Part III: Crystal Structures

What happens when we have a large (meaning, comparable to Avogadro's) number of atoms together? We already started getting some anwers to this in the previous section: in the ground-state, it is usually favorable for them to group into ordered structures, i.e. crystals. However, the ground-state is only fully occupied at $T = 0$. At finite temperatures, higher energy states are populated with finite probability $\propto \exp(-\beta E)$, which increases as the temperature increases.

As a function of temperature T, what typically happens if we have a large number of atoms is sketched below:

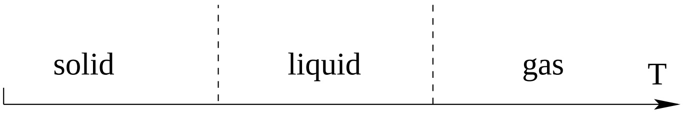

Figure 1: Rough phase diagram for a collection of many atoms, showing a gas phase at high-T, a liquid phase at intermediary T, and a solid phase at low-T.

At high temperatures, each atom has a large kinetic energy (remember that classically, we expect the average kinetic energy to be $3k_BT/2$). So they are whizzing by at fast speeds, and spend very little time in each other's vicinity. As a result, the weak interactions between atoms are nearly neglegible compared to these large kinetic energies, so the result is a gas phase. This is a very simple state: it is homogeneous (i.e., there is no difference between any two spatial points, in average things look precisely the same anywhere you look); it is isotropic (things look the same in any direction); it has very low density, since in average atoms are far from each other, etc. For an ideal gas, i.e., where we completely ignore interactions, we describe this equilibrium state by $pV = Nk_BT$ – this is a very simple system to understand, as you'll learn in a stat-mech course. If we treat the weak Lennard-Jones interactions perturbationally, then we get the Van der Waals equation of state. In any event, this state of matter is quite easy to understand, but it is not "condensed".

If we lower the temperature, the average kinetic energy becomes smaller and the contributions from the potential energies are no longer negligible. As a result, atoms/molecules prefer to stay fairly close to each other, to lower their potential energies, but at the same time are still mobile, moving around because of their still considerable kinetic energy. This is achieved in the liquid state, which is also homogeneous and isotropic, but much more dense than the gas phase, since atoms are much closer together. In most cases, a liquid is just slighly less dense than its corresponding solid (water/ice is one of the more famous exceptions to this rule).

Understanding the liquid state is a very difficult enterprise, because as I just said, average kinetic and potential energies are comparable. We understand what happens when one dominates: we get the gas phase if we can ignore the potential energy, vs. ordered crystal if we ignore the kinetic energy. These two states are very different from one another, but a liquid has some similarities with both (it is homogeneous like a gas, but dense like a solid) so it's hard to know where to start. There are macroscopic theories to describe a liquid (hydrodynamics, for example you should remember Bernoulli's equation from introductory physics), but connecting these to microscopic theories is a very difficult enterprise. We will completely ignore the liquid phase in this course.

Finally, if we cool T even more, there is a transition to a solid phase. In this limit, we can (mostly) ignore the motion of the atoms, so the way to minimize their total energy to get to the ground-state is to arrange them in ordered structures, or crystals, as we have just discussed. Note that for a crystal, the system is no longer homogeneous. The atoms are arranged in a certain pattern, so at some points in space we have atoms and at all other points there is nothing – no longer do things look the same at all points. The system is not isotropic, either: along some directions there are lines of atoms, along other ones there are none, etc. However, we still have a very ordered structure, which clearly still has quite a lot of symmetries, as we'll discuss in a bit. It is these discrete symmetries that will allow us to make quite a lot of progress in understanding the microscopic state of a crystal, and how it determines its macroscopic properties.

Before moving on, let me make two comments:

a) it is not always the case that upon cooling a liquid, one gets a perfect crystal. In fact, this is rather the exception, and only occurs if the cooling is very very slow, so that each atom has time to find its "proper" location before they freeze into place.

If the cooling is fast, one is usually left with an amorphous solid, or even a glass. The former you can think of as a state where locally, on very short scale, the atoms have arranged themselves like in a proper crystal – but different pieces don't quite fit together properly, for example, they're not oriented quite in the same direction, so there are lots and lots of defects in this system. Such disorder makes this quite hard to study and understand. A glass is, in fact, more like a very viscous liquid – atoms are still re-arranging themselves and trying to minimize their potential energies, but this happens on an extremely long timescale because they have very little kinetic energy. On a short time-scale it looks like everything is frozen, i.e. like a very disordered crystal. The glass in our windows is in such a state. Apparently, if you look in old churches/monasteries in Europe, which have some of the oldest glass panels ever made and not yet broken, you see the "pooling" of the glass at the bottom due to this very slow, very viscous flow. Glasses also have all sorts of fascinating properties but we will not deal with them, either. Instead, we will limit ourselves to perfectly ordered crystals, where we can take advantage of their discrete order and make some progress.

b) the phase diagram shown above is, of course, just a sketch. To be more precise, one should also monitor the pressure – you can think of what I showed above as what happens at ambient pressure, for example. If you change the pressure, the temperatures where one goes from gas-liquid and liquid-solid will change. In some cases, you can even skip the liquid altogether and go directly from solid-gas. I'm sure you'll study (or already studied) these sorts of things and all the various possible exceptions in a thermodynamics course.

