High-Cycle Fatigue Analysis

Stress-Life Made Easy

Engineering Mechanics White Paper

Endeavor Analysis, LLC
11033 Marine View Dr., SW
Seattle, WA 98146
PH: 206.805.9030
www.endeavoranalysis.com

Predictive Engineering, Inc.
2512 SE 25th Ave, Suite 205
Portland Oregon 97202
PH: 503.206.5571
www.predictiveengineering.com

Table of Contents

List of Figures

1. INTRODUCTION

A brief walk-through is given on how a fatigue analysis works and a bit of foundation knowledge to guide a new user through this process. It should be mentioned that we are focusing on the high-cycle fatigue of metals. Just to ensure that we are all on the same page, the difference between low-cycle and high-cycle fatigue is briefly summarized in Table 1.

Table 1: A quick summary of the difference between Low-Cycle and High-Cycle Fatigue

Low-Cycle Fatigue (Strain-Life)	High-Cycle Fatigue (Stress-Life)
Stress > 80% σ_{Yield}	Stress < 80% σ_{Yield}
Cycles < 10,000	Cycles > 10,000

Since most design work focuses on structures with near infinite life, the stress target is typically 80% of the material's yield strength (σ_{ys}) or lower. This requirement makes the stress-life approach a natural fit.

As a side note, one should not consider this article as the "last word" or even a "complete word" about the fatigue process. There are dozens of handbooks on fatigue analysis, and if one would like to become proficient in this branch of engineering, it can take years of study and perhaps a master's or a Ph.D. of engineering along the way. Heretofore, our objective in this note is just to provide a common foundation of understanding from which to launch more complete discussions.

1.1 THE PROCESS

For clarity, the fatigue process is broken down into five sequential steps:

(i) Stress calculations, whether by hand or turning the FEA crank;
(ii) Sketching out the load events to create load cycles;
(iii) Form logical pairings of maximum and minimum stresses between load sets (Rainflow);
(iv) Calculate damage for each load pairing from fatigue curve;
(v) Sum damage using Miner's rule.

1.2 WHAT IS FATIGUE IN METALS?

This is not meant to be a treatise but just enough to whet your interest in the theme of fatigue theory, and its application.

Fatigue starts with the movement of dislocations within the metal's crystal lattice. These dislocations pile up along grain boundaries, impurities (i.e., oxides), secondary hard phases (e.g., the silicon network within A356 cast aluminum alloys) and interstitial compounds or just in general, anything that is not part of the pure crystalline metallic matrix. Over thousands and thousands of cycles, these dislocations pile up to such an extent that a network of microscopic cracks is created within the material. Once this network of cracks has formed, the fatigue process speeds up significantly with these small cracks bridging together into larger cracks and finally zipping along to form a final large massive crack where the structure unexpectedly fails. The failure is termed unexpected since nobody thought that the stresses were excessive since they had designed to 50% of the yield/ultimate strength of the material or some other "rule-of-thumb".

If one is of the curious sort, it begs to question how the 50% rule-of-thumb got started. In many handbooks, Figure 1 aptly describes the relationship between alternating stress and the number of cycles to failure for ferrous and non-ferrous materials. Ferrous materials (steel) exhibit a plateau while non-ferrous materials (aluminum, brass, magnesium, etc.), will eventually fail given billions and billions of cycles. Of course, exceptions occur and in practice, here is a short list:

- Occasional overloads or impacts can destroy the ability of the material to have a fatigue limit;
- Corrosion (a common rational offered-up sometimes by metallurgists to explain unexpected failures);
- High-temperatures that can introduce microstructural changes.

In the special case of non-ferrous materials, it is more common to specify fatigue strength (S_f) as stress per number of cycles to failure. As example, a manufacturer of aluminum A356-T6 truck hubs uses a design limit of S_f = 100 MPa with an estimated 1e10^8 cycles to failure. For a standard commercial long-haul truck, it's enough to satisfy their clients' fatigue requirements.

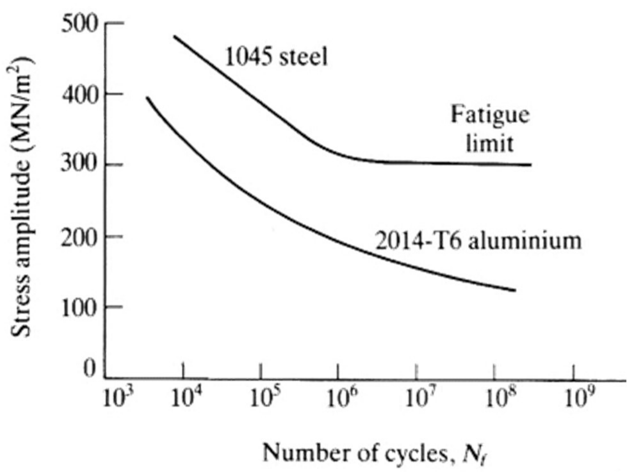

Figure 1: A standard representation of an S-N (stress-cycles) curve for typical ferrous and non-ferrous materials. The fatigue limit for ferrous materials is roughly ½ the material's ultimate strength.

