Reflection Groups

Notes for a short course at the Ithaca High School Senior Math Seminar

Balázs Elek
Cornell University

December 5, 2017

Contents

1 Three examples

We introduce three seemingly different kinds of symmetry and try to find something common in them.

1.1 Permutations

Consider the permutations of the numbers $1, 2$, and 3. We will use so-called one-line notation to refer to them, i.e. writing 213 means that $1, 2, 3$ have been rearranged so that 2 is first, 1 is second and 3 is third (Note that we could have chosen a, b, c, x, y, z, or any other placeholders for this, the fact that $1, 2, 3$ are the first four positive integers is irrelevant here).

Exercise 1.1. *How many permutations of* $1, 2, 3$ *are there?*

The above permutations can be viewed as symmetries of the set $\{1, 2, 3\}$, e. g. the set $\{2, 1, 3\}$ is still the same (since it has the same elements) as $\{1, 2, 3\}$. If that is true, however, one ought to be able to compose these symmetries to obtain new ones. How to do that with permutations? Let $\sigma = 132$ and $\pi = 213$ be two permutations of $1, 2, 3$. One can also think of them as functions from $\{1, 2, 3\}$ to itself, where $\sigma(1) = 1, \sigma(2) = 3$, and $\sigma(3) = 3$, and similarly for π. Then, in effect, π swaps 1 and 2, whereas σ swaps 2 and 3. So if we want to see what happens when we perform σ first, and then π, then we should do these in order:

$$\begin{tikzcd} 123 \arrow[d, "\sigma"] \\ 132 \arrow[d, "\pi"] \\ 231 \end{tikzcd}$$

We will write this as $\pi\sigma$. Now, you might wonder, why are we writing π first if it is σ that is happening first. The reason for this is that we are thinking of both σ and π as *functions*, and $\pi\sigma$ refers to their composition. For instance, to find where 3 is sent when we perform σ first then π, we should compute $\sigma(3) = 2$, then $\pi(\sigma(3)) = \pi(2) = 1$.

Exercise 1.2. *What happens if we compose* 321 *with itself?*

Exercise 1.3. *We know* $\pi\sigma = 231$*, what is* $\sigma\pi$*?*

Exercise 1.4. *The permutations* π*,* σ*, and* 321 *all satisfy* $\pi^2 = \sigma^2 = (321)^2 = 123$ *is this true for any permutation?*

1.2 Symmetries of the triangle

Take an equilateral triangle in the plane. What sort of symmetries does it have? We can rotate it by integral multiples of $120°$, reflect it across three lines of symmetry, or maybe do many of these in a row.

Exercise 1.5. *How many symmetries does the triangle have? How about the* n*-gon?*

In this case, it is straightforward how to compose symmetries, we just apply them one after another, the result is clearly another symmetry. Call reflection across the red line of symmetry r and reflection across the green line g.

Exercise 1.6. *Compare what happens if we*

(a) Do r *first then* g*.*

(b) Do g *first then* r*.*

1.3 Words in the alphabet $\{a, b\}$.

A *word* is a sequence of letters in an alphabet. Let's start modestly with an alphabet consisting of only two letters a and b (In fact, from a mathematical perspective, there is little difference between alphabets with two or twenty-six letters). We will call a and b generators, the "empty word" e, and say that some words have the "same meaning", in particular $aa = bb = ababab = e$ (these are called relations), and this should hold even when words are embedded in other words. To clarify, for example

$$abba = a(bb)a = aa = e.$$

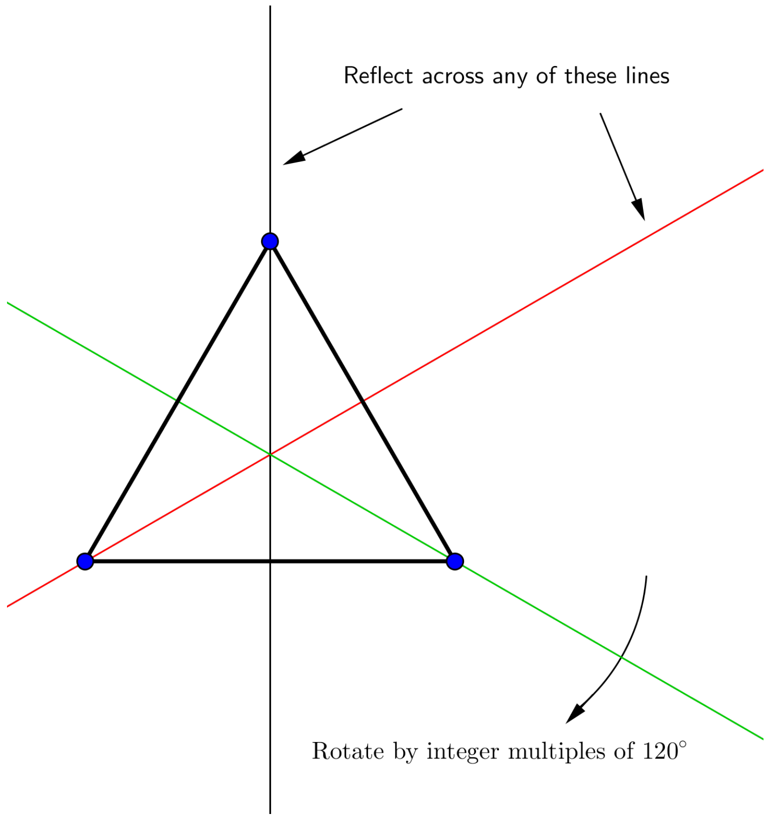

Figure 1: Symmetries of the triangle

Exercise 1.7. *How many words can you find in this "language"? I.e. how many words are there with distinct "meaning"?*

How should we define the composition (which seems to be a theme now) of words? Simply concatenate, keeping in mind the relations.

Exercise 1.8. *Is $aba = bab$?*

1.4 Why these are the same symmetries

All of the examples above really quantify the same type of symmetry (one which we will call type A_2 soon). If we correspond $\sigma \leftrightarrow r \leftrightarrow a$, and $\pi \leftrightarrow g \leftrightarrow b$, then all the symmetries of the above three objects behave the same way! That is, if some equality holds in one of the cases, like $\sigma\pi\sigma = \pi\sigma\pi$, then in the other cases, the same equality must hold ($rgr = grg$ and $aba = bab$). This shows that these are *isomorphic*, which means that they quantify the same kind of symmetry.

1.5 Groups

When an object has any kind of symmetry, there is always a *group* that describes that symmetry. We will be studying *reflection groups*, which are symmetries generated by reflections. Not all groups are reflection groups, here are some examples:

1. Rotational symmetries of the hexagon.
2. Rational numbers with composition being addition.
3. Symmetries of the circle[1].

[1] Technically, the symmetry group of the circle *is* generated by reflections, but not by finitely many ones. We will focus our attention on *finitely generated* groups.

2 Reflection Groups

2.1 Reflections on the plane

So what is a reflection after all? We will start by examining reflections in $\mathbb{R}^2$, where it is clear what they should be, and try to extend this concept further. So take a line l across the origin in $\mathbb{R}^2$ with a vector $\mathbf{v}$ on it that is not the zero vector, and a nonzero vector $\mathbf{w}$ that lies on the line through the origin perpendicular to l. We will call the

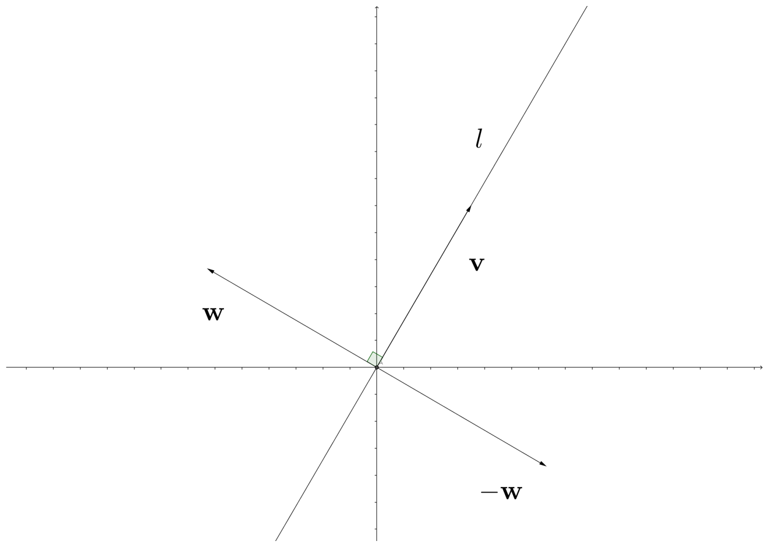

Figure 2: Reflection across a line

reflection $s_l : \mathbb{R}^2 \to \mathbb{R}^2$. We know that

$$s_l(\mathbf{v}) = \mathbf{v}, \quad s_l(\mathbf{w}) = -\mathbf{w}.$$

How do we figure out what happens to other vectors?

Exercise 2.1. *Find $s_l(\mathbf{v} + \mathbf{w})$ in terms of $\mathbf{v}$ and $\mathbf{w}$.*

Exercise 2.2. *Find $s_l(c_1\mathbf{v} + c_2\mathbf{w})$ where c_1, c_2 are scalars in terms of $\mathbf{v}, \mathbf{w}, c_1$ and c_2.*

Remark 2.3. *A transformation satisfying the above property is called a linear transformation.*

So we know how to compute $s_l(\mathbf{q})$ for any vector $\mathbf{q}$ in the plane! Just find c_1 and c_2 such that $\mathbf{q} = c_1\mathbf{v} + c_2\mathbf{w}$, and use the previous result.

Exercise 2.4. *Check that for any $\mathbf{q} \in \mathbb{R}^2$,*

$$s_l(\mathbf{q}) = \mathbf{q} - \frac{2(\mathbf{q} \cdot \mathbf{w})}{(\mathbf{w} \cdot \mathbf{w})}\mathbf{w},$$

where $\cdot$ denotes the dot product.

Exercise 2.5. *Use the previous result to compute*

$$s_l(s_l(\mathbf{q})).$$

2.2 Interacting reflections

Now we will take a look at how two reflections interact with each other. For simplicity, assume that we have chosen the line l to be the x-axis, and pick up another line e with any angle θ between e and l. We will chose two vectors normal to the lines. Let $\mathbf{v} = (0, 1)$, which is perpendicular to l, and choose a unit length vector $\mathbf{w}$ normal to e such that the angle between $\mathbf{w}$ and $\mathbf{v}$ is obtuse.

Exercise 2.6. *Express $\mathbf{w}$ in coordinates.*

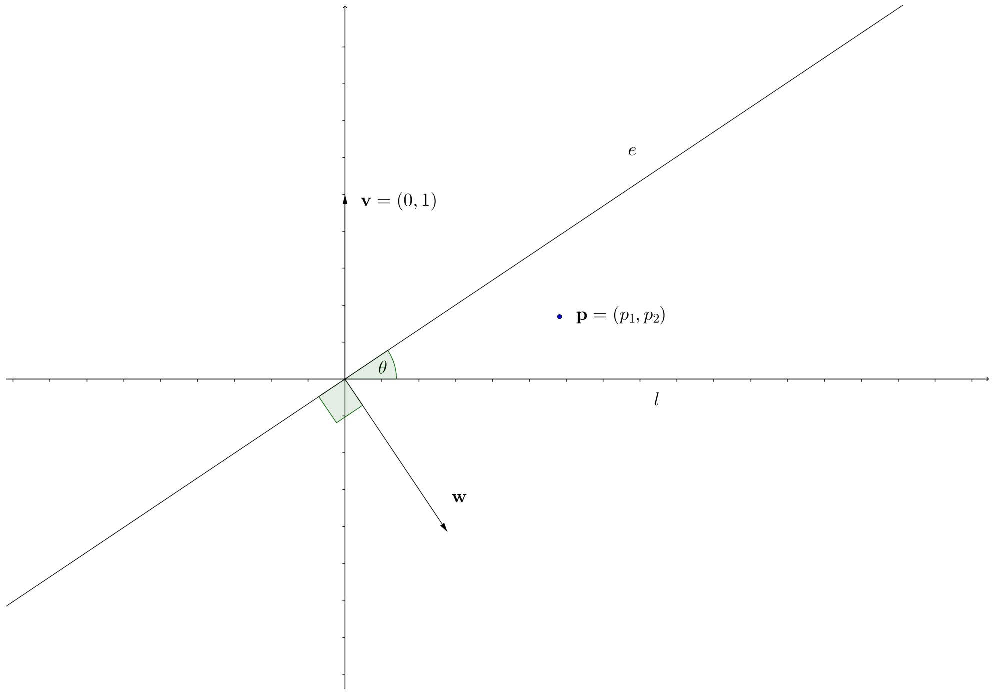

Figure 3: Reflections across two lines

Theorem 2.7. *Let $\mathbf{p} = (p_1, p_2)$ be an arbitrary vector in the plane, $l, e, \mathbf{v}, \mathbf{w}$ as above. Then*

$$s_e(s_l(\mathbf{p})) = ((\cos 2\theta)p_1 - (\sin 2\theta)p_2, (\sin 2\theta)p_1 + (\cos 2\theta)p_2),$$

i.e. $s_e(s_l(\mathbf{p}))$ is $\mathbf{p}$ rotated anticlockwise by an angle 2θ.

