

The NAFEMS Benchmark Challenge Volume 1

Angus Ramsay

08/Nov/19

Contents

1	Introduction	5
2	CS1: Design of a Bandsaw Pulley	7
3	CS2: Torsion of Curved Beams	9
4	CS3: Design of Steel Silos	11
5	CS4: Fatigue Assessment in Nuclear Piping	13
5.1	From college theory to an actual engineering problem	13
5.1.1	Tips and tricks	15
5.2	Case study: reactors, pipes and fatigue	16
5.2.1	Nuclear reactors	17
5.2.2	Pressurised pipes	18
5.2.3	Fatigue	18
5.3	Solid mechanics, or what we are taught at college	20
5.3.1	The stress tensor	20
5.3.2	An infinitely-long pressurised pipe	21
5.4	Finite elements, or solving actual problems	24
5.4.1	The name of the game	24
5.5	Piping in nuclear reactors	25
5.5.1	Thermal transient	26
5.5.2	Seismic loads	30
5.5.3	Linearity (not yet linearisation)	31
5.5.4	ASME stress linearisation (not linearity!)	38
5.6	The infinite pipe revisited after college	40
5.7	Adding complexity: the truth is out there	42
5.7.1	Two (or more) materials	42

CONTENTS

5.7.2	A parametric tee	45
5.7.3	Bake, shake and break	49
5.8	Cumulative usage factors	56
5.8.1	In air (ASME's basic approach)	57
5.8.2	In water (NRC's extension)	59
5.9	Conclusions	63
5.9.1	Online stuff	66
6	Concluding Remarks	67

Chapter 1

Introduction

CHAPTER 1. INTRODUCTION

Chapter 2

CS1: Design of a Bandsaw Pulley

Chapter 3

CS2: Torsion of Curved Beams

Chapter 4

CS3: Design of Steel Silos

Chapter 5

CS4: Fatigue Assessment in Nuclear Piping

5.1 From college theory to an actual engineering problem

First of all, please take this text as a written chat between you and me, i.e. an average engineer that has already taken the journey from college to performing actual engineering works using finite element analysis and has something to say about it. Picture yourself in a coffee bar, talking and discussing concepts and ideas with me. Maybe needing to go to a blackboard (or notepad?). Even using a tablet to illustrate some three-dimensional results. But always as a chat between colleagues.

Please also note that I am not a mechanical engineer, although I shared many undergraduate courses with some of them. I am a nuclear engineer with a strong background in mathematics and computer programming. I went to college between 2002 and 2008. Probably a lot of things have changed since then—at least that is what these “millennial” guys and girls seem to be boasting about—but chances are we all studied solid mechanics and heat transfer with a teacher using a piece of chalk on a blackboard while we as students wrote down notes with pencils on paper sheets. And there is really not much that one can do with pencil and paper regarding mechanical analysis. Any actual case worth the time of an engineer needs to be more complex than an ideal canonical case with analytical solution.

Whether you are a student or a seasoned engineer with many years of experience, you might recall from first year physics courses the introduction of the [simple pendulum](#) as a case study (fig. 5.1a). You learned that the period does not depend on

the hanging mass because the weight and the inertia exactly cancelled each other. Also, that [Galileo](#) said (and [Newton](#) proved) that for small oscillations the period does not even depend on the amplitude. Someone showed you why it worked this way: because if $\sin \theta \approx \theta$ then the motion equations converge to an [harmonic oscillator](#). It might have been a difficult subject for you back in those days when you were learning physics and calculus at the same time.

But it was probably after college, say when you took your first son to a swing on a windy day (fig. 5.1b), that you were faced with a real pendulum worth your full attention. The very same distance between what I imagined studying a pendulum was and what I saw that day at the swing (namely that the period *does* depend on the hanging mass) is the same distance between the mechanical problems studied in college and the actual cases encountered during a professional engineer's lifetime (fig. 5.2). In this regard, I am referring only to technical issues. The part of dealing with clients, colleagues, bosses, etc. which is definitely not taught in engineering schools (you can get a heads up in business schools, but again it would be a theoretical pendulum) is way beyond the scope of both this article and my own capacities.




Figure 5.1: A [simple pendulum from college physics courses](#) and a real-life pendulum. Hint: the swing's period *does* depend on the hanging mass as shown in <https://youtu.be/Q-LKK4A20zA>

Again, take all this information as coming from a fellow that has already taken such a journey from college's pencil and paper to real engineering cases involving complex numerical calculations. And developing, in the meantime, both an actual working finite-element [back-end](#) and [front-end](#) from scratch.

5.1. FROM COLLEGE THEORY TO AN ACTUAL ENGINEERING PROBLEM

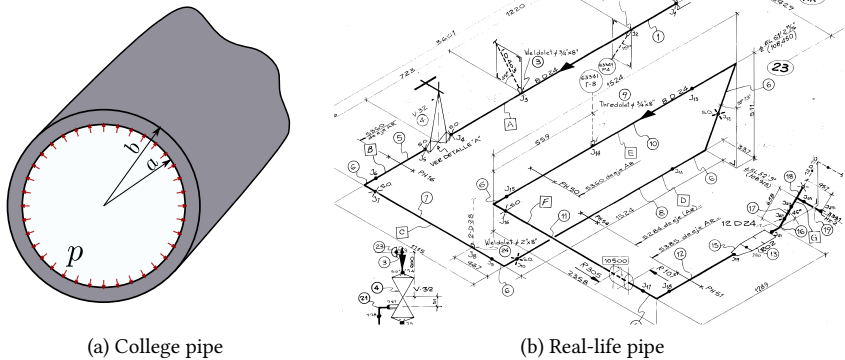


Figure 5.2: An infinitely-long pressurised thick pipe as taught in college and an isometric drawing of a section of a real-life piping system. The details of the isometric are not important individually but they are included to emphasize the fact that a real engineering drawing is far more complex and has far more information than what a professor can draw in a class blackboard.

5.1.1 Tips and tricks

There are some useful tricks that come handy when trying to solve a mechanical problem. Throughout this text, I will try to tell you some of them.

One of the most important ones to use your *imagination*. You will need a lot of imagination to “see” what it is actually going on when analysing an engineering problem. This skill comes from my background in nuclear engineering where I had no choice but to imagine a [positron-electron annihilation](#) or an [spontaneous fission](#). But in mechanical engineering, it is likewise important to be able to imagine how the loads “press” one element with the other, how the material reacts depending on its properties, how the nodal displacements generate stresses (both normal and shear), how results converge, etc.

This journey will definitely need your imagination. We will peek a little bit into equations, numbers, plots, schematics, CAD geometries, 3D views, etc. Still, when the theory says “thermal expansion produces normal stresses” you have to picture in your head three little arrows pulling away from the same point in three directions, or whatever mental picture you have about what you understand thermally-induced stresses are. What comes to your mind when someone says that out of the nine elements of the [stress tensor](#) (sec. 5.3.1) there are only six that are independent? Whatever it is, try to practice that kind of graphical thoughts with every new concept. Nevertheless, there will be particular locations of the text where

imagination will be most useful. I will bring the subject up then and again.

Another heads up is that we will be digging into some mathematics. Probably they would be simple and you would deal with them very easily. But chances are you do not like equations. No problem! Just ignore them for now. Read the text skipping them, it should work as well. Anyhow, here comes another experience tip: keep exercising mathematics. You have used [differences of squares](#) in high school, didn't you?. You know (or at least knew) how to [integrate by parts](#). Do you remember what [Laplace transforms](#) are used for? Once in a while, perform a division of polynomials using [Ruffini's rule](#). Or compute the second [derivative of the quotient of two functions](#). Whatever. It should be like doing crosswords on the newspaper. Grab those old physics college books and solve the exercises at the end of each chapter. All the effort will pay off later on.

One final comment: throughout the text I will be referring to “your favourite FEM program.” I bet you do have one. Mine is [CAEplex](#) (it works on top of [Fino](#), which is free and open source). We will be using it to perform some tests and play a little bit. And we will also use it to think about what it means to use a FEM program to generate results that will eventually end up in a written project with your signature. Keep that in mind.

5.2 Case study: reactors, pipes and fatigue

Piping systems in sensitive industries like nuclear or oil & gas should be designed and analysed following the recommendations of an appropriate set of codes and norms, such as the [ASME Boiler and Pressure Vessel Code](#). This code of practice was born during the late 19th century, before finite-element methods for solving partial differential equations were even developed. And much longer before they were available for the general engineering community. Therefore, much of the code assumes design and verification is not necessarily performed numerically but with paper and pencil (yes, like in college). However, it still provides genuine guidance in order to ensure pressurised systems behave safely and properly without needing to resort to computational tools. Combining finite-element analysis with the ASME code gives the cognisant engineer a unique combination of tools to tackle the problem of designing and/or verifying pressurised piping systems.

In the years following [Enrico Fermi's](#) demonstration that a self-sustainable [fission reaction](#) chain was possible (actually, in fact, after WWII was over), people started to build plants in order to transform the energy stored within the atoms nuclei into usable electrical power. They quickly reached the conclusion that high-pressure heat exchangers and turbines were needed. So they started to follow codes

of practise like the aforementioned ASME B&PVC. They also realised that some requirements did not fit the needs of the nuclear industry. But instead of writing a new code from scratch, they added a new chapter to the existing body of knowledge: the celebrated [ASME Section III](#).

After further years passed by, engineers (probably the same people that forked section III) noticed that fatigue in nuclear power plants was not exactly the same as in other piping systems. There were some environmental factors directly associated to the power plant that were not taken into account by the regular ASME code. Again, instead of writing a new code from scratch, people decided to add correction factors to the previously-amended body of knowledge. This is how (sometimes) knowledge evolves, and it is this kind of complexities that engineers are faced with during their professional lives. We have to admit it, it would be a very difficult task to re-write everything from scratch every time something changes.




Figure 5.3: Three-dimensional CAD model for the piping system of fig. 5.2b

5.2.1 Nuclear reactors

In each of the countries that have at least one nuclear power plant there exists a national regulatory body who is responsible for allowing the owner to operate the reactor. These operating licenses are time-limited, with a range that can vary from 25 to 60 years, depending on the design and technology of the reactor. Once expired, the owner might be entitled to an extension, which the regulatory authority can accept provided it can be shown that a certain (and very detailed) set of safety criteria are met. One particular example of requirements is that of fatigue in pipes, especially those that belong to systems that are directly related to the reactor safety.

5.2.2 Pressurised pipes

How come that pipes are subject to fatigue? Well, on the one hand and without getting into many technical details, the most common nuclear reactor design uses liquid water as coolant and moderator. On the other hand, nuclear power plants cannot by-pass the thermodynamics of the [Carnot cycle](#), and in order to maximise the efficiency of the conversion of the energy stored in the uranium nuclei into electricity we need to reach temperatures as high as possible. So, if we want to have liquid water in the core as hot as possible, we need to increase the pressure. The limiting temperature and pressure are given by the [critical point of water](#), which is around 374°C and 22 MPa. It is therefore expected to have temperatures and pressures near those values in many systems of the plant, especially in the primary circuit and those that directly interact with it, such as pressure and inventory control system, decay power removal system, feedwater supply system, emergency core-cooling system, etc.

Nuclear power plants are not always working at 100% of their maximum power capacity. They need to be maintained and refuelled, they may undergo operational (and some incidental) transients, they might operate at a lower power due to load following conditions, etc. These transient cases involve changes both in temperatures and in pressures that the pipes are subject to, which in turn give rise to changes in the stresses within the pipes. As the transients are postulated to occur cyclically during a number of times during the life-time of the plant (plus its extension period), mechanical fatigue in these piping systems may arise, especially at the interfaces between materials with different thermal expansion coefficients.

An important part of the analysis that almost always applies to nuclear power plants but usually also to other installations is the consideration of a possible seismic event. Given a postulated design earthquake, both the civil structures and the piping system itself need to be able to withstand such a load, even if it occurs at the moment of highest mechanical demand during one of the operational transients.

5.2.3 Fatigue

Mechanical systems can fail due to a wide variety of reasons. The effect known as fatigue can create, migrate and grow microscopic cracks at the atomic level, called [dislocations](#). Once these cracks reach a critical size, then the material fails catastrophically even under stresses much lower than the tensile strength limits. There are not complete mechanistic models from first principles which can be used in general situations, and those that exist are very complex and hard to use. There are two main ways to approach practical fatigue assessment problems using experimental

data very much like the [Hooke's Law](#) experiment:

1. stress life, or
2. strain life

The first one is suitable for cases where the loads are nowhere near the yield stress of the material. When plastic deformation is expected to occur, strain-life methods ought to be employed.

