


Coxeter and Dynkin Diagrams

John Baez
Department of Mathematics
University of California, Riverside


based on “week62–week65” and “week182” of
This Week's Finds



I want to explain about a fascinating subject of huge importance in both mathematics and physics: Coxeter and Dynkin diagrams. Not to keep you in suspense, I've already shown you the Dynkin diagrams on the front page of these notes. Coxeter diagrams are a bit different, but similar.

I'll start by explaining how Coxeter diagrams can be used to classify "finite reflection groups": that is, finite groups of linear transformations of \mathbb{R}^n generated by reflections. Then I'll explain how Dynkin diagrams classify lattices in \mathbb{R}^n having finite reflection groups as symmetries. Next, I'll explain how we use Dynkin diagrams to classify compact simple Lie groups and their Lie algebras—and what those things are! Finally, I'll say how the "simply-laced" Dynkin diagrams—the ones without arrows on them—can be used to classify integral lattices with a basis of vectors v with $\|v\|^2 = 2$, quivers with tame representation theory, and finite subgroups of the group of rotations of 3-dimensional Euclidean space.

In short, Coxeter and Dynkin diagrams show up all over the place when you start trying to classify beautiful and symmetrical things. In fact, we'll see that the Platonic solids:



are connected to Coxeter and Dynkin diagrams in *two separate ways!*

Coxeter diagrams and finite reflection groups

Okay, so what the heck is a Coxeter diagram? We get these when we try to classify "finite reflection groups." Say we are in n -dimensional Euclidean space. Then given any nonzero vector v , there is a reflection that takes v to $-v$ and doesn't do anything to the vectors orthogonal to v . Let's call this a **reflection through v** . A finite reflection group is a finite group of transformations of Euclidean space such that every element is a product of such reflections. For example, the group of symmetries of a regular n -gon is a finite reflection group. Showing this is a useful exercise if you don't see it right off the bat.

Note that if we do two reflections, we get a rotation. In particular, suppose we have vectors v and w at an angle θ from each other, and let r and s be the reflections through v and w , respectively. Then rs is a rotation by the angle 2θ . Draw a picture and check it! This means that if $\theta = \pi/m$, then $(rs)^m$ is a rotation by the angle 2π , which is the same as no rotation at all, so $(rs)^m = 1$. On the other hand, if θ is not a rational number times π , we never have $(rs)^n = 1$, so r and s can not both be in some *finite* reflection group.

So, if r and s are two distinct reflections in a finite reflection group we must have

$$r^2 = s^2 = 1$$

since doing a reflection twice gets you back where you started, but also

$$(rs)^m = 1$$


for some $m = 2, 3, \dots$

Using this, we can see that any finite reflection group has a presentations described by a **Coxeter diagram**. The idea is that the group is generated by reflections through vectors that are at angles of π/m from each other for various choices of $m = 2, 3, \dots$. To keep track of this, we draw a dot for each one of these vectors. Then, suppose two of the vectors are at an angle π/m from each other. If $m = 2$, the reflections must commute, and we don't bother drawing a line between the two dots. Otherwise we draw a line between them and label it with the number m .

Conversely, if someone hands us a Coxeter diagram we get a group called a **Coxeter group** with

- one generator r for each dot,
- one relation $r^2 = 1$ for each generator,
- one relation $(rs)^2 = 1$ for each pair of generators with no line connecting their dots,
- one relation $(rs)^m = 1$ for each pair of generators with a line labeled by the number m connecting their dots.


For example, the Coxeter diagram



gives us the Coxeter group with presentation

$$\begin{aligned} r^2 = s^2 = t^2 = u^2 = 1 \\ (rs)^3 = (st)^7 = (ru)^4 = (ut)^3 = 1 \\ (rt)^2 = (tu)^3 = 1. \end{aligned}$$

However, in this game a lot of edges wind up being labeled with the number 3. People usually leave out the label when this happens. Then we draw the above diagram like this:




Now for some big theorems. First, it turns out that if a Coxeter group is *finite*, it's a finite reflection group! Not every diagram yields a finite group. But we can classify all possible Coxeter

diagrams giving finite groups! They have names, and they are famous: they're like the rock stars of finite group theory.

First there is A_n , which has n dots in a row, like this:




For example, the group of symmetries of the equilateral triangle is the Coxeter group A_2 . The group of symmetries of a regular tetrahedron is A_3 .




More generally, A_n is the group of symmetries of a regular n -dimensional simplex —which is just the group of all permutations of the $n + 1$ vertices.

Then there is BC_n , which has n dots:




Only the last edge is labeled with a 4. Later we'll see how BC_n spawns two Dynkin diagrams called B_n and C_n ; this accounts for its strange name.


The Coxeter group BC_2 is the symmetry group of a square, while BC_3 is the symmetry group of the cube or its dual, the regular octahedron.




In general, BC_n is the group of symmetries of an n -dimensional hypercube, or its dual, whose vertices lie at the centers of the faces of this hypercube. The dual of a hypercube could be called a “regular hyperoctahedron”, but it's actually called a **cross-polytope** or **orthoplex**. Here's how the duality works for $n = 3$:



The next infinite series of finite reflection groups is D_n , which has n dots arranged in a row that branches at the end. Here is D_6 :




Since $D_n \cong A_n$ for $n < 4$, the first really exciting case is D_4 :




and the symmetry of this diagram, called **triality**, gives birth to many remarkable things.

The Coxeter group D_n is the symmetry group of the n -dimensional **demihypercube**, whose vertices are gotten by taking every other vertex of a hypercube: that is, half the hypercube's vertices, with no two right next to each other. This is easiest to visualize when $n = 3$. Here are the two demicubes in a cube:



They're just regular tetrahedra! We might have expected this, since $D_3 \cong A_3$. A 4-dimensional demihypercube has one of these tetrahedra "on top" and the other "on the bottom".

Then there are E_6 , E_7 , and E_8 :




Interestingly, this series does *not* go on. Thus the A_n , B_n , and D_n Coxeter diagrams are called “classical” while the E series is called “exceptional”.

