MSE307 Engineering Alloys 2014-15 L1: Introduction, Alloy Selection and Lifing

D. Dye[a]

[a] Rm G03b, Department of Materials, Imperial College London, London SW7 2AZ, UK. david.dye@imperial.ac.uk

---

Abstract

The Aim of the course is to build a synthesised understanding of the major classes of engineering alloys, and how these are developed, made and used. In this first lecture, we first introduce the course, and then turn to look briefly about alloy selection; our conclusion will be that it is more complicated than a simple Ashby-map type approach would suggest. Then, we turn to think about how 'synthesis' or 'survey' courses differ from 'theory,' 'mechanisms,' or 'concept' types of courses. In the second part of the lecture, we then turn to consider gas turbine engineering as a case study, how to find the safe fatigue life for a part in the presence of defects, and how such defects distributions might be measured. This will motivate some of the future lectures.

MSE307 is a course that aims to build a synthetic understanding of engineering alloys; how they work, how to design alloys, and how this interacts with how machine designers design with them. Most of the first and second year content concerns physical principles and phenomena; MSE307 seeks to apply these and place them in the context of making, designing and using structural metallic materials. Einstein said *"[I do not] carry such information in my mind since it is readily available in books. …The value of a college education is not the learning of many facts but the training of the mind to think."* Today, all the world's knowledge is available, very nearly, through the internet. But, being able to assemble this into a coherent whole and practically apply it, the skill of *synthesis*, is a critical skill for professionals, for knowledge workers in the modern era.

And so, it is clearly wrong to think of this course as a dull procession of facts to be memorised and regurgitated in an exam, then quickly forgotten as expeditiously as possible. There *are* many facts to be placed together in building a picture of 'how alloys work;' your study to date should provide the glue to hold the pieces together and weave them into a beautiful tapestry. If you have the glue and the ability to build the picture, then it will be relatively easy to hold a coherent picture in your mind, whereas straight memorisation will be both arduous and of little lasting value to you.

The course is led by two overarching goals for the metallurgist; (i) to understand how to design alloys and microstructures to order, and (ii) to manufacture components with these that don't break unexpectedly in service. We want to connect the topics in second year together in order to do this. So solubility (MSE204) connects to how much solution strengthening we can achieve (MSE203); the intermetallic precipitates we can produce and their interfaces govern the particle size distributions we can have and therefore precipitate strengthening; the processing we are allowed to have governs the textures and grain sizes we can achieve. And the processing governs the defect distributions. Its also important to note that yield strength may not be the over-riding goal; fracture toughness may be more important, or the allowable cyclic stress for a fatigue-limited part, or some other property of the material. Overall, we are aiming for a practical orientation (a concern to make materials that are realisable in real components in service) overlaying a strong theoretical foundation (a love of the insight that physical science provides).

The course is arranged, in 2014-15, as follows. One series of 8 lectures is provided by Dr David Dye, and introduces the course, followed by four lectures on titanium alloys. The second series of 8 lectures is provided by Dr Vassili Vorontsov, and concerns nickel base superalloys and aluminium alloys. The final series of 8 is provided by Prof Tony Paxton, on the physical metallurgy of steels. Steels are an enormous topic which in some universities provide the foundation for the entire materials science degree programme; they of course are by far the highest tonnage engineered material used on Earth after concrete. But, in the present implementation of the course ironmaking and steelmaking, and the mechanical behaviour of steels, will not be discussed in detail. For Dr Vorontsov and Dr Dye's lectures the major focus will be aerospace materials and particularly jet engines, which will be introduced later in this lecture.

Syllabus

The material discussed in the 8 lectures by David Dye is as follows

1 Introduction to MSE307; materials selection; jet engines and lifing
2 Alloy design, the Hume Rothery rules and elastic moduli.

3 Casting and forging of wrought nickel and titanium alloys, and speciality steels; VIM, VAR, ESR, EBCHR and open and closed die forging.
4 The Sioux City air accident.
5 Phase metallurgy of titanium alloys.
6 Titanium alloy microstructure engineering.
7 Micromechanics of titanium alloys I.
8 Micromechanics of titanium alloys II; near-α, metastable β and CP Ti.

Prerequisites and Reading

The graduate, having completed their formal education, must be able to extend their knowledge from textbooks; indeed, the ability to self-educate based on the foundations provided by a degree is the primary distinguishing feature of the graduate compared to the school-leaver. And, in a course of this nature, the lectures can only be the beginnings of the study. Indeed, for a full time degree over 32 weeks with each lecture courses comprising 1/8th of the year and 24 lectures, each lecture can fairly claim $32 \times 40/(8 \times 24) = 6.7\,\mathrm{h}$ of a student's time. Therefore even after accounting for revision and examination practice, and the lecture itself, significant out-of-class study should be the norm.

For the titanium lectures the primary reference is Gerd Lütjering and Jim Williams excellent book, *Titanium.* In addition, the standard metallurgy references from first and second year remain of great utility. For the Ni, Al and steels lectures then those lecturers will provide their own guidance.

1 A Cottrell. An Introduction to Metallurgy, 2nd Ed., Institute of Materials, London, UK, 1995.
2 RE Smallman. Modern Physical Metallurgy, 4th Ed., Butterworths, 1985.
3 DA Porter and KE Easterling. Phase Transitions in Metals and Alloys, 2nd Ed., Chapman and Hall, 1992.
4 GE Dieter. Mechanical Metallurgy, 3rd Ed., McGraw Hill, 1988.
5 D Hull and DJ Bacon. Introduction to Dislocations, 4th Ed., Butterworth-Heinemann, 2001.
6 G Lütjering and JC Williams. Titanium, Springer, 2003.

Towards the end of the degree, students should also be starting to prepare for postgraduate study, where reading and critically assessing research articles becomes normative. So in addition to the notes and lectures above, for many of the lectures I will also set a journal article to be read beforehand.