Another way to think about this 50% rule is to look at the mechanics of a void within a large body. Standard mechanics calculate the stress concentration (K_t) of a spherical void within a large body as 2.0. The FEA model in Figure 2 provides a visualization of these mechanics in color.

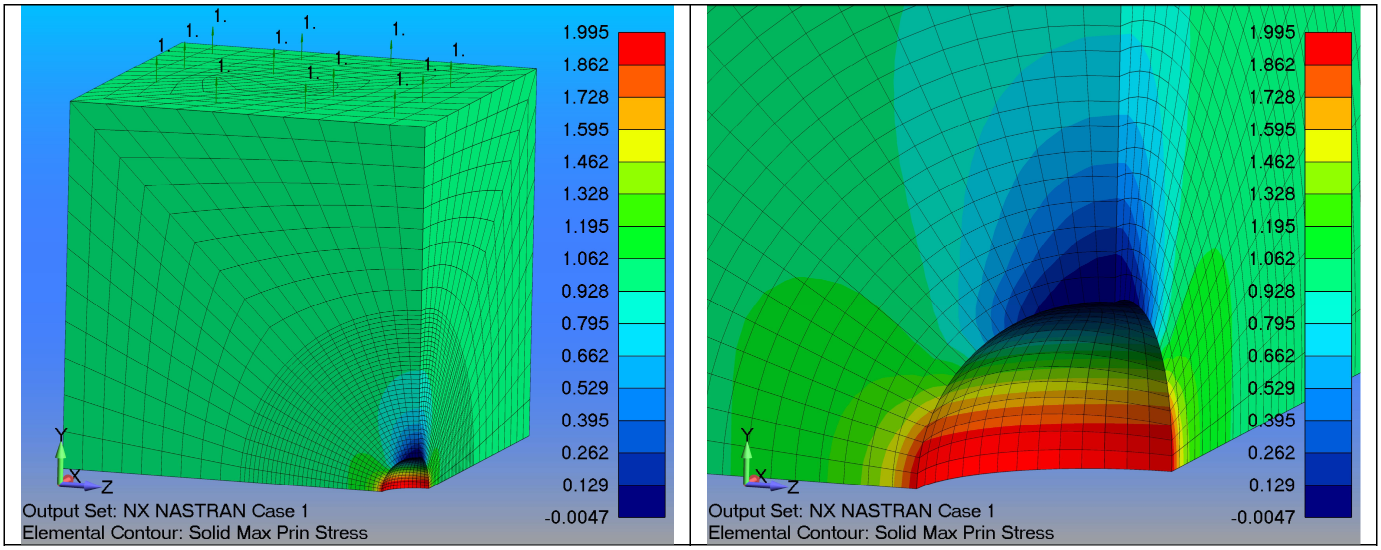

Figure 2: A symmetric block is given a uniform pressure load of -1.0. The FEA model shows a maximum stress of 1.995.

Since all materials contain small defects, it is easy to imagine that when designing to 50% of the yield strength, the true stress at microstructural defects is at 100% of the material's yield strength. Not to belabor this point but since this is a material's discussion and the yield strength of a ferrous/non-ferrous material is based on the empirical observation that when the load is released, no observable plastic deformation is noted, but in reality, extensive dislocation movement occurs at stresses greater than 50% of the yield strength of the material (σ_y). Hence, even before the material reaches its σ_y dislocations are moving, combining, clustering and causing nano-sized cracks in the crystalline structure. Given this basis, whenever the load is greater than 50% σ_y, we have dislocations moving through-out the material and near defects, causing rather massive localized plasticity. This is the essence of material fatigue and why every test sample will fail at a different number of cycles due to metallurgical imperfections.

1.2.1 Stress Modification Factors

Of course, all these defects can be present on the surface of the material and given that normal usage of most engineered structures, scratches, dings and other damage is a common occurrence. Additionally, bending loads always focus the stress on the outer fibers or surfaces of the structure and it is normal and to be expected that for most loading scenarios, the highest stress will be on the surface. To account for these unknowns within the material and on the surface of the structure, the industry has relied upon the usage of stress modification factors. For example, in many aerospace fatigue analyses, a stress concentration of 1.4 is used on all hand and FEA calculated stress values. If one calculates 100 MPa, then a stress value of 1.4*100 (140 Mpa) would be used to calculate the number of cycles to failure from the fatigue curve. Although not exactly physics based it allows one to account for the differences between smooth and polished fatigue test coupons and that of the normal surface finish found on most engineered structures.

1.2.2 Summary

What this means for the designer responsible for the survivability of structures subject to high-cycle loads is that one should be a bit nervous about how your material is going to react after a few millions of cycles and whether or not your loads are going to be well-behaved or subjected to periodic high-load excursions. When viewed from this perspective, the unexpected failure in a high-cycle environment often has its roots in not considering the coupling between material variability and fatigue mechanics. But this is getting a bit ahead and subsequent sections will delve a bit deeper into material variability as it relates to fatigue life and how one can safely hedge your bet.

