

Chapter 5

Designing for Service in a Complex Product Service System-Civil Aerospace Gas Turbine Case Study

Andrew Harrison

Abstract Gas turbine engines used in commercial aviation are a complex engineering product for which the cost of maintenance over their full operational life can exceed the initial cost of manufacture by a factor of 3–4. With an increasing shift to contracting fixed price maintenance services at the point of sale of the engine, it has become an imperative that the cost of maintenance support is a prime design attribute. This chapter provides a case study on the process of designing for an optimal maintenance cost level from the earliest stages of conceptual design through to product entry into service. The nature of an optimal maintenance cost in terms of customer service value is also explored.

5.1 Case Study Context: The Rolls-Royce Aero Engine/TotalCare[®] Service

Rolls-Royce is a global company with a portfolio of products and services predominantly in power generation¹ systems across a range of applications.

The majority of our products have a long lifespan of typically 25 years or more and many of our product lines will have an in-service span (1st production to final disposal) in excess of 50 years.

Rolls-Royce is widely credited with bringing the *Power by the Hour*[®] service contract to the civil aviation market with its TotalCare[®] service support contract, in which the engine operator contracts a fixed price suite² of support services on a \$

¹Rolls-Royce has 49,900 employees worldwide and is active in the Commercial Aerospace, Defence Aerospace, Commercial and Naval Marine and Energy sectors (as of 2016) [1].

²TotalCare[®] comprises a suite of engine support services including; proactive in-service engine management, overhaul shop maintenance and the collection and management of engine data.

A. Harrison (✉)

Engineering Associate Fellow (Life Cycle Engineering), Rolls-Royce Plc,
PO Box 31, Derby DE24 8BJ, UK
e-mail: andrew.harrison@rolls-royce.com

per flying hour basis. TotalCare[®] contracts can cover just a few years of operation up to the full operational life of the product. This contrasts with the traditional contracting model in which the operator commissions and pays for individual maintenance events as they are required (commonly known as Time and Material contracting).

With a significant proportion of engine sales and support contracts being contracted during the product design and development phase³ there is a significant transfer of cost management risk from the operator to the supplier. Maintenance contract prices are often established a decade ahead of statistically significant service cost data being available.

The advantages of TotalCare[®] are the alignment of operator and engine supplier benefits and incentives. Under TotalCare[®] it is clearly in Rolls-Royces interests to design a product that not only meets the operational guarantees but also delivers a cost of support in line with the contracted rate. The return for Rolls-Royce is the stability of committed future service revenue streams.

It is this transfer of long term support cost risk to the supplier that brings the need to design for service into such sharp focus.

5.2 The Costs of Owning a Gas Turbine Aero Engine

The following case study is based around a gas turbine engine for the civil aviation market (visualize the engines hanging from the wings of the last commercial aircraft you travelled on). Whilst the examples may be specific to this application many of the features of the scenario are common to the majority of complex, high value, long life products.

One of the major difficulties with designing for service contract risk management is the sheer complexity of the cost structure and the number of degrees of freedom that impact its evaluation. The following sections describe in outline the key elements of the cost breakdown and some of the primary driving factors that influence them.

In general all costs associated with the ownership of an engine fall into one of two categories:

- Direct consumables—costs incurred as a consequence of normal operation and performance of the engine (fuel, oil).
- Loss of functionality costs—costs associated with the prevention or recovery from loss of engine functionality (due to degradation, damage, or other causes).

³Rolls-Royce announced further sales of Trent XWB engines with contracted TotalCare[®] services at the 2009 Paris air-show 4 years ahead of the expected entry into service date of the aircraft.

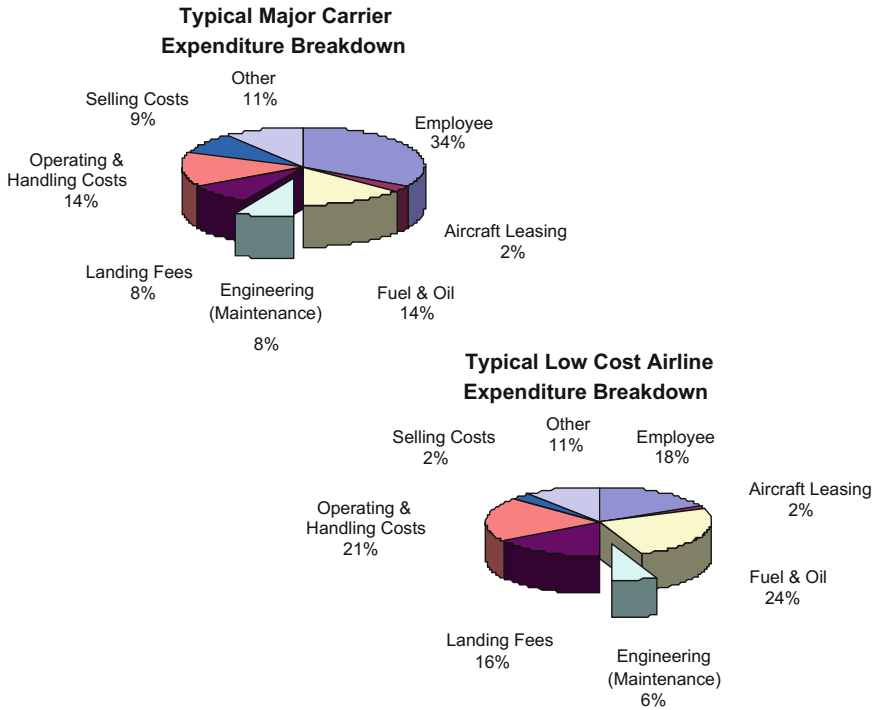


Fig. 5.1 Maintenance costs as a proportion of overall operating cost for typical low cost and major carrier airlines circa 2005

Of these, fuel typically dominates the cost of engine maintenance by as much as 5–8 times in some customer operations [2, 3].⁴ Maintenance and support activity clearly has a significant impact on all of the above costs as it impacts the engines health in terms of both functionality and performance efficiency (fuel consumption) (Fig. 5.1).

