Prologue: Regular polytopes

Current Version
January 1, 2005

To motivate the study of finite reflection groups and Coxeter groups more generally, I'll begin by briefly sketching the classification of regular polytopes. Convex polytopes are fundamental objects in mathematics which can be viewed in a number of equivalent ways: as the convex hull of a finite set of points in $\mathbb{R}^n$, as the intersection of a finite number of half-spaces whose intersection is compact, or as the image of a high-dimensional simplex under a linear transformation. Within the class of convex polytopes, those which are "completely symmetric" are particularly beguiling; they also have a tendency to play a major role in seemingly disparate areas of mathematics. These highly symmetric polytopes are more commonly known as *regular polytopes*. Before giving a precise definition of a regular polytope, let's consider some familiar, low-dimensional examples.

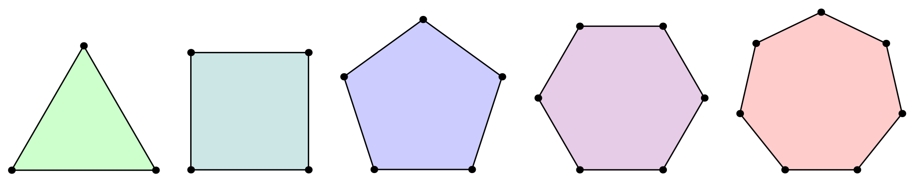

FIGURE 1. Regular m-gons for $m = 3, 4, 5, 6, 7$.

EXAMPLE 0.1 (Regular polytopes in low dimensions). Regardless of the exact definition, it is clear that the class of regular polytopes should at the very least include the regular m-gons (i.e. those m-sided polygons which are equi-lateral and equi-angular) and the 5 Platonic solids (i.e. the regular tetrahedron, cube, octahedron, dodecahedron, and icosahedron). See Figures 1 and 2. Technically, a closed interval of the reals will qualify as the only 1-dimensional regular polytope. That these examples are the only regular polytopes in dimensions up to 3 is one consequence of the classification theorem established below.

Associated to any convex polytope is a natural simplicial subdivision called its barycentric subdivision. Since we are primarily interested in those polytopes which are highly symmetric, we define this subdivision using the circumcenters of the faces of the polytope.

DEFINITION 0.2 (Circumcenters). Given a bounded set A in $\mathbb{R}^n$ and a point $x \in \mathbb{R}^n$ there is some minimum radius, $\text{rad}_A(x)$, such that the closed ball around x of radius $\text{rad}_A(x)$ contains all of A. The collection of all such minimal radii has an infimum and any point $x \in \mathbb{R}^n$ such that $\text{rad}_A(x)$ attains this infimum is called a *circumcenter* of A.

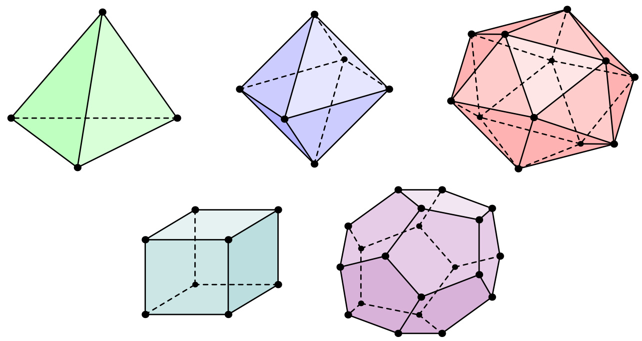

FIGURE 2. The 5 Platonic solids.

That unique circumcenters exist is a consequence of the non-positively curved nature of $\mathbb{R}^n$. We will return to this theme in Chapter 2.

PROPOSITION 0.3 (Circumcenters exists). *Any bounded set in $\mathbb{R}^n$ has a unique circumcenter.*

Proposition 0.3 is a special case of a theorem which holds much more generally. Since the full version is proved later, the argument here is merely sketched. Existence essentially follows from the completeness of $\mathbb{R}^n$ and uniqueness is immediate once it is noticed that for any three distinct collinear points x, y and z with y between x and z, $\text{rad}_A(y) < \max\{\text{rad}_A(x), \text{rad}_A(z)\}$. In particular, supposing x and z to be distinct circumcenters leads to an immediate contradiction. Finally, notice that this uniqueness argument only relies on a very elementary fact about the curvature of $\mathbb{R}^n$.

