UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING AND THE ENVIRONMENT

Computational Engineering Design Research Group

PhD Thesis

Ultrasonic Analysis: A Key Driver in Manufacturing Design Cost Optimisation for High Pressure Turbine Discs

Author:

James W C Meas

Supervisors:

Prof. James Scanlan

Dr. Steve Wiseall

Thesis for the degree of Doctor of Philosophy

December 2015

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

Faculty of Engineering and the Environment
Computational Engineering Design Research Group

Doctor of Philosophy

Ultrasonic Analysis: A Key Driver in Manufacturing Design Optimisation for High Pressure Turbine Discs

by James Meas

This thesis investigated the best strategy to calculate the unit cost of isothermally forged discs used within large civil gas turbine engines. Research has shown that ultrasonic examination constraints required during the production of forged parts strongly influence the manufacturing design. These constraints have traditionally been ignored during preliminary design analysis due to the unavailability of suitable models. Cost reduction objectives are also commonly simplified using mass as a direct surrogate. These oversights can lead to significantly inaccurate predictions of both cost and material properties.

This thesis assessed the impact of satisfying the ultrasonic constraint and quantified potential cost savings through direct unit cost optimisation. This work created a fast automated method for evaluating ultrasonic examination limitations, and created an optimisation workflow where a detailed analytical cost model can be rapidly evaluated. The inclusion of ultrasonic constraint analysis within an automated design loop increased both the precision of manufacturing designs and unit cost estimates. This improvement has led to detailed manufacturing designs being assessed within the preliminary design environment enabling true concurrent engineering. These new advancements enable a disc designer to obtain manufacturing designs alongside a detailed cost breakdown of the optimised manufacturing route. This approach is now routinely used within Rolls-Royce who have also patented the modelling technique.

Results show that including the ultrasonic constraint increased initial manufacturing geometry mass by at least 3.8% with costs increasing by up to 2.4%. Subsequent unit cost optimisation showed cost increases of between -1.6% and +4.3% with maximum heat treatment depth increasing up to 4.6% compared to the baseline non-ultrasonic constraint scenario. These results prove that ultrasonic constraints are a necessity to predict accurate manufacturing data. Introduction of a detailed analytical cost model also found that direct cost optimisation improved savings by up to 1.9% compared to traditional mass reduction objectives. This confirms the need for cost estimation to achieve effective cost reduction.

This page intentionally left blank

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

AFRC	Advanced Forming Research Centre
APU	Auxiliary Power Unit
ASM	American Society for Metals
BF	Black Forging
BSI	Borescope Inspection
CAD	Computer Aided Design
CNC	Computer Numerical Control
COE	Collaborative Optimisation Environment
COS	Condition of Supply
CSSO	Concurrent Subspace Optimisation
DOE	Design of Experiments
DUNDE	Dummy Ultrasonic Non-Destructive Examination
EPSRC	Engineering and Physical Sciences Research Council
FAA	Federal Aviation Administration
FCC	Face-Centre-Cubic
FPI	Fluorescent Penetrant Inspection
GA	Genetic Algorithm
GEAE	General Electric Aero Engines
HP	High Pressure
HPC	High Pressure Compressor
HPT	High Pressure Turbine
HT	Heat Treatment
HTC	Heat Transfer Coefficients
ID	Inner Diameter
IPT	Intermediate Pressure Turbine
LHS	Left Hand Side
LPC	Low Pressure Compressor
LPT	Low Pressure Turbine
MDO	Multi-Disciplinary Optimisation
ME	Manufacturing Engineer
MPI	Magnetic Particle Inspection
NDE	Non-Destructive Examination
NDT	Non-Destructive Testing
OCR	Overall Compression Ratio
OD	Outer Diameter
PL	Parting Line
PSEF	Peened Surface Extrusion Folds
RHS	Right Hand Side
RI	Rectilinear Inspection
SAMULET	Strategic Affordable Manufacturing in the UK with Leading Environmental Technology
SF	Shape Factor

TSB	Technology Strategy Board
VF	Volume Factor

This page intentionally left blank

DECLARATION OF AUTHORSHIP

I, James Meas, declare that the thesis entitled 'Ultrasonic Analysis: A Key Driver in Manufacturing Design Optimisation for High Pressure Turbine Discs' and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University;
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- Where I have consulted the published work of others, this is always clearly attributed;
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- None of this work has been published before submission.

Signed: ……..

Date:……

This page intentionally left blank

INTELLECTUAL PROPERTY NOTICE

Some of the data contained within this thesis is proprietary to Rolls-Royce Plc; therefore, this thesis may not be distributed without permission unless the commercially sensitive – proprietary data appendices are removed.

LIST OF PUBLICATIONS

Due to an ongoing patent application at Rolls-Royce related to topics in this thesis, no publications have been submitted to date.

FUNDING SOURCES

This work is jointly funded by Rolls-Royce Plc, the Engineering and Physical Sciences Research Council (EPSRC) and the Technology Strategy Board (TSB) as part of the Strategic Affordable Manufacturing in the UK with Leading Environmental Technology (SAMULET) programme.

This page intentionally left blank.

ACKNOWLEDGEMENTS

Firstly, I would like to express my sincere gratitude to my academic advisor Prof. Jim Scanlan for the continuous support throughout my PhD study and for giving me the freedom in the direction of my research. Thank you also for your understanding and patience whilst I worked then travelled during the writing of this thesis. I would also like to thank my industrial supervisor Steve Wiseall who sadly passed away last year for his help, advice and for linking me to many of the helpful staff at Rolls Royce. Thank you also to all the staff of Rolls-Royce who frequently gave up their time to support me in my research, with a special mention to Stan Nikov who always had time to help me.

I would like to thank everyone who made my time at Southampton an amazing experience this includes the CED research group as well as many of the societies I joined. I'd like to include a special thank you to the engineering hockey team who provided endless laughs, stories as well as lifetime friends. I'd also like to thank Michael Gedge for providing me with endless stimulating discussions over numerous coffee breaks during our PhD studies.

Particular thanks go to my Mum and Dad for always being there to support and encourage me in all my ventures throughout my life and for never standing in my way.

Last but not least, thank you to my fiancée Alison who has also stood by me during the stressful time of writing my PhD thesis. Thank you for always being there.

This page intentionally left blank.

Chapter 01 INTRODUCTION

1.1. TURBINE DISC BACKGROUND

Discs are considered to be one of the most critical components in an aircraft engine in terms of safety and performance. A typical High Pressure Turbine (HPT) disc is designed to operate at rotational speeds of 10,500 rpm[1] with rim temperatures approaching melting point (≈1000 K) (Reed 2006). This hostile environment causes several design issues, which include creep and fatigue due to cyclic overstressing. 65% of the tensile stresses are caused by the rotational mass of the disc itself, whilst the remainder of the tensile stresses result from the turbine blades (Cristinacce 2010) which can experience as much as 6,600g creating radial forces of 18 tonnes per blade (Rolls-Royce plc 2005). The creep results from prolonged radial forces at elevated temperatures. The fatigue cracks arise from alternating radial forces (due to throttle alterations) which make discs a life limited component (Bamberg and Spies 2007).

Engine manufacturers are constantly trying to improve engine technology to produce higher performance engines; whilst reducing specific fuel consumption, CO_2 emissions and noise levels to stay ahead of competitors. To achieve these requirements:

- Overall Compression Ratios (OCR) need to rise to improve engine efficiency.
- Materials and/or cooling need to be improved to withstand higher turbine entry temperatures and compression ratios.
- Increased shaft speeds are required to achieve the OCR requirements.
- Bypass ratios needs to be increased to reduce noise and improve the propulsive efficiency.
- Shaft diameters need to increase due to shaft power required to rotate the larger fans.

  (Wagner 2004)

These evolving requirements greatly affect the discs within the engine; greater rotational speeds will increase centrifugal stresses within the disc reducing fatigue life. Greater temperatures will increase creep effects and larger low pressure shafts will increase turbine disc bore diameters, further increasing stresses on the disc (Wagner 2004). It is these three requirements that are driving the development of materials and pushing them to their

[1] Rolls-Royce Trent 800

limits. Modern superalloys are generally more expensive the greater the maximum working temperature, for this reason the material can make up around 50% of the total forging cost of a disc (S.K. Srivatsa 2003). The cost breakdown to manufacture a disc is shown in Figure 1, here it can be seen that the remaining 50% of costs are value added.

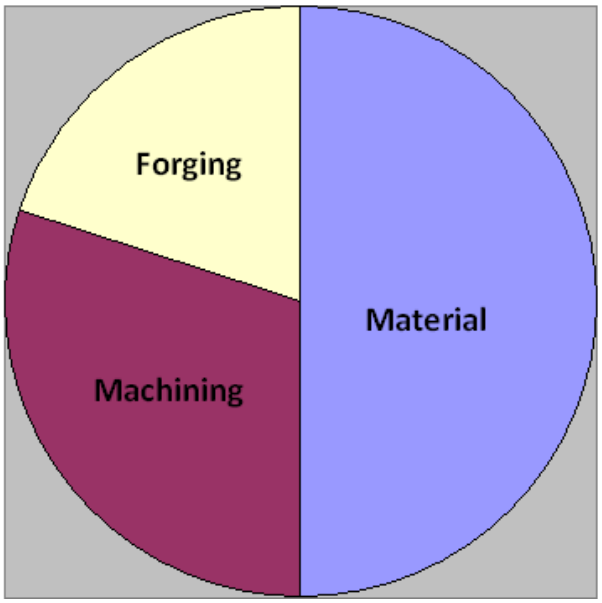

Figure 1 – Generic cost breakdown for gas turbine engine discs (Wagner 2004).

Typical Buy-to-fly mass ratios are in the order of 10-25% (S. Wiseall 2009), resulting in substantial material waste. Considering that discs make up approximately 10% (Reed 2006) of the total unit cost of an engine, as much as 4% of the total engine cost is directly spent on raw disc material which ends up as waste material.

1.2. Disc Design Background

A gas turbine engine operates using the Brayton Cycle. In order to operate the Brayton cycle a compressor, a method of heat addition and an expander is required. In a gas turbine engine, the amount of energy required to compress the air at the start of the cycle is provided directly by a turbine. This is possible as the amount of energy produced from the combustion process is greater than that used by the compressor.

In large civil engines energy is extracted from the expander section using axial turbines which convert axial flow into circumferential rotation. This rotational energy is transferred to the compressor where an axial compressor is used to increase the pressure across the compressor stage.

Discs play a critical role in the energy transfer between the turbine and compressor by transferring torque created by the turbine blades through to the shaft then back into compressor discs which then convert the rotational energy into compression using the compressor blades.

The discs have many roles, but the primary ones are to:

- Maintain the radial position of the blades.
- Maintain the axial position of the blades.
- Transfer circumferential loads between the blades at the rim of the disc to the main shaft.

To achieve these key roles, the discs have the following key features to enable them to perform their role:

- Rim Blade Slot – To transfer radial forces from the blade into the body of the disc.
- Axial blade clamping features – To axially position the blades to the rim of the disc.
- Diaphragm – To transfer radial forces at the rim down towards the bore of the disc.
- Cob – To structurally maintain the radial position of the diaphragm of the disc.
- Drivearm – To transfer circumferential loads between the disc and shaft and maintain concentricity.
- Rear drivearm or stubshaft – To transfer loads between discs and to maintain concentricity.

Typical features of a high pressure turbine disc are shown in Figure 2 below.

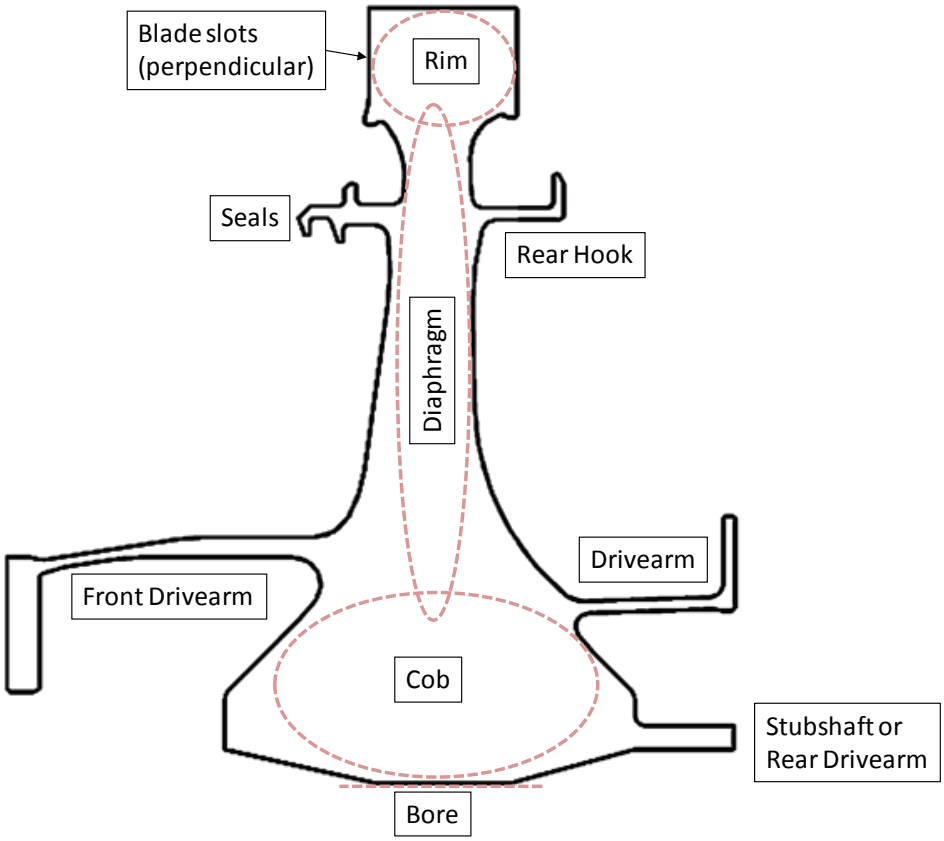

Figure 2 – A generic high pressure turbine disc with its key features labelled.

1.3. Manufacturing Design Optimisation

This thesis is concerned with reducing the unit cost of a disc whilst maintaining disc integrity for safety. It has been established that there is a large amount of waste material during the manufacture of the HPT disc which increases the cost in two ways. The first is the initial extra cash required to purchase the superalloy material and the second is the additional value added cost required to form and to machine away excess material. In order to accurately reduce costs a detailed cost model will be created. To make the cost model accurate manufacturing models will be created to ensure upstream manufacturing designs are estimated correctly to ensure the best possible cost estimate. This will involve the creation of the following geometry:

- Rectilinear inspection geometry
- Forging geometry
- Heat treatment geometry
- Billet geometry

To ensure the manufacturing geometry created is feasible all the geometry must be created to a set of rules to ensure they can all be manufactured. The rectilinear inspection geometry must also be designed to meet an additional safety requirement involving ultrasonic non-destructive examination.

The rectilinear inspection geometry is the first geometry to be created around the disc profile and what all subsequent manufacturing geometry is based on; it is therefore important to get this geometry both manufacturable and fully ultrasonically inspected in order to avoid design and cost prediction errors.

Cost optimisation on the manufacturing design is useful for part cost reduction once the disc design has been completed, but this method can also become an essential tool during design concept analysis. If the method is sufficiently quick, concurrent engineering can happen. The concurrent availability of accurate cost estimates is vital to help aid disc design decisions.

1.4. SUMMARY

This introduction has covered the background of this thesis and the projected investigation that will follow. The thesis will involve elements of cost, design and analysis to successfully reduce the manufacturing unit cost of a disc. To understand how these main points can be achieved, the literature review will cover in depth background knowledge and reveal previous research focused on investigations related to this study. The literature review will also clearly state the objectives of this thesis.

This page intentionally left blank.

Chapter 02 LITERATURE REVIEW

2.1. INTRODUCTION

This literature review concentrates on current technologies available for disc design and manufacture to better understand current cost factors. Material cost has already been highlighted as a key cost driver; therefore, the details behind nickel-based superalloys are covered first in Section 2.2. Disc manufacture in Section 2.3 covers disc production starting from the initial material production and forging through to heat treatment and machining. Section 2.4 investigates the requirement for safety; including ultrasonic non-destructive examination procedures performed on the part during manufacture. Section 2.5 investigates the importance of modelling cost during the design process. Section 2.6 briefly describes current manufacturing design practices and Section 2.7 reveals the intention of this thesis before performing further in depth research into recently published work. The key findings from the literature review are described in Section 2.8 and the chapter ends with the thesis aims and objectives in Section 2.9.

2.2. MATERIALS

2.2.1. MATERIAL BACKGROUND

Most gas turbine discs experience large tensile stresses at high temperatures. Compressor discs operate in a relatively cool environment compared to turbine discs; therefore they can be made out of titanium alloys, but once the disc experiences temperatures greater than 540 degrees Celsius, titanium is no longer strong enough (Donachie and Donachie 2002). This is because once the temperature has reached 50% of the material's melting point; creep becomes a problem. It is in these conditions that superalloys are required.

2.2.2. SUPERALLOYS

Superalloys are materials that exhibit excellent mechanical properties, creep resistance and resist chemical degradation at temperatures close to their melting point (Donachie and Donachie 2002; Jeniski Jr and Kennedy 2006; Reed 2006). There are three possible contenders that could be used for turbine disc applications; platinum, cobalt and nickel.

The desired atomic lattice layout is Face-Centre-Cubic (FCC) which makes the material ductile (important for forging) and tough (Reed 2006). Many of the platinum group metals have an FCC layout, but also have high densities and costs associated with them (Reed 2006). Cobalt superalloys possess satisfactory density and material characteristics but are expensive. Nickel-based superalloys are an efficient compromise between performance and

economics (Furrer and Fecht 1999), this is because of its high melting point, low cost and relative low density compared to platinum (Reed 2006). Nickel also possesses the favourable FCC layout so maintains its phase from room temperature right through to its melting point eliminating issues related to expansions and contractions (Reed 2006). It is therefore nickel-based superalloys that are most commonly used for turbine and hot end compressor discs.

2.2.3. NICKEL-BASED SUPERALLOYS:

It was the gas turbine engine that drove the development of nickel-based superalloys which began in the 1940's (Donachie and Donachie 2002). This group of materials possess high temperature strength, toughness, and resistance to degradation in corrosive and oxidizing environments (Pollock and Tin 2006) making it a good application for turbine technology. To obtain the specific material properties of nickel-based superalloys, as many as 14 different alloy elements can be added (see Table 1) (Donachie and Donachie 2002). For example, titanium and aluminium are added to the nickel to produce a gamma-prime phase which further strengthens the material at high temperatures (Schafrik and Sprague 2008).

Table 1 – Chemical composition of alloys discussed in this report (Reed 2006)

Alloys	Alloying Elements (%)													
	Cr	Co	Mo	W	Nb	Al	Ti	Ta	Fe	Hf	C	B	Zr	Ni
IN 718	19.0	---	3.0	---	5.1	0.5	0.9	---	18.5	---	0.04	---	---	Bal
Waspaloy	19.5	13.5	4.3	---	---	1.3	3.0	---	---	---	0.08	0.006	---	Bal
Rene 95	14.0	8.0	3.5	3.5	3.5	3.5	2.5	---	---	---	0.15	0.010	0.05	Bal
RR1000	15.0	18.5	5.0	---	1.1	3.0	3.6	2.0	---	0.5	0.027	0.015	0.06	Bal
Udimet 720Li	16.0	15.0	3.0	1.25	---	2.5	5.0	---	---	---	0.025	0.018	0.05	Bal

Depending on the specific composition of the alloying elements, different materials are created as shown in Table 1. Inconnel 718 (or Alloy 718) is arguably the most successful of these superalloys, as it makes up 65% of metals used in General Electric engines in the year 2000 (Schafrik and Sprague 2008).

Other common superalloys are Waspaloy, Udimet 720 and RR1000. Waspaloy is relatively low strength, but possesses a stable phase providing good processing, low cost and provides high resistance to fatigue crack propagation (Manning, Knowles et al. 2003). Udimet 720 has a higher strength than Waspaloy, but has lower crack propagation resistance (Manning, Knowles et al. 2003). Rolls-Royce developed a patented material, RR1000, which has a

similar strength to Udimet 720 but maintains higher fatigue crack propagation resistance and creep rupture life (Manning, Knowles et al. 2003). Although RR1000 is the most expensive out of the three, it has the highest ultimate tensile strength as shown in Table 2.

Table 2 - Comparison of ultimate tensile strength (Manning, Knowles et al. 2003)

Alloy	Typical Ultimate Tensile Strength MPa Standard Commercial Heat Treatment
RR1000	>1500
Udimet 720Li	>1450
Waspaloy	>1100

It has been shown that the nickel-based superalloys used in gas turbine engines are very complex and expensive with no feasible alternatives available. Heat treatments can alter the material properties which allow materials to be tuned to their specific application. This is discussed in greater detail in Section 2.3.4.

Material Summary

To enable the use of cheaper metals would require the reduction in either rim temperature of the disc, or the centrifugal loading. In either case, these options are beyond the scope of this project and are unlikely to occur due to engine performance requirements. Research and development of materials possessing greater strength and creep resistance will continue to meet the demands of tomorrow's engines. It is highly likely that the costs of these materials are going to increase further highlighting the importance of reducing the buy-to-fly ratios to save unnecessary costs.

2.3. DISC MANUFACTURE

The manufacturing of a disc is a very complex process with numerous constraints. The disc starts off as the required raw elements (as shown previously in Table 1) which are then combined through several treatment and manufacturing processes to produce the final disc geometry. The main steps involved during the manufacture of a disc are as follows:

1. Alloy production – Ingot production.
2. Cogging (Upset and draw) / Extrude – to make the final billet diameter and to homogenise and create the desired grain size.
3. Ultrasonic Non-Destructive Examination (ultrasonic NDE) – to detect any internal defects within the billet.
4. Sawing – to divide the billet into required billets sizes.
5. Forging – to make the near-rectilinear geometry.
6. Heat treatment – to form the correct material properties.
7. Turning – to create the rectilinear geometry.
8. Ultrasonic NDE – to detect any internal defects.
9. Turning – to produce the final axi-symmetric disc profile.
10. Milling & Broaching – to create non axi-symmetric features.
11. Cleaning and Surface Finishing and checks – to remove contaminants and improve surface strength and heat resistance.

Each one of these steps is important in creating a disc which will function as required whilst ensuring high safety standards for the required number of flight cycles. Step 1 was researched in the previous section (Section 2.2); steps 1, 2, 5, 6 7 and 9 will be covered next in the remainder of this section. Steps 3 and 8 relating to ultrasonic testing for safety are covered later in Section 2.4. Step 4 involves sawing the billet into the required lengths so therefore will not be covered in any detail nor will Steps 10 and 11 which involve 3D feature creation and surface finishing and checks which are beyond the scope of this thesis due to time constraints.

2.3.1. INGOT AND BILLET CREATION

The creation of the billet is a complicated and a highly controlled process where one of two methods can be performed depending on the material:

Ingot Metallurgy	**Powder Metallurgy**
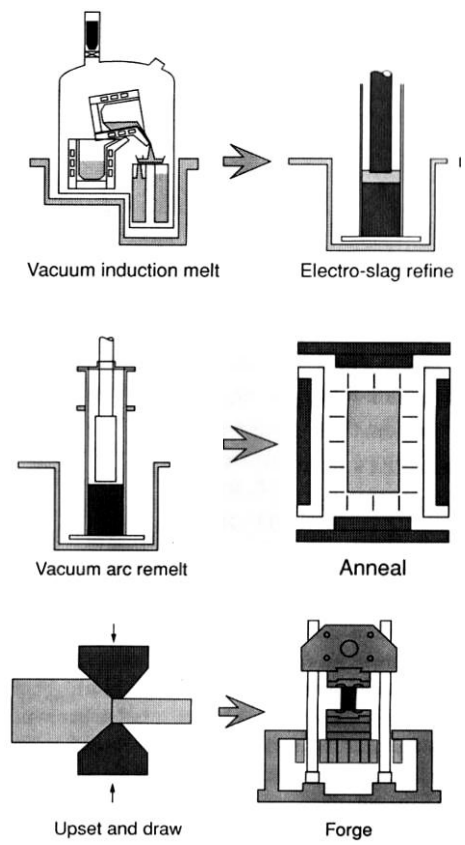 (Reed 2006)	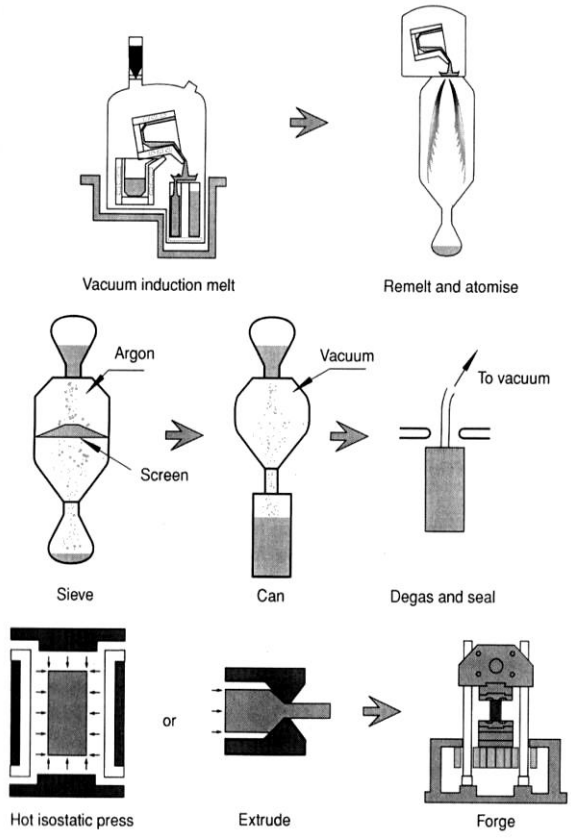 (Reed 2006)
1. Vacuum induction melt – prepare material composition 2. Electro-Slag Refine – remove oxygen containing inclusions 3. Vacuum Arc Re-Melt – further improve the quality of the metal 4. Anneal – homogenise the material 5. Upset and draw –create a finer grain 6. Forge (Benz 1999)	1. Vacuum induction melt – prepare material composition 2. Re-Melt and Atomise – turn the alloy into powder 3. Sieve – remove oxygen-containing inclusions 4. Canning – Pour into billet mould. 5. Degas and Seal 6. Hot Isostatic Press/Extrude – Reconsolidate the powder 7. Forge (Benz 1999)

Figure 3 - Illustrations and process descriptions of ingot and powder metallurgy

The actual manufacturing sequence used during the production of superalloys depends upon the alloy, for instance Waspaloy and IN718 are normally prepared using ingot metallurgy due to low volumes of strengthening elements; whereas superalloys such as Rene 95 and RR1000 which are heavily alloyed need to be prepared using powder metallurgy (Reed 2006). In theory inclusions within powder metallurgy are lower than those produced by ingot metallurgy, but this process comes with greater complexity and cost (Reed 2006).

It is important that no impurities exist within the alloys as they can severely hinder the following material properties:

- Tensile strength
- Stress rupture
- Creep resistance
- Fatigue life
- Oxidation resistance
- Corrosion resistance
- Hot workability
- Weldability performance

  (Benz 1999)

It is for this reason materials are formed in either vacuums or in an inert gas.

2.3.2. COGGING

Ingots produced through the remelting process lack the acceptable grain size (Reed 2006). To set the grain structure up for isothermal forging, the large ingots require cogging.

Cogging is a form of open die forging used when a component is too large for closed die forging. It is used to reduce the grain size down from the as cast structure (Reed 2006). The process involves rotating the ingot in discrete steps between forging strokes to reduce the diameter and laterally moving the ingot to achieve a consistent diameter along the full length of the workpiece (Valberg 2010) as illustrated in Figure 4.

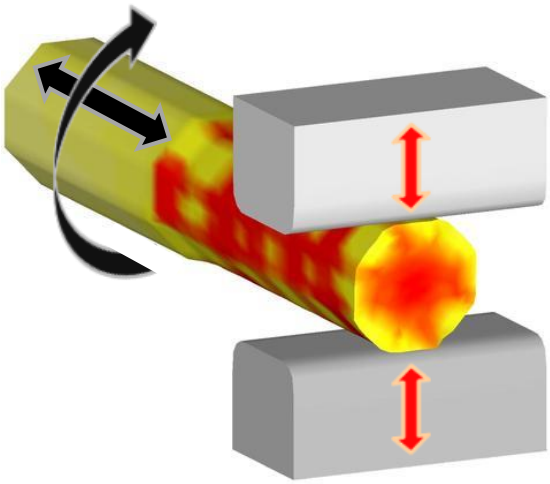

Figure 4 – Illustration of Cogging. Red arrows represent the periodic die movement; the black arrows show the ingot manipulation. (Adapted from (Wilde Analysis Ltd 2011))

2.3.3. FORGING

Forging is essential in the manufacturing of a disc as it has several advantages over other techniques (S.K. Srivatsa 2003). Wrought materials created through forging possess greater ductility which allow larger shapes such as discs to be created (Donachie and Donachie 2002). Casting and machining can produce net/near-net shapes but neither possess both low material waste and favourable material properties (Donachie and Donachie 2002; Semiatin 2005). The forging process can provide better tensile properties, stress rupture behaviour, creep strength and low cycle fatigue life (Donachie and Donachie 2002). These properties can be controlled by aligning the grain in a favourable direction, adjusting the second phase morphology and adjusting grain size (Donachie and Donachie 2002).

The forging method used for turbine discs is isothermal closed-die forging which was a technology specifically developed for the aerospace industry (Semiatin 2005). Isothermal forging allows difficult materials such as nickel-based superalloys to be forged which are difficult or near impossible to forge by conventional means (Semiatin 2005). This forging method also boasts the ability of forging closer to net-shape, this has the advantage of lower material input, hence reduced machining (Semiatin 2005) as illustrated in Figure 5. Closed die forging is used as it produces near-net shape profiles with relatively high tolerances compared to open die forging reducing the need for excess material. To improve die fill, superalloys are isothermally forged to reduce forging loads (Donachie and Donachie 2002), it is also utilised to eliminate die chilling by eradicating thermal gradients (Shen and Furrer 2000; Semiatin 2005). The forging operation is performed using a hydraulic press to provide tight control over the press speed which ensures acceptable strain rates are not

exceeded. This is useful as it minimises deformation heating, flow stresses and allows superplastic behaviour in certain materials (Semiatin 2005). Due to the high temperatures involved, dies are generally made out of molybdenum due to their high melting point, high strength and good machinability. However at isothermal conditions, the die requires a vacuum to avoid corrosion (Semiatin 2005), this adds overhead costs to the process.

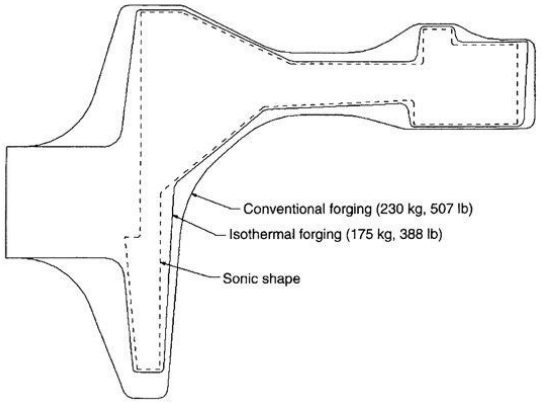

Figure 5 - Example showing input material for a jet engine turbine disc using both conventional and isothermal forging. Isothermal forging provides a 55kg saving in this case (Semiatin 2005). The 'sonic shape' or rectilinear geometry is the desired output of the forging therefore the forging profile must be greater than this.

Forging makes up a large proportion of the total cost of a disc, therefore it is important that the forging die is designed right first time.

Problem Areas in the Analysis of Metal Forming

Industrial experience has shown that forging issues can be split up into several "problem areas" (Valberg 2010). Analysing each area of a process individually may allow bottle neck situations to be solved during optimisation exercises. The different problem areas are as follows:

- Plastic Zone – This is the area in which the workpiece is being plastically deformed. In the case of roll extrusion, the plastic deformation occurs between the two rollers. In the case of closed-die forging, the whole workpiece is the plastic zone.
- Workpiece at Initial Stage – This is the state the workpiece is in before undergoing plastic deformation despite any heat treatment or annealing that has taken place beforehand.
- Workpiece at final Stage – This is the workpiece after the plastic deformation process has been completed. This is the wrought product ready for machining or just the finished product

- Interface between Workpiece and Die – This contact interface may be direct contact between the die and workpiece, but is normally separated by at least an oxide film layer. This interface can be improved through the use of lubricants which help reduce surface roughness and prevent localised pressure-welds of material onto the die.
- Dies – Forge dies are generally very expensive, therefore great care needs to be focused on designing a die which can achieve a satisfactory die life with limited wear to maintain component tolerances and batch cost.
- Surface Layer of the Workpiece – The surface texture of the workpiece is important for many reasons. For instance, it may affect how effective the lubricant will be during the forging process or in a more extreme case an abrasive texture may directly cause excessive wear to the dies. Any roughness may be exaggerated through the forming process.
- Forming Machine – Metal forming machines are normally complex and expensive to purchase. Forge machines should be designed specific to their task to ensure production of parts goes smoothly and they operate to the required tolerances.
  (Valberg 2010)

This list highlights areas that need to be considered when designing a new forged part. For instance, the maximum forging load to make the part must not exceed that which the machine can produce, and to reduce the number of dies required die ware must be minimised.

Important Process Parameters to Maintain Final Part Quality

Some of the important outcomes that need to be avoided throughout the forging process are as follows:

- Poor grain size control
- Grain size banded areas
- Poor carbide or second phase morphology/distribution
- Internal cracking
- Surface cracking
  (Donachie and Donachie 2002)

To avoid cracks and other such defects it is important that several parameters are controlled within the plastic zone of the workpiece, these include:

- Strain

- Strain rate
- Flow stress
- Temperature
- Contact pressures
- Shear stress
  (Valberg 2010)

Many of these parameters control the material properties of the workpiece which are also influenced by the lubricant used. Material properties which cannot be achieved through forging may be solved through heat treatments stages(Valberg 2010).

Forging Models

Forging experts are highly experienced therefore they can use their knowledge to avoid such problems, but due to the high costs associated with the isothermal forging of nickel-based superalloys, accurate predictions are required to aid their engineering judgement. Predictions can be obtained through either modelling or experimentation. Experiments provide accurate results, but add cost (wages, equipment and plant space) and lead time, therefore computer modelling is the preferred method to explore the design space (Srivatsa 2003).

Modern Computing power is able to provide predictions of material microstructure within a component (Donachie and Donachie 2002). The method used is Finite Element (FE) analysis which has been used to model forging since the early 70's (Chenot, Massoni et al. 1996). Modern simulations are becoming increasingly more complete and realistic as material properties and coupling effects are being included (Chenot, Massoni et al. 1996). The modelling takes out the risk associated with educated guessing, this eliminates trial and error techniques which ultimately saves costs and design time. "Understanding the process methods, equipment and materials are vital to the successful implementation of computer process modelling" (Shen and Furrer 2000). Therefore, to be able to model this complex operation everything needs to be considered. Several process simulations have been developed, the more common process modelling software include DEFORM, ABAQUS, Forge3 and MSC.Marc (Semiatin 2005).

The Isothermal Forging Process

Forging involves compressing a billet between two dies 'squashing' material out radially filling out into the upper and lower dies. The maximum load experienced by the press usually occurs at the end of the forging stroke (He, Rohl et al. 1998) where the greatest surface contact exists between the part and die. If this maximum force is greater than the press rating, then complete forging will not be possible. If this is the case a die redesign is required.

With the exception of the forge press, the most expensive and critical tools used in the process are the dies (Kulon, Broomhead et al. 2006). Therefore, part of creating a successful forging design is to ensure the die life is greater than the full production run of components whilst remaining within a satisfactory tolerance. This is currently achieved using inefficient 'trial and error' techniques (Kulon, Broomhead et al. 2006). In most cases, the forging process route relies upon a forging experts knowledge and intuition (Kulon, Mynors et al. 2006). Capture of this knowledge is extremely difficult due to the large number of variables involved, as well as lack of cooperation by experts (Srivatsa 2003). Figure 6 shows how a knowledge based approach can be used to produce the approximate black forging geometry from a given part; where the given part is usually the Rectilinear Inspection (RI) geometry.

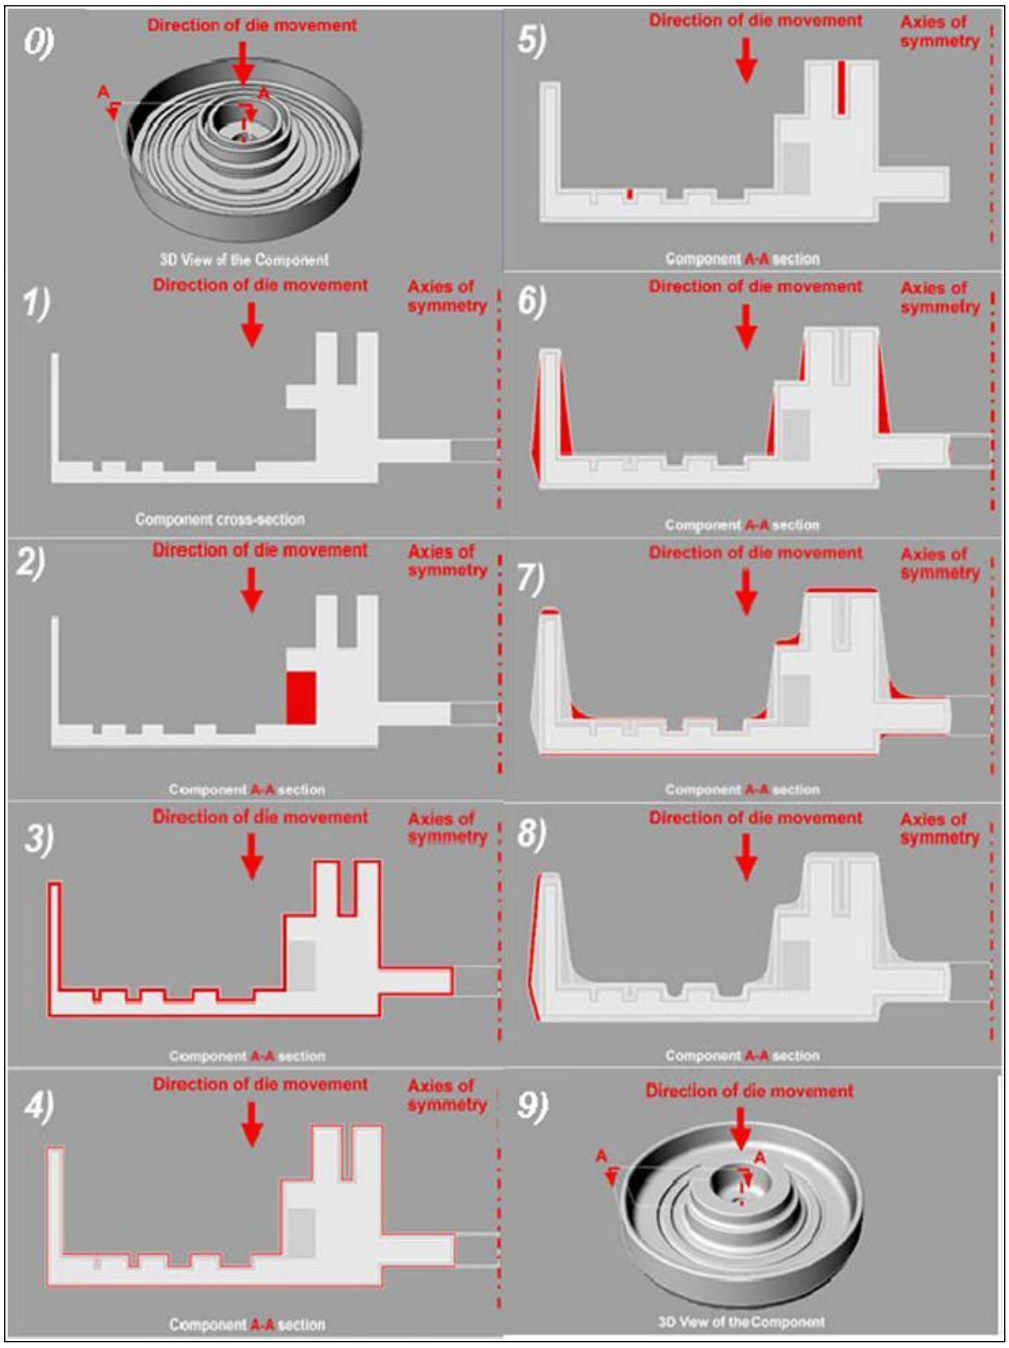

Figure 6 - General steps required to create a forgeable geometry (Kulon, Mynors et al. 2006)

The black forging design stage is the transformation of the required component geometry to a forgeable shape (Kulon, Mynors et al. 2006). The steps in this forging design stage as shown in Figure 6 are as follows:

1. Finished component cross section
2. Fill in undercuts
3. Add in machining tolerances
4. Add in forging tolerances

5. Fill in narrow features
6. Include draft angles
7. Include fillets and corners of suitable radii
8. Assess rib thickness

After following these steps a theoretical forgeable shape is produced, with a high degree of confidence that can then be further analysed or verified by FE analysis (Srivatsa 2003; Kulon, Mynors et al. 2006). This rule based method is a very good starting point for creating the forging geometry. The rules are looked at in greater detail in Section 3.4.2.1. It should be noted that the process of converting the given geometry into forgeable geometry adds additional material to the part. The initial shape of the part will determine how much extra material is required.

New Technology – Advanced Roll Forming

Advanced roll forming is a new method under development which will be capable of creating complex axi-symmetric disc shapes. This technology could provide reduced costs over conventional forging because of its net shape capability and lower tooling costs (Bewlay, Gigliotti et al. 2003). The technology works by rotating the workpiece about its symmetry axis whilst two or more rollers work radially to shape the material as shown in Figure 7. This has been used to make several different parts including automotive wheels and engine discs (Semiatin 2005); it has been demonstrated using different metal alloys including several titanium alloys and nickel-based superalloys (Semiatin 2005). It was found that discs which required an ultrasonic inspection envelope could be manufactured with a 10-30% reduction in forging weight (Bewlay, Gigliotti et al. 2003). The reduction in total input material will evidently reduce part unit cost. Further benefits arise from the elimination of the dies (used in closed die forging) and the versatility of the rollers which can be used for several geometries (Semiatin 2005). This is a promising new technology, but is not currently fully developed.

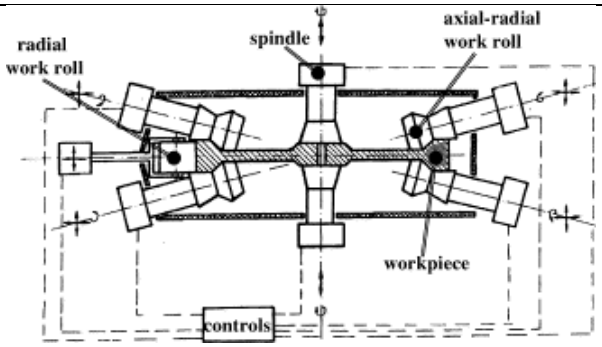

(Bewlay, Gigliotti et al. 2003)

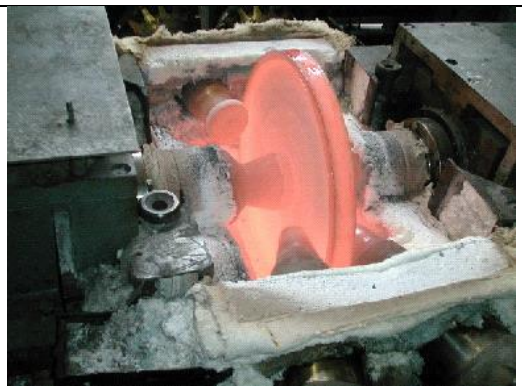

(Institute for Metals Superplasticity Problems of RAS 2011)

Figure 7 - SRD800 roll-forming/reeling mill

Forging Summary

The detailed forging design is ultimately out of the control of Rolls-Royce and hence this project. Research and development in these areas are being under taken by companies and research departments, for example the Advanced Forming Research Centre (AFRC) at the University of Strathclyde whose core research is aimed at "developing forming and forging into scientifically understood and repeatable processes" (University of Strathclyde, 2015), and the Engineering Research Centre for Net Shape Manufacturing at The Ohio State University have research related to forging die design, prediction and elimination of forging defects and wear (The Ohio State University, 2015). Despite the lack of involvement in the forging process, the forging geometry could still be estimated using methods described by *Kulon et al. (Kulon, Mynors et al. 2006)* which would be essential to calculate the cost of a part. Knowing the approximate geometry of the forged part could be vital in allowing the RI geometry to be designed more efficiently to the upstream manufacture which has the potential to reduce costs.

2.3.4. HEAT TREATMENT

Heat treatment takes place after the forging process has been completed; the black forging is normally altered to obtain a suitable heat treatment shape prior to the heat treatment process. Heat treatment geometry is required to create the necessary mechanical properties of a component (Furrer, Shankar et al. 2003). The goal of modern heat treatment is to achieve the best balance of mechanical properties and residual stresses (Shen and Furrer 2000). To achieve this, nickel-based superalloys will normally involve one or more of the following cycles:

- Solution heat treating

- Stabilising
- Aging

  (Furrer, Shankar et al. 2003)

Each type manipulates specific metallurgical structures within the component which is necessary to obtain the correct properties. For example, heat treatment can be used to improve at least one of the following:

- Reduce residual stresses
- Allow the atoms to diffuse within the material
- Promote grain growth
- Promote grain re-crystallisation
- Dissolve phases
- Produce new phases from precipitation
- Cause new phases to form from the introduction of foreign atoms

  (Donachie and Donachie 2002)

While the heating of the component is relatively straight forward, the cooling of the component is key to obtaining the desired properties (DeAntonio, Duhl et al. 1991; Furrer, Shankar et al. 2003). In order to increase the mechanical properties a fast cooling rate is required (Röhl, Srivatsa et al. 1997; Shen and Furrer 2000), but this can lead to high residual stresses, quench cracking, subsequent machining complications and operational life reduction (Furrer, Shankar et al. 2003). Therefore, a lot of attention is focused on the exact cooling path seen across the entire component. Predominantly superalloys are strengthened by the production of a secondary phase (Donachie and Donachie 2002). By adjusting the heat treatment, large yield strength improvements can be achieved although at the expense of creep rupture strength (Donachie and Donachie 2002). Traditionally oil quenching has been used to rapidly cool the part, but there is very little control over the process due to the limited number of variables (Röhl, Srivatsa et al. 1997; Rohl 1997). Fan cooling is increasingly taking over the quenching process as it allows greater control over the cooling coefficients around the surface of the disc. An advance on fan cooling is ATI Ladish's[2] patented super cooler which offers much higher heat transfer coefficients (Bunge and Furrer 2003). With several jets positioned around a single disc, a greater number of

[2] ATI Ladish Co., Inc - Forging supplier of high pressure turbine discs.

controlled variables become available enabling an opportunity to improve the resultant residual stresses through optimisation.

Furrer et al. (Furrer, Shankar et al. 2003) explains how oil or polymer quenching heat treatment would often dictate the size and shape of the forging geometry. The forging geometry was selectively increased in size to ensure sufficient heat treatment geometry was available to optimise the cooling rates within the part. Utilising new controlled quenching techniques, the addition of material is no longer required to optimise the heat treatment process for material properties. Therefore, the heat treatment geometry can be configured near-net shape; not affecting the forging geometry design.

Advances to the grain structure have been worked into the process by improving the grain size at the two extreme ends of the disc. Recent work on dual grain structures has been successfully achieved where fine grains are formed at the bore to create high strength and coarse grains are grown at the rim to provide high creep resistance (Shen and Furrer 2000; Gayda, Gabb et al. 2006). Utilising this property could allow the mass of the final disc to be reduced or for the same mass; a greater life or capability can be achieved.

2.3.5. MACHINING – TURNING

Discs due to their axi-symmetric profile are machined using a technique called 'Turning'. Turning is where the disc is revolved about its axis and a rigid single-point turning tool is forced into the part to remove layers of material.

Despite the advances in near-net forging the need for turning is inevitable during the manufacture of a disc. In practice there are three stages of turning required; these are to get from:

1. The black forging geometry to the desired heat treatment shape
2. The heat treatment shape to the rectilinear inspection geometry
3. The RI geometry to the final disc 2D-axisysmetric profile.

The machining process for superalloys is considerably more costly than for other metals such as steel due to cutting speeds being around 90-95% slower (Donachie and Donachie 2002). Their poor machinability is due to their material properties. This is because superalloys:

- Retain strength at high temperatures
- Have unusually high dynamic shear strength
- Contain hard carbides in their microstructure making them abrasive

- Work harden during metal cutting
- Have poor thermal diffusivity which leads to high cutting tool tip temperatures
- Form a tough continuous chip during metal cutting
  (Donachie and Donachie 2002)

Nickel alloys work harden rapidly due to the high pressures required during machining (Davis 1989), and because of the harsh cutting conditions carbide coated tools are generally used due to their high resistance to wear. The most influential factor in determining tool life is the surface speed, although feed rate and depth of cut are also important (Donachie and Donachie 2002). Because of the importance of tolerances and machine utilisation, the discs are first rough machined, then 'semi' machined, before being finish machined. Rough machining removes the vast majority of the material at a fast rate. The time to machine is unlikely to improve through tool path optimisation, as the time taken to machine is very closely related to the material removal rate of the tool. Tool path routes require greater consideration when large residual stresses exist within the part to ensure final part geometry remains within tolerance. Specific research in this area has been performed in the past by *Urresti et al.* (Urresti, Nikov et al. 2009) and recent advances in the modelling of metal machining processes has been extensively covered by *Arrazola et al* (Arrazola, Özel et al. 2013).

2.3.5.1. Factors Affecting Turning

There are also several other variables affecting the turning process; feed rates are important for productivity but have to be weighed against tool wear. A sufficient feed rate is required to prevent burnishing or glazing, which will work harden the surface making cuts progressively more difficult increasing the tool wear rate (Davis 1989). It is important that the rake angle (tool angle relative to the component surface) of the turning tool is positive to ensure the cut metal is forced away from the finished surface of the disc (Davis 1989).

The systematic process a CNC (Computer Numerical Control) manufacturing expert should follow is briefly looked at below:

- Drawing Evaluation – 'What is the best way to machine the part?'
- Identify Part Material and Size – To locate specific material properties and necessary equipment
- Part Holding – A fixed, safe stable method of holding the work piece in place
- Tooling Selection:
  1. Cutting tool

2. Cutting tool holder

- Cutting Conditions – tooling selection, setup method, depth of cut, width of cut, spindle speed and cutting federates.
- Program Writing – Manual Programming, Macro Programming, Computer Programming and Conversational Programming.
- Program Verification – Simulation Tools, Manual checks
- Program Transfer – Flash Drive, Remote computer
- Work Completion:
    1. Evaluate the part program
    2. Check supplied material
    3. Prepare required tools
    4. Setup and register tools
    5. Setup part in a fixture
    6. Load Program
    7. Set various offsets
    8. Run first part
    9. Optimise program if necessary
    10. Run production
    11. Inspect frequently

    (Smid 2010)

There are three steps which directly influence the start and finish geometry of the part. These are: Drawing Evaluation, Part holding and Tooling Selection, the influence of which will be explained shortly. In terms of cost, the 'cutting conditions' will impact the most, this is because tool wear and cutting rates will determine the majority of the machining operating time.

Listed below are rules which may need to be considered during the design of rectilinear inspection geometry or during the manufacturing procedure:

1. Clamping Surfaces – determined by 'Part Holding'
2. Narrow and deep grooves – determined by 'Tool Selection'
3. Minimum radius – determined by 'Tool Selection'
4. Additional material thickness to allow finishing procedure – determined by the tolerances found in the 'Drawing Evaluation'
5. Undercuts – determined by 'Tool Selection'

Clamping the part is critical during turning to prevent unwanted movement during the turning process. Depending on the turning process, the part may be clamped in only one or two places during the turning operation or several clamping locations may be required throughout. The design of the rectilinear inspection geometry must at least consider the locations that the first operation will use to ensure additional turning operations are not required.

The depth and width of grooves that can be machined depends upon the tooling used. Details for standard tooling and insert shapes can be obtained in a similar format to Figure 8.

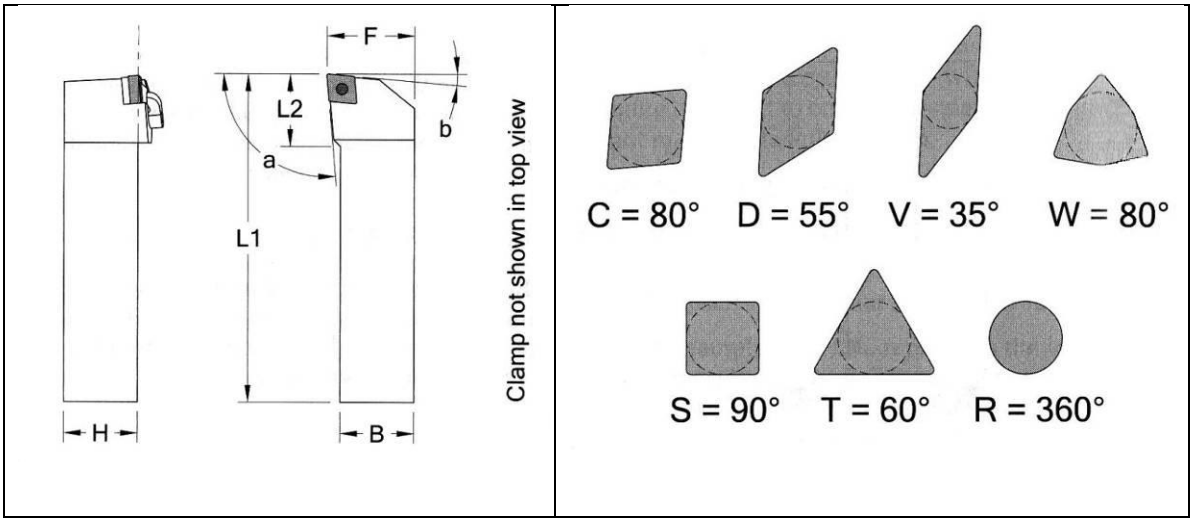

Figure 8 - Turning tool profile view displaying the variables on the most common turning tool in CNC lathes, and examples of the most common insert shapes. (Smid 2010)

For grooves where standard tooling is to large, special grooving tools such as the ones shown in Figure 9 can be used. As well as these groove tools being able to create the narrow deep features required, there will inevitably be speed and feed rate implications ultimately increasing machine time and costs.

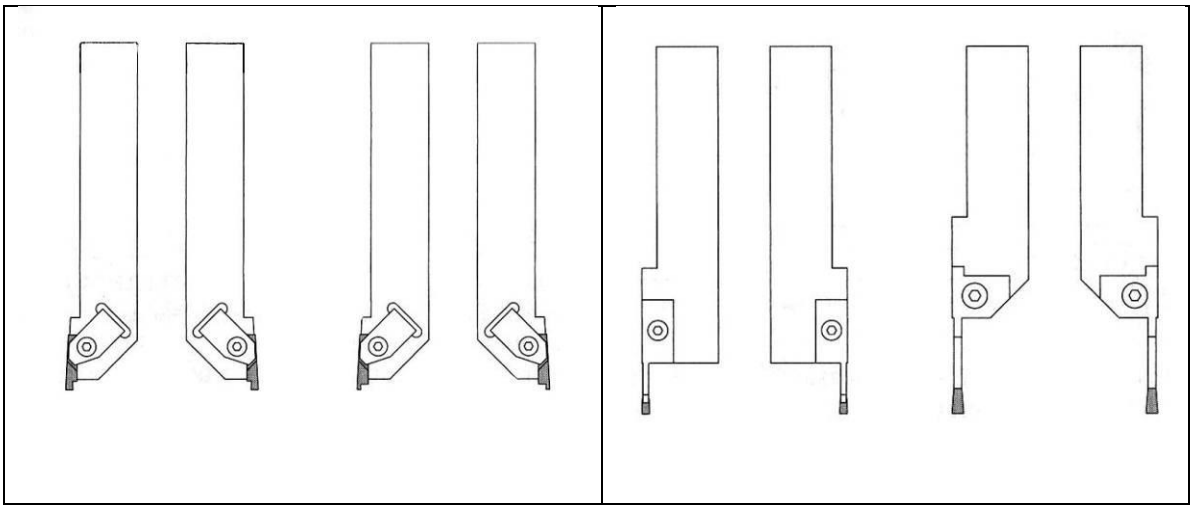

Figure 9 - Examples of grooving tools used to cut deep and narrow channels. (Smid 2010)

Along with speed and feed rates, the tool insert radius is also directly related to the manufacturing cost. A large radius can produce high volumetric removal rates due its larger surface contact; it will also require fewer insert replacements due to low tool wear. On the other hand, a smaller radius will have a lower volumetric removal rate and will require a greater number of insert replacements due to the higher wear rate. Although tool replacements due to wear will add cost, the majority arises from machine overhead costs, therefore to minimise costs a high volumetric removal rate is desired.

The finishing procedure when the material is close to the required contour is to perform a G70[3] Finishing Procedure or G73[3] Pattern Repeating Cycle. In the G73 pattern repeating cycle the last contour is repeated a number of times until the final contour has been achieved as illustrated in Figure 10 (Smid 2010).

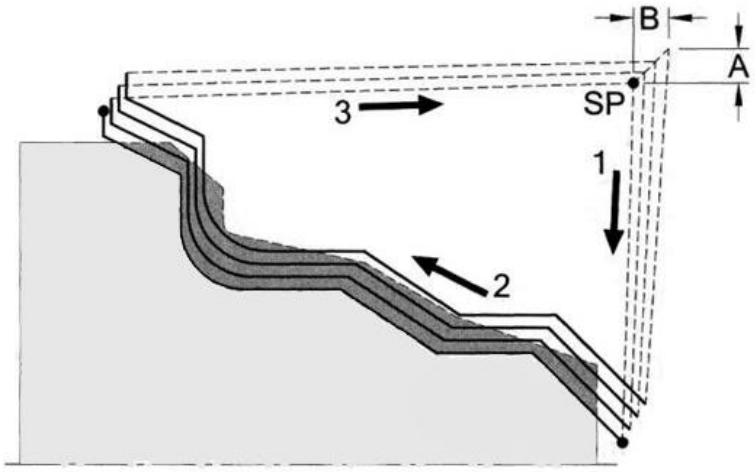

[3] Programming cycle identifier used in CNC machines.

Figure 10 - Pattern repeating cycle to machine the finish contour. (Smid 2010)

Manufacture Summary

The manufacturing process involves very complex processes requiring great expertise in order to obtain the geometry and properties required. It has been shown that the forging shape can be derived through a sequence of rules which should theoretically produce a design of least waste, but unless further FE analysis is performed, the design may not be permissible. Heat treatment is a very specialised topic where material experts are required to obtain optimal properties which will be out of the scope of this project. Machining appears to be an area where time and cost could be shaved off, but the actual machining process is highly dependent upon tooling. One area that could be improved is machine utilisation where cheaper machines can be used for rough cutting leaving accurate CNC machines to perform the final cuts to the high tolerances required. This investigation would Require Discrete Event Simulation which is not currently considered and is beyond the scope of this research.

2.4. SAFETY

2.4.1. HISTORY OF DISC FAILURE IN SERVICE

Discs are classed as a safety-critical component (Hersman 2010) as well as a life-limited component. Therefore to prevent uncontained failures of a disc they are subjected to special engineering, manufacturing and maintenance restrictions (Hersman 2010). Despite all the Federal Aviation Administration's (FAA) stringent safety regulations, catastrophic disc failures have occurred in the past with sometimes fatal consequences. For example the Sioux City disaster in 1989 where 111 out of the 296 on board lost their lives (National Transportation Safety Board 1990). After conducting further research in to catastrophic disc failures, 12 different cases of disc failure were found involving these disc stages:

- Low Pressure Compressor (LPC) / Fan – 1 type, 1 case. (National Transportation Safety Board 1990)
- High Pressure Compressor (HPC) – 1 type, 1 case. (FSF Editorial Staff 1996; National Transportation Safety Board 1996)
- High Pressure Turbine (HPT) – 2 types, 6 cases. (Kolstad 1991; Ausralian Transport Safety Bureau 2004; Rosenker 2006)
- Intermediate Pressure Turbine (IPT) – 1 type, 1 case. (Australian Transport Safety Bureau 2010)
- Low Pressure Turbine (LPT) – 1 type, 4 cases. (Hersman 2010)

It appears that all disc stages throughout the engine are vulnerable. This list captures most of the incidents over the last 25 years and shows that disc failures are infrequent incidents.

Understanding how these discs failed could provide important knowledge required to prevent such events occurring in new designs. Next is a discussion on which mechanisms of failure caused such incidents.

The Sioux City disaster mentioned above has one of the highest numbers of deaths associated with a disc failure. The McDonnell Douglas DC-10-10 was over an hour into the flight when a fan disc ruptured in one of its GEAE[4] CF6-6 engines located in the tail (number 2 engine) of the tri-engine aircraft. The disc split into two large very high energy projectile segments which passed through the containment ring[5] and subsequently severed all three of the triple redundant hydraulics, locking all control surfaces for the remainder of the flight. There were no direct casualties associated with the disc rupture, but due to all direct control[6] of the aircraft having been made inoperable, an unsuccessful emergency landing at Sioux City was the direct cause of casualties.

