BRIDGING THE RELIABILITY GAP FROM HEAT EXCHANGER DESIGN TO OPERATION

Sean A Hennigan [1*]

[1]INEOS Acetyls, Saltend, Hull, HU12 8DS, United Kingdom

ABSTRACT

As an industry of users of heat transfer equipment, we are always interested in the next big thing that will help us transfer heat more effectively, save energy, minimize our environmental impact, save us money and emit less carbon. At the same time, we want projects on the ground to maximise our current potential, employing much effort from professional process and mechanical engineers within the organisation and through contractors / equipment suppliers to get designs of conventional heat exchanger equipment correct and optimised.

Within INEOS we tend to use more exotic metallurgy as the norm for key heat exchangers, so pressure is always on to reduce excess heat transfer area which directly gives capital savings. We employ the "workhorse" of heat transfer - *the shell and tube heat exchanger* - in its many guises and rely on understanding its performance.

Safe and reliable operation is paramount, and we have an array of tools to model, design and assess exchangers as we strive for frugal design and reliability.

Unfortunately, the real world has other plans. At times the industry has installed units that will not do what we want, that leak, fail, block up or otherwise not deliver on the design promise. This paper will illustrate a few examples from our experience in three main areas of where and why this has happened.

1. THE SHELL AND TUBE HEAT EXCHANGER

The shell and tube exchanger is an old friend when it comes to design and operation, dating back to the steam engine and probably well before.

Methods to predict heat transfer coefficients and pressure drops are continually improving; modelling tools have regularly updated correlations for the trickiest (usually two phase) situations and even throw new light on the simpler systems that we can directly plug into our "workhorses" or their more compact cousins. Combine this with deep insights and novel developments from the universities and academic institutions and we are in a good position to get increasingly better process performance out of our units.

We can design it to take it apart, have a good inspection and retain the convenience of being able to plug a few tubes off if they fail or leak and carry on without really noticing. Thousands of easily made tubes packed into a shell offer a bafflingly huge amount of area for the volume we are looking at and can be enhanced even further by grooving fins into the tube surface. The entire unit itself is really a set of cylinders all built inside each other, great at containing pressure and resisting vacuum, with predictable thermal and pressure stresses that are calculated by respected mechanical methods. We have detailed advice based on years of experience and other guides that code up the best way to build an exchanger going into detail that simulation models cannot.

*Corresponding Author: seanhennigan@ineos.com

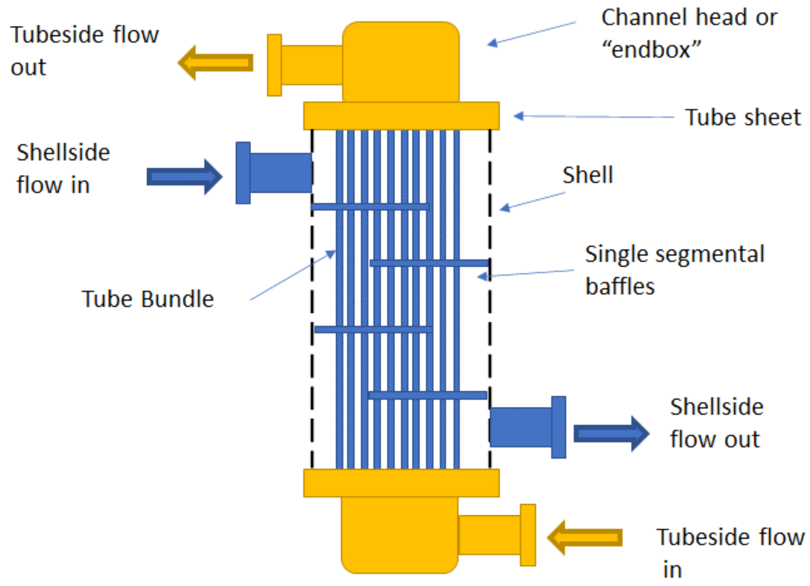

Figure 1: A typical layout of the shell and tube heat exchanger.

This example shows a vertical unit with fixed tubesheets in that the tubes are rigidly connected between drilled metal plates bolted or welded to the shell; if the tubes and shell are operated at very different temperatures there can be built up stress as hot tubes tend to thermally expand into cold shell side components or equally cold tubes can tend to pull hot shell components together. This can be relieved by making the shell flexible through an expansion joint. Some exchangers employ a U-Tube design where the bundle is only attached to one tubesheet and is therefore free to grow or shrink within the shell.

The flow on the shell side is directed by baffles to increase local velocity to enhance heat exchange on the shell side (at the cost of pressure drop). This same action also increases a potentially wearing velocity over the tubes and fluid vortex shedding in the shadow of the tube flow will induce movement of the tubes to a greater or lesser degree.

We can enhance heat transfer on the shell side with baffles and naturally we tend to get a high Nusselt number with the tubes as they have a small diameter. However, we still need to get heat through the metal wall of the tube and need to choose a material that allows this to happen easily whilst still having the required robust mechanical properties.

There are clues emerging here within these three areas particularly as to where we can still fail.

