5

NbTi: The Wonderful Workhorse

Niobium-titanium (NbTi) does not, at first sight, appear to be a particularly special superconductor. It does not have a high critical temperature and it cannot withstand super high magnetic fields. It does not have weird chemistry or exhibit an exotic flavour of superconductivity. Nevertheless, it is the single most important superconducting material we have. It is used in all of the MRI magnets worldwide. It is what is used in the magnets that bend the proton beams in the Large Hadron Collider (LHC) at CERN. In fact, the vast majority of commercial large scale applications of superconductors use NbTi. Even if ultra high fields are required, NbTi magnets are typically used to generate the background field and coils made from ‘better’ superconductors are placed inside to boost the field. So why is it so used so ubiquitously? It has one major selling point that the other technological superconductors can only dream of, and it has nothing to do with its superconducting properties. The material is *ductile*, which means that we can permanently change its shape without it breaking. This is super important because we need to be able to make many kilometres of the stuff as thin wires. We also need to be able to wind them into magnets, and the product needs to be mechanically robust. Amazingly, as we will discuss in Chapter 8, it is also possible to make wires out of brittle superconductors, but the processing is much more expensive and difficult. NbTi is nearly always the answer to the question: Which superconductor should I use? For this reason, it is often referred to as *the workhorse superconductor*, but as you will see in this chapter, it is by no means a boring material. In fact it is an exquisite example of top class materials science.

5.1 The Nb-Ti phase diagram

Before we get into the nitty-gritty of how to optimise the performance of NbTi wires, first we need to get to grips with the chemistry of the Nb-Ti alloy system. One of the most important tools that materials scientists rely on for understanding a chemical system is the phase diagram. The term *phase* relates to how the atoms or molecules are arranged within the material, and phase diagrams show what chemical phase (or phases) are most energetically stable under certain conditions. For simple systems like water, we saw in Section 3.2 that we generally only need to consider three different phases relating to the three different states of matter—solid, liquid and gas.[1] However, in general it is possible for a system in a given state to exist as several different phases.

---

[1] However, there are multiple different structures of ice, each of which is formally a different phase.

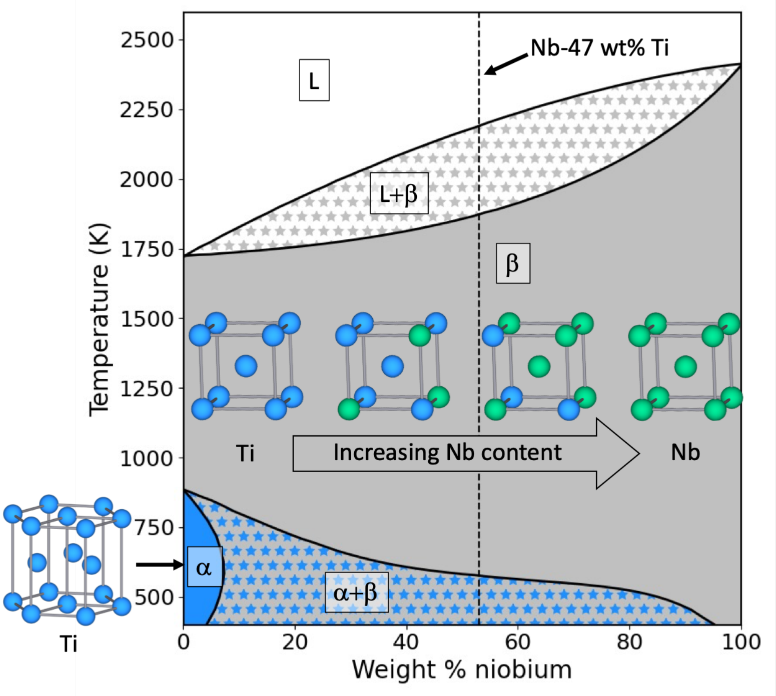

Fig. 5.1: Nb-Ti phase diagram. Insets show the crystal structures of the α-Ti and β-NbTi phases with Ti atoms shown in blue and Nb atoms shown in green.

For example, there is a whole zoo of different solid forms of carbon including graphite, diamond, buckyballs, nano-tubes and graphene. Each of these allotropes[2] is a different solid phase because the carbon atoms are arranged differently, resulting in them having different physical properties.

Water and carbon are both examples of *single component systems*. These are systems that consist of a single chemical substance. This substance could be a pure element made up of one type of atom (like the carbon case) or could be a compound consisting of small molecules (like the H_2O molecules in water), which do not break up into their constituent elements when melted or boiled. Phase diagrams get more complicated when we start looking at systems with multiple components, because each additional component adds an extra dimension to the phase diagram. Niobium-titanium is an example of a binary system because it has two components: niobium and titanium atoms. Fortunately, since we are only interested in standard room pressure (1 atmosphere (atm) = 100 kilopascals), we can draw a useful two-dimensional diagram with alloy composition on the x-axis and temperature on the y-axis, as shown in Fig. 5.1. At high temperature, the liquid phase (L) extends across the whole composition range.

[2] Allotropes are different solid forms of a pure element. More generally, for materials that are not chemical elements we use the term *polymorph* instead.

This is not surprising because we would expect to be able to melt alloys of any composition if we go to high enough temperature. The melting point of pure niobium is higher than that of pure titanium, so the lower edge of the liquid phase field is higher at the Nb-rich end of the diagram.

Now cast your eye downwards on the diagram, ignoring for the moment the region labelled 'L + β', until you get to the region labelled 'β'. Here the greek letter 'β' (beta) refers to one of the solid phases that is present in the Nb-Ti system.[3] Like the liquid phase, the β phase region extends across the whole width of the phase diagram—all the way from pure titanium to pure niobium. This essentially means that the atoms inside the material are arranged in the same way regardless of how many of them are Nb atoms and how many are Ti atoms. Materials with atoms arranged locally in a regular, repeating pattern are described as being *crystalline*. Nearly all materials are naturally crystalline, including metals and alloys like niobium-titanium. If the basic repeating unit of the pattern—its *unit cell*—is the shape of a cube, the material is said to have a cubic crystal structure. There are seven different crystal systems in total, each with a different unit cell shape.