For the case study, as the loads come principally from operational loads, the ASME stress-life approach should be used. The stress amplitude of a periodic cycle can be related to the number of cycles where failure by fatigue is expected to occur. For each material, this dependence can be computed using normalised tests and a family of “fatigue curves” like the one depicted in fig. 5.4 (also called $S-N$ curve) for different temperatures can be obtained.




Figure 5.4: A fatigue or $S-N$ curve for two steels.

It should be noted that the fatigue curves are obtained in a particular load case, namely purely-periodic and one-dimensional, which cannot be directly generalised to other three-dimensional cases. Also, any real-life case will be subject to a mixture of complex cycles given by a stress time history and not to pure periodic conditions. The application of the curve data implies a set of simplifications and assumptions that are translated into different possible “rules” for composing real-life cycles. There also exist two safety factors which increase the stress amplitude and reduce the number of cycles respectively. All these intermediate steps render the analysis

of fatigue into a conservative computation scheme. Therefore, when a fatigue assessment performed using the fatigue curve method arrives at the conclusion that “fatigue is expected to occur after ten thousand cycles” what it actually means is “we are sure fatigue will not occur before ten thousand cycles, yet it may not occur before one hundred thousand or even more.”

5.3 Solid mechanics, or what we are taught at college

So, let us start our journey. Our starting place: undergraduate solid mechanics courses. Our goal: to obtain the internal state of a solid subject to a set of movement restrictions and loads (i.e. to solve the solid mechanics problem). Our first step: [Newton’s laws of motion](#). For our purposes, we can recall them here like this:

1. a solid is in equilibrium if it is not moving in at least one inertial coordinate system,
2. in order for a solid not to move, the sum of all the forces ought to be equal to zero, and
3. for every external load there exists an internal reaction with the same magnitude but opposite direction.

We have to accept that there is certain intellectual beauty when complex stuff can be expressed in such simple terms. Yet, from now on, everything can be complicated at will. We can take the mathematical path like [D’Alembert](#) and his virtual displacements ideas (in his mechanical treatise, he brags that he does not need to use a single figure throughout the book). Or we can go graphical following [Culmann](#). Or whatever other logic reasoning to end up with a set of actual equations which we need to solve in order to obtain engineering results.

5.3.1 The stress tensor

In any case, what we should understand (and imagine) is that external forces lead to internal stresses. And in any three-dimensional body subject to such external loads, the best way to represent internal stresses is through a 3×3 *stress tensor*. This is the first point in which we should not fear mathematics. Trust me, it will pay back later on.

Does the term *tensor* scare you? It should not. A **tensor** is a general mathematical object and might get complex when dealing with many dimensions (as those encountered in weird stuff like [string theory](#)), but we will stick here to second-order tensors. They are slightly more complex than a vector, and I assume you are not

afraid of vectors, are you? If you recall fresh-year algebra courses, a **vector** somehow generalises the idea of a **scalar** in the following sense: a given vector \mathbf{v} can be projected into any direction \mathbf{n} to obtain a scalar p . We call this scalar p the “projection” of the vector \mathbf{v} in the direction \mathbf{n} . Well, a tensor can be also projected into any direction \mathbf{n} . The difference is that instead of a scalar, a vector is now obtained.

Let me introduce then the three-dimensional **stress tensor** (a.k.a **Cauchy tensor**):

$$\begin{bmatrix} \sigma_x & \tau_{xy} & \tau_{xz} \\ \tau_{yx} & \sigma_y & \tau_{yz} \\ \tau_{zx} & \tau_{zy} & \sigma_z \end{bmatrix}$$

It looks (and works) like a regular 3×3 matrix. Some brief comments about it:

- The σ s are normal stresses, i.e. they try to stretch or tighten the material.
- The τ s are shear stresses, i.e. they try to twist the material.
- Due to rotational equilibrium requirements the conjugate shear stresses should be equal: $\tau_{xy} = \tau_{yx}$, $\tau_{yz} = \tau_{zy}$, and $\tau_{zx} = \tau_{xz}$. Therefore, the stress tensor is symmetric i.e. there are only six independent elements.
- The elements of the tensor depend on the orientation of the coordinate system.
- There exists a particular coordinate system in which the stress tensor is diagonal, i.e. all the shear stresses are zero. In this case, the three diagonal elements are called the **principal stresses**, which happen to be the three **eigenvalues** of the stress tensor. The basis of the coordinate system that renders the tensor diagonal are the **eigenvectors**.

What does this all have to do with mechanical engineering? Well, once we know what the stress tensor is for every point of a solid, in order to obtain the internal forces per unit area acting in a plane passing through that point and with a normal given by the direction \mathbf{n} , all we have to do is “project” the stress tensor through \mathbf{n} . In plain simple words:

If you can compute the stress tensor at each point of your geometry, then... Congratulations! You have solved the solid mechanics problem.

5.3.2 An infinitely-long pressurised pipe

Let us proceed to our second step, and consider the infinite pipe subject to uniform internal pressure already introduced in fig. 5.2a. Actually, we are going to solve the mechanical problem on an infinite hollow cylinder, which looks like pipe. This case is usually tackled in college courses, and chances are you already solved it. In

fact, the first (and simpler) problem is the “thin cylinder problem.” Then, the “thick cylinder problem” is introduced (the one we solve below), which is slightly more complex. Nevertheless, it has an analytical solution.¹ For the present case, let us consider an infinite pipe (i.e. a hollow cylinder) of internal radius a and external radius b with uniform mechanical properties—Young’s modulus E and Poisson’s ratio ν —subject to an internal uniform pressure p .

What follows is more or less what we are taught in school: some equations with a brief explanation of the results. And then we move on to the next subject.

5.3.2.1 Displacements

Remember that when any solid body is subject to external forces, it has to react in such a way to satisfy the equilibrium conditions. The way solids do this is by deforming a little bit in such a way that the whole body acts as a compressed (or elongated) spring balancing the load. So it is worth to ask how a pressurised pipe deforms to counteract the internal pressure p .

- There are no longitudinal displacements u_t because the pipe is infinite in the axial direction.
- There are no azimuthal displacements u_θ because the pipe is fully symmetric around the axis.
- There are only radial displacements u_r and they depend only on the radial coordinate r and not on the axial position z or on the azimuthal angle θ . These displacements are

$$u_r(r) = p \cdot \frac{1 + \nu}{E} \cdot \frac{a^2}{b^2 - a^2} \cdot \left[1 - 2\nu + \frac{b^2}{r^2} \right] \cdot r$$

What does this mean? Well, that overall the whole pipe expands a little bit radially with the inner face being displaced more than the external surface (use your imagination!). How much?

1. linearly with the pressure, i.e. twice the pressure means twice the displacement, and
2. inversely proportional to the Young’s Modulus E divided by $1 + \nu$, i.e. the more resistant the material, the less radial displacements.

¹A detailed analysis of the analytical solution and how results obtained with finite elements converge with respect to the mesh size can be found in the report “On convergence of linearized stresses in an infinite pipe computed using the finite element method.” See sec. 5.9.1.

That is how an infinite pipe withstands internal pressure. And that is what we are taught in college, which is actually true by the way!

5.3.2.2 Stresses

As the solid is deformed, that is to say that different parts are relatively displaced one from another, strains and stresses appear. When seen from a cylindrical coordinate system, the stress tensor (recall sec. 5.3.1) has these features.

- There are no shear stresses as there is no bending due to the fact that the pipe is infinite (so it cannot bend in the axial direction) and azimuthally symmetric (there is no particular direction so circles must remain circles).
- The normal stresses depend only on the radial coordinate r and are

- the radial stress $\sigma_r(r) = \frac{p \cdot a^2}{b^2 - a^2} \cdot \left(1 - \frac{b^2}{r^2}\right)$
- the azimuthal (or hoop) stress $\sigma_\theta(r) = \frac{p \cdot a^2}{b^2 - a^2} \cdot \left(1 + \frac{b^2}{r^2}\right)$, and
- the longitudinal (or axial) stress $\sigma_l(r) = 2\nu \cdot \frac{p \cdot a^2}{b^2 - a^2}$

We can note that

1. The stresses do not depend on the mechanical properties E and ν of the material (the displacements do).
2. All the stresses are linear with the pressure p , i.e. twice the pressure means twice the stress.
3. The axial stress is uniform and does not depend on the radial coordinate r .
4. As the stress tensor is diagonal, these three stresses happen to also be the principal stresses.

That is all what we can say about an infinite pipe with uniform material properties subject to an uniform internal pressure p . If

- the pipe was not infinite (say any real pipe that has to start and end somewhere), or
- the cross-section of the pipe was not constant along the axis (say there is an elbow or even a reduction), or
- there was more than one pipe (say there is a tee), or
- the material properties were not uniform (say the pipe does not have an uniform temperature but a distribution), or
- the pressure was not uniform (say because there is liquid inside and its weight cannot be neglected),

then we would no longer be able to fully solve the problem with paper and pencil and draw all the conclusions above. However, at least we have a start because we know that if the pipe is finite but long enough or the temperature is not uniform but almost, we still can use the analytical equations as approximations. After all, [Enrico Fermi](#) managed to reach criticality in the [Chicago Pile-1](#) with paper and pencil. But what happens if the pipe is short, there are branches and temperature changes like during a transient in a nuclear reactor? Well, that is why we have finite elements. And this is where what we learned at college pretty much ends.

5.4 Finite elements, or solving actual problems

Besides infinite pipes (both thin and thick), spheres and a couple of other geometries, there are no other cases for which we can obtain analytical expressions for the elements of the stress tensor. To get results for a solid with real engineering interest, we need to use numerical methods to solve the equilibrium equations. It is not that the equations are hard *per se*. It is that the mechanical parts we engineers like to design (which are of course more complex than cylinders and spheres) are so intricate that render simple equations into monsters which are unsolvable with pencil and paper. Hence, finite elements enter into the scene.

5.4.1 The name of the game

But before turning our attention directly into finite elements (and leaving college, at least undergraduate) it is worth some time to think about other alternatives. Are we sure we are tackling your problems in the best possible way? I mean, not just engineering problems. Do we take a break, step back for a while and see the whole picture looking at all the alternatives so we can choose the best cost-effective one?

There are literally dozens of ways to numerically solve the equilibrium equations, but for the sake of brevity let us take a look at the three most famous ones. Coincidentally, they all contain the word “finite” in their names. We will not dig into them, but it is nice to know they exist. We might use

1. [Finite differences](#)
2. [Finite volumes](#)
3. [Finite elements](#)

Before proceeding, I would like to make two comments about common nomenclature. The first one is that if we exchanged the words “volumes” and “elements” in all the written books and articles, nobody would notice the difference. There is

nothing particular in both theories that can justify why FVM use “volumes” and FEM use “elements”. Actually volumes and elements are the same geometric constructions. As far as I know, the names were randomly assigned.

The second one is more philosophical and refers to the word “simulation” which is often used to refer to solving a problem using a numerical scheme such as the finite element method. [I am against at using this word for this endeavour](#). The term simulation has a connotation of both “pretending” and “faking” something, that is definitely not what we are doing when we solve an engineering problem with finite elements. Sure, there are some cases in which we simulate, such as using the [Monte Carlo method](#) (originally used by Fermi as an attempt to understand how neutrons behave in the core of nuclear reactors). But when solving deterministic mechanical engineering problems I would rather say “modelling” than “simulation.”

5.5 Piping in nuclear reactors

So we need to address the issue of fatigue in nuclear reactor pipes that

1. are not infinite and have cross-section changes, branches, valves, etc.
2. are made of different materials,
3. are fixed at different locations through piping supports,
4. are subject to
 - a. pressure transients,
 - b. heat transients, and
 - c. seismic loads.

As a nuclear engineer, I learned (theoretically in college but practically after college) that there are some models that let you see some effects and some that let you see other effects (please [say “modelling” not “simulation.”](#)). And even if, in principle, it is true that more complex models should allow you compute more stuff, they definitely might show you nothing at all if the model is so big and complex that it does not fit into a computer (say because it needs hundreds of gigabytes of RAM to run) or because it takes more time to compute than you may have before the final report is expected.