Corrosion of course is also a major source of failure, and we rely on metallurgical expertise to advise the correct material choice. This is a vast topic itself and this paper will concentrate more on the three areas above.

2. RELIABILITY PITFALLS AT THE DESIGN STAGE

2.1 The Cost of Unreliability

Without going into specifics, we can still get an idea of the costs involved here. A new heat exchanger for a world scale plant with exotic metallurgy may cost several million dollars to buy and install. We could work our optimisation magic; modern design tools allow you to simultaneously tweak the process design, calculate the mechanical requirements, metal thickness and hence total metal cost. Once set up this can be done very quickly, and we can directly link improved correlations and frugal process design decisions to real cash savings. It is a logical strategy to use up all your available pressure drop to maximise heat transfer coefficient and pay for less area.

The value of the capex saving is maybe of the order of millions of dollars for a highly optimised design, but if we have gone too far and this optimisation introduces reliability risks, then a world scale plant down for a few days to repair a tube could cost the same or more in lost revenue, and that is just the first failure. If multiple failures occur, we could lose all this capex saving. One benefit of shell and tube design is the ability to plug tubes, if these failures need extensive tube plugging we might be starting to compromise the whole plant performance. What if we come out of the tubes slightly hotter or colder because we are now blocking off tubes and heat transfer area? We have probably minimised excess area design margin in the capex reduction, so removing even more area is rarely good. Can the downstream plant handle the new temperature? What gives up first? We can't keep blocking tubes forever.

Getting this wrong at the design stage is clearly storing up trouble and costs for operations later, so where is the risk of failure creeping in?

2.2 Three Design Questions

Let's look at the process that drives design, building and installation of exchangers on site. Heat exchangers are unique in some respects as the thermal and mechanical aspects of design and operation are closely entwined. Process and mechanical engineers can approach a design from different angles, leaving behind a "reliability gap" on the handover from one to another where it may not be clear who is addressing the issue.

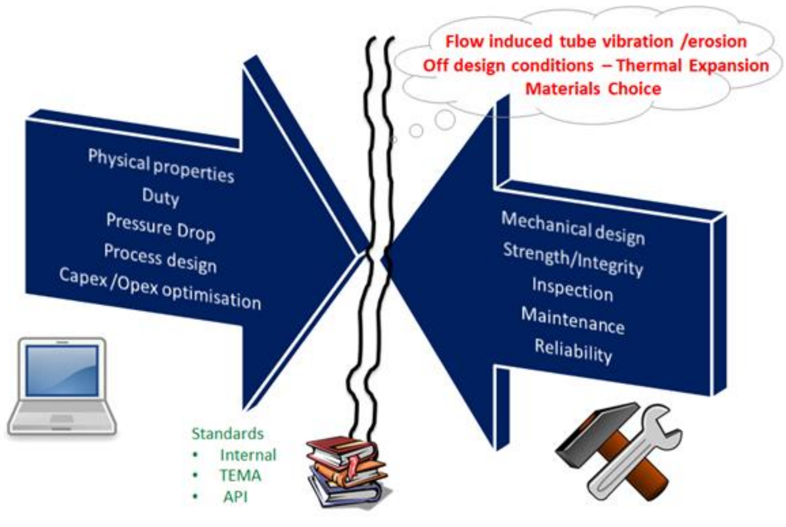

Figure 2: The reliability gap

A rigorous and holistic approach can bridge this gap by looking at three key areas

- **Erosion & Vibration:** Who will ensure the internals do not erode and tubes do not vibrate? This should be rigorously designed by the thermal designer and is as important as getting the heat transfer right. Unfortunately process engineers are not very comfortable with vibration and can be seen as a "mechanical" problem, but vibration is driven by the interaction of fluid flow patterns against the stiffness and support pattern of the tubes. The best modern heat exchanger design tools will be able to model this simultaneously with heat transfer. Sometimes the baffle layout may be set to avoid vibration as much as it is set to improve heat transfer. Only when we have a non-eroding, vibration free design should this be passed this to mechanical design.
- **Tube Material Properties:** Who decides correct materials of construction to get the right balance between heat transfer, strength and corrosion protection? This has got to be an interactive conversation. Materials good for heat transfer may be weak mechanically and vice versa, but very often the tube metal is not the dominant resistance for transferring heat. Giving the mechanical design flexibility on the choice of tube material may indeed have negligible impact, but not always! A stark example of this is shared further on.
- **Thermal Expansion:** Who ensures that the off-design conditions are addressed such as startup and shutdown? How do we ensure that both normal operation and off design conditions do not introduce extreme different metal temperatures between connected parts that build up stress? This should be at the forefront of the thermal engineers' mind as the design is finalised, but again is an interactive conversation that should involve mechanical engineers and operations specialists. Before any design is signed off, we need to know if any conditions of operation will introduce stress and if we need to use U-tube design or shell expansion joints to relieve it. In a sense it should be reviewed almost like a Hazard and Operability Study (HAZOP) to make sure all extreme conditions are captured.

2.3 Concept to Hardware

The process should not stop after these questions are addressed. There are a chain of activities taking a concept to hardware summarised in the schematic below.