5.2.1 Indirect Operational Disruption Costs

Disruption costs are incurred as a result of additional work or activity to mitigate the operational impact of loss of functionality of an engine.

Unless covered by warrantee these costs typically fall directly to the aircraft operator.

⁴Fuel price increases since these 2003/2006 conference presentations/papers have significantly increased the ratio of fuel price to overall aircraft maintenance cost. Also engine maintenance cost is only a sub-set of the total Aircraft maintenance cost.

For example:

- Lease costs for a replacement engine/aircraft
- Additional Airport landing fees
- Cost of compensation of passengers for delay or cancellation
- Cost of accommodation of passengers for overnight cancellations
- Repositioning of flight crew
- Recovery of aircraft requiring remote location engine change

Such costs are highly variable and dependent upon the location, severity and timing of disruptive events.

Example: Imagine a large civil aircraft carrying 250 passengers that experiences a cockpit warning message during engine startup. The aircraft is scheduled to fly in the last 30 min of the airports operational day before night time noise curfews come into force. The passengers have started the embarkation process and are taking their seats on the aircraft.

- If the message requires a change of an oil or fuel filter (high filter pressure drop detected):
 - A filter change may be possible via a small engine cowl access door within 10 min. The aircraft possibly leaves a few minutes late.
 - If the filter is not accessible via an access door it may be necessary to raise the engine cowls (large doors covering the majority of the side of the engine). Question: as a passenger how would you feel entering an aircraft where the mechanics appear to be ‘fixing’ the engine? This alone can cause a policy of no ‘open cowl’ activity within sight of the passengers or during embarkation/debarkation (approx 20 min after landing or before take-off):— The aircraft is in danger of failing to depart before the curfew, incurring overnight accommodation costs for 250 passengers, flight schedule impacts to address out of position aircraft and crew (the stranded aircraft/crew should have been somewhere else in the morning requiring a spare or lease aircraft to fulfill its schedule), missed passenger connections and the re-seating administrative burden. Additionally there are intangible costs of loss of customer confidence and lost future revenue.⁵
- If the message requires a more invasive maintenance activity it is often more cost effective to remove the engine and replace it to allow the aircraft to return to service. Depending upon the location of the stranded aircraft this may involve shipping of a spare engine and work crew to perform the change, leaving the aircraft out of service for 24 h or more.

⁵In Japan domestic flights compete with the Bullet Train which enjoys a near 100% on time departure record. Even minor delays (minutes) in aircraft departure schedules have a significant impact in customer perception and loyalty.

5.2.2 *Installed Costs*

The installed costs are those direct costs incurred on the engine whilst it remains attached to the aircraft. This is dominated by the consumable cost of fuel and oil but does include a level of installed maintenance, both preventative and reactive.

5.2.2.1 **Consumables**

Fuel is the dominant cost in the consumables arena. Whilst it is not a cost that falls to the engine supplier it is typically a major factor in the customers' engine and aircraft selection.⁶

Fuel cost is dominated by the efficiency characteristics of the basic engine design and the fuel price. Up to 3–5% of additional fuel consumption rate may occur prior to major overhaul as a result of engine deterioration however. This *performance deterioration* fuel burn penalty is an integral function of the rate of performance loss and the frequency and quality of performance recovery at overhaul. Lower loss rates, more frequent overhaul and better recovery at overhaul all benefit through life fuel burn.

5.2.2.2 **Installed Maintenance Burden**

Installed maintenance is a mix of preventative activity such as routine borescope⁷ inspection, filter changes, engine washing (fuel burn recovery) and engine health condition monitoring. The costs are generally related to the frequency, manpower and equipment required but are massively complicated by the mechanics and constraints of operating out of commercial airfields. These costs are often dictated by some fairly simple considerations.

For example consider engine maintenance access. It could be considered good practice to place the fuel filter at the top of the engine where the fuel enters from the aircraft wing immediately before it is split off to feed the fuel injection system as this minimizes the length of pipe run, weight and ultimately fuel burn cost. However next time you are at an airport have a look out of the window and see if you can see any maintenance stands. They are typically a simple selection of working platforms sufficient for one or two mechanics to stand on. They come in

⁶For the majority of large civil aircraft the customer airline contracts purchase of the basic aircraft from an airframe manufacturer (Boeing, Airbus and several others) but then has a choice of engine suppliers. Each engine has to have been designed and certified for the Aircraft type, which due to the huge costs involved typically limits the choice to just two engine options.

⁷A borescope is a rigid or flexible optical inspection device that allows internal inspection of the engine through relatively small access holes strategically placed within the engines external and internal casing structure.

fairly standard heights up to 1.5 ms and are often placed at all the designated maintenance stands around most airports.⁸

What happens if your engine design places its filter 4 meters off the ground? Are you going to demand 2.5 m tall mechanics or a special stand just for your engine? A special stand means either a huge infrastructure cost (stand cost * airport count * average number of maintenance bays) or poor health and safety practices (stands on stands). Either way it is far better to consider the ergonomics and practicalities of the maintenance activity at the design stage.

The frequency and duration of maintenance is the other key consideration. Aircraft have natural windows of maintenance opportunity. There is the 20 min after the passengers offload and the aircraft is being cleaned (before the next set of passengers board), but remember that the passengers don't want to see the sides of the engine exposed and maintenance crews in action.

There is the overnight airport flight curfew when many aircraft are unable to fly for noise reasons. This can provide a 6–8 h maintenance window.

There are scheduled aircraft maintenance windows (like the 10,000 mile or annual service of your car) that can provide even longer maintenance windows.

Quite reasonably aircraft operators want no additional burden or disruption from installed engine maintenance. They want all installed maintenance needs to match the available windows of opportunity, and as infrequently as possible to minimize mechanic and equipment costs. This raises the issue of advanced warning of maintenance needs. The longer the maintenance activity takes the less frequent the opportunity for disruption free maintenance is. Therefore:

- Actions that take 5 min can be accommodated with only a few hours notice (time to schedule the mechanics).
- Actions that take >20 min (overnight window at an airport with suitable maintenance facilities available to the airline) require several flights notice. Note that an individual aircraft may have up to 8 flights scheduled in the day or have a route structure that returns it to its home airport only once a week.
- Actions requiring longer than the overnight window may require several months of continued safe operational time in order to avoid a special maintenance *out of service* event.