Proof. We know that $s_l(\mathbf{p}) = s_l(p_1, p_2) = (p_1, -p_2)$. Then,

$$\begin{aligned}
s_e(s_l(\mathbf{p})) &= s_l(p_1, -p_2) \\
&= (p_1, -p_2) - \frac{2\left((p_1, -p_2) \cdot (\sin\theta, -\cos\theta)\right)}{((\sin\theta, -\cos\theta) \cdot (\sin\theta, -\cos\theta))}(\sin\theta, -\cos\theta) \\
&= (p_1, -p_2) - 2\left(p_1 \sin\theta + p_2 \cos\theta\right)(\sin\theta, -\cos\theta) \\
&= (p_1, -p_2) + (-2p_1 \sin^2\theta - 2p_2 \sin\theta\cos\theta, 2p_1 \sin\theta\cos\theta + 2p_2 \cos^2\theta) \\
&= ((\sin^2\theta + \cos^2\theta)p_1 - 2p_1 \sin^2\theta - p_2 \sin 2\theta, -(\sin^2\theta + \cos^2\theta)p_2 + p_1 \sin 2\theta + 2p_2 \cos^2\theta) \\
&= (\cos^2\theta - \sin^2\theta)p_1 - \sin 2\theta p_2, \sin 2\theta p_1 + (\cos^2\theta - \sin^2\theta)p_2) \\
&= ((\cos 2\theta)p_1 - (\sin 2\theta)p_2, (\sin 2\theta)p_1 + (\cos 2\theta)p_2)
\end{aligned}$$

Q.E.D.

Exercise 2.8. *Find $s_l(s_e(\mathbf{p}))$.*

Recall that $s_l^2 = s_e^2 = e$, the identity transformation which leaves the whole plane fixed.

Exercise 2.9. *How many different transformations are there which can be written down as successive applications of s_l and s_e?*

2.3 Reflection Groups

The above is the notion of a *group generated by* the reflections s_l and s_e. We say that a group of symmetries is a *reflection group* if it can be generated by finitely many reflections in some Euclidean space.

Before moving on, we should discuss what reflections are not in the plane, but space, or higher-dimensional Euclidean spaces. The key here is to find the correct generalization of the concept of reflections in the plane. In $\mathbb{R}^2$, a reflection fixes a line through the origin, and sends any vector perpendicular to the line to its negative. But there is more room in 3-dimensional space than in the plane, and the above two properties do not determine a transformation yet. So we have to make a decision: do we want to fix a plane containing the origin and send any of its normal vectors to their negatives, or do we want to fix a line through the origin, and send every vector perpendicular to it to its negative? Notice that in the plane, reflections and rotations are distinct, in particular, we can never produce a reflection by composing different rotations. We would like to keep this distinction, and then it becomes clear that a reflection in $\mathbb{R}^3$ should fix a plane through the origin, and send only a single line's worth of vectors to their negatives (or, in higher dimensions, it should fix a *hyperplane*, and send only a line's worth of vectors to their negatives).

2.4 Back to our three examples

2.4.1 Permutations

We will use the notation S_n to refer to the group of permutations of $12\dots n$. Let $\mathbf{e_1}, \mathbf{e_2}, \mathbf{e_3}$ be the unit coordinate vectors in $\mathbb{R}^3$, i.e.

$$\mathbf{e_1} = (1,0,0), \dots, \mathbf{e_3} = (0,0,1).$$

Now we will apply the permutations to permute the *indices* $1,2,3$. So, if, as before $\sigma = 132, \pi = 213$, then

$$\sigma(a,b,c) = (a,c,b),$$
$$\pi(a,b,c) = (b,a,c).$$

We will call a permutation π satisfying (for $i \neq j$)

$$\pi(k) = \begin{cases} i & \text{if } k = j, \\ j & \text{if } k = i, \\ k & \text{otherwise} \end{cases}$$

a *transposition*. So transpositions are nontrivial permutations that send all but two numbers to themselves (for example, σ, π and 321 are transpositions but $\sigma\pi$ and $\pi\sigma$ are not)

Now we claim that if we use permutations to permute the indices of the unit coordinate vectors, then transpositions correspond to reflections.

Exercise 2.10. *Find a vector* $\mathbf{v}$ *sent to* $-\mathbf{v}$ *by* σ*.*

Now to show S_3 is a reflection group, we need to write down any element of it as a product (composition) of reflections.

Exercise 2.11. *Let* $\sigma = \sigma(1)\sigma(2)\sigma(3)$ *be a permutation. Write it down as a product of transpositions.*

Exercise 2.12. *Consider* S_3*, i.e. permutations of* 123*. Similarly as above, we can view these as transformations of* $\mathbb{R}^3$*. Using that realization, find a regular polygon somewhere in* $\mathbb{R}^3$ *whose symmetry group is precisely* S_3*.*

Exercise 2.13. *Try to generalize Exercise 2.12 to the case of* S_4*. What is the geometric object whose symmetries correspond precisely to* S_4*?*

2.4.2 Symmetries of the triangle or words in the alphabet $\{a,b\}$

Consider the following figure: Let a, b denote reflections across the red and green lines, respectively. Notice that $a^2 = b^2 = e$, the identity transformation. Moreover, as ab is rotation by $\frac{2\pi}{3}$, $(ab)^3$ is rotation by $360°$, i.e. $(ab)^3 = e$ (it is important to note that (ab) and $(ab)^2$ both do not equal e). Since any word in our alphabet can be, by definition, written down using a and b, and we have realized these as reflections in $\mathbb{R}^2$, this makes the set of our words a reflection group.

Exercise 2.14. *Can you find a "language" in the alphabet* $\{a,b\}$ *whose words correspond to symmetries of the hexagon?*

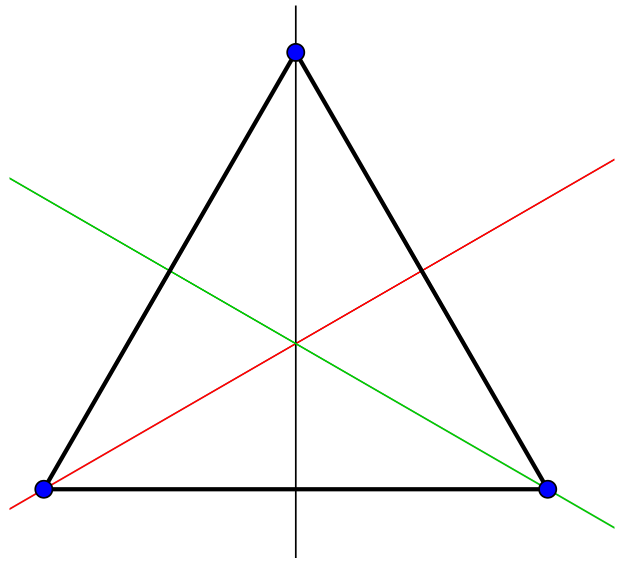

Figure 4: Generating symmetries of the regular triangle

2.5 Moral of the story

So what to take away from this? Many groups of symmetries can be reflection groups, even if they seemingly have nothing to do with Euclidean space. So the important thing when dealing with groups is to realize that it's only the symmetries that count, not the object which has those symmetries (for example, the sets $\{1, 2, 3, 4\}$ and $\{\mathbf{e_1}, \mathbf{e_2}, \mathbf{e_3}, \mathbf{e_4}\}$ both have symmetry groups S_4, but they are very different).

3 Regular Polytopes

3.1 Regular Polygons

Before jumping in the general theory of reflection groups, by now we suspect that symmetry groups of regular polygons are reflection groups. Indeed, if we draw all the symmetry lines of our n-gon and take two adjacent ones,

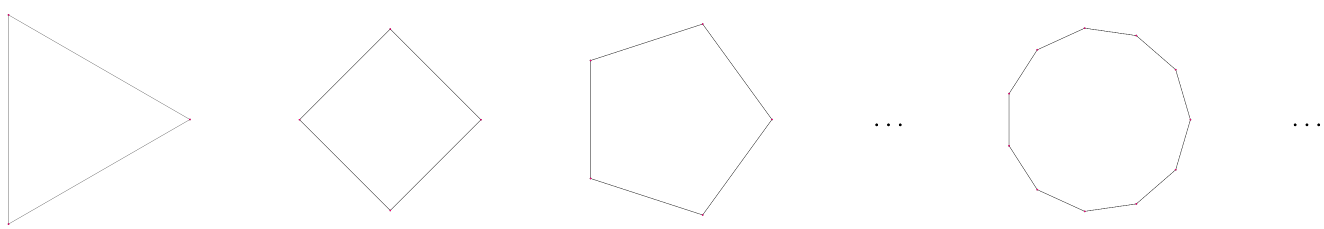

Figure 5: Regular polygons

the reflections across them will generate all the symmetries.

Exercise 3.1. *Let a and b denote reflections across the two adjacent lines mentioned above. Write down the relations for the symmetry group (as in Exercise 2.14).*

Symmetry groups of regular polygons are called *dihedral groups*.

3.2 Regular Polyhedra

How would you generalize the concept of a regular polygons in the plane to solids in space? We want a 3-dimensional convex[2] object whose faces (edges) of regular polygons are all the same, and they always meet at the same angle

[2] There are also nonconvex ones, but we'll politely ignore them for now.

at a vertex. If we go along these lines, the 3-dimensional analogues of regular polygons are the Platonic solids:

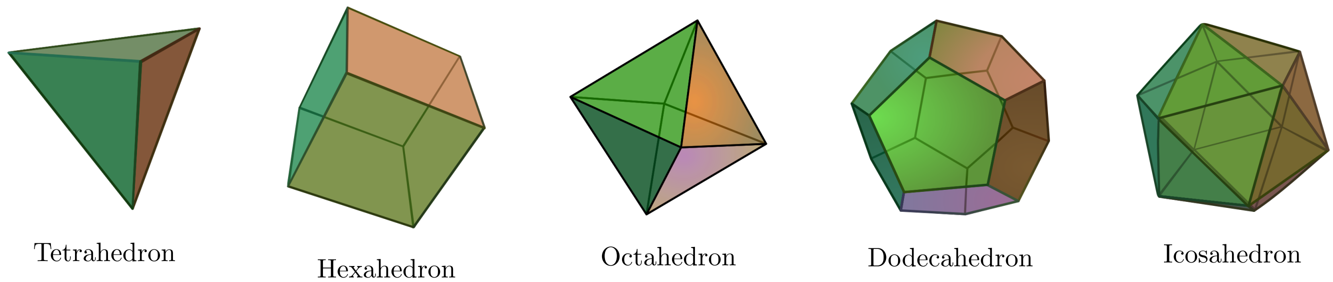

Figure 6: Convex regular polyhedra

One curious thing about them is that while there are infinitely many regular polygons, there are only five regular polyhedra (meaning that their faces all have to be the same regular polygons meeting at the same angles at every vertex). This mesmerized mathematicians as early as ancient Greece, and in fact, the final book of Euclid's Elements constructs the Platonic solids inscibed in a sphere and proves that besides these five, there are no more regular solids.

Let's take a look at the symmetry group of the dodecahedron[3].

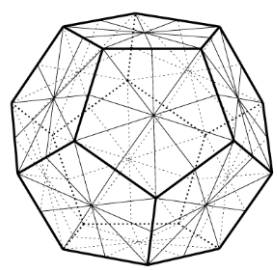

Figure 7: Symmetry planes of the dodecahedron

The dodecahedron has 15 planes of symmetry (can you count them on the picture?), which subdivide its faces to 120 congruent triangles.

Exercise 3.2. *How many symmetries does the dodecahedron have? (Hint: what happens to the triangles if we apply a symmetry of the dodecahedron?)*

It turns out that the symmetry groups of regular polyhedra are also all reflection groups. In fact, the reason why there are only five Platonic solids is closely related to the fact that there are only finitely many finite reflection groups of rank three (which, roughly speaking means that they act as symmetries of objects in 3-dimensional space)! Later we are going to determine them and see how they correspond to the regular polyhedra. But why should we stop in three dimensions?

3.3 Regular Polychora

We can just carry our previous definition of a regular polyhedron one step further, let's say we want to build a 4-dimensional convex object whose "faces" (remember, these should be 3-dimensional) should be regular *polyhedra*, all fitting together nicely and meeting at the same angles at every vertex. Then we arrive at the notion of a regular *polychoron*, or 4-*polytope*.

[3]Picture from [3].