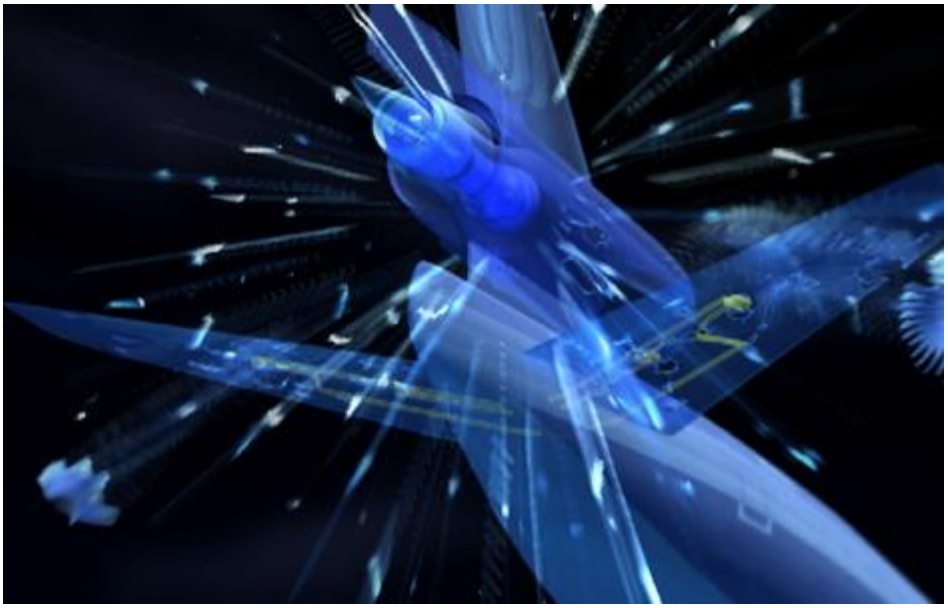

Figure 11 – Illustration showing how the fragments from the catastrophic failure of the fan disc penetrated the tails hydraulic lines. (National Geographic 2011)

[4] General Electric Aircraft Engines

[5] A stainless steel hoop surrounding the engines stage one fan disc blades. This hoop is designed to absorb enough energy of a single fan blade

[6] Limited control of the aircraft was maintained through use of engine biasing and balancing the centre of gravity of the aircraft through strategic pumping of fuel around the aircraft

The failure of the fan disc was due to the growth of a fatigue crack that grew over the majority of the disc's life (41,009 hours and 15,503 cycles). The disc had a predicted service life of at least 54,000 cycles, with an FAA approved service life of 18,000 cycles (using a 1/3 safety factor multiplier). The crack initiated due to the presence of a type 1 hard alpha metallurgical defect which formed during the initial ingot production. Hard alpha defects are caused by excessive amounts of nitrogen which is unintentionally introduced whilst in a molten state. Two vacuum melting procedures were performed on the disc material providing an opportunity to dissolve any hard alpha inclusions. Several types of inspection were performed on the disc at several stages during the manufacture, but no defects were detected. A triple vacuum melt process and improvements in furnace cleaning at the time of manufacture would have prevented the defect early on. The 'probable cause' was likely due to human factors during quality-control inspection procedures. It was United Airlines engine overhaul facility which failed to detect the 0.476 inch surface crack using a Fluorescent Penetrant Inspection (FPI) procedure. (National Transportation Safety Board 1990)

On June 8th 1995 during takeoff on a Mc Donnell Douglas DC-9-32 the 7th HPC disc in the right hand engine catastrophically failed, resulting in an aborted takeoff and a cabin fire. Despite the majority of passengers and crew being unharmed, there was one serious injury to a flight attendant who suffered 2nd degree burns from the onboard fire as well as two puncture wounds resulting from flying shrapnel. The disc failed due to a fatigue crack propagating to beyond the critical crack length. This crack is calculated to have been 100% detectable using the two inspection techniques[7] at its previous overhaul. (National Transportation Safety Board 1996)

In the last two decades, all[8] commercial HPT disc failures in service have occurred on five GEAE CF6-80 engines on Boeing 767s. Details of the stage 1 HPT disc failures are as follows:

1. 1991 – failure due to cracking within rim bolt holes (Ausralian Transport Safety Bureau 2004)
2. 1991 (during climb) – Failure occurred due to fracture post chrome plate repair on the disc shaft section (Ausralian Transport Safety Bureau 2004)

[7] Florescent Penetrant Inspection (FPI) and Magnetic Particle Inspection (MPI)

[8] with exception of a small HPT disc from an APU Kolstad, J. L. (1991). Safety Recommendation A-91-085. National Transportation Safety Board. H. J. S. Busey, Administrator, F. A. Administration, Washington and D. C. 20591. Washington, National Transportation Safety Board**:** 1-4.

3. 2000 (during maintenance)– failure resulting from fatigue cracking within a disc fir tree slot (Ausralian Transport Safety Bureau 2004; Rosenker 2006)
4. 2002 (during climb) – Failure resulting from fatigue cracking within two disc fir tree slots (Ausralian Transport Safety Bureau 2004)
5. 2006 (during maintenance) – Failure resulting from fatigue cracking within two disc fir tree slots (Rosenker 2006) (Figure 12)

Figure 12 – part of the port HPT disc embedded in the starboard nacelle (Blakey 2006)

With the exception of the first two disc rupture incidents, the most common failure mechanism has arisen from fatigue cracking arising from the disc's fir tree slots. The reason for repetitive occurrence of this problem was first put down to non-conformance to the engineering drawing requirements. The defective region of the slot was the chamfer's radii variance resulting from a manual finishing process (Rosenker 2006). Further finishing issues were observed in the 2002 disc where shot peening in the fir tree slots caused Peened Surface Extrusion Folds (PSEF) creating local fatigue crack initiation points. Although the PSEF has not been ruled as the primary mode to failure, it is identified as being a significant contributory factor of the premature fatigue cracking (Ausralian Transport Safety Bureau 2004). Further identifiers to potential fatigue crack initiation points were identified during handling of the disc during manufacture and maintenance due to the "sensitivity" of the slot bottoms (Rosenker 2006). It was concluded that the CF6-80 1st stage

HPT disc material may be operating close to its limit; therefore a design review was to be conducted to prevent the occurrence of slot cracks (Rosenker 2006).

1991 saw an uncontained failure of an Auxiliary Power Unit (APU)[9] located in the tail of an Airbus A300. The root cause of the failure was due to a 0.316 inch long piece of stainless steel weld rod which was dropped into the melt during manufacture of the disc ingot. It was this 'freak' inclusion that initiated the fatigue crack.

The most recent uncontained disc event (to date [11/01/2011]) was that of the RR Trent-900 on the new Airbus A380. The incident occurred on the #2 engine which failed shortly after takeoff, causing considerable damage and fuel leaks. The Preliminary report which has been conducted by the Australian Transport Safety Bureau suggests that the stub pipe feeding oil to the HP/IP bearing structure had a manufacturing defect causing a fatigue crack to form. Subsequent failure due to the fatigue crack resulted in an oil leak. There has currently been no evidence of pre-existing defects within the located disc section and the fractures found on the disc are consistent with a ductile overstress failure mechanism (Australian Transport Safety Bureau 2010).

The most frequent occurrence of in-flight disc failure is that of the GEAE CF6-45/-50[10] where four instances of partial low pressure turbine ejections have occurred. These incidents occurred in a short time period:

1. July 2008
2. March 2009
3. December 2009
4. April 2010
   (Hersman 2010)

The cause of the LPT stage 3 disc failure was directly linked to a High Pressure (HP) rotor unbalance, where vibrations are transferred through a common bearing support exciting the stage 3 blisk (bladed-disc) at its natural frequency. This vibration coupling resulted in the forward spacer arm exceeding designed bending loads causing circumferential high-amplitude fatigue cracks to form. Once cracks link together around the complete circumference of the disc drive arm, the aft-section of the low pressure rotor separates

[9] Garrett Auxiliary Power Division – TSCP700-5

[10] General Electric Aircraft Engine CF6-45/-50 is found on Boeing 747s, McDonnell Douglas DC-10s and Airbus A300s

penetrating the engine casing. The imbalance of the HP rotor occurred due to turbine blade damage (loss of material). Four instances of this very same problem had previously occurred in the late 70's where a different model (GEAE CF6-6) failed in the same way. The severity of this problem resulted in the FAA/GEAE limiting engine operations by requiring Borescope Inspections (BSI) once every 15 flight cycles as shown in Figure 13 (Hersman 2010).

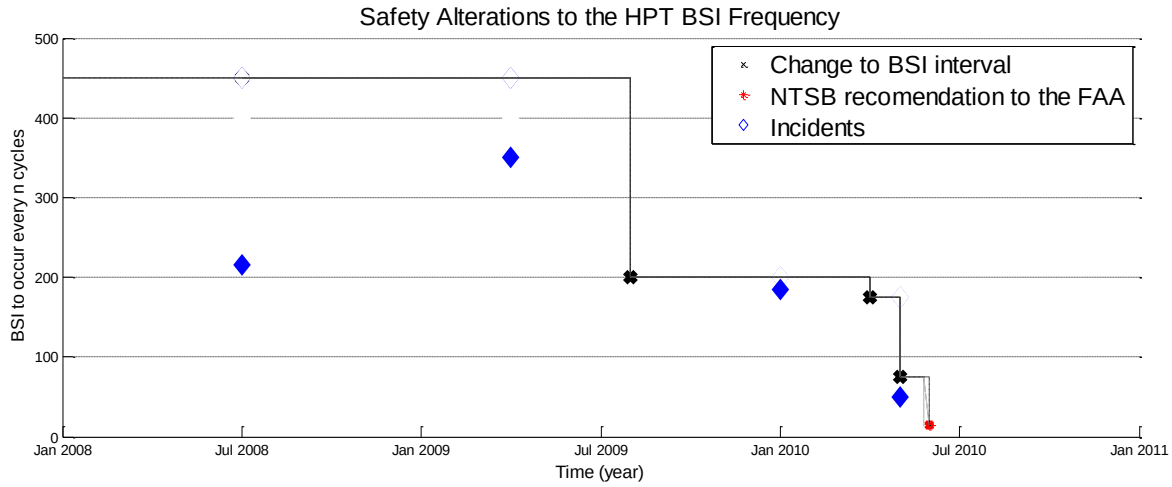

Figure 13 - Safety alterations to the HPT disc Bore scope inspection frequency (created from data collected in ref (Hersman 2010))

Due to the increased frequency at which BSI needs to be performed, maintenance costs and turnaround time no doubt increases. Failures affect the air travel, airline and engine manufacturers' reputation.

Table 3 shows the predominant factors which caused in-service discs to fail. Development in manufacturing and NDE techniques has reduced the possibility of defects whilst increasing the likelihood of detecting faults. Disc finishing, maintenance and handling should be improved through better procedures but could also be improved through greater design robustness.

Table 3 - Causes of engine disc failure over the past 25 years (data collected from: (National Transportation Safety Board 1990; Kolstad 1991; National Transportation Safety Board 1996; Ausralian Transport Safety Bureau 2004; Rosenker 2006; Australian Transport Safety Bureau 2010; Hersman 2010)

Engine Type	Predominant factor causing disc failure					
	Metallurgy	Finishing	Maintenance	Handling	Design	Other
GEAE – CF6-6	1					
P&W – JT8D			1			
GEAE CF6-80		1	1	1	1	1
GAPD	1					
RR T800						1
GEAE CF6-						4

Although Table 3 shows that causes of disc failure are thinly spread across many disciplines; it should be noted that a single reduction in any of these causes could save hundreds of lives by preventing a catastrophic incident similar to that seen in Sioux City in 1989. Defects in the metallurgy is an area that is extensively checked within Rolls-Royce where they perform a wide range of non-destructive testing techniques to ensure safety risks are kept to a minimum.

2.4.2. NON-DESTRUCTIVE EXAMINATION

Non-Destructive Examination (NDE) also known as Non-Destructive Testing (NDT) does not add value to the disc in any form other than increasing confidence that there are no inclusions or cracks of a detectable size within the material. As the disc is a critical component, the integrity of the disc must be certain to avoid catastrophic failures, similar to those described in the previous section. The discs are required to operate at their redline speed (~120% of the max operating speed) for five minutes (JAR-E840 2001) and withstand terminal speed in the event of a shaft break. To ensure part reliability is maximised, integrity and lifing calculations assume microstructure anomalies within the final component are below a threshold value. Microstructure anomalies above this threshold value must be detected using ultrasonic NDE.

To pick up any possible microstructure anomalies several types of examination are performed during the manufacture of a disc, these include:

- Visual inspection –A human inspector analyses the surface of the component for any surface defects which could lead to part rejection.

- Penetrant inspection – Penetrant fluid is applied to the surface of the component and left to dwell allowing the fluid to move into cracks, once the excess fluid has been removed a developer is applied to provide contrast highlighting surface cracks with the aid of a UV light.
- Eddy current inspection – An operator moves an electric coil around the surface of the part to detect inclusions beneath the surface.
- Ultrasound inspection – An operator moves an ultrasonic probe over the surface to detect inclusions deep within the part.
  (Raj, Jayakumar et al. 2002)
- Sample analysis – A sample is removed from the material around the part to check the material microstructure.

All are NDE techniques with the exception of the sample analysis where testing is performed on a sample taken from the waste material which should closely resemble the material properties within the disc, this is normally an integral test ring taken from redundant material designed to be extracted between heat treatment and the turning of the rectilinear inspection geometry stage.

To ensure any possible material anomalies within the finished disc are detected, ultrasonic NDE is performed on every disc at two stages during manufacture (Benz 1999), on the Billet (prior to forging) and RI geometry (post heat treatment). The parts are ultrasonically inspected at the latest possible stage during the manufacturing process to ensure defects are identified. Rolls-Royce performs a detailed ultrasonic inspection pre forging and post heat treatment, just prior to final machining. The most recent ultrasonic NDE specification[11] at Rolls-Royce requires scans in at least four directions over the entire disc to be classed as fully inspected (Rolls-Royce 2008). Common consensus shows that the more scans performed on the disc, the greater the certainty that there are no internal defects greater than the detectable size (Wright 2010).

There are several reasons why ultrasonic inspection should be performed thoroughly at the RI geometry stage:

- To pick up defects which were possibly missed during the initial billet inspection stage.

[11] RRP-58008 - Rolls-Royce (2008). RRP 58008 - Methodology for Consitruction of Ultrasonic Inspection Scan Plans and Ultrasonic Limitation Diagrams, Rolls-Royce - Aerospace Process Specification.

- To pick up the small defects which were previously too small to pick up during the initial billet inspection stage as during billet deformation these defects may have been exaggerated.
- To pick up defects caused during the forging process. This could be caused for example by excessive strain rates.
- To pick up defects caused during the Heat treatment process, predominantly the quenching stage which may have caused defects.

The ultrasonic NDE specification within Rolls-Royce simplifies the complex behaviour of the ultrasonic scans to a set of rules that need to be followed. This ensures the disc profile within the RI inspection geometry is confidently inspected. To ensure a disc is fully inspected by the ultrasonic scans, the rules are applied to every surface to locate areas where the scan should detect an internal defect. To achieve a fully inspected part, all areas of the disc must contain at least four scans made up of normal and angled scans (See Appendix II). If any part of the disc is not fully scanned, the part is reviewed by the 'Materials' and 'Critical Parts Lifing' departments to check if limitations are of concern. The reason not all parts are fully inspected is due to the time it takes to create the inspection diagrams and to solve the limitation problem. Current ultrasonic NDE inspection diagrams take up to 8 hours to draw and analyse, and as long as 10 days lead time is required prior to that (Wright 2010). It is for these reasons that only a few design iterations are performed. The process requires Manufacturing Engineers (ME) and NDE experts to use their knowledge to design the rectilinear inspection geometry. Software developed by Rolls-Royce to replace manual methods has been trailed but with limited success (further details covered in Section 3.3).

In the future, ultrasonic inspection technology could move towards phased arrays (Wright 2010). Phased arrays involve a linear array of small simple ultrasonic probes which allow electronic scanning, focusing and deflections (Poguet, Garcia et al. 2002) by creating an artificial wave created by changing the timings of the individual probes. This makes the array probe far more versatile than a single ultrasonic transducer probe. The phased array probe can produce a clear cross section of the inspected region where as the single transducer must be translated across the surface but can only detect the depth of the ultrasound reflection as illustrated in Figure 14.

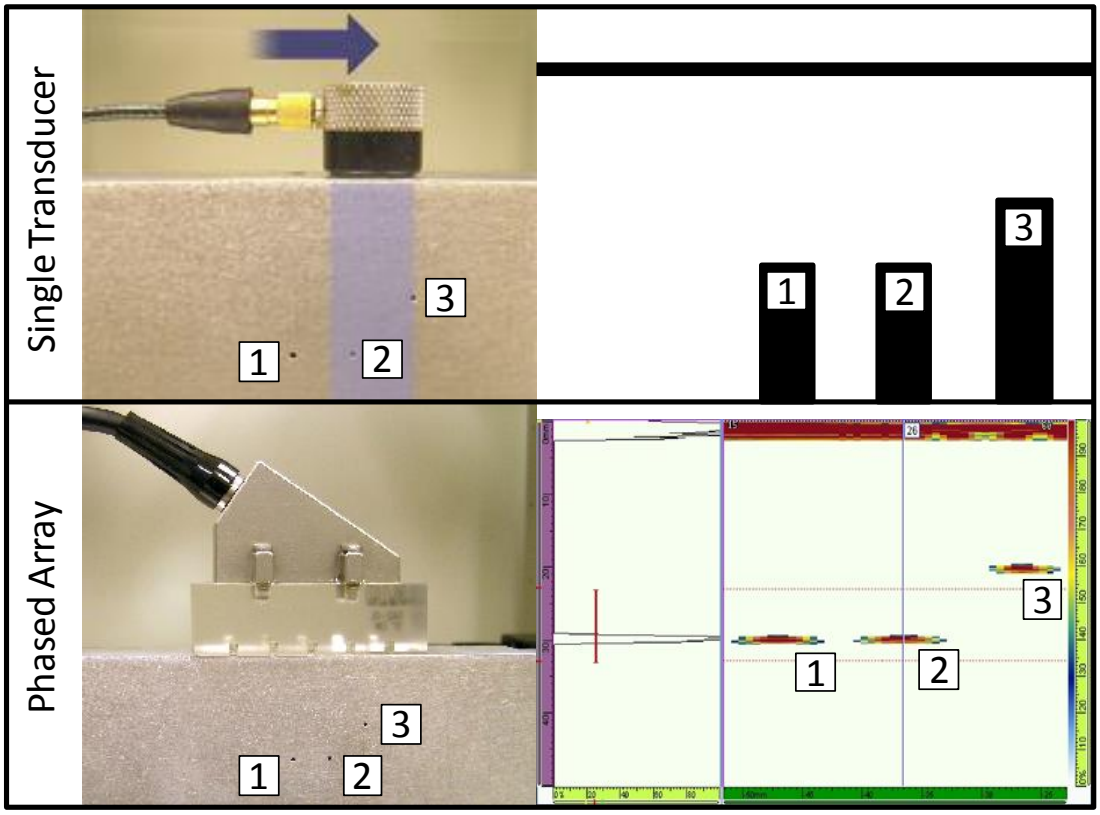

Figure 14 – Single transducer Vs. Phased array B-Scan output. Pictures adapted from (Olympus Industrial 2011)

Non-Destructive Examination Summary

As mentioned, NDE is a constraint during the production of the disc, and it is this constraint which adds considerable volume to the disc which creates a heavier forging. It is also noticed that very few design iterations are performed on the rectilinear shape due to the time it takes to create limitation diagrams. This is an area that could be looked into to help improve the rectilinear geometry generation around the disc.

2.5. MODELLING COST

Early cost estimation of a concept design is important for the overall cost of the disc. It is found that 70% of the total cost is determined during the initial design stage, although this stage only accounts for a small proportion of the total development costs (Tammieni, Scanlan et al. 2007). Any alterations to the design after this point will become increasingly more expensive as the part goes into production. Therefore the best opportunity to reduce manufacturing cost is during the design stage (Schreve, Schuster et al. 1999). So it is important to get the design right first time to reduce potential costs later on. Designers are not experts in manufacturing so they are unlikely to choose a design of minimum cost as their priority is to ensure performance requirements are met (Masel, Young II et al. 2010). A

good way for designers to understand how the design impacts upon the cost would be to provide a live cost estimate. This would allow designers to systematically see how their design decisions impact on the manufacturability of the product (Schreve, Schuster et al. 1999). During the early stages of design, the cost estimates should only be as detailed as the designer has information, therefore costs should become more accurate as the design detail increases. Uncertainty will exist during the early stages of design so it is important that a statistical confidence should be part of the model (Schreve, Schuster et al. 1999).

Cost estimation is an important part of a business as over estimating the cost of a product can result in loss of business but underestimating costs may lead to profit losses *(Winchell 1989; Niazi, Dai et al. 2005)*. Cost estimation itself does not add any value to the product nor will it affect the actual finished cost, but it is a valuable tool for evaluating different design proposals (Winchell 1989). Cost estimates will also provide means of justification to why a design should or shouldn't be selected (Ou-Yang and Lin 1997; Duverlie and Castelain 1999; Layer, Brinke et al. 2002).

A key concept of concurrent engineering is to analyse factors which could affect subsequent production processes during the design stage (Ou-Yang and Lin 1997). This early analysis reduces the likelihood of needing large costly design changes towards the end or after the design stage has finished (Ou-Yang and Lin 1997; Layer, Brinke et al. 2002). The estimating of cost early in the design process will enable concept design decisions to be made with cost as a key factor; this information is vital as 80% of the cost of a product is committed during the design process (Layer, Brinke et al. 2002) therefore late design changes will typically cost more. A rapid cost estimate turnaround is also an important requirement as cost feedback need to be provided concurrently before design alterations are implemented (Layer, Brinke et al. 2002).

Typically finished designs are sent out for quote to get the most accurate part cost, but there are several cost estimation techniques available which can predict part costs from concept right through to detailed design with varying accuracy. The main techniques for product cost estimation has been reviewed in a paper by *Niazi et al.(Niazi, Dai et al. 2005);* the distinct approaches highlighted in the paper are shown in Figure 15 below.

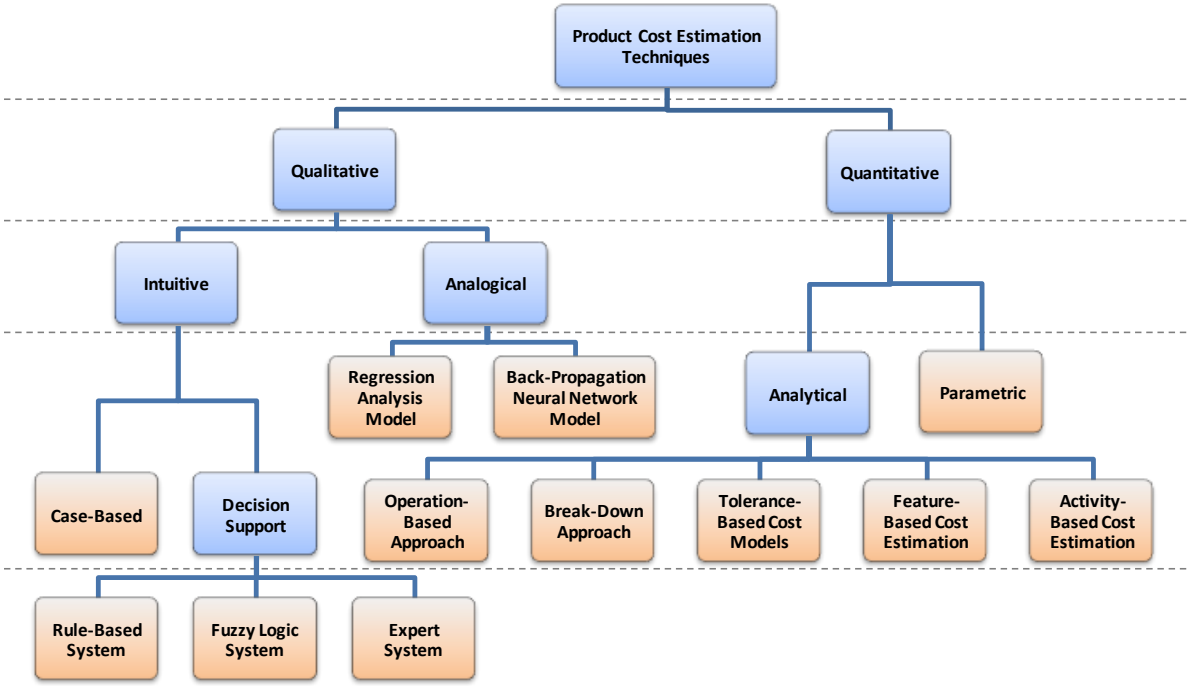

Figure 15 - Classification of product cost estimation. Taken from (Niazi, Dai et al. 2005)

The product cost estimation techniques are split into two high level techniques: qualitative and quantitative.

Qualitative cost estimation techniques are primarily based on comparison analysis between the new product and existing products to identify similarities. The identified similarities can then be used to help incorporate the past data into the new product to avoid gathering this data from scratch. The technique is useful to generate rough cost estimates or to serve as a decision aid during the early stages of the design process, however when a sufficiently detailed design is available a quantitative technique can be used to produce more accurate estimates (Niazi, Dai et al. 2005).

The Quantitative cost estimation technique is based on detailed analysis of a products design, features, and manufacturing processes rather than relying on past data or knowledge. It is noted that although this technique is "known to provide more accurate results, their use is normally restricted to the final phase in the design cycle" *(Niazi, Dai et al. 2005), Niazi et al.* states that its use is normally restricted to the final design phase is due to the requirement of a detailed product definition (Niazi, Dai et al. 2005).

Niazi et al. splits the quantitative technique into two categories, the first quantitative technique discussed is the parametric cost estimation technique. This technique uses statistical methodologies to express a part costs as a function of its variables or derived

properties. This method is useful where known cost drivers can be easily identified. Although easy to develop, one of the main criticisms highlighted by *Duverlie et al.* is that it works similar to a "black-box"[12] (Duverlie and Castelain 1999). Although a cost change can be easily identified, the specific details of the manufacturing process that caused the cost change is unknown.

The second quantitative technique discussed by *Niazi et al.* is the analytical cost estimation technique. This technique requires decomposing a product into elementary units, operations and activities that represent the different resources required during production; the total cost is expressed as the sum of all these components. This technique can be further broken down into specific methods; these are listed in Table 4 below.

Table 4 - High level analytical cost estimation techniques (Niazi, Dai et al. 2005)

Technique	Key Advantages	Limitations
Operation-based cost models	Alternative process plans can be evaluated to get optimized results.	Time-consuming, require detailed design and process planning data.
Break-down cost models	Easier method.	Detailed cost information required about the resources consumed.
Cost tolerance models	Cost effective design tolerances can be identified.	Require detailed design information.
Feature-based cost models	Features with higher costs can be identified.	Difficult to identify costs for small and complex features.
Activity-based cost models	Easy and effective method using unit activity costs.	Require lead-times in the early design stages.

In the paper written by *Caputo et al.(Caputo, Pelagagge et al. 2016)* the idea that statistical analogical cost estimation techniques which reside under qualitative methods do not usually allow cost based comparisons between differing designs due to the cost model's low level of detail. The Analytical methods are thought to be the most accurate technique in the sense that they can try to portray the actual production process by breaking the cost down into each individual manufacturing step. The sum of each monetary manufacturing and processing step will yield the theoretical unit cost of the product. Detail such as material quantity, labour rates, time, and prices are used to estimate the direct costs of an activity or product. The paper also highlights that the analytical models requires a large amount of

[12] A "black-box" methodology is where contents of the system is unknown, only the inputs and outputs can be analysed.

information and are more time consuming due to the requirement of a detailed design and knowledge of the manufacturing process; this last point is also highlighted by *Layer et al.* who also suggests computer support is necessary for it to work effectively (Layer, Brinke et al. 2002).

2.6. CURRENT DISC DESIGN AND MANUFACTURE PRACTICE

The disc design process is heavily constrained and to a large extent performed manually greatly restricting exploration of the design space. There are several research papers and projects pushing towards fully integrated disc design environments where basic 2D axi-symmetric design can be thermo-mechanically optimised to minimise part mass (Farshi, Jahed et al. 2004; Kesseler and van Houten 2007; Brujic, Ristic et al. 2009; Huang, Wang et al. 2011). The process after the disc design has been fixed is briefly described next. After a disc design has been commissioned for production, the design is passed on to Manufacturing Engineers (ME). The aim of the ME is to create rectilinear inspection geometry of minimum volume whilst maintaining full ultrasonic inspectability to the required standard. This design stage might undergo 2 or 3 design iterations until the lifing experts are satisfied that all requirements are met. Upon agreement on the rectilinear inspection geometry profile, the design is then sent to the forgers. The forgers will apply their own forging rules, knowledge and modelling to come up with forging designs of least cost, which adhere to the required material properties. Once the forgers have established a suitable forging geometry the total costs can be found, this process is illustrated in Figure 16 below. This manufacturing design route is similar to that shown by *Srivatsa* (Srivatsa 2003).

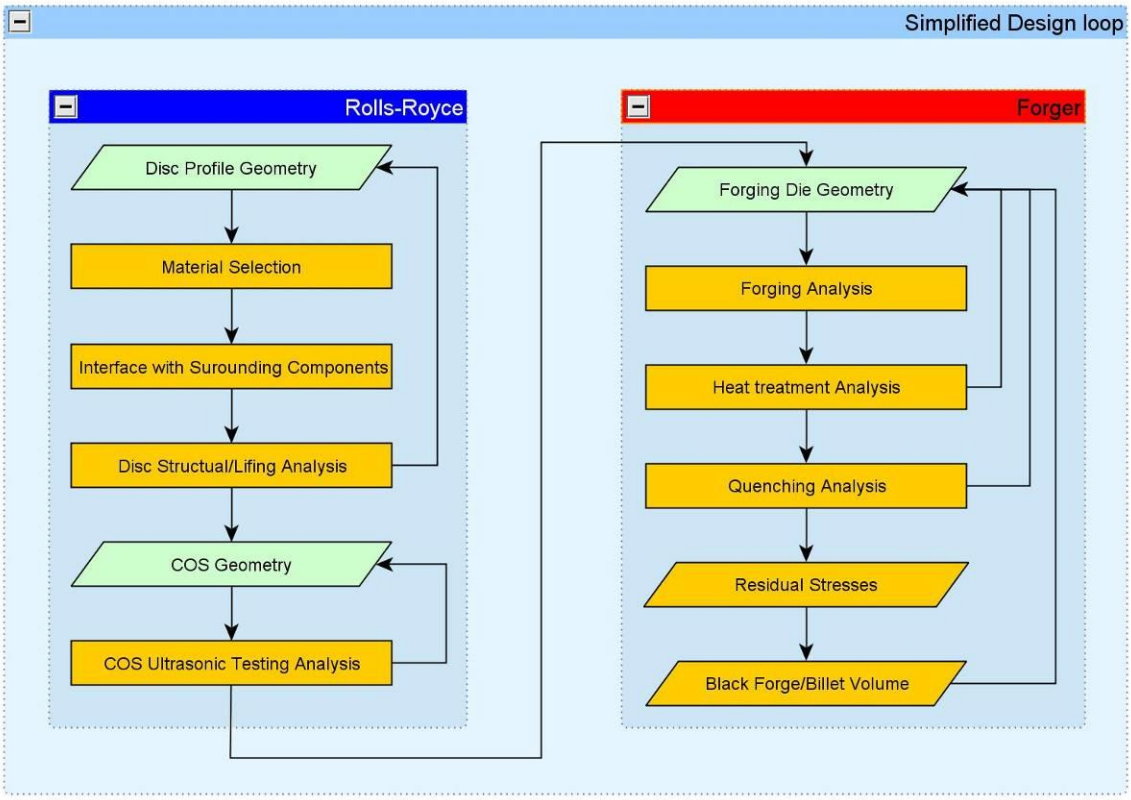

Figure 16 - Simplified disc design and manufacturing design process

2.7. THESIS FOCUS

It has been established that in order to reduce the unit cost of a disc, the optimal disc design needs to be found before the tooling and manufacturing has started. Cost estimates for a single design may take weeks to be provided and require expensive human resources to interoperate the design, create manufacturing layouts, and perform any necessary analysis. It would therefore be extremely costly to request detailed cost estimates for each manufacturing process for multiple concept designs. To remove these timely and costly overheads a method of achieving cost estimates for each manufacturing process is required. An automated approach could be the answer if all aspects of the problem can be modelled efficiently. This will not only save on human resources, but could enable optimisation to be performed.

With the required manufacturing and analysis models in place the rectilinear inspection geometry could be manipulated by an optimiser with the objective of minimising part unit cost. This could then be evolved to allow disc designers to rapidly evaluate the unit cost of a disc to allow relative comparisons of designs. The disc designs should theoretically have the same stress and life properties, but the cost of the discs will be different due to the volume and manufacturing processes. In order for such a tool to exist all the disciplines required

following completion of a disc design need to be captured and collated. The current list of data that needs to be collated or created is as follows:

- Preliminary disc design variables
- Simple lifing calculation
- Rectilinear inspection geometry generation and parameterisation
- Ultrasonic NDE rules
- Ultrasonic NDE cost model
- Machining rules
- Manufacturing cost model
- Heat treatment rules
- Heat treatment cost model
- Black forging rules
- Forging cost model
- Billet material costs

The following subsections analyse recently published work on similar or related problems.

2.7.1. MULTI-DISCIPLINARY DISC OPTIMISATION

There are relatively few published papers on disc manufacturing optimisation; this is due to there being only a handful of gas turbine engine manufacturers in the world as well as the confidential nature that surrounds such technology. The complexity of Multi Disciplinary Optimisation (MDO) has evolved in parallel with the development of computational capabilities. This has allowed ever more accurate models to be run within iterative loops in reasonable time periods. The studied papers related to disc optimisation have all been written within the last 20 years. The availability of batch capable software has been the main cause of hindrance as shown in General Electric's published papers (Rohl 1997; He, Rohl et al. 1998; Rohl, He et al. 1998; Srivatsa 2003)

The earliest published disc MDO was by *Fisher et al* In 1997 to try and find the lowest cost solution in as short a time as possible (Fisher, Gunasekera et al. 1997). A Discrete Event Dynamic System is used to model the problem to find the cheapest manufacturing method to produce a part. *Fischer et al* knew that to avoid excessive computational time, the design space first needs to be narrowed down. This will free up more computational time which could be spent in the feasible regions to find the best solution; to achieve this, models that provide fast analysis to produce good estimates are required. An example talks about using a generalised hill climbing algorithm which requires 10^2 – 10^5 function evaluations. At that

time finite element analysis forging models took between 15 minutes to several hours, which would make this method highly infeasible in terms of time. To improve run times, low fidelity models are used. Once the design space has narrowed down to a "good enough" region higher fidelity models can be used to produce accurate results. At the time there was no consistent software available which integrated individual modules for automated analysis, so a system was developed which produced consistent forging, machining, extrusion and heat treatment models which executed "several orders of magnitude faster than finite element analysis" (Fisher, Gunasekera et al. 1997) enabling an optimisation to be run. This paper uses a very simplified cost model as the surrogate for optimisation rather than just using the material weight which is the most common technique. Costs include equipment, labour rates and tooling. The accuracy of these costs are unknown and the fidelity of the models is also not specified. The idea of using simplified models worked well in this paper by enabling fast convergence into feasible design space, but details of how the system can know it is within the top 5% of the absolute optimum is not explained (as the optimal point has not first been identified).

General Electric is the first engine manufacturer to publish a paper on the disc optimisation subject. Written by *Rohl et al* their objective is to find the minimum weight forgeable disc (Rohl, He et al. 1998). This is achieved by performing several optimisation loops within a bigger optimisation loop to satisfy many important constraints such as residual stresses, distortions and lifing. These outputs are found by running the following analysis models:

1. Mechanical
2. Forging
3. Heat Treatment
4. Machining
5. Lifing

Figure 17 below illustrates the key activities which were the focus for the developing workflows. It can be noted that Ultrasonic inspection of the forged part does not feature at this time.

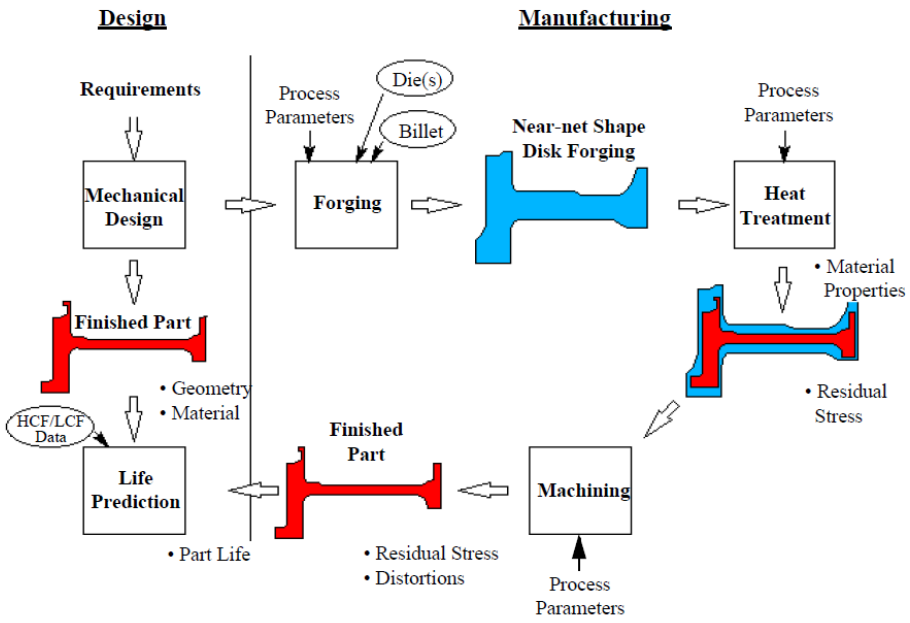

Figure 17 - Illustration taken from Rohl's paper displaying what are perceived to be key activities to produce a forged component (Rohl, He et al. 1998).

The 1998 paper written by (Rohl, He et al. 1998) is preceded by two papers (Röhl, Srivatsa et al. 1997; He, Rohl et al. 1998). These papers all describe the ideal of an integrated procedure to achieve an optimal design but describe only one analysis of that shown in Figure 17. These two papers will be covered first.

In the paper by *Rohl et al.* (Röhl, Srivatsa et al. 1997) the focus is on optimising the heat treatment model required for a larger complete optimisation model. The need for a fast analysis tool is emphasised in this paper which is understandable considering the potential size of optimisation. In this paper *Rohl et al.* uses commercial software DEFORM™ for the analysis and Isight[13] as the optimisation driver. The analysis performed avoided using time consuming heat transfer-stress analysis for each point; instead only a heat transfer analysis was used. To take into account the stress analysis a simple quadratic penalty function was used, penalising non-uniform cooling. Each iteration took around 13 minutes to compute and took around four hours to converge using a gradient based optimisation method. To quickly narrow the design space down, for the first 6 iterations only one variable was used. It was concluded that considerably lower residual stresses were possible without the use of

[13] Software used to combine multiple cross-disciplinary models and applications together to enable an automated workflow.

time consuming coupled thermo-stress analysis. The simplification of the heat treatment problem is shown to be very effective on actual turbine discs and is said to be applicable to many forged discs, but there are no cost or performance implications described for differing levels of-non uniformity.

In 1998, *He et al* published work creating an automated isothermal closed die disc forging optimiser (He, Rohl et al. 1998), which again uses DEFORM and Isight. The objective is to minimise the cover over the geometry whilst not exceeding the maximum press load. In this example a disc with 6 radii parameters is used. A 10% perturbation of the design variables is performed during sensitivity analysis which is used to smooth out the design space for the gradient based optimisation algorithm. After 50 hours of user time and 15 simulation runs the press load dropped by 20%, however the forged part weight increased by 12%. The paper shows how forging shape optimisation can improve the forgability of a part significantly. This paper refers to the same integrated procedure as described in the papers written by *Rohl et al.* (Röhl, Srivatsa et al. 1997; Rohl, He et al. 1998), a single optimisation run of 50 hours could be excessive if this method is to be run within a larger optimisation design architecture.

The 1998 paper written by *Rohl et al.* (Rohl, He et al. 1998) combines the knowledge described in the previous two papers as well as a more detailed explanation of the integrated procedure to optimise a manufacturing route for minimum part and billet weight. The need for detailed simulation of the manufacturing process is required in order to find the residual stresses and subsequent distortions of the finished part. This paper "demonstrates an integrated procedure that addresses both mechanical design and manufacturing processes"; this is necessary due to the iterative nature of the process. *Rohl et al.* notes that individual simulation tools of individual stages exist but they have never been used in an integrated fashion (Rohl 1997). DEFORM™ [14] is used for both the forging and heat treatment modelling. The optimisation strategy looked into was Concurrent Subspace Optimisation (CSSO), but this was not performed due to the integration of the mechanical design and engineering analysis to the system was lagging behind in development. Response surface models and other MDO techniques are discussed as potential methods once the system was fully coupled. The paper describes a near complete optimisation environment and the current issues that need to be solved in order to achieve

[14] An engineering software that enables designers to analyze metal forming, heat treatment, machining and mechanical joining processes.

the aim. The individual discipline constraints and objectives appear well thought out and assembled in a logical order, but the process misses out physical rectilinear inspection geometry which should exist between the heat treatment and machining geometry. The major optimisation objective is initially to minimise weight and/or cost, but later in the paper a combination of part and billet weights are set as the objective. With such a detailed manufacturing model envisaged a cost model could be effectively used in place of weights.

In *Srivatsa*'s paper, he brings the three papers written by *Rohl et al.* together to form the multi-disciplinary optimiser (Srivatsa 2003). Actual discs were run through the optimiser creating disc forgings with material reduction around 5-10% less than the manual process. If this reduction is compared to the minimum material usage, then material saving is 50-80%. Several constraints are considered which include, Equipment, Material, Geometrical, die stresses and time to forge with an overall optimisation objective to minimise material weight. This project combines multiple detailed manufacturing models which could be effectively utilised within a detailed cost model but the focus is still primarily on achieving minimum forge/billet weight. Figure 18 below is the re-engineered diagram of the one shown previously in Figure 17 above. It shows that that ultrasonic inspection has now been considered one of the key activities during the manufacture of a forged component. Ultrasonic analysis is also now mentioned in the future work section where it states that sonicability should be included as a geometric constraint during forging optimisation. Despite highlighting the existence of the ultrasonic constraint, the paper fails to demonstrate or even consider the impact it could have on upstream manufacturing geometry. Failing the insertion of this constraint could impact the accuracy of the analysis within the paper.

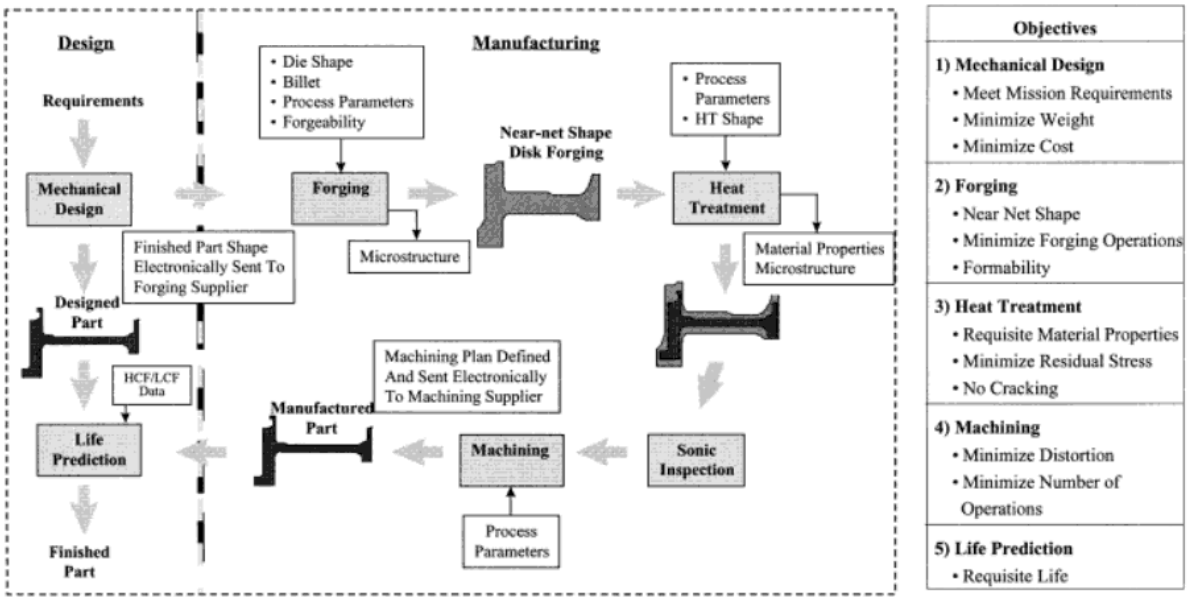

Figure 18 – Illustration taken from Srivatsa's paper displaying what are perceived to be key activities to produce a forged component (Srivatsa 2003).

2.7.2. DISC OPTIMISATION

Rao et al. noticed why MDO hasn't been used much in the past, for many disc optimisation cases weight has been used as the surrogate for optimisation. However in virtually all cases it is the cost that needs to be reduced and not the weight, as there isn't always a positive relationship between the two (Rao, J.P.Scanlan et al. 2007). It was identified that there are a lack of detailed cost models to allow cost to be made the surrogate for optimisation; this is made more difficult when relating the manufacturing costs to the geometry parameters of a disc. In this paper 10 variables were optimised using a genetic algorithm (8 geometry parameters, surface finish parameter and shot peening intensity). After approximately 2500 evaluations it was found that a design of the same weight (whilst satisfying the same constraints), could vary in cost by up to 11% showing that modelling disc manufacturing cost is very important. Despite using the cost model to prove that weight and cost are not directly linked, the method for obtaining the rectilinear inspection geometry to calculate costs from is not explained therefore the cost result accuracy could be questioned.

Tammieni et al. looked at applying a cost model to the design of a disc in order to support design decisions (Tammieni, Scanlan et al. 2007). The use of spreadsheets to interact with cost models are widely used, but it was noticed that the use of spreadsheets resulted in many different errors. The DATUM project unearthed a tool called DecisionPro (now Vanguard Studio) where object-orientated methods can be employed for the model and simulation. The software allows the cost to be broken down into smaller portions making

the cost structure more transparent. The method and software described in this paper is very insightful and could prove useful when creating and analysing detailed cost models.

Masel et al. looked into creating a methodology to predict the material volume input required to forge a finished disc during the preliminary disc design stage (Masel, Young II et al. 2010). He states that a real-time method for estimating forging volume is required to aid the designer make design decisions based on performance and production cost. This not only will reduce design time, but also yield a design of low manufacturing cost. There are four different approaches to estimating the cost of a disc forging: intuitive, analogical, parametric and analytical. The parametric method was brought forward as it was considered to be more robust due to the fact that the costs can be calculated using attributes of the part's features. The methodology used to create the black forging shape is rule based, where an intuitive step by step addition of material is performed similar to that shown by *Kulon et al* as seen in Figure 6. The software developed for the purpose of cost estimation was shown to be robust, but the method did not turn out to be very accurate with only 59% of designs producing a result within 10% accuracy. This paper talks about the desire to estimate part cost, but only part weight is discussed in the method and results. The method for obtaining the black forging weight is very simplified, but it does contain an attempt to cater for the ultrasonic analysis constraint by adding additional mass to the component; an important step which *Srivatsa*'s (Srivatsa 2003) method was missing.

Brujic et al. has looked into disc optimisation in terms of final weight through varying the geometry of the disc and the Fur tree features (Brujic, Ristic et al. 2009). This was achieved whilst constraining maximum stresses in throughout the disc. This paper is purely disc design so does not consider manufacturing costs nor material properties developed during manufacture.

Nagendra et al. have created a strategy to produce a rapid evaluation of designs using Neural Networks (Nagendra, Staubach et al. 2005). The example given has ten variables which took 7 hours to compute, but once trained new evaluations could be computed in a tenth of a second with errors less than 4% of the actual values. This work was based on optimising a turbine disc design for minimum weight with thermal and mechanical constraints. The use of neural networks could have also been trained to estimate costs if the data was available increasing the value of the work.

2.7.3. *FORGING OPTIMISATION*

Like many manufacturing industries, forgers retain most of their intellectual property. Material specific processes and design rules are for this reason extremely difficult to obtain. Despite this fact, there are still papers on disc forging optimisation such as the paper by *Castro et al* who creates a method to find the optimal preform shape and temperature which is free of barrelling[15] (Castro, Antonio et al. 2004). In this case they look at the inverse problem using FE analysis and a genetic algorithm to minimise energy. Although the example given has a small number of elements for a small aluminium billet, modelling the mechanics and heat transfer takes only a few seconds to compute, but takes 2-4 hours to optimise. Future investigations into more complicated preform dies is noted to take longer compute times.

2.7.4. *ALTERNATIVE APPLICATIONS*

The use of an ultrasonic inspection constraint to design axi-symmetric forged parts is not limited to discs within gas turbine engines; ultrasonic inspection is also used during the production of land based turbines (Azuma, Tanaka et al. 1997) including steam turbines. Another forged component which also undergoes ultrasonic inspection are turbine shafts which are used in gas and steam turbines. Other applications include the rail industry where train wheel axles which are classed as 'safety-critical' (Zerbst, Beretta et al. 2013) are ultrasonically inspected throughout their lives to avoid catastrophic incidents from occurring (Wise 1968; Szelazek 1987; Lonsdale and Swarzell 2000; Klinger and Bettge 2013; Zerbst, Beretta et al. 2013). Another potential application could be for high performance wheels which are used between a vehicle and the ground such as racing and aircraft wheels which operate at high speeds.

2.8. LITERATURE REVIEW KEY FINDINGS

This literature review initially broke the disc manufacture and design down into five main topics; material, production, safety, cost and design. From this it became apparent that to reduce the unit cost of a disc, both production and safety needed to be considered. Section 2.7 looked at previous research in disc design and manufacture of which none of the reviewed research had covered these key points.

The papers reviewed give a general overview on what has been achieved in terms of disc and manufacturing optimisation. General Electric has demonstrated advancements in their

[15] Barrelling or bulging is where friction exists between the die and work piece preventing the contact surface from expanding out radially. This results in a greater radial expansion half way between the two contact surfaces.

Collaborative Optimisation Environment (COE) laying the path for a fully automated and synchronised disc manufacturing optimiser. Although noted that ultrasonic inspectability could be added into the optimisation loop, there has been no work in utilising or optimising this constraint as yet. There are also no fast automated programs which can generate ultrasonic limitation analysis. The creation and addition of ultrasonic NDE analysis is an area that could potentially produce further savings therefore should be investigated further. Most disc manufacturing research focused on optimisation primarily to improve material properties in the part and only one had an objective of reducing part cost. There was no research related to manipulating the rectilinear inspection geometry to either reduce cost or improve material properties. The research also showed that investigations were primarily performed at the detailed design stage.

2.9. THESIS AIMS AND OBJECTIVES

This thesis sets out to link cost to both safety and manufacture within a design optimisation study. This should be suitable for use at the detailed and preliminary design stages enabling concurrent engineering within the design process. The deliverable is to create an automated manufacturing design process which aims to optimise the unit cost of an HPT disc as a primary objective through the manipulation of the Rectilinear Inspection (RI) geometry. This process should also maintain full ultrasonic inspection for safety and be able to relate the impact of the design on material properties.

In order to create the rectilinear inspection geometry optimiser for unit cost, there are several problems that need to be solved:

- Current RI geometry creation methods are basic and are unable to find fully inspectable designs every time.
- There is no suitable ultrasonic analysis software available which can assess the ultrasonic limitations of a design to the required practices.
- No automated methods available to this project which automate the generation of Isothermal forging die designs for use in estimating cost.
- No simple methods to relate heat treatment design to the life or material properties of the disc
- There are no cost models which utilise manufacturing design optimisation to generate cost.

By solving each one of these problems and piecing them all together the following could be available to the disc designer within a matter of hours enabling concurrent engineering design practices:

- An automated RI geometry design optimisation tool
- Full detailed ultrasonic inspection analysis.
- Estimate of the isothermal forging die and black forging design.
- Estimated heat treatment design and comparative measure for life/material properties.
- Mature unit cost estimation and unit cost breakdown.

These points all build upon each other and could be used to create an optimisation loop that is practical and useful. It will tell the disc design engineer the potential costs of the disc as well as enlightening the designer with cost breakdowns.

There are additional benefits which could come out of this project. Due to the lead time required in creating RI geometry, the design envelope of the disc is normally fixed early which can be frustrating to the disc design team. Fixing boundaries early also affect additional design team's components which interface with the disc as these normally have a shorter manufacturing lead time. Successfully creating an optimisation tool which will create a fully inspected RI geometry for a given disc design within a shorter time period. This will benefit both manufacturing engineers and ultrasonic NDE experts who will save several hours to days of time which can be better spent performing their primary roles. The design teams will also gain additional time due to the reduced RI geometry lead-time increasing timescales in which to lock down parts of their design.

Chapter 3 will focus on developing and testing these methods discussed climaxing with a fully automated manufacturing design optimisation for unit cost.

This page intentionally left blank.

Chapter 03 RECTILINEAR INSPECTION GEOMETRY OPTIMISATION DEVELOPMENT

3.1. CHAPTER DIRECTION

3.1.1. INTRODUCTION

This chapter is based around the development and testing of methods to enable a turbine disc to be optimised for unit cost. The chapter will build the optimisation process up one method at a time so that the performance and benefits can be observed and discussed clearly.

3.1.2. OBJECTIVES

The objective which was set out in the literature review (Section 2.9) stated that the manufacturing design process needed to take into account both cost and safety of the product; therefore, the deliverable was to create an automated manufacturing design process which aims to optimise the unit cost of an HPT disc whilst maintaining full ultrasonic inspection whilst also considering the impact on material properties. For this optimisation to be useful within a concurrent engineering context, the time to produce an acceptable result should be achievable within 48 hours or less (to simulate a weekend). In order to successfully achieve this goal, the following automated methods need to be developed:

- A universal parameterised rectilinear inspection geometry which can adapt to any given disc design.
- An ultrasonic NDE analysis program which can perform an inspection analysis on any given disc using any given rectilinear inspection geometry.
- A method to estimate the black forging mass required to manufacture any given rectilinear inspection geometry.
- A method to convert any given black forging geometry into suitable heat treatment geometry for simplistic residual stress and minimum material property assessment.
- A detailed analytical unit cost model for use with any given HPT disc design.
- An optimisation strategy which can take in a given HPT disc design and generate an optimal solution within the timeframe.

One additional key element which needs to be incorporated into all six of these methods is speed. All methods need to be programmed efficiently to maximise the optimisation potential for the given time frame.

The flow of work undertaken in the following chapter will start by automating the generation of RI geometry through the use of optimisation. Section 3.2 uses the RI geometry mass as a basic surrogate for unit cost, but the optimisation results are somewhat flawed as they miss out critical safety analysis. Therefore, in Section 3.3 new high performance ultrasonic NDE limitation software is developed which allows the RI geometry to be optimised for unit mass whilst achieving the inspection requirement constraint. Section 3.4 looks into improving the 'cost' surrogate by replacing the optimisation objective with a new sophisticated isothermal forging die model to estimate the black forging mass. This takes into account the upstream manufacturing process creating a better cost estimate of the part. With a good model for the black forging geometry, it is possible to easily estimate the heat treatment geometry to obtain a good surrogate for disc material properties, the method by which this is achieved is shown in Section 3.5. Section 3.6 completes the chapter by converting all the newly acquired data created in the previous sections to populate a detailed cost model which becomes a fitting surrogate for cost. This means that when optimising the RI geometry, the unit cost will be directly minimised whilst conforming to inspection constraints. By keeping the heat treatment material property surrogate combined with multi objective optimisation, the designs will also be tailored to help improve the minimum material properties of the finished disc.

3.1.3. CASE STUDY EXAMPLES AND NORMALISATION OF DATA

To help illustrate and verify results in this chapter, four HPT disc case study examples are used:

- Case Study Disc 1 – Referred to as 'CS-1' in this chapter. Note the current RI geometry design has two stages for ultrasonic inspection analysis. This means the manufacturing process for this part is more complex. It is for this reason only the final RI geometry design is considered for comparisons, but in reality the manufacturing masses are considerably higher.
- Case Study Disc 2 – Referred to as 'CS-2' in this chapter.
- Case Study Disc 3 – Referred to as 'CS-3 in this chapter.
- Case Study Disc 4 design envelope – Referred to as 'CS-4 in this chapter. The design envelope refers to the area in which the finished disc will be manufactured as the design is not complete.

Many of these disc designs have existing RI geometry designs which are in production (See Appendix VI for part references). Some examples are relatively new so they do not accurately represent mature production designs of the RI geometry and or forgings. The 4 case study examples were selected as they are all manufactured using the same processes and are designed to production standards to enable a like-for-like comparison.

All case studies in this thesis are based on existing part data for model verification. Data which is potentially commercially sensitive will therefore be normalised. The methods for which the data will be normalised is as follows. The mass for each case study will be divided by the existing rectilinear inspection geometry mass (referred to as "Existing RI geometry design" in results tables) for that particular engine with the exception of CS-4 where no data exists so will therefore use data from the CS-3 design. The normalised mass calculation is as follows:

$$\text{Mass}_{\text{norm}} = 100 \times \frac{\text{Mass}_{\text{obj}}}{Mass_{\text{RI\_geom}}}$$

Where:

- $\text{Mass}_{\text{norm}}$ is the normalised mass relative to the engine part's existing rectilinear inspection geometry calculated data.
- Mass_{obj} is the calculated mass of the part being studied.
- $\text{Mass}_{\text{RI\_geom}}$ is the calculated mass of the respective engine part's existing rectilinear inspection geometry.

Costs values will be normalised using the respective engine part's existing rectilinear inspection geometry to estimate the unit cost (using the cost model specified). The normalised cost calculation is as follows:

$$\text{Cost}_{\text{norm}} = 10{,}000 \times \frac{Cost_{\text{obj}}}{\text{Cost}_{\text{RI\_geom\_pred}}}$$

Where:

- $\text{Cost}_{\text{norm}}$ is the normalised cost relative to the engine part's existing rectilinear inspection geometry.
- $\text{Costs}_{\text{obj}}$ is the calculated cost of the part being studied.
- $\text{Cost}_{\text{RI\_geom\_pred}}$ is the calculated unit cost of the respective engine part's existing rectilinear inspection geometry (using the cost model specified).

Heat treatment depth values will be normalised using the respective engine part's existing rectilinear inspection geometry to estimate the heat treatment depth. The normalised heat treatment depth is as follows:

$$\text{HT Depth}_{\text{norm}} = 4 \times \frac{HT\ Depth_{\text{obj}}}{\text{HT Depth}_{\text{RI\_geom\_pred}}}$$

Where:

- $\text{HT Depth}_{\text{norm}}$ is the normalised heat treatment depth relative to the engine part's existing rectilinear inspection geometry.
- $\text{HT Depth}_{\text{obj}}$ is the calculated heat treatment depth of the part being studied.
- $\text{HT Depth}_{\text{RI\_geom\_pred}}$ is the calculated heat treatment depth of the respective engine part's existing rectilinear inspection geometry.

Details of the specific rectilinear inspection geometry's used for each disc can be found in Appendix VI.

3.2. AUTOMATED RECTILINEAR INSPECTION GEOMETRY GENERATION

3.2.1. INTRODUCTION

This section focuses on the generation of Rectilinear Inspection (RI) geometry, often known as the Condition of Supply (COS) within Rolls-Royce (see Figure 19). The RI geometry generation will be required in every investigation within this chapter, it is therefore important that this geometry is generated efficiently and in an automated manner to maximise the productivity of all future studies. All studies will investigate the impact of RI geometry against a cost surrogate, for this reason it is important that a method is created to maximise the exploration of the available RI geometry design space. This will ensure the best possible designs are found during optimisation studies.

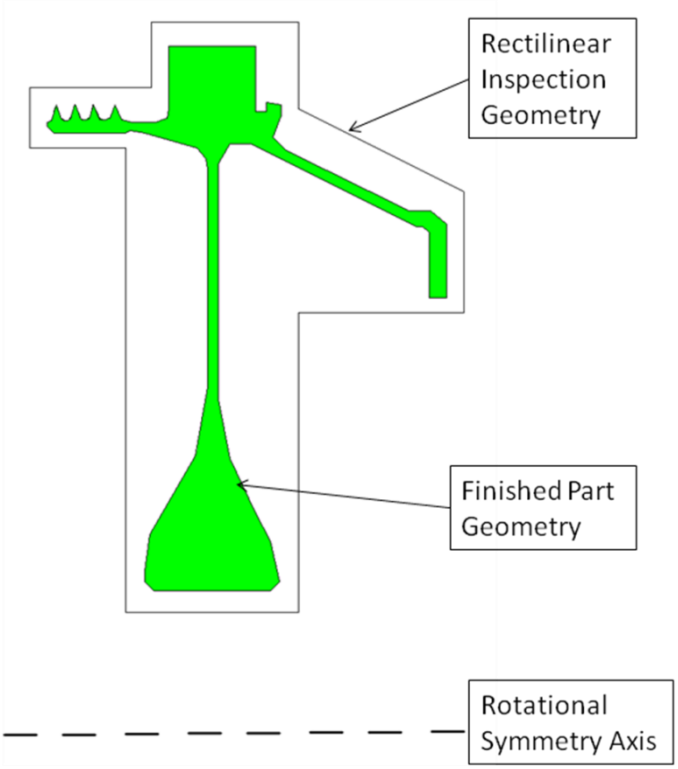

Figure 19 - Diagram showing both the finished geometry and the rectilinear inspection geometry. Note the example disc is a compressor disc.

The RI geometry can be thought of as the shape that will define the unit cost of a finished component for a given disc design. Normally once RI geometry has been defined and is in production it is rarely changed due to the high cost associated with changing the manufacturing production line.

Once the RI geometry has been defined a workflow will then commence which will result in a production line for the finished disc design. This work flow takes several months to years and will cost hundreds of thousands of pounds to complete; a manufactured to tolerance RI geometry will be the output.

This intermediate manufactured geometry holds two purposes; the first is the state in which the engine manufacturer will receive the part before further in-house manufacturing takes place. The second is the shape for which Ultrasonic Non-Destructive Examination (ultrasonic NDE) will be performed.

The ultrasonic NDE is necessary for lifing to check that there are no detectable inclusions within the finished disc profile. In order to ultrasonically inspect a component, the part must be manufactured with a RI geometry that fully encloses the finished disc part design. The design of the RI geometry can be extremely difficult to design due to the complexity of ultrasonic rules and the requirement to minimise cost.