Monitoring the engines health to provide such warning is a major requirement to avoid the costs of disruption.

We also have to consider the practicalities of airport life. Some maintenance activity requires the engine to be tested (potentially up to take off power levels) after completion to ensure absolute safety. E.g. after changing some control or fuel system elements. At many airfields this may require the aircraft to be moved to a special location for noise or airport personnel health and safety reasons. It may even

⁸The author vividly recalls a personal blinding realization of the complete disconnect between the well lit, CADD enabled design office and the reality of Narita Airports maintenance stands (on a cold wet windy afternoon in failing light) as our airline host politely pointed out the practicality of some of some early conceptual design features.

require the allocation of a normal take off slot on the runway with the loss of airport capacity. This immediately runs into scheduling issues; can the aircraft be moved or does it need a break in other aircraft traffic movements? Can such tests be conducted in the night time noise curfew? How long will it take to reposition the aircraft, test it and return it to the departure gate? These considerations can turn an apparent 30 min maintenance activity into a 24 h elapsed time event.

5.2.2.3 Monitoring Costs

Modern engines are generally equipped with some level of health monitoring capability. At the basic level this may be limited to self checking algorithms embedded within the electronic engine controllers that can display a cockpit message when an error state is detected. At the highest end of sophistication health monitoring becomes a major feat of technology. For example operators of the latest generation Rolls-Royce Trent Engine family which power the most recent Airbus and Boeing aircraft are supported by a 24/7 operations centre based in Derby in the United Kingdom with access to near real-time engine health status.

The health monitoring system is embedded in the architecture of the engine with access to the standard engine control parameter sensor data as well as some dedicated sensors specifically for health detection. Onboard algorithms constantly check these measurements against expected levels in order to detect the earliest indications of abnormal behavior. Detection methodologies vary but tend to centre around absolute deviations from expectation, comparative differences (detector A versus detector B) or rate of change of parameters over time.

Anomaly detection action is then graded according to the type of anomaly:

- Set a maintenance message that will be available at the next maintenance inspection (mechanic plugs in a message downloader).
- Send a message to the pilots display providing them with information to which they may need or choose to react.
- Send a message to a ground station by satellite (during flight) providing information for more investigation by ground based support systems.⁹

Additionally a certain amount of summary data of the parameters recorded during each flight are routinely collected and provided for automatic or requested download when the aircraft reaches an enabled airport.

Once data is available to the ground based systems it is run through a barrage of analysis that establishes a measure of the engines health. Any anomalies that are detected can then either be automatically triaged and the appropriate people

⁹Imagine the equivalent whilst driving your car. The onboard monitoring notices an approaching low brake pad thickness indication and it texts a warning message to your local garage. When you arrive at your home that evening the local mechanic is there waiting with a set of replacement pads and all the equipment required to change them for you. You are never exposed to the potential safety issue of poor braking or the additional cost of metal on metal damage to your brake discs.

informed (e.g. e-mail/text message to airline maintenance staff), or can be referred to the 24/7 help desk and their supporting expert support for decisions on the most appropriate reaction to recommend.

Consider the previous example of the filter change being required. Imagine the benefits of in flight or ground analysis based detection of a gradually increasing filter pressure drop. Weeks in advance of the ‘do not dispatch aircraft’ message the maintenance team is contacted requesting procurement and replacement of the filter. The activity can then be scheduled for a convenient location and overnight maintenance slot with no risk of customer and operational disruption.

There is of course a cost associated with provision of this type of monitoring service in the form of communications costs (satellite and ground based), data warehousing and analysis, continuous development and refinement of algorithms for detection and the staffing of help desk and supporting technical staff. However this generally proves to be low relative to the major benefits in operator schedule adherence and reduction in indirect disruption costs (individual disruption events can be valued in \$ millions).

5.2.3 *Uninstalled Costs—Maintenance*

We can breakdown the uninstalled costs further to the following set. These broadly align with the mechanics of the maintenance process:

- Removal of the engine
- Replacement of the engine (substitute engine asset)
- Transport of the engine
- Disassembly into major sub assemblies (Modules)
- Module disassembly
- Inspection
- Repair or replace
- Rebuild engine
- Test engine
- Transport engine
- Re-install engine

5.2.3.1 Removal of the Engine

Removal costs are broadly associated with the dimensions of time, manpower and equipment.

The time element impacts the costs associated with all the assets occupied during the period of removal. This includes the aircraft itself.

If the engine change (removal and installation of replacement) can be completed in a 6–8 h shift then it can probably be achieved during a scheduled overnight

downtime of the aircraft with no disruption to its planned operational schedule. If it can't then a replacement aircraft will have to be utilized with lease costs (or spare aircraft capacity inventory costs). Alternatively the replacement will be constrained to coincide with other planned aircraft maintenance that already requires the additional downtime.

Manpower and equipment costs relate to the number of events that occur and where they occur.

Take for example the case of a maintenance requirement that causes an aircraft to be stranded away from its home base, at an airport without its own maintenance facilities. A *remote site rescue team* and a spare engine may have to be dispatched to travel to the stranded aircraft, conduct the engine change and then recover both the aircraft and the removed engine back to the home base. The costs then include the travel costs for the maintenance team, spare engine, spare engine stand, removal equipment etc. The travel time alone probably demands the use of a lease aircraft to support the stranded aircrafts schedule commitments.

Removal costs therefore range from a team of 2 mechanics for 8 h through to the costs of leasing an Antonov transport aircraft to ship the spare engine and equipment.

5.2.3.2 Replacement of the Engine (Substitute Asset)

Maintenance of the engine may take as little as a day or two or may extend to months for a major refurbishment overhaul.

During this time the aircraft will either remain out of service or will require a spare engine asset to be installed. Whilst many larger airlines carry their own stock of spare engines to cover this requirement many smaller ones rely on access to lease engine pools. In either case there is a substantial cost of purchasing, storing and maintaining the spare engine pools that is borne ultimately by the airline.