1.3 Basic Fatigue Mechanics Terminology

If you master this section, you'll know more about interpreting stress data and S-N fatigue curves than the majority of your colleagues in the engineering profession. These concepts seem simple enough but are actually quite difficult to implement in practice given the general noise inherent in real-world load cases. But before we sprint, let's crawl through the classic fatigue schematic shown in Figure 3. What the schematic is telling us is that the applied load is creating an elevated mean stress (σ_m) in our structure while the measured stress cycles up and down between a max (σ_{max}) and a min (σ_{min}). This all seems quite logical mathematically but maybe a bit hard to visualize in engineering practice. One example of a structure that experiences high mean stress is a helicopter rotor blade. During operation, the centripetal force pulls on the blade creating a field of constant high tension. As the blade rotates, the drag force switches sign every 180 degrees. This creates the perfect sinusoidal σ_{max}, σ_{min} under a high σ_m.

What causes material damage is the alternating stress component ($\Delta\sigma$) and accelerated by the magnitude of the mean stress (σ_m). This accelerator effect is shown schematically in Figure 4. As the mean stress increases, the alternating stress (σ_a) to initiate fatigue failure decreases.

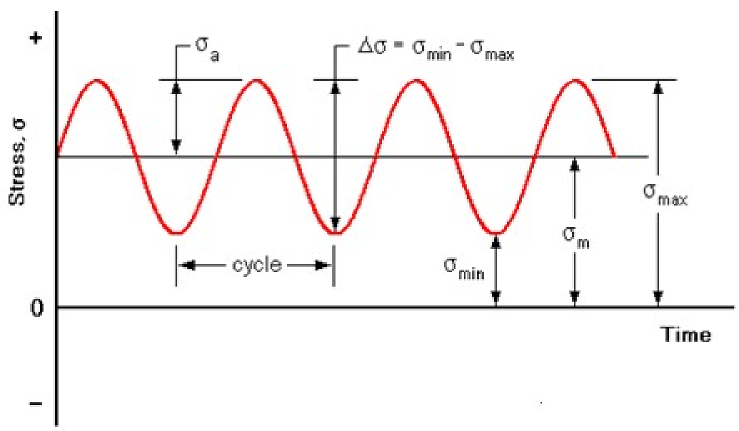

$$\Delta\sigma = \sigma_{max} - \sigma_{min} = Stress\ Range$$

$$\sigma_a = \frac{\sigma_{max} - \sigma_{min}}{2} = Stress\ Amplitude$$

$$\sigma_m = \frac{\sigma_{max} + \sigma_{min}}{2} = Mean\ Stress$$

$$R = \frac{\sigma_{min}}{\sigma_{max}} = Stress\ Ratio$$

Figure 3: This sketch lays down the foundation of how stresses within a cyclic event are described within the world of fatigue terminology.

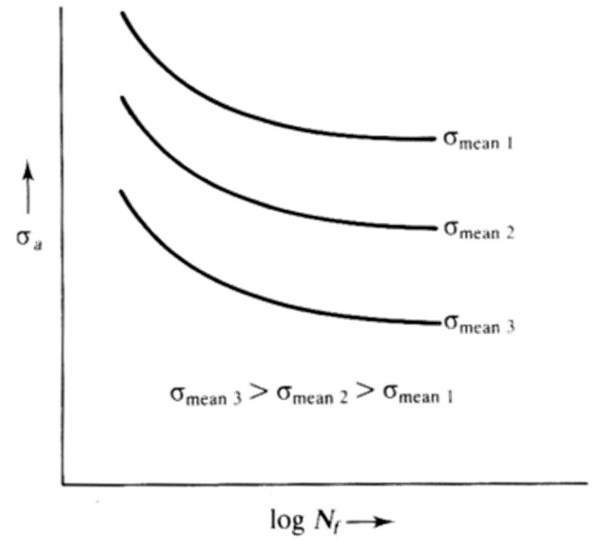

Figure 4: As the mean stress (σ_m) increases, the alternating stress (σ_a) to failure decreases.

In the majority of S-N curves presented in the literature, a common denominator is the use of stress amplitude (σ_a) to indicate the driving stress to failure with no mention of the mean stress (σ_m). This can cause a few problems for someone new to the field in trying to decipher the utility of the presented data since σ_a by itself doesn't paint a very complete picture. The reality is that if no other information is presented, then an S-N curve showing σ_a versus cycles (S-N) is always at a stress ratio of R = -1.0. The reason that some S-N curves typically only provide σ_a with no mention of stress ratio is that generating fatigue data is very expensive and requires a large data set for good statistical accuracy. One of easiest methods to generate fatigue data is the ASTM rotating beam test where a cylindrical test specimen is polished and mounted as shown on the left-hand side of Figure 5. This type of test is easy to operate, and since the stress ratio is fixed at R = -1.0, only one data set is generated. Hence, when no other information is presented, it is highly likely that the given data is at R = -1.0 and that the mean stress (σ_m) is zero. If the effect of σ_m is required, one needs to use a more complex setup, as shown on the right-hand side in Figure 5. As one can imagine, the resulting data set is much larger and more cumbersome to process.