It turns out that the β phase of NbTi has the *body-centred cubic (BCC)* crystal structure shown schematically in the inset of Fig. 5.1. At the titanium side of the phase diagram, you can think of the structure as consisting of a regular array of cubic unit cells with a Ti atom at each corner and right in the centre of the cube (at the position referred to as the body-centre). As we move across to the right on the phase diagram, we gradually replace some of the Ti atoms, at random, with Nb atoms. It is possible to carry on replacing Ti atoms with Nb atoms until there are no Ti atoms left (pure Nb) and the crystal still has the same body-centred cubic structure. The β phase is an example of something called a *substitutional solid solution* because we are essentially dissolving Nb in Ti by substituting Nb for Ti atoms (or vice versa). Because the β phase extends across the whole width of the phase diagram, Nb-Ti is an example of a system showing *complete solid solubility* at elevated temperatures. There is a set of empirical rules, known as the Hume-Rothery rules, that can help us to predict how soluble one element will be in another.[4] It turns out that elements tend to show good solid solubility over a wide composition range if, like titanium and niobium, the atoms have a similar size and similar chemistry (i.e. are close to each other on the periodic table).

In between the β phase region and the liquid phase region, there is a region labelled 'L + β' where both the liquid and β phases coexist under equilibrium conditions. You will probably be aware that a pure substance (like elemental Nb or Ti at the extreme edges of the phase diagram) melts at a single temperature—the melting point. However, impure substances melt over a range of temperatures, leading to the opening of this two phase region on the phase diagram. The upper boundary of the $L + \beta$ region is called the *liquidus* line and represents the temperature at which the liquid of a certain

---

[3] By convention we refer to different solid phases in a system by greek letters to distinguish them from chemical elements.

[4] William Hume-Rothery founded the Department of Metallurgy (now the Department of Materials) at the University of Oxford in the 1950s.

The Wider View

Crystal systems

There are seven different crystal systems, each with different symmetries. The unit cells are defined using three lengths called the lattice parameters (a, b and c) and three angles (α, β and γ). By convention, α is the angle between the sides b and c, β is the angle between sides c and a and γ is the angle between a and b.

Table 5.1: Unit cells of the seven crystal systems.

<table>
<tr>
<td colspan="2">Cubic
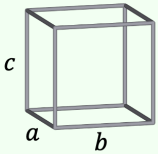
$a = b = c$
$\alpha = \beta = \gamma = 90^o$</td>
<td colspan="2">Tetragonal
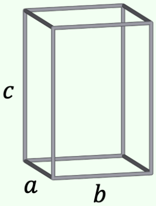
$a = b \neq c$
$\alpha = \beta = \gamma = 90^o$</td>
<td colspan="2">Orthorhombic
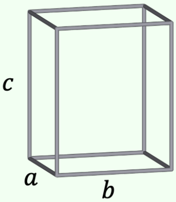
$a \neq b \neq c$
$\alpha = \beta = \gamma = 90^o$</td>
</tr>
<tr>
<td colspan="3">Rhombohedral
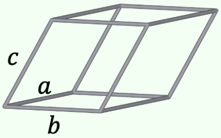
$a = b = c$
$\alpha = \beta = \gamma \neq 90^o$</td>
<td colspan="3">Hexagonal
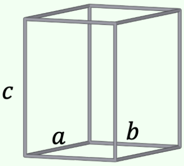
$a = b \neq c$
$\alpha = \beta = 90^o$, $\gamma = 120^o$</td>
</tr>
<tr>
<td colspan="3">Monoclinic
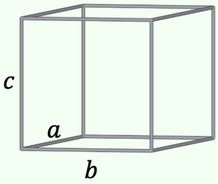
$a \neq b \neq c$
$\alpha = \beta = 90^o$, $\gamma \neq 90^o$</td>
<td colspan="3">Triclinic
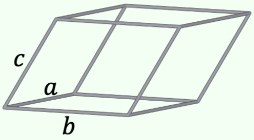
$a \neq b \neq c$
$\alpha \neq \beta \neq \gamma$</td>
</tr>
</table>

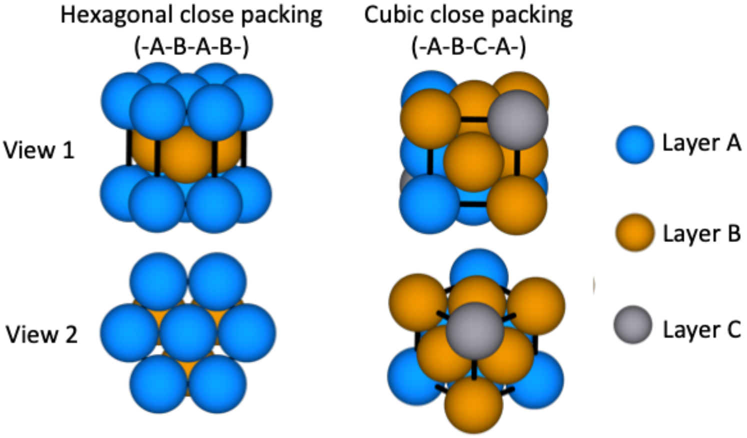

Fig. 5.2: Hexagonal close packed and cubic close packed structures.

alloy composition starts to solidify. The lower boundary is known as the *solidus* and gives the temperature at which the last remaining liquid solidifies.

