal complement v_i^\perp : suppose that $W \subseteq \mathbb{R}^d$ is T^i -irreducible and not contained in the orthogonal v_i^\perp . In other words, W is non-orthogonal to $\text{span}\{v_i\}$. If two T^i -irreducible subspaces are non-orthogonal, then T^i acts isomorphically on them (see [Lemma A.9](#)). Since T^i acts trivially on $\text{span}\{v_i\}$, it must then also act trivially on W , which means $W \subseteq \text{Fix}(T, \Sigma_i) = \text{span}\{v_i\}$.

We can now show that v is irreducible: let $W \subseteq \mathbb{R}^d$ be T -invariant. Then W is also invariant *w.r.t.* T^i for all $i \in V$. By what we have shown before, for each $i \in V$ either $\text{span}\{v_i\} \subseteq W$, or $W \subseteq v_i^\perp$. Because of vertex-transitivity, if $\text{span}\{v_i\} \subseteq W$ for some $i \in V$, then for all $i \in V$. Since v is full-dimensional, this implies $W = \mathbb{R}^d$. Likewise, if $W \subseteq v_i^\perp$ for some $i \in V$, then

for all $i \in V$. Since v is full-dimensional, this implies $W = \{0\}$. Thus, W is a trivial invariant subspace, and v is irreducible. This proves (ii).

Since v is rigid and irreducible, it follows that v is balanced by [Theorem 2.12](#), which proves (iii). \square

Corollary 2.28. *A reducible arc-transitive realization is not of full local dimension.*

[Theorem 2.27](#) applies to arc-transitive polytopes (that is, polytopes with an arc-transitive skeleton, e.g. regular polytopes): their skeleta are rigid, irreducible and balanced.

We give a more general argument in [Chapter 3](#), but already now we can use [Theorem 2.27](#) to show that (the skeleta of) regular polytopes are not only balanced, but even spectral. This complements the findings of Licata and Powers [[47](#)], who have shown this, on a case-by-case basis, for all regular polytopes excluding the 4-dimensional exceptions (the 24-cell, 120-cell and 600-cell).

Example 2.29. Let P be an arc-transitive polytope. By [Theorem 2.27](#) its skeleton is balanced and irreducible, and thus also spherical (by [Lemma 1.13](#)). We can then apply [Corollary 2.21](#). It remains a matter of checking tabulated values for circumradius and edge-length, as well as the eigenvalues of the edge-graph, to decide whether the skeleton of P is spectral (the radii and edge-lengths of regular polytopes can be derived from the vertex coordinates given in [[18](#), Section 8.7]). The spectra of the edge-graphs have been computed in [[12](#)]).

We demonstrate this on one of the 4-dimensional exceptions not dealt with by Licata and Powers. Let v be the skeleton of the 24-cell, and let G be its edge-graph. The coordinates of the vertices of the 24-cell are all coordinate permutations and sign selections of the following vector (see [[18](#), Section 8.7]):

$$(\pm 1, \pm 1, 0, 0) \in \mathbb{R}^4. \tag{2.9}$$

We find the ratio $\ell(v)/r(v) = 1$. Since v is arc-transitive and balanced, the formula for relative edge-length in [\(2.6\)](#) can be rearranged for θ and yields

$$\theta = \deg(G) \left(1 - \frac{1}{2} \cdot \left[\frac{\ell(v)}{r(v)} \right]^2 \right) = 8 \cdot \left(1 - \frac{1}{2} \right) = 4. \tag{2.10}$$

Indeed, the spectrum of G is $\{8^1, 4^4, 0^6, (-4)^4, (-8)^1\}$ with eigenvalue $\theta_2 = 4$ of multiplicity four [[12](#)]. This matches the dimension of P , and we find that (the skeleton of) the 24-cell is indeed θ_2 -spectral.

The reader can use [Table 2.1](#) and [\(2.6\)](#) to convince himself that all regular polytopes are θ_2 -spectral. In each case $\dim \text{Eig}_G(\theta_2)$ matches d .


Even though we can verify it on a case-by-case basis, we still seem to miss a fundamental reason for why these skeleta are spectral (or stronger, θ_2 -spectral). One such reason for most regular polytopes will be given already in the next section, still not using any polytope structure (see [Example 2.36](#)). This approach still fails for the 4-dimensional exceptions, though the reasons are unrelated to the ones of Licata and Powers.

We close this section by constructing an irreducible arc-transitive realization of full local dimension that is *not* spectral, showing that [Theorem 2.27](#) cannot be improved in general.

polytope	d	$\deg(G)$	$\ell(v)$	$r(v)$	θ_2
n -gon	2	2	$2 \sin(2\pi/n)$	1	$2 \cos^2(2\pi/n)$
d -simplex	d	d	$\sqrt{2(d+1)}$	\sqrt{d}	-1
d -cube	d	d	2	\sqrt{d}	$d-2$
d -crosspolytope	d	$2(d-1)$	$\sqrt{2}$	1	$d-1$
dodecahedron	3	3	4	$3+3\sqrt{5}$	$\sqrt{5}$
icosahedron	3	5	4	$10+2\sqrt{5}$	$\sqrt{5}$
24-cell	4	8	1	1	4
120-cell	4	4	$3-\sqrt{5}$	$\sqrt{8}$	$(1+3\sqrt{5})/2$
600-cell	4	12	2	$1+\sqrt{5}$	$3+3\sqrt{5}$

Table 2.1. Metric and spectral data for the edge-graphs of the regular polytopes.

Example 2.30. Consider the graph $G := C_6 \times C_6$. This graph can be found as the edge-graph of the 4-dimensional (6,6)-duoprism, or alternatively, as the edge-graph of the (non-convex) hexagonal torus:



The graph is arc-transitive with spectrum

$$\text{Spec}(G) = \{4^1, 3^4, 2^4, 1^4, 0^{10}, (-1)^4, (-2)^4, (-3)^4, (-4)^1\}.$$

Note in particular the eigenvalue zero of multiplicity ten. Let $V = V(C_6) \times V(C_6) = \{(i, j) \mid i, j \in V(C_6)\}$ be the vertex set of G . Then

$$v_{(i,j)} := \begin{pmatrix} (-1)^i \\ (-1)^j \end{pmatrix} \in \mathbb{R}^2$$

defines a 2-dimensional irreducible and 0-balanced $\text{Aut}(G)$ -realization. This realization is of full local dimension, but is not spectral, since its dimension is smaller than $\dim \text{Eig}_G(0) = 10$.

2.4 Distance-transitive realizations

After various failures in pinning down “sufficient symmetry”, we turn to a particularly strong symmetry, called *distance-transitivity*.

Distance-transitive graphs, and their generalizations, the *distance-regular graphs*, form a class of graphs especially accessible by the methods of spectral graph theory. The standard

source is the monograph by Brouwer, Cohen and Neumaier [11]. The generic distance-regular graph has a trivial automorphism group, which makes it less relevant to this discussion, and we focus on distance-transitive graphs only.

Let $\text{dist}(i, j)$ denote the (graph theoretical) *distance* between any two vertices $i, j \in V$, i.e., the length of a shortest path between i and j . The *diameter* of a connected graph is

$$\text{diam}(G) := \max_{i, j \in V} \text{dist}(i, j),$$

the maximal distance between any two of its vertices. In particular, we shall assume that all graphs in this section are connected.

Definition 2.31. A group $\Sigma \subseteq \text{Aut}(G)$ acts *distance-transitively* on G if it acts transitively on each of the sets

$$D_\delta := \{(i, j) \in V \times V \mid \text{dist}(i, j) = \delta\}, \quad \text{for each } \delta \in \{0, \dots, \text{diam}(G)\}.$$

Distance-transitive graphs and realizations are defined parallel to Definition 2.14. Note that being arc-transitive is equivalent to being transitive on the set D_1 . In particular, distance-transitive graphs are arc-transitive.

Example 2.32. Complete graphs and cycle graphs are distance-transitive. More generally, the edge-graphs of the regular polytopes are distance-transitive, with the usual 4-dimensional exceptions (the 24-cell, 120-cell and 600-cell). Even stronger, the skeleton of any of these regular polytopes is a distance-transitive realization of the edge-graph (this can be checked by hand, but will also follow from our results in Chapter 3).

Other examples of distance-transitive graphs that are not necessarily edge-graphs are the Petersen graph (see Figure 2.2) and all complete r -partite graphs $K_{k, \dots, k}$. Further examples are listed in Theorem 3.3 in the next chapter.




Figure 2.2. The Petersen graph.

For us, the most relevant observation concerning distance-transitive symmetry was already made in [11]:

Theorem 2.33 ([11], Proposition 4.1.11). *If $\Sigma \subseteq \text{Aut}(G)$ is distance-transitive, then the Σ -irreducible subspaces of \mathbb{R}^n are exactly the eigenspaces of G .*

We shall give a partial proof for Theorem 2.33 in Section 2.5, that is, we show that the eigenspaces are indeed irreducible (see Lemma 2.40).

As we mentioned in the introduction of this chapter, the observation of [Theorem 2.33](#) was made by Du and Fan in [[21](#), Section 4] specifically for the Petersen graph. This observation now also follows from the fact that the Petersen graph is distance-transitive.

[Theorem 2.33](#) in the form of realizations reads as follows (cf. [Definition 2.8](#)):

Theorem 2.34. *The irreducible distance-transitive realizations of G are spectral.*

In other words, distance-transitivity is a “sufficient symmetry” in the sense of [Question 6](#). One can wonder whether there is any weaker (and still easily definable) form of symmetry with the same property. We do not have an answer to this.

Corollary 2.35. *Let $\nu \in \mathcal{A}_d(G, \Sigma)$ be a (not necessarily irreducible) distance-transitive realization. Then there holds:*

- (i) ν is rigid.
- (ii) ν is an $\text{Aut}(G)$ -realization (i.e., Σ cannot be geometrically separated from $\text{Aut}(G)$).
- (iii) the following are equivalent: ν being balanced, spectral and irreducible.

Proof. The graph G has only finitely many eigenspaces. By [Theorem 2.33](#) there are only finitely many Σ -irreducible subspaces. All Σ -realizations are then rigid by [Theorem 1.24](#). This proves (i).

By [Theorem 2.33](#) the Σ -irreducible subspaces are eigenspaces of G , which are $\text{Aut}(G)$ -invariant. It follows that all Σ -invariant subspaces are $\text{Aut}(G)$ -invariant. By [Theorem 1.19](#) all Σ -realizations are then $\text{Aut}(G)$ -realizations. This proves (ii).

Finally, let $U \subseteq \mathbb{R}^n$ be the arrangement space of ν . If ν is irreducible, then ν is spectral by [Theorem 2.34](#). If ν is spectral, then it is balanced by definition. If ν is θ -balanced, then $U \subseteq \text{Eig}_G(\theta)$. But since $\text{Eig}_G(\theta)$ is Σ -irreducible (by [Theorem 2.33](#)), and U is Σ -invariant (by [Lemma 1.12](#)), $U = \text{Eig}_G(\theta)$ and ν is spectral and irreducible. This proves (iii). \square

Example 2.36. Let P be a polytope with a distance-transitive skeleton (e.g. a regular polytope other than a 4-dimensional exception, cf. [Example 2.32](#)). By [Theorem 2.27](#) the skeleton of P is irreducible, and by [Theorem 2.34](#) it is then spectral.

This gives a very short proof for the results Licata and Powers [[47](#)], except that it does not yet show that these skeleta are spectral to θ_2 .

In [Chapter 4](#) we give a complete classification of “distance-transitive polytopes” (see [Theorem 4.18](#)).

2.5 Cosine vectors and cosine sequences

In this final section of the chapter we discuss an idea for extending results like [Theorem 2.34](#) beyond distance-transitivity. We also give a partial proof of [Theorem 2.34](#).

Definition 2.37. For a vertex-transitive realization ν and some vertex $i \in V$, the *cosine vector* $u \in \mathbb{R}^n$ of ν is the vector with components $u_j := \langle \nu_i, \nu_j \rangle$ for all $j \in V$.

Vertex-transitivity ensures that the cosine vector is independent of the choice of $i \in V$ (up to some coordinate permutation). We can therefore assume that $u_i := \langle v_1, v_i \rangle$. Also, we can write $u = \Phi v_1$ and see that the cosine vector $u \in \text{span } \Phi =: U$ is contained in the arrangement space of v .

Our goal is to develop a technique to show that a particular given θ -balanced realization is actually θ -spectral. This is trivial if we know the multiplicity of $\theta \in \text{Spec}(G)$, but harder if we have only access to geometric data. Consider the following idea:

Observation 2.38. Let v be a θ -balanced Σ -realization with arrangement space $U \subseteq \mathbb{R}^n$. In particular, U is Σ -invariant (by [Lemma 1.12](#)) and $U \subseteq \text{Eig}_G(\theta)$.

If U is a proper subspace of $\text{Eig}_G(\theta)$, then the subspace $U' := U^\perp \cap \text{Eig}_G(\theta) \subseteq \text{Eig}_G(\theta)$ is non-zero and Σ -invariant. We can choose a θ -balanced Σ -realization v' with arrangement space U' (via [Construction 1.2](#)). If $u, u' \in \mathbb{R}^n$ are the cosine vectors of v and v' respectively, then $\langle u, u' \rangle = 0$ because they are contained in the orthogonal subspaces U and U' .

The idea is to show that, in the right setting, being balanced Σ -realizations to the same eigenvalue is already so restrictive, that the corresponding cosine vectors have no chance to be orthogonal. If this is the case, then we reached a contradiction and found that $\text{Eig}_G(\theta)$ must have been Σ -irreducible and v therefore θ -spectral.

This applies most directly in the case of distance-transitivity: one can show that the cosine vectors of distance-transitive realizations depend only on G and θ , and thus two distance-transitive θ -balanced realizations of G cannot have orthogonal cosine vectors.

This is well-known, but we included a proof below (see [Lemma 2.40](#)).

Observation 2.39. If v is distance-transitive, then the value of $u_i = \langle v_1, v_i \rangle$ depends only on $\delta := \text{dist}(1, i)$. One therefore groups all entries with the same distance to $1 \in V$ into a single entry u_δ , for all $\delta \in \{0, \dots, \text{diam}(G)\}$. The sequence $u_0, \dots, u_{\text{diam}(G)}$ is called *cosine sequence* of v and is a well-established terminology in the theory of distance-regular graphs.

Clearly the cosine sequence and cosine vector of a distance-transitive realization determine each other. We show the following:

Lemma 2.40. *The cosine sequence of a θ -balanced distance-transitive realization of radius $r(v) = 1$ only depends on G and the eigenvalue θ .*

Proof. Let $N_\delta(i) := \{j \in V \mid \text{dist}(i, j) = \delta\}$ denote the set of vertices at distance δ from i . In a distance-transitive graph, the cardinality of the intersection $N_{\delta_1}(i) \cap N_{\delta_2}(j)$ only depends on δ_1, δ_2 and $\text{dist}(i, j)$. The following parameters are therefore independent of the exact choice of $i, j \in V$, but only depend on $\delta := \text{dist}(i, j)$:

$$c_\delta := \underbrace{|N_{\delta-1}(i) \cap N_1(j)|}_{=: N_c}, \quad a_\delta := \underbrace{|N_\delta(i) \cap N_1(j)|}_{=: N_a}, \quad b_\delta := \underbrace{|N_{\delta+1}(i) \cap N_1(j)|}_{=: N_b},$$

where N_a, N_b and N_c depend implicitly on δ, i and j . The order of the parameter names a, b and c might appear counter-intuitive, but is standard in the literature. In the literature, the list of parameters a_δ, b_δ and c_δ is called the *intersection array* of G . Note also that $N(j) = N_a \cup N_b \cup N_c$.

Suppose now that v is a θ -balanced distance-transitive realization with cosine sequence u_δ . For all $\delta \in \{0, \dots, \text{diam}(G)\}$ there is an $j \in N_\delta(1)$, and thus

$$\begin{aligned} \theta u_\delta = \langle v_1, \theta v_j \rangle &\stackrel{(2.2)}{=} \left\langle v_1, \sum_{i \in N(j)} v_i \right\rangle = \sum_{i \in N_c} \overbrace{\langle v_1, v_i \rangle}^{u_{\delta-1}} + \sum_{i \in N_a} \overbrace{\langle v_1, v_i \rangle}^{u_\delta} + \sum_{i \in N_b} \overbrace{\langle v_1, v_i \rangle}^{u_{\delta+1}} \\ &= c_\delta u_{\delta-1} + a_\delta u_\delta + b_\delta u_{\delta+1}. \end{aligned}$$

Rearranging for $u_{\delta+1}$ yields a three-term recurrence for the components of the cosine sequence that only involves θ and the intersection array:

$$u_{\delta+1} = \frac{1}{b_\delta} ((\theta - a_\delta)u_\delta - c_\delta u_{\delta-1}). \quad (2.11)$$

We assumed $r(v) = 1$. Since v is arc-transitive, we obtain initial conditions

$$u_0 = [r(v)]^2 = 1, \quad u_1 = \omega(v) \stackrel{(2.6)}{=} \frac{\theta}{\text{deg}(G)}.$$

These initial conditions only depend on θ and G (its degree, and intersection array), and so by (2.11) the whole cosine sequence only depends on θ and G . \square

Lemma 2.40 together with **Observation 2.38** shows that the eigenspaces of a graph are Σ -irreducible for all distance-transitive $\Sigma \subseteq \text{Aut}(G)$ (this is also part of **Theorem 2.33**).

We do not yet have a systematic way to extend this reasoning to a large class of graphs that are not distance-transitive. However, the following example describes a general procedure that can be applied on a case-by-case basis.

Example 2.41. We show that the skeleton of the 24-cell is a spectral realization without using knowledge of its spectrum (showing that was easy when knowing the eigenvalue multiplicities, cf. **Example 2.29**). The (skeleton of the) 24-cell is arc-transitive (but not distance-transitive) and therefore balanced by **Theorem 2.27**.

From the vertex coordinates given in (2.9), we can derive the cosine vector

$$u = (2^1, 1^8, 0^6, (-1)^8, (-2)^1)$$

(we ignore the exact ordering of the entries of the vector and only list the multiplicities).

The single entry of value 2 in u belongs to the radius $r(v) = \sqrt{2}$ of this realization. Also, the eight entries with value $\langle v_1, v_i \rangle = 1$ belong to the neighbors $i \in N(1)$, and so this value is determined by (2.6). In conclusion, any other balanced arc-transitive realization v' to the same eigenvalue and of the same radius must have a cosine vector of the form

$$u' = (2^1, 1^8; x_1, \dots, x_6; y_1, \dots, y_8; z),$$

where the x_i match up with the 0-entries in u , the y_i match up with the -1 -entries in u , and z matches with the -2 -entry.

As discussed in [Observation 2.38](#), we can assume $\langle u, u' \rangle = 0$, which expands to

$$0 = \langle u, u' \rangle = 4 + 8 - y_1 - \cdots - y_8 - 2z. \quad (2.12)$$

We may further assume that v' is centered at the origin, that is, $v'_1 + \cdots + v'_n = 0$ (otherwise, translate it; this does not affect any relevant properties). Applying $\langle v'_1, \cdot \rangle$ we get


$$0 = \sum_{i \in V} \langle v'_1, v'_i \rangle = \sum_{i \in V} u'_i = 2 + 8 + x_1 + \cdots + x_6 + y_1 + \cdots + y_8 + z \quad (2.13)$$

The sum of [\(2.12\)](#) and [\(2.13\)](#) yields $z = 22 + x_1 + \cdots + x_6$.

Note that each component of u' satisfies $u'_i = \langle v'_1, v'_i \rangle = [r(v)]^2 \cos \angle(v'_1, v'_i) \in [-2, 2]$. But this is incompatible with $z = 22 + x_1 + \cdots + x_6$. We then found that no second arc-transitive θ -balanced realization besides v can exist, and that the skeleton of the 24-cell is spectral by [Observation 2.38](#).

[Example 2.41](#) makes clever use of the shape of the cosine vector of the 24-cell. The same argument works essentially unchanged for skeleta with a similar distribution of entries in the cosine vector. For example, the skeleton of the *cubeoctahedron* (also arc-transitive) has cosine vector

$$u = (2^1, 1^4, 0^2, (-1)^4, (-2)^1).$$



We do not yet know whether this idea can be applied systematically to a wide variety of cosine vectors.

Summary

In this chapter we asked whether symmetry can be sufficient to ensure that a graph realization is spectral. Using the arrangement space dictionary, the question has been reformulated in terms of subspaces, and we found us asking about the relation between eigenspaces of G and invariant/irreducible subspaces of $\Sigma \subseteq \text{Aut}(G)$.

We have seen that this connection between symmetry and spectrum is still quite weak for vertex-, edge- and arc-transitive realizations. But we also identified distance-transitivity as a “sufficient symmetry” in the sense of [Question 6](#).

Returning to our initial motivation, how far have these investigations brought us in better understanding spectral polytopes (*i.e.*, polytopes with spectral skeleta)?

- [Theorem 2.5](#) applies directly to spectral polytopes and shows that they are as symmetric as their edge-graphs.

- With [Theorem 2.27](#) we conclude that the skeleta of arc-transitive polytopes are always rigid, irreducible and balanced (to some eigenvalue θ). It follows from [Lemma 2.20](#) that many metric properties of these polytopes can then be computed from θ .
- In [Example 2.29](#) we gave a procedure to decide whether a given arc-transitive polytope is spectral by just assuming knowledge of its edge-length, its circumradius, and the spectrum of its edge-graph. This was sufficient to reproduce and generalize the results of Licata and Powers [[47](#)], that all the regular polytopes are θ_2 -spectral, including the 4-dimensional exceptions.
- As a consequence of [Theorem 2.34](#), “distance-transitive polytopes” (polytopes with a distance-transitive skeleton, see [Section 4.6](#)) are spectral.

But we have also learned that for a complete understanding of spectral polytopes using the polytope structure is inevitable:

- All polytopes that we have identified as spectral turned out to be θ_2 -spectral. However, the meaning of θ_2 as the relevant eigenvalue for polytopes is still mysterious. At least for “distance-transitive polytopes” there are some not too deep arguments to show that they must be θ_2 -spectral (using arguments about the sign changes in eigenvectors), but in view of the more general theory of [Chapter 3](#), this proof is omitted.
- We cannot yet conclude that a sufficiently symmetric polytope is *uniquely* determined by its edge-graph. It is conceivable that there are distinct eigenvalues $\theta, \theta' \in \text{Spec}(G)$ so that both, the θ -realization and the θ' -realization, are the skeleta of polytopes. These polytopes would have the same edge-graph and would both be spectral.

In the positive, the results of this chapter have a larger applicability, being valid for all realizations and not just polytope skeleta. We have seen that arc-transitivity combined with full local dimension comes with nice properties, and we have seen that distance-transitive symmetries cannot be geometrically separated from $\text{Aut}(G)$.

We are also left with several open questions about graph realizations:

Question 2.42. Are irreducible arc-transitive realizations always balanced?


Equivalently, if $\Sigma \subseteq \text{Aut}(G)$ is arc-transitive, are then all Σ -irreducible subspaces contained in eigenspaces?

Question 2.43. Are arc-transitive $\text{Aut}(G)$ -realizations rigid?

Equivalently, if G is arc-transitive, are there only finitely many $\text{Aut}(G)$ -invariant subspaces?

[Questions 2.42](#) and [2.43](#) can also be asked for combined vertex- and edge-transitivity or half-transitivity. Likewise, we wonder whether [Theorem 2.27](#) can be proven for these other symmetry classes too.

3 Eigenpolytopes and Spectral Polytopes



Our impression so far is that the structure of graph realizations is insufficient for acquiring a deeper understanding of how symmetry and spectral properties of polytopes are related. In this chapter we apply results from convex geometry that make explicit use of the polytope structure and we derive geometric criteria for spectral polytopes. In particular, we shall see that polytopes of combined vertex- and edge-transitivity are θ_2 -spectral.

For this chapter let $P \subset \mathbb{R}^d$ denote a convex polytope, *i.e.*, the convex hull of finitely many points (see also [Appendix C](#) for a reminder on the basic terminology for polytopes). We shall assume that P is of full dimension. By $\mathcal{F}_\delta(P)$ we denote the set of δ -dimensional faces of P .

We fix an enumeration $v_1, \dots, v_n \in \mathcal{F}_0(P)$ of the vertices of P . Let $G_P = (V, E)$ denote its edge-graph on the vertex set $V = \{1, \dots, n\}$, so that $i \in V$ corresponds to vertex $v_i \in \mathcal{F}_0(P)$. The *skeleton* of P is the map

$$\text{sk}_P : V \rightarrow \mathbb{R}^d, i \mapsto v_i,$$

which is a graph realization of the edge-graph. This definition of skeleton enables us to reuse the terminology of graph realizations. For example, a *spectral polytope* can now be simply defined as a polytope with a spectral skeleton.

A symmetry of a d -dimensional polytope is an isometry of \mathbb{R}^d that fixes P set-wise. For convenience we shall assume that P is centered at the origin, so that the symmetries are orthogonal transformations. The (*Euclidean*) *symmetry group* of P is then

$$\text{Aut}(P) := \{T \in \text{O}(\mathbb{R}^d) \mid TP = P\} \subseteq \text{O}(\mathbb{R}^d).$$

If $\text{Aut}(P)$ acts transitively on the vertices of P , edges of P , etc., then P is said to be *vertex-transitive*, *edge-transitive*, etc. This is compatible with the terminology inherited from graph realizations: *e.g.* P is vertex-transitive (in this new sense) if and only if its skeleton is vertex-transitive in the sense of [Definition 2.14](#).

Chapter overview

Our definition of “spectral polytope” deviates from the literature definition, which is based on *eigenpolytopes*. For this reason we start in [Section 3.1](#) with a formal introduction of eigenpolytopes supplemented with an extensive literature overview also addressing previous work on spectral polytopes.

From [Section 3.2](#) on we focus on *spectral polytopes*. We show that our definition of spectral polytope (via the skeleton) is aligned with the literature definition based on eigenpolytopes. We give an overview of examples and non-examples and describe the general properties of spectral polytopes.

In [Section 3.3](#) we present the *Theorem of Izmestiev*, our as of yet most powerful tool for the identification of θ_2 -spectral polytopes. We derive geometric criteria for θ_2 -spectral polytopes and show that polytopes of combined vertex- and edge-transitivity are of this form.

3.1 Eigenpolytopes

The historical origin of spectral polytopes is in the study of *eigenpolytopes*. In this section we briefly introduce eigenpolytopes and review their literature.

Eigenpolytopes were first introduced by Godsil in 1978 [\[31\]](#). They are obtained as convex hulls of spectral graph realizations.

Definition 3.1. For a graph G and an eigenvalue $\theta \in \text{Spec}(G)$, the θ -*eigenpolytope* is

$$P_G(\theta) := \text{conv}\{v_i \mid i \in V\},$$


where v_i is the θ -realization of G (see [Construction 2.3](#)).

For later use, the θ -realization used in the above definition will be denoted eig_G^θ .

Example 3.2. In the introduction we computed the θ_2 -eigenpolytope of the edge-graph of the cube. The spectrum of this edge-graph G is

$$\text{Spec}(G) = \{3^1, 1^3, (-1)^3, (-3)^1\}.$$

The corresponding eigenpolytopes are depicted below:



Only for θ_2 we obtain the cube as eigenpolytope. The θ_1 -eigenpolytope of a connected regular graph is always a single point. Likewise, whenever a regular graph is bipartite, the eigenpolytope to the smallest eigenvalue is 1-dimensional and thus a line segment.

To compute further examples, the reader can find a Mathematica script for that purpose in [Appendix F](#).

Literature overview

Godsil introduced eigenpolytopes to study symmetry groups of graphs. In [\[31\]](#) he proved the existence of a group homomorphism $\text{Aut}(G) \rightarrow \text{Aut}(P_G(\theta))$. The existence of this homomorphism also follows from the symmetry properties of θ -realizations (see [Theorem 2.5](#)).

In general, this is not an isomorphism, that is, $\text{Aut}(G) \not\cong \text{Aut}(P_G(\theta))$ (consider the eigenpolytopes in [Example 3.2](#)).

The combinatorial structure of an eigenpolytope encodes properties of the original graph. For example, Rooney [[64](#)] used the size of the facets of the $P_G(\theta)$ to deduce statements about the size of cocliques in G .

Padrol and Pfeifle [[59](#)] investigate how common graph operations translate to operations on their (Laplacian) eigenpolytopes.

Several authors became interested in constructing and investigating eigenpolytopes of particular graphs and graph families. Powers [[62](#)] studied the eigenpolytopes of the Petersen graph, which he termed the *Petersen polytopes* (one of which appears as a distance-transitive polytope in [Section 4.6](#)). In [[57](#)] Mohri described the face structure of the *Hamming polytopes*, the θ_2 -eigenpolytopes of the Hamming graphs. These polytopes turned out to be the cartesian powers of regular simplices and will also reappear as distance-transitive polytopes in [Section 4.6](#).

The author of [[8](#)] provides a careful enumeration of the eigenpolytopes (actually, spectral graph realizations) of the edge-graphs of uniform polyhedra (see [Appendix E](#)). Unfortunately, this write-up has never been published formally. The listing provides empirical evidence that every (irreducible) uniform polyhedron (including the Platonic and Archimedean solids) has a realization as an eigenpolytope of its edge-graph (see also [Example 3.11](#)).

Particular attention was given to the eigenpolytopes of distance-regular graphs [[29,32,63](#)]. For example, in [[32](#)] Godsil shows that if G is a *distance-regular* graph, then every edge of G is mapped into an edge of $P_G(\theta_2)$ (via the θ_2 -realization of G). In other words, G appears as a spanning subgraph of the edge-graph of $P_G(\theta_2)$. The same is not necessarily true for eigenvalues other than θ_2 , and it is an open question how much structure is necessary to obtain the same result for more general graph classes.

In the extreme case of this “spanning subgraph” phenomenon, G becomes isomorphic to the edge-graph of the eigenpolytope (which led to the literature definition of what we call a *spectral polytope*). This observation has been made repeatedly: we previously mentioned Licata and Powers [[47](#)], who conjectured that all regular polytopes are θ_2 -eigenpolytopes of their edge-graph. They proved their conjecture in all cases excluding the 4-dimensional exceptions. We have shown in [Example 2.29](#) that their result extends to all regular polytopes.

A major result for spectral polytopes was obtained by Godsil in [[32](#)], where he obtained a complete classification in the case of distance-regular graphs.

Theorem 3.3 ([[32](#)], Theorem 4.3). *Let G be distance-regular. If G is isomorphic to the edge-graph of its θ_2 -eigenpolytope, then G is one of the following:*

- (i) a cycle graph C_n , $n \geq 3$,
- (ii) the edge-graph of the dodecahedron,
- (iii) the edge-graph of the icosahedron,
- (iv) the complement of a disjoint union of edges (also known as a crown graph),
- (v) a Johnson graph $J(n, k)$,
- (vi) a Hamming graph $H(d, q)$,

- (vii) a halved d -cube $\frac{1}{2}Q_d$,
- (viii) the Schläfli graph, or
- (ix) the Gosset graph.




Figure 3.1. The Schläfli graph (left) and the Gosset graph (right).


A second look at this list reveals a remarkable “coincidence”: while the generic distance-regular graph has few or no symmetries, all graphs in this list are highly symmetric, in fact, *distance-transitive* (cf. Section 2.4). It is this enumeration in Theorem 3.3 that enables us to give a complete classification of “distance-transitive polytopes” in Section 4.6.

3.2 Spectral polytopes

The literature defines *spectral polytope* as a “polytope that is the eigenpolytope of its edge-graph”¹. Particular focus was thereby on the case θ_2 .

We point to some subtleties when extending this definition to arbitrary eigenvalues.

Example 3.4. Consider the regular pentagon. Its edge-graph, the 5-cycle, has the following two spectral realizations:




The left realization is constructed from the second-largest eigenvalue θ_2 and looks expectedly like the skeleton of a regular pentagon. In contrast, the θ_3 -realization resembles a *pentagram*.

¹Note that the term “spectral polytope” was just introduced in this thesis and that it is our interpretation which previous literature aims to address this phenomenon.

In either case, the corresponding eigenpolytope (the convex hull of the graph realization) is a *regular pentagon*. Should we therefore consider the pentagon as both θ_2 -spectral and θ_3 -spectral? We argue that this is undesirable, because the θ_3 -realization of the 5-cycle is not an embedding into the skeleton of the respective eigenpolytope.

This problem is not unique to dimension two (cf. [Example 2.4](#)):



For both eigenvalues, the eigenpolytope is the regular dodecahedron.

Our preferred definition of “spectral polytope” is derived from spectral graph realizations:

Definition 3.5. A polytope is called θ -spectral (or just *spectral*) if its skeleton is θ -spectral.

The connection of [Definition 3.5](#) with the literature definition is captured in the following technical proposition:

Proposition 3.6. P is θ -spectral if and only if there is a linear transformation $X \in \text{GL}(\mathbb{R}^d)$ so that $P_{G_P}(\theta) = XP$ and for which the following diagram commutes:

$$\begin{array}{ccc}
 P & \xrightarrow{X} & P_{G_P}(\theta) \\
 \swarrow \text{sk}_P & & \uparrow \text{eig}_{G_P}^\theta \\
 & & G_P
 \end{array} \tag{3.1}$$

Recall that G_P is the edge-graph of P , sk_P its skeleton and $\text{eig}_{G_P}^\theta$ denotes the θ -realization of G_P as used in defining eigenpolytopes (see [Definition 3.1](#)).

Corollary 3.7. If P is θ -spectral then it is the θ -eigenpolytope of its edge-graph (up to invertible linear transformation).

Proof of Proposition 3.6. If (3.1) commutes for $X \in \text{GL}(\mathbb{R}^d)$ then this means $\text{eig}_{G_P}^\theta = X \text{sk}_P$. By [Corollary 1.4](#), both realizations have then the same arrangement space. Since $\text{eig}_{G_P}^\theta$ is θ -spectral, sk_P is then θ -spectral too. Therefore P is θ -spectral.

For the converse, assume that P is θ -spectral. Then sk_P has arrangement space $\text{Eig}_G(\theta)$. By definition, $\text{eig}_{G_P}^\theta$ has the same arrangement space and by [Corollary 1.4](#) there is a linear transformation $X \in \text{GL}(\mathbb{R}^d)$ so that (3.1) commutes. Since P is the convex hull of sk_P , and $P_G(\theta)$ is the convex hull of $\text{eig}_{G_P}^\theta$, X relates P and $P_G(\theta)$ as required. \square

We still have the following pitfalls:

- Even if P is the θ -eigenpolytope of its edge-graph (up to invertible linear transformation), it is not necessarily θ -spectral.
- Even if the edge-graph of $P_G(\theta)$ is isomorphic to G , $P_G(\theta)$ is not necessarily θ -spectral.