For the rest of this course, we will be interested in seeing how much we can understand of the properties of a fully ordered crystal – I hope I convinced you that this is the condensed phase for which we have the best chance of success.

The first question to be answered is: what kind of crystals are possible, and how can we characterize them? To answer this, let me start with the simplest case, and then we'll briefly discuss the more complicated ones:

1 One-dimensional crystals

First, the disclaimer: you should know that purely 1D crystals do not exist in nature. There is a famous Mermin-Wagner theorem, which you will study in a more advanced cond-mat course, which proves that purely 1D and 2D crystals are unstable to disordering at any finite T.

However, this doesn't mean that there are no 1D-like ordered structures. They do exist, but are embedded in larger systems, for examples, as chains inside a 3D crystal. Weak interactions with the rest of the crystal are, in such cases, sufficient to stabilize the chain, however precisely because these interactions are quite weak, it is often a good approximation to ignore the rest of the crystal when studying the properties of the chain. We will also do this from now on, in such cases, and we will not worry as to what precisely makes the chain stable.

Case 1: suppose all atoms in the 1D chain are of the same type. Then, by symmetry, we would expect them to be at equal distances from one another so that each is in the same environment as any other one. The only numbers we need in order to know where all atoms are, are (i) **the lattice constant** a, defined as the distance between two consecutive atoms; and (ii) where is one of them located. Let me call this location $\vec{R}_0$.

If the chain is oriented along the x axis, as sketched below, from this information we can find the position of all the other atoms to be

$$\vec{R}_n = \vec{R}_0 + na\hat{x}$$

where $n = 0, \pm1, \pm2, \ldots$. This is the beauty of a crystal, we need very little information to know where everything is. By contrast, in a liquid or a gas, we need to individually measure the position of each atom; knowing where one is doesn't tell us much about where the other ones are.

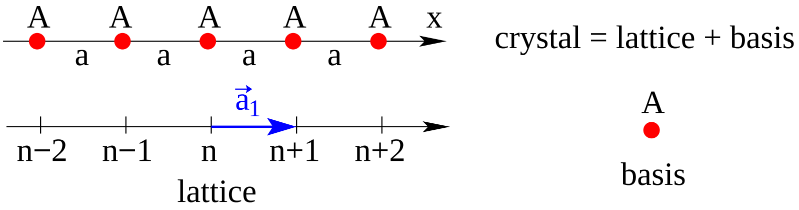

Figure 2: (top) 1D crystal made of identical atoms A; (bottom) corresponding lattice and the basis

To avoid worrying about the orientation of the chain, it is customary to define a **lattice vector**

$$\vec{a}_1 = a\hat{x}$$

(more generally, a vector whose length is equal to the lattice constant, and whose direction shows the orientation of the chain). Once we know $\vec{a}_1$ and the location of one atom $\vec{R}_0$, we know the location of all atoms: $\vec{R}_n = \vec{R}_0 + n\vec{a}_1$.

For reasons which will become clear soon, it turns out that it is convenient to think of the crystal as the sum of a **lattice** and a **basis**. The lattice is the collection of geometric points from which the crystal looks identical. For this simple example, the lattice is the collection of points where we place the atoms (see bottom of figure above), since sitting at any such point and looking around, we see precisely the same thing as if we sat at any other lattice point. In this case, the basis consists of a single atom A. You can see that by placing a basis at each lattice point, we regain the whole crystal.

If we had another crystal with B atoms instead of A, but somehow managed to keep the lattice constant the same, then the crystal made of B atoms would have the same lattice, but a different basis (the basis would now consist of a single B atom).

To better start getting a feel as to why we need to think in terms of bases and lattices, let's look at a more complicated model:

Case 2: Suppose the chain is made of alternating A and B atoms, for example as we would expect to get in a ionic crystal. Let's say the distance between consecutive A and B atoms is again a.

In this case, I cannot say that the lattice is just as above, i.e. the ordered collection of all points where we have atoms. This is because if I sit at an A site, I see quite a different view than if I sit at a B site, so these points are no longer equivalent and they can't all be in the lattice. The lattice only contains equivalent points!

So I have to make a choice: let's say I keep only the A sites to define the lattice (sitting at any A site, I see precisely the same thing as at any other A site). The lattice is then as sketched below, with a lattice constant, this time: $\vec{a}_1 = 2a\hat{x}$! Once we know $\vec{a}_1$ and the location of one lattice point $\vec{R}_0$, the other lattice points are at $\vec{R}_n = \vec{R}_0 + n\vec{a}_1$, just as before. In fact, apart from the specific value of $|\vec{a}_1|$, all the 1D lattices look the same.

What should the basis be? Clearly, it can't only have an A atom like before, since if I decorate the lattice with this basis, I only get the As but no Bs. To get the full crystal, the basis must contain one A and one B, at the proper distance a from each other, as shown below.