DEFINITION 0.4 (Barycentric subdivision). The *barycentric subdivision* of a convex polytope P introduces a new vertex at the circumcenter of each i-dimensional face and then subdivides appropriately. More specifically, there is a simplex in the subdivision if and only if the faces to which the vertices correspond form a *partial flag*, i.e. given any two faces in the list, one is contained in the boundary of the other. For convenience later, we think of every vertex of the subdivision as having an integer assigned which records the dimension of the cell of which it is the circumcenter. Notice that under this scheme, distinct integers are assigned to each of the vertices in a simplex of the subdivision. The barycentric subdivision of a regular pentagon is shown in Figure 3.

The regularity of the regular m-gons and the Platonic solids is quite apparent when we barycentrically subdivide. In each case, all of the top-dimensional simplices in the subdivision are isometric. This is, in fact, almost exactly the definition of a regular polytope.

DEFINITION 0.5 (Regular polytopes). Let P be an n-dimensional convex polytope. A (complete) *flag* in P is a sequence of i-dimensional faces in P, one for each $i = 0, \ldots, n$, ordered by inclusion. In other words, a flag consists of a vertex which is contained in an edge which is contained in a 2-cell, etc. The polytope P is called *regular* if its isometry group acts transitively on its flags. An alternative definition can be given using the barycentric subdivision of P. First note that there is a one-to-one correspondence between the top dimensional cells in the subdivided

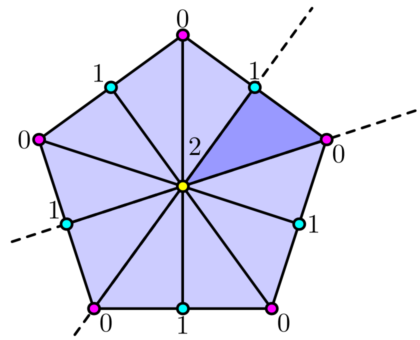

FIGURE 3. The barycentric subdivision of a regular pentagon with a fundamental chamber highlighted. The basic reflections are marked with dashed lines.

complex (sometimes called *chambers*) and the complete flags in P. Since any isometry of P must take the circumcenter of a face to the circumcenter of its image, all of the isometries of P induce simplicial maps from the barycentric subdivision of P to itself. The integers assigned to the vertices are, of course, preserved under these maps. As a consequence, P is regular if and only if its isometry group acts transitively on the chambers of the subdivision.

It is straightforward to check that any particular example is a regular polytope according to this definition. The difficult part of the classification theorem (as in any classification theorem) is to show that we have found a complete list of examples. The easy part will be done first. In order to ease the introduction of high-dimensional examples, we disgress for a moment to discuss the dual of a polytope.

DEFINITION 0.6 (Dual polytopes). If P is a n-dimensional regular polytope, then the convex hull of the circumcenters of its $(n-1)$-dimensional faces will be another regular polytope called its *dual*.

Although it is not clear from the definition, the dual of a regular polytope is, as claimed, once again regular and the dual of the dual gives a rescaled version of the original. Both of these facts are readily checked in the concrete examples we are presenting so the general theory of polytopal duals will not be rigorously developed. We should note, however, that the polar dual of a generic polytope is defined using linear functionals rather than circumcenters; it is only in the case of regular polytopes, that the general definition agrees with the one given above. Among the examples already introduced, a regular m-gon is self-dual (i.e. dual to itself), the tetrahedron is self-dual, the cube and the octahedron are dual to each other, and the dodecahedron and icosahedron are dual to each other.