First of all, we should note that we need to solve

- i. the heat transfer equation to get the temperature distribution within the pipes,
- ii. the natural frequencies and oscillation modes of the piping system to obtain the pseudo-accelerations generated by the design earthquake, and finally

- iii. the elastic problem to obtain the stress tensor needed to compute the alternating stress to enter into the fatigue curve.

So for each time t of the operational transient, the pipes are subject to


- a. an uniform internal pressure $p_i(t)$ that depends on time,
- b. a non-uniform internal temperature $T_i(t)$ that gives rise to a non-trivial time-dependent temperature distribution $T(\mathbf{x}, t)$ in the bulk of the pipes, and
- c. internal distributed forces $\mathbf{f} = \rho \cdot \mathbf{a}$ at those times where the design earthquake is assumed to occur.

5.5.1 Thermal transient


Let us invoke our imagination once again. Assume in one part of the transients the temperature of the water inside the pipes falls from say 300°C down to 100°C in a couple of minutes, stays at 100°C for another couple of minutes and then gets back to 300°C. The temperature within the bulk of the pipes changes as times evolves. The internal wall of the pipes follow the transient temperature (it might be exactly equal or close to it through the [Newton's law of cooling](#)). If the pipe was in a state of uniform temperature, the ramp in the internal wall will start cooling the bulk of the pipe creating a transient thermal gradient. Due to thermal inertia effects, the temperature can have a non-trivial dependence when the ramps start or end (think about it!). So we need to compute a real transient heat transfer problem with convective boundary conditions because any other usual tricks like computing a sequence of steady-state computations for different times would not be able to recover these non-trivial distributions.

Remember the main issue of the fatigue analysis in these systems is to analyse what happens around the location of changes of piping classes where different materials (i.e. different expansion coefficients) are present, potentially causing high stresses due to differential thermal expansion (or contraction) under transient conditions. Therefore, even though we are dealing with pipes we cannot use beam or circular shell elements, because we need to take into account the three-dimensional effects of the temperature distribution along the pipe thickness. And even if we could, there are some tees that connect pipes with different nominal diameters that have a non-trivial geometry, such as the weldolet-type junction shown in [figs. 5.5a, 5.5b](#). In this case, there are a number of SCLs (Stress Classification Lines) that go through the pipe's thickness at both sides of the material interface as illustrated in [fig. 5.6](#). It is in these locations that fatigue is to be evaluated.

On the one hand, a reasonable number of nodes (it is the number of nodes that defines the problem size, not the number of elements) in order to get a decent grid



(a) Overall view



(b) Unstructured tetrahedra-based grid

Figure 5.5: CAD model of a piping system with a 3/4-inch weldolet-type fork (stainless steel) from a main 12-inch pipe (carbon steel)




Figure 5.6: Location of six SCLs defined to analyse fatigue around a junction.

is around 200k for each system. On the other hand, solving a couple of dozens of transient heat transfer problems (which we cannot avoid due to the large thermal inertia of the pipes) during a few thousands of seconds over a couple hundred of thousands of nodes might take more time and storage space to hold the results than we might expect.

There is a wonderful essay by [Isaac Asimov](#) called “[The Relativity of Wrong](#)” where he introduces the idea that even if something cannot be computed exactly, there are different levels of error. For instance, believing that the Earth is a sphere is less wrong than believing that the Earth is flat, but wrong nonetheless, since it really deviates from a perfect sphere and resembles more an oblate spheroid.

We can then merge this idea by Asimov with an adapted version of the [Saint-Venant’s principle](#) and note that the detailed transient temperature distribution is important only around the location of the SCLs. We can then make an engineering approximation and

1. compute the transient thermal problem using a reduced mesh around the SCLs, and
2. assume the part of the full system which is not contained in the reduced mesh is at an uniform (though not constant) temperature equal to the average of the inner and outer temperatures at each side of the reduced mesh.




Figure 5.7: An example case where the SCLs are located around the junction between stainless-steel valves and carbon steel pipes at both sides of the material interface in the vertical plane both at the top and at the bottom of the pipe.

As an example, let us consider the system depicted in fig. 5.7 where there is a stainless-carbon steel interface at the discharge of the valves. Instead of solving the transient heat-conduction problem with the internal temperature of the pipes equal to the temperature of the water in the reference transient condition of the power plant and an external condition of natural convection to the ambient temperature in the whole mesh, a reduced model consisting of half of one of the two valves and a small length of the pipes at both the valve inlet and outlet is used. Once the temperature distribution $\hat{T}(\mathbf{x}, t)$ for each time is obtained in the reduced mesh (fig. 5.8, which has the origin at the centre of the valve), the actual temperature distribution $T(\mathbf{x}, t)$ is computed by an algebraic generalisation of $\hat{T}(\mathbf{x}, t)$ in the full coordinate system. As stated above, those locations which are not covered by the reduced model are generalised with a time-dependent uniform temperature which is the average of the inner and outer temperatures at the inlet and outlet of the reduced mesh.

Note that there is no need to have a one-to-one correspondence between the elements from the reduced mesh with the elements from the original one. Actually, the reduced mesh contains first-order elements whilst the former has second-order elements. Also the grid density is different. Nevertheless, the finite-element solver *Fino* used to solve both the heat and the mechanical problems, allows to read functions of space and time defined over one mesh and continuously evaluate and use them into another one even if the two grids have different elements, orders or even




Figure 5.8: Mesh of half-a-symmetric-valve refined around the interface where the transient heat conduction problem is solved.

dimensions. In effect, in the system from fig. 5.3 the material interface is between a orifice plate made in stainless steel that is welded to a carbon-steel pipe (fig. 5.9). The thermal problem can be modelled using a two-dimensional axi-symmetric grid and then generalised to the full three-dimensional mesh using the algebraic manipulation capabilities provided by *Fino* (actually by *wasora*) as shown in fig. 5.9.

5.5.2 Seismic loads

Every nuclear power plant is designed to withstand earthquakes. Of course, not all plants need the same level of reinforcements. Those built in large quiet plains will be, seismically speaking, cheaper than those located in geologically active zones. Keep in mind that all the 54 Japanese nuclear power plants did structurally resist the 2011 earthquake, and all of the reactors were safely shut down. What actually happened in *Fukushima* is that one hour after the main shake, a 14-metre tsunami splashed on the coast, jumping over the 9-metre defences and flooding the emergency *Diesel generators* that provided power to the pumps in charge of removing the remaining *decay power* from the already-stopped *reactor core*.

Since the computation of the loads that a certain earthquake gives rise to would add a significant amount of complexity to this already-complex case study, the details are skipped. In any case, we need to compute the first natural modes of oscillation of the piping section under study (fig. 5.10) and then, after some cumbersome maths, obtain an statically-equivalent load distribution. In effect, fig. 5.11 shows a sample distribution of acceleration distribution within a certain piping system which, when multiplied by the density, give a load distribution which is statically




Figure 5.9: Temperature distribution for a certain time within the transient computed on a reduced two-dimensional axi-symmetric mesh modelling half the orifice plate and a length of the carbon pipe and generalisation of the temperature to the full three-dimensional grid.

equivalent to the dynamic solicitations created by the earthquake. The ASME code says that these accelerations ought to be applied twice: once with the original sign and once with all the elements with the opposite sign, as SRSS “looses signs.” The application of each of these equivalent loads should last two seconds in the original time domain.

5.5.3 Linearity (not yet linearisation)

Even though we did not yet discuss it in detail, we want to solve an elastic problem subject to an internal pressure condition, with a non-uniform temperature distribution that leads to both thermal stresses and variations in the mechanical properties of the materials. And as if this was not enough, we want to add during a couple of seconds a statically-equivalent distributed load arising from a design earthquake. This last point means that at the transient instant where the stresses (from the fatigue’s point of view) are maximum we have to add the distributed loads that we computed from the seismic spectra to the other thermal and pressure loads. But we have a linear elastic problem (well, we still do not have it but we will in sec. 5.7.3), so we might be tempted to exploit the problem’s linearity and compute all the ef-







Figure 5.10: First six natural oscillation modes for a piping section between axial supports. Blue-red colour scale shows magnitude of deformation, as compared to the original geometry (grey). Full animations available online (sec. 5.9.1).



(a) a_x



(b) a_y



(c) a_z

Figure 5.11: Equivalent accelerations for a certain piping section

fects separately and then sum them up to obtain the whole combination. We may thus compute just the stresses due to the seismic loads and then add these stresses to the stresses at any time of the transient, in particular to the one with the highest ones. After all, in linear problems the result of the sum of two cases is the result of the sum of the cases, right? Not always.

Let us jump out of our nuclear piping problem and step back into the general finite-element theory ground for a moment (remember we were going to jump back and forth). Assume you want to know how much your dog weighs. One thing you can do is to weigh yourself (let us say you weigh 81.2 kg), then to grab your dog and to weigh both yourself and your dog (let us say you and your dog weigh 87.3 kg). Would you swear your dog weighs 6.1 kg plus/minus the scale’s uncertainty? I can tell you that the weight of two individual protons and two individual neutrons is not the same as the weight of an α particle. Would not there be a master-pet interaction that renders the weighting problem non-linear?

Time for both of us to make an experiment. Grab your favourite FEM program for the first time (remember mine is CAEplex, which can also be accessed through Onshape) and create a $1\text{mm} \times 1\text{mm} \times 1\text{mm}$ cube. Set any values for the Young’s Modulus and Poisson ratio as you want. I chose $E = 200 \text{ GPa}$ and $\nu = 0.28$. Restrict the three faces pointing to the negative axes to their planes, i.e.

- in face “left” ($x < 0$), set null displacement in the x direction ($u = 0$),
- in face “front” ($y < 0$), set null displacement in the y direction ($v = 0$),
- in face “bottom” ($z < 0$), set null displacement in the z direction ($w = 0$).

Now we are going to create and compare three load cases:²

- a. Pure normal loads (<https://caeplex.com/p/d8fe>)
- b. Pure shear loads (<https://caeplex.com/p/b494>)
- c. The combination of A & B (<https://caeplex.com/p/9899>)

The loads in each cases are applied to the three remaining faces, namely “right” ($x > 0$), “back” ($y > 0$) and “top,” ($z > 0$). Their magnitude in Newtons are:

	“right”			“back”			“top”		
	F_x	F_y	F_z	F_x	F_y	F_z	F_x	F_y	F_z
Case A	+10	0	0	0	+20	0	0	0	+30
Case B	0	+15	-15	+25	0	-5	-15	+25	0
Case C	+10	+15	-15	+25	+20	-5	-15	+25	+30

²The provided links lead to FEA cases which are fully usable directly from the web browser, i.e. “finite elements in the cloud.”

In the first case, the principal stresses are uniform and equal to the three normal loads. As the forces are in Newton and the area of each face of the cube is 1 mm^2 , the usual sorting leads to

$$\sigma_{1A} = 30 \text{ MPa}$$

$$\sigma_{2A} = 20 \text{ MPa}$$


$$\sigma_{3A} = 10 \text{ MPa}$$

In the second case, the principal stresses are not uniform and have a non-trivial distribution. Indeed, the distribution of σ_3 obtained by CAEplex is shown in fig. 5.12a. Now, if we indeed were facing a fully linear problem then the results of the sum of two inputs would be equal to the sum of the individual inputs. And fig. 5.12b, which shows the principal stress 3 of case C is not the result from case B plus any of the three constants from case A. Had it been, the colour distribution would be *exactly* the same as the scale goes automatically from the most negative value in blue to the most positive value in red. And $7+30 \neq 33$. Alas, it seems that there exists some kind of unexpected non-linearity (the feared master-pet interaction?) that prevents us from fully splitting the problem into simpler chunks.

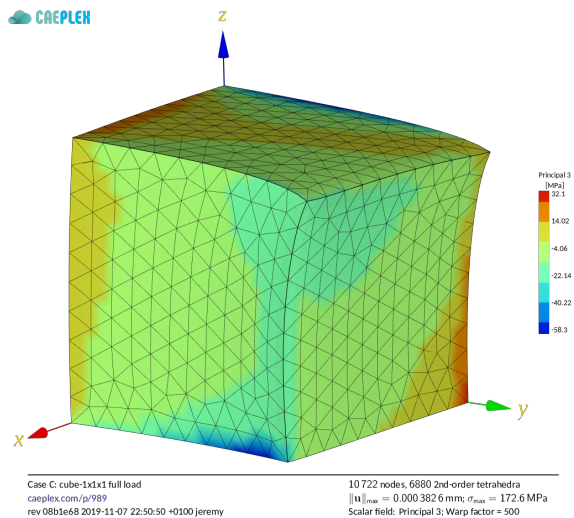
So what is the source of this unexpected non-linear effect in an otherwise nice and friendly linear formulation? Well, probably you already know it because after all it is almost high-school mathematics. But I learned long after college, when I had to face a real engineering problem and not just back-of-the-envelope pencil-and-paper trivial exercises.

