9.1 Introduction and synopsis

German engineering was not always what it is today. The rapidly expanding railway system of the mid-19th century was plagued, in Germany and elsewhere, by serious accidents caused by the failure of the axles of the coaches. The engineer August Wöhler[1] was drafted in to do something about it. It was his systematic tests that first established the characteristics of what we now call *fatigue*.

Repetition is tiring, the cause of many human mistakes and accidents. Materials, too, grow 'fatigued' if repeatedly stressed, with failure as a consequence. This chapter is about the damage and failure that can result when materials are loaded in a cyclic, repetitive way. Oscillating stresses cause the slow accumulation of damage, a little on each cycle, until a critical level is reached at which a crack forms. Continued cycling causes the crack to grow until the component suddenly fails. Fatigue failure is insidious e the stresses are often well below the elastic limit, and there is little sign that anything is happening until, bang, it fails. So when the clip breaks off your pen or your office chair collapses, it is probably fatigue that is responsible (cover picture). Even when the amplitude of the cycles is very small, some energy dissipation, or *damping*, occurs, which has consequences for applications involving vibration. We start with this low-amplitude cyclic loading and damping. We then turn to the accu- mulation of damage and cracking and associated design rules that are associated with several different regimes of fatigue failure.

9.2 Vibration: the damping coefficient

Bells, traditionally, are made of bronze. They can be (and sometimes are) made of glass, and they could (if you could afford it) be made of silicon carbide. Metals, glasses and ceramics all, under the right circumstances, have low material damping or 'internal friction', an important material property when structures vibrate. By contrast, lead, cast iron, wood (particularly when green or wet) and most foams, elastomers and polymers have high damping, useful when you want to absorb vibration.

We are speaking here of an elastic response. Until now, we have thought of the elastic part of the stressestrain curve as linear and completely reversible so that the elastic energy stored on loading is all recovered when the load is removed. No material is so perfect e some energy is always lost in a loadeunload cycle. If you load just once, you might not notice the loss, but in vibration at acoustic frequencies, the material is loaded between 20 and 20,000 times per second, as in the stress cycle of Figure 9.1(a). Then the energy loss becomes important.

The *mechanical loss coefficient* or *damping coefficient,* η (a dimensionless quantity), measures the degree to which a material dissipates vibrational energy. If an elastic material is loaded, energy is stored (Chapter 5). If it is unloaded, the energy is returned e it is how springs work. But the amount of energy returned is slightly less and varies between materials (from 10^{-5} in some glasses and ceramics, to close to 1 in elastomers). The *mechanical loss* or

[1] August Wöhler (1819e1914), German engineer, and from 1854 to 1889, director of the Prussian Imperial Railways. It was Wöhler's systematic studies of metal fatigue that first gave insight into design methods to prevent it.

damping coefficient, η, is the fraction of the stored elastic energy that is not returned on unloading. If you seek materials for bells, you choose those with low η; to absorb vibration, you choose material with high η.

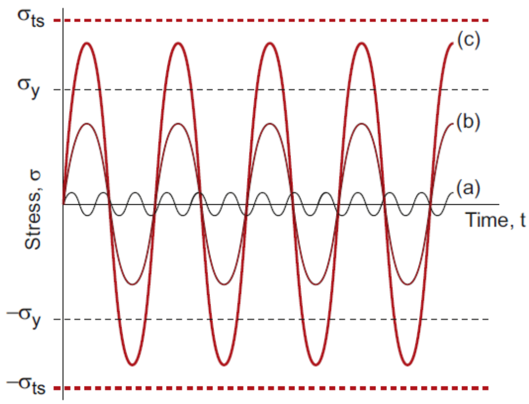

Figure 9.1 Cyclic loading: (a) very low amplitude acoustic vibration; (b) high cycle fatigue: cycling below general yield, σ_y; (c) low cycle fatigue: cycling above general yield (but below the tensile strength σ_{ts}).