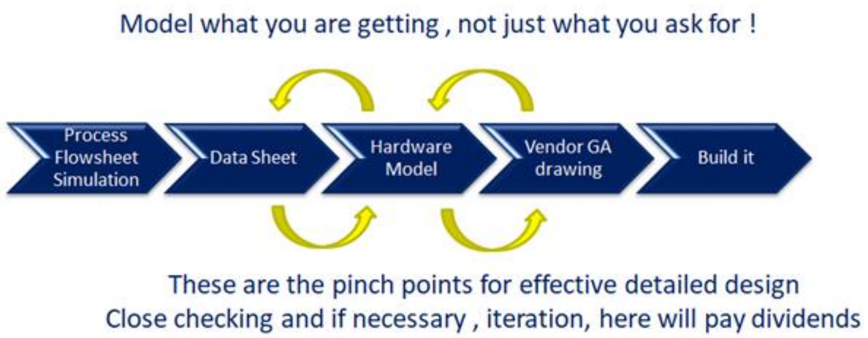

Figure 3: Reliability from concept to hardware

Key to this is the idea of understanding what you are getting, not what you asked for. The basic model of the unit can easily be tweaked once the detailed drawings and specifications emerge, there may be no appreciable difference, but you will have more details to work with. Mechanical requirements such as tube material choice, baffle to tube and baffle to shell clearances will feedback on the process design.

3. THE PERILS OF FLUID KINETIC ENERGY

The shell and tube exchanger has at first glance no moving parts therefore nothing to wear or erode, but the fluids themselves are the moving parts. Unlike rotors, stators or other solid objects they have the potential to become very unpredictable in terms of speed and location. We can and do wear and erode the internals (inside the shells and outside the tubes) through the sheer force of these fluids.

Maximising pressure drop usually means speeding things up, particularly if it is gases or vapours on the shellside, we are also always pushing for mega-size units to a scale we have not experienced before. There are advised limits on the fluid kinetic energy value (a function of fluid density and velocity) within the nozzles and the shell itself, based on years of experience. If this number is too high, the internals will see long term wear.

3.1 The Erosion and Vibration protection dilemma

Increased sizes of heat exchangers demand increased inlet and outlet nozzle sizes, driven by general good practice on line velocity but also the nozzle kinetic energy limits; basically we are setting a velocity limit to avoid erosion throughout the unit, but particularly at the shell entrance and exit where these velocities tend to be highest.

We support the tubes in the shell and increase the shell side velocity to get a good shell side coefficient using baffles. We want the feed to the exchanger to all go into the first baffle space otherwise we will not maximise shell side fluid contact at the correct velocity. By extension therefore the inlet baffle spacing must be wider than the inlet nozzle and ideally the nozzle should sit central in the inlet baffle spacing. As we make bigger exchangers eventually the tube length between the tubesheet and the first baffle must increase leading to a second problem.

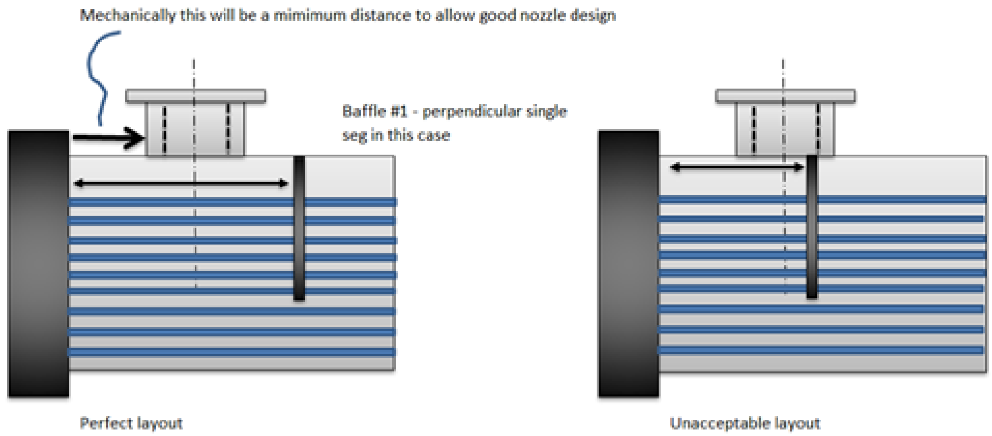

Figure 4: The preferred inlet baffle layout

As we go wider in nozzle size to manage erosion risk we get longer in unsupported tube length which introduces a flow induced tube vibration risk through fluid vortex shedding.