The course is primarily designed for students who have already done the first and second year of the Imperial materials degrees; therefore that background knowledge is assumed. Erasmus students and MSc students lacking this background from their previous studies can find assistance in the books above and in the lecture notes for first and second year.

Delivery Mode

David Dye's lectures will be delivered by peer instruction, as for MSE104 and MSE203. That is, students will first read the notes, watch videos of the exposition of the material online, and read around the subject. Then, in the class session we will discuss the topics raised by looking at simple conceptual problems. Evidence suggests that this results in greater understanding of the material when the time comes for revision and then better exam performance. Prof Paxton and Dr Vorontsov will lecture their parts of the course in a conventional fashion.

Course Support and Assessment

This section of the course is composed of 8 lectures. It is examined in the summer exam (approx. one third of the exam). The materials provided consist of the lectures and accompanying notes. Reading of the literature and textbooks is encouraged.

1. Introduction

1.1. Materials Selection

First, lets consider how a designer chooses the optimal material to use for their application; understanding the thought process will help the metallurgist to advise them, and to design more appealing alloys for future use. We will consider first airframe materials, before turning to gas turbines.

Aeroplanes have, in one sense, remained very similar since the dawn of the jet age in 1958; long pressurised tubes with wings, onto which external podded gas turbine are mounted. But, over the last 20 years, there has been a slow revolution in aircraft manufacture as aluminium alloys have been replaced, at least in part, by composites and titanium. Aircraft being introduced in the 2010s consume around 25% less fuel (and hence produce 25% less CO_2) per passenger mile than the 1980s aircraft they replace; *e.g.* the B767 by the B787. Around half of this improvement is due to the airframe and half to improvements in the engines, and in each of those cases around half of those improvements are due to the materials. Thus, materials contribute more than improved aerodynamics to the emissions reductions, and so materials understanding is a core competence for a successful aircraft or jet engine manufacturer.

A new class of aircraft has also emerged over the last 20 years, principally in the military sphere, of the unmanned aerial vehicle or UAV. Lacking pilots, it might be possible that such vehicles might be able to tolerate a greater accident rate and hence, use less extensively qualified materials. In practice, the cost of avionics and design limit

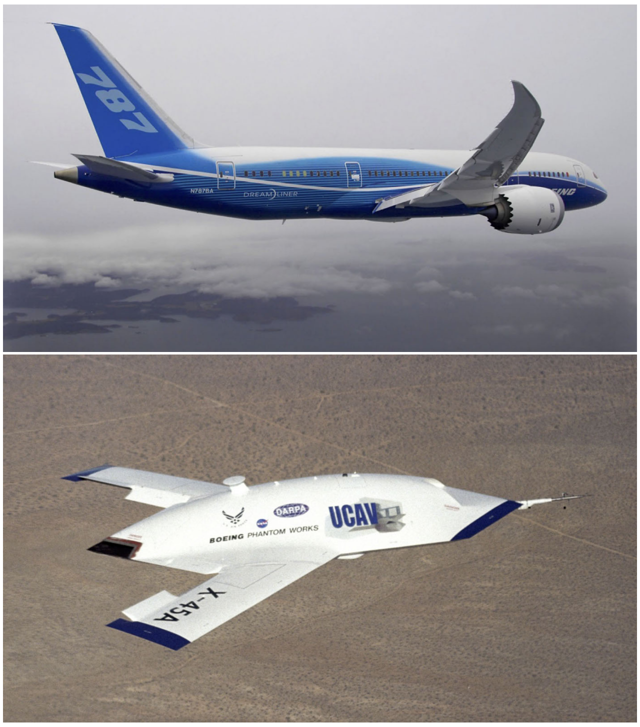

Figure 1: The Boeing 787 and X-45A (a demonstrator unmanned aerial vehicle) in flight.

the utility of scrimping on material quality, but it should still be understood that cost is a prime consideration for material insertion right across the aerospace industry, even though jet engines cost about as much, per unit weight, as silver.

Turning to jet engines, the thermal efficiency of a turbine increases quite dramatically with temperature and follows the form of $\Delta T/T_{\max}$, where $T_{\max}$ is the turbine entry temperature and the base temperature must be greater than the environmental temperature of the heat sink, so approx. $300K$. Therefore the efficiency increases with $T_{\max}$. However, the maximum temperature is very often limited by the materials capability that can be provided in the operating environment. And, we want to use the lowest mass material we can, to avoid having to move all of that unnecessary weight. Therefore, we use different materials in different parts of the engine that experience different temperatures.

At the front of the engine, we suck in cold air from the air stream impinging on the engine. In a modern *high bypass* turbofan, 80% of this air is accelerated by the first few rows of blades - called *stages* - and is then ejected around the rest of the engine. Over time, fan diameters have tended to increase and the bypass ratio has increased, which improves the engine efficiency. The remaining air is then compressed, in the succeeding stages, increasing its pressure and temperature. At some point, in the compressor, the temperature becomes so high that nickel superalloys must be used, for both the blades and the discs that connect the blades to the shaft that drives them.

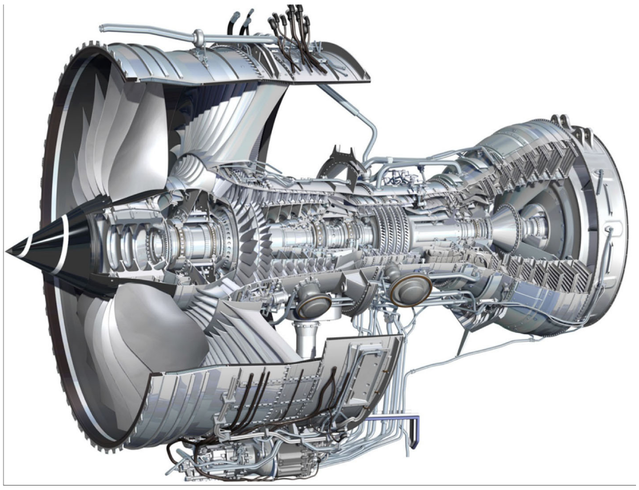

Figure 2: Schematic cutaway of the Rolls-Royce TrentXWB, the engine for the A350, which entered service in Dec 2014.