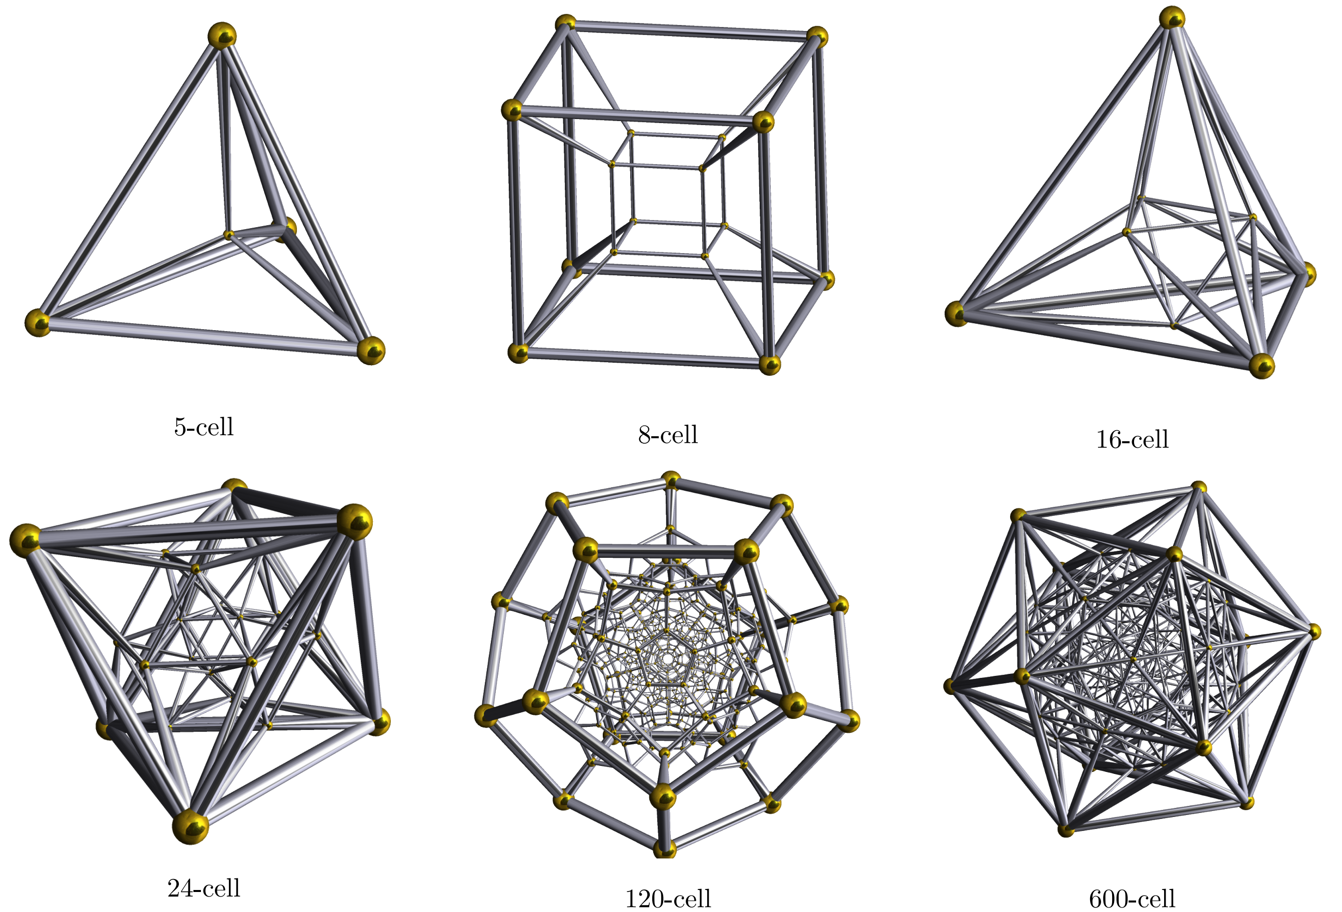

Figure 8: Schlegel projections of convex regular polychora

You might be puzzled when looking at Figure 8 above, and how they are supposed to represent these polychora which are in reality, 4-dimensional. If you imagine standing outside a big cube whose 2-dimensional faces are made of glass, you might see something like this:

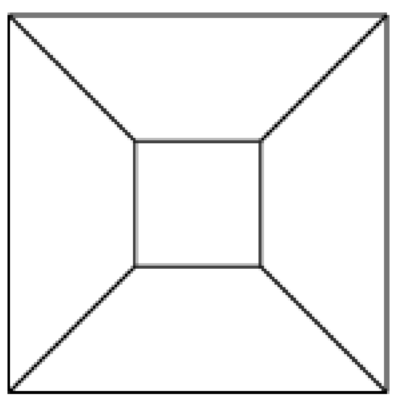

Figure 9: Schlegel projection of a cube

This is called the *Schlegel projection* of the cube. Note that we can recover all the combinatorial information from this projection (which of the faces contain or are adjacent to other faces).

Exercise 3.3. *Identify which regular polyhedron we have projected to the plane:*

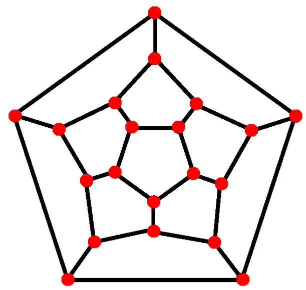

Figure 10: A regular polyhedron projected to the plane

Now we can do the exact same thing to our regular polychora: Imagine standing outside of them in 4-sphace, then project down to $\mathbb{R}^3$! The images in Figure 8 are obtained this way.

Another frequently used method is the following: Just as regular polygons can be inscribed in a circle, or polyhedra in a sphere, regular polychora can be inscribed in a 3-sphere. Then we can "inflate" our polytope so that it lies on the surface of the sphere. Now in the cases of the circle and the ordinary sphere, we can pretend to put our circle on a line, or our sphere on the plane, and project down from the North pole (this is called *stereographic projection*). This way, every point in our circle/sphere except the North pole corresponds to exactly one point of the line/plane. The edges of the polytope will look curved after this projection, for example, the 8-cell projected to $\mathbb{R}^3$ will look like this:

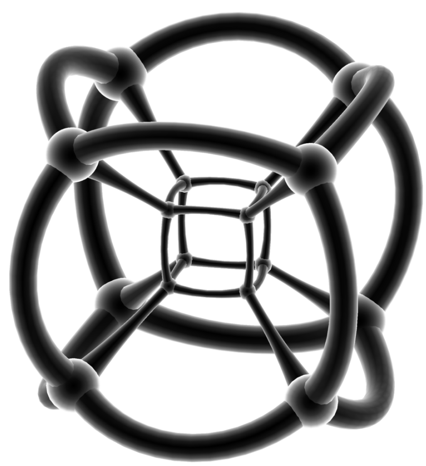

Figure 11: Stereographic projection of the 8-cell

Exercise 3.4. *Identify the building blocks of each of the convex regular polychora in Figure 8 (i.e. what regular polyhedra they are made of).*

Exercise 3.5. *Which convex regular 4-polytope has been "folded out" to $\mathbb{R}^3$ in Figure 12 below?*

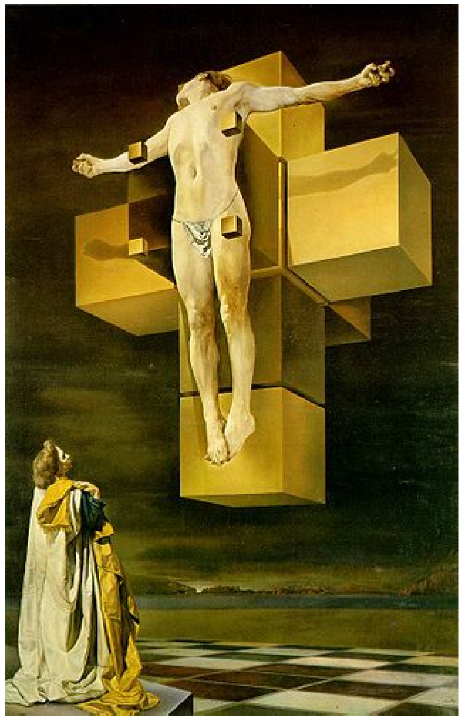

Figure 12: Salvador Dalí : Crucifixion

As expected, the symmetry groups of all the convex regular polychora are reflection groups.

3.4 And so on ...?

Now it seems there is nothing stopping us from moving higher and higher in dimensions, and construct more regular polytopes. There are a couple straightforward ones. One is the n-simplex, which is defined recursively by taking the $n-1$ simplex, and connecting all vertices to a new vertex. The 0-simplex is taken to be a point.

Exercise 3.6. *Identify the simplices in Figures 5, 6, and 8.*

The other family which is easily constructed is the family of the n-cubes. They are also defined recursively: Take the $n-1$ cube, double it, and connect the corresponding edges, again, the 0-cube is taken to be a point.

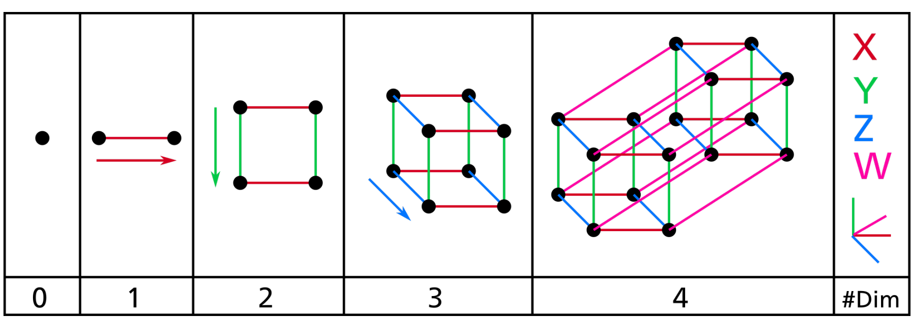

Figure 13: Constructing the first four hypercubes

Having seen Figure 13 above, there really is no point giving you the exercise to locate them in Figures 5, 6, and 8. However, there is another family of regular polytopes closely related to the n-cubes, the *n-cross-polytopes.* They are related to the cubes like this: Take the n-cube, and turn each of its faces (remember, its faces are $n-1$-cubes) to a vertex. Connect two vertices with an edge if the faces they correspond to are adjacent (that is, they intersect along an $n-2$-cube). Now, whenever three faces intersect in an $n-3$-cube, fill in the triangle corresponding to them in the cross-polytope, and continue this process. This procedure is called *dualizing* the polytope.

Exercise 3.7. *Identify the cross-polytopes in Figures 5, 6, and 8.*

Exercise 3.8. *What happens if you dualize the n-simplex?*

Exercise 3.9. *What happens if you dualize the dodecahedron? The icosahedron?*

Exercise 3.10. *Show that dual polytopes have identical symmetry groups.*

So far we have seen three infinite families of regular polytopes: the n-simplices, the n-cubes and their duals, the n-cross-polytopes. But how about the other guys? The dodecahedron, icosahedron, 24-cell, 120-cell and 600-cell? It turns out that it is dimension 4 where their story ends. They have absolutely no analogues in higher dimensions, moreover, from dimensions 5 onwards, the only regular polytopes are the n-simplices, the n-cubes and the n-cross-polytopes. The reason behind this fascinating result is the lack of exceptional finite reflection groups that would serve as these polyotope's symmetry groups.

4 Root systems

We have remarked previously that the reason for the lack of regular polytopes in higher dimensions was the lack of suitable symmetry groups. We are now going to classify finite reflection groups. The way we are going to do this is to find a nice combinatorial gadget that we can associate to a reflection group, and classify those.

4.1 A_2 and B_2

We have defined reflectiong groups to be groups generated by reflections in Euclidean space. However, looking at the examples in Section 2.4, it seems that all the stuff in $\mathbb{R}^n$ is not really necessary to describe our reflection group W. For instance, if we know s sends a vector $\mathbf{v}$ to $-\mathbf{v}$, then s sends every vector on the line through the origin and $\mathbf{v}$ to their negatives. It would be nice if we could think about what W does without concerning ourselves with all of $\mathbb{R}^n$. In particular, it would be great if we could find a *finite* set of vectors in $\mathbb{R}^n$, which are preserved by W, and are sufficient to describe its behavior.

Let's see what we can do with the symmetry group of the regular triangle (a.k.a. S_3, but later we'll call it A_2).

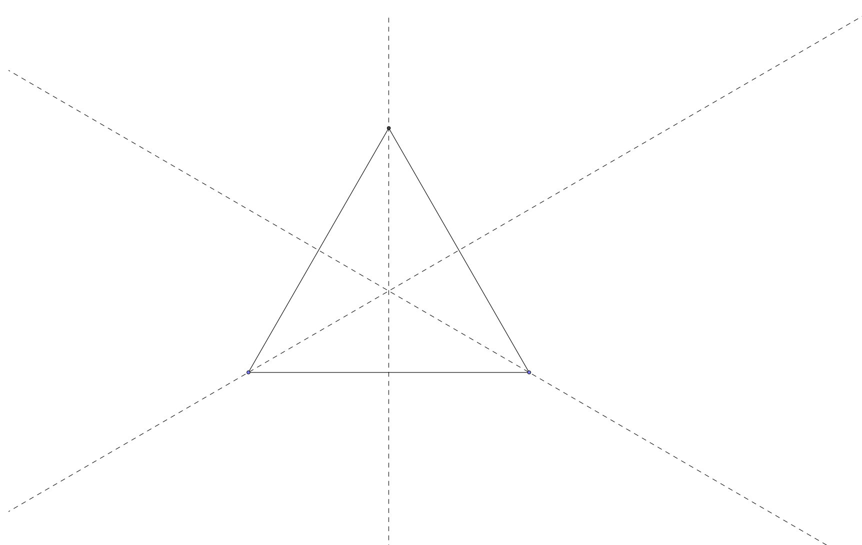

Figure 14: Symmetry lines of the regular triangle

Since any element of S_3 will preserve the triangle, it will also preserve these lines. However, we would like to keep track of the reflections sending certain vectors to their negatives, so let's pick two vectors of equal length that are perpendicular to each line

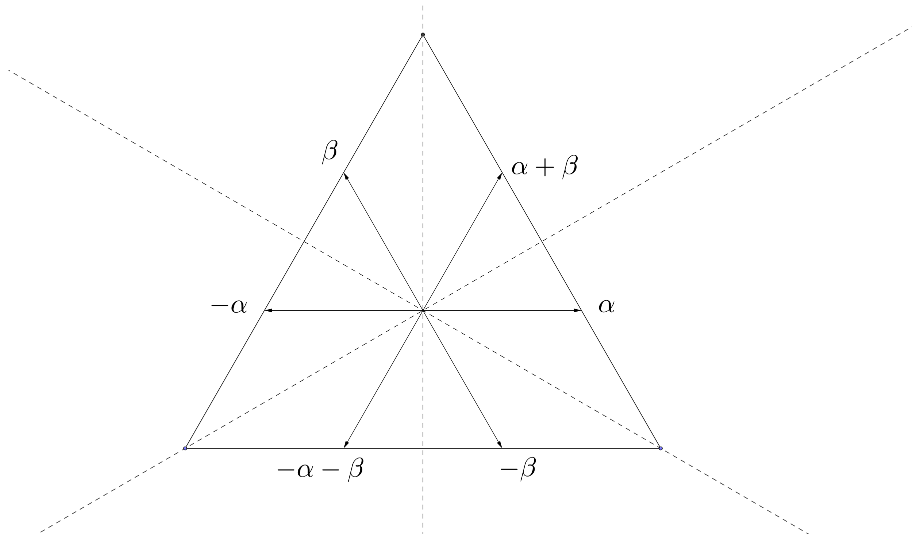

Figure 15: The root system of A_2

Now notice that any reflection along the symmetry lines will permute the set $\{\alpha, \beta, \alpha+\beta, -\alpha, -\beta, -\alpha-\beta\}$. So we have managed to turn this geometric action of a reflection group into combinatorics. We can write down in a table what happens to the roots if we apply elements of A_2 to them.