Recall that principal stresses are the eigenvalues of the stress tensor. And the fact that in a linear elastic formulation the stress tensor of case C above is the sum of the individual stress tensors from cases A and B does not mean that their eigenvalues can be summed (think about it!). Again, imagine the eigenvalues and eigenvectors of cases A & B. Got it? Good. Now imagine the eigenvalues and eigenvectors for case C. Should they sum up? No, they should not! Let us make another experiment, this time with matrices using Octave or whatever other matrix-friendly program you want (try to avoid black boxes as explained in sec. 5.7.1).

First, let us create a 3×3 random matrix R and then multiply it by its transpose R^T to obtain a symmetric matrix A (recall that the stress tensor from sec. 5.3.1 is symmetric):



(a) Case B, pure-shear loads



(b) Case C, normal plus shear loads

Figure 5.12: Spatial distribution of principal stress 3 for cases B and C. If linearity applied, case C would be equal to case B plus a constant

```
octave> R = rand(3); A = R*R'
A =
    2.08711    1.40929    1.31108
    1.40929    1.32462    0.57570
    1.31108    0.57570    1.09657
```

Do the same to obtain another 3×3 symmetric matrix B:

```
octave> R = rand(3); B = R*R'
B =
    1.02619    0.73457    0.56903
    0.73457    0.53386    0.37772
    0.56903    0.37772    0.53141
```

Now compute the sum of the eigenvalues first and then the eigenvalues of the sum:

```
octave> eig(A)+eig(B)
ans =
    0.0075113
    0.8248395
    5.7674016

octave> eig(A+B)
ans =
    0.049508
    0.782990
    5.767255
```

Did I convince you? More or less, right? The third eigenvalue seems to fit. Let us not throw all of our beloved linearity away and dig in further into the subject. There are still two important issues to discuss which can be easily addressed using fresh-year linear algebra (remember not to fear maths!). First of all, even though principal stresses are not linear with respect to the sum they are linear with respect to pure multiplication. Once more, think what happens to the the eigenvalues and eigenvectors of a single stress tensor as all its elements are scaled up or down by a real scalar. They are the same! So, for example, the **Von Mises stress** (which is a combination of the principal stresses) of a beam loaded with a force $\alpha \cdot \mathbf{F}$ is α times the stress of the beam loaded with a force \mathbf{F} . Please test this hypothesis by playing with your favourite FEM solver. Or even better, take a look at the stress invariants I_1 , I_2 and I_3 (you can search online or peek into the source code of **Fino**, grep for

the routine called `fino_compute_principal_stress()` and see (using paper and pencil!) how they scale up if the individual elements of the stress tensor are scaled by a real factor α .

The other issue is that even though in general the eigenvalues of the sum of two matrices are not the same as the eigenvalues of the matrix sum, there are some cases when they are. In effect, if two matrices A and B commute, i.e. their product is commutative

$$A \cdot B = B \cdot A$$

then the sums (in plural because there are three eigenvalues) of their eigenvalues are equal to the eigenvalues of the sums. In order for this to happen, both A and B need to be diagonalisable using the same basis. That is to say, the diagonalising matrix P such that $P^{-1}AP$ is diagonal should be the same that renders $P^{-1}BP$ also diagonal. What does this mechanically mean? Well, if case A has loads that are somehow “independent” from the ones in case B, then the principal stresses of the combination might be equal to the sum of the individual principal stresses. A notable case is for example a beam that is loaded vertically in case A and horizontally in case B. I dare you to grab your FEM program one more time, run a test and picture the eigenvalues and eigenvectors of the three cases in your head.

The moral of this fable is that if we have a case that is the combination of two other simpler cases (say one with only surface loads and one with only volumetric loads), in general we cannot just add the principal stresses (or Von Mises) of two cases and expect to obtain a correct answer. We have to solve the full case again (both the surface and the volumetric loads at the same time). In case we are stubborn enough and still want to stick to solving the cases separately, there is a trick we can resort to. Instead of summing principal stresses, what we can do is to sum either displacements or the individual stress components, which are fully linear. So we might pre-deform (or pre-stress) case B with the results from case A and then expect the FEM program to obtain the correct stresses for the combined case. However, this scheme is actually far more complex than just solving the combined case in a single run and it also needs an educated guess and/or trial and error to know at what time the pre-deformation or pre-stressing should be applied to take into account the seismic load.

5.5.4 ASME stress linearisation (not linearity!)

After discussing linearity, let us now dig into linearisation. The name is similar but these two animals are very different beasts. We said in sec. 5.2 that the ASME

Boiler and Pressure Vessel Code was born long before modern finite-elements methods were developed and of course being massively available for general engineering analysis (democratised?). Yet the code provides a comprehensive, sound and, more importantly, a widely and commonly-accepted body of knowledge as for the regulatory authorities to require its enforcement to nuclear plant owners. One of the main issues of the ASME code refers to what is known as “membrane” and “bending” stresses. These are defined in [section VIII annex 5-A](#), although they are widely used in other sections, particularly [section III](#). Briefly, they give the zeroth-order (membrane) and first-order (bending) [moments](#) (in the statistical sense) of the stress distribution along a so-called Stress Classification Line or SCL, which should be chosen depending on the type of problem under analysis.

The computation of these membrane and bending stresses are called “[stress linearisation](#)” because it is like computing the [Taylor expansion](#) (or for the case, expansion in [Legendre polynomials](#)) of an arbitrary stress distribution along a line, and retaining the first two terms. That is to say, to obtain a *linear approximation*. There are physical interpretations for both membrane and bending stresses, which are beyond the scope of this chapter. As for the ASME requirements, they are a way of having the mean and linear contributions of a certain stress distribution along the pipe’s wall thickness.

So what about the SCLs? Well, the ASME standard says that they are lines that go through a wall of the pipe (or vessel or pump, which is what the ASME code is for) from the inside to the outside and ought to be normal to the iso-stress curves. Stop. Picture yourself a stress field, draw the iso-stress curves (those would be the lines that have the same colour in your picture) and then imagine a set of lines that travel in a perpendicular direction to them. Finally, choose the one that seems the prettiest (which most of the time is the one that seems the easiest). There you go! You have an SCL. But there is a catch. So far, we have referred to a generic concept of “stress.” Which of the several stresses out there should you picture? One of the three normals, the three shear, Von Mises, Tresca? Well, actually you will have to imagine tensors instead of scalars. And there might not be such a thing as “iso-stress” curves, let alone normal directions. So pick any radial straight line through the pipe wall at a location that seems relevant and now you are done. In our case study, there will be a few different locations around the material interfaces where high stresses due to differential thermal expansion are expected to occur. Just keep this though with you: it is very important to define where the SCLs are located, as they will define the “quality” of the obtained results.

5.6 The infinite pipe revisited after college

Let us now make a (tiny) step from the general and almost philosophical subject from the last section down to the particular case study, and reconsider the infinite pressurised pipe once again. It is time to solve the problem with a computer using finite elements and to obtain some funny coloured pictures instead of just equations (like we did in sec. 5.3.2).

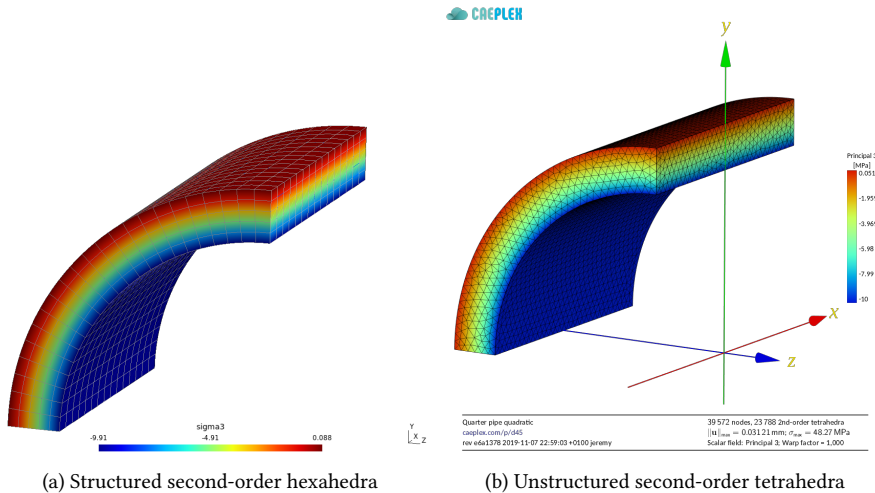
The first thing that has to be said is that, as with any interesting problem, there are literally hundreds or different ways of solving it, each of them throwing particular conclusions. For example, one can:

1. solve a real 3D problem or a 2D axi-symmetric case (or even a 1D case using the [collocation method](#) to solve the radial dependence),
2. have a full cylindrical geometry or just a half (180°), or a quarter (90°), or a thin slice (a small amount of degrees),
3. use a structured or an unstructured grid,
4. uniformly or locally-refine the mesh (with several choices of refinement),
5. use first or second-order (or higher) elements,
6. use tetrahedra or hexahedra,
7. in the case of second-order hexahedra, use complete (i.e. 27-node hexahedron) or incomplete (i.e. 20-node hexahedron) elements,
8. have different mesh sizes from very coarse to very fine,
9. solve the same problem in a few different solvers,
10. etc.

You can get both the exponential nature of each added bullet and how easily we can add new further choices to solve a FEM problem. And each of these choices will reveal you something about the nature of either the mechanical problem or the numerical solution. It is not possible to teach any possible lesson from every outcome in college, so you will have to learn them by yourself getting your hands at them. I have already tried to address the particular case of the infinite pipe in a [recent report](#)³ that is worth reading before carrying on with this article. The main conclusions of the report are:

- For the same number of elements, second-order grids need more nodes than linear ones, although they can better represent curved geometries.
- The three stress distributions computed with the finite-element give far more reasonable results for the second-order case than for the first-order grid. This

³The aforementioned “On convergence of linearized stresses in an infinite pipe computed using the finite element method.”



(a) Structured second-order hexahedra

(b) Unstructured second-order tetrahedra

Figure 5.13: Two of the hundreds of different ways the infinite pressurised pipe can be solved using FEM. The axial displacement at the ends is set to zero, leading to a [plane-strain](#) condition

is qualitatively explained by the fact that first-order tetrahedra have uniform derivatives and such the elements located in both the external and external faces represent the stresses not at the actual faces but at the barycentre of the elements.

- Membrane stresses converge well for both the first and second-order cases because they represent a zeroth-order moment of the stress distribution and the excess and defect errors committed by the internal and external elements approximately cancel out.
- Membrane plus bending stresses converge very poorly with linear elements because the excess and defect errors do not cancel out because it is a first-order moment of the stress distribution.
- The computational effort to solve a given problem, namely the CPU time and the needed RAM (fig. 5.14) depend non-linearly on various factors, but the most important one is the problem size which is three times the number of nodes in the grid * The error with respect to the analytical solutions as a function of the CPU time needed to compute the membrane stress is similar for both first and second-order grids. But for the computation of the membrane plus bending stress (fig. 5.14b), first-order grids give very poor results

compared to second-order grids for the same CPU time.

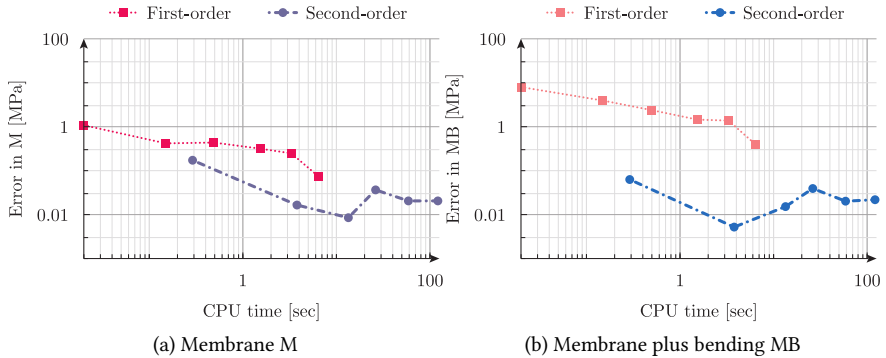


Figure 5.14: Error in the computation of the linearised stresses vs. CPU time needed to solve the infinite pipe problem using the finite element method

An additional note should be added. The FEM solution, which not only gives the nodal displacements but also a method to interpolate these values inside the elements, does not fully satisfy the original equilibrium equations at every point (i.e. the strong formulation). It is an approximation to the solution of the **weak formulation** that is close (measured in the vector space spanned by the **shape functions**) to the real solution. Mechanically, this means that the FEM solution leads only to nodal equilibrium but not point-wise equilibrium.

