Vertex-Transitive Polytopes

Vivian

July 13, 2017

1 Preliminaries

The first question to get our hands on is, what is a polytope? The most handwavy definition is that it's an n-dimensional version of a polygon or polyhedron. Here's a definition that is somewhat less handwavy:

Definition 1.0.1. A subset $A \subseteq \mathbb{R}^n$ is *convex* if for any two points x, y in A, the line segment from x to y is also contained in A.

Definition 1.0.2. A *convex polytope* in $\mathbb{R}^n$ is the convex hull of finitely many points in $\mathbb{R}^n$, or the smallest convex set containing finitely many points in $\mathbb{R}^n$. It is *n-dimensional* or a *convex n-polytope* if it's not contained in a copy of $\mathbb{R}^{n-1}$.

In this class, all polytopes will be convex and connected. The boundary of an n-polytope is a collection of $(n-1)$-polytopes.

Examples.

- Point. Boundary is a point.
- Line segment. Boundary is two points.
- Polygons (convex hulls of sets points in $\mathbb{R}^2$, including one example where one of the points isn't a vertex). Boundary is a collection of line segments.
- Polyhedra. Boundary is a collection of polygons.
- n-simplex. Boundary is the union of $(n+1)$ many $(n-1)$-simplices.

The collection of maximal $(n-1)$-polytopes on the boundary of an n-polytope is the set of its $(n-1)$-*facets*. Those facets are polytopes in their own right, so they have $(n-2)$-facets; these are also considered to be $(n-2)$-facets of the original polytope. By induction, we get k-facets for every $k \leq n$. The 0-facets are the *vertices* of the polytope.

Looking at the vertices and facets of a polytope gives us a fairly combinatorial intuition for looking at symmetries. There's a set of geometric symmetries of the polytope, called its *automorphism group* (I promise no group theory, only that one word). One way for a polytope to be very symmetric is if we can pick any two k-facets, and there's a symmetry taking one to the other.

Example. Hexagons vs. hexagons with three sides big and three sides small (truncated triangle). In the first case, there's a symmetry taking any edge to any other, and taking any vertex to any other. In the second, we can take any vertex to any other, but there's no symmetry that takes short edges to long edges. TODO: Add tikZ here.

$\triangle$

But we can do even better, if we want our polytope to be super-symmetric, by looking at flags.

2 Regular Polytopes

Definition 2.0.1. A *flag* in an n-polytope P is a sequence of facets $a_0, a_1, \ldots, a_{n-1}$, where a_i is an i-facet and $a_i \subseteq a_{i+1}$.

Definition 2.0.2. A polytope is *regular* if there's a symmetry taking any flag to any other flag.

Examples.

- Line segments and points!
- Polygons (regular vs. non)
- Platonic solids
- n-simplex, n-cube

An important concept for us when discussing regular polytopes is going to be the *dual* of a polytope. If we start with an n-polytope, the dual polytope is constructed by placing a vertex in the center of every $(n-1)$-facet, then taking the convex hull. Any two vertices in the dual polytope are contained in a k-facet if and only if the corresponding $(n-1)$-facets in the original shared a common $(n-k-1)$-facet.

Example. Cube - Octahedron; Dodecahedron - Icosahedron. Add tikZ!

$\triangle$

Definition 2.0.3. The *n-crosspolytope* is the dual of the n-dimensional cube.

So the big shiny question is, what regular polytopes are there? And the big shiny answer is this theorem.

Theorem 2.0.4.

- *In 3D, there are five regular polytopes (tetrahedron, cube, octahedron, dodecahedron, icosahedron).*
- *In 4D, there are six regular polytopes (5-cell, 8-cell, 16-cell, 24-cell, 120-cell, 600-cell).*
- *If $n \geq 5$, there are precisely three n-dimensional polytopes. These are the n-simplex, the n-cube, and the n-crosspolytope (dual to the n-cube).*

This one, we prove!

Proof: 3- and 4- dimensional cases. (add pictures) First, let's talk briefly about three dimensions. Say we're trying to build a regular 3-polytope. Well, the facets have to be regular 2-polytopes, or regular polygons. For starters, let's pick equilateral triangles. Our next choice is then how many to fit around a vertex. If we pick three, four, or five, we automatically get distinct platonic solids. If we try to fit six, we get a flat tiling, which will never curve around to close up, and we can't fit more than six. So these are our only options for triangles.