Positioning (global placement) and availability of these spare assets clearly has a major impact on the potential for aircraft to remain out of service longer than is absolutely necessary for the engine change. Failure to have a spare engine available often, inevitably leads to the need for a spare aircraft at a much higher lease cost or cancelled revenue generating operations.

5.2.3.3 Transport of the Engine

The fuel efficiency of a modern gas turbine aircraft engine is related to its bypass ratio.¹⁰ This has the tendency to push engines to ever greater diameters in order to generate a required level of thrust at the lowest fuel burn (and environmental) cost.

¹⁰In a bypass engine only a proportion of the total airflow pulled into the engine passes through the engine core. This is mixed with fuel and is ultimately expelled as a high temperature, high speed

For the highest thrust applications this represents a challenge in transportation. There are three basic modes of transportation for an engine. One is in place on the wing of the aircraft, second as freight on an aircraft and third as ground freight.

The first of these is rarely utilized as it generally precludes the aircrafts use during the flight for commercial purposes and it presupposes that a replacement engine already exists at the destination.

The second is subject to the availability of transport aircraft. Military transporters tend to have the highest payload capacity (size and weight) but are not generally available for commercial airline hire. Some specialized commercial cargo aircraft are available for large loads but these are expensive to hire and have limited availability. Therefore the desire is to be able to transport the engine in a widely available commercial cargo aircraft such as the Boeing 747. This has a prescribed limitation on the size of the freight due to the dimensions of the loading door and internal airframe capacity. Failure to design within this size window dramatically increases transportation costs. Some designs of engine require the bypass fan assembly to be removed from the rest of the engine in order to be able to meet the transportation size constraints (although there are also other considerations of modular overhaul frequency requirements that impact on this design decision).

The third option of ground based transport has its own issues. The first and foremost is clearly one of speed. There are however also practical constraints. For example road freight out of some major cities requires the engine to be transported through road tunnels (e.g. New York) or similar with dimension constraints. These result in time window limitations (you can only transport loads above these dimensions overnight) or size limitations (road freight in Tokyo restricted to 3 m square section). Both can impact design considerations.¹¹

5.2.3.4 Modular Disassembly

The engine has now made it to the overhaul facility and undergoes the first stage of disassembly.

The important thing to recognize is that most complex mechanical designs have a degree of modularity where the product can be progressively broken down into smaller subassemblies until it eventually reaches the level of individual components. This is true of modern gas turbine engine design.¹²

(Footnote 10 continued)

exhaust (the traditional Jet engine). The jet core flow in a bypass engine is kept to a minimum and is primarily used to generate the power necessary to turn a much larger set of fan blades that push a slower but much higher volume of *bypass* air from the rear of the engine. The higher the bypass ratio the more efficient the engine becomes in simplistic terms.

¹¹The Trent 1000 engine for the Boeing 787 Dreamliner aircraft has a special flat on its fan case flange that allows it to meet the 3 m transportation requirement for Japan.

¹²See the Rolls-Royce publication [4] or the web sites of most of the major Gas Turbine Engine manufacturers for examples of engine modularity.

In this first phase of disassembly the engine is stripped down to a level that allows the modules requiring maintenance to be accessed. Modules that do not require maintenance (at this point in time) may be stored awaiting the refurbished modules to be completed or may be re-allocated to another engines rebuild.

Importantly at this stage the modularity dictates how much of the engine will be subject to the next level of disassembly activity. Take the hypothetical example of an engine with just two modules. The first is the main jet gas path which sees the highest temperatures, pressures, speeds and loads within the engine. The other is the bypass system which lives in a relatively benign environment (speeds, temperatures, pressures). You would rightly expect that the gas path components would deteriorate at a higher rate than the bypass and would therefore require more frequent refurbishment maintenance. If all the high deterioration parts are contained in the jet gas path module then you would expect to overhaul it more frequently than the bypass.

e.g. a maintenance policy such as.

Overhaul jet gas path → overhaul both → overhaul jet gas path → overhaul both.

But what happens if you include just one high deterioration part in the bypass module as a result of poor selection of the module interface line? Either the part does not have the life expectancy to reach every second overhaul event as planned (disruption from failures, reduction in function or unplanned premature maintenance) or the module has to be planned for overhaul at every occasion (additional strip, inspection and build costs for all the components that did not required overhaul at this interval.

The decision on which modules to expose for module disassembly is one based upon both the timing and reason for the overhaul event as well as the basic modular architecture and maintenance policy for the product.

As already discussed when an engine is put into service there will be an expectation of a *reasonable* operational life before major overhaul maintenance is required. If an unexpected maintenance requirement arises well before this life a decision is required on the purpose of the maintenance that is done. Take the example of an engine that has achieved just 40% of the expected time to overhaul before one of its sub-systems develops the early symptoms of a fault. There are two choices:

1. Do sufficient maintenance to address the immediate need with the intention of allowing the engine to complete the other 60% of the intended life to its refurbishment overhaul (*check and repair*).
2. Pull forward the planned refurbishment and undertake sufficient work to return the engine to service with a full 100% life before the next refurbishment expectation (*refurbishment*).

This is typically an economic question based upon the cost of the *check and repair* option versus a full *refurbishment* overhaul. If the *check and repair* costs 10% of the full *refurbishment* then it might be reasonable to undertake this

approach at a life up to 90% of the planned overhaul interval but not beyond. In practice this type of logic is normally embodied in a set of module overhaul policies which often employ a concept of *soft lives*. The *soft life* is the life beyond which the module would normally be refurbished if a maintenance opportunity is presented.

At induction to the overhaul facility a typical modular overhaul decision making process tends to be:

- Is the module beyond the policy soft life? If yes then refurbish.
- If below the soft life conduct basic modular inspection to ensure its health is in line with expectation for its life. If inspections failed then refurbish.

5.2.3.5 Module Disassembly

Modular disassembly can be simple (strip to individual components), or complex (strip to sub assemblies and then apply logic as for module strip decision). In either case it may involve a large amount of labour and time to reduce the modules down to the level of piece parts or subassemblies that can be inspected and refurbished. The design of the product can have a profound effect on the costs and complexity of disassembly. For example interference fit spigot joints¹³ may be easy to make on assembly but provide a host of difficulties on disassembly that result in additional damage or even component scrap risk.