	α	β	$\alpha+\beta$	$-\alpha$	$-\beta$	$-\alpha-\beta$
e	α	β	$\alpha+\beta$	$-\alpha$	$-\beta$	$-\alpha-\beta$
s_α	$-\alpha$	$\alpha+\beta$	β	α	$-\beta-\alpha$	$-\beta$
s_β	$\alpha+\beta$	$-\beta$	α	$-\alpha-\beta$	β	$-\alpha$
$s_\alpha s_\beta$	β	$-\alpha-\beta$	$-\alpha$	$-\beta$	$\alpha+\beta$	α
$s_\beta s_\alpha$	$-\alpha-\beta$	α	$-\beta$	$\alpha+\beta$	$-\alpha$	β
$s_\alpha s_\beta s_\alpha$	$-\beta$	$-\alpha$	$-\alpha-\beta$	β	α	$\alpha+\beta$

Table 1: The action of A_2 on its root system

Exercise 4.1. *If someone covers the first column of Table 1, how would you try to identify which rows correspond to reflections and which to rotations?*

Let's try to replicate this for the symmetry group of the square (which later will be known as B_2).

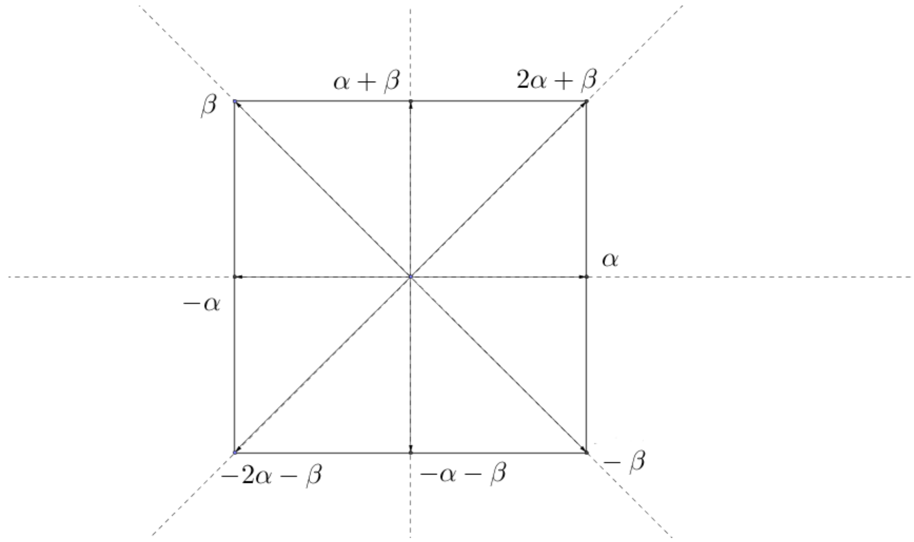

Figure 16: The root system of B_2

Notice that we have chosen our roots to be of two different lengths so we can write them nicely as nonnegative integer sums of α and β, just like we did with A_2. It is important to realize that this won't cause a problem, since the short roots will always be sent to the short roots, and long roots to long roots by B_2. At the moment, we are only concerned about the stability of our roots under the reflection group in question, not about the relative lengths of the roots. We will return to this integral feature later, but now we could have chosen all our roots to be, for instance, unit length, since we are only concerned about their combinatorics.

Exercise 4.2. *Write down a table like Table 1 for B_2.*

4.2 General root systems

In our discussion, we are following [2] quite closely. Motivated by the previous two examples, we are going to find root systems for an arbitrary finite reflection group. Let W be a finite reflection group acting on $\mathbb{R}^n$. Recall from Section 2.3 that each reflection $s_\alpha \in W$ determines a reflecting hyperplane H_α (which is fixed pointwise by s_α) and a line L_α perpendicular to it (whose points are sent to their negatives by s_α). As in the case of A_2 and B_2, we hope that the set of these lines $\{L_\alpha | s_\alpha \in W \text{ is a reflection}\}$ is preserved (i.e. permuted) by W. The following theorem does just that:

Theorem 4.3. *Let W be a finite reflection group acting on $\mathbb{R}^n$. Let $w \in W$, α any nonzero vector in $\mathbb{R}^n$. Then $s_\alpha \in W \Leftrightarrow s_{w(\alpha)} \in W$.*

Proof. For any nonzero $\beta \in \mathbb{R}^n$, we will denote by L_β the line through the origin and β, and by H_β the hyperplane perpendicular to L_β. First we claim that $s_{w(\alpha)} = ws_\alpha w^{-1}$. For this, we need to show that it fixes $H_{w(\alpha)}$ pointwise and sends every vector in $L_{w(\alpha)}$ to its negative.

Since

$$ws_\alpha w^{-1}(w(\alpha)) = ws_\alpha(\alpha) = w(-\alpha) = -w(\alpha),$$

the transformation $ws_\alpha w^{-1}$ indeed sends vectors in $L_{w(\alpha)}$ to their negatives.

And note that since w preserves dot products, $\lambda \cdot \alpha = (w(\lambda)) \cdot (w(\alpha))$, we have $\lambda \in H_\alpha \Leftrightarrow w\lambda \in H_{w(\alpha)}$, and

$$(ws_\alpha w^{-1})(w\lambda) = ws_\alpha \lambda = w\lambda,$$

so $ws_\alpha w^{-1}$ fixes $H_{w(\alpha)}$ pointwise.

Q.E.D.

Now we are ready to state the definition of a root system[4]

[4] This definition differs from the one most widely used in the literature, but it is best for our purposes. We have taken it verbatim from [2], p6.

Definition 4.4. *A* ***Root system*** Φ *is a finite set of vectors in* $\mathbb{R}^n$ *not containing the origin such that*

- $\Phi \cap L_\alpha = \{\alpha, -\alpha\} \quad \forall \alpha \in \Phi,$
- $s_\alpha \Phi = \Phi \quad \forall \alpha \in \Phi.$

If we now define W_Φ to be the group generated by the reflections $\{s_\alpha | \alpha \in \Phi\}$ then W_Φ is a reflection group by construction.

Exercise 4.5. *Prove that the group* W_Φ *is in fact finite (this is on the harder side if you have not taken some abstract algebra).*

Now if we have any finite reflection group W in mind, we can build a root system out of it by taking all the reflecting hyperplanes for the reflections in W, and picking two opposite unit vectors perpendicular to each of them. Using Theorem 4.3, it is clear that this will give us a root system Φ, with $W_\Phi = W$. So we now have a way to go from reflection group to root system and back (of course, many different root systems give us the same reflection group).

Now we are going to use root systems to study our reflection groups.

4.3 Positive and simple roots

Remember, our goal with the general definition of a root system was to create a gadget that looked like Figures 15 and 16. Those two had two more nice properties we have not achieved yet with our definition of the root system, and they were:

- We could write any root as a linear combination (i.e. a sum of vectors with real coefficients) of two distinguished vectors, α and β.
- In those linear combinations, all the coefficients had the same sign (either all nonnegative, or all nonpositive).

To formalize the coefficients having the same sign, we need to talk about *total orders* on $\mathbb{R}^n$.

Definition 4.6. *A* total order *on* $\mathbb{R}^n$ *is a relation on* $\mathbb{R}^n$*, denoted as* $<$ *satisfying*

1. *If* $\lambda < \mu$ *and* $\mu < \nu$ *then* $\lambda < \nu$*.*
2. *For each pair* λ, μ*, exactly one of* $\lambda < \mu, \lambda = \mu, \mu < \lambda$ *holds.*
3. *For all* λ, μ, ν*, if* $\lambda < \mu$*, then* $\lambda + \nu < \mu + \nu$*.*
4. *If* $\lambda < \mu$ *and* c *is a nonzero real number then* $c\lambda < c\mu$ *if* $c > 0$ *and* $c\mu < c\lambda$ *if* $c < 0$*.*

Given a total ordering on $\mathbb{R}^n$, we say λ is *positive* if $\lambda > 0$.

Exercise 4.7. *Prove that this notion of positivity has the properties that the sum of positive vectors is positive, and a positive multiple of a positive vector is positive.*

The next part in our discussion is the first part when we are going to be doing some "real math". So, instead of looking at pretty examples, and trying to intuitively feel the structures beneath them, we are going to abstract away from them, and ignore everything but the most fundamental properties of the objects we want to study. While this may seem dry compared to what we have been doing before, you should keep in mind the examples we looked at before and relate them to the coming definitions and theorems.

Definition 4.8. *Let* $\Phi \subset \mathbb{R}^n$ *be a root system, and* $<$ *some total order on* $\mathbb{R}^n$*. Define* $\Pi = \{\lambda \in \Phi | \lambda > 0\}$*. We refer to* Π *as the set of* ***positive roots*** *(or a positive system).*

Exercise 4.9. *Let* $\Pi \subseteq \Phi$ *be the set of positive roots in* Φ *with respect to some ordering* $<$*. Prove that* $-\Pi = \{\lambda \in \Phi | -\lambda \in \Pi\}$ *is the system of positive roots with respect to a total order on* $\mathbb{R}^n$*. Moreover, show that* $\Phi = -\Pi \cup \Pi$*.*

Definition 4.10. *We will call a subset* $\Delta = \{\alpha_1, \ldots, \alpha_n\} \subseteq \Phi$ *a set of* ***simple roots*** *(or a simple system) if*

1. *Any vector* λ *in* Φ *can be expressed as a sum* $\lambda = \sum_{i=1}^n c_i \alpha_i$ *(i.e. a linear combination of* $\alpha_1, \ldots, \alpha_n$*) with coefficients* c_i *all of the same sign.*

2. *The equation*

$$\sum_{i=1}^{n} c_i \alpha_i = \mathbf{0} \qquad \textit{(the zero vector)}$$

has only the trivial solution $c_1 = c_2 = \ldots = c_n$ *(this condition is usually referred to as the set of vectors* $\{\alpha_1, \ldots, \alpha_n\}$ *being linearly independent).*

At this moment, it is not clear why simple systems should exist. However, consider figures 15 and 16. In each case, take $\Delta = \{\alpha, \beta\}$, and for A_2, let $\Pi = \{\alpha, \beta, \alpha + \beta\}$, and for B_2, let $\Pi = \{\alpha, \beta, \alpha + \beta, 2\alpha + \beta\}$. Then the Δ's are simple, and the Π's are positive systems in A_2 and B_2.

Theorem 4.11. *(Theorem 1.3. in [2]) Let* Φ *be a root system in* $\mathbb{R}^n$.

(a) If Δ *is a simple system in* Φ*, then there is a unique positive system* Π *containing it.*

(b) Every positive system Π *contains a unique simple system* Δ.

Proof. (also from [2])

(a) Suppose Δ is contained in some positive system Π. Then Π must contain all roots which are nonnegative linear combinations of the vectors $\alpha \in \Delta$. The negatives of these roots obviously cannot be in Π, so Π is determined by Δ. To see that such a Π exists, we just need to find an ordering $<$ with respect to which all $\alpha \in \Delta$ are positive. It is not very hard to show that such an ordering exists.

(b) For the uniqueness part, note that if a positive system Π contains a simple system Δ, then Δ may be characterized as the set of roots in Π which cannot be written as the sum of two or more vectors in Π. This determines Δ, so if Δ is contained in Π, then Δ is unique.