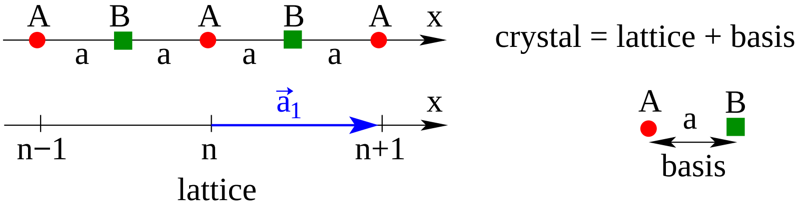

Figure 3: (top) 1D crystal made of alternating A and B atoms ; (bottom) corresponding lattice and the basis

Again, knowing the lattice (which is an ordered geometric structure) and the basis (which contains very few atoms) allows us to know where everything is.

Before moving on, you might ask me why not choose $\vec{a}_1 = 4a\vec{x}$, i.e. only keep every second A site in the lattice (that is a valid lattice, since those points are equivalent); and increase the basis to contain two A and two B atoms – this combo will give us the same crystal. There is nothing wrong with this choice, it is a correct choice, however as we'll see once we start doing some calculations, it is convenient to have to deal with the simplest basis possible (well, like any rule, this one also has some exceptions). Let's just say that it is usually convenient to make the simplest and most obvious choice – and if that is not convenient for calculations, we can always choose a lattice with fewer points and increase the basis accordingly. In practice, after you solve a few problems, you start to know intuitively which is the best choice for the problem at hand.

For your entertainment, consider how to find a lattice and a basis for case 3, sketched below. We'll practice a few more cases in class.

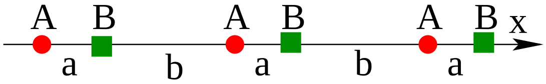

Figure 4: Another 1D crystal made of alternating A and B atoms

Before moving on, let's consider briefly the symmetries of these 1D crystals. By construction, each lattice (and therefore, its crystal as well) is **invariant to translations by any multiple of the lattice vector $\vec{a}_1$**. What this means is that if you draw the 1D crystal in some location; and then move everything either to the left, or to the right, by a multiple of $\vec{a}_1$, and draw it again, the final picture will be identical to the original picture. Note that this is not true for other translations, for example by $\vec{a}_1/2$ - in this case, for example 1, the new picture will have atoms half-way in between where the old one had them, so the two cases are distinguishable. For case 2, doing this would give us a picture with atoms at the same locations, but where we used to have As now we have Bs, and viceversa. Again, we can distinguish one from the other, so these are **not** symmetries of the problem. A symmetry must lead to a final state indistinguishable from the original one.

While all lattices are invariant to translations by their lattice vector (oftentimes, people use this property to define what the lattice is), sometimes they can have additional symmetries, for example to rotations. Look at Fig. 5. If I take case 1, and rotate everything by 180^o with respect to a transversal axis going through (any) one of the A atoms, the end result is identical to the original one – so this is a symmetry of the lattice. Such an axis is labelled by C_2, where the 2 shows that the rotation is by $360^o/2$. I could also place the axis of rotation half-way between two atoms, and this would again work, so that's another symmetry C_2. By contrast, in case 2, the first axis of rotation still works, but the one between atoms doesn't (in this case, after the rotation As are replaced by Bs, and viceversa). In case 3, there is no such rotational symmetry; you should convince yourself that no matter where you place an axis, after a 180^o rotation you get a different result.

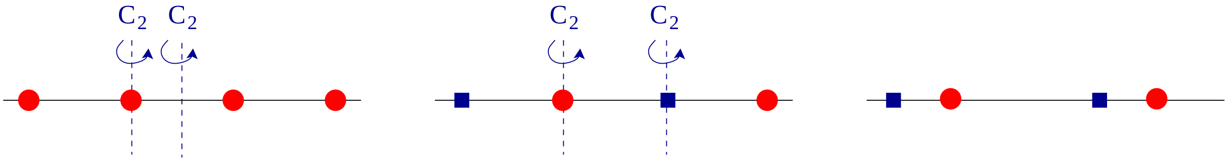

Figure 5: Valid choices for axes of rotation by 180^o for case 1. For case 2, there is no axis centered between atoms. For case 3, there is no symmetry to such rotations.

Unlike translational symmetries, which are always guaranteed, these other symmetries may or may not appear, depending on the particular crystal. In fact, as you'll learn when you'll take a more advanced course, it is then convenient to distinguish between additional symmetries of the lattice vs. symmetries of the crystal. For example, in all the above cases, the lattices themselves have 2 C_2 axes, just like in the first case (the lattices are all similar, the only difference is their lattice constant). It is the presence of the basis that spoils the symmetry to rotations in the latter cases, so one has to be a bit careful when going more deeply into this topic. For us, this will not really be necessary.

One final comment before we move on to higher dimensions.

As we've just seen, for a 1D crystal, once we're given the lattice vector $\vec{a}_1$, we know that properties of the crystal are the same if we translate by $n\vec{a}_1$ – for example, the density of electrons $n(x) = n(x + na_1)$, for any x and any integer n. This is the periodicity in real space, and what it means is that it suffice to figure out what happens inside one **unit cell**, which is any segment of length a_1. If we know the density, or any other property of interest, inside the unit cell, then by periodicity we know it everywhere in the crystal.