5.7 Adding complexity: the truth is out there

Let us review some issues that appear when solving our case study and that might not have been thoroughly addressed back during our college days.

5.7.1 Two (or more) materials

The main issue with fatigue in nuclear piping during operational transients is that at the welds between two materials with different thermal expansion coefficients there can appear potentially-high stresses during temperature changes. If these transients are repeated cyclically, fatigue may occur. We already have risen a warning flag about stresses at material interfaces in sec. 5.7.1. Besides all the open

questions about computing stresses at nodes, this case also adds the fact that the material properties (say the Young's Modulus E) is different in the elements that are at each side of the interface.

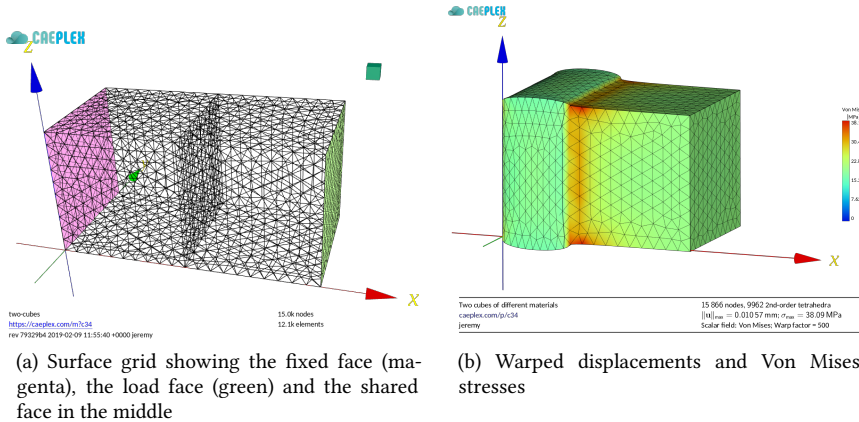


Figure 5.15: Two cubes of different materials (the one in the left soft, the one in the right hard) share a face and a pressure is applied at the right-most face

To simplify the discussion that follows, let us replace for one moment the full 3×3 tensor and the nine partial derivatives of the displacement by just one strain ϵ and let the [linear elastic strain-stress relationship](#) to be the simple scalar expression

$$\sigma = E \cdot \epsilon$$

Faced with the problem of computing the stress σ at one node shared by many elements, we might:

1. compute the (weighted?) averages $\langle E \rangle$ and $\langle \epsilon \rangle$ and then compute the stress as $\langle \sigma \rangle = \langle E \rangle \cdot \langle \epsilon \rangle$. This would be like having a meta-material at the interface with average properties, or
2. compute the stress as the (weighted?) average of the product $E \cdot \epsilon$ in each node $\langle \sigma \rangle = \langle E \cdot \epsilon \rangle$. This would be like forcing a non-differentiable function to behave smoothly, or
3. internally duplicate the nodes at the interface and compute one stress for each material. This would result in a stress distribution which is not a real function because the same location \mathbf{x} will be associated to more than one stress value, or

4. duplicate the surface elements at the interfaces and solve the problem using a contact formulation. This would render the problem non-linear and add the complexity of having to find appropriate penalty coefficients.

There might be other choices as well. Do you know what your favourite FEM program does? Now follow up with these questions:

- a. Does the manual say?
- b. Does it tell you the details like how it weighs the averages?
- c. Does it discard values that are beyond a number of standard deviations away?
- d. How many standard deviations?
- e. ...

You can still add a lot of questions that you should be having right now. If you cannot get a clear answer for at least one of them, then start to worry. After you do, add the following question:

Do you believe your favourite FEM program's manual?

What we as responsible engineers who have to sign a report stating that a nuclear power plant will not collapse due to fatigue in its pipes, is to fully understand what is going on with our stresses. [Richard Stallman](#) says that the best way to solve a problem is to avoid it in the first place. In this case, we should avoid having to trust a written manual and rely on software whose [source code](#) is available. What we need is the capacity (RMS calls it *freedom*) to be able to see the detailed steps performed by the program so we can answer any question we (or other people) might have.

Without resorting into philosophical digressions about the difference between [free and open-source software](#) (not because it is not worth it, but because it would take a whole book), the programs that make their source code available for their users are called [open-source software](#). If the users can also modify and re-distribute the modified versions, they are called [free software](#). Note that the important concept here is freedom, not price. In Spanish (my native language) it would have been easier because there are two separate words for free as in freedom (“libre”) and for free as in price (“gratis”).

In effect, a couple of years ago Angus Ramsay noted [a weird behaviour](#) in the results given by a certain commercial [non-free](#) FEA software regarding the handling of expansion coefficients from ASME data. To understand what was going on, Angus and I had to guess what the program was doing to [reproduce the allegedly weird results](#). Finally, it was a [matter how the data was rounded up to be](#)

presented in a paper table rather than a programming flaw. Nevertheless, we were lucky our guesses lead us to a reasonable answer. If we had access to the program's source code, we could have thoroughly analysed the issue in a more efficient way. Sure, we might not have the same programming skills the original authors of the software have, but if it had been *free software* we would have had the *freedom* to hire a programmer to help us out. That is what *free* means. In *Eric Raymond's* words, "given enough eyeballs, all bugs are shallow." This is rather important in engineering software where *verification and validation* is a must, especially in regulated fields like the nuclear industry. First, think how can a piece of software be verified if the source code is not available for independent analysis. And then, ask yourself another question:

Do you trust your favourite FEM program?

Back to the two-material problem, all the discussion above in sec. 5.7.1 about non-continuous derivatives applies to a sharp abrupt interface. In the study case the junctions are welded so there is a *heat-affected zone* with changes in the material micro structure. Therefore, there exists a smooth transition from the mechanical properties of one material to the other one in a way that is very hard to predict and to model. In principle, the assumption of a sharp interface is conservative in the sense that it is expected the computed stresses to be larger than the actual ones. There cannot be an SCL exactly on a material interface so there should be at least two SCLs, one at each side of the junctions as fig. 5.6 illustrates. The actual distance would have to be determined first as an educated guess, then via trial and error and finally in accordance with the regulator.

5.7.2 A parametric tee

Time for another experiment. We know (more or less) what to expect from an infinite pressurised pipe from sec. 5.3.2 and sec. 5.6. What if we added a branch to such pipe? Even more, what if we successively varied the diameter of the branch to see what happens? This is called parametric analysis, and sooner or later (if not now) you will find yourself performing this kind of computations more often than any other one.

Let us solve the following mock-up case: a long main 12-inch schedule 80 pipe has an orthogonal branch of a certain nominal diameter of d_b inches (it seems that the *SI* did not do well amongst the piping engineering community). Both the main line and the branch are pressurised with $p = 10$ MPa. The main pipe is aligned with the x axis and the branch coincides with the z axis. Thus, the x - z plane acts as a symmetry plane and we only need to model two octants of the full geometry, as

shown in fig. 5.16. Note that the tee is modelled as the boolean intersection of two cylinders. There are no filleted edges nor rounded corners or any other smoothing. A real CAD file containing the appropriate geometry needs to be built for the real case study.




Figure 5.16: Geometry of the parametric 12-inch tee for the particular case of a 4-inch branch

The boundary conditions are

- Magenta: axial symmetry $u = 0$
- Salmon: plane symmetry $v = 0$
- Yellow: axial $u = 0$ and radial symmetry $0 = vz - wy$
- Dark green: axial $w = 0$ and radial symmetry $0 = vx - uy$
- Light green: internal pressure $p = 10$ MPa

Three radial SCLs—namely cyan, yellow and green—are located at a distance of one thickness of the main line from the external wall of the branch into the main direction x (fig. 5.17). In fig. 5.18 we can see the results of the parametric computation, namely the linearised membrane and bending stresses computed in each of the three SCLs as a function of d_b . Note that d_b is a continuous variable and even unreal values (such as 9-inches) were used to perform the parametric run.

The thickness of such cases was interpolated using actual schedule-80 geometrical data from ASME B36.310M.




Figure 5.17: Location of the three radial SCLs: cyan, yellow and green

Fig. 5.19 illustrates how the pipes deform when subject to the internal pressure. When the branch is small, the problem resembles the infinite-pipe problem where the main pipe expands radially outward and there is only traction. For large values of d_b , the pressure in the branch bends down the main pipe, generating a complex mixture of traction and compression. The tipping point seems to be around a branch diameter $d_b \approx 5$ in.

Do you now see the added value of training throw-ins with watermelons? We might go on...

- studying how the maximum Von Mises stress as a function the radius r_f of an hypothetical fillet operation on the tee's sharp edges,
- making the branch another material and adding thermal expansion,
- adding piping supports to restrain the degrees of freedom of the pipes,
- using distributed forces from earthquakes,
- etc.

Most of the time at college we would barely do what is needed to be approved in one course and move on to the next one. If you have the time and consider a career

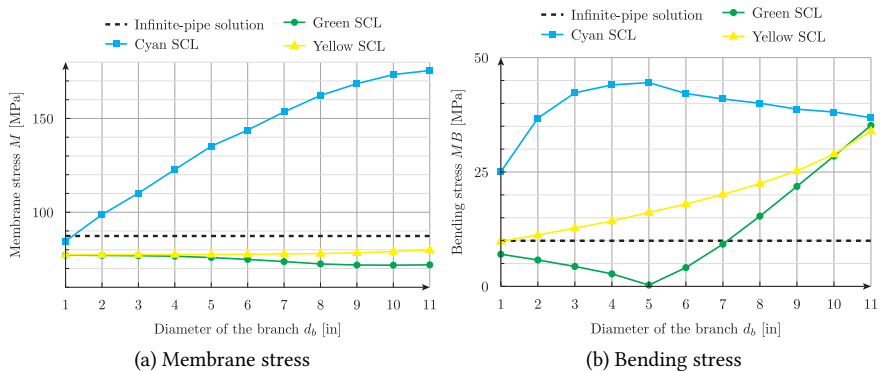


Figure 5.18: Parametric stresses as a function of the nominal diameter d_b of the branch




Figure 5.19: Von Mises stress and 400x warped displacements for three values of d_b

related to finite-element analysis, please do not. Clone the repository (sec. 5.9.1) with the input files for [Fino](#) and start playing. If you are stuck, do not hesitate asking for help in [wasora's mailing list](#).

One further detail: it is always a sane check to try to explain the numerical results based on physical reasoning (i.e. “with your fingers”) as we did two paragraphs above. Most of the time you will be solving problems whilst already knowing what the result would (or ought to) be.

5.7.3 Bake, shake and break

A fellow mechanical engineer who went to the same high school I did, who went to the same engineering school I did and who worked at the same company I did, but who earned a PhD in Norway once told me two interesting things about finite-elements graduate courses. First, that in Trondheim the classes were taught by faculty from the the mathematics department rather than from the mechanical engineering department. It made complete sense to me, as I always have thought finite elements mainly as a maths subject. And even though engineers might know some maths, it is nothing compared to actual mathematicians. Secondly, that they called the thermal, natural oscillations and elastic problems as the rhyming trio “bake, shake and break” (they also had “wake” for fluids, but that is a different story). These are just the three problems listed in section sec. 5.5 that we need to solve in our nuclear power plant.

So here we are again with the case study where we have to compute the linearised stresses at certain SCLs located near the interface between two different kinds of steels during operational and incidental transients of the plant. The first part is then to “bake” the pipes, modelling a thermal transient with time-dependent temperature or convection boundary conditions (it depends on the system). This steps gives a time and space-dependent temperature $T(x, y, z, t)$.

We then proceed to “shake” the pipes. That is to say, we obtain a distributed load vector $\mathbf{f}(x, y, z)$ which is statically equivalent to the design earthquake.