9.3 Fatigue

The problem of fatigue Low-amplitude vibration causes no permanent damage in materials. Increase the amplitude, however, and the material starts to suffer fatigue. You will, at some time or other, have used fatigue to flex the lid off a sardine can or to break an out-of-date credit card in two. In metals, the cyclic stress hardens the material and causes damage such as dislocation tangles to accumulate, from which a crack nucleates and grows until it reaches the critical size for fracture. There are many potential candidates for fatigue failure — anything that is repeatedly loaded (like an oil rig loaded by the waves, or a pressure vessel that undergoes pressure cycles), rotates under bending loads (like an axle), reciprocates under axial load (like an automobile connecting rod), or vibrates (like the rotor of a helicopter).

Figure 9.2 shows in a schematic way how the stress on the underside of an aircraft wing might vary during a typical flight, introducing some important practicalities of fatigue. First, the amplitude of the stress varies as the aircraft takes off, cruises at altitude, bumps its way through turbulence and finally lands. Second, on the ground, the weight of the engines (and the wings themselves) bends them downward, putting the underside in compression; once in flight, aerodynamic lift bends them upward, putting the underside in tension. So during a flight, the amplitude of the stress cycles varies, but so does the mean stress. Fatigue failure depends on both of these and of course on the total number of cycles. Aircraft wings bend to and fro at a frequency of a few hertz. In a trans-Atlantic flight, tens of thousands of loading cycles take place; in the lifetime of the aircraft, it is millions. For this reason, fatigue testing needs to apply tens of millions of fatigue cycles to provide meaningful design data.

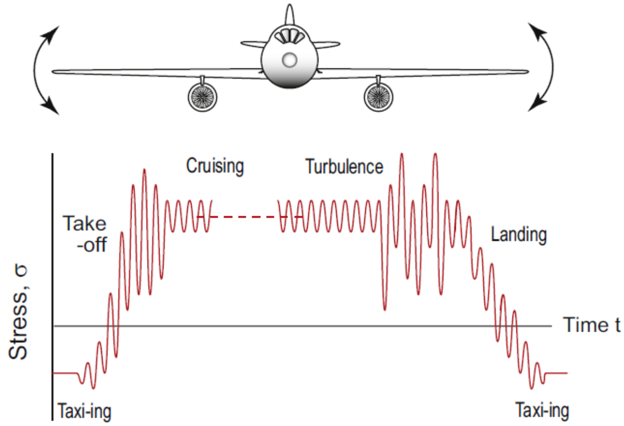

Figure 9.2 Schematic of stress cycling on the underside of an aircraft wing during a flight.

The food can and credit card are examples of *low-cycle fatigue*, meaning that the component survives for only a small number of cycles. It is typical of cycling at stresses above the yield stress, σ_y, like that shown in Figure 9.1(c). More significant in engineering terms is *high-cycle fatigue*: here the stresses remain generally elastic and may be well below σ_y, as in cycle (b) of Figure 9.1; cracks nonetheless develop and cause failure, albeit taking many more cycles to do so, as in Wöhler's railway axles.

In both cases, we are dealing with components that are initially undamaged, containing no cracks. In these cases, most of the fatigue life is spent generating the crack. Its growth to failure occurs only at the end. We call this *initiation-controlled* fatigue. Some structures contain cracks right from the word go, or are so safety-critical (like aircraft) that they are assumed to have small cracks. There is then no initiation stage — the crack is already there — and the fatigue life is *propagation-controlled*; that is, it depends on the rate at which the crack grows. A different approach to design is then called for — we return to this later.

High-cycle fatigue and the S–N curve Figure 9.3 shows how fatigue characteristics are measured and plotted. A sample is cyclically stressed with amplitude $\Delta\sigma/2$, about a mean value σ_m, and the number of cycles to cause fracture is recorded. The data are presented as $\Delta\sigma - N_f$ ('S–N') curves, where $\Delta\sigma$ is the peak-to-peak range over which the stress varies, and N_f is the number of cycles to failure. Most tests use a sinusoidally varying stress with an amplitude σ_a of

$$\sigma_a = \frac{\Delta\sigma}{2} = \frac{\sigma_{max} - \sigma_{min}}{2} \tag{9.1}$$