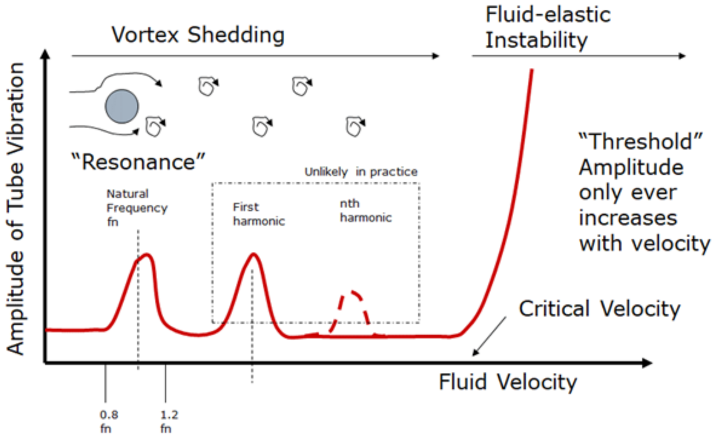

Figure 5: The vortex shedding phenomenon over heat exchanger tubes

The shedding of fluid vortices behind the tube induces an alternating force at a predictable frequency, if this frequency gets close to the natural frequency of the tube a resonant "lock -in" will make the tube move. With most exchangers in the author's experience large scale plant shell flows often exceed the tube natural frequency but usually the resulting tube deflection is small. The severity of movement is proportional to the unsupported length to the fourth power.

If we go too far in fluid velocity, we hit a condition called fluid-elastic instability where the tube movement itself interacts with the fluid to produce chaotic negative damping effect that can instantly break the tube.

Eventually we hit a paradox in that we want a slow nozzle and inlet flow and the inlet baffle spacing will reduce the size of that problem in proportion to the length squared, but the price to pay is a long unsupported length which introduces vibration problems that increase with length to fourth power. Ironically one potential damage mechanism is slowly solved whilst a second one is quickly exacerbated.

3.2 Impingement protection and tube support

Flow induced tube vibration is extremely complex but can be predicted using modelling tools if we have an idea of tube and support mechanical properties and fluid properties and flows.

This is further confused by the use of impingement plates to protect the top row of tubes from direct impact of fluid (and maybe incoming debris) from the nozzle, in some instances this can actually increase local internal velocity as shell inlet area is closed off and make vibration and long-term erosion worse, even though nominally it is a protection. Carefully considered design and some degree of judgement as to the size and location of the worst possible velocity is key. Modern designs favour the use of impingement rods rather than plates, where the first one or two rows of tubes are replaced with solid rods supported at the first baffle that follow the same or similar tube layout pattern. These offer a good compromise between protection from the incoming fluid without introducing new fluid accelerations into the shell.

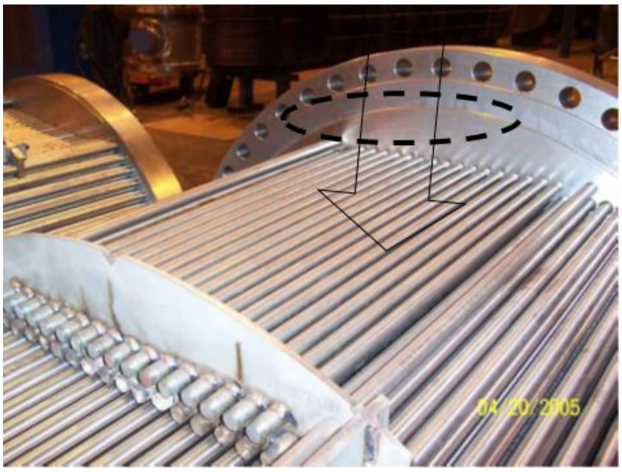

Figure 6: Typical layout of impingement rods in a heat exchanger bundle. (INEOS internal image file [1])

Impingement rods however will not stop vortex shedding and the potential for tube vibration needs to be dealt with at a fundamental level in the design.

The issue can be solved many ways; the simplest is by a support plate which reduces the unsupported length on the most vulnerable tubes. This is like an extra baffle but does not take part in the enhancement of heat transfer on the shell side and should not impede flow.

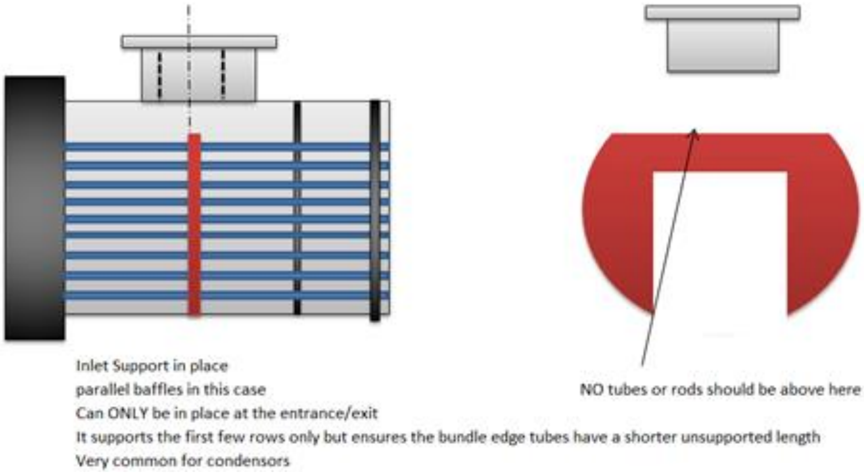

Figure 7: Example of an inlet support plate for a double segmental baffle

At even bigger scale even this may not be enough, and we are driven to a "No Tubes in Window" or NTIW design where we completely avoid tubes in the potentially long unsupported length necessary to fit tubes in the baffle window. In doing this we can add several extra supports with the baffle spacing without blocking the main flow path.