Finally we attempt to “break” the pipes successively solving many steady-state elastic problems for different times t of the operational transient. Some remarks about this step:

1. The material properties are temperature-dependent (we use data from [ASME II](#) part D).
2. Thermal expansion is taken into account. The reference temperature (i.e. the temperature at which there is no expansion) is 20°C that coincides with ASME's decision of the reference temperature for the mean thermal expansion coefficients.

3. The temperature distribution $T(x, y, t, z)$ for bullets 1 & 2 is the generalisation of the reduced-model procedure explained in sec. 5.5.1.
4. The internal faces of the pipes are subject to an uniform pressure $p(t)$ given by the definition of the transient.
5. There are mechanical supports throughout the piping system. Depending on the type of the support (i.e. vertical, lateral, axial, full, etc.) one or more degrees of freedom (i.e. u, v and/or w) are fixed to zero. The ends of the CAD models are chosen always to have axially-null displacements.
6. The earthquake-equivalent volumetric force $\mathbf{f}(x, y, z)$ is only be applied at the time t where the maximum stresses occur. Note that due to the discussion from sec. 5.5.3 we cannot compute the stresses raised just by $\mathbf{f}(x, y, z)$ and then add them to the main stresses. The force has to be included into the “break” step. An educated guess of the time where the maximum stress occur is usually enough. Anyway, it might be necessary a trial and error scheme to find the sweet spot.
7. According to ASME III, the seismic load has to be applied during two seconds with the two possible signs. That is to say, apply $\mathbf{f}(x, y, z)$ during two seconds and then $-\mathbf{f}(x, y, z)$ during two further seconds when the main stresses are maximums.
8. A number of stress classification lines have to be defined. The locations should be previously accorded with the plant owner and/or the regulator.
9. The stress linearisation has to be performed individually for each principal stress σ_1, σ_2 and σ_3 to fulfil the requirements ASME III NB-3126 (see sec. 5.8.1 below).
10. This “break” step is linear.

Is the last bullet right? [Surely you’re joking, Mr. Theler!](#) Linear problems are simple and almost useless. How can this extremely complex problem be linear? Well, let us see. First, there are two main kinds of non-linearities in FEM:

1. Geometrical non-linearities
2. Material non-linearities

The first one is easy. Due to the fact that the pipes are made of steel, it is expected that the actual deformations are relatively small compared to the original dimensions. This leads to the fact that the mechanical rigidity (i.e. the stiffness matrix) does not change significantly when the loads are applied. Therefore, we can safely assume that the problem is geometrically linear.

Let us now address material non-linearities. On the one hand we have the temperature-dependent issue. According to ASME II part D, what depends on temperature T is the Young’s Modulus E . But the stress-strain relationship is yet

$$\sigma = E(T) \cdot \epsilon$$

What changes with temperature is the slope of σ with respect to ϵ (think and imagine!), but the relationship between them is still *linear*.

On the other hand, we have a given non-trivial temperature distribution $T(\mathbf{x}, t)$ within the pipes that is a snapshot of a transient heat conduction problem at a certain time t (think and picture yourself taking photos of the temperature distribution changing in time). Let us now forget about the time, as after all we are solving a steady-state elastic problem. Now you can trust me or ask a FEM teacher, but the continuous displacement formulation can be loosely written as

$$K[E(T(\mathbf{x})), \mathbf{x}] \cdot \mathbf{u}(\mathbf{x}) = \mathbf{b}(\mathbf{x})$$

If you notice, the complex dependence of the stiffness matrix K can be reduced to just the spatial vector \mathbf{x} :


$$K(\mathbf{x}) \cdot \mathbf{u}(\mathbf{x}) = \mathbf{b}(\mathbf{x})$$

And this last expression is linear in \mathbf{u} ! In effect, the discretisation step means to integrate over \mathbf{x} . As K , \mathbf{u} and \mathbf{b} depend only on \mathbf{x} , then after integration one gets just numbers inside K and \mathbf{b} . Again, you can either, in increasing order of recommendation:


1. trust me,
2. ask a teacher, or
3. go through with the maths.

To recapitulate, the figures in this section show some partial non-dimensional results of an actual system of a certain nuclear power plant. The main issues to study were the interfaces between a carbon-steel pipe and a stainless-steel orifice plate used to measure the (heavy) water flow through the line. The steps discussed so far include

1. building a CAD model of the piping section under study, which will be the main domain (fig. 5.20a or figs. 5.3, 5.5a, 5.7)
2. creating a mesh for the main domain refining locally around the material interfaces (fig. 5.20b or figs. 5.5b, 5.9)
3. defining the number and locations of the SCLs (fig. 5.21 or figs. 5.6, 5.7)



(a) General view. Carbon steel is grey and stainless steel is green.



(b) Detail of the mesh refinement around the interface.

Figure 5.20: A section of a piping system in a nuclear power plant






Figure 5.21: Location of the 15 SCLs



(a) Simplified axi-symmetric domain



(b) Generalisation

Figure 5.22: Temperature distribution for a certain instant of the transient, computed in the simplified two-dimensional axi-symmetric domain and its generalisation to the three-dimensional mechanical domain as discussed in sec. 5.5.1

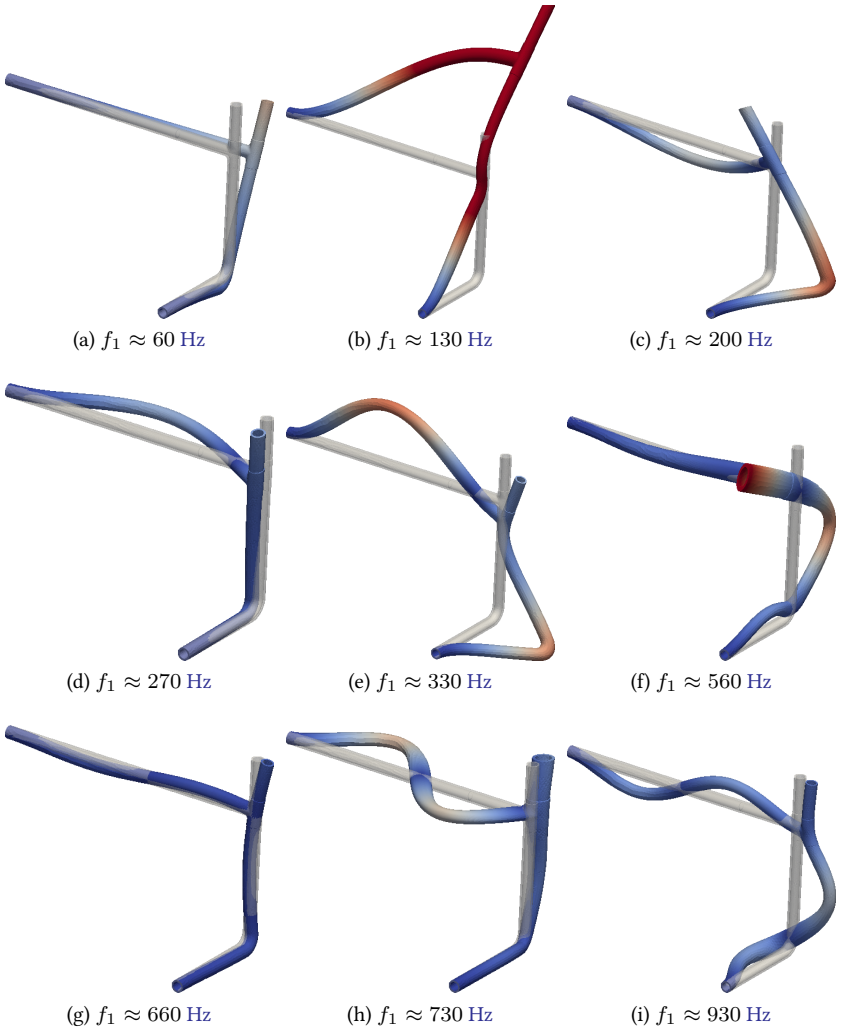


Figure 5.23: First nine natural modes of oscillation of the piping system subject to the boundary conditions the supports provide. Animations are available online (sec. 5.9.1).

4. computing a heat conduction (bake) transient problem with temperatures as a function of time from the operational transient in a simple domain using temperature-dependent thermal conduction coefficients from ASME II part D (fig. 5.22 or fig. 5.8)
5. generalising the temperature distribution as a function of time to the general domain (fig. 5.22b or fig. 5.9)
6. performing a modal analysis (shake) on the main domain to obtain the main oscillation frequencies and modes (fig. 5.23
7. obtaining a distributed force statically-equivalent to the earthquake load (shake, fig. 5.24 or fig. 5.11)
7. successively solving the linear elastic problem for different times using the generalised temperature distribution taking into account
 - a. the dependence of the Young's Modulus E and the thermal expansion coefficient α with temperature,
 - b. the thermal expansion effect itself
 - c. the instantaneous pressure exerted in the internal faces of the pipes at the time t according to the definition of the operational transient
 - d. the restriction of the degrees of freedom of those faces, lines or points that correspond to mechanical supports located both within and at the ends of the CAD model
 - e. the earthquake load, which according to ASME should be present only during four seconds of the transient: two seconds with one sign and the other two seconds with the opposite sign. This period should be selected to coincide with the instant of highest mechanical stress (conservative computation)
8. computing the linearised stresses (membrane and membrane plus bending) at the SCLs combining them as Principal 1, Principal 2, Principal 3 and Tresca⁴
9. juxtaposing these linearised stresses for each time of the transient and for each transient so as to obtain a single time-history of stresses including all the operational and/or incidental transients under study, which is what stress-based fatigue assessment needs (recall sec. 5.2.3 and go on to sec. 5.8).

A pretty nice list of steps, which definitely I would not have been able to tackle when I was in college. Would you?

⁴A deep technical note should be added for discussing the validity of linearising a stress invariant as the principal stresses individually. It should be enough for the present study to keep in mind that in sec. 5.8.1 it is the difference of these linearised principal stresses (i.e. Tresca-like) are used to compute the stress amplitude of the periodic load for fatigue assessment.




Figure 5.24: The static equivalent accelerations computed using the SRSS method

5.8 Cumulative usage factors

Strictly speaking, finite elements are not needed anymore at this point of the analysis. But even though we are (or want to be) FEM experts, we have to understand that if the objective of a work is to evaluate fatigue (or fracture mechanics or whatever), finite elements are just a mean and not an end. If we just mastered FEM and nothing else, our field of work would be highly reduced. We need to use all of our computational knowledge to perform actually engineering tasks and to be able to tell our bosses and clients whether the pipe would fail or not. This tip is induced in college but it is definitely reinforced afterwards when working with actual clients and bosses.

Another comment I would like to add is that I had to learn fatigue practically from scratch when faced with this problem for the first time in my engineering career. I remembered some basics from college (like the general introduction from sec. 5.2.3), but I lacked the skills to perform a real computation by myself. Luckily there still exist books, there are a lot of interesting online resources (not to mention Wikipedia) and, even more luckily, there are plenty of other fellow engineers that are more than eager to help you. My second tip is: when faced to a new challenging problem, read, learn and ask for guidance to real people to see if you read and learned it right.

Back and distantly, in sec. 5.2 we said that people noticed there were some environmental factors that affected the fatigue resistance of materials. The basic ASME approach does not take care of these factors, and it is regarded as fatigue “in air.” We are interested in taking them into account, so we follow the US Nuclear Regulatory Commission guidelines to evaluate fatigue “in water.”

5.8.1 In air (ASME’s basic approach)

We already said in sec. 5.2.3 that the stress-life fatigue assessment method gives the limit number N of cycles that a certain mechanical part can withstand when subject to a certain periodic load of stress amplitude S_{alt} . If the actual number of cycles n the load is applied is smaller than the limit N , then the part is fatigue-resistant. In our case study there is a mixture of several periodic loads, each one expected to occur a certain number of times. ASME’s way to evaluate the resistance is to break up the stress history into partial stress amplitudes $S_{\text{alt},j}$ between a “peak” and a “valley” and to compute individual usage factors U_j for the j -th amplitude (which does not need to coincide with one of the k transient loads) as

$$U_j = \frac{n_j}{N_j}$$

The overall cumulative usage factor is then the algebraic sum of the partial contributions, a.k.a. [Miner’s rule](#) as learned in college:

$$\text{CUF} = U_1 + U_2 + \dots + U_j + \dots$$

When $\text{CUF} < 1$, then the part under analysis can withstand the proposed cyclic operation. Now, if the extrema of the partial stress amplitude correspond to different transients, then the following note in ASME III’s NB-3224(5) should be followed:

In determining $n_1, n_2, n_3, \dots, n_j$ consideration shall be given to the superposition of cycles of various origins which produce a total stress difference range greater than the stress difference ranges of the individual cycles. For example, if one type of stress cycle produces 1,000 cycles of a stress difference variation from zero to +60,000 psi and another type of stress cycle produces 10,000 cycles of a stress difference variation from zero to –50,000 psi, the two types of cycle to be considered are defined by the following parameters:

- (a) for type 1 cycle, $n_1 = 1,000$ and $S_{\text{alt},1} = (60,000 + 50,000)/2$;
- (b) for type 2 cycle, $n_2 = 9,000$ and $S_{\text{alt},2} = (50,000 + 0)/2$.