Now if we pick squares, we can fit three around a vertex, giving us a cube, but four will tile and five or more wont' fit. If we pick pentagons, we can fit three around a vertex, but no more. And if we pick hexagons or bigger, we can't even fit three around a vertex (in th case of hexagons, we end up tiling the plane). So our facets have to be pentagons or smaller, and we've exhausted the possibilities there. See the chart for a summary!

sides per facet	facets per vertex	polyhedron
3	3	tetrahedron
3	4	octahedron
3	5	icosahedron
4	3	cube
5	3	dodecahedron

Remark. Swapping the first two columns of this table dualizes the table.

OK, so now for the 4-dimensional case. Now, instead of asking how many regular polygons each vertex is contained in, we can increase all dimensions, and ask how many regular polyhedra each edge is contained in.

In order to do this, we need to determine for each polyhedron the *dihedral angle*.

Definition 2.0.5. The *dihedral angle* in a regular polytope is the angle between two faces around and $(n-2)$-face. (Draw picture)

This is a (rote) exercise in geometry, so at least for now I'm just going to drop this table here.

polyhedron	dihedral angle
tetrahedron	$\cos^{-1}(\frac{1}{3}) \approx 1.23$
cube	$\frac{\pi}{2}$
octahedron	$\pi - \cos^{-1}(\frac{1}{3}) \approx 1.91$
dodecahedron	$\pi - \tan^{-1}(2) \approx 2.03$
icosahedron	$\pi - \cos^{-1}(\frac{\sqrt{5}}{3}) \approx 2.41$

Now, we want to fit at least three polyhedra around each edge, but the sum of all dihedral angles around each edge should be less than 2π; if it's equal, that's a perfect tiling that will not close up, and if it's more, those polyhedra just won't fit. We have six options, and they all end up being geometrically realizable.

polyhedral facets	facets per edge	polytope
tetrahedron	3	5-cell (simplex)
tetrahedron	4	16-cell (crosspolytope)
tetrahedron	5	600-cell
cube	3	8-cell (cube)
octahedron	3	24-cell
dodecahedron	3	120-cell

Bring pictures! Maybe go through argument for why they're all realizable, but I'm comfortable black boxing that and saying that like listing out the 4-vectors constituting the vertices of this thing has been done. □

OK! So now that we've addressed the 4-polytope case, we have an idea of where to go in general. Our strategy for classifying n-polytopes is basically:

- Classify $n-1$ polytopes.
- Find dihedral angles (angles around an $(n-3)$-facet in the $(n-1)$-polytopes).
- We have to fit at least 3 $(n-1)$-polytopes around an $(n-3)$-facet, and not so many that the total dihedral angle exceeds 2π.

Tackling n-polytopes for $n \geq 5$ smells a lot like induction. And it is! Let's do the inductive step first.

Inductive Step. Assume that the only $(n-1)$-polytopes are an $(n-1)$-simplex, an $(n-1)$-cube, and an $(n-1)$-crosspolytope, or the dual to an $(n-1)$-cube, for $n \geq 5$. We will show that the only n-polytopes are an n-simplex, an n-cube, and an n-crosspolytope.

First of all, we need some dihedral angles. The dihedral angle for an n-cube is always $\frac{\pi}{2}$. Considering the crosspolytope, we can imagine an n-cube with vertices at each of the points with all ± 1 coordinates in $\mathbb{R}^n$. Those vertices are the centers of the facets in the crosspolytope, and two of these facets are adjacent if they differ by exactly one coordinate. In that case, the dot product of these vectors gives us (add drawing here, to clarify) that

$$n - 2 = n \cos\theta,$$

where θ is the central angle, so the dihedral angle is $\pi - \cos^{-1}(\frac{n-2}{n})$.

Lastly, for the n-simplex, the dihedral angle is $\cos^{-1}(\frac{1}{n})$. You can see this by standing an n-simplex on one of its facets as a base and looking at it from above (agh, this really needs a tikZ picture too). The total volume of this picture from above is V, the $(n-1)$-volume of one facet. But we've split it into n portions, each of which contributes $V \cos\theta$, for θ the dihedral angle. So

$$nV \cos\theta = V,$$

and thus $\theta = \cos^{-1}(\frac{1}{n})$.