Counterexamples have been provided in [Example 3.4](#).

Spectral polytopes have properties not shared by most polytopes:

Corollary 3.8. *Let P be θ -spectral. Then*

- P can be reconstructed from its edge-graph (up to invertible linear transformations).*
- if P is spherical, then P realizes all the symmetries of its edge-graph (i.e., its skeleton is an $\text{Aut}(G_P)$ -realization).*

Proof. The reconstruction (i) is obtained via the θ -eigenpolytope, which is a linear transformation of P by [Corollary 3.7](#). Part (ii) follows via [Theorem 2.5](#). \square

The introduction to this thesis already contains a brief discussion about why the properties in [Corollary 3.8](#) are relevant and generally non-trivial.

For the rest of this section we discuss examples and non-examples for spectral polytopes. Being spectral is still considered as a quite rare property. For most polytopes it is easy to see that they cannot be spectral or combinatorially equivalent to any spectral polytope (often their skeleta can also not be balanced).


Example 3.9. Let P be a *neighborly polytope* other than a simplex (*neighborly* means that its edge-graph is the complete graph K_n). The spectrum of K_n is

$$\text{Spec}(K_n) = \{(n-1)^1, (-1)^{n-1}\}.$$

The corresponding eigenpolytopes are either a point or the regular simplex. P can therefore not be spectral.

Example 3.10. For many prisms the edge-graph has no eigenvalue of multiplicity three. For example, if G is the edge-graph of the hexagonal prism, then its spectrum is

$$\text{Spec}(G) = \{3^1, 2^2, 1^1, 0^4, (-1)^1, (-2)^2, (-3)^1\}$$




We see that none of its eigenpolytopes is 3-dimensional, and so the 6-prism cannot be spectral. Even stronger, the 6-prism has no realization as a spectral polytope.

Example 3.11. The class of *uniform polytopes* is defined recursively: a uniform polygon is a regular polygon, and a uniform polytope is a vertex-transitive polytope all whose facets are uniform (see [Appendix E](#)). This class of polytopes has received much attention in the past, especially in the work of Coxeter [18, 19] and Johnson [40], but a complete classification has not been achieved to this date. This class contains the regular polytopes, Archimedean solids, permutahedra, etc.

Most uniform polytopes seem to be “essentially θ_2 -spectral”, in the sense that they are combinatorially equivalent to a θ_2 -spectral polytope (we have seen in [Example 3.10](#) that this is not true for some prisms, and likewise, we do not expect this to be true for other “reducible” uniform polytopes). We know this to be true in dimension two (for regular polygons). The informal enumeration in [8] suggests this to be true in dimension three as well. Some own numerical experiments led to no counterexamples in higher dimensions.

The image below shows the *cubeoctahedron* (left) and a spectral version of the *truncated tetrahedron* (right):




The truncated tetrahedron, as a uniform polytope, has all edges of the same length. But its unique spectral realization is not of this form (recall [Example 2.6](#), where we have shown that most realizations of the truncated tetrahedron are not even balanced).

The cubeoctahedron on the other hand is edge-transitive, and so is its edge-graph. A realization of the cubeoctahedron as a spectral polytope must therefore be edge-transitive as well (by [Corollary 3.8 \(ii\)](#)) and must have all edges of the same length. In fact, the cubeoctahedron (and as we shall see, almost every edge-transitive polytope) has a unique edge-transitive realization, and only this realization can be (and actually is) θ_2 -spectral (the hard part is to see that this realization is not just θ_2 -balanced).

All spectral polytopes that we have encountered so far have been vertex-transitive. Non-vertex-transitive spectral polytopes can be found among the *Catalan solids* (the duals to the Archimedean solids), but not all of them are spectral.


We discuss two especially relevant instances:

Example 3.12. It is known that there exist exactly two edge-transitive polyhedra that are *not* vertex-transitive ([36]; see also [Chapter 4](#), and especially [Chapter 6](#)): the *rhombic dodecahedron* (the polar dual of the cubeoctahedron from [Example 3.11](#)) and the *rhombic triacontahedron*.



Both have a unique edge-transitive realization that is depicted in the figure.

It turns out that the rhombic dodecahedron is θ_2 -spectral (this was already noted by Licata and Powers in [47]), but the rhombic triacontahedron is not spectral, and in fact, is not even balanced. For example, the θ_2 -realization of its edge-graph looks as follows:



We argue that the case of the rhombic dodecahedron is quite exceptional, and one would not have expected to find it being spectral. Recall Figure 1.8 (on page 46), which depicts a deformation of the skeleton of the rhombic dodecahedron that preserves its symmetry. Similar to the argument in Example 2.6, only finitely many of the realizations in the deformation can be spectral. But in contrast to Example 2.6, this time *only one* of the realizations in the deformation is the skeleton of a polytope. The fact that the unique realization that yields a skeleton is also one of the few realizations that are spectral (and even θ_2 -spectral) appears completely accidental (this accident fails to happen for the rhombic triacontahedron).

3.3 The Theorem of Izmestiev

In this section we discuss a powerful tool that can be used to identify spectral polytopes by geometric means – the *Theorem of Izmestiev*. Especially remarkable, this approach automatically yields θ_2 -spectral polytopes.

Recall that for a polytope P with $0 \in \text{int}(P)$, the *polar dual* P° is defined as


$$P^\circ := \{x \in \mathbb{R}^d \mid \langle x, v_i \rangle \leq 1 \text{ for all } i \in V\},$$

where the v_i are an enumeration of the vertices of P .

We generalize this notion: for a vector $c = (c_1, \dots, c_n) \in \mathbb{R}^n$ define the generalized polar

$$P^\circ(c) := \{x \in \mathbb{R}^d \mid \langle x, v_i \rangle \leq c_i \text{ for all } i \in V\}.$$

Then $P^\circ(1, \dots, 1) = P^\circ$, and $P^\circ(c)$ results from P° by shifting facets along their normal vectors.



In the following, $\text{vol}(C)$ denotes the *relative volume* (relative to the affine hull) of a compact convex set $C \subset \mathbb{R}^d$ (which will always be a polytope). One can show that $\text{vol}(P^\circ(c))$ is two times continuously differentiable in c .

Theorem 3.13 (Izmestiev [39], Theorem 2.4). *For a polytope $P \subset \mathbb{R}^d$ with $0 \in \text{int}(P)$ consider the matrix $M \in \mathbb{R}^{n \times n}$ (which we shall call the Izmestiev matrix of P) with components*

$$M_{ij} := \left. \frac{\partial^2 \text{vol}(P^\circ(c))}{\partial c_i \partial c_j} \right|_{c=(1, \dots, 1)}.$$

M has the following properties:

- (i) $M_{ij} < 0$ whenever $ij \in E$.
- (ii) $M_{ij} = 0$ whenever $ij \notin E$ and $i \neq j$.
- (iii) M has a unique negative eigenvalue of multiplicity one.
- (iv) $M\Phi = 0$, where Φ is the arrangement matrix of (the skeleton of) P .
- (v) $\dim \ker M = d$.

In the language of [39], the matrix M constructed in **Theorem 3.13** is a *Colin de Verdière matrix* of G_P , that is, a matrix satisfying a list of properties, among these (i), (ii) and (iii), and another one known as the *strong Arnold property*.

Among the Colin de Verdière matrices, one usually cares about the ones with the largest possible kernel. The dimension of this largest kernel is known as the *Colin de Verdière graph invariant* $\mu(G_P)$ [71], and in this sense, **Theorem 3.13** proves $\mu(G_P) \geq d$. This is not too surprising and was known before. However, the result of Izmistiev is remarkable for a different reason: it shows that there is a Colin de Verdière matrix whose corank (the dimension of the kernel) is *exactly* d (property (v)) and that is compatible with the geometry of P (property (iv)).

Izmestiev furthermore shows that the matrix M can be expressed in terms of simple geometric properties of the polytope: for $ij \in E$ let $f_{ij} \in \mathcal{F}_{d-2}(P^\circ)$ be the dual face to the edge $\text{conv}\{v_i, v_j\} \in \mathcal{F}_1(P)$. Then

$$M_{ij} = -\frac{\text{vol}(f_{ij})}{\|v_i\| \|v_j\| \sin \angle(v_i, v_j)}. \quad (3.2)$$

The entries M_{ii} , $i \in V$ on the diagonal can be computed from property (iv), $M\Phi = 0$.

The proof of the Theorem of Izmistiev uses techniques from convex geometry, in particular, mixed volumes, and lies beyond the scope of this thesis.

If Φ is the arrangement matrix of P , and since P is full-dimensional, we have $\text{rank } \Phi = d$. From **Theorem 3.13** (iv) and (v) then follows

$$\text{span } \Phi = \ker M. \quad (3.3)$$

Theorem 3.14. *Let P be a polytope with $0 \in \text{int } P$ and $M \in \mathbb{R}^{n \times n}$ its Izmistiev matrix. If*

- (i) M_{ii} is the same for all $i \in V$, and

(ii) M_{ij} is the same for all $ij \in E$,

then P is θ_2 -spectral.

Proof. By assumption, there are $\alpha \in \mathbb{R}$ and $\beta > 0$ (using [Theorem 3.13 \(i\)](#)) so that

$$M = \alpha \text{Id} - \beta A \quad \implies \quad A = -\frac{1}{\beta} M + \frac{\alpha}{\beta} \text{Id}, \quad (3.4)$$

where A is the adjacency matrix of G_P .

The arrangement space of P is $U := \text{span } \Phi \stackrel{(3.3)}{=} \ker M$. By [Theorem 3.13 \(iii\)](#), $\ker M$ is the eigenspace to the *second-smallest* eigenvalue of M . From (3.4) then follows that U is the eigenspace to the *second-largest* eigenvalue $\theta_2 = \alpha/\beta$ of A . Thus, P is θ_2 -spectral. \square

By (3.2) the entries $M_{ii}, i \in V$ and $M_{ij}, ij \in E$ are determined by local geometric properties of the vertices and edges of P . If these vertices and edges are identical under symmetry then they are identical geometrically and we obtain

Corollary 3.15. *If P is vertex- and edge-transitive, then P is θ_2 -spectral.*

Note that we made the implicit assumption $0 \in \text{int}(P)$. This is justified if we assume that P is centered at the origin, that is

$$\sum_{i=1}^n v_i = 0.$$

When rescaled with $1/n$, this equation becomes a convex combination of the vertices that yields zero. Using full dimension one can then show that 0 lies indeed in the interior of P .

We have then established that combined vertex- and edge-transitivity is a “sufficient symmetry” in the sense of [Question 3](#). Unfortunately, θ_2 emerges still somehow “magically” from the use of the Theorem of Izmistiev as a black box. All we can say at this level is that θ_2 is linked to the convexity of P , and its specialness follows from the properties of the “Hessian matrix of the volume” (see [\[39\]](#) for details).

[Corollary 3.15](#) applies to all regular polytopes and thus explains the finding of Licata and Powers [\[47\]](#) and our finding concerning the remaining 4-dimensional exceptions (see [Example 2.29](#)). The consequences of [Corollary 3.15](#) will be explored in detail in [Chapter 4](#) (see specifically [Theorem 4.5](#)).

It remains the question whether [Theorem 3.14](#) already characterizes spectral polytopes, or at least θ_2 -spectral polytopes. If it turns out to characterize general spectral polytopes, then this shows that spectral polytopes are always θ_2 -spectral. All of this is widely open.

We close this chapter with a theorem about edge-transitive spectral polytopes. For its proof we have to anticipate a result from [Chapter 6](#).

Theorem 3.16. *If P is an edge-transitive θ -spectral polytope, then*

(i) $\theta = \theta_2$.

(ii) if P is not vertex-transitive, then it is the rhombic dodecahedron (cf. [Example 3.12](#)).

Proof. Suppose that P is vertex-transitive. By [Corollary 3.15](#) it is then θ_2 -spectral and $\theta = \theta_2$. This proves (i) if P is vertex-transitive.

Suppose now that P is *not* vertex-transitive. [Theorem 6.1](#) in [Chapter 6](#) states that an edge-transitive polytope that is not vertex-transitive is either a polygon, the rhombic dodecahedron or the rhombic triacontahedron (introduced in [Example 3.12](#)). Since a spectral polytope has all the symmetries of its edge-graph (by [Corollary 3.8 \(ii\)](#)), a polygon is spectral if and only if it is vertex-transitive. Of the both remaining polyhedra, only the rhombic dodecahedron is spectral (as discussed in [Example 3.12](#)). This proves (ii). The rhombic dodecahedron is also θ_2 -spectral, which proves (i) also in the not vertex-transitive case. \square

Summary

In this final chapter of Part I we gave an overview of eigenpolytopes and spectral polytopes. We arrived at a geometric criterion for being θ_2 -spectral (see [Theorem 3.14](#)) and applied it to polytopes of combined vertex- and edge-transitivity. This implies, via [Corollary 3.8](#), that a polytope of this symmetry is uniquely determined by its edge-graph and realizes all its symmetries (for details, see [Theorem 4.5](#) in Part II). We thereby demonstrated that spectral graph theory adds to the toolbox of proof techniques to be of use in polytope theory.

This chapter still left us with many unanswered questions and we repeat two such questions from the introduction to which no final answer has been found:

Question 3.17 (*cf.* [Question 1](#)). Can we classify spectral polytopes?

Question 3.18 (*cf.* [Question 2](#)). Is there a spectral polytope for an eigenvalue other than θ_2 ?

Some considerations on nodal domains make it plausible that spectral polytopes exist only for θ_2 . We have shown this for edge-transitive polytopes. But no proof is known for the general case.

Question 3.19. Are spectral polytopes (or at least θ_2 -spectral polytopes) already characterized by [Theorem 3.14](#)?

If [Theorem 3.14](#) characterizes general spectral polytopes, then this implies a negative answer to [Question 3.18](#). One can verify that the rhombic dodecahedron (*cf.* [Example 3.12](#)) satisfies the conditions of [Theorem 3.14](#). This is clear for (ii). The fact that (i) holds appears purely accidental.

We have found that many uniform polytopes are spectral, among them the regular polytopes and edge-transitive polytopes. Some uniform polytopes are not spectral which seems to result from their symmetry group not being irreducible (*e.g.* most prisms and anti-prisms, see [Example 3.10](#)). Others are irreducible and not spectral, but still combinatorially equivalent to a spectral polytope (*e.g.* the truncated tetrahedron in [Example 3.11](#)). All “irreducible” uniform polytopes that we have tested (numerically or analytically) turned out to be of this kind.

Question 3.20. If P is an “irreducible” uniform polytope (or a Wythoffian polytope, see [Appendix E](#)), then is it combinatorially equivalent to a θ_2 -spectral polytope?

All the spectral polytopes that we have encountered so far had a reasonably large symmetry group. We have seen that symmetry can be a sufficient condition for being spectral, but is it also necessary?


Question 3.21. How symmetric must a spectral polytope be? Can it have a trivial symmetry group?

Having many symmetries is one reason for a graph to have large eigenspaces, which then give rise to eigenpolytopes of reasonably large dimensions. But if the graph has a trivial symmetry group, there have to be other structural reasons for large eigenspaces. For example, distance-regular graphs have large eigenspaces while often being without any symmetry. However, as mentioned below [Theorem 3.3](#), distance-regular graphs give rise to spectral polytopes only when highly symmetric, namely, distance-transitive.

Part II

Transitivities in Convex Polytopes

4 Edge-Transitive Polytopes



Over the course of Part I we have seen that polytopes with sufficiently many symmetries can have remarkable properties, such as unique reconstruction from the edge-graph and a strong relation between geometric and combinatorial properties.

Motivated by these findings, Part II takes a deeper look at the symmetry that was recognized to be at the core of this phenomenon – *edge-transitivity*.

As in [Chapter 3](#), $P \subset \mathbb{R}^d$ shall always denote a convex d -dimensional polytope with edge-graph G_P . We assume that P is full-dimensional.

From this chapter on we shall assume some familiarity with the (*finite*) *reflection groups* as well as their orbit polytopes, the *Wythoffian polytopes*. A sufficient introduction to both can be found in [Appendix D](#) and [Appendix E](#) respectively. In brief, a reflection group $\Gamma \subseteq O(\mathbb{R}^d)$ is a matrix group generated by reflections, that is,

$$\Gamma := \langle \rho_r \mid r \in R \rangle,$$

where $R \subseteq \mathbb{R}^d \setminus \{0\}$ and $\rho_r \in O(\mathbb{R}^d)$ is the reflection on the hyperplane r^\perp . A Wythoffian polytope is then an orbit polytope of a *finite* reflection group, that is,

$$P = \text{Orb}(\Gamma, x) := \text{conv}\{Tx \mid T \in \Gamma\},$$

for some point $x \in \mathbb{R}^d$. The finite reflection groups admit a surprisingly concise classification, and so do the Wythoffian polytopes.

Of particular importance are furthermore the so-called Wythoffian *uniform* polytopes, that is, Wythoffian polytopes in which all edges are of the same length (see [Appendix E.3](#)). These polytopes are our most important source for examples. Among others, they include the Platonic and Archimedean solids, prisms and permutahedra. To denote specific Wythoffian uniform polytopes we make use of their *Coxeter-Dynkin diagrams* (see [Appendix E.2](#)).

Chapter overview

In [Section 4.1](#) we take a general look at transivities in convex polytopes. We frame the state of the knowledge by discussing the two currently best understood extreme cases – the regular polytopes and the vertex-transitive polytopes.

We then shift our attention to the next “simplest” kind of transitivity – *edge-transitivity*. [Section 4.2](#) provides an overview of the known edge-transitive polytopes in lower dimensions.

In the subsequent sections we introduce and explore a hierarchy of edge-transitive polytopes. In [Section 4.3](#) we consider edge-transitive polytopes that are additionally vertex-transitive. We explore the consequences of their spectral properties. In [Section 4.4](#) we consider the sub-class of arc-transitive polytopes, for which we conjecture a classification in terms of Wythoffian polytopes. In [Section 4.5](#) we briefly address the potential symmetry class of half-transitive polytopes. Finally, in [Section 4.6](#) we apply a classification of Godsil ([Theorem 3.3](#)) to give a complete list of distance-transitive polytopes.

4.1 An overview of transitivity in polytopes

Recall that $\text{Aut}(P) \subseteq O(\mathbb{R}^d)$ denotes the *Euclidean symmetry group* of $P \subset \mathbb{R}^d$, i.e., the group of orthogonal matrices that fix P set-wise.

Definition 4.1.

- (i) P is δ -transitive for some $\delta \in \{0, \dots, d-1\}$ if $\text{Aut}(P)$ acts transitively on the δ -dimensional faces of P .
- (ii) P is Δ -transitive for some $\Delta \subseteq \{0, \dots, d-1\}$ if P is δ -transitive for all $\delta \in \Delta$.

For example, 0-transitive polytopes have been previously called *vertex-transitive*. Likewise, 1-transitive polytopes have been called *edge-transitive*.

Although the study of such transivities has the appeal of a classical question, general δ -transitive polytopes seem to be still badly understood. For example, in the 2003 edition of his book “Convex Polytopes”, Branko Grünbaum writes [[35](#), Section 19.1, p. 413]

No serious consideration seems to have been given to polytopes in dimension $d \geq 4$ about which transitivity of the symmetry group is assumed only for faces of suitably low dimensions, [...].

In his article “A Hierarchical Classification of Euclidean Polytopes with Regularity Properties” ([\[52\]](#), or as a chapter in [[5](#), Section 3, p. 74]), Horst Martini notes

More generally, one can consider k -transitivity for $k \in \{0, 1, \dots, d-1\}$. Among the various questions concerning this notion, the relation between the transivities of different dimensions deserve to be investigated.

In essence, the following question appears widely open:

Question 4.2. For which $\delta \in \{0, \dots, d-1\}$ resp. $\Delta \subseteq \{0, \dots, d-1\}$ can we classify the δ -transitive resp. Δ -transitive polytopes?

As of yet, a full classification is known only in the case $\Delta = \{0, \dots, d-1\}$. The Δ -transitive polytopes of this kind are transitive on faces of all dimensions. They are better known as the *regular polytopes* and were completely classified by the Swiss mathematician Ludwig Schläfli already in the 19th century [[66](#)]. The 3-dimensional regular polytopes, known as *Platonic solids*, have been known since antiquity (see [Figure 4.1](#)).




Figure 4.1. The five 3-dimensional regular polytopes (better known as *Platonic solids*). From left to right: tetrahedron, cube (or hexahedron), octahedron, dodecahedron and icosahedron.

It turns out that being regular is more restrictive in higher dimensions. This is little surprising given that the set $\Delta = \{0, \dots, d-1\}$ of symmetry restrictions grows in size with increasing d . One finds that there are only *three* regular polytopes in any dimension $d \geq 5$, namely, the d -simplex, the d -cube and the d -crosspolytope (the dual of the d -cube).

In contrast, the class of 0-transitive (or vertex-transitive) polytopes is of the quite opposite flavor. Those polytopes can most likely not be classified in a satisfying manner: it turns out that almost every finite group \mathcal{G} is isomorphic to the symmetry group of a vertex-transitive polytope [2, 25] (recall also Section 1.3). The exceptions to this rule are most of the cyclic and dicyclic groups: any polytope that attempts to have such a group of symmetries turns out to have additional symmetries (similar to how distance-transitivity cannot be geometrically separated from additional symmetries, cf. Section 2.4).

We have then seen that the presence of geometric transivities is no guarantee for either a lot of structure or richness. For every other class of transivities we should therefore ask the following question: do we expect this class to be more like the “regular polytopes” (quite restricted, can be completely classified), or more like the “vertex-transitive polytopes” (wild, probably not subject to any manageable classification).


We close this brief overview with a note on two related research directions that not quite fit our roadmap, but which are still of major interest:

- We can drop geometric constraints and move to a more abstract combinatorial setting. There we can find the *abstract regular polytopes*, the theory of which has experienced an immense development over the last decades (see the standard books by Schulte and McMullen [55, 56]). As it turns out, dropping convexity gives rise to a much richer family of objects with still enough structure for a fruitful investigation.
- Instead of transitivity on some or all faces, one can ask for polytopes with *few* orbits on these substructures. This path has been followed by Matteo [53]. He found that when setting an upper bound on the number of so-called “flag-orbits”, we are still left with only the regular polytopes as soon as we hit sufficiently high dimension.

4.2 Known edge-transitive polytopes

We shall now focus on 1-transitive (or edge-transitive) polytopes. Those are well-understood up to dimension three.

Enumerating the edge-transitive *polygons* is straightforward: for odd n , there exists only the regular n -gon, and for even n , there is a 1-dimensional continuous family of edge-transitive n -gons, including the regular n -gon, but also non-regular polygons with alternating interior angles:



The edge-transitive *polyhedra* (that is, the 3-dimensional edge-transitive polytopes) have also been completely enumerated (e.g. see [36]). Besides the five Platonic solids, this list contains the following four other polyhedra (nine in total):





Figure 4.2. From left to right: the cuboctahedron, icosidodecahedron, rhombic dodecahedron and rhombic triacontahedron. The first two polyhedra are uniform and are known as the *quasi-regular* polyhedra. The latter two belong to the *Catalan solids* and are the polar duals of the former two.

Among the nine edge-transitive polyhedra only the rhombic dodecahedron and the rhombic triacontahedron are *not* also vertex-transitive (cf. [Example 3.12](#)). We shall see that this is quite special (see [Section 4.3](#) and [Chapter 6](#)).

And this is how far we have come: there has not been obtained a complete enumeration of edge-transitive polytopes in any dimension $d \geq 4$. This is not to say that there are no known examples of such polytopes. In fact, many examples can be constructed from the regular polytopes. We describe two such constructions without proof.

Example 4.3. Let $P \subset \mathbb{R}^d$ be a regular polytope, and let $w_1, \dots, w_m \in \mathbb{R}^d$ be an enumeration of its *edge midpoints*. The convex hull $\text{conv}\{w_1, \dots, w_m\}$ turns out to be an edge-transitive polytope. These polytopes are known as the *rectified regular polytopes*.



In dimension three, this is one way to obtain two of the four non-regular edge-transitive polytopes (cf. [Figure 4.2](#)): the cuboctahedron (from the cube or octahedron), and the icosidodecahedron (from the dodecahedron or icosahedron). In particular, rectifications of dual Platonic solids result in the same polyhedron.

This is not the case in four dimensions: for $d = 4$ rectification leads to five new edge-transitive polytopes, one for each of the six regular polytopes, excluding the 4-dimensional cross-polytope, whose rectification is the regular 24-cell.

Example 4.4. Given a regular polytope $P \subset \mathbb{R}^d$, the *cartesian powers*

$$P^k := \underbrace{P \times \cdots \times P}_{k \text{ times}} \subset \mathbb{R}^{kd}$$

are edge-transitive.

With P being a regular n -gon, P^2 is a so-called (n, n) -*duoprism*. The higher powers we shall call (n, \dots, n) -*hyperprisms*. The cartesian powers of simplices were previously described as the *Hamming polytopes* (θ_2 -eigenpolytopes of the Hamming graphs [57], see Section 3.1).

In fact, all the polytopes constructed in the previous examples are instances of *Wythoffian uniform polytopes*. A look at their Coxeter-Dynkin diagrams hints to a possible generalization that we discuss further below (see Conjecture 4.11).

4.3 Simultaneously vertex- and edge-transitive polytopes

Our goal for this chapter is to bring structure into the class of edge-transitive polytopes. For this purpose, we developed the hierarchical classification scheme presented in Figure 4.3.

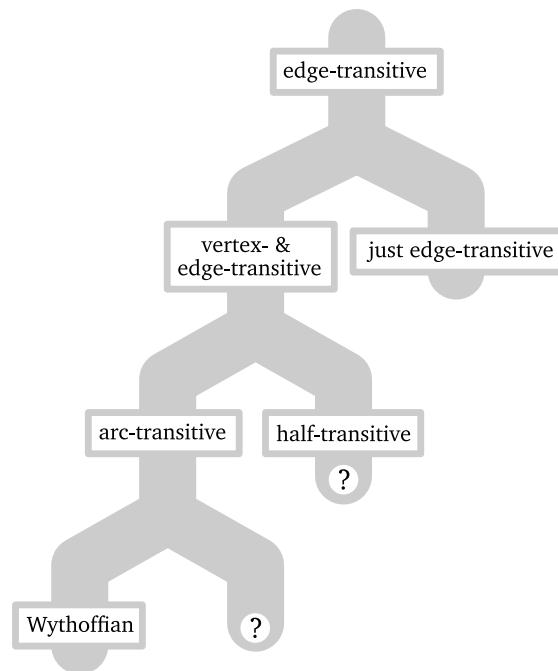


Figure 4.3. A hierarchy of edge-transitive polytopes.

At the first branching the class of edge-transitive polytopes divides into two sub-classes: the edge-transitive polytopes that are also vertex-transitive (e.g. the regular polytopes), and those

that are not vertex-transitive (e.g. the rhombic dodecahedron and the rhombic triacontahedron, shown in [Figure 4.2](#)).

Naively, one might assume that the class without enforced vertex-transitivity is larger since we impose fewer symmetry constraints. To our surprise, this intuition is misleading. In fact, in [Chapter 6](#) we prove that all edge-transitive polytopes in four or more dimensions are also vertex-transitive. The proof is quite technical and requires a full chapter. The focus of this section is then on the remaining polytopes of combined vertex- and edge-transitivity.

Under the consideration that almost all edge-transitive polytopes are vertex-transitive, our approach to higher-dimensional edge-transitivity via spectral polytopes (specifically [Corollary 3.15](#)) appears now to provide a much larger coverage of this class.

It is worthwhile to recollect what our achievements of Part I can tell us about the simultaneously vertex- and edge-transitive polytopes:

Theorem 4.5. *If P is vertex- and edge-transitive (aka. $\{0, 1\}$ -transitive), then*

- (i) P is θ_2 -spectral,
- (ii) P is uniquely determined by its edge-graph (up to scale and orientation),
- (iii) $\text{Aut}(P)$ is irreducible,
- (iv) P realizes all the symmetries of its edge-graph, in particular, $\text{Aut}(P) \cong \text{Aut}(G_P)$,
- (v) if P has edge-length ℓ and circumradius r , then

$$\frac{\ell}{r} = \sqrt{\frac{2\lambda_2}{\deg(G_P)}} = \sqrt{2 - \frac{2\theta_2}{\deg(G_P)}},$$

where θ_2 (resp. λ_2) is the second-largest eigenvalue (resp. second-smallest Laplacian eigenvalue) of the edge-graph G_P ,

- (vi) if the polar dual P° has dihedral angle α , then

$$\cos(\alpha) = -\frac{\theta_2}{\deg(G_P)}.$$

Proof. Part (i) was proven in [Corollary 3.15](#).

From [Corollary 3.8](#) (i) follows that P is uniquely determined up to invertible linear transformation. This already determines the combinatorial type of P . Since P is inscribed (all its vertices lie on a common sphere) and has all edges of the same length, the same holds for its 2-faces, which are then regular polygons. Knowing the combinatorial type of P and the shape of all 2-faces, *Cauchy's rigidity theorem* (in the form of [Corollary C.3](#)) implies that P is determined up to orientation. The possible variation in scale comes from the variation in edge length. This proves (ii).

Suppose that $\mathbb{R}^d = W_1 \oplus W_2$ is a decomposition of the ambient space into $\text{Aut}(P)$ -invariant subspaces. Consider the linear transformation $X := 2\pi_{W_1} + \pi_{W_2} \in \text{GL}(\mathbb{R}^d)$. If W_1 and W_2 are non-trivial invariant subspaces, then X is neither orthogonal nor a pure rescaling. But $\text{Aut}(P)$ commutes with X , hence acts vertex- and edge-transitively on $P' := XP$. But P and P' have

the same edge-graph, contradicting (ii). In conclusion, $\text{Aut}(P)$ must have been irreducible, which proves (iii).

Since $\text{Aut}(P)$ is irreducible, the skeleton of P is an irreducible Σ -realization of G_P for some $\Sigma \subseteq \text{Aut}(G_P)$. The skeleton is then spherical by Lemma 1.12. Hence, P is a spherical θ_2 -spectral polytope, and by Corollary 3.8 (ii) it realizes all the symmetries of its edge-graph. This proves (iv).

Since P is θ -spectral and spherical, its skeleton satisfies the conditions of Corollary 2.21, which proves (v) and (vi) (for the latter, recall that the dihedral angle in the dual is π minus the angle between vertices in P , cf. Example 2.23). \square

4.4 Arc-transitive polytopes

In Section 2.2 we introduced *arc-transitive graphs* as a sub-class of the simultaneously vertex- and edge-transitive graphs. We also briefly touched on the complementary sub-class, the half-transitive graphs (see Remark 2.16). Both symmetries, arc- and half-transitivity, can be defined for polytopes as well. This distinction is the next branching point in Figure 4.3.

This section is devoted to the discussion of arc-transitive polytopes. The half-transitive polytopes will be discussed in Section 4.5.

Definition 4.6. A polytope $P \subset \mathbb{R}^d$ is *arc-transitive* if its symmetry group acts transitively on incident vertex-edge pairs (also known as *arcs* or *1-flags*).

In addition to being vertex- and edge-transitive, in an arc-transitive polytope an edge can also be mapped onto itself with “flipped orientation” (if the edge $e \in \mathcal{F}_1(P)$ has end-vertices $v, w \in \mathcal{F}_0(P)$, then we can map the arc (e, v) onto the arc (e, w)).

One might assume that this edge-flip is performed by a reflectional symmetry of P . But this is not necessarily true. One can imagine other symmetries that perform such a flip, and we shall discuss such further below.

However, all simultaneously vertex- and edge-transitive polytopes encountered so far can indeed realize their edge flips by reflections. This turns out to be in line with the observation that all our examples so far have been Wythoffian:

Lemma 4.7. *If each edge $e \in \mathcal{F}_1(P)$ of a polytope P can be flipped by a reflection $\rho_e \in \text{Aut}(P)$, then P is a Wythoffian polytope.*

Proof. Let $\Gamma := \langle \rho_e \mid e \in \mathcal{F}_1(P) \rangle$ be the matrix group generated by the edge-flipping reflectional symmetries. Since $\Gamma \subseteq \text{Aut}(P)$, this is a *finite* reflection group.

It remains to show that Γ acts transitively on the vertices of P . For this, fix two adjacent vertices $v, w \in \mathcal{F}_0(P)$ and let $e := \text{conv}\{v, w\} \in \mathcal{F}_1(P)$ be the incident edge. The map $\rho_e \in \Gamma$ exchanges the end vertices of e , that is, maps v onto w . Since the edge-graph of P is connected, an appropriate concatenation of such reflections can map any vertex of P onto any other vertex (along a path between v and w). Hence, Γ acts vertex-transitively. \square

Thus, if we want to go beyond Wythoffian polytopes, we need to consider edge-flips that are not reflections. Recall that a reflection is a linear transformation $T \in \text{O}(\mathbb{R}^d)$ with spectrum $\{1^{d-1}, (-1)^1\}$. Consider the following generalizations:

Definition 4.8.

- (i) A k -reflection $T \in O(\mathbb{R}^d)$ is a linear transformation with spectrum $\{1^{d-k}, (-1)^k\}$.
- (ii) A k -reflection group $\Gamma \subseteq O(\mathbb{R}^d)$ is a matrix group generated by k -reflections.

Clearly, the 1-reflections are just the usual reflections, and the 1-reflection groups are the usual reflection groups.

Theorem 4.9. *An arc-transitive polytope P is the orbit polytope of a k -reflection group for some $k \geq 1$.*

Proof. Let $e \in \mathcal{F}_1(P)$ be an edge of P . By arc-transitivity, there exists a symmetry $T_e \in \text{Aut}(P)$ that flips e . Since T_e exchanges the end-vertices of e , the smallest number $m \geq 1$ for which $T_e^m = \text{Id}$ must be even.

The map $T'_e := T_e^{m/2}$ is then a symmetry of P , but is also an involution (i.e., $T_e'^2 = \text{Id}$). As a non-identity involutory isometry, its spectrum is of the form $\{(-1)^k, 1^{d-k}\}$ for some $k \geq 1$, and T'_e is a k -reflection. By edge-transitivity, T'_e can be chosen as a k -reflection with the same value k for every $e \in \mathcal{F}_1(P)$.