5.2.3.6 Inspection

There are three basic outcomes of the refurbishment decision making process at piece part level.

1. The part is acceptable to be refitted without any other work being done. i.e. the part has sufficient residual life to allow the next planned overhaul period to take place without significant risk of the part failing to perform its function.
2. The part does not have sufficient residual life to meet the criteria for 1) but can be economically repaired and would then meet the residual life criteria.
3. The part has insufficient residual life and is beyond economic or effective repair limits.

¹³Some joints need to be extremely tight to provide mechanical stability of the assembly. One such joint type involves two concentric rings on mating components in which the inner rings outer diameter is larger than the outer rings inner diameter. In order to assemble the two parts the inner ring is shrunk by super cooling, often with liquid nitrogen, in order to temporarily remove the interference. As the inner part warms, it grows back into contact with the outer part. Once the two are in close thermal contact the reverse operation becomes much more difficult. It may prove impossible to shrink the inner component leaving mechanical force as the only method of disassembly.

The purpose of inspection is therefore to make this determination in the most cost effective manner. This is often accomplished by working through a series of inspections that determine the acceptability of the functional features of the part. Careful consideration of the order of such inspections can optimize the costs of rejecting a part. E.g. if 90% of all rejections take place at inspection number 10 then can the inspection be done as step 1 saving the costs of the previous 9 steps?

The counter argument to the optimized inspection order is the one of providing information to enable product improvement. If 90% of parts are rejected at inspection step 1) is it cost effective to develop a repair for this feature? If only 2% of parts would be rejected for the other feature inspections then the answer is possibly yes. If however 80% would also be rejected for subsequent inspections then the single repair would yield few benefits.

Other issues that can occur are the result of the differences of inspecting for refurbishment versus new product assembly. Simple issues such as part marking and measurement datum arise. If either is placed on a surface or feature that is subject to wear then you may find that a perfectly functional part cannot be refitted because it can no longer be identified with 100% confidence or the reference point for a critical dimension has been compromised. Often all these issues require is some consideration of the build, strip, inspect; build, strip, inspect cycle at the point of design.

5.2.3.7 Repair or Replace

So we now have a part or assembly that has been rejected for refitting. The decision is do we repair it or replace it?

If we replace it we invest the full cost of a replacement part and the full environmental cost of disposal or recycling of the rejected part. However we will be obtaining a part that has the fullest possible residual life (as it is new).

If we repair it then in all likelihood we are only repairing certain features. Whilst these may be returned to the as new life, those features that we don't repair will not. Dependent upon the combination of features and their expected residual lives, a repaired part may have considerably shorter residual life expectancy than a new part. Additionally, repair often involves additional processing relative to the new production part. For example repair of coatings often requires the original coating to be removed before a new coating can be applied. This involves additional costs both financially and environmentally.

The repair or replace decision is therefore a complex one that is ultimately best made based on the detail of the individual component or assembly. However, at the point of design it is essential to have a view of the likelihood of repair being the viable economic and environmental option. If it is, then there are design decisions that can be taken that maximize the probability of a successful repair whilst reducing the cost and environmental impacts.

5.2.3.8 Rebuild Engine

Rebuilding the engine after refurbishment should be a repeat of the new product assembly process. However because the engine may not have been fully stripped to individual piece parts there may be some differences.

5.2.3.9 Test Engine

Dependent upon the nature of the maintenance work undertaken a *pass-off* test is typically required to demonstrate that the engine is fully functional and meets guaranteed performance levels prior to return to the customer. Such testing is generally conducted at a dedicated test cell at or near the overhaul facility.

5.2.3.10 Transport Engine

A repeat of the previous transportation considerations.

5.2.3.11 Re-install Engine

The reverse of the previous removal process.

5.3 Is a Gas Turbine Engine Unique in Its Service Support Complexity?

As can be seen from the previous sections the costs associated with operating a gas turbine engine for large aircraft is complex, with a multitude of degrees of freedom that impact those costs.

However none of this is particularly novel to the gas turbine example. The majority of considerations apply equally to a range of product sectors. For example take a washing machine. Other than the presence of rotating and static elements in the mechanics it apparently has little similarity to a gas turbine. However many of the considerations for service cost are the same:

- Do I react to the strange noise the drum is making or wait to see if it stops working? (*health monitoring*)
- When it fails will it just stop or will it ruin my clothes/flood the floor? (*indirect disruption costs*)

- Will I call out the repair man or is it old enough to make it cost effective to invest in a new washer? (*check and repair* or *replace*)
- If I go for repair will they replace the timer module or just the obviously burnt out capacitor? (*modular* versus *component* overhaul).

5.4 The Nature of Cost

In order to be able to deal with the complexity of the cost, it is helpful to be able to construct a simplified mental model.

$$\text{Total cost} = \text{intrinsic cost} + \text{inefficiency cost}$$

5.4.1 *Intrinsic Cost (Minimum Cost with Perfect Decision Making)*

- All mechanical systems exhibit wear out as a result of use
- With perfect visibility and understanding of this wear out there is a minimum fundamental cost of maintaining a functional product.
 - Maintenance undertaken at exactly the optimum time
 - Precisely the right quantity and location of maintenance resources to meet the demand
 - Just the right level of maintenance done to achieve optimum functionality/cost of support.