To locate a simple system Δ in Π, define a subset $\Delta' \subseteq \Pi$ which is minimal (with respect to inclusion) among subsets subject to the requirement that every element of Π can be written as a nonnegative linear combination of Δ'. Now we have to show that this set is linearly independent. We will need a key geometric observation. Take a look at figures 15 and 16 notice that the angle between α and β in both cases is obtuse. This is crucial if we want all our positive roots to be nonnegative linear combinations of our simple roots. It turns out this is true in general, i.e. if we let α, β be two nonequal roots in Δ', then

$$\alpha \cdot \beta \leq 0. \tag{1}$$

We will argue by contradiction. Let α and β to be nonequal roots in Δ' with $\alpha \cdot \beta > 0$. Then

$$s_\alpha(\beta) = \beta - \frac{2\alpha \cdot \beta}{\alpha \cdot \alpha}\alpha$$

with $\frac{2\alpha\cdot\beta}{\alpha\cdot\alpha} > 0$. Since $s_\alpha(\beta) \in \Phi$, either it is a positive root (i.e. $s_\alpha(\beta) \in \Pi$), or it is a negative root, in which case $-s_\alpha(beta) \in \Pi$. We will show that it cannot be positive. Say $s_\alpha(\beta) = \sum_\gamma c_\gamma \gamma$ for $c_\gamma \geq 0$. If $c_\beta < 1$, then

$$s_\alpha(\beta) = \beta - \frac{2\alpha \cdot \beta}{\alpha \cdot \alpha}\alpha = c_\beta \beta + \sum_{\gamma \neq \beta} c_\gamma \gamma,$$

which would imply that we can write $(1 - c_\beta)\beta$ as a nonnegative linear combination of vectors in $\Delta' \setminus \{\beta\}$, so we could discard β, contradicting the minimality of Δ'. If $c_\beta \geq 1$, then we instead get

$$0 = (c_\beta - 1)\beta + \frac{2\alpha \cdot \beta}{\alpha \cdot \alpha}\alpha + \sum_{\gamma \neq \beta} c_\gamma \gamma,$$

but this is impossible, since we are forming nonnegative linear combinations of positive vectors, and at least the coefficient (on the right hand side) of α is strictly positive. We can argue similarly to show that $s_\alpha(\beta)$ cannot be negative either.

Now, if the vectors in Δ' were not linearly independent, then $\sum_{\gamma \in \Delta'} c_\gamma \gamma = \mathbf{0}$ would have a solution with not all the c_γ's equal to zero. Then we could reorganize the terms in a form

$$\sum c_\gamma \gamma = \sum c_\eta \eta$$

with all the coefficients positive, and none of the elements of Δ' appearing on both sides of the equation. Let $\sigma = \sum c_\gamma \gamma = \sum c_\eta \eta$, then, by (1),

$$0 \leq \sigma \cdot \sigma = \left(\sum c_\gamma \gamma\right) \cdot \left(\sum c_\eta \eta\right) \leq 0,$$

which would mean $\sigma = \mathbf{0}$, a contradiction. So Δ' must be a linearly independent set.

Q.E.D.

Note that (1) implies that the angle between two simple roots in any root system is obtuse (or a right angle).

Now we have established a correspondence between positive and simple systems in a reflection group. But there are many possible positive (and thus simple) systems, and a natural next question to ask is, in what way do they differ from each other? It turns out, they are not very different at all. For one thing, for any simple system Δ and for any w, the set of vectors $w\Delta = \{w(\alpha) | \alpha \in \Delta\}$ is again a simple system.

Exercise 4.12. *Can you describe the positive system containing $w\Delta$?*

The following two results are also from [2]:

Proposition 4.13. *Let Δ be a simple system, contained in the positive system Π. If $\alpha \in \Delta$, then $s_\alpha(\Pi \setminus \{\alpha\}) = \Pi \setminus \{\alpha\}$, or, in other words, s_α permutes the positive roots other than α.*

Proof. Let $\beta \in (\Pi \setminus \alpha)$. We will show $s_\alpha(\beta) \in (\Pi \setminus \alpha)$. Write $\beta = \sum_{\gamma \in \Delta} c_\gamma \gamma$ with $c_\gamma \geq 0$. Since

$$s_\alpha(\beta) = \beta - \frac{2\alpha \cdot \beta}{\alpha \cdot \alpha}\alpha,$$

only the coefficient c_α may change as a result of applying s_α. Since $\beta \neq \alpha$, there exists some γ such that $\gamma \neq \alpha$ and $c_\gamma > 0$. Therefore $s_\alpha(\beta) \in \Pi$. We only need to show that $s_\alpha(\beta) \neq \alpha$, but this is obvious, as if $s_\alpha(\beta) = \alpha$ then

$$\beta = s_\alpha(s_\alpha(\beta)) = s_\alpha(\alpha) = -\alpha,$$

which is absurd.

Q.E.D.

Theorem 4.14. *Fix any positive system Π in Φ. Then all positive systems in Φ are of the form $w\Pi$ for $w \in W$.*

Proof. Let Π and Π' be positive systems. We will proceed by induction on $r = |\Pi \cap -\Pi'|$. If $r = 0$ then there is nothing to prove. If $r > 0$, then consider the unique simple system Δ contained in Π. By Theorem 4.11 it cannot be contained in Π' as well. So pick an $\alpha \in (\Delta \setminus \Pi')$. Now Proposition 4.13 implies that $|s_\alpha \Pi \cap -\Pi'| = r - 1$, and we are done by induction.

Q.E.D.

After these results, we pretty much know what simple and positive systems look like. Let's take another look at figures 15 and 16. These correspond to dihedral groups, so we know that they are generated by two reflections, corresponding to adjacent symmetry lines. For instance, we could take the ones perpendicular to α and β (the simple roots) in both figures. In fact, this is true in full generality:

Theorem 4.15. *(Theorem 1.5. in [2]) Fix any simple system Δ in Φ. Then the reflection group W is generated by the reflections s_α for $\alpha \in \Delta$.*

Proof. Let W' denote the group generated by the simple reflections, i.e. s_α for $\alpha \in \Delta$. Quite clearly $W' \subseteq W$, so we just need to show $W \subseteq W'$. If we are able to show that any root β can be written in the form $w(\alpha)$ for some $w \in W'$ and $\alpha \in \Delta$, then we would be done by Theorem 4.3. To do that, we introduce the *height* of a positive root. Let $\beta \in \Pi$, and write

$$\beta = \sum_{i=1}^{n} c_i \alpha_i.$$

Note that there is a unique way to do this. then define

$$\mathrm{ht}(\beta) = \sum_{i=1}^{n} c_i.$$

We will proceed by induction on $\mathrm{ht}(\beta)$. The base case is easy, as if $\mathrm{ht}(\beta) = 1$, then $\beta \in \Delta$, and we may take $w = e$, the identity element. Now assume that $\mathrm{ht}(\beta) > 1$, and that any root of height less than $\mathrm{ht}(\beta)$ can be written of the form $w(\alpha)$ for some $w \in W'$ and $\alpha \in \Delta$. We have

$$0 < \beta \cdot \beta = \beta \cdot \left(\sum_{j=1}^{n} c_j \alpha_j \right) = \sum_{j=1}^{n} c_j (\beta \cdot \alpha_j).$$

So for at least one $1 \leq k \leq n$, we must have $\beta \cdot \alpha_k > 0$. Since $\beta \neq \alpha_k$, by Proposition 4.13, $s_{\alpha_k}(\beta)$ is still a positive root. We want to show that $\mathrm{ht}(s_{\alpha_k}(\beta)) < \mathrm{ht}(\beta)$ to complete the induction. Since

$$s_{\alpha_k}(\beta) = \sum_{j \neq k} c_j \alpha_j + \left(c_k - \frac{2\beta \cdot \alpha_k}{\alpha_k \cdot \alpha_k} \right) \alpha_k,$$

and $\beta \cdot \alpha_k > 0$, we are done.

Q.E.D.

5 Coxeter graphs

5.1 A presentation of W

Having Theorem 4.15 at hand, we now seek to put our reflection group to some standard form, called a *presentation* of our group W. We want to arrive at a form as in section 1.3, with generators the simple reflections, and some relations. Fix a reflection group W with a positive system $\Delta = \{\alpha_1, \ldots, \alpha_n\}$. For any $1 \leq i, j \leq n$, let $m_{i,j}$ denote the smallest positive integer for which

$$(s_{\alpha_i} s_{\alpha_j})^{m_{i,j}} = e.$$

Theorem 5.1. *The reflection group W is generated by $\{s_{\alpha_i}\}_{i=1}^{n}$ subject only to the relations*

$$(s_{\alpha_i} s_{\alpha_j})^{m_{i,j}} = e.$$

We will omit the proof, which has a similar combinatorial flavor as the proofs we were doing above. The interested reader should look up [2], sections 1.7-1.9 for details.

What Theorem 5.1 is telling us is that if we know the $m_{i,j}$'s, we can reconstruct our reflection group W. Remember, initially, the data determining W was some finite set of reflections in Euclidean space, so using the root systems, and their simple roots, we have achieved a radical reduction in the information necessary to determine W.

5.2 Coxeter graphs of familiar reflection groups

To record the information determining a finite reflection group, let's draw a graph, whose vertices correspond to simple reflections, and connect two simple reflections s_{α_i} and s_{α_j} by an edge if $m_{i,j} \geq 3$, and if in fact $m_{i,j} > 3$, then label the edge with the number $m_{i,j}$. Let's explore what sort of graphs we get if we do this for our familiar examples. For all dihedral groups, we should get a graph with two connected vertices, with possibly a label, e.g.

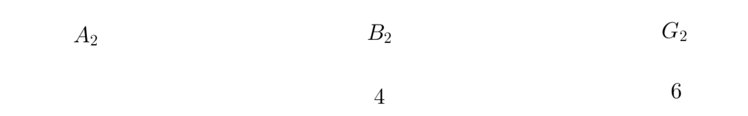

Figure 17: Coxeter graphs of A_2, B_2, G_2

We also have a good description of the reflection group A_n. Recall that it is just the group of permutations of $1, \ldots, n+1$, and we realized that it is a reflection group by making it act by permuting the coordinates of $\mathbb{R}^{n+1}$. One thing that we have not done yet is to identify a set of simple reflections in A_n. Let's look for a convenient root system for A_n. Let Φ be the set of vectors in $\mathbb{R}^{n+1}$ of the form $\mathbf{e_i} - \mathbf{e_j}$ for $i \neq j$ (recall that $\mathbf{e_i}$ is the vector with i-th coordinate equal to 1 and all other coordinates zero).

Exercise 5.2. *Prove that Φ is a root system, with $W_\Phi = A_n$.*

Now choose $\Pi = \{\mathbf{e_i} - \mathbf{e_j} \mid i < j\}$ and $\Delta = \{\mathbf{e_i} - \mathbf{e_{i+1}} \mid i = 1, \ldots n\}$.

Exercise 5.3. *Prove that Π is a positive, and Δ is a simple system for Φ.*

So, a set of simple reflections for A_n is *adjacent transpositions.*

Exercise 5.4. *Let $s_i \in A_n$ denote the transposition swapping i and $i+1$. Show that, for $i < j$,*

$$m_{i,j} = \begin{cases} 3 & \textit{if } j = i+1, \\ 2 & \textit{otherwise} \end{cases}$$

So the Coxeter graph of A_n is the "line" with n vertices.

Figure 18: The Coxeter graph of A_n

Note that the Coxeter graph of a finite reflection group need not be connected. For instance, consider the symmetry group of the rhombus below:

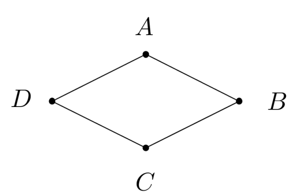

Figure 19: A rhombus

We see that we can reflect across the AC and BD symmetry lines, swapping A and C, or B and D respectively, but these symmetries are independent of each other. This group is called $A_1 \times A_1$, since it is built up of two non-interacting A_1's, and its Coxeter graph is, therefore

$A_1 \times A_1$

Figure 20: The Coxeter graph of $A_1 \times A_1$

Incidentally, the same group is (up two musical transposition[5]) the symmetry group that is allowed to act in the twelve-tone composition technique devised by Arnold Schoenberg. Note the $A_1 \times A_1$ symmetry in the following figure (P denotes the *prime* form, the original tune, R denotes it in *retrograde*, and I denotes its *inversion*):

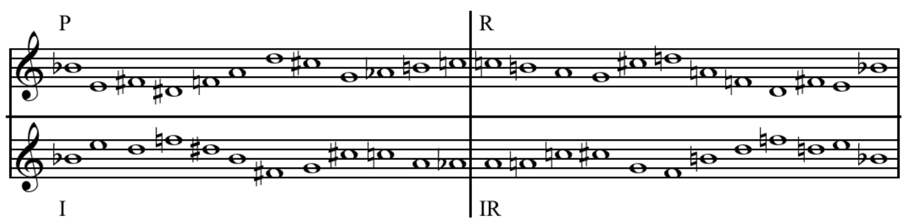

Figure 21: Two bars from Schoenberg's Variations for Orchestra (op. 31)

5.3 Crystallographic reflection groups and Dynkin diagrams

For a while we are going to restrict our attention to finite *crystallographic* reflection groups. It turns out that "most" finite reflection groups are crystallographic, but we'll discuss this issue later. A reflection group is said to be crystallographic, if it is capable of stabilizing a lattice in the $\mathbb{R}^n$ it is acting on. The details are not very important for us, but the name comes from crystallography, where the possible crystal patterns are governed by the available crystallographic reflection groups.

Figure 22: Pyrite cubes

[5]Not to be confused with transpositions in the symmetric group, these are essentially translations by an integer.

