8.1 Introduction and synopsis

It is easy to set a value on the engineering science that enables success, that makes things happen, but much harder to value engineering science that prevents failure, that *stops* things happening. One of the great triumphs of recent engineering science has been the development from the 1960s onward of a rigorous mechanics of material fracture. We have no numbers for the money and lives it has saved by preventing failures; all we know is that, by any measure, it is enormous. This chapter is about the ways in which materials fail when loaded progressively, and design methods to ensure that fracture won't happen unless you want it to.

Sometimes, of course, you do. Aircraft engines are attached to the wing by shear-bolts, designed to fail and shed the engine if it suddenly seizes. At a more familiar level, peel-top cans, seals on food containers, and many other safety devices rely on controlled tearing or fracture. And processes like machining and cutting use a combination of plasticity and fracture.

We start by distinguishing *strength* from *toughness*. Toughness — resistance to fracture — requires a new material property to describe it: the *fracture toughness*, developed in Section 8.3. Data for this new property is explored in Section 8.4 using charts like those we have already seen for modulus and strength. The underlying science mechanisms (Section 8.5) give insight into ways in which toughness can be manipulated (Section 8.6). The chapter ends in the usual way with a Summary, suggestions for Further Reading, and Exercises.

8.2 Strength and toughness

Strength and toughness? Why both? What's the difference? *Strength*, when speaking of a material, is short-hand for its resistance to plastic flow. Think of a sample loaded in tension. Increase the stress until dislocations sweep right across the section, meaning the sample just yields, and you measure the initial *yield strength*. Strength generally increases with plastic strain because of work hardening, reaching a maximum at the *tensile strength*. The area under the whole stress—strain curve up to fracture is the *work of fracture*. We've been here already — it was the subject of Chapter 5.

Toughness is the resistance of a material to the propagation of a *crack*. Suppose that the sample of material contained a small, sharp crack, as in Figure 8.1(a). The crack reduces the cross-section A and, since stress σ is F/A, it increases the stress. But suppose the crack is small, hardly reducing the section, and the sample is loaded as before. A tough material will yield, work harden, and absorb energy as before — the crack makes no significant difference. But if the material is *not* tough (defined in a moment) then the unexpected happens; the crack suddenly propagates and the sample fractures at a stress that can be far below the yield strength. Design based on yield is common practice. The possibility of fracture at stresses below the yield strength is really bad news. And it has happened, on spectacular scales, causing boilers to burst, bridges to collapse, ships to break in half, pipelines to split, and aircraft to crash. We get to that in Chapter 10.

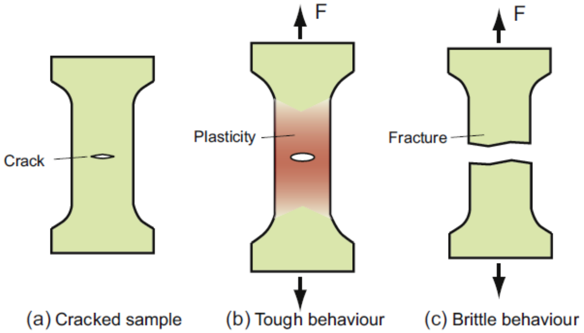

(a) Cracked sample (b) Tough behaviour (c) Brittle behaviour

Figure 8.1 Tough and brittle behaviour. The crack in the tough material, shown at (b), does not propagate when the sample is loaded; that in the brittle material, at (c), propagates without general plasticity, and thus at a stress less than the yield strength.

So what is the material property that measures the resistance to the propagation of a crack? And just how concerned should you be if you read in the paper that cracks have been detected in the track of the railway on which you commute, or in the pressure vessels of the nuclear reactor of the power station a few miles away? If the materials are tough enough, you can sleep in peace. But what is 'tough enough'?

This difference in material behaviour, once pointed out, is only too familiar. Buy a CD, a pack of transparent folders, or even a toothbrush: all come in perfect transparent packaging. Try to get them out by pulling and you have a problem: the packaging is strong. But nick it with a knife or a key or your teeth and suddenly it tears easily. That's why the makers of shampoo sachets do the nick for you. What they forget is that the polymer of the sachet becomes tougher when wet, and that soapy fingers can't transmit much force. But they had the right idea.