As in the proof of [Lemma 4.7](#), we show that the group $\Gamma := \langle T'_e \mid e \in \mathcal{F}_1(P) \rangle$, which is a finite k -reflection group, acts transitively on the vertices of P . Thus, P is an orbit polytope of Γ . □

A sufficient understanding of k -reflection groups might therefore help in classifying the arc-transitive polytopes¹. Unfortunately, we have not yet constructed any new arc-transitive polytopes using generalized reflection groups. The question is then whether this is possible at all.

Question 4.10. Are there non-Wythoffian arc-transitive polytopes?

Our lack of examples (despite quite some effort²) makes us believe that the answer is *No*. A careful examination of the known Wythoffian examples and their Coxeter-Dynkin diagrams led us to the following conjecture:

Conjecture 4.11. *The following are equivalent:*

- (i) P is arc-transitive.
- (ii) P is a Wythoffian uniform polytope represented by a “transitive Coxeter-Dynkin diagram”, i.e., the symmetry group of the diagram acts transitively on the ringed vertices.

For a reminder on the representation of Wythoffian polytopes by Coxeter-Dynkin diagrams see [Appendix E.2](#). Examples of such diagrams are shown in [Figure 4.4](#) further below.

¹A complete classification of general k -reflection groups seems unlikely: every non-abelian simple group has a faithful representation as a k -reflection group for a sufficiently large k (namely, its left-regular representation). It might be feasible to classify the k -reflection groups for small fixed values of k . For example, the 2- resp. 4-reflection groups are related to (though distinct from) the complex and quaternionic reflection groups.

²We checked (numerically) all arc-transitive graphs on $n \leq 30$ vertices, and also all arc-transitive graphs in the Mathematica graph library on $n \leq 1000$ vertices. For this, we computed the θ_2 -eigenpolytope and compared its edge-graph to the original graph.

The hard part of [Conjecture 4.11](#) is $(i) \implies (ii)$. In contrast, $(ii) \implies (i)$ seems tractable (but a proof is not part of this thesis). If true, [Conjecture 4.11](#) enables a systematic enumeration of arc-transitive polytopes. For the rest of this section we shall take a look at this potential classification.

A potential classification of arc-transitive polytopes

Suppose that [Conjecture 4.11](#) is true. We explore the potential classification of arc-transitive polytopes that would follow from that.

Recall that the (irreducible) finite reflection groups admit a surprisingly compact classification into the four infinite families $I_2(n)$, A_d , B_d and D_d ($n, d \geq 3$), and the seven exceptional groups I_1 , H_3 , H_4 , F_4 , E_6 , E_7 and E_8 (the subscripts denote dimension). The classification of Wythoffian uniform polytopes follows from this by a purely combinatorial enumeration of the possible placements of rings in their Coxeter-Dynkin diagrams. There are some subtleties because different diagrams can describe the same polytope. We discuss this below (without any further proofs).

We shall take a look at the placements of rings that give rise to transitive diagrams.

Observation 4.12. A Coxeter-Dynkin diagram describes a *full-dimensional* polytope if and only if each connected component contains a ringed vertex (see [Observation E.3](#)). Thus, if a transitive diagram is *not* connected, then all its components must be identical. We can therefore focus on a single component.

Such a component is certainly transitive if one of the following holds:

- (i) it has only a single ringed vertex.
- (ii) it is one of the Coxeter-Dynkin diagrams depicted in [Figure 4.4](#).




Figure 4.4. Transitive Coxeter-Dynkin diagrams with more than one ringed vertex.

There are other transitive diagrams besides the ones depicted in [Figure 4.4](#). However, it turns out that the other diagrams do not describe any new polytopes: any other transitive diagram generates a polytope that is also generated by a diagram with a single ring or a diagram listed in [Figure 4.4](#).

Example 4.13. Each regular polytope is represented by a Coxeter-Dynkin diagram which is a path graph with a single ring on one of its end vertices (cf. [Definition E.11](#)). By shifting the ring by one vertex we obtain the *rectified* regular polytopes discussed in [Example 4.3](#). More general configurations with a single ring generalize on this construction.

For example, the polytopes of the A_d -family obtained from a diagram with a single ringed vertex are called *hyper-simplices*. Geometrically, the hyper-simplex $\Delta_{d,k}$ can be constructed as the convex hull of all 01-vectors in $\{0, 1\}^{d+1}$ with exactly k many 1-entries.

Example 4.14. If P is represented by a transitive Coxeter-Dynkin diagram, so are its cartesian powers. The diagram of P^k is the disjoint union of k copies of the diagram of P . Thus **Example 4.4** works not just with regular polytopes, but with any arc-transitive Wythoffian uniform polytope.

Polytopes that can be obtained as cartesian products are called *prismatic*. We shall exclude them from the further discussion as their classification follows immediately from the classification of the *non-prismatic* polytopes. One exception to this rule are the hypercubes, which are technically prismatic, but will still be considered.

As a general rule of thumb, by **Conjecture 4.11**, each irreducible finite reflection group in dimension d should be expected to yield exactly d unique arc-transitive polytopes. This is most plausible if the Coxeter-Dynkin diagram of the group has no symmetries: the d distinct polytopes result from the d ways to place a single ring. If there are symmetries, then there are two opposing effects: on the one hand, there are now symmetric configurations with multiple rings resulting in new polytopes. On the other hand, several of the placements with a single ring now represent the same diagram up to symmetry. It turns out that these two effects cancel each other out and we still get d polytopes (consider the example in **Figure 4.5**).




Figure 4.5. Left: the four transitive diagrams for B_4 (the diagram has no symmetries because of the asymmetric placement of the edge-label “4”). Right: the four transitive diagrams for A_4 (the diagram has symmetries).

By this reasoning, if there are $\kappa(d)$ irreducible reflection groups in dimension d , we would expect to find exactly $d\kappa(d)$ (non-prismatic) arc-transitive polytopes in dimension d . However, this reasoning fails because distinct diagrams can describe the same polytope. These “coincidences” are manageable:

- All but one polytope of the D_d -family are also generated by the group $B_d \supset D_d$. Thus, D_d provides only one polytope rather than d (further exceptions happen for $d \in \{3, 4\}$, see below). This unique polytope is called the *demi-cube* and can be obtained as the convex hull of all vectors in $\{-1, 1\}^d$ with an even number of 1-entries (half of the vertices of the cube, therefore the name).
- The version of $I_2(n)$ with two rings yields the regular $2n$ -gon, which is also generated by $I_2(2n)$. Thus, $I_2(n)$ generates only a single unique polytope rather than two.
- In dimension three holds $A_3 \cong D_3$ (one can check that their Coxeter-Dynkin diagrams coincide). Thus, the pair A_3/D_3 accounts only for a single polytope (the tetrahedron, which is the demi-cube of this dimension).

- In dimension four two exceptions occur: one of the B_4 -polytopes coincides with an F_4 -polytope (the rectified 4-crosspolytope coincides with the 24-cell) and the 4-dimensional demi-cube from D_4 coincides with the 4-crosspolytope from B_4 . Thus, D_4 provides no unique polytope and B_4 yields only three rather than four unique polytopes.

Considering all these exceptions we arrive at the following conjectured number of (non-prismatic) arc-transitive polytopes by dimension:

d	1	2	3	4	5	6	7	8	≥ 9
#	1	∞	7	15	11	19	22	25	$2d + 1$

In the dimensions “without exceptions”, that is, $d \geq 9$ and $d = 5$, we obtain exactly $2d + 1$ irreducible arc-transitive polytopes (d from A_d , d from B_d and one from D_d). In dimension $d \in \{6, 7, 8\}$ there are $3d + 1$ polytopes (d additional polytopes from the group E_d).

For $d \in \{3, 4\}$ various exceptions occur. The number seven for $d = 3$ agrees with the number given in Section 4.2 (excluding the two polyhedra that are not vertex-transitive). For $d = 4$ we obtain the following 15 (non-prismatic) arc-transitive polytopes:

- the six 4-dimensional regular polytopes,
- five rectifications of regular polytopes (the rectification of the 4-crosspolytope was already counted as the 24-cell, cf. Example 4.3),
- two “bitruncations” (of simplex and 24-cell), and
- two “runcinations” (of simplex and 24-cell).

The terminology for the various modifications of the regular polytopes goes back to Norman Johnson [40]. Their Coxeter-Dynkin diagrams are depicted in Figure 4.6.

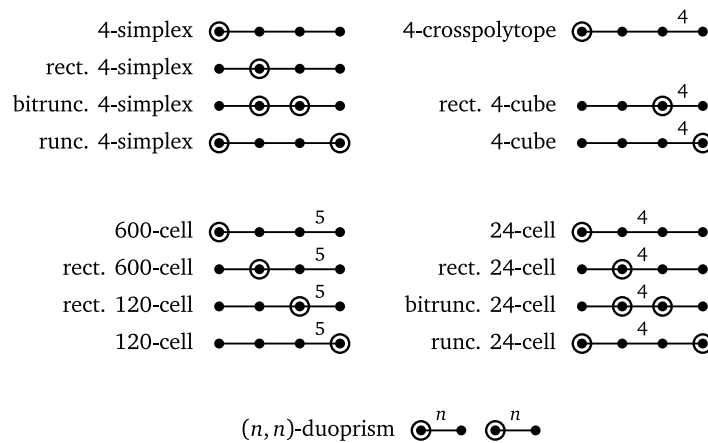


Figure 4.6. The Coxeter-Dynkin diagrams of the (conjectured) 4-dimensional non-prismatic arc-transitive Wythoffian uniform polytopes (and the duoprisms).

4.5 Half-transitive polytopes

A *half-transitive polytope* is a polytope that is vertex- and edge-transitive, but that is not arc-transitive. That is, its edges cannot be flipped in orientation.

By [Theorem 4.5](#) (iv), the edge-graph of a half-transitive polytope must be half-transitive. Since half-transitive graphs are already quite rare, the existence of half-transitive polytopes seems rather unlikely. We make a bold claim:

Conjecture 4.15. *There are no half-transitive polytopes.*

Example 4.16. Recall the *Holt graph* (see [Figure 4.7](#)), the smallest half-transitive graph. We show that it is *not* the edge-graph of a half-transitive polytope.




Figure 4.7. The Holt graph – the smallest half-transitive graph.

Suppose that $P \subset \mathbb{R}^d$ is a half-transitive polytope and has the Holt graph as its edge-graph. P is then θ_2 -spectral by [Theorem 4.5](#) (i). In particular, its dimension matches the multiplicity of θ_2 , which is six. The edge-graph of a polytope in dimension six has a degree of at least six. But the vertex degree of the Holt graph is only *four*.

Observation 4.17. The argument of [Example 4.16](#) can be applied more generally. In order for a graph G to be the edge-graph of a simultaneously vertex- and edge-transitive polytopes, the multiplicity of θ_2 must be at most $\deg(G)$.

Half-transitive graphs G with $\dim \text{Eig}_G(\theta) \leq \deg(G)$ exist, but are even rarer. Numerical experiments on such candidates have not resulted in any half-transitive polytopes.

4.6 Distance-transitive polytopes

A polytope is distance-transitive if its symmetry group acts distance-transitively on its edge-graph (cf. [Definition 2.31](#)). Note that “distance” in the definition of distance-transitivity still refers to the graph-theoretical distance along the edge-graph rather than a Euclidean distance between vertices.

Godsil [[32](#)] classified all distance-regular graphs that appear as the edge-graphs of their θ_2 -eigenpolytopes (see [Theorem 3.3](#)). We noted that all graphs in this list are in fact distance-

transitive. By [Corollary 3.8 \(ii\)](#), their eigenpolytopes are therefore distance-transitive polytopes. Conversely, by [Theorem 4.5 \(ii\)](#), every distance-transitive polytope can be obtained in this way. In other words, we obtain a complete classification of distance-transitive polytopes:

Theorem 4.18. *If $P \subset \mathbb{R}^d$ is distance-transitive then P is one of the following:*

- (i) a regular polygon ($d = 2$),
- (ii) the regular dodecahedron ($d = 3$),
- (iii) the regular icosahedron ($d = 3$),
- (iv) the d -dimensional crosspolytope,
- (v) a Hamming polytope (aka. a cartesian power of a regular simplex; this includes d -simplices and d -cubes, see [Example 4.4](#)),
- (vi) a hyper-simplex $\Delta_{d,k}$ (see [Example 4.13](#)),
- (vii) a d -dimensional demi-cube (the unique D_d -polytope, see [Example 4.14](#)),
- (viii) the 2_{21} -polytope (also known as Schläfli polytope, $d = 6$),
- (ix) the 3_{21} -polytope (also known as Gosset polytope, $d = 7$).

The ordering of the polytopes in this list agrees with the ordering of graphs in [Theorem 3.3](#).

The latter two polytopes (viii) and (ix) were first constructed by Gosset in [34]. Distance-transitive polytopes are arc-transitive, and in fact, all of the polytopes in [Theorem 4.18](#) are Wythoffian uniform polytopes in line with [Conjecture 4.11](#) (see [Figure 4.8](#) for their Coxeter-Dynkin diagrams)

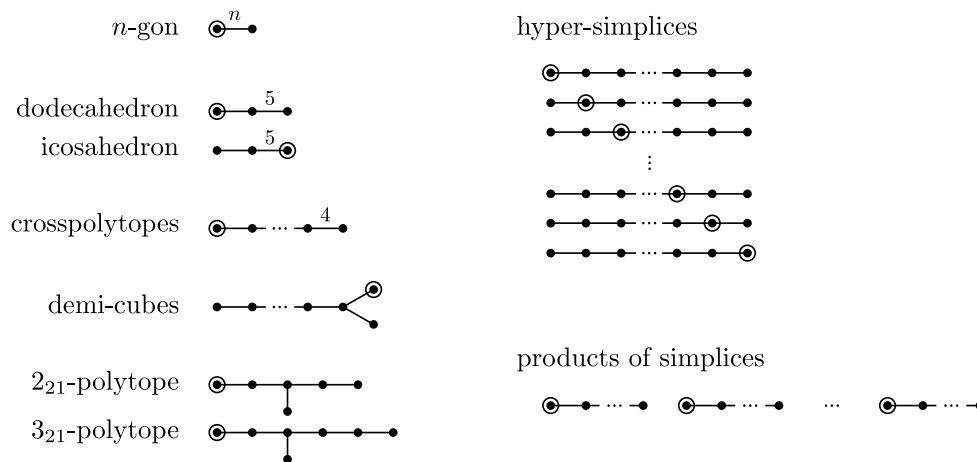


Figure 4.8. The Coxeter-Dynkin diagrams of the distance-transitive polytopes.

[Theorem 4.18](#) lists all the regular polytopes, excluding the 4-dimensional exceptions: the 24-cell, 120-cell and 600-cell are *not* distance-transitive. The distance-transitive polytopes therefore form a distinct class of highly symmetric polytopes that is not immediately related to the class of regular polytopes.

Finally, the 4-dimensional *Petersen polytope* as studied by Powers in [62] (θ_3 -eigenpolytope of the Petersen graph, see also Section 3.1) is contained in this list as the hyper-simplex $\Delta_{4,2}$.

Summary


In this chapter we presented a hierarchical classification scheme for edge-transitive polytopes that is meant to help in the process of classification. As of now, it seems that essentially only a single branch of the hierarchy is populated:

- Almost all edge-transitive polytopes are vertex-transitive. A proof was postponed until Chapter 6.
- All known polytopes of combined vertex- and edge-transitivity are arc-transitive. We do not know any half-transitive polytopes.
- All known arc-transitive polytopes are Wythoffian. We made a precise conjecture about the shape of the Coxeter-Dynkin diagrams of arc-transitive Wythoffian polytopes (Conjecture 4.11).

Even though the edge-transitive polytopes feel “closer to” the vertex-transitive polytopes than to the regular polytopes (in terms of the restrictions on structure), our current picture of the situation seems to suggest that they behave more like the latter. It might well be that the class of edge-transitive polytopes has a quite tractable classification in terms of Wythoffian polytopes and a few lower dimensional exceptions that are not vertex-transitive.

Conjecture 4.19. *All edge-transitive polytopes in dimension $d \geq 4$ are Wythoffian uniform polytopes.*

5 Vertex-Transitive Zonotopes



For this chapter we take a pause from our investigation of edge-transitivity and instead consider *vertex-transitivity* for a special class of polytopes – so-called *zonotopes*.

The purpose of this investigation in the larger picture of this thesis is its application in the upcoming [Chapter 6](#). In [Theorem 6.10](#) we need to use that certain geometrically defined polytopes are always vertex-transitive, which is one of the main results of this section.

The literature defines zonotopes in multiple equivalent ways, depending on the subfield of geometry, combinatorics or algebra from which they are approached. We present the most relevant definitions for our cause (other definitions are given in [Appendix C.1](#))

Definition 5.1 (cf. [Definition C.4](#)). A zonotope $Z \subseteq \mathbb{R}^d$ is a polytope that satisfies any (and then all) of the following equivalent conditions:

- (i) Z is the Minkowski sum of (finitely many) line segments.
- (ii) Z has only centrally-symmetric faces.
- (iii) Z has only centrally-symmetric 2-faces.

The equivalence of these definitions is well-established, but some directions are far from obvious (see [\[54\]](#) for the direction (iii) \implies (i), (ii), or the references in [\[79, Section 7.3\]](#) for a general overview).




Figure 5.1. Some examples of 3-dimensional zonotopes.

The goal of this chapter is to obtain a complete classification of *vertex-transitive* zonotopes. That this is possible in a concise manner is quite surprising: both classes, the zonotopes and the vertex-transitive polytopes, are intractably rich in structure. It was therefore unexpected to obtain the following result:

Theorem 5.2. *A vertex-transitive zonotope is a Γ -permutahedron, where $\Gamma \subseteq O(\mathbb{R}^d)$ is a finite reflection group.*

A permutahedron is a *generic* Wythoffian polytope (more precisely defined in [Definition 5.9](#), or [Definition E.6](#) in the appendix). Permutahedra are sufficiently well understood, so that a complete classification can be derived.

On our way we shall furthermore prove the following:

Theorem 5.3. *If a zonotope is inscribed (that is, all its vertices are on a common sphere) and all its edges are of the same length, then it is a (uniform) Γ -permutahedron.*

Note that [Theorem 5.3](#) is neither stronger nor weaker than [Theorem 5.2](#): it is not clear a priori that all vertex-transitive zonotopes have edges of the same length (in fact, this is not true), nor that the zonotopes addressed by [Theorem 5.3](#) are vertex-transitive (this is true, but will only follow from our results).

Chapter overview

In [Section 5.1](#) we introduce the relevant terminology and discuss some fundamental results concerning zonotopes and Γ -permutahedra. In particular, we discuss the *generators* of a zonotope, explain how they determine its faces and how they characterize permutahedra via root systems.

The proofs of the main theorems ([Theorem 5.2](#) and [Theorem 5.3](#)) are presented over the course of [Section 5.2](#) and [Section 5.3](#). In [Section 5.2](#) we deal with the 2-dimensional case and explain how this serves as a base case for higher dimensions. In [Section 5.3](#) we first prove [Theorem 5.3](#) (which is surprisingly easy) and then use this to prove [Theorem 5.2](#). In [Section 5.4](#) we then derive an explicit classification of the relevant zonotope classes from the classification of finite reflection groups.

We close in [Section 5.5](#) with some miscellaneous results and notes, such as a new characterization of root systems, and some comments on general *inscribed zonotopes* and *hyperplane arrangements*.

5.1 Generators, faces, symmetry and permutahedra

In this section we recollect the most important properties and terminology for zonotopes and introduce relevant sub-classes. Consider also [Appendix C.1](#) for proofs of some of the less obvious claims.

For this chapter let $Z \subset \mathbb{R}^d$ be a full-dimensional zonotope of dimension $d \geq 2$.

Generators

Zonotopes are always centrally symmetric, and we may assume $Z = -Z$. By [Definition 5.1](#) (ii) Z can then be written as a Minkowski sum of line segments, that is,

$$Z = \text{Zon}(R) := \sum_{r \in R} \text{conv}\{0, r\},$$

for some finite centrally symmetric set $R \subset \mathbb{R}^d$. We set $\text{Zon}(\emptyset) := \{0\}$. We say that Z is generated by R . A zonotope $Z \neq \{0\}$ can be generated by many different sets.

We say that $R \subset \mathbb{R}^d$ is *reduced* if $R \cap \text{span}\{r\} = \{\pm r\}$ for all $r \in R$. Z is generated by a unique reduced set, the elements of which are called (*standard*) *generators* of Z (cf. [Proposition C.7](#)). This set is explicitly given by

$$\text{Gen}(Z) := \{r \in \mathbb{R}^d \mid \text{conv}\{\pm r\} \text{ is the translate of an edge of } Z\}.$$

Then $\text{Zon}(\text{Gen}(Z)) = Z$. Moreover, if $R \subset \mathbb{R}^d$ is finite, centrally symmetric and reduced, then $\text{Gen}(\text{Zon}(R)) = R$.

Faces

The faces of Z are again zonotopes (cf. [Corollary C.6](#)), that is, they can be written in terms of generators. In fact, they can be written in terms of generators of Z . We need the following:

Definition 5.4. Let $R \subset \mathbb{R}^d$ be a finite centrally-symmetric set:

- (i) a subset $S \subset R$ is called *semi-star* of R if it is the intersection of R with a half-space that contains exactly half the elements of R .

In particular, S contains exactly one element from each subset $\{\pm r\} \subseteq R$.

- (ii) a subset $F \subseteq R$ is called a *flat* of R if it is the intersection of R with a linear subspace, or equivalent, if $F = R \cap \text{span } F$.




Figure 5.2. Visualization of a semi-star $S \subset R$ and a flat $F \subseteq R$.

Lemma 5.5 (see also [Lemma C.10](#)). For $F \subseteq \text{Gen}(Z)$ the following are equivalent:

- (i) F is a flat,
- (ii) $F = \text{Gen}(f)$ for some non-empty face $f \in \mathcal{F}(Z)$.

Lemma 5.6 (see also [Lemma C.11](#)). The vertices of Z are in one-to-one correspondence with the semi-stars of $\text{Gen}(Z)$: for each semi-star $S \subset \text{Gen}(Z)$

$$v_S := \sum_{r \in S} r \in \mathcal{F}_0(P)$$

is a vertex of Z . Conversely, for $v \in \mathcal{F}_0(P)$ there is a unique semi-star $S_v \subset \text{Gen}(Z)$ with $v = v_{S_v}$.

As indicated, proofs can be found in [Appendix C.1](#).

Symmetries

The operations Zon and Gen commute with multiplication by invertible matrices, *i.e.*, $\text{Zon}(XR) = X \text{Zon}(R)$ and $\text{Gen}(XZ) = X \text{Gen}(Z)$ for all $X \in \text{GL}(\mathbb{R}^d)$. As a consequence, the zonotope and its generators have the same Euclidean symmetries (orthogonal transformations that fix them set-wise):

Proposition 5.7. $\text{Aut}(Z) = \text{Aut}(\text{Gen}(Z))$.

For later use we characterize vertex-transitivity in terms of generators. Recall that being *congruent* (or *isometric*) means “related by an orthogonal transformation”.

Proposition 5.8. Z is vertex-transitive if and only if all semi-stars of $\text{Gen}(Z)$ are congruent.

Proof. By [Lemma 5.6](#) there is a one-to-one relation between the vertices of Z and the semi-stars of $\text{Gen}(Z)$. Note that the given map $\text{Gen}(Z) \ni S \mapsto v_S$ commutes with multiplication by invertible matrices, *i.e.*, $Xv_S = v_{XS}$ for all $X \in \text{GL}(\mathbb{R}^d)$.

Suppose that Z is vertex-transitive and fix two semi-stars $S, S' \subset \text{Gen}(Z)$. Then there is a map $X \in \text{Aut}(Z)$ with $Xv_S = v_{S'}$. By [Proposition 5.7](#) we have $X \in \text{Aut}(\text{Gen}(Z))$, and XS is a semi-star as well. Thus, $v_{XS} = Xv_S = v_{S'}$, and since the relation is one-to-one, $XS = S'$ and the semi-stars are congruent.

Conversely, suppose that all semi-stars are congruent and fix two vertices $v, v' \in \mathcal{F}_0(P)$. Since $S_v, S_{v'} \subset \text{Gen}(Z)$ (as defined in [Lemma 5.6](#)) are congruent, there is a map $X \in \text{O}(\mathbb{R}^d)$ with $XS_v = S_{v'}$. Then $v' = v_{S_{v'}} = v_{XS_v} = Xv_S = Xv$. It remains to show $X \in \text{Aut}(Z)$. By [Proposition 5.7](#) we can show $X \in \text{Aut}(\text{Gen}(Z))$ instead. Since $\text{Gen}(Z)$ is finite and X is invertible, it suffices to show $Xr \in \text{Gen}(Z)$ for all $r \in \text{Gen}(Z)$. Fix a generator $r \in \text{Gen}(Z)$. Then either $r \in S_v$ and $Xr \in XS_v = S_{v'} \subset \text{Gen}(Z)$. Or $r \notin S_v$, but by definition of semi-star $-r \in S_v \implies Xr = -X(-r) \in -S_{v'} \subset \text{Gen}(Z)$. \square

Permutahedra and root systems

Our main results require the notion of the Γ -permutahedron. A permutahedron is a “generic” Wythoffian polytope (see also [Definition E.6](#)).

Definition 5.9. Let $\Gamma \subseteq \text{O}(\mathbb{R}^d)$ be a finite reflection group. A Γ -permutahedron P is a polytope that satisfies any (and then all) of the following equivalent conditions:

- (i) P is an orbit polytope $\text{Orb}(\Gamma, x)$ of a generic point $x \in \mathbb{R}^d$, that is, x is not fixed by any non-identity element of Γ .
- (ii) $\Gamma \subseteq \text{Aut}(P)$ and Γ acts regularly (*i.e.*, transitively and freely) on the vertices of P .

Most permutahedra are not zonotopes, but each permutahedron has one or more realizations as a zonotope. This includes the unique realization in which all edges are of the same length (the *uniform* permutahedron) (see [Corollary 5.11](#) below).

The zonotopes among the Γ -permutahedra (we shall call them the Γ -zonotopes) can be characterized in terms of so-called *root systems*. We recommend [Appendix D.2](#) for a reminder on root systems and their connection to reflection groups. In short, a *root system* $R \subset \mathbb{R}^d \setminus \{0\}$




Figure 5.3. The 3-dimensional uniform permutahedra to the reflection groups (from left to right) A_3 , B_3 and H_3 .

is a finite set of non-zero vectors with $\rho_r R = R$ for all $r \in R$, where $\rho_r \in O(\mathbb{R}^d)$ denotes the reflection on the hyperplane r^\perp . In particular, root systems are centrally symmetric.

Each centrally symmetric set $R \subset \mathbb{R}^d \setminus \{0\}$ induces a reflection group

$$\Gamma(R) := \langle \rho_r \mid r \in R \rangle,$$

the so-called *Weyl group* of R . Root systems are characterized as those centrally symmetric sets for which $\Gamma(R)$ is a *finite* reflection group (this is non-trivial but well-known).

Lemma 5.10. *Given a zonotope $Z \subset \mathbb{R}^d$ and a finite reflection group $\Gamma \subseteq O(\mathbb{R}^d)$, the following are equivalent:*

- (i) Z is a Γ -permutahedron,
- (ii) $\text{Gen}(Z)$ is a root system with Weyl group Γ .

The proof of **Lemma 5.10** requires the concept of the *Weyl chamber* of Γ . For this, consider the hyperplane arrangement

$$\mathcal{H}(\Gamma) := \{r^\perp \mid \rho_r \in \Gamma\}.$$

The *Weyl chambers* are the connected components of the complement $\mathbb{R}^d \setminus \mathcal{H}(\Gamma)$.

It is known that Γ acts regularly on its Weyl chambers. In a sense, the Weyl chambers contain the “generic points” of Γ . In particular, $\text{Orb}(\Gamma, x)$ is a Γ -permutahedron if and only if x is an element of a Weyl chamber. We conclude that a Γ -permutahedron has exactly one vertex per Weyl chamber.

Proof of Lemma 5.10. Suppose that Z is a Γ -permutahedron. We show that $\Gamma(\text{Gen}(Z)) = \Gamma$. Since then $\Gamma(\text{Gen}(Z))$ is finite, $\text{Gen}(Z)$ must be a root system.

Let $\rho_r \in \Gamma$ be some generating reflection of Γ . Since Z is a “generic” orbit polytope of Γ , no vertex of Z is fixed by ρ_r , that is, no vertex lies on r^\perp . But since r^\perp is a reflection hyperplane of Z , there are vertices on both sides of r^\perp , and there must be an edge $e \in \mathcal{F}_1(Z)$ crossing the hyperplane. Since ρ_r is a symmetry of Z , this edge e must be perpendicular to r^\perp , that is, $e = \text{conv}\{\pm ar\}$ for some $\alpha > 0$. Then $\alpha r \in \text{Gen}(Z)$ and $\rho_r \in \Gamma(\text{Gen}(Z))$. We established $\Gamma \subseteq \Gamma(\text{Gen}(Z))$.

For the other inclusion let $\rho_r \in O(\mathbb{R}^d)$ be some generating reflection of $\Gamma(\text{Gen}(Z))$, i.e., $\text{conv}\{\pm r\}$ is the translate of an r -parallel edge $e \in \mathcal{F}_1(P)$. Since there is a single vertex of Z per Weyl chamber, the end vertices of e are in different Weyl chambers and e crosses a

reflection hyperplane. Since reflection on this hyperplane is a symmetry of Z , we find that e is perpendicular to the hyperplane, which must then be r^\perp . Thus, $\rho_r \in \Gamma$ as well and we found $\Gamma(\text{Gen}(Z)) \subseteq \Gamma$.

We then arrive at $\Gamma = \Gamma(\text{Gen}(Z))$ and proved (i) \implies (ii).

For the other direction, assume that $\text{Gen}(Z)$ is a root system and set $\Gamma := \Gamma(\text{Gen}(Z)) \subseteq \text{Aut}(\text{Gen}(Z))$. By [Proposition 5.7](#) we have $\Gamma \subseteq \text{Aut}(Z)$. It remains to show that Γ acts regularly on the vertices of Z , to show that Z is a Γ -permutahedron in the sense of [Definition 5.9](#) (ii).

Showing that Γ acts transitively proceeds similarly to the proof of transitivity in [Lemma 4.7](#): fix adjacent vertices $v, w \in \mathcal{F}_0(Z)$ and let $e := \text{conv}\{v, w\} \in \mathcal{F}_1(Z)$ be the incident edge. Then e is a translate of $\text{conv}\{\pm r\}$ with $r := (v - w)/2$. Thus $r \in \text{Gen}(Z)$ and $\rho_r \in \Gamma \subseteq \text{Aut}(Z)$. In particular, ρ_r exchanges the end vertices of e and maps v to w . Since the edge-graph of Z is connected, an appropriate chain of such reflections can map any vertex to any other.

To show that Γ acts regularly, recall the one-to-one correspondence between vertices and semi-stars described in [Lemma 5.6](#). Suppose that $v \in \mathcal{F}_0(P)$ is a vertex fixed by a non-identity element of Γ and let S_v be the associated semi-star. Then v is not in a Weyl chamber of Γ , in particular, v is on a reflection hyperplane $r^\perp, r \in \text{Gen}(Z)$. That is, $\rho_r \in \Gamma$ fixes v . Since exactly one of $\pm r$ is in S_v , the other one must be in $\rho_r S_v$, and $S_v \neq \rho_r S_v$. But by linearity $v_{\rho_r S_v} = \rho_r v_{S_v} = \rho_r v = v$, contradicting uniqueness of the semi-star associated with v .

This finalizes the proof of (ii) \implies (i) □

Corollary 5.11. *For each finite reflection group $\Gamma \subseteq \text{O}(\mathbb{R}^d)$ there is a unique Γ -permutahedron (up to scale and orientation) with all edges of the same length. This Γ -permutahedron is a zonotope.*

Proof. Consider the zonotope generated from the root system $R := \{r \in S^{d-1} \mid \rho_r \in \Gamma\}$. □

5.2 The case $d = 2$

The proof of the main theorems starts with the case of dimension two. Showing that the 2-dimensional vertex-transitive zonotopes are permutahedra is rather straightforward: a 2-dimensional zonotope is a centrally symmetric $2n$ -gon, which is vertex-transitive if

- (i) it is a regular $2n$ -gon, or
- (ii) n is even and has alternating edge lengths as seen in [Figure 5.4](#).

This list is complete: every vertex-transitive polygon is an orbit polytope to a dihedral group. The list contains all these orbit polytopes that are centrally symmetric.

If Z is a $2n$ -gon as listed in (i) or (ii), $\text{Gen}(Z)$ consists of $2n$ vectors in \mathbb{R}^2 , equally spaced by an angle of π/n . In the case (ii) these vectors alternate in length (cf. [Figure 5.4](#)). These are exactly the root systems that corresponds to the reflection groups $I_2(n)$ (if $n \geq 3$) and $I_1 \oplus I_1$ (if $n = 2$). Applying [Lemma 5.10](#), we obtain the characterization in dimension two.

Corollary 5.12. *A 2-dimensional vertex-transitive zonotope is a Γ -permutahedron.*




Figure 5.4. Some 2-dimensional vertex-transitive zonotopes and their generators.

Now, the true main result of this section is the following tool for extending this characterization to general dimension:

Theorem 5.13. *If all 2-faces of Z are vertex-transitive, then Z is a Γ -permutahedron.*

Proof. Choose generators $r, s \in \text{Gen}(Z)$. We show $\rho_{r,s} \in \text{Gen}(Z)$, establishing that $\text{Gen}(Z)$ is a root system. Z is then a Γ -permutahedron by [Lemma 5.10](#).

The case $r = \pm s$ is trivial. We therefore assume that

$$R := \text{Gen}(Z) \cap \text{span}\{r, s\}$$

is 2-dimensional. In particular, $R \subseteq \text{Gen}(Z)$ is a 2-dimensional flat. By [Lemma 5.5](#) there exist a 2-face $f \in \mathcal{F}_2(Z)$ with $\text{Gen}(f) = R$. By assumption, f is vertex-transitive and R therefore a root system (by [Corollary 5.12](#) and [Lemma 5.10](#)). In conclusion, $\rho_{r,s} \in R \subseteq \text{Gen}(Z)$. \square

To apply [Theorem 5.13](#) we would need to know that certain 2-faces of Z are vertex-transitive, and this is often not evident from the situation (we start from a vertex-transitive zonotope, but know nothing about its faces). We need to use the following auxiliary result:

Proposition 5.14. *If $Z \subset \mathbb{R}^2$ is a 2-dimensional zonotope which*

- (i) *is inscribed (i.e., has all vertices on a common circle), and*
- (ii) *has the same edge directions as a regular $2n$ -gon,*

then Z is vertex-transitive.