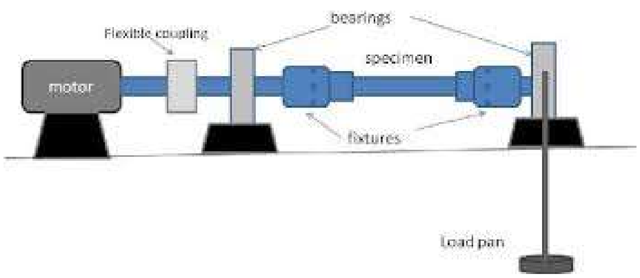

Figure 5: Classic ASTM rotating beam fatigue test providing a R = -1.0 and the more modern suite of fatigue test machines that can cycle the sample at nearly any stress ratio.

1.4 CORRECTING FOR MEAN STRESS

Traditionally, the industry has lacked arrays of instrumented testing machines as shown in Figure 5 and had to rely on the basic rotating beam test where the data was always at R = -1.0 and σ_m = 0.0. This presented a rather serious problem since it was well known that σ_m would significantly lower the fatigue life of the structure. To leverage the large and economical database of R = -1 fatigue data, several scientists over the years have developed empirical relationships that allow the correction of fatigue data at other σ_m values. The most popular of these corrections is the Modified-Goodman developed in the early 1900's (see Dowling's paper in Section 2, Suggested Readings). More recent work in the 1970's by Walker and Smith-Watson-Topper (SWT) provide formulas that are considered more accurate (see Dowling's paper). Figure 7 shows an application of the Goodman and SWT σ_m correction based on an original data set of R=-1.0. The process is to start with a base equation relating alternating stress $\sigma_a^{R=-1.0}$ to cycles to failure (N_f). The equation format is generic and provides a nice fit to most metallic fatigue data up to the point of the material's endurance limit. The mean stress corrections (σ_a^*) are then inserted as the corresponding value of $\sigma_a^{R=-1.0}$. As shown in Figure 6, as the mean stress increases from σ_m =0.0 (R=-1.0) to σ_m =σ_a/2 (R=0.0), the fatigue curve shifts downward as shown schematically in Figure 4.

$$\sigma_a^{R=-1.0} = \sigma_f \left(2N_f\right)^b$$

$$\sigma_f = 1758 \quad b = -0.098$$

$$Goodman\ \sigma_a^* = \left(\frac{\sigma_a \sigma_u}{\sigma_u - \sigma_m}\right)$$

$$SWT\ \sigma_a^* = \left((\sigma_m + \sigma_a/2)\sigma_a\right)^{0.5}$$

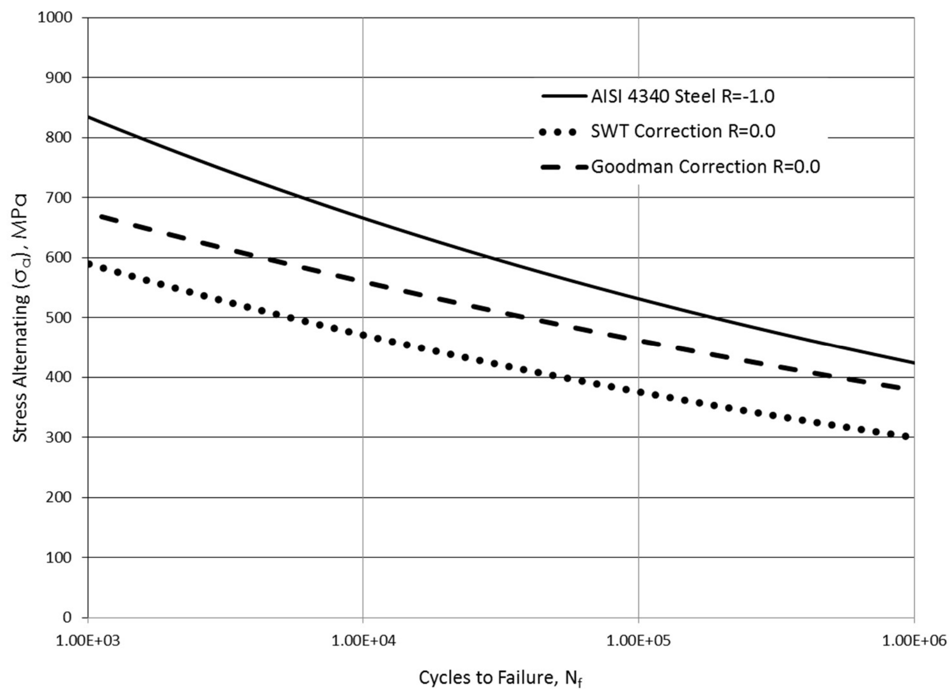

Figure 6: Starting with experimental data at R=-1, a mean stress correction to R=0.0 is done using the SWT and Goodman equations.