EXAMPLE 0.7 (High dimensions). There are several high-dimensional regular polytopes with which the reader is probably already familar, including n-simplices and n-cubes. The dual of an n-cube is a generalized version of an octahedron called the n-dimensional *cross-polytope*. All three families of examples are easy to describe explicitly using coordinate systems. Let $\vec{e}_i$, $i \in \{1, \ldots, n\}$ denote the standard orthonormal basis of $\mathbb{R}^n$. The convex hull of the tips of the vectors $\vec{e}_i$

form a regular $(n-1)$-simplex. The space $[-1,1]^n$ is an n-cube, and its dual is the convex hull of the vectors $\pm\vec{e}_i$.

There are two slightly more exotic examples of regular polytopes in dimension 4 which are closely related to the Poincaré homology 3-sphere.

Example 0.8 (120-cell and 600-cell). The *Poincaré homology sphere* is the name given to the counterexample which Poincaré found that violated his famous conjecture in its original form. Poincaré originally suggested that any 3-diimensional manifold with the homology groups of the 3-sphere might be homeomorphic to the 3-sphere. After finding a counterexample, he reformulated the conjecture with "homotopy groups" in place of "homology groups". The construction of his counterexample goes as follows. Start with a solid dodecahedron and identify antipodal 2-cells with a slight clockwise twist (a $\pi/5$ twist to be precise). The result is a 3-manifold which can be given a metric with constant curvature $+1$ and a universal cover which is isometric to $\mathbb{S}^3$. Since the fundamental group of the original 3-manifold has size 120, the universal cover is tiled with 120 regular (spherical) dodecahedra. Thinking of $\mathbb{S}^3$ as sitting inside of $\mathbb{R}^4$ we can take the convex hull of the 600 vertices of this tiling and get a regular 4-polytope known as the 120-*cell*, named after its 120 dodecahedra. Its dual is another regular 4-polytope with 120 vertices and 600 regular tetrahedra. It is called, of course, the 600-*cell*.

Our last example is easiest to describe via direct construction.

Example 0.9 (24-cell). The sphere of radius 2 centered at the origin in $\mathbb{R}^4$ contains exactly 24 vectors whose coordinates are integers. There are 16 vectors of the form $(\pm1,\pm1,\pm1,\pm1)$ and 8 vectors which are ±2 times a standard basis vector. The regularity of the convex hull of these 24 vectors is hinted at once it is observed that the 16 vectors of the form $(\pm1,\pm1,\pm1,\pm1)$ can be split into two groups with 8 vectors each so that any two vectors in the same group are either orthogonal or parallel. Moreover, it can also be checked that these three groups of 8 vectors which look like the vertices of a 4-cross-polytope are symmetrically arranged with respect to one another. At this point, it should at least seem plausible that these 24 points form the vertices of a regular (and self-dual) 4-polytope.

Perhaps surprisingly, the examples given above form a complete list of regular polytopes in all dimensions. Verification that these examples are indeed regular polytopes is left to the reader. As was mentioned earlier, the real difficulty is showing that no other examples exist. The trick in this case is to shift our attention from the polytope itself to its isometry group and a fundamental domain of its action.

As a first step we show that the isometry group of a regular polytope is always a finite group generated by reflections. Recall that a *reflection* is an isometry of $\mathbb{R}^n$ which fixes an $(n-1)$-dimensional hyperplane H and sends vectors perpendicular to H to their negatives. We begin by establishing some additional notation.

Definition 0.10 (Fundamental chamber). Let P be a regular n-dimensional polytope which has been barycentrically subdivided. One of its chambers C (i.e. a top-dimensional simplex) is selected and called the *fundamental chamber* of P. Let v_i, $i \in \{0,\ldots,n\}$ be the unique vertex of C labeled i. Without loss of generality we may assume that P is situated in $\mathbb{R}^n$ so that the circumcenter of P itself, v_n, is located at the origin.

Notice that the vectors from the origin out to the other n vertices in C form a basis of $\mathbb{R}^n$. Thus, their image under an isometry determines that isometry uniquely. Because each isometry of P sends chambers to chambers and the isometry is completely determined by the image of the fundamental chamber, the number of chambers in the barycentric subdivision of P are in one-to-one correspondence with the elements of the isometry group.