Now see what happens if we continue to lower the temperature from the β phase region towards room temperature. At the titanium-rich side of the diagram we see a new solid phase labelled 'α' (alpha). In the α phase, the Ti atoms are arranged in a different crystal structure called *hexagonal close-packed* (*HCP*) . This structure can be described as stacks of two-dimensional layers with the atoms arranged in a hexagonal pattern within the layers. If you model the atoms as hard spheres like ping-pong balls, each atom touches six other atoms within the layer. There are two different ways in which these layers can be stacked efficiently on top of each other to produce close-packed structures that minimise empty space. Hexagonal close packing is one of these stacking arrangements, the other being *cubic close packing* (also known as *face-centred cubic* (*FCC*), as shown in Fig. 5.2. The percentage of the volume occupied by atoms as opposed to free space is higher in the close-packed structures than in the body-centred cubic structure of the NbTi β phase. It is, again, possible to replace some of the Ti atoms in the hexagonal close-packed α structure with Nb atoms, but the structure quickly becomes unstable—the close-packed structure simply cannot accommodate very many Nb atoms (which are slightly larger than Ti atoms). Instead, it becomes energetically favourable for the alloy to separate into a mixture of Ti-rich α phase and Nb-rich β phase. Under equilibrium conditions, the compositions of the two phases in the mixture at a certain temperature (i.e. the concentrations of Nb and Ti in each phase) are given by the lines that bound the $\alpha + \beta$ region on the phase diagram—the *solvus* lines. As we cool further, the solid solubilities of niobium in titanium and vice versa change a bit, but we would still expect these same two solid phases to coexist across a large range of alloy compositions.

Under the Lens

The Lever rule

The phase diagram can be used to figure out, not only which phases we would expect to find in a particular alloy composition and temperature, but also how much of each phase should be present under equilibrium conditions in a two phase mixture. The method is as follows, using the $\alpha + \beta$ phase region of NbTi as an example.

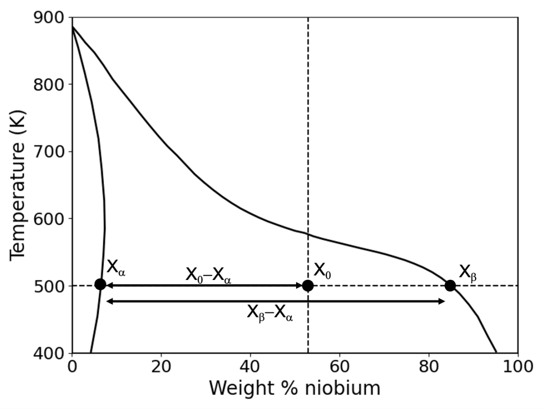

Fig. 5.3: Section of the Nb-Ti phase diagram.

- Draw a horizontal line across the $\alpha + \beta$ at the temperature of interest. This is called a *tie line*.

- Read off the compositions of the two phases: X_α and X_β. These are the positions where the tie line reaches the boundary with the single phase α and β regions either side, respectively. In this example X is given as a weight fraction of Nb, but the same procedure can be followed with atomic fractions if preferred.

- We can then write the overall alloy composition X_0 in terms of the fractions of each phase: f_α and f_β where $f_\alpha + f_\beta = 1$.

$$X_0 = f_\alpha X_\alpha + f_\beta X_\beta = f_\alpha X_\alpha + (1 - f_\alpha) X_\beta \tag{5.1}$$

- Rearranging the equation we can find expressions for the fractions of each phase.

$$f_\alpha = \frac{X_\beta - X_0}{X_\beta - X_\alpha} f_\beta = \frac{X_0 - X_\alpha}{X_\beta - X_\alpha} \tag{5.2}$$

This equation is known as the *Lever rule*.

If we take the standard NbTi alloy with 47 weight percent Ti (53 weight percent Nb), we can use the Lever rule to estimate the fraction of α and β phases at a chosen temperature. Let's take 500°C because we can easily read the compositions off the phase diagram: $X_\alpha \approx 0.06$ and $X_\beta \approx 0.85$.

$$f_\beta = \frac{X_0 - X_\alpha}{X_\beta - X_\alpha} = \frac{0.53 - 0.06}{0.85 - 0.06} = \frac{0.47}{0.79} \approx 0.6 \tag{5.3}$$

So under equilibrium conditions at 500°C, we would expect our alloy to contain 60% of the β phase by weight. If we want to know what the volume fraction of each phase is, we would need to know the density of each phase. In practice, the fraction of the superconducting β phase will tend to be higher than this because it is difficult to reach equilibrium.

5.2 Choosing the alloy composition

The beta phase (β) is the superconducting phase we refer to when we talk of niobium-titanium. The alpha phase (α) also superconducts, but because it has a critical temperature below 4 K it is non-superconducting when operating at liquid helium temperature. From the phase diagram, we have seen that at low temperatures (anywhere below about 600°C), we would expect most compositions of NbTi alloy to contain two phases: superconducting β and non-superconducting α. This may at first appear to be a disadvantage for making superconducting wires because the regions of non-superconducting α phase will reduce the cross-section of the wire that can carry superconducting current. On the contrary though, it actually turns out to be very fortunate because the α particles can act as excellent flux pinning centres. The key is to tailor the way we process the material to control how much α phase we end up with and how it is distributed microscopically throughout the alloy.