A more common way to present fatigue data obtained at different σ_m levels is by plotting the maximum stress (σ_m) against N_f. Figure 7 presents data from the MMPDS based on this format. When data is available at different σ_m levels (or stress ratios), one can obtain a better correction using the Walker equation by fitting the exponent to the data set. For example, the SWT equation uses a fixed exponent of 0.5 while the Walker equation presented in Figure 7 is derived from the data set and is 0.63. The fit to the data is obviously better using the Walker equation but for legacy reasons the Goodman correction is still prevalent.

$$LogN_f = 20.68 - 9.84Log(\sigma_{max}^*)$$

$$\sigma_u = 45ksi$$

$$Goodman\ \sigma_{max}^* = \left(\frac{\sigma_{max}\sigma_u(1-R)}{2\sigma_u - \sigma_{max}(1+R)}\right)$$

$$Walker\ \sigma_{max}^* = \sigma_{max}(1-R)^{0.63}$$

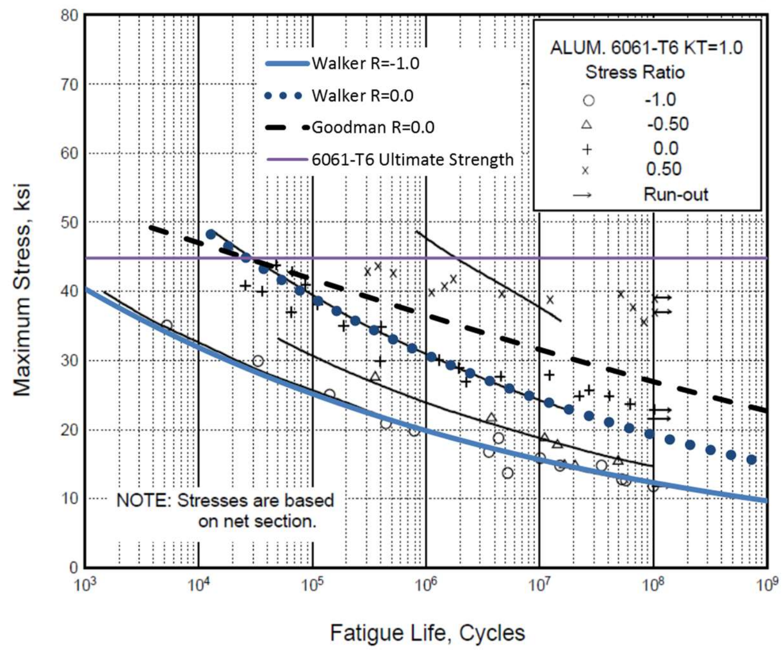

Figure 7: Working with data presented in the MMPDS, the Goodman and Walker mean stress correction is overlaid the experimental data for R = 0.0. One will note that the ultimate strength of the material σ_u=45 ksi is logically never exceeded by the experimental data but that the fitted curves will incorrectly bump above this limit.

Let's now solve a more fundamental fatigue analysis problem where the designer only has the most basic of mechanical steel property data, e.g., the ultimate strength (σ_u) of the steel and needs "quick and mostly accurate" assessment of fatigue life at a stress ratio R = 0.0. For a broad range of steels, it is reasonable to assume that at R=-1, one can say that the fatigue life N_f at 1,000 cycles is 0.9σ_u and that at N_f=1e6 it is 0.5 σ_u (see Bannantine, Section 2, Suggested Readings and note that this is only for polished samples and real structures rarely get this lucky). This curve is given in Figure 8 along with the

SWT mean stress correction at R=0.0. The correction is straightforward since it is only necessary to reformulate the SWT formula as given in Figure 8.

$$SWT\ \sigma_a^* = \sigma_{max}\left(\frac{1-R}{2}\right)^{0.5}$$

$$SWT\ \sigma_a^* = \sigma_{max}\left(\frac{1}{2}\right)^{0.5}$$

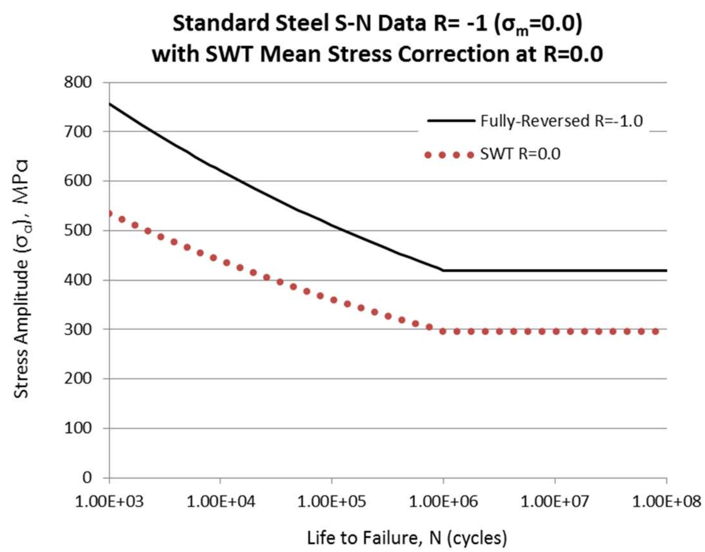

Figure 8: Starting with basic steel data where the σ_u=840 MPa, a simple fatigue curve can be constructed with an assumed stress ratio R=-1.0. The SWT correction is also given for R=0.0.