Then, the air is mixed with fuel and burnt in the combustor. This comprises an inner chamber or flame tube, which is insulated by compressor air from the pressure casings on the inside and outside. There are also fuel injectors, etc. Burning of the fuel results in a temperature and pressure increase, taking the gas stream temperatures to above 1600 °C at entry to the turbine. The turbine then gradually increases in diameter as the gas cools and increases in volume, and power is extracted. Finally a high temperature supersonic jet of air exits the rear of the engine. Mixing of this air with the bypass air helps reduce the jet noise.

Thus, for different sorts of components different properties or degradation mechanisms end up being the limiting factor. For the wings or aeroplanes, often stiffness is the limiting factor - the total deflection that can be tolerated. Landing gear are strength limited. Jet engine turbine blades are limited by creep - continuous extension at constant load that would, left unchecked, result in contact with the casing. Fan blades, on the other hand, are fatigue limited. This can be either low cycle fatigue (*i.e.* flight cycles) or high cycle fatigue ('flutter' or vibrational modes at transient conditions on run-up or run-down). Nuclear power station reactor pressure vessels are, absent irradiation embrittlement, LCF limited (the number of fuel power operational cycles).

Turning back to a table from MSE104, recall that the very best steels, aluminium alloys and conventional titanium alloys all have the same density-corrected (specific) Young's modulus and strength. Therefore, at least to a simplistic analysis, its not strength or stiffness that leads to the selection of one material over the other. It might be the specific properties achievable at an acceptable cost - so in the automotive industry there has been a continuous innovation tussle between the steel and aluminium industries over the major structural components. Commonly, it might be the fatigue or other degradation mechanism-

Table 1: Specific properties for some typcial aerospace materials.

Material	Young's Modulus E GPa	Yield Stregnth σ_y MPa	Density ρ kg.m^{-3}	E/ρ MJ.kg^{-1}	σ_y/ρ kJ.kg^{-1}
CFRP	120	1200	1600	75	**750**
Ti-5553	110	1400	4500	24	310
Ti-64	105	900	4500	23	200
7075-T6 (Al)	70	550	2700	26	200
A300M steel	210	2050	7830	27	262

limited performance that can be achieved, with an affordable manufacturing route and using existing equipment. After all, for existing equipment the marginal cost of production os nearly zero, and this factor inhibits the adoption of new production technology.

But, two stand-out factors from this table are apparent; high strength titanium alloys do have a higher specific strength than any other metallic material, and are corrosion resistant. The other is that carbon fibre reinforced polymer composites can have triple the specific strength and stiffness of conventional metallic materials.

But, the march of composites hasn't proceeded without upset. The F22 was supposed to be a mostly polymeric airplane, but it turned out that over the development phase much of the composite had to be replaced with titanium, with the same modulus and hence the potential for replacement without redesign. And so the F22 ended up being an astounding 42% titanium (by mass), mostly comprised of forged and machined components. Similarly, in the B787 development, Boeing had problems with the wing spar root fixturing, which reveals the main design problem with composites. Even if long fibre carbon composites can reliable be manufactured in such large pieces, they only show their outstanding performance in tension along the fibre direction. Their shear and transverse properties are quite poor and so in triaxial stress situations, such as joints and fittings, problems can occur. Mostly this is resolved by making such fittings out of titanium, which is corrosion resistant and modulus matches the composites.

Figure 3: The F22 fighter aircraft, the airframe of which is 42% Ti, 24% polymer matrix composite (PMC), 24% aluminium alloy and 10% steels. At the beginning of manufacturing development a majority PMC share was planned.

The other thing to remember about metals is that, more than strength or stiffness, their amazing property is their damage tolerance - their ability to resist catastrophic crack growth and instead to fail gradually, allowing time for inspection to discover growing crack and retire components. Most ceramics (including carbon fibres themselves) have a fracture toughness K_{1c} of around $1\,\text{MPa}\,\sqrt{\text{m}}$, whereas even an ordinary steel has a K_{1c} of $90\,\text{MPa}\,\sqrt{\text{m}}$.

We can also, to an extent, often optimise the material processing to, at minimum, trade strength for toughness, allowing optimisation relatively late in the development phase. So, we see that the plane strain fracture toughness of both Al alloys and β-Ti alloys shows this relationship. We also see that, for aluminium alloys, controlling the tramp Fe and Si contents that result in crack-initiating inclusions allows an increase of around $10\,\text{MPa}\,\sqrt{\text{m}}$, which is very significant for the strongest alloys used in aircraft.

All of these discussions lead to the main point of this section: there is more to materials selection than Ashby