Tests for toughness If you were asked to devise a test to characterise toughness, you might dream up something like those of Figure 8.2: notch the material, then yank it or whack it until it breaks, measuring the energy to do so. Tests like these (there are many variants) in fact are used for ranking and as an acceptance procedure when taking delivery of a new batch of material. The problem is that they do not measure a true material property, meaning one that is independent of the size and shape of the test sample, so the energy measurements do not help with design. To get at the real, underlying, material properties we need the ideas of *stress intensity* and *fracture toughness*.

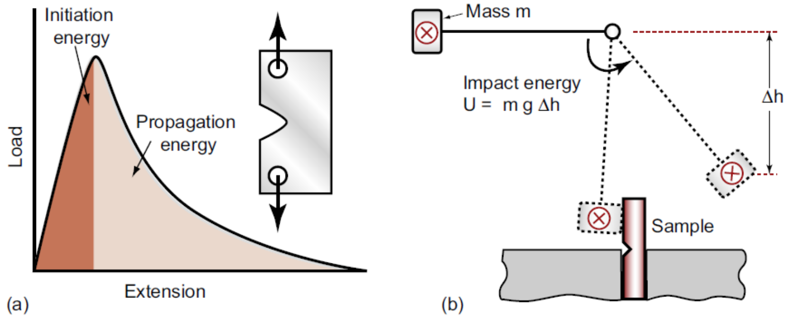

Figure 8.2 (a) The tear test, and (b) the impact test. Both are used as acceptance tests and for quality control, but neither measures a true material property.

8.3 The mechanics of fracture

Stress intensity K_1 and fracture toughness K_{1c} Cracks and notches concentrate stress. For notches we defined (in Chapter 7) a 'stress concentration factor', which tells us how much greater the peak local stress is compared to the remote stress, for changes in cross-section with a well-defined size and radius, such as a circular hole. Now we consider how cracks affect the stress field. Figure 8.3 shows a remote stress σ applied to a cracked material. We can envisage 'lines of force' that are uniformly spaced in the remote region but bunched together around the crack, giving a stress that rises steeply as the crack tip is approached. Referring to equation (7.12), we can see that the stress concentration factor for a notch does not help here — it is

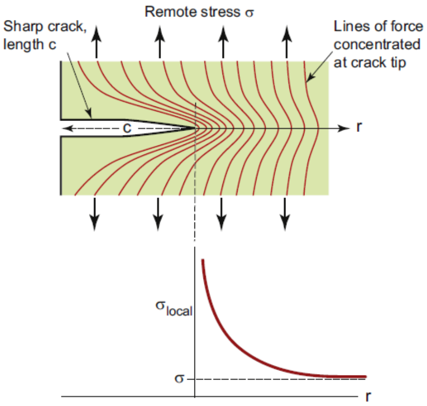

Figure 8.3 Lines of force in a cracked body under load; the local stress is proportional to the number of lines per unit length, increasing steeply as the crack tip is approached.

relevant only for features with a finite radius of curvature. Cracks are sharp: the radius at the tip is essentially zero, giving (notionally) an infinite stress by equation (7.12). A different approach is needed for cracks. Analysis of the elastic stress field ahead of a sharp crack of length c shows that the local stress at a distance r from its tip, caused by a remote uniform tensile stress σ, is

$$\sigma_{\text{local}} = \sigma\left(1 + Y\sqrt{\frac{\pi c}{2\pi r}}\right) \tag{8.1}$$

where Y is a constant with a value near unity that depends weakly on the geometry of the cracked body. Far from the crack, where $r \gg c$, the local stress falls to the value σ; near the tip, where $r \ll c$, it rises steeply (as shown in Figure 8.3) as

$$\sigma_{\text{local}} = Y\frac{\sigma\sqrt{\pi c}}{\sqrt{2\pi r}} \tag{8.2}$$

So for any given value of r the local stress scales as $\sigma\sqrt{\pi c}$, which therefore is a measure of the 'intensity' of the local stress field (the inclusion of the π is a convention used universally). This quantity is called the *mode 1 stress intensity factor* (the 'mode 1' means tensile loading perpendicular to the crack), and given the symbol K_1 (with units of MPa.m$^{1/2}$):

$$K_1 = Y\sigma\sqrt{\pi c} \tag{8.3}$$

The stress intensity factor is thus a measure of the elastic stress field near the tip of a sharp crack, equation (8.2). Note that as $r \to 0$, this equation predicts an infinite stress. In practice the material will yield in a small contained zone at the crack tip (or some other localised form of damage will occur, such as micro-cracking); this is discussed further later. The important point for now is that the loading on the crack tip region that drives potential failure is an elastic stress field that scales with K_1. As a result, for reasons again explored later, cracks propagate when the stress intensity factor exceeds a critical value. This critical value is called the *fracture toughness*, K_{1c}.