Consider an alloy with about 50% Nb and 50% Ti. If we heat it for long enough at 1500°C the atoms will move about by a process called *solid state diffusion* until the composition is uniform throughout—it will be single phase β NbTi. Now, if we cool it really slowly to a temperature in the two phase $\alpha + \beta$ region, some of the β phase will transform to α phase. We call the α phase a *precipitate* and the $\beta \rightarrow \alpha$ phase transformation is called *precipitation*. The precipitation process occurs by a two-stage mechanism called *nucleation and growth*. In the nucleation stage, Ti atoms which are constantly jiggling around, sometimes come together and form a cluster. If the cluster is very small, it is most likely that it will simply disappear again. However, if by chance the cluster manages to get large enough we refer to it as a *nucleus* and it becomes energetically more favourable for it to grow in size rather than to shrink. In the subsequent growth stage, Ti atoms diffuse and attach to these α nuclei until the equilibrium phase fractions and compositions are reached. As we will see in Chapter 8 the solid state diffusion process required for the Ti and Nb atoms to redistribute themselves can happen much faster at higher temperatures. This means that if we cool the alloy from the β phase region quickly the atoms do not have time to get to their equilibrium positions. In fact, if we cool fast enough, we can suppress the α

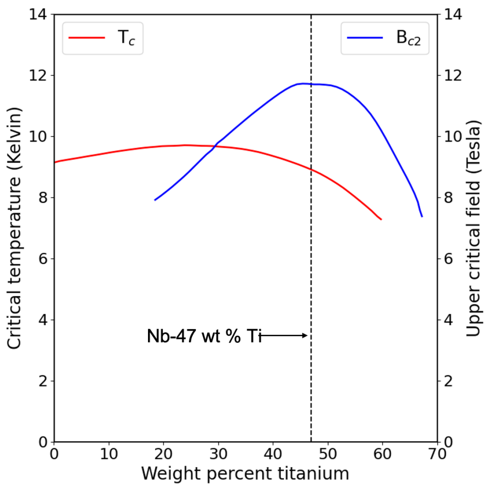

Fig. 5.4: Effect of composition on critical temperature and critical field in Nb-Ti.

precipitation process altogether and keep the whole material in the superconducting β phase. We call this a *supersaturated solid solution* because it is a solid solution that has more of the impurity component (in this case Ti) dissolved in it than is possible according to the equilibrium phase diagram. Supersaturated solid solutions could lower their free energy by separating into two phases, but the atoms are not mobile enough to redistribute themselves.

This means that we can actually make β NbTi alloys with the body-centred cubic crystal structure over a large range of chemical compositions, all the way from pure Nb to alloys with as much as 70% Ti by weight, that do not, in practice, spontaneously decompose into the two separate solid phases at room temperature. This allows us to investigate how the superconducting properties of β NbTi vary with Ti content (see Fig. 5.4). As you can see, the superconducting transition temperature, T_c, does not change very much when we increase titanium content—at least until we near the Ti-rich compositions. In contrast, the upper critical field (B_{c2}) is quite strongly influenced by composition, with a maximum value of around 12 tesla at around 50% by weight of Ti. It is for this reason that the standard alloy composition used for NbTi superconducting wires has 47% titanium by weight (Nb-47 wt % Ti).

Under the Lens

Nucleation theory

Many phase transformations happen by a process called *nucleation and growth*. The nucleation stage involves atoms clustering together. The simplest situation to think about is the formation of a solid phase particle from a liquid. Atoms are in constant motion and they will periodically come together to form a small cluster of solid. If that cluster gets large enough, it will tend to grow further, but if it is too small it will most likely shrink again. This behaviour is governed by the interplay between two different energy terms: volume free energy and surface energy. The first of these is the difference in Gibbs free energy between the solid state and the liquid state and provides the driving force for solidification. Below the melting point, the solid state is more stable, and so the free energy of the solid will be lower than the free energy of the liquid. The amount of free energy saved by a particle solidifying scales with its volume. However, we also have to take into account the fact that the boundary between the solid and the liquid will have an energy cost associated with it that scales with surface area of the particle. The origin of this surface energy can be understood by thinking about the surface as a region where the bonding is not ideal—there are broken bonds. When the particle is very small, the surface area to volume ratio is high and so it costs energy overall and the solid particle will be unstable, but as the particle grows, the volume free energy term starts to dominate, and the particle will become a stable nucleus.

Assuming the growing solid particle is spherical in shape, the total Gibbs free energy change associated with converting a particle of radius r from liquid to solid is given by

$$\Delta G = -\frac{4}{3}\pi r^3 \Delta G_v + 4\pi r^2 \gamma \tag{5.4}$$

where $\frac{4}{3}\pi r^3$ is the volume and $4\pi r^2$ is the surface area of the spherical particle, ΔG_v is the free energy saved by solidification per unit volume and γ is the surface energy per unit area. The first term has a negative sign because it is the energy saving associated with the particle solidifying, whereas the second term is positive because it is the energy cost associated with the surface. This is shown graphically in Fig. 5.5.

The particle will become stable when growing larger starts to decrease its total energy—that is, as soon as it passes the peak of the energy curve. Our standard mathematical method for finding a peak—a turning point—is to find where the first derivative (gradient) of the curve becomes zero. Here that means that we need to differentiate the expression for ΔG as a function of r and set it equal to zero.

$$\frac{d\Delta G}{dr} = -4\pi r^2 \Delta G_v + 8\pi r\gamma = 0$$

$$4\pi r^2 \Delta G_v = 8\pi r\gamma$$

$$r = \frac{2\gamma}{\Delta G_v} \tag{5.5}$$

We call this value the *critical radius* and give it the symbol r^*. By substituting r^* back into equation 5.4, the activation barrier for nucleation (ΔG^*) is found to be

$$\Delta G^* = \frac{16\pi\gamma^3}{3\Delta G_v^2} \tag{5.6}$$

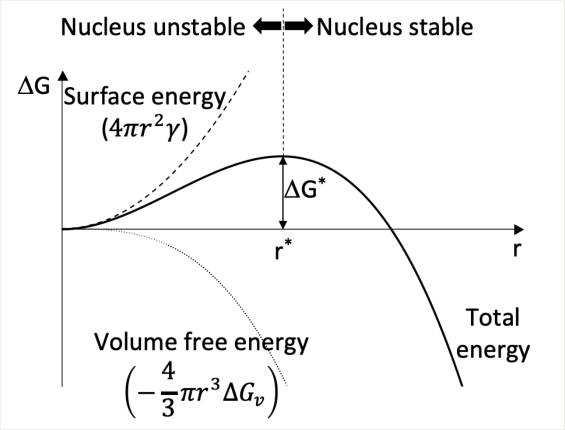

Fig. 5.5: Free energy as a function of particle size.