DEFINITION 0.11 (Basic reflections). Since the boundary of P is (topologically) an $(n-1)$-sphere and, in particular, an $(n-1)$-dimensional manifold, for each $(n-1)$-dimensional face of C there is a unique chamber C' distinct from C which contains this face. Because P is a regular polytope, by definition there is an isometry of P, which we call r_i, which takes the fundmental chamber C to the unique chamber containing all of the vertices of C except v_i. Since r_i must preserve the labeling of the vertices, the image of v_i under r_i is another vertex labeled i. The thing to notice is that r_i is a reflection. To see this, note that it fixes the $(n-1)$-dimensional hyperplane H spanned by the vectors from the origin to v_j, $j \neq i$. Moreover, being an isometry, it sends vectors perpendicular to H to vectors perpendicular to H and since it isn't the identity, it sends them to their negative. See Figure 3. The reflections r_i, $i = \{0, 1, \ldots, n-1\}$ are called the *basic reflections* of P with respect to C.

As promised, the isometry group is generated by these basic reflections.

LEMMA 0.12 (Generators). *If P is a regular n-dimensional polytope, then the isometry group of P is a finite group generated by n reflections. More precisely the basic reflections of P with respect to any fundamental chamber C generate the isometry group.*

PROOF. Because there is a bijection between the isometry group of G and the chambers of P, it is sufficient to show that products of the basic reflections can send C to any chamber of P. Consider the orbit of C under the group of isometries generated by the reflections r_i. Let C' be a chamber in this orbit and let w be a sequence of basic reflections which move C to C'. Applying r_i prior to applying this sequence will send C to the neighbor of C' which shares all of its vertices with C' except for the vertex labeled i. Doing to for all i shows that all of the chambers which share an $(n-1)$-dimensional face with C' also belong to the orbit of C. This completes the proof since the only subcomplex of P which contains C and is closed under the taking of neighbors which share a codimension 1 face is all of P. $\square$

Because the basic reflections generate the isometry group, the entire regular polytope P can be reconstructed from the shape of the fundamental chamber C by simply iteratively reflecting in its maximal proper faces. The fundamental chamber, in turn, can be reconstructed from the collection of dihedral angles between the basic reflections. It might seem that these angles only encode the shape of the polyhedral cone emanating from the origin, but the final side is an affine hyperplane perpendicular to the vector from v_n to v_{n-1}. One fact which makes regular polytopes easy to analyze is that most pairs of basic reflections have orthogonal normal vectors.

LEMMA 0.13 (Orthogonality relations). *Let P be a regular polytope with fundamental chamber C and let r_i, $i = \{0, \ldots, n-1\}$, be its basic reflections with respect*

Common name	**Schläfli symbol**	**Cartan-Killing type**
n-simplex	$\{3^{n-1}\}$	A_n
n-cross-polytope	$\{3^{n-2},4\}$	B_n
n-cube	$\{4,3^{n-2}\}$	B_n
4-simplex	$\{3,3,3\}$	A_4
4-cross-polytope	$\{3,3,4\}$	B_4
4-cube	$\{4,3,3\}$	B_4
24-cell	$\{3,4,3\}$	F_4
600-cell	$\{3,3,5\}$	H_4
120-cell	$\{5,3,3\}$	H_4
tetrahedron	$\{3,3\}$	A_3
octohedron	$\{3,4\}$	B_3
cube	$\{4,3\}$	B_3
icosahedron	$\{3,5\}$	H_3
dodecahedron	$\{5,3\}$	H_3
m-gon	$\{m\}$	$I_2(m)$

Table 1. Translating between the various notations for the regular polytopes.

to C. If $|i-j|>1$ then the reflections r_i and r_j commute and their normal vectors are orthogonal.

Proof. Let k be an integer with $i<k<j$ and let F be the face of P whose circumcenter is v_k. The basic reflection r_i sends the vertex v_i in the fundmental chamber C to the circumcenter of a different i-dimensional face of F. In particular, the line segment joining v_i and its image (whose direction is parallel to the normal vector of r_i) lies in the face F. On the other hand, the basic reflection r_j fixes the face F pointwise and hence fixes its entire span. As a consequence, the normal vector for r_j is perpendicular to the face F. Since the normal vectors of r_i and r_j are perpendicular, the reflections commute. □

The converse of the above statement is also true (and an easy proof of this will be added at some point).