This statement is elementary. We sketch its proof:

Proof of Proposition 5.14. A centrally symmetric polygon has exactly twice as many edges as edge-directions (the edges come in parallel pairs). By (ii), Z and a regular $2n$ -gon have the same number of edge directions, and thus Z must be a $2n$ -gon as well.

Let $\alpha_i \in \mathbb{R}$ be the exterior angle at the i -th vertex of Z (see [Figure 5.5](#)). By (ii) we have $\alpha_i = k_i \pi / n$, where $k_i \in \mathbb{N}$ is an integer ≥ 1 . The exterior angles of a (convex) polygon add up to 2π , and so we estimate

$$2\pi = \sum_{i=1}^{2n} \alpha_i = \sum_{i=1}^{2n} \frac{k_i \pi}{n} \geq 2n \cdot \frac{1 \cdot \pi}{n} = 2\pi.$$




Figure 5.5. The exterior angles of a (convex) polygon always add up to 2π .

We conclude $k_i = 1$ and all exterior angles are equal. It follows that also all *interior* angles are equal.

For simplicity, assume that Z , and all the polygons mentioned below, are of circumradius one. Let ℓ be the length of the shortest edge of Z , which cannot be longer than an edge of a regular $2n$ -gon. There exists a vertex-transitive $2n$ -gon with an edge of this length: imagine a continuous transition from a regular $2n$ -gon to a regular n -gon by shortening every second edge. At some point the shortening edges attain length ℓ .



But an inscribed polygon with prescribed identical interior angles at every vertex is already uniquely determined by placing a single edge: the placement of the incident edges follows uniquely from the given restrictions, which, when applied repeatedly, determines the whole polygon. Therefore, Z must be this vertex-transitive polygon. \square

5.3 The general case

Definition 5.15. A zonotope is called *uniform*, if it is inscribed (*i.e.*, all its vertices are on a common sphere) and all its edges are of the same length.

It is surprisingly easy to show that uniform zonotopes are indeed Γ -permutahedra, which is one of our main results:

Proof of Theorem 5.3. Faces of zonotopes are zonotopes, and faces of inscribed polytopes are inscribed. It follows easily that the faces of uniform zonotopes are uniform zonotopes again. This holds in particular for the 2-faces.

An inscribed polygon with all edges of the same length must be a regular polygon. It follows that the 2-faces of a uniform zonotope Z are regular $2n$ -gons, in particular, vertex-transitive. By Theorem 5.13, Z is a Γ -permutahedron. \square

Note that the identifier “uniform” was chosen appropriately. By Theorem 5.3 uniform zonotopes belong to the Wythoffian *uniform* polytopes (they are Γ -permutahedra for which all edges are of the same length).

The technique of the proof of [Theorem 5.3](#) cannot be immediately applied to vertex-transitive zonotopes, as their 2-faces are not necessarily regular. We need the following construction:

Definition 5.16. The *normalization* of Z is the zonotope

$$Z^* := \text{Zon} \left(\left\{ \frac{r}{\|r\|} \mid r \in \text{Gen}(Z) \right\} \right).$$

The normalization has the same edge directions as Z , but all edges are of the same length. This suffices to prove that a vertex-transitive zonotope Z is a Γ -permutahedron.

Proof of [Theorem 5.2](#). We first show that Z^* is a Γ -permutahedron.

Since Z is vertex-transitive, by [Proposition 5.8](#) all semi-stars of $\text{Gen}(Z)$ are congruent. Normalizing the generators does not change this, and all semi-stars of $\text{Gen}(Z^*)$ are congruent too. By [Proposition 5.8](#), Z^* is then vertex-transitive, in particular, inscribed. As a normalization, it has also all edges of the same length, thus is uniform. By [Theorem 5.3](#), Z^* is a Γ -permutahedron and $\text{Gen}(Z^*)$ therefore a root system by [Lemma 5.10](#).

We now translate this result to Z . For this, choose a 2-face $f \in \mathcal{F}_2(Z)$. The set

$$R^* := \left\{ \frac{r}{\|r\|} \mid r \in \text{Gen}(f) \right\} = \text{Gen}(Z^*) \cap \text{span}(\text{Gen}(f))$$

is a 2-dimensional flat in the root system $\text{Gen}(Z^*)$, and is therefore a root system itself. A 2-dimensional root system consists of vectors that are equally spaced by an angle π/n for some $n \in \mathbb{N}$, or in other words, the elements of R^* are the edge directions of a regular $2n$ -gon. The 2-face f has the edge directions (but not necessarily the edge lengths) contained in R^* , hence, the same edge direction as a regular $2n$ -gon. Also, as the face of a vertex-transitive polytope, f is inscribed. By [Proposition 5.14](#) f is therefore vertex-transitive.

We found that all 2-faces of Z are vertex-transitive. [Theorem 5.13](#) then proves that Z is a Γ -permutahedron. \square

We obtained the following equivalences:

Theorem 5.17. *The following are equivalent:*

- (i) Z is vertex-transitive.
- (ii) Z is a Γ -permutahedron.
- (iii) all semi-stars of $\text{Gen}(Z)$ are congruent.
- (iv) $\text{Gen}(Z)$ is a root system.

Proof. (iii) $\stackrel{5.8}{\iff}$ (i) $\stackrel{5.2}{\iff}$ (ii) $\stackrel{5.10}{\iff}$ (iv), where the numbers over the arrows denote an application of [Proposition 5.8](#), [Theorem 5.2](#) and [Lemma 5.10](#) respectively. \square

Corollary 5.18.

- (i) *The faces of a vertex-transitive zonotope are vertex-transitive.*

- (ii) If a zonotope is inscribed and has all edges of the same length (that is, uniform), then it is vertex-transitive.
- (iii) A zonotope in which all faces (in fact, all 2-faces) are uniform (resp. vertex-transitive) is itself uniform (resp. vertex-transitive).

Proof. If Z is vertex-transitive then $\text{Gen}(Z)$ is a root system by [Theorem 5.17](#). For each face $f \in \mathcal{F}(P)$ its set $\text{Gen}(f)$ of generators is a flat in the root system $\text{Gen}(Z)$ (by [Lemma 5.5](#)). Hence $\text{Gen}(f)$ a root system too. Thus, f is a Γ -permutahedron by [Theorem 5.17](#), and therefore vertex-transitive. This proves (i).

Part (ii) follows immediately from [Theorem 5.3](#). Part (iii) follows from [Theorem 5.13](#) (a uniform 2-face is regular, hence vertex-transitive). \square

Remark 5.19. None of the points of [Corollary 5.18](#) holds for more general polytopes:

- (i) An “elongated anti-prism” (see [Figure 5.6](#)) is vertex-transitive, but most of its faces are not vertex-transitive.
- (ii) The square pyramid (half of an octahedron) is inscribed and has all edges of the same length, but is not vertex-transitive.
- (iii) If we stack the square pyramid on top of a cube, then all faces of the resulting polyhedron are vertex-transitive and uniform, but the polyhedron itself is neither.




Figure 5.6. An “elongated anti-prism” and a “pyramid on top of a cube”.

5.4 The classification

The classification of the vertex-transitive and uniform zonotopes now follows from the classification of finite reflection groups resp. root systems.

We shall focus on the “irreducible” (or non-prismatic) zonotopes – those that result from *irreducible* reflection groups. The “reducible” (or prismatic) zonotopes can be obtained as cartesian products of the irreducible ones. The most prominent reducible examples are probably the prisms (from $I_1 \oplus I_2(n)$) and hypercubes (from $I_1 \oplus \dots \oplus I_1$).

Uniform zonotopes

We obtain one uniform zonotope per finite reflection group. They are classically known as the *omnitruncated uniform polytopes*, see [\[40\]](#). Each is uniquely determined up to scale and orientation. We have the following numbers per dimension:

d	#	groups
2	∞	$I_2(n), n \geq 2$
3	3	$A_3 \cong D_3, B_3, H_3$
4	5	A_4, B_4, D_4, F_4, H_4
5	3	A_5, B_5, D_5
6	4	A_6, B_6, D_6, E_6
7	4	A_7, B_7, D_7, E_7
8	4	A_8, B_8, D_8, E_8
≥ 9	3	A_d, B_d, D_d

Table 5.1. Number of (irreducible) uniform zonotopes per dimension.

Some of these have special names:

- The uniform zonotopes in dimension two are the regular $2n$ -gons.
- The uniform zonotopes of type A_d are also known as the *standard permutahedra*. The A_d -permutahedron can be obtained as the convex hull of the coordinate permutations of the vector $(1, \dots, d+1) \in \mathbb{R}^{d+1}$ (the resulting polytope is contained in a d -dimensional affine subspace). The A_3 -permutahedron is also known as the *truncated octahedron* as it can be obtained from the octahedron by cutting off its vertices (see [Figure 5.3](#)).
- The uniform zonotope of type D_4 is also known as the *truncated 24-cell* (because it can be obtained from the 24-cell by cutting off its vertices).

By definition, a Γ -permutahedron has $|\Gamma|$ vertices. Each Γ -permutahedron is furthermore a simple polytope (*i.e.*, the vertex degree of the edge-graph equals d), and therefore has $d|\Gamma|/2$ edges. Consider [Table 5.2](#) for the precise numbers.

Γ	#vertices	#edges
$I_2(n)$	$2n$	$2n$
A_d	$(d+1)!$	$d(d+1)!/2$
B_d	$d! \cdot 2^d$	$d! \cdot d2^{d-1}$
D_d	$d! \cdot 2^{d-1}$	$d! \cdot d2^{d-2}$
H_3	120	180
F_4	1,152	2,304
H_4	14,400	28,800
E_6	25,920	77,760
E_7	2,903,040	10,160,640
E_8	696,729,600	2,786,918,400

Table 5.2. Vertex and edge count for the Γ -permutahedra.




Figure 5.7. Samples of B_3 -zonotopes and the degenerated limit case, the unique uniform A_3 -zonotope (right). The middle image depicts the *uniform* B_3 -zonotope.

Vertex-transitive zonotopes

Each uniform zonotope is vertex-transitive, but the converse is not true. The zonotopes to the reflection groups $I_2(2n)$, B_d and F_4 are *not* uniquely determined up to scale and orientation, but each case forms a continuous 1-dimensional family of combinatorially equivalent zonotopes (consider the $4n$ -gons in Figure 5.4, or Figure 5.7 for the case B_3). Table 5.1 can still be understood as listing combinatorial types of vertex-transitive zonotopes.

The reason that these cases are different is as follows: typically, the vectors of a root system form a single orbit under the action of the associated Weyl group, except in the cases $I_2(2n)$, B_d and F_4 in which they form two orbits [38, Section 2.11]. The common length of the vectors in each orbit can then be chosen independently from each other, giving each such family of zonotopes one degree of freedom that manifests itself in two potentially different edge lengths.

Such degrees of freedom are also present in all reducible vertex-transitive zonotopes. For example, the d -cubes belong to the family of d -orthotopes with d degrees of freedom.

Example 5.20. The vertices of a general B_d -permutahedron are formed by the coordinate permutations and sign selections of some vector

$$(\pm x_1, \dots, \pm x_d) \in \mathbb{R}^d, \quad \text{with } x_1, \dots, x_d > 0. \quad (5.1)$$

One can show that this results in a zonotope if and only if the x_i form a linear sequence $x_i = x_0 + \epsilon(i - 1)$ for some $x_0, \epsilon > 0$. The quotient ϵ/x_0 parametrizes the 1-dimensional family, and $\epsilon/x_0 = \sqrt{2}$ corresponds to the uniform representative (see Figure 5.7, the polyhedron in the middle).

5.5 Related topics

The results of this chapter touch on several other topics that we now discuss briefly.

Characterizing root systems

Given a root system $R \subset \mathbb{R}^d \setminus \{0\}$, it is well-known that its Weyl group $\Gamma(R)$ acts transitively on the semi-stars (also known as the sets of *positive roots* of R). This is also a consequence of Proposition 5.8 and Lemma 5.10. In particular, all semi-stars of a root system are congruent.

Our results allow to formulate a converse, resulting in a new characterization of root systems:

Theorem 5.21. *If $R \subset \mathbb{R}^d \setminus \{0\}$ is finite and reduced¹, then the following are equivalent:*

- (i) R is a root system,
- (ii) all semi-stars of R are congruent.

Proof. Since R is reduced, it is the set of generator of $\text{Zon}(R)$ and we can apply the equivalence (ii) \Leftrightarrow (iii) from [Theorem 5.17](#). \square

Let the *norm* of a semi-star be the norm of the sum of its elements. It follows from [Theorem 5.21](#) that a root system has all semi-stars of the same norm. We give another criterion by which to recognize a root system:

Theorem 5.22. *Let $R \subset S^{d-1}$ be a finite centrally symmetric set of unit vectors. If all semi-stars of R have the same norm, then R is a root system.*

Proof. By [Lemma 5.6](#) each vertex of $\text{Zon}(R)$ can be written as the sum of the vectors in some semi-star of R . The norm of that semi-star therefore equals the distance of that vertex from the origin. Thus, if all semi-star have the same norm, $\text{Zon}(R)$ is inscribed. Since R is centrally symmetric and all vectors in R are of the same length, all edges of $\text{Zon}(R)$ are of the same length too. Thus, $\text{Zon}(R)$ is uniform and a Γ -permutahedron by [Theorem 5.3](#).

Since R is a centrally symmetric set of unit vectors, it is reduced. Thus, $R = \text{Gen}(\text{Zon}(R))$ is the set of generators of a Γ -permutahedron and therefore a root system by [Lemma 5.10](#). \square

The condition $R \subset S^{d-1}$ (i.e., that R is a set of unit vectors) is necessary for [Theorem 5.22](#): there are inscribed zonotopes with edges of distinct lengths that are not permutahedra (see the next subsection), the generators of which have therefore all semi-stars of the same norm, but are not root systems.

Likewise, central symmetry is necessary for [Theorem 5.22](#): let $R \subset \mathbb{R}^2$ be the set of vertices of a regular triangle centered at the origin. All semi-stars of R (intersections with generic half-spaces) have the same norm, but it is clearly not a root system. Still, $R \cup -R$ is the root system $I_2(3)$.

Question 5.23. What other non-centrally symmetric sets $R \subset S^{d-1}$ have all semi-stars of the same norm? Is $R \cup -R$ always a root system?

Inscribed zonotopes

All vertex-transitive zonotopes are inscribed. However, not all inscribed zonotopes are Γ -permutahedra. Examples are easiest to find in dimension two, but exist in all dimensions. An inscribed zonotope needs also not to be combinatorially equivalent to a Γ -permutahedron. Examples were provided by Raman Sanyal and Sebastian Manecke (personal communication): the orthogonal projection of a Γ -permutahedron along one of its edge directions is

¹The assumption “reduced” is not necessary, but then the proof becomes more cumbersome.

again an inscribed zonotope, but not necessarily combinatorially equivalent to a Γ -permutahedron (see Figure 5.8).




Figure 5.8. Three inscribed zonotopes obtained as projections of higher-dimensional Γ -zonotopes. Left: the projection of the A_4 -zonotope is combinatorially equivalent to the A_3 -permutahedron, but is *not* vertex-transitive. Middle: a projection of the D_4 -zonotope. Right: a projection of the uniform F_4 -zonotope. The latter two are not combinatorially equivalent to a Γ -permutahedron.

There are further known examples which cannot be obtained as such repeated projections of Γ -permutahedra. All of them are still combinatorially equivalent to one which was obtained as a projection.

Question 5.24 (by Sanyal and Manecke). Are there inscribed zonotopes which are not combinatorially equivalent to a Γ -permutahedron or the repeated projection of a Γ -permutahedron along edge directions?

Hyperplane arrangements with congruent chambers


Zonotopes have a known relation to real hyperplane arrangements (the vertices of the zonotope correspond to the chambers of the hyperplane arrangement). Our results on vertex-transitive zonotopes translate as follows: a hyperplane arrangement whose symmetry group acts transitively on its chambers must be a reflection arrangement (the set of reflection hyperplanes of a finite reflection group). A more general question was asked by Caroline J. Klivans and Ed Swartz [44, Problem 13]:

Question 5.25 (by Klivans and Swartz). If all chambers of a real hyperplane arrangement are congruent, is it a reflection arrangement?

One finds that such an arrangement must be central and simplicial. The answer to [Question 5.25](#) is known to be affirmative in dimensions $d \in \{2, 3\}$ [22], but is open in $d \geq 4$. Dualizing again, the analogous question for zonotopes is the following:


Question 5.26. If all vertices of a zonotope are locally identical (they have the identical vertex-figures), is it combinatorially equivalent to a Γ -permutahedron, or more precisely, is its normalization (see [Definition 5.16](#)) a Γ -permutahedron?

6 Purely Edge-Transitive Polytopes




For this final chapter we return to the investigation of edge-transitivity for convex polytopes. In [Section 4.3](#) we claimed that most edge-transitive polytopes are also vertex-transitive. We even stated that there are no edge-transitive polytopes in dimension $d \geq 4$ that are not also vertex-transitive. In this chapter we provide the proof for this claim, which turns out to be quite technical.

Over the course of this thesis we mentioned repeatedly that there are exactly two edge-transitive polyhedra that are not vertex-transitive: the *rhombic dodecahedron* and the *rhombic triacontahedron* (from left to right):



Further examples with this symmetry can be found in dimension two: for every $k \geq 2$ there exists an infinite 1-parameter family of edge-transitive $2k$ -gons that are not vertex-transitive:



They are characterized by alternating interior angles, which is only possible for even-sided polygons. An odd-sided edge-transitive polygon is always vertex-transitive.

The surprising observation, and the main result of this chapter, is then that this is already the complete list:

Theorem 6.1. *If a polytope is edge- but not vertex-transitive, then it is one of the following:*

- (i) *a non-regular $2n$ -gon with alternating interior angles,*
- (ii) *the rhombic dodecahedron, or*
- (iii) *the rhombic triacontahedron.*

In particular, every edge-transitive polytope in dimension $d \geq 4$ is vertex-transitive.

A natural idea for proving [Theorem 6.1](#) would be by induction: we try to classify the edge-transitive polytopes that are not vertex-transitive by first classifying their potential facets, which are hopefully of the same type of symmetry and have been classified in a previous step. Unfortunately, this approach fails as the facets do not necessarily satisfy the induction hypothesis: facets of edge-transitive polytopes are not necessarily edge-transitive (for example, the facets of the (k, k) -duoprism from [Example 4.4](#) are k -gonal prisms, which are not edge-transitive if $k \neq 4$). Likewise, if a polytope is not vertex-transitive, its facets can still be vertex-transitive (as mentioned in [Remark 5.19 \(iii\)](#)).

As it turns out, we can still use this idea if we first embed the class of edge- but not vertex-transitive polytopes in the larger class of so-called *bipartite polytopes*, defined by geometric and combinatorial constraints instead of by symmetry (see [Section 6.1](#)). A core feature of a bipartite polytope is that its faces are again bipartite.

Chapter overview

In [Section 6.1](#) we introduce *bipartite polytopes*, a geometric generalization of edge-transitive polytopes that are not vertex-transitive. We show that faces of bipartite polytopes are bipartite and we explain how their classification helps in achieving the main result ([Theorem 6.1](#)). From that point on we focus on the classification of bipartite polytopes.

Some “easy” cases are taken care of in [Section 6.2](#), where we address dimension $d = 2$ as well as *inscribed* bipartite polytopes (this makes use of a result from [Chapter 5](#)). We reduce the classification of the remaining bipartite polytopes to the classification of bipartite *polyhedra* that are not inscribed (so-called *strictly* bipartite polyhedra).

This turns out to be the main technical part of the proof. It is presented over the course of [Sections 6.3](#) to [6.5](#). In [Section 6.3](#) we derive first properties of the geometric and combinatorial structure of strictly bipartite polyhedra. In [Section 6.4](#) we introduce *adjacent pairs*, a tool for systematically narrowing down the possible bipartite polytopes. This section ends in a larger case analysis, at the end of which we are left with a final candidate polyhedron. This polyhedron is then investigated in [Section 6.5](#), in which it is proven to be not strictly bipartite. It still turns out to be a remarkable near-miss.

The proof of the classification makes use of various classical geometric techniques, such as spherical polyhedra, the classification of the rhombic isohedra and the geometric realization of polyhedral graphs. We recall each briefly at the point of first use.

6.1 Bipartite polytopes

Let $P \subset \mathbb{R}^d$ denote a convex full-dimensional polytope.


Definition 6.2. P is called *bipartite*, if

- (i) all its edges are of the same length ℓ ,
- (ii) its edge-graph is bipartite, inducing partition $\mathcal{F}_0(P) = V_1 \cup V_2$ on the vertex set, and

(iii) there are radii $r_1 \leq r_2$ so that $\|v\| = r_i$ for all $v \in V_i$.

If $r_1 < r_2$, then P is called *strictly bipartite*. A vertex $v \in V_i$ is called an *i-vertex*. The numbers r_1, r_2 and ℓ are called the *parameters* of P .

Remark 6.3. An alternative definition of bipartite polytopes would replace (iii) by the condition that P has an *edge in-sphere*, that is, a sphere to which each edge of P is tangent to. The configuration depicted below (an edge of P connecting two vertices $v_1 \in V_1, v_2 \in V_2$) allows an immediate computation of the radius ρ of this in-sphere by solving the given system of equations:



This characterization will become relevant in [Section 6.5](#). But [Definition 6.2 \(iii\)](#) is the more convenient version to work with for the larger part and so we shall stick with it.

By a slight abuse of notation we call a polytope still bipartite even if it is only the translate of a bipartite polytope. This allows us to prove that faces of P are bipartite, even if they are not centered at the origin (cf. [Proposition 6.6](#)).

As advertised, bipartite polytopes generalize edge- but not vertex-transitive polytopes:

Proposition 6.4. *If P is edge- but not vertex-transitive, then P is bipartite.*

[Proposition 6.4](#) is a geometric analogue of the well known theorem in graph theory, that every edge- but not vertex-transitive graph is bipartite. This also served as the motivation for the term “bipartite polytope”. A proof of the graph version can be found in [\[31\]](#). Our proof of the geometric version proceeds analogously:

Proof of Proposition 6.4. The goal is to establish parts (i), (ii) and (iii) of [Definition 6.2](#). Part (i) follows trivially from edge-transitivity. To prove (ii) and (iii) fix some edge $e \in \mathcal{F}_1(P)$ with end vertices $v_1, v_2 \in \mathcal{F}_0(P)$. Let V_i denote the orbit of v_i under $\text{Aut}(P)$. We prove that $V_1 \cup V_2 = \mathcal{F}_0(P)$, $V_1 \cap V_2 = \emptyset$ and that the edge-graph is bipartite inducing partition $V_1 \cup V_2$.

Let $v \in \mathcal{F}_0(P)$ be some vertex and $\tilde{e} \in \mathcal{F}_1(P)$ an incident edge. By edge-transitivity, there is a symmetry $T \in \text{Aut}(P)$ that maps \tilde{e} onto e , and therefore maps v onto v_i for some $i \in \{1, 2\}$. Thus, v is in the orbit V_i . This holds for all vertices of P , and therefore $V_1 \cup V_2 = \mathcal{F}_0(P)$.

In general, orbits of group actions are either identical or disjoint. Since $V_1 \cup V_2 = \mathcal{F}_0(P)$, from $V_1 = V_2$ would follow $V_1 = \mathcal{F}_0(P)$, stating that P has a single orbit of vertices. But since P is *not* vertex-transitive, this cannot be. Thus, $V_1 \cap V_2 = \emptyset$, and therefore $V_1 \cup V_2 = \mathcal{F}_0(P)$.

Let $\tilde{e} \in \mathcal{F}_1(P)$ be an edge with end vertices \tilde{v}_1 and \tilde{v}_2 . By edge-transitivity, \tilde{e} can be mapped onto e by some symmetry $T \in \text{Aut}(P)$. Equivalently $\{T\tilde{v}_1, T\tilde{v}_2\} = \{v_1, v_2\}$. Since v_1 and v_2

belong to different orbits under $\text{Aut}(P)$, so do \tilde{v}_1 and \tilde{v}_2 . Hence, \tilde{e} has one end vertex in V_1 and one end vertex in V_2 . This holds for all edges $\tilde{e} \in E$ and so G_P is bipartite with partition $V_1 \cup V_2$. This proves (ii).

Finally, it remains to determine the radii $r_1 \leq r_2$ and prove (iii). Set $r_i := \|v_i\|$ (assuming w.l.o.g. that $\|v_1\| \leq \|v_2\|$). Then for every $v \in V_i$ there is a symmetry $T \in \text{Aut}(P) \subset O(\mathbb{R}^d)$ so that $Tv_i = v$, and thus $\|v\| = \|Tv_i\| = \|v_i\| = r_i$. \square

We shall see that the bipartite polytopes are still a surprisingly small class of polytopes:

Theorem 6.5. *If P is a bipartite polytope, then it is one of the following:*

- (i) *an edge-transitive polygon,*
- (ii) *the rhombic dodecahedron,*
- (iii) *the rhombic triacontahedron, or*
- (iv) *a Γ -permutahedron for some finite reflection group $\Gamma \subseteq O(\mathbb{R}^d)$ (see [Definition 5.9](#)).*

Presupposing [Theorem 6.5](#) we can proof our main result:

Proof of [Theorem 6.1](#). If P is edge-transitive but not vertex-transitive, then it is bipartite by [Proposition 6.4](#) and must be one of the polytopes listed in [Theorem 6.5](#). But it cannot be a regular polygon or a Γ -permutahedron as those are vertex-transitive. Thus, P is one of the remaining polytopes, all of which are edge-transitive. This matches the claim. \square

It then remains to prove [Theorem 6.5](#). This is “easier” because faces of bipartite polytopes are bipartite, as we shall prove now:

Proposition 6.6. *Let $f \in \mathcal{F}(P)$ be a face of P . It holds:*

- (i) *if P is bipartite, so is f .*
- (ii) *if P is strictly bipartite, then so is f , and $v \in \mathcal{F}_0(f) \subseteq \mathcal{F}_0(P)$ is an i -vertex in P if and only if it is an i -vertex in f .*
- (iii) *if $r_1 \leq r_2$ are the radii of P and $\rho_1 \leq \rho_2$ are the radii of f , then*

$$h^2 + \rho_i^2 = r_i^2,$$

where h is the height of f , that is, the distance of $\text{aff}(f)$ from the origin.

Proof. Properties clearly inherited by f are that all edges are of the same length and that the edge graph is bipartite. It remains to show that the radii $\rho_1 \leq \rho_2$ exist and are distributed compatibly with the bipartition of the edge graph of f .

Let $c \in \text{aff}(f)$ be the orthogonal projection of the origin onto $\text{aff}(f)$. Then $h := \|c\|$ is the height of f mentioned in (iii). For any vertex $v \in \mathcal{F}_0(f)$ which is an i -vertex in P , the triangle $\Delta := \text{conv}\{0, c, v\}$ has a right angle at c . Set $\rho_i := \|v - c\|$ and observe

$$\rho_i^2 := \|v - c\|^2 = \|v\|^2 - \|c\|^2 = r_i^2 - h^2. \tag{6.1}$$

In particular, the value ρ_i does not depend on the choice of the vertex, but only on i . In other words, translating f so that c becomes the origin gives a bipartite polytope according to [Definition 6.2](#). This proves (i), and (6.1) is equivalent to the equation in (iii). From (6.1) also follows $r_1 < r_2 \Leftrightarrow \rho_1 < \rho_2$, which proves (ii). \square

Observation 6.7. Given two adjacent vertices $v_1, v_2 \in \mathcal{F}_0(P)$ so that $v_i \in V_i$. If P has parameters r_1, r_2 and ℓ , then

$$\ell^2 = \|v_1 - v_2\|^2 = \|v_1\|^2 + \|v_2\|^2 - 2\langle v_1, v_2 \rangle = r_1^2 + r_2^2 - 2r_1r_2 \cos \angle(v_1, v_2),$$

which rearranges to

$$\cos \angle(v_1, v_2) = \frac{r_1^2 + r_2^2 - \ell^2}{2r_1r_2}.$$

In particular, the angle can be computed from the parameters alone and does not depend on the choice of the adjacent vertices. This will be of use later on.


6.2 Some easy cases

6.2.1 Bipartite polygons

The members easiest to describe (and to explicitly construct) are the bipartite *polygons*.

Foremost, the edge-graph is bipartite, and thus, a bipartite polygon must be a $2k$ -gon for some $k \geq 2$. One can show that the bipartite polygons are exactly the edge-transitive $2k$ -gons (cf. [Section 4.2](#)), and that such one is *strictly* bipartite if and only if it is *not* vertex-transitive (or equivalently, not regular).

The parameters r_1, r_2 and ℓ uniquely determine a bipartite polygon, as can be seen by explicit construction:



One starts with an arbitrarily chosen 1-vertex $v \in V_1$ placed on the circle $S_{r_1}(0)$. Its neighboring vertices are then uniquely determined as the intersections $S_{r_2}(0) \cap S_\ell(v)$. The procedure is repeated with the new vertices until the edge cycle closes (which only happens if the parameters are chosen appropriately).

The procedure also shows that the interior angle $\alpha_i \in (0, \pi)$ at an i -vertex only depends on i (and the parameters), but not on the chosen vertex $v \in V_i$. We summarize these insights in a corollary:

Corollary 6.8. *A bipartite polygon $P \subset \mathbb{R}^2$ is a $2k$ -gon with alternating interior angles $\alpha_1, \alpha_2 \in (0, \pi)$ (α_i being the interior angle at an i -vertex) and its shape is uniquely determined by its parameters (up to orientation).*

The exact value for the interior angles is not of much importance (though they could be obtained with [Observation 6.7](#)). However, the following will be of use:

Proposition 6.9. *The interior angles $\alpha_1, \alpha_2 \in (0, \pi)$ of a $2k$ -gonal bipartite polygon satisfy*

$$\alpha_2 \leq \alpha_{\text{reg}}^k \leq \alpha_1, \quad \text{with } \alpha_{\text{reg}}^k := \left(1 - \frac{1}{k}\right)\pi, \quad (6.2)$$

with equality in either part if and only if $r_1 = r_2$. Note also that α_{reg}^k is the interior angle of a regular $2k$ -gon.

Proof. The sum of interior angles of a $2k$ -gon is $2(k-1)\pi$, and thus

$$k\alpha_1 + k\alpha_2 = 2(k-1)\pi \quad \implies \quad \alpha_1 + \alpha_2 = 2\left(1 - \frac{1}{k}\right)\pi. \quad (6.3)$$

For two adjacent vertices $v_1, v_2 \in \mathcal{F}_0(P)$ (where $v_i \in V_i$), consider the triangle with vertices $0, v_1$ and v_2 , whose edge lengths are r_1, r_2 and ℓ , and whose interior angles at v_1 resp. v_2 are $\alpha_1/2$ resp. $\alpha_2/2$. By the law of sine holds

$$\frac{\sin(\alpha_1/2)}{\sin(\alpha_2/2)} = \frac{r_2}{r_1},$$

and from $r_1 \leq r_2$ (resp. $r_1 < r_2$) follows $\alpha_1 \geq \alpha_2$ (resp. $\alpha_1 > \alpha_2$). Together with [\(6.3\)](#) we obtain [\(6.2\)](#). \square

6.2.2 Inscribed bipartite polytopes

Because we are well-prepared (by [Chapter 5](#)), the next easiest case are the *inscribed* bipartite polytopes. Those are characterized by parameter identity $r_1 = r_2$. As we shall see, such bipartite polytopes must be zonotopes.

Theorem 6.10. *If $P \subset \mathbb{R}^d$ is an inscribed bipartite polytope, then it is a Γ -permutahedron.*

Proof. By [Proposition 6.6](#), the 2-dimensional faces of P are inscribed bipartite polygons. In particular, they have all edges of the same length. An inscribed polygon with all edges of the same length is regular. By [Corollary 6.8](#) the 2-faces are regular $2k$ -gons, therefore centrally symmetric.

But if all 2-faces of P are centrally symmetric, then P is a zonotope (by [Definition 5.1 \(iv\)](#)). Thus, P is an inscribed zonotope with all edges of the same length, and thus a Γ -permutahedron by [Theorem 5.3](#). \square

Γ -permutahedra are vertex-transitive by definition and do not provide examples of edge-transitive polytopes that are *not* vertex-transitive.

6.2.3 Strictly bipartite polytopes for $d \geq 4$

It remains to classify the *strictly* bipartite polytopes. This problem will be dealt with in two steps: dimension $d = 3$ and dimension $d \geq 4$. The detailed investigation of the case $d = 3$ (which turns out to be the actual hard work) will happen over the course of [Sections 6.3 to 6.5](#), the result of which is the following theorem:

Theorem 6.11. *If $P \subset \mathbb{R}^3$ is a strictly bipartite polyhedron (i.e., 3-dimensional polytope), then P is the rhombic dodecahedron or the rhombic triacontahedron.*

Presupposing [Theorem 6.11](#), the second step is surprisingly easy and we can classify the bipartite polytopes immediately:

Proof of [Theorem 6.5](#). The cases $d \in \{2, 3\}$ are dealt with in [Corollary 6.8](#) and [Theorem 6.11](#) respectively. It remains to show that there are no further strictly bipartite polytopes for $d \geq 4$. It suffices to show that there are no strictly bipartite polytopes in dimension $d = 4$ as any higher-dimensional example has a strictly bipartite 4-face (by [Proposition 6.6](#)).

Let $P \subset \mathbb{R}^4$ be a strictly bipartite 4-polytope. Let $e \in \mathcal{F}_1(P)$ be an edge of P . Then there are $s \geq 3$ cells (aka. 3-faces) $f_1, \dots, f_s \in \mathcal{F}_3(P)$ incident to e , each of which is again strictly bipartite (by [Proposition 6.6](#)). By [Theorem 6.11](#) each f_i is a rhombic dodecahedron or rhombic triacontahedron.

The dihedral angle of the rhombic dodecahedron resp. triacontahedron is 120° resp. 144° at every edge [[18](#)]. However, the dihedral angles meeting at e must sum up to less than 2π . With the given dihedral angles, this is impossible. \square

The remainder of this chapter is devoted to the proof of [Theorem 6.11](#).

6.3 Strictly bipartite polyhedra

Over the course of the next three sections we derive the classification of *strictly bipartite polyhedra*. The goal is to show that there are only the rhombic dodecahedron and the rhombic triacontahedron. In this section we gather general information about the geometric and combinatorial structure of strictly bipartite polyhedra to be used in the later sections.