Figure 8.4 shows two sample geometries used to measure K_{1c} (there are others, described in Chapter 10). A sample containing a sharp crack of length c (if a surface crack) or $2c$ (if a contained crack) is loaded, recording the tensile stress σ^* at which it suddenly propagates. It is essential that the crack be sharp – not an easy thing to achieve – because if it is not, the part of the stress field with the highest stresses, where the lines of force in Figure 8.3 are closest together, is changed. There are ways of making sharp cracks for doing this test: growing a crack by cyclic loading is one, a process called *fatigue* that is described in the next chapter. Indeed, fracture mechanics based on the stress intensity factor is particularly relevant to design against fracture due to cracks that developed by fatigue. The value of Y for the centre-cracked plate in Figure 8.4 is 1, provided $c \ll w$; for other geometries, there is a small correction factor. Then the quantity K_{1c} is given by

$$K_{1c} = K_1 = Y\sigma^*\sqrt{\pi c} \approx \sigma^*\sqrt{\pi c} \tag{8.4}$$

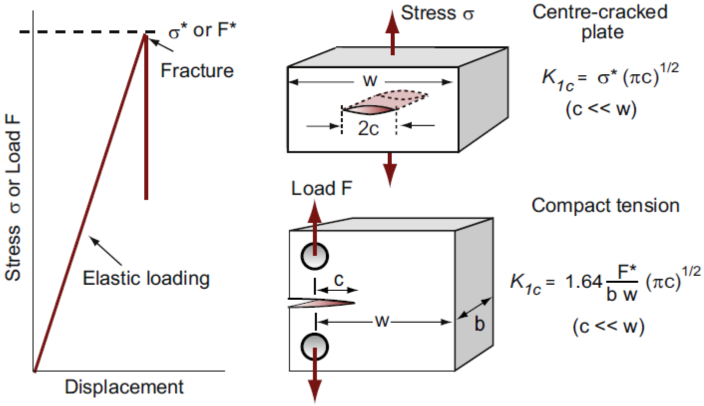

Figure 8.4 Measuring fracture toughness, K_{1c}. Two test configurations are shown here; others are described in Chapter 10.

Fracture toughness is a *material property*, and this means two things. The first is that its value is independent of the way it is measured: tests using different geometries, if properly conducted, give the same value of K_{1c} for any given material (Figure 8.4). The second is that it can be used for design, in ways described in Chapter 10. For now, we note two important ways that we could use equation (8.4): to find the failure stress, if we know that there is a crack of a given size present; and the reverse, to find the 'critical crack length', being the maximum crack size that we can tolerate without failure, for a given stress.

Energy release rate G and toughness G_c When a sample fractures, a new surface is created. Surfaces have energy, the *surface energy* γ, with units of joules[1] per square metre (typically $\gamma = 1$ J/m^2). If you fracture a sample across a cross-section area A you make an area $2A$ of new surface, requiring an energy of at least $2A\gamma$ joules to do so. Consider first the question of the *necessary condition for fracture*. It is that sufficient external work be done, or elastic energy released, to at least supply the surface energy, γ per unit area, of the two new surfaces that are created. We write this as

$$G \geq 2\gamma \tag{8.5}$$

where G is called the *energy release rate*. In practice, it takes much more energy than 2γ because of plastic deformation round the crack tip. But the argument still holds: growing a crack costs energy G_c J/m^2 for the two surfaces — a sort of 'effective' surface energy, replacing 2γ. It is called, confusingly, the *toughness* (or the *critical strain energy release rate*). This toughness G_c is related to the *fracture toughness* K_{1c} in the following way.