This cryptic paragraph can be better explained by using a clearer example. To avoid using actual sensitive data from a real power plant, let us use the same test case used by both the [US Nuclear Regulatory Commission](#) (in its report NUREG/CR-6909) and the [Electric Power Institute](#) (report 1025823) called “EAF (Environmentally-Assisted Fatigue) Sample Problem 2-Rev. 2 (10/21/2011)”.




Figure 5.25: A low-alloy steel vessel nozzle (blue) welded to a stainless steel pipe (grey)

It consists of a typical vessel nozzle with attached piping as illustrated in fig. 5.25. These components are subject to four transients $k = 1, 2, 3, 4$ that give rise to linearised stress histories (slightly modified according to NB-3216.2) which are given as individual stress values juxtaposed so as to span a time range of about 36,000 seconds (fig. 5.26a). As the time steps is not constant, each stress value has an integer index i that uniquely identifies it:

k	Time range [s]	Index range	Cycles n_k
1	0–3210	1–523	20
2	3210–6450	524–959	50
3	6450–9970	960–1595	20
4	9970–35971	1596–2215	100

A design-basis earthquake was assumed to occur briefly during one second (sic) at around $t = 3470$ seconds, and it is assigned a number of cycles $n_e = 5$. The detailed stress history for one of the SCLs including both the principal and linearised stresses, which are already offset following NB-3216.2 so as to have a maximum stress equal to zero, can be found as an appendix in NRC’s NUREG/CR-6909 report, or in the repository with the scripts I prepared for you to play with this problem.

To compute the global usage factor, we first need to find all the combinations of local extrema pairs and then sort them in decreasing order of stress difference. For example, the largest stress amplitude is found between $i = 447$ and $i = 694$. The second one is 447–699. Then 699–1020, 699–899, etc. For each of these pairs, defined by the indexes $i_{1,j}$ and $i_{2,j}$, a partial usage factor U_j should be computed. The stress amplitude $S_{alt,j}$ which should be used to enter into the $S-N$ curve is

$$S_{alt,j} = \frac{1}{2} \cdot k_{e,j} \cdot \left| MB'_{i_1,j} - MB'_{i_2,j} \right| \cdot \frac{E_{SN}}{E(T_{max_j})}$$

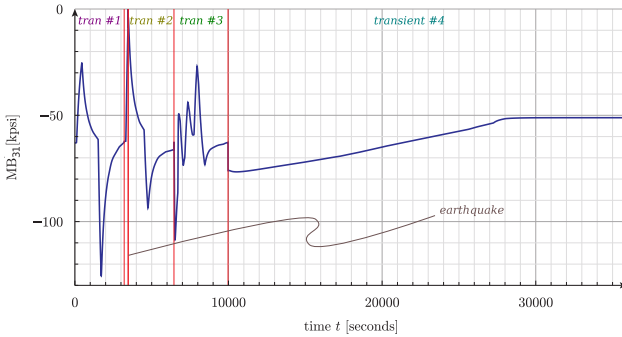
where k_e is a plastic correction factor for large loads (NB-3228.5), E_{SN} is the Young's Modulus used to create the $S-N$ curve (provided in the ASME fatigue curves) and $E(T_{max_j})$ is the material's Young's Modulus at the maximum temperature within the j -th interval.

We now need to comply with ASME's obscure note about the number of cycles to assign a proper value of n_j . Back to the largest pair 447-694, we see that 447 belongs to transient #1 which has assigned 20 cycles and 694 belongs to the earthquake with 5 cycles. Therefore $n_1 = 5$, and the cycles associated to each index are decreased in five. So $i = 694$ disappears and the number of cycles associated to $i = 447$ are decreased from 20 to 15. The second largest pair is now 447-699, with 15 (because we just spent 5 in the first pair) and 50 cycles respectively. Now $n_2 = 15$, point 447 disappears and 699 remains with 35 cycles. The next pair is 699-1020, with number of cycles 35 and 20 so $n_3 = 20$, point 1020 disappears and 699 remains with 15 cycles. And so on, down to the last pair.

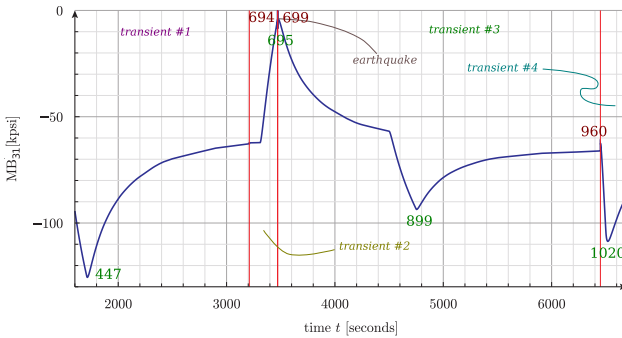
Why all these details? Not because I want to teach you how to perform fatigue evaluations just reading this section without resorting to ASME, fatigue books and even other colleagues. It is to show that even though these computation can be made "by hand" (i.e. using a calculator or, God forbids, a spreadsheet) when having to evaluate a few SCLs within several piping systems it is far (and I mean really far) better to automate all these steps by writing a set of scripts. Not only will the time needed to process the information be reduced, but also the introduction of human errors will be minimised and repeatability of results will be assured—especially if working under a [distributed version control](#) system such as [Git](#). This is true in general, so here is another tip: learn to write scripts to post-process your FEM results (you will need to use a script-friendly FEM program) and you will gain considerable margins regarding time and quality.

5.8.2 In water (NRC's extension)

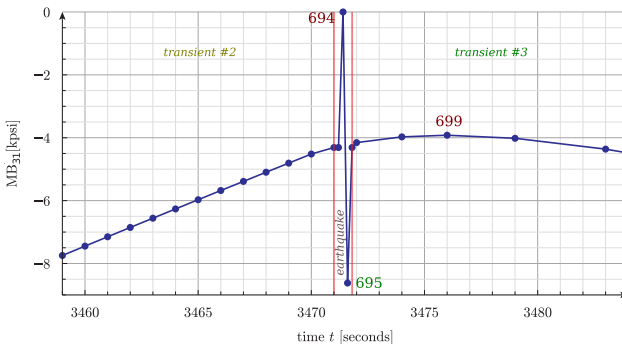
The fatigue curves and ASME's (both III and VIII) methodology to analyse cyclic operations assume the parts under study are in contact with air, which is not the case of nuclear reactor pipes. Instead of building a brand new body of knowledge to replace ASME, the NRC decided to modify the former adding environmentally-assisted fatigue multipliers to the traditional usage factors, formally defined as



(a) Full time range of the juxtaposed transients spanning $\approx 36,000$ seconds



(b) Detail of transients showing the ids of some extrema



(c) Detail of the earthquake (it does not follow the ASME two-second rule)

Figure 5.26: Time history of the linearised stress MB_{31} corresponding to the example problem from NRC and EPRI. The indexes i of the extrema are shown in green (minimums) and red (maximums)

5.8. CUMULATIVE USAGE FACTORS

Pair A	Pair B	Applied Cycles A	Applied Cycles B	M+B STRESS (psi)	K_e	Total Stress (psi)	S_{alt} (psi)	N_h	n_h	U_h	Max. Metal Temp. (°F)	DO (ppm)
694	447	5	20	125542.9	2.580	144164.4	220490.4	140.005	5	0.0357	566.6	0.150
699	447	50	15	121622.8	2.405	139047.0	198300.6	178.958	15	0.0838	566.6	0.550
699	1021	35	20	104691.5	1.653	126037.5	124507.0	582.468	20	0.0343	600.4	0.550
699	899	15	50	89695.4	1.000	102302.8	57864.5	6339.47	15	0.0024	336.1	0.550
695	899	5	35	84993.9	1.000	98798.6	55882.4	7027.83	5	0.0007	336.1	0.550
185	899	20	30	68222.2	1.000	76465.1	43250.2	15549.1	20	0.0013	336.1	0.550
1432	899	20	10	66665.7	1.000	83098.8	47002.3	11892.7	10	0.0008	336.1	0.550
1432	1653	10	100	49437.0	1.000	61950.9	33687.5	35734.8	10	0.0003	103.0	0.522
1296	1653	20	90	32478.6	1.000	38719.1	22025.4	154852	20	0.0001	366.2	0.522
1136	1653	20	70	27045.6	1.000	33751.1	19388.7	258499	20	0.0001	417.7	0.522
2215	1653	100	50	25255.9	1.000	25668.1	15147.6	1.15E+06	50	0.0000	547.0	0.522
2215	1213	50	20	22343.7	1.000	25298.3	14929.4	1.30E+06	20	0.0000	547.0	0.050
2215	1562	30	20	22047.7	1.000	24970.1	14735.7	1.46E+06	20	0.0000	547.0	0.050
2215	1	10	20	11956.0	1.000	12255.6	7232.5	1.00E+11	10	0.0000	547.0	0.150
1347	1	20	10	3786.5	1.000	4173.0	2412.1	1.00E+11	10	0.0000	450.0	0.150
1347	1595	10	20	3408.0	1.000	3430.2	1963.3	1.00E+11	10	0.0000	398.7	0.050
960	1595	20	10	241.8	1.000	259.9	146.0	1.00E+11	10	0.0000	299.5	0.050
960	960	5	5	0.0	1.000	0.0	0.0	1.00E+11	10	0.0000	299.5	0.050

TOTAL CUF = 0.1596

(a) Reference from NUREG/CR6909

j	A_j	B_j	$n(A_j)$	$n(B_j)$	MB_j [ksi]	$k_{e,j}$	S'_j [ksi]	$S_{alt,j}$ [ksi]	N_j	n_j	U_j	$T_{max,j}$ [°F]
1	447	694	20	5	125.5	2.580	144.2	220.400	1.40×10^2	5	3.57×10^{-2}	566.6
2	447	699	15	50	121.6	2.405	139	198.300	1.79×10^2	15	8.38×10^{-2}	566.6
3	699	1020	35	20	104.7	1.653	126.5	124.900	5.77×10^2	20	3.47×10^{-2}	599.2
4	699	899	15	50	89.7	1.000	102.3	62.640	5.02×10^3	15	2.99×10^{-3}	336.1
5	695	899	5	35	84.99	1.000	98.8	59.750	5.77×10^3	5	8.67×10^{-4}	336.1
6	899	1432	30	20	66.67	1.000	83.1	50.360	9.56×10^3	20	2.09×10^{-3}	634.2
7	184	899	20	10	68.23	1.000	76.76	46.440	1.24×10^4	10	8.09×10^{-4}	600.0
8	184	1641	10	100	51.22	1.000	55.83	33.630	3.59×10^4	10	2.78×10^{-4}	634.2
9	1296	1641	20	90	32.69	1.000	38.94	22.110	1.53×10^5	20	1.31×10^{-4}	366.2
10	1134	1641	20	70	27.31	1.000	34.49	19.800	2.34×10^5	20	8.53×10^{-5}	419.2
11	1641	2215	50	100	25.47	1.000	25.89	15.270	1.07×10^6	50	4.66×10^{-5}	547.0
12	1213	2215	20	50	22.34	1.000	25.3	14.930	1.31×10^6	20	1.53×10^{-5}	547.0
13	1630	2215	100	30	24.88	1.000	25.2	14.870	1.35×10^6	30	2.22×10^{-5}	547.0
14	1347	1630	20	70	16.71	1.000	17.12	9.798	3.72×10^9	20	5.38×10^{-9}	398.7
15	960	1630	20	50	13.54	1.000	13.95	8.405	7.76×10^{10}	20	2.58×10^{-10}	634.2
16	1595	1630	20	30	13.3	1.000	13.69	7.690	1.00×10^{11}	20	2.00×10^{-10}	299.4
17	1	1630	20	10	12.92	1.000	12.95	7.469	1.00×10^{11}	10	1.00×10^{-10}	450.0
18	1	1596	10	100	12.92	1.000	12.95	7.469	1.00×10^{11}	10	1.00×10^{-10}	450.0
19	1562	1596	20	90	2.829	1.000	0.2345	0.132	1.00×10^{11}	20	2.00×10^{-10}	299.4