How one chooses the unit cell is a matter of convenience. For example, we could choose it as the segment between any two lattice points. Or, we could make a symmetric choice, and include in the unit cell all the points that are closest to one lattice point than any other one (see fig. below). This latter choice is more symmetric, and is known as a **Wigner-Seitz cell**. Other choices would be equally valid, but as usual, the more symmetric a choice is, the easier the resulting calculations.

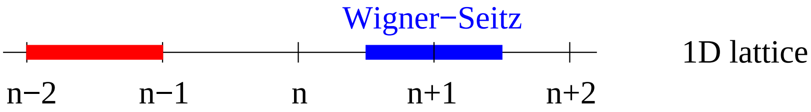

Figure 6: Two different choices for the unit cell of a 1D lattice. The one on the right is the Wigner-Seitz unit cell; it includes all points that are closest to one lattice point than to any other ones.

This should hopefully ring a bell. If you remember, when we discussed the tight-binding model for the 1D chain, we found something similar in the momentum space. To briefly rehash the arguments: because of the symmetry of the problem in real space, we decided that the eigenstates must be plane-wave like, and therefore they are indexed by a momentum k. However, we then saw that only momenta inside a finite interval are distinct, and everything else is just a periodic repetition. For example, the eigenenergies were $\epsilon(k) = -2t\cos(ka)$, periodic with a period $b = 2\pi/a$. This is why we decided that it's enough to know what happens in an interval of length b in the momentum space, which we called the Brillouin zone. We also said that we could choose the Brillouin zone in various ways, but the symmetric choice is the best.

This parallel is not accidental. In general, for any 1D lattice defined by a lattice vector $\vec{a}_1$ with the corresponding $\vec{R}_n = \vec{R}_0 + n\vec{a}_1$, and such that any spatial property is periodic $f(x) = f(x + na_1)$, we will find that the momentum space is also periodic, i.e. there exists a **reciprocal lattice vector** $\vec{b}_1$ and a result reciprocal lattice $\vec{G}_n = n\vec{b}_1$ so that any k-space property is also periodic, for example $\epsilon(k) = \epsilon(k + G_n)$. If you go back through the calculations, you should be able to convince yourselves that we must choose these vectors such that:

$$\vec{a}_1 \cdot \vec{b}_1 = 2\pi$$

Clearly they must be parallel since the momentum is along the axis of the chain, and the size of $\vec{b}_1$ is the correct one, at least for the simple example we considered where $a_1 = a$. From this point of view, the Brillouin zone is simply the k-space analog of the Wigner-Seitz unit cell in real space.

2 Two-dimensional crystals

As before, these will consist of a 2D lattice and a basis. The lattice is now defined by two non-collinear lattice vectors $\vec{a}_1$ and $\vec{a}_2$, so that if we know one point on the lattice, say $\vec{R}_{0,0}$, any other point is given by $\vec{R}_{n,m} = \vec{R}_{0,0} + n\vec{a}_1 + m\vec{a}_2$, where n, m are integers. As before, we must choose the lattice points as being those who are identical in the crystal structure. The basis, then, is the collection of a few atoms which, when placed at each lattice site, reproduces the crystal. A couple of examples are below.

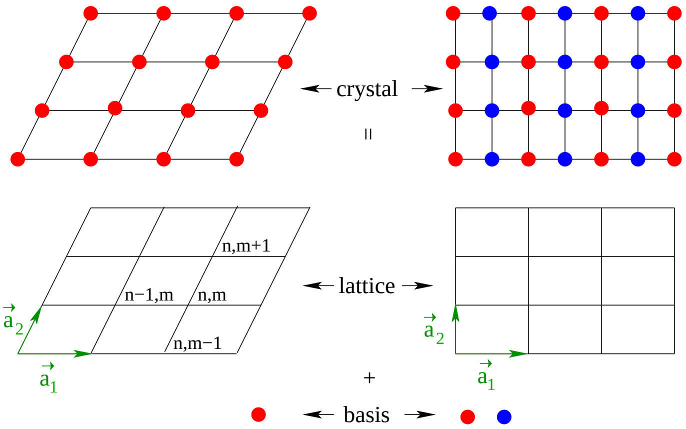

Figure 7: Two examples of 2D crystals (above) and their corresponding lattices, lattice vectors, and bases (below).

While $\vec{a}_1$ and $\vec{a}_2$ can be anything (so long as they are not parallel to each other), there are a few special choices which lead to lattices with their own special names. Let's list them:

2.1 Simple rectangular lattice

In this case, $\vec{a}_1$ and $\vec{a}_2$ are perpendicular to one another, but $|\vec{a}_1| \neq |\vec{a}_2|$. Below I sketched one such lattice, and two possible choices for its unit cell (the one on the right is the Wigner-Seitz choice).