Think of a slab of material of unit thickness carrying a stress σ. The elastic energy stored in it (Chapter 4) is

$$U_v = \frac{\sigma^2}{2E} \tag{8.6}$$

---

[1] James Joule (1818–89), English physicist, did not work on fracture or on surfaces, but his demonstration of the equivalence of heat and mechanical work linked his name to the unit of energy.

per unit volume. Now put in a crack of length c, as in Figure 8.5. The crack relaxes the stress in a half-cylinder of radius about c — the reddish half-cylinder in the figure — releasing the energy it contained:

$$U(c) = \frac{\sigma^2}{2E} \times \frac{1}{2}\pi c^2 \tag{8.7}$$

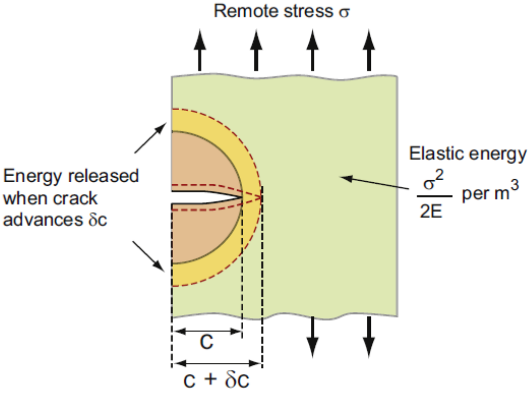

Figure 8.5 The release of elastic energy when a crack extends.

Suppose now that the crack extends by δc, releasing the elastic energy in the yellow segment. This energy must pay for the extra surface created, and the cost is $G_c\,\delta c$. Thus differentiating the last equation, the condition for fracture becomes

$$\delta U = \frac{\sigma^2 \pi c}{2E}\delta c = G_c \delta c \tag{8.8}$$

But $\sigma^2 \pi c$ is just K_{1c}^2, so from equation (8.4), taking $Y = 1$,

$$\frac{K_{1c}^2}{2E} = G_c \tag{8.9}$$

This derivation is an approximate one. A more rigorous (but much more complicated) one shows that the form of equation (8.9) is right, but that it is too small by exactly a factor of 2. Thus, correctly, the result we want (taking the square root) is:

$$K_{1c} = \sqrt{EG_c} \tag{8.10}$$

Toughness G_c is also therefore a material property, and is perhaps more obviously related to the physics of fracture: a tough material requires more energy to be dissipated when a crack propagates. But quantifying the energy release rate G for a loaded crack is difficult, so for design it is preferable to use $K = K_{1c}$ as the failure criterion. The elastic stress field is something we will be analysing anyway, and the crack length is a physical quantity that can be directly observed and measured. The physical equivalence of the two criteria, expressed by equation (8.10), is straightforward: crack growth is driven by the release of elastic stored energy (G), and this scales with the crack tip elastic stress field (K).

The crack tip plastic zone The intense stress field at the tip of a crack generates a *process zone*: a plastic zone in ductile solids, a zone of micro-cracking in ceramics, a zone of delamination, debonding, and fibre pull-out in composites. Within the process zone, work is done against plastic and frictional forces; it is this that accounts for the difference between the measured fracture energy G_c and the true surface energy 2γ. We can estimate the size of a plastic zone that forms at the crack tip as follows. The stress rises as $1/\sqrt{r}$ as the crack tip is approached (equation (8.2)). At the point where it reaches the yield strength σ_y the material yields (Figure 8.6) and – except for some work hardening – the stress cannot climb higher than this. The distance from the crack tip where $\sigma_{local} = \sigma_y$ is found by setting equation (8.2) equal to σ_y and solving for r. But the truncated part of the elastic stress field is redistributed, making the plastic zone larger. The analysis of this is complicated but the outcome is simple: the radius r_y of the plastic zone, allowing for stress redistribution, is twice the value found from equation (8.2), giving

$$r_y = 2\left(\frac{\sigma^2 \pi c}{2\pi\sigma_y^2}\right) = \frac{K_1^2}{\pi\sigma_y^2} \tag{8.11}$$

(taking $Y = 1$). Note that the size of the zone shrinks rapidly as σ_y increases: cracks in soft metals have large plastic zones; those in ceramics and glasses have small zones or none at all.

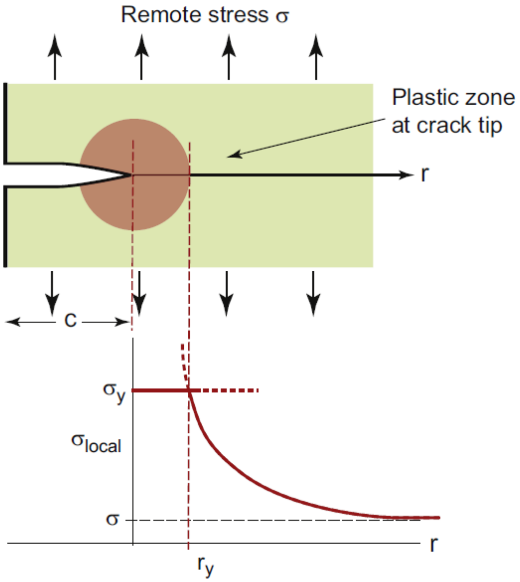

Figure 8.6 A plastic zone forms at the crack tip where the stress would otherwise exceed the yield strength σ_y.