Common methods for creating RI geometry involve the direct parameterisation between the geometry and the finished disc profile; these parameters can then be varied until a fully inspectable polygon is created. These RI geometry designs generally use simplistic perpendicular geometry created from approximately a dozen or more lines to fit around the disc profile. Although simple, this method is very quick to setup, but requires an experienced user with ultrasonic NDE analysis experience to set this up efficiently for both weight and inspection.

It is desirable from an ultrasonic NDE expert's point of view to design the RI geometry with minimal edges, as this reduces the complexity in drawing the required inspectability diagrams and simplifies the problem. From a cost point of view, the manufacturing engineer will want to increase the number of edges on the RI geometry as more corners will allow the RI geometry to wrap itself around the disc profile tighter. This will mean less material removal during in-house manufacturing and is also perceived to produce a cheaper component from the assumed reduction of upstream raw material costs and manufacture. Not only does this make the scan profile more complex, but it also reduces the amount of 'sound-energy' into the component making it more difficult to achieve the full inspection requirement. What makes the process even more challenging is that the inspection diagrams are currently drawn manually either by hand or using suitable CAD software all of which can take an expert 4 to 8 hours to solve; this is on top of 5 working days of lead time required to schedule the job in.

Manual implementation of ultrasonic NDE analysis can take as long as a day to complete depending on the complexity of the RI geometry, but if the part does not meet the strict requirements for inspection, an altered or new design will need to be drawn up and re-evaluated. The current process within Rolls-Royce to create the RI geometry is a process that goes through several iterations (approximately 2-5) so can take several weeks before a suitable design is selected. The reason for this is due to the time it takes for a level 3 NDE expert to find free time as well as completing the ultrasonic NDE inspection analysis. If the finished disc design within the RI geometry is found to be not fully inspected at the end of the analysis, then a lifing expert can asses if the features in question require the full inspection requirement. If the non-fully inspected region is a critical feature, then the engineer and or ultrasonic NDE expert must alter the RI geometry to try and achieve full inspection. This is no easy feat and requires knowledge of the rules as well as a logical mind in order to solve this sometimes frustrating puzzle. The iterative process which takes place during the design of the RI geometry is illustrated in Figure 20 below.

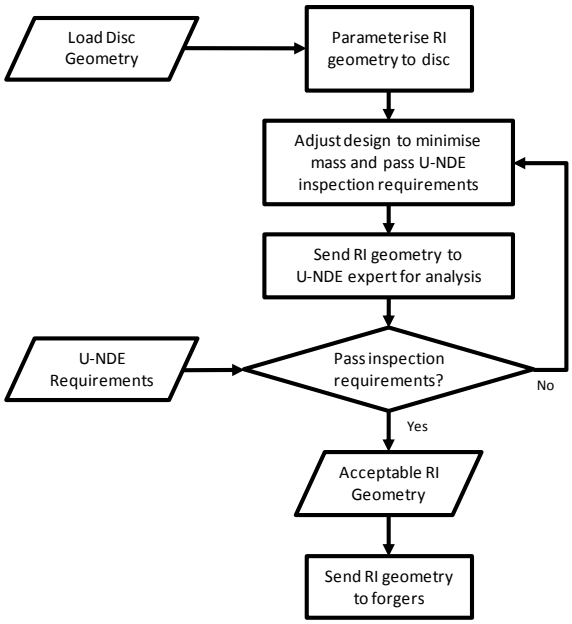

Figure 20 - Flow diagram of the current RI goemtry design process

With the development of fast automated ultrasonic NDE inspection software (described in detail in Section 3.3), as many as 60,000 RI geometries can be ultrasonically analysed within the same time period that the manual process takes just to analyse one[16]. With this in mind it has recently become possible to automate the RI geometry design process shown in

[16] Based on a one week turnaround from the submission of the RI geometry to the NDE expert.

Figure 20. Automating the design of the RI geometries not only will allow effective optimisation procedures to be run potentially improving designs, but will also require less of the engineer's time.

The design of the parameterisation is extremely important as the variety of shapes created is entirely dependent on this and will determine how effective an optimised geometry can become. Maximising the flexibility of the design space will increase the likelihood of improving a design. Typically, the parameterisation of the RI geometry will be created using suitable CAD software. An experienced design engineer could create a basic parameterised RI geometry in under an hour. Creating more flexible parameterised geometry is possible, however the designer would require additional knowledge, programming skills and an additional week to develop. Although this is a possibility, from experience complex parameterisation within CAD programs can result in tight variable constraints, and multiple reliability issues. One way to remove some of the complex CAD parameterisation is to have a library of standard RI geometry designs, each one having a unique design. Utilising a large library of parametric RI geometry should improve the likelihood of finding an improved design, but will require considerably more setup time. The utilisation of a library will still constrain the design space as it is still limited by the designer's ingenuity; this could mean that the potential optimal design will be missed as complex or even simple designs do not have the opportunity to be analysed.

There has been very little research on the area of Ultrasonic inspection shapes. New processes have been created such as that created by Rolls-Royce, who have used some of the rules described in an internal specification RRP58008 (See Appendix II) to intelligently create a starting RI geometry. This approach is interesting as it frequently generates geometry similar to existing designs which can then be manually adjusted to improve both inspection and design. The improvement of these designs still requires a manual knowledge based approach which may still not be solved in all cases.

Using methods such as a rule based RI approach severely limits the potential benefits of any optimisation as once the baseline geometry has been set, only local (perturbation of nodes) optimisation can be performed. For this reason, it is desirable to move away from this single rule based method and generate a system which will provide acceptable RI geometry, but is parameterised such that a large portion of the design space can be explored.

As explained previously, the benefits of optimising the RI geometry is entirely dependent on the initial geometry used. The number of edges used is extremely important as it provides a

trade-off between the complexity of solving the inspection problem and the amount of material required to manufacture the part. As well selecting the number of edges, locations of these edges could lie in an infinite number of positions, therefore selecting the right parameters for a given disc is extremely difficult. Creating good starting geometry that will give the optimisation a head start would be hugely beneficial in terms of savings and inspection.

The objective of this section is to automate the creation of rectilinear geometry around any given disc design. Due to the complex nature of the ultrasonic NDE inspection rules, a "right first time" approach is unlikely to succeed especially if minimum cost or weight is desired. Therefore, the rectilinear geometry should have sufficient parameters to enable a full but efficient exploration of the design space. The RI geometry could take an infinite number of shapes, therefore to ensure efficiency, inspection and manufacturing requirements must be taken into account.

The key outcomes of this section will be:

- To create an automated method of generating RI geometry for a given disc design. This method must include parameters to enable the creation of a wide variety of RI geometry for use within an optimisation loop.
- To design and setup an optimisation strategy to enable the RI geometry to be optimised for minimum mass.

The scope of this section is limited to the automated generation of RI geometry which will ultimately be used in conjunction with ultrasonic limitation software created later in Section 3.3.

Once the global design space has been searched it would be desirable to further optimise the output designs by improving them through local node perturbation. This stage will become more valuable in Section 3.3 when ultrasonic NDE requirements need to be satisfied.

Although the application of this work can be used for all generic gas turbine discs which receive a similar process, there is insufficient time to develop and test the software on all possible designs. Therefore, out of scope designs are listed:

- Large discs which require concaved bore cut-outs to allow ultrasonic NDE of difficult to reach regions of the disc due to ultrasonic limitations (for example fan discs).
- Large discs/design envelopes which would require a two stage inspection procedure to pass ultrasonic requirements.

The optimisation will have a similar philosophy to that of a selfish manufacturing engineer where the RI geometry will be minimised to reduce the amount of work required within his factory therefore the ultrasonic inspection requirement will not be considered. This means the RI geometry must be made manufacturable with minimal material removal to make the finished disc profile. To achieve this, the objective is to minimise the mass of the RI geometry.

The method described next will consist of four parts. First the high level method to generate initial RI geometry is down selected, and then an optimisation loop is setup to down select the parameterisation method to create RI geometry. The third part creates another new method to perform local optimisation. The final part combines the two geometry creation methods to create an optimisation process which will be used on four case studies.

3.2.2. METHOD – AUTOMATED RECTILINEAR INSPECTION GEOMETRY GENERATION

The philosophy chosen for this thesis has been to design the RI geometry generator to be as flexible as possible to increase the likelihood of finding the optimum design. In order to do this effectively it was decided to break the creation of the RI geometry into two key stages, a global search of the design space and then a local exploration of the design space with the selected best designs.

The global exploration of the design space will allow the RI geometry to be created from nothing but the input disc geometry. The input disc geometry will guide the RI geometry generator in an effective manner such that it will always create acceptable geometry from a manufacturing point of view. For example, the RI geometry will not overlap the input disc geometry, and will attempt to wrap itself as close to the disc geometry as possible to help minimise mass.

The second stage local optimisation will run off the output of the global optimisation. Using this starting geometry the nodes of the RI geometry can be perturbed with the objective of further minimising mass improving on the previous design.

Two stages are required as a compromise between searching the entire design space efficiently and locating the optimal design. A large number of parameters are required to explore the design space fully, and most of the output geometry will produce poor results. Adding a few rules to position the nodes close to the finished component reduces the number of parameters in half and will produce sensible geometry. At the end of this global search a new technique will be used on the current best designs which has far more

parameters, but will be used more efficiently as the starting positions will already be in good positions.

Tools Used

To generate the initial disc geometry Computer Assisted Design (CAD) software was required; in this thesis Siemens NX6 CAD software was used due to availability. The geometry was exported using a steriolithography format (surface triangulation) by extruding the 2D disc profile by a finite distance on the X-Y plane as illustrated in Figure 21 below. The steriolithography format was used primarily because the files are easy to decode and import into MATLAB® as a series of coordinates. The steriolithography format has an additional benefit of breaking curves and splines into a series of discrete points reducing the amount of programming required.

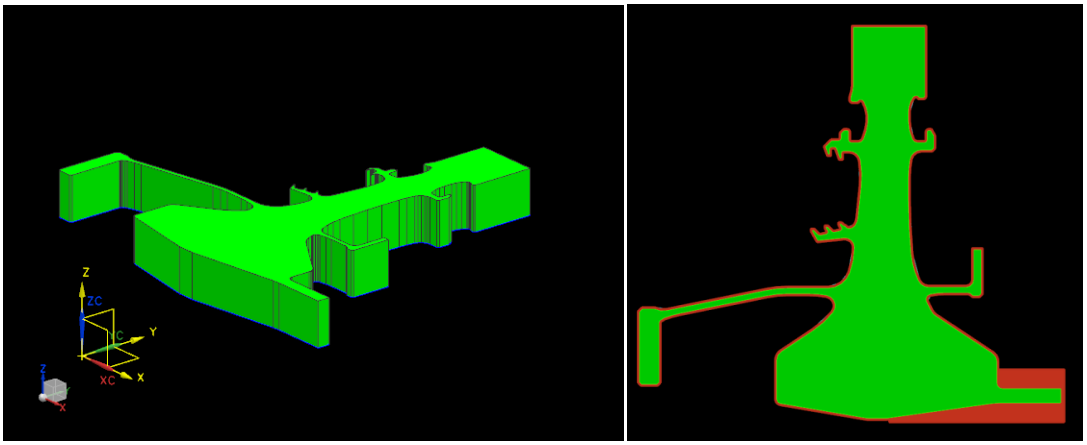

Figure 21 – Left: Extruded 2D disc profile to be exported to steriolithogical format. Right: additional material satisfying manufacturing requirements is shown highlighted in red.

The inclusion of manufacturing requirements is also key to the design of the RI geometry. In order to ensure any required additional material is included within the finished RI geometry, a supplementary model can also be created within NX6 (as shown in Figure 21 – Right) and exported in the same manner as described above.

The remainder of the software used throughout this thesis was programmed in MATLAB® although any programming language could be utilised. The MATLAB® programming language was used due to the experience held. MATLAB® also contained several optimisation algorithms such as the 'genetic algorithm' for both single and multi-objective problems as well as several gradient based optimisation algorithms such as the 'interior point' and 'trust region reflective' algorithms (found within the 'fmincon' function). The MATLAB® function 'fminbnd' was also used which is based on golden section search and

parabolic interpolation. These algorithms are used throughout this thesis where optimisation was required. The digital location of the code written for this chapter can be found in Appendix VII.

To simplify the code, all input and output geometry in this software is 2D axi-symmetric. Although the area of the 2D geometry can be calculated quickly, it cannot be directly linked to the volume/mass of the part; therefore, it is vital to calculate the volume of the geometry which can then be used to calculate mass. To calculate the volume of the RI geometry a 2D to 3D volume calculator was used (see Appendix I).

The performance of the RI geometry generator will be assessed by comparing the output geometry verses the minimum potential volume possible. These results will then be assessed against the number of optimisation evaluations to obtain these results.

Visual examination of the RI geometry was used to assess suitability of the outputted geometry. Design of Experiments (DOE) were performed on the RI geometry generator's parameters using several case study discs as a form of black box testing.

3.2.2.1. RI Geometry Creation Techniques

Two methods were initially considered to generate the initial RI geometry; these were:

- Pure knowledge based method – By using a set of rules, for a given disc geometry a repeatable RI geometry will be produced.
- Exhaustive minimal-rule-based method – This method uses a small set of rules to define the no-go areas around the disc, and then uses a technique to generate RI geometry using a set of input parameters to allow the creation of unlimited designs which should also be suitable.

A knowledge based method was developed by Rolls-Royce to achieve a similar task. The software could successfully generate good baseline RI geometry to initiate an iterative design process, but the range of designs were severely limited. The software utilises ultrasonic inspection knowledge (or rules) to help define the new RI geometry. Although this method intelligently uses the rules, it does not take into account the actual inspection results therefore does not guarantee a right first time result or any confidence of being close to a fully inspected solution; nor does the software actively seek to minimise the mass of the geometry.

Directly using some inspection rules which were used to improve the likelihood of right first time inspection design could actually prevent a solution from being found in some instances.

For example, the knowledge based approach will limit the minimum length of all RI geometry edges, as it knows that if it is too small then the edge will not be ultrasonically scanned. This is a very logical approach to creating the RI geometry, but there are several occasions when a small line would be extremely helpful in linking two adjacent edges which are both positioned in an ideal position as illustrated in Figure 22. In this case edges A and C are angled to be most effective and edge B joins the two edges together, although no ultrasonic inspection is applied to edge B it enables the optimal design. If the small edge was omitted (as illustrated in Figure 22B) then the two edges A and C must be joined only allowing one or the other to be positioned optimally. In this case, areas which are highlighted Red have been lost and may not be able to be replaced by scans from different edges (depending on the geometry).

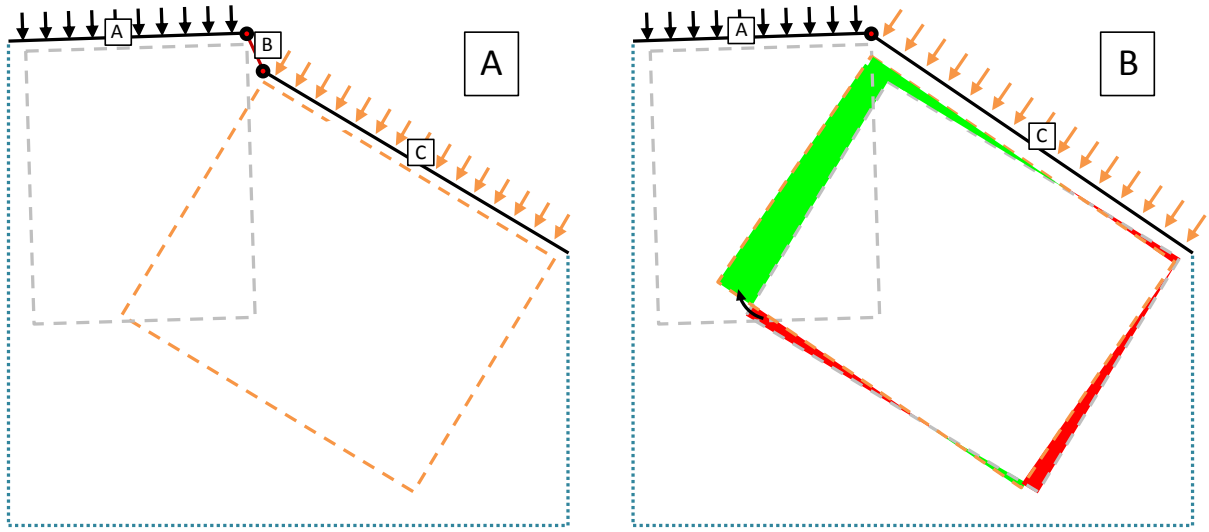

Figure 22 - Example geometry modification when removing small edge. Green is additional scans created and red is where the scan has been removed. The blue dotted edges could take any shape in this example.

It was clear that using a knowledge based approach would limit the design space exploration and also one of the key objectives set out is to maximise the design space exploration. Creating a similar pure knowledge based approach from experience will likely hinder the discovery of the optimum design so therefore will not be pursued.

An exhaustive method was first considered to allow full exploration of the design space. This method used parameters which directly controlled the coordinates of the RI geometry nodes. It was quickly realised that the vast majority of time spent exploring the design space either produced infeasible designs or highly inefficient geometry. The clear way forward from this was to maintain control over the RI geometry node points and to embed sufficient

knowledge such that generated designs would always produce feasible results. Embedding the knowledge ensured the design space was still sufficiently explored and meant all designs were sensible from a manufacturing and inspection point of view.

It was for this reason that the decision was taken to create a RI geometry generator, which has great freedom but was controlled sufficiently to allow RI geometry designs to be optimised efficiently. This decision led the RI geometry generator down another rule based method but using the rules to ensure manufacturability and design efficiency as opposed to setting the design.

3.2.2.2. Global Exploration Method

It was identified earlier that simplicity in RI geometry usually yields simpler effective inspection results which are straight forward to solve. On the other hand, complex RI geometry will result in a lower mass, but results in a highly complex inspection problem which are difficult to solve. It is this trade-off which is difficult to quantify. Simplicity in the number of parameters is also key to the success of this software. The fewer the parameters the faster an optimised solution can be found.

The global exploration of the RI geometry design space method requires a 2D axi-symmetric disc profile and manufacturing rules so that it can output an RI geometry for a given set of parameters. The development of this tool went through several iterations of maturity. The first iteration consisted of a simple stack of parameterised boxes by directly controlling the radial coordinate of each node (see Figure 23).

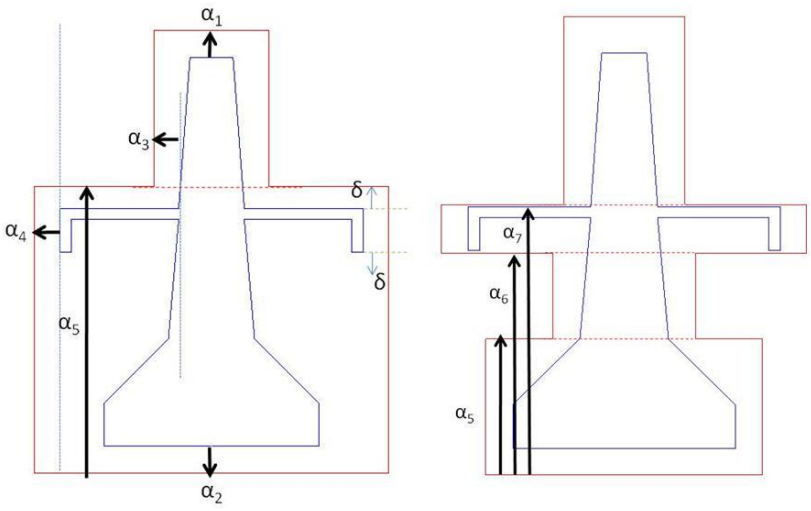

Figure 23 - Early development of the RI geometry parameter setup

This basic setup had minimal parameters but turned out to be unsatisfactory from an inspection point of view, primarily due to the number of sharp right angled concave corners. The second iteration introduced improved flexibility in RI geometry by removing the horizontal restriction on the stacked boxes allowing the shape to flex more like an independent polygon. This improved both mass and ultrasound coverage in the finished part with only a small increase in parameters. The third iteration added knowledge to the geometry generator to improve both manufacturability and inspection compatibility. The parameter control over the node coordinate system was also altered to improve the performance within the optimisation runs. It is this third iteration that will be described in detail.

There are three steps to the selected method. The first step finds the minimum RI geometry possible, the second step creates the RI geometry using the parameter values, and the third step post-processes the geometry to ensure manufacturability.

Step 1 – Minimum Possible RI Geometry

The first step in this method is to identify the smallest possible geometry that the RI geometry can occupy.

The minimum RI geometry is defined using the following set of rules:

- No undercut geometry to be allowed – Due to the difficulty in ultrasonic scanning in hard-to-reach areas, these areas are filled in to improve the penetration of scans from other edges. This is defined as there is to be no material removal from areas which cannot be seen from the horizontal directions to the centreline of the disc, i.e. left hand side to the disc centreline and the right hand side to the centreline of the disc. This is illustrated in Figure 24a as the area highlighted in orange.
- Disc Perimeter Offset Required – To avoid ultrasonic NDE inspection dead zones near to the surface of the RI geometry, an offset should be applied to ensure the surface of the part is always within the inspection limits. The area highlighted in yellow on Figure 24b illustrates this offset.
- Additional Manufacturing Requirements – In some cases, additional material is required around the finished part profile to ensure manufacturing process requirements are met. I.e. The addition of rotational friction weld stock material as illustrated on Figure 24c highlighted in blue.

The flow diagram shown in Figure 24 displays the order in which these steps are performed. The disc is broken down into two sides, the Right Hand Side (RHS) and the Left Hand Side

(LHS). Splitting the disc in this way helps simplify the problem and can be done as the LHS geometry is independent of the RHS.

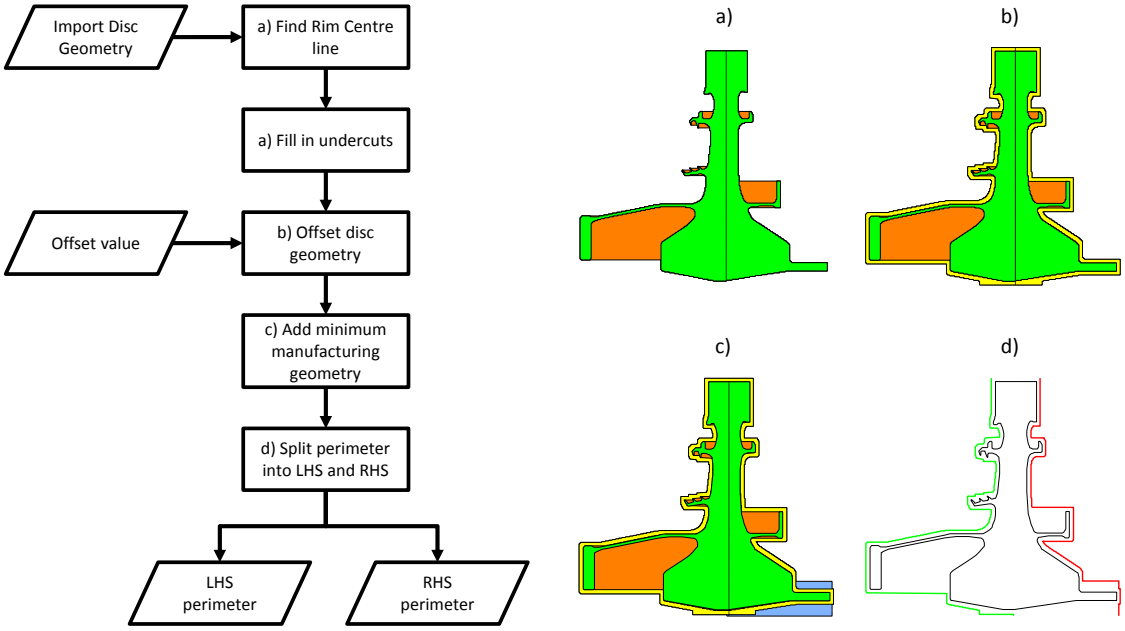

Figure 24 - Flow diagram with example showing how the RI geometry perimeter is created starting from the original disc geometry

The output of the LHS and RHS perimeters are two of the core inputs to the next step in this method.

Step 2 – Converting Input Parameters into RI Geometry

Step 2 of this method focuses on converting parameter values into feasible RI geometry. It is this part of the method which will run for every parameter design.

One of the first parameters created controls the maximum number of edges the RI geometry will have. With the exception of the Inner Diameter (ID) and Outer Diameter (OD) edges which will always be horizontal, the number of parameters to control the remaining nodes will be the maximum number of edges (N) minus four (the fixed OD and ID positions).

To increase the flexibility of the RI geometry the ratio of edges split evenly between the LHS and RHS, will be controlled by a single parameter. This parameter sets the percentage of nodes on each side.

To minimise the complexity of parameterisation, each node will only have one unique parameter to set each node position around the disc. There will be one additional global parameter which will control an offset distance to the minimum possible RI geometry. This

parameterisation simplification roughly halves the required number of parameters for the optimiser to solve, whilst still giving the optimiser control to solve inspection problems.

All parameters will range between the values of 0 and 1. The perimeters can be easily converted to 1D lines as shown in Figure 25a and then scaled such that the parameter points can be translated onto the appropriate LHS or RHS minimum RI geometry perimeters as shown in Figure 25b.

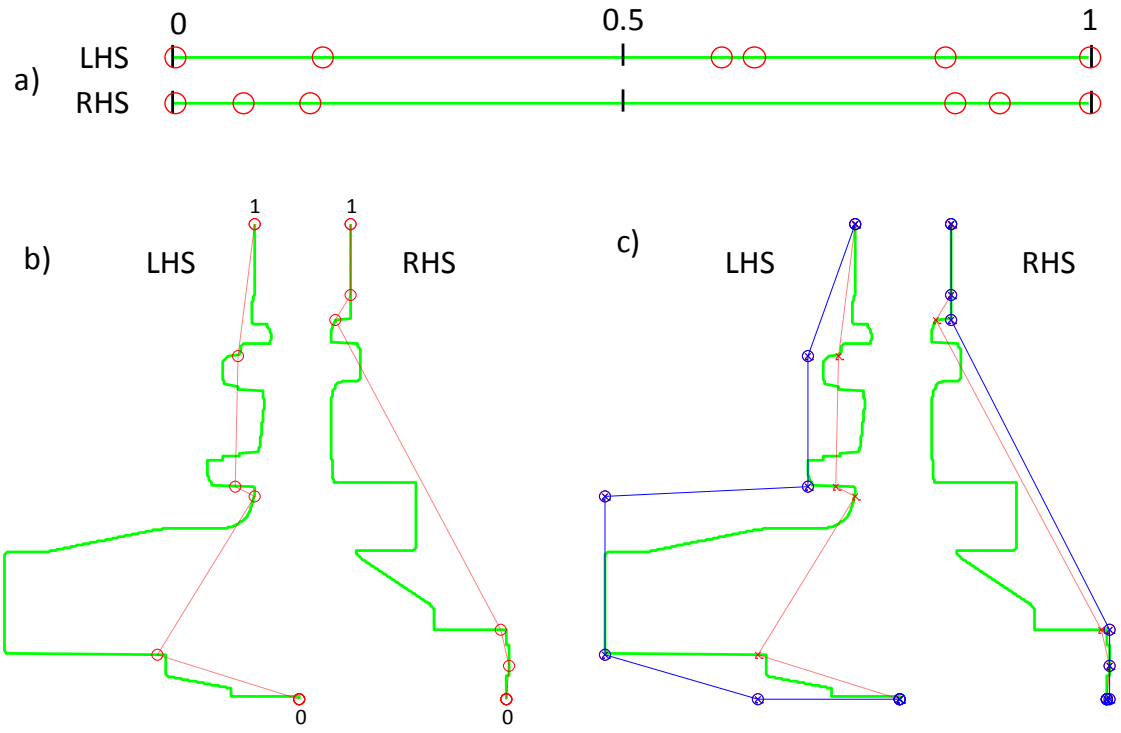

Figure 25 - Diagram showing how the model converts the parameter values into the RI geometry.

Joining these corner points together (as illustrated by the red dashed line in Figure 25b) produces a result which cuts through the minimum RI geometry. To convert this result into acceptable RI geometry a new set of rules needs to be applied. The workflow shown in Figure 26 below converts the results of Figure 25b into a realistic design as shown in Figure 25c. The following additional law applies to the workflow: All intersect points on the perimeter boundary can only move horizontally outward from the centreline:

- I.e. The outwards direction for the LHS is left and the inwards direction is to the right.
- I.e. The outwards direction for the RHS is right and the inwards direction is to the left.

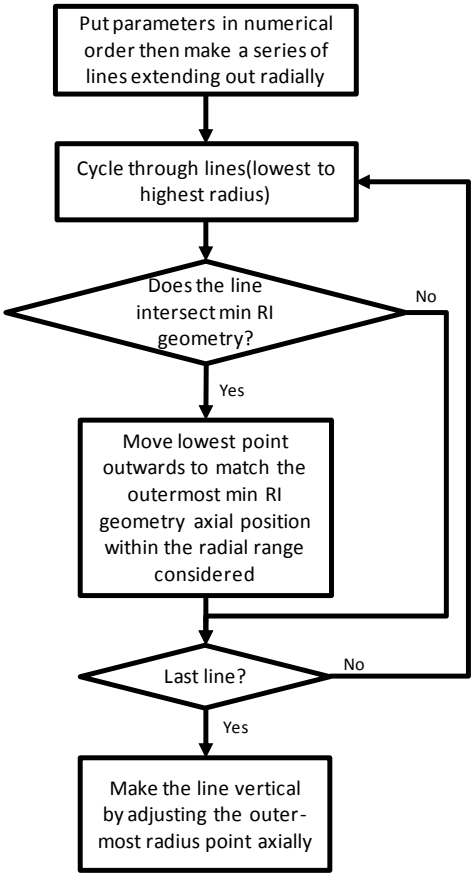

Figure 26 – Flow diagram illustrating sequence required to convert the raw parameters into acceptable geometric coordinates.

To limit the number of parameters required to create the RI geometry, it was decided to only have one parameter which controlled additional offset to the perimeters outputted from stage 1. In later sections where ultrasonic NDE constraints are difficult to satisfy, some flexibility in the overall RI geometry offset is required to ensure maximum inspection is achieved. By limiting the offset to one parameter which controls the offset around the entire perimeter simplifies the optimisation problem, which is deemed more beneficial at this stage than minimising mass. Further mass reductions can be achieved using the local improvements described in Section 3.2.2.4.

Step 3 – Post Processing RI Geometry

After experimenting with the output of Step 2 results during the development of the RI geometry generator; it was decided that more knowledge was required to improve the inspectability of the RI geometry. This involved simplifying the geometry in certain cases where it was deemed that the removal of nodes would have very little impact on mass but should improve the inspection of the component.

To control whether a node was removed from the RI geometry, the angle between every adjacent line was calculated. If this angle was below the specified angle, then the node was removed. After the removal of nodes, the geometry is not checked for intersections (although this could be easily done). Intersections are not checked at this point as the cut-off-angles are very shallow (less than 2 degrees). It was deemed that if a convex portion of the RI geometry was removed, firstly the offset from the disc would be able to absorb a large proportion of the material lost. Also if the geometry becomes too close to the disc geometry then the later inspection analysis results will be impacted which the optimiser would consider within its algorithm.

Nodes are also removed if they are repeated coordinates as this does not impact either the inspection or the mass of the component.

Although the Ultrasonic NDE software developed later in Section 3.3 is designed to take into account the external RI geometry to evaluate inspection limitations. It is favourable from a manufacturing point of view to avoid acute angles on the surface; as any potential mass savings are likely to be removed due to tool access restrictions during turning, as well as forging limitations.

The required actions performed during the post processing of the RI geometry are listed here:

- Manufacturing requirements (Illustrated by the red lines in Figure 27b):
  - **Horizontal bore –** Used for clamping during manufacture.
  - **Horizontal rim –** Used for clamping during manufacture.
  - **Vertical clamping faces at the rim –** Used for clamping during manufacture.
  - **Minimum clamping length at the rim face** – To ensure fixture clamps have enough land to securely hold the part.
  - **Vertical faces on the far LHS and far RHS –** Used for clamping, measurement calibration and to avoid damage when resting on a surface.
  - **Limit on acute angles –** (This is covered in the NDE requirement below) acute angles can cause potential manufacturing issues depending on angle and access.
- Ultrasonic NDE requirements:
  - **Parallel vertical edges at the rim** – (This is covered by the manufacturing requirements) Used to calibrate background attenuation noise.
- Ultrasonic NDE simplification (illustrated by the red cross and circle in Figure 27b):

- **No acute angles –** Acute angles are difficult for the ultrasonic probe to access. This simplification removes geometry which can severely reduce inspection due to large portions of the surface being un-inspectable.
- **Remove shallow obtuse angles –** A marginal angle (e.g. two degrees off a concentric line) will cause a discontinuity in the scan removing ultrasonic inspection over a significant portion of a scanned edge. In this case, it is better to add the marginal mass to remove as many discontinuities as possible.

An example of a few of these additional rules being applied is shown in Figure 27 below.

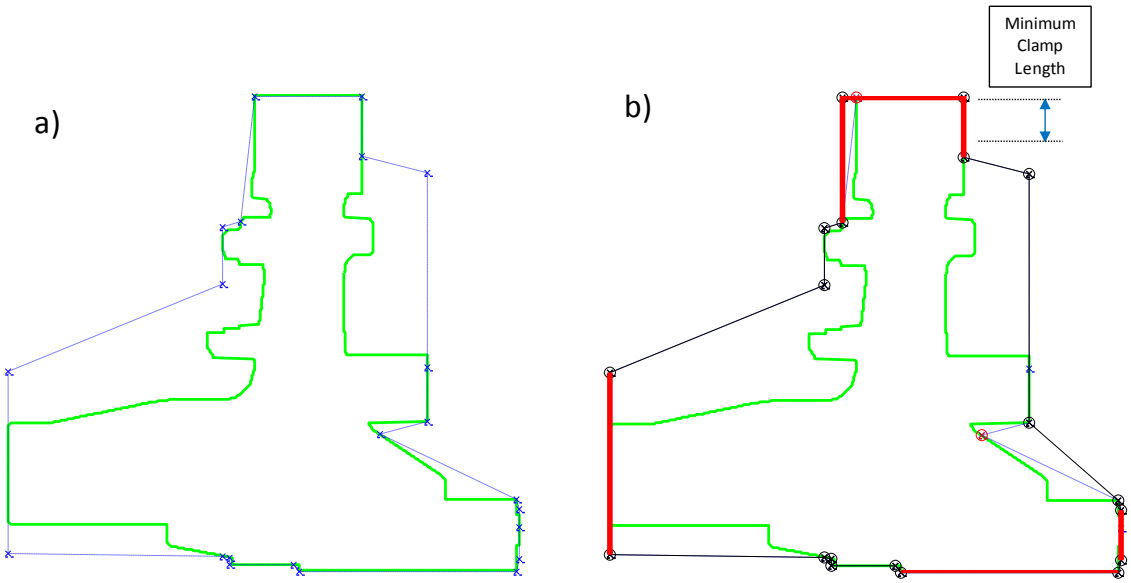

Figure 27 - Post processing example. a) Blue dashed line shows the output geometry after Stage 2 of RI geometry creation process. B) black/red lines show the post processed RI geometry after additional rules are satisfied.

The three steps described above when run consecutively resulted in manufacturable RI geometry being successfully created 100% of the time. Further exploration of the design space is shown in Figure 28 below. Running both a Latin Hypercube Design of Experiment (DOE) and a Genetic Algorithm (GA) optimisation method shows how many possible RI geometry designs there are and how effectively the design space is covered.

DOE Design Results	Optimisation Design Results
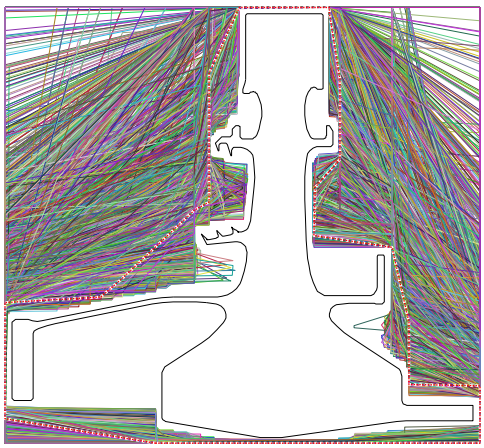	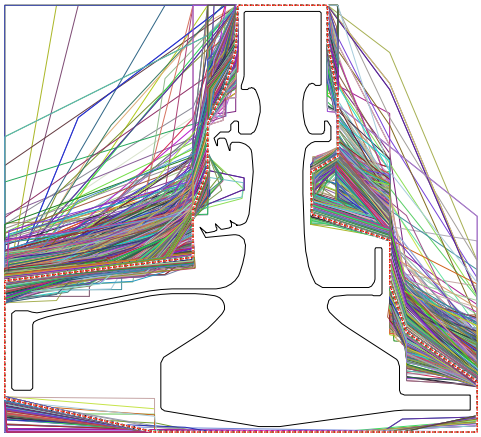

Figure 28 - Displayed RI geometry designs resulting from performing a latin hypercube over the parameters (Left) and and performing the genetic algorithm on the parameters (Right).

The Latin Hypercube DOE results shown in Figure 28 illustrate how well the parameterisation is setup covering a large portion of the design space. The GA optimisation results show how quickly suitable parameters are selected, quickly narrowing down the design space where more efficient designs lie.

Additional Detailed Parameter Setup

Three independent parameter setups were considered to understand the most effective technique to aid the genetic algorithm on converging.

It was decided that the RI geometry generator should be designed to cope with input parameters independent of each other. This would later improve the optimisation of the component as constraint functions slow down the optimisation process.

Type 1 Parameterisation

The first method was to use the input parameters as cumulative percentage cuts for example: [0.5, 0.5, 0.5] would create splits at [0.5, 0.75, 0.875]. As each consecutive cut will have a decreasing amount of space to divide.

Type 2 Parameterisation

The second idea was to convert the input parameters into ratios. i.e. [0.3, 0.2, 0.4, 0.1] would create splits at [0.3, 0.5, 0.9]. Each parameter has an equal weighting therefore the cuts will split equally. The down side is that you need one more parameter (per side) to allow this flexibility.

Type 3 Parameterisation

The third kept all parameters independent of each other i.e. [0.4, 0.7, 0.3, 0.9] would split at [0.3, 0.4, 0.7, 0.9] in order.

To test the methods which are explored above an optimisation was set up to investigate the benefits of each method discussed. The optimisation work flow is shown in Figure 29 below.

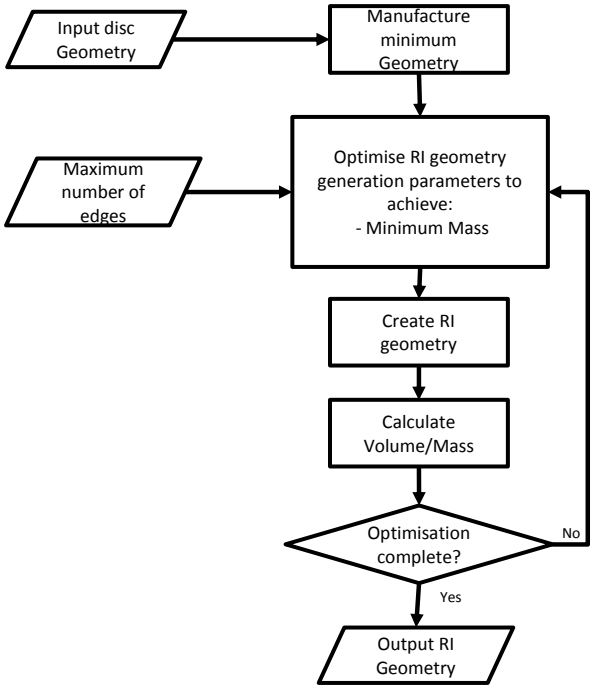

Figure 29 – Section 3.2 Optimisation workflow

3.2.2.3. Parameterisation method - Down Selection

In order to down select the parameterisation method, the optimisation algorithm first needs to be set up to ensure efficient and consistent results are produced within a sensible time frame.

The global optimisation algorithm 'genetic algorithm' (GA) has several input parameters to tune the algorithm. The key input parameters are listed here:

- Population size – The number of individual design analysis which will be performed within each generation.
- Generation size – The maximum number of generations which the optimisation run will perform.
- Elite Count – The number of individuals in the current generation that are guaranteed to survive to the next generation.

- Crossover Fraction – The fraction of the next generations' population that will be created by the crossover function.
- Migration Fraction – The fraction of individuals in each subpopulation that will migrate to a different subpopulation.
- Migration Interval – The number of generations that take place between migrations of individuals between subpopulations.

All of these GA input parameters will have an effect on the efficiency of the optimisation to some extent. Due to time limitations not all of the input parameters to the GA can be tuned to this specific problem, therefore they will be set to the default values supplied by MATLAB® with the exception of two, the population and the maximum number of generations. These two parameters could severely affect the repeatability of the optimisation if set too low, but if set too high, the optimisation may take an infeasible amount of time to complete.

All of the following studies were performed using an old RI geometry generator at the level of step 2 (described above); it therefore lacked the design efficiency improvements which were applied in step 3. The results therefore do not represent the current status of the model but were used as an initial guide to setup the GA input parameters as the trends are expected to be similar.

To understand the effect, the 'population' and 'generations' have on the effectiveness of the GA optimisation an experiment was set up. The experiment varied the population value from 50 to 1000, but a fixed value of 200 was used for the parameter 'generations'. The total number of evaluations is set by the initial population multiplied by the number of generations. This means that the larger the parameter 'generations' the longer the optimisation will run unless the optimisation 'stalls'. Therefore, by using a large value for 'generations' it will be possible to find out how many generations are required to converge to its optimum result. Every population optimisation run was repeated 50 times to enable the results to be compared fairly.

This investigation would ideally be performed for every problem investigated in the later sections of this chapter, but due to time and computing resource limitations this will not be possible. Analysis performed at this stage in this section provides the only opportunity where analysis time is feasibly short enough to investigate these parameters within the time frame of this thesis. As the RI geometry generation method will remain the same for the

remainder of this chapter, the result of this investigation will be used to set all minimum population sizes as well as the upper limit of the number of generations.

Figure 30 compares the average result of the 50 repeat runs to gauge how the population size relates to the minimum mass of the part using Type 3 parameterisation. The 'excess material' is the additional mass required above the minimum possible RI geometry (described previously).

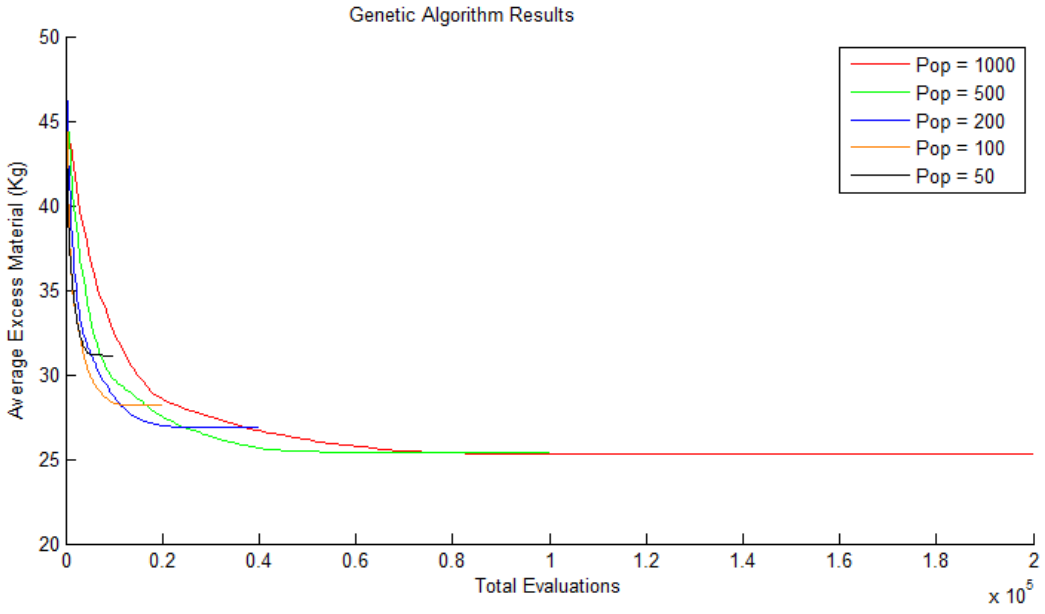

Figure 30 - Average results for 200 Generations (repeated 50 times)

The results clearly show that improvements trail off at around 100 generations (half of the total evaluations) for population sizes 100 and above. However, in order to get a substantial benefit from the optimisation, the plot shows that a population of 500 produces a result very similar to that of 1000, but is achieved after almost half the total number of evaluations.

To compare the different parameterisation types a study was performed to directly compare their performance. Using a population of 1000, the three parameterisation types have been run 50 times to get a direct comparison of performance; the results are shown in Figure 31 below.

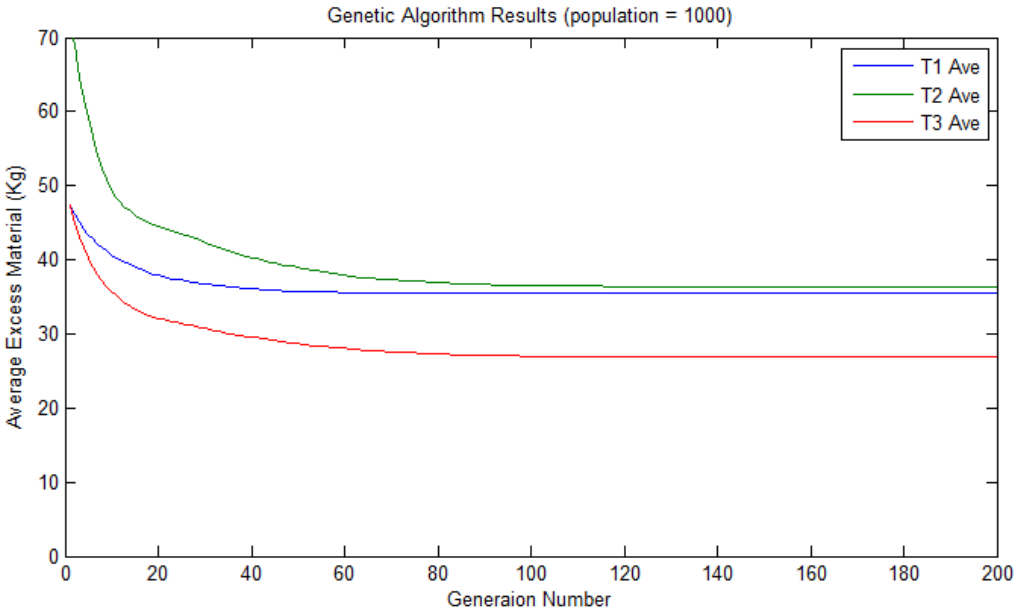

Figure 31 - Results produced from averaging 50 independent runs of the genetic algorithm optimisation for each generation.

The results in Figure 31 show the clear benefit of using Type 3 parameterisation where on average a 25% reduction in waste efficiency is seen, which amounts to approximately 9 kg in material savings.

To check the consistency of each parameterisation, type the range of results for each generation is shown in Figure 32 below. The range of results is found by subtracting the lowest mass value from the highest selected out of the 50 repeat results for each generation.

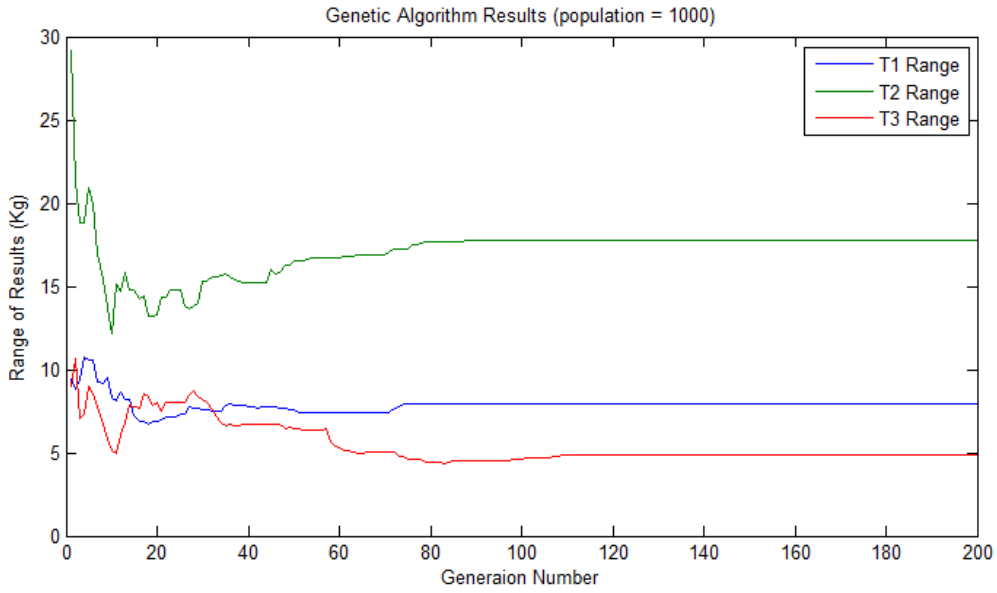

Figure 32 - Range of results at each generation from 50 repeated GA runs.

The range shown in Figure 32 was greater than expected showing ±2.5 kg for Type 3 and as much as 8 ¾ kg for Type 2.

Despite greater than expected range of the results, Type 3 was chosen due to the relatively high perfomance and higher consistancy shown from the repeated optimsation runs. Both a population of 500 and 1000 produce simlar superiour results. As a population of 500 yields a greater optimisation efficiency this would be favourable parameter value when optimisation time is of concern.

3.2.2.4. Local Improvement Method

It was appreciated that the initial parameters of the RI geometry generator may not be set up perfectly for all given problems. For example, to avoid having an additional offset parameter for every node, a universal parameter was used which would ensure that ultrasonic inspection near the surface of the disc would have a reduced risk of not receiving inspection from perpendicular scans[17]. This improves the optimiser's chance of finding a design which should be fully inspectable, but comes at a cost of universally adding mass to the design around the surface.

To ensure the outputted RI geometry from the global optimisation is optimal, a further local optimisation is performed on this geometry. This will remove all redundant excess material produced during the initial global optimisation, whilst maintaining or improving the weight of the disc. This method will become even more relevant later in Section 3.3 where it is used to maintain or improve the inspection of the disc and RI geometry mass.

To maintain the manufacture and inspection requirements of the RI design, all the RI geometry's handling, clamping and measuring surfaces will be automatically identified and parameterised. This will ensure all six essential edges are constrained to being horizontal and vertical as discussed previously. By having these additional constraint limits the amount the design can perturb therefore reduces the number of required parameters by 6 which should improve optimisation performance.

The method used to perturb the RI geometry during this local optimisation is shown in Figure 33 below.

[17] The 3-4-5 rule described in APPENDIX II describes how an edge may need to be up to 5 mm from the part to ensure attenuation from the wall does not interfere with detecting a defect.

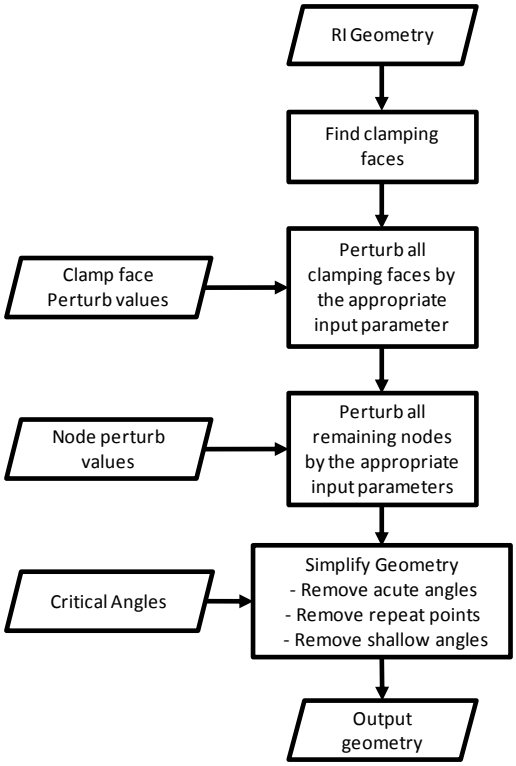

Figure 33 – Flow diagram showing how the RI geometry improvement tool works.

The main aim of the RI geometry improvement tool was to help reduce mass whilst still being able to achieve the goal/constraint of achieving 100% ultrasonic NDE inspection. This is discussed in greater detail in Section 3.3.

3.2.2.5. Optimisation

The full optimisation loop shown in Figure 34 below will become the backbone for all future optimisation in the ensuing sections. This optimisation loop covers both a global exploration of the initial design space and then a second 'local' optimisation to further improve the design through small adjustments of the first design. As this optimisation only has one objective to minimise mass, then there will be only one design for which to perform local optimisation. In later sections there may be multiple designs outputted from the global optimisation which can all be run in parallel through a local optimisation loop.

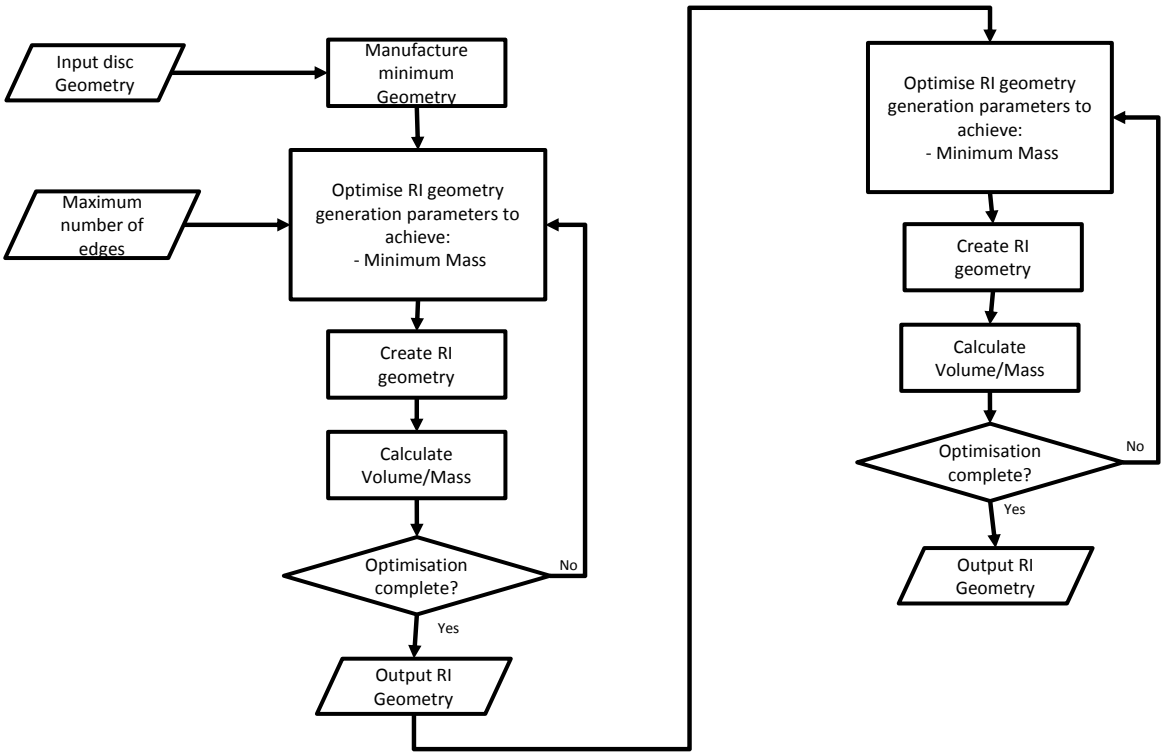

Figure 34 - Multi stage optimisation loop

The optimisation was performed on the 4 case study discs which were described in Section 3.1. For each of these case studies the optimisation was repeated 5 times using an identical setup with the exception of the random seed which creates the initial population of parameters for investigation within the first generation of optimisation. Repeat experiments were performed to improve the likelihood of finding near optimal designs. The number of repeat experiments was set to 5 due to a combination of time and computing constraints.

To ensure the minimum manufacturing requirements were enforced, a penalty function was introduced called the 'Objective Volume', it is this objective which the optimiser works to minimise:

$$\text{Objective Volume} = \text{RI geometry Volume} \times \left(1 + \frac{\text{Manu Penalty Volume}}{\text{RI Geometry Volume}} + \text{Clamp Penalty}\right)$$

Where the manufacturing clamp length penalty 'Clamp Penalty' is calculated using the following equation:

$$\text{Clamp Penalty} = |\text{LHS Length - Req Length}| + |\text{RHS Length - Req Length}|$$

And the manufacture penalty volume 'Manu Penalty Volume' is calculated using the following equation:

$$\text{Manu Penalty Volume} = 1 \times 10^4 \times \left(f_{\text{Volume}}(\text{Min RI Geom}) - f_{\text{Volume}}(\text{Min RI Geom} \cap \text{Output RI Geom}) \right)$$

Where 'Min RI Geom' is the minimum possible manufacturing geometry allowed (as described earlier in Figure 24); 'Output RI Geom' is the newly generated RI geometry and 'f_{Volume}' is a function which calculates the volume of an axi-symmetric 2D geometry. If manufacture penalty volume variable becomes less than zero; the variable is made to equal zero. The weighting was of the manufacturing penalty was increased by a factor of 4 to ensure any designs with manufacturing volume issues stood out clearly to be easily identified by the optimiser as a bad design.

The 'Objective Volume' is calculated in this way to factor in objective penalties which are required to handle the constraints. The clamp penalty will multiply the RI geometry volume by the magnitude of the total clamp length missing (measured in millimetres) ensuring that the optimiser will always ensure that the clamp lengths are sufficient. The manufacturing penalty volume will ensure that the RI geometry will always be greater than the minimum manufacturing geometry.

3.2.3. RESULTS

Table 5 shows the best results for each of the 7 different optimisation runs. It was found that increasing the number of edges from 22 to 32 showed a noticeable reduction in RI geometry mass of between 3 and 6 percent. The RI geometry has been optimised to only take into account manufacturing requirements but has shown mass reductions of up to 19.8% compared to the current designs.

Table 5 – Optimisation results showing the optimised mass (normalised to existing RI geometry) for the given disc and parameters.

	Result Source	No. of Edges	RI Geom $Mass_{norm}$
CS-1	Existing RI Geometry Design*	-	100.00*
	RI Geometry Optimisation	22 Edges	92.29
	RI Geometry Optimisation	32 Edges	86.84
CS-2	Existing RI Geometry Design	-	100.00
	RI Geometry Optimisation	22 Edges	87.31
	RI Geometry Optimisation	32 Edges	83.09
CS-3	Existing RI Geometry Design	-	100.00
	RI Geometry Optimisation	22 Edges	82.31
	RI Geometry Optimisation	32 Edges	79.96
CS-4	Existing RI Geometry Design	-	N/A
	RI Geometry Optimisation	22 Edges	83.86

* The CS-1's existing RI geometry mass is based only on the final stage RI geometry design, the first stage is 36.2% heavier

Figure 35 shows the RI geometry for both 22 and 32 edges side-by-side for comparison on the CS-2 and CS-3 HP turbine discs. The results for 22 edges are similar in simplicity to current existing RI geometry, whereas the 32 edge geometry manages to achieve greater weight reduction in hard to reach areas such as in the rear C-groove. The CS-2 RI geometry achieves a better weight reduction from 22 to 32 edges due to the position of the middle seal which doesn't appear on the CS-3 disc.

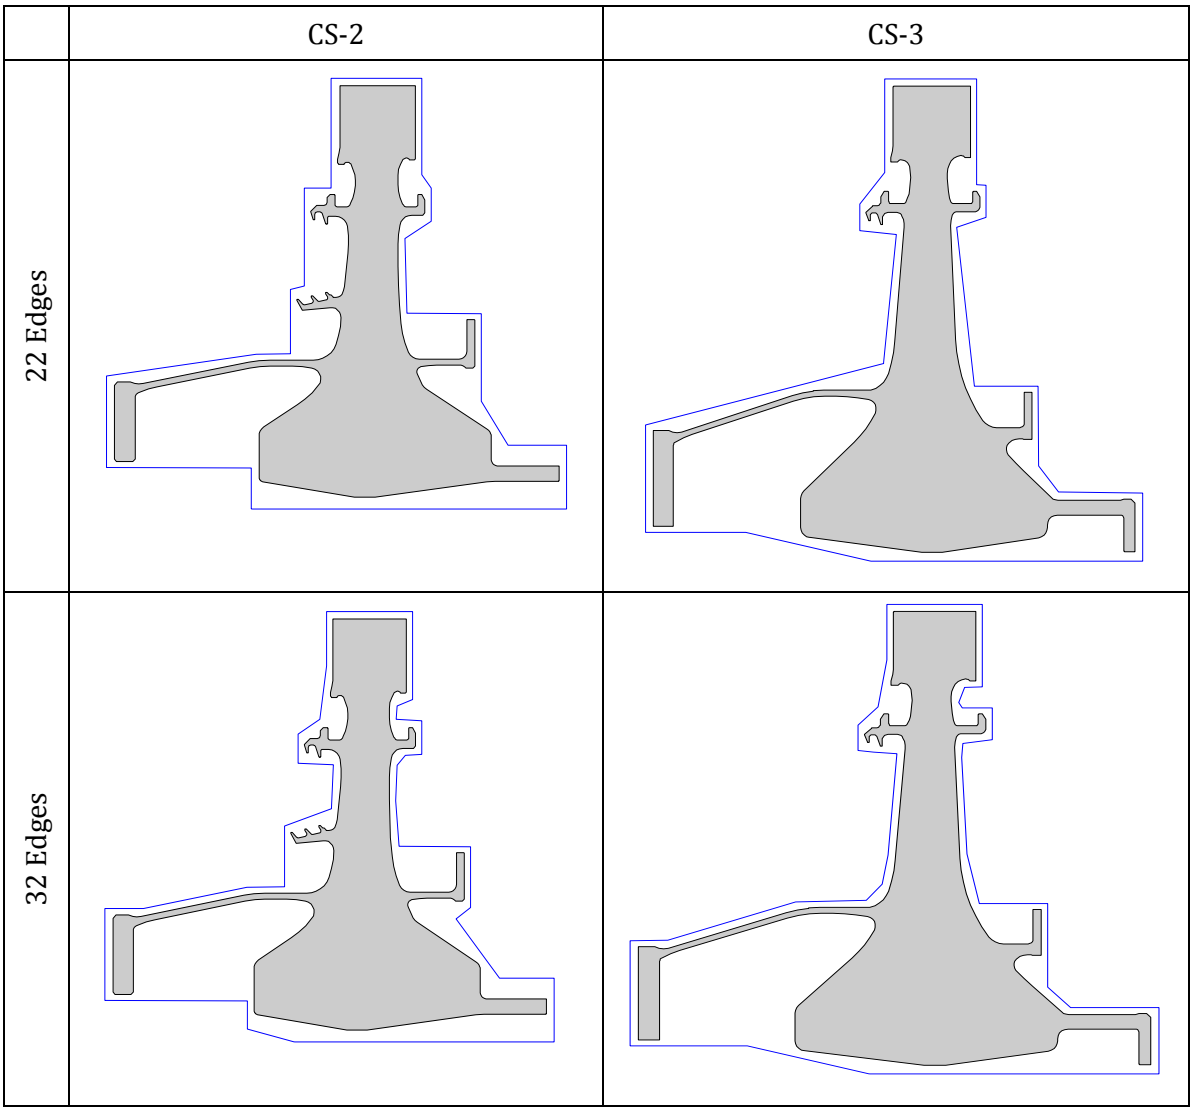

Figure 35 – Outputted RI geometry designs for CS-2 and CS-3.