Great! So, now we have some dihedral angles. Building an n-polytope with an $(n-1)$-cube, then, we always have dihedral angle $\frac{\pi}{2}$, so we can fit three but not four around an

$(n-3)$-facet. This is the n-cube. If we start with an $(n-1)$-crosspolytope, for $n \geq 5$, the dihedral angle $\pi - \cos^{-1}(\frac{n-3}{n-1})$ satisfies

$$\pi - \cos^{-1}(\frac{n-3}{n-1}) \geq \frac{2\pi}{3},$$

so it is too large to use as a facet.

Meanwhile, $\cos^{-1}(\frac{1}{n-1}) > \frac{2\pi}{5}$ for $n \geq 5$, but it approaches $\frac{\pi}{2}$ from below as n approaches infinity. So we can always build an n-polytope with 3 $(n-1)$-simplices around each vertex (an n-simplex), and one with 4 $(n-1)$-simplices around each vertex (an n-crosspolytope), but no more.

This completes the proof. □

Now, all that remains is checking that in 5 dimensions, we really only get the simplex, the cube, and the crosspolytope. At this point, this comes down to checking dihedral angles for our regular 4-polytopes. (I'd like to drop a table here, but I don't have all the exact angles). The above argument addresses the 5-polytopes with cubes, simplices, or crosspolytopes as facets. The 600-, 120-, and 24-cells all have dihedral angles at least that of the crosspolytope, so the above argument shows that there aren't any 5-polytopes with any of these as facets.

3 Vertex-Transitive 3-Polytopes

So at this point we have a pretty good idea of what regular polytopes are out there. But there aren't so many, and there are still a lot of shapes that we'd like to call pretty symmetric that aren't regular. Going back to the definition of regular polytopes, we can relax our constraint a bit. Instead of requiring that there be a symmetry taking each flag of a polytope to each other flag, we'll just require that a symmetry take each vertex to each other vertex.

Definition 3.0.1. A polytope is *vertex-transitive* if there is a symmetry taking any vertex to any other vertex.

Examples.

- All regular polytopes.
- $2n$-gon with alternating long and short side lengths.
- Prisms:

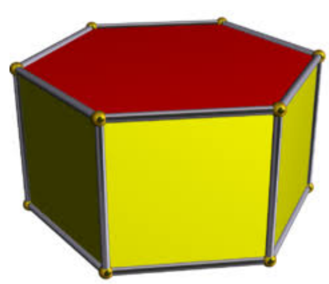

- Antiprisms:

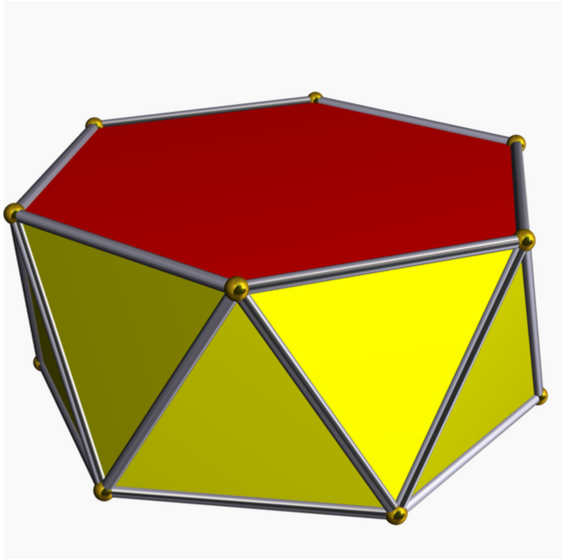

- This thing (truncated octahedron):

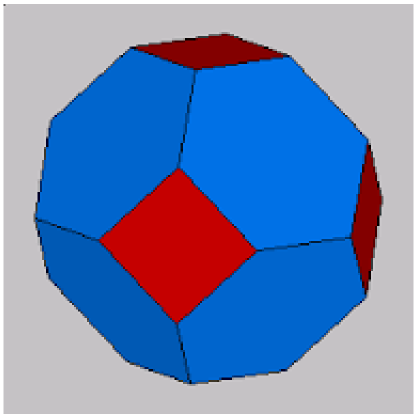

This last one is also an example where it is *not* possible to swap any two vertices.

- Soccer ball!