The property K_{1c} has well-defined values for brittle materials and for those in which the plastic zone is small compared to all dimensions of the test sample so that most of the sample is elastic. When this is not so, a more complex characterisation is needed. In very ductile materials the plastic zone size exceeds the width of the sample; then the crack does not propagate at all—the sample simply yields.

When cracks are small, materials yield before they fracture; when they are large, the opposite is true. But what is ‘small’? Figure 8.7 shows how the tensile failure stress varies with crack size. When the crack is small, this stress is equal to the yield stress; when large, it falls off according to equation (8.4), which we write (taking $Y = 1$ again) as

$$\sigma_f = \frac{K_{1c}}{\sqrt{\pi c}} \tag{8.12}$$

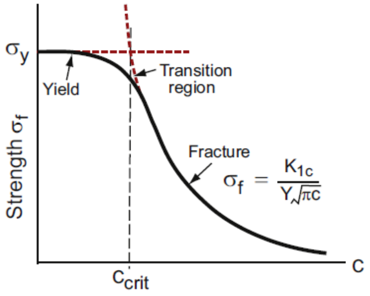

Figure 8.7 The transition from yield to fracture at the critical crack length c_{crit}.

The transition from yield to fracture is smooth, as shown in the figure, but occurs around the intersection of the two curves, when $\sigma_f = \sigma_y$, giving the *transition crack length*

$$c_{crit} = \frac{K_{1c}^2}{\pi \sigma_y^2} \tag{8.13}$$

Note that this is the same as the plastic zone size at fracture (equation (8.11)), when $K_1 = K_{1c}$.

Table 8.1 lists the range of values of c_{crit} for the main material classes. These crack lengths are a measure of the *damage tolerance* of the material. Tough metals are able to contain large cracks but still yield in a predictable, ductile, manner. Ceramics (which always contain small cracks) fail in a brittle way at stresses far below their yield strengths. Glass can be used structurally, but requires careful treatment to prevent surface flaws developing. Polymers are perceived as tough, due to their resistance to impact when they are not cracked. But the table shows that defects less than 1 mm can be sufficient to cause some polymers to fail in a brittle manner.

Table 8.1 Approximate crack lengths for transition between yield and fracture

Material class	Transition crack length, c_{crit} (mm)
Metals	1–1000
Polymers	0.1–10
Ceramics	0.01–0.1
Composites	0.1–10

8.4 Material property charts for toughness

The fracture toughness–modulus chart The fracture toughness K_{1c} is plotted against modulus E in Figure 8.8. The range of K_{1c} is large: from less than 0.01 to over 100 MPa.m$^{1/2}$. At the lower end of this range are brittle materials, which, when loaded, remain elastic until they fracture. For these, linear elastic fracture mechanics works well, and the fracture toughness itself is a well-defined property. At the upper end lie the super-tough materials, all of which show substantial plasticity before they break. For these the values of K_{1c} are approximate but still helpful in providing a ranking of materials. The figure shows one reason for the dominance of metals in engineering; they almost all have values of K_{1c} above 15 MPa.m$^{1/2}$, a value often quoted as a minimum for conventional design.

The log scales of Figure 8.8 allow us to plot contours of *toughness*, G_c, the apparent fracture surface energy (since $G_c \approx K_{1c}^2/E$). The diagonal broken lines on the chart show that the values of the toughness start at 10^{-3} kJ/m^2 (about equal to the surface energy, γ) and range through almost five decades to over 100 kJ/m^2. On this scale, ceramics (10^{-3}–10^{-1} kJ/m^2) are much lower than polymers (10^{-1}–10 kJ/m^2); this is part of the reason polymers are more widely used in engineering than ceramics, a point we return to in Chapter 10.

The fracture toughness–strength chart Strength-limited design relies on the component yielding before it fractures. This involves a comparison between strength and toughness – Figure 8.9 shows them on a property chart. Metals are both strong and tough – that is why they have become the workhorse materials of mechanical and structural engineering.