Neither disc has material removed neither under the drive-arm nor behind the rear flanges which adds considerable volume to the RI geometry. Removing this material was made out of scope in the method section due to the upstream manufacturing considerations, as these areas will not add material costs during the forging process (covered in greater detail later in Section 3.5).

3.2.4. DISCUSSION

This section set out to create a method to produce RI geometry in such a way that the coverage of the design space was maximised and could be explored fully using global optimisation methods.

This section lays the foundation for the remainder of the chapter defining the method of RI geometry generation, which will have a direct impact on the efficiency of the optimisation and subsequent results.

Summary of Findings

It has been shown that the method of generating the RI geometry works and explores the design space sufficiently. The method described can create geometry which is able to surpass current designs when not considering ultrasonic requirements.

The results showed material savings of as much as 20% compared to current pre-existing designs. This may be partially due to inefficient existing designs, but is more likely to come from the requirement of ultrasonic NDE limitation constraints. The effect of the ultrasonic NDE requirements will be covered extensively in Section 3.3 later in this chapter.

It was found that RI geometries could be created with the potential of obtaining improved near-net-shapes. Increasing the number of edges led to a greater reduction in RI geometry mass. Comparing the optimised results against the actual mass is an unfair comparison considering that the actual RI geometries have also been through an iterative loop using the ultrasonic NDE results as a constraint to the designs. Visually looking at the geometry created it is possible to find areas which have been missed, but such narrow areas as seen in Figure 35 are unlikely to provide any cost benefits when considering the entire manufacturing process.

The second objective of this section was to create a standard template for which optimisation will be performed. This hybrid optimisation method not only analyses tens of thousands of varying designs, it also picks out the best designs for further local improvements to ensure that the best design can be found.

Parameter setup was shown to be a very important factor in allowing the optimiser to find not only the best solution, but also the most efficient route to it. Three different methods of parameterisation using the same basic concept caused drastic changes in the effectiveness of the optimisation results which was unexpected.

The optimisation parameters for the Genetic Algorithm showed that the population parameter should be 500 or greater, and that the maximum number of generations should be at least 100 to enable an efficient design to be found. Although experiments with a population parameter set at 1000 showed slightly better results, the required doubling of function evaluations would result in the optimisation taking twice as long.

Academic Implications

The method created in this section has shown that RI geometry no longer has to be defined parametrically in relation to a disc, but can be automatically generated through the setting of a limited number of parameters. This new method is considerably more complex, but enables a substantially larger portion of the design space to be explored improving potential gains during optimisation.

Business Implications

The existence of this automated RI geometry generator not only provides the Manufacturing Engineers with a tool which can be used to produce initial RI geometry, but it also provides a tool for cost modelling engineers to use. The tool requires minimal input from the user as it works by applying pre-defined rules to the input disc geometry which allows the RI geometry to be optimised using the parameters defined.

From a factory perspective the less material removal that needs to be done 'in house', the cheaper and faster the finishing operations will be. This automated process follows this behaviour well.

Limitations and Directions for Future Research

The number of iteration to perform the optimisation is in the order of 100,000 iterations. This results in a computation time of a few hours for the current analysis therefore when further analysis is performed within each iteration, run times will increase further.

The main limitations on this study is that the solutions to the optimisation problems found do not take into account the Ultrasonic limitations which are a key constraint, and should be a part of the RI geometry optimisation. Creating a suitable analysis tool is vital to the creation of the RI geometry which will in turn define the cost of the product. This limitation is addressed in Section 3.3 and covers the ultrasonic NDE constraint in great detail.

It is anticipated that the greater the number of edges the less sound energy will go into the RI geometry resulting in compromises to the inspectability of the RI geometry. For this reason, it is clear that optimising the RI geometry for just minimal mass is not an ideal objective. Ultrasonic NDE will play a large part in the suitability of the RI geometry; therefore, further ultrasonic NDE analysis which will be investigated in Section 3.3 would greatly improve the RI geometry design process.

This page intentionally left blank.

3.3. BATCH CAPABLE ULTRASONIC NDE SOFTWARE FOR USE IN OPTIMISATION

3.3.1. INTRODUCTION

Safety is paramount in the aerospace industry therefore all safety critical discs are manufactured using tightly controlled processes. One of these processes is ultrasonic inspection of the finished forged part. This analysis is carried out to detect potentially catastrophic defects within the finished part. In order to ensure these defects are detected the RI geometry is designed and analysed using a set of rules which was previously performed by hand. In this section a new automated ultrasonic NDE analysis tool will be described which is currently being patented by Rolls-Royce (Wiseall, Phipps et al. 2014).

Ultrasonic inspection is an important analysis for Rolls-Royce rotating critical parts. This is due to assumptions used to life the highly stressed components. Not detecting anomalies within the finished component could lead to tragic events similar to that seen in the past, most notable Sioux City (National Transportation Safety Board 1990). To prevent such events occurring, it is imperative that the finished part be fully ultrasonically inspected.

As manufacturing and cost optimisation takes a foothold in the part manufacturing process it is paramount that the ultrasonic examination stage is not overlooked. To maintain the high standards of inspection within Rolls-Royce it is important that the same if not higher constraints are applied to the inspectability of rectilinear inspection geometry and for this to happen, a manual process needs to be converted into an automated system.

The examples given in *Rohl et al. (Rohl, He et al. 1998)* and *Srivatsa*'s (Srivatsa 2003) paper showing disc optimisation show no details of their ultrasonic requirements and go into no further detail other than stating that there is a 'sonic shape'. Although some research has included the RI geometry within part of their workflow, there is no sign that any ultrasonic inspection analysis has been performed on the geometry. Within Rolls-Royce the ultrasonic inspection is an important mile stone in the buy-off process to issuing the RI geometry for production. It was shown in Section 3.2 that it is possible to design RI geometry very close to the final disc (near net shape), but without an ultrasonic NDE assessment being performed on this shape it may fail ultrasonic NDE requirements. Therefore, the step of performing an ultrasonic inspection assessment is a key part in defining the RI geometry. Without a valid RI geometry design any resultant cost or heat treatment shapes created early to assess residual stresses will have no practical use, as these designs are likely to be invalid.

To identify whether the design of the RI geometry can be sufficiently inspected the polygon geometry is analysed using a set of inspection rules. The series of rules used to analyse the inspection geometries have been created using experiments and allow a high level of confidence that defects of a particular size can be detected within the inspected region.

The ultrasonic NDE rules are relatively straight forward to apply so are frequently performed manually either by hand or electronically using CAD packages, but this comes at a cost as a single analysis can take several man hours to complete.

Rolls-Royce started developing automated ultrasonic inspection software in 2010 to be run within Siemens NX CAD software. At the time, this software proved temperamental, slow to compute and was non-batch[18] capable. The approximate run times for this software are as shown in Table 6 below.

Table 6 - Resolution Vs Time

Resolution (mm)	50	5	1
Approximate Time (s)	240 (4 minutes)	1800 (0.5 hours)	7200 (2 hours)

The requirement for the software was to have a 0.1 [mm] resolution, but even with a 1 mm resolution it would take in the order of 2 hours to compute. Such a long run time would make any RI geometry optimisation a time consuming process especially for the optimisation task outlined in Section 3.2 where up to 100,000 design iterations were run. No other software was found which could perform this bespoke task.

Although the series of rules are relatively straight forward, software to date has not shown to be efficient at processing the rules to identify the inspectable areas. The software described used an exhaustive discrete method to identify the inspectable regions which proved to be extremely inefficient. This exhaustive method could be crudely described as asking someone to survey an obstacle course blind using only a stick and a set of rules to locate keep out areas. Every step that is taken would mean re-evaluating the surroundings. This blind approach requires a lot of probing at intervals equivalent to the required resolution as illustrated in Figure 36 below.

[18] Unsuitable for use within an automated environment.

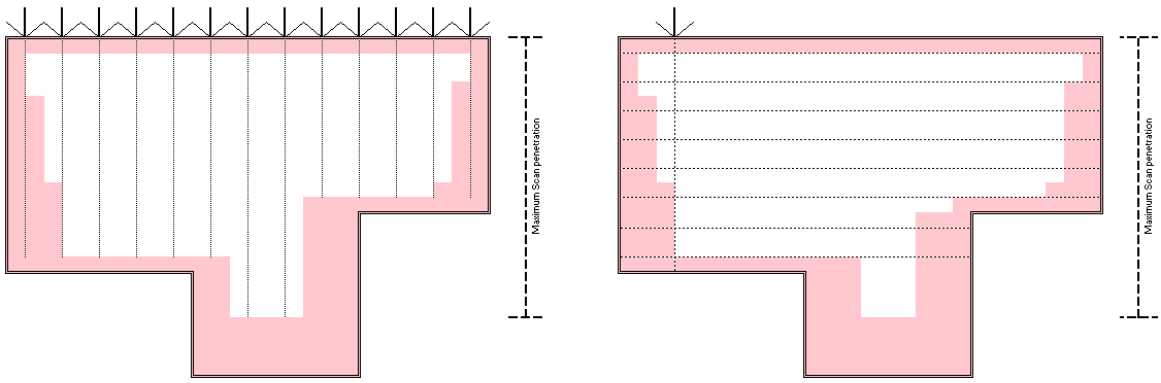

Figure 36 – Blind approach to solving the problem using discrete intervals requires several analyses being performed on a single edge, then each individual scan requires several more analyses along its whole length

Every time the resolution or object's perimeter is doubled, the number of probing analyses needs to increase four fold. Figure 37 shows how computationally expensive this method is revealing the number of nested loops that is processed in order to locate the inspectable regions within the part.

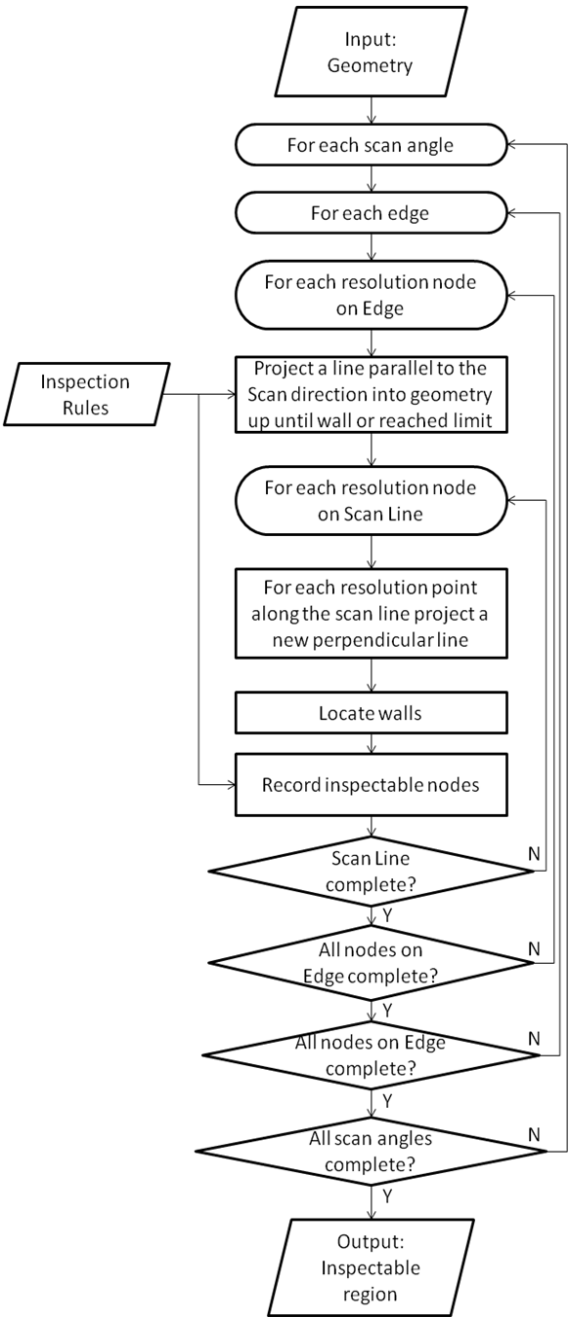

Figure 37 - Blind man and stick approach

Due to the number of nested loops involved in this discrete method the resolution required has a considerable impact on run time, this is why the Rolls-Royce developed software took a couple of hours to run at a 1mm resolution. By considering the performance trends shown in Table 6 a 0.1mm resolution run time could take as little as 14 hours, but in the worst case it could take as long as 200 hours. The calculated number of function evaluations performed within Figure 37's flow diagram is displayed in Table 7. It can be seen how the number of evaluations changes with resolution and the perimeter of a part.

Table 7 - Computation inefficiency described using resolution and perimeter

Required Loops		Resolution [mm]			
		5	2	1	0.1
Perimeter [mm]	1000	40,000	250,000	1,000,000	100,000,000
	800	25,600	160,000	640,000	64,000,000
	500	10,000	62,500	250,000	25,000,000
	300	3,600	22,500	90,000	9,000,000
	200	1,600	10,000	40,000	4,000,000
	100	400	2,500	10,000	1,000,000

Despite the software still being in development, it is clear that the software will not increase in speed by the orders of magnitude required for use in this study. Using the optimisation setup from Section 3.2 a total of 50,000 evaluations would ideally need to be performed. That would take 11.5 years to complete using the substandard 1 mm resolution using the current developed software within Siemens NX6; this clearly is not practical.

After solving the ultrasonic inspection problem, it is often the case that the RI geometry does not fulfil the inspection requirements. If this is the case, then the RI geometry must be suitably altered to correct the problem areas.

Adjusting the inspection geometry is an extremely complex task requiring great spatial awareness. Altering an edge not only has a primary effect on the scan passing through that edge, it also has a secondary impact on the adjacent edges and in many cases a tertiary impact on the inspection scan as illustrated in Figure 38.

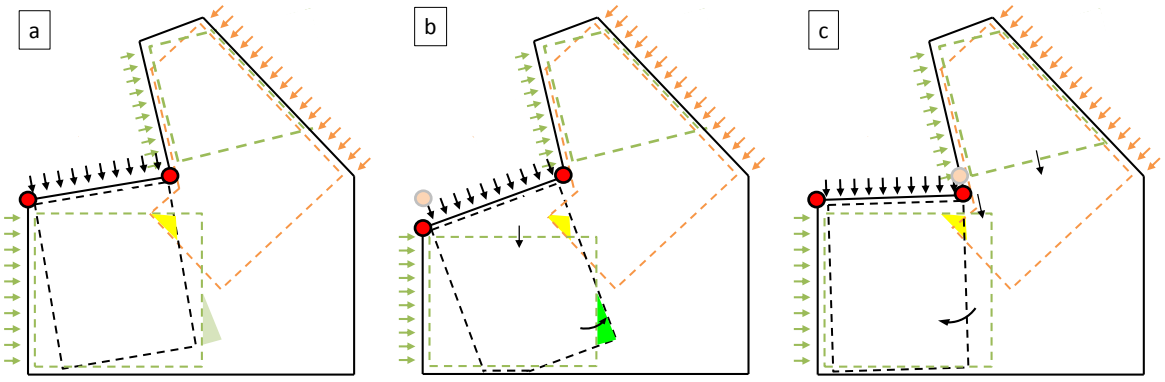

Figure 38 - Diagram showing the primary and secondary implications of changing a scan edge.

The black arrows and box show the primary edge and scan area. The green arrows and boxes are scans which are also directly affected by the primary edge (Figure 38b). The orange arrows and box show an example of a non-connected edge and how its scan can be influenced through secondary effects. This is why the problem can be quite complex.

Past research overlooked ultrasonic analysis as a key driver in the creation of RI geometry. Rolls-Royce have attempted to advance from manual to automated ultrasonic analysis, but have missed an opportunity to use this software within an RI geometry optimisation loop, resulting in maintaining the manual aspect of RI geometry creation.

In the literature review it was shown that work performed by *Srivatsa et al* (Srivatsa 2003) that the forging shape is dependent upon the RI geometry, but no work was performed to include this step within the optimisation. This might be due to one of the following reasons:

- Lack of appropriate software
- Lack of foresight into the benefits of including RI geometry optimisation.
- A lower focus on ultrasonic NDE

The objective of this section is to create a fast reliable method of analysing geometry for its ultrasonic inspectability through an automated system. Due to the high number of optimisation evaluations required to find a low RI geometry mass in Section 3.2 the software needs to be run in seconds rather than minutes to allow an optimiser to be run in a suitable time frame. The aim is to run this software within an optimisation loop in conjunction with the RI geometry generator created in Section 3.2.

To get one step closer to improving the design process through improving the RI geometry inspection analysis. The key research objectives are as follows:

- To develop a faster ultrasonic limitations analysis method suitable for batch use within an optimisation loop.
- To investigate the impact of the RI geometry design when considering the ultrasonic limitations in the design.
- To compare results between pre-existing RI geometry designs with new optimised designs for both mass/volume and percent fully-inspected.
- To answer the hypothesis “Is it possible to create suitable RI geometry without analysing the geometry for ultrasonic limitations?”

These tests were conducted on the initial example as well as various case study parts.

The ultrasonic limitation analysis method is a piece of software which will require RI geometry as well as the finished component geometry to inspect. The software should analyse the geometry for ultrasonic limitations then returns the results in an automated manner.

The method development contains four parts. First the pre-existing software is analysed to establish what methods need to be avoided. The second part verifies that the software creates true ultrasonic limitation results. Once the software is ready it is then inserted into the optimisation loop to optimise the RI geometry for both minimum mass and full inspectability. The final section analyses the results from the optimiser to conclude the hypothesis.

3.3.2. METHOD – THE DEVELOPMENT OF ULTRASONIC NON-DESTRUCTIVE EXAMINATION SOFTWARE

The core of this section is to create an automated method of taking in RI geometry then applying an ultrasonic NDE rule set; this will produce the ultrasonic limitation data required to assess the designs suitability. The high level method used to realise this is shown in Figure 39 below. It is this process which is contained within the developed software.

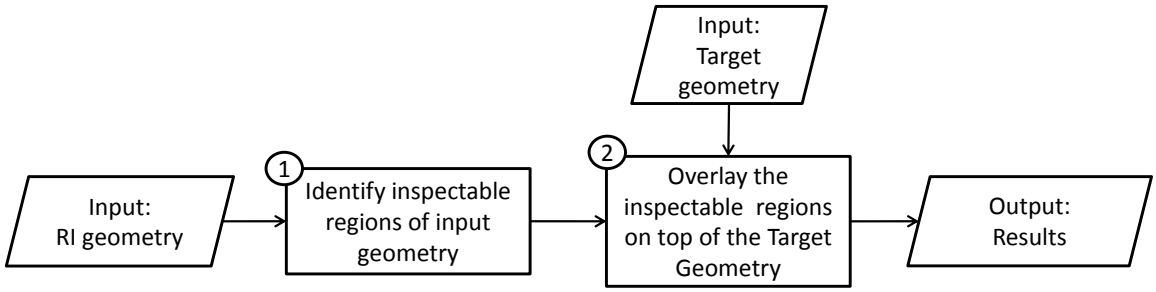

Figure 39 – Process to check how inspectable a finished geometry is for a given rectilinear inspection geometry.

The development version of this software was named DUNDE which stands for Dummy Ultrasonic Non-Destructive Examination. The word 'dummy' is used to convey the fact that rules are used to locate inspectable regions as opposed to performing physics based methods such as ray-tracing.

The software is made up of two key processes, the first identifies all inspectable regions contained within the RI geometry, and the second post processes these results combined with the target (i.e. finished disc) geometry to produce the inspection results.

The software has approximately two dozen rules which it runs to produce the same diagrams that were being achieved with the previous human method. These rules include:

- Near/dead zones
- Far zones
- Maximum penetration depth due to attenuation
- Wall effects
- Line of sight requirement
- Interference due to internal geometry
- External geometry influences on surface visibility

The speed of this software is proportional to the number of edges which makes up the rectilinear inspection geometry during the rule implementation stages and is independent of the resolution required. Typical times for a 20 edged geometry takes approximately 1-3 seconds per analysis.

3.3.2.1. DUNDE – Method to Identify Inspectable Regions of a Given Rectilinear Inspection Geometry

Using the same analogy as described in the introduction of this section, the approach taken is to ask a blind man to survey a large room but in this scenario the walls have been displaced to provide a definitive edge of the keep in region. Therefore, to identify the inspectable region, one only needs to follow the wall around the whole perimeter until you have returned back to the starting point. Using this method removes the 'blind man' from having to perform any rule checks during the survey and only needs to record the walk around the perimeter until he has returned back to where he started.

This simplified problem works because of one assumption and one novel method. The assumption is that there are no hollow voids within the RI geometry. As this will always be true for a forging, this is a valid assumption. The novel method used is to remove the need for the 'blind man' to perform any rule checks as he explores the RI geometry. The way this is done is to displace the original RI geometry and create new bounding geometry to exclude the non-inspected regions as described by the rules. In effect you are using the ultrasonic NDE rules in reverse. Therefore, instead of blindly locating keep out zones using the original RI geometry and ultrasonic NDE rules, the new geometry created reflects the inspection limits (as shown in Figure 40). Applying the ultrasonic NDE rules to the geometry cuts the number of calculations down by several orders of magnitude. Using this method removes

the ‘blind man’ from having to perform any rule checks during the survey and only needs to record his walk around the perimeter until he is back where he started.

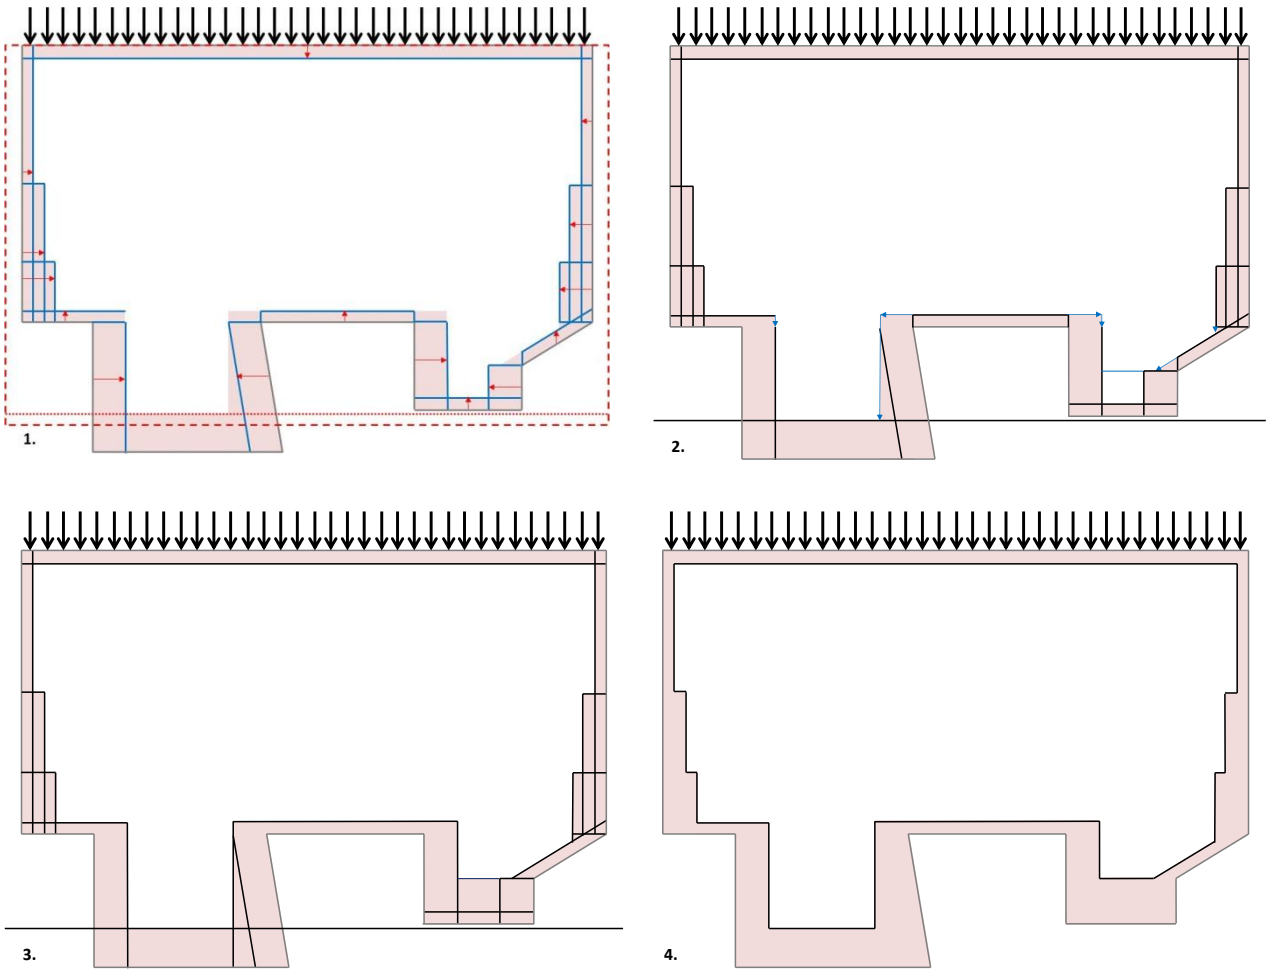

Figure 40 – Creation of the limitation geometry. Grey Lines – Original Geometry. Black arrows – Scan direction. Pink shaded areas – non-inspectable region. Red dashed lines – Region that may interact with the scan. Red dotted line – Maximum scan penetration depth. Red Arrows – Displacing edge geometry. Solid blue/black lines – New scan boundaries.

This new method will produce results of infinite resolution as there is no discretisation of the RI geometry required. This is because only the inspectable region’s perimeter coordinates are required as opposed to discreet inspectable points contained within the perimeter. This means that the higher the required resolution, the more efficient this method is compared to the original blind man with a stick approach.

The high level method to identifying the inspectable regions of RI geometry is shown in Figure 41.

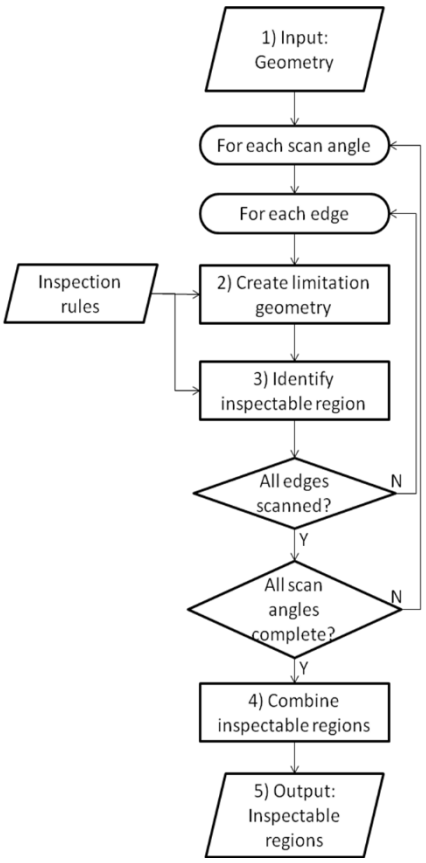

Figure 41 - New method to locate the inspectable region of a given rectilinear geometry

The details of this flowchart are explained here:

1. Read in the rectilinear inspection geometry
2. Using the NDE rules, new 'limitation geometry' is created through displacements of the original geometry. Additional geometry is created from the new limitation geometry to satisfy further NDE rules. Figure 42 expands on this process:
   2.1. The rules provide information to identify the maximum distance geometry can interfere with a scan, removing all unnecessary geometry is used to speed up the process (Figure 40.1 Dashed red box).
   2.2. Displace all of the remaining edges to satisfy the rules as shown by the red arrows in Figure 40.1.
   2.3. Extend the displaced lines as appropriate (interpreting the rules) until all edges are connected as illustrated in Figure 40.2
3. Identifying the inspectable region can be done in two ways:
   a. A rule based method – Uses the two known starting points of the scan as starting positions for an algorithm to run. The algorithm follows the perimeter of the inspectable region until a closed loop is created defining the inspectable region.

b. Applying Inequalities to the limitation geometry – by turning the 'limitation geometry' into bounded or grouped inequalities, all of the necessary nodes which define the inspectable region can be identified and combined to define the inspectable region. Figure 43 expands on this process:
    b.1. Split all lines where intersected (Figure 40.3).
    b.2. Remove all lines which do not define the inspectable region (Figure 40.4).
    b.3. Remove any repeated lines.
4. Combine the inspectable regions from each scan to create either a matrix or polygons of varying degrees of inspection within the rectilinear inspection geometry.
5. Output the results of the inspection analysis.

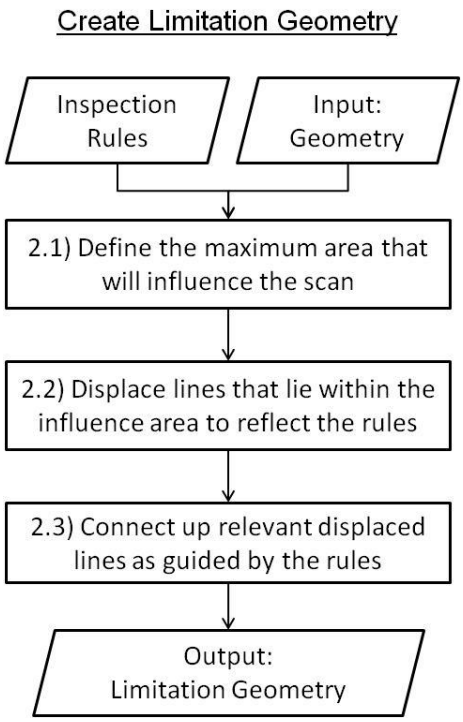

Figure 42 - Creation of the limitation geometry

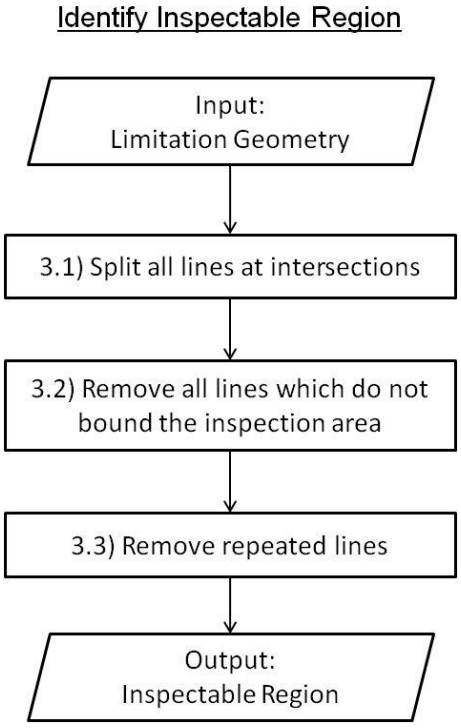

Figure 43 - Identify the inspectable region

3.3.2.2. DUNDE - Post Processing the Ultrasonic Scan Regions to Locate Areas of Ultrasonic Limitation.

Two methods were considered to post process the results from Step 03 (of Figure 41), the first was simple and robust but produces a discrete array of results. The second method was more challenging to create a robust technique but provided a precise result which would be crucial for gradient based optimisation later.

Method 1 was initially the preferred method as it was easy to setup and was very robust. This 'discrete' method required a matrix the size of the RI geometry with a mesh controlled by the required resolution. Every node within the matrix is then interrogated to determine if it lay within each scan. Adding each matrix from each scan produced the total number of received scans for each node. By comparing this final matrix with the disc geometry it was then possible to determine how inspectable the geometry was. Figure 44 below shows the workflow used to determine the inspectability of a part.

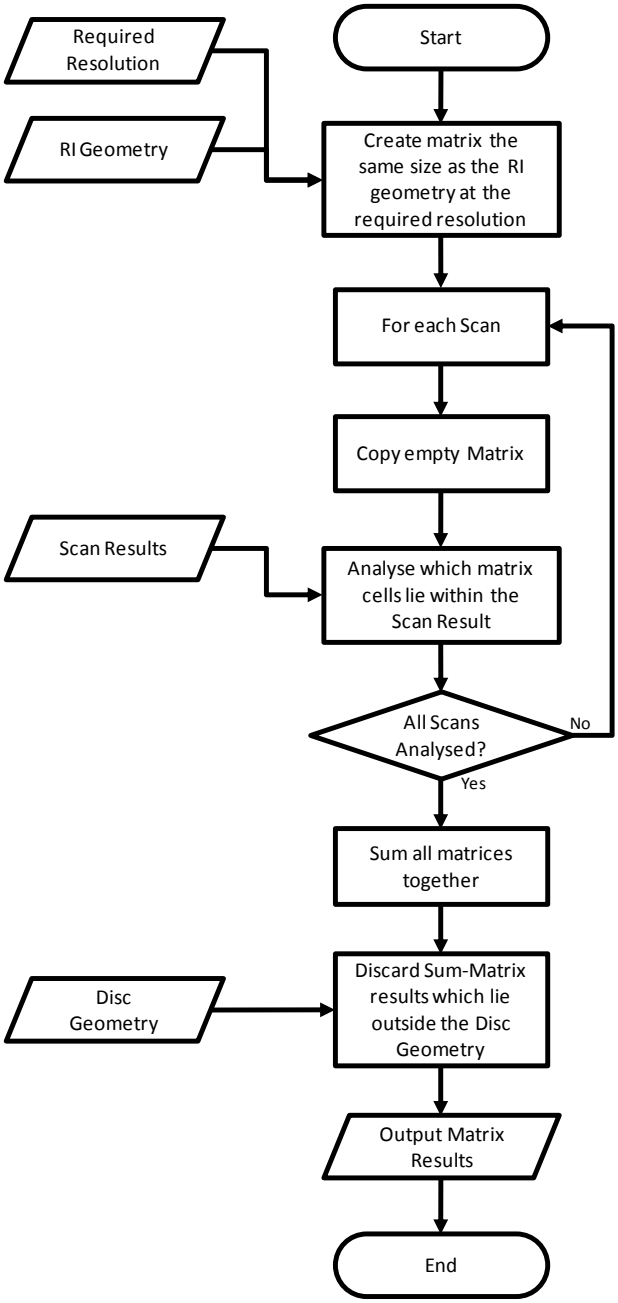

Figure 44 - Flow Diagram showing how the inspectability is calculated discretely

This method was very easy to set up but the time to run this post-processing analysis was proportional to the inverse of the resolution squared. Figure 45 below shows how the matrix for each scan is summed to provide the cumulative result.

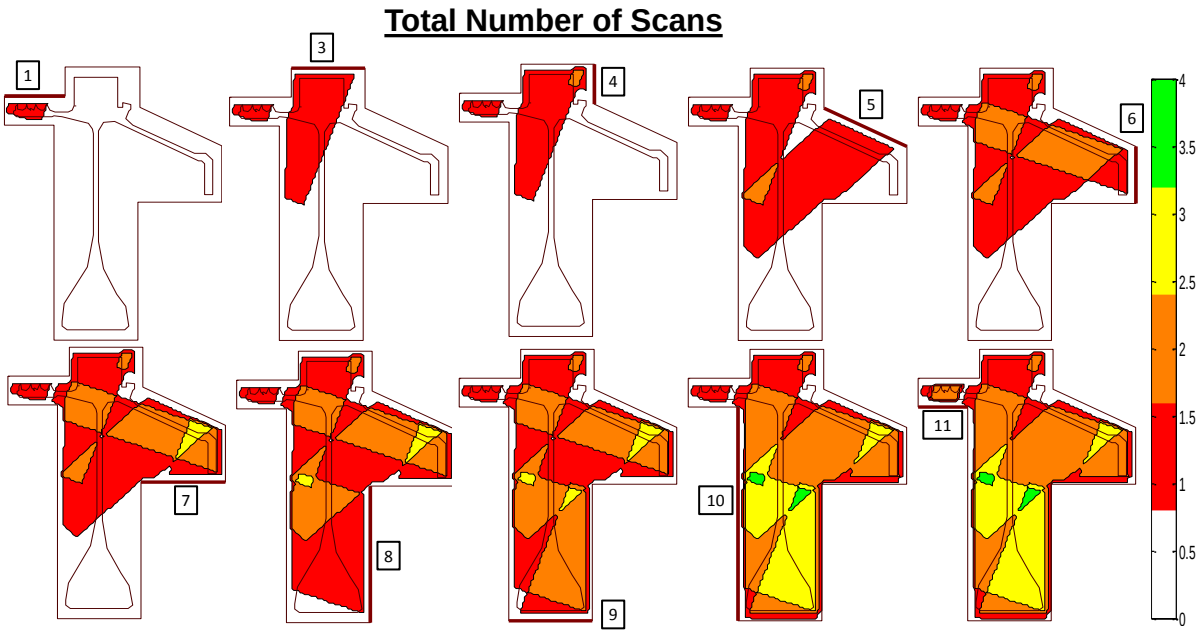

Figure 45 – Step by step process of adding each edge's inspectable polygon region to obtain the total number of scans for a +20 degree scan angle. The total number of scans (Note, edges 2 and 12 are below the minimum length to be scanned therefore have been omitted)

For the purposes of the rules used at Rolls-Royce the number of scans is capped at four as this is the minimum number of scans required to pass the inspection criteria. By removing all the values outside the disc geometry the relevant points for the disc's inspection can then be found. Using the results of Figure 45 plus additional scans from -20 degrees and 0 degrees, the relevant results are displayed in Figure 46. The fully inspectable region can be defined as a percentage as shown in the table on the left hand side of Figure 46.

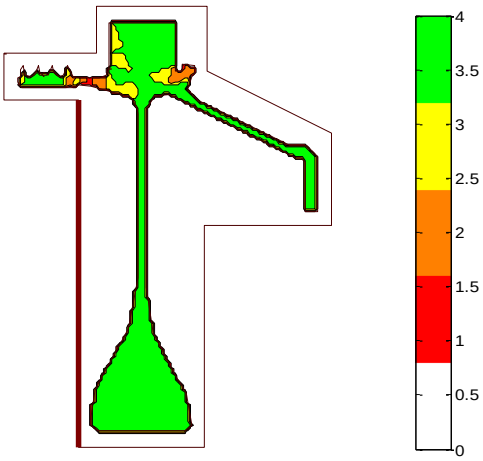

Scans	Percentage Covered
4+	91.0 %
3	06.0%
2	02.7 %
1	00.3 %
0	00.0 %

Figure 46 - Total number of scans only within the disc region (capped at four scans)

The current compute times are between one to five seconds for the majority of discs (at a resolution of 1 millimetre). Generally, a third of the time is spent finding the inspectable regions; the remaining two thirds is spent converting each inspectable region into a 2D matrix of discreet points. The method works well for low resolution arrays taking approximately 1-2 seconds for a 1 mm resolution, but the specified resolution required to satisfy the rules is 0.1 mm. If the arrays are run at this resolution, then the run times would be 100 times longer taking minutes rather than seconds. This is solely due to the fact that the compute time to calculate these matrices is directly proportional to one over the square of the resolution required. For an initial exploration of the design space using the Genetic Algorithm a low resolution could be used but to run a gradient based algorithm would require a much finer mesh to provide the necessary fidelity required to avoid a zero gradient surface by the optimiser.

With potentially a very large design space, a run time of over a minute would be unacceptable. The resolution could be reduced to improve the run time, but a lower resolution on components with high precisions is unacceptable.

The performance of the ultrasonic NDE software was being severely hindered by the method used to sum the inspectable regions within the disc. Like the original ultrasonic NDE method created, it was exhaustive, performing mostly unnecessary calculations.

3.3.2.3. Post Improving the Post Processing of the Scanned Regions

To improve the performance of this software a new technique was looked into which no longer required discretisation of the RI geometry area to find the total number of scans. The method is described next.

Method 2 was investigated further to improve the performance of the post-processing. By converting from a discrete method to a continuous one will have three benefits:

- Infinite resolution
- Low memory usage
- Consistent run times

The new method considered was based on geometric Boolean algebra. By comparing each scan with the previous then locating the 'union' and 'intersection' region allows you to build up the total scanned regions as shown in Figure 47 below.

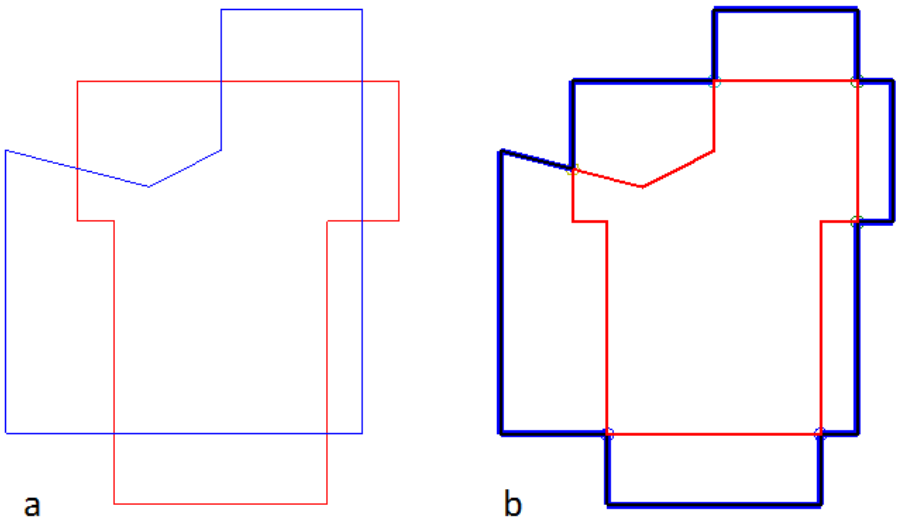

Figure 47 – a) The two original polygons to be summed. b) The resulting polygons; Blue region contains at least one scan and red region contains two.

By automating these Boolean operators in a loop, all the scans can be summed to create a library of closed polygons describing the number of scans it represents. This flow diagram is shown in Figure 48 below.

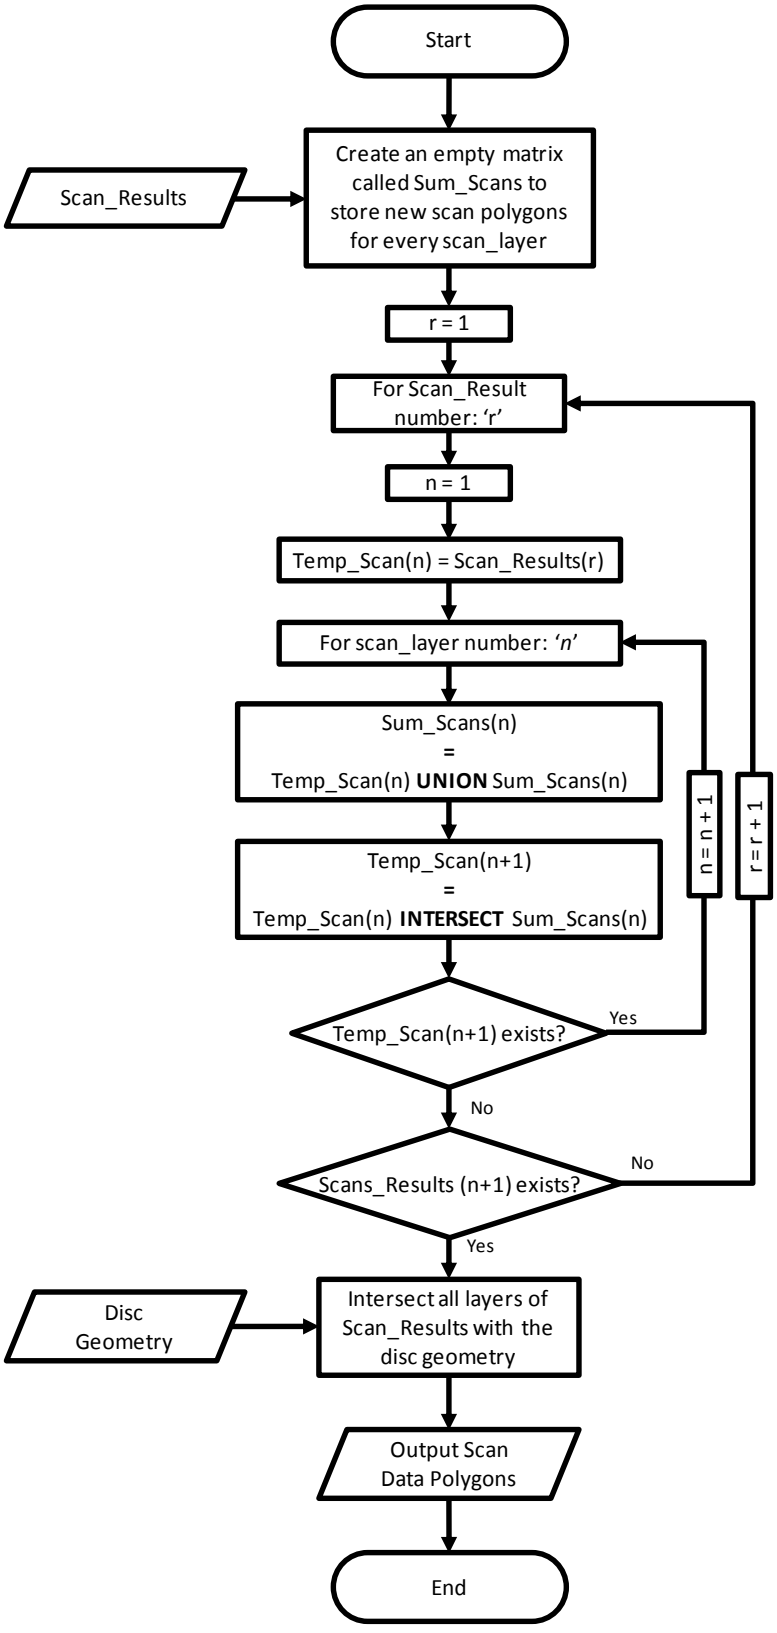

Figure 48 - Flow diagram showing the method used to accumulate scan results using Boolean operators. Variables 'Scan_Results', 'Temp_Scan' and 'Sum_Scans' are matrices containing polygon coordinates for each scan layer in the corresponding 3rd dimension of the matrix.

By uniting each scan, you can create a polygon showing an area which contains at least one scan, to find the area which contains at least 2 scans you need to check for polygon intersections. By following the flow diagram above, a complete map of scan coverage can be created as shown in Figure 49 below.

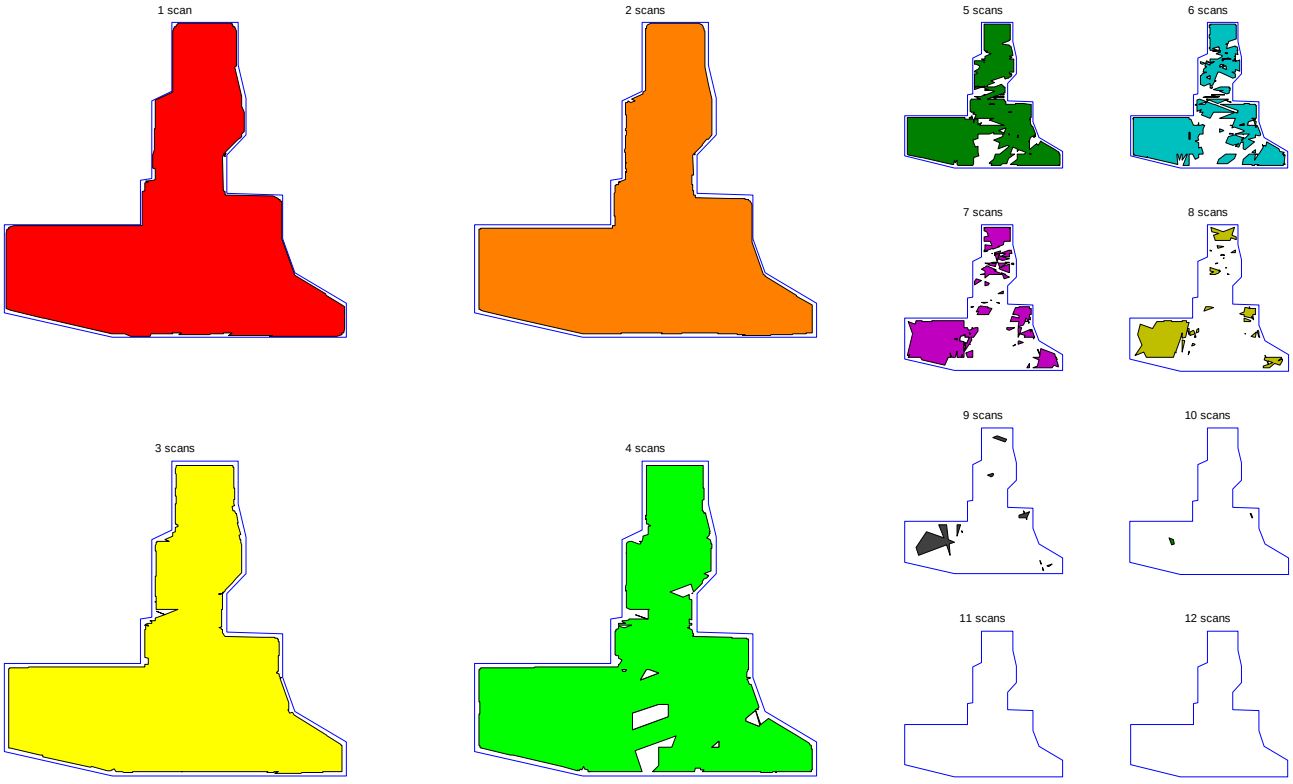

Figure 49 – Visualisation of at least the number of scans received (i.e. greater than 1 scan received, greater than 2 scans received, etc…)

3.3.3. RESULTS – SECTION 3.2 RE-ANALYSED

Now that we have the capability to perform ultrasonic inspection on RI geometry, we can now look back at Section 3.2 to see how the optimised RI Geometry designs have performed from an ultrasonic inspection point of view.

Figure 50 and Figure 51 below show the inspection results performed on CS-1, CS-3 and CS-2 discs retrospectively.

	CS-1					CS-2					CS-3				
32 Edge Limit	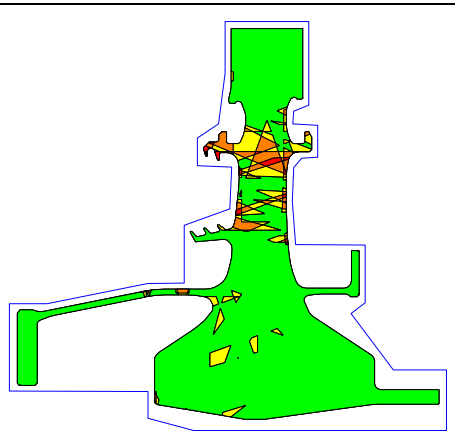					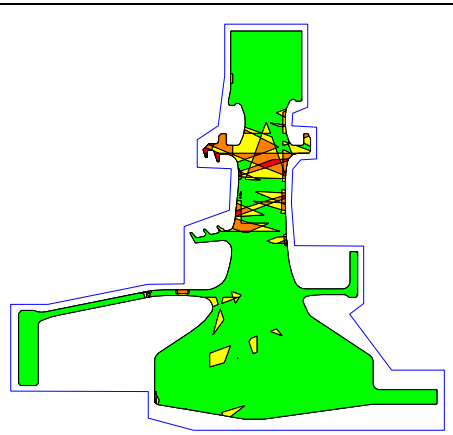					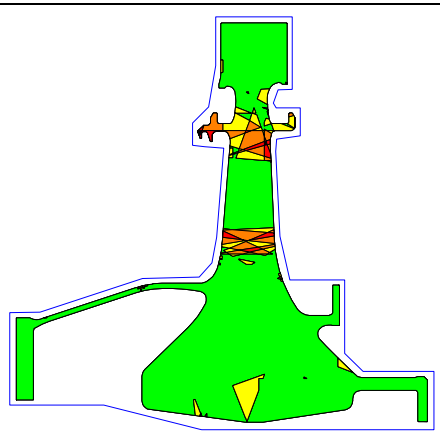				
	4+	3	2	1	0	4+	3	2	1	0	4+	3	2	1	0
	85.29 %	8.50 %	4.24 %	1.46 %	0.51 %	88.95 %	6.75 %	3.67 %	0.60 %	0.02 %	88.66 %	6.10 %	4.17 %	1.02 %	0.05 %
22 Edge Limit	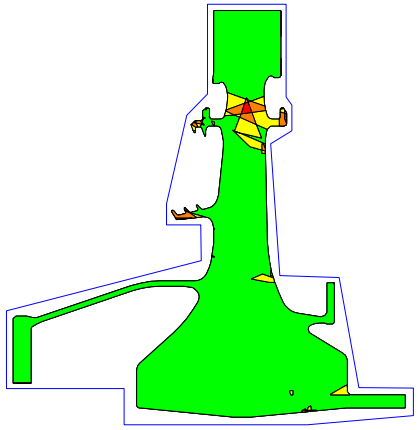					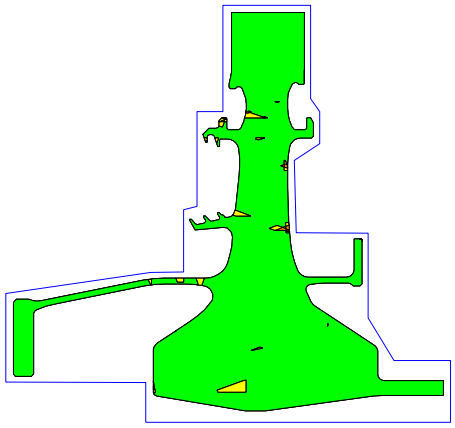					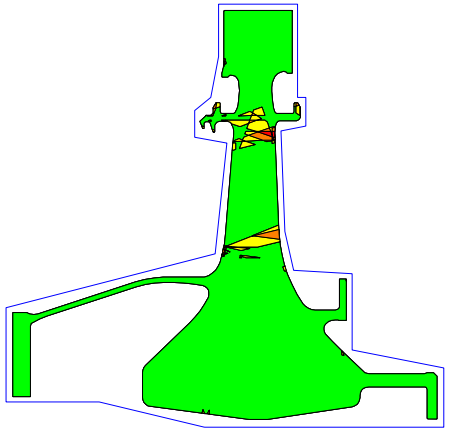				
	4+	3	2	1	0	4+	3	2	1	0	4+	3	2	1	0
	94.78 %	3.51 %	1.45 %	0.026 %	0.00 %	98.56 %	1.32 %	0.12 %	0.01 %	0.00 %	96.10 %	2.96 %	0.79 %	0.14 %	0.00 %

Figure 50 - CS-3 HPT Disc – Results of the ultrasonic analysis when performed on Section 3.2 designs.

The results in Figure 50 show clear evidence that the greater the number of edges the worse the inspection results despite a reduction in RI geometry mass. The reason for the significant reduction in inspection is down to a corner offset rule. The rules state that the scan cannot commence until it is 3 mm away from a convex corner and 5 mm from a concave corner (see Appendix II for greater detail). Therefore, the greater the number of corners on the RI geometry, the greater the inspection area will be lost.

None of the results from Section 3.2 meet the target of 100% full inspection using the methodology of minimising RI geometry mass whilst ignoring the inspection criteria. The 32 edge limit designs missed full inspection by at least 11% whereas the 22 edge limit designs all missed full inspection by less than 5.25%. None of the results have a fully inspected 'cob' nor at least 3 scans on the 'diaphragm' which would fail more detailed requirements dictated by materials and lifing experts in the field (see Figure 2 in Section 1.2 for feature locations). From these results a 22 edge limit looks favourable for use in future investigations.

3.3.4. ULTRASONIC NDE WITHIN THE OPTIMISATION LOOP

The ultrasonic NDE requirement would ideally be set as a constraint within an optimisation algorithm. For this type of problem applying a 100% inspection constraint within the current optimisation is likely to yield zero results. The reason for this is due to the initial starting parameters which are unlikely to yield fully inspectable designs first time. As none of the designs would meet the constraints, none of the results can be used to guide the optimisation, this is likely cause the optimisation to fail or severely underperform. To solve this issue, the 100% inspection constraint was converted into one of the optimiser's objectives. This multi objective optimisation would be setup to both maximise the inspection and minimise RI geometry mass.

Building on the methods created in Section 3.2 we now are able to edit the optimisation work flow as described previously in Figure 34 (in Section 3.2) to allow us to include inspectability of the part.

The genetic algorithm was used to explore the design space to locate good designs. Optimisation methods cope better with multi objectives rather than constraints which are difficult to achieve. As the constraint function (ultrasonic NDE) is a very difficult constraint to satisfy, the problem has become a multi-objective optimisation to help aid the optimisation. One of the products of running a multi-objective optimisation is that a Pareto front is created. This will output several optimal designs where inspection and mass can be

traded. We therefore have two objectives. The first objective is to minimise the RI geometry mass and the second objective will minimise the percentage area which does not receive four scans.

To emphasise the focus on achieving full inspection the following equations were used to calculate the constraint objective: (note: 'PerScan4x' and 'AreaScan4x' is the 'percentage area' and 'area' respectively with at least 4 scans, 'PerScan5x' is with at least 5 scans etc.)

$$\text{U-NDE Objective} = \mathbf{100} \times \left((\mathbf{1} - \text{PerScan4x}) + \text{ScanPenalty}\right)$$

If the disc is 99.99% fully inspected (4 scans) or greater the following equation is used to calculate PerScan4x:

$$\text{PerScan4x} = 1 + \frac{\text{AreaScan4x} + \text{AreaScan5x} + \text{AreaScan6x} + \text{AreaScan7x}}{1\text{e}^5}$$

Otherwise if the disc is less than 99.99% fully inspected the following equation is used:

$$\text{PerScan4x} = \frac{AreaScan4x}{\text{AreaDisc}}$$

The Scan Penalty is made up of the sum of two penalties:

$$\text{ScanPenalty} = \text{ScanPenalty1} + \text{ScanPenalty2}$$

Where ScanPenalty1 is a penalty for not achieving a full scan in critical areas such as the cob of the disc:

$$\text{ScanPenalty1} = 100 \times (\text{Percentage of the critical areas not fully scanned})$$

And ScanPenalty2 is a penalty for not achieving at least 3 scans, 2 scans or 1 scan within the disc profile:

$$\text{ScanPenalty2} = 100 \times (\text{PerScan0x} + \text{PerScan1x} + \text{PerScan2x})$$

Using these formulae, the designs which have less than 99.99% fully scanned areas will create a value based on the percentage not fully scanned but may have an additional penalty resulting from non fully scanned critical areas or by having results with fewer than 3 scans.

Results which have 99.99% or greater fully scanned areas will create values greater than 1 as it is these results which are favourable due to the high likelihood of being close to a fully inspected design. The actual result is based on the non-dimensionalised summation of fully scanned areas. This method has been applied to draw out additional results through the use

of the Pareto front. These results with greater than 99.99% fully inspected may still be classed as a poor design if the scan penalty score is sufficiently high.

The manufacture minimum area penalty which was included in Section 3.2 unintentionally did not feature in this or in the remainder of the optimisation loops although ideally it should be included. This penalty could not be re-introduced into this chapter due to the lack of time and computing resources required. In the examples shown later in this section, 2 out of the 4 examples are marginally effected, so will be corrected where results are compared between the sections.

Stage 2 starts off by post processing the results of the initial global optimisation. The constraint of full inspection can now be applied, but as you may not necessarily get a 100% inspected design, the following criteria was setup:

1. Discard results which do not have at least a 99.5% inspected cob. The cob of the disc must be fully inspected, therefore removing designs which are not close are better off discarded
2. Discard results which are not at least 99.5% 3 scan inspected. Keeps only the top designs.
3. Discard results which are not at least 99.0% 4 scan inspected. Keeps only the top designs.
4. Discard results which are greater than 5% of the current minimum volume. Removes all the inspectable designs which are considerably heavier than designs already found
5. Discard results which are similar in design (all nodes lie within a 5mm radius of each other). Removes one of the designs which are almost identical as local optimisation are likely to make them the same which would be a waste of resources.

This removes a large number of non feasible designs and ensures only unique designs continue to the next stage where local optimisation is performed on the remaining RI designs.

Stage 3 performs local optimisation on all designs which passed through Stage 2's down selection process. The optimisation algorithm 'fmincon' the MATLAB® optimisation toolbox is not a multi objective optimiser, therefore a single objective function was created to take into account both the mass and the inspection constraint:

$$\text{Objective} = \text{RI Geom Mass} \times \left(1 + 100 \times \left(\sqrt[3]{1 - \text{PerScan4x}} + \text{Scan Penalty}\right) + \text{Manu Penalty}\right)$$

Where 'Manu Penalty' is the same as in Section 3.2. Here the objective is to minimise the mass of the RI geometry but in order to satisfy the 100% inspectable constraint, a large penalty function has been added to ensure inspection takes priority. To make sure this only takes over the objective when the part is not fully inspected the cubed route is used.