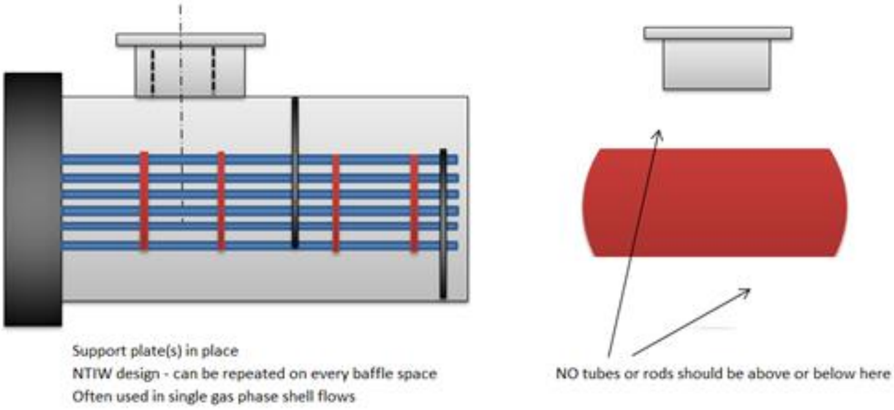

Figure 8: Example of several inlet support plates for NTIW design.

The cost of this is that we do not fill the shell with tubes efficiently, but we do get a very low unsupported length and problematic tube vibration can be eliminated.

3.3 Chronic and acute vibration

There are chronic and acute vibration scenarios driven by different mechanisms. Chronic failure is driven simply by day-in / day-out microscopic tube movement of gentle resonance. This is often seen as circumferential wear on the tube in the baffle hole which becomes self-exacerbating as the baffle hole is made bigger and eliminates the support point for the tube leading to effectively longer unsupported length.

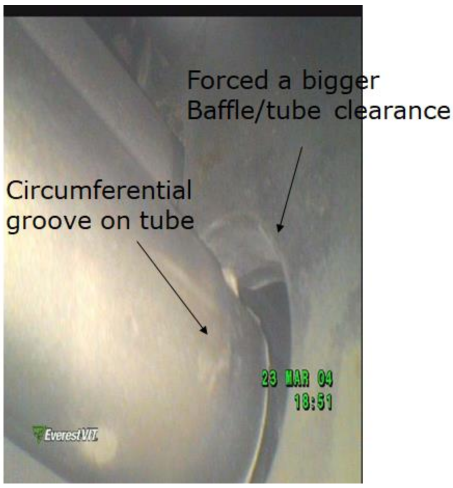

Figure 9: Example of observed external tube wear and baffle hole wear [2]

This kind of damage may not serious enough to be a problem for decades (it can be checked using in-tube Eddy Current Testing and other techniques) but is insidious and can be hastened as plants debottleneck and push more material through the shell, increasing fluid vortex shedding. Fluid-elastic instability failure is an acute event that will quickly show itself if the design detail has been missed.

3.4 An example of fluid-elastic failure

An example of this within INEOS was a gas/gas intercooler installed as part of a new plant. Vibration problems were anticipated as the exchanger was an NTIW design with 4 intermediate support plates per baffle space. Shortly after commissioning distinct and intermittent sounds were heard from the unit at around 36 dBA and 9 Hz frequency. Acoustic vibration is a separate (equally problematic but also solvable) phenomenon but would produce much louder noise than this so was quickly ruled out, leaving the mystery of how internal damage could be occurring. No obvious process upset was seen at the time.

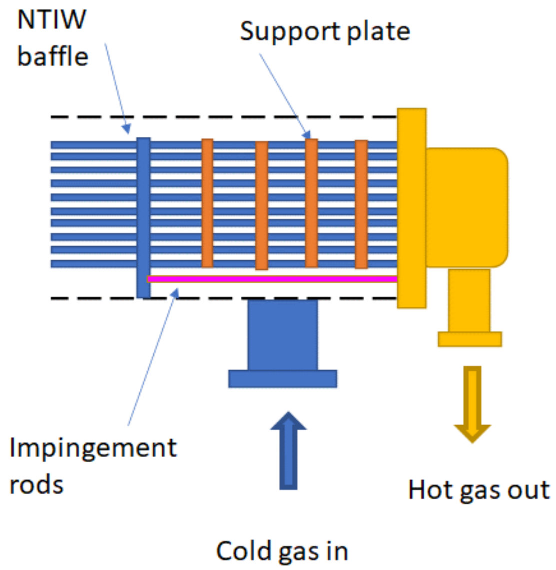

Figure 10: Example of gas cooler layout with NTIW baffles.

The actual cause of the problem was only obvious after careful analysis of the as-built drawings which revealed that the protective impingement rods on the inlet were only supported between the tubesheet and the first baffle, not on the 4 support plates. As a result of these rods have 4 times the unsupported length of the tubes they were 4^4 =256 times more susceptible to damage than the tubes and failed probably from day 1.

Figure 11: Impingement rod damage as inspected [2]

The damage is clear from this picture and pretty much every rod was broken whereas the well supported tubes remained intact. Fast action detected this early enough before the loose rods could cause more damage, the industry has many examples where tubes themselves have broken; usually showing up as unacceptable contamination on the low pressure side that needs a shutdown to find and plug the broken tube.