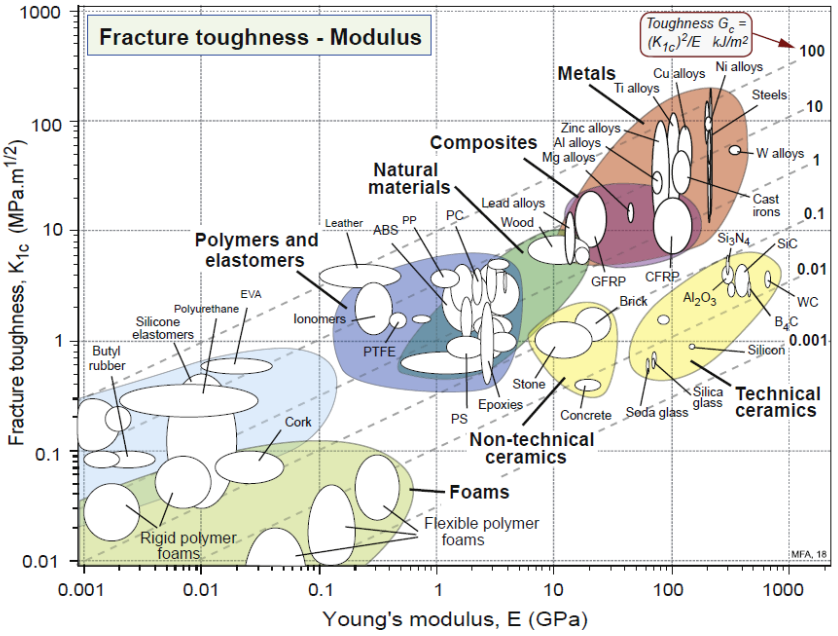

Figure 8.8 A chart of fracture toughness K_{1c} and Young's modulus E. The contours show the toughness G_c.

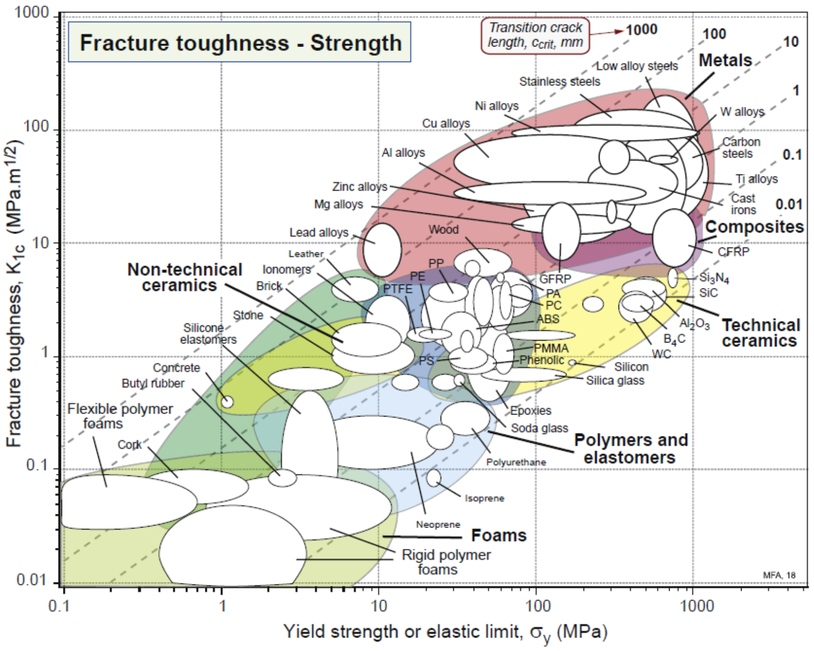

Figure 8.9 A chart of fracture toughness K_{1c} and yield strength σ_y. The contours show the transition crack size, c_{crit}.

The stress at which fracture occurs depends on both K_{1c} and the crack length c (equation (8.12)). The transition crack length c_{crit} at which ductile behaviour is replaced by brittle is given by equation (8.13). It is plotted on the chart as broken lines labeled 'Transition crack length'. The values vary enormously, from near-atomic dimensions for brittle ceramics and glasses to almost a metre for the most ductile of metals like copper or lead. Materials toward the bottom right have high strength and low toughness; they *fracture before they yield*. Those toward the top left do the opposite; they *yield before they fracture*.

The diagram has application in selecting materials for the safe design of load-bearing structures (Chapter 10). The strength–fracture toughness chart is also useful for assessing the influence of composition and processing on properties.