From this section on, let $P \subset \mathbb{R}^3$ denote a fixed strictly bipartite polyhedron with radii $r_1 < r_2$ and edge length ℓ . The 2-faces of P will be shortly referred to as just *faces* of P . Since they are bipartite, they are necessarily $2k$ -gons. We will then use the following terminology:

- a face of P is of *type $2k$* (or a *$2k$ -face*) if it is a $2k$ -gon.
- an edge of P is of *type $(2k_1, 2k_2)$* (or a *$(2k_1, 2k_2)$ -edge*) if the two incident faces are of type $2k_1$ and $2k_2$ respectively.
- a vertex of P is of *type $(2k_1, \dots, 2k_s)$* (or a *$(2k_1, \dots, 2k_s)$ -vertex*) if its incident faces can be enumerated as f_1, \dots, f_s so that f_i is a $2k_i$ -face (note, the order of the numbers does not matter).

We write $\tau(v)$ for the type of a vertex $v \in \mathcal{F}_0(P)$.

In a given bipartite polyhedron, the type of a vertex, edge or face already determines much of its metric properties. For faces we can prove this already now:

Proposition 6.12. *For a face $f \in \mathcal{F}_2(P)$, any one of the following three properties determines the other two:*

- (i) *its type $2k$.*
- (ii) *its interior angles $\alpha_1 > \alpha_2$.*
- (iii) *its height h (that is, the distance of $\text{aff}(f)$ from the origin).*

Proof. Fix a face $f \in \mathcal{F}_2(P)$.

Suppose that the height h of f is known. By **Proposition 6.6**, a face of P of height h is bipartite with radii $\rho_i^2 := r_i^2 - h^2$ and edge length ℓ . By **Corollary 6.8**, these parameters then uniquely determine the shape of f , which includes its type and its interior angles. This shows (iii) \implies (i), (ii).

Suppose now that we know the interior angles $\alpha_1 > \alpha_2$ of f (it actually suffices to know one of these, say α_1). Fix a 1-vertex $v \in V_1$ of f and let $w_1, w_2 \in V_2$ be its two adjacent 2-vertices in f . Consider the simplex $\Delta := \text{conv}\{0, v, w_1, w_2\}$. The length of each edge of Δ is already determined, either by the parameters alone, or by additionally using the known interior angles. This is only non-obvious for the edge $\text{conv}\{w_1, w_2\}$:


$$\begin{aligned} \|w_1 - w_2\|^2 &= \|w_1 - v\|^2 + \|w_2 - v\|^2 - 2\langle w_1 - v, w_2 - v \rangle \\ &= 2\ell^2(1 - \underbrace{\cos \angle(w_1 - v, w_2 - v)}_{\alpha_1}). \end{aligned}$$

Thus, the shape of Δ is determined. In particular, this determines the height of the face $\text{conv}\{v, w_1, w_2\} \subset \Delta$ over the vertex $0 \in \Delta$. Since $\text{aff}\{v, w_1, w_2\} = \text{aff}(f)$, this determines the height of f in P . This proves (ii) \implies (iii).

Finally, suppose that the type $2k$ is known. We want to show that the height h is uniquely determined.¹ For the sake of contradiction, suppose that the type $2k$ does *not* uniquely determine the height of the face. Then there is another $2k$ -face $f' \in \mathcal{F}_2(P)$ of some height $h' \neq h$. W.l.o.g. assume $h' < h$.

Consider both faces as convex polygons embedded in \mathbb{R}^2 , centered at the origin. The vertices of a bipartite polygon are equally spaced by an angle of π/k (e.g. seen via **Observation 6.7**). We can therefore assume that the vertex v_i of f (resp. v'_i of f') is a positive multiple of $(\cos(i\pi/k), \sin(i\pi/k)) \in \mathbb{R}^2$ for $i \in \{1, \dots, 2k\}$ (see the figure below). In particular, there are factors $\delta_i \in \mathbb{R}$, so that $v'_i = \delta_i v_i$.

¹The reader motivated to prove this himself should know the following: it is indeed possible to write down a polynomial in h of degree four whose coefficients involve only r_1, r_2, ℓ and $\cos(\pi/k)$, and whose zeroes include all possible heights of any $2k$ -face of P . However, it turns out to be quite tricky to work out which zeroes correspond to feasible solutions. For certain values of the coefficients, there are multiple positive solutions for h , some of which correspond to *non-convex* $2k$ -faces. There seems to be no easy way to tell them apart.



The norms of the vectors v_1 , v_2 , $\delta_1 v_1$ and $\delta_2 v_2$ are the radii of the bipartite polygons f and f' . With **Proposition 6.6 (iii)** from $h > h'$ follows $\|v_1\| < \|\delta_1 v_1\|$ and $\|v_2\| < \|\delta_2 v_2\|$, and thus, (*) $\delta_1, \delta_2 > 1$.

Since both faces have edge length ℓ , we have $\|v_1 - v_2\| = \|\delta_1 v_1 - \delta_2 v_2\| = \ell$. We are going to derive the following contradiction:

$$\ell = \|v_1 - v_2\| \stackrel{(*)}{<} \delta_1 \|v_1 - v_2\| = \|\delta_1 v_1 - \delta_1 v_2\| \stackrel{(**)}{\leq} \|\delta_1 v_1 - \delta_2 v_2\| = \ell,$$

It remains to prove inequality (**). This inequality is trivially satisfied if $\delta_1 = \delta_2$. W.l.o.g. assume $\delta_1 < \delta_2$. We now provide a chain of equivalence transformations of (**) (note the use of $\delta_1 - \delta_2 < 0$ in the third step to reverse the inequality):

$$\begin{aligned} \|\delta_1 v_1 - \delta_1 v_2\|^2 &\leq \|\delta_1 v_1 - \delta_2 v_2\|^2 \\ \delta_1^2 \|v_2\|^2 - 2\delta_1^2 \langle v_1, v_2 \rangle &\leq \delta_2^2 \|v_2\|^2 - 2\delta_1 \delta_2 \langle v_1, v_2 \rangle \\ (\delta_1^2 - \delta_2^2) \|v_2\|^2 &\leq 2\delta_1 (\delta_1 - \delta_2) \langle v_1, v_2 \rangle \\ (\delta_1 + \delta_2) \|v_2\|^2 &\geq 2\delta_1 \langle v_1, v_2 \rangle \\ \bar{\delta} \|v_2\|^2 &\geq \delta_1 \langle v_1, v_2 \rangle, \end{aligned}$$

where $\bar{\delta} := (\delta_1 + \delta_2)/2$. Since $\bar{\delta} > \delta_1$, it suffices to check $\|v_2\|^2 \geq \langle v_1, v_2 \rangle$ in order to conclusively prove (**).

Note that $\|v_2\|^2 \geq \langle v_1, v_2 \rangle$ is equivalent to $\langle v_2, v_2 - v_1 \rangle \geq 0$, which is equivalent to the statement that the angle α (see figure above) is at most 90° . This is true since f is convex (the interior angle is $2\alpha \leq 180^\circ$).

We therefore found a contradiction to the assumption that there are two non-congruent $2k$ -faces and this proves $(i) \implies (ii), (iii)$. \square

Corollary 6.13. *Any two faces of P of the same height, or the same type, or with the same interior angles, are congruent.*

As a consequence of **Proposition 6.12** the notion of the interior angle $\alpha_i^k \in (0, \pi)$ of a $2k$ -face at an i -vertex is well-defined. If we define $\epsilon_k := (\alpha_1^k - \alpha_2^k)/2\pi$, then $\epsilon_k > 0$ by **Proposition 6.9**, and

$$\alpha_1^k = \left(1 - \frac{1}{k} + \epsilon_k\right)\pi, \quad \alpha_2^k = \left(1 - \frac{1}{k} - \epsilon_k\right)\pi.$$

Definition 6.14. If $\tau = (2k_1, \dots, 2k_s)$ is the type of a vertex, then define

$$K(\tau) := \sum_{i=1}^s \frac{1}{k_i}, \quad E(\tau) := \sum_{i=1}^s \epsilon_{k_i}.$$

Both quantities are strictly positive.

Proposition 6.15. Let $v \in \mathcal{F}_0(P)$ be a vertex of type $\tau = (2k_1, \dots, 2k_s)$.

- (i) If $v \in V_1$, then $E(\tau) < K(\tau) - 1$ and $s = 3$.
- (ii) If $v \in V_2$, then $E(\tau) > s - 2 - K(\tau)$.

Proof. Let $f_1, \dots, f_s \in \mathcal{F}_2(P)$ be the faces incident to v , so that f_j is a $2k_j$ -face. The interior angle of f_j at v is $\alpha_i^{k_j}$, and the sum of these must be smaller than 2π . In formulas

$$2\pi > \sum_{j=1}^s \alpha_i^{k_j} = \sum_{j=1}^s \left(1 - \frac{1}{k_j} \pm \epsilon_{k_j}\right)\pi = (s - K(v) \pm E(v))\pi,$$

where \pm is the plus sign for $i = 1$, and the minus sign for $i = 2$. Rearranging for $E(v)$ yields $(*) \mp E(v) > s - 2 - K(v)$. If $i = 2$, this gives the statement in (ii). For $i = 1$ note that from $k_j \geq 2 \implies K(v) \leq s/2$ follows

$$s < -E(v) + K(v) + 2 \leq 0 + \frac{s}{2} + 2 \implies s < 4.$$

The minimal degree of a vertex in a polyhedron is at least three, hence $s = 3$, and $(*)$ yields (i). □

This allows us to exclude all but a manageable list of possible types for 1-vertices. By **Proposition 6.15** (i) a vertex $v \in V_1$ has degree three and a type of the form $(2k_1, 2k_2, 2k_3)$.

Corollary 6.16. For a 1-vertex $v \in V_1$ of type τ holds $K(\tau) > 1 + E(\tau) > 1$. One checks that this leaves exactly the options in **Table 6.1**.

τ	$K(\tau)$	Γ
(4, 4, 4)	3/2	$I_1 \oplus I_1 \oplus I_1$
(4, 4, 6)	4/3	$I_1 \oplus I_2(3)$
(4, 4, 8)	5/4	$I_1 \oplus I_2(4)$
(4, 4, 10)	6/5	$I_1 \oplus I_2(5)$
(4, 4, 12)	7/6	$I_1 \oplus I_2(6)$
\vdots	\vdots	\vdots
(4, 4, 2k)	$1 + 1/k$	$I_1 \oplus I_2(k)$
(4, 6, 6)	7/6	$A_3 \cong D_3$
(4, 6, 8)	13/12	B_3
(4, 6, 10)	31/30	H_3

Table 6.1. Possible types of 1-vertices, their K -values and the Γ of the Γ -permutahedron in which all vertices have this type.


The types in Table 6.1 are called the *possible types* of 1-vertices. Each of the possible types is realizable in the sense that there exists a bipartite polyhedron in which all 1-vertices have this type. Examples provide the Γ -permutahedra (the Γ of that Γ -permutahedron is listed in Table 6.1). These are not *strictly* bipartite though.

The convenient feature of Γ -permutahedra is that all their vertices are of the same type. We cannot assume this for general strictly bipartite polyhedra, not even for their 1-vertices. To determine which types of vertices can appear together in the same polyhedron, we have to turn to *spherical interior angles* defined via *spherical polyhedra*.

A *spherical polyhedron* is an embedding of a planar graph into the unit sphere, so that all edges are embedded as great circle arcs and each region (spherical face) is convex (all interior angles are less than 180°). If $0 \in \text{int}(P)$ then we can define a spherical polyhedron P^S by applying central projection

$$\mathbb{R}^3 \setminus \{0\} \rightarrow S^2, \quad x \mapsto \frac{x}{\|x\|}$$

to all vertices and edges of P .



The vertices, edges and faces of P have spherical counterparts in P^S obtained as projections onto the unit sphere. Those will be denoted with a superscript “S”. For example, if $e \in \mathcal{F}_1(P)$ is an edge of P , then e^S denotes the corresponding “spherical edge”, which is a great circle arc obtained as the projection of e onto the sphere.

We show that the spherical polyhedron of a strictly bipartite polyhedron is well-defined:

Proposition 6.17. $0 \in \text{int}(P)$.

Proof. The proof proceeds in several steps.

Step 1: Fix a 1-vertex $v \in V_1$ with neighbors $w_1, w_2, w_3 \in V_2$, and let $u_i := w_i - v$ be the direction of the edge $\text{conv}\{v, w_i\}$ emanating from v . Let $f_{ij} \in \mathcal{F}_2(P)$ denote the face of P that contains v, w_i and w_j . The interior angle of f_{ij} at v is then $\sphericalangle(u_i, u_j)$, which by Proposition 6.9 and $k \geq 2$ satisfies

$$\sphericalangle(u_i, u_j) > \left(1 - \frac{1}{k}\right)\pi \geq \frac{\pi}{2} \implies \langle u_i, u_j \rangle < 0.$$

Step 2: Besides v , the polyhedron P contains another 1-vertex $v' \in V_1$. It holds $v' \in v + \text{cone}\{u_1, u_2, u_3\}$, which means that there are non-negative coefficients $a_1, a_2, a_3 \geq 0$, at least one positive, so that $v + a_1u_1 + a_2u_2 + a_3u_3 = v'$. Rearranging and applying $\langle v, \cdot \rangle$ yields

$$\begin{aligned} a_1\langle v, u_1 \rangle + a_2\langle v, u_2 \rangle + a_3\langle v, u_3 \rangle &= \langle v, v' \rangle - \langle v, v \rangle \\ &= r_1^2 \cos \sphericalangle(v, v') - r_1^2 < 0. \end{aligned} \tag{6.4}$$

The value $\langle v, u_i \rangle$ is independent of i (see [Observation 6.7](#)). Since there is at least one positive coefficient a_i , from [\(6.4\)](#) follows $\langle v, u_i \rangle < 0$.

Step 3: By the previous steps, $\{v, u_1, u_2, u_3\}$ is a set of four vectors with pair-wise negative inner product. The convex hull of such an arrangement in 3-dimensional Euclidean space is known to contain the origin in its interior (a proof is given in [Proposition G.1](#)), that is, there are positive coefficients $a_0, \dots, a_3 > 0$ with $a_0v + a_1u_1 + a_2u_2 + a_3u_3 = 0$. In other words: $0 \in v + \text{int}(\text{cone}\{u_1, u_2, u_3\})$.

Step 4: If $H(f)$ denotes the half-space that defines the face $f \in \mathcal{F}_2(P)$, then

$$0 \in v + \text{int}(\text{cone}\{u_1, u_2, u_3\}) = \bigcap_{f \sim v} \text{int}(H(f)).$$

Thus, $0 \in \text{int}(H(f))$ for all faces f incident to v . But since every face is incident to a 1-vertex, we obtain $0 \in \text{int}(H(f))$ for all $f \in \mathcal{F}_2(P)$, and thus $0 \in \text{int}(P)$ as well. \square

The main reason for introducing spherical polyhedra is that we can talk about the *spherical interior angles* of their faces.

Let $f \in \mathcal{F}_2(P)$ be a face, and $v \in \mathcal{F}_0(f)$ one of its vertices. Let $\alpha(f, v)$ denote the interior angle of f at v , and $\beta(f, v)$ the spherical interior angle of f^S at v^S . It only requires a relatively straightforward computation (involving some spherical geometry) to establish a direct relation between these angles: e.g. if v is a 1-vertex, then

$$\sin^2(\ell^S) \cdot (1 - \cos \beta(f, v)) = \left(\frac{\ell}{r_2}\right)^2 \cdot (1 - \cos \alpha(f, v)),$$

where ℓ^S denotes the arc-length of an edge of P^S (indeed, all edges are of the same length). An equivalent formula exists for 2-vertices. The details of the computation and the actual relation are not of importance and can be found in [Appendix G.2](#). The core message is that the value of $\alpha(f, v)$ uniquely determines the value of $\beta(f, v)$ and vice versa. Since the value of $\alpha(f, v) = \alpha_i^k$ only depends on the type of the face and the partition class of the vertex, the same then holds for $\beta(f, v)$. Therefore, the notion β_i^k for the spherical interior angle of a $2k$ -gonal spherical face of P^S at (the projection of) an i -vertex is well-defined. We have

$$\beta_i^{k_1} = \beta_i^{k_2} \iff \alpha_i^{k_1} = \alpha_i^{k_2} \stackrel{6.12}{\iff} k_1 = k_2, \quad (6.5)$$

where we use [Proposition 6.12](#) in the right-most equivalence.

Observation 6.18. The spherical interior angles β_i^k have the following properties:

- (i) The spherical interior angles surrounding a vertex add up to exactly 2π . That is, for an i -vertex $v \in \mathcal{F}_0(P)$ of type $(2k_1, \dots, 2k_s)$ holds

$$\beta_i^{k_1} + \dots + \beta_i^{k_s} = 2\pi.$$

- (ii) The sum of interior angles of a spherical polygon exceeds the interior angle sum of a respective flat polygon. That is,

$$k\beta_1^k + k\beta_2^k > 2(k-1)\pi \implies \beta_1^k + \beta_2^k > 2\left(1 - \frac{1}{k}\right)\pi.$$

These simple observations have several relevant consequences for the structure of a strictly bipartite polyhedron P :

Corollary 6.19. P contains at most two different types of 1-vertices, and if there are two different types, then one is of the form $(4, 4, 2k)$ and the other one is of the form $(4, 6, 2k')$ for distinct $k \neq k'$ and $2k' \in \{6, 8, 10\}$.

Proof. Each possible type listed in Table 6.1 is either of the form $(4, 4, 2k)$ or of the form $(4, 6, 2k')$ for some $2k \geq 4$ resp. $2k' \in \{6, 8, 10\}$.

If P contains simultaneously 1-vertices of type $(4, 4, 2k_1)$ and $(4, 4, 2k_2)$, apply Observation 6.18 (i) to see

$$\beta_1^2 + \beta_1^2 + \beta_1^{k_1} \stackrel{(i)}{=} \beta_1^2 + \beta_1^2 + \beta_1^{k_2} \implies \beta_1^{k_1} = \beta_1^{k_2} \stackrel{(6.5)}{\implies} k_1 = k_2.$$

If P contains simultaneously 1-vertices of type $(4, 6, 2k'_1)$ and $(4, 6, 2k'_2)$, then

$$\beta_1^2 + \beta_1^3 + \beta_1^{k'_1} \stackrel{(i)}{=} \beta_1^2 + \beta_1^3 + \beta_1^{k'_2} \implies \beta_1^{k'_1} = \beta_1^{k'_2} \stackrel{(6.5)}{\implies} k'_1 = k'_2.$$

Finally, if P contains simultaneously 1-vertices of type $(4, 4, 2k)$ and $(4, 6, 2k')$, then

$$\beta_1^2 + \beta_1^2 + \beta_1^k \stackrel{(i)}{=} \beta_1^2 + \beta_1^3 + \beta_1^{k'} \implies \beta_1^k - \beta_1^{k'} = \underbrace{\beta_1^3 - \beta_1^2}_{\neq 0 \text{ by (6.5)}} \stackrel{(6.5)}{\implies} k \neq k'.$$

□

Since each edge of P is incident to a 1-vertex, we obtain the following:

Corollary 6.20. If P has only 1-vertices of types $(4, 4, 2k)$ and $(4, 6, 2k')$ then each edge of P is of one of the types

$$\underbrace{(4, 4), (4, 2k)}_{\text{from a } (4, 4, 2k)\text{-vertex}}, \underbrace{(4, 6), (4, 2k') \text{ or } (6, 2k')}_{\text{from a } (4, 6, 2k')\text{-vertex}}.$$

Corollary 6.21. The dihedral angle of an edge $e \in \mathcal{F}_1(P)$ of P only depends on its type.

Proof. Suppose that e is a $(2k_1, 2k_2)$ -edge. Then e is incident to a 1-vertex $v \in V_1$ of type $(2k_1, 2k_2, 2k_3)$. By Observation 6.18 (i) holds $\beta_1^{k_3} = 2\pi - \beta_1^{k_1} - \beta_1^{k_2}$, which further determines k_3 . By Proposition 6.12 we have uniquely determined interior angles $\alpha_1^{k_1}$, $\alpha_1^{k_2}$ and $\alpha_1^{k_3}$.

It is known that for a simple vertex (that is, a vertex of degree three) the interior angles of the incident faces already determine the dihedral angles at the incident edges (a proof can be found in Proposition G.2). Consequently, the dihedral angle at e is determined. □

The next result shows that Γ -permutahedra are the only bipartite polyhedra in which all vertices can have the same type.

Corollary 6.22. P cannot contain a 1-vertex and a 2-vertex of the same type.

Proof. Let $v \in \mathcal{F}_0(P)$ be some vertex of type $(2k_1, 2k_2, 2k_3)$. The incident edges are of type $(2k_1, 2k_2)$, $(2k_2, 2k_3)$ and $(2k_3, 2k_1)$ respectively. By [Corollary 6.21](#) the dihedral angles of these edges are uniquely determined, and since v is a simple vertex, the interior angles of its incident faces are also uniquely determined (see [Proposition G.2](#)). In particular, we obtain the same angles independent of whether v is a 1-vertex or a 2-vertex.

Suppose now that $v_i \in V_i, i \in \{1, 2\}$ are vertices of the same type τ . Since 1-vertices are always simple, the 2-vertex must be simple too, and their incident faces share the same interior angles. That is, $\alpha_1^k = \alpha_2^k$ for each $2k \in \tau$. But this is not possible if P is strictly bipartite (using [Proposition 6.6 \(ii\)](#) and [Proposition 6.9](#)). \square

6.4 Adjacent pairs

In this section we prepare and perform a larger case analysis, at the end of which we will be left with the description for a single remaining candidate polyhedron that has to be refuted. We achieve this by showing that almost all types of 1-vertices lead to contradictions. Our main tool will be so-called *adjacent pairs*.

Definition 6.23. Given a 1-vertex $v \in V_1$ of type $\tau_1 = (2k_1, 2k_2, 2k_3)$, for any two distinct $i, j \in \{1, 2, 3\}$, v has a neighbor $w \in V_2$ of type $\tau_2 = (2k_i, 2k_j, *, \dots, *)$, where $*$ are placeholders for unknown entries. The pair of types

$$(\tau_1, \tau_2) = ((2k_1, 2k_2, 2k_3), (2k_i, 2k_j, *, \dots, *))$$

is called an *adjacent pair* of P .

We proceed proving that most adjacent pairs cannot occur in P . Excluding enough adjacent pairs for fixed τ_1 then proves that the type τ_1 cannot occur as the type of a 1-vertex.

Our main tools for achieving this will be the inequalities established in [Proposition 6.15 \(i\)](#) and [\(ii\)](#), that is,

$$E(\tau_1) \stackrel{(i)}{<} K(\tau_1) - 1 \quad \text{and} \quad E(\tau_2) \stackrel{(ii)}{>} s - 2 - K(\tau_2),$$

where s is the number of elements in τ_2 . For a warmup, and as a template for further calculations, we prove that the adjacent pair $(\tau_1, \tau_2) = ((4, 6, 8), (6, 8, 8))$ cannot occur in P .

Example 6.24. By [Proposition 6.15 \(i\)](#) we have

$$\epsilon_2 + \epsilon_3 + \epsilon_4 = E(\tau_1) \stackrel{(i)}{<} K(\tau_1) - 1 = \frac{1}{2} + \frac{1}{3} + \frac{1}{4} - 1 = \frac{1}{12}. \quad (6.6)$$

On the other hand, by [Proposition 6.15 \(ii\)](#) we have

$$\begin{aligned} \frac{2}{12} &= 3 - 2 - \left(\frac{1}{3} + \frac{1}{4} + \frac{1}{4}\right) = s - 2 - K(\tau_2) \\ &\stackrel{(ii)}{<} E(\tau_2) = \underbrace{\epsilon_3 + \epsilon_4}_{<1/12} + \underbrace{\epsilon_4}_{<1/12} < \frac{2}{12}, \end{aligned} \quad (6.7)$$

which is a contradiction. Hence this adjacent pair cannot occur. Note that we used [\(6.6\)](#) to upperbound certain sums of ϵ_i in [\(6.7\)](#).

An adjacent pair excluded by using the inequalities from [Proposition 6.15](#) as demonstrated in [Example 6.24](#) will be called *infeasible*.

The argument applied in [Example 6.24](#) will be repeated many times for many different adjacent pairs in the upcoming case analysis, and we shall therefore use a tabular form to abbreviate it. After fixing $\tau_1 = (4, 6, 8)$, the argument to refute the adjacent pair $(\tau_1, \tau_2) = ((4, 6, 8), (6, 8, 8))$ will be abbreviated as shown in the first row of the following table:

τ_2	$s - 2 - K(\tau_2)$	$\stackrel{?}{<}$	$E(\tau_2)$
$(6, 8, 8)$	2/12	$\not<$	$(\epsilon_3 + \epsilon_4) + \epsilon_4 < 2/12$
$(6, 8, 6, 6)$	9/12	$\not<$	$(\epsilon_3 + \epsilon_4) + \epsilon_3 + \epsilon_3 < 3/12$

The second row displays the analogue argument for another adjacent pair, namely, the pair $((4, 6, 8), (6, 8, 6, 6))$, showing that it is infeasible as well. Both rows will reappear in later tables when we exclude $\tau_1 = (4, 6, 8)$ entirely. Note that the terms in the column $E(\tau_2)$ are grouped by parenthesis to indicate which subsums are upper bounded via [Proposition 6.15](#) (i). In this example, if there are n groups, then the sum is upper bounded by $n/12$.

The placeholders in an adjacent pair $((2k_1, 2k_2, 2k_3), (2k_i, 2k_j, *, \dots, *))$ can, in theory, be replaced by an arbitrary sequence of even numbers, and each such pair has to be refuted separately. The following proposition makes this task more tractable. We write $\tau \subset \tau'$ if τ is a *subtype* of τ' , that is, a vertex type that can be obtained from τ' by removing some of its entries.

Proposition 6.25. *If (τ_1, τ_2) is an infeasible adjacent pair, then so is (τ_1, τ'_2) for every $\tau'_2 \supset \tau_2$.*

Proof. Suppose $\tau_2 = (2k_1, \dots, 2k_s)$, $\tau'_2 = (2k_1, \dots, 2k_s, 2k_{s+1}, \dots, 2k_{s'}) \supset \tau_2$, and that the pair (τ_1, τ'_2) is not infeasible. Then τ'_2 satisfies [Proposition 6.15](#) (ii):

$$\begin{aligned} & E(\tau'_2) > s' - 2 - K(\tau'_2) \\ \implies & E(\tau_2) > s - 2 - K(\tau_2) + \sum_{i=s+1}^{s'} \overbrace{\left(1 - \frac{1}{k_i} - \epsilon_{k_i}\right)}^{\alpha_2^{k_i}/\pi > 0} > s - 2 - K(\tau_2). \end{aligned}$$

But this is exactly the statement that τ_2 satisfies [Proposition 6.15](#) (ii), i.e., the pair (τ_1, τ_2) is also not infeasible. \square

By [Proposition 6.25](#) it is sufficient to exclude so-called *minimal infeasible adjacent pairs*, that is, infeasible adjacent pairs (τ_1, τ_2) for which (τ_1, τ'_2) is not infeasible for any $\tau'_2 \subset \tau_2$.


A second potential problem is that we know little about the values that might replace the placeholders in $\tau_2 = (2k_i, 2k_j, *, \dots, *)$. For our immediate goal, dealing with the following special case is sufficient:

Proposition 6.26. *The placeholders in an adjacent pair $((4, 6, 2k'), (6, 2k', *, \dots, *))$ can only contain 4, 6 and $2k'$.*

Proof. Suppose that P contains an adjacent pair

$$(\tau_1, \tau_2) = ((4, 6, 2k'), (6, 2k', 2k, *, \dots, *))$$

induced by a 1-vertex $v \in V_1$ of type τ_1 with neighbor $w \in V_2$ of type τ_2 . Suppose further, that $2k \notin \{4, 6, 2k'\}$. The vertex w is then incident to a $2k$ -face, and therefore also to a 1-vertex $u \in V_1$ of type $(4, 4, 2k)$ (u cannot be of type $(4, 6, 2k)$ because of $k \neq k'$ and [Corollary 6.19](#)). This configuration is depicted below:



Note that w is also incident to a 4-face, and thus $(6, 2k', 2k, 4) \subseteq \tau_2$.

By [Proposition 6.15](#) (i) the existence of 1-vertices of type $(4, 4, 2k)$ and $(4, 6, 2k')$ yields inequalities

$$\epsilon_2 + \epsilon_2 + \epsilon_k < \frac{1}{k} \quad \text{and} \quad \epsilon_2 + \epsilon_3 + \epsilon_{k'} < \frac{1}{k'} - \frac{1}{6}. \quad (6.8)$$

Since τ_2 has $\tau := (6, 2k', 2k, 4)$ as a subtype, by [Proposition 6.25](#) it suffices to show that the pair $((4, 6, 2k'), (6, 2k', 2k, 4))$ is infeasible. This follows via [Proposition 6.15](#) (ii):

$$\begin{aligned} \frac{7}{6} - \frac{1}{k} - \frac{1}{k'} &= 4 - 2 - \left(\frac{1}{3} + \frac{1}{k'} + \frac{1}{k} + \frac{1}{2} \right) = 4 - 2 - K(\tau) \\ &\stackrel{(ii)}{<} E(\tau) = \underbrace{\epsilon_2 + \epsilon_3 + \epsilon_{k'}}_{< 1/k' - 1/6} + \underbrace{\epsilon_k}_{< 1/k} \stackrel{(6.8)}{<} \frac{1}{k} + \frac{1}{k'} - \frac{1}{6}, \end{aligned}$$

which rearranges to $1/k + 1/k' > 2/3$. Recalling $2k' \in \{6, 8, 10\} \implies k' \geq 3$ (from [Corollary 6.19](#)) and $2k \notin \{4, 6, 2k'\} \implies k \geq 4$ shows that this is not possible. \square

6.4.1 The case $\tau_1 = (4, 6, 10)$

If P contains a 1-vertex of type $(4, 6, 10)$, then it contains an adjacent pair of the form

$$(\tau_1, \tau_2) = ((4, 6, 10), (6, 10, *, \dots, *)).$$

We proceed as demonstrated in [Example 6.24](#). [Proposition 6.15](#) (i) yields $\epsilon_2 + \epsilon_3 + \epsilon_5 < 1/30$. By [Proposition 6.26](#) the placeholders can only take on values 4, 6 or 10. The following table lists the minimally infeasible adjacent pairs and proves their infeasibility.

τ_2	$s - 2 - K(\tau_2)$	$\stackrel{?}{<}$	$E(\tau_2)$
$(6, 10, 6)$	$4/30$	$\not<$	$(\epsilon_3 + \epsilon_5) + \epsilon_3 < 2/30$
$(6, 10, 10)$	$8/30$	$\not<$	$(\epsilon_3 + \epsilon_5) + \epsilon_5 < 2/30$
$(6, 10, 4, 4)$	$14/30$	$\not<$	$(\epsilon_2 + \epsilon_3 + \epsilon_5) + \epsilon_2 < 2/30$

By [Proposition 6.25](#) we conclude: the placeholders in $\tau_2 = (6, 10, *, \dots, *)$ can contain no 6 or 10, and at most one 4. This leaves us with the option $\tau_2 = (4, 6, 10) = \tau_1$ which is not possible by [Corollary 6.22](#). Therefore, P cannot contain a 1-vertex of type $(4, 6, 10)$.

6.4.2 The case $\tau_1 = (4, 6, 8)$

If P contains a 1-vertex of type $(4, 6, 8)$, then it also contains an adjacent pair of the form

$$(\tau_1, \tau_2) = ((4, 6, 8), (6, 8, *, \dots, *)).$$

We proceed as demonstrated in [Example 6.24](#). [Proposition 6.15](#) (i) yields $\epsilon_2 + \epsilon_3 + \epsilon_4 < 1/12$. By [Proposition 6.26](#) the placeholders can only take on values 4, 6 or 8. The following table lists the minimally infeasible adjacent pairs and proves their infeasibility.

τ_2	$s - 2 - K(\tau_2)$	$\stackrel{?}{<}$	$E(\tau_2)$
$(6, 8, 8)$	$2/12$	$\not<$	$(\epsilon_3 + \epsilon_4) + \epsilon_3 < 2/12$
$(6, 8, 4, 4)$	$5/12$	$\not<$	$(\epsilon_2 + \epsilon_3 + \epsilon_4) + \epsilon_2 < 2/12$
$(6, 8, 4, 6)$	$7/12$	$\not<$	$(\epsilon_2 + \epsilon_3 + \epsilon_4) + \epsilon_3 < 2/12$
$(6, 8, 6, 6)$	$9/12$	$\not<$	$(\epsilon_2 + \epsilon_3 + \epsilon_4) + \epsilon_3 + \epsilon_3 < 3/12$

By [Proposition 6.25](#) we conclude: the placeholders in $\tau_2 = (6, 8, *, \dots, *)$ can contain no 8, and at most one 4 or 6, but not both at the same time.

This leaves us with the options $\tau_2 = (4, 6, 8)$ and $\tau_2 = (6, 6, 8)$. In the first case, $\tau_1 = \tau_2$ which is not possible by [Corollary 6.22](#). In the second case, there would be two adjacent 6-faces, but P does not contain $(6, 6)$ -edges by [Corollary 6.20](#) with $2k' = 8$. Therefore, P cannot contain a 1-vertex of type $(4, 6, 8)$.

6.4.3 The case $\tau_1 = (4, 6, 6)$

If P contains a 1-vertex of type $(4, 6, 6)$, then it also contains an adjacent pair of the form

$$(\tau_1, \tau_2) = ((4, 6, 6), (6, 6, *, \dots, *)).$$

We proceed as demonstrated in [Example 6.24](#). [Proposition 6.15](#) (i) yields $\epsilon_2 + \epsilon_3 + \epsilon_3 < 1/6$. By [Proposition 6.26](#) the placeholders can only take on values 4 or 6. The following table lists the minimally infeasible adjacent pairs and proves their infeasibility.

τ_2	$s - 2 - K(\tau_2)$	$\stackrel{?}{<}$	$E(\tau_2)$
$(6, 6, 4, 4)$	$2/6$	$\not<$	$(\epsilon_2 + \epsilon_3 + \epsilon_3) + \epsilon_2 < 2/6$
$(6, 6, 6, 4)$	$3/6$	$\not<$	$(\epsilon_2 + \epsilon_3 + \epsilon_3) + \epsilon_3 < 2/6$
$(6, 6, 6, 6)$	$4/6$	$\not<$	$(\epsilon_3 + \epsilon_3) + (\epsilon_3 + \epsilon_3) < 2/6$

By [Proposition 6.25](#) we conclude: the placeholders in $\tau_2 = (6, 6, *, \dots, *)$ can contain at most one 4 or 6, but not both at the same time.