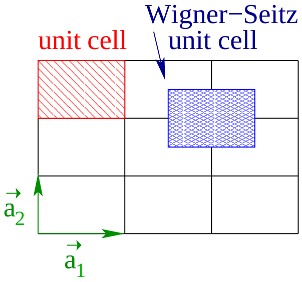

Figure 8: Simple rectangular lattice, its lattice vectors, and two possible choices of its unit cell.

2.2 Centered rectangular lattice

You can think of this as two simple rectangular lattices, interpenetrated. It is obtained by adding a lattice point in the center of each rectangle, see below:

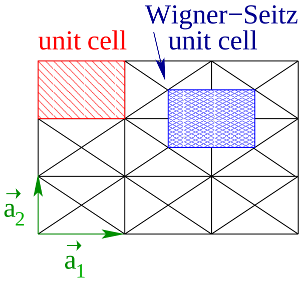

Figure 9: Centered rectangular lattice, its lattice vectors, and two possible choices of its unit cell.

Note that in this case, there are **two lattice points per unit cell**, whereas in all previous cases we had one lattice point per unit cell. By the way, unit cells with one lattice point per unit cell are called **primitive unit cells**, because they are the smallest possible. As we discussed in the 1D case, we can always make the unit cell bigger (i.e., not primitive), but that is usually not convenient.

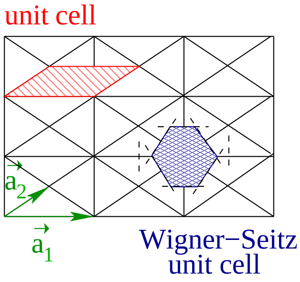

Figure 10: Same lattice as above, but we use its primitive unit cell instead of the more symmetric choice above.

The example here is one example where this is convenient, the reason being that the bigger unit cell is more symmetric. There is nothing to stop us from choosing the lattice vectors like above, corresponding to a primitive unit cell, but you have to admit that its unit cell is much less symmetric – and as we will see, this lack of symmetries will make calculations more difficult. So this is an example where we prefer to "artificially" make the unit cell larger than its absolute minimum, just to make it more symmetric.

2.3 Square lattice

As you would guess, in this case $\vec{a}_1 \perp \vec{a}_2$ and $|\vec{a}_1| = |\vec{a}_2|$, as shown below:

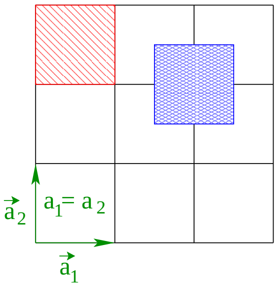

Figure 11: Square lattice, its lattice vectors, and two possible choices of its unit cell.

Question for you: why don't we define a centered square lattice, as well?

2.4 Triangular lattice

In this case $|\vec{a}_1| = |\vec{a}_2|$ and the angle between the two is 60^o (or 120^0, since here we can choose the lattice vectors in different ways, see below):

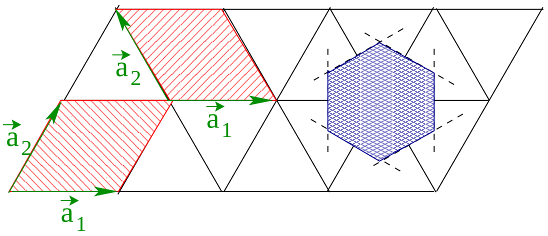

Figure 12: Triangular lattice, two possible choices for its lattice vectors, and the corresponding choices for its unit cell plus the Wigner-Seitz unit cell.

2.5 Oblique lattice

This is the general (and least symmetric) case where the two lattice vectors have different lengths and the angle between them is neither 90^o, like for square/rectangular lattices, nor 60^o or 120^o like for the triangular lattice. The example in Fig.10 shows the lattice vectors and unit cells in this case. As already discussed, when possible it is more convenient to instead double the unit cell and use a centered rectangular lattice, with its increased symmetry.

By construction, these lattices are invariant to translations by any lattice vector $n\vec{a}_1 + m\vec{a}_2$. Just like in the 1D case, that means that it's enough to know the spatial properties inside a unit cell, to know them everywhere. Also like in 1D, we expect eigenstates and eigenenergies to depend on a

2D momentum $\vec{k}$. The k-space should also be periodic, so that if we know what happens inside the Brillouin zone, we know everything.

Indeed, one now defines the reciprocal lattice vectors $\vec{b}_1, \vec{b}_2$ so that

$$\vec{a}_i \cdot \vec{b}_j = 2\pi\delta_{ij}$$

where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise. These vectors define the **reciprocal lattice** $\vec{G}_{n,m} = n\vec{b}_1 + m\vec{b}_2$, and the Brillouin zone is the Wigner-Seitz unit cell in this space. We will have some homework problems with some examples, and we'll also encounter this later on.

Before moving on to the basis and the crystal, let me mention that these 2D lattices also can have additional symmetries besides translations. We can again have rotations, for example we can find many C_2 axes for a simple rectangular lattice, as shown below. A square lattice has higher symmetry, and also has symmetries to rotations by 90^o (these are C_4 axes. Remember that in general, C_n is for rotation by $360^0/n$). Of course, a C_4 axis is, at the same time, also a C_2 axis – if the system is invariant to rotations by 90^o, then it is also invariant to two such rotations which give one by 180^o. One usually labels the axis of rotation by its largest n. For triangular lattices, one can find C_6 axes, also shown below.