A **polytope** is the higher-dimensional generalization of a polygon or polyhedron. The E_6 , E_7 and E_8 Coxeter groups are symmetries of some remarkably subtle polytopes:

- The E_8 Coxeter group is the symmetry group of an 8-dimensional polytope with 240 vertices called the **E_8 root polytope**. To build this, take a sphere 8 dimensions and get as many equal-sized spheres as possible to touch it. There will be 240. The centers of these spheres are the vertices of an E_8 root polytope.
- The E_7 Coxeter group is the symmetry group of a 7-dimensional polytope with 126 vertices called the **E_7 root polytope**. To build this, pick any vertex of the E_8 root polytope. Draw a line through it and the center of that polytope. The vertices in the 7d space orthogonal to this line are the vertices of a E_7 root polytope.
- The E_6 Coxeter group is the symmetry group of a 6-dimensional polytope with 72 vertices called the **E_6 root polytope**. Pick any 6 vertices of the E_8 root polytope forming a regular hexagon in some plane containing the center of that polytope. The vertices in the 6d space orthogonal to this plane are the vertices of an E_7 root polytope.

We will eventually give more manageable descriptions of these polytopes.

Then there is F_4 :



The F_4 Coxeter group is the symmetry group of a 4-dimensional polytope with 24 vertices and 24 octahedral faces, called the **24-cell**. To build a 24-cell, first draw a 4-dimensional hypercube centered at the origin. Put a point in the center of each face to get the vertices of a 4d orthoplex. Then expand the orthoplex until its vertices are as far from origin as the hypercube’s vertices. Then the vertices of the hypercube and orthoplex, taken together, are the vertices of a 24-cell!

Regular polytopes are the most obvious higher-dimensional generalization of Platonic solids. In all dimensions $n > 4$ there are just three regular polytopes: the n -simplex, the n -dimensional hypercube and the n -dimensional orthoplex. In dimension 4 there are three more! One is the 24-cell. We shall meet the other two soon.

Then there is G_2 :




The G_2 Coxeter group is just the symmetry group of the regular hexagon. There are infinitely many regular polygons; we'll see later why the equilateral triangle (with symmetry group A_2 , the square (with symmetry group B_2 and the hexagon (with symmetry group G_2 are singled out as special. If you've ever tried to tile your floor with regular pentagons or heptagons maybe you can guess!

Then there are H_3 and H_4 :



H_3 is the group of symmetries of the regular dodecahedron or icosahedron.



H_4 is the group of symmetries of the remaining two 4d regular polytopes. They could be called the "hyperdodecahedron" and "hypericosahedron", but in fact they are called the 120-cell and 600-cell.

To get your hands on these things, it is quickest to start with a regular dodecahedron. It has 60 rotational symmetries since you can rotate your favorite faces to any of the 12 faces, and in 5 ways. These rotational symmetries form a group Γ that's a subgroup of $SO(3)$, the group of all rotations in 3d space. There's a 2-1 map from $SU(2)$ to $SO(3)$, so two elements of $SU(2)$ map to each element of Γ . These element of $SU(2)$ form a 120-element subgroup of $SU(2)$. But geometrically $SU(2)$ is a sphere in 4d space: for example, it's isomorphic to the group of unit quaternions. So, Γ consists of 120 points on a sphere in 4d space... and these are the vertices of the **600-cell**.

The 600-cell has 120 vertices and 600 tetrahedral faces. If we put a point at the center of each face, we get the vertices of the dual polytope, which has 600 vertices and 120 tetrahedral faces. This is called the **120-cell**.

Finally, there is another very simple infinite series, I_m :



The Coxeter group coming from this diagram is the symmetry group of the regular m -gon. Notice that

$$I_3 \cong A_2, \quad I_4 \cong B_2, \quad I_6 \cong G_2.$$

So, the only *new* Coxeter diagrams here are I_m for $m = 5$ and $m \geq 7$.


We can get other finite reflection groups from disjoint unions of the Coxeter diagrams we've already seen: they're just products of the Coxeter groups we've already seen. But our list of *connected* Coxeter diagrams giving finite reflection groups is *done!*

Let's list them without repetitions:


• $A_n, n > 0$




• $BC_n, n > 1$




• $D_n, n > 3$




• E_6, E_7, E_8




• F_4




• G_2



• H_4, H_5



• I_m




The proof that these are the only possibilities is a rather elaborate inductive argument. For details, see for example *Finite Reflection Groups and Coxeter Groups* by Humphreys.

Note that the symmetry groups of the Platonic solids and their higher-dimensional relatives fit in nicely into this classification. Later we shall see a more mysterious relation between Platonic solids and these diagrams. But first we need to introduce Dynkin diagrams! These show up when we think about lattices.

Dynkin diagrams and lattices

We get a **lattice** by taking n linearly independent vectors in n -dimensional Euclidean space and forming all linear combinations with integer coefficients.



Sometimes lattices have interesting symmetry groups. Every lattice in n dimensions has \mathbb{Z}^n acting as translation symmetries, so let's focus on symmetries that fix the origin: that is, combinations of rotations and reflections that map the lattice to itself. To do this, we can exploit the classification of finite reflection groups.

So, suppose we have a connected Coxeter diagram that gives a finite reflection group. If our diagram has n dots, this group acts on \mathbb{R}^n . *When is there a lattice in \mathbb{R}^n having this group as symmetries?*

If one exists, we say our group satisfies the **crystallographic condition**. The only ones that do are

$$A_n, B_n, D_n, E_6, E_7, E_8, F_4, \text{ and } G_2!$$

In other words, Coxeter diagrams with any edges labeled by numbers $m = 5$ or $m \geq 7$ are ruled out. So, we can't find lattices whose symmetries include the symmetry groups of the regular pentagon, the regular heptagon, or regular polygons with more sides. This has big implications for crystals, but also for pure math.


Say we have a finite reflection group Γ obeying the crystallographic condition. How can get a lattice with Γ as symmetries? It turns out we can always do it by taking linear combinations of vectors, one for each dot in the Coxeter diagram. The reflections generating Γ will be reflections through these vectors.