8.5 Drilling down: the origins of toughness

Surface energy The surface energy of a solid is the energy it costs to make it. It is an energy per unit area, units J/m^2. Think of taking a 1 m cube of material and cutting it in half to make two new surfaces, one above and one below, as in Figure 8.10. To do so we have to provide the cohesive energy associated with the bonds that previously connected across the cut. The atoms are bonded on all sides so the surface atoms lose one-sixth of their bonds when the cut is made. This means that we have to provide one-sixth of the cohesive energy H_c (an energy per unit volume) to a slice $4r_o$ thick, were r_o is the atom radius, thus to a volume $4r_o$ m^3. So the surface energy should be

$$2\gamma = \frac{1}{6}H_c \cdot 4r_o \quad \text{or} \quad \gamma = \frac{1}{3}H_c \cdot r_o$$

with H_c typically 3×10^{10} J/m^3 and r_o typically 10^{-10} m, so surface energies are around 1 J/m^2.

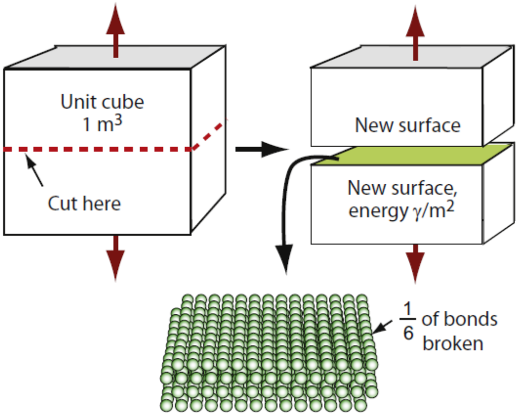

Figure 8.10 When new surface is created as here, atomic bonds are broken, requiring some fraction of the cohesive energy, H_c.

The toughness G_c cannot be less than 2γ. The chart of Figure 8.8 shows contours of G_c; for most materials its value is hundreds of times larger than 2γ. Where is the extra energy going? The answer is: into plastic work. We will examine that in more detail in a moment. First, let's examine cleavage fracture.

Brittle 'cleavage' fracture Brittle fracture is characteristic of ceramics and glasses. These have very high yield strengths, giving them no way to relieve the crack tip stresses by plastic flow. This means that, near the tip, the stress reaches the ideal strength (about $E/15$, Chapter 6). That is enough to tear the atomic bonds apart, allowing the crack to grow as in Figure 8.11. And since $K_1 = \sigma\sqrt{\pi c}$, an increase in c means an increase in K_1, causing the crack to accelerate until it reaches the speed of sound – that is why brittle materials fail with a bang. Some polymers are brittle, particularly the amorphous ones. The crack tip stresses unzip the weak van der Waals bonds between the molecules.

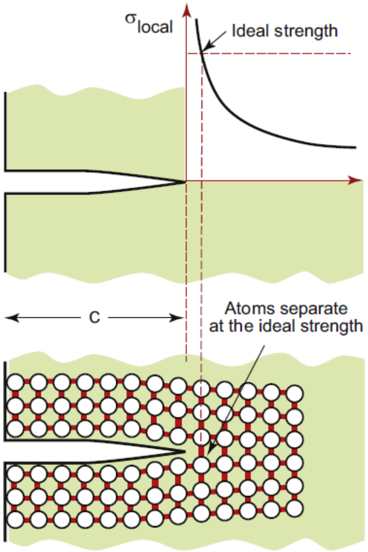

Figure 8.11 Cleavage fracture. The local stress rises as $1/\sqrt{r}$ toward the crack tip. If it exceeds that required to break inter-atomic bonds (the "ideal strength") they separate, giving a cleavage fracture. Very little energy is absorbed.

Tough 'ductile' fracture To understand how cracks propagate in ductile materials, think first of pulling a sample with no crack, as in Figure 8.12. Ductile metals deform plastically when loaded above their yield strength, work hardening until the tensile strength is reached. Thereafter, they weaken and fail. What causes the weakening? If ultra-pure, the metal may simply thin down until the cross-section goes to zero. Engineering alloys are not ultra-pure; almost all contain inclusions – small, hard particles of oxides, nitrides, sulfides, and the like. As the material – here shown as a test specimen – is stretched, it deforms at first in a uniform way, building up stress at the inclusions, which act as stress concentrations. These either separate from the matrix or fracture, nucleating tiny holes. The holes grow as strain increases, linking and weakening the part of the specimen in which they are most numerous until they finally coalesce to give a *ductile fracture*. Many polymers, too, are ductile. They don't usually contain inclusions because of the way in which they are made. But when stretched they *craze* – tiny cracks open up in the most stretched regions, whitening them if the polymer is transparent, simply because the van der Waals bonds that link their long chains to each other are weak and pull apart easily. The details differ but the results are the same: the crazes nucleate, grow, and link to give a ductile fracture.