and a mean stress σ_m of

$$\sigma_m = \frac{\sigma_{max} + \sigma_{min}}{2} \tag{9.2}$$

as illustrated in the figure. Fatigue data are usually reported for a specified R-value:

$$R = \frac{\sigma_{min}}{\sigma_{max}} \tag{9.3}$$

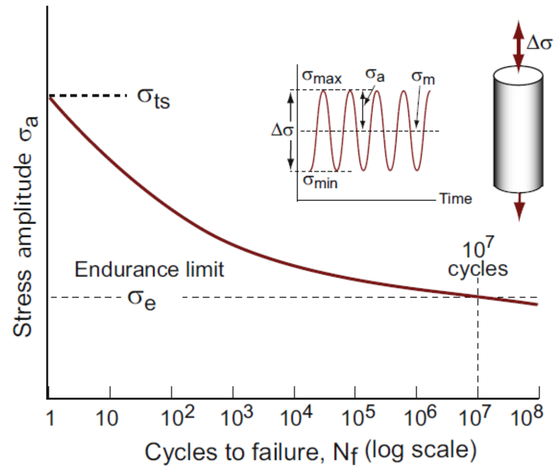

Figure 9.3 An S–N curve, with the fatigue strength at 10^7 cycles defining the endurance limit, σ_e.

An R-value of -1 means that the mean stress is zero; an R-value of zero means the stress cycles from zero to σ_{max}. For many materials, there exists a *fatigue* or *endurance limit*, σ_e (units: MPa). It is the stress amplitude σ_a, about zero mean stress, below which fracture does not occur at all or occurs only after a very large number ($N_f > 10^7$) of cycles. Design against high-cycle fatigue is therefore very similar to strength-limited design, but with the maximum stresses limited by the endurance limit σ_e rather than the yield stress σ_y.

Experiments show that the high-cycle fatigue life is approximately related to the stress range by what is called Basquin's law:

$$\Delta\sigma N_f^b = C_1 \tag{9.4}$$

where b and C_1 are constants; the value of b is small, typically 0.07 and 0.13. Dividing $\Delta\sigma$ by the modulus E gives the strain range $\Delta\varepsilon$ (since the sample is elastic):

$$\Delta\varepsilon = \frac{\Delta\sigma}{E} = \frac{C_1/E}{N_f^b} \tag{9.5}$$

or, taking logs,

$$\log(\Delta\varepsilon) = -b\log(N_f) + \log(C_1/E)$$

This is plotted in Figure 9.4, giving the right-hand, high-cycle fatigue part of the curve with a slope of $-b$.

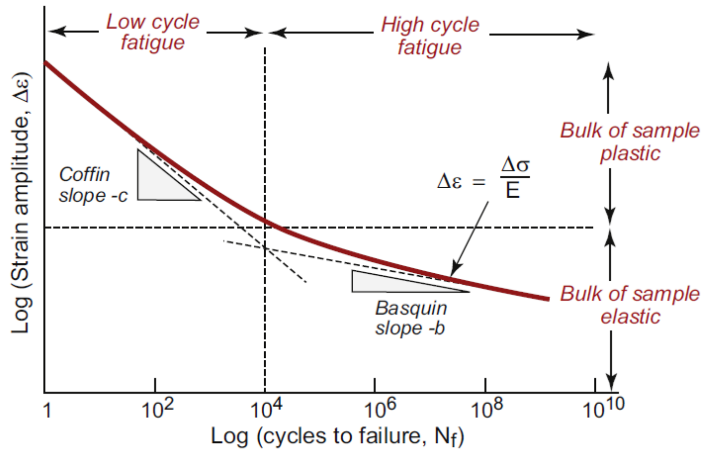

Figure 9.4 The low- and high-cycle regimes of fatigue and their empirical description.