Unlike with regular polytopes, there are infinitely many of these in three dimensions! (For now, we're going to be sticking to three dimensions). What are some ways we can make really big vertex-transitive polytopes (where by really big, we mean "having arbitrarily many vertices")?

Well, there's prisms and antiprisms.

Theorem 3.0.2. *All sufficiently large vertex-transitive 3-polytopes are either prisms or antiprisms.*

To prove this, we're going to need a discrete version of Gauss-Bonnet.

3.1 Discrete Gauss-Bonnet

Definition 3.1.1. The *angle defect* $d(v)$ of a vertex v in a 3-polytope is the amount by which the angles incident to v fail to add to 2π. For F_v the set of faces containing v and $\alpha_f(v)$ the interior angle of the face f at v, we have

$$d(v) = 2\pi - \sum_{f \in F_v} \alpha_f(v).$$

Theorem 3.1.2 (Discrete Gauss-Bonnet)**.** *Let P be a 3-polytope, with V the set of vertices of that polytope. Then*

$$\sum_{v \in V} d(v) = 4\pi.$$

Proof. We'll start by proving that the angle defect at a vertex v is equal to its *discrete Gaussian curvature*, the area on the unit sphere bounded by a spherical polygon whose vertice are the unit normals of the faces around v. To do that, we're going to need a little bit of spherical geometry.

Lemma 3.1.3. *For a spherical triangle on the unit sphere with interior angles $\alpha_1, \alpha_2, \alpha_3$, the area of the spherical triangle is $A = \alpha_1 + \alpha_2 + \alpha_3 - \pi$.*

Proof of Lemma. Let A_1 be the subset of the sphere swept out by the angle α_1; let A_2 and A_3 be defined similarly. (TikZ Picture). Then

$$2A_1 + 2A_2 + 2A_3 = 4\pi + 4A.$$

But each A_i has area $2\alpha_i = 4\pi \cdot \frac{\alpha_i}{2\pi}$; viewed from the top, the fraction $\frac{A_i}{4\pi}$ is going to be $\frac{\alpha_i}{2\pi}$, since it sweeps out that percentage of the total angle. Thus

$$\begin{aligned} 4\pi + 4A &= 2A_1 + 2A_2 + 2A_3 \\ &= 4\alpha_1 + 4\alpha_2 + 4\alpha_3 \\ \Rightarrow \pi + A &= \alpha_1 + \alpha_2 + \alpha_3 \\ \Rightarrow A &= \alpha_1 + \alpha_2 + \alpha_3 - \pi. \end{aligned}$$

□

Lemma 3.1.4. *The area of a polygon with consecutive interior angles $\beta_1, \ldots, \beta_n$ on the unit sphere is*

$$A = (2-n)\pi + \sum_{i=1}^{n} \beta_i.$$

Proof of Lemma. Triangulate the polygon (without adding extra vertices). This takes $n-2$ triangles, and the total angle sum of these triangles is $\sum_{i=1}^{n} \beta_i$. Then by the previous lemma, the area of all these triangles is

$$\left(\sum_{i=1}^{n} \beta_i\right) - (n-2)\pi.$$

□

Now, consider a vertex v with faces N_1, N_4, and N_5, as in the following picture; we can treat this as an arbitrary vertex (we'll never use that there are five faces).

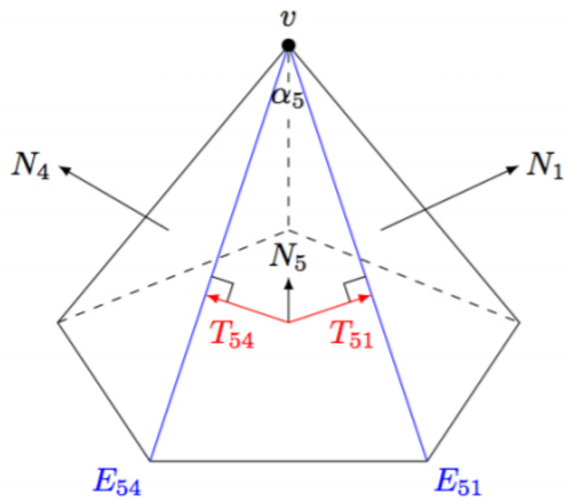

Let T_{54} be a tangent vector to the sphere at N_5 pointing towards N_4, and let T_{51} be a tangent vector to the sphere at N_5 pointing towards N_1. Then the angle between T_{54} and T_{51} is precisely the interior angle at N_5 on the spherical polygon. Both T_{54} and T_{51} are orthogonal to N_5, so they can be interpreted as lying on the face containing N_5. Note that E_{54} is orthogonal to both N_4 and N_5, so T_{54}, which is in the span of N_5 and N_4, is orthogonal to E_{54}. Similarly T_{51} is orthogonal to E_{51}. So we have in the face of N_5 the following quadrilateral:

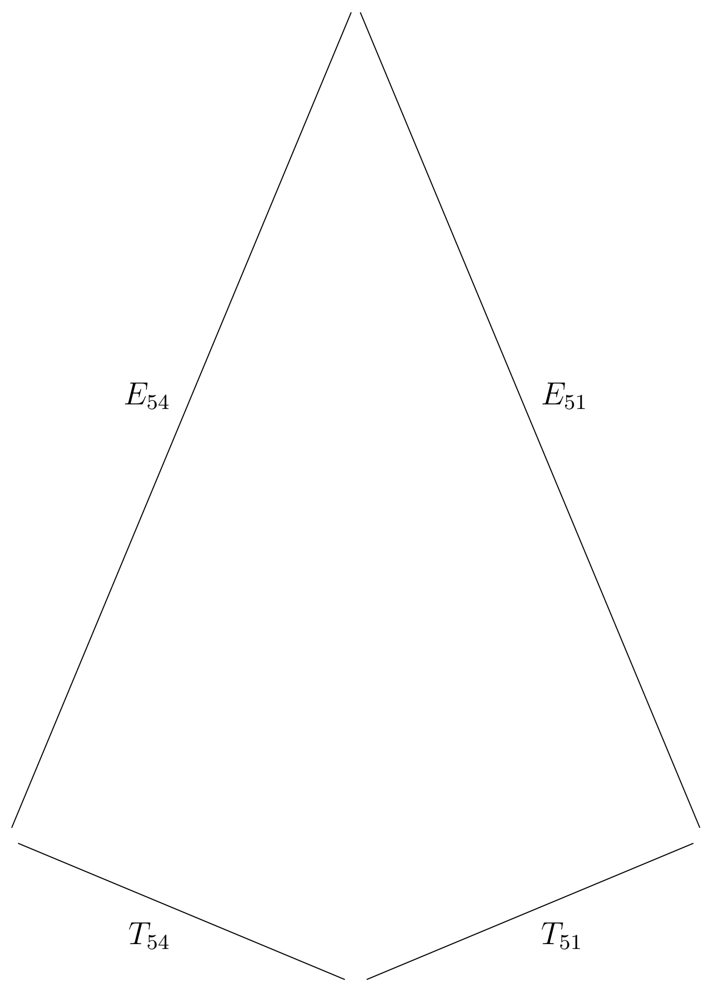

The angle at the top is α_1, and the two side angles are each $\pi/2$, so the angle at the bottom - or the angle of the spherical polygon - is $\pi - \alpha_1$. Thus the discrete Gaussian curvature at v is given by

$$\begin{aligned}(2-n)\pi + \sum_{f \in F_v} (\pi - \alpha_f(v)) &= 2\pi - n\pi + n\pi - \sum_{f \in F_v} \alpha_f(v) \\ &= 2\pi - \sum_{f \in F_v} \alpha_f(v) \\ &= d(v),\end{aligned}$$

so the angle defect at a vertex v is equal to its discrete Gaussian curvature.

But, now let's consider adding up the total discrete Gaussian curvature across the whole polytope. Well, (ADD TIKZ HERE) imagine placing a unit ball within the polytope. Then the spherical polygons with area the discrete Gaussian curvature at each vertex span the surface of the sphere, so the total discrete Gaussian curvature over all vertices must be 4π, the surface area of a unit sphere.