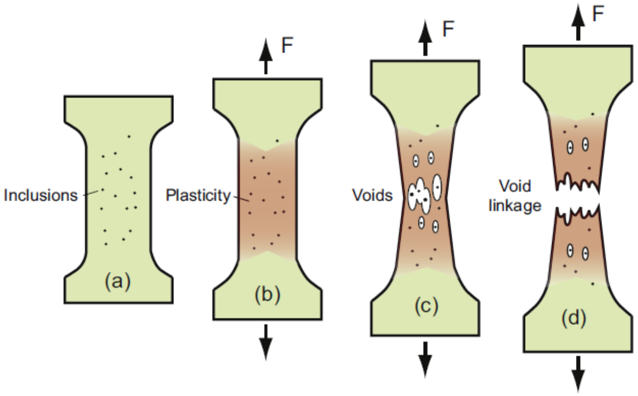

Figure 8.12 Ductile fracture. Plasticity, shown in brown, concentrates stress on inclusions that fracture or separate from the matrix, nucleating voids that grow and link, ultimately causing fracture.

Return now to the cracked sample, shown in Figure 8.13. The stress still rises as $1/\sqrt{r}$ as the crack tip is approached, but at the point that it exceeds the yield strength σ_y the material yields and a plastic zone develops. Within the plastic zone the same sequence as that of Figure 8.12 takes place: voids nucleate, grow, and link to give a *ductile fracture*. The crack advances and the process repeats itself. The plasticity blunts the crack and the stress-concentrating effect of a blunt crack is less severe than that of a sharp one, so that at the crack tip itself the stress is just sufficient to keep plastically deforming the material there. This plastic deformation absorbs energy, increasing the toughness G_c.

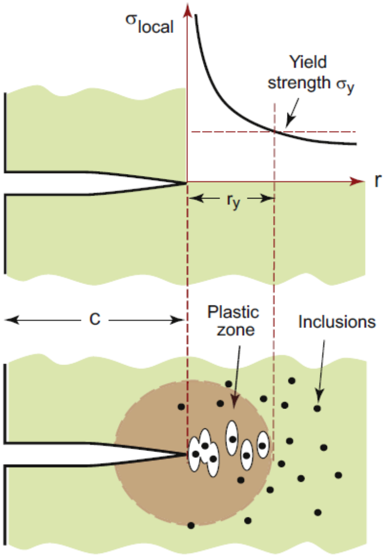

Figure 8.13 If the material is ductile, a plastic zone forms at the crack tip. Within it, voids nucleate, grow, and link, advancing the crack in a ductile mode, absorbing energy in the process.

The ductile-to-brittle transition A cleavage fracture is much more dangerous than one that is ductile: it occurs without warning or any prior plastic deformation. At low temperatures some metals and all polymers become brittle and the fracture mode switches from one that is ductile to one of cleavage — in fact only those metals with a face-centred cubic structure (copper, aluminium, nickel, and stainless steel, for example) remain ductile to the lowest temperatures. All others have yield strengths that increase as the temperature falls, with the result that the plastic zone at any crack they contain shrinks until it becomes so small that the fracture mode switches, giving a *ductile-to-brittle transition*. For some steels that transition temperature is as high as 0°C (though for most it is considerably lower), with the result that steel ships, bridges, and oil rigs are more likely to fail in winter in high latitudes, and in polar regions all year round. Polymers, too, have a ductile-to-brittle transition, a consideration in selecting those that are to be used in freezers and fridges.

Embrittlement of other kinds Change of temperature can lead to brittleness; so, too, can chemical segregation. When metals solidify, the grains start as tiny solid crystals suspended in the melt, and grow outward until they impinge to form grain boundaries. The boundaries, being the last bit to solidify, end up as the repository for the impurities in the alloy (more details in *Guided Learning Unit 2*). This grain boundary *segregation* can create a network of low-toughness paths through the material so that, although the bulk of the grains is tough, the material as a whole fails by brittle intergranular fracture (Figure 8.14). The locally different chemistry of grain boundaries causes other problems, such as corrosion (Chapter 17) — one way in which cracks can appear in initially defect-free components.