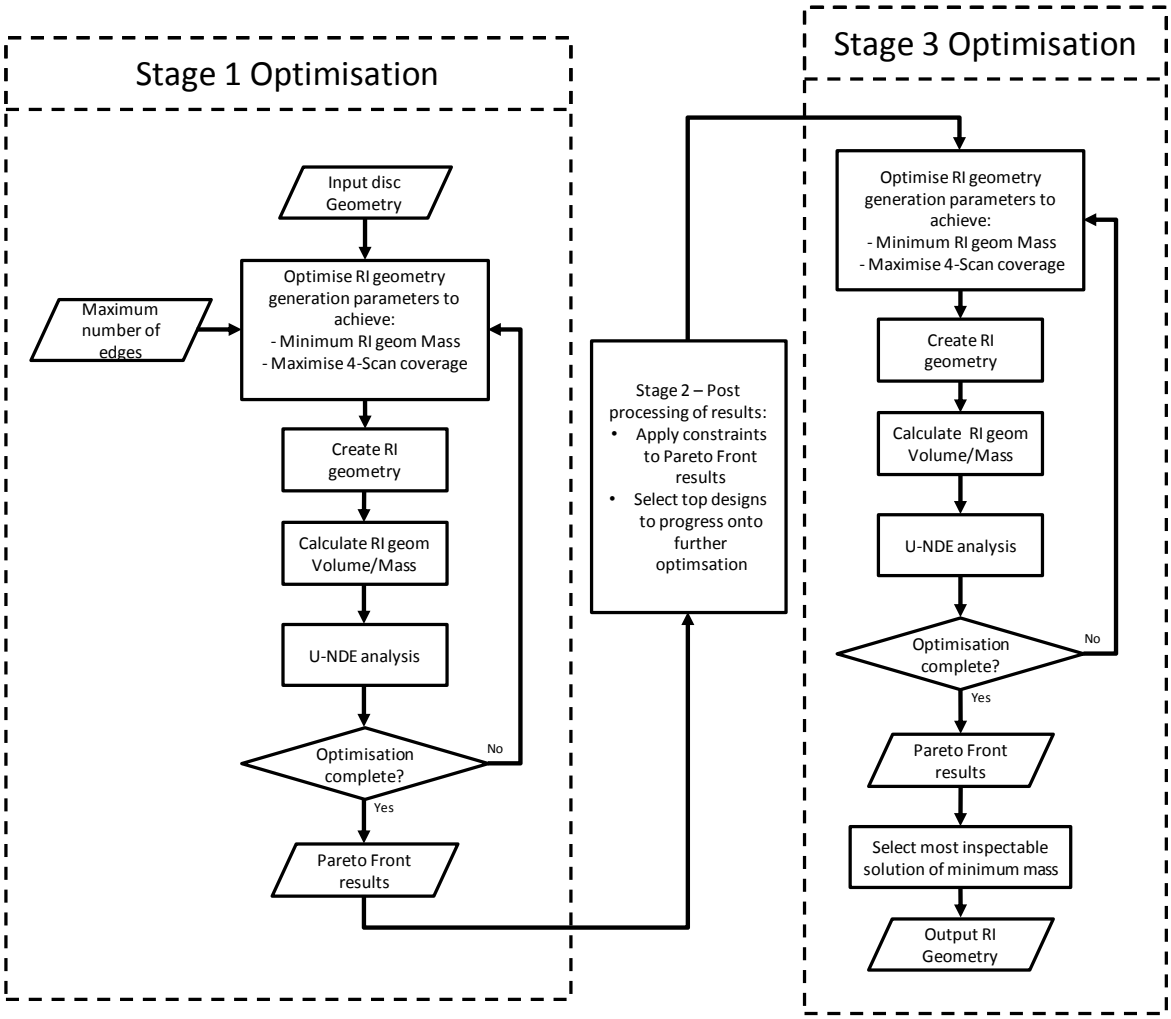

Figure 51- Workflow to optimise RI geometry for both mass and ultrasonic NDE inspection

As this is faster than all other known software it means that running the same optimisation within Rolls-Royce is now practical and possible.

With this new workflow in place a full optimisation on the RI geometry can be performed to find the minimum RI geometry mass is to achieve a 100% fully inspected design.

As one of the parameters for the number of RI geometry edges was kept external to the optimisation loop, it was found that running the geometry for a large number of edges frequently had solutions that contained a fewer number of edges due to the added code for simplifying the geometry. Using this knowledge, it was decided that running the RI

geometry for a variety of number of edges was not required as the optimisation algorithm would automatically discover parameter values, which would result in fewer RI geometry edges if it improved the mass or inspection. As a compromise between optimisation time and flexibility 22 edges was selected to be the default number for all subsequent optimisation runs. By not having to run the optimiser for a variety of number of edges reduced the optimisation time considerably with no loss to the flexibility of the optimiser.

3.3.5. RESULTS – OPTIMISATION WORKFLOW

The results of the optimisation on 4 case study discs are shown in Table 8 below. Each disc shows the best result for each of the 5 repeat runs along with the results produced in Section 3.2 including the ultrasonic NDE result.

Table 8 – Mass and ultrasonic NDE results of the optimised designs.

	Result Source		RI geometry	U-NDE
	Description	Section	$Mass_{norm}$	Coverage (%)
CS-1	Existing RI geometry design*	-	100.00*	93.67^
	RI geometry optimisation without U-NDE	3.2	92.29	94.78
	RI geometry optimisation with U-NDE	3.3	98.38	100.00
CS-2	Existing RI geometry design	-	100.00	100.00
	RI geometry optimisation without U-NDE	3.2	87.31	98.56
	RI geometry optimisation with U-NDE	3.3	90.64 [a] | 90.78 [b]	100.00
CS-3	Existing RI geometry design	-	100.00	99.88
	RI geometry optimisation without U-NDE	3.2	82.31	96.10
	RI geometry optimisation with U-NDE	3.3	86.18	100.00
CS-4	Existing RI geometry design	-	-	-
	RI geometry optimisation without U-NDE	3.2	83.86	98.68
	RI geometry optimisation with U-NDE	3.3	94.26	100.00

* The CS-1's actual RI geometry mass is based only on the final stage RI geometry design, the first stage is 36.2% heavier

^ This result becomes 100% inspected when the two stage RI geometry ultrasonic NDE results are combined, but it is not considered in this comparison.

[a] Result produced though optimisation ignoring the manufacturing minimum material requirement.

[b] Mass of the RI geometry post design correction to take into account the manufacturing minimum material requirement.

The results in Table 8 show that complying with the ultrasonic NDE requirement increases the RI geometry mass for all engine designs compared to the previous section. The example designs have increased by 6.6, 4.0 and 4.7 percent for CS-1, CS-2 and CS-3 respectively. Despite the increase in overall RI geometry mass, each result is still below existing RI geometries by 9.2 and 13.6 percent for CS-2 and CS-3 respectively. CS-1 has an RI geometry mass reduction of 1.6 percent, but the result from the existing single RI geometry stage[19] does not meet the inspection requirements.

It was explained previously in Section 3.3.4 that one of the constraints was unintentionally missed out within the optimisation loop. Due to time constraints these experiments could not be repeated. Due to the unique setups of each disc design the manufacturing minimum geometry constraint would only impact CS-1 and CS-2 RI geometry designs due to the welded stubshaft design. After analysing the RI geometry results it was found that CS-1 was unaffected as the constraint was already satisfied, however CS-2 did not satisfy the constraint. In order to obtain valid side-by-side comparisons between results the design was altered such that the manufacturing constraint was met whilst maintaining full inspectability. Figure 52 illustrates the necessary changes performed to satisfy the missed constraint. One edge was manually moved resulting in a 0.16 mass increase in RI geometry mass.

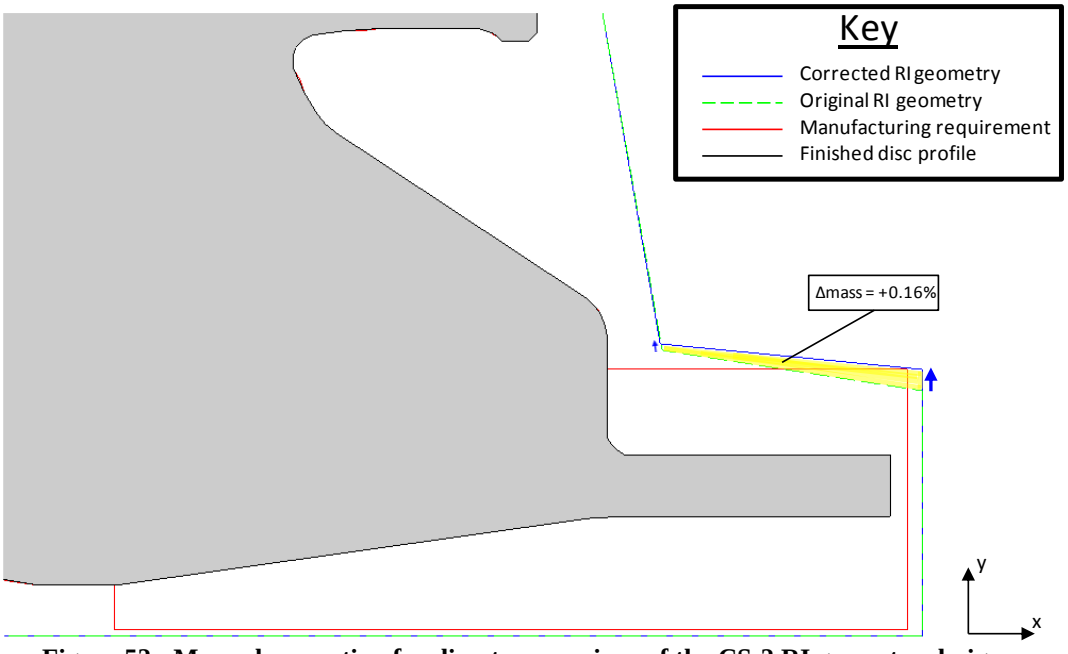

Figure 52 - Manual correction for direct comparison of the CS-2 RI geometry design.

[19] The CS-1 RI geometry is inspected in two stages to achieve 100 percent inspection, the first RI geometry stage is 36.2 percent heavier than the final RI geometry stage shown.

Figure 53 below shows that the optimal designs do not differ considerably from those analysed in Figure 50 from Section 3.2.

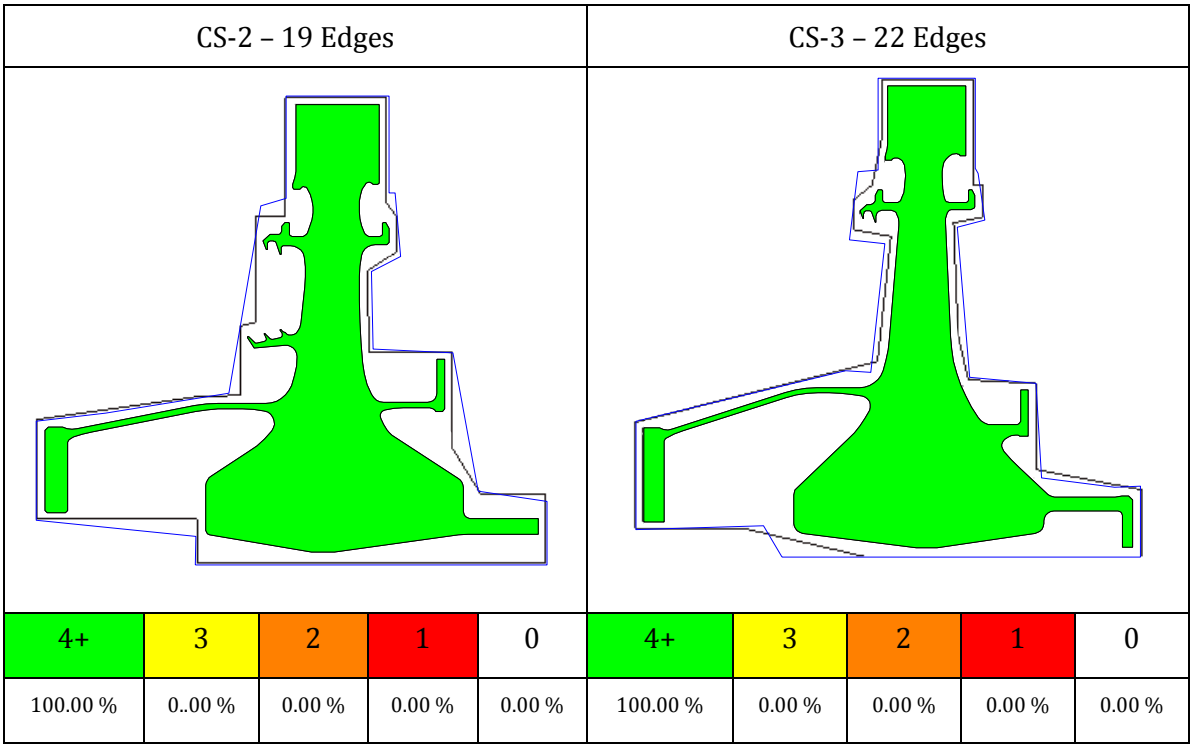

Figure 53 – Optimal RI geometry designs generated by the optimiser. The blue RI geometry outline represents this sections results where ultrasonic NDE is considered; the black RI geometry outline represents the designs created in the previous section (Section 3.2) where no ultrasonic NDE is considered.

The complexity of the new CS-2 RI geometry design has reduced as the optimal number of edges selected is 19 out of the possible 22 allowed. The reduction in edges is a sign that the inspection constraint is difficult to achieve due to the reduced sound going into the part. The way in which it has increased the inspectability of the component is to reduce the number of edges. The CS-3 RI geometry on the other hand is using the maximum number of edges allowed showing that there are areas which have plenty of ultrasound. The mass of the CS-3 RI geometry has increased by 4.6% primarily due the inspection requirement around the rim and hooks. As the rim and hooks are at a high radius, the amount of mass put on for a given cross sectional area is 2 to 3 times that of the bore.

CS-4 failed to find a good result where all five optimisations produced similar styled forgings as shown in Figure 54.

Section 3.2 result – 22 Edges	Section 3.3 result – 19 Edges
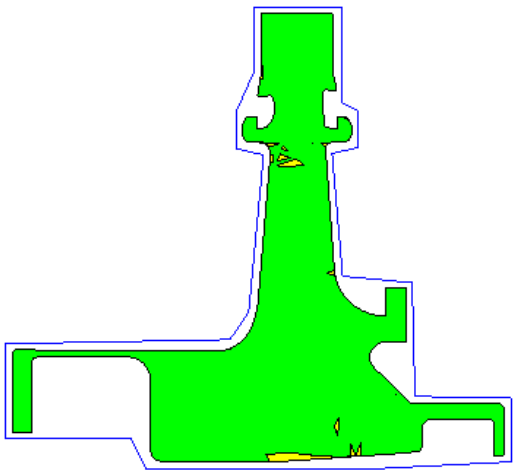	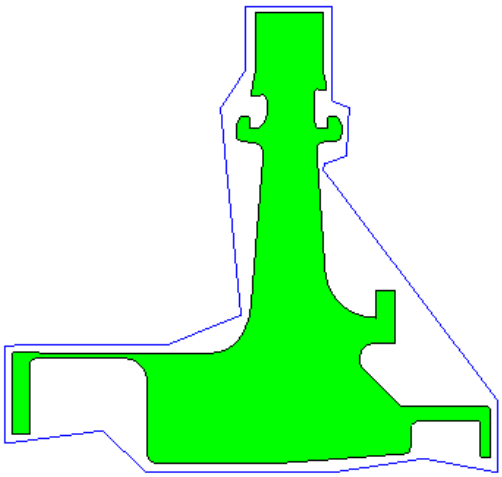

Figure 54 - Optimised RI geometry designs for the CS-4 disc for both this Section 3.3 and the previous Section 3.2.

The results shown in Figure 54 show a considerable change in RI geometry. The reason for this is again the new inspection requirement. Due to the difficulty in finding geometry which results in 100 percent inspection the optimiser has again simplified the geometry by reducing the number of edges to 19. After obtaining the inspection by creating one single straight line at the rear of the disc, the local note perturber has pushed the nodes at the bore outward causing slight undercuts reducing both the mass of the RI geometry and improving the inspection within areas of the bore and diaphragm.

The CS-1 result is shown in Figure 55 below. The comparison between Section 3.2 and the latest optimisation results show that the vast majority of additional material added is around the rear portion of the cob above the stubshaft (see Figure 2 in Chapter 01 terminology).

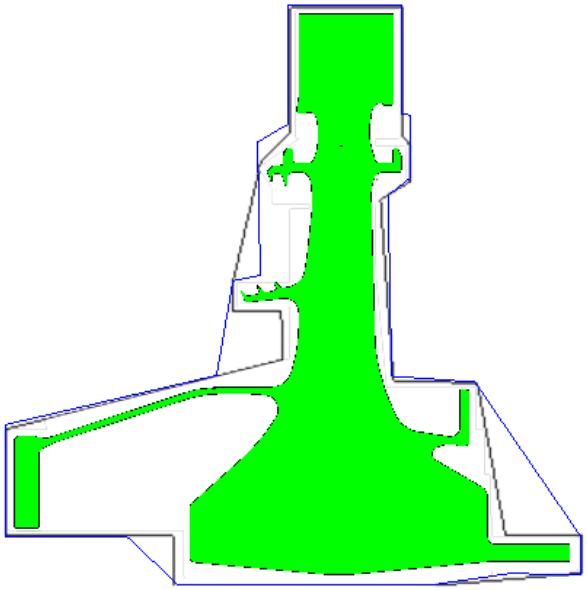

Figure 55 - Optimised CS-1 RI geometry. The blue RI geometry outline represents this sections results where ultrasonic NDE is considered; the black RI geometry outline represents the designs created in the previous section (Section 3.2) where no ultrasonic NDE is considered.

Figure 56 below displays the trade-off between mass and inspection after performing Stage 3 optimisation.

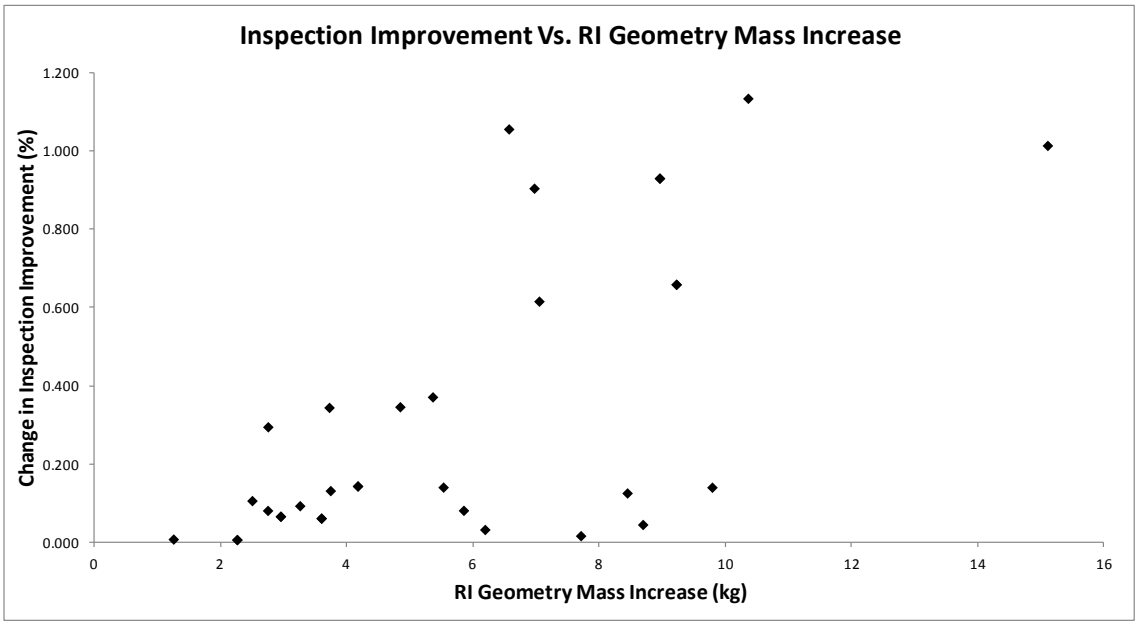

Figure 56 - Graph displaying the cost on RI geometry mass by increasing the inspection levels up to 100% using the Stage 3 local improvement optimisation method. The results are taken from the CS-3 case study.

The results of Figure 56 show how the ultrasonic inspection can be improved significantly but typically at the cost of increasing the RI geometry mass. A reduction in mass and or cost is normally desired, but not at the cost of inspection.

3.3.6. DISCUSSION

This section set out to create a piece of software which was able to quickly analyse a given geometry to a greater accuracy than has previously done before. It also fills an important gap in the research.

Utilising this software to analyse a design in such a short time span provides a big opportunity in terms of optimisation for unit cost. Due to the cost impact from the ultrasonic analysis that has to be performed, estimates of the total unit cost of the component can now be improved as a dummy representation of RI geometry is worthless if it does not accurately represent practical and verified geometry.

Summary of Findings

The study found that using intelligent manipulation of geometry to locate the inspectable regions of the disc proved to be substantially quicker than previous discrete methods which allowed the software to be used in batch to analyse geometry quicker than ever before.

The inclusion of the inspection constraint on the design of RI the geometry was shown to have a considerable impact on the RI geometry mass. This mass increase ranged from 4.0% for CS-2 through to an increase of 12.4% produced by the CS-4 optimisation. The increase in RI geometry mass has been shown to be considerable proving that the RI geometry cannot be optimised without sufficient ultrasonic analysis being performed during the optimisation process.

This automated process has shown that the RI geometry design can be run autonomously without the need of a manufacturing engineer or NDE expert. Once all the key data from the NDE and Manufacturing experts are input then the optimiser can be left alone and has shown to produce a result 9.2% more efficient than the existing CS-2 RI geometry design, and 13.6% more efficient than the current CS-3 RI geometry design when looking at RI geometry mass alone. The CS-1 case study has shown the performance of the optimisation workflow creating a workable design 1.6% lighter than the final RI geometry stage (of two stages) but with a 100% inspectable design. This would mean the CS-1 RI geometry could be produced without the additional 36.2% mass required to make the existing first stage RI geometry design necessary to achieve 100% inspection making savings on both material and manufacture.

It was demonstrated that it is not possible to create suitable RI geometry without designing it around ultrasonic limitation analysis. Designs created in the previous section (Section 3.2) were found to have failed the ultrasonic NDE requirement. 32 edge RI geometry designs had less than 89% fully scanned, where as the 22 edge designs achieved greater than 94.75% fully inspected. All three designs had limitations in the cob region of the discs and also had fewer than 3 scans in places. Results showed that to achieve the inspection requirement between 6 and 26 kilograms of material needed to be added. This demonstrates that performing preliminary analysis on designs which have not taken into account ultrasonic NDE requirements may not reflect production designs potentially rendering the results useless.

The missed out constraint in this section only impacted one of the four case studies. The impact of missing the constraint only equated to a 0.16% mass increase in RI geometry. Due to the small effect on mass, future results which also miss out this constraint will be analysed without the need to modify the design to validate it assuming visual checks show similar discrepancies.

Academic Implications

The creation of the (patent pending) Ultrasonic NDE analysis software has bridged the gap which published research was missing. This new analysis software will enable a clear and accurate understanding of how the RI geometry can impact upstream manufacturing processes and costs.

Business Implications

Running the workflow illustrated in Figure 51 has shown to be extremely effective and would allow the designer of the component to create a RI geometry in under a day as opposed to weeks in the traditional fashion. This saves time and resource with only a small time penalty to the person who sets up the optimisation workflow. The running of the software also impacts the RI geometry cost and using the RI geometry mass as a surrogate for cost, the raw material savings add up to a considerable saving.

Limitations and Directions for Future Research

The main limitations of this study is the NDE and critical parts lifing experts' knowledge, which would allow in certain circumstances, the acceptance of RI geometry designs which are not fully inspected in specific areas of the disc. Utilising more detailed knowledge within this software could achieve even greater savings. The aim of the method was to create a 100% fully inspected design, but when this becomes a struggle for the optimiser an NDE expert is required to asses if the non-fully inspected zones are acceptable. In the case

studies shown here all outputted designs passed ultrasonic NDE assessment therefore this was not a problem.

This method ultimately tries to minimise the sound going into the disc whilst just maintaining full inspectability. From a risk point of view, it would be ideal to maximise the number of scans into the component from as many different directions as possible. This technique was utilised within the objective function for the ultrasonic inspection, but was used to allow a greater number of near fully inspected designs to be considered rather than being used to directly improve the inspection. The opportunity could have been explored further in this section, but would have steered away from the primary objectives set out in Section 2.9. Another possible improvement would be to re-evaluate the ultrasonic NDE rules due to the impact the fully inspected requirement has on cost. If the ultrasonic NDE rules were more flexible it could allow the ultrasonic NDE rule set to be optimised for a particular design allowing greater freedom in the RI geometry design.

The biggest limitation still within this optimisation loop is the objective to minimise mass as a surrogate for unit cost. Without improving on the cost estimation method it may not be worth minimising the RI geometry for mass in all areas. This is because there may not be a strong impact on the cost of the disc when considering the manufacturing process. The cost of a disc is heavily linked to the input mass of the material at forging, the RI geometry mass is not necessarily proportional to the weight of the forgings. Therefore, to get a better surrogate for cost it is not the RI geometry mass that should be minimised but the black forging mass. The next section will look into upstream material costs as a method of improving the cost surrogate.

3.4. FORGING MODELS TO BETTER PREDICT COST

3.4.1. INTRODUCTION

This section looks in more detail at the manufacture of disc components, more specifically the forging. The manufacturing method for isothermally forged discs is an expensive route and makes up a substantial portion of the cost. Understanding the manufacturing process allows the manufacturing design route and original design to fully exploit any upstream benefits, such as cost savings. Typically, forger's Black Forging (BF) designs are proprietary information and can take many weeks to produce given fixed RI geometry. Creating these black forging designs to a good standard during the design stage in an accurate and efficient manner could have substantial cost benefits which are not currently realised.

The original aim of this project was to optimise the RI geometry without trying to model the forging geometry; this decision was made to avoid what was perceived to be complex and time consuming modelling. There was also a lack of forging knowledge within the research group. It was eventually realised that the forging geometry is an essential necessity to the RI geometry optimisation procedure when trying to minimise part unit cost. Without an estimate of the forging, the RI geometry would almost certainly be optimised to the smallest possible volume whilst just maintaining full inspectability; but this does not necessarily minimise cost.

Early attempts at optimising the RI geometry for unit volume found correctly that the RI geometry would be made as close to the disc as possible, whilst just maintaining inspection criteria. In Chapter 02's literature review it was noted that discs are manufactured with considerable excess material in the forging process, which is paid for then later machined away. This material which is machined away could be used to design a larger RI geometry which could improve ultrasonic inspection with no overall increase in material input and subsequent cost. This could allow the RI geometry to utilise the unused forging material whilst reducing expensive material from other areas of the forging. This will ultimately improve the component cost and the RI geometry design efficiency. The improvement in cost assumes that turning material removal cost rates are the same for the Heat Treatment (HT) geometry to the RI geometry, and the RI geometry to the 2D axi-symmetric disc profile. Therefore, modelling the black forging geometry will not only provide further inputs to improve the accuracy of a cost model, but will also provide far more freedom to the ultrasonic NDE expert by not necessarily having to worry about spending more on material.

Solving these upstream questions on manufacture during RI design phase will provide knowledge to the design which would ordinarily be missed.

Introducing a simple forging model will provide invaluable feedback to the RI design process. The RI geometry in the previous section was costed using its own mass, but by estimating the forging mass around the RI geometry a better cost estimate will be found. Minimising the RI geometry volume causes a dispute between inspectability and weight. However, by introducing the forging cost as the new target it may be possible to provide a better inspection per unit cost as not all of the geometry will directly impact the cost as shown in Figure 57. This allows RI geometry to be created which will be better value for money. Exploiting the forging shape to provide a better inspected disc improves the overall quality/safety of the product at no extra (direct) cost in material.

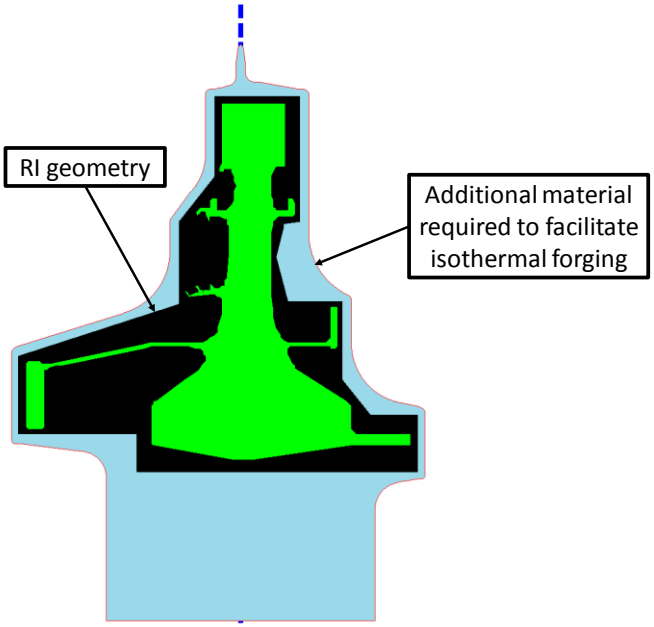

Figure 57 - Forging geometry based around an optimised RI geometry from Section 3.3.

To enable us to further explore how the forging geometry can be exploited, this section will look at efficient and effective methods of generating forging geometries for HPT discs.

The key objectives of this section are as follows:

- To generate a piece of software to model the black forging design around a given RI geometry to primarily produce a forging mass estimate. This should run in the time frame comparable to the DUNDE software created previously

- To investigate the impact of using Black Forging mass as a new surrogate for cost as opposed to the RI geometry mass.

The study will be run on the same case study discs which were shown previously.

This section first looks at how a forging geometry can be created around RI geometry. An automated method is then created and tested. The final part of this section is to incorporate this new forging and material estimation tool within the optimisation loop to see the benefits of optimising the geometry for minimum input mass.

3.4.2. METHOD – BLACK FORGING GEOMETRY GENERATOR

The method used in this section to generate the Black Forging mass is meant to provide a good estimate of the forging geometry and is not an attempt at producing a finished forging/die design.

In order to obtain a good cost estimate of the black forging geometry more in depth research was undertaken to gauge the practicality of creating a forging model. Early research found how complex the art of forging is due to several dozen parameters which influence the actual forging process and its interactions with the detailed material properties. A heuristic approach described by *Kulon* (Kulon, Mynors et al. 2006) shows a very good graphical method of designing a forging geometry using some basic rules. Although these rules appeared to make a good estimate of the forging profile, it was appreciated that this would yield the minimum possible volume of the final forging. To check that the design is forgeable with no defect, a forging FE simulation would have to be performed, and then altered accordingly. Enacting this methodical approach would add considerable time and complexity to the optimisation making it unsuitable for this thesis. It was also appreciated that such detail may be unnecessary for the preliminary design stage. Separate discussions with experienced forgers (Bryant 2010; Brooks 2010a) agreed that a forging profile would be a benefit to the design of the RI geometry, and even an extremely basic model would be sufficient to gauge the increase in volume following changes to the RI geometry design. Further discussions established that applying a more specific rule set would improve the accuracy to the extent that it would be capable of creating geometries which may be very similar to the actual designs (assuming realistic values for the rules can be extracted or found). The ASM (American Society for Metals) Handbook (Semiatin 2005) provided an abundant resource of knowledge with far more detail that could be used to create forging designs with many general 'rule of thumb' rules. The next sub-section describes how the rules have been interpreted from *Kulon et al* paper (Kulon, Mynors et al. 2006) and the ASM

handbook (Semiatin 2005), to produce forging profiles from generic shapes with minimal computational time.

3.4.2.1. Further Research On Forging Rules

Using the ASM international handbook on forming (Semiatin 2005); the forging is looked at in greater depth to understand the specific rules which will influence the black forging geometry design.

3.4.2.1.1. Draft

Draft is the term used to describe the taper applied to vertical walls to allow easy removal of the workpiece from a die (Semiatin 2005). The absence of draft leads to the forged component tending to 'stick' to the die surface despite the presence of lubricants and highly polished surfaces (Semiatin 2005). In general draft angles of ~3-7 degrees are applied to hammer and press forgings. Zero draft forgings can be produced but require 'strippers' and/or 'knockout pins' to eject the parts from the die cavities (Semiatin 2005). It is common practice to apply draft to steel, nickel and titanium alloys as they generally require a machined finish for an improved surface finish (Semiatin 2005).

The draft values used directly affect producibility, die filling, withdrawal and die wear (Semiatin 2005) and the location of maximum draft is directly linked to the parting line discussed later. The greater the draft angle the greater the material waste and machining costs.

There are several types of basic draft as shown in Figure 58. The 'outside draft' which would be found on the outside of ribs and bosses (material generally shrinks away from this surface as it cools). 'Inside draft' is found on the inside of ribs and bosses (material will contract towards the die plugs). 'Blend draft' is draft necessary to blend the top and bottom dies together and can take three different scenarios as shown in Figure 58. 'Natural draft' is draft that is already incorporated in the original design, i.e. surfaces which have angles from the vertical greater than the required draft angle. Other types of draft include:

- 'Shift draft' – This is where the draft is applied to only the inner part of the rib but double that of the required draft angle.
- 'Back draft' – This is a more extreme version of 'shift draft' where the rib now undercuts relative to the die motion, but because the part can be withdrawn at an angle it will not get stuck.

It is unlikely these two draft types are required for this problem though due to the axi-symmetric nature of the problem.

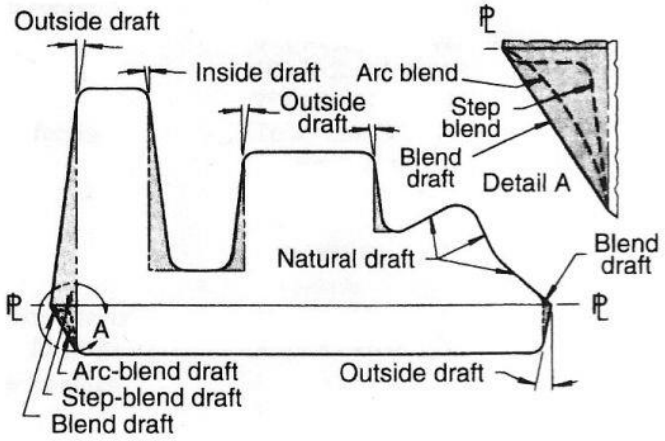

Figure 58 - Illustration of draft types (Semiatin 2005)

According to the 'ASM handbook' (Semiatin 2005) for forging, heat-resisting alloys and refractory metals (which include nickel-based superalloys) a draft angle of 7-10 degrees is required with a tolerance of +2, -0 degrees (Semiatin 2005). The draft angles applied to both Titanium, Nickel-based alloys and refractory metals are generally analogous to steels (Semiatin 2005). It is common practice to apply a single draft angle across an entire forging even to both internal and external surfaces (Semiatin 2005).

Conclusion

As you would expect an increase in draft angle does not tend to enhance the component properties or microstructure, and adds unnecessary weight which increases the machining costs. Therefore, the best way to reduce cost is to use natural/inherent draft in the design where the draft angle is implemented into the component design.

3.4.2.1.2. Corners and Fillets

The corners and fillets provide a smooth gradual interface between two flat surfaces removing abrupt angular junctions. They are required to improve metal flow characteristics as well as grain direction. Forging energy can also be reduced as well as the inherent die wear and subsequent risk of cracking.

Corners are defined as a convex arc where two intersecting faces have an angle greater than 180 degrees. Fillets are located at joints where two intersecting faces have an angle of less than 180 degrees such as where a rib intersects a web.

The size of corners and fillets depend on several factors including rib height, the material properties, the forging process, factors associated with die filling and producibility (Semiatin 2005). The radii associated with corners are directly related to the height of the

rib which is measured from the parting line. For multiple forging stages the size of the corner radii is inversely related to the refinement of the forging. Therefore, blockers will start with large radii then will become progressively smaller in subsequent blockers towards the final forging design.

The radii associated with fillets are also directly related to the height of the adjacent rib or boss, but are also related to the forging process (Semiatin 2005). Again for multiple forging stages, fillet radii decrease with increasing refinement of the forging process. Fillet radii are affected by confined webs and if the web is less than the rib width (Semiatin 2005). The design of corners and fillets must satisfy both the requirements of metal flow in forging and considerations of cost arising from metal usage and the subsequent material removal.

Larger corner radii reduce stress concentrations in a die impression and thus promote die life. Smaller radii will reduce material and machining costs, but increase the likelihood of die breakage as well as increasing the forging pressure requirement (Semiatin 2005).

Fillets are also used to prevent 'laps' or folds in the material flow. Having a greater radius reduces the lateral movement of the flow preventing a large separation from the wall. Use of a blocker can also solve lap problems (Semiatin 2005). For any given rib height, the size of a fillet is invariably larger than the size of a corner as seen in Figure 59.

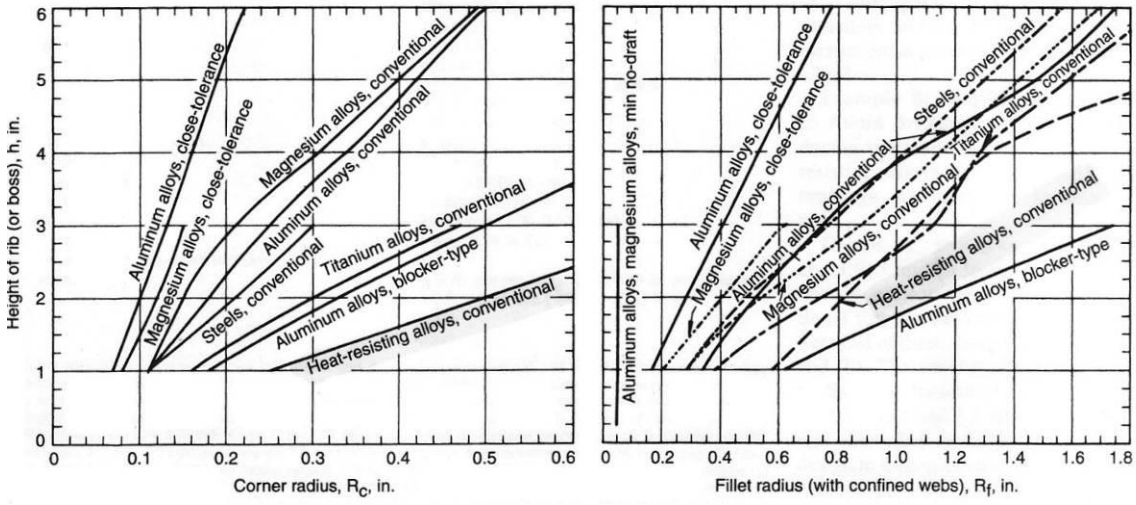

Figure 59 - How radii are related to the height of the rib (Semiatin 2005)

Conclusion

The minimum radius on a fillet or corner is strongly dependent upon the rib height, but to achieve die life the radii need to be maximised, therefore the trade-off between cost and die life first needs to be considered.

3.4.2.1.3. Webs and Ribs

The web is a relatively thin plate like element of a forging which generally connects ribs and bosses. In general, a web lies at an angle less than 45 degrees relative to the forging plane where features that lie at an angle greater than 45 degrees are described as ribs. The location of a web is thought-out along with the design of any ribs, bosses and the location of the parting line. (Semiatin 2005)

Forgability issues arise when the web thickness to rib height ratio is too large, as material flow will first miss the rib, then reverse direction as the pressure builds up potentially causing folds. This issue can be solved using greater fillet radii or by creating a preform die.

Conclusion

The issues arising with the webs and ribs will require consideration; a possible solution to forging problems is to increase the rib thickness to achieve the correct ratios.

3.4.2.1.4. Flash and Trim

Flash is all the excess metal which is forced outward from the workpiece during a closed die operation. It is a necessary part of the design to ensure full filling of the die impression and to avoid a 'fin' which would stop complete closure of the dies in the event of oversized billets being used.

A 'fin' is where the excess material flows in between the flats of the two die surfaces at the parting line. Once material enters the parting line the dies are unable to completely close. To prevent a fin from occurring, a suitable size 'flash gutter' is required.

Billet volumes tend to have a relatively large tolerance (currently unknown), due to inaccuracies in billet diameter. Therefore, the flash gutter needs to be large enough to allow the largest billets to be processed, but to ensure there is enough back pressure to allow the die impression to be fully filled a 'flash saddle' is required as illustrated in Figure 60. Nickel-based superalloy forgings can require a flash thickness of over 5.1 mm. Common 'width of land' to 'land thickness' ratios range from 3-to-1 to 4-to-1. For further choking of the material flow through the flash saddle, use of a 'corrugated land' can be utilised. (Semiatin 2005)

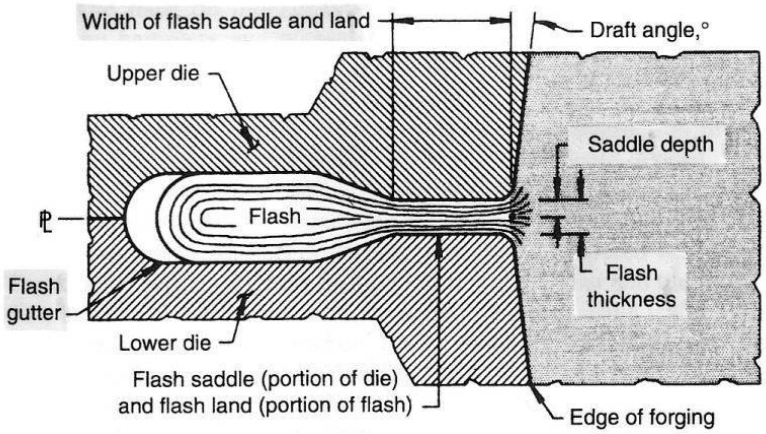

Figure 60 - Illustration showing terms related to the flash (Semiatin 2005).

The relief from the flash land to the gutter is commonly 15 to 30 degrees and for simple configurations the flash can make up 5% of the total weight of the forging. (Semiatin 2005)

The flash will always be located at the parting line of the dies so three options are available to where the flash will end up; either in the upper, lower or both dies.

Conclusion

The flash provides a convenient means of disposing of excess metal allowing the use of slightly oversized billets, saving costs in terms of scraped parts. The second use for the flash is to control the metal flow in the die impression ensuring complete die filling. As 5% of the total mass may end up as flash it is necessary to model the exact amount which may be required.

3.4.2.1.5. Parting Line

One important design feature of the forging is the Parting Line (PL). The parting line is the line which divides the two die halves as shown in Figure 58. The parting line position can be positioned and orientated in an infinite number of positions, but the final decision is generally influenced by the performance of the forged part (due to the material flows) and the impact on unit costs of the forgings (Semiatin 2005) (material costs, die life and batch number).

The performance of the forging in service is significantly dependent upon grain direction; therefore the parting line will normally conform to the direction of greatest strength within the component (Semiatin 2005) which for a disc would be along the plane of the web. This for instance can be achieved by ensuring the material flow vectors conform to optimum grain flow direction suitable for in-service operation. This can involve a straight parting line

for straight forward components or a broken parting line to obtain a desired grain flow through the component (Semiatin 2005).

Although aerospace components in general have a greater emphasis on performance, the design for producibility takes priority: In general the parting line will lie at the largest cross section to avoid forcing material into deep die impressions (Semiatin 2005). The parting line can significantly affect the input material required. Ribs, bosses, flash, corners, trim and draft will all influence the location of the parting line.

Conclusion

The parting line location is another optimisation problem in itself with no single hard and fast rule associated with it. In terms of producibility, the parting line needs to be positioned taking into account draft angles, corners, etc to obtain a forging of least volume and least die cost, and in terms of performance the parting line should lie on or close to a key web. Other general rules to research further such as equal volume in each half of the die (Brooks 2010b) and validations are to follow later.

3.4.2.1.6. Preform Blocker

For complicated forging designs it is necessary to use a blocker as an additional step between the billet and finished forged geometry. The blocker is a simple die design used to create favourable billet geometry for subsequent forging. The rules for designing the blocker geometry is relatively ambiguous compared to the finished geometry, this can be overcome using a Shape Factor (SF) and Volume Factor (VF) concept (Kim and Park 1998). The SF is defined as the ratio of difference between the blocker and finished geometry, so varies between 0 and 1 and the VF is defined as the ratio of the blocker volume to the finished forge volume and can take values greater or equal to 1.

In general, the blocker draft angles are greater by 1-2 degrees for external draft and 2-3 degrees for internal draft up to 10 degrees. Also the radii of the blocker are larger than that of the final die. The fillet radius is about 25% smaller than the corner radii.(Kim and Park 1998)

Conclusion

It is currently understood that the isothermal forging process does not normally use blockers due to the time and cost associated with changing the dies; therefore, a preform blocker is out of scope for this project.

3.4.2.2. Method for Generating a Forging Profile from Generic Geometry

As explained earlier, the rules described by *Kulon et al* (Kulon, Mynors et al. 2006) were used as a starting point together with the ASM handbook (Semiatin 2005) for more detailed knowledge. Table 9 shows the procedure compared to the base line model:

Table 9 - Heuristic approach to generating black forging geometry for a given parting line

Step No.	**New heuristic approach**	**heuristic approach *by Kulon et al.*** (Kulon, Mynors et al. 2006)
1	Fill in undercuts (up to parting line)	Fill in undercuts (up to parting line)
2	Apply offsets to factor in Machining Tolerances (also creates corners)	Machining Tolerance
3	Apply offsets to factor in Forging Tolerances (also creates corners)	Forging Tolerance
4	Remove Narrow Features	Remove Narrow Features
5	Insert Draft (up to parting pine)	Insert Draft (up to parting line)
6	Remove newly created Narrow Features (this step is omitted in the loop due to computation time and the fact that the fillet radii sufficiently apply this step)	Fillets and offset corners
7	Create Flash around Parting Line	Ribs (interpreted as: Parallel to plane offsets)
8	Form Fillets (both inside and outside of billet diameter)	
9	Remove excess flash– smooth the top and cut at axis of symmetry at the bore	

It was noticed that Step 1 ('fill in undercuts') was not required as the later Step 5 ('insert draft') would fill in all undercuts eventually. Despite this finding 'fill in undercuts' was still included; by keeping this step redundant lines are removed which simplify the geometry. This improves performance because it reduces the process time for all subsequent steps up to and including Step 5, the 'insert draft' step. The new developed process is illustrated in Figure 61 below.

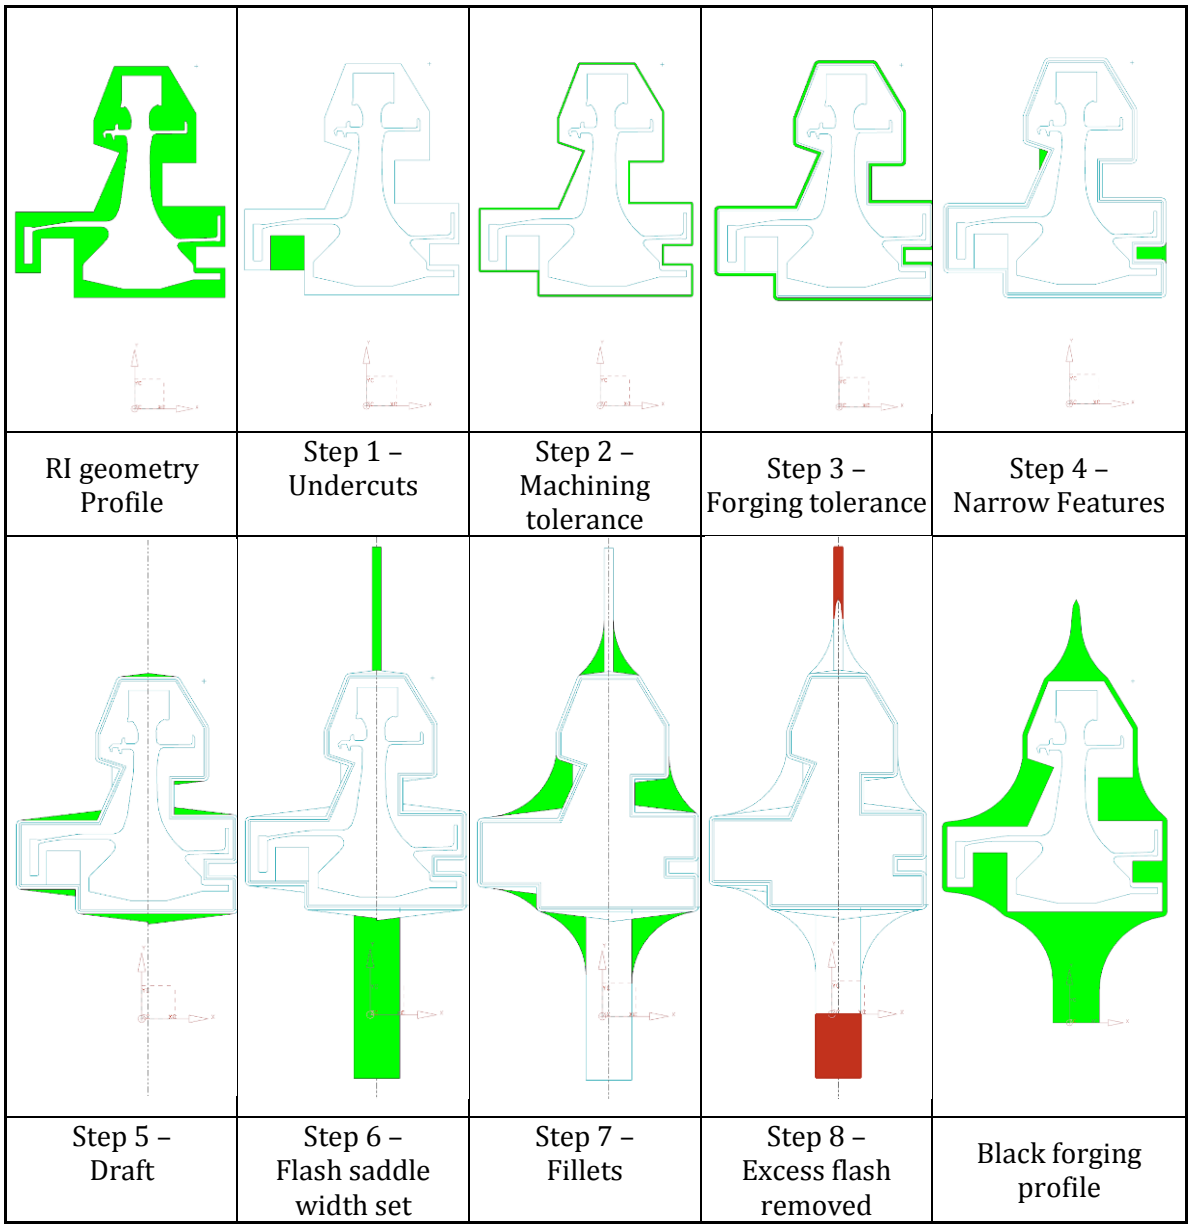

Figure 61 – Current heuristic approach to generating black forging geometry.

This combines all the key rules required to design a forging. It is appreciated that the rules actually produce the smallest possible forging volume required to make the part and may not actually be forgeable in practice. A discussion with Prof. Brooks of Strathclyde University suggested that only minor alterations to knowledge based design are required to make a design forgeable from this stage (Brooks 2010b).

Coding each one of these operations has been carried out in MATLAB® (The digital location of the code can be found within Appendix VII) as separate functions with exception of the 'machining Tolerance' and 'forging tolerance' which uses the same displacement function

with different input values. A basic flow diagram of the functions to perform this task is shown in Figure 62.

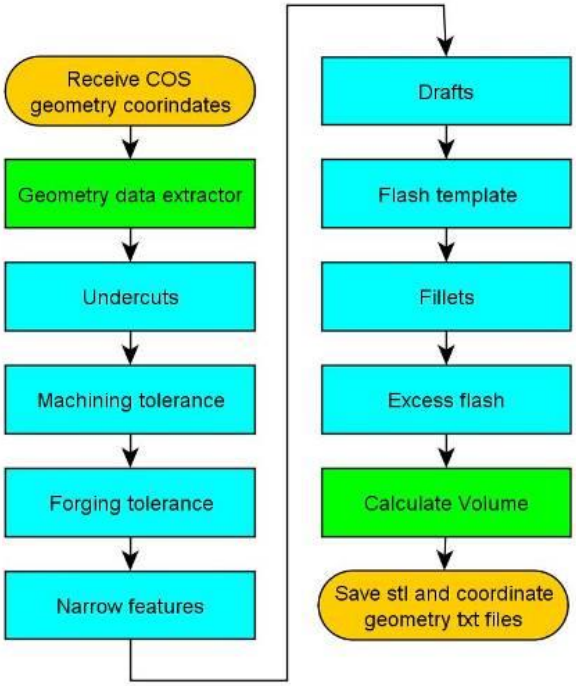

Figure 62 – Programmatic route in creating the black forging geometry.

3.4.2.3. Parting Line Position

The parting line position discussed in Section 3.4.2.1.5 is one step which is not explicit but could strongly influence the cost of a component if positioned incorrectly. Therefore, it is a parameter which can be optimised to provide the design of minimum volume as shown in Figure 63. The current optimisation method uses a golden search algorithm to find the minimum volume which takes approximately a dozen iterations to solve.

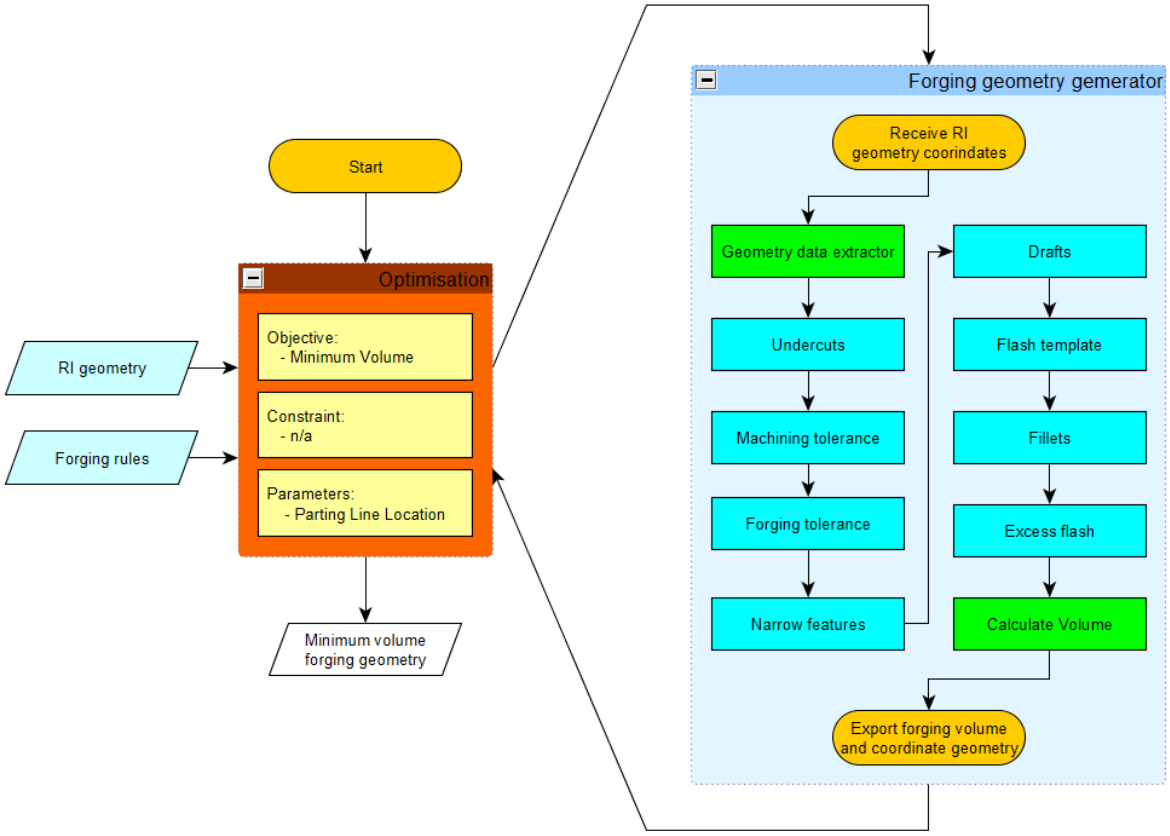

Figure 63 - Method used to find the parting line value thus the minimum forging mass.

An example of how the parting line effects the forging geometry for a given RI geometry can be seen in Figure 64a to Figure 64c with the corresponding impact on the forging volume shown in Figure 64d.

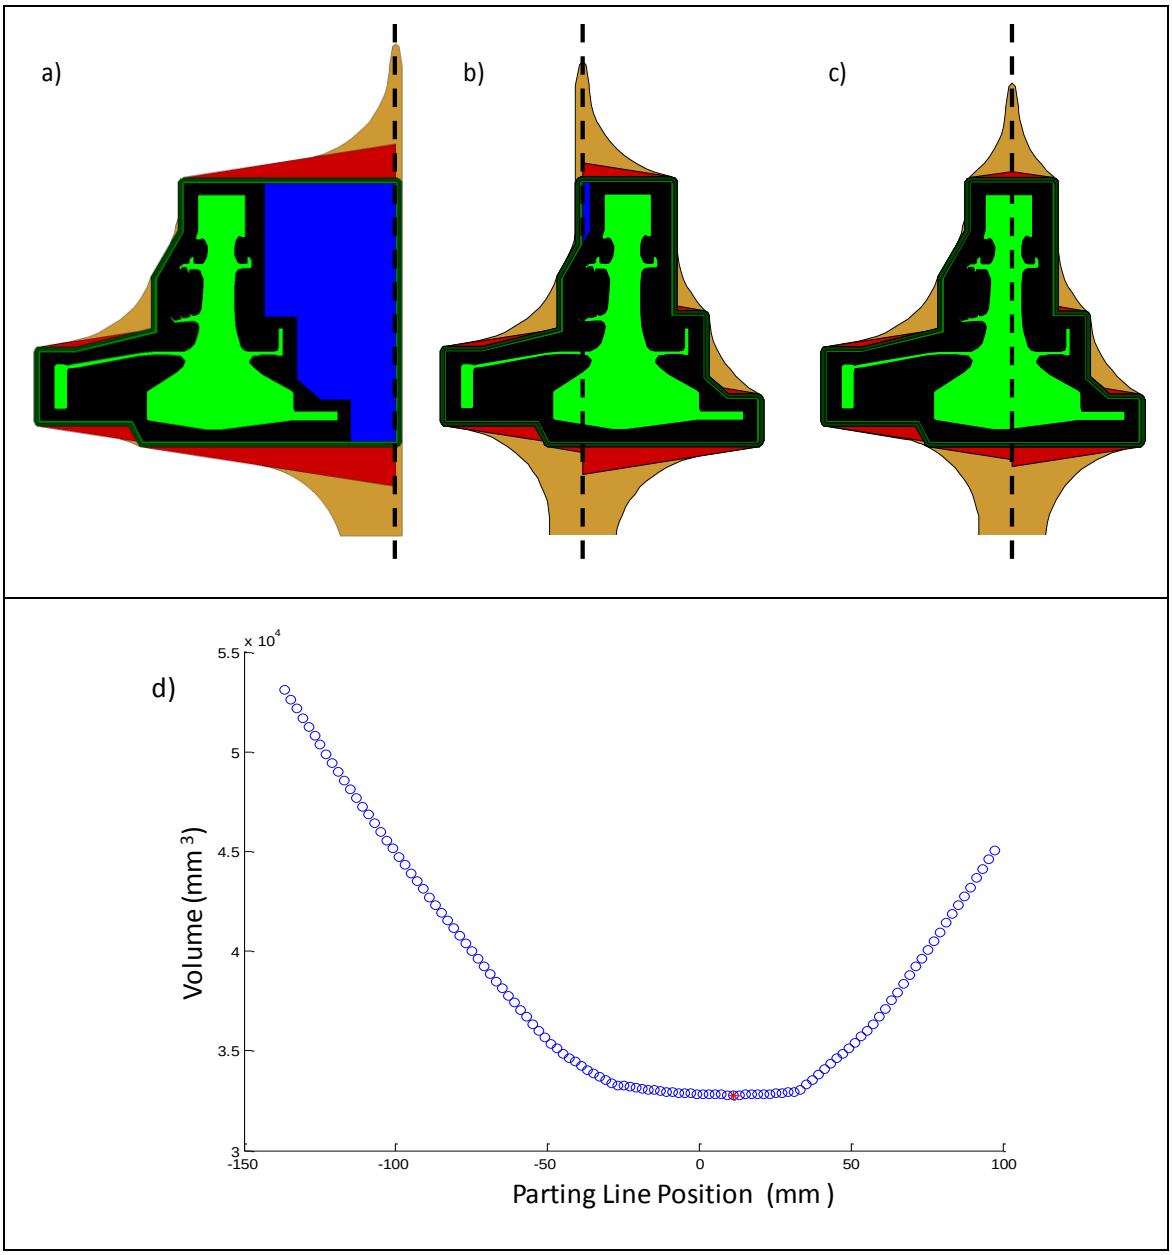

Figure 64 – Example illustrating how the parting line influences the volume/mass of the forging. Diagrams A to C show the parting line in three different positions and diagram D shows the relationship between the black forging volume and the position of the parting line relative to the centre line of the disc.

Initial RI geometry design optimisation loops used the parting line optimisation within each design iteration, but it was later found that the parting line moved negligibly for a given disc design. Therefore, to save computational time the parting line position is optimised around the RI geometry minimum design (as shown in Figure 24 in Section 3.2), then it is fixed for the remainder of the optimisation. The flow diagram for this alternative workflow is shown in Appendix III.

3.4.2.4. Black Forging Geometry Validation

The actual forging geometry used to produce the case study parts were not available to directly overlay and compare. One way in which the forging estimations can be directly compared is using the billet input mass. Figure 65 below shows the output of the black forging geometry generation software and Table 10 compares the results of the software against the actual input billet mass values.