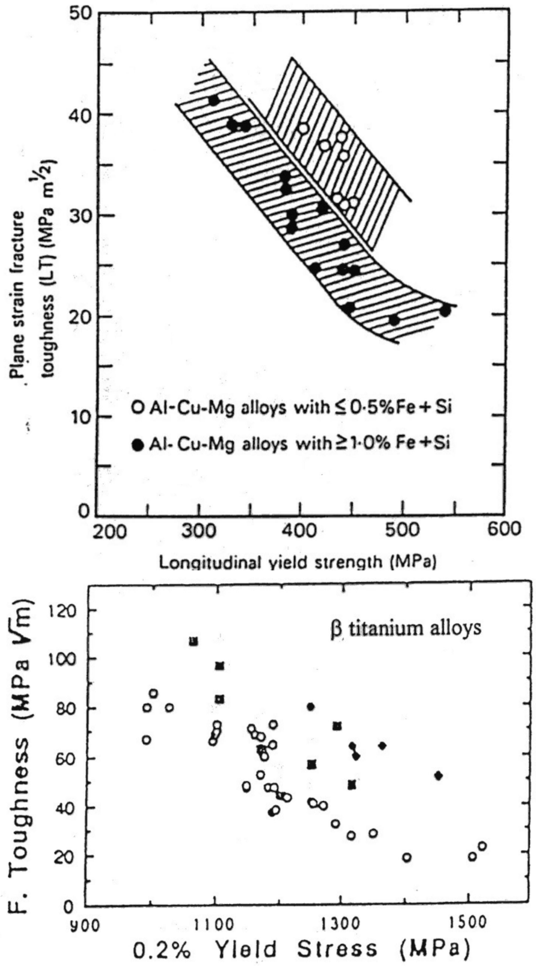

Figure 4: Variation in fracture toughness with strength achievable by processing for (top) aluminium alloys and (bottom) β-Ti alloys.

diagrams. In reality, for the highest performance, most safety-critical materials we use in society today, alloy design, manufacturing, component design and in-service support are tightly interlinked. These four, quite distinct areas of expertise must both be individually excellent and relate well together for a world-leading engineering organisation to develop market-leading products. In addition, often a *balance of properties* must be maximised - including poorly-understood phenomena such as corrosion, erosion, oxidation and sulphidation performance, or weldability and manufacturability, or fatigue behaviour, or the manufacturing value chain, together with the 'simple' properties of strength and stiffness.

1.2. Jet engines

In the titanium and nickel parts of the course, much of the discussion will be around jet engine materials, for which is it helpful to have both some technical and industrial background. Rolls-Royce is a large UK corporate, with a market capitalisation of £14.3bn (Jan '11 figures), placing it in the middle of the FTSE 100. In 2010, its earnings were £10.9bn (0.5% of UK GDP), approximately 2/3rds of which were from civil aerospace, with a 10.2% operating margin and a 4.9% net profit margin. It holds around 40% of the market for large civil jet engines serving the twin aisle market, in a near-duopoly with GE, the world's largest industrial conglomerate and a major US defence contractor. In 2011, it had 40,300 employees, around 22,000 of whom were in the UK, with a revenue of around £260k per employee per year; 50% of this revenue was in services. It is the largest UK exporter on any measure, with a £900m annual R&D spend. It delivers around 900 engines per year, with over half its installed base being under 10 years old, guaranteeing a service revenue stream for decades to come; the order book backlog also stretches nearly out of sight (e.g. 10+ years). And it has nearly zero net debt. And that's before we start considering the supply chain. If there is going to be a UK manufacturing renaissance, we need Rolls-Royce to grow, and we need more firms like it.

It is particularly important to the materials R&D community because its products have that combination of strong drivers to improve efficiency and low tolerance for failure that can sustain high enough margins to drive research. So, a 0.5% improvement in operating efficiency is make-or-break for an engine order, costing an airline £m in fuel per year. In contrast as a society we tolerate production failures in the automotive sector that kill thousands each year, and many will deliberately set out to purchase heavy, inefficient and slow vehicles for fashion (SUVs in Kensington, anyone?) And so in the steel and aluminium industry, there is relatively little investment in R&D, despite their high production volumes, because there is so little profit margin to pay for innovation.

It is also important to realise that bringing a jet engine to market that is a close evolution of an existing engine costs on the order of £0.5bn. There is serious doubt as to whether the US will ever develop a new high temperature 'core' of a fighter aircraft engine again, for example. For a new entrant to the market, the cost to develop an entirely new engine architecture would undoubtedly be in the billions of pounds, and then the manufacturer would need to tolerate a poor safety record and uncompetitive fuel efficiency for 5 years of fleet operations to gain enough experience to make an iterative modification. On the third attempt, after an almost unimaginable outlay, a competitive engine might be developed. So this is a market with very high barriers to entry.

Turning to the engine architecture, as we said earlier, essentially the gas flow proceeds first through a fan, then the non-bypass air goes through the compressor, into the combustor and then out through the turbine. The turbine extracts the energy to drive the fan and compressor. Two or three concentric shafts are used, with the widest diameter, shortest shaft connecting the hottest, fastest spinning, highest pressure part of the turbine to the similarly hot, high pressure inner compressor. There are then the Intermediate (IP) compressor and turbine, and the low pressure (LP) fan/compressor and turbine. Having three shafts instead of two adds weight but allows for better thermofluid optimisation of the stages.

As the gas heats up, there is a transition from a titanium construction to nickel superalloys, at around 550 °C, while the shafts are made of steel. The fan casing is also

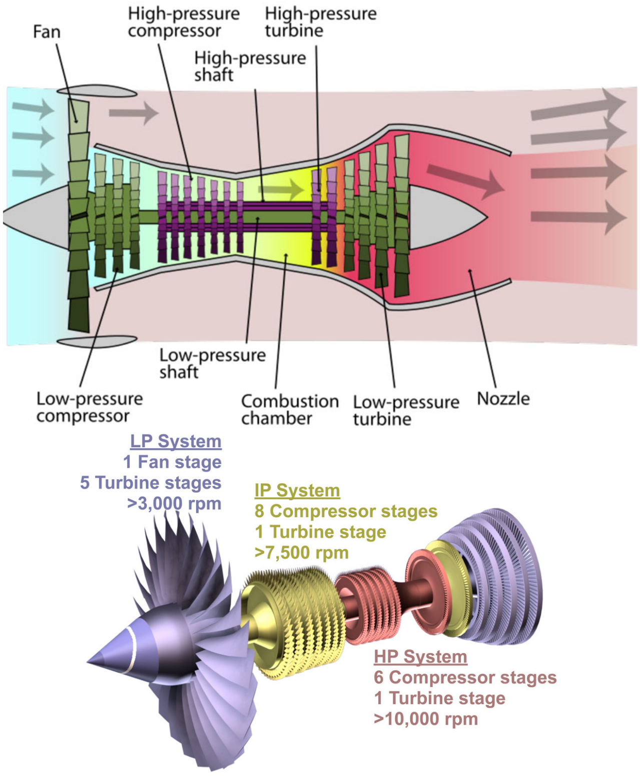

Figure 5: (top) Schematic showing the different parts of a jet engine, (bottom) numbers of stages and rotation speeds for the three shafts in the Trent XWB.