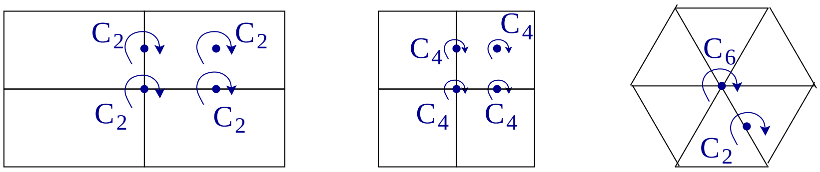

Figure 13: C_2 and C_4 rotation axes for rectangular and square lattices, respectively, and C_2 and C_6 axes for a triangular lattice.

There are other symmetries possible, besides translations and rotations. The examples below show mirror planes (left) and an inversion point (right). The former is self-explanatory, hopefully. The latter is a point with the property that if we use it as the reference point, what we see at any point $\vec{r}$ is identical to what we see at $-\vec{r}$. More symmetric lattices have more mirror planes and/or inversion points.

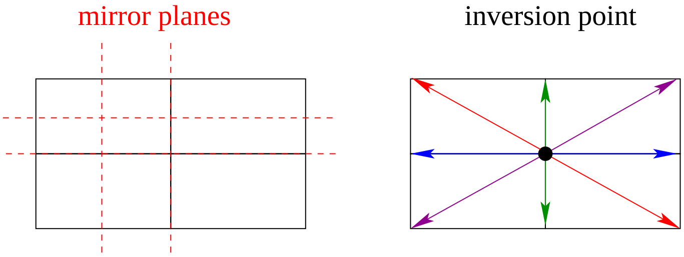

Figure 14: Example of mirror planes (left) – the two sides of the mirror are identical images of each other; and an inversion point (right) – each lattice point at $\vec{r}$ has a counterpart at $-\vec{r}$, as measured from the inversion point.

Before moving on, let me just mention that there is a branch of mathematics, called group theory, which formalizes all these symmetry operations and derives their consequences. This is very powerful stuff; basically, if a symmetry exists, it is usually bound to have certain consequences which must appear independent of any gorry details. For example, because of symmetry to rotations by 90^0, for a square lattice we expect that $\epsilon(k_x, k_y) = \epsilon(k_y, k_x)$. Different crystals with a square lattice may have very different eigenstates $\epsilon(k_x, k_y)$, but the symmetry enforces this equality no matter what. Understanding and taking advantage of the power of these symmetries is highly recommended. If you plan to study solid state in more detail, you'll definitely need to learn some group theory.

So far we have discussed all possible (distinct) 2D lattices. To get the full crystal, we also need a basis. In reality, of course, we know the crystal and we have to decide how to choose the basis and the lattice. Below are a couple of famous 2D examples from real life, which we will discuss together in class:

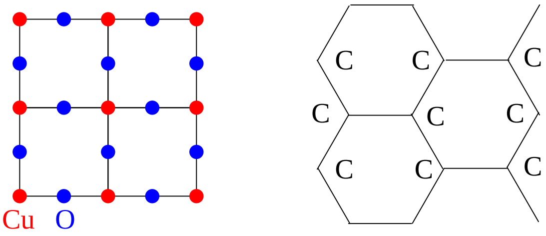

Figure 15: CuO_2 plane from high-temperature superconducting cuprates (left) and graphene (right). These are 2D crystals whose decomposition into a lattice +basis will be discussed in class. Based on what we discussed in the previous chapter, do you expect graphene to be a metal or an insulator?

3 3D crystals

Guess what: here we need three non-coplanar lattice vectors $\vec{a}_1, \vec{a}_2, \vec{a}_3$ to define the lattice by using $\vec{R}_{n_1,n_2,n_3} = \vec{R}_{0,0,0} + n_1\vec{a}_1 + n_2\vec{a}_2 + n_3\vec{a}_3$. The least symmetric case is when these vectors have different lengths and the angles they make with one another are neither 90^0 nor 60^0 (this would be the equivalent of the oblique lattice, in 2D). This is known as the **triclinic** lattice.

If some of the lattice vectors are equal to one another and/or are at nice angles with respect to one another, then we have more symmetric lattices. Like in 2D, there is a limited number of such options. Let me quickly list them (you will discuss them extensively in a more advanced course, for sure). For simplicity, let α_{ij} be the angle between vectors $\vec{a}_i$ and $\vec{a}_j$.

3D lattices:

- **triclinic**, if $a_1 \neq a_2 \neq a_3$, and all $\alpha \neq 90^o$.
- **monoclinic**, if $a_1 \neq a_2 \neq a_3$ and $\alpha_{12} = \alpha_{13} = 90^o$ while $\alpha_{23} \neq 90^o$.
- **orthorombic**, if $a_1 \neq a_2 \neq a_3$ and $\alpha_{12} = \alpha_{13} = \alpha_{23} = 90^o$. This is the analog of the rectangular 2D lattice. Like there, it also comes in several flavors: **simple, base-centered, body-centered and face-centered**. A unit cell for each one of these is sketched below. Only in the first case is the unit cell primitive, in the other cases it has more than one point per unit cell (figure out how many!).