5.4.2 *Inefficiency Cost (Additional Costs from Imperfect Decision Making)*

- Visibility and understanding of wear out or other causes of damage is imperfect
- Decision making is constrained by factors such as understanding, training, time available to decide etc.
 - Maintenance undertaken too late—disruption or secondary damage incurred
 - Maintenance undertaken too early—useable operational life is lost for the same maintenance costs
 - Insufficient maintenance resources—maintenance delayed (as above) or queued requiring additional spare assets

- Excessive maintenance resources—additional inventory and employment costs
- Insufficient maintenance done when opportunities arise—short life to next maintenance event equals loss of useable operational life
- Excessive maintenance done when opportunity arises—additional costs for no additional operational life benefit

Intrinsic cost and inefficiency costs are both fundamental. Unless the product is never used it is impossible to reduce either to zero. The trick is to do what we can to minimize both elements. Consideration of both allows us to simplify the discussion of how to minimize service maintenance costs whilst considering both the product design and service delivery elements:

1. **MINIMISE INTRINSIC COST:** optimize the design to maximise the operational life and minimise maintenance requirements.
2. **PREDICTABILITY:** optimize the understanding of maintenance requirements to maximize the probability of undertaking the right maintenance at the right time with the right amount of available resources.
3. **MONITOR:** measure and record actual behaviour to refine understanding and improve the quality of decision making.
4. **REACT:** address opportunities to improve either the intrinsic cost or the inefficiency cost as appropriate.

Methodologies for addressing these are discussed in subsequent sections.

5.5 The Nature of Deterioration

Before discussing opportunities to improve intrinsic and efficiency costs it is worth a quick examination of the root cause of all of them all; deterioration.

The only reason for undertaking maintenance is that products deteriorate with use (or over time). If no deterioration took place the product would function for ever exactly as it did when first manufactured and the maintenance industry would all be out of a job. However there are no perpetual motion machines and neither is there a deterioration free product (even the most passive products can be subject to external sources of damage).

So let us examine the nature of deterioration. There are a number of facets to the process of deterioration all of which impact our methods of tackling the subsequent maintenance costs.

- Progressive versus instantaneous (e.g. gradual wear versus impact damage from a bird flying into the engine)
- Short versus long time-base (e.g. does significant deterioration and loss of function potentially occur within a fraction of a standard overhaul interval or over the period of a number of overhauls. In the latter case there will be one or

Fig. 5.2 Illustration of just one classification method for deterioration mechanisms

Progressive	Oxidation Corrosion Elemental diffusion	Wear LCF
	Melting	HCF Buckling Impact
	Chemical	Physical

more opportunities to observe the progression of deterioration and take remedial action)

- Monitorable versus hidden (e.g. slow measurable change in gas temperatures over months versus sudden failure of a part with no prior warning)
- Utilization sensitive versus utilization insensitive (e.g. thermal degradation of parts closely related to the number of take offs of the aircraft versus corrosion of external parts which may progress even if the engine remains unused)
- Fundamental versus probabilistic (e.g. any two mating parts will wear to some level [100% probability] but significant wear may require some other factor that is not always present [less than 100% probability]) (Fig. 5.2)

A further key assertion is that in the vast majority of cases deterioration impacts features of components or systems rather than the whole physical entity. Whilst we may talk about failure of component X or repair of component Y what we actually mean is that a feature of component X has deteriorated to a point where the component no longer serves its function or a feature of component Y has deteriorated to the point where it needs repair (but other features of the part may remain perfectly functional).

The reality of any physical system is that it is a collection of systems, subsystems, parts and part features that together combine to perform some overall function. Deterioration principally occurs at the feature level. Each feature may suffer from many or no significant deterioration mechanisms, each of which competes over time to undermine or reduce the functionality of the feature. When sufficient loss of feature functionality occurs (either individually or in combination) then the overall system function is impacted requiring either proactive (avoiding) or reactive (recovering) maintenance.

Managing maintenance costs (direct and indirect) is all about learning to predict, avoid, mitigate and recover from feature based deterioration.

5.6 The Design and Verification Process—Mitigating Deterioration

So how do we tackle the problem of optimizing maintenance cost when we are faced with such complexity of causes and potential methods of treatment?

The answer is to reduce the problem to its fundamental elements as far as possible.

1. Do we have a clear understanding of the requirement? Do we want the minimum possible direct maintenance costs (overhaul bills), or do we want the minimum overall cost of operation (fuel costs, installed maintenance and overhaul cost)?
2. Do we have a clear strategy for achieving this at the whole system level? Do we have clear product capability requirements (time between overhaul, operational reliability levels, cost of typical overhaul) and a corresponding support policy and infrastructure (fix when non-functional vs fix before loss of function, overhaul whilst in situ or return to overhaul facility)?
3. Do we understand the key success drivers of this strategy (physical deterioration mechanisms, optimum maintenance policies etc.)?
4. Can we mitigate the risks of loss of function at the optimal point in the process (avoid, detect, contain, recover)?
5. Can we predict expected behaviour of the product/service and monitor reality against it (early identification of issues)?

This amounts to the classical V of requirements cascade, design action and solution verification (Fig. 5.3).

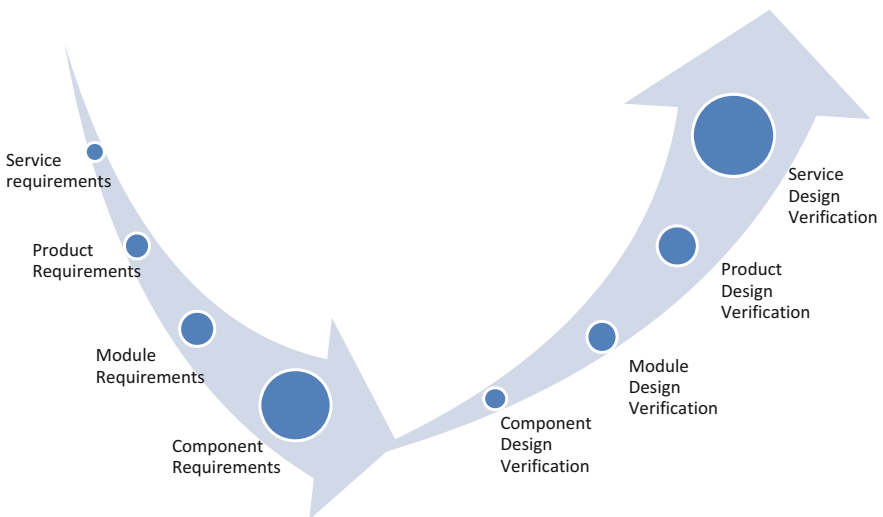


Fig. 5.3 The V model of requirements breakdown, design solution and verification

We will consider each of these stages in turn to examine the types of action that can be taken to optimize product design and maintenance service offerings.