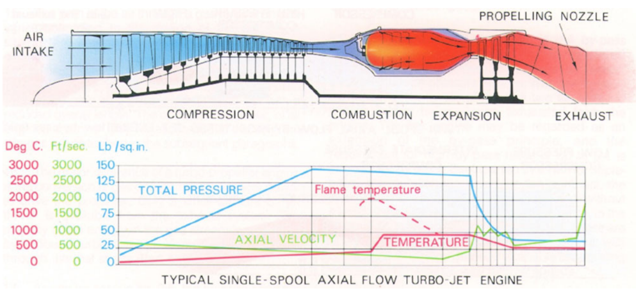

Figure 6: Evolution of pressure, temperature and gas velocity through the core of a jet engine. From Rolls Royce, 'The Jet Engine.'

important, as it must contain the fan blades in bird strike / blade off scenarios. As important as the blades, the rotor discs that attach them to the shafts are so heavy that their fragments cannot be contained in the event of failure and therefore they cannot be allowed to fail and hazard the airframe. A disc failure due to a generic materials problem would be very serious for the credibility of a gas turbine manufacturer. Some approximate operating temperatures, stresses and lives are given in the Table.

In a turbine blade, the temperature varies very significantly from the core of the blade, which is cooled, to the hot skin. The blade is also coated with a ceramic thermal barrier; together the cooling air and thermal barrier coating bring the temperature down anywhere up to 500 °C. As the blade creeps, the applied stresses redistribute such that the strong, cool core supports the majority of the load.

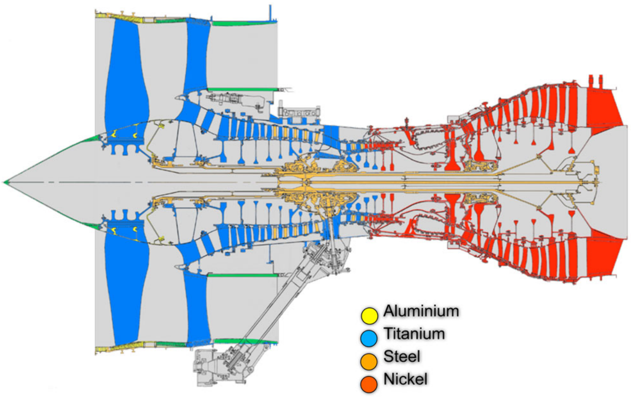

Figure 7: Simplified section through a jet engine, showing the materials used in different locations.

Table 2: Some approximate maximum operating temperatures and stresses for different components in a large civil jet engine.

Component	T °C	Peak Stress MPa	Replacement Interval	Material
Fan blade	40	550	N=10,000	Ti-6Al-4V
IP compressor blade	450	450	20,000 h	Ti-6246
HP compressor disc rim	750	400	20,000 h	RR1000
HP compressor disc bore	600	600	20,000 h	RR1000
HP turbine blade core	700	700	10,000 h	CMSX-4
HP turbine blade skin	1100	200	10,000 h	CMSX-4

Figure 8: B52 on takeoff. For a modern aircraft, no smoke is visible.

One problem in aviation is that as we make planes more efficient (at maybe 2% p.a.), and as the world economy grows, demand increases by about 5% p.a. (tripling in 20 years). Therefore, if we successfully complete an energy transition to non-emitting sources in land-based transport, electricity production and space heating, aviation could grow from 12% of UK emissions today to be a major contributor. And the problem is that kerosene has such a high energy density that its almost impossible to foreseen substitution (e.g. with hydrogen or batteries) on the grounds of physics. Using crop-derived biofuel may be the best we can do. In any case, as well as on industrial and economic grounds, there is a strong moral imperative to improve jet engine operating efficiency.

And to date, there has been continuous improvement. Since 1950, fuel burn per passenger km has decreased by 70% and takeoff noise by a factor of 4 - an in addition we don't see trails of black soot any more. Seriously, check out video of a B52 or a croaky old 747 and compare to a new airplane. So there is hope.

1.3. *Lifing*

We turn now to lifing, the process of establishing the safe life of a component subject to fatigue loading. Recall the form of a fatigue crack growth curve, with a threshold cyclic stress intensity factor ΔK_{th} below which crack growth will not occur, where $\Delta K = \Delta\sigma\sqrt{\pi a}$, $\Delta\sigma$ is the applied cyclic stress and a_0 is the initial flaw size. Once the crack begins to grow, a Paris Law type relationship with $\frac{da}{dN} = C\Delta K^m$ is found, and the crack grows until the fast fracture stress intensity $K_{1,c}$ is approached.

Deterministically Safe Lifing.