Just to close on this rather important subject and to start the introduction of load cycles, Figure 9 shows common examples of structures that often experience stress ratios of R=-1, R=0.0 and R=0.8. In summary, the process for correcting experimental R=-1 (σ_m=0.0) to different stress ratios ($\sigma_m \neq 0.0$) is not difficult but one should not lose sight that these corrections are based on empirical curve fits to experimental data and contain their own statistical uncertainties that are not easily quantifiable. If accuracy is paramount, then it is best to generate the data directly from coupons.

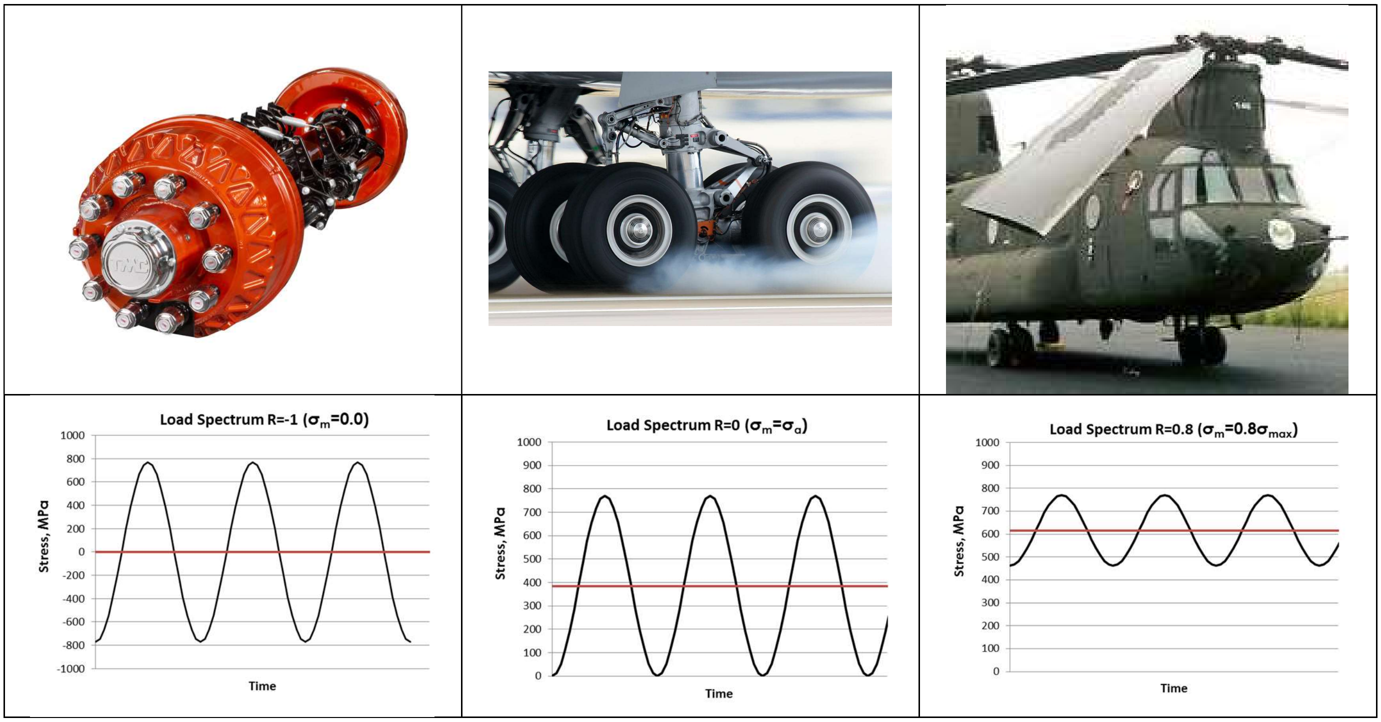

Figure 9: Examples of equipment where major components often experience stress ratios of R=-1.0 (truck hub), R=0.0 (aircraft landing gear) and R=0.8 (helicopter rotor blades).

1.5 WORKING CURVE OR STATISTICAL ANALYSIS OF FATIGUE DATA

Fatigue analysis is statistically messy. The accuracy of the process is sensitive to variation in material data, surface finish of the structure and calculated stress data.

Let’s start simple and take a look at some publically available fatigue data published in the MMPDS handbook as given in Figure 10 for 6061-T6 aluminum. Focus just on the fatigue data presented on the graph (the symbols mark experimentally determined data) and notice the scatter for each grouping. We’ll cover what each symbol means later on, but to gain a sense of how accurate the fatigue data might be, just focus on how the symbols move up and down in relation to their respective curve. It is not an exact fit and it is the general problem that everyone faces: how to safely use fatigue data to count cycles to failure for their design.

This concept of adjusting the provided fatigue data to create a statistically safe curve is commonly referred to as creating a “working curve”. The typical challenge is that the end user is often faced without having access to the raw experimental data or more likely, due to the expense of obtaining the fatigue data, only the most limited amount of data was collected. In an ideal situation, specific data would be created at the exact σ_m that the structure experiences and one could adjust the fatigue data as shown in Figure 11.