5.6.1 Requirements Cascade

This is arguably the most critical of all the stages in terms of determining overall success. It is worth at this point considering the plight of the design engineer charged with the definition of one component within a complex system such as a gas turbine engine.

The design engineer will be faced with a series of requirements that they are expected to meet. These will come from a variety of sources; some may be legal requirements to meet safety regulations, some may be derived from customer requirements and others will be internal business requirements from their own company. In all cases the likelihood that the requirement will be met has three primary dimensions:

- Do I understand what I am being asked to deliver?
- How important is it relative to other requirements?
- How easy is it to achieve?

Maintenance cost traditionally suffers on all three counts:

- How do you express the maintenance cost requirement at the level of the component? The costs are typically billed at the whole system level, or if broken down typically refer only to the direct costs incurred in repair or replacement of the part. How do you define the requirement for a component that costs just \$5 to replace, but causes premature overhaul of the entire system at a cost of \$1 m? How do you remove the ambiguity of the frequency of overhaul of the part the designer is accountable for being determined predominantly by the lives of other parts of the system?

Example A component has a value of \$5 and an engine contains a set of 100 parts. All parts are replaced at routine overhaul at a cost of $\$5 * 100$ parts + \$50 labour cost. The typical cost of refurbishing the rest of the engine is \$1 m per overhaul. The frequency of overhaul is entirely determined by the part in question (it consistently has the shortest life of all the parts in the system). If the part averages a life of 10,000 h before overhaul is the costs per operational hour due to the part:

- (a) $\$550/10,000 \text{ h} = \$0.055/\text{hr}$ (direct cost expended on the part)
- (b) $(\$550 + \$1 \text{ m})/10,000 = \$100.055/\text{hr}$ (direct costs *caused* by the part)

What would the cost be if the engine had a second component that would have caused overhaul at 11,000 h?

What would the cost be if each of the 100 parts in the engine had a distribution of life around 10,000 h such that on average the *set* required overhaul more frequently than the 10,000 individual part average life?

What would the cost be if the mode of failure caused operational disruption with additional costs to the customers operation over and above the pure maintenance recovery cost?

- Safety considerations are absolute requirements for product certification. Weight is a hard customer requirement (as the wing and pylon are certified to a maximum engine weight) and unit cost is a key requirement for business viability. All have feedback loops within the timescale of the design process. Any failure to achieve requirements then becomes apparent and subject to escalation of priorities. Even when requirements start with equal priority, those that have an effective feedback loop of status versus requirement quickly get brought into much sharper focus and generally get higher levels of attention.

Example Component x needs to achieve (a) a weight of 1.5 lb, (b) be purchased for less than \$650 and (c) have a maintenance cost of less than \$0.09 per hour of service operation.

Weight is predicted with high confidence from the CADD drawing package in real time.

Supplier quotations are available within weeks

Maintenance costs may become apparent 5–10 years into service (but unlikely to be broken down to the level of individual components).

- The design changes necessary to impact maintenance cost are in reality no more or less difficult to accomplish than for the other requirements. The difficulty appears in the ability to understand the relationship between the options available to the designer and the impact they are likely to have on the maintenance cost. The major source of this complexity is the statistical nature of deterioration and the many uncertainties in the driver; e.g. *if* the operator uses it as we expect; *if* it is maintained as we expect etc.

Example Selecting a more expensive and denser material with superior properties *may* increase the life for a particular deterioration mode but is this the only deterioration mode that is significant in determining the parts ultimate life? The change in material will *definitely* impact the component cost and weight.

What the design engineer therefore needs is a set of clear and unambiguous requirements. The role of the requirement is to frame the problem in such a way that it reduces the issues above to the minimum whilst ensuring that all the designers and other stakeholders are aligned to delivering a total system solution that works. It is no use having a system of 20,000 parts where 19,999 meet the 10,000 h life requirement and the remaining one is only capable of 8,000 h. Neither is it acceptable to design a product for 10,000 h maintenance interval with a maintenance network sized for a 12,000 h interval.

5.6.2 Prediction and Design Action

The absolute key here is our ability to generate sufficient understanding of the impact of our design options on the maintenance cost to be able to guide effective decision making. If we are uncertain of the outcome of choices we will, as a direct result of human nature tend to prevaricate, postpone or otherwise avoid making the decision. Gambling ability is rarely part of the designer recruitment process!

The single biggest thing we can do to improve the consideration of maintenance costs in the product design is to make it feel real to the design population. It is the authors experience that the design community likes nothing better than the opportunity to solve a well stated and justified problem. Provided you give them the clear problem statement, sufficient understanding of the issues and the time, freedom and tools to explore the problem they will come up with truly great solutions.

So what are the magic ingredients?

- Education: make the problem *feel* real and important.
- Education: explain the complexity but bring it back to simple achievable goals that drive >80% of the costs.
- Make the requirement numerate or a clear pass/fail criterion.
- Predict the maintenance cost—waiting for reality is doomed to failure, having a mechanism to raise priority when failure to achieve requirement is predicted is key.
- Treat the prediction as reality; accept that it is less certain than other attribute predictions and measurements but don't use that as a rationale to ignore it when it looks hard to address.

Have you noticed the trend in the items above? With the exception of the maintenance prediction capability, the majority of issues are actually related to influencing human behaviour. If you want to be effective in addressing intangible hard to define, statistically difficult requirements with exceptionally long lead times to the emergence of proof of success or failure then it is far more a behavioural problem than a process or tools issue.

You can't design for maintenance cost successfully unless you address the human behaviour issues. Tools and processes alone are doomed to fail.

Some of the practical techniques available are therefore:

- Expose the design teams to the realities of service life. Ideally directly through attachment, secondment or visits but at the very least through direct contact with the front line staff.
- Recognise that the customer often has the most experience, so find a way to access it.
- Work through historical issues and understand the root causes and drivers of service impact. Take this learning through into the earliest stages of design consideration.¹⁴
- Conduct failure mode analysis and design to avoid all the associated consequences including maintenance cost.
- Enumerate and communicate the predicted status against the requirement and react to predicted shortfalls.