This leaves us with the options $\tau_2 = (4, 6, 6)$ and $\tau_2 = (6, 6, 6)$. In the first case we have $\tau_1 = \tau_2$, which is not possible by [Corollary 6.22](#).

Excluding $\tau_2 = (6, 6, 6)$ needs more work: fix a 6-face $f \in \mathcal{F}_2(P)$. Each edge of f is either of type $(4, 6)$ or of type $(6, 6)$ (by [Corollary 6.20](#)). Each 1-vertex of f is incident to *exactly one* of the $(6, 6)$ -edges of f (since its type is $(4, 6, 6)$). Each 2-vertex of f is incident to *either exactly zero or exactly two* $(6, 6)$ -edges of f (since if there is one $(6, 6)$ -edge, then its type

must be of type $(6, 6, 6)$). Such a configuration of edge types around a 6-gon is not possible: the conditions from the 1-vertices imply that there are *exactly* three $(6, 6)$ -edges around f , but the conditions from the 2-vertices imply that the number of $(6, 6)$ -edges is *even* (see also [Figure 6.1](#)).




Figure 6.1. Possible distributions of $(4, 6)$ -edges (black) and $(6, 6)$ -edges (gray) around a 6-face as discussed in [Section 6.4.3](#). The top row shows configurations compatible with the conditions imposed by the 1-vertices (black), and the bottom row shows the configurations compatible with the conditions imposed by the 2-vertices (white).

Therefore, P cannot contain a 1-vertex of type $(4, 6, 6)$.

Observation 6.27. It is a consequence of [Sections 6.4.1](#) to [6.4.3](#) that P cannot contain a 1-vertex of a type $(4, 6, 2k')$ for a $2k' \in \{6, 8, 10\}$. By [Corollary 6.19](#) this means that *all* 1-vertices of P are of the same type $\tau_1 = (4, 4, 2k)$ for some fixed $2k \geq 4$.

We distinguish the case $(4, 4, 4)$ from the cases $(4, 4, 2k)$ with $2k \geq 6$.

6.4.4 The case $\tau_1 = (4, 4, 4)$

In this case, all 2-faces are 4-gons, and all 4-gons are congruent by [Proposition 6.12](#). A 4-gon with all edges of the same length is known as a *rhombus*, and the polyhedra with congruent rhombic faces are known as *rhombic isohedra* (from german *Rhombenisoeder*). These have a known classification:

Theorem 6.28 (Bilinski [[4](#)]). *If P is a polyhedron with congruent rhombic faces, then P is one of the following (cf. [Figure 6.2](#)):*

- (i) *a member of the infinite family of rhombic hexahedra, i.e., P can be obtained from a cube by stretching or squeezing it along a long diagonal,*
- (ii) *the rhombic dodecahedron,*
- (iii) *the Bilinski dodecahedron,*
- (iv) *the rhombic icosahedron, or*
- (v) *the rhombic triacontahedron.*

Corollary 6.29. *If P is strictly bipartite with all 1-vertices of type $(4, 4, 4)$, then P is one of the following:*




Figure 6.2. From top to bottom, from left to right: two versions of the “stretched cube” (rhombic hexahedron), the rhombic dodecahedron, the Bilinski dodecahedron, the rhombic icosahedron, and the rhombic triacontahedron.

- (i) the rhombic dodecahedron,
- (ii) the rhombic triacontahedron.

Proof. The polytopes (i) and (ii) are edge-transitive but not vertex-transitive, and they are not inscribed. By [Proposition 6.4](#) they are therefore *strictly* bipartite.

It remains to exclude the other polyhedra listed in [Theorem 6.28](#). The rhombic hexahedra include the cube, which is inscribed, hence not strictly bipartite. The other polyhedra are not inscribed, and so [Proposition 6.9](#) with $2k = 4$ yields

$$\alpha_2 < \pi/2 < \alpha_1.$$

That is, the interior angles at 1-vertices are obtuse, while the interior angles at 2-vertices are acute. Visual inspection of the remaining polyhedra (see [Figure 6.2](#)) shows that each has a vertex in which acute and obtuse angles meet. These vertices can neither be assigned to V_1 nor to V_2 , and the polyhedron cannot be bipartite. \square

Since we expect that the polyhedra listed in [Corollary 6.29](#) are the only two strictly bipartite polyhedra, the goal of the remaining sections is to show that all other possible configurations must lead to a contradiction.

6.4.5 The case $\tau_1 = (4, 4, 2k)$, $2k \geq 6$

If P contains a 1-vertex of type $(4, 4, 2k)$ with $2k \geq 6$, then it also has an adjacent pair of the form

$$(\tau_1, \tau_2) = ((4, 4, 2k), (4, 2k, *, \dots, *)).$$


We proceed as demonstrated in [Example 6.24](#). [Proposition 6.15](#) (i) yields $\epsilon_2 + \epsilon_2 + \epsilon_k < 1/k$. Since $(4, 4, 2k)$ is the only type of 1-vertex of P , there are only 4-faces and $2k$ -faces and the placeholders can only take on the values 4 and $2k$ (note that we do not use [Proposition 6.26](#)). The following table lists some inequalities derived for infeasible pairs:

τ_2	$s - 2 - K(\tau_2)$	$\stackrel{?}{<}$	$E(\tau_2)$
$(4, 2k, 4, 4, 4)$	$1 - 1/k$	$<$	$(\epsilon_2 + \epsilon_2 + \epsilon_k) + (\epsilon_2 + \epsilon_2) < 2/k$
$(4, 2k, 4, 4, 2k)$	$3/2 - 2/k$	$<$	$(\epsilon_2 + \epsilon_2 + \epsilon_k) + (\epsilon_2 + \epsilon_k) < 2/k$

One checks that these inequalities are not satisfied for $2k \geq 6$. **Proposition 6.25** then states that the placeholders can contain at most two 4-s, and if exactly two, then nothing else. Moreover, τ_2 must contain at least as many 4-s as it contains $2k$ -s, as otherwise we would find two adjacent $2k$ -faces while P cannot contain a $(2k, 2k)$ -edge by **Corollary 6.20**. We are therefore left with the following options for τ_2 :

$$(4, 4, 2k), (4, 4, 4, 2k) \text{ and } (4, 2k, 4, 2k).$$

The case $\tau_2 = (4, 4, 2k)$ is impossible by **Corollary 6.22**. We show that $\tau_2 = (4, 4, 4, 2k)$ is also not possible: consider the local neighborhood of a $(4, 4, 4, 2k)$ -vertex (the highlighted vertex):



Since the 1-vertices (black dots) are of type $(4, 4, 6)$, this configuration forces on us the existence of the two gray 6-faces. These two faces intersect in a 2-vertex, which is then incident to two $2k$ -faces and must be of type $(4, 2k, 4, 2k)$. But we can show that the types $(4, 4, 4, 2k)$ and $(4, 2k, 4, 2k)$ are incompatible by **Observation 6.18 (i)**:

$$\beta_2^2 + \beta_2^2 + \beta_2^2 + \beta_2^k \stackrel{(i)}{=} \beta_2^2 + \beta_2^k + \beta_2^2 + \beta_2^k \implies \beta_2^2 = \beta_2^k \stackrel{(6.5)}{\implies} 4 = 2k \geq 6.$$

Thus, $(4, 4, 4, 2k)$ cannot occur.

We conclude that every 2-vertex incident to a $2k$ -face must be of type $(4, 2k, 4, 2k)$. Consider then the following table:

τ_2	$s - 2 - K(\tau_2)$	$\stackrel{?}{<}$	$E(\tau_2)$
$(4, 2k, 4, 2k)$	$1 - 2/k$	$<$	$(\epsilon_2 + \epsilon_2 + \epsilon_k) + \epsilon_2 < 2/k$

The established inequality yields $2k \leq 6$, and hence $2k = 6$. We found that then all 1-vertices must be of type $(4, 4, 6)$, and all 2-vertices incident to a 6-face must be of type $(4, 6, 4, 6)$.

6.4.6 The case $\tau_1 = (4, 4, 6)$

At this point we can now assume that all 1-vertices of P are of type $(4, 4, 6)$ and that each 2-vertex of P that is incident to a 6-face is of type $(4, 6, 4, 6)$. In particular, P contains a 2-vertex $w \in V_2$ of this type. Since there is no $(6, 6)$ -edge in P , the two 6-faces incident to w cannot be adjacent. In other words, the faces around w must occur alternatingly of type 4 and type 6, which is the reason for writing the type $(4, 6, 4, 6)$ with alternating entries.

On the other hand, P contains $(4, 4)$ -edges, and none of these is incident to a $(4, 6, 4, 6)$ -vertex surrounded by alternating faces. Thus, there must be further 2-vertices of a type other than $(4, 6, 4, 6)$, necessarily *not* incident to any 6-face. These must then be of type

$$(4^r) := \underbrace{(4, \dots, 4)}_r, \quad \text{for some } r \geq 3.$$

Proposition 6.30. $r = 5$.

Proof. If there is a (4^r) -vertex, **Observation 6.18** (i) yields $\beta_2^2 = 2\pi/r$. Analogously, from the existence of a $(4, 6, 4, 6)$ -vertex follows

$$2\beta_2^2 + 2\beta_2^3 \stackrel{(i)}{=} 2\pi \implies \beta_2^3 = \frac{2\pi - 2\beta_2^2}{2} = \left(1 - \frac{2}{r}\right)\pi.$$

Recall $k\beta_1^k + k\beta_2^k > 2\pi(k-1)$ from **Observation 6.18** (ii). Together with the previously established values for β_2^2 and β_2^3 , this yields

$$\beta_1^2 > \frac{2\pi(2-1) - 2\beta_2^2}{2} = \left(1 - \frac{2}{r}\right)\pi \quad \text{and} \quad \beta_1^3 > \frac{2\pi(3-1) - 3\beta_2^3}{3} = \left(\frac{1}{3} + \frac{2}{r}\right)\pi. \quad (6.9)$$

Since the 1-vertices are of type $(4, 4, 6)$, **Observation 6.18** (i) yields

$$2\pi \stackrel{(i)}{=} 2\beta_1^2 + \beta_1^3 \stackrel{(6.9)}{>} 2\left(1 - \frac{2}{r}\right)\pi + \left(\frac{1}{3} + \frac{2}{r}\right)\pi = \left(\frac{7}{3} - \frac{2}{r}\right)\pi.$$

And one checks that this rearranges to $r < 6$.

This leaves us with the options $r \in \{3, 4, 5\}$. If $r = 4$, then $\beta_2^3 = \pi/2 = \beta_2^2$, which is impossible by equation (6.5). And if $r = 3$, then (6.9) yields $\beta_1^3 > \pi$, which is also impossible for a convex face of a spherical polyhedron. We are left with $r = 5$. \square

6.5 An almost strictly bipartite polyhedron

It turns out that the restrictions obtained so far leave us with a unique candidate polyhedron left to be investigated. It is the purpose of this section to prove that this polyhedron is *not* strictly bipartite. Instead it will turn out to be a remarkable near-miss.

To summarize, we found that P is strictly bipartite with all 1-vertices of type $(4, 4, 6)$ and all 2-vertices of types $(4, 6, 4, 6)$ or (4^5) , and both types actually occur. This information uniquely determines the edge-graph of P , which is shown in **Figure 6.3**.





Figure 6.3. The edge-graph of the final candidate polyhedron.

This graph can be constructed by starting from a hexagon in the center with vertices of alternating colors (indicating the partition classes). One then successively adds further faces (according to the structural properties determined above), layer by layer. This process involves no choice and thus the result is unique.

As mentioned in Remark 6.3, a bipartite polyhedron has an edge in-sphere. Thus, P is a polyhedral realization of the graph in Figure 6.3 with an edge in-sphere. It is known that any two such realizations are related by a projective transformation [65]. One representative from this class (which we do not yet know to coincide with P) can be constructed by applying the following operation \star to each vertex of the regular icosahedron:



The operation is performed so that one vertex of each new 4-gon is positioned in the center of an edge of the icosahedron, and so that its edges are tangent to a common sphere centered at the origin (assuming that the icosahedron is centered at the origin). This results in the polyhedron shown in Figure 6.4.




Figure 6.4. The final candidate polyhedron.


One can check, that this polyhedron has indeed the edge-graph depicted in Figure 6.3.

Evidently, all 4-gonal faces of this polyhedron are mutually congruent, and so are the 6-gonal faces (as we would expect for a bipartite polyhedron). Since P has an edge in-sphere and the same edge-graph as this “modified icosahedron”, it must be a projective transformation thereof. However, any projective transformation, that is not just a re-orientation or a uniform rescaling will inevitably destroy the property of congruent faces. In conclusion, P must be the polyhedron in Figure 6.4.

By construction, P has a bipartite edge-graph and an edge in-sphere. Moreover, Figure 6.4 gives the impression that P has all edges of the same length and should therefore be bipartite. However, we shall see that the edges must have tiny differences in their lengths that cannot be spotted by visual inspection. The author thanks Frank Göring who provided the argument that is presented in the remainder of this section.


We proceed by contradiction: we assume that P has all edges of the same length and show that its 2-vertices $v \in V_2$ must have different values for $\|v\|$ depending on their type. That is, we cannot choose a consistent value for r_2 and P cannot be bipartite.

Consider the following well-known construction of the regular icosahedron from the cube of edge-length 2 (centered at the origin):




In words: insert a line segment in the center of each face of the cube as shown in the image. Each line segment is of length 2ϕ , where $\phi \approx 0.61803$ is the positive solution of $\phi^2 = 1 - \phi$ (one of the numbers commonly known as the *golden ratio*). The convex hull of these line segments results in a regular icosahedron with edge length 2ϕ .

It is now sufficient to consider a single vertex of the icosahedron together with its incident faces. The image below shows this vertex after we applied the operation \star .



The image on the right is the orthogonal projection of the configuration on the left onto the yz -plane. In this projection we can give explicit 2D-coordinates to several important points:



The points **A** and **C** are 2-vertices of P of type (4^5) and $(4, 6, 4, 6)$ respectively. Both points and the origin **O** are contained in the yz -plane onto which we projected, and so the distances between these points are preserved during the projection. Assuming that P is bipartite, we would expect to find $|\overline{OA}| = |\overline{OC}| = r_2$. We shall see that this is not the case, by explicitly computing the coordinates of **A** and **C** in the new coordinate system (y, z) .

By construction, $C = (0, 1)$ and $|\overline{OC}| = 1$. Other points with easily determined coordinates are **P**, **Q**, **R**, **S**, **T** (the midpoint of **R** and **S**) and **U** (the midpoint of **Q** and **S**).

By construction, the point **B** lies on the line segment \overline{QT} . The parallel projection of a rhombus is a (potentially degenerated) parallelogram, and thus, opposite edges in the projection are still parallel. Hence, the gray edges in the figure are parallel. For that reason, the segment \overline{UB} is parallel to \overline{PQ} . This information suffices to determine the coordinates of **B**, which is now the intersection of \overline{QT} with the parallel of \overline{PQ} through **U**. The coordinates of the intersection are given in the figure.

The rhombus containing the vertices **A**, **B** and **C** degenerated to a line. Its fourth vertex is also located at **B**. Therefore, the segments \overline{CB} and \overline{BA} are translates of each other. Since the point **B** and the segment \overline{CB} are known, this allows the computation of the coordinates of **A** as given in the figure.

We can finally compute $|\overline{OA}|$. For this, recall $(*) \phi^2 = 1 - \phi$, or more generally $(**) \phi^{2n} = F_{2n-2} - \phi F_{2n-1}$, where F_n denotes the n -th Fibonacci number with initial conditions $F_0 = F_1 = 1$. Then

$$\begin{aligned} |\overline{OA}|^2 &= (4\phi - 3)^2 + (1 - 3\phi)^2 \\ &= 25\phi^2 - 30\phi + 10 \\ &\stackrel{(*)}{=} 25(1 - \phi) - 30\phi + 10 \\ &= 35 - 55\phi \\ &= 1 + (34 - 55\phi) \stackrel{(**)}{=} 1 + \phi^{10} > 1, \end{aligned}$$

and thus, P cannot be bipartite. Remarkably, we find that

$$|\overline{\mathbf{OA}}| = \sqrt{1 + \phi^{10}} \approx 1.00405707$$

is only about 0.4% larger than $|\overline{\mathbf{OC}}| = 1$. This explains why the polyhedron in [Figure 6.4](#) cannot be visually distinguished from a bipartite polyhedron. And so, while P is not bipartite, it is a remarkable near-miss.

This finalizes the proof of [Theorem 6.11](#) and establishes the main results of this chapter.

Conclusions and Outlook

This thesis explored two large topics: first, the interplay between symmetry and spectrum for polytopes and graph realizations, and second, transitivity phenomena in convex polytopes.

The common theme over the first part was our search for an appropriate formalization of “sufficient symmetry” in order to answer [Question 3](#) resp. [Question 6](#): is every “sufficiently symmetric” graph realization resp. polytope automatically spectral?

The machinery for formalizing and answering these questions has been developed in stages. We first investigated point arrangements via the so-called arrangement space, which we subsequently linked to concepts like symmetry (cf. [Question 8](#)) and rigidity (cf. [Question 7](#)). We then considered graph realizations and found that in this case distance-transitivity is a sufficient symmetry ([Theorem 2.34](#)). For polytopes we found that combined vertex- and edge-transitivity is sufficient ([Corollary 3.15](#)). In fact, every polytope that is both vertex- and edge-transitive turned out to be θ_2 -spectral, which answered [Question 4](#) in the affirmative. Besides this special case, [Question 1](#) and [Question 2](#) are still widely open, but the Theorem of Izmistiev ([Theorem 3.13](#)) points towards the right tools for attacking these problems.

In the second part we took a closer look at edge-transitive polytopes. We explored their properties via the previously established spectral techniques ([Theorem 4.5](#)), introduced a hierarchical classification scheme for approaching [Question 5](#) and classified several sub-classes, such as the distance-transitive polytopes ([Theorem 4.18](#)) and the edge-transitive polytopes that are not vertex-transitive ([Theorem 6.1](#)). In preparation for these results we obtained a classification of the vertex-transitive zonotopes ([Theorem 5.2](#)) and the uniform zonotopes ([Theorem 5.3](#)).

Among the many new and old open questions that we have addressed in this thesis, we want to close with the following two, which best reflect the progress that we made:

Question 9. Are spectral polytopes characterized by [Theorem 3.14](#)?

Question 10. Are all edge-transitive polytopes in dimension $d \geq 4$ Wythoffian with transitive Coxeter-Dynkin diagrams (cf. [Conjecture 4.11](#))?

Outlook

One of our declared goals was to demonstrate the usefulness of the techniques from spectral graph theory in polytope theory. We believe to have only scratched the surface of this connection. Below we describe two further potential applications.

In a third section we briefly address the many open questions we face in the context of transitivity phenomena in convex and non-convex polytopes.

Graph colorings and polytope symmetries

Every symmetry of a polytope induces a symmetry of its edge-graph, that is,

$$\text{Aut}(P) \subseteq \text{Aut}(G_P).$$

The converse is not always true. But one might suggest to enrich the edge-graph G_P , e.g. by coloring vertices and/or edges, so that the two symmetry groups become isomorphic. Is this possible?




Figure 2. Three 4-gons with a respective coloring of the edge-graph, so that $\text{Aut}(P) \cong \text{Aut}(G_P^*)$. It can be necessary to color vertices and edges.

The Theorem of Izestiev can provide such a coloring: if each edge $ij \in E$ of G_P is colored with X_{ij} (where X is the Izestiev matrix) and each vertex $i \in V$ is colored with X_{ii} , then one can show that the resulting colored edge-graph G_P^* has $\text{Aut}(G_P^*) \cong \text{Aut}_{\text{GL}}(P)$ (where $\text{Aut}_{\text{GL}}(P) \subseteq \text{GL}(\mathbb{R}^d)$ is the *linear* symmetry group of P rather than the Euclidean symmetry group). This is surprising, since there is no a priori reason to believe that the linear symmetries of a polytope can be captured combinatorially by the edge-graph (there are graph realizations whose full symmetry cannot be captured in this way).

Question 11. Can we extend this to a coloring G_P^{**} so that $\text{Aut}(G_P^{**}) \cong \text{Aut}(P)$ (the Euclidean symmetry group)? How about projective symmetries?

Question 12. Do we need to color both vertices and edges if we restrict to dimension $d \geq 3$? That this is necessary in dimension two can be seen in [Figure 2](#).

Many polytopes are eigenpolytopes

We have seen that most polytopes are not spectral. In contrast, many naturally occurring polytopes can be represented as eigenpolytopes.

For example, consider the graph $G = (V, E)$ with $V = \{\text{permutations of } (0, \dots, n)\}$, where two such permutations are adjacent in G if and only if they differ in a transposition of two neighboring entries, e.g.

$$(0, 1, \underline{2}, 3, 4) \text{ is adjacent to } (0, 1, \underline{3}, 2, 4).$$

This graph is also known as the *Bruhat graph* or the *edge-graph of the n -dimensional permutahedron*. In fact, its θ_2 -eigenpolytope is a (non-uniform) realization of the permutahedron.

Things change if we add further edges between any two permutations that are related by a transposition of the “outer entries”, e.g.

$$(\underline{0}, 1, 2, 3, \underline{4}) \text{ is adjacent to } (\underline{4}, 1, 2, 3, \underline{0}).$$

This graph is now arc-transitive and numerical experiments (up to $n = 7$) suggest that the *traveling salesman polytope* is one of its eigenpolytopes (but it is not the θ_2 -eigenpolytope).

Other experiments suggest that the *Birkhoff polytope* can be reproduced in a similar fashion.

Question 13. What other practically relevant polytopes can be represented as eigenpolytopes? What is the practical relevance of such a representation?

Question 14. Can the representation of the traveling salesman polytope be used to give a formulation of the traveling salesman problem in terms of spectral graph theory? What can be learned from this?

General transitivity in polytopes

We noted that general transitivity in convex polytopes is still badly understood and that no classification of δ -transitive polytopes is known for any $\delta \in \{1, \dots, d - 2\}$. It is conceivable that each such class consists only of Wythoffian polytopes (and their polar duals), but it would be completely mysterious for why these degrees of symmetry are related to reflection groups.

Question 15. Is there a δ -transitive polytopes $P \subset \mathbb{R}^d$ for some $\delta \in \{1, \dots, d - 2\}$ that is not Wythoffian or the polar dual of a Wythoffian polytope?

Similar phenomena are expected for the “wild” classes of 0-transitive and $(d - 1)$ -transitive polytopes if we impose additional weak constraints. As shown in [Chapter 5](#), the restriction to centrally symmetric 2-faces results in a collapse of complexity to Γ -permutahedra (which are Wythoffian). Other interesting sub-classes are vertex-transitive polytopes that are

- facet-transitive,
- circumscribed,
- self-dual, or
- simple.

In each case, the only known examples are Wythoffian polytope, their duals and orbit polytopes that are constructed from groups that are closely related to reflection groups (e.g. positive determinant subgroups). But it is rather unclear why this should be the case. Convexity seems to impose stark restrictions, even under comparatively weak symmetry assumptions, and in each case it leads us back to the Wythoffian polytopes:

Question 16. What is the deeper connection of the reflection groups with convexity?

The classification of highly symmetric polytopes can also lead to classification results in less geometric domains. For example, matroids allow several representations by polytopes, one of which is called the *matroid base polytope*, spanned by the characteristic vectors of its bases. The classification of edge-transitive polytopes provides a classification of *base-exchange-transitive matroids*, that is, matroids in which any two base-exchanges are related by a matroid symmetry. The conjectured classification of edge-transitive polytopes would imply that the only matroids of this degree of symmetry are uniform matroids. It might be worthwhile to identify other classes of combinatorial objects that have an intrinsic “convex nature” like matroids and therefore lend themselves to such an investigation.

Convexity might not be the essential restriction in all of this. It is unclear how much richness we gain by switching to the combinatorial domain (e.g. to simplicial or polytopal complexes, or even certain types of partially ordered sets). Replacing convexity by mere compactness allows for still wild but highly symmetric objects, such as the *abstract regular polytopes*. Does this change if we restrict to homology spheres or simplicial spheres?

Question 17. Can we classify the edge-transitive (or for that matter, δ -transitive) simplicial complexes under any of these restrictions?

Theses

1. The existence of spectral polytopes is a curious phenomenon in the intersection of combinatorics, geometry and spectral graph theory that has since evaded a satisfying explanation. This phenomenon can be addressed on multiple levels, most satisfyingly with the techniques of convex geometry.
2. The arrangement space $U := \text{span } \Phi$ of a family of points $v_1, \dots, v_n \in \mathbb{R}^d$, where $\Phi^\top := (v_1, \dots, v_n) \in \mathbb{R}^{d \times n}$, is a unifying tool in the investigation of symmetry, rigidity and spectral properties of point arrangements, graph realizations and polytopes.
3. Symmetry restrictions can be used to develop a fruitful theory of rigidity for point arrangements, graph realizations and polytopes.
4. There is a rich interplay between symmetry properties of graphs and their spectral properties. These are expressed in the theory of spectral graph realizations. The spectrum of a graph influences rigidity properties of graph realizations with prescribed symmetries.
5. There is a connection between geometric convexity and the second-largest eigenvalue θ_2 of a graph. This is most evident from the fact that θ -spectral convex polytopes are known only if $\theta = \theta_2$.
6. Polytopes with combined vertex- and edge-transitivity have outstanding properties: they are always θ_2 -spectral, uniquely determined by their edge-graph, they realize all the symmetries of the edge-graph, have an irreducible symmetry group and metric properties that are tightly linked to the second-largest eigenvalue θ_2 of the edge-graph.
7. Classifying edge-transitive polytopes is a non-trivial task. Still, it is possible to organize edge-transitive polytopes in a hierarchy that makes their structural variety (or lack thereof) more evident. This hierarchical organization allows the identification of several sub-classes for which a classification can be achieved, such as distance-transitive polytopes and edge-transitive polytopes that are not vertex-transitive.
8. The class of edge-transitive polytopes is probably less rich than one would expect from its comparison to the wild class of vertex-transitive polytopes. Edge-transitivity seems to have a deeper connection to the reflection groups and their associated orbit polytopes – the Wythoffian polytopes. The known examples allow the formulation of a precise conjecture for their classification.
9. The investigation of edge-transitivity leads to questions about polytopes with a bipartite edge-graph, such as zonotopes or permutahedra. Having a bipartite edge-graph or even centrally symmetric 2-faces is an unexpectedly strong structural restriction when combined with other properties, such as being vertex-transitive. For example, the vertex-transitive zonotopes can be completely enumerated.

Appendix

A Matrix Groups and Representations

We follow no particular literature on representation theory, but a classical reference on linear representations of finite groups is [69]. All non-trivial statements without a reference are proven below.

A *matrix group* is a set $\Gamma \subseteq \text{GL}(\mathbb{R}^d)$ of invertible matrices closed under matrix multiplication and inverses. In other words, it is a group whose elements are matrices and whose group operation is matrix multiplication. For a permutation group $\Sigma \subseteq \text{Sym}(V)$ on a set $V = \{1, \dots, n\}$, a (*linear*) Σ -*representation* is a group homomorphism $T : \Sigma \rightarrow \text{GL}(\mathbb{R}^d)$. Both concepts, matrix groups and representations, are closely related, use very similar terminology and yield equivalent theorems (after all, the image of a representation is a matrix group). We discuss matrix groups first, and then transfer the terminology and results to representations.

A.1 Matrix groups

Most matrix groups that we discuss are finite subgroups of $\text{O}(\mathbb{R}^d)$, that is, are finite *orthogonal* matrix groups. Likewise, most of our representations $T : \Sigma \rightarrow \text{O}(\mathbb{R}^d)$ are *orthogonal* representations. This comes with no loss of generality.

Lemma A.1. *Every finite matrix group $\Gamma \subset \text{GL}(\mathbb{R}^d)$ is conjugate (hence isomorphic) to an orthogonal matrix group $\Gamma' \subseteq \text{O}(\mathbb{R}^d)$. That is, there is $X \in \text{GL}(\mathbb{R}^d)$ so that*

$$\phi : \Gamma \rightarrow \Gamma', \quad T \mapsto T' := XTX^{-1}$$

defines a group isomorphism.

Proof. If ϕ is well-defined then it is a group isomorphism. It therefore suffices to construct $X \in \text{GL}(\mathbb{R}^d)$ so that XTX^{-1} is orthogonal for all $T \in \Gamma$. Since Γ is finite, we can set

$$\bar{X} := \sum_{T \in \Gamma} T^{\top} T.$$

For all $x \in \mathbb{R}^d \setminus \{0\}$ holds

$$x^{\top} \bar{X} x = \sum_{T \in \Gamma} x^{\top} T^{\top} T x = \sum_{T \in \Gamma} \|Tx\|^2 > 0.$$

That is, \bar{X} is positive definite, and we can define $X := \bar{X}^{1/2} \in \text{GL}(\mathbb{R}^d)$ as the unique positive

definite matrix with $X^\top X = \bar{X}$. We can then show that XTX^{-1} is orthogonal for all $T \in \Gamma$:

$$\begin{aligned} (XTX^{-1})^\top (XTX^{-1}) &= X^{-\top} T^\top X^\top X T X^{-1} \\ &= X^{-\top} T^\top \left(\sum_{\bar{T} \in \Gamma} \bar{T}^\top \bar{T} \right) T X^{-1} \\ &= X^{-\top} \left(\sum_{\bar{T} \in \Gamma} (\bar{T} T)^\top (\bar{T} T) \right) X^{-1} \\ &= X^{-\top} \left(\sum_{\bar{T} \in \Gamma} \bar{T}^\top \bar{T} \right) X^{-1} = X^{-\top} X^\top X X^{-1} = \text{Id}. \end{aligned}$$

□

From now on we focus on finite orthogonal matrix groups.

Definition A.2. Let $\Gamma \subseteq O(\mathbb{R}^d)$ be a matrix group.

(i) A subspace $U \subseteq \mathbb{R}^d$ is called Γ -invariant if it is set-wise fixed by all $T \in \Gamma$, that is,

$$Tx \in U, \quad \text{for all } x \in U \text{ and } T \in \Gamma.$$

The subspaces $\{0\}, \mathbb{R}^d \subseteq \mathbb{R}^d$ are always invariant, and are called *trivial* invariant subspaces of Γ .

(ii) A Γ -invariant subspace $U \subseteq \mathbb{R}^d$ is called Γ -irreducible if it has no non-zero Γ -invariant subspace. Otherwise it is called Γ -reducible.

(iii) Γ is called *irreducible* if there are no non-trivial Γ -invariant subspaces, or equivalently, if \mathbb{R}^d is Γ -irreducible. Otherwise it is called *reducible*.

We often write invariant, irreducible, etc. if the group is clear from context.

Observation A.3. Whenever $U, U' \subseteq \mathbb{R}^d$ are Γ -invariant, so is $U \cap U'$. Since $U \cap U' \subseteq U, U'$, if at least one of U or U' is irreducible then either $U = U'$ or $U \cap U' = \{0\}$.

Observation A.4. If $U \subseteq \mathbb{R}^d$ is Γ -invariant (and Γ is an orthogonal matrix group), then the orthogonal complement $U^\perp \subseteq \mathbb{R}^d$ is Γ -invariant as well. This is special to *orthogonal* matrix groups and is not true for general matrix groups $\Gamma \subseteq \text{GL}(\mathbb{R}^d)$.

Applied recursively, we can see that \mathbb{R}^d decomposes as a direct sum

$$\mathbb{R}^d = U_1 \oplus \cdots \oplus U_K$$

of pair-wise orthogonal Γ -irreducible subspaces $U_i \subseteq \mathbb{R}^d$. Note however that this decomposition is not necessarily unique.

Observation A.5. If $U, U' \subseteq \mathbb{R}^d$ are Γ -invariant and π_U denotes the orthogonal projection onto U , then

$$\pi_U U' := \{\pi_U u \mid u \in U'\} \subseteq U$$

is again Γ -invariant. This follows essentially from π_U commuting with all $T \in \Gamma$ (one checks easily that they commute on U and U^\perp respectively).

If U is irreducible and since $\pi_U U' \subseteq U$, we have either $\pi_U U' = U$ or $\pi_U U' = \{0\}$. Likewise, if U' is irreducible, so is $\pi_U U'$, as proven in [Proposition A.6](#) below.

Proposition A.6. *If $U, U' \subseteq \mathbb{R}^d$ are Γ -invariant and U' is irreducible, then $\pi_U U'$ is irreducible as well.*

Proof. Let $U'' \subseteq \pi_U U'$ be a Γ -invariant subspace. Since U' is irreducible, the projection $\pi_{U'} U''$ is either U' or $\{0\}$ (by [Observation A.5](#)). In the first case we find $\dim U' = \dim \pi_{U'} U'' \leq \dim U'' \leq \dim \pi_U U' \leq \dim U'$, and hence $U'' = \pi_U U'$. In the second case we have $U'' \perp U'$, and if $u \in U'' \setminus \{0\}$, then $u \perp U'$, in contradiction to $u \in \pi_U U'$, and therefore $U'' = \{0\}$. In any case, U'' is a trivial invariant subspace of $\pi_U U'$, and $\pi_U U'$ is irreducible. \square

A.2 Representations

The image of a representation $T: \Sigma \rightarrow \text{O}(\mathbb{R}^d)$ (we restrict to orthogonal representations) is a matrix group. If $\Sigma \cong \text{im } T$ (that is, if T is injective), the representation is said to be *faithful*.

The terminology introduced for matrix groups translates to representations: a subspace $U \subseteq \mathbb{R}^d$ is T -invariant resp. T -irreducible if it is invariant resp. irreducible w.r.t. the image $\text{im } T$. T is said to be irreducible if its image is irreducible as a matrix group.

Every permutation group $\Sigma \subseteq \text{Sym}(V)$ has a canonical orthogonal representation $\Sigma \ni \sigma \mapsto \Pi_\sigma \in \text{Perm}(\mathbb{R}^n)$ by permutation matrices. Invariant and irreducible subspaces w.r.t. this representation are called Σ -invariant resp. Σ -irreducible.

Definition A.7. Given two representations $T: \Sigma \rightarrow \text{O}(\mathbb{R}^d)$ and $T': \Sigma \rightarrow \text{O}(\mathbb{R}^{d'})$.

(i) An *equivariant map* (or *intertwining map*) is a transformation $R \in \mathbb{R}^{d \times d'}$ with

$$T_\sigma R = R T'_\sigma, \quad \text{for all } \sigma \in \Sigma,$$

or $TR = RT'$ for short.