Definition 0.14 (Schläfli symbols). Since almost all of the dihedral angles between codimension 1 faces in the fundamental chamber are $\pi/2$, it makes sense to only record the remaining angles. In other words, we should record the dihedral angles between the basic reflections r_{i-1} and r_i for $i=1,\ldots,n-1$. Since each of these angles is π/m for some integer m, it makes sense to encode all of the necessary information into a short sequence of positive integers. In preparation for the general situation introduced in Chapter 1, each hyperplane containing a codimension 1 face of C will be replaced with its unit normal vector which selects the side of the hyperplane containing C. If the dihedral angle between to faces is π/m, then the angle between their inward points normal vectors will be $\pi-\pi/m$. The *Schläfli symbol* for a regular n-dimensional polytope is the sequence $\{m_1,m_2,\ldots,m_{n-1}\}$ where the dihedral angle between the inward pointing normal vectors of the basic reflections r_i and r_{i-1} is $\pi-\pi/m_i$. The *Schläfli symbol* for a cube, for example, is

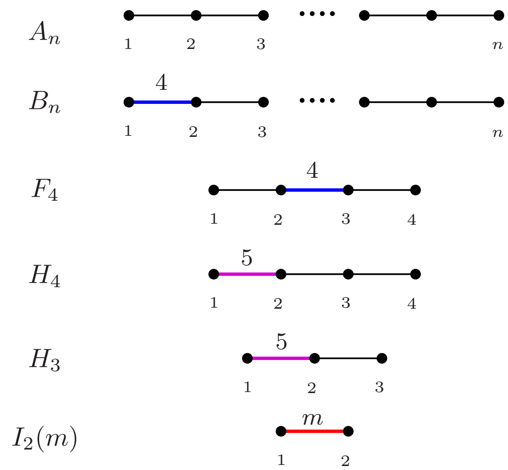

FIGURE 4. Dynkin diagrams for the regular polytopes.

$\{4, 3\}$ since there is a $\pi - \pi/4$ angle between r_0 and r_1 and a $\pi - \pi/3$ angle between r_1 and r_2.

A more flexible notation containing essentially the same information is the Dynkin diagram. One key advantage of Dynkin diagrams over Schläfli symbols is that they retain their usefullness even after we leave the world of regular polytopes.

DEFINITION 0.15 (Dynkin diagrams). Let P be a regular polytope with fundamental chamber C. The Dynkin diagram of P records the angles between the inward-pointing unit normal vectors of the codimension 1 faces of C in a finite labeled graph. The vertices correspond to the basic reflections. If two basic reflections commute, then no edge is drawn connecting the corresponding vertices. If the angle between them is $\pi - \pi/m$ for $m > 2$ then an edge labeled m is drawn between their vertices. Because edges labeled 3 are quite common, these particular labels are usually suppressed. The Dynkin diagrams for the isometry groups of the regular polytopes are shown in Figure 4. The conversions between their common names, their Schläfli symbols and the Cartan-Killing type of their associated Dynkin diagrams are given in Table 1

The main difficulty of the classification theorem can now be restated using Dynkin diagrams. Every regular polytope is completely encoded in the geometry of its fundamental chamber, which is determined by the dihedral angles between its codimension 1 faces containing the central vertex. These angles can be encoded in a Dynkin diagram which, by Lemma 0.13 is a linear string of edges. The main question is which sequences of edge labels are possible? The answer uses linear algebra.

DEFINITION 0.16 (Positive definite matrices). Recall that if M is a real symmetric matrix, then all of its eigenvalues are real and it has an orthonormal basis of eigenvectors. Such a matrix is called *positive definite* when all of its eigenvalues are positive.

Positive definite matrices are relevant because of their close connection with arrangements of vectors in space. The key result we need is the following.