Understanding of vibration issues is a relatively recent development (E.A.D. Saunders *Heat Exchangers* 1988 Chapter 11 [2]) as plants and exchangers simply got physically so big to the point where this critical point is now reached regularly with world scale designs.

4. A MATERIAL CHOICE NIGHTMARE

This is a further example of where the ongoing conversation between the process and mechanical engineers pays dividends.

Uncertainty is a common factor in heat exchanger design and we have found the clearest way to deal with this is by defining the Area Design Margin (ADM).

$$ADM = \left[\frac{U_C}{U_o} - 1\right] \times 100\% \qquad (1)$$

Uc = clean coefficient (W/m2K)

Uo = overall coefficient in service (W/m2K)

For example ADM = 30% implies that the exchanger has 30% more area than it needs to satisfy the design duty at the design temperature driving force. This allows for a degree of fouling, or other uncertainty in, say, correlation accuracy to give the design some headroom to deliver what it should. Importantly it does *not* mean we will get 30% more duty from the unit, if we need extra duty out of the unit it should be designed in from the outset.

We cannot rely on margin to mean a "free" debottleneck of the unit in the future, particularly in a world where we are building large units with exotic materials and capital is expensive; there is continual pressure to "right-size" exchangers so that they deliver the requirement to do the job and no more. Again, this plays into the reliability issue, we need to ensure the unit will perform its function, but we can be in a situation where even defining the uncertainty can be confusing! The concepts of fouling factor, design margin and overdesign are often conflated.

4.1 Fouling factors and area design margin

Traditionally the margin would be addressed with fouling factors, essentially a notional fouling thickness divided by the foulant's conductivity, added as extra resistances in the overall heat transfer coefficient equation.

$$\frac{1}{Uo} = \frac{1}{h_{io}} + f_{io} + \frac{1}{h_0} + f_0 + \frac{do \ln\left(\frac{d_0}{d_i}\right)}{2k_w} \qquad (2)$$

U_o = overall heat transfer coefficient (W/m2K)

h_{io} = inside coefficient adjusted* to outside area (W/m2K)

h_o = outside coefficient (W/m2K)

f_o = outside fouling resistance (Km2/W)

f_{io}= inside fouling resistance adjusted to outside area (W/m2K)

k_w = wall thermal conductivity (W/mK)

d_o = outside tube diameter (m)

d_i = inside tube diameter (m)

Arranging equation (2) in the form below allows one to view the sum of resistances to heat transfer, the relative amount of each term can be thought of as the fraction of total exchanger area allocated to overcoming this overall resistance.

$$R_T = R_f + (R_t + R_S) + R_w \qquad (3)$$

Total resistance = fouling resistance + convective resistances + metal wall resistance (all m2W/K)

The notional example below is one way to look it; this assumes we have water in the tubes and an organic in the shell with stainless steel tube material and typical convective coefficients.

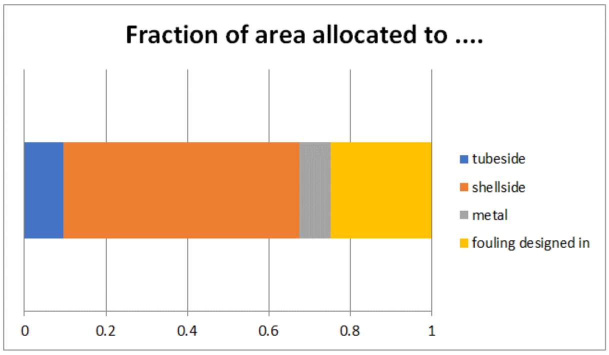

Tubeside h	14000	W/m2K
Shellside h	2000	W/m2K
metal k	20	W/mK
Rf	0.000215	Km2/W
ADM	33%	

Figure 12: A typical fractional area allocation of resistances

Clearly the low organic coefficient on the shell is dominating the heat transfer and the fraction of area allocated to fouling is a reasonable 25% by applying a fouling resistance of 0.0001 Km2/W on each side. The tube wall metal resistance is very small to the point where we could almost choose any vaguely similar metal and would have little impact on the ADM.

Note the fraction of area assigned to deal with fouling (FR) is not the same as the ADM.

When FR is very small (around 10% or less) then ADM approximates FR. ADM looks at overall fouled to clean ratio, FR tells you how much area is allocated relatively to fouling. If the fouling resistance accounts for half the area (FR =50%) then the exchanger is twice as big as the no fouling case so ADM =200%

This is a very useful way to look at the choice of margin via ADM as the fouling factor is a confusingly small number. If we find ourselves building an exchanger with say twice the clean area requirement we have either hopelessly overestimated the amount of fouling or are not dealing with it in the correct way.

In many cases we can take a very pragmatic view and simply ignore fouling factors. If we have a similar process duty and exchanger type on a similar plant we can use plant data to imply U_o, use models of the unit to calculate U_c and simply plug into equation (1) to get the ADM for use in the next design.