To get this to work, the angle between two of these vectors needs to be π/m when the edge between the dots is labeled by m . But remember: in this game an unlabeled edge counts as an edge labeled by 3, so then the vectors need to be at a $\pi/3$ angle from each other. No edge at all counts as an edge labeled by 2, so then the vectors need to be at an angle of $\pi/2$ from each other—that is, at right angles.

For example, take A_2 :




where I'm writing the 3 to clarify some patterns. To get a lattice with this Coxeter group as symmetries, take two vectors of the same length at an angle of $\pi/3$ from each other. Their integer linear combinations form the lattice we want. It's the lattice of vertices of the **triangular tiling**:



Or consider B_2 :




To get a lattice with this Coxeter group as symmetries, we can use two vectors at an angle of $\pi/4$ from each other. But they can't be of equal length! One must be $\sqrt{2}$ times as long as the other. Then their integer linear combinations form the lattice we want. It's the lattice of vertices of the **square tiling**:



Similarly, for G_2 :



To get a lattice with this Coxeter group as symmetries, we can use two vectors at an angle of $\pi/6$ from each other. One must be $\sqrt{3}$ times as long as the other. Then their integer linear combinations form a lattice with the symmetries we want. It's the lattice of vertices of the **hexagonal tiling**:




No, it's not! I was just checking to see if you're paying attention! The vertices of the hexagonal tiling don't form a lattice. To get a lattice you can throw in an extra point at the center of each hexagon, but then you get the triangular tiling again. In fact the G_2 lattice looks just like the A_2 lattice, at least up to a rotation and rescaling. The difference is that now we're considering it with a bigger symmetry group: A_2 gives the symmetry group of an equilateral triangle, while G_2 gives that of a regular hexagon.

In general, a Coxeter diagram can only give a group preserving a lattice if all its edges are labelled by $m = 3, 4$ or 6 . To build this lattice we need to pick a basis of vectors, one for each dot in the diagram. Now suppose two dots are connected by an edge.


- If $m = 3$, we choose vectors for them with the same length.
- If $m = 4$, we choose one vector to be $\sqrt{2}$ times as long as the other.
- If $m = 6$, we choose one vector to be $\sqrt{3}$ times as long as the other.

When $m = 4$ or 6 we need to decide which vector is longer. So, we draw an arrow on the edge from the dot with the longer vector to the dot with the shorter vector. These arrows are enough to specify a lattice on which our Coxeter group acts as symmetries!—at least up to rotations, reflections and rescalings.

For example, in the diagram F_4 we have two choices:




and




However, in this case the resulting diagrams with arrows are isomorphic: you can turn one around and get the other.

In fact, you can check that for every type of Coxeter diagram except BC_n with $n \geq 3$, you get isomorphic diagrams with arrows no matter how you point the arrows! So, except in those cases, there is only *one* lattice having that group of symmetries—up to rotations, reflections and rescalings. But for BC_n with $n \geq 3$ there are two! Recall that the diagram BC_n looks like this, with n dots:



with n dots. The two ways to draw arrows on the last edge are called B_n and C_n :


- B_n :



- C_n :



The points near the origin in the B_3 lattice lie on a cube, while those in the C_3 lie on an octahedron:



That's how it works in higher dimensions too, with a hypercube and orthoplex.

People also use another style of diagram to describe what's going on here: "Dynkin diagrams". For these, we take our Coxeter diagram with arrows on it, and:

- replace any edge labeled with a 4 by two parallel edges;
- replace any edge labeled with a 6 by three parallel edges.


Yes, I know this sounds confusing at first! It's just a different way to draw the same information. I guess people like it because you can draw pictures without any numbers. Let me do a couple of examples. Our friend G_2 , with an arrow on it:



gets drawn as this Dynkin diagram:



while our friend B_3 :



gets drawn as this Dynkin diagram:




The upshot is that any Dynkin diagram with n dots describes a collection of vectors $v_1, \dots, v_n \in \mathbb{R}^n$ such that:

1. integer linear combinations of these vectors form a lattice $L \subset \mathbb{R}^n$
2. reflections through these vectors generate a finite reflection group Γ
3. the action of Γ on \mathbb{R}^n preserves the lattice L .

People usually normalize these vectors so that the shortest ones have $v_i \cdot v_i = 2$, for reasons that should become clear later. This determines them up to rotations and reflections. Furthermore, if the action of Γ on \mathbb{R}^n is **indecomposable**, meaning we can't chop \mathbb{R}^n into a direct sum of two nontrivial subspaces both preserved by Γ , then any collection of vectors obeying 1–3 comes from a *connected* Dynkin diagram!

And here are all the connected Dynkin diagrams, drawn in a more artistic style by R. A. Nonemacher. The dots are drawn as big circles:



Compact simple Lie algebras

Now let us turn to the theory of Lie groups. Lie groups are the most important “continuous” (as opposed to discrete) symmetry groups. Examples include the real line with addition as the group operation, the circle with addition modulo 2π , and the so-called “classical groups”, which include:

- The **general linear group** $GL(N, \mathbb{C})$, consisting of all invertible linear transformations of \mathbb{C}^N , or in other words, all $N \times N$ complex matrices with nonzero determinant.
- The **special linear group** $SL(N, \mathbb{C})$, consisting of all linear transformations of \mathbb{C}^N with determinant 1.
- The **unitary group** $U(N)$, consisting of all unitary linear transformations of \mathbb{C}^N .
- The **special unitary group** $SU(N)$, consisting of all unitary linear transformations of \mathbb{C}^N with determinant 1.

All these Lie groups are incredibly important in both physics and mathematics. Thus it is wonderful, and charmingly ironic, that the same Dynkin diagrams that classify the oh-so-discrete lattices with finite reflection groups as symmetries also classify some of the most beautiful of Lie groups: the “simple” Lie groups.

There is a vast amount known about semisimple Lie groups, and everyone really serious about mathematics winds up needing to learn some of this stuff. I took courses on Lie groups and their Lie algebras in grad school, but it was only later that I really came to appreciate the beauty of the simple Lie groups. One reason I found them frustrating was that the work involved in their classification was so algebraic, and I preferred the more geometrical aspects of Lie groups. This algebraic approach de-emphasized the Lie groups themselves, and emphasized an tool for working with Lie group: namely, Lie *algebras*.