5.6.3 *Feedback Loops*

Service feedback is critical to success in managing maintenance costs. The more complex the product becomes the more fiendish the combination of deterioration mechanisms leading to service costs appears to be. It is practically impossible to predict all possible causes of service deterioration, rate of functional impact and the subsequent cost of management and recovery.

Feedback loops are therefore vital to obtaining the earliest possible identification of an emerging cause of cost and allowing the maximum reaction time to address it. Multiple feedback loops are possible within the life cycle of the product:

1. Feedback from prior experience (review of similar product histories).¹⁵

Note however that differences in environment, design or utilisation may impact the relevance of this experience.

2. Peer or customer review (identification of potential service cost drivers by those with more direct service experience and knowledge than the design team themselves).
3. Simulation and prediction (generation of a predicted cost using best available predictive methods).

¹⁴The majority of organisations will conduct a FMECA or Failure Mode Effect and Criticality Analysis as part of their design/safety process. In your organisation is this done before design starts as a guide to good design practice or after design as a check for safety compliance? One of these approaches has a high probability of reducing maintenance costs relative to the other.

¹⁵With the time from product design launch to statistically significant service maintenance data being measured in decades it is often the case that the only solid data comes from a design of engine 2 or 3 generations older than the current project.

4. Component or system test. This can be to demonstrate functionality at a specific set of conditions or often of more value is the test to failure that demonstrates the limits of capability and level of design margin.
5. Service health monitoring to detect the initiation or occurrence of loss of service function.
6. Service failure and maintenance cost statistics.

Each of these has a valid role in the provision of feedback to the design process. In general the later in the list an item appears the more definitive or real the feedback it provides. The counter point is that it also occurs later and therefore provides much less opportunity to minimise the costs. Items 1–3 generally allow you to impact the intrinsic cost of maintenance as they allow the opportunity to change the basic product design. Item 4 may offer the same opportunity but quite often test opportunities occur late in the design process when the ability to introduce change is constrained. Hence items 4–7 predominantly allow you to impact the inefficiency cost through provision of improved understanding of intrinsic deterioration costs and therefore the opportunity to manage them more effectively.

Also consider the customer impact. The final two items are the point at which it's the customer who feels the pain of unexpected service deterioration. Finding and resolving a potential issue from an earlier feedback loop can eliminate it before it ever gets seen by the customer¹⁶ avoiding customer dissatisfaction and the inevitable impact that has on future revenue opportunities.

5.7 Whose Costs Should We Design for, Ours or Our Customers?

I am often asked which of the service costs should we actually design for? After all it's the customer who will pay for fuel and indirect disruption costs but we will pay for the costs of maintaining to avoid them (assuming a fixed price maintenance contract has been agreed).

There are two ways to look at this, both of which lead to the same basic conclusion:

- Unless you intend to go out of business you must be hoping that your current customers will in the future continue to buy your products and services. In their shoes would you prefer to buy from a supplier who actively helps you reduce all

¹⁶Whilst customers can be very vocal about the positive benefits that health monitoring brings it is only ever minimising the pain that is inherently built into the product design. Most customers would much rather have you design that pain out before it ever gets to them. Rolls-Royces major customers all agree that whilst rapid and effective response to their emerging issues is an admirable trait in a supplier they would rather operate boring and predictable products that never present issues in the first place.

your costs or one who makes decisions in their own interest even if it's at your expense?

- Most businesses exist to make money. The opportunity is represented by the difference between the revenue available (which is a measure of the value it brings to the ultimate customer) and the costs of providing the product or service. In the civil aviation market that means the revenue depends on the ability to get passengers and cargo reliably and comfortably from point A to point B. The costs include acquisition of aircraft and engines, fuel, crew, disruption recovery etc. The bigger the gap between the revenue generating value and the true costs the more profit is available to share between all the parties involved in the value chain.

It may be slightly heretical but my personal view is that the engineering task is to create the largest possible profit pot; the commercial teams' task is to agree how it gets shared between the organisations involved.

So the answer I inevitably come to is that it's in the organisations long term interest to maximise the value the product and service brings whilst minimising the total cost of their provision. If this is apparently at odds with the organisations immediate best interests then set the commercial team the challenge of capturing a fair proportion of the profit margin that is being generated.

5.8 Conclusions

In this case study of a gas turbine engine what have we concluded about the process of designing for service maintenance costs?

Firstly the very nature of service adds many dimensions of uncertainty to the problem, such as how will the product be used? How will it be maintained? How will two seemingly identical products deteriorate when exposed to the same environment and operation? This complexity and uncertainty is at the root of why maintenance costs traditionally get less attention in the design process than other more tangible product attributes.

Secondly, the problem is essentially one of human behaviour. It is the difficulty we as human beings face in dealing with this complexity that inhibits our ability and willingness to tackle the problems effectively.

We have however discovered that there are ways to almost artificially simplify the situation and reduce the intangibility that inhibits effective design. Clear objective requirements, effective education (and simplification) of the service realities, creation of successive feedback loops all allow our design communities to engage more effectively with the issues. It is not an easy transition to make from a product design company that also sells support services to a service company that happens to design the underlying product, but it is rewarding.

The final conclusion is that by looking beyond our own revenues and costs to those of our customer we may discover opportunities to generate additional value generation opportunities.

Acknowledgements Some text from this chapter is reproduced with the permission of Rolls-Royce plc from a future planned issue of *The Jet Engine* book.

References

1. Rolls-Royce plc Investor news (2009) Paris Air Show orders reflect strength of Rolls-Royce portfolio In: Rolls-Royce corporate website, <http://www.rolls-royce.com>
2. Horder P, Senior Vice President, SH&E Ltd (2003) Managing aircraft maintenance costs conference (Brussels)
3. Harrison A (2006) GT2006-90570 Design for service—harmonizing product design with a services strategy. In: Proceedings of GT2006 ASME turbo expo 2006: power for land, Sea and Air, May 8-11, 2006, Barcelona, Spain
4. Rolls Royce plc, The Jet Engine book, The Jet Engine—a complete overview of the modern gas turbine. Key Publishing Ltd