Thus

$$\sum_{v\in V} d(v) = \sum_{v\in V}(2\pi - \alpha(v)) = 4\pi.$$

□

3.2 Proof of Theorem 3.0.2

So, equipped with discrete Gauss-Bonnet, let's talk about vertex-transitive polyhedra. We'll use one last couple of observations. First, the fact that the average angle in an n-gon is $\frac{n-2}{n}\pi$. For the second observation, for any angle θ in a face of our vertex-transitive polyhedron, let $\overline{\theta}$ be the average angle in that face. So, if θ is an angle in a pentagon, then $\overline{\theta} = \frac{3\pi}{5}$. For each face, the sum of all angles θ is the same as the sum of all angles $\overline{\theta}$. So then

$$\sum_{\theta \text{ any angle in } P} \theta = \sum_{\theta \text{ any angle in } P} \overline{\theta}.$$

Let's fix a vertex x in our polytope. Then we can list the polygonal faces around x. We'll say that x is contained in an n_1-gon, an n_2-gon, ..., and an n_k-gon, where

$$n_1 \geq n_2 \geq \cdots \geq n_k.$$

If α_i is the angle in the n_i-gon at v, then by the above angle equation, we get

$$\begin{aligned} |V|\sum_{i=1}^{k}\alpha_i &= |V|\sum_{i=1}^{k}\frac{n_i-2}{n_i}\pi \\ \Rightarrow \sum_{i=1}^{k}\alpha_i &= \sum_{i=1}^{k}\frac{n_i-2}{n_i}\pi. \end{aligned}$$

An important comment is that $k < 6$. Why? Well, the total angle sum around a vertex must be strictly less than 2π, and each $\frac{n_i-2}{n_i}\pi \geq \frac{\pi}{3}$. So, we can have at most five of these terms.

By discrete Guass-Bonnet,

$$\begin{aligned} &\sum_{v\in V} d(v) = 4\pi \\ \Rightarrow &\sum_{v\in V}\left(2\pi - \sum_{i=1}^{k}\frac{n_i-2}{n_i}\pi\right) = 4\pi \\ \Rightarrow &\left(2\pi - \sum_{i=1}^{k}\frac{n_i-2}{n_i}\pi\right)|V| = 4\pi \\ \Rightarrow &\sum_{i=1}^{k}\frac{n_i-2}{n_i}\pi = 2\pi - \frac{4\pi}{|V|} < 2\pi. \end{aligned}$$

Yes, it's a little frustrating that we needed all that with Gauss Bonnet for this little thing (nb: this means we need discrete Gauss-Bonnet for the reg polytopes thing, if this is something we really wanna prove (but, maybe dGB is cool and worth doing?); if not we should sweep it under the rug here as well).

Our goal is to classify the tuples $(n_1, \ldots, n_k)$. Note that for any fixed tuple, we know exactly how many vertices are on our polyhedron by Gauss-Bonnet. Similarly, for any fixed n_1, we get a bound on the number of vertices per Gauss-Bonnet. So as our polyhedra become large, the number n_1 must approach infinity. But as n_1 grows quite large, $\frac{n_1-2}{n_1}$ approaches 1. So for large n_1,

$$\sum_{i=2}^{k} \frac{n_i - 2}{n_i}\pi < \pi + \varepsilon.$$

Each $\frac{n_i-2}{n_i}$ term is bounded below by $\frac{1}{3}$, so here are our options in terms of k:

$\underline{k = 5}$: Not possible for sufficiently large n_1, since $\frac{4}{3}\pi > \pi + \frac{1}{3}\pi$, so there are ε's small enough that this doesn't work.

$\underline{k = 4}$: Our only choice is $n_2 = n_3 = n_4 = 3$; everything else is larger, and in fact larger than $\pi + \frac{1}{6}\pi$.

$\underline{k = 3}$: The options that are small enough are $n_2 = n_3 = 4$, $n_2 = 4$ and $n_3 = 3$, and $n_2 = n_3 = 3$.

So we have four options for our choice of $(n_1, \ldots, n_k)$: namely $(n_1, 3, 3, 3)$, $(n_1, 4, 4)$, $(n_1, 4, 3)$, and $(n_1, 3, 3)$. The last two are not physically constructible (homework). The first forms an antiprism (drawing), and the second forms a prism (drawing). This completes the proof.

4 Vertex-Transitive Simplicial Polytopes

So that's a healthy discussion of vertex-transitive 3-polytopes. But, what about in higher dimensions? Well, to help our approach, we're going to need to talk about Euler characteristic.

4.1 Euler Characteristic

Let's start by thinking in 3D. The *Euler characteristic* of a polyhedron with V vertices, E edges, and F faces is $\chi = V - E + F$.