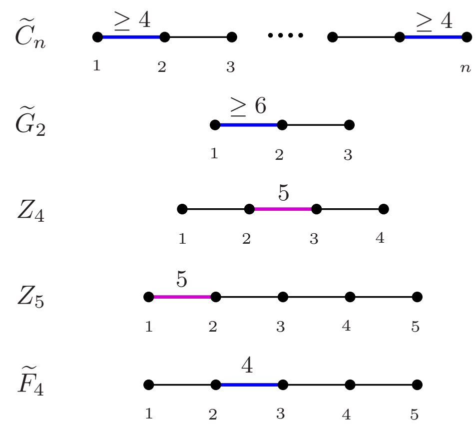

FIGURE 5. Some Dynkin diagrams which are not positive definite.

THEOREM 0.17 (Vector arrangements and positive definite matrices). *If $\vec{v}_i$, $i = 1, \ldots, n$ is a set of linearly independent vectors in $\mathbb{R}^n$, then the real symmetric matrix M whose (i,j)-entry is $\vec{v}_i \cdot \vec{v}_j$ is positive definite. Conversely, given a real symmetric positive definite matrix M, there exist an ordered n-tuple of linearly independent vectors in $\mathbb{R}^n$ whose dot products are described by M.*

There is an easy and well-known criterion which determines whether or not a matrix is positive definite.

PROPOSITION 0.18 (Positive definite test). *An $n \times n$ matrix is positive definite if and only if all of its principal minors has positive determinants.*

In our case, we are starting with a Dynkin diagram and we are trying to create an arrangement of vectors in space with the right angles. The prescribed angles between the inward pointing normal vectors will be $\pi - \pi/m$ for $m \geq 2$ so the values of cosine we need to consider are $2\cos(\pi/2) = 0$, $2\cos(\pi/3) = 1$, $2\cos(\pi/4) = \sqrt{2}$, $2\cos(\pi/5) = \tau$ (where τ is the golden ratio, $(1+\sqrt{5})/2)$, and $2\cos(\pi/6) = \sqrt{3}$. Because most of these values have 2 in the denominator, it is easier to test whether $2M$ is positive definite, rather than M itself.

EXAMPLE 0.19 (Determinant calculations). Using these values it is easy to calculate the determinants of the matrices corresponding to the graphs in Figure 4 and see that they really are positive. As an example, consider the matrix $2M$ for the Dynkin diagram of type A_4. The calculation is as follows.

$$\det \begin{bmatrix} 2 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ 0 & 0 & -1 & 2 \end{bmatrix} = 2 \cdot \det \begin{bmatrix} 2 & -1 & 0 \\ -1 & 2 & -1 \\ 0 & -1 & 2 \end{bmatrix} - (-1)^2 \det \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix}$$

which simplifies to $2(4) - 3 = 5$. More generally, an easy induction shows that for the diagram of type A_n the determinant of $2M$ is $n+1$. As a consequence, the matrix associated to A_n is positive definite and there exist vectors arranged with the necessary angles. Using these values (and the fact that $\tau^2 = \tau + 1$), the determinants assoicated with H_3, H_4 and Z_5 simplify to $4-2\tau$, $5-3\tau$, and $6-4\tau$, respectively. Since $\tau \sim 1.618$, the first two are positive while the third is negative, thereby establishing that the H_3 and H_4 describe a possible arrangement of vectors in space, but that Z_5 does not.

In some sense, we already knew these Dynkin diagram produced positive definite matrices since they were derived from the shapes of fundamental chambers for the explicit regular polytopes constructed earlier. More importantly, it is also easy to calculate the determinants associated with the 5 Dynkin diagrams shown in Figure 5 and verify that they are not positive definite, and thus do not describe any arrangement of vectors in space.

Corollary 0.20 (Forbidden subgraphs). *If X_n is the Dynkin diagram of a regular polytope, then X_n cannot contain any of the graphs shown in Figure 5 as a subgraph.*

Using Corollary 0.20, it is now straightforward to complete the classification of regular polytopes.