The value of ΔG_v is proportional to the latent heat of the material (how much heat is given out when it solidifies) and how far the liquid has been cooled below the melting point—the *undercooling*. The bigger the latent heat and the undercooling, the larger ΔG_v will be. This translates to a smaller critical radius (r^*) and a smaller energy barrier for nucleation (ΔG^*), indicating that the nucleation process is easier. In practice, the barrier to nucleation of a spherical particle within the bulk of a liquid is quite high, so nucleation tends to happen at surfaces such as the solid wall of a mould because this effectively reduces the energy barrier for nucleation and speeds up the process.

5.3 Mechanical properties of metals

Niobium-titanium is the workhorse superconductor, not because it has amazing superconducting properties, but because it has much better mechanical properties than the other technological superconductors. So it is worth looking at what gives metals their excellent mechanical properties and allows them to be easily made into wires. When a rod of metal is stretched by pulling the ends in opposite directions, the initial response is elastic. This means that if we stop pulling on it—remove the tensile force we are applying—it will return to its original length. What is actually happening is that the bonds between the metal atoms are being stretched. The stronger the bonds, the more force we need to stretch them by a certain amount, and the stiffer the metal rod will be. In this elastic regime, if the force is removed the metal rod will go back to its original length. It essentially behaves in the same way as a spring and follows Hooke's law: the extension of the rod (how much longer it gets) is proportional to the force that is applied. The constant of proportionality is called the spring constant and it relates to how stiff the sample (or spring) is. We can write Hooke's law as $F = kx$ where F is the force, x is the extension and k is the spring constant.

Ideally, materials scientists like to define things like stiffness as *materials properties*—quantities that do not change with the dimensions of the object that is being tested. An example of a materials property that we have already encountered is *resistivity* which depends on the material and temperature but (unlike resistance) is not influenced by the size and shape of the wire. Intuitively we know that the thicker the diameter of a metal rod, the more force is needed to extend it by the same amount. Essentially this is because the force we apply gets distributed over more bonds and so the force we are applying per bond has decreased. If you model the original rod as a single spring, doubling the cross-sectional area of the rod is like putting another identical spring in parallel with the first and pulling on both together, as shown in Fig. 5.6(a). The force pulling each spring is half of the original amount. Therefore, it

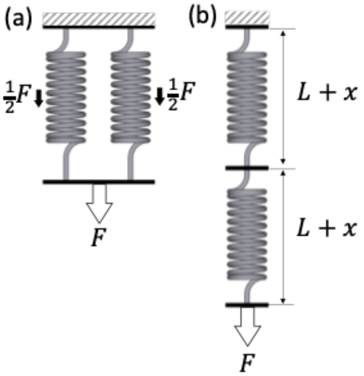

Fig. 5.6: Spring model for the elastic properties of a material. (a) Two springs connected in parallel are equivalent to doubling the cross-sectional area of the sample. (b) Two springs connected in series are equivalent to doubling the sample length.

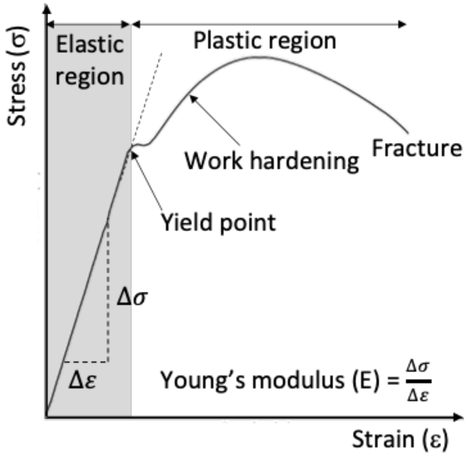

Fig. 5.7: Typical stress–strain curve of a ductile metallic material.

is useful to introduce a quantity called *stress* (σ), which is defined as the force per unit cross-sectional area (A) of the sample: $\sigma = \frac{F}{A}$. We also have to consider the extension term in a similar way. If we double the length of the rod, we would expect the extension to double. This is because each bond along the length of the rod is being stretched by the same amount. The longer the rod, the more bonds there are along its length and so the greater the total extension will be. This is simply like connecting two identical springs end-to-end in series (Fig. 5.6(b)). When we apply a force, F, we would expect each spring to extend by a distance x and so the total extension will be $2x$. So, instead of using the total extension of the rod, we use a quantity called *strain* (ε) which expresses the change in length of the rod x as a fraction of its original length L: $\varepsilon = \frac{x}{L}$.[5] We can then rewrite Hooke's law in terms of stress and strain: $\sigma = E\varepsilon$, Here, we have defined a new quantity E, which we call the *Young's modulus*. Just like resistivity, it is a materials property that does not depend on the dimensions of the sample and is a measure of the material's stiffness.[6] Figure 5.7 shows a typical stress–strain curve for a metal. The elastic region is where Hooke's law is obeyed and the stress increases linearly with strain. The gradient of the stress–strain curve in the elastic region gives the Young's modulus of the material.

So what happens if we keep pulling on the metal rod? Eventually the material will *yield* and something called *plastic deformation* starts to occur. When this happens it gets easier to extend the rod and the stress–strain curve flattens out. Inside the material, instead of the bonds just carrying on stretching, a different mechanism called *slip* is activated. Crystal defects called *dislocations* start to move under the influence of the applied stress. Applying a tensile force along the rod produces a shear (sliding) force

---

[5] Formally this definition of strain is known as the *engineering strain.*

[6] It is simple to show that the spring constant, k, is related to Young's modulus, E, by the equation $k = \frac{EA}{L}$.