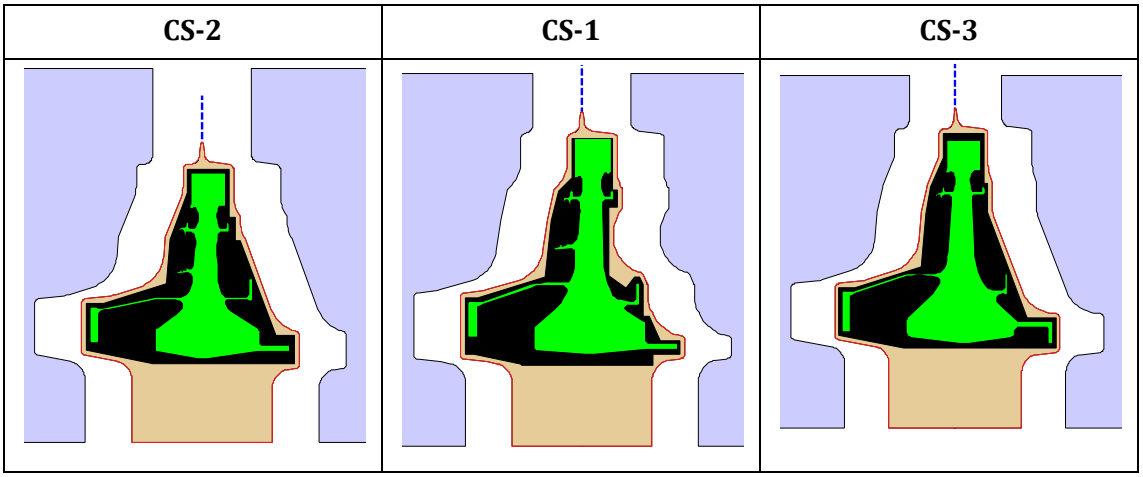

Figure 65 - Production RI geometry with the generic forging rules applied to produce a prediction of the input billet mass.

Table 10 - Table with comparisons between actual billet input mass and predicted forging mass.

	CS-2	CS-1	CS-3
Existing Billet Mass$_{norm}$	139.55	Not known	131.55
Predicted Forging Mass$_{norm}$	134.86	225.49	132.46
Percentage Error	3.4%	Not known	0.7 %

The results shown in Table 10 above demonstrate acceptable accuracy being within 3.5% of the actual mass value. This comparison assumes that the billet mass is equal to that of the forging volume as the flash is included within the forging design, but dependent on the forgers, the billet volume may be approximately 0.5% larger than the forging volume to factor in billet and forging die errors. Other errors may have arisen due to differing parameter values compared to those used by the forgers as they will have their own parameters based on experience. Another cause of error could be from forging geometry design tweaks; these would have been done to improve the material properties after running finite element analysis.

3.4.2.5. Optimisation

With a functioning forging geometry generator, we are now able to calculate the actual material input mass, as opposed to a scaled mass from the RI geometry as used in Sections 3.2 and 3.3. The running time of this software is similar to the ultrasonic NDE software; therefore, this material estimator is able to improve the accuracy of surrogate for mass/cost in Section 3.3. By adding this to the loop it is expected that the RI geometry can be designed more effectively, conforming to the forging rather than the disc. The work flow from Section 3.3 has been adapted to include the forging geometry and is shown in Figure 66 below.

Note to improve the computation time of the Forging estimator, the Parting Line is set by running the parting line optimiser around the minimum disc geometry as discussed earlier. The value outputted by this is used for all the runs within the loop. An actual parting line optimisation is performed on the final RI geometry to get the true result. Details of these adaptations to the flow diagram shown in Figure 66 are covered in Appendix III.

.

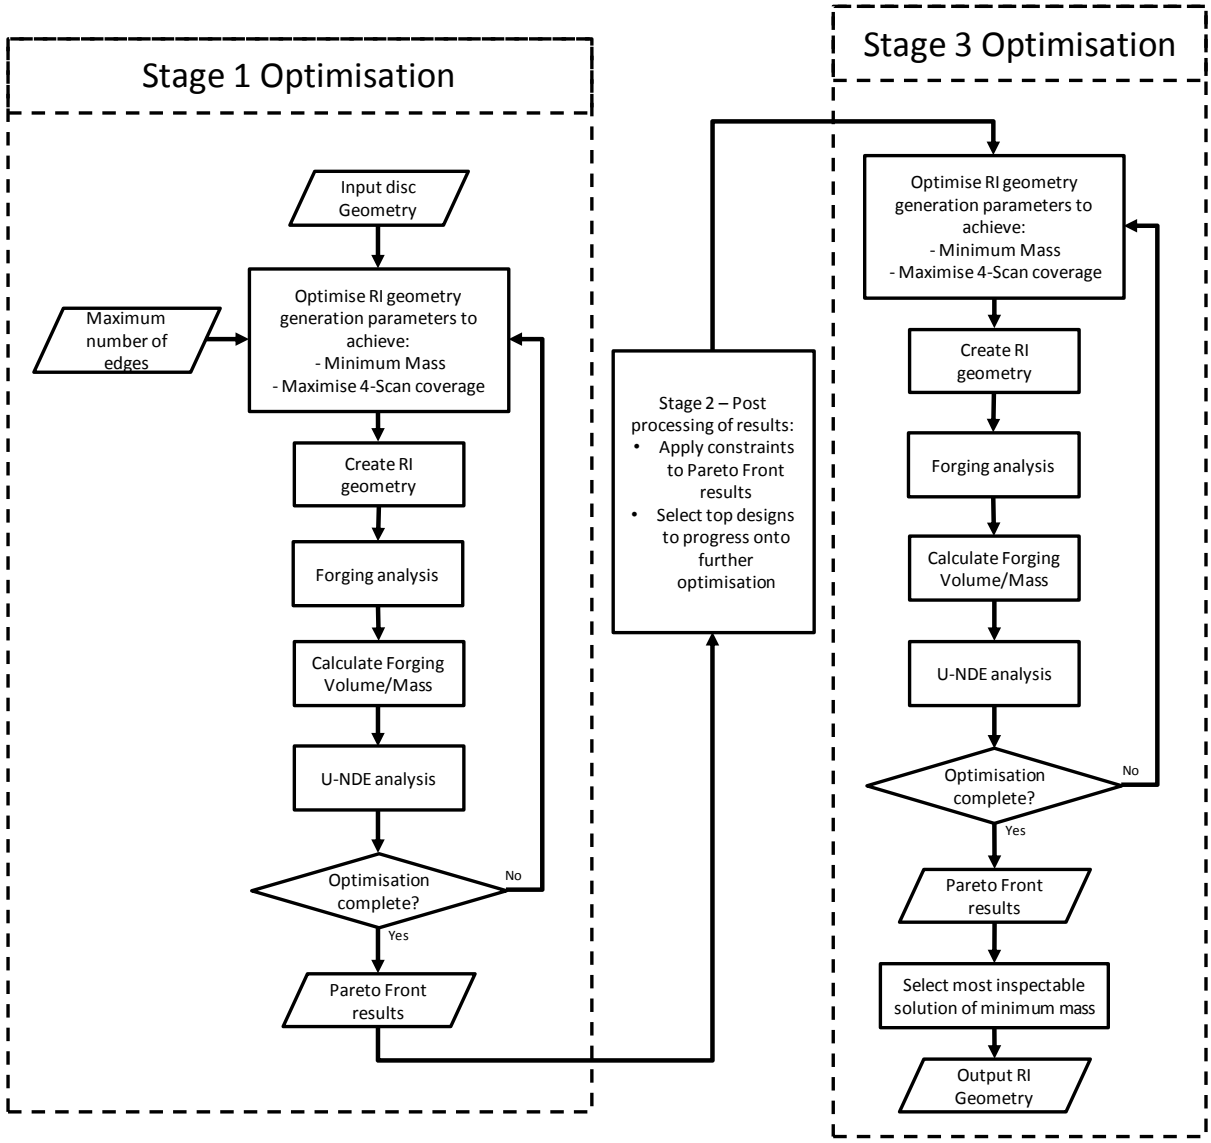

Figure 66 – Workflow for Section 3.4 optimisation to minimise black forging mass.

The workflow illustrated will be run using the same optimisation parameters as used in the previous section to give a like for like comparison.

The objective functions for both the ultrasonic constraint and mass minimisation remain the same as the previous section with the exception that the new mass objective is to minimise the black forging mass instead of the RI geometry mass.

3.4.3. RESULTS

This new software module can now be run on the results from both Section 3.2 and 3.3 to see how they compare when evaluated using the forging geometry.

Table 11 – Results of optimisation including previous section results with new predicted black forging mass estimation. The results displayed are the best results out of 5 independent optimisation runs. (Values in brackets are actual billet mass values).

	Result Source		RI geometry	U-NDE	BF
	Description	Section	$Mass_{norm}$	Coverage (%)	$Mass_{norm}$
CS-1	Existing RI geometry design*	-	100.00	93.67	135.62^
	Min RI geometry Mass without U-NDE	1.1	92.29	94.78	133.36
	Min RI geometry mass with U-NDE	3.3	98.38	100.00	136.27
	Min BF mass	3.4	99.14	100.00	**136.06**
CS-2	Existing RI geometry design	-	100.00	100.00	134.86
	Min RI geometry Mass without U-NDE	1.1	87.31	98.56	128.23
	Min RI geometry mass with U-NDE	3.3	90.64	100.00	**129.49**
	Min BF mass	3.4	96.95	100.00	131.95
CS-3	Existing RI geometry design	-	100.00	99.88	132.46
	Min RI geometry Mass without U-NDE	1.1	82.31	96.10	118.92
	Min RI geometry mass with U-NDE	3.3	86.18	100.00	**123.66**
	Min BF mass	3.4	89.61	100.00	124.19
CS-4	Existing RI geometry design	-	-	-	-
	Min RI geometry Mass without U-NDE	1.1	83.86	98.68	123.75
	Min RI geometry mass with U-NDE	3.3	94.26	100.00	132.55
	Min BF mass	3.4	89.95	100.00	**126.69**

* The CS-1's existing RI geometry is based only on the final stage RI geometry design

^ Considering the first stage RI geometry the forging mass estimation is 25 percent greater

Table 11 shows how changing the objective to forging mass produces a change in trend for the RI geometry mass.

For CS-1 the optimisation has managed to find a forging design which has reduced the input billet mass by 0.15% despite the RI geometry mass increasing by 0.78%. CS-2 has unsuccessfully managed to find a reduced forging mass but has instead increased by 1.9%; this increase is low considering that the RI geometry has increased by 7.0%. CS-3 also hasn't found a better forging mass solution, but the increase in forging mass is only 0.4% compared to a RI geometry mass increase of 4.0%. Optimising CS-4 has located a step improvement in design which has produced an overall mass saving of 4.4% on the forging and 4.6% on the RI geometry.

All of these results show that the RI geometry mass is only a portion of the forging mass so an increase in RI geometry may or may not increase the overall forging mass as illustrated in Figure 67.

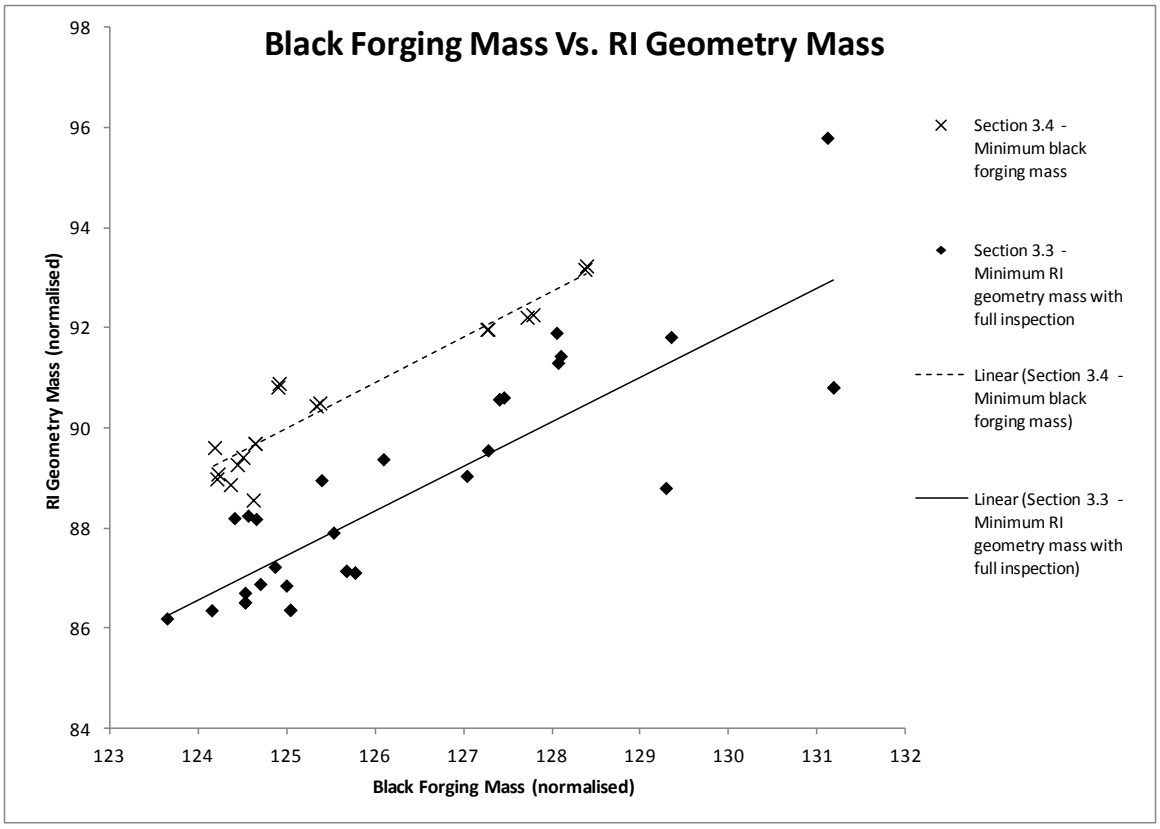

Figure 67 - Plot of all the CS-3 analysed results from Section 3.3 and 3.4 gathered from every local optimisation run for each of the 5 optimisation runs.

Figure 67 shows that for a given forging mass there are a number of feasible RI geometry solutions. For instance, for a normalised forging mass of around 125 it is possible to have a

normalised RI geometry mass ranging from 86 to 91. The reason for this is that for a given forging geometry you could get a RI geometry which either hugs the finished component or one which hugs the forging geometry. The graph also shows how the objective of minimising the black forging geometry has meant that for a given mass you always get a larger RI geometry suggesting that the objective of minimising the RI geometry mass is a misguided choice.

In Section 3.3's results section the effectiveness of the Stage 3 local improving optimisation was analysed. Using new results from this section overlaid with the previous results the comparison between mass and ultrasonic inspection improvement can be seen in Figure 68.

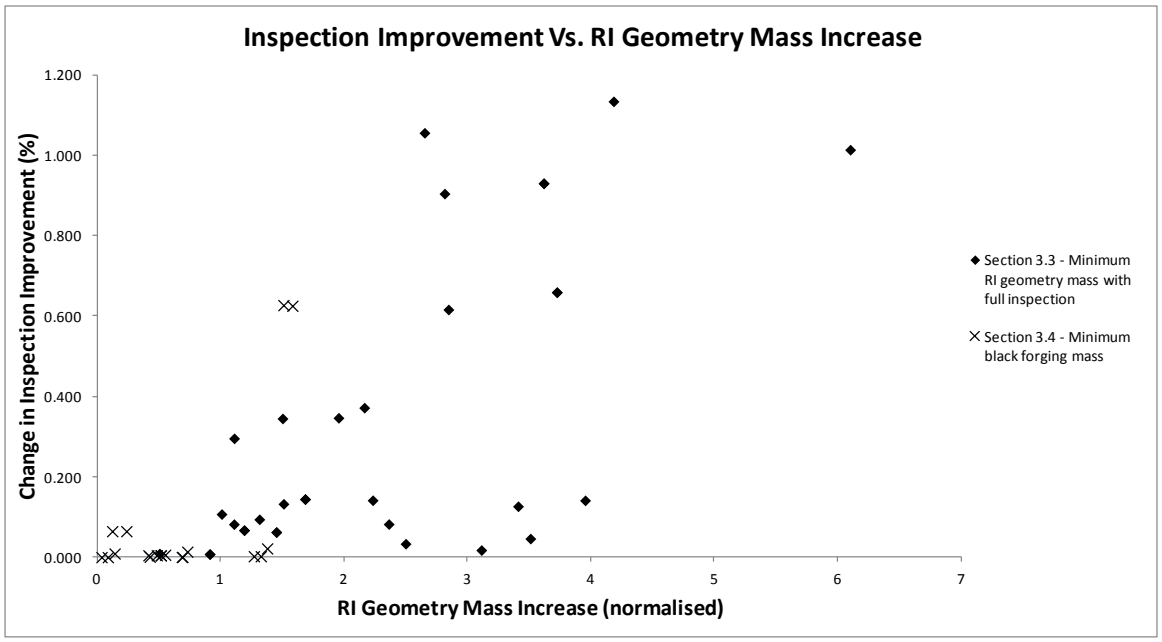

Figure 68 - Graph displaying the cost on RI geometry mass by increasing the inspection levels up to 100% using the Stage 3 local improvement optimisation method. The results are taken from the CS-3 case study for Sections 3.3 and 3.4.

The results show that this section's optimisation has a higher initial inspection coverage; an average of 99.917% compared to 99.669% for Section 3.3 for the CS-3 results. The differences between the two sections also show that the mass increase in the RI geometry is considerably lower on average.

The change in mass differences is exaggerated further when the black forging mass is considered instead of the RI geometry mass; Figure 69 displays these results.

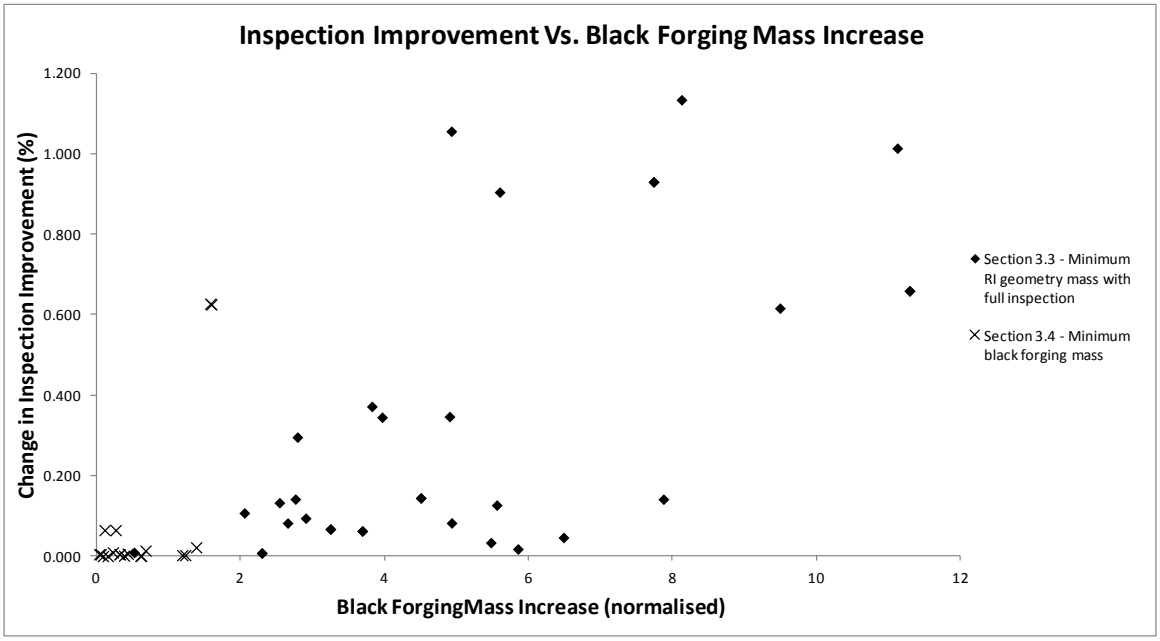

Figure 69 - Graph displaying the cost on black forging mass by increasing the inspection levels up to 100% using the Stage 3 local improvement optimisation method. The results are taken from the CS-3 case study for Sections 3.3 and 3.4.

It is again clear that this sections objective of minimising black forging mass instead of RI geometry mass produces greater improvement efficiency as locally improving the inspection has a much reduced effect on the objective of minimising black forging mass.

3.4.4. DISCUSSION

This section set out to investigate the impact of forging mass on the overall RI geometry design. It is the first time the forging models have been used to autonomously drive the design of rectilinear inspection geometry which ultimately produces a cheaper product with the equivalent functionality of previous designs. Previously unseen forging material which is often overlooked during the design of the RI geometry can actually positively influence the design of the RI geometry. Utilising this Black forging geometry generating software provides a big opportunity in terms of optimisation for unit mass and will become a vital input into cost models looked at later in this chapter.

Summary of Findings

It has been shown that it is possible to create a good estimate of the black forging design using only a set of rules. This approach and method allows the black forging geometry to be generated in a matter of seconds. This would avoid lengthy and costly communication between the RI geometry designer and the forgers.

Setting the optimisation objective from RI geometry mass to black forging mass minimising has shown promising results. It was expected that the RI geometry mass would increase in mass whilst the black forging mass would decrease beyond that seen in the previous section. This did happen in two of the four case studies, but in the other two case studies where the black forging mass did increase slightly, it showed that the RI geometry mass was allowed to increase significantly. On average the RI geometry mass increased by 1.79% with only a 0.19% increase in black forging mass shared across the four case studies.

Results also showed better ultrasonic inspection efficiency during the initial global search of the design space when the optimisation objective was changed from RI geometry mass to the black forging mass. The reason for this is due to a lower weight penalty when the RI geometry increases in size to improve ultrasonic inspection. The optimiser is able to find better inspected designs with a lower weight increase. CS-4 showed to be the most successful example where an overall better design was found compared to Section 3.2's methods.

Using the forging geometry to improve the surrogate for cost was expected to improve the forging mass for each component, but the results have shown that changing the objective of the optimisation from minimum RI geometry mass to BF mass improved results for only half the cases. The reason for the poor performance is highly likely to be down to a combination of the following:

- The constraint of full ultrasonic inspection is a difficult constraint to achieve.
- The limited number of optimisation evaluations needed to converge to an optimal result.
- The RI geometry parameters do not always have a direct impact on the black forging mass which may cause confusion within the optimisation algorithm.
- The efficiency of the optimisation algorithm for this problem.
- The parameterisation method of the RI geometry creation.

Despite the slight inefficiency of the optimisation and method, the results still prove that the RI geometry design has far greater freedom when optimising for minimum BF mass as opposed to RI geometry mass. This is due to the amount of forging material which could be utilised to enlarge the RI geometry at no extra cost.

Running the workflow described in Figure 66 has shown to be effective as a new objective for optimisation. This would allow the designer of the component to create the RI geometry in a matter of hours or days as opposed to weeks in the traditional fashion utilising the

forging mass in the most efficient manner. This saves time and resource with only a small time penalty to the person who sets up the optimiser. The running of the software also impacts the RI geometry cost and using the BF mass as a surrogate for cost, the raw material savings add up to a considerable saving.

Academic Implications

The implementation of the generic forging geometry generator was shown to give good results creating forging estimates of within 3.4% for the two case study examples shown. The use of this useful mass estimation tool will be vital for estimating part unit cost in later sections.

Business Implications

The main limitation on this study is that the forging geometry generated can vary from forger to forger depending on their own forging parameters which are found from trial and error and modelling. These forger specific parameters have not been investigated due to software, time and expertise limitations. During the creation of the forging geometry generator several more detailed questions arose when programming the forging model, for instance:

- Draft angle: should there be independent inner and outer draft angles?
- Offset: is there vertical offset to implement which varies inversely with distance from the centre of rotation axis?
- Minimum groove width: would this vary with depth of grove rather than being a fixed value?
- Corners: do corners always add material horizontally or can you add the material radially if it saves material? (Should this be implemented so that it is close to what the forgers currently do or what makes better financial sense?)
- Fillets: should the fillet radii be related to the corner radii through results shown in the ASM Forging Handbook, or be related to the thickness of the web at that point, or just use a fixed value for the entire part? Is the fillet at the bore and rim flash the same as the rest?
- Flash: Is there a simple way to calculate how thick the bore will be? How long will the flash be after the fillet? How big should the saddle be? What percentage of the input material will end up in the flash gutter?
- Ribs: currently we do not check to see if the rib needs to be made wider to ensure that it is forgeable. According to ASM Handbook (Semiatin 2005), the width of the

rib should conform to a ratio of the web. Does this need to be considered? The fillet radii also cover the same issue.

- Parting Line: What other rules are there in determining where the parting line should lay? Is minimum volume a valid reason to position the parting line? Also methods such as it should lie along the main web, or equal material should lie in both halves of the die? Would you use a broken/stepped parting line which might be useful for the front cover plate[20]?
- Test Ring location: Does a test ring need to be incorporated into the Forging geometry design, or are integral test rings in the bore sufficient?

The answers to these questions would be useful in improving the accuracy of the black forging models, but for the purposes of this thesis, where relative comparisons of different designs are required they were not required.

Limitations and Directions for Future Research

This method ultimately tries to generate a good approximation of the forging geometry, even if the geometry is not the exact geometry which a forger may use, it does allow a comparison of different RI geometry designs for relative comparison.

In creating the black forging geometry additional data has been created opening the possibility of populating a detailed analytical cost model. Aspects such as die size and billet length are all inputs to an ideal cost model which would otherwise be unavailable. This is explored further in Section 3.6.

[20] In the case where a seal cannot be integrated onto a disc, it is integrated on to a cover plate. These are similar to discs and fasten onto the side of the main disc.

3.5. MATERIAL PROPERTIES FROM HEAT TREATMENT OPERATIONS

3.5.1. INTRODUCTION

This section looks at creating a surrogate model to assess the potential impact on material properties within the finished component.

Material properties are an important input into the design of a component, therefore it is a necessary requirement to achieve at least the minimum specification of a material, but it would be advantageous to increase the material properties above the minimum. This would increase the safety margin in the component or could be used to enable further optimisation of the finished component geometry using the improved properties. This section has been included to demonstrate the impact that the different optimisation objectives have on the heat treatment geometry. Although the direct impact that the heat treatment geometry will have on life is not investigated in much depth, being aware of the potential material property improvements could be advantageous to the disc design team.

Heat Treatment Quenching

There are two key desired requirements which can be influenced at Heat Treatment:

- Residual Stress
- Minimum Material Properties

Residual stress's can improve the life of the disc by providing compressive residuals in the disc where high stresses are experienced throughout the cycle, but the disc must have balanced residuals in the finished part otherwise it will distort. Therefore, any areas where compressive residuals are specified on the part, it must be balanced out by accepting tensile stresses in other features. To control the residual stresses around the disc and heat treatment shape, the Heat Transfer Coefficients (HTCs) around the profile of the heat treatment shape need to be controlled independently. Effective control on these HTCs can be very powerful and new technology has been created such as the Super Cooler © (Bunge 2002; Bunge and Furrer 2003) patented by Ladish and a similar technology which has been patented by Firth Rixson (Hebert, Friedman et al. 2014). In both these technologies the HTCs are tightly controlled around the component during quenching as illustrated in Figure 70.

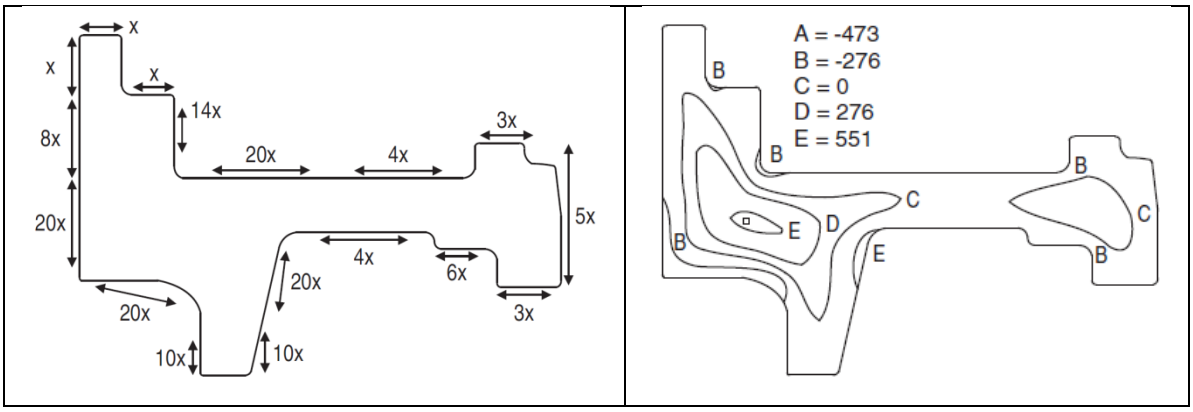

Figure 70 - Optimisation results of non-dimensionallised HTCs (left) with subsequent results of residual stresses (right) (images taken from (Furrer, Shankar et al. 2003)).

The HTCs only directly affect the surface of the heat treatment shape; this means the cooling rates diminish as you get deeper into the component. It is therefore beneficial to have a slim heat treatment shape to allow more control of the cooling rates throughout the component.

Standard Nickel alloys were traditionally oil quenched which involves a very high HTC at the surface and can potentially causing quench cracking. Although the whole heat treatment surface is in contact with the liquid, the convection currents which are created through the heat transfer into the oil cannot be controlled, therefore it is more likely to have non-uniform cooling rates. This method in the past has created manufacturing issues due to the sometimes drastically differing residual stresses from part to part as well going from fat forgings to slim-line forgings. The differing residual stresses across the component cause the part to move during manufacture, initially making the manufacturing process very difficult. A large percentage of discs can be scrapped during the development stage of the manufacturing process. This scrap percentage reduces quickly as manufacturing engineers learn how the component reacts to the cutting sequence. This issue has been reduced drastically through the use of process modelling, but part to part differences are more difficult to account for. At the other extreme the components can be left to air cool. This has a very low HTC which for large nickel based components is unacceptable.

Linked to the creation of residual stresses are the material properties. Part of the material properties are controlled through the quenching stage of the disc. The cooling path of the material can drastically alter the material structure and subsequent properties. To lock in the material properties created during the solution heat treatment a fast cooling rate is required. Unfortunately to control the residual stresses of the component the HTCs might be reduced in areas to balance out the residual stresses across the disc. If it is considered that the minimum material properties within the disc are a limiting factor in the lifing of the disc,

it is desirable to improve this limiting area. As the material properties are related directly to the cooling path it takes during the quenching cycle, the locations which lie furthest away from the surface of the disc are the greatest effected. With the use of controlled HTCs around the surface of the component's heat treatment shape, it is desirable to make the heat treatment geometry as close to the finished part as possible, this increases the minimum possible cooling rate at the centre of the component which should increase control over the components minimum material properties.

To enable both residual stresses and improved minimum material properties to be improved; it is beneficial to make the heat treatment shape as small as possible. As the cost of the component is related to the amount of material used. Minimising the RI geometry allows the black forging and subsequent heat treatment shape to be reduced improving both cost and material properties. The cob of the disc is a highly stressed feature therefore requires good properties, but due to the size and location it normally has the slowest cooling rate.

Catering for the heat treatment process would allow RI geometry to be created which will ensure material properties for life are also considered in the selection of RI geometry design. Also having an early estimation of the heat treatment geometry may allow a preliminary assessment of the residual stresses and material properties. This analysis could be vital information for the design engineering team enabling them to perform initial life assessments.

The life of a component is as important as the cost of the HPT disc. Not only is it important to keep the unit cost of the disc down, but if there are opportunities to improve the life cycle cost of the disc at the same time then this should also be exploited. Therefore, in this section we will be looking at efficient and effective methods of generating forging geometries as well as creating a surrogate for life for HPT discs. Solving these upstream questions during the RI geometry design phase will provide knowledge to the design which would ordinarily be missed.

We have already tried to optimise the RI geometry to minimise the black forging shape but what this doesn't take into account is the minimum depth at the deepest part in the heat treatment shape. We can assume the Heat Treatment geometry is very similar to the black forging geometry as typically only the rim and bore is removed to avoid additional cost.

The greater the minimum depth at the deepest point, the worse the material properties will be or the greater the cooling rates are required at the surface to achieve a faster cooling rate

at this point. However, increasing the cooling rates at the surface can cause surface cracking and unwanted residual stress gradients throughout the part.

This section primarily looks at creating a simple surrogate for material properties; once a method has been created for achieving this then the module will be incorporated into the optimisation loop, where we can see how the consideration of the HT shape will influence the design and cost of the component.

The key objectives for this section are to:

- Generate a typical heat treatment geometry using the Forging geometry as the foundation.
- Calculate the minimum depth at the deepest point within the heat treatment geometry.

In this section we first look at how we can convert the black forging geometry created in Section 3.4 into a typical heat treatment geometry, after this we then develop a method to find the minimum depth at the deepest point within the heat treat shape. The final part is to include this additional analysis into the optimisation loop to see how this influences the design of the RI geometry.

3.5.2. METHOD – CREATION OF HEAT TREATMENT GEOMETRY

The creation method of typical heat treatment geometry is to convert the black forging design into a workable heat treatment shape at minimal cost. The method used typically removes both the rim and bore of the forging and nothing else. This is done because normally the forgings are near-net designs around the RI geometry leaving little material to remove. If further turning stages were added, they would still have to be re-machined post heat treatment doubling the turning financial cost.

The steps to achieving heat treatment geometry are straight forward and are described next:

1. Import black forging geometry.
2. Remove the rim of the black forging geometry *x* mm from the outer diameter of the RI geometry parallel to the axis.
3. Remove the bore of the black forging geometry *x* mm from the inner diameter of the RI geometry parallel to the axis.

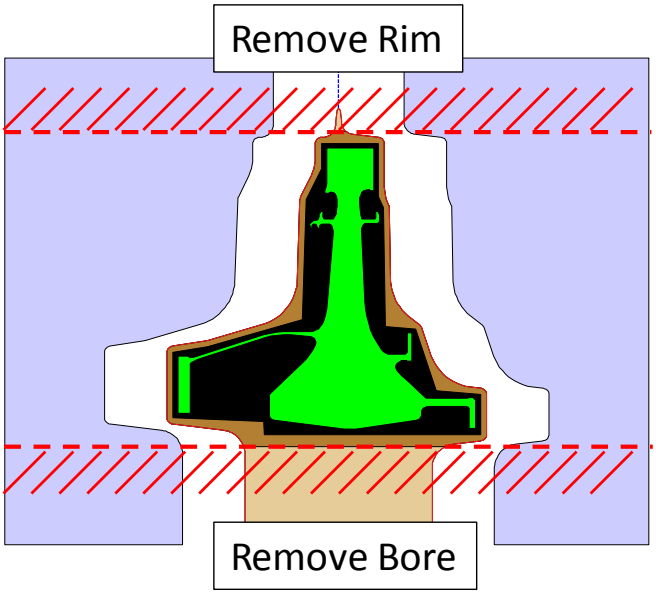

Figure 71 - Illustration to show the sections of the forging which should be removed to convert the black forging geometry into the heat treatment geometry

3.5.2.1. Locating The Minimum Depth of the Deepest Point Within Heat Treatment Geometry.

It was established in the literature review that increasing the control over the cooling rate within the component should improve the desired material properties and residual stresses. The easiest way in which we can assess the control of cooling rates of the deepest point in the heat treatment geometry is to calculate the depth of the deepest point of the heat treatment geometry.

The deepest location in the heat treatment geometry will be defined as the point which lies furthest away from the perimeter of the 2D heat treatment geometry. This depth can then be used as an objective to minimise within the optimisation which will in effect be a surrogate for minimum material property values.

In order to locate the deepest point within the geometry a method for finding the deepest point is required.

There were two methods considered, these were:

- Expanding balloon – This simulates a circle growing within the HT geometry profile. As the circle makes contact with the profile, it will then grow away until the circle can no longer expand.

- Curve offset – By offsetting the perimeter of the HT geometry inwards in incremental steps (similar to contours), eventually the deepest point can be located.

The curve offset method was originally the preferred method as it would locate the global deepest location within the HT geometry. Code written for offsetting the forging geometry already existed but proved to be unstable when offsetting internally due to the complex intersections which were created. The current design of the software was slow (0.2 seconds per offset) and would have required around 10 to 50 iterations within an optimiser to pinpoint the centre. In order to improve the robustness of the curve offset code, a greater number of analyses would have needed to be performed and would therefore take longer to run.

With compute time of the entire optimisation being an objective to ensure the optimisation takes hours rather than weeks, the expanding balloon method was looked into as a potential quick method, but with a risk that only the local minimum would be found. For this specific problem, the forgings are generally considerably larger at the bore than the rim. For this reason, with a sensible starting location this method would almost always find the cob's deepest location. This meant that the expanding balloon method would be most suited.

3.5.2.2. Expanding Balloon Method

The method is based on trying to expand a circle within an object which is constrained by a perimeter of the HT geometry. The method is described next with accompanying illustrations provided in Figure 72:

1. Position the centre of the circle at the centroid of the component (2D axi-symmetric) (Figure 72A). Note the example uses a random start point (not the centroid) to help illustrate this method.
2. Increase the diameter of the object until it contacts the perimeter (Figure 72B)
3. Keeping the contact point of the circle fixed, un-constrain the centre of the circle to the circle diameter to continue growing until it makes contact for a second time. (Figure 72C & Figure 72F)
4. To allow the circle to continue growing the point contacts must change to sliding joints with the freedom of moving along the tangencies of the original contact points. (Figure 72D, Figure 72E & Figure 72G)
5. When a third contact is made we need to check if the circle is still contacting the original two edges (as the circle was given the freedom to follow the tangent edges).

If it is concluded that only two points remain in contact, then go back to step 4 or if only the new point is in contact return back to step 3. (Figure 72F)

6. Before accepting these three points there are three rules which must be obeyed to fully constrain the circle:
    a. There are at least three points contacting the circle.
    b. The angles between all three points must be less than 180 degrees from each other.

    If two consecutive points have an angle greater or equal to 180 degrees then the third point does not provide any restraining force so is therefore discarded, when this occurs you move back to step 4. (Figure 72D)
7. The circle is fully constrained therefore the local minimum has been located. (Figure 72H)

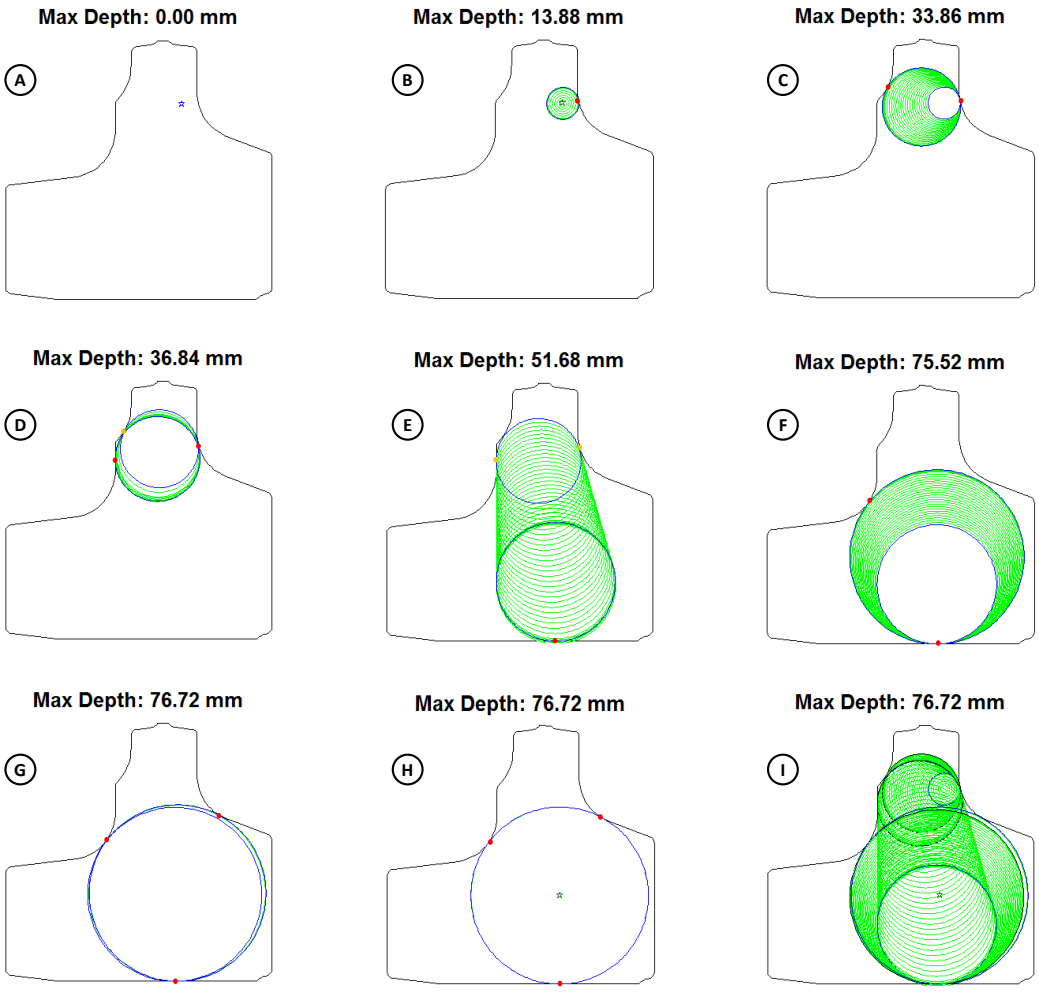

Figure 72 – Locating the deepest point within a geometry from a random starting point.

The method described locates the minimum depth of the deepest point within the heat treatment shape for typical forged disc geometries.

3.5.2.3. Optimisation

To see the impact of designing the RI geometry to minimise the 'maximum depth' (deepest point) of the heat treatment shape we adapt the optimisation loop created in Section 3.4 to include an additional objective of minimising the 'maximum heat treatment depth' ('Max HT Depth'). This additional objective has been added to the global optimisation stage of the genetic algorithm and has been added as a penalty function within the local gradient based optimisation as shown in the following equation:

$$\text{Objective} = \text{RI Geom Mass} \times \left(1 + 100 \times \left(\sqrt[3]{1 - \text{PerScan4x}} + \text{Scan Penalty} + \frac{\text{Max HT Depth}}{60}\right) + \text{Manu Penalty}\right)$$

This equation is the same as was described in Section 3.3 but with the addition of a new heat treatment depth penalty. The maximum HT depth has been non-dimensionalised and scaled to make it a significant weighting for the optimisation objective.

The adapted workflow can be seen in Figure 73 below.

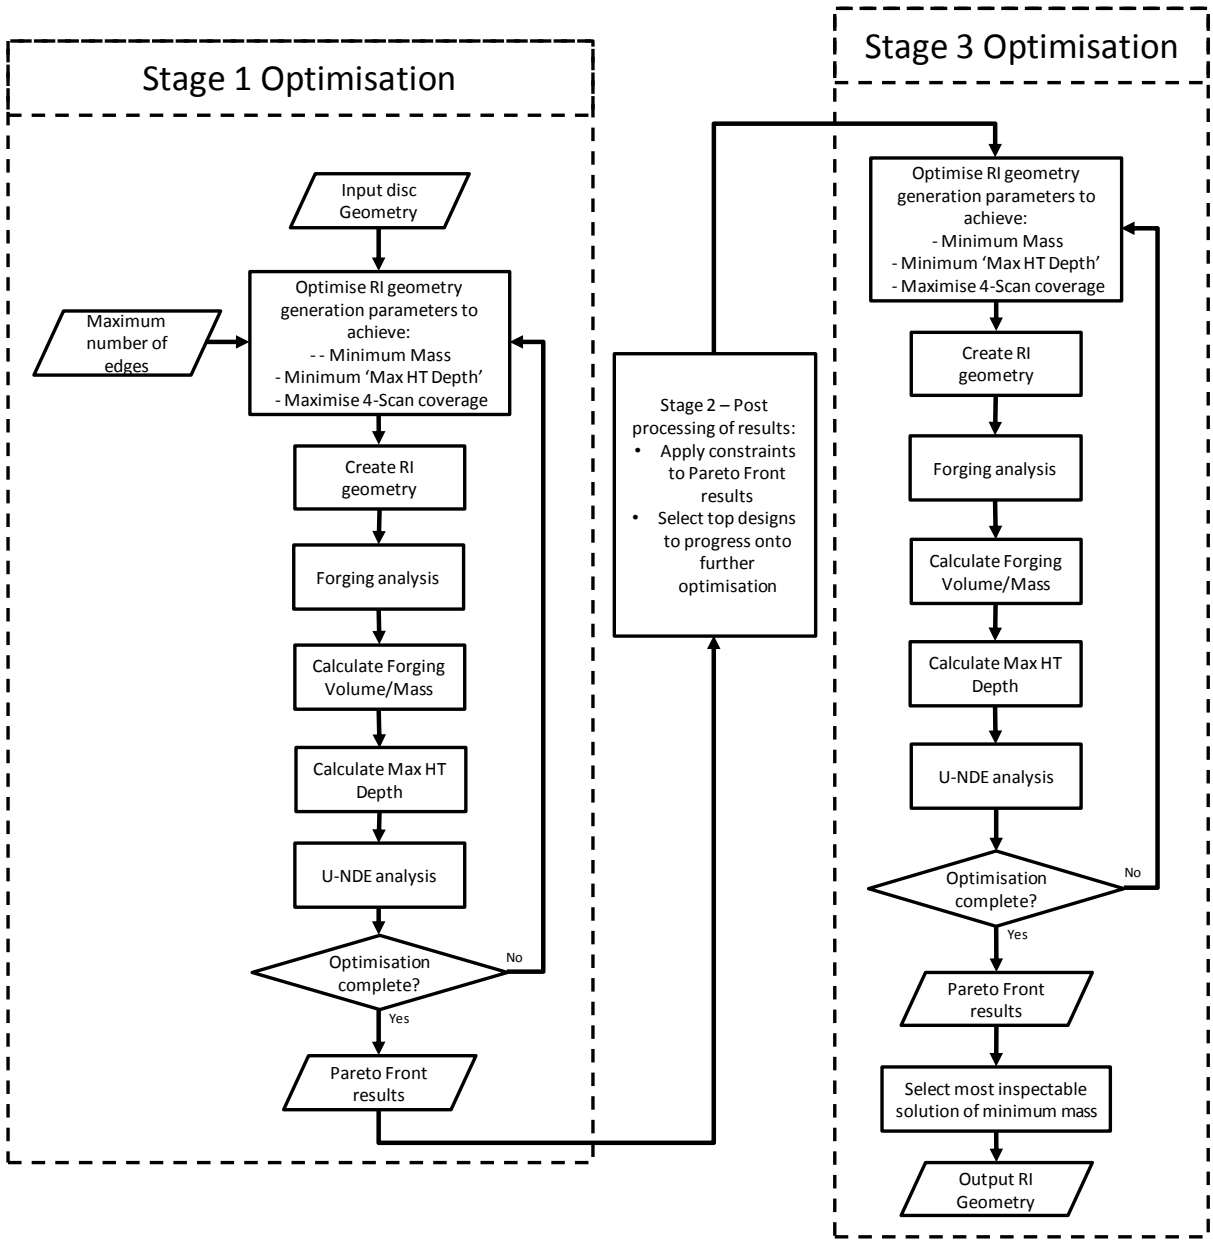

Figure 73 - Workflow used for Section 3.5 optimisation.

3.5.3. RESULTS

The results of the multi-objective optimisation for minimum 'max HT depth' are shown in Table 12 alongside the results from previous sections. The results shown are the best out of 5 repeat optimisation runs.

Table 12 – Results of Optimisation including results of the previous section for comparison.

	Result Source		RI Geom	U-NDE	BF	HT
	Description	Section	$Mass_{norm}$	(%)	$Mass_{norm}$	$Depth_{norm}$
CS-1	Existing RI geometry design*	-	100.00	93.67	135.62	100.00
	Min RI geometry Mass without U-NDE	3.2	92.29	94.78	133.36	94.16
	Min RI geometry mass with U-NDE	3.3	98.38	99.99	136.27	**94.88**
	Min BF mass	3.4	99.14	100.00	**136.06**	95.75
	Min BF mass with min HT depth	3.5	101.65	100.00	138.81	96.31
CS-2	Existing RI geometry design	-	100.00	100.00	134.86	100.00
	Min RI geometry Mass without U-NDE	3.2	87.31	98.56	128.23	94.76
	Min RI geometry mass with U-NDE	3.3	90.64	100.00	**129.49**	96.43
	Min BF mass	3.4	96.95	100.00	131.95	101.21
	Min BF mass with min HT depth	3.5	95.45	100.00	129.97	**94.48**
CS-3	Existing RI geometry design	-	100.00	99.88	132.46	100.00
	Min RI geometry Mass without U-NDE	3.2	82.31	96.10	118.92	90.26
	Min RI geometry mass with U-NDE	3.3	86.18	100.00	**123.66**	92.15
	Min BF mass	3.4	89.61	100.00	124.19	94.03
	Min BF mass with min HT depth	3.5	91.12	100.00	126.33	**91.15**
CS-4	Existing RI geometry design	-	-	-	-	-
	Min RI geometry Mass without U-NDE	3.2	83.86	98.68	123.75	90.15
	Min RI geometry mass with U-NDE	3.3	94.26	100.00	132.55	99.39
	Min BF mass	3.4	89.95	100.00	**126.69**	**90.78**
	Min BF mass with min HT depth	3.5	90.72	100.00	127.22	90.82

* The CS-1 existing RI geometry is based only on the final stage RI geometry design

The results in Table 12 show that by including the heat treatment geometry's maximum depth value as one of the objectives has in general had a positive effect on the max HT depth. The CS-3's maximum HT depth has improved upon all previous sections and has come within 0.5% of the minimum found for Section 3.2 (Minimum RI geometry with no inspection) which could be considered as the minimum theoretical HT depth.

This additional objective has meant that the minimum black forging geometry has increased relative to Section 3.4 as shown in Figure 74.

CS-3 Section 3.4 result for minimum forge mass	CS-Section 3.5 result for minimum HT depth & mass
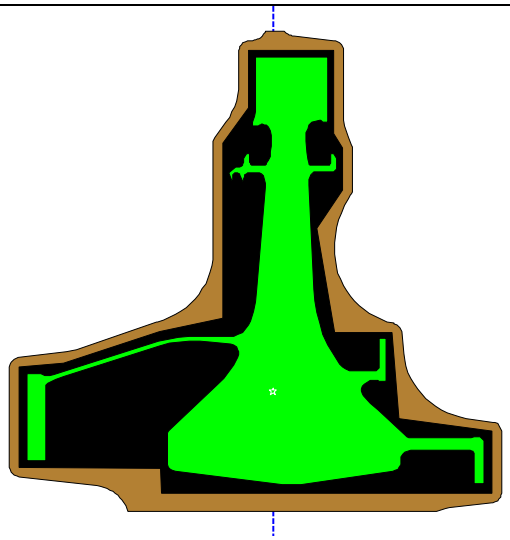	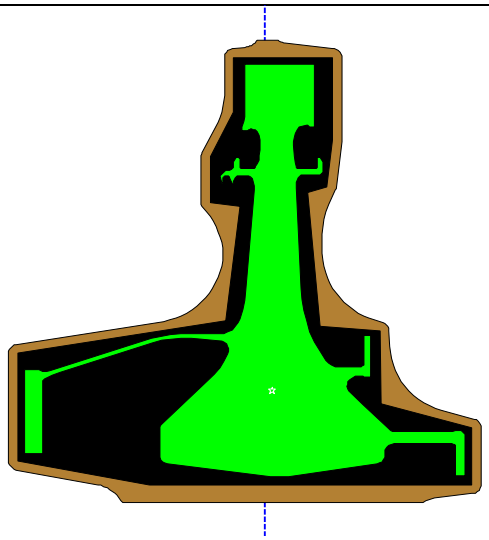

Figure 74 - Comparison of results from Sections 3.4 and 3.5. The white star in the centre of each heat treatment geometry represents the point of maximum depth.

In trying to minimise the HT shape, the RI geometry needs to get as close to the finished part at the intersection between the front drive arm and diaphragm and the rear part of the diaphragm.

It appears that in general the minimising the HT depth is having a positive effect on the Forging mass; although it isn't performing as well as minimising the RI geometry mass, with the exception of CS-4 Section 3.3 results.

3.5.4. DISCUSSION

This section was created to look at how the design of the RI geometry could affect the life and efficiency of the final part design. This is important because without knowing the impact on life of the component you may optimise to minimise the cost of the component but are having a detrimental impact on life cycle cost or weight of the finished component.

Summary of Findings

Minimisation of the maximum HT depth has a general positive impact on keeping the mass of the forging down. In most cases it hasn't reduced the mass of the forging down below that of results found in Section 3.3, suggesting that optimisation for minimum RI geometry mass is more effective at this task in its current setup.

Although minimisation of the HT depth may not produce a reduced forging mass, it does provide an additional objective which will encourage designs to be explored of minimum HT depth. As the optimisation will create a Pareto front of minimum forging mass against HT depth within the component, it is certain that for every forging mass outputted by the optimiser a design of minimum heat treatment depth will be selected. This will ensure the heat treatment quenching strategy has the opportunity to be better optimised to improve the desired material properties.

In reality, the differences in the HT depths produced optimisation have shown relatively small benefits. This is due to the objective of minimising the mass of the forging geometry is having a positive impact on the maximum HT depth. It is possible that the few percent of depth reduction could improve the minimum material properties sufficiently to pass an improved lifing threshold making it still worthwhile.

Academic Implications

The use of such a simple surrogate model to represent material properties has the potential to quantify material properties and worst case residual stress distributions. Despite not having a direct relationship for this surrogate, the method is still an important output and objective to the optimisation ensuring that these properties are still considered.

Business Implications

The method developed in this section is a very simplistic surrogate for heat treatment properties and only provides an insight into the benefits of the heat treatment process and potential improvements, but will not directly impact further optimisation.

Limitations and Directions for Future Research

The method developed in this section assumes a single zone in which we would like to minimise the maximum heat treatment depth, and that the local maximum HT depth is the global maximum depth within the forging, which is likely to be located near the centroid of the heat treatment shape.

In this section we have looked at creating heat treatment geometry in the cheapest possible way which was to only remove the bore and rim of the black forging. If is possible to improve material properties further by reducing the heat treatment geometry by turning the black forging geometry down to a geometry greater than the RI geometry. With additional optimisation and heat transfer coefficients the entire heat treatment geometry could be optimised in-between the forging and RI geometry to maximise both the components minimum material specs and improve residual stresses. Including greater

detail into this part of the loop would allow the finished component material properties to be assessed and fed directly back into the disc design. This could be accomplished before the forgers or even the manufacturing engineers have looked at the geometry therefore cutting out several months of uncertainty on the potential material properties and residual stresses. This step is becoming increasingly more important during the disc design stage as without knowing the residual stress on the finished part then the disc is not able to be optimised to its greatest potential.

This page intentionally left blank.

3.6. ANALYTICAL COST MODEL CREATION

3.6.1. INTRODUCTION

The final section in the chapter unites all the models created in Sections 3.2 to 3.5 to produce a detailed monetary cost for use within the optimisation loop. This is a big change from the mass estimates which were used as a surrogate for cost throughout the early part of this chapter.

RI geometry mass is generally a good parameter to generate a parametric cost model on as the parts mass is typically one of the major cost drivers (Masel, Young II et al. 2010). Basic surrogates are used because the cost models do not require a detailed definition of the manufacturing process (Caputo, Pelagagge et al. 2016). Optimising for minimum mass of the disc's RI geometry, and black forging geometry cannot give any certainty that the actual unit cost of the disc will be reduced. Furthermore low detail parametric cost models may not allow cost based comparisons between alternative designs (Caputo, Pelagagge et al. 2016). In general, the more manufacturing processes that have been modelled and costed accurately, the better the cost implications can be understood. Not only should a cost model provide a better cost estimate, but it will also allow both the disc design and manufacturing engineer to understand the cost drivers for their component. Knowing the cost break down of manufacture will allow them to further justify or discredit designs based on cost, which is especially useful when unit cost is of increasing importance.

Cost models within Rolls-Royce typically have an intense focus on the in house manufacturing costs as these are the costs which can be driven down. External costs which are either bought or subcontracted out have little understanding and knowledge in order to drive the cost models. Therefore, prices are generally scaled based on similar products which have a cost. Contractors and supplier also generally do not like to disclose detailed cost breakdowns to ensure their competitive edge is maintained.

Current cost predictions for forgings within Rolls-Royce use a historic database of trends that enable cost to be predicted from a RI geometry size and mass. Although this gives a good estimate it doesn't provide the detail in which to reverse engineer the RI geometry design to minimise cost. This also means that the designer has no direct guidelines on cost reduction on which to drive this design.

During the design phase the cost predictions are again based on similarities between previous designs, but until slim line RI geometry has been designed and issued an accurate mature cost estimate cannot be created. Slim line RI geometry designs are normally created

after the final production disc design has been issued which could be up to two to three years after the initial preliminary design started.

None of the research covering disc manufacturing optimisation used should-be-cost models with actual production predictions to estimate production cost whilst utilising ultrasonic NDE to ensure a workable design.

In order to create an analytical cost model, it is imperative that enough detail is included to ensure all cost drivers are covered.

This section aims to cover the following:

- To create a generic disc forging analytical cost model.
- To utilise the analytical cost model to drive the RI geometry design.
- To determine the benefits of using a detailed cost model for optimisation.

Another key aim for this section is to find out if analytical cost models are required to find cheaper RI geometry designs. Is it possible to design a part cheaper than using a far simpler surrogate model such as forging mass?

The cost model will be setup for a generic isothermally forged HP turbine disc and will be applied to the same examples that have been shown throughout this thesis.

This section is made up of four parts. First we look at the generic cost breakdown of the HP Turbine disc and identify the key processes and data which are required to create a 'should-be' analytical cost model. A skeleton cost model is then created. The next part goes into details of how the skeleton cost model can be populated using data produced by models developed within the section and reveals new models to help setup the required inputs. The penultimate part will validate the cost model before utilising this unique cost model to perform optimisation on the RI geometry to achieve the ultimate goal, minimise the RI geometry for unit cost.

3.6.2. METHOD – CREATING A DETAILED COST MODEL

Current Cost Models

Minimum cost will be the objective for the optimisation experiment at the end of this section therefore getting an accurate cost estimate is very important. The cost model needs to be sufficiently sensitive to ensure relevant parameters have a realistic impact on cost to produce an accurate result.

A Rolls-Royce cost model for use during design was provided as an initial starting point to build upon. The vast majority of the cost model deals with all of the in house manufacturing such as machining from the 2D axi-symmetric profile down to the final finished disc geometry. The cost estimate for the operations between the billet and the RI geometry is a parametric model based on historical data. Cost break downs from suppliers may be of use to back calculate costs of manufacturing operations, but the provided data may be distorted to conceal true operational costs to help their competitive edge in the market. The billet mass is calculated using crude forging geometry model which scales with the component. This model was not used directly within this section as it is parameterised around a particular disc design; therefore, it was used only a guide for all the processes performed from the 2D axi-symmetric profile down to the finished 3D part design. The turning cost rates also guided the new turning models created later.

New Cost Model

The cost model has been developed using known cost drivers and key processes which occur during the production of the HPT Disc. Table 13 shows how the cost model will be broken down into key manageable sections.

Table 13 - Cost model breakdown.

#	**Process**			**Cost Model**
1	Raw Material	→	Billet	Material
2	Billet	→	Black Forging profile	Forging
3	Black Forging profile	→	Heat treatment profile	Turning
4	Perform heat treatment on heat treatment profile			Heat Treatment
5	Heat treatment profile	→	Rectilinear profile	Turning
6	Perform ultrasonic NDE on rectilinear profile			Ultrasonic NDE
7	Rectilinear profile	→	Disc profile	Turning
8	Disc profile	→	Disc with 3D features	Milling, Broaching etc.

Processes 1 to 6 are all costs which are managed by the forging supplier and are the most difficult to create and verify. Costs for process 7 already exists within the Rolls-Royce cost model therefore could be verified.

Process 8 costs are also known in great detail within Rolls-Royce cost models. Fixed finished disc geometry will not change depending upon the manufacturing route; because of this the actual cost value for process 8 will be set at a constant value.

The cost model shown in Figure 75 below is the high level breakdown of the cost. All the cost models were created using Vanguard Studio due to its clear and simple graphical user interface. These cost models were then manually programmed into MATLAB® for speed and simplicity in interfacing with the existing software modules.

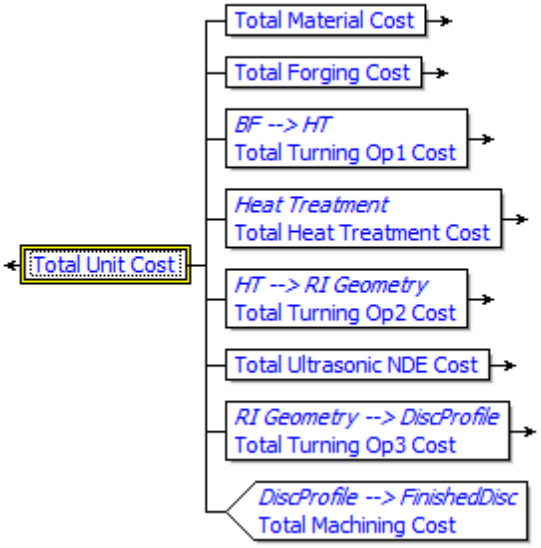

Figure 75 – High level cost model structure.

Graphically Vanguard Studio displays the cost tree very well, but diagrams which include all of the equations are too large to show within this thesis. For this reason, a graphical representation of the cost structure has been created using Vanguard Studio and will be shown in this section. The digital cost model location can be found within Appendix VII or the equations used to calculate the costs can be found in the cost model code provided in Appendix IV. The lower level cost breakdowns are covered next.

3.6.2.1. Billet Cost

The billet cost model can be created to reflect the reality of billet production. Stock billets which are produced in a mill are generally delivered and cut to length as required. For rare materials where few applications exist the amount of stock left over on the end of the stock billet may not be utilised and would need to be sold as scrap back to the mill. An illustration to show how the stock billet is split into the required stock billet sizes is shown in Figure 76 and the cost model accompanying it is shown in Figure 77. The problem found with this method is that a small change in forging mass may alter the number of billets that can be produced from a single stock resulting in a step change in price. Therefore, for optimisation purposes a straightforward price per kg is used. This is also illustrated in Figure 77.