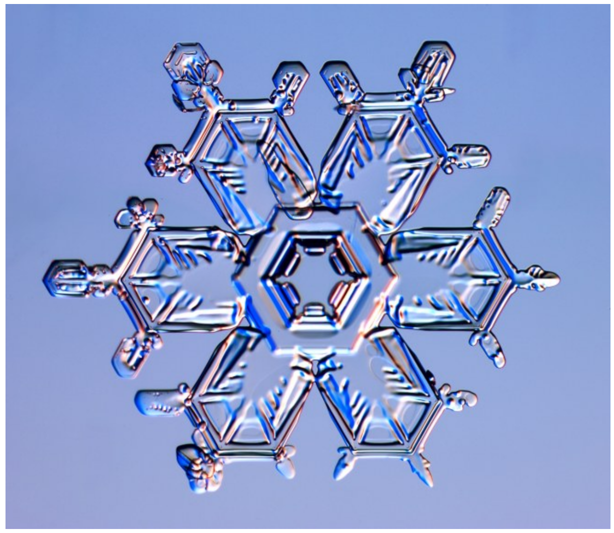

Figure 23: A snowflake

Exercise 5.5. *What is the symmetry group of the snowflake displayed in figure 23?*

It turns out that a necessary and sufficient condition for W_Φ of some root system to stabilize a lattice is that any root can be written as an *integral* linear combination of simple roots (Notice that this is the case for A_2, B_2, G_2). In this case, the *root lattice* (i.e. arbitrary integral linear combinations of the simple roots) is stabilized by W_Φ. Recall that if α and β are roots with θ being the angle between them, then

$$s_\alpha(\beta) = \beta - \frac{2\beta \cdot \alpha}{\alpha \cdot \alpha}\alpha.$$

So the numbers $\frac{2\beta\cdot\alpha}{\alpha\cdot\alpha}$ better be integers for all pairs of roots (it is actually sufficient to require it to be integers for pairs of simple roots). We also know that

$$\frac{2\beta \cdot \alpha}{\alpha \cdot \alpha} = 2\frac{|\beta|}{|\alpha|}\cos\theta,$$

so if $\frac{2\beta\cdot\alpha}{\alpha\cdot\alpha}$ and $\frac{2\alpha\cdot\beta}{\beta\cdot\beta}$ are both integers, then so is

$$\left(\frac{2\beta \cdot \alpha}{\alpha \cdot \alpha}\right)\left(\frac{2\alpha \cdot \beta}{\beta \cdot \beta}\right) = 4\cos^2\theta. \tag{2}$$

And this limits our options severely, as $0 \leq \cos^2\theta \leq 1$. So if we assume that $\beta \neq \pm\alpha$ then the value of (2) can only be $0, 1, 2$, or 3.

Exercise 5.6. *Show that this limits the possible labels $m_{i,j}$ on the edges of the Coxeter graph to* $2, 3, 4, 6$.

Also it is important to note that not only $4\cos^2\theta$, but $2\frac{|\beta|}{|\alpha|}\cos\theta$ must be integers too. So now we have to fix the relative lengths of simple roots in a compatible way, e.g. the way we did it for B_2 in Figure 16. To be more precise, whenever we have $m_{i,j} = 3, 4, 6$, the ratio of squared lengths of the long root over the short root must be $1, 2, 3$, respectively. Inspired by this, we will replace the 4-s and 6-s of our Coxeter graph by double/triple edges, and we will orient multiple edges to account for the relative lengths, with long $>$ short, to obtain the *Dynkin diagram* of a root system of a crystallographic reflection group.

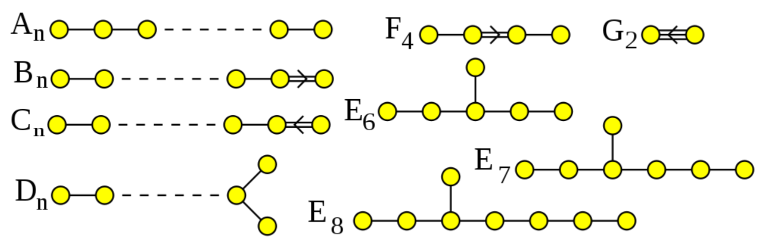

Figure 24: Dynkin diagrams

We usually insist that $n \geq 1$ for B_n, $n \geq 2$ for C_n, and $n \geq 4$ for D_n to avoid repetitions. Note that since we fixed the relative lengths of the roots, a Dynkin diagram is associated to a root system, not quite the reflection group per se. For instance, the groups B_n and C_n are the same, they correspond to the symmetry groups of the n-cube. The root systems B_n and C_n are dual to each other, much like the n-cube and the n-cross-polytope.

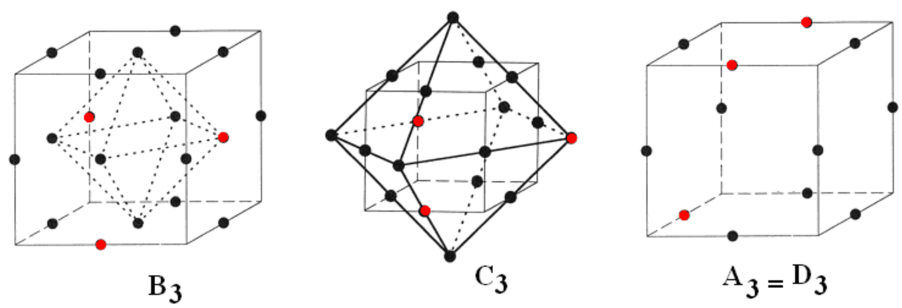

Figure 25: The root systems $B_3, C_3, A_3 = D_3$ with a possible choice of simple roots is displayed in red. You should imagine that the origin is at the center of the cubes.

The amazing thing is that the Dynkin diagrams displayed in figure 24 are the *only* connected Dynkin diagrams that produce finite reflection groups.

5.4 Kostant's find the highest root game

I have learnt this (among countless other things) from Allen Knutson. Let's play a (solitaire) game on a simple graph G, i.e. assume that we have no double nor triple edges. In the language of root systems, this means that all the roots are of the same length, so we might as well assume them to be unit length. The technical term for this is a *simply laced* root system. Also, assume that our graph is connected, if it isn't, we can play this game on each connected component of it, so this isn't a big deal. We will put numbers on the vertices. Start by assigning a one to an arbitrary vertex, and zero to all others. Now say that a vertex is *unhappy* if its value is less than half the sum of its neighbors', *happy* if it is equal to that, and *manic* if it is more. To play the game, find an unhappy vertex, and replace its value by the sum of the values of its neighbors, minus its old value. Repeat until there are no unhappy vertices. We say that the graph is of *finite type* if the game ends[6].

For example, let's play this game on the A_n diagram. If we start with a 1 on the left end vertex, its neighbor will be unhappy. There is always only one move we can make, which "propagates" the 1 to the right, and in the end, all vertices will have a 1 on them, with the two terminal vertices being manic, and everyone else happy.

Exercise 5.7. *Show that (assuming the graph is connected) the choice of the starting vertex does not matter.*

Exercise 5.8. *Show that if H is a subgraph of G and H is not of finite type, then G is not of finite type.*

Let's use our knowledge of root systems to interpret this game. One may think of picking the initial 1 as picking a simple root in the root system. Moreover, the moves of the game correspond to replacing a root by another root

[6]There is some sleight of hand here. A priori, it seems possible that the game may or may not end depending on the sequence of moves we make, but this does not happen, so we will politely ignore this subtlety.

obtained from it by a simple reflection! Let $\beta \in \Phi$, write $\beta = \sum_{i=1}^n c_i\alpha_i$, and say that for some $k \in 1, \ldots, n$, we have

$$c_k < \frac{1}{2}\left(c_{l_1} + \ldots + c_{l_m}\right), \tag{3}$$

where $\alpha_{l_1}, \ldots, \alpha_{l_m}$ are the simple roots which are joined by an edge to α_k. Then

$$s_{\alpha_k}(\beta) = \sum_{i=1}^n c_i\alpha_i - 2\left(\sum_{i=1}^n c_i\alpha_k \cdot \alpha_i\right)\alpha_k = \sum_{i\neq k} c_i\alpha_i + (-c_k + (c_{l_1} + \ldots + c_{l_m}))\alpha_k$$

and as a consequence of the assumption (3), $c_k < c_{l_1} + \ldots + c_{l_m} - c_k$, so $\mathrm{ht}(s_{\alpha_k}(\beta)) > \mathrm{ht}(\beta)$. And, since there are only finitely many roots, the height cannot increase indefinitely, so the game will terminate.

So it is clear that the game will terminate on a simply laced Dynkin diagram, so now we are going to look for graphs where the game terminates. Paradoxically, a really good way to do this is to find some graphs that are comparatively simple for which the game *does not* terminate, then use these and exercise 5.8 to eliminate the potential finite type graphs.

Theorem 5.9. *Let G be a graph of type A_n, for $n \geq 2$,D_n, E_6, E_7, or E_8. Play the game on the graph until the end. Bring in a new vertex v_0 and attach it to the manic vertices, and call the resulting graph $\widetilde{A_n}, \widetilde{D_n}, \widetilde{E_6}, \widetilde{E_7}$, or $\widetilde{E_8}$ (these are called affine diagrams). None of these new graphs is of finite type*

Proof. First play the game without touching v_0 until the end. Now bring in v_0 with value 1 on it (this is not a legal move of the game on the affine diagram, we are cheating here), and call this configuration R. Check all of the cases, and notice that in each of them, everyone is happy, so if we are playing on the affine diagram, all our moves will commute with adding or subtracting multiples of R, since in the state R, all the vertices have values equal to exactly half of the sum of their neighbors'.

Next, play the game on the affine diagram the following way: first play without touching v_0 until there are no moves left, then fire v_0. Notice that in each case, its value switches from 0 to 2, so we may subtract R, and continue playing as if we just started from having 1 on v_0 and 0's everywhere else. Using exercise 5.7, we may pretend we have a 1 on some other vertex, and play the game as if we just started. We end up with arbitrarily high multiples of R, so the affine diagram cannot be of finite type.

Q.E.D.

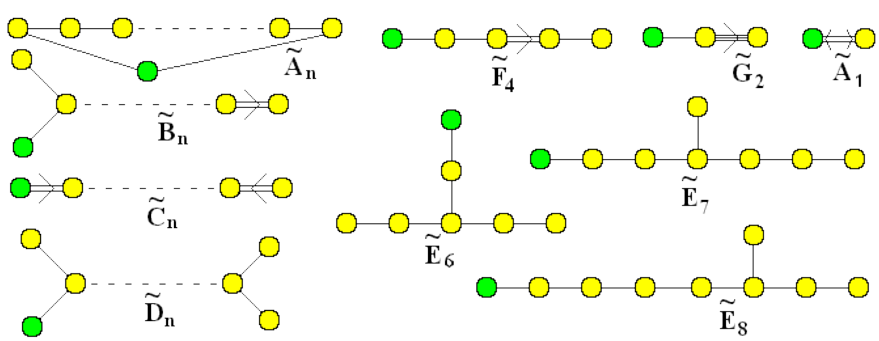

Figure 26: Affine Dynkin diagrams

Theorem 5.10. *The type A, D, E diagrams are the only connected simply laced diagrams of finite type.*

Proof. We will be using exercise 5.8 extensively. Let G be a simply laced Dynkin diagram which is not A_n, D_n, E_6, E_7, or E_8. Note that $\widetilde{A_n}$ can not be a subgraph of G, so G is a tree. Since G is not A_n, it has at least one branch point. Let's investigate these. Since $\widetilde{D_4}$ is not a subdiagram of G, G cannot have any 4-valent vertices, so all branch points are at most trivalent. Moreover, since $\widetilde{D_3}$ is not a subgraph, G has exactly one branch point. Now let's see how far we can stretch the arms out from the branch point. Since $\widetilde{E_6}$ is not a subgraph, at least one branch is at most length 1. Since $\widetilde{E_7}$ is not a subgraph, the length of the second shortest branch is at most 2. And finally, since $\widetilde{E_8}$ is not a subgraph, the length of the third branch is at most 4, but then our graph is one of E_8, E_7, E_6, or D_n for $n = 4, \ldots, 8$, a contradiction.

Q.E.D.

One may modify the rules of the game so that it can be played on a non-simply laced Dynkin diagram, and complete a similar classification for those to reduce the possible Dynkin diagrams to the ones displayed in figure 24.

5.5 Non-crystallographic finite reflection groups

It turns out we haven't missed out on too many finite reflection groups when restricting to the crystallographic case (except for the infinitely many dihedral groups, of course), since the list of all reflection groups is barely larger:

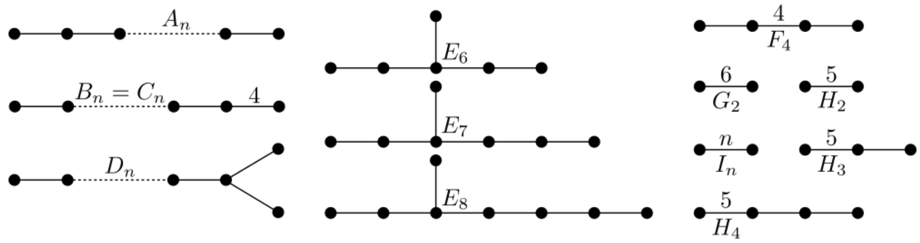

Figure 27: Coxeter graphs of all finite reflection groups

Notice that (as we discussed before) the actual reflection group is insensitive to the relative lengths of roots in the root system, so the Coxeter graphs of types B_n and C_n are the same.