Theorem 0.21 (Classification of regular polytopes). *Every regular polytope is*

1. *a closed interval,*
2. *a regular m-gon with $m \geq 3$,*
3. *one of the 5 platonic solids,*
4. *one of the 6 regular 4-polytopes, or*
5. *an n-dimensional simplex, cube or cross-polytope with $n > 4$.*

Proof. Let P be a regular polytope and X_n be its Dynkin diagram. Since there exist regular polytopes for each of the Dynkin diagrams listed, it only remains to show that this list is complete. By Corollary 0.20, it is sufficient to show that the only linear Dynkin diagrams which avoid the 5 types of graphs shown in Figure 5 are the ones we have listed. The outline of the proof is given in Figure 6.

If X_n diagram has at most 2 vertices then X_n is either a trivial graph (and P is an interval) or X_n is of type $I_2(m)$ (and P is a regular m-gon). Thus we may assume $n > 2$. If X_n has no edges with a label larger than 3, then X_n is of type A_n (and P is a regular n-simplex). On the other hand, if X_n has more than one such edge label, then it contains $\widetilde{C}_n$ as a subgraph, contradiction. Thus we may assume that X_n contains exactly one label bigger than 3. If this label is 6 or more, then X_n contains $\widetilde{G}_2$ as a subgraph, contradiction, so we may assume the label is either 4 or 5.

Suppose the label is 4. If this label occurs at either end of X_n, then X_n is of type B_n (and P is either an n-cube or an n-cross-polytope). If it does not occur at an end, then either X_n is F_4 (and P is the 24-cell), or it contains $\widetilde{F}_4$ as a subgraph, contradiction. Finally, consider the case where the label is 5. If it does not occur at an end of X_n, then X_n contains Z_4 as a subgraph, contradiction. On the other hand, if it occurs at one end, then X_n is either H_3 (making P a dodecahedron or an icosahedron), H_4 (making P a 120-cell or a 600-cell), or it contains Z_5, contradiction. $\square$

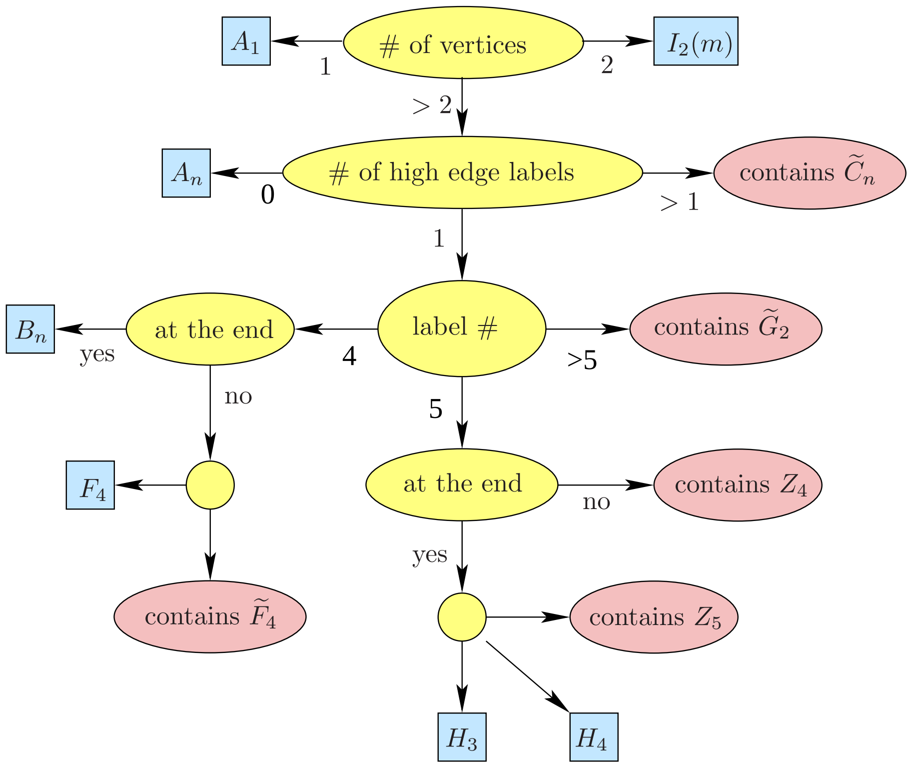

FIGURE 6. Outline of the proof of Theorem 0.21.