Low-cycle fatigue In low-cycle fatigue, the peak stress exceeds yield, so at least initially (before work hardening raises the strength), the entire sample is plastic. Basquin is no help to us here; we need another empirical law, this time that of Dr Lou Coffin:

$$\Delta\varepsilon^{pl} = \frac{C_2}{N_f^c} \tag{9.6}$$

where $\Delta\varepsilon^{pl}$ means the plastic strain range — the total strain minus the (usually small) elastic part. For our purposes, we can neglect that distinction and plot it in Figure 9.4 as well, giving the left-hand branch (and a smooth transition to high-cycle behaviour). Coffin's exponent, c, is much larger than Basquin's; typically it is 0.5.

High-cycle fatigue: mean stress and variable amplitude These laws adequately describe the fatigue failure of uncracked components cycled at constant amplitude about a mean stress of zero. But as we saw, real loading histories are often much more complicated (Figure 9.2). How do we make some allowance for variations in mean stress and stress range? Here we need yet more empirical laws, courtesy this time of Goodman and Miner. Goodman's rule relates the stress range $\Delta\sigma_{\sigma_m}$, for failure under a mean stress σ_m to that for failure at zero mean stress $\Delta\sigma_{\sigma_o}$:

$$\Delta\sigma_{\sigma_m} = \Delta\sigma_{\sigma_o}\left(1 - \frac{\sigma_m}{\sigma_{ts}}\right) \tag{9.7}$$

where σ_{ts} is the tensile strength. So Goodman's rule applies a correction to the stress range to give the same number of cycles to failure. A tensile mean stress σ_m means that a smaller stress range $\Delta\sigma_{\sigma_m}$ is as damaging as a larger $\Delta\sigma_{\sigma_o}$ applied with zero mean (Figure 9.5(a)). The corrected stress range may then be plugged into Basquin's law.

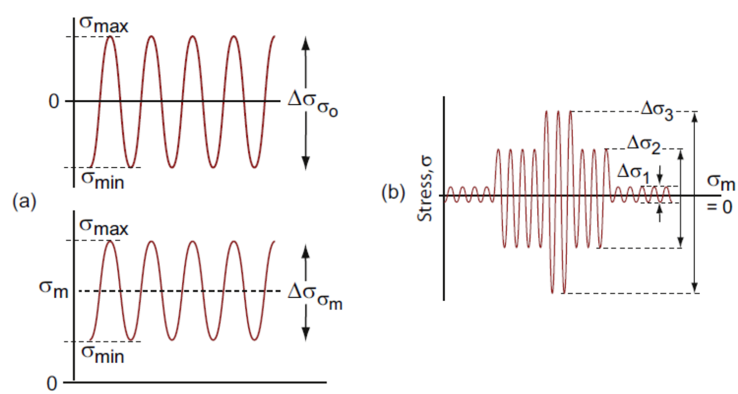

Figure 9.5 (a) S–N curves refer to cyclic loading under a zero mean stress. Goodman's rule scales the amplitude to an equivalent value under a mean stress σ_m. (b) When the cyclic stress amplitude changes, the life is calculated using Miner's cumulative damage rule.

The variable amplitude problem can be addressed approximately with Miner's *rule of cumulative damage*. Figure 9.5(b) shows an idealised loading history with three stress amplitudes (all about zero mean). Basquin's law gives the number of cycles to failure if each amplitude was maintained throughout the life of the component. So if N_1 cycles are spent at stress amplitude $\Delta\sigma_1$, a fraction N_1/N_{f1} of the available life has been used up, where N_{f1} is the number of cycles to failure at that stress amplitude. Miner's rule assumes that damage accumulates in this way at each level of stress. Then failure will occur when the sum of the damage fractions reaches 1 — that is, when

$$\sum_{i=1}^{n} \frac{N_i}{N_{f,i}} = 1 \tag{9.8}$$

Goodman's rule and Miner's rule are adequate for preliminary design, but they are approximate; in safety-critical applications, tests replicating service conditions are essential. It is for this reason that new models of cars and trucks are driven over rough 'durability tracks' until they fail — it is a test of fatigue performance. The discussion so far has focused on initiation-controlled fatigue failure of uncracked components. Now it is time to look at those containing cracks.