So what’s the basic idea? Let me summarize two semesters of grad school, and tell you the basic stuff about Lie groups and the classification of simple Lie groups. Forgive me if it’s a bit rushed and sketchy: hopefully the main ideas will shine through the murk better this way.

A Lie group is a group that’s also a manifold, for which the group operations (multiplication and taking inverses) are smooth functions. This lets you form the tangent space to any point in the group, and the tangent space at the identity plays a special role. It’s called the **Lie algebra of the group**. If we have any element x in the Lie algebra, we can exponentiate it to get an element $\exp(x)$ in the group, and we can keep track of the noncommutativity of the group by forming the **Lie bracket** of elements x and y in the Lie algebra:

$$[x, y] = \left. \frac{d}{dt} \frac{d}{ds} \exp(sx) \exp(ty) \exp(-sx) \exp(-ty) \right|_{s,t=0}$$

where s and t are real numbers. Note that $[x, y] = 0$ if the group is commutative. This bracket operation satisfies some axioms, and algebraists call anything a Lie algebra that satisfies those axioms. Just to satisfy your curiosity, a **Lie algebra** is a vector space \mathfrak{g} with a bracket operation such that

- the bracket is linear in the second argument:

$$[x, ay + bz] = a[x, y] + b[x, z] \quad \text{for all } x, y, z \in \mathfrak{g}, a, b \in \mathbb{R}$$

- the bracket is antisymmetric:

$$[x, y] = -[y, x] \quad \text{for all } x, y \in \mathfrak{g}$$

and thus linear in the first argument as well,

- bracketing with any element x satisfies a version of the product rule

$$[x, [y, z]] = [[x, y], z] + [y, [x, z]] \quad \text{for all } x, y, z \in \mathfrak{g}$$

called the **Jacobi identity**

All our Lie algebras will be finite-dimensional real vector spaces. For example, you could take $n \times n$ real matrices and let $[x, y] = xy - yx$.

I will not actually do anything with these axioms, since this is not really a course on Lie algebras. But these axioms are part of why Lie algebras are so wonderful. They're all about linear algebra—they're just a vector space with a bracket operation obeying some axioms—and yet we can *almost* recover any Lie group from its Lie algebra.

Why “almost”? Well, first of all, we can't do it unless we require that our Lie group is connected. The reason is that we could add extra connected components without changing the tangent space at the identity or its bracket operation. For example, every finite group is a Lie group with Lie algebra $\{0\}$. And second of all, we can't do it unless we require a bit more. The reason is that any connected covering space of a connected Lie group is another Lie group with the same Lie algebra. Luckily, every connected Lie group has a “universal cover” which has no covering spaces except itself. This guy is both connected and simply connected.

With these caveats we're okay: we can recover a connected and simply connected Lie group from its Lie algebra. Even better, *every* finite-dimensional Lie algebra comes from a connected and simply connected Lie group!

This reduces the problem of classifying connected and simply connected Lie groups to a problem in linear algebra. Unfortunately it's an incredibly hard problem in linear algebra. It appears to be hopeless unless we stick to low dimensions. There are just too many ways to build bigger Lie groups, or Lie algebras, from smaller ones. But the problem becomes much easier if we stick to *compact* Lie groups—so that's what we will do.

The Lie algebra of a compact Lie group is called a **compact Lie algebra**. It's not that the Lie algebra is literally compact: it's a vector space. But compact Lie algebras are very nice. It turns out that we can take direct sums of Lie algebras by defining operations componentwise, and a compact Lie algebra is always the direct sum of an “abelian” Lie algebra and a “semisimple” one. These lie at opposite extremes.

To understand these it helps to think about the “Killing form” of a Lie algebra \mathfrak{g} . For any $x \in \mathfrak{g}$ there's a linear operator on \mathfrak{g} given by bracketing with x ; it's usually called $\text{ad}_x : \mathfrak{g} \rightarrow \mathfrak{g}$.

$$\text{ad}_x(y) = [x, y].$$

The **Killing form** is a bilinear form on \mathfrak{g} given by

$$B(x, y) = \text{tr}(\text{ad}(x)\text{ad}(y)).$$

By the cyclic property of the trace,

$$B(x, y) = B(y, x).$$

This is great: it's a bit like getting an inner product for free! But stay tuned: it's a bit subtler than that.

Here are the two kinds of Lie algebras I mentioned:

- A Lie algebra is **abelian** if $[x, y] = 0$ for all x and y .

- A Lie algebra is **semisimple** if the Killing form is nondegenerate: if $\langle x, y \rangle = 0$ for all y then $x = 0$.

Note that the Killing form of an abelian Lie algebra is zero, which is at the opposite extreme from being nondegenerate. The Killing form of a semisimple Lie algebra is not really an inner product—it doesn't need to be positive definite. But it comes very close for a **compact semisimple Lie algebra**—that is, a semisimple Lie algebra that's the Lie algebra of a compact Lie group. Namely: for such a Lie algebra, the Killing form is *negative* definite, so

$$\langle x, y \rangle = -B(x, y)$$

is an inner product.

There's one abelian Lie algebra of each dimension, and every abelian Lie algebra is the Lie algebra of a **torus**, which is a product of finitely many circles. So all the hard work lies in understanding the semisimple Lie algebras. A connected Lie group whose Lie algebra is semisimple is called—surprise!—a **semisimple Lie group**.

Now let's start with a compact semisimple Lie group G and see how there's a Dynkin diagram hiding inside it. For starters, inside G there is always some subgroup T that's a torus not contained in any larger torus. Such a subgroup is called a **maximal torus** of G . It's basically unique: there are a bunch of maximal tori, but any two are conjugate to each other in G . So, people often sloppily talk about “the” maximal torus.

Let $\text{Lie}(T)$ stand for the Lie algebra of a maximal torus T . The inner product on $\text{Lie}(G)$ restricts to an inner product on $\text{Lie}(T)$. So, $\text{Lie}(T)$ is isomorphic to good old \mathbb{R}^n with its usual inner product. Furthermore, sitting inside $\text{Lie}(T)$ there is a lattice L , consisting of all elements x with $\exp(x) = 1$. This is how lattices in \mathbb{R}^n sneak into the picture!