Now, say that we know the the probability distribution function of defects in the material will be. In most materials there will be some population of gas pores from casting and ceramic inclusions from the moods used for the liquid metal. We could then estimate, for (i) a given desired failure rate, say 1 per 10 million takeoffs, (ii) a given material replacement rate (say 60,000 h or 8,000 flight cycles), (iii) a volume of material in the engine (say 2 tonnes), what the largest inclusion would be that would have a 50% chance of appearing at that frequency. Given that inclusion size

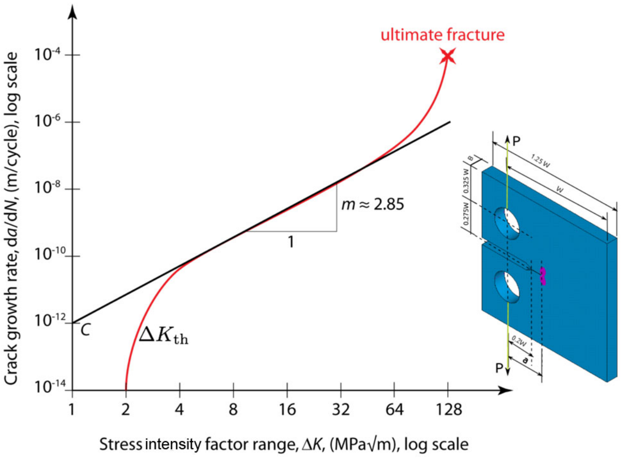

Figure 9: Schematic of a fatigue crack growth curve for an alloy showing Paris Law behaviour.

($\sim$ 1 mm), we could then calculate the maximum stress that could be applied without driving that flaw into fatigue crack growth at ΔK_{th}. The problem would be that no airplane could be made that could fly under such a conservative regime. Notice that, despite the term, this approach still requires knowledge of the (necessarily rare) occurrence of small defects in material produced by the tonne, including all the possible manufacturing variances.

The historic, empirical approach that is actually used is based on spin testing. Here, it is assumed that the distribution of component lives is log-normal. Then, a series of spin tests are performed to grow cracks, with periodic inspections to stop the test once a 0.75 mm crack is found by non-destructive inspection. The mean and standard deviation of these lives is then found, and then mean decremented by 3σ to get a 1/750 life. For 3-8 components in the test series, this is around 1/3rd of the expected mean disc life. Further safety factors are supplied by the time for the 0.75 mm crack to grow to failure, and by periodic inspections mandated at overhaul intervals. Therefore, for 20,000 cycle lives, and $\sim$ 6 critical assemblies in the engine, this approach gives a probability of loss of around 0.4 per million flights.

The problem with this approach is that first that it has no mechanistic basis - a statistical assumption is made, and then empirical testing is used to establish a life. The other problem is that the 2/3 of components are discarded long before their actual service life. This is handy for the jet engine manufacturers, as it guarantees a nice service business, but otherwise just makes flying more expensive that it needs to be (jet engine sales are a bit like razor blades - there isn't much money in the engine, people make money on the parts and service contract). Seocndly, its disheartening - implicitly, an acceptable level of risk is tolerated. Since the Vietnam war, the notion of corporations and governments taking risks with people's lives has become intolerable, particularly when the failures are 'newsworthy' - intermittent and spectacular, rather than everyday and mundane. This is true even if those risks are tiny compared to everyday hazards we willingly accept, like having our children killed by careless drivers (see C. Perrow, *Normal Accidents*).

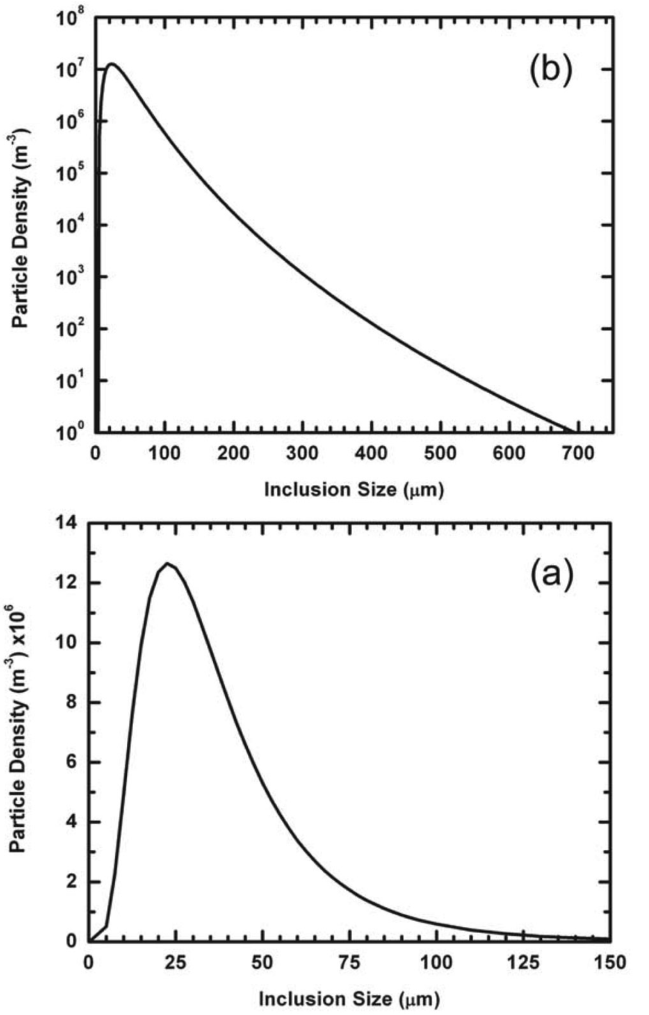

Figure 10: Typical inclusion size distributions in a super clean, powder metallurgy nickel-base superalloy such as Rene 95, on (a) linear and (b) log-linear axes. Note that, for the volumes in a typical jet engine, the probability of a 0.5 mm size inclusion is actually quite high. From Reed, The Superalloys, p.265.