Now, for convex polyhedra, the Euler characteristic is always $\chi = 2$. Why? Well, if we pump up our polyhedra with extra air to make them spherical, they just become graphs on the sphere. Then after that we can fold them out around one face, so this becomes Euler's formula for planar graphs $V - E + F$, but with an extra face for the outside face. (Add drawing!)

Lemma 4.1.1 (Euler's formula for connected graphs)**.** *Given a connected planar graph G, with V vertices, E edges, and F planar faces, $V - E + F = 1$.*

Proof. We can build up the graph from nothing, while keeping it connected at every step. We begin with a graph that is just one vertex; in this case $E = F = 0$, and $V = 1$, so Euler's formula holds. Then to expand the graph we either add a vertex and an edge connecting it to the rest of the graph, or add an edge between existing vertices.

Adding a vertex and an edge: In this case, we're not closing up a new face, but we are adding an edge and a vertex. So if V, E, and F were the old vertex, edge, and face numbers, with $V-E+F = 1$, the new ones are $V+1$, $E+1$, and F, and $V+1-E-1+F = V-E+F = 1$.

Adding an edge between existing vertices: In this case, we've either created a new face or split an old face in two. So the number of edges increases by one and the number of faces increase by one, whereas the number of vertices remains constant. But $V - E - 1 + F + 1 = V - E + F = 1$. $\square$

So on the sphere, we always get $V - E + F = 2$ for any polytope.

This is a much harder result, which we will black-box, but it turns out that this holds in all dimensions. If we have an n-polytope P, where the number of i-faces is f_i, then

$$\sum_{i=0}^{n-1} (-1)^i f_i$$

is a constant depending only on n. Moreover (I can talk about this in terms of cell complexes a bit, but I don't know if it's worth it), if n is odd, this number is 2, and if n is even, this number is 0.

4.2 Vertex-transitive simplicial polytopes

So, using the concept of this generalized Euler characteristic, we can prove the following result.

Definition 4.2.1. An n-polytope is *simplicial* if all of its $(n-1)$-facets are simplices. These are not necessarily regular simplices; a *simplex* is a polytope in which every subset of vertices are the vertices of some face.

Theorem 4.2.2. *For n odd, the number of vertices in a vertex-transitive simplicial n-polytope is bounded by a function of n. In other words, an odd-dimensional vertex-transitive simplicial n-polytope can't have arbitrarily many vertices.*

Proof. We consider the Euler characteristic. Let P be a vertex-transitive simplicial n-polytope, for n odd. Let f_i be the number of i-faces of P. Let v be an arbitrary vertex of P. Let v_i be the number of i-faces containing v. So, $v_0 = 1$, and v_1 is the number of edges with v as an endpoint, and so on.

Now, we'll count "the number of pairs (x, f) of vertices x and i-faces f, where $x \in f$" in two ways. In the first way, for each vertex, we can count the number of i-faces containing that vertex. There are v_i i-faces containing v by definition; by vertex transitivity, there are v_i i-faces containing each vertex x.

In the second way, for each face, we can count the number of vertices it contains. There are f_i i-face total, and each one is a simplex, so it must contain $i+1$ vertices. Thus

$$v_i f_0 = f_i(i+1),$$

so

$$f_i = f_0 \frac{v_i}{i+1}.$$

Now let's look at the Euler characteristic. Since this is an n-polytope for n odd,

$$\begin{aligned}
\sum_{i=0}^{n} (-1)^i f_i &= 2 \\
\Rightarrow \sum_{i=0}^{n} (-1)^i f_0 \frac{v_i}{i+1} &= 2 \\
\Rightarrow f_0 \left(\sum_{i=0}^{n} (-1)^i \frac{v_i}{i+1} \right) &= 2 \\
\Rightarrow f_0 \left(\sum_{i=0}^{n} (-1)^i v_i \frac{n!}{i+1} \right) &= 2n!
\end{aligned}$$

Every term in the sum on the LHS is an integer, so the sum on the LHS is an integer. But then f_0 divides $2n!$, so it must be bounded. □

Homework exercises: What's the largest number of vertices on a vertex-transitive simplicial 3-polytope? 5-polytope (omg is this possible)? Can you prove the same statement weakened to vertex-regular (i.e. all vertices have the same star, or all vertices "look the same")? Can you prove the same statement weakening simplicial to the condition that the number of vertices per facet is bounded by a constant? What fails in even dimensions?