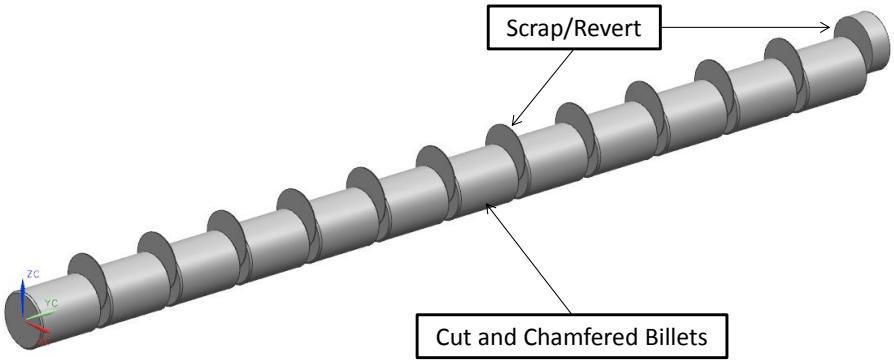

Figure 76 - Example stock billet distribution to produce a batch of disc forgings.

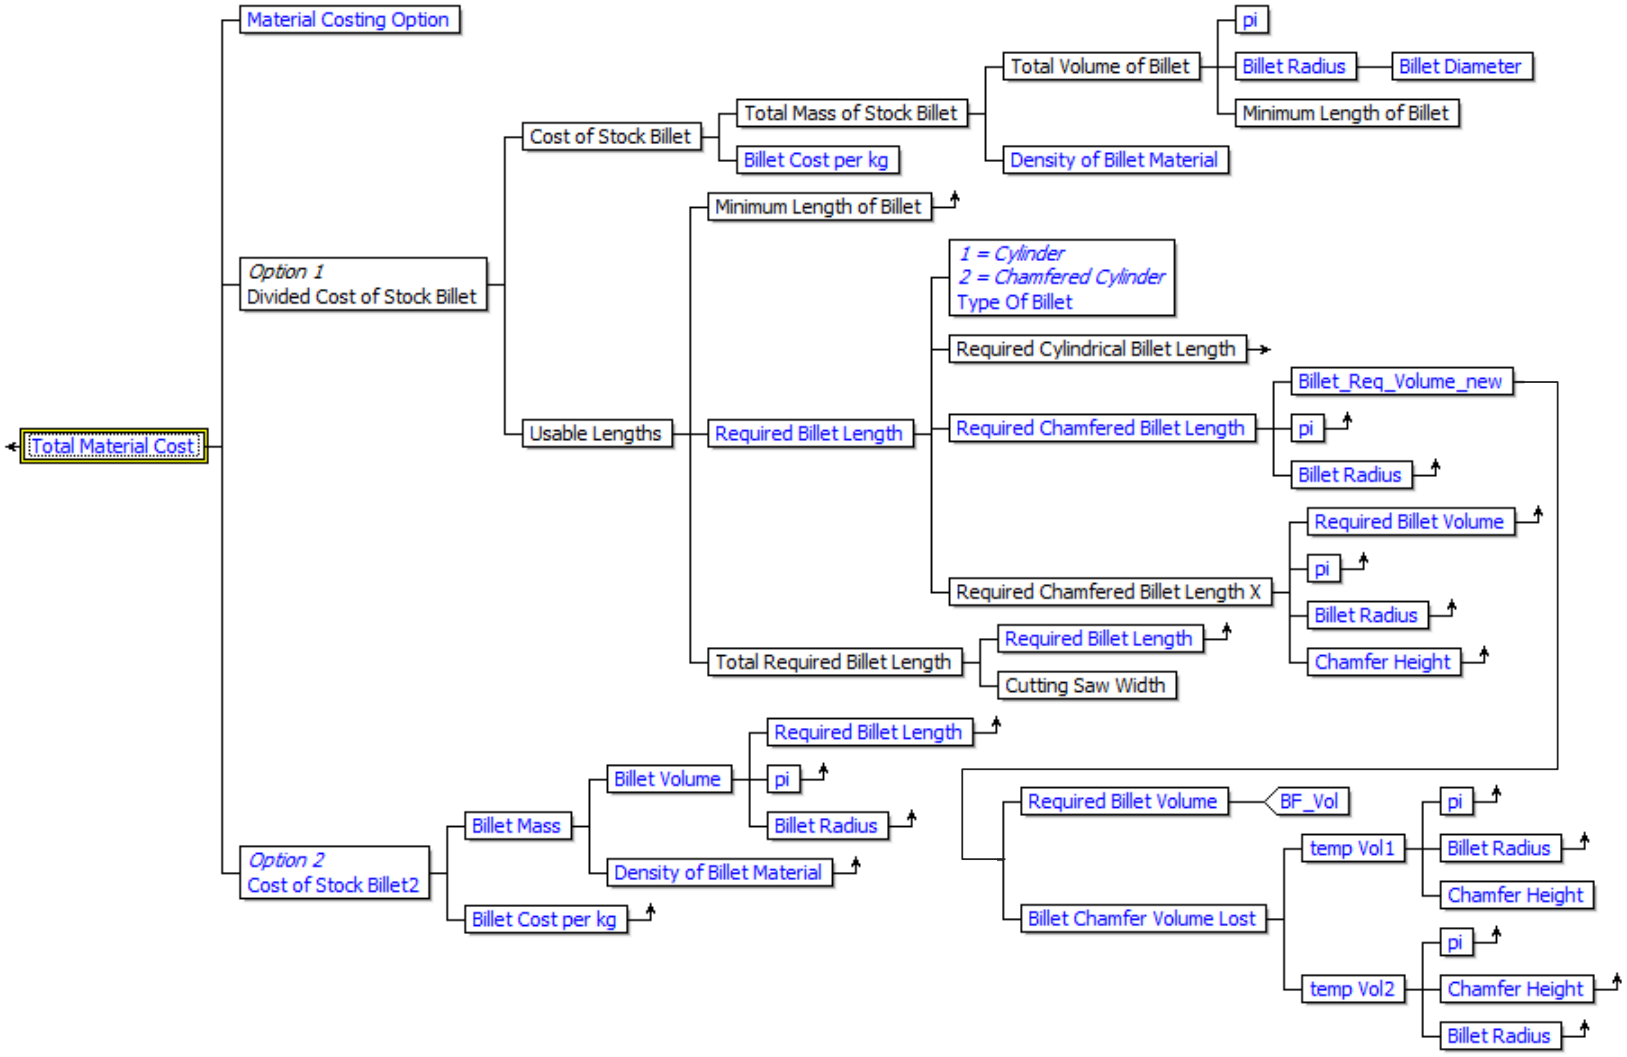

Figure 77 – Cost model structure for billet material costs.

The billet mass is calculated using the forging model which was created in Section 3.2. By using the input mass required to produce the forging, the billet length can be back calculated from the volume, stock billet radii and the material lost from chamfering the billet.

3.6.2.2. Forging cost

Forging cost can be one of the largest initial start-up costs for producing a disc. For discs which are isothermally forged, a Molybdenum based alloy is used to make the dies and could set a project back between a quarter and half a million dollars. Using the models created in Section 3.4, the top and bottom die material costs can be calculated using the cylindrical dimensions produced (see Figure 78 for an example die design output). Additional turning and machining costs can be calculated using a material removal rate and the finished die mass.

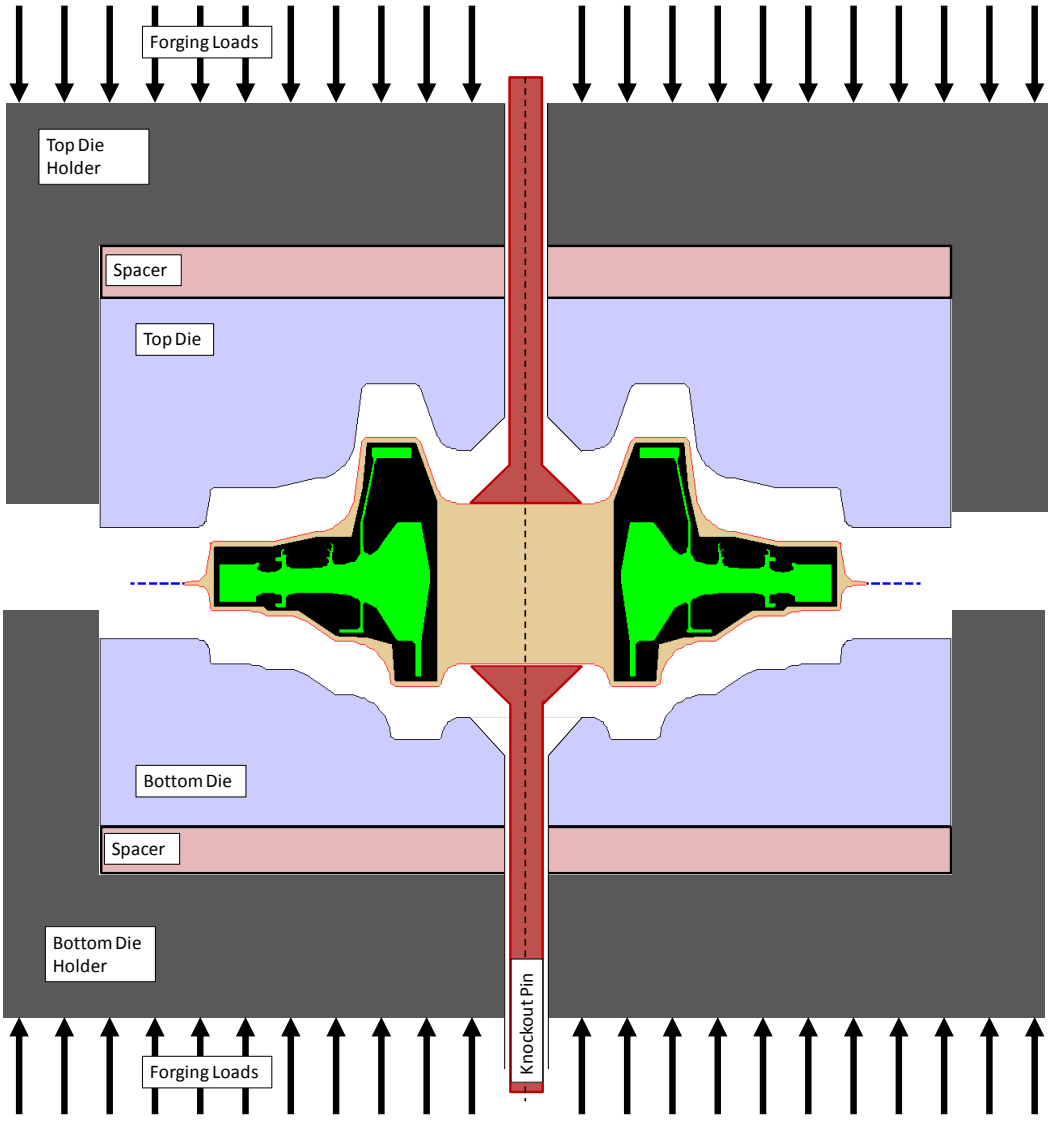

Figure 78 - Example forging dies and forging with additional context of an example isothermal forge layout.

Dies are generally designed to last the production life of the part; therefore, it is important to know the expected production run for an engine. The detailed cost model is shown in Figure 79 below and shows how the die overheads make up a considerable portion of the forging with an assumed production run of 150 discs.

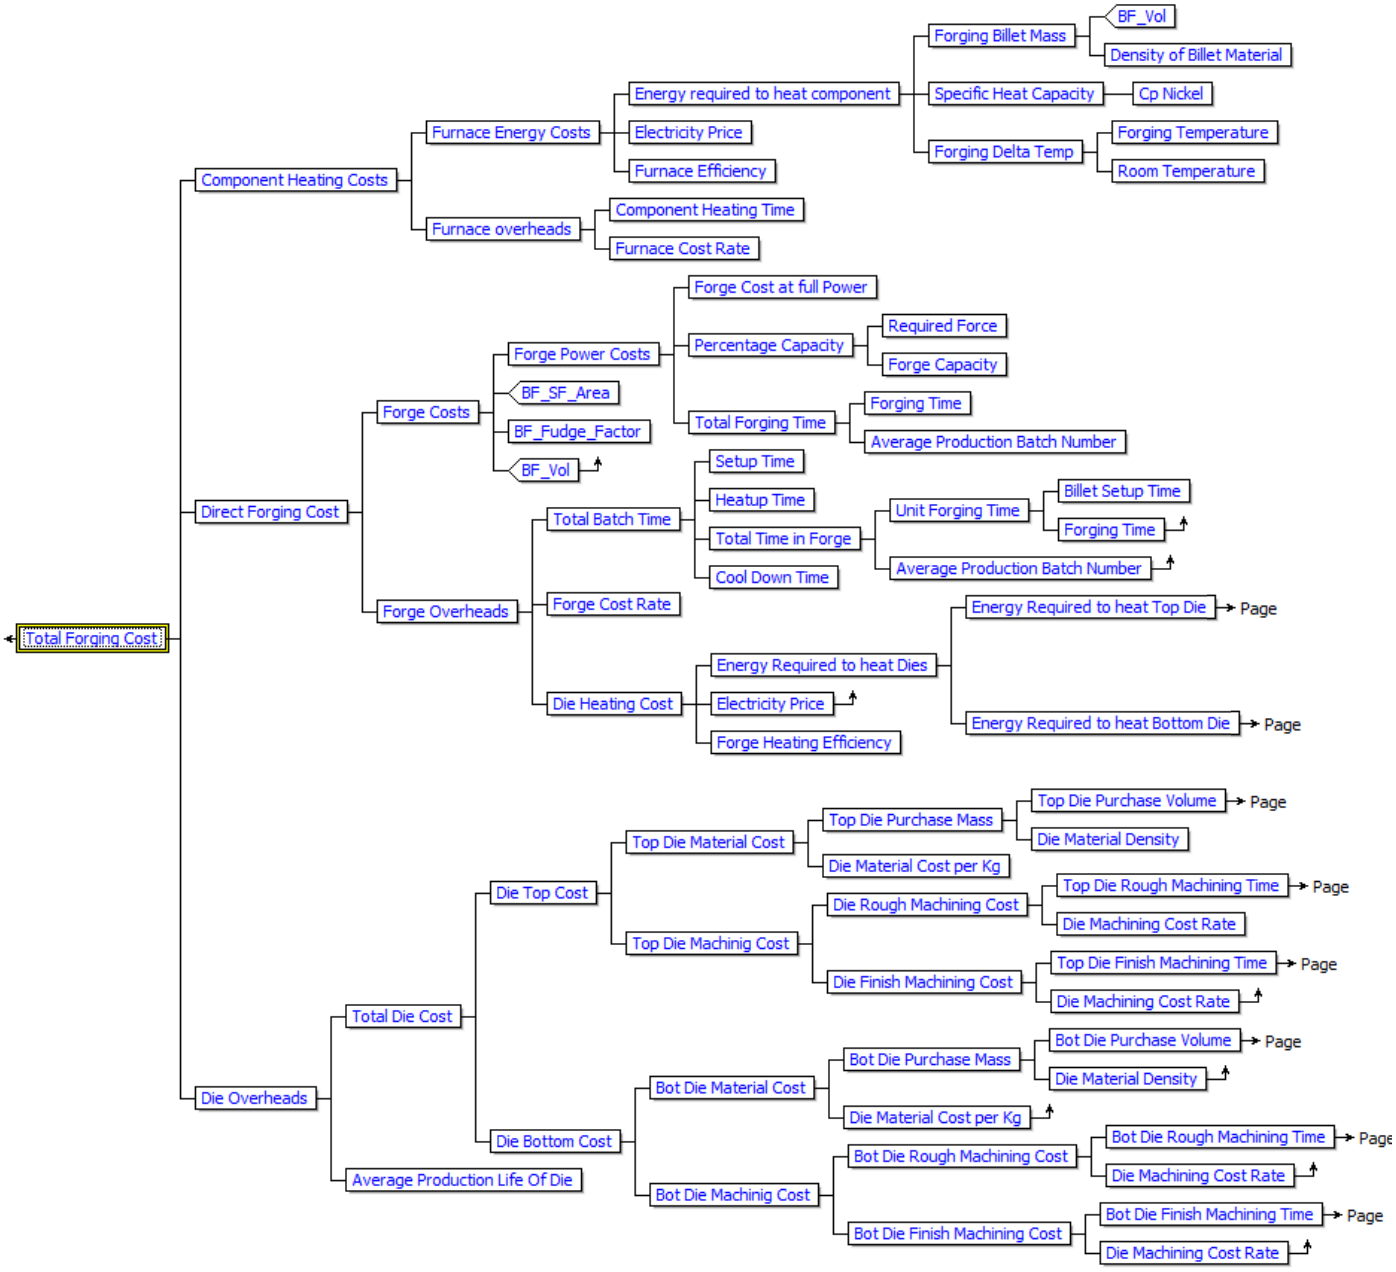

Figure 79 – Cost model structure for forging cost.

The total cost has been made up out of three sub costs:

- Heating
- Direct forging costs
- Forging overheads

There are likely to be more costs which need to be included to improve the accuracy of the cost estimate, but with the absence of the detailed forge cost data it is not possible. The forging costs which have been chosen are costs which can be estimated using data created in earlier sections of this Thesis with values approximated to produce realistic values.

3.6.2.3. Heat Treatment Costs

The heat treatment cost portion of the cost model will cover the following processes:

- Solution Heat Treat
- Aging
- Quenching

The actual changes in Heat Treatment shape have little impact on the heat treatment costs as changes in energy required to heat and cool the geometry is trivial compared to the machinery usage time. The depth of the component which was calculated in Section 3.5 is not actually required for this cost model, as the tightly controlled ramp in temperature increase is slow enough such that the outside temperature of the component isn't drastically greater than the deepest part of the component.

The cost model to calculate the cost of the Heat Treatment is as shown in Figure 80 below.

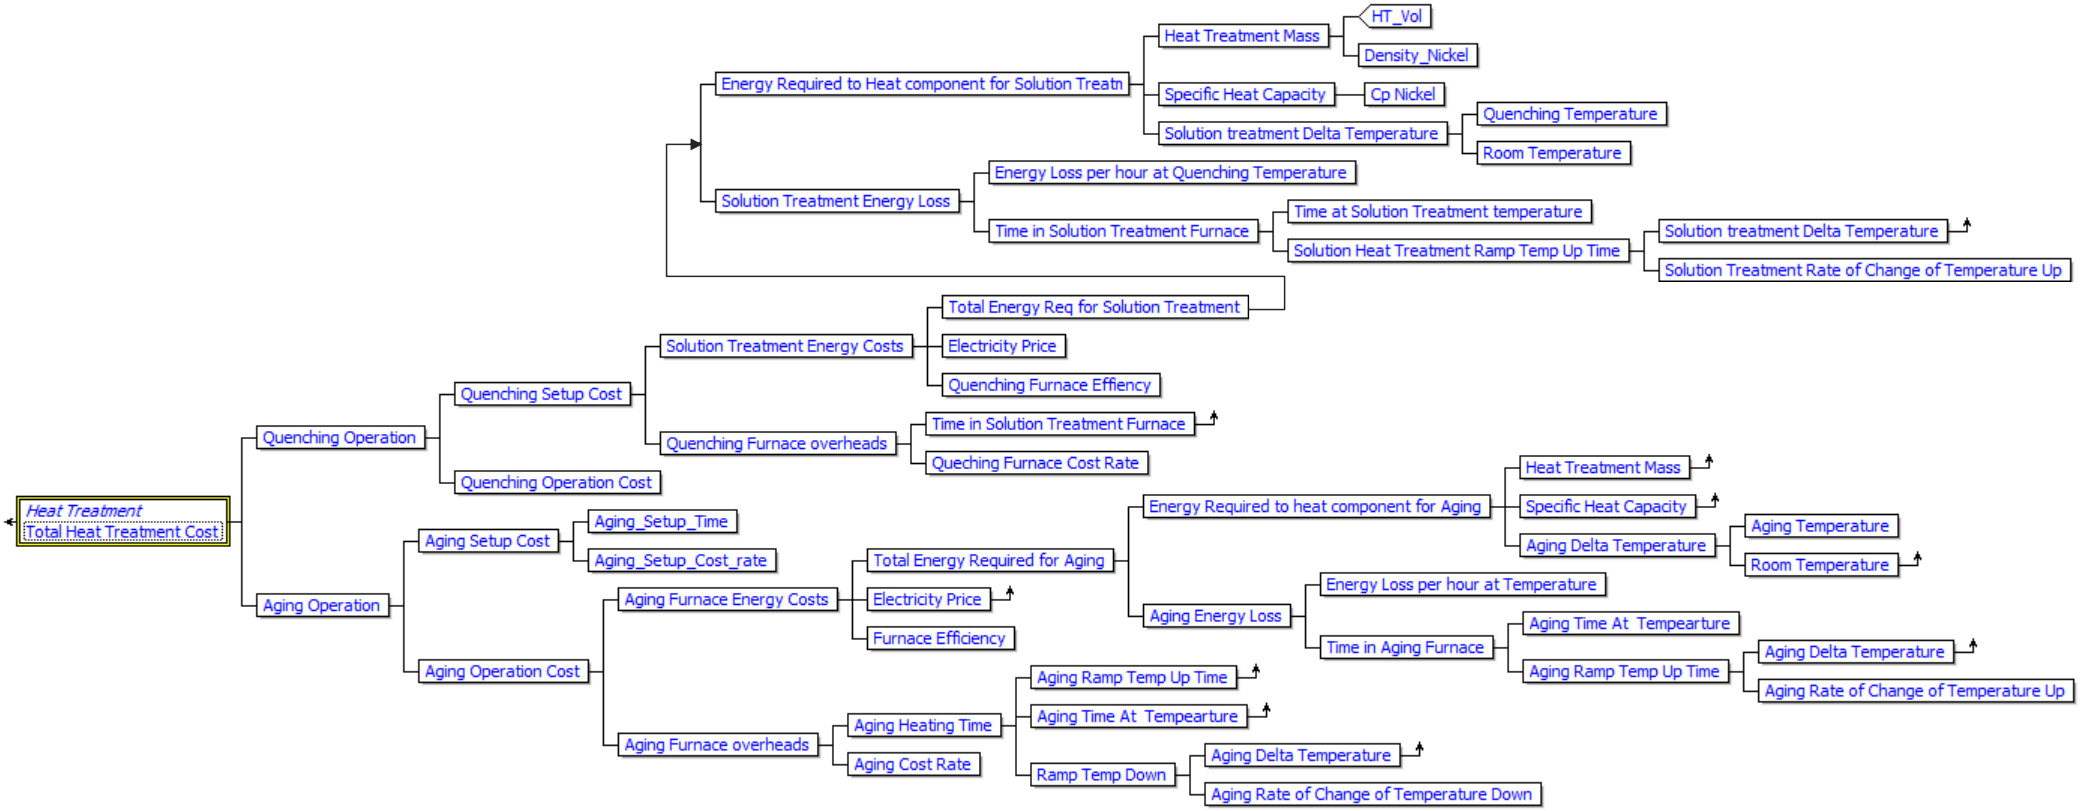

Figure 80 – Cost model structure for heat treatment cost.

3.6.2.4. Turning Costs

It was decided that a simple cost per unit volume of removal should be used which best calculates the cost of turning without having to simulate complex cutting paths. The turning of components is usually broken down into three stages:

- Roughing
- Semi-finishing
- Finishing

The roughing stage is where the bulk of the excess material is removed and is normally the most cost effective material removal rate.

Semi-finishing does not always take place, but is common during the manufacture of discs. This is because discs can move excessively during material removal resulting from the residual stresses formed during the heat treatment process. The semi-finishing cut will aim to remove material down to approximately the last millimetre so that the part in the best possible condition to perform the final finishing cut which will normally be to tight tolerances.

Finishing is the last cut performed and typically only removes the last half to one millimetre. This is process has the highest cost rate associated with it despite removing the least amount of material. However, it is the most important cut as any surface anomalies or non-conformance could result in the part to be scrapped.

Assuming these three turning operations have a single cost rate would be understandable if there were no complicated features on the components. In hard to reach locations and tightly confined areas it is increasingly difficult to manufacture as tools become more susceptible to vibrations and distortion. Therefore, another cost rate would be suitable to improve the accuracy of turning operations.

Small radii are also a cost increase for the component as to gain the required surface finish with a the smaller turning tool tip radius will result in a turning time which is inversely proportional

The cost rates for the generic turning stages aren't fixed rates as they depend on the tooling and tool fixturing available. In reality there are an infinite number of cost rates depending on the feature, surface finish requirement, material and ease of access. To work this out would require a highly complex knowledge database. Due to this complexity, only two

different cost rates will be created for each turning stage and will depend upon tooling access to features on the disc.

The first cost rate will be when preferred standard tooling and tool fixturing is utilised for material removal.

The second cost rate will be when special fixturing and or tooling are required in difficult to access locations. This cost rate will cover a broad range of situations but will simplify them as just the one cost rate. The method in which these areas of complexity will be located is by using tooling line of sight access and narrow feature areas. This is simplistic, but it does show to reveal many areas of complexity without having to create any new programs. These 'difficult' to access areas are shown later in Figure 81.

Basing the manufacturing cost in this simplified but effective manner allows the cost model to be built and populated easily using the already existing models.

In order to locate the regions of complexity a method similar to that used in the post processing of the ultrasonic scan results from Section 3.3 is used. This geometric Boolean algebra allows all six of the simplified regions to be located. Simple offsets are used to locate the finish and semi-finish zones, and then these are compared to the undercut regions to provide a series of cuts.

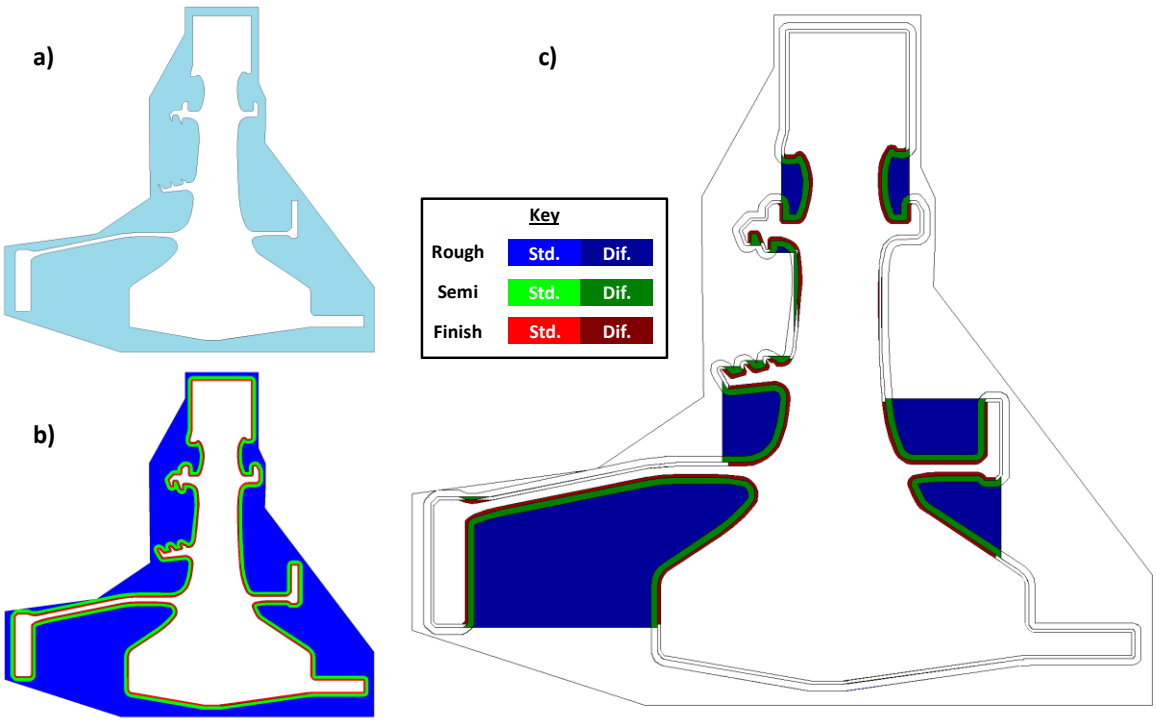

Figure 81 - a) Material highlighted in blue is the material to be removed from the RI geometry to the finished disc profile. b) Finish, Semi-finish and Roughing zones defined through offset curves; c): Difficult to reach, undercut regions are highlighted due to the potential access issues.

The undercut regions are highly simplified at this stage and could be made more accurate by determining the tooling access availability or by finding a model of complexity in the geometry, but for the purpose of this cost model it is considered sufficient.

Turning operation 1 contains both a turning operation to remove the rim of the forging and the trepan operation to remove the bore. The areas which are turned and trepanned are illustrated in Figure 82. The regions removed for this operation are created using a parameter to control a vertical offset from the RI geometry. A simple horizontal cut at both the Rim and Bore offset parameter locations of the disc is used.

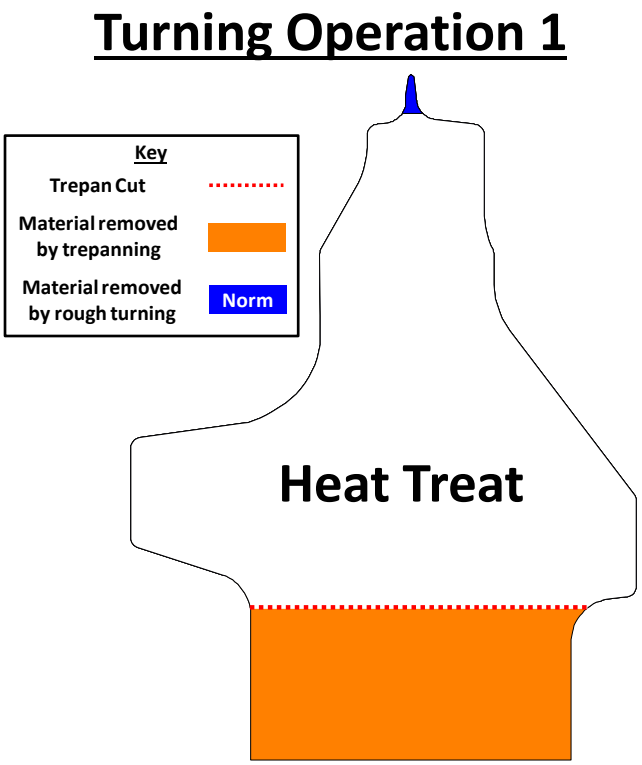

Figure 82 – Machined areas of the black forging to produce the basic heat treatment geometry.

The diagram shows the trepan location, the bore portion which is removed by the trepan operation and the location of the rough machining operation. The results of these two operations produce the heat treatment geometry (described in Section 3.5).

The cost model which calculates the cost of these operations are shown in Figure 83 below

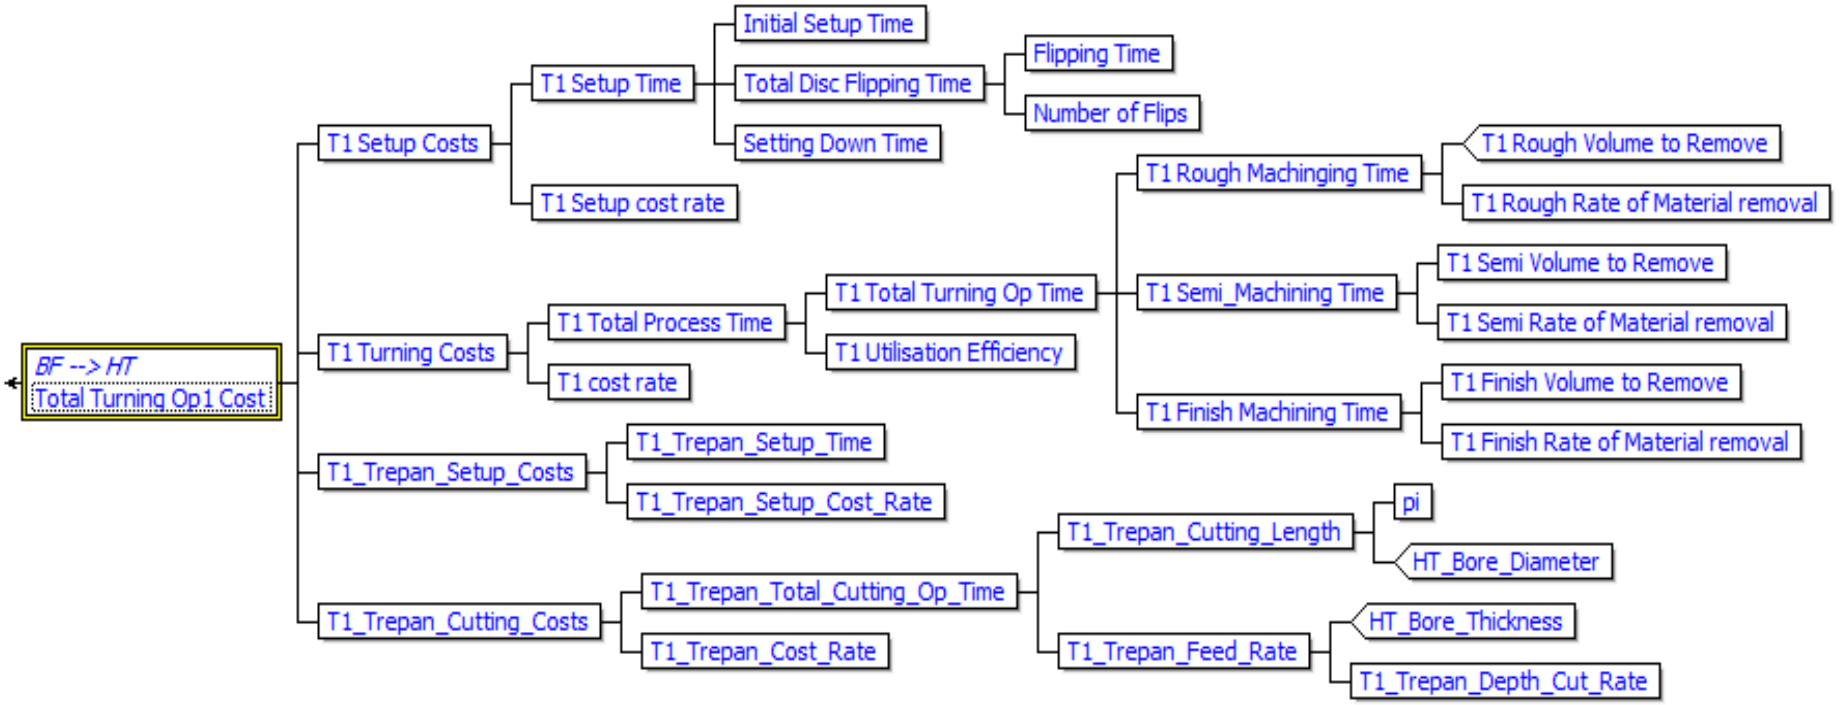

Figure 83- Cost model structure for turning operation 1 cost.

Both turning operation 2 and 3 do not require trepanning; only turning. Turning operation 2 which goes from the heat treatment geometry to the RI geometry is shown in the left illustration of Figure 84. This operation has clear access from both the left and right hand sides of the part therefore there are no hard to reach areas which would require special tooling rates for this turning operation. The illustration on the right within Figure 84 shows the final turning operation removing the material from the RI geometry to the final 2D axi-symmetric disc profile. This turning stage requires all six cost rates due to narrow access regions and non-direct access from the horizontal directions.

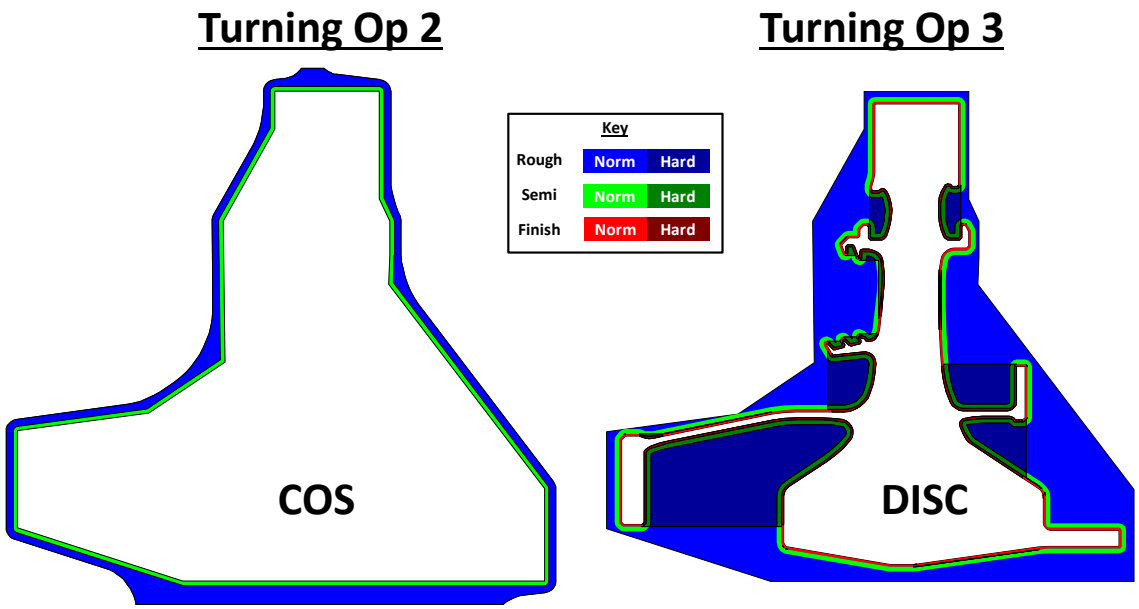

Figure 84 –Diagrams for both turning operations 2 and 3 showing the material removal rates applied to the turning cost model (note the 'COS' or Condition of Supply' is the RI geometry).

Turning operation 2 generally does not require any special tooling due to the simplicity of the RI geometry, therefore after the bulk of the heat treatment skin has been removed; a semi-finished condition is then required to obtain the surface finish necessary for ultrasonic inspection.

The remaining two turning operation's detailed cost breakdowns are shown in Figure 85 below.

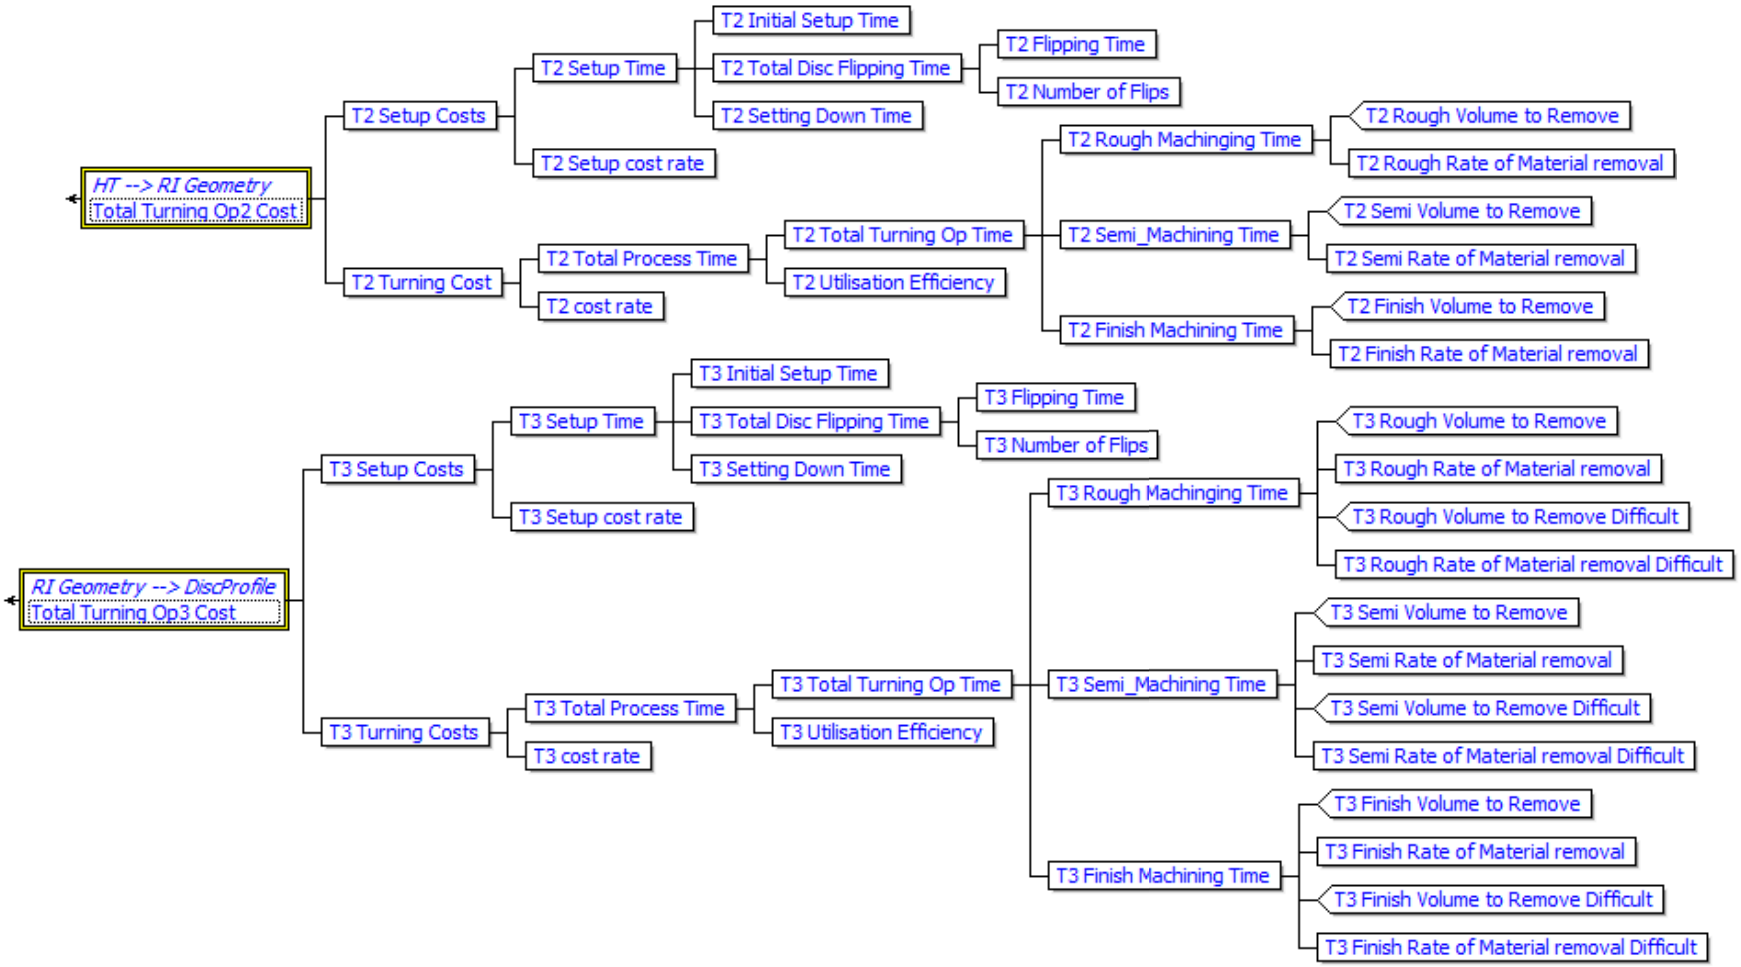

Figure 85 - Cost model structures for turning operations 2 and 3 costs.

In both cost models, the cost is split into two contributions, the turning costs and the setup costs. The turning costs are based on a single machine with a fixed cost rate. The usage of the machine is dependent upon the operation being performed.

Due to the way in which the RI geometry is created, no geometry is created which will require special tooling to remove difficult to access regions. For this reason, only standard removal rates are modelled in this cost model for turning operation 2.

The material removal rates are not only dependent upon the turning operation (i.e. roughing, semi-finishing or finishing) it is also dependent upon the tool and tool fixture being used. Difficult to reach areas of the disc such as undercuts are likely to require special kinked tool fixtures to avoid contact with the part. Narrow features may also require special fixtures. Despite utilising special tool fixtures, the dynamics of the tool will be different due to the reduced rigidity of the tool fixture and therefore would usually require a reduced material removal rate.

3.6.2.5. Ultrasonic Costs

Despite Section 3.3 suggesting only portions of the RI geometry are inspected, in reality the entire part is scanned to improve the likelihood of detecting defects. Although there are surface areas of the RI geometry which are unable to be inspected from certain angles (due to geometry obstructions), the cost model does not factor these in as these non-inspectable surfaces are insignificant compared to the inspectable surfaces. Figure 86 shows the cost model used to calculate the ultrasonic NDE inspection costs.

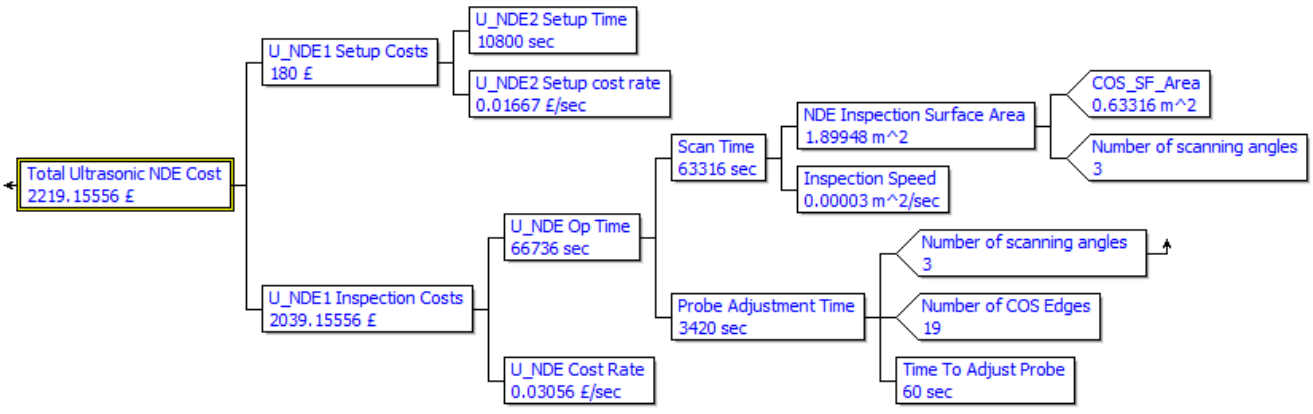

Figure 86 – Cost model structure for ultrasonic non-destructive examination cost

3.6.2.6. Finishing Costs

The remainder of the value-added costs to take the HPT disc down from a 2D disc profile to a finished disc is made up of the following main costs:

- Inspection (visual, binocular, die-penetrant, X-ray, magnetic, etc.)
- Machining (bolt holes and other 3D features)
- Broaching (blade slots)
- Surface Treatments (chemical cleans, washing, drying, etc.)
- Surface Finish (polishing, shot peening)
- Additional processes (inertia welding)

These manufacturing processes which are performed from the 2D axi-symmetric profile down to the finished part have not been included in the cost calculation. They have not been included because for a given 2D disc profile, the manufacturing method will not change, therefore this cost can be considered fixed for the purposes of the RI geometry optimisation.

3.6.2.7. Cost model Analysis

For a given disc design, three main cost drivers stand out which will vary considerably from the RI geometry parameters:

- Turning operation 3
- Turning operation 2
- Input billet mass

To lower the total cost all three of these sub-costs need to be minimised. Minimising the Input billet material requires the billet forging mass to be reduced and to lower the two turning costs the total amount of material to be removed needs to be reduced. All three of these cost drivers complement each other, giving the impression that lowering the forging mass should be the main objective. With an assumption that the cost rates are identical to perform turning operations 2 & 3, it would not matter which turning operation removed the most material. In reality, this is not the case as these two turning operations would generally be performed by different companies therefore it would be highly unlikely that the cost rates would be identical. For many nickel based superalloys, the turning speeds and feeds for optimum material removal could be proprietary; therefore, the material removal cost rates are likely to vary considerably. To ensure this is taken into account, two different cost models have been created to highlight differences between cost break downs between the two turning operations.

For cost model A the highest material removal rate is during turning operation 3 and in cost model B, the highest material removal rate is during turning operation 2. Experimenting with both of these cost models will allow us to see how the cost model setup could be vital in correctly cutting costs.

3.6.2.8. Optimisation

The cost model is the final addition to the optimisation loop which has been developed throughout this chapter. The objective previously (in Sections 3.2, 3.3, 3.4 and 3.5) was to minimise mass which was used as a surrogate for cost. Now that that a cost model has been created we can now replace the 'cost' surrogate for the actual predicted unit cost.

The work flow from Figure 73 Section 3.5 has been adapted to include the unit cost prediction and is shown in Figure 87 below. The three objectives for Stage 1 global optimisation are:

- Minimise manufacturing unit cost.
- Maximise ultrasonic NDE inspection.
- Minimise the maximum heat treatment depth.

Although the maximum heat treatment depth has not yet been quantified, it plays an important role in obtaining a greater variety of results to build the Pareto front. We can be confident that for a given low cost design with full inspection, the smallest maximum HT depth will be selected.

Stage 3's local optimisation objective has also changed relative to Section 3.5's objective function. It now reflects the new cost objective in place of the old black forging mass objective:

$$\text{Objective} = \text{Disc Unit Cost} \times \left(1 + 100 \times \left(\sqrt[3]{1 - \text{PerScan4x}} + \text{Scan Penalty} + \frac{\text{Max HT Depth}}{60{,}000}\right) + \text{Manu Penalty}\right)$$

Due to the fact that the heat treatment function has not yet been directly related to the life of the part, the optimisation weighting of the maximum HT depth has been reduced for this final Stage 3 local optimisation run (relative to that used in Section 3.5).

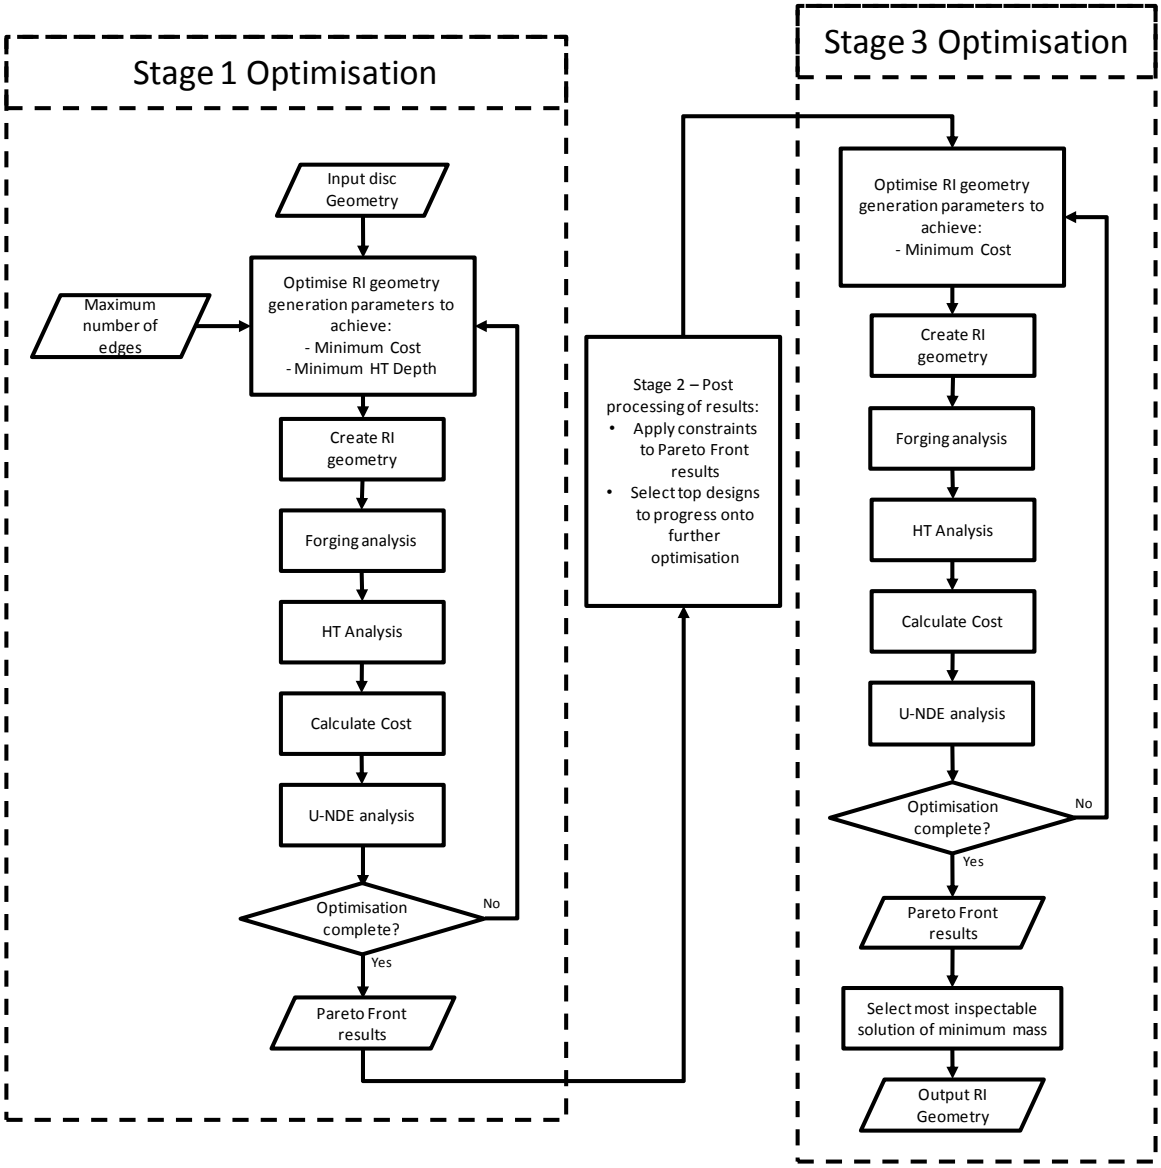

Figure 87 – Optimisation workflow with the objective of minimising component unit cost.

Stage 1 optimisation was run with a population set of 1000 and a generation number of 80. The local optimisation performed in Stage 3 allowed a maximum of 2400 evaluations.

Two experiments were run, 'cost model A' had a lower cost rate at the sub-contractors relative to Rolls-Royce and 'cost model B' had a lower cost rate at Rolls-Royce relative to the sub-contractors to see if this had a considerable impact on the results.

3.6.3. RESULTS

The total time to run this loop on a 12-core machine took between 16 and 24 hours and involved no interaction during the run. Each case study was run 5 times so that the best results could be used for comparison.

2.1.1.1. Results Overview

The results for CS-1, CS-2, CS-3 and CS-4 HPT discs are shown in Table 14 below. All designs from previous sections are also shown in Table 14 to allow a direct comparison between the different RI geometry design strategies.

Table 14 - Results of the optimisation for both 'cost rate A' and 'cost rate B'. Cells with a grey background depict results which were not available at the time of the original section optimisation so has been calculated post optimisation to allow direct comparison of results.

	Result Source		RI Geom	U-NDE	BF	HT	Rate_A	Rate_B
	Description	Section	$Mass_{norm}$	Coverage (%)	$Mass_{norm}$	$Depth_{norm}$	$Cost_{norm}$	$Cost_{norm}$
CS-1	Actual RI geometry design	-	100.00	93.67	135.62	100.00	10,000	10,000
	Min RI geometry mass without U-NDE	3.2	92.29	94.78	133.36	94.16	9,838	10,009
	Min RI geometry mass with U-NDE	3.3	98.38	99.99	136.27	94.88	9,979	10,000
	Min BF mass	3.4	99.14	100.00	136.06	95.75	**9,970**	**9,960**
	Min BF mass with min HT depth	3.5	101.65	100.00	138.81	96.31	10,078	10,051
	Min cost for cost rate A	3.6A	102.53	100.00	140.09	98.15	10,266	10,181
	Min cost for cost rate B	3.6B	104.16	100.00	139.57	97.59	10,236	10,076
CS-2	Existing RI geometry design	-	100.00	100.00	134.86	100.00	10,000	10,000
	Min RI geometry mass without U-NDE	3.2	87.31	98.56	128.23	94.76	9,735	9,994
	Min RI geometry mass with U-NDE	3.3	90.64	100.00	129.49	96.43	**9,800**	9,970
	Min BF mass	3.4	96.95	100.00	131.95	101.21	9,896	9,929
	Min BF mass with min HT depth	3.5	95.45	100.00	129.97	94.48	9,817	9,844
	Min cost for cost rate A	3.6A	94.96	100.00	131.10	97.12	9,867	9,936
	Min cost for cost rate B	3.6B	97.21	100.00	130.70	94.35	9,850	**9,836**
CS-3	Existing RI geometry design	-	100.00	99.88	132.46	100.00	10,000	10,000
	Min RI geometry mass without U-NDE	3.2	82.31	96.10	118.92	90.26	9,400	9,662
	Min RI geometry mass with U-NDE	3.3	86.18	100.00	123.66	92.15	9,627	9,870
	Min BF mass	3.4	89.61	100.00	124.19	94.03	9,645	9,775
	Min BF mass with min HT depth	3.5	91.12	100.00	126.33	91.15	9,735	9,873
	Min cost for cost rate A	3.6A	87.78	100.00	123.03	92.11	**9,443**	**9,682**
	Min cost for cost rate B	3.6B	90.41	100.00	125.00	93.45	9,525	9,723
CS-4	Existing RI geometry design	-	-	-	-	-	-	-
	Min RI geometry mass without U-NDE	3.2	83.86	98.68	123.75	90.15	9,158	9,793
	Min RI geometry mass with U-NDE	3.3	94.26	100.00	132.55	99.39	9,552	10,044
	Min BF mass	3.4	89.95	100.00	126.69	90.78	9,302	9,789
	Min BF mass with min HT depth	3.5	90.72	100.00	127.22	90.82	9,325	9,797
	Min cost for cost rate A	3.6A	91.58	100.00	127.31	91.48	**9,185**	9,692
	Min cost for cost rate B	3.6B	96.75	100.00	129.67	90.81	9,287	**9,673**

* The CS-1 existing RI geometry is based only on the final stage RI geometry design

When considering designs which are only 100 percent inspected, the results shown in Table 14 show that 5 out of the 8 case studies have improved on total unit cost after optimising solely on cost for the given cost rates. In both CS-3 and CS-4 the best designs were found for their respective cost rates with the exception of the CS-3 cost rate B where, cost rate A found the best solution. CS-2 had a successful result for the cost rate B case study, but the cost rate A case study only found the fourth best result. This was 0.68 percent higher than that found when the minimum RI geometry mass was investigated in Section 3.3. CS-1 had its best cost design found in Section 3.4 where minimum forging mass was targeted; the minimum cost objective from this section yielded results 2.87% and 0.53% higher for cost rates A and B respectively.

With the cost model now in place, a more detailed analysis can be performed between RI geometry design with (Section 3.3) and without (Section 3.2) ultrasonic NDE constraints. Ignoring the CS-4 result (which had an unusually high mass in Section 3.3 compared to all later sections) the mass increase from without, to with an ultrasonic NDE constraint was between 3.8% and 6.6%. The cost increase on the other hand varied between -0.2% and +2.4% for the first 3 case study discs (covering both cost rates).

Comparing the optimised unit cost values to Section 3.2 where no ultrasonic constraint was applied, it can be seen that the price increases between +0.3% and +4.3% for cost rate A and between -1.6% and +0.7% for cost rate B. Looking at heat treatment depth (as a surrogate for material properties) Section 3.2 underestimated the depth by between +1.5% and +4.2% for cost rate A and -0.4% and +3.6% for cost rate B.

3.6.3.1. Cost Rate A Results

Greater detail of the cost break down can be seen in Table 15 below.

Table 15 – Breakdown of costs for cost rate A (normalised to 'existing' geometry).

	Section	Total	Mat	BF	T Op 1	HT	T Op 2	U-NDE	T_ Op 3	Finish
CS-1	Existing	10,000	3,390	951	21	287	469	364	2,072	2,446
	3.2	9,838	3,334	952	22	287	525	347	1,925	2,446
	3.3	9,979	3,406	952	22	287	473	347	2,046	2,446
	3.4	9,970	3,401	951	22	287	457	348	2,058	2,446
	3.5	10,078	3,470	951	22	287	460	344	2,098	2,446
	3.6A	10,266	3,502	951	22	287	469	349	2,240	2,446
	3.6B	10,236	3,489	951	22	287	432	343	2,266	2,446
CS-2	Existing	10,000	3,112	1,044	23	317	404	301	2,090	2,709
	3.2	9,735	2,959	1,044	22	317	496	305	1,882	2,709
	3.3	9,800	2,988	1,044	23	317	465	302	1,953	2,709
	3.4	9,896	3,045	1,046	22	317	411	300	2,045	2,709
	3.5	9,817	2,999	1,043	23	317	401	301	2,023	2,709
	3.6A	9,867	3,025	1,044	23	317	425	300	2,024	2,709
	3.6B	9,850	3,016	1,043	23	317	385	300	2,057	2,709
CS-3	Existing	10,000	3,458	914	21	273	432	328	2,248	2,327
	3.2	9,400	3,105	916	21	273	497	325	1,936	2,327
	3.3	9,627	3,228	915	21	273	511	327	2,024	2,327
	3.4	9,645	3,242	916	21	273	462	324	2,081	2,327
	3.5	9,735	3,298	915	21	273	473	322	2,106	2,327
	3.6A	9,443	3,212	916	21	273	474	324	1,896	2,327
	3.6B	9,525	3,264	916	21	273	462	323	1,940	2,327
CS-4	Existing	-	-	-	-	-	-	-	-	-
	3.2	9,158	3,231	925	20	273	562	334	1,487	2,327
	3.3	9,552	3,460	923	21	273	533	339	1,677	2,327
	3.4	9,302	3,308	925	20	273	511	334	1,606	2,327
	3.5	9,325	3,321	924	20	273	506	334	1,619	2,327
	3.6A	9,185	3,324	924	20	273	492	333	1,492	2,327
	3.6B	9,287	3,385	924	20	273	446	335	1,577	2,327

It can be seen that the main source of cost reduction has come from material reduction which has a double impact on cost due to the reduction in value added cost in the combined turning operations.