Figure 27 tells us that any finite reflection group with a connected Coxeter graph is one of the types $A - I$, there is still the question whether these groups actually exist. We have met the dihedral groups, the groups of type $A_n, B/C_n$, as symmetry groups of simplices and hypercubes. We have seen the groups H_3 and H_4 before (albeit in different context), H_3 is the symmetry group of the regular dodecahedron, and its dual, the icosahedron, and H_4 is the symmetry group of the 120-cell, and its dual, the 600-cell. The group F_4 is the symmetry group of the 24-cell. So the only ones remaining are the types D and E groups. They are also symmetry groups of polytopes, not regular ones, but so called *uniform* polytopes. If you think that's not as nice as being a symmetry group of a convex regular polytope, take a look at the following figure, which is the projection of the 4_{21}-polytope to the plane, whose vertices correspond to the roots of the root system E_8.

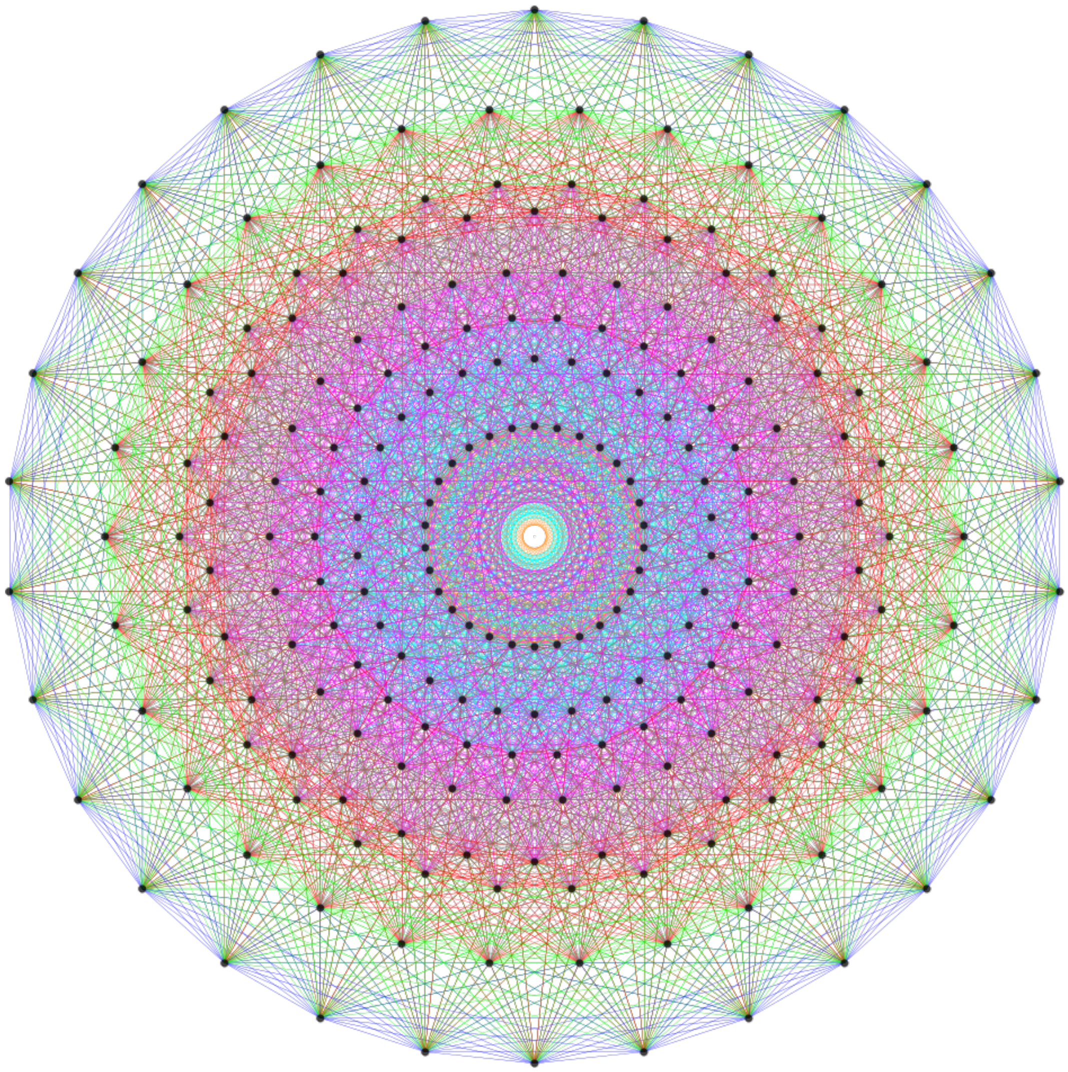

Figure 28: The 4_{21}-polytope's projection to the plane.

6 Appendix: Reflection groups in Matrix groups

In this section we will explore how reflection groups arise when thinking about matrix groups. We will call an n by n matrix A *invertible* if there is a matrix B such that $AB = BA = I$, the n by n identity matrix. If this is the case, we usually write A^{-1} for the matrix B.

Exercise 6.1. *Show that if A and B are invertible, then so is AB.*

We will refer to the set of invertible n by n matrices as GL_n. We want to find symmetries of GL_n. One natural thing to look at when we have a bunch of matrices at hand is row operations. They come in three types. If we have a matrix A, then we can

1. Multiply a row of A by a nonzero number.
2. Add a multiple of a row to another row.
3. Exchange two rows.

Remark 6.2. *Notice that each of these can be "undone" by performing another row operation.*

Exercise 6.3. *Find matrices E such that EA corresponds to each of the types of row operations.*

Exercise 6.4. *Interpret Remark 6.2 in terms of the matrices of Exercise 6.3*

So if we perform row operations on an invertible matrix, it stays invertible. So row operations are symmetries of GL_n! Now one thing one likes to do when dealing with large geometric objects (like GL_n for instance) is to use some symmetries to simplify it, and hopefully make it more comprehensible. For instance, we might naively declare two matrices to be "the same" when there is a sequence of row operations leading from one to the other. Unfortunately this is a little too much, as any invertible matrix can be reduced to the identity matrix by a sequence of row operations. What happened here is that we used too much of the symmetry of our big object, and squashed it into a single point, which does not really enhance our understanding. In fact, GL_n is a group with the group operation corresponding to matrix multiplication. And any group acts as symmetries of itself (think about what happens if you multiply any $g \in G$ to every element of your group). So we need to be a little more clever, as general row operations are too powerful (by Exercise 6.4 sequences of row operations correspond to multiplication by an arbitrary matrix in GL_n on the left).

Let B be the set of upper triangular matrices in GL_n, i.e. matrices that look like this:

$$\begin{pmatrix} * & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \\ 0 & 0 & 0 & * \end{pmatrix}$$

and are also invertible. Actually, upper triangular matrices are invertible if and only if none of the diagonal entries are zeros. We take time to point out the obvious thing again, namely that if we consider two matrices $A, B \in GL_n$ to be equivalent, or, to represent the same point if there is some $g \in GL_n$ such that $gA = B$, then all the matrices are equivalent. Now let's say we want to declare two matrices $A, B \in GL_n$ to be equivalent if there are two matrices $b_1, b_2 \in B$ such that

$$b_1 A b_2 = B.$$

What can we make of this relation? We want to look for matrices that represent the (possibly many) equivalence classes. We may try to proceed in a way we did before, i.e. try to reduce our matrix to the identity matrix I and see if there is a problem.

Exercise 6.5. *Try to describe the row operations that correspond to multiplying a matrix A by $b \in B$ on the left. Then try to describe the* column *operations that correspond to multiplying a matrix A by $b \in B$ on the* right.

Now let's restrict our attention to 3 by 3 matrices for a moment. Let's try to see how we can simplify the matrix

$$\begin{pmatrix} 6 & 15 & 12 \\ 4 & 8 & 6 \\ 1 & 2 & 1 \end{pmatrix}$$

Using upward row and rightward column operations. Look at the 1 in the southwest corner. We can use upward row operations to "clear out" all the entries of the first column above it, leading to

$$\begin{pmatrix} 0 & 3 & 6 \\ 0 & 0 & 2 \\ 1 & 2 & 1 \end{pmatrix}$$

and now rightward column operations to clear out the last row

$$\begin{pmatrix} 0 & 3 & 6 \\ 0 & 0 & 2 \\ 1 & 0 & 0 \end{pmatrix}.$$

Now let's move on to the second column. We can certainly make the 3 to be a 1 using row operations

$$\begin{pmatrix} 0 & 1 & 2 \\ 0 & 0 & 2 \\ 1 & 0 & 0 \end{pmatrix}$$

then clear out the first row using rightward column operations

$$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 1 & 0 & 0 \end{pmatrix}$$

and finally make the remaining 2 into a 1 by multiplying the second row by $\frac{1}{2}$, and what we are left with is:

$$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix}.$$

A matrix with a 1 in each row and each column. A matrix like that is called a *permutation matrix*. In fact, for any matrix in GL_n, you can use upward row and rightward column operations to make it into a permutation matrix.

Exercise 6.6. *Explain the choice of name for these matrices.*

7 Project idea: Non-simply laced classification

In class, we have completed the classification of simply laced (i.e. type A, D, E) Dynkin diagrams, and hence, of crystallographic reflection groups. This project is about finishing the classification of *all* Dynkin diagrams. You will

- Extend the rules of Kostant's find the highest root game to diagrams with oriented double/triple edges. Caution: the rules for vertex happiness will be asymmetrical now, roots will care about the length of their neighbors, not just the values those neighbors have.
- Similarly to the simply laced case, find a list of "small" diagrams that are forbidden, i.e. are not of finite type.
- Use the list to eliminate all possibilities other than $A_n, B_n, C_n, D_n, E_6, E_7, E_8, F_4$ and G_2.

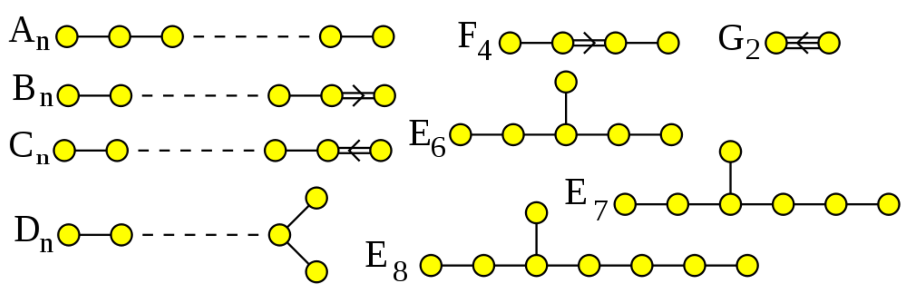

Figure 29: Dynkin diagrams

8 Project idea: Length function and the Bruhat order

The length function and the Bruhat order are two closely related combinatorial concepts on a reflection group W. Recall that any element $w \in W$ can be expressed (possibly in many ways) as a product of simple reflections

$$w = s_1 s_2 \dots s_k.$$

The *length* $l(w)$ of an element w of a reflection group is defined to be the smallest k among all of the possible expressions.

A partial order on a set S is a relation denoted $\leq$ satisfying:

$$\begin{gathered} u \leq u \\ u \leq v \text{ and } v \leq w \Rightarrow u \leq w \\ u \leq v \text{ and } v \leq u \Rightarrow u = v. \end{gathered}$$

Note that two elements of the set may be incomparable (i.e. it is possible that neither of $u \leq v$ and $v \leq u$ holds.

On W, let's denote by T the set of reflections in W (you may think of it like the reflections across any of the roots in a root system for W). If u, v are two elements of W, then by $u \xrightarrow{t} v$ we mean that $l(u) < l(v)$ and $u^{-1}w = t \in T$, and by $u \to v$, we mean $u \xrightarrow{t} v$ for some $t \in T$. Then the *Bruhat order* on W is defined to be the partial order where

$$u \leq w \Leftrightarrow u = u_0 \to u_1 \to \dots \to u_{k-1} \to u_k = w.$$

There are *lots* of easy and interesting theorems one can prove about the length function and the Bruhat order, for example:

- The length $l(w)$ of an element is also equal to the number of positive roots moved to a negative root by w.
- For type A_n, the length $l(w) =$ inversions of the permutation $w = |\{(i,j) : i < j, w(i) > w(j)\}|$.
- In the Bruhat order $v \leq w \Leftrightarrow$ in any word $s_1 s_2 \dots s_k = w$ of minimal length for w, we can find a subword of minimal length for v.
- In the Bruhat order, $v \leq w$ implies that there is a chain

  $$v = u_0 < u_1 < \dots < u_k = w$$

  with $l(u_i) + 1 = l(u_{i+1})$.
- In any finite reflection group, there is a unique longest element.

In this project, you will familiarize yourselves with the length function and the Bruhat order, and present some of the interesting theorems you choose.