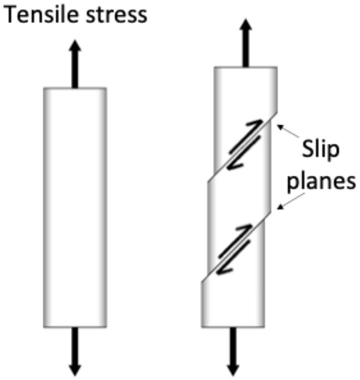

Fig. 5.8: Slip along planes tilted at 45° to the applied tensile stress.

along the diagonal direction (45° to the axis of the rod), as shown in Fig. 5.8. This has the effect of the diagonal crystal planes trying to slide over one another. If the crystal was perfect (no dislocation defects) it would take an enormous shear stress to shift the planes over one another in one go, but dislocations provide a much easier way for this shear to take place. A good way of thinking about how they work is by taking the example of trying to move large carpet across the floor. If we just hold it at one edge and give it a tug, we need to pull quite hard to get it to move because of friction. An easier way would be to make a ruck at one edge of the carpet and then push it across the carpet to the other side. The force we have to apply to move the ruck is much lower than the force we would need to move the whole carpet in one go, but we would only have succeeded in moving it by a short distance. We would need to repeat the process over and over again, moving the carpet a little bit sideways each time.

This is exactly what happens with dislocations. They are linear crystal defects—like the ruck in the carpet—and each time one moves across from one side of the sample to the other, the crystal planes shift by a small distance. The offset produced by each dislocation is called the *Burgers vector* and it depends on what kind of crystal structure it has and what kind of dislocation it is, but it is of the order of the spacing between atoms. Millions of dislocations need to move through the material to get appreciable extensions to our sample, so if we looked inside a material that has been plastically deformed, using a powerful electron microscope, we would see a high density of these linear defects. Interestingly, once we have enough shear stress to get the dislocations to start moving, they will carry on moving if we keep the stress roughly constant, which is why the graph flattens off. If we remove the applied stress, the dislocations do not move back, so the extension produced by slip is permanent (unlike the reversible behaviour of the elastic region). The sample is said to have undergone plastic deformation. Materials in which it is easy for dislocations to move are said to be *ductile*, whereas those in which dislocation motion is difficult will be *brittle*. Metals tend to be ductile because metallic bonding is not directional and it is relatively easy for the atoms to move their positions slightly as a dislocation passes through. In contrast, it is much more difficult to generate and move dislocations in ceramic materials with ionic or covalent bonding.

Yield—the onset of plastic deformation—is not the end of the story. When we have millions of dislocations all moving through the crystal on different slip planes, they end up getting tangled up with each other. This stops them from moving and there becomes a backlog of dislocations piling up. To free them up and enable them to move again, we typically need to increase the stress that we apply. This is because each dislocation will have a preferred crystal plane that it travels along, and we need to get it to switch to another (less favourable) plane to get past the obstacle. The material therefore gets harder—the gradient of the stress–strain curve increases again. We call this *work-hardening* because it happens as a result of mechanically 'working' the material. Eventually the dislocations are so tangled up inside the material that fracture occurs instead (the bonds break).

5.4 Thermomechanical processing of NbTi

Now we will turn our attention to how we optimise the critical current density in a niobium-titanium alloy. We know from the discussion in Section 4.1 that in type II superconductors like NbTi we need to pack the material full of structural defects to be able to get high currents without generating resistance. In fact, if we had perfect single phase NbTi, it would carry almost no resistance-less current because there would be nothing to stop flux lines moving and dissipating energy. It turns out that there is a really neat way of doing this in NbTi, and the phase diagram gives us the clue. We know that one of the ways to introduce pinning centres in superconductor is to introduce non-superconducting secondary phase particles. As luck would have it, a β NbTi alloy can actually lower its free energy by forming some particles of non-superconducting α Ti. By controlling this precipitation process we can obtain fantastic flux pinning without having to add in anything else. In principle this can be done by taking the supersaturated β alloy (the one with too much Ti dissolved in it) and heating it up a bit so that the atoms can move (diffuse) fast enough for the α phase to nucleate and grow. By tweaking the exact temperature and length of time of this heat treatment (called an *anneal*) we can control the number and size of the α precipitates to optimise pinning.

In practice, in the NbTi system it is actually quite difficult to persuade the α phase to precipitate even though it is thermodynamically favourable for it to do so. One effective way we can speed up the process is by putting in some mechanical deformation. Because the work-hardened material contains a high density of dislocations which raise its free energy, there is a big driving force for it to regrow as defect-free material if we heat it up enough to allow atoms to diffuse. This process is known as *recrystallisation* because new crystals nucleate and grow to consume the defective material. As this recrystallisation process happens, α phase precipitation is facilitated, so by mechanically working (physically deforming) and then annealing (heating up) the alloy we can increase the fraction of non-superconducting α phase.