4.2 An example of a very wrong margin

The importance of fully understanding the concept of margin is demonstrated in this real example of a new heater design. The heater itself is almost a standard design with a well understood performance, a reasonable and “safe” ADM requirement of around 25% was built in (FR ~ 20%) to accommodate any fouling beyond the norm.

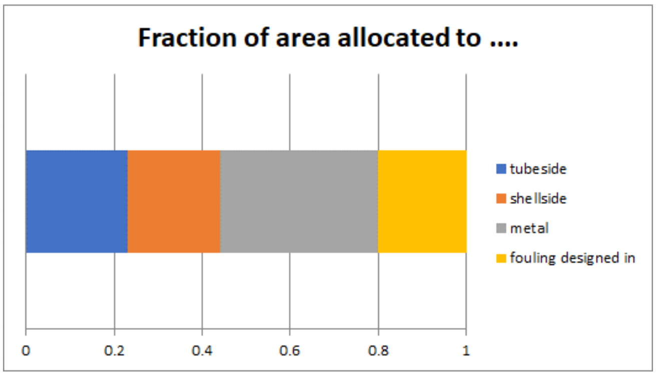

Figure 13: Fractional area allocation for this example

The area allocation picture shows that unusually the metal is around 36% of the total resistance with fouling built in, but more importantly almost half of the clean resistance.

The tube material in this case is exotic, costly and also has challenging mechanical properties. During detailed mechanical design a decision was taken to use a different grade of material which doubled the tensile strength allowing modest savings on the shell and tubesheet thickness costs. On startup of this unit it became clear that the unit was not delivering the correct outlet temperature and a set of operational attempts to resolve this came to the same conclusion that it was not delivering enough heat transfer performance, despite the known ADM built in.

Deeper analysis of the properties of this new grade of tube material revealed the fundamental issue; the new material was highly alloyed and its thermal conductivity was about half that of the standard grade, thus doubling the metal resistance.

This increased resistance therefore added an extra 50% to the *clean* area requirement even before any consideration of fouling, thus comfortably wiping out all the ADM for fouling and still leaving us 25% short on area, even if clean!

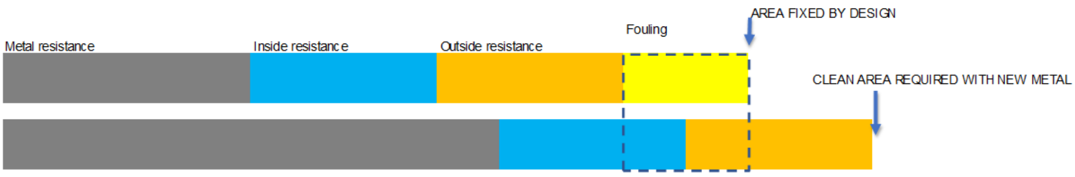

Figure 13: Build-up of resistance area requirement picture with standard and alloy grade tubes

This shortfall was resolved by relocating a spare exchanger of a similar design after this unit, but the plant suffered many months struggling to get to design rate.

This kind of situation where the heat transfer coefficients are very high on both sides of the exchanger and the metal resistance becomes very important also occurs with water / water systems such as surface condensors, chillers and water interchangers. This has traditionally led to choosing very high conductivity copper-based metals to get the total exchanger size lower.

5. THERMAL EXPANSION PITFALLS

As heat transfer relies on a temperature gradient to work all shell and tube heat exchangers will naturally have different temperatures of metal between the shell and tube materials. In normal operation this could be significant and built up tensile or compressive stress from thermal expansion of one metal component into the other needs to be managed. Less obvious is that we need to go from an exchanger where all the metal is at the same ambient temperature to a condition where one component is significantly hotter than the other. Literally something has "got to give" and we often go for a flexible element in the shell known as an expansion joint that allows the shell to grow or shrink as temperatures change.

Figure 14: Typical shell expansion joint (bellows type) [2]

Small temperature differences of say tens of degrees may not build up too much stress, but more extreme conditions will lead to buckling of tubes. If the tubes are hotter than the shell then the tubes expand and push towards tubesheets; if the shell cannot expand, then the shell feels a tensile stress (pulled apart) with no immediate issues unless a material dependent critical stress is reached and the shell could be damaged. The tubes will feel a compressive stress (crushing) and at a critical stress the tubes will buckle in complex modal shapes. In compressive stress, even without buckling the tubes have a lower natural frequency hence more likely to vibrate. If the tubes are colder than the shell, the tubes will contract and tend to pull the tubesheets together. If the shell cannot contract then the shell will feel a compressive stress and the tubes a tensile stress. In the mechanical design a margin is set to avoid exceeding critical stresses but crucially all the scenarios for possible thermal differences need to be understood.

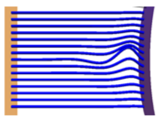

Figure 15: Tube buckling (tubes hotter than shell)

5.1 Setting all thermal scenarios

We need to have a list of all shell and tube metal temperatures under all scenarios to design out excessive stress. Normal operation is an obvious starting point and generally speaking the shell temperatures will be approximately the average of the shell fluid inlet and shell fluid outlet temperatures, the temperature of the tubes themselves do not directly affect the shell metal temperature. Tube metal temperature is more complex as is depends on all the shell and tube fluid temperatures in and out. Design software will calculate this directly.