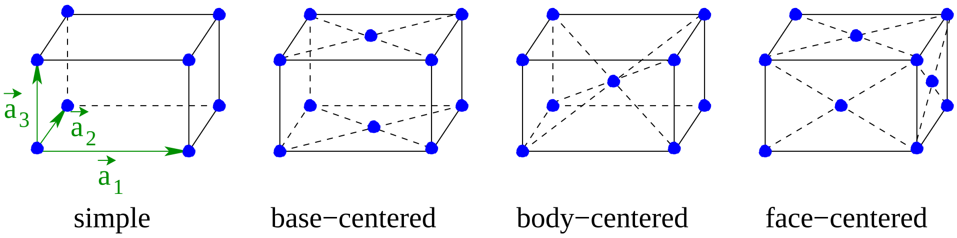

Figure 16: Orthorombic lattices: simple, base-centered, body-centered and face-centered.

• **tetragonal**, if $a_1 = a_2 \neq a_3$ and $\alpha_{12} = \alpha_{13} = \alpha_{23} = 90^o$. This only has the **simple** and the **body-centered** versions. Try to figure out why.

• **cubic**, if, as you would guess, $a_1 = a_2 = a_3$ and $\alpha_{12} = \alpha_{13} = \alpha_{23} = 90^o$. This can be **simple (abbreviation is SC from simple cubic), body-centered (BCC) or face-centered (FCC)**.

• **hexagonal**, if $a_1 = a_2 \neq a_3$ and $\alpha_{12} = 60^0, \alpha_{13} = \alpha_{23} = 90^o$. In other words, it has a triangular lattices in the basis plane, see below:

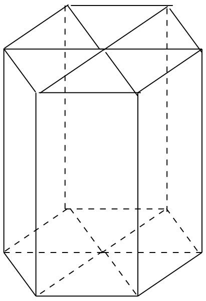

Figure 17: Hexagonal lattice.

• **trigonal**, if $a_1 = a_2 = a_3$ and $\alpha_{12} = \alpha_{13} = \alpha_{23} \neq 90^o$.

As you can easily imagine, there are a lot of additional rotational symmetries, mirror points and inversion points and even some other, more complicated operations, for the more symmetric lattices. These can all be investigated using group theory, and that also allows one to identify all the possible symmetries of the crystal (i.e., after you add a basis to the lattice). For your entertainment only, I will post notes prepared by N. Ingle a few years ago, which list all of these (the notes also have professional pictures, unlike my notes, so have a look there if you're not quite sure what I drew). I do not expect you to learn the additional stuff in those notes – again, if you will continue to persevere in this field, and in particular if you will need to deal with crystallography, you will certainly learn all this and a lot more. If not, I just want you to know that we can identify all possible structures and give them names, etc.

By analogy with the other cases, we can also define here three reciprocal lattice vectors $\vec{b}_1, \vec{b}_2, \vec{b}_3$, so that

$$\vec{a}_i \cdot \vec{b}_j = 2\pi\delta_{ij}$$

With them, we can construct the reciprocal lattice defined by $\vec{G}_{m_1,m_2,m_3} = m_1\vec{b}_1 + m_2\vec{b}_2 + m_3\vec{b}_3$, and the Brillouin zone as its unit cell. Any momentum-dependent property is uniquely known if we know its values inside the Brillouin zone.

The equations above are a bit cumbersome to solve by brute force. It turns out that its solutions are given by:

$$\vec{b}_1 = 2\pi \frac{\vec{a}_2 \times \vec{a}_3}{\vec{a}_1 \cdot (\vec{a}_2 \times \vec{a}_3)}$$

$\vec{b}_2$ and $\vec{b}_3$ solutions are obtain from cyclic permutations $1 \to 2 \to 3 \to 1$.

Here's an interesting fact, that you should be able to prove: the quantity in the denominator, i.e. $\vec{a}_1 \cdot (\vec{a}_2 \times \vec{a}_3) = V$, is the volume of the real space unit cell. Note that it doesn't matter how we choose the unit cell, whether using Wigner-Seitz or anything else, the volume is always the same (why?).

We can also calculate the volume of the reciprocal lattice unit cell, i.e. the volume of the Brillouin zone, and this will obviously equal $\vec{b}_1 \cdot (\vec{b}_2 \times \vec{b}_3) = \Omega$, since this lattice is defined by the b vectors. From the above link between a's and b's, you can show that $\Omega = (2\pi)^3/V$ (you might want to try this for an orthogonal lattice, maybe, although with some goodwill you can show this for the general case). In other words, the bigger the volume of the unit cell in real space, the smaller that of the Brillouin zone and viceversa. This holds not only in 3D, but in all dimensions. We already saw that in 1D $b_1 = 2\pi/a_1$. Similarly, in 2D the areas of the real space unit cell and the Brillouin zone are inversely proportional, the factor being $(2\pi)^2$.