CUF total = 0.1615

(b) Results reproduced by the author

Figure 5.27: Tables of individual usage factors for the NRC/EPRI “EAF Sample Problem 2-Rev. 2 (10/21/2011).” One table is taken from a document issued by almost-a-billion-dollar-year-budget government agency from the most powerful country in the world and the other one is from a third-world engineering startup. Guess which is which.

$$F_{\text{en}} = \frac{N_{\text{air}}}{N_{\text{water}}} \geq 1$$

This way, the environmentally-assisted usage factor for the j -th load pair is $\text{CUF}_{\text{en},j} = U_j \cdot F_{\text{en},j}$ and the global cumulative usage factor in water is the sum of these partial contributions

$$\text{CUF}_{\text{en}} = U_1 \cdot F_{\text{en},1} + U_2 \cdot F_{\text{en},2} + \cdots + U_j \cdot F_{\text{en},j} + \cdots \quad (5.1)$$

In EPRI's words, the general steps for performing an EAF analysis are as follows:

1. perform an ASME fatigue analysis using fatigue curves for an air environment
2. calculate F_{en} factors for each transient pair in the fatigue analysis
3. apply the F_{en} factors to the incremental usage calculated for each transient pair (U_j), to determine the CUF_{en} , using eq. 5.1

Again, if $\text{CUF}_{\text{en}} < 1$, then the system under study can withstand the assumed cyclic loads. Note that as $F_{\text{en},j}$, we can have $\text{CUF} < 1$ and $\text{CUF}_{\text{en}} > 1$ at the same time. The NRC has performed a comprehensive set of theoretical and experimental tests to study and analyse the nature and dependence of the non-dimensional correction factors F_{en} . They found that, for a given material, they depend on:

- a. the concentration $O(t)$ of dissolved oxygen in the water,
- b. the temperature $T(t)$ of the pipe,
- c. the strain rate $\dot{\epsilon}(t)$, and
- d. the content of sulphur $S(t)$ in the pipes (only for carbon or low-alloy steels).

Apparently it makes no difference whether the environment is composed of either light or heavy water. There are somewhat different sets of non-dimensional analytical expressions that estimate the value of $F_{\text{en}}(t)$ as a function of $O(t)$, $T(t)$, $\dot{\epsilon}(t)$ and $S(t)$, both in the few revisions of NUREG/CR-6909 and in EPRI's report #1025823. Although they are not important now, the actual expressions should be defined and agreed with the plant owner and the regulator. The main result to take into account is that $F_{\text{en}}(t) = 1$ if $\dot{\epsilon}(t) \leq 0$, i.e. there are no environmental effects during the time intervals where the material is being compressed.

Once we have the instantaneous factor $F_{\text{en}}(t)$, we need to obtain an average value $F_{\text{en},j}$ which should be applied to the j -th load pair. Again, there are a few

different ways of lumping the time-dependent $F_{\text{en}}(t)$ into a single $F_{\text{en},j}$ for each interval. Both NRC and EPRI give simple equations that depend on a particular time discretisation of the stress histories that, in my view, are all ill-defined. My guess is that they underestimated their audience and feared readers would not understand the slightly-more complex mathematics needed to correctly define the problem. The result is that they introduced a lot of ambiguities (and even technical errors) just not to offend the maths illiterate. A decision I do not share, and a another reason to keep on learning and practising math.

When faced for the first time with the case study, I have come up with a weighting method that I claim is less ill-defined (it still is for border-line cases) and which the plant owner accepted as valid. Fig. 5.28 shows the reference results of the problem (based on computing two correction factors and then taking the maximum) and the ones obtained with the proposed method (by computing a weighted integral between the valley and the peak). Note how in fig. 5.28a, pairs 694-447 and 699-447 have the same F_{en} even though they are (marginally) different. The results from fig. 5.28b give two (marginally) different correction factors.

5.9 Conclusions

Back in college, we all learned how to solve engineering problems. And once we graduated, we felt we could solve and fix the world (if you did not graduate yet, you will feel it shortly). But there is a real gap between the equations written in chalk on a blackboard (now probably in the form of beamer slide presentations) and actual real-life engineering problems. This chapter introduces a real case from the nuclear industry and starts by idealising the structure such that it has a known analytical solution that can be found in textbooks. Additional realism is added in stages allowing the engineer to develop an understanding of the more complex physics and a faith in the veracity of the finite-element results where theoretical solutions are not available. Even more, a brief insight into the world of evaluation of stress-life fatigue using such results further illustrates the complexities of real-life engineering analysis, even though the presented case was simplified for the sake of clearness. Here is a list of the tips that arose throughout the text:

- use and exercise your imagination
- practise math
- start with simple cases first
- grasp the dependence of results with independent variables
- remember there are other methods beside finite elements

CHAPTER 5. CS4: FATIGUE ASSESSMENT IN NUCLEAR PIPING

PAIR A	PAIR B	Fen A	Fen B	USAGE FACTOR	MAXIMUM Fen	CUFen
694	447	15.520	15.520	0.0357	15.520	0.5543
699	447	1.517	15.520	0.0838	15.520	1.3009
699	1021	1.517	5.065	0.0343	5.065	0.1739
699	899	1.517	6.429	0.0024	6.429	0.0152
695	899	1.517	6.429	0.0007	6.429	0.0046
185	899	4.298	6.429	0.0013	6.429	0.0083
1432	899	2.009	6.429	0.0008	6.429	0.0054
1432	1653	2.009	11.465	0.0003	11.465	0.0032
1296	1653	4.184	11.465	0.0001	11.465	0.0015
1136	1653	5.065	11.465	0.0001	11.465	0.0009
2215	1653	11.465	11.465	0.0000	11.465	0.0005
2215	1213	11.465	4.184	0.0000	11.465	0.0002
2215	1562	11.465	1.886	0.0000	11.465	0.0002
2215	1	11.465	4.298	0.0000	11.465	0.0000
1347	1	2.009	4.298	0.0000	4.298	0.0000
1347	1595	2.009	1.886	0.0000	2.009	0.0000
960	1595	6.429	1.886	0.0000	6.429	0.0000
960	960	6.429	6.429	0.0000	6.429	0.0000
TOTAL CUFen =						2.0689

(a) Results according to the author of NUREG/CR-6909 corresponding to the latest draft of the document

j	A_j	$t(A_j)$	B_j	$t(B_j)$	U_j	$F_{en,j}$	$CUF_{en,j}$
1	447	1716.0	694	3471.4	3.57×10^{-2}	15.550	5.55×10^{-1}
2	447	1716.0	699	3476.0	8.38×10^{-2}	15.160	1.27
3	699	3476.0	1020	6518.0	3.47×10^{-2}	6.582	2.28×10^{-1}
4	699	3476.0	899	4754.0	2.99×10^{-3}	1.000	2.99×10^{-3}
5	695	3471.6	899	4754.0	8.67×10^{-4}	1.000	8.67×10^{-4}
6	899	4754.0	1432	7955.0	2.09×10^{-3}	4.461	9.34×10^{-3}
7	184	451.0	899	4754.0	8.09×10^{-4}	15.170	1.23×10^{-2}
8	184	451.0	1641	10519.0	2.78×10^{-4}	8.755	2.44×10^{-3}
9	1296	7356.0	1641	10519.0	1.31×10^{-4}	1.971	2.58×10^{-4}
10	1134	6764.0	1641	10519.0	8.53×10^{-5}	2.866	2.44×10^{-4}
11	1641	10519.0	2215	35971.0	4.66×10^{-5}	11.260	5.25×10^{-4}
12	1213	7059.0	2215	35971.0	1.53×10^{-5}	4.627	7.08×10^{-5}
13	1630	10071.0	2215	35971.0	2.22×10^{-5}	11.260	2.51×10^{-4}
14	1347	7689.0	1630	10071.0	5.38×10^{-9}	1.971	1.06×10^{-8}
15	960	6450.0	1630	10071.0	2.58×10^{-10}	3.834	9.88×10^{-10}
16	1595	9969.0	1630	10071.0	2.00×10^{-10}	1.000	2.00×10^{-10}
17	1	0.0	1630	10071.0	1.00×10^{-10}	8.300	8.30×10^{-10}
18	1	0.0	1596	9971.0	1.00×10^{-10}	8.300	8.30×10^{-10}
19	1562	8461.0	1596	9971.0	2.00×10^{-10}	1.886	3.77×10^{-10}
CUF _{en} total =						2.0827	

(b) Results reproduced by the author using his own weighting scheme

Figure 5.28: Tables of individual environmental correction and usage factors for the NRC/EPRI “EAF Sample Problem 2-Rev. 2 (10/21/2011).” The reference method assigns the same F_{en} to the first two rows whilst the proposed lumping scheme does show a difference

- take into account that even within the finite element method, there is a wide variety of complexity in the problems that can be solved
- play with your favourite FEM solver (mine is [CAEplex](#)) solving simple cases, trying to predict the results and picturing the stress tensor and its eigenvectors in your imagination
- prove (with pencil and paper) that even though the principal stresses are not linear with respect to summation, they are linear with respect to multiplication
- grab any stress distribution from any of your FEM projects, compute the iso-stress curves and the draw normal lines to them to get acquainted with SCLs
- search online for “stress linearisation” (or “linearization” if you want) and then get a copy of ASME VIII Div 2 Annex 5-A
- take into account that FEM solutions lead only to nodal equilibrium but not point-wise equilibrium
- remember that welded materials with different thermal expansion coefficients may lead to fatigue under cyclic temperature changes
- try to find an explanation of the results obtained, just like we did when we explained the parametric curves from [fig. 5.18](#) with two opposing effects which were equal in magnitude around $d_b = 5$ inches
- think thermal-mechanical plus earthquakes cases as “bake, break and shake” problems
- understand why the elastic problem of the case study is still linear after all
- keep in mind that finite-elements are a means to get an engineering solution, not an end by themselves
- learn to write scripts to post-process FEM results (from an script-friendly open-source FEM program)
- work under a [distributed version control](#) system such as [Git](#), even when just editing input files or writing reports
- do not write ambiguous reports by replacing appropriate mathematical formulae with words just not to offend the illiterate
- try to avoid [proprietary](#) programs and favour [free and open source](#) ones.

About your favourite FEM program, ask yourself these two questions:

1. Does your favourite FEM program’s manual say what the program does?
2. Do you believe your favourite FEM program’s manual?
3. Do you trust your favourite FEM program?

And finally, make sure that at the end of the journey from college theory to an actual engineering problem your conscience is clear knowing that there exists

a report with your signature on it. That is why we all went to college in the first place.

5.9.1 Online stuff

Here is a list of sub-problems and stuff to play with.

- The pendulum-swing video from sec. 5.1
 - <https://youtu.be/Q-lKK4A20zA>
- “On convergence of linearized stresses in an infinite pipe computed using the finite element method” from secs. 5.3.2, 5.5.3, 5.3.2, 5.6
 - <https://www.seamplex.com/fino/doc/pipe-linearized/>
- The three cubes from sec. 5.5.3
 - case A: pure normal loads (<https://caeplex.com/p/d8fe>)
 - case B: pure shear loads (<https://caeplex.com/p/b494>)
 - case C: the combination of A & B (<https://caeplex.com/p/9899>)
- The animations of natural oscillations in fig. 5.10
 - <https://www.seamplex.com/docs/nafe4/mode1.webm>
 - <https://www.seamplex.com/docs/nafe4/mode2.webm>
 - ...
- The animations of natural oscillations in fig. 5.23
 - <https://www.seamplex.com/docs/nafe4/case-mode1.webm>
 - <https://www.seamplex.com/docs/nafe4/case-mode2.webm>
 - ...
- The parametric tee repository from sec. 5.7.2
 - <https://github.com/seamplex/tee>
- The environmental fatigue sample problem repository from secs. 5.8.1, 5.8.2
 - <https://github.com/seamplex/cufen>

See <https://www.seamplex.com/nafe4> for new material, updated links and the full version of this case with many more details about the case and the associated mathematics.

Chapter 6

Concluding Remarks