(ii) The set of equivariant maps between T and T' is denoted $\text{Hom}(T, T')$.

If $T = T'$ we define the so-called *endomorphism ring* $\text{End}(T) := \text{Hom}(T, T)$. This set is naturally equipped with the structure of an \mathbb{R} -algebra, that is, addition, multiplication and \mathbb{R} -scalar multiplication.

(iii) If there exists an invertible $R \in \text{Hom}(T, T')$, then T and T' are said to be isomorphic representations.

One of the most important tools in representation theory is known as *Schur's lemma*. We state a version of Schur's Lemma for real orthogonal representations:

Theorem A.8 (Schur's Lemma). *Suppose that $T: \Sigma \rightarrow \text{O}(\mathbb{R}^d)$ is an irreducible representation.*

(i) *If $T': \Sigma \rightarrow \text{O}(\mathbb{R}^d)$ is another Σ -representation with $R \in \text{Hom}(T, T')$, then $R = \alpha R'$ for some $\alpha \in \mathbb{R}$ and $R' \in \text{O}(\mathbb{R}^d)$. In particular, either $\text{Hom}(T, T') = \{0\}$, or T and T' are isomorphic.*

(ii) *$\text{End}(T)$ is a division algebra over \mathbb{R} , that is, isomorphic to either \mathbb{R} , \mathbb{C} or \mathbb{H} (the quaternions).*

Proof. Since T' is an orthogonal transformation, we have $T_\sigma^\top = T_\sigma^{-1} = T_{\sigma^{-1}}$. Then

$$T_\sigma R R^\top = R T'_\sigma R^\top = R (R T'_{\sigma^{-1}})^\top = R (T_{\sigma^{-1}} R)^\top = R R^\top T_\sigma.$$

That is, $R R^\top$ commutes with T_σ for all $\sigma \in \Sigma$, and the T_σ must preserve the eigenspaces of $R R^\top$. But since T is irreducible, there are no non-zero proper T -invariant subspaces, and $R R^\top$ must have a single eigenspace to some eigenvalue $\theta \in \mathbb{R}$. Thus, $R R^\top = \theta \text{Id}$. Note that $R R^\top$ is positive semi-definite, hence $\theta \geq 0$. Set $\alpha = \theta^{1/2}$. If $\alpha = 0$, then (i) follows trivially. If $\alpha > 0$, then set $R' = \alpha^{-1} R$, so that $R' R'^\top = \text{Id}$ and R' is orthogonal. Then $R = \alpha R'$, which proves (i).

We know that $\text{End}(T)$ is an \mathbb{R} -algebra. It remains to show that it permits division by $X \in \text{End}(T) \setminus \{0\}$. But given that all elements in $\text{End}(T) \setminus \{0\}$ are positive multiples of orthogonal matrices (by (i)), this is obvious, and (ii) follows. The fact that a real division algebra is isomorphic to either \mathbb{R} , \mathbb{C} or \mathbb{H} is a standard result and known as the *Frobenius theorem* (see e.g. [61]). \square

We close this section with a result about non-orthogonal irreducible subspaces. Two subspaces $U, U' \subseteq \mathbb{R}^d$ are said to be *non-orthogonal* if $U' \not\subseteq U^\perp$ and vice versa. Furthermore, if $U \subseteq \mathbb{R}^d$ is an invariant subspace of a representation T , then we can consider the restricted representation $T|_U: \Sigma \rightarrow \text{O}(U)$, mapping $\sigma \in \Sigma$ to T_σ interpreted as a map $U \rightarrow U$.

Lemma A.9. *Given a Σ -representation T , if $U, U' \subseteq \mathbb{R}^d$ are non-orthogonal T -irreducible subspaces, then*

- (i) $\dim(U) = \dim(U')$, and
- (ii) *the restrictions $T|_U$ and $T|_{U'}$ are isomorphic representations.*

Proof. As mentioned in **Observation A.5**, $\pi_U: U' \rightarrow U$ commutes with T_σ for all $\sigma \in \Sigma$, and so does $\pi_{U'}: U \rightarrow U'$. Then $\pi_U \pi_{U'} \in \text{End}(T|_U)$, and by Schur's lemma (**Theorem A.8**), this map is either zero or invertible. Since U and U' are non-orthogonal it cannot be zero and must be invertible, and so must be $\pi_U: U' \rightarrow U$. Thus, $\dim(U) = \dim(U')$, proving (i), and π_U is an invertible equivariant map between $T|_U$ and $T|_{U'}$, proving (ii). \square

B Spectral Graph Theory

A general introduction into spectral graph theory can be found in [15]. Aspects of spectral graph theory that are particularly focused on highly symmetric or otherwise highly structured graphs can be found in [30, Chapter 8 - 13].

Let $G = (V, E)$ be a finite simple graph on the vertex set $V = \{1, \dots, n\}$. *Spectral graph theory* studies graphs via the spectral properties of associated matrices, such as the *adjacency matrix* $A \in \{0, 1\}^{n \times n}$ with

$$A_{ij} := [ij \in E] = \begin{cases} 1 & \text{if } ij \in E \\ 0 & \text{if } ij \notin E \end{cases},$$

or the *Laplace matrix* $L \in \mathbb{Z}^{n \times n}$ with

$$L_{ij} := \begin{cases} -1 & \text{if } ij \in E \\ \deg_G(i) & \text{if } i = j \\ 0 & \text{otherwise} \end{cases},$$

where $\deg_G(i)$ denotes the vertex degree of $i \in V$.

In this context, notions like spectrum, eigenvalue, eigenvector and eigenspace of a graph are then meant to refer to the respective quantities of an associated matrix. For our purpose this will mostly be the adjacency matrix. The respective quantity of the Laplace matrix shall be called *Laplacian eigenvalue*, *Laplacian eigenvector*, etc.

Other matrices have been considered in this regard, such as the *signless* and *normalized Laplace matrix*, or more general matrices $M \in \mathbb{R}^{n \times n}$ with $M_{ij} = 0$ whenever $ij \notin E$ (a precursor of so-called *discrete Schrödinger operators*).

Of particular interest are *symmetric* matrices (such as adjacency and Laplace matrix), as for these all eigenvalues are real and can be enumerated in a standard order. For example, the eigenvalues of the adjacency matrix are denoted $\theta_1 > \theta_2 > \dots > \theta_m$ in decreasing order. The set of all these eigenvalues (the *spectrum*) is denoted $\text{Spec}(G)$. An arbitrary element of $\text{Spec}(G)$ will be denoted θ (without a subscript). In contrast, the Laplacian eigenvalues are enumerated in increasing order: $0 = \lambda_1 < \lambda_2 < \dots < \lambda_m$ (where m might be different from the adjacency case), and an arbitrary Laplacian eigenvalue is denoted λ . Note in particular that all eigenvalues of L are non-negative and that L is therefore a positive semi-definite matrix. Moreover, the smallest eigenvalue is always zero and its multiplicity equals the number of connected components of G .

Matters further simplify if we restrict to regular graphs (all vertices are of the same degree, which we denote by $\deg(G)$). Then, adjacency matrix and Laplace matrix are related by

$$L = \deg(G)\text{Id} - A,$$

and their eigenvalues satisfy $\lambda_i = \deg(G) - \theta_i$. Moreover, each eigenvector of A to θ_i is an eigenvector of L to λ_i . Since we often restrict to regular graphs, all construction that are based on the eigenvectors (such as the spectral realizations) do therefore not depend on the choice of matrix.

C Polytope Theory

We follow the modern treatment in [79]. For a more classical reference consider [35].

A d -dimensional (*convex*) *polytope* $P \subset \mathbb{R}^d$ is the convex hull of finitely many points in Euclidean space. In particular, P is a compact convex set. We often assume that P is *full-dimensional*, that is, $\text{aff} P = \mathbb{R}^d$.

A convex subset $F \subseteq P$ is called *face* of P , if for any two $x, y \in P$ holds

$$\text{if } tx + (1-t)y \in F \text{ for some } t \in (0, 1), \text{ then } x, y \in F.$$

The faces of P are themselves polytopes, and $\mathcal{F}(P)$ denotes the set of all faces. The faces are partially ordered by inclusion and $\mathcal{F}(P)$ can be considered as a partially ordered set. In fact, $\mathcal{F}(P)$ carries the structure of a lattice (each subset has a least upper bound and a greatest lower bound) which justifies the name *face lattice*. It is a bounded lattice with a least and a greatest element $\emptyset, P \in \mathcal{F}(P)$ respectively (the *trivial faces*). Two polytopes are said to be *combinatorially equivalent* if their face lattices are isomorphic.

A face is said to be δ -dimensional (or is called a δ -*face* of P) if its affine hull is of dimension δ (the dimension of \emptyset is defined to be -1). By $\mathcal{F}_\delta(P) \subseteq \mathcal{F}(P)$ we denote the set of δ -dimensional faces of P . The 0-faces are called *vertices*, the 1-faces are called *edges*, and the $(d-1)$ -faces are called *facets* of P .

Except for the empty face, all faces of a polytope can be defined by *supporting vectors*:

Proposition C.1. *For $f \subseteq P$ the following are equivalent:*

- (i) f is a non-empty face of P .
- (ii) f is the set of maximizers of $\langle c, \cdot \rangle$ in P , where $c \in \mathbb{R}^d$ is called a supporting vector of f .

The *edge-graph* G_P of P is (isomorphic to) the graph with vertex set $\mathcal{F}_0(P)$, where $v, w \in \mathcal{F}_0(P)$ are adjacent if they are the end vertices of an edge of P , that is, if $\text{conv}\{v, w\} \in \mathcal{F}_1(P)$. If G_P is a regular graph and its degree matches $\dim \text{aff} P$, then P is called a *simple polytope*.

The *skeleton* (or *1-skeleton* of P) usually denotes the union of the elements of $\mathcal{F}_1(P)$. We shall use a different definition: the skeleton is a map $\text{sk}_P: G_P \rightarrow \mathcal{F}_0(P)$ mapping each vertex of the edge-graph to its respective point in space (this is the identity if the edge-graph is defined as above, but the edge graph can also be defined on a different vertex set).

If $0 \in \text{int} P$, then the *polar dual* of P ,

$$P^\circ := \{x \in \mathbb{R}^d \mid \langle x, v \rangle \leq 1 \text{ for all } v \in \mathcal{F}_0(P)\},$$

is again a polytope. In fact, the face lattice $\mathcal{F}(P^\circ)$ is isomorphic to the inverse lattice of $\mathcal{F}(P)$ (that is, with reversed lattice order). Each δ -dimensional face $f \in \mathcal{F}(P)$ has an associated $(d-\delta-1)$ -dimensional *dual face* in P° .

The (Euclidean) symmetry group $\text{Aut}(P) \subseteq \text{O}(\mathbb{R}^d)$ is the set of orthogonal matrices that fix P set-wise. It holds $\text{Aut}(P) \cong \text{Aut}(P^\circ)$.

One of the earliest results on polytopes concerns their rigidity. The 3-dimensional version of the following theorem was proven by Cauchy (a version of the proof can be found in [1, Section 12]). The higher dimensional analog is a result by Alexandrov (proven e.g. in [60, Theorem 27.2]).

Theorem C.2 (Cauchy's rigidity theorem). *Let $P_i \subset \mathbb{R}^{d_i}$, $i \in \{1, 2\}$, $d_i \geq 3$ be two combinatorially equivalent polytopes, and $\phi : \mathcal{F}(P_1) \rightarrow \mathcal{F}(P_2)$ a corresponding face lattice isomorphism. If each (proper) face $\sigma \in \mathcal{F}(P_1) \setminus \{P_1\}$ is congruent to the corresponding face $\phi(\sigma) \in \mathcal{F}(P_2)$, then P_1 and P_2 are congruent.*

Two subsets of \mathbb{R}^d are *congruent* (or *isometric*) if there exists a distance-preserving bijection (or *isometry*) between them.

The following is an equivalent version that follows by induction on the dimension:

Corollary C.3. *Let $P_i \subset \mathbb{R}^{d_i}$, $i \in \{1, 2\}$ be two combinatorially equivalent polytopes, and let $\phi : \mathcal{F}(P_1) \rightarrow \mathcal{F}(P_2)$ be a corresponding face lattice isomorphism. If*

- (i) *each edge $e \in \mathcal{F}_1(P_1)$ has the same length as $\phi(e) \in \mathcal{F}_1(P_2)$, and*
- (ii) *each 2-face $f \in \mathcal{F}_2(P_1)$ is congruent to $\phi(f) \in \mathcal{F}_2(P_2)$,*

then P_1 and P_2 are congruent.

C.1 Zonotopes

Zonotopes are a special class of polytopes that have importance in various subfields of geometry, combinatorics and algebra. Their omnipresence is partially explained by their many equivalent definitions:

Definition C.4. A *zonotope* $Z \subseteq \mathbb{R}^d$ is a polytope that satisfies any (and then all) of the following equivalent conditions:

- (i) Z is the projection of a cube.
- (ii) Z is the Minkowski sum of (finitely many) line segments.
- (iii) Z has only centrally-symmetric faces.
- (iv) Z has only centrally-symmetric 2-faces.

The equivalence of these definitions is well-established, but some directions are far from obvious (see [54] for the direction (iv) \implies (i),(iii), or the references in [79, Section 7.3] for a general overview).

Some properties of zonotopes are immediately evident from these definitions. For example, by Definition C.4 (iii) (and since a polytope is a face of itself), zonotopes are centrally symmetric.

Let us assume that $Z \subset \mathbb{R}^d$ is a full-dimensional zonotope. By central symmetry, we may assume $Z = -Z$. From **Definition C.4** (ii) then follows that Z can be written as

$$Z = \text{Zon}(R) := \sum_{r \in R} \text{conv}\{0, r\} \quad (\text{C.1})$$

for some finite centrally symmetric set $R \subseteq \mathbb{R}^d$. This set might not be unique. It is clear from **(C.1)** that $\dim \text{span } Z = \dim \text{span } R$. We also see that every point $v \in Z$ can be written as

$$v = \sum_{r \in R} a_r r, \quad \text{with } a \in [0, 1]^R. \quad (\text{C.2})$$

This representation is, in general, not unique.

The combinatorial structure of a zonotope is clearly encoded in its generating sets and it can be easily extracted.

Proposition C.5. *If $Z = \text{Zon}(R)$ and $f \in \mathcal{F}(Z)$ is a non-empty face of Z , then there is a unique partition $R = R_- \cup R_0 \cup R_+$ into disjoint sets, so that*

$$f = \text{Zon}(R_0) + \sum_{r \in R_+} r. \quad (\text{C.3})$$

Moreover, if $c \in \mathbb{R}^d$ is a supporting vector for f , then

$$R_0 := \{r \in R \mid \langle r, c \rangle = 0\} = R \cap c^\perp \quad \text{and} \quad R_\pm := \{r \in R \mid \pm \langle r, c \rangle > 0\}. \quad (\text{C.4})$$

In particular, these sets are independent of the choice of the supporting vector c .

Proof. By **(C.2)** each point $v \in Z$ can be written in the form

$$v = \sum_{r \in R} a_r r, \quad \text{for some coefficients } a_r \in [0, 1].$$

Fix a partition $R = R_- \cup R_0 \cup R_+$. Consider the following two possible sets of constraints for the coefficients:

$$(*) a_r = \begin{cases} 0 & \text{if } \langle r, c \rangle < 0 \\ 1 & \text{if } \langle r, c \rangle > 0, \\ \text{arbitrary} & \text{if } \langle r, c \rangle = 0 \end{cases}, \quad (**) a_r = \begin{cases} 0 & \text{if } r \in R_- \\ 1 & \text{if } r \in R_+ \\ \text{arbitrary} & \text{if } r \in R_0 \end{cases}.$$

Since c is a supporting vector of the face f , v lies in f if and only if v maximizes the functional $\langle c, \cdot \rangle$, which is the case if and only if the coefficients a_r satisfy **(*)**.

On the other hand, a_r satisfying **(**)** is equivalent to

$$v \in \left\{ \sum_{r \in R_0} \bar{a}_r r + \sum_{r \in R_+} r \mid \bar{a} \in [0, 1]^{R_0} \right\} = \text{Zon}(R_0) + \sum_{r \in R_+} r.$$

Thus, **(C.3)** holds if and only if the two constraints describe the same set of coefficients. Comparing the constraints **(*)** and **(**)** shows that this happens exactly when the partition is chosen as in **(C.4)**. \square

Corollary C.6. *The (non-empty) faces of a zonotope are zonotopes.*

We mentioned previously that there can be multiple sets that generate the same zonotope. However, there is only a single *reduced* generating set: a centrally symmetric set $R \subseteq \mathbb{R}^d$ is *reduced* if $R \cap \text{span}\{r\} = \{\pm r\}$ for all $r \in R$.

Proposition C.7. *Given a zonotope Z , there is a unique finite, centrally symmetric and reduced set $R \subset \mathbb{R}^d$ with $Z = \text{Zon}(R)$. This set can be explicitly constructed as*

$$R = \text{Gen}(Z) := \{r \in \mathbb{R}^d \mid \text{conv}\{\pm r\} \text{ is the translate of an edge of } Z\}.$$

The elements of $\text{Gen}(Z)$ are called the (standard) generators of Z .

Proof. It is easy to see that $\text{Gen}(Z)$ is finite, centrally symmetric and reduced.

Suppose that $R \subset \mathbb{R}^d$ is another finite, centrally symmetric and reduced set which satisfies $\text{Zon}(R) = Z$. Since R is reduced, $r \in R$ if and only if there is a vector $c \in \mathbb{R}^d$ with $R \cap c^\perp = \{\pm r\}$. By **Proposition C.5** this is equivalent to the face of Z with supporting vector c being a translate of $\text{Zon}(\{\pm r\}) = \text{conv}\{\pm r\}$, which by dimension consideration must be an edge. And this is equivalent to $r \in \text{Gen}(Z)$. This shows $R = \text{Gen}(Z)$. \square

The generators of a zonotope are invariant under translation of the zonotope and they are well-defined for zonotopes that are not centered at the origin (but $\text{Zon}(\text{Gen}(Z)) = Z$ might no longer apply). In particular, they are well-defined for the faces of a zonotope.

Corollary C.8. *If $f \in \mathcal{F}(Z)$ is a non-empty face of Z with support vector $c \in \mathbb{R}^d$, then*

$$\text{Gen}(f) = \text{Gen}(Z) \cap c^\perp.$$

Proof. By **Proposition C.5** f is a translate of $\text{Zon}(R_0)$ with $R_0 := \text{Gen}(Z) \cap c^\perp$. \square

Definition C.9. Let $R \subset \mathbb{R}^d$ be a finite centrally-symmetric set:

(i) a subset $S \subset R$ is called *semi-star* of R if it is the intersection of R with a half-space that contains exactly half the elements of R .

In particular, S contains exactly one element from each subset $\{\pm r\} \subseteq R$.

(ii) a subset $F \subseteq R$ is called a *flat* of R if it is the intersection of R with a linear subspace, or equivalently, if $F = R \cap \text{span } F$.

Lemma C.10. *For $F \subseteq \text{Gen}(Z)$ the following are equivalent:*

- (i) F is a flat,
- (ii) $F = \text{Gen}(f)$ for some non-empty face $f \in \mathcal{F}(Z)$.

Proof. If F is a flat, then it can be written as $F = \text{Gen}(Z) \cap c^\perp$ for some $c \in \mathbb{R}^d$. Let $f \in \mathcal{F}(Z)$ be the face with support vector c . Then $F = \text{Gen}(f)$ by **Corollary C.8**.

Conversely, let $f \in \mathcal{F}(Z)$ be a non-empty face with $F = \text{Gen}(f)$, and let $c \in \mathbb{R}^d$ be a support vector of f . Then $F = \text{Gen}(f) = \text{Gen}(Z) \cap c^\perp$ by **Corollary C.8**, and F is a flat. \square

Lemma C.11. *The vertices of Z are in one-to-one correspondence with the semi-stars of $\text{Gen}(Z)$: for each semi-star $S \subset \text{Gen}(Z)$*

$$v_S := \sum_{r \in S} r \in \mathcal{F}_0(P)$$

is a vertex of Z . Conversely, for $v \in \mathcal{F}_0(P)$ there is a unique semi-star $S_v \subset \text{Gen}(Z)$ with $v = v_{S_v}$.

Proof. Given $v \in \mathcal{F}_0(P)$, by **Proposition C.5** there are unique sets $R_0, R_+ \subseteq \text{Gen}(Z)$ with

$$v = \text{Zon}(R_0) + \sum_{r \in R_+} r.$$

Since v is a single point, we necessarily have $R_0 = \emptyset$. The set R_+ as defined in **Proposition C.5** is then clearly a semi-star. \square

D Reflection Groups and Root Systems

An introduction to *finite reflection groups* and *root systems* can be found in [38, 43].

D.1 Reflection groups

Definition D.1. A reflection $\rho \in O(\mathbb{R}^d)$ is an orthogonal transformation that satisfies any (and then all) of the following equivalent conditions:

- (i) ρ is a non-identity transformation that fixes a hyperplane point-wise.
- (ii) ρ can be written in the form

$$\rho = \rho_r := \text{Id} - \frac{rr^\top}{\|r\|^2}$$

for some non-zero vector $r \in \mathbb{R}^d$.

- (iii) ρ has spectrum $\{1^{d-1}, (-1)^1\}$.

Definition D.2. A reflection group $\Gamma \subseteq O(\mathbb{R}^d)$ is a matrix group generated by reflections, that is, it can be written in the form

$$\Gamma = \langle \rho_r \mid r \in R \rangle$$

for some set $R \subseteq \mathbb{R}^d \setminus \{0\}$.

We are specifically interested in the *finite* reflection groups. For such, a complete classification is available. The irreducible finite reflection groups are listed below. For now, we list only their standard names, a description for how to construct them will follow in [Appendix D.3](#):

- a 1-dimensional group $I_1 = \{\text{Id}, -\text{Id}\}$,
- for each $n \geq 3$, a 2-dimensional group $I_2(n)$,
- for each $d \geq 3$, three d -dimensional groups A_d, B_d and D_d (with exception $A_3 \cong D_3$),
- further six exceptional groups H_3, H_4, F_4, E_6, E_7 and E_8 in dimensions $d \in \{3, 4, 6, 7, 8\}$.

The subscripts always denote the dimension of the group. The reducible finite matrix groups are obtained as direct sum of the irreducible ones.

Many of these groups are related to known geometric structures, such as the regular polytopes (cf. [Appendix E.3](#)):

D.2 Root systems

Definition D.3. A root system $R \subseteq \mathbb{R}^d$ is a non-empty finite set of vectors, fixed set-wise under reflections ρ_r for all $r \in R$, i.e.,

$$\rho_r R = R, \quad \text{for all } r \in R.$$

The elements of R are called *roots*.

Note that the literature sometimes requires a further “crystallographic condition” that we not include in our definition. Root systems are always centrally symmetric, that is, $R = -R$. Finite reflection groups and root systems are two sides of the same coin:

Definition D.4.

- (i) For a set $R \subset \mathbb{R}^d \setminus \{0\}$ the group $\Gamma(R) := \langle \rho_r \mid r \in R \rangle$ is called the *Weyl group* of R .
- (ii) If $\Gamma \subseteq O(\mathbb{R}^d)$ is a finite reflection group, then $R(\Gamma) := \{r \in S^{d-1} \mid \rho_r \in \Gamma\}$ is called the *root system* of Γ (where $S^{d-1} \subset \mathbb{R}^d$ is the unit sphere in \mathbb{R}^d).

Clearly, $\Gamma(R)$ is a reflection group.

Theorem D.5.

- (i) The Weyl group $\Gamma(R)$ of a root system $R \subset \mathbb{R}^d$ is a finite reflection group.
- (ii) If $\Gamma \subseteq O(\mathbb{R}^d)$ is a finite reflection group, then $R(\Gamma)$ is a root system with Weyl group Γ .

Because of this connection, root systems admit a similar classification as the finite reflection groups, though it is not one-to-one. It becomes one-to-one if we require the root systems to be a set $R \subset S^{d-1}$ of unit vectors.

Most finite reflection groups act transitively on their root systems, that is, for any two $r, s \in R(\Gamma)$ we find a $T \in \Gamma$ with $Tr = s$. Few reflection groups are exceptions to this rule: each of $I_2(2n)$, B_d and F_4 has two orbits on its roots [38, Section 2.11].

Root systems are naturally associated with hyperplane arrangements. To each root system $R \subset \mathbb{R}^d$ we can assign the hyperplane arrangement

$$\mathcal{H} := \{r^\perp \mid r \in R\}.$$

The connected components of $\mathbb{R}^d \setminus \mathcal{H}$ are known as *Weyl chambers* of R resp. of $\Gamma(R)$. The Weyl group $\Gamma(R)$ acts transitively, in fact, regularly, on the Weyl chambers.

D.3 Coxeter-Dynkin diagrams

Coxeter-Dynkin diagrams provide a compact graphical notation for reflection groups (and root systems) that contains enough information to reconstruct the group (or the root system) from the notation alone.

Given angles $\theta_{ij} \in \mathbb{R}$ for $i, j \in \{1, \dots, d\}, i \neq j$, there exists, up to orthogonal transformations, at most one set of unit vectors $r_1, \dots, r_d \in \mathbb{R}^d$ with

$$\angle(r_i, r_j) = \theta_{ij} \quad \text{for all } i \neq j.$$


These vectors can then be used to define a reflection group $\Gamma := \langle T_{r_i} \mid i \in \{1, \dots, d\} \rangle \subseteq O(\mathbb{R}^d)$, and from this, a root system. One can show that in order for this reflection group to be a finite group, the angles must satisfy $\theta_{ij} = \pi/m_{ij}$ for some integers $m_{ij} \geq 2, i \neq j$. Each such set of integers then determines a reflection group, but most of them are still not finite.

An ingenious graphical representation of these integers is given by an edge-labeled complete graph $G = (V, E)$ on d vertices $V = \{1, \dots, d\}$, for which the edge $ij \in E$ is labeled with m_{ij} . For ease of interpretation one usually applies the following two rules:

- if $m_{ij} = 2$, then we do not draw an edge (the graph is not complete anymore).
- if $m_{ij} = 3$, then we do not draw the edge label.

This decorated graph is known as a *Coxeter-Dynkin diagram*, and each one describes a unique reflection group or root system.

Below you can find depictions of the standard diagrams that generate the irreducible finite reflection groups from [Appendix D.1](#):



A Coxeter-Dynkin diagram with d vertices induces a reflection group in d -dimensional space, and so the dimension is immediately evident from the representation. Also, note that every finite d -dimensional reflection group can therefore be generated from d reflections.

The Coxeter-Dynkin diagram of a direct sum of reflections groups Γ and Γ' is the disjoint union of the diagrams of Γ and Γ' respectively. Every reducible finite reflection group can be represented in this way. A reflection group is irreducible if and only if it can be represented by a connected Coxeter-Dynkin diagram.

E Wythoffian, Uniform and Regular Polytopes

E.1 Orbit polytopes and Wythoffian polytopes

Definition E.1. Given an orthogonal matrix group $\Gamma \subseteq O(\mathbb{R}^d)$ and a point $x \in \mathbb{R}^d$ (called the *generator*), the *orbit polytope* $\text{Orb}(\Gamma, x)$ is the convex hull of the orbit Γx , that is

$$\text{Orb}(\Gamma, x) := \text{conv}\{Tx \mid T \in \Gamma\}.$$

Some notes on orbit polytopes:

- In general, the combinatorial type of $\text{Orb}(\Gamma, x)$ depends on the choice of x .
- An orbit polytope is vertex-transitive. In fact, every vertex-transitive polytope is an orbit polytope of its symmetry group.
- It always holds $\Gamma \subseteq \text{Aut}(\text{Orb}(\Gamma, x))$, but we do not always have equality.

An especially important class of orbit polytopes is generated by the finite reflection groups introduced in [Appendix D](#).

Definition E.2. A *Wythoffian polytope* is an orbit polytope of a finite reflection group.

E.2 Coxeter-Dynkin diagrams for Wythoffian polytopes


In [Appendix D.3](#) we introduced a graphical notation, the Coxeter-Dynkin diagram, to denote reflection groups. We can extend this notation to denote combinatorial types of Wythoffian polytopes.

Recall that a Coxeter-Dynkin diagram with d vertices uniquely determines a set of d unit vectors $r_1, \dots, r_d \in \mathbb{R}^d$ (up to orthogonal transformation), and by this, a set of hyperplanes r_i^\perp . Let Γ be the reflection group determined by the diagram. Suppose that it is a finite group and consider the Wythoffian polytopes $\text{Orb}(\Gamma, x)$ with generator $x \in \mathbb{R}^d$.

The hyperplanes r_i^\perp bound one of the Weyl chambers of Γ , and since Γ acts transitively on these chambers, we can assume that x was chosen from this chamber.

It then turns out that the combinatorial type of the Wythoffian polytope $\text{Orb}(\Gamma, x)$ depends only on the set $\{i \mid x \in r_i^\perp\} \subseteq \{1, \dots, d\}$, that is, the set of hyperplanes the generator is contained in.

We can “enrich” Coxeter-Dynkin diagrams to denote this choice of hyperplanes by adding rings around vertices: a vertex $i \in V$ of the diagram is *ringed* if and only if $x \notin r_i^\perp$ (note the negation). One says that the reflection on r_i^\perp is *active*.



Observation E.3. If the Coxeter-Dynkin diagram has *no* ringed vertices, then the generator is on *all* hyperplanes. But the intersection of *all* hyperplanes is only the origin. The resulting Wythoffian polytopes is therefore a single point.

Likewise, if not every connected component of the diagram has at least one ringed node, then the polytope collapses in some dimensions and is no longer full-dimensional.

Observation E.4. If the Coxeter-Dynkin diagram has a single ringed vertex, then the generator is on all hyperplanes but one. The intersection of these $d-1$ hyperplanes in d -dimensional space results in a 1-dimensional subspace. The choice of the generator from this 1-dimensional subspace (minus the origin) leaves us to control the scale of the resulting orbit polytope, but gives no further freedom. The resulting Wythoffian polytope is therefore uniquely determined up to scale and orientation.

“Single-ringed Wythoffian polytopes” include the regular polytopes (see [Appendix E.3](#)).

The diagram notation for Wythoffian polytopes has other remarkable properties.

Theorem E.5. *The faces of a Wythoffian polytope $P \subset \mathbb{R}^d$ are themselves Wythoffian.*

The Coxeter-Dynkin diagrams of the $(d-k)$ -dimensional faces of P are obtained by deleting k vertices from the diagram, so that each resulting connected component has at least one ringed vertex.

The “generic” Wythoffian polytopes are a form of generalized *permutahedra*:

Definition E.6. Given a finite reflection group $\Gamma \subseteq O(\mathbb{R}^d)$, a Γ -*permutahedron* is a polytope $P \subset \mathbb{R}^d$ that satisfies any (and then all) of the following equivalent conditions:

- (i) P is a Wythoffian polytope described by a Coxeter-Dynkin diagram in which *all* vertices are ringed.
- (ii) $P = \text{Orb}(\Gamma, x)$, where $x \in \mathbb{R}^d$ is a *generic* generator, that is, x is not fixed by any non-identity transformation of Γ .
- (iii) Γ acts regularly on the vertices of P .

In the literature the term “permutahedron” is often reserved for the *standard permutahedron*, which is the convex hull of the coordinate permutation of $(1, 2, \dots, d+1) \in \mathbb{R}^{d+1}$ (which lives in a d -dimensional affine subspace of \mathbb{R}^{d+1}). Considered as a d -dimensional polytope, this actually coincides with the A_d -permutahedron from [Definition E.6](#).

E.3 Uniform and regular polytopes

Each Wythoffian polytope has a unique realization in which all edges are of the same length. Often the ringed Coxeter-Dynkin diagrams are meant to denote only this special realization, which is then a unique polytope up to scale and orientation.

Wythoffian polytopes of this form are called *Wythoffian uniform polytopes*, and they belong to the larger and historically relevant class of *uniform polytopes*:

Definition E.7. A vertex-transitive polytope $P \subset \mathbb{R}^d$ is called *uniform*, if either

- (i) $d = 2$, and it is a regular polygon, or
- (ii) $d > 2$, and all its facets are uniform polytopes.

The uniform polytopes have been most famously studied by Coxeter [18, 19], as well as by Johnson [40], who coined a lot of their terminology.

Example E.8. The uniform polytopes include the regular polytopes.

In three dimensions they encompass the Platonic and Archimedean solids, the prisms and anti-prisms, most of which are Wythoffian polytopes (cf. **Observation E.10**).

Every Wythoffian polytope described by a connected Coxeter-Dynkin diagram with a *single* ringed vertex is already uniform (cf. **Observation E.4**).

Observation E.9. Given any two uniform polytopes, their cartesian product is again a uniform polytope. Uniform polytopes that can be obtained in this way are called *prismatic*.

The distinction prismatic/non-prismatic can be compared to the distinction reducible/irreducible. However, the analogy is not perfect, as there are non-prismatic uniform polytopes with reducible symmetry groups (e.g. the anti-prisms).

Observation E.10. As mentioned previously, every Wythoffian polytope has a unique realization as a uniform polytope. Only few non-Wythoffian non-prismatic uniform polytopes are known:

- for $d = 3$, the *snub cube* and the *snub dodecahedron* (certain orbit polytopes of the positive determinant subgroups $B_3^+ \subset B_3$ and $H_3^+ \subset H_3$).
- also for $d = 3$, the *anti-prisms* (certain orbit polytopes of the positive determinant subgroups of $I_2(2n) \oplus I_1$).
- for $d = 4$, the *grand anti-prism*, discovered by Conway [17] (a specific orbit polytope of the positive determinant subgroup of $I_2(10) \oplus I_2(10)$).

The classification of uniform polytopes is complete only up to dimension $d = 4$, but open in all higher dimensions. In particular, it is unknown whether there exist non-prismatic non-Wythoffian uniform polytopes in dimension five or above.

Regular polytopes

The famously known *regular polytopes* (including the Platonic solids) emerge as special cases among the uniform polytopes. We call a Coxeter-Dynkin diagram a *string diagram* if it is a single path without any branches.

Definition E.11. A *regular polytope* is a polytope $P \subset \mathbb{R}^d$ that satisfies any (and then all) of the following equivalent conditions:

- (i) P is a Wythoffian (uniform) polytope represented by a string diagram in which exactly one of the end-vertices is ringed.
- (ii) $\text{Aut}(P)$ acts transitively on the δ -dimensional faces for all $\delta \in \{0, \dots, d - 1\}$.