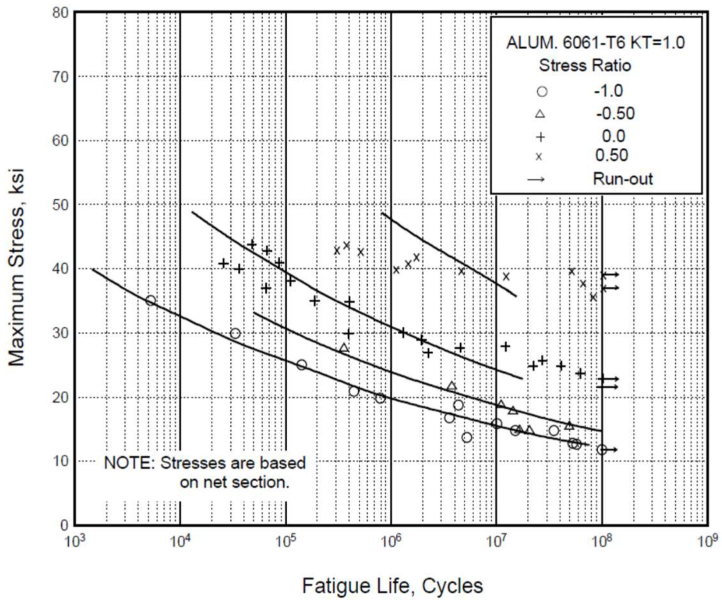

Figure 10: MMPDS data for 6061-T6 aluminum.

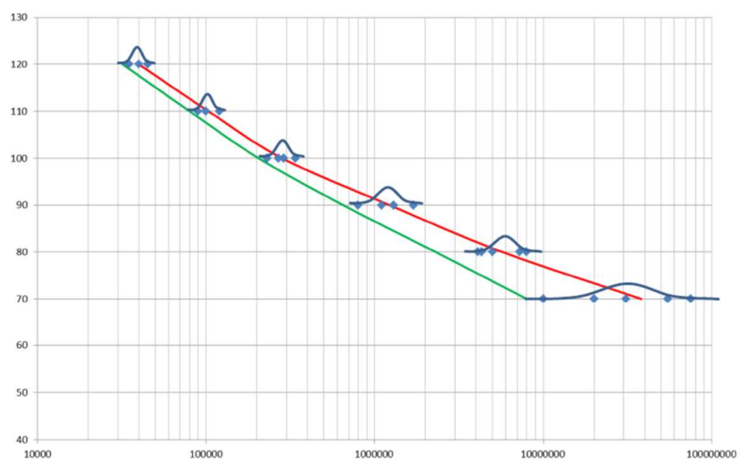

Figure 11: An example of perfectly obtained experimental data where the working curve (the line on the far left) can be created from a statistical fit of the raw data. The central line is termed the 50/50 line.

In statistical terms, the curves presented in the MMPDS and in S-N fitted data are termed 50/50 curves where one has a 50% chance that the N_f calculation will be within one standard deviation of error. The obvious challenge to this approach is that most fatigue data sets are limited and that statistical information is often lacking. Given this challenge, we have three general approaches:

- Increase the stress value (i.e., stress modification factor) used to calculate fatigue damage;
- Divide the number of calculated cycles to failure N_f by some scatter factor;
- Or a combination of the above two methods.

Figure 12 shows these approaches using the R=0.0 data within the MMPDS 7050 data set. This approach can also be directly applied to data at different stress ratios. It is not a straightforward subject since some companies prefer decreasing the number of calculated N_f by a certain number varying from 5 to 20 while other companies prefer to just increase the calculated stress by some factor prior to the calculation of N_f. To keep things interesting, some companies combine both techniques with a stress reduction at low cycles and a scatter factor at high cycles.