So, with the advent of fracture mechanics in the 1970s, it is worth asking if something better cannot be done. For a given flaw size and stress, we can integrate the Paris Law to obtain the life

$$\frac{da}{dN} = C\Delta K^m = C_1\left(\Delta\sigma\sqrt{\pi a}\right)^m \tag{1}$$

We then separate and integrate:

$$\int_0^{N_f} dN = \int_{a_0}^{a_{K_{1c}}} \frac{da}{C_1 \Delta\sigma^m \left(\pi a\right)^{m/2}} \tag{2}$$

So the life N_f is given by

$$N_f = \frac{1}{C_2 \Delta\sigma^m}\left[a^{\left(1-\frac{m}{2}\right)}\right]_{a_0}^{a_{K_{1c}}} \tag{3}$$

The first approach we consider is called damage tolerant lifing or retirement for cause. He, there initial flaw size is set to the limit detectable by non-destructive inspection, $a_{\rm NDI}$. Then, the toughness and applied stress are used to define the critical flaw size at failure, $a_{\rm crit}$. The Paris Law is then used to find out how long a crack would take to grow from $a_{\rm NDI}$ to $a_{\rm crit}$, and inspection is then mandated more frequently than this interval. If, on inspection, a crack is found, then the part is retired, otherwise it is returned to service. Thus, discs are only retired when the is cause to do so.

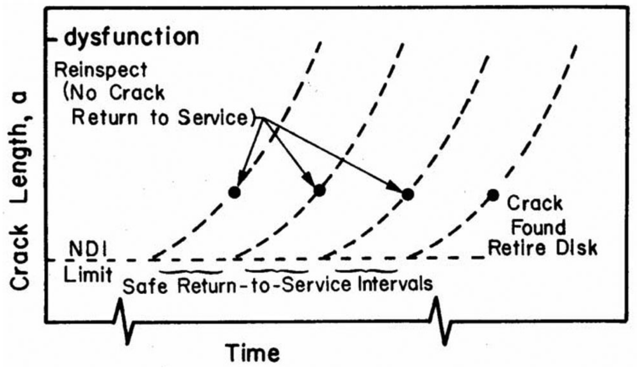

Figure 11: Illustration of the retirement for cause approach. The NDI inspection limit and fracture mechanics are used to establish the inspection interval; if no crack is detected at inspection then the assumed maximum crack size is re-set to the NDI limit and the material allowed to continue in service.

This procedure subtly shifts the emphasis in designing an alloy towards increasing its fatigue crack growth resistance (the Paris Law coefficients), as for a given inspection interval determined by operational requirements this will be the determinant of the operating stress than can be tolerated, and hence the part weight and engine efficiency. It is argued that improving this so-called 'damage tolerance' is preferable to a futile effort to eliminate defects and hence stay below $\Delta K_{\rm th}$, which will tend to result in an unaffordable spiral of increasing costs.

The problem, which we shall return to later in the course, if that it requires the inspection procedure to be executed flawlessly. It has been widely adopted in the USAF and deployed now for nearly 20 years, and so there is at least some evidence, in small engine fleets, that this can work. In the US military, there is therefore evidence that with strong indoctrination, training and leadership, such risks can be mitigated by operator training and oversight. The US Navy has a similarly impressive record of nuclear reactor operational safety, which many argue has not been achieved by US civil power plant operators. Therefore there is at least some reason to doubt that retirement for cause is a practical lifing strategy for mass application. The way to address this criticism is to incorporate a probability of detection at each inspection interval and therefore to turn this into a probabilistic lifing strategy.

A final reason why this strategy is more appropriate for combat aircraft is that they - except in wartime - don't operate for that many hours compared to civil aircraft, because of the costs of fuel and training. Therefore a relatively short inspection interval can be tolerated.

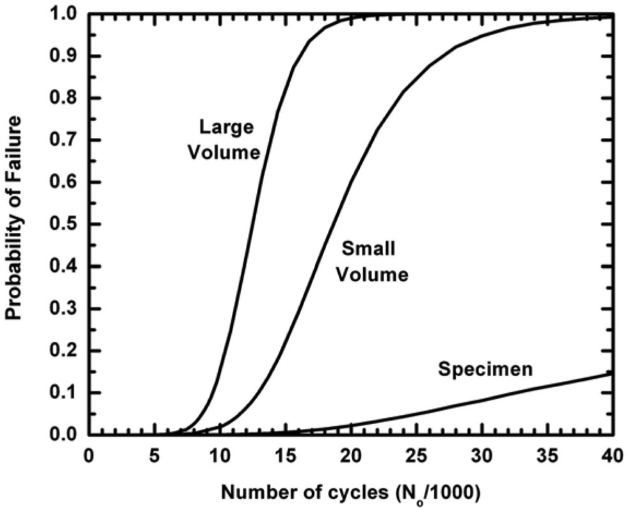

Figure 12: Variation in the expected cyclic life for different stressed volumes, using a probabilistic approach to lifing for Rene 95. From Reed, The Superalloys, p.268.

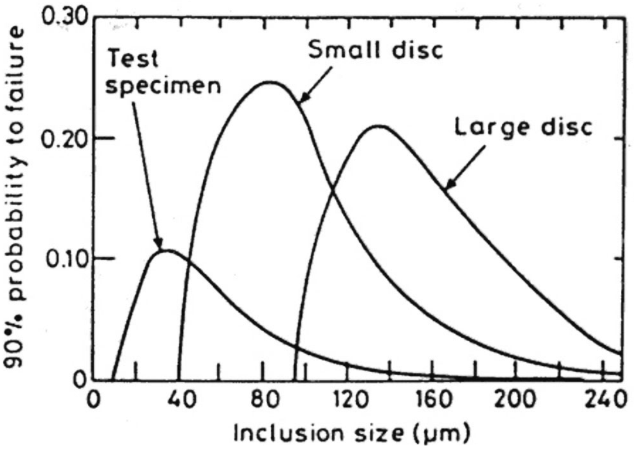

Figure 13: Probability density function for the size of the initiating inclusio,n, for three different geometries in IN718 at 1050 MPa. From Flower, High Performance Materials in Aerospace, Fig 4.12.