- (iii) $\text{Aut}(P)$ acts transitively on the *flags* of P , where a flag is a chain $f_0 \subset f_1 \subset \dots \subset f_d$ of faces with $f_\delta \in \mathcal{F}_\delta(P)$ being a δ -dimensional face of P .
- (iv) all facets of P are the same regular polytope, and at each vertex of P meet the same number of facets.

The historically first definition of regular polytope is captured in **Definition E.11** (iv). The 2- and 3-dimensional regular polytopes have been known since antiquity (as *regular polygons* and *Platonic solids*). A classification in general dimension was obtained by the Swiss mathematician Ludwig Schläfli already in the 19th century [66]. He obtained the following list, which, given the alternative definitions in **Definition E.11**, can be easily derived from the classification of finite reflection groups:

d	
1	I_1 line segment
2	$I_2(n)$ regular n -gon, $n \geq 3$
3	A_3 tetrahedron B_3 cube B_3 octahedron H_3 dodecahedron H_3 icosahedron
4	A_4 4-simplex (5-cell) B_4 4-cube (8-cell) B_4 4-crosspolytope (16-cell) F_4 24-cell H_4 120-cell H_4 600-cell
≥ 5	A_d d -simplex B_d d -cube B_d d -crosspolytope

Note that in every dimension $d \geq 5$ there exist only three regular polytopes: the simplex, the cube and the crosspolytope.

F Mathematica Scripts

The following Mathematica script takes as input a graph G (in the example below, this is the edge-graph of the dodecahedron, but can be replaced by an arbitrary graph), and an index k of an eigenvalue. It then compute the θ_k -realization (vertex coordinates stored in `vert`). If the dimension turns out to be appropriate, the spectral embedding of the graph, as well as the eigenpolytope are plotted.

```
(* Input:
 * the graph G, and
 * the index k of an eigenvalue (k = 1 being the largest eigenvalue).
*)
G = GraphData["DodecahedralGraph"];
k = 2;

(* Computation of vertex coordinates 'vert' *)
n = VertexCount[G];
A = AdjacencyMatrix[G];
eval = Tally[Sort@Eigenvalues[A//N], Round[#1-#2,0.00001]==0 &];
d = eval[[-k,2]]; (* dimension of the eigenpolytope *)
vert = Transpose@Orthogonalize@
  NullSpace[eval[[-k,1]] * IdentityMatrix[n] - A];

(* Output:
 * the graph G,
 * its eigenvalues with multiplicities,
 * the spectral embedding, and
 * its convex hull (the eigenpolytope).
*)
G
Grid[Join[{{ $\theta$ ,"mult"}}, eval], Frame→All]
Which[
  d<2 , Print["Dimension too low, no plot generated."],
  d==2, GraphPlot[G, VertexCoordinates→vert],
  d==3, GraphPlot3D[G, VertexCoordinates→vert],
  d>3 , Print["Dimension too high, 3-dimensional projection is plotted."];
  GraphPlot3D[G, VertexCoordinates→vert[;;,1;3]] ]
]
If[d==2 || d==3,
  Region`Mesh`MergeCells[ConvexHullMesh[vert]]
]
```


G Additional proofs and computations

We collect some straightforward computations and proofs of well-known facts that would have interrupted the flow of the main text.

G.1 Two geometric statements

Proposition G.1. *Given a set $x_0, \dots, x_d \in \mathbb{R}^d \setminus \{0\}$ of $d+1$ vectors with pair-wise negative inner product, then there are positive coefficients $\alpha_0, \dots, \alpha_d > 0$ with*

$$\alpha_0 x_0 + \dots + \alpha_d x_d = 0.$$

Proof. We proceed by induction. The induction base $d = 1$ is trivially satisfied.

Suppose now $d \geq 2$, and, w.l.o.g. assume $\|x_0\| = 1$. Let π_0 be the orthogonal projection onto x_0^\perp , that is, $\pi_0(u) := u - x_0 \langle x_0, u \rangle$. In particular, for $i \neq j$ and $i, j > 0$

$$\langle \pi_0(x_i), \pi_0(x_j) \rangle = \underbrace{\langle x_i, x_j \rangle}_{<0} - \underbrace{\langle x_0, x_i \rangle}_{<0} \underbrace{\langle x_0, x_j \rangle}_{<0} < 0.$$

Then $\{\pi_0(x_1), \dots, \pi_0(x_d)\}$ is a set of d vectors in $x_0^\perp \cong \mathbb{R}^{d-1}$ with pair-wise negative inner product. By induction assumption there are positive coefficients $\alpha_1, \dots, \alpha_d > 0$ so that $\alpha_1 \pi_0(x_1) + \dots + \alpha_d \pi_0(x_d) = 0$.

Set $\alpha_0 := -\langle x_0, \alpha_1 x_1 + \dots + \alpha_d x_d \rangle > 0$. We claim that $x := x_0 \alpha_0 + \dots + \alpha_d x_d = 0$. Since $\mathbb{R}^d = \text{span}\{x_0\} \oplus x_0^\perp$, it suffices to check that $\langle x_0, x \rangle = 0$ as well as $\pi_0(x) = 0$. This follows:

$$\begin{aligned} \langle x_0, x \rangle &= \alpha_0 \underbrace{\langle x_0, x_0 \rangle}_{=1} + \underbrace{\langle x_0, \alpha_1 x_1 + \dots + \alpha_d x_d \rangle}_{=-\alpha_0} = 0, \\ \pi_0(x) &= \alpha_0 \underbrace{\pi_0(x_0)}_{=0} + \underbrace{\alpha_1 \pi_0(x_1) + \dots + \alpha_d \pi_0(x_d)}_{=0} = 0. \end{aligned}$$

□

An alternative proof of **Proposition G.1** is based on the Perron-Frobenius theorem about matrices with positive entries.

Proposition G.2. *Let $P \subset \mathbb{R}^3$ be a polyhedron and $v \in \mathcal{F}_0(P)$ a vertex of degree three. The interior angles of the faces incident to v determine the dihedral angles at the edges incident to v and vice versa.*

Proof. For $w_1, w_2, w_3 \in \mathcal{F}_0(P)$ the neighbors of v , let $u_i := w_i - v$ denote the direction of the edge e_i from v to w_i . Let f_{ij} be the face that contains v, w_i and w_j . Then $\sphericalangle(u_i, u_j)$ is the interior angle of f_{ij} at v .

The set $\{u_1, u_2, u_3\}$ is uniquely determined (up to some orthogonal transformation) by the angles $\sphericalangle(u_i, u_j)$. Furthermore, since P is convex, $\{u_1, u_2, u_3\}$ forms a basis of \mathbb{R}^3 , and this uniquely determines the dual basis $\{n_{12}, n_{23}, n_{31}\}$ for which $\langle n_{ij}, u_i \rangle = \langle n_{ij}, u_j \rangle = 0$. In other words, n_{ij} is a normal vector to f_{ij} . The dihedral angle at the edge e_j is then $\pi - \sphericalangle(n_{ij}, n_{jk})$, hence uniquely determined. The other direction is analogous, via constructing $\{u_1, u_2, u_3\}$ as the dual basis to the set of normal vectors. \square

G.2 A proof for spherical interior angles

The edge lengths in a spherical polyhedron (see page 125) are measured as angles between its end vertices. Consider adjacent vertices $v_1^S, v_2^S \in \mathcal{F}_0(P^S)$, then the incident edge has (arc-)length $\ell^S := \sphericalangle(v_1^S, v_2^S) = \sphericalangle(v_1, v_2)$.

It follows from **Observation 6.7** that these angles are completely determined by the parameters, hence the same for all edges of P^S .

Proposition G.3. *For a face $f \in \mathcal{F}_2(P)$ and a vertex $v \in \mathcal{F}_0(f)$, there is a direct relationship between the value of $\alpha(f, v)$ and the value of $\beta(f, v)$.*

Proof. Let $w_1, w_2 \in V_2$ be the neighbors of v in the $2k$ -face f , and set $u_i := w_i - v$. Then $\sphericalangle(u_1, u_2) = \alpha(f, v)$. W.l.o.g. assume that v is a 1-vertex (the argument is equivalent for a 2-vertex).

For convenience, we introduce the notation $\chi(\theta) := 1 - \cos(\theta)$. We find that

$$\begin{aligned}
 (*) \quad 2\ell^2 \cdot \chi(\alpha(f, v)) &= \ell^2 + \ell^2 - 2\ell^2 \cos(\sphericalangle(u_1, u_2)) \\
 &= \|u_1\|^2 + \|u_2\|^2 - 2\langle u_1, u_2 \rangle \\
 &= \|u_1 - u_2\|^2 = \|w_1 - w_2\|^2 \\
 &= \|w_1\|^2 + \|w_2\|^2 - 2\langle w_1, w_2 \rangle \\
 &= r_2^2 + r_2^2 - 2r_2^2 \cos \sphericalangle(w_1, w_2) = 2r_2^2 \cdot \chi(\sphericalangle(w_1, w_2)).
 \end{aligned}$$

The side lengths of the spherical triangle $w_1^S v^S w_2^S$ are $\sphericalangle(w_1, w_2)$, ℓ^S and ℓ^S . By the spherical law of cosine¹ we obtain

$$\begin{aligned}
 \cos \sphericalangle(w_1, w_2) &= \cos(\ell^S) \cos(\ell^S) + \sin(\ell^S) \sin(\ell^S) \cos(\beta(f, v)) \\
 &= \cos^2(\ell^S) + \sin^2(\ell^S)(\cos(\beta(f, v)) - 1 + 1) \\
 &= [\cos^2(\ell^S) + \sin^2(\ell^S)] + \sin^2(\ell^S)(\cos(\beta(f, v)) - 1) \\
 &= 1 - \sin^2(\ell^S) \cdot \chi(\beta(f, v)) \\
 \implies \sin^2(\ell^S) \cdot \chi(\beta(f, v)) &= \chi(\sphericalangle(w_1, w_2)) \stackrel{(*)}{=} \left(\frac{\ell}{r_2}\right)^2 \cdot \chi(\alpha(f, v)).
 \end{aligned}$$

¹ $\cos(c) = \cos(a)\cos(b) + \sin(a)\sin(b)\cos(\gamma)$, where a, b and c are the side lengths (arc-lengths) of a spherical triangle, and γ is the interior angle opposite to the side of length c .

□

Bibliography

- [1] M. Aigner, G. M. Ziegler, K. H. Hofmann, and P. Erdős. *Proofs from the Book*, volume 274. Springer, 2010.
- [2] L. Babai. Symmetry groups of vertex-transitive polytopes. *Geometriae Dedicata*, 6(3):331–337, 1977.
- [3] G. Berkolaiko and W. Liu. Eigenspaces of symmetric graphs are not typically irreducible. *Letters in Mathematical Physics*, 108(8):1825–1835, 2018.
- [4] S. Bilinski. Über die Rhombenisoeder. *Glasnik*, 15:251–263, 1960.
- [5] T. Bisztriczky, P. McMullen, R. Schneider, and A. I. Weiss. *Polytopes: abstract, convex and computational*, volume 440. Springer Science & Business Media, 2012.
- [6] A. Björner, P. H. Edelman, and G. M. Ziegler. Hyperplane arrangements with a lattice of regions. *Discrete & Computational Geometry*, 5(3):263–288, 1990.
- [7] R. Blind and P. Mani-Levitska. Puzzles and polytope isomorphisms. *Aequationes Mathematicae*, 34(2-3):287–297, 1987.
- [8] “Blue”. Spectral realizations of graphs. Website: daylateanddollarshort.com/mathdocs/Spectral-Realizations-of-Graphs.pdf (August 2020).
- [9] J. Bokowski, G. Ewald, and P. Kleinschmidt. On combinatorial and affine automorphisms of polytopes. *Israel Journal of Mathematics*, 47(2-3):123–130, 1984.
- [10] S. Boyd, P. Diaconis, and L. Xiao. Fastest mixing Markov chain on a graph. *SIAM review*, 46(4):667–689, 2004.
- [11] A. Brouwer, A. Cohen, and A. Neumaier. Distance-regular graphs. 1989. *Ergeb. Math. Grenzgeb.*(3), 1989.
- [12] F. Buekenhout and M. Parker. The number of nets of the regular convex polytopes in dimension ≥ 4 . *Discrete Mathematics*, 186(1-3):69–94, 1998.
- [13] P. G. Casazza, G. Kutyniok, and F. Philipp. Introduction to Finite Frame Theory. In *Finite frames*, pages 1–53. Springer, 2013.
- [14] A. Chirvasitu, F. Ladisch, and P. Soberón. Non-commutative groups as prescribed polytopal symmetries. *arXiv preprint arXiv:1907.10022*, 2019.

- [15] F. R. Chung and F. C. Graham. *Spectral graph theory*. Number 92. American Mathematical Soc., 1997.
- [16] H. Cohn, N. Elkies, A. Kumar, and A. Schürmann. Point configurations that are asymmetric yet balanced. *Proceedings of the American Mathematical Society*, 138(8):2863–2872, 2010.
- [17] J. H. Conway, H. Burgiel, and C. Goodman-Strauss. *The Symmetries of Things*. CRC Press, 2016.
- [18] H. S. M. Coxeter. *Regular polytopes*. Courier Corporation, 1973.
- [19] H. S. M. Coxeter, M. S. Longuet-Higgins, and J. C. Miller. Uniform polyhedra. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 246(916):401–450, 1954.
- [20] R. Diestel. Graph theory, volume 173 of. *Graduate Texts in Mathematics*, page 7, 2012.
- [21] W. X. Du and Y. Z. Fan. Graph automorphisms from the geometric viewpoint. *arXiv preprint arXiv:1312.2778*, 2013.
- [22] R. Ehrenborg, C. Klivans, and N. Reading. Coxeter arrangements in three dimensions. *Beiträge zur Algebra und Geometrie*, 57(4):891–897, 2016.
- [23] E. Friese. Unitary groups as stabilizers of orbits. *Archiv der Mathematik*, 109(2):101–103, 2017.
- [24] E. Friese. *Generic symmetries of group representations*. PhD thesis, Universität Rostock, 2018.
- [25] E. Friese and F. Ladisch. Affine symmetries of orbit polytopes. *Advances in Mathematics*, 288:386–425, 2016.
- [26] E. Friese and F. Ladisch. Classification of affine symmetry groups of orbit polytopes. *Journal of Algebraic Combinatorics*, 48(3):481–509, 2018.
- [27] R. Frucht. Herstellung von Graphen mit vorgegebener abstrakter Gruppe. *Compositio Mathematica*, 6:239–250, 1939.
- [28] G. Gévay. On perfect 4-polytopes. *Beiträge zur Algebra und Geometrie*, 43(1):243–259, 2002.
- [29] C. Godsil. Euclidean geometry of distance regular graphs. *London Mathematical Society Lecture Note Series*, pages 1–24, 1995.
- [30] C. Godsil and G. F. Royle. *Algebraic Graph Theory*, volume 207. Springer Science & Business Media, 2013.
- [31] C. D. Godsil. Graphs, groups and polytopes. In *Combinatorial Mathematics*, pages 157–164. Springer, 1978.

-
- [32] C. D. Godsil. Eigenpolytopes of distance regular graphs. *Canadian Journal of Mathematics*, 50(4):739–755, 1998.
- [33] F. Göring, C. Helmberg, and M. Wappler. Embedded in the shadow of the separator. *SIAM Journal on Optimization*, 19(1):472–501, 2008.
- [34] T. Gosset. On the regular and semi-regular figures in space of n dimensions. *Messenger of Mathematics*, 29:43–48, 1900.
- [35] B. Grünbaum. *Convex Polytopes*, volume 221. Springer Science & Business Media, 2013.
- [36] B. Grünbaum and G. C. Shephard. Edge-transitive planar graphs. *Journal of Graph Theory*, 11(2):141–155, 1987.
- [37] D. F. Holt. A graph which is edge transitive but not arc transitive. *Journal of Graph Theory*, 5(2):201–204, 1981.
- [38] J. E. Humphreys. *Reflection groups and Coxeter groups*, volume 29. Cambridge University Press, 1990.
- [39] I. Izmetiev. The Colin de Verdière number and graphs of polytopes. *Israel Journal of Mathematics*, 178(1):427–444, 2010.
- [40] N. W. Johnson. *The theory of uniform polytopes and honeycombs*. PhD thesis, University of Toronto, 1967.
- [41] V. Kaibel and A. Wamer. Automorphism groups of cyclic polytopes. 2010.
- [42] G. Kalai. A simple way to tell a simple polytope from its graph. *Journal of combinatorial theory, Series A*, 49(2):381–383, 1988.
- [43] R. Kane. *Reflection groups and invariant theory*. Springer Science & Business Media, 2001.
- [44] C. J. Klivans and E. Swartz. Projection volumes of hyperplane arrangements. *Discrete & Computational Geometry*, 46(3):417, 2011.
- [45] Y. Koren. On spectral graph drawing. In *International Computing and Combinatorics Conference*, pages 496–508. Springer, 2003.
- [46] P. W. Lemmens, J. J. Seidel, and J. Green. Equiangular lines. In *Geometry and Combinatorics*, pages 127–145. Elsevier, 1991.
- [47] C. Licata and D. L. Powers. A surprising property of some regular polytopes. Technical report, Clarkson University Potsdam NY Department of Mathematics and Computer Science, 1986.
- [48] L. Lovász. Spectra of graphs with transitive groups. *Periodica Mathematica Hungarica*, 6(2):191–195, 1975.

- [49] L. Lovász. On the Shannon capacity of a graph. *IEEE Transactions on Information theory*, 25(1):1–7, 1979.
- [50] J. Malestein and L. Theran. Generic rigidity with forced symmetry and sparse colored graphs. In *Rigidity and symmetry*, pages 227–252. Springer, 2014.
- [51] P. Mani. Automorphismen von polyedrischen graphen. *Mathematische Annalen*, 192(4):279–303, 1971.
- [52] H. Martini. A hierarchical classification of Euclidean polytopes with regularity properties. In *Polytopes: abstract, convex and computational*, pages 71–96. Springer, 1994.
- [53] N. Matteo. *Convex polytopes and tilings with few flag orbits*. Northeastern University, 2015.
- [54] P. McMullen. Polytopes with centrally symmetric faces. *Israel Journal of Mathematics*, 8(2):194–196, 1970.
- [55] P. McMullen. *Geometric Regular Polytopes*, volume 172. Cambridge University Press, 2020.
- [56] P. McMullen and E. Schulte. *Abstract Regular Polytopes*, volume 92. Cambridge University Press, 2002.
- [57] T. Mohri. The θ_1 -eigenpolytopes of the Hamming graphs (groups and combinatorics). 1997.
- [58] A. Nixon and B. Schulze. Symmetry-forced rigidity of frameworks on surfaces. *Geometriae Dedicata*, 182(1):163–201, 2016.
- [59] A. Padrol Sureda and J. Pfeifle. Graph operations and Laplacian eigenpolytopes. In *VII Jornadas de Matemática Discreta y Algorítmica*, pages 505–516, 2010.
- [60] I. Pak. Lectures on discrete and polyhedral geometry. *Manuscript* (<http://www.math.ucla.edu/~pak/book.htm>), 2010.
- [61] R. S. Palais. The classification of real division algebras. *The American Mathematical Monthly*, 75(4):366–368, 1968.
- [62] D. L. Powers. The Petersen polytopes. Technical report, Clarkson University Potsdam NY Department of Mathematics and Computer Science, 1986.
- [63] D. L. Powers. Eigenvectors of distance-regular graphs. *SIAM journal on Matrix Analysis and Applications*, 9(3):399–407, 1988.
- [64] B. Rooney. *Spectral Aspects of Cliques in Graphs*. PhD thesis, University of Waterloo, 2014.
- [65] H. Sachs. Coin graphs, polyhedra, and conformal mapping. *Discrete Mathematics*, 134(1-3):133–138, 1994.

-
- [66] L. Schläfli and H. Wild. *Theorie der vielfachen Kontinuität*, volume 38. Birkhäuser, 1901.
- [67] E. Schulte and G. I. Williams. Polytopes with preassigned automorphism groups. *Discrete & Computational Geometry*, 54(2):444–458, 2015.
- [68] B. Schulze. Symmetry as a sufficient condition for a finite flex. *SIAM Journal on Discrete Mathematics*, 24(4):1291–1312, 2010.
- [69] J.-P. Serre. *Linear representations of finite groups*, volume 42. Springer, 1977.
- [70] E. Steinitz. Polyeder und Raumeinteilungen. *Encyklopädie der mathematischen Wissenschaften*, 12:38–43, 1922.
- [71] H. Van Der Holst, L. Lovász, A. Schrijver, et al. The Colin de Verdière graph parameter. *Graph Theory and Computational Biology (Balatonlelle, 1996)*, pages 29–85, 1999.
- [72] S. F. Waldron. Group frames. In *An Introduction to Finite Tight Frames*, pages 209–243. Springer, 2018.
- [73] S. F. Waldron. *An Introduction to Finite Tight Frames*. Springer, 2018.
- [74] M. Winter. Classification of vertex-transitive zonotopes. *arXiv preprint arXiv:1912.10469*, 2020.
- [75] M. Winter. Eigenpolytopes, spectral polytopes and edge-transitivity. *arXiv preprint arXiv:2009.02179*, 2020.
- [76] M. Winter. Geometry and topology of symmetric point arrangements. *Linear Algebra and its Applications*, 2020.
- [77] M. Winter. On polytopes that are edge-transitive but not vertex-transitive. *arXiv preprint arXiv:2005.14285*, 2020.
- [78] M. Winter. Symmetric and spectral realizations of highly symmetric graphs. *arXiv preprint arXiv:2009.01568*, 2020.
- [79] G. M. Ziegler. *Lectures on polytopes*, volume 152. Springer Science & Business Media, 2012.

Index

- $(2k_1, \dots, 2k_s)$ -vertex, 122
- $(2k_1, 2k_2)$ -edge, 122
- (n, \dots, n) -hyperprism, *see* hyperprism
- (n, n) -duoprism, *see* duoprism
- 2_{21} -polytope, 99
- $2k$ -face, 122
- 3_{21} -polytope, 99
- P° , *see* polar dual
- $P^\circ(c)$, 80
- $P_G(\theta)$, *see* θ -eigenpolytope
- $\Delta_{d,k}$, *see* hyper-simplex
- $\mathcal{F}_\delta(P)$, 73
- Γ -arrangement, 33
 - irreducible, 33
 - reducible, 33
- Γ -permutahedron, 104
 - uniform, 104
- Γ -realization, 52
 - irreducible, 52
 - reducible, 52
- Γ -zonotope, 104
- $\text{Gen}(Z)$, 103
- Π_σ , *see* permutation matrix
- Σ -arrangement
 - Σ -flexible, 41
 - Σ -rigid, 41
 - exceptionally rigid, 47
 - transitive, 45
- Σ -deformation, 41
 - proper, 41
- Σ -invariant subspace, *see* invariant subspace
- Σ -irreducible subspace, *see* irreducible subspace
- Σ -realization
 - Σ -flexible, 52
 - Σ -rigid, 52
- $\text{Zon}(R)$, 102
- $\ell(v)$, 60
- $\mathcal{A}_d(G, \Sigma)$, 52
- $\mathcal{A}_d(\Sigma)$, 41
- $\omega(v)$, 60
- sk_p , *see* skeleton
- $\tau(v)$, *see* type of a vertex
- θ -eigenpolytope, 15, 74
- θ -realization, 53
- $\text{vol}(C)$, *see* relative volume
- i -vertex, 117
- k -reflection, 93
- k -reflection group, 93
- 1-flag, *see* arc
- 1-reflection, *see* reflection
- 24-cell, 64, 69
- arc, 58
- arrangement, *see* point arrangement
- arrangement matrix, 28
- arrangement space, 28
- arrangement space dictionary, 32
- automorphism group, *see* symmetry group
- barycenter, 57
- binary polyhedral group, 48
- bitruncation, *see* bitruncated regular polytope
- cartesian power, 91
- Catalan solid, 79
- central symmetry, 102
- circumradius, 60
- Colin de Verdière graph invariant, 81
- Colin de Verdière matrix, 81

- congruent, 104, *see also* isometric
- cosine, 60
- cosine sequence, 68
- cosine vector, 67
- Coxeter-Dynkin diagram
 - transitive, 94
- cubeoctahedron, 70, 79, 90
- deformation, *see* Σ -deformation
- demi-cube, 96, 99
- demi-hypercube, *see* demi-cube
- diameter
 - of a graph, 65
- dihedral angle, 61
- dihedral group, 37
- discrete Schrödinger operator, 59
- distance
 - in a graph, 65
- dodecahedral graph, 54
- dodecahedron
 - regular, *see* regular dodecahedron
 - rhombic, *see* rhombic dodecahedron
- duoprism, 91
- edge in-sphere, 117
- edge length, 60
- edge midpoint, 90
- edge-flip, 93
- eigenpolytope, *see* θ -eigenpolytope
- eigenspace
 - of a graph, 53
- eigenvalue
 - of a graph, 53
- eigenvector
 - of a graph, 53
- equivariant map, 33
- Euclidean symmetry
 - of an arrangement, 32
- exterior angle, 107
- flag-orbit, 89
- flat, 103
- flipped orientation, 93
- generator
 - of a zonotope, 103
- generic point, 104
- geometric separation, *see* separated geometrically
- Gosset graph, 76
- Gosset polytope, *see* 3_{21} -polytope
- graph
 - arc-transitive, 59
 - distance-regular, 65, 75
 - distance-transitive, 66
 - edge-transitive, 59
 - half-transitive, 59
 - vertex-transitive, 58
- graph realization, 51
 - θ -balanced, 56
 - θ -spectral, 53
 - arc-transitive, 59
 - balanced, *see* θ -balanced graph realization
 - distance-transitive, 66
 - edge-transitive, 59
 - full-dimensional, 52
 - half-transitive, 59
 - normalized, 52
 - of full dimension, *see* full-dimensional graph realization
 - of full local dimension, 62
 - spectral, *see* θ -spectral graph realization
 - spherical, 52
 - symmetric, *see* Γ -realization
 - vertex-transitive, 58
- halved n -cube, 76
- Hamming graph, 76
- Hamming polytope, 75, 91, 99
- height
 - of a face, 118, 122
- Holt graph, 98
- hyper-simplex, 95, 99
- hyperplane arrangement, 105
 - chamber-transitive, 114
 - with congruent chambers, 114
- hyperprism, 91

- icosahedron
 - regular, *see* regular icosahedron
- icosidodecahedron, 90
- invariant subspace, 33
 - irreducible, *see* irreducible subspace
- irreducible subspace, 34
- Izmestiev matrix, 81
- Johnson graph, 76
- Minkowski sum, 102
- norm
 - of a semi-star, 113
- normalization
 - of a zonotope, *see* normalized zonotope
- orbit polytope, 87
- orthogonal projection, *see* projection
- parameters
 - of a bipartite polytope, 117
- pentagram, 77
- permutahedron, *see* Γ -permutahedron
- permutation group
 - Σ -rigid, 44
 - arc-transitive, 59
 - distance-transitive, 66
 - edge-transitive, 59
 - half-transitive, 59
 - regular, 47
 - rigid, *see* Σ -rigid permutation group
 - transitive, 45
 - vertex-transitive, 58
- permutation matrix, 33
- Petersen graph, 66
- Petersen polytope, 75, 99
- Platonic solid, *see* regular polytope, 88
- point arrangement, 27
 - equivalent, 29
 - essentially transitive, 45
 - flexible, *see* Σ -flexible point arrangement
 - full-dimensional, 28
 - normalized, 30
 - of full dimension, *see* full-dimensional
 - rigid, *see* Σ -rigid point arrangement
 - spherical, 30
 - symmetric, *see* Γ -arrangement
- polar dual, 80
- polygon
 - bipartite, 119
 - edge-transitive, 89, 115
 - regular, 89, 106
 - strictly bipartite, 119
- polyhedron
 - edge-transitive, 90
 - strictly bipartite, 121
- polytope, 73
 - Δ -transitive, 88
 - δ -transitive, 88
 - θ -spectral, 15, 77
 - 0-transitive, *see* vertex-transitive
 - 1-transitive, *see* edge-transitive
 - abstract regular, 89
 - arc-transitive, 93
 - bipartite, 117
 - bitruncated regular, 97
 - centered at the origin, 73, 82
 - convex, *see just* polytope
 - distance-transitive, 98
 - dual, *see* polar dual
 - edge-transitive, 73, 88, 89
 - generic Wythoffian, *see* Γ -permutahedron
 - half-transitive, 97
 - inscribed bipartite, 120
 - neighborly, 78
 - non-prismatic, 96
 - omnitruncated uniform, 110
 - perfect, 17
 - polar, *see* polar dual
 - prismatic, 96
 - purely edge-transitive, 115
 - rectified regular, 90, 97
 - regular, 64, 88
 - runcinated regular, 97
 - simple, 111

- simultaneously vertex- and edge-transitive, 91
- spectral, *see* θ -spectral
- strictly bipartite, 117, 121
- uniform, 78
- vertex-transitive, 73, 88
- Wythoffian, 87
- Wythoffian uniform, 87
- projection, 29
- radius, *see also* circumradius
 - of a bipartite polytope, 117
 - of a graph realization, 59
 - of a point arrangement, 32
- realization, *see* graph realization
- rectification, *see* rectified regular polytope
- reduced set of vectors, 103
- reflection
 - generalized, *see* k -reflection
- reflection group
 - finite, 87
 - generalized, *see* k -reflection group
- regular dodecahedron, 54
- regular icosahedron, 61
- representation
 - of a Σ -arrangement, 33
- rhombic dodecahedron, 79, 90, 115
- rhombic triacontahedron, 79, 90, 115
- root system, 104, 112
 - reduced, *see* reduced set of vectors
- runcination, *see* runcinated regular polytope
- Schläfli graph, 76
- Schläfli polytope, *see* 2_{21} -polytope
- semi-star, 103
- separated geometrically, 20, 37
- skeleton
 - of a polytope, 73
- spectrum
 - of a graph, 53
- standard permutahedron, 111
- sufficiently symmetric, 17, 56
- symmetry, *see also* Euclidean symmetry
 - of a polytope, 73
 - sufficient, *see* sufficiently symmetric
- symmetry group
 - combinatorial, 52
 - Euclidean, 73
- Theorem of Izmestiev, 80
- truncated 24-cell, 111
- truncated octahedron, 111
- truncated tetrahedron, 54, 79
- type
 - of a vertex, 122
- volume
 - relative, 81
- Weyl chamber, 105
- Weyl group, 105
- zonotope, 101
 - inscribed, 102, 113
 - irreducible, 110
 - normalized, 109
 - reducible, 110
 - uniform, 108, 110
 - vertex-transitive, 101, 112

Erklärung gemäß §6 der Promotionsordnung

Hiermit erkläre ich an Eides Statt, dass ich die von mir eingereichte Arbeit

„Spectral Realizations of Symmetric Graphs, Spectral Polytopes and Edge-Transitivity“

- selbstständig und nur unter Benutzung der in der Arbeit angegebenen Quellen und Hilfsmittel angefertigt habe,
- in dieser oder ähnlicher Form an keiner anderen Stelle zum Zwecke eines Promotionsverfahrens vorgelegt habe.

Hiermit erkläre ich an Eides Statt, dass ich keine weiteren Promotionsverfahren bei anderen Stellen beantragt hatte bzw. beantragt habe.

Chemnitz, den 30.03.2021

Martin Winter

Curriculum Vitae

Martin Winter

born 12/07/1992 in Räckelwitz, Germany

Studies and education

- | | |
|-------------------------|--|
| Since 01/03/2018 | PhD program in Mathematics at TU Chemnitz
Supervisor: Prof. Christoph Helmberg
Title of thesis: <i>Spectral Realizations of Symmetric Graphs, Spectral Polytopes and Edge-Transitivity</i> |
| 01/10/2013 – 30/09/2015 | Master course in <i>Mathematics with Computer Science</i> (minor in Physics) at TU Chemnitz
Title of thesis: <i>Notes on $P(k)$-Graphs and the Conjecture of Kotzig</i>
Final grade: 1.1*, with distinction |
| 01/10/2010 – 30/09/2013 | Bachelor course in <i>Computational Science</i> at TU Chemnitz
Title of thesis: <i>Image Alignment</i>
Final grade: 1.0*, with distinction |
| 07/2010 | Abitur at Goethe Gymnasium Bischofswerda, Germany |

*Grades are from 1.0 (best) to 5.0 (not passed).

Published articles

1. T. Jahn and M. Winter, *Vertex-Facet Assignments for Polytopes*. Contributions to Algebra and Geometry (2020), doi.org/10.1007/s13366-020-00504-9.
2. M. Winter, *Geometry and Topology of Symmetric Point Arrangements*. Linear Algebra and its Applications (2021), doi.org/10.1016/j.laa.2020.11.017.
3. M. Winter, *Classification of Vertex-Transitive Zonotopes*. Discrete & Computational Geometry (2021), doi.org/10.1007/s00454-021-00303-6.

Presentations at conferences and workshops

1. *Capturing Polytopal Symmetries in the Edge-Graph*. 06/04/2021, Polytopics Conference (Max Planck Institute for Mathematics in the Sciences, Leipzig).

2. *Spectral Graph Theory for Polytopes.*
19/09/2020, DMV Annual Conference 2020 (Chemnitz).
3. *Inscribed Zonotopes and Hyperplane Arrangements with Congruent Chambers.*
13/03/2020, Combinatorial Coworkspace (Kleinwalsertal).
4. *Edge-Transitive Polytopes.*
 - 06/12/2019, Geometry Day (Jena).
 - 08/11/2019, Colloquium on Combinatorics (Paderborn).
 - 22/10/2019, South-Eastern Graph Theory Workshop (Chemnitz).
5. *Spectral Methods for Symmetric Polytopes.*
 - 10/07/2019, Discrete Geometry Day 2 (Budapest).
 - 13/03/2019, Polytope Spring School (Bochum).
 - 24/08/2018, FRICO (Chemnitz).

Invited talks

1. *Capturing Polytopal Symmetries in the Edge-Graph.*
24/03/2021, combinatorics seminar at KTH Royal Institute of Technology in Stockholm (Prof. Katharina Jochemko).
2. *On Symmetry, Rigidity and Spectrum.*
04/03/2021, research seminar at National University of Galway.
3. *Edge-Transitive Polytopes.*
 - 06/05/2020, research seminar at TU Berlin (Prof. Michael Joswig).
 - 04/02/2020, research seminar at Goethe University Frankfurt (Prof. Raman Sanyal).
 - 22/01/2020, research seminar University of Rostock (Prof. Achill Schürmann).