Fatigue loading of cracked components In fabricating large structures like bridges, ships, oil rigs, pressure vessels and steam turbines, cracks and other flaws cannot be avoided. Cracks in castings appear because of differential shrinkage during solidification and entrapment of oxide and other inclusions. Welding, a cheap, widely used joining process, can introduce both cracks and internal stresses caused by the intense local heating. If the cracks are sufficiently large that they can be found, it may be possible to repair them, but finding them is the problem. All non-destructive testing (NDT) methods for detecting cracks have a resolution limit; they cannot tell us that there are no cracks, only that there are none longer than the resolution limit, c_{lim}. Thus, it is necessary to assume an initial crack exists and design the structure to survive a given number of loadings. So how is the propagation of a fatigue crack characterised?

Fatigue crack growth is studied by cyclically loading specimens containing a sharp crack of length c like that shown in Figure 9.6. We define the cyclic stress intensity range, ΔK, using equation (8.4), as

$$\Delta K = K_{max} - K_{min} = \Delta\sigma\sqrt{\pi c} \tag{9.9}$$

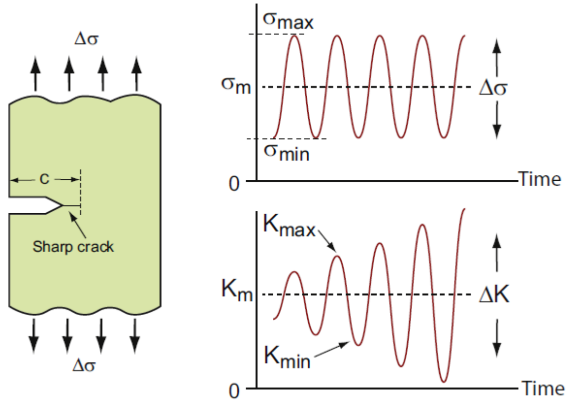

Figure 9.6 Cyclic loading of a cracked component. A constant stress amplitude $\Delta\sigma$ gives an increasing amplitude of stress intensity, $\Delta K = \Delta\sigma\sqrt{\pi c}$ as the crack grows in length.

The range ΔK increases with time under constant cyclic stress because the crack grows in length: the growth per cycle, dc/dN, increases with ΔK in the way shown in Figure 9.7. The rate is zero below a threshold cyclic stress intensity ΔK_{th}, useful if you want to make sure it does not grow at all. Above it, there is a steady-state regime described by the Paris law:

$$\frac{dc}{dN} = A\Delta K^m \tag{9.10}$$

where A and m are constants. At high ΔK, the growth rate accelerates as the maximum applied K approaches the fracture toughness K_{1c}. When it reaches K_{1c}, the sample fails at the peak of the final load cycle.

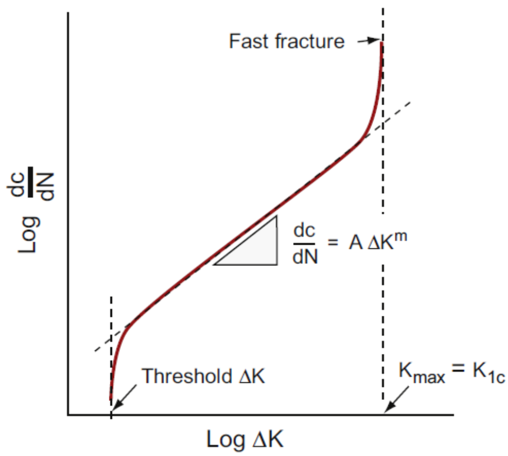

Figure 9.7 Crack growth rate during cyclic loading.

Safe design against fatigue failure in potentially cracked components means calculating the number of loading cycles that can safely be applied without the crack growing to a dangerous length. We return to this in Chapter 10.

9.4 Charts for endurance limit

The most important single property characterising fatigue strength is the endurance limit, σ_e, at 10^7 cycles and zero mean stress (an R-value of -1). Given this and the ability to scale it to correct for mean stress, and to sum contributions when stress amplitude changes (equations (9.7) and (9.8)), enables design to cope with high cycle fatigue.