How do finite reflection groups sneak into the picture? Note that for some elements g in G , if we conjugate T by g , that is, form the set of all elements gtg^{-1} where t is in T , we get T back. In other words, these elements of G act as symmetries of the torus T . But some elements of $g \in G$ act trivially on T : they have $gtg^{-1} = t$ for all $t \in T$. So, the quotient group

$$W(G) = \frac{\{g \in G \mid gtg^{-1} \in T \text{ for all } t \in T\}}{g \in G \mid gtg^{-1} = t \text{ for all } t \in T}$$

acts on our maximal torus T . And $W(G)$ is called the **Weyl group** of G . (In reality it also depends on our choice of maximal torus, but changing our maximal torus to another one gives an isomorphic Weyl group, so our notation ignores the dependence on T .)


If a group acts as symmetries of something, it also acts as symmetries of everything naturally cooked up from that thing. For this reason, the Weyl group of G also acts as symmetries of $\text{Lie}(T)$ and of the lattice L sitting inside $\text{Lie}(T)$. So we have a lattice together with a group of symmetries acting on it.

And here come two nonobvious theorems. First, the Weyl group is actually a finite reflection group acting on $\text{Lie}(T)$! From what we learned in the previous section, this means our lattice and its Weyl group of symmetries is determined, up to isomorphism, by some Dynkin diagram. And second, it turns out that the Lie algebra of G is determined, up to isomorphism, by the lattice L and the finite reflection group acting on it!

Putting everything together, we get one-to-one correspondences between the following three things, each considered up to the relevant sort of isomorphism:

- Dynkin diagrams
- compact semisimple Lie algebras
- connected and simply connected compact Lie groups.

Furthermore, it turns out that the operation of taking direct sums of Lie groups of this sort corresponds to taking direct sums of their Lie algebras, which corresponds to taking disjoint unions of Dynkin diagrams. So to get the “building blocks”, from which everything else can be built via direct sums, we only need to worry about the *connected* Dynkin diagrams, which we have completely classified:



The compact semisimple Lie algebras coming from *connected* Dynkin diagrams are called **compact simple Lie algebras**. But what are they actually like? People have figured them out. Amaz-

ingly, those that come in infinite series are all related to rotations in real, complex or quaternionic vector spaces. So let me give you a crash course on those:

1. The real-linear transformations of \mathbb{R}^n that preserve its usual inner product

$$\langle v, w \rangle = \sum_{i=1}^n v_i w_i$$

form a compact Lie group called the **orthogonal group** $O(n)$. The subgroup consisting of transformations with determinant 1 is a compact Lie group called the **special orthogonal group** $SO(n)$. Its Lie algebra is called $\mathfrak{so}(n)$, and it consists of $n \times n$ real matrices that have trace zero and are minus their own transpose. This is a compact simple Lie algebra when $n \geq 3$. What is it when $n = 1$ or 2 ?

2. The complex-linear transformations of \mathbb{C}^n that preserve its usual inner product

$$\langle v, w \rangle = \sum_{i=1}^n \bar{v}_i w_i$$

form a compact Lie group called the **unitary group** $U(n)$. The subgroup consisting of transformations with determinant 1 is a compact Lie group called the **special unitary group** $SU(n)$. Its Lie algebra is called $\mathfrak{su}(n)$, and it consists of $n \times n$ complex matrices that have trace zero and are minus their own conjugate transpose. This is a compact simple Lie algebra when $n \geq 2$. What is it when $n = 1$?

3. The quaternion-linear transformations of \mathbb{H}^n that preserve its usual inner product

$$\langle v, w \rangle = \sum_{i=1}^n \bar{v}_i w_i$$

form a Lie group called the **quaternionic unitary group** $Sp(n)$. (With Dieudonné's definition of the quaternionic determinant, all matrices in this group have determinant 1.) The Lie algebra of $Sp(n)$ is called $\mathfrak{sp}(n)$, and it consists of $n \times n$ quaternionic matrices that have trace zero and are minus their own conjugate transpose. This is a compact simple Lie algebra when $n \geq 1$.

The third item looks a lot like the first two, but you may be unfamiliar with the quaternions \mathbb{H} . For now I'll just say that there are three finite-dimensional associative algebras over \mathbb{R} equipped with a norm obeying

$$|ab| = |a||b|.$$

These are the real numbers \mathbb{R} , the complex numbers \mathbb{C} , and most excitingly the **quaternions**

$$\mathbb{H} = \{a + bi + cj + dk \mid a, b, c, d \in \mathbb{R}\}$$

where multiplication is determined by the equations Hamilton carved into a wall on on the 16th of October in 1843:

$$i^2 = j^2 = k^2 = ijk = -1.$$

We can define conjugation for quaternions by

$$q = a + bi + cj + dk \implies \bar{q} = a - bi - cj - dk$$

and then the norm is given by

$$|q| = \sqrt{q\bar{q}} = \sqrt{\bar{q}q}.$$

The space \mathbb{H}^n acts like a “quaternionic vector space” even though \mathbb{H} is not a field: we can multiply vectors on the left or the right by quaternions, and we say that $T: \mathbb{H}^n \rightarrow \mathbb{H}^n$ is **quaternion-linear** if

$$T(vq) = T(v)q \quad \text{for all } v \in \mathbb{H}^n, q \in \mathbb{H}.$$

With this convention, left multiplication by any $n \times n$ matrix of quaternions gives a quaternion-linear map $T: \mathbb{H}^n \rightarrow \mathbb{H}^n$. While the quaternions are a fascinating subject, this is all you need to know for now.

Now we are ready to list the four infinite series of compact simple Lie algebras:

- **A_n**: This Dynkin diagram gives the compact simple Lie algebra $\mathfrak{su}(n+1)$.
- **B_n**: This Dynkin diagram gives the compact simple Lie algebra $\mathfrak{so}(2n+1)$.
- **C_n**: This Dynkin diagram gives the compact simple Lie algebra $\mathfrak{sp}(n)$.
- **D_n**: This Dynkin diagram gives the compact simple Lie algebra $\mathfrak{so}(2n)$.