This really plays into our hands because we want to manufacture thin wires of NbTi. This involves starting with a large billet of the material (a big cylindrical piece) and

The Wider View

Precipitation hardening

Solid state precipitation is used in structural alloys known as precipitation hardened alloys. The archetypal alloy of this type is Al–4 wt% Cu. Aluminium is a useful metal for automotive and aerospace applications because it has a low density, good ductility and good corrosion resistance. However, pure Al is far too soft to be useful. By adding a few weight percent of Cu and carrying out a heat treatment similar to the one described for NbTi, we can form secondary phase precipitates that obstruct the passage of dislocations. Because dislocations motion is impeded, the material work-hardens more effectively, enhancing its mechanical strength. In this case, we need to control the heat treatment to optimise the size and spacing of the precipitates, as summarised in Fig. 5.9(a). In the early stages of the precipitation anneal, metastable phases form first. These are not actually the thermodynamically most stable phases but they form very quickly because their crystal lattice closely matches the host crystal—they are said to be *coherent* and they have a much lower interface energy than incoherent precipitates. However these fine, coherent precipitates can easily be sliced through by dislocations because the crystal structure is not very distorted and so they are not very good at strengthening the material. If, on the other hand, we let the precipitates grow too big, they become less coherent with the host lattice. That makes it more difficult for dislocations to cut through them. Unfortunately, as they grow, the precipitates get further apart and the dislocations can then bow around them instead (Fig. 5.9(b)). For this reason we need to strike a happy medium to optimise the hardness of the alloy.

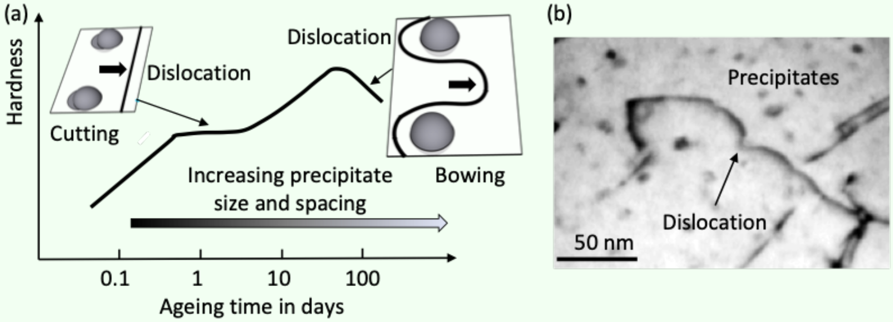

Fig. 5.9: (a) Optimising strength of Al-Cu alloy by ageing. (b) TEM of dislocations in steel bowing around copper precipitates (courtesy of S. Lozano-Perez, University of Oxford).

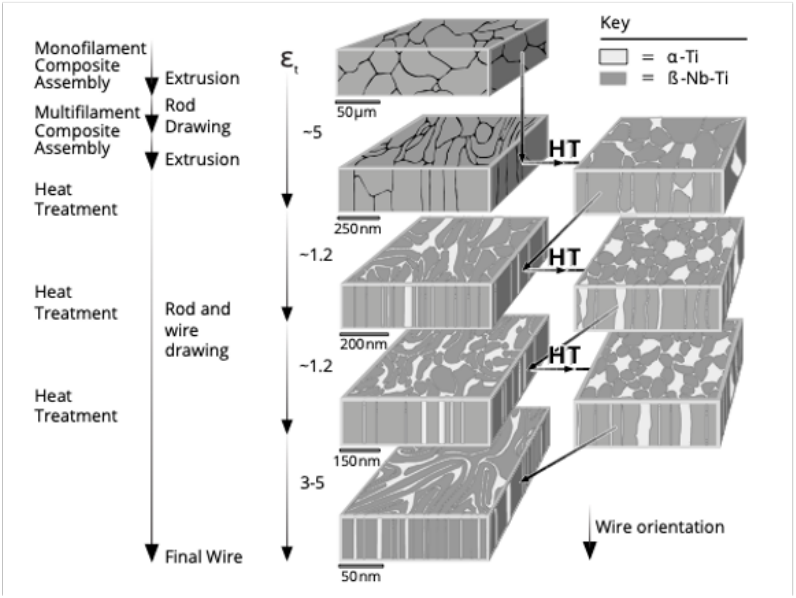

Fig. 5.10: Thermomechanical processing of NbTi. (Courtesy of P. Lee, NHMFL.)

extruding it or drawing it through a series of increasingly narrow nozzles called dies. Extrusion is a process just like pushing toothpaste out of a tube, and it is typically used in the early stages of the process. Wire-drawing is used later in the process and involves grabbing hold of the end of the wire and pulling it through the die under tension. To get a thin wire we cannot just keep pulling it through narrower and narrower dies because the metal would become so work-hardened that it would break. Instead, we draw it down a bit, putting in some mechanical damage, and then we do a heat treatment to recover the ductility and encourage some α phase to precipitate. Then we draw it some more, and heat it again. This sequential process is repeated until the desired final wire diameter is obtained. Figure 5.10 shows a schematic diagram of the microstructure that is obtained at different stages of the process. As you can see, not only do you end up increasing the total volume of α phase, but the process also has the effect of decreasing the size of the α precipitates by a factor of about 1000. (Note that the scale bar starts off at 50 micrometres and finishes off at 50 nanometres.) The final precipitate size can be controlled by selecting the original billet dimensions, final wire diameter and heat treatment parameters to end up with the desired microstructure. This is what materials science is all about—using our understanding of materials systems to design materials (and strategies for processing them) to optimise their microstructure and hence their performance.

5.5 Large Hadron Collider wires

So far, we have talked about how we can engineer a two phase microstructure in NbTi to optimise flux pinning by exploiting the thermomechanical process that we use for making thin wires. But for real applications, we cannot use single strands of NbTi wire. We actually need many fine filaments of superconductor, all embedded in plenty of copper to keep the currents (and magnetic fields they produce) as stable as possible. Mechanical vibrations (for example) can lead to flux lines jumping from one pinning site to another and dissipating some heat. If we cannot get that heat out of the place where it has been generated, an avalanche effect happens. Hot spots generated by these flux jumps create small regions that stop superconducting so that the supercurrent has to divert around them. Effectively the cross-sectional area of superconductor for the current to flow through is reduced so the current gets concentrated in the vicinity of a hot spot. If the current locally exceeds the critical current density, the rest of the wire will lose superconductivity and a huge amount of energy suddenly has to go somewhere. The copper surrounding the filaments of superconductor is essential for extracting the heat generated and stopping these quenches. The copper also produces a good electrical shunt for current to pass through if the superconductor does become resistive. In the magnets used to bend and focus the proton beam in the Large Hadron Collider (LHC), the 1 mm diameter wire strands contain no fewer than 6300 separate filaments of NbTi ($\sim$ 0.006 mm in diameter), each surrounded by a layer of copper as shown in Fig. 5.11.