Finally, to get a crystal we must also have a basis. Below are some examples of real materials which I have already mentioned in the previous section. I also drew the lattice vectors and indicated the basis nearby.

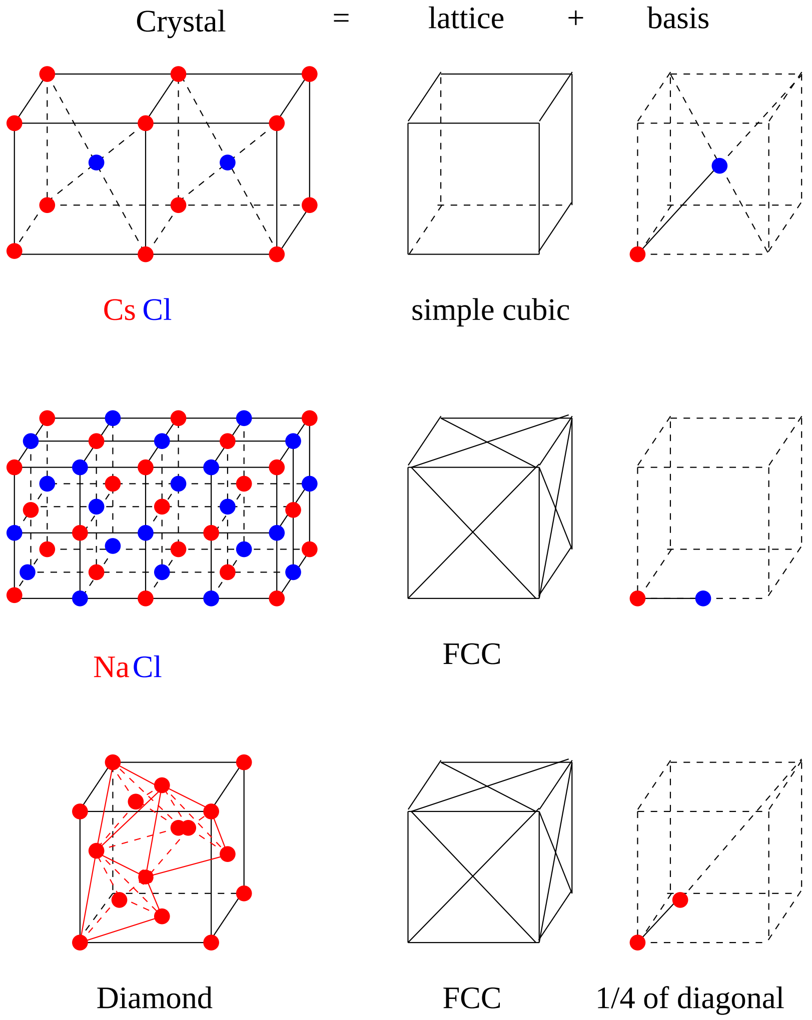

Figure 18: Crystal structures for (i) CsCl – a simple cubic lattice with a basis consisting of one Cs and one Cl. Other materials with similar structure are CsBr, CsI, TlCl, TlBr, TlI; (ii) salt, NaCl – it has an FCC lattice with a basis consisting of one Na and one Cl. Other materials with similar structure are LiF, LiCl, LiBr, LiI, NaF, NaCl, NaBr, NaI, etc etc.; (iii) C, diamond structure – an FCC lattice with a basis consisting of two C atoms. I marked with red lines some of the regular tetragons (not all since the figure would be too complicated. I didn't show all C atoms for the same reason). Many other compounds, like GaP, GaAs, GaSb, CuF, CuCl, CuBr, CuI, etc have similar structure, with the basis containing one of each. This is known as the zinc-blende structure, because ZnS also has it.

One thing you may notice here is that FCC appears very often. The reason for this is that it is a so-called "close-packed" structure. The way to think about this, is what would happen if we had identical hard sphere (as an approximation to our atoms with their electronic clouds) and we tried to order them so as they occupy as little as possible space, in other words we minimize the empty space in between the various spheres. In such a situation, they are brought as close together as possible, and therefore can lower their potential energy by interacting with as many neighbors as possible.

It turns out that FCC is one of these very efficient types of packing. Another one is the hexagonal close-packed structure, which consists of two interpenetrating simple hexagonal lattices. It turns out that there is a nice link between these two, and that one can generate other close-packed structures. Depending on how we're doing for time, we may discuss this further in class. Again, those who will take a more advanced course will certainly see this discussion later on. For those who don't, I hope this brief note has already given you enough of a flavor.

Ok, so now we have a better idea of how to classify crystals in terms of their lattice and their basis, that there are only so many possible types of lattices, and that using symmetries allows us to properly study and find all of them.

The next question is: how do we know what is the actual crystal structure of a material? When we discussed the various mechanisms for bonding, we saw how one could go about making a theoretical model and therefore a theoretical prediction for what structure one would expect. But how do we measure what is the actual structure? We need to do that to verify our models, to see if our understanding of what is going on is correct.

This is what we will discuss next.