Not surprisingly, endurance limit and strength are related. The strongest correlation is with the tensile strength σ_{ts}, shown in the chart of Figure 9.8. The data for metals and polymers cluster around the line

$$\sigma_e \approx 0.33\sigma_{ts}$$

shown on the chart. For ceramics and glasses

$$\sigma_e \approx 0.9\sigma_{ts}$$

In the next section, we examine why.

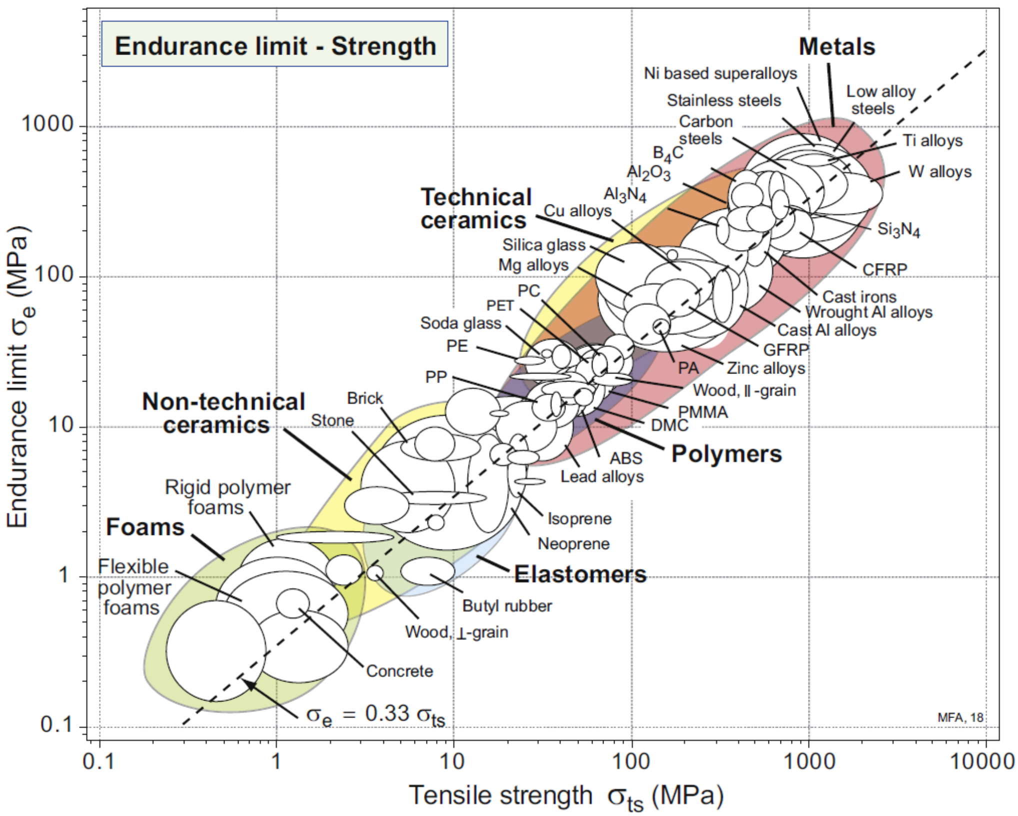

Figure 9.8 The endurance limit plotted against the tensile strength. Almost all materials fail in fatigue at stresses well below the tensile strength.