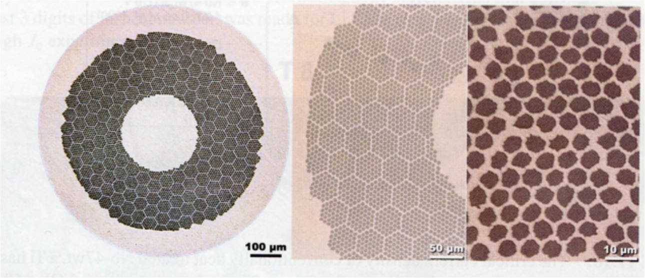

Fig. 5.11: LHC wire cross section. Reproduced from (Rogalla and Kes, 2012) with permission of Taylor and Francis Group, LLC, a division of Informa plc.

So how on earth do we go about making wires many kilometres long with thousands of really tiny but continuous filaments all embedded in copper? Luckily, copper and NbTi have rather similar mechanical properties. To make the wire, the starting material is a cylindrical billet made by assembling a rod of NbTi inside a Cu tube, sometimes with an extra thin layer of Nb in between. When this composite billet is extruded, the separate components extrude together, so the final object is just a narrower diameter version of the original—NbTi surrounded by a layer of Nb and then a layer of Cu (Fig. 5.12). This co-extrusion method is also used for getting stripy toothpaste and it is how sweets like seaside rock are made, with letters or patterns that stretch throughout the whole length. After initial extrusion that reduces its diameter, the superconducting

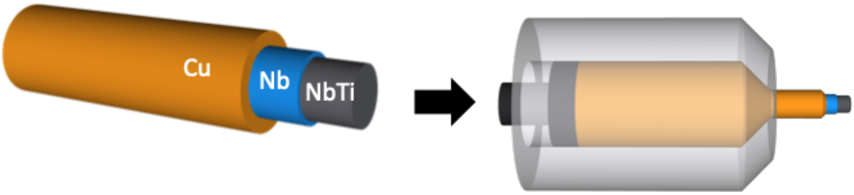

Fig. 5.12: Initial stages of processing NbTi wire.

wire is passed through a hexagonal die. The resulting hexagonal cross-section rods can then be packed together in a honeycomb arrangement, placed inside another copper tube, and the extrusion process can be repeated to create a wire with a bundle of filaments neatly arranged inside. These multifilamentary wires are then again passed through a hexagonal die, stacked together in the same way as before, and extruded again to produce a wire with a hierarchy of hexagonal structures inside. This is then drawn down to the final diameter required (Fig. 5.13).

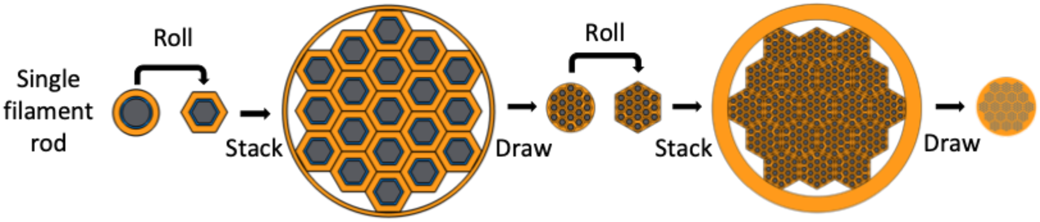

Fig. 5.13: Multifilamentary wire processing method.

This still is not the end of the story. We also often need to twist each strand and weave multiple strands together to make a cable. This is essentially to minimise the electromagnetic cross-talk between the filaments that can introduce unwanted losses. Rutherford cable is a particular type of cable that is widely used in accelerator magnets, where the magnetic field quality needs to be excellent. The Rutherford cable in the Large Hadron Collider magnets (shown in Fig. 5.14) contains 36 strands, and each strand contains 6300 NbTi filaments. In total the LHC uses 7500 km of Rutherford cable, which translates to over 250,000 km of strand (wire). If we lined up all of the individual NbTi filaments, they would stretch for about 1.5 billion kilometres—10 times the distance from the earth to the sun!

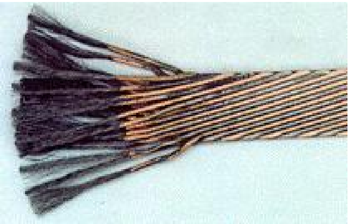

Fig. 5.14: Rutherford cable. (Courtesy of CERN).

Chapter summary

- NbTi is called the workhorse superconductor because it is by far the most commonly used superconducting material in large scale applications, but the materials science that has gone into optimising the wires is fascinating. Not only does it need to be made in complex multifilamentary wire geometries to get the required stability, but every strand has to be engineered on the nanoscale to optimise its current carrying capacity at the operating fields.
- The precise geometry and architecture of the wire—number and diameter of filaments, how much copper is used, size and distribution of the α Ti precipitates—is tailored to the specific application.
- Large Hadron Collider wires contain far more filaments than the wires used for MRI magnets and their microstructures are optimised for operation at lower temperatures and higher magnetic fields.
- The Nb-Ti phase diagram has been examined in some detail because phase diagrams are key tools that materials scientists use for designing alloys and processing strategies. The idea of controlling precipitation of secondary phases in alloys is a major theme in materials science.