These are called the **classical** compact simple Lie algebras, and they would be pretty easy to reinvent for yourself, or get interested in for all sorts of reasons. It may seem weird that $\text{SO}(2n)$ is so different from $\text{SO}(2n+1)$, but it’s true! For example, can put n orthogonal planes in \mathbb{R}^{2n} , and by doing a rotation in each of these planes you get an element of $\text{SO}(2n)$ that fixes only the origin. But in odd dimensions there’s one dimension left over, so any rotation must fix some nonzero vector.

The remaining five compact simple Lie algebras are called **exceptional**, and they are much more mysterious. They were only discovered when people like Killing and Cartan figured out the classification of simple Lie algebras. And as it turns out, they are all related to the octonions! The octonions \mathbb{O} are the only finite-dimensional *nonassociative* algebra over \mathbb{R} that is equipped with a norm obeying

$$|ab| = |a||b|.$$

They are an amazing freak of nature, which begets many other strange things. I am somewhat obsessed with them, but I won’t say much about them here: see the references for more.

Here are the 5 exceptional compact simple Lie algebras. It is quickest to describe them using compact Lie groups. I’ll list them in order of dimension, not alphabetical order:

1. **G₂**: This Dynkin diagram gives a 14-dimensional compact simple Lie algebra called \mathfrak{g}_2 . The automorphism group of the octonions is a compact Lie group whose Lie algebra is \mathfrak{g}_2 .
2. **F₄**: This Dynkin diagram gives a 52-dimensional compact simple Lie algebra called \mathfrak{f}_4 . The octonions give a 16-dimensional projective plane called $\mathbb{O}P^2$. This is a Riemannian manifold, and its **isometry group**—the group of diffeomorphisms preserving the Riemannian metric—is a compact Lie group whose Lie algebra is \mathfrak{f}_4 .

3. E_6 : This Dynkin diagram gives a 78-dimensional compact simple Lie algebra called \mathfrak{f}_4 . The octonions tensored with the complex numbers give a 32-dimensional projective plane called $(\mathbb{C} \otimes \mathbb{O})P^2$. This is a Riemannian manifold, and its isometry group is a compact Lie group whose Lie algebra is \mathfrak{e}_6 .
4. E_7 : This Dynkin diagram gives a 133-dimensional compact simple Lie algebra called \mathfrak{e}_7 . The octonions tensored with the quaternions give a 64-dimensional Riemannian manifold called $(\mathbb{H} \otimes \mathbb{O})P^2$, even though it is not technically a projective plane. Its isometry group is a compact Lie group whose Lie algebra is \mathfrak{e}_7 .
5. E_8 : This Dynkin diagram gives a 248-dimensional compact simple Lie algebra called \mathfrak{e}_8 . The octonions tensored with the octonions give a 128-dimensional Riemannian manifold called $(\mathbb{O} \otimes \mathbb{O})P^2$, even though it is not technically a projective plane. Its isometry group is a compact Lie group whose Lie algebra is \mathfrak{e}_8 .

A tragic fact is that currently nobody know how to get their hands on the Riemannian manifolds $(\mathbb{H} \otimes \mathbb{O})P^2$ and $(\mathbb{O} \otimes \mathbb{O})P^2$ without building their isometry groups first, or at the same time. Thus the explanations we have given of \mathfrak{e}_7 and \mathfrak{e}_8 , while true, are not as useful as we might like. This is especially frustrating for \mathfrak{e}_8 , since this is the “king” of simple Lie algebras, connected to many other amazing exceptional structures in mathematics.

We’ve listed the compact simple Lie algebras, but what about their Lie groups? This is an important subject, since Lie algebras are ultimately just a tool for working with Lie groups. But we need to be a bit careful, since a Lie algebra may be the Lie algebra of several nonisomorphic connected Lie groups. Remember, taking a covering space of a Lie group does not change its Lie algebra. For example, the Dynkin diagram coincidence $A_1 \cong B_1$ implies that $\mathfrak{su}(2) \cong \mathfrak{so}(3)$, but the Lie group $SU(2)$ is not isomorphic to $SO(3)$: it’s a double cover of $SO(3)$.

People usually take a relaxed attitude and call *any* connected Lie group whose Lie algebra is a compact simple Lie algebra a **compact simple Lie group**. This is true even though such a group may have nontrivial normal subgroups, so it is not simple in the usual sense of group theory. For example $SU(2)$ has a 2-element normal subgroup, its center, containing ± 1 . Yet we call it a compact simple Lie group.

With this terminology in place, we get a one-to-one correspondence between these three things, each considered up to the relevant sort of isomorphism:

- connected Dynkin diagrams
- compact simple Lie algebras
- simply connected compact simple Lie groups.

Among the the classical compact simple Lie groups it turns out that $SU(n)$ and $Sp(n)$ are simply connected, while $SO(n)$ is not: for $n \geq 3$, which is all that matters here, $SO(n)$ has a simply connected double cover called the **spin group** $Spin(n)$. This group is very important in physics and also differential geometry, since it has important representations called “spinors” that are not representations of $SO(n)$. I have written quite a bit about the exceptional compact simple Lie groups, but here I’ll just mention that any connected Lie group with Lie algebra $\mathfrak{g}_2, \mathfrak{f}_4, \mathfrak{e}_6, \mathfrak{e}_7$ or \mathfrak{e}_8 is called G_2, F_4, E_6, E_7 or E_8 . This is uncreative and a bit ambiguous, but that’s that way it is.

Simply-laced Dynkin diagrams

We've seen that any Dynkin diagram with n dots describes a collection of vectors $v_1, \dots, v_n \in \mathbb{R}^n$ such that:

1. integer linear combinations of these vectors form a lattice $L \subset \mathbb{R}^n$
2. reflections through these vectors generate a finite reflection group Γ
3. the action of Γ on \mathbb{R}^n preserves the lattice L .


Unfortunately, we can't always choose all the vectors v_i to have the same length. But we've seen that this problem doesn't happen when all these vectors are at an angle of $\pi/2$ or $\pi/3$ from each other. This happens when our Dynkin diagram is **simply laced**: all its edges are unlabeled.