Tube-only startup and shutdown, shell-only startup and shutdown and other non-normal conditions all need to be considered. The worst case is usually a hot tube startup within a cold empty shell; the tubes are unable to transmit temperature to the shell quickly so will expand instantly within the cold shell and all the tube metal will potentially rise to the fluid inlet temperature. A shell start scenario assumes that the shell fluid will quickly get all the shell and tube metals to the same (shell fluid inlet) temperature so no temperature difference (but there could be pressure difference). An easy to miss example of this can be seen during plant turnarounds where often tubes are steamed out when offline to clean them with the shell at ambient conditions. This may give a much bigger temperature difference between the tube and shell metal than normal operation.

A review is required with all process, mechanical and operation engineers and each credible scenario listed in table that will go with the exchanger data sheet. This ensures that the mechanical design is aware of all possible metal temperature differences not just those in normal operation. Each scenario should include the temperature extremes but also the pressure extremes which can also affect stress. Table 1 below is an example of conditions to consider as a starting point at least.

Table 1: Thermal Expansion Scenarios

Scenario	Shell metal temperature	Shell pressure	Tube metal temperature	Tube Pressure
Normal Operation	Average shell fluid in/out	Design	Consider all fluid temperatures	Design
Shell side start	Shell fluid in	Design	Shell fluid in	Ambient
Tubeside start	Ambient	Ambient	Tube fluid in	Design
Shell shutdown	Normal	Ambient	Normal	Design
Tube shutdown	Normal	Design	Normal	Ambient

5.2 Problems with expansion joints

If we have no choice but to add a flexible element to the design, there is a trade-off to consider. The joint itself is, almost by design, a weak point. It has to be flexible enough to move in line with the expansion but needs to be strong enough to contain the shell side material. If that material is toxic this may not even be an option and we may have to go for a U tube design. U tubes sound like generally preferable option but can have problems in fouling duties as the U bend is difficult to clean, plus we are immediately forced to a multi-pass design which can have some downsides in reduced temperature driving force.

If these limitations force us down the expansion joint route there are further issues to consider, as well as the joint moving the entire exchanger also has to move and layout of the exchanger supports and pipework is also critical.

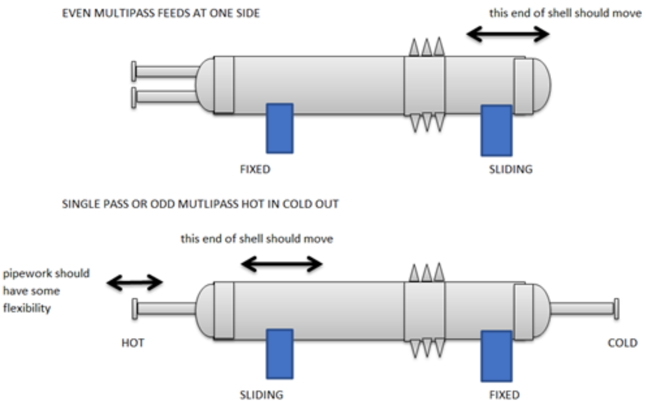

Figure 16: Exchanger supports and flexibility

As mentioned earlier this degree of flexibility brings inherent weakness, but also specifying the incorrect requirements of what you want the do can also cause damage, so there is a tension between overdesigning the joint making it weak and trying to eliminate it risking tube failure. These are several examples where the expansion joints have failed.

Figure 17: Buckled tubes, corroded expansion joint and "squirmed" expansion joint [2]

The left picture is an example of under specifying the thermal expansion requirements where the tubes, baffles and supporting bundle tie rods have all buckled (the joint has been cut out to access tubes). The centre picture is an example of a joint failing under corrosion. The thinner material and presence of dead spots in the joint makes is more susceptible to corrosion. In this case normal operation of the shell was not problematic but was washed with a chemical during a turnaround which was incompatible with the joint material. The final right picture is an example of a joint with too much flexibility to the point where the two rigid parts of the shell become misaligned undergoing a phenomenon known as "squirm" which can further weaken the integrity of pressure containment and prevent the joint from moving to the degree required.

6. CONCLUSIONS

Design of reliable heat exchangers is not easy and there are many potential pitfalls. Three areas of concern are highlighted in this paper where bridging potential gaps between process design and mechanical design issues can avoid failure later-on. Early understanding of the erosive and vibrational potential of incoming fluid and designing out high velocities and long unsupported lengths will help keep the unit running as it is pushed hard. Keeping on top of material choices and other mechanical details as they emerge and modelling what you think you are getting not just what you asked for will avoid potentially severe under-sizing of a unit. Finally spelling out the full range of metal temperatures and corresponding stresses the unit will be asked to see and making a carefully considered choice as to the best way to relieve this will save your unit from damage, especially during plant excursions, in the future.

ACKNOWLEDGEMENTS

I would like to acknowledge the hard work of the many people I have consulted and been consulted by who work at the frontline of plant operations, engineering and technology within INEOS that have helped build up this bank of knowledge.

REFERENCES

[1] INEOS Internal Image file.
[2] E.A.D. Saunders *Heat Exchangers* (1988)