The last lifing methodology we shall consider is termed probabilistic lifing. We start with the observation that, when we test lab specimens, we test only $\sim 0.5\,{\rm cm}^3$ at a time, whereas in service we test maybe $\sim 0.5 \times 10^6\,{\rm cm}^3$ of titanium per engine. So the stress-life (S-N) curves we generate in the lab will be based on initiating defects that are unrepresentative (too small). But, if we know the defect size distribution and the fatigue growth curve, we can estimate a safe total life. This is particularly attractive when ceramic inclusions might be present. Essentially it is similar to the so-called 'deterministically safe' approach, but it provides a life rather than a guarantee of survival.

So, for each initial defect size, at each location in each part, the time is found for it to grow to the critical size $a_{\rm crit}$

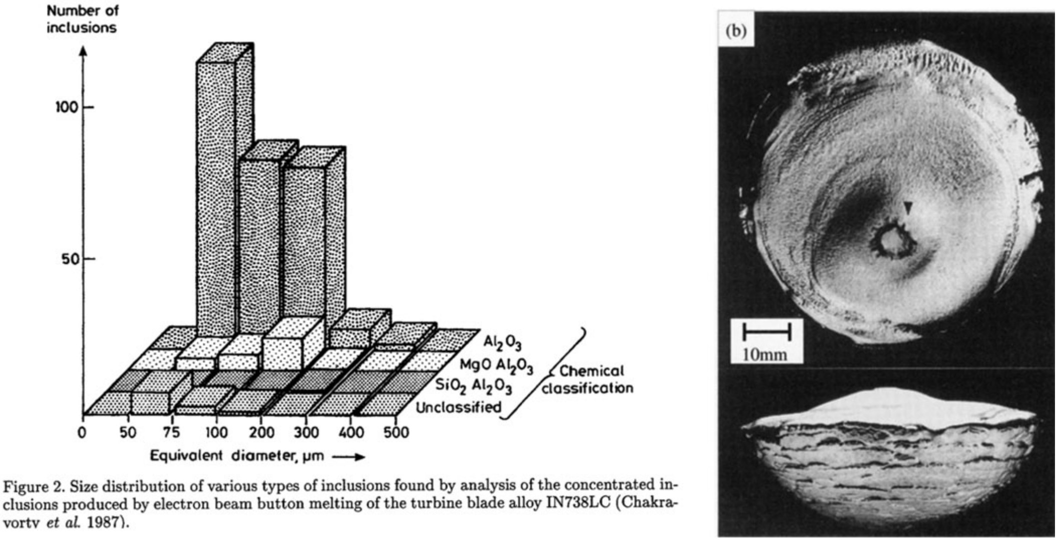

Figure 2. Size distribution of various types of inclusions found by analysis of the concentrated inclusions produced by electron beam button melting of the turbine blade alloy IN738LC (Chakravorty *et al.* 1987).

Figure 15: (left) Size distribution of different types of inclusions found in a nickel superalloy, after concentration by electron beam button melting. (right) Images of the melted button.

for failure, and then the probability distribution, susceptible volume and failure frequency are integrated to provide a life.

The attraction of this is that it can allow for approaches like variations in the defect density around a casting - for instance, casting pores in non-highly stressed locations might not be critical. Similarly, mitigation actions such as initial and periodic NDI that reduce the frequency of occurrence of large defects can be incorporated, without assuming that operators are perfect.

Some Consequences The above discussion will have illuminated certain points that it is worth now discussing, with the aid of data. Firstly, for a given defect distribution inherited from manufacture, large volumes will fail earlier than testpieces, and at lower stresses.

Secondly, the defect giving rise to failure will be small in a testpiece than in a disc, because the largest defect present is likely to be smaller. So, testpeice tests may not be representative of real parts. The attached Figure for IN718 illustrates this point.

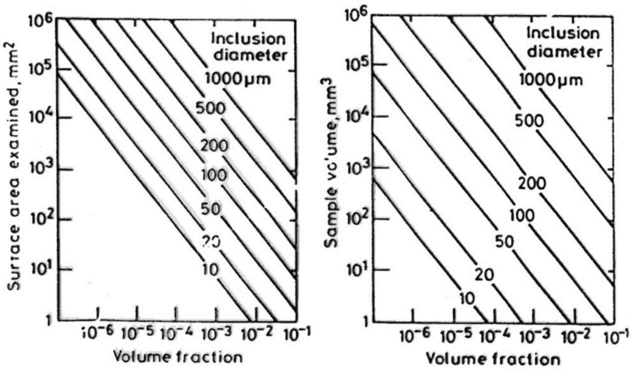

Figure 14: Surface area and volume that must be examined to give an 80% confidence in measuring the volume fraction of inclusions of various diameters, assuming a random distribution. From Flower, High Performance Materials in Aerospace, Fig 4.14.

There is also an inspection challenge. The adjacent Figures show the amount of area or volume that must be examined to establish the inclusion volume fraction, for different inclusion sizes. For optical examination of surfaces, these will often result in unfeasible demands - for example, in 400 micrographs $0.5 \times 0.5\,\mathrm{mm}$ in size (sufficient to measure 10 µm particles), one would only examine $100\,\mathrm{mm}^2$ of material, sufficient merely to measure a volume fraction of 10^{-4} or above. Volumetric inspections, such as radiography and ultrasonics, of course enable more material to be examined and so can make this challenge less difficult, but only if the contrast mechanism is strong enough to detect the defect.

For ceramic inclusions resulting from oxides in the melt or erosion to ceramic molds, another way to measure the defect distribution is to concentrate them somehow for analysis. Ceramic inclusions that have lower density than the metal and that are large enough not to stay suspended in the solution can, for example, be floated to the top of a button that is remelted, and then the surface inspected in the SEM. For some materials, such as nickel superalloys, this is often used to characterise the inclusion distribution for a given manufacturing supply chain route.

— END OF LECTURE 1 —