Here are all the connected simply-laced Dynkin diagrams:


- A_n , which has n dots like this:



- D_n , which has n dots, where we can assume $n > 3$ since $D_n \cong A_n$ for $n = 1, 2, 3$:



- $E_6, E_7,$ and E_8 :



These diagrams are ubiquitous in mathematics. But before getting into that I should describe the corresponding lattices more explicitly, to make it clear how simple they really are.

So, what are the A, D, and E lattices?

- A_n : We can describe the A_n lattice as the set of all $(n + 1)$ -tuples of integers (x_1, \dots, x_{n+1}) such that

$$x_1 + \dots + x_{n+1} = 0.$$

It's a fun exercise to show that A_2 is a 2-dimensional hexagonal lattice, the sort of lattice you use to pack pennies as densely as possible. Similarly, A_3 gives a standard way of packing cannon balls, which is the densest lattice packing of equal-sized spheres in 3 dimensions. A much harder fact, due to Hales, is that no non-lattice packing of equal-sized spheres can beat the density of the A_3 lattice.

- D_n : We can describe the D_n lattice as the set of all n -tuples of integers (x_1, \dots, x_n) such that

$$x_1 + \dots + x_n \text{ is even.}$$

Or, if you like, you can imagine taking an n -dimensional checkerboard, coloring the cubes alternately red and black, and taking the center of each red cube. In four dimensions, D_4 gives a denser packing of spheres than A_4 ; in fact, it gives the densest lattice packing possible. Moreover, D_5 gives the densest lattice packing of in dimension 5. However, in dimensions 6, 7, and 8, the E_n lattices give the densest lattice packings. In fact Viasovska showed that in 8 dimensions, no non-lattice packing of equal sized spheres can beat the density of the E_8 lattice!

- E_6, E_7, E_8 : We can describe the E_8 lattice as the set of 8-tuples (x_1, \dots, x_8) such that the x_i are either all integers or all integers plus $1/2$ and

$$x_1 + \dots + x_8 \text{ is even.}$$

Each point in this lattice has 240 nearest neighbors. For example, the nearest neighbors of the origin have length $\sqrt{2}$, and you can check there are 240 of them.

Alternatively, if you take the D_8 lattice and use it to pack equal-sized spheres that just touch each other, there is actually just enough room to slip in another D_8 lattice of equal-sized spheres in the remaining space, doubling the density! And if you do this, your spheres will be centered at points in the E_8 lattice.

Once you have E_8 in hand, you can get its little pals E_7 and E_6 as follows. To get E_7 , just take all the vectors in E_8 that are perpendicular to one lattice vector of length $\sqrt{2}$. To get E_6 , find a copy of the lattice A_2 in E_8 generated by 2 vectors of length $\sqrt{2}$, and then take all the vectors in E_8 perpendicular to everything in that copy of A_2 .

The A, D and E Dynkin diagrams show up in many places throughout mathematics, in a spooky sort of way. Let me sketch three of the most famous.

First, Witt's theorem says that the A, D, and E lattices and their direct sums are the only integral lattices having a basis consisting of vectors v with $\|v\|^2 = 2$. Here a lattice is **integral** if the dot product of any two vectors in it is an integer. In fact, any integral lattice having a basis consisting of vectors with $\|v\|^2$ equal to 1 or 2 is a direct sum of copies of A, D, and E lattices and the integers, thought of as a 1-dimensional lattice.

Second, a **quiver** is just some dots with arrows between them. A **representation** of a quiver is a way of assigning a finite-dimensional complex vector space to each dot and a linear map between these vector spaces to each arrow. There's an obvious category of representations $\text{Rep}(Q)$ of any quiver Q . Gabriel proved an astounding result about these categories $\text{Rep}(Q)$. We say a quiver Q has **finite representation type** if $\text{Rep}(Q)$ has finitely many isomorphism classes of **indecomposable** objects: objects that aren't direct sums of others. And, it turns out the quivers of finite representation type are just those coming from simply-laced Dynkin diagrams!


Actually, for this to make sense, you need to take your Dynkin diagram and turn it into a quiver by putting arrows along the edges. If you have a simply-laced Dynkin diagram, you get a quiver of finite representation type no matter which way you let the arrows point.

Third, there is a cool relationship between the ADE diagrams and the symmetry groups of the Platonic solids, called the **McKay correspondence**. Here is one way to get it. First, take the rotational symmetry group of a Platonic solid, not including reflections, or more generally any finite subgroup G of $SO(3)$. Since $SO(3)$ has $SU(2)$ as a double cover, you can get a double cover of G , say \tilde{G} , sitting inside $SU(2)$. Since \tilde{G} is finite, it has finitely many irreducible representations on complex vector spaces (up to isomorphism). Draw a dot for each of these. One comes from the obvious representation of $SU(2)$ on \mathbb{C}^2 . When you tensor this one with any other irreducible representation R , you get a direct sum of irreducible representations. Draw one line from the dot for R to another dot for each time that other irreducible representation appears in your direct sum. What do you get?


You get an “affine Dynkin diagram”, which is like a usual Dynkin diagram but with an extra dot thrown in—corresponding to the trivial rep of \tilde{G} . And if you throw out that extra dot, you get a simply laced Dynkin diagram! In fact you get all all the connected simply laced Dynkin diagrams this way!

The correspondence goes like this:


- A_n : this corresponds to the cyclic group sitting inside $SO(n)$ as the rotational symmetries of a regular n -gon, where we don't let ourselves flip this polygon over.
- D_n : this corresponds to the dihedral group sitting inside $SO(n)$ as the symmetries of a regular n -gon, where we *do* let ourselves flip this polygon over.
- E_6 : this corresponds to the rotational symmetry group of the regular tetrahedron.



- E_7 : this corresponds to the rotational symmetry group of the regular cube, or octahedron.



- E_8 : this corresponds to the rotational symmetry group of the regular dodecahedron, or icosahedron.



We saw another relation between Platonic solids and Coxeter diagrams near the start of this paper, but that one made sense. This one is black magic.