9.5 Drilling down: the origins of damping and fatigue

Material damping: the mechanical loss coefficient There are many mechanisms of material damping. Some are associated with a process that has a specific time constant; then the energy loss is centred about a characteristic frequency. Others are frequency independent; they absorb energy at all frequencies. In metals, a large part of the loss is caused by small-scale dislocation movement: it is high in soft metals like lead and pure aluminium. Heavily alloyed metals like bronze and high-carbon steels have low loss because the solute pins the dislocations; these are the materials for bells. Exceptionally high loss is found in some cast irons, in manganese–copper alloys and in magnesium, making them useful as materials to dampen vibration in machine tools and test rigs. Engineering ceramics have low damping because the dislocations in them are immobilised by the high lattice resistance (which is why they are hard). Porous ceramics, on the other hand, are filled with cracks, the surfaces of which rub, dissipating energy, when the material is loaded. In polymers, chain segments slide against each other when loaded; the relative motion dissipates energy. The ease with which they slide depends on the ratio of the temperature T to the polymer's glass temperature, T_g. When $T/T_g < 1$, the secondary bonds are 'frozen'; the modulus is high, and the damping is relatively low. When $T/T_g > 1$, the secondary bonds have melted, allowing easy chain slippage; the modulus is low, and the damping is high.

Fatigue damage and cracking A perfectly smooth sample with no changes of section, and containing no inclusions, holes or cracks, would be immune to fatigue provided neither σ_{max} nor σ_{min} exceeds its yield strength. But that is a vision of perfection that is unachievable. Blemishes, small as they are, can be deadly. Rivet holes, sharp changes in section, threads, notches and even surface roughness concentrate stress in the way described in Chapter 7, Figure 7.7. Even though the general stress levels are below yield, the locally magnified stresses can lead to reversing plastic deformation. Dislocation motion is limited to a small volume near the stress concentration, but that is enough to cause damage that develops into a tiny crack.

In high cycle fatigue, once a crack is present, it propagates in the way shown in Figure 9.9(a). During the tensile part of a cycle, a tiny plastic zone forms at the crack tip, stretching it open and thereby creating a new surface. On the compressive part of the cycle, the crack closes again, and the newly formed surface folds forward, advancing the crack. Repeated cycles make it inch forward, leaving tiny ripples on the crack face marking its position on each cycle. These 'striations' are characteristic of a fatigue failure and are useful, in a forensic sense, for revealing where the crack started and how fast it propagated.

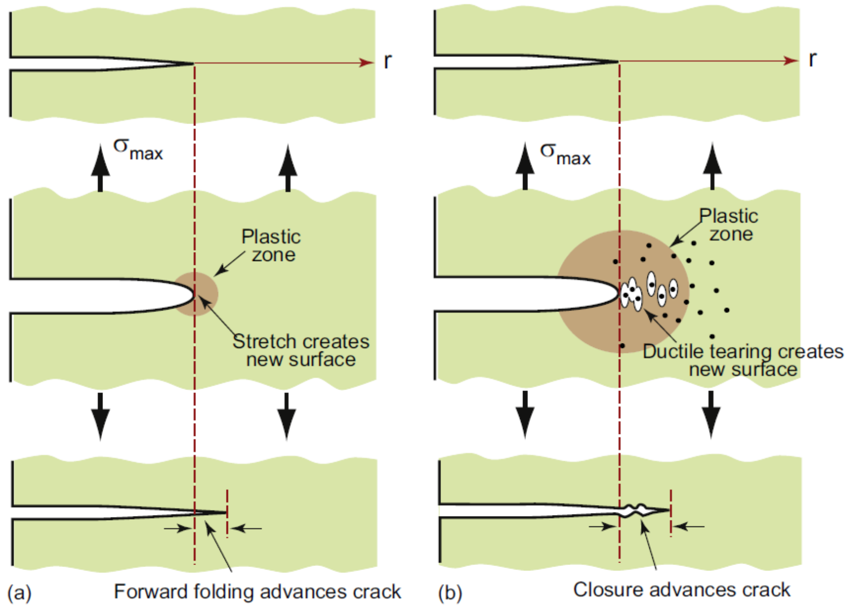

Figure 9.9 (a) In high cycle fatigue, small-scale crack tip plasticity opens and advances the crack. (b) In low cycle fatigue, a larger plastic zone nucleates voids which coalesce to advance the crack.

In low cycle fatigue, the stresses are higher and the plastic zone larger, as in Figure 9.9(b). It may be so large that the entire sample is plastic, as it is when you flex the lid of a tin to make it break off. The largest strains are at the crack tip, where plasticity now causes voids to nucleate, grow and link, just as in ductile fracture (Chapter 8).