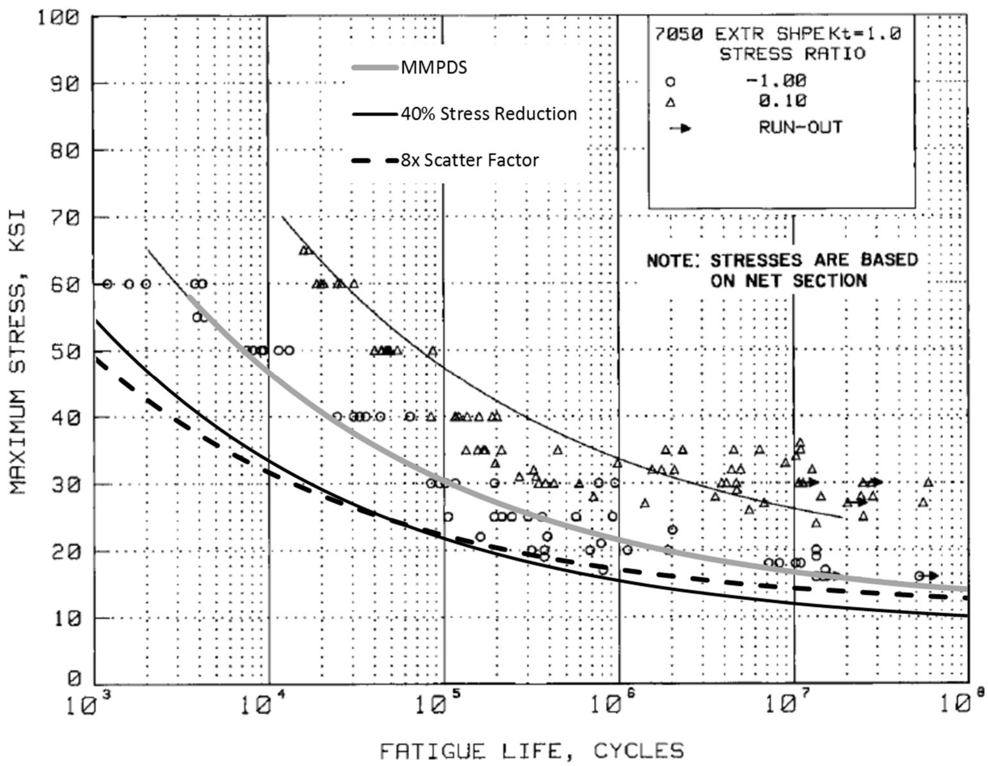

Figure 12: Examples of how to create a working curve using a stress reduction and an 8x scatter factor.

The reality of having to create a working curve is to account for the statistical uncertainty of the experimental fatigue data. As mentioned, the gold standard is just to create your own data and perform a complete statistical analysis on the data. However, this is often quite expensive since if the structure has load cycles that create different stress ratios, then the fatigue data set can get quite large as shown in Figure 10, and still not quite provide deep statistical accuracy.

Another reason why the industry and many certifying organizations insist upon a statistically safe working curve is that the base R=-1.0 data is often empirically adjusted to other stress ratios and that most end-users do not have complete statistical data sets. Given these reasons, most working curves represent a significantly decreased curve as that compared to the original fatigue data.

2. SUGGESTED READINGS

- DOT/FAA/AR-MMPDS-01, *Metallic Materials Properties Development and Standardization (MMPDS)*, Office of Aviation Research, Washington, D.C., January 2003.
- Bannantine, J., Comer, J. and Handrock, J., *Fundamentals of Metal Fatigue Analysis*, PrenticeHall (New Jersey), 1990.
- Dowling, N.E., *Mean Stress Effects in Stress-Life and Strain-Life Fatigue*, SAE, F2004/51, 2004.
- Faupel, J. H., Fisher, F. E., *Engineering Design*, Wiley-Interscience (New York), 1981.
- Forrest, P.G., *Fatigue of Metals*, Pergamon Press Ltd. (London), 1952.
- Juvinall, R. C., Marshek, K. M., *Fundamentals of Machine Component Design - 2nd Edition*, John Wiley and Sons (New York), 1991.
- Neuber, H., *Theory of Notch Stresses: Principle for Exact Stress Calculations*, Edwards (Ann Arbor,MI), 1946.
- Osgood, C. C., *Fatigue Design - 2nd Edition*, Pergamon Press (Oxford), 1982.
- Peterson, R. E., *Analytical Approach to Stress Concentration Effects in Aircraft Materials,* Technical Report 59-507, U.S. Air Force - WADC Symp. Fatigue Metals (Dayton, OH), 1959.
- Peterson, R. E., *Relation Between Life Testing and Conventional Tests of Materials*, Bulletin ASTM No. 133, March, 1945.
- Sines, G., Waisman, J. L., *Metal Fatigue*, McGraw-Hill (New York), 1959.
- Graham, J.A., *Fatigue Design Handbook*, Vol. 4, SAE (Warrendale, PA), 1968.
- Stephens, R.I. et al., *Metal Fatigue in Engineering, 2nd Ed.*, 2001.
- Madayag, A.F., *Metal Fatigue: Theory and Design*, 1969
- Dowling, N.E., *Mechanical Behavior of Materials*, 1993

3. AUTHOR BIOGRAPHY

Brian Reiling
Partner / Principal Engineer
Endeavor Analysis, LLC

Brian Reiling graduated from University of Portland in Mechanical Engineering. He has over 20 years of aerospace structural analysis experience working on Airbus, Boeing, Bombardier, and Sikorsky aircraft along with several others. Currently a founding partner at Endeavor Analysis the focus is on landing gear, hydraulic actuation, and test rig design and analysis. Endeavor Analysis is considered one of the premier independent consultants for landing gear structural analysis. For FE analysis Endeavor uses Simcenter Nastran and ADINA pre- and post-processed with FEMAP.

George Laird, Ph.D., P.E.
Principal Mechanical Engineer
Predictive Engineering, Inc.

George Laird earned his doctorate in philosophy while a scientist with the U.S. Bureau of Mines (USBM). In this prior position, he was an expert in the mechanics of materials for the mining and mineral processing industries. After the closure of the USBM in 1996, Dr. Laird started Predictive Engineering with a focus on finite element analysis. Over the years, Predictive has steadily grown and it is now recognized as one of the more capable engineering service firms that can perform analyses on structures ranging from satellites, submarines, ASME vessels, large transmissions for the oil and gas industry, turbine burst containments, large-strain impact analyses or something as simple the dynamic simulation of French fries moving down a conveyor. For analysis, Predictive uses Simcenter Nastran, LS-DYNA and STAR-CCM+.