References

For a gentle introduction to symmetry groups in 2 and 3 dimensions and a valuable warmup for higher-dimensional considerations, see:

- D. L. Johnson, *Symmetries*, Springer, Berlin, 2001.

For the classification of finite reflection groups in terms of Coxeter diagrams, see:

- James E. Humphreys, *Finite Reflection Groups and Coxeter Groups*, Cambridge U. Press, Cambridge, 1990.

For their appearance as symmetry groups of polytopes, try this endlessly entertaining book:

- Harold Scott Macdonald Coxeter, *Regular Polytopes*, Dover, Mineola, 1974.

For a huge amount of information on lattices, see this book:

- John Horton Conway and Neil James Alexander Sloane, *Sphere Packings, Lattices and Groups*, Springer, Berlin, 2013.

Gian-Carlo Rota said of this book, “This is the best survey of the best work in the best fields of combinatorics written by the best people. It will make the best reading by the best students interested in the best mathematics that is now going on.” Chapter 4 discusses the lattices coming from Dynkin diagrams.

For the classification of compact Lie groups, see:

- John Frank Adams, *Lectures on Lie Groups*, Benjamin, New York, 1969.

This is nice because it uses topology and handles compact Lie groups that are not semisimple on a more or less equal footing with the semisimple ones. There are many texts that take a more algebraic approach to classifying semisimple Lie algebras and/or Lie groups: Here are a few:

- William Fulton and Joe Harris, *Representation Theory — a First Course*, Springer, Berlin, 1991.

- Sigurdur Helgason, *Differential Geometry, Lie Groups, and Symmetric Spaces*, Academic Press, New York, 1979.
- James Humphreys, *Introduction to Lie Algebras and Representation Theory*, Springer, Berlin, 2012.
- Daniel Bump, *Lie Groups*, Springer, Berlin, 2004.

Often such books start by classifying semisimple Lie algebras over \mathbb{C} , and then show each one is the complexification of a unique compact semisimple Lie algebra. Each such complex Lie algebra is also the Lie algebra of a unique connected and simply connected **complex Lie group**, meaning a group in the category of complex manifolds and holomorphic maps. So, there is actually a one-to-one correspondence between five things, each considered up to isomorphism:

- connected Dynkin diagrams
- compact simple Lie algebras
- connected and simply connected compact simple Lie groups
- complex simple Lie algebras
- connected and simply connected complex simple Lie groups.

In the body of this paper I assume connectedness as part of the definition of “simple Lie group”, but now I’m making it explicit.

For more on the exceptional Lie groups and their connection to octonions, see:

- John Frank Adams, *Lectures on Exceptional Lie Groups*, eds. Zafer Mahmud and Mamoru Mimura, U. Chicago Press, Chicago, 1996.
- John C. Baez, **The octonions**, *Bull. Amer. Math. Soc.* **39** (2002), 145–205. Errata in *Bull. Amer. Math. Soc.* **42** (2005), 213.
- Ichiro Yokota, **Exceptional Lie groups**.

For more on the ubiquitous appearance of Coxeter and Dynkin diagrams, and especially the ADE Dynkin diagrams, see:

- M. Hazewinkel, W. Hesselink, D. Siermsa, and F. D. Veldkamp, **The ubiquity of Coxeter–Dynkin diagrams (an introduction to the ADE problem)**, *Nieuw. Arch. Wisk.* **25** (1977), 257–307.
- John McKay, **A rapid introduction to ADE theory**.
- Vladimir I. Arnol’d, “Problems of Present Day Mathematics” in *Mathematical Developments Arising from Hilbert’s Problems*, ed. F. E. Browder, Proc. Symp. Pure Math. **28**, AMS, Providence, Rhode Island, 1976.

Arnol’d lists a lot of important math problems, following up on Hilbert’s famous turn-of-the-century listing of problems. Problem VIII in this book is the “ubiquity of ADE classifications”.

Witt showed that all the integral lattices generated by vectors v with $v \cdot v = 2$ come from simply laced Dynkin diagrams in this paper:

- Ernst Witt, Spiegelungsgruppen und Aufzählung halbeinfacher Liescher Ringe, *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* **14** (1941), 289–322.

There should be some more useful modern reference! For a proof of Gabriel's theorem that only simply laced Dynkin diagrams give quivers of tame representation type, see:

- William Crawley-Boevey, [Lectures on representations of quivers](#).

For more discussions of Gabriel's theorem, see:

- Harm Derksen and Jerzy Weyman, [Quiver representations](#), *Notices Amer. Math. Soc.* **52** (2005), 200–206.
- Idun Reiten, Dynkin diagrams and the representation theory of algebras, *Notices Amer. Math. Soc.* **44** (1977), 546–558.
- Alistair Savage, [Finite-dimensional algebras and quivers](#).

For the McKay correspondence, see:

- John McKay, Graphs, singularities and finite groups, in *Proc. Symp. Pure Math.* vol **37**, AMS, Providence, Rhode Island, 1980, pp. 183–186.
- D. Ford and John McKay, Representations and Coxeter graphs, in *The Geometric Vein Coxeter Festschrift* (1982), Springer, Berlin, pp. 549–554.
- Pavel Etinghof and Mikhail Khovanov, [Representations of tensor categories and Dynkin diagrams](#).
- Joris van Hoboken, [Platonic solids, binary polyhedral groups, Kleinian singularities and Lie algebras of type A,D,E](#), Master's Thesis, University of Amsterdam, 2002.
- Klaus Lamotke, *Regular Solids and Isolated Singularities*, Vieweg & Sohn, Braunschweig, 1986.
- Peter Slodowy, Platonic solids, Kleinian singularities, and Lie groups, in *Algebraic Geometry*, Springer, Berlin, 1983, pp. 102–138.

As you can see, some of these describe the connection between the ADE diagrams and singularity theory, which I did not touch on here.

The green images of Platonic solids were created by [Kjell André](#) and converted into SVG files by [DTR](#). They are available at Wikicommons under a CC BY-SA 3.0 license. The images of tilings and the fancy images of Dynkin diagrams were created by [R. A. Nonemacher](#) and are available at Wikicommons under a CC BY-SA 4.0 license. The other images were either created by myself, or I've been unable to track down their source.