9 Project idea: Permutohedra

We realized reflection groups as symmetry groups of regular polytopes, like the n-simplex or the n-cube. That is, a given reflection group was the *complete group of symmetries* of these polytopes. There is another family of polytopes, called *permutohedra*, whose vertices correspond to elements of reflection groups. They are also pretty:

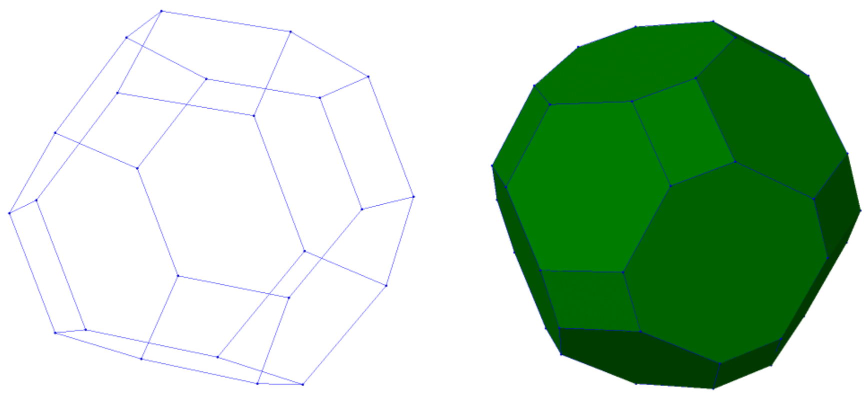

This project will explore some of the properties of permutohedra, for example:

- How they can tesselate space:

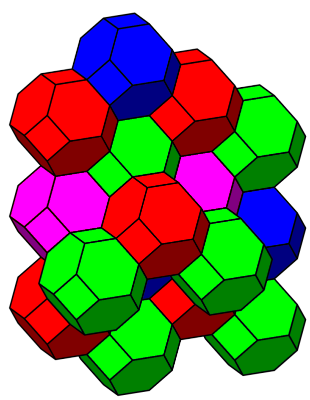

- How they can be built from line segments.
- We might look at several results of the paper "Permutohedra, associahedra and beyond" by Postnikov [4], which explores several interesting formulas involving the volumes of these polytopes, for example:

$$\mathrm{Vol} P_W(x) = \frac{1}{r!} \sum_{w \in W} \frac{(\lambda, w(x))^r}{(\lambda, w(\alpha_1)) \cdots (\lambda, w(\alpha_r))},$$

  where $\alpha_1, \ldots, \alpha_r$ are the simple roots of the reflection group W.

10 Project idea: Schläfli symbols and Regular polytopes

The Schläfli symbol is a notation describing regular polytopes (and possibly tesselations of space) in any dimension. In this project, you will explore how the Schläfli symbol describes:

- The *facets*, which are the maximal-dimensional faces of your polytope.
- The *vertex figures*, which are the polytopes you get by taking a cross-section of your polytope close to a vertex.
- The Coxeter graph of the symmetry group of your polytope.
- How certain Schläfli symbols determine tesselations instead of regular polytopes.

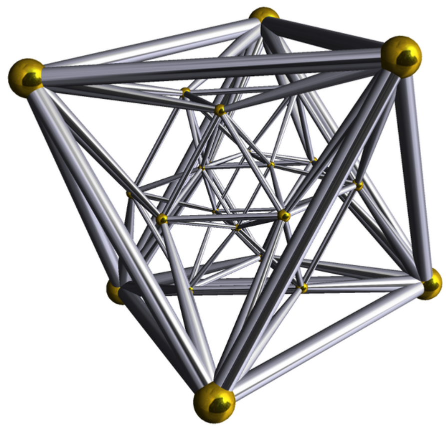

Figure 30: The 24-cell has Schläfli symbol $\{3, 4, 3\}$

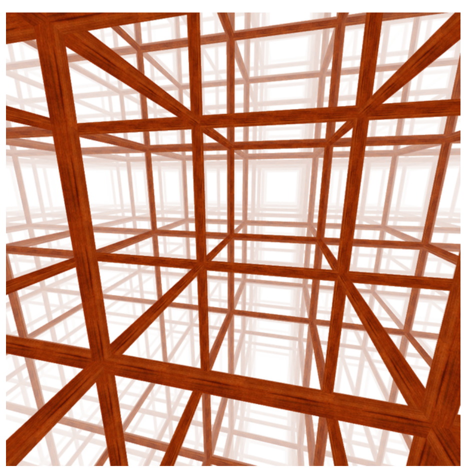

Figure 31: The *cubic honeycomb*, a tesselation of $\mathbb{R}^3$, has Schläfli symbol $\{4, 3, 4\}$

11 Project idea: Reflection Groups in real life

During the course, we have seen many examples of reflection groups occurring outside of mathematics. In this open-ended project, you will look everywhere in nature and human art for anything who has symmetries that are reflection groups, recognize the Coxeter diagrams of the reflection groups present, and possibly use some of the other things we learnt during the semester.

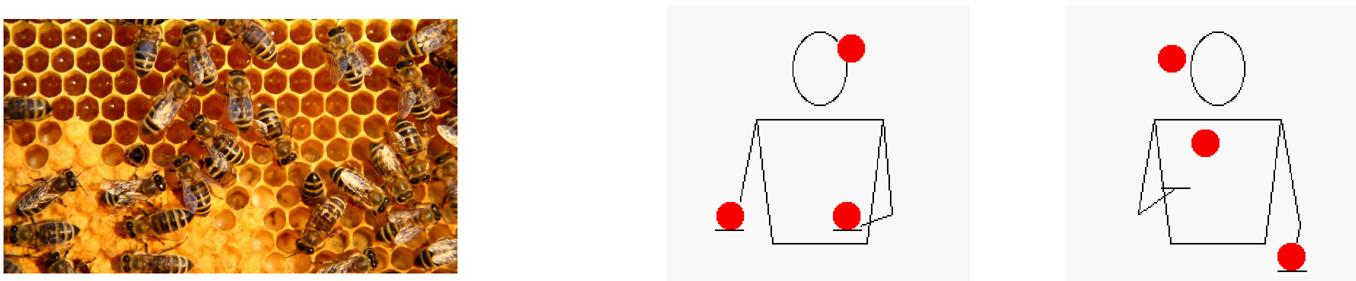

Figure 32: Reflection groups in real life

12 "Solutions" to the exercises

I put solutions to quotes since they are not full solutions, more like strong hints.

Exercise 1.1 There are $3! = 6$ of them.

Exercise 1.2 We get 123.

Exercise 1.3 $\sigma\pi = 312$. Note that it is not equal to $\pi\sigma$.

Exercise 1.4 No, for example, $(\pi\sigma)^2 = 231$.

Exercise 1.5 The triangle has 6 symmetries. The n-gon has $2n$.

Exercise 1.6 We get rotations of the triangle to opposite directions.

Exercise 1.7 6.

Exercise 1.8 Yes.

Exercise 2.1 $\mathbf{v} - \mathbf{w}$.

Exercise 2.2 $c_1\mathbf{v} - c_2\mathbf{w}$.

Exercise 2.4 Just write $\mathbf{q} = c_1\mathbf{v} + c_2\mathbf{w}$ and expand both sides.

Exercise 2.5 We get $\mathbf{q}$ back.

Exercise 2.6 $(\sin\theta, -\cos\theta)$.

Exercise 2.8 $s_l(s_e(\mathbf{p})) = ((\cos -2\theta)p_1 - (\sin -2\theta)p_2, (\sin -2\theta)p_1 + (\cos -2\theta)p_2)$, i.e. $s_l(s_e(\mathbf{p}))$ is $\mathbf{p}$ rotated clockwise by an angle 2θ.

Exercise 2.9 If θ is not a rational multiple of 2π, then there are infinitely many. Otherwise, write $\theta = \frac{2m\pi}{n}$, where m and n are relatively prime. Then there are $2n$ such transformations. Does this $2n$ remind you to something?

Exercise 2.10 $\mathbf{e_2} - \mathbf{e_3}$.

Exercise 2.11 Think of a sorting algorithm, we want to sort 1234 as $\sigma(1)\sigma(2)\sigma(3)\sigma(4)$. So what should the first transposition be? (Hint: where should $\sigma(1)$ end up eventually?)

Exercise 2.12 The regular triangle whose vertices are $(1,0,0), (0,1,0)$, and $(0,0,1)$.

Exercise 2.13 The regular tetrahedron whose vertices are $(1,0,0,0), (0,1,0,0), (0,0,1,0)$ and $(0,0,0,1)$.

Exercise 2.14 Take a and b as generators with relations $a^2 = b^2 = (ab)^6 = e$.

Exercise 3.1 The relations are $a^2 = b^2 = (ab)^n = e$.

Exercise 3.2 There are at most 120 symmetries, since putting one of the triangles in a place determines the location of all the others, the only question is, can we put a distinguished one to any of the 120 places? Certainly we can put it anywhere on the same pentagonal face, since that just involves moving our triangle to any other using the symmetries of the pentagon. We can also move from one face to any neighboring one using the edges of the dodecahedron. So we can move our original triangle to any other, so there are in fact 120 symmetries of the dodecahedron.

Exercise 3.3 The dodecahedron.

Exercise 3.4 From left to right, first row: tetrahedron, cube, tetrahedron. Second row: octahedron, dodecahedron, tetrahedron.

Exercise 3.5 It is the hypercube, the 8-cell.

Exercise 3.6 They are the triangle, tetrahedron and 5-cell.

Exercise 3.7 They are the square, the octahedron and the 16-cell.

Exercise 3.8 You get the n-simplex. It is self-dual.

Exercise 3.9 They are dual to each other.

Exercise 3.10 One can dualize the polytope, apply a symmetry of the dual, then dualize again. This is clearly a symmetry of the polytope. To see that all symmetries arise this way, we can do the same thing the opposite direction.

Exercise 4.1 One thing you can try is to see which rows correspond to elements which are order 2, i.e. if you apply them twice, you get the identity.

Exercise 4.5 Let S_Φ denote the permutation group on the roots. Since $w\Phi = w \quad \forall w \in W_\Phi$, we may define a group homomorphism $\iota : W_\Phi \to S_\Phi$ by sending w to the element of S_Φ which permutes the roots the same way as w. In order to show W_Φ is finite, we will show this homomorphism has trivial kernel. Note that any reflection $s_\alpha \in W_\Phi$, and hence every $w \in W_\Phi$ fixes the orthogonal complement (in $\mathbb{R}^n$) of the subspace spanned by the vectors in Φ, therefore only the identity element can fix all elements of Φ.

Exercise 4.7 Let $\lambda, \mu > 0$, then

$$\begin{aligned} \lambda &> 0 \\ \lambda + \mu &> \mu \\ \lambda + \mu &> 0 \end{aligned}$$

And similarly for the other property.

Exercise 4.9 Since roots come in pairs $\{-\alpha, \alpha\}$, and exactly one of $-\alpha, \alpha$ is positive, this is clear.

Exercise 4.12 It is $w\Pi$.

Exercise 5.2 Notice that $s_{(\mathbf{e_i}-\mathbf{e_j})}$ swaps the i-th and j-th coordinate while leaving the others fixed, so it is a transposition, moreover, we can form any transposition this way. Showing that Φ is invariant under the action of A_n is not hard, and it is obvious that $\Phi \cap L_{(\mathbf{e_i}-\mathbf{e_j})} = \{(\mathbf{e_i} - \mathbf{e_j}), (\mathbf{e_j} - \mathbf{e_i})\}$.

Exercise 5.3 It is not hard to see that Δ is a simple system and then, using Theorem 4.11, it is clear that Π is a positive system.

Exercise 5.4 The fact that s_i and s_j commute when $i+1 < j$ is clear. For the other case, you can just check this directly using the permutations, or notice that the transformation $s_i s_{i+1}$ fixes $\mathbf{e_1}, \dots, \mathbf{e_{i-1}}, \mathbf{e_{i+3}}, \mathbf{e_{n+1}}$, and it rotates the $\mathbf{e_i}, \mathbf{e_{i+1}}, \mathbf{e_{i+2}}$-space by the angle $120°$ around the line through $\mathbf{e_i} + \mathbf{e_{i+1}} + \mathbf{e_{i+2}}$.

Exercise 5.5 It is G_2 (at least it's close enough).

Exercise 5.6 It limits the possible pairwise angles between simple roots to be one of $90°, 60°, 45°, 30°$, which makes the reflection group generated by the two simple reflections in question to be isomorphic to $A_1 \times A_1, A_2, B_2$, or G_2, respectively.

Exercise 5.7 Since the graph is connected, there is a path between any two vertices. We can start the game by propagating the 1's through the path, arriving at a configuration which can result from both starting positions.

Exercise 5.8 We can just keep playing the game on H indefinitely, effectively ignoring other parts of G.

Exercise 6.1 You are looking for $(AB)^{-1}$. Try $B^{-1}A^{-1}$.

Exercise 6.3 This should give you an idea:

$$E_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \alpha & 0 \\ 0 & 0 & 1 \end{pmatrix} \qquad E_2 = \begin{pmatrix} 1 & 0 & \alpha \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \qquad E_3 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}$$

Exercise 6.2 It means that these matrices are invertible.

Exercise 6.5 Multiplying by b on the left corresponds to "upward" row operations, i.e. we are only allowed to multiply rows by nonzero scalars and add multiples of rows to earlier rows. We also can't swap rows. Multiplying b on the right correspond analogously to "rightward" column operations.

Exercise 6.6 They correspond to permutations. In fact, a good way to see this is to multiply these matrices to n by 1 vectors, and notice that we end up with the representation of S_n that we used to make it clear it was a reflection group.

References

[1] Humphreys, James E., *Introduction to Lie Algebras and Representation Theory*, GTM 9, Springer 1972.

[2] Humphreys, James E., *Reflection Groups and Coxeter Groups*, Cambridge studies in advanced mathematics 29, Cambridge University Press 1990.

[3] Björner, Anders, Brenti, Francesco, *Combinatorics of Coxeter Groups*, GTM 231, Springer 2005

[4] Postnikov, A., *Permutohedra, associahedra, and beyond.* Int. Math. Res. Not. IMRN 2009, no. 6, 10261106. 05E30