In order to reduce the overall cost in the cost rate A scenario the RI geometry mass should be minimised as this would result in less cost during turning operation 3 which turns out to be expensive due to the high cost rate and poor turning efficiency compared to turning operation 2.

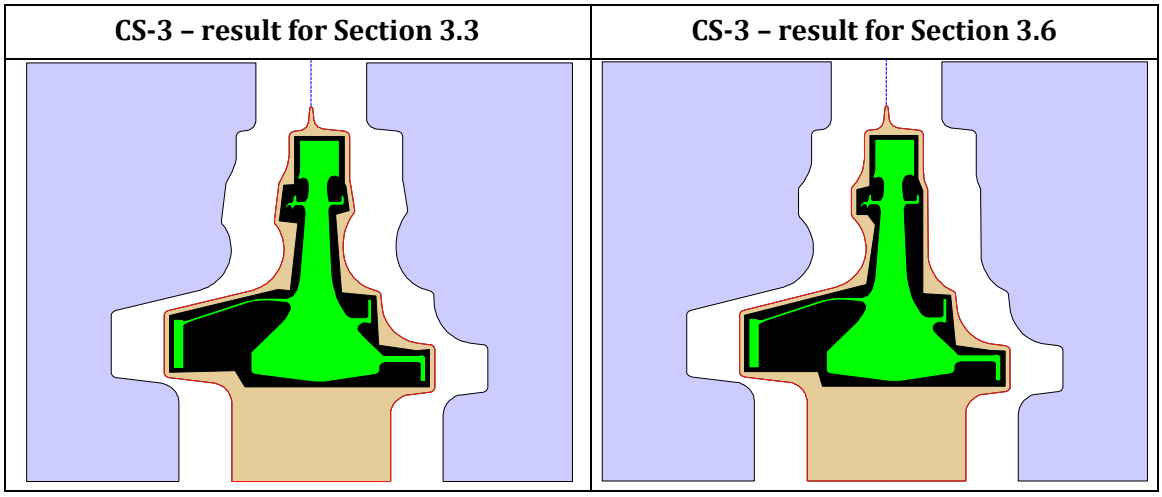

Figure 88 - Results of CS-3 for Section 3.3's minimise forging mass and Section 3.6's minimise cost optimisation using cost rate A.

Figure 88 shows how cost rate A has managed to save cost over Section 3.3 which aimed to minimise the RI Geometry mass. By having an objective for cost the RI geometry has been simplified requiring less material around the front side of the diaphragm due to the improved inspectability for a minimal cost increase. Figure 89 below shows how Section 3.6's cost rate A objective compares against Section 3.3's minimise RI geometry mass.

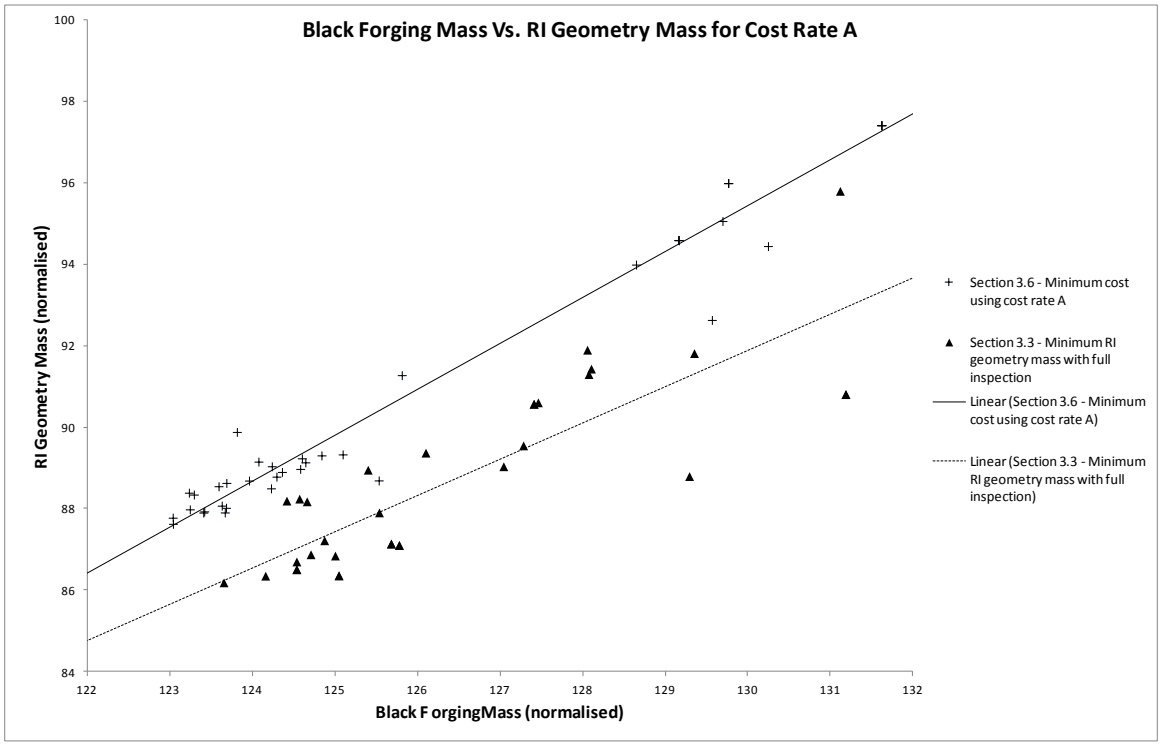

Figure 89 – Graph showing the linear fit (extrapolated) for all Section 3.3 and Section 3.6a acceptable results for CS-3.

The graph shows that all the lowest cost solutions for cost rate A in most cases have a greater RI geometry mass relative to Section 3.3 for a given black forging mass, a result very similar to that found in Section 3.4.

3.6.3.2. Cost Rate B Results

The results in Table 16show just how powerful changing the Turning cost rates are on the size of the RI geometry. Both CS-2 and CS-4 found the lowest cost design by optimising for minimum cost and shows a completely different result compared to all previous sections.

Table 16 – Breakdown of costs for cost rate B (normalised to 'existing' geometry).

	Section	Total	Mat	Forge	T Op 1	HT	T Op 2	U-NDE	T_ Op 3	Finish
CS-1	Existing	10,000	3,612	1,013	22	305	1,163	388	891	2,606
	3.2	10,009	3,552	1,014	23	305	1,305	370	833	2,606
	3.3	10,000	3,629	1,014	23	305	1,173	370	881	2,606
	3.4	9,960	3,623	1,013	23	305	1,133	371	885	2,606
	3.5	10,051	3,696	1,014	24	305	1,139	366	901	2,606
	3.6A	10,181	3,730	1,013	24	305	1,162	372	967	2,606
	3.6B	10,076	3,717	1,013	24	305	1,069	365	977	2,606
CS-2	Existing	10,000	3,351	1,125	24	342	1,007	324	910	2,917
	3.2	9,994	3,186	1,124	24	342	1,244	328	828	2,917
	3.3	9,970	3,218	1,124	24	342	1,164	325	857	2,917
	3.4	9,929	3,279	1,127	24	342	1,025	323	893	2,917
	3.5	9,844	3,230	1,124	24	342	999	325	884	2,917
	3.6A	9,936	3,258	1,124	24	342	1,061	323	887	2,917
	3.6B	9,836	3,248	1,123	24	342	959	323	900	2,917
CS-3	Existing	10,000	3,749	991	23	296	1,088	356	976	2,523
	3.2	9,662	3,366	993	23	296	1,257	352	853	2,523
	3.3	9,870	3,500	992	23	296	1,294	355	888	2,523
	3.4	9,775	3,515	993	23	296	1,166	351	910	2,523
	3.5	9,873	3,575	992	23	296	1,194	349	920	2,523
	3.6A	9,682	3,482	993	23	296	1,198	351	817	2,523
	3.6B	9,723	3,538	993	23	296	1,166	350	835	2,523
CS-4	Existing	-	-	-	-	-	-	-	-	-
	3.2	9,793	3,502	1,003	22	296	1,426	362	659	2,523
	3.3	10,044	3,751	1,001	22	296	1,350	367	735	2,523
	3.4	9,789	3,586	1,002	22	296	1,292	362	707	2,523
	3.5	9,797	3,601	1,002	22	296	1,280	362	712	2,523
	3.6A	9,692	3,603	1,002	22	296	1,245	361	641	2,523
	3.6B	9,673	3,670	1,001	22	296	1,123	363	675	2,523

Figure 90 shows a visual example of the optimised output for the CS-4 disc for the two cost rates. The key difference between the two designs is that the RI geometry (coloured in black in Figure 90) is visibly larger meaning that the RI geometry mass will be higher than that shown in cost rate A.

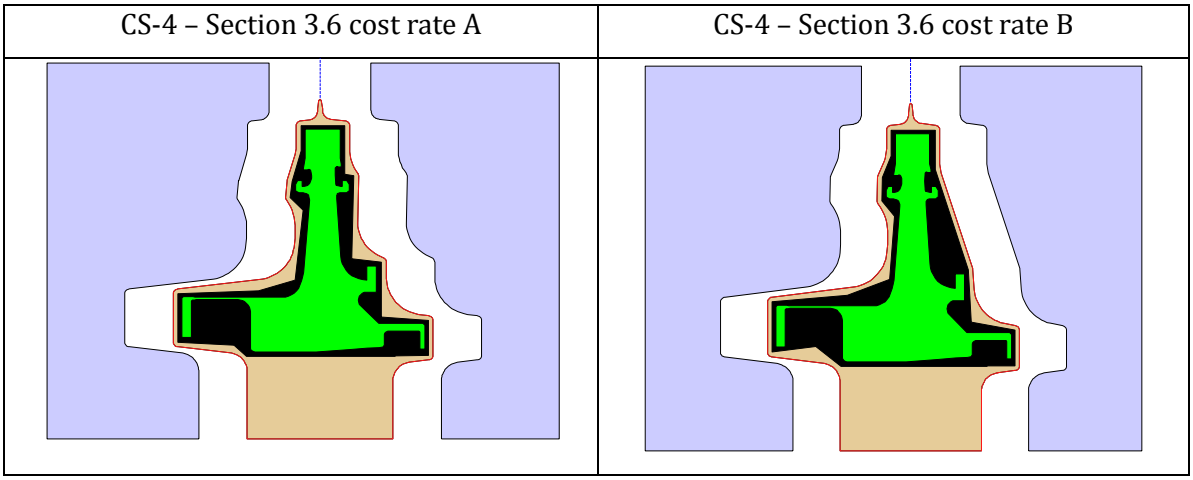

Figure 90 – Visualisation of results for the two different cost rate settings.

Such is the cost rate of turning operation 2, It is cheaper to add additional mass to the forging to enable the RI geometry to easily fill into the forging with cost rate B applied.

In Section 3.4 the black forging mass was minimised and the RI geometry was manipulated to maintain full inspection. Cost rate A requires the black forging mass to be minimised as well as the RI geometry as the more mass that is removed during turning operation 2, the cheaper the overall cost will be. On the other hand, cost rate B has the opposite cost rate and therefore would like to minimise the black forging mass but maximise the RI geometry such that a greater proportion of mass is removed during turning operation 3.

3.6.4. DISCUSSION

This section set out to determine the impact of detailed cost analysis had on the design of the RI geometry to determine the impact on the design. Detailed manufacturing cost models driven off the design during the design process is very rare and very approximate within Rolls-Royce. With a growing concern for cost and how to reduce unit cost in new projects it is important new tools become available as there is no cost estimator within the designer's tool kit for use in trade studies let alone for use within optimisation. With the creation of a cost estimator which produces the analytical cost using feasible RI geometry designs we now have a tool for which to provide the key cost drivers of a disc.

Summary of Findings

In this section we have identified the three largest drivers for cost for a given disc design:

- Input billet mass (based on black forging mass)
- Turning operation 2 (Heat treatment geometry to RI geometry)
- Turning operation 3 (RI geometry to finished disc profile)

The reduction in input billet mass has a twofold effect, firstly due to the high cost of purchasing the material and secondly due to additional turning it requires either during Turning operation 2 or turning operation 3.

The two remaining key cost drivers are the two same turning operations; turning operation 2 and 3. For a given input billet mass and a final disc profile, there is a fixed amount of material which needs to be removed. The way in which the material is distributed to the two turning operations is highly influential if the cost rates for material removal between the two operations are significant.

Cost rate A which assumed a higher cost rate during turning operation 3 relative to turning operation 2 showed that as well as trying to minimise the black forging mass, it also tried to minimise the RI geometry mass to avoid the higher cost rate.

Cost rate B which assumed a lower cost rate during turning operation 3 relative to turning operation 2 showed again that a reduction in black forging mass is key to reducing cost. It also tried to maximise the RI geometry within the forging to minimise the turning cost during turning operation 2.

If the cost rate between A and B was equivalent, then the optimisation would likely have produced a result similar to that shown in Section 3.4 where the billet input mass is the key cost driver.

Analysing the difference in designs between minimum unit cost and minimum RI geometry mass neglecting ultrasonic constraints as performed in Section 3.2 showed that unit cost inaccuracies were large ranging between -1.6% and 4.3%. The maximum heat treatment depth change was up to 4.2% greater when the RI geometry was primarily optimised for unit cost. This shows that the material properties and residual stresses could be significantly different to what was assumed in Section 3.2.

A direct cost comparison between Section 3.2 and Section 3.3 showed that part cost estimates increased in two-thirds of the cases but could also decrease in cost by up to 0.2% (ignoring the CS-4 anomaly result). This happened despite RI geometries consistently

increasing in mass from 3.7% to 6.6% for all cases (ignoring CS-4 anomaly result). This shows that considering including the ultrasonic limitation constraint does not necessarily add cost.

It was expected that optimising for unit cost should produce the cheapest RI geometry designs compared to previous sections. The optimisation was successful producing the cheapest RI geometry designs in 4 out of the 8 case studies shown when only considering designs which were 100 percent inspected. The results of the combined cost optimisation results produced 5 out of the 8 case studies showing that on average optimising for unit cost will typically produce a better result. Despite this success, the section had the same issues as shown in all of the previous three sections relating to optimisation performance. Each section should theoretically find the superior result for their own objective target, but this did not always occur. This issue is likely to be related to the efficiency of the optimiser where it has struggled to satisfy the ultrasonic NDE requirement of full inspection. Changing some of the GA's untouched parameters (described in Section 3.2.2.3) could potentially improve the optimisation's performance; also the number of function evaluations could be increased to permit the optimiser the freedom to explore the design space further. Unfortunately, due to time constraints both the additional GA parameter settings could not be investigated nor could the optimisation be run for additional generations.

Academic Implications

The workflow established in this section has shown that the RI geometry can be optimised beyond what has been achieved before. The RI geometry can now be optimised to minimise unit cost whilst maintaining manufacturability, full ultrasonic NDE inspection, and ensuring a minimum heat treatment depth solution is used when possible. Out of the relevant studies investigated, direct optimisation of unit cost was shown to be the most effective method of minimising disc unit cost when compared to mass minimisation objectives.

The study also showed that the direct impact of including the ultrasonic constraint does not necessarily increase the part cost regardless of any mass increase.

Business Implications

This workflow isn't limited for use by manufacturing engineers who rationally design the RI geometry; it can be applied by anyone, most notably the design engineers of the discs in question. This tool can be used throughout the design process as a cost prediction aid and will enable a disc designer to engage with forging suppliers earlier in the process. This would improve visibility of the design to broader design team and enable early provisioning of long lead time die materials.

Limitations and Directions for future research

The cost model developed within this section has been developed with limited access to cost breakdowns, therefore provides a representative model which could be improved in several areas, most importantly around forging, heat treatment and machining operations performed outside of Rolls-Royce.

It has been established that the key cost drivers are material input mass to produce the black forging and the turning costs during operations 2 and 3. Improving these three models further would improve the optimisations performance the most.

Further investigations on forging rules for isothermally forged HPT discs would ensure accurate forging mass productions

Turning was simplified to establish the different turning material removal rates. Two simple rules where used to define the difficult to turn zones which included non-visible regions when viewed along the disc's axis and narrow slots. More research is required on turning tools and feed and speed rates to determine zones of complexity in which to improve the total turning time which could be utilised for all turning operations.

All other aspects of the cost model have minimal impact on the relative costs during optimisation, but will influence the final costs accuracy which would be required for quoting. Therefore, all aspects of the cost model should be further researched, verified and implemented to establish a new reliable cost estimating tool.

Chapter 04 SUMMARY, CONCLUSIONS AND FUTURE WORK

4.1. CONCLUSIONS AND REFLECTIONS

This thesis set out to create an effective method to reduce the unit cost of manufacturing an isothermally forged disc. This was to be achieved through the use of a detailed cost model which could be used in both preliminary and detailed design. The thesis also set out to determine the impact an ultrasonic NDE constraint had on the design and cost of the RI geometry.

With tough competition in the gas turbine industry, companies are continuously being pushed to cut costs as well as improving their impact on the environment. This study has shown its importance to the industry by proving that both the unit cost of HPT discs can be reduced but it also decreases manufacturing waste reducing the environmental impact.

The first model created in Section 3.2 in this thesis was the automated rectilinear inspection geometry generator. This took a disc design then optimised RI geometry around it to produce a weight efficient design meeting the manufacturing requirements. This proved that there was plenty of potential to reduce weight/cost, but ignored ultrasonic NDE analysis which is the key constraint required to successfully life the part.

In order to incorporate the ultrasonic NDE analysis into the optimisation workflow a new approach was needed to cut analysis time down by several orders of magnitude compared to the current methods. This was achieved in Section 3.3 by creating a new piece of software which utilised a novel method to efficiently analyse the RI geometry. This method proved to be a success reducing analysis time from several hours down to only a few seconds. This method was passed to Rolls-Royce Plc. who productionised and introduced the software into the working environment replacing both the manual methods and the underperforming Rolls Royce developed software. The methods for this software were also submitted for patent approval by Rolls-Royce (Wiseall, Phipps et al. 2014) in 2014. Upon ultrasonic analysis of Section 3.2's optimised results, it was found that full inspection was missed by up to 5.22% proving that RI geometry needs to be designed with 100% ultrasonic inspection set a constraint.

The ultrasonic NDE analysis software when coupled with the RI geometry generator optimisation workflow successfully proved that RI geometry needs to be designed with the ultrasonic NDE set as a constraint to ensure the required inspection standards are met. Using this newly developed optimisation loop, RI geometry designs were created with

superior weight savings compared to the manual designs already in production. The weight savings ranged between 1.6 and 15.7 percent which prove that software can efficiently design the RI geometry despite the complex nature of the ultrasonic inspection rules. The software could also take the design of the RI geometry out of the manufacturing engineer's hands and straight into the design engineers to allow a relatively instant RI geometry design reducing the lengthy lead times related to the manual process.

The success of Section 3.3 replicating the RI geometry design process within a single automated process brought the software in line with current manual practices within Rolls-Royce, but the key optimisation objective used was the RI geometry mass used to represent cost. Although the cost of the disc is heavily linked to the input material mass to produce the forging, the RI geometry mass is not proportional to the weight of the forgings. Therefore, to get a better surrogate for cost the black forging mass was estimated using models created in Section 3.4. Models which were aimed at replicating the forging geometry design process proved to be an excellent approach. The technique used did not need finite element analysis to create an acceptable black forging design making it extremely fast to model. Like the ultrasonic NDE analysis speed was key in connecting this model to Section 3.3's automated optimisation loop. The transition from changing the objective from minimising the RI geometry mass to the forging mass was drastic. It was clear that the old requirement to minimise RI geometry mass should be relaxed and the focus should be on the reduction of the estimated black forging mass. By changing the objective to minimise the black forging mass resulted in a trend of RI geometry mass increasing relative to that found in Section 3.3, but with an overall reduction in black forging mass. By optimising the RI geometry for minimum black forging mass gave the RI geometry more freedom to potentially better the RI geometry design for ultrasonic inspection. This new design technique betters the current Rolls Royce manual methods as it uses known forging design rules to pre-empt what the forgers will do with the actual finished RI geometry.

The heat treatment surrogate created in Section 3.5 was not an initial objective to the optimisation but it was investigated so that the effects of the different studies could keep a track on how it was impacted. It was deemed important to track as the heat treatment process locks in the material properties into the part which are linked to both the strength and life of the component. The optimisation results showed that when the heat treatment maximum depth was minimised alongside the reduction of black forging mass, the forging mass tended to increase relative to Section 3.4's black forging mass optimisation, with only marginal to negative gains in the heat treatment performance.

The creation of the cost model in Section 3.6 formed the pinnacle of this thesis combining all the developed manufacturing model data and converting it directly into a monetary value. This allowed the optimiser to directly optimise the RI geometry for minimum unit cost without having to use mass as a surrogate to represent cost. The results of the optimisation yielded promising results, further improving the unit cost to manufacture discs in 5 of the 8 case study examples. The method successfully demonstrated further potential cost savings of up to 1.91% on one of the CS-3 case studies compared to the next cheapest design from a previous section.

Although a big cost driver was the input billet material, another large portion of the cost came from turning. Turning cost rates turned out to be one of the most important factors in determining the design of the RI geometry. Depending on which stage the turning operation is cheaper will dictate whether the RI geometry should be maximised within the black forging geometry or minimised around the disc profile.

Previous research which neglected the ultrasonic examination constraint was proven to produce wrong estimates. The cost analysis showed that including an ultrasonic NDE constraint did not necessarily add cost to the part despite a direct 6.6% mass increase to the RI geometry. Comparisons between Section 3.2 and Section 3.6 showed that unit cost inaccuracies were up to +1.6% different. Even more worrying is the potential impact on the assumed material properties which had a maximum heat treatment depth of up to 4.2% greater with the ultrasonic examination constraint applied and optimised primarily for unit cost. This proves that the ultrasonic constraint must be included in all manufacturing investigation studies in order to create feasible designs similar to what would be produced in reality.

The literature review unveiled three key gaps:

- No ultrasonic NDE analysis has been performed as part of a disc manufacturing design optimisation loop – Section 3.2 when compared with Section 3.3 revealed how much different the rectilinear inspection geometry could become. Sections 3.4 and 3.5 revealed the potential upstream impacts on detailed forging and heat treatment designs, thus proving that the ultrasonic NDE constraint must be included when designing the rectilinear geometry.
- No fast ultrasonic NDE analysis tool for use with 2D axi-symmetric disc designs – The bulk of Section 3.3 was dedicated to developing a new patentable method of

creating ultrasonic NDE analysis. The software created takes only a few seconds to perform an analysis which would otherwise take hours using pre-existing methods.

- There was no detailed cost model for use in optimisation as typically mass was regarded as a good enough surrogate for cost in many examples – Section 3.6 was dedicated to creating a detailed cost model, which effectively took advantage of data produced by models created in Sections 3.2 to 3.5 to generate an actual calculated cost. Using this cost model, the optimisation study in Section 3.6 produced the most successful optimum design out of all the studies in the chapter.

Throughout all of the sections the optimisation underperformed compared to expectations. It was expected that the optimisation objective in each section would outperform previous sections, as the optimisation had clear links with the objectives. Only around half of the results in each section improved on earlier designs using their respective objectives. The reasoning for this is most likely due to the optimisation setup. The optimiser was calibrated in Section 3.2 when the RI geometry generator was first created. Later sections saw the optimisation problem expand and included multiple objectives. This may be partially why optimisation performance was poor. In early optimisations where just the RI geometry was to be minimised, there was a strong link between the objective and the optimisation parameters. These parameters were all linked directly to the geometry which directly affected the objective function. When the upstream manufacturing models were created and the objectives changed, the parameter links to the objective function became less direct. For example, the BF geometry was strongly linked to parts of the RI geometry in places, but was weakly linked in areas where the RI geometry did not impact the forging geometry at all. By the time the cost model was introduced, the unit cost objective function was linked to a wide range of information all linked in some way to the input parameters. This would cause the objective function to react irrationally making it very difficult for an optimiser to find the optimal point. This would make it exceptionally difficult for the local gradient based optimiser to function effectively.

What was not investigated further was the parameter setup for the optimiser as this was only performed in the initial stages for the initial RI geometry optimisation. Due to the potential increase in complexity the population and maximum generation parameters should have been re-assessed to ensure the design space was effectively explored; the optimisation time might have needed to increase, but the concept of optimising for minimum unit cost could have been realised.

HT optimisation – The heat treatment objective of minimising the maximum depth within the heat treatment geometry had only a small impact on the designs. The results had little relevance to the unit cost optimisation, but could have been used to great effect if residual stress, material properties and life surrogate models were effectively incorporated.

Academic Implications

This study has shown the importance of modelling the complete manufacturing design process in order to obtain viable non-misleading data for which further analysis can be performed. The results have shown the impact of neglecting an ultrasonic analysis constraint during investigative manufacturing design optimisation. Mass, material properties and unit cost were all estimated incorrectly. This could lead to late design changes that would be required to resolve the repercussions of using inaccurate data. Including the ultrasonic constraint mitigates these concerns as the produced manufacturing designs are now more realistic than without.

The study has also shown the power of using a detailed analytical cost model as an objective for the optimisation of manufacturing design parts. Simplistic part mass minimisation as a surrogate for cost was proved to be a flawed objective. In many areas RI geometry had no impact on forge mass so waste material could therefore be used more efficiently to solve the ultrasonic constraint problem. The desired size of the RI geometry was actually dependent upon the cost rates of removing material. This showed that unit cost optimisation was the most effective method for manufacturing design optimisation for unit cost.

One of the most important methods produced in this study was the ultrasonic analysis methodology which was designed to speed up the analysis process. The novel method developed to implement the ultrasonic rule base reduced ultrasonic analysis time down from several hours to a few seconds. This enabled tens of thousands of RI geometry designs to be analysed within the same time frame as previous methods allowing large explorative optimisation to occur. The results of this now possible optimisation enabled substantial design improvements within days which were not previously possible.

Business Implications

Due to the lead time required in creating RI geometry, the design envelope of the disc is normally fixed early. This can be frustrating to the disc design team and to all the parties of components which interface with the disc which have much later manufacturing lead times. Using this optimisation loop, a valuable couple of weeks could be recovered reducing delays or enabling the disc design to improve.

Improving both the optimisation method and or parameters has the potential to realise the ideal of optimising the RI geometry for unit cost. This would give the designer a higher confidence of the cost for the assumed cost rates provided. In a business which is dependent upon external companies to provide quotes for manufacturing costs; this would provide a highly insightful asset to the designer.

Conventional methods of generating costs could take weeks to produce requiring both external company quotes using a single fixed RI geometry. Now for a given design, a clear cost breakdown can be generated within seconds. There are further benefits from having this additional information

- External contract negotiators could use the detailed cost breakdown to check contract details.
- Material property expectations could be shared with forgers after in-house material properties have been investigated.
- A clear mature unit cost could be provided weeks or even years prior to the actual manufacturing design process commencing.

Even discounting the speed at which the cost estimate can be produced, the cost reduction has already proved to be a significant amount, potentially saving millions of dollars of material waste over the production life of an engine. With an assumption that an average disc costs approximately $30k and there are as many as 27 forged discs in an engine, just a 1% reduction in unit cost to each could save up to $1,000,000 over 124 engines or 62 twin engine aircraft. Not only does this optimisation benefit the profit margins of the engine manufacturers, but it will also benefit the environment in terms of waste and energy reduction.

This workflow not only provides a method in which to design efficient RI geometries, it also provides the disc design engineer with a mature cost estimate in which to compare different concepts throughout the development of the disc. This tool has never been at the designer's disposal due to the fact that a RI geometry would normally take weeks if not months to be developed into an acceptable shape as well as taking up valuable time. Utilising this optimisation method within a disc optimiser could further improve the disc design process so that discs are always designed with a good cost estimate available which can be considered against both performance and weight.

Limitations and Directions for Future Research

The black forging model was created using generic rules for creating good initial black forging geometry to help design dies prior to physical experiments or computational finite element testing. No forging experiments were performed to check the forgability of the geometry created, but the questioning of experienced forgers confirmed that the designs created by this rule based model is likely to be very similar to the actual design as only small alterations are normally required to fix any flow problems (Bryant 2010; Brooks 2010a). Part of the reason why the rules create such a good first attempt at the forging geometry is because most of the discs are created isothermally, if these were hot die forged or cold die forged, then the forging process becomes a lot more complicated.

The main limitation of this developed workflow is the cost model which has been generated. The data used are estimates and have been set from experts' experience and judgement, but may not accurately represent how the cost is calculated in the actual external factories. In an ideal world, the supplier would incorporate a detailed cost model around what has been created to make use of the fully integrated connections to the geometry designs. However, in reality the factory costs are normally a secretive and proprietary to the company with quotes being more of an art than a calculated value. Improvements to this model should only improve the performance, but may never accurately predict the actual contractual unit cost of a part, as these will come down to agreement from both parties. It gives a good way to methodically compare similar RI geometry designs and give reasons why costs are different, and how a disc can be altered to minimise the cost further.

As discussed, the time to converge on a good design can take several thousand iterations. For this study the number of evaluations is not an issue as we are trying to achieve a solution within a 24 hour time frame. To improve the practicality of this RI geometry optimisation, we would ideally have a result within minutes or hours as opposed to a day. Therefore, to improve the time taken a more intelligent method of parameterisation or optimisation could be used to quickly converge on the optimum result.

4.2. FUTURE WORK AND INVESTIGATIONS

There has been considerable progress since this project began in 2009; all the software was successfully created with considerable robustness to allow optimisation investigations to be performed. One of the key areas of improvement should be on the analytical cost model, which requires greater research into accurate cost rate values as well as additional cost factors which need to be included. Although the cost model provides detailed trends in which to compare different designs, the models have not been verified. Further

development is required to tune the model to the different factories to get more accurate cost rates and accurate absolute unit cost results.

Optimisation strategies should be investigated further where genetic algorithms, hill climbing algorithms and response surface models can be looked at in greater detail.

If the optimisation efficiency can be improved to bring the solve time down to only an hour or two, this optimisation loop could be integrated into a disc optimisation loop so that the cost can be used as one of the objectives.

There are several other investigations that could be pursued to further understand the input and output variables associated with the complete process. Listed below are potential investigations to look into in further detail related to the ultrasonic inspection requirement:

- Finding relationships between cost and inspectability. We hope to reduce the overall cost of the disc but also improve the confidence in the disc design in a lifing sense by increasing the total volume of ultrasound within the part.
- Improving the inspection of RI geometry designs through refinement of the inspection rules. The ultrasonic NDE rules are currently fixed from a specification derived over 30 years ago. There is some flexibility in some of the rules which may allow RIG designs to be improved by changing a rule for that instance.
- Investigate integrating probabilistic lifing into the ultrasonic NDE analysis. This is currently performed manually after a RI geometry design fails to achieve 100% inspectability to check if partially inspected area is in a no risk region.
- Apply weightings to ultrasound scans to improve inspection robustness. Currently it does not matter what direction the four ultrasound scans come from. Introducing weighted scans would encourage RI geometry designs to scan using a good distribution of angles which increase the likelihood of finding inclusions.
- Zoned inspection requirements to relax the requirement of 100% inspection from 4 directions in non-critical areas.

The ultimate aim once all the models have been completed and verified is to run an optimisation algorithm on a pre-existing disc design to compare the actual and optimised solutions. This will show the extent in which this workflow can reduce cost, improve inspection procedures as well as the time that can be saved compared to current standards.

The surrogate for life has not been utilised to its full potential in this thesis. Understanding the relationship further between the heat treatment depth and material properties could lead to a more beneficial use of the output.

APPENDIX I - AXI-SYMMETRIC VOLUME CALCULATION

To calculate the volume of a 2D axi-symmetric polygon a function was created which converts 3D volume into a 2D polygon where the area of the 2D polygon is the equivalent of the 3D volume. This method was used to utilise the Matlab function called 'polyarea'. The function 'polyarea' calculates the area of a polygon.

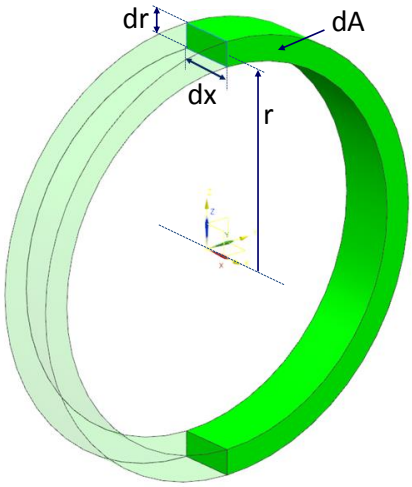

It can be proved that the volume of a polygon has a linear relationship to the area times the radius.

$$dA = \pi\left[\left(r + \frac{1}{2}dr\right)^2 - \left(r - \frac{1}{2}dr\right)^2\right]$$

Expand the equation:

$$dA = \pi\left[\mathrm{r}^2 + rdr + \frac{dr^2}{4} - \mathrm{r}^2 + rdr - \frac{dr^2}{4}\right]$$
$$= 2\pi[rdr]$$

The volume of the object can be calculated using the following equation:

$$V = \iint dx\, dA$$

Substitute the equation for *dA* into the volume calculation to find the volume:

$$V = \iint 2\pi r\, dx\, dr$$

We are interested in using this result to produce a 2D polygon whose area will represent the volume of the axi-symmetric revolved shape. The area of a polygon can be calculated using the following formula where *y* and *r* are interchangeable:

$$A = \iint f(x,r)\, dx\, dr$$

As we require the volume and area to be equal we can make the equations equal each other:

$$A = \iint f(x,r)\, dx\, dr = \iint 2\pi r\, dx\, dr = V$$

This gives the result:

$$f(x,r) = 2\pi r$$

The function is only a function of *r*; therefore, for every finite element of a given radius, the *x* value can be multiplied by the function $2\pi r$ to calculate a finite elements volume instead of area. This could be visualised by multiplying a finite slice at a given radius by the circumference to produce an area of 'd*r*' thickness.

Knowing these two results, a 3D axi-symmetric geometry can be represented in 2D to visualise the 3D volume in a 2D area as illustrated in Figure 91 below.

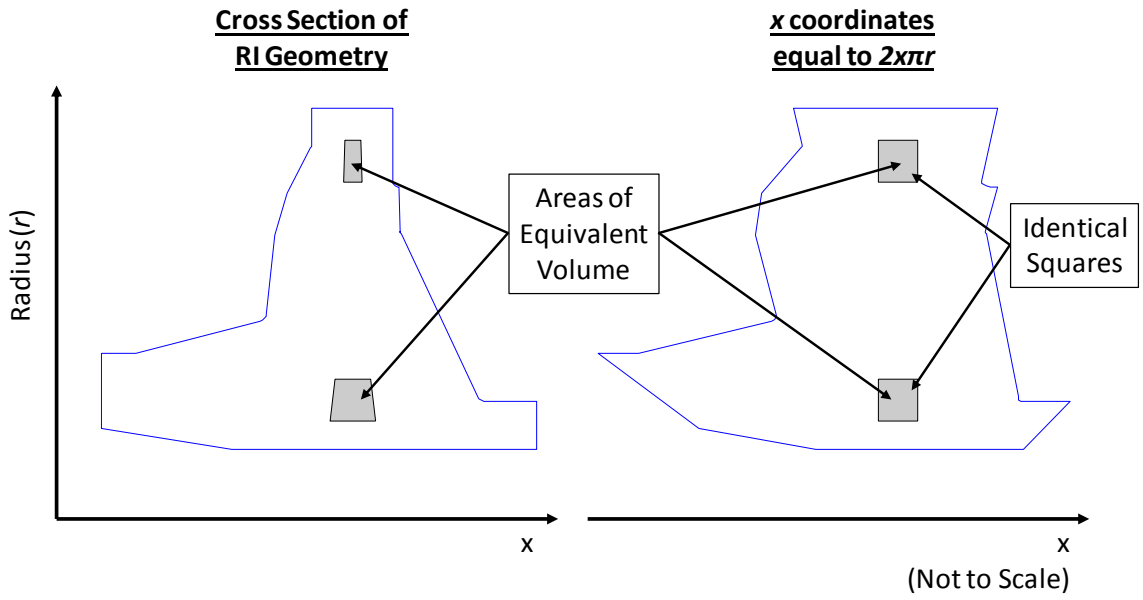

Figure 91 - Visualisation of the conversion between 2D axi-symmetric 3D volume to equivalent volume presented in a 2D area.

The Matlab ‘polyarea’ function can therefore be used to quickly calculate the area which is the equivalent of the volume.

If the 2D axi-symmetric polygon has its axial coordinates (x-axis) replaced with $2\pi rx$ the volume can be easily calculated using the same ‘polyarea’ function with the following parameters:

$$V = f([2\pi xy], r)$$

This page intentionally left blank.

APPENDIX II – EXPANDED ULTRASONIC NON-DESTRUCTIVE EXAMINATION RULES (RRP 58008) (CONTAINS PROPRIETARY DATA)

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

This page intentionally left blank.

APPENDIX III – GENERIC OPTIMISATION WORK FLOW

Flow diagram of the generic optimisation work flow. Note that the parting line position is calculated using the disc geometry and not the RI geometry to save significant computational time with minimal impact on accuracy.

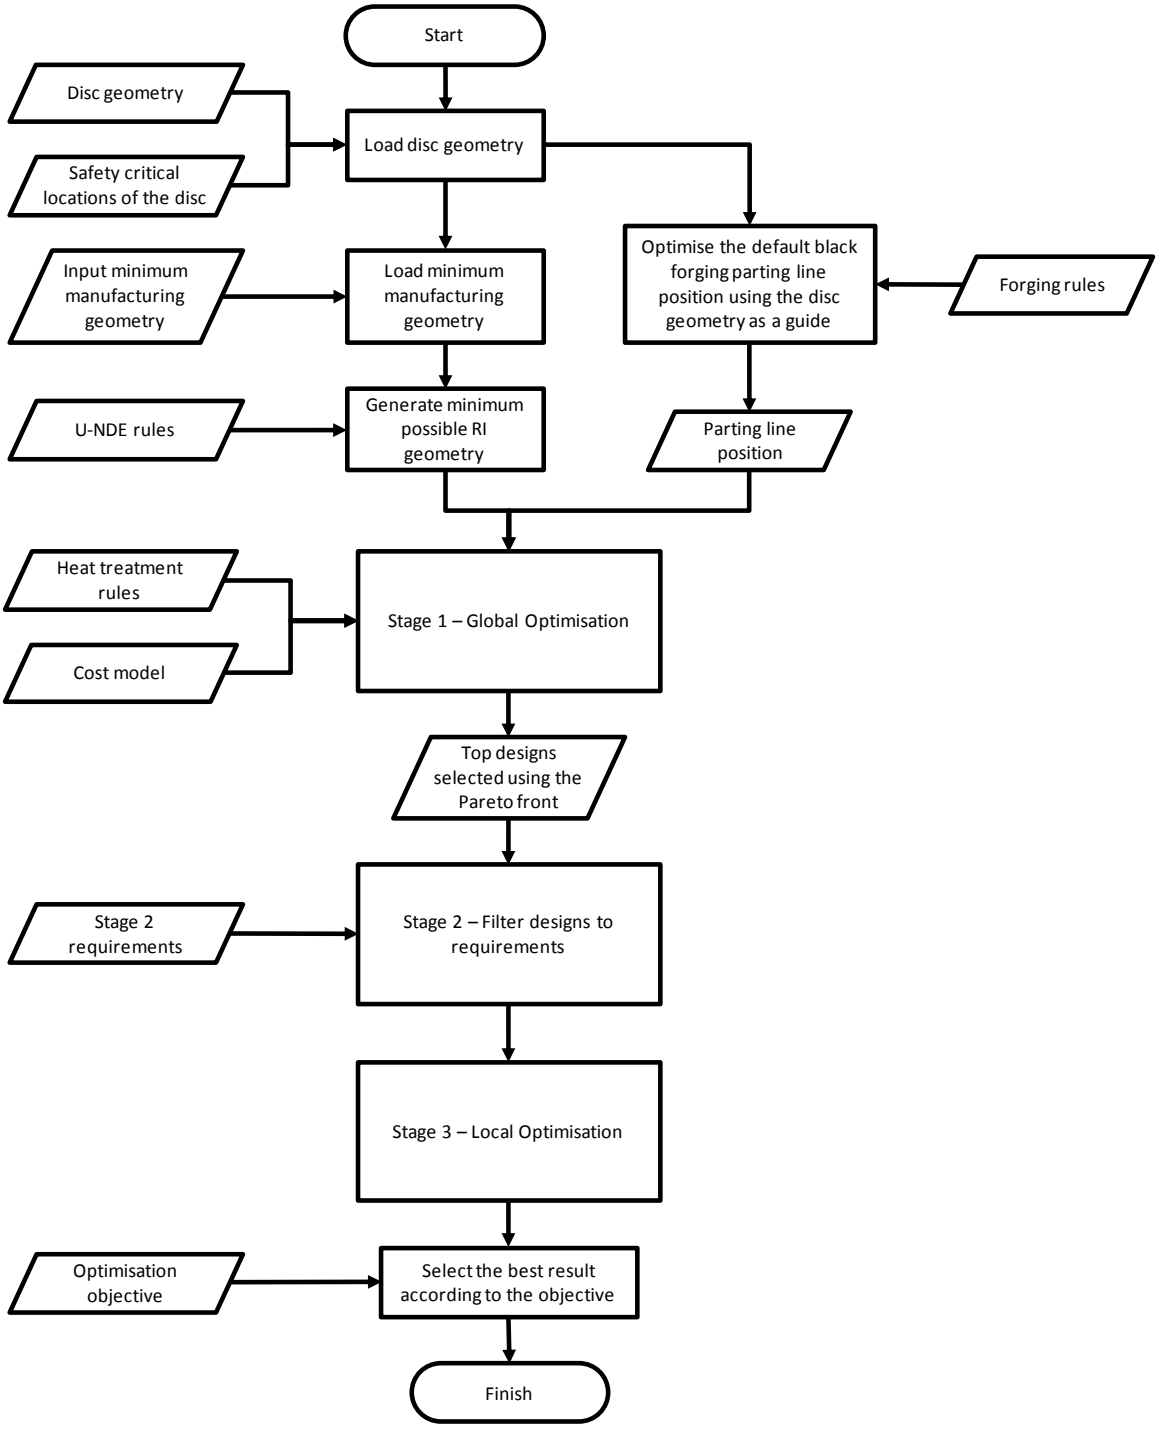

This page intentionally left blank.

APPENDIX IV – COST MODEL DETAILED STRUCTURE

The programmatic code used in Chapter 3.6 to calculate the cost is provided here for reference. Please note that some equations span across two or more lines due to 'word wrap'. Lines which end in '...' is an intentional cue the code continues on to the following line.

STD – Standard fixed values used in Chapter 3.6 are provided in Appendix V
INPUT – Values found during manufacturing design analysis which are part specific
OUTPUT – Values which are outputted

```
function OUTPUT = HPTD_UNIT_COST(INPUT,STD,OUTPUT)

OUTPUT = Total_Material_Cost(INPUT,STD,OUTPUT);
OUTPUT = Total_Forging_Cost(INPUT,STD,OUTPUT);
OUTPUT = Total_Turning_Op1_Cost(INPUT,STD,OUTPUT);
OUTPUT = Total_Heat_Treatment_Cost(INPUT,STD,OUTPUT);
OUTPUT = Total_Turning_Op2_Cost(INPUT,STD,OUTPUT);
OUTPUT = Total_Ultrasonic_NDE_Cost(INPUT,STD,OUTPUT);
OUTPUT = Total_Turning_Op3_Cost(INPUT,STD,OUTPUT);
OUTPUT = Total_Machining_Cost(INPUT,STD,OUTPUT);

HPTD_Unit_Cost = OUTPUT.Total_Material_Cost...
               + OUTPUT.Total_Forging_Cost...
               + OUTPUT.Total_Turning_Op1_Cost...
               + OUTPUT.Total_Heat_Treatment_Cost...
               + OUTPUT.Total_Turning_Op2_Cost...
               + OUTPUT.Total_Ultrasonic_NDE_Cost...
               + OUTPUT.Total_Turning_Op3_Cost...
               + OUTPUT.Total_Machining_Cost;

OUTPUT.HPTD_Unit_Cost = HPTD_Unit_Cost;
```

```
function OUTPUT = Total_Material_Cost(INPUT,STD,OUTPUT)

Billet_Total_Volume = (pi() * STD.Billet_Radius^2)*STD.Billet_Min_Length;
Billet_Total_Stock_Mass = Billet_Total_Volume / STD.Nickel_Density;
Billet_Stock_Cost = STD.Billet_Cost_Per_Kg/Billet_Total_Stock_Mass;

if STD.Billet_Chamfered_Req == 0
    Billet_Req_Length = INPUT.BF_Vol.*(pi().*INPUT.(Billet_Diameter/2).^2);
elseif STD.Billet_Chamfered_Req == 1
    temp_Vol1 = pi() * STD.Billet_Radius^2 * STD.Billet_Chamfer_Height;
    temp_Vol2 = pi*(...
                    (STD.Billet_Chamfer_Height^3)/3 ...
                    +       STD.Billet_Chamfer_Height^2       *       (STD.Billet_Radius-
               STD.Billet_Chamfer_Height)...
                    +        STD.Billet_Chamfer_Height        *       (STD.Billet_Radius-
               STD.Billet_Chamfer_Height)^2 ...
                    );
    Billet_Chamfer_Volume_Lost = temp_Vol1 - temp_Vol2;
    Billet_Req_Volume_new = INPUT.BF_Vol...
                                  + Billet_Chamfer_Volume_Lost;
    Billet_Req_Length =  Billet_Req_Volume_new/(pi()*STD.Billet_Radius^2);
else
    error('Chamfered_Billet_Req must equal 0/false or 1/true')
end

if (STD.Billet_Utilisation_Req == 1) || (STD.Billet_Utilisation_Req == 2)
    Billet_Usable_Lengths = floor(STD.Billet_Stock_Min_Length/Billet_Req_Length);
    Total_Material_Cost = Billet_Stock_Cost/Billet_Usable_Lengths;
end


if (STD.Billet_Utilisation_Req == 1)
```

```
    Billet_Volume = Billet_Req_Length*pi()*STD.Billet_Radius^2;
elseif (STD.Billet_Utilisation_Req == 2)
    Billet_Leftover_Stock_Percent         =         (STD.Billet_Stock_Min_Length         -
            (Billet_Req_Length*Billet_Usable_Lengths)) / STD.Billet_Stock_Min_Length;
    Billet_Revert               =               STD.Billet_Revert_Percent               *
            Billet_Stock_Cost*Billet_Leftover_Stock_Percent;
    Billet_Revert_Per_Part = Billet_Revert/Billet_Usable_Lengths;
    Total_Material_Cost = Total_Material_Cost - Billet_Revert_Per_Part;
    Billet_Volume = Billet_Req_Length*pi()*STD.Billet_Radius;
elseif STD.Billet_Utilisation_Req == 3 % Straight Forward Cost per kg
    Billet_Volume = Billet_Req_Length*pi()*STD.Billet_Radius^2;
end
Billet_Mass = Billet_Volume * STD.Nickel_Density;
OUTPUT.Total_Material_Cost = STD.Billet_Cost_Per_Kg * Billet_Mass;
```

```
function OUTPUT = Total_Forging_Cost(INPUT,STD,OUTPUT)

% Level 10
Forging_Billet_Mass = INPUT.BF_Vol * STD.Nickel_Density;
Top_Die_Mass = INPUT.Top_Die_Volume * STD.Die_Material_Density;
Bot_Die_Mass = INPUT.Bot_Die_Volume * STD.Die_Material_Density;

Required_Die_Temp = STD.Forging_Temperature;

Top_Die_Rough_Machining_Volume      =      INPUT.Top_Die_Material_To_Remove_Volume      *
            INPUT.Top_Die_Rough_Machining_per;
Bot_Die_Rough_Machining_Volume      =      INPUT.Bot_Die_Material_To_Remove_Volume      *
            INPUT.Bot_Die_Rough_Machining_per;
Top_Die_Finish_Machining_Volume     =      INPUT.Top_Die_Material_To_Remove_Volume      *
            INPUT.Top_Die_Finish_Machining_per;
Bot_Die_Finish_Machining_Volume     =      INPUT.Bot_Die_Material_To_Remove_Volume      *
            INPUT.Bot_Die_Finish_Machining_per;

% Level 9

Die_Material_Specific_Heat_Capacity = STD.Molybdenum_Specific_Heat_Capacity;
Die_Delta_Temperature = Required_Die_Temp...
                      - STD.Room_Temperature;

Top_Die_Purchase_Volume = pi() * (INPUT.Top_Die_Diameter/2)^2 * INPUT.Top_Die_Depth;
Bot_Die_Purchase_Volume = pi() * (INPUT.Bot_Die_Diameter/2)^2 * INPUT.Bot_Die_Depth;

Top_Die_Rough_Machining_Time          =          Top_Die_Rough_Machining_Volume          /
            STD.Die_Rough_Machining_Material_Removal_Rate;
Bot_Die_Rough_Machining_Time          =          Bot_Die_Rough_Machining_Volume          /
            STD.Die_Rough_Machining_Material_Removal_Rate;

Top_Die_Finish_Machining_Time         =         Top_Die_Finish_Machining_Volume         /
            STD.Die_Finish_Machining_Material_Removal_Rate;
Bot_Die_Finish_Machining_Time         =         Bot_Die_Finish_Machining_Volume         /
            STD.Die_Finish_Machining_Material_Removal_Rate;

% Level 8

Unit_Forging_Time = STD.Billet_Setup_Time...
                  + STD.Forging_Time;

Energy_Required_to_heat_Top_Die  =  Top_Die_Mass  *  Die_Material_Specific_Heat_Capacity  *
            Die_Delta_Temperature;
Energy_Required_to_heat_Bottom_Die = Bot_Die_Mass * Die_Material_Specific_Heat_Capacity *
            Die_Delta_Temperature;

Top_Die_Purchase_Mass = Top_Die_Purchase_Volume * STD.Die_Material_Density;
Bot_Die_Purchase_Mass = Bot_Die_Purchase_Volume * STD.Die_Material_Density;

Top_Die_Rough_Machining_Cost          =          Top_Die_Rough_Machining_Time          *
            STD.Die_Rough_Machining_Cost_Rate;
Bot_Die_Rough_Machining_Cost          =          Bot_Die_Rough_Machining_Time          *
            STD.Die_Rough_Machining_Cost_Rate;

Top_Die_Finish_Machining_Cost         =         Top_Die_Finish_Machining_Time         *
            STD.Die_Finish_Machining_Cost_Rate;
Bot_Die_Finish_Machining_Cost         =         Bot_Die_Finish_Machining_Time         *
            STD.Die_Finish_Machining_Cost_Rate;
```

```
% Level 7

Total_Forging_Time = STD.Forging_Time * STD.Average_Production_Batch_Number;
Percentage_Capacity = INPUT.Required_Force / STD.Forge_Capacity;

Specific_Heat_Capacity = STD.Nickel_Specific_Heat_Capacity;
Forging_Delta_Temp = STD.Forging_Temperature-STD.Room_Temperature;

Total_Time_in_Forge = Unit_Forging_Time * STD.Average_Production_Batch_Number;

Energy_Required_to_heat_Dies = Energy_Required_to_heat_Top_Die...
                             + Energy_Required_to_heat_Bottom_Die;

Top_Die_Material_Cost = Top_Die_Purchase_Mass * STD.Die_Material_Cost_per_Kg;
Bot_Die_Material_Cost = Bot_Die_Purchase_Mass * STD.Die_Material_Cost_per_Kg;
Top_Die_Machinig_Cost = Top_Die_Rough_Machining_Cost...
                      + Top_Die_Finish_Machining_Cost;
Bot_Die_Machinig_Cost = Bot_Die_Rough_Machining_Cost...
                      + Bot_Die_Finish_Machining_Cost;

% Level 6
Energy_required_to_heat_component            =            Forging_Billet_Mass            *
              Specific_Heat_Capacity*Forging_Delta_Temp;

Forge_Power_Costs     =     STD.Forge_Cost_at_full_Power     *     Percentage_Capacity     *
              Total_Forging_Time;

Total_Batch_Time = STD.Setup_Time...
                 + STD.Heatup_Time...
                 + Total_Time_in_Forge...
                 + STD.Cool_Down_Time;
Die_Heating_Cost     =     Energy_Required_to_heat_Dies     *     STD.Electricity_Price     *
              STD.Forge_Heating_Efficiency;

Die_Top_Cost = Top_Die_Material_Cost...
             + Top_Die_Machinig_Cost;
Die_Bottom_Cost = Bot_Die_Material_Cost...
             + Bot_Die_Machinig_Cost;

% Level 5
Furnace_Energy_Costs   =   Energy_required_to_heat_component   *   STD.Electricity_Price   /
              STD.Furnace_Efficiency;
Furnace_overheads = STD.Component_Heating_Time * STD.Furnace_Cost_Rate;

Forge_Costs = Forge_Power_Costs...
            + (INPUT.BF_SF_Area * STD.BF_Fudge_Factor) / INPUT.BF_Vol;
Forge_Overheads = (Total_Batch_Time * STD.Forge_Cost_Rate)...
                + Die_Heating_Cost;

Total_Die_Cost = Die_Top_Cost...
               + Die_Bottom_Cost;


% Level 4
Component_Heating_Costs = Furnace_Energy_Costs...
                        + Furnace_overheads;
Direct_Forging_Cost = Forge_Costs...
                    + Forge_Overheads;
Die_Overheads = Total_Die_Cost / STD.DIE_Average_Production_Life_Of_Die ;

% Level 3
Pancake_Operation_Cost = 0; % could expand but is not needed for current models
Forging_Operation_Cost = Component_Heating_Costs...
                       + Direct_Forging_Cost...
                       + Die_Overheads;

% Level 2
Should_be_Forging_cost = Pancake_Operation_Cost + Forging_Operation_Cost;
Forge_Profit = Should_be_Forging_cost * STD.Forge_profit_margin;

%Level 1
OUTPUT.Total_Forging_Cost = Should_be_Forging_cost + Forge_Profit;
```

```
function OUTPUT = Total_Turning_Op1_Cost(INPUT,STD,OUTPUT)
```

```
% Level 6
T1_Rough_Volume_to_Remove = INPUT.BF_Rim_Flash_Volume;
T1_Trepan_Cutting_Length = pi() * INPUT.HT_Bore_Diameter;
T1_Trepan_Feed_Rate = 1/(INPUT.HT_Bore_Thickness/STD.T1_Trepan_Depth_Cut_Rate);

% Level 5
T1_Total_Disc_Flipping_Time = STD.T1_Flipping_Time * STD.T1_Number_of_Flips;
T1_Rough_Machinging_Time             =             T1_Rough_Volume_to_Remove             /
               STD.T1_Rough_Rate_of_Material_removal;
T1_Trepan_Total_Cutting_Op_Time = T1_Trepan_Cutting_Length / T1_Trepan_Feed_Rate;

% Level 4
T1_Setup_Time = STD.T1_Initial_Setup_Time...
                + T1_Total_Disc_Flipping_Time...
                + STD.T1_Setting_Down_Time;

T1_Total_Turning_Op_Time = T1_Rough_Machinging_Time;
T1_Trepan_Setup_Costs = STD.T1_Trepan_Setup_Time * STD.T1_Trepan_Setup_Cost_Rate;
T1_Trepan_Cutting_Costs = T1_Trepan_Total_Cutting_Op_Time * STD.T1_Trepan_Cost_Rate;

T1_Total_Process_Time = T1_Total_Turning_Op_Time / STD.T1_Utilisation_Efficiency;

% Level 3
T1_Setup_Costs = T1_Setup_Time * STD.T1_Setup_cost_rate;
T1_Turning_Costs = (T1_Total_Process_Time)* STD.T1_cost_rate;

% Level 2
T1_Total_should_be_cost = T1_Setup_Costs...
                          + T1_Turning_Costs...
                          + T1_Trepan_Setup_Costs...
                          + T1_Trepan_Cutting_Costs;
T1_Profits = T1_Total_should_be_cost * STD.T1_Profit_Margin;

% Level 1
OUTPUT.Total_Turning_Op1_Cost = T1_Total_should_be_cost...
                              + T1_Profits;
```

```
function OUTPUT = Total_Heat_Treatment_Cost(INPUT,STD,OUTPUT)

% Layer 10
Solution_treatment_Delta_Temperature = STD.Quenching_Temperature - STD.Room_Temperature;

% Layer 9
Aging_Delta_Temperature = STD.Aging_Temperature - STD.Room_Temperature;
Aging_Ramp_Temp_Up_Time              =              Aging_Delta_Temperature              /
               STD.Aging_Rate_of_Change_of_Temperature_Up;
Solution_Heat_Treatment_Ramp_Temp_Up_Time    =    Solution_treatment_Delta_Temperature    /
               STD.Solution_Treatment_Rate_of_Change_of_Temperature_Up;

% Layer 8
Specific_Heat_Capacity = STD.Nickel_Specific_Heat_Capacity;
Heat_Treatment_Mass = INPUT.HT_Vol * STD.Nickel_Density;

Time_in_Aging_Furnace = STD.Aging_Time_At_Tempearture...
                        + (0.5 * Aging_Ramp_Temp_Up_Time);
Time_in_Solution_Treatment_Furnace = STD.Time_at_Solution_Treatment_temperature...
                                   + Solution_Heat_Treatment_Ramp_Temp_Up_Time;

% Layer 7
Energy_Required_to_heat_component_for_Aging = Heat_Treatment_Mass * Specific_Heat_Capacity
               * Aging_Delta_Temperature;
Aging_Energy_Loss = STD.Energy_Loss_per_hour_at_Aging_Temperature * Time_in_Aging_Furnace;
Ramp_Temp_Down = Aging_Delta_Temperature / STD.Aging_Rate_of_Change_of_Temperature_Down;

Energy_Required_to_Heat_component_for_Solution_Treatment     =     Heat_Treatment_Mass     *
               Specific_Heat_Capacity * Solution_treatment_Delta_Temperature;
Solution_Treatment_Energy_Loss    =    STD.Energy_Loss_per_hour_at_Quenching_Temperature    *
               Time_in_Solution_Treatment_Furnace;

% Layer 6
Total_Energy_Required_for_Aging     =     Energy_Required_to_heat_component_for_Aging     +
               Aging_Energy_Loss;

Aging_Heating_Time = Aging_Ramp_Temp_Up_Time...
                     + STD.Aging_Time_At_Tempearture...
```

```
                    + Ramp_Temp_Down;
Total_Energy_Required_for_Solution_Treatment                                    =
              Energy_Required_to_Heat_component_for_Solution_Treatment...
                                                 + Solution_Treatment_Energy_Loss;

% Layer 5
Aging_Furnace_Energy_Costs  =  Total_Energy_Required_for_Aging  *  STD.Electricity_Price  *
              STD.Furnace_Efficiency;
Aging_Furnace_overheads = Aging_Heating_Time * STD.Aging_Cost_Rate;
Solution_Heat_Treatment_Energy_Costs   =   Total_Energy_Required_for_Solution_Treatment   *
              STD.Electricity_Price * STD.Quenching_Furnace_Effiency;
Quenching_Furnace_overheads         =         Time_in_Solution_Treatment_Furnace         *
              STD.Sln_HT_Furnace_Cost_Rate;

% Layer 4
Aging_Setup_Cost = STD.Aging_Setup_Time * STD.Aging_Setup_Cost_rate;
Aging_Operation_Cost = Aging_Furnace_Energy_Costs...
                     + Aging_Furnace_overheads;
Sln_HT_and_Quenching_Setup_Cost       =       Solution_Heat_Treatment_Energy_Costs       +
              Quenching_Furnace_overheads;
% Layer 3
Aging_Operation_Total_Cost = Aging_Setup_Cost...
                           + Aging_Operation_Cost;

Sln_HT_and_Quenching_Operation_Total_Cost = Sln_HT_and_Quenching_Setup_Cost...
                                         + STD.Quenching_Operation_Cost;

% Layer 2
should_be_heat_threatment_cost           =           Aging_Operation_Total_Cost           +
              Sln_HT_and_Quenching_Operation_Total_Cost;
Heat_Treatment_profit = should_be_heat_threatment_cost * STD.Heat_Treatment_Profit_Margin;

% Layer 1

OUTPUT.Total_Heat_Treatment_Cost = should_be_heat_threatment_cost...
                           + Heat_Treatment_profit;
```

```
function OUTPUT = Total_Turning_Op2_Cost(INPUT,STD,OUTPUT)

% Level 6
T2_Rough_Volume_to_Remove = (INPUT.HT_Vol - INPUT.COS_Vol) * INPUT.T2_Per_Rough;
T2_Semi_Volume_to_Remove = (INPUT.HT_Vol - INPUT.COS_Vol) * INPUT.T2_Per_Semi;
T2_Finish_Volume_to_Remove = (INPUT.HT_Vol - INPUT.COS_Vol) * INPUT.T2_Per_Finish;

% Level 5
T2_Total_Disc_Flipping_Time = STD.T2_Flipping_Time * STD.T2_Number_of_Flips;

T2_Rough_Machinging_Time             =             T2_Rough_Volume_to_Remove             /
              STD.T2_Rough_Rate_of_Material_removal;
T2_Semi_Machining_Time = T2_Semi_Volume_to_Remove / STD.T2_Semi_Rate_of_Material_removal;
T2_Finish_Machining_Time             =             T2_Finish_Volume_to_Remove             /
              STD.T2_Finish_Rate_of_Material_removal;

% Level 4
T2_Setup_Time = STD.T2_Initial_Setup_Time...
               + T2_Total_Disc_Flipping_Time...
               + STD.T2_Setting_Down_Time;
T2_Total_Turning_Op_Time = T2_Rough_Machinging_Time...
                         + T2_Semi_Machining_Time...
                         + T2_Finish_Machining_Time;
T2_Total_Process_Time = T2_Total_Turning_Op_Time / STD.T2_Utilisation_Efficiency;

% Level 3
T2_Setup_Costs = T2_Setup_Time * STD.T2_Setup_cost_rate;
T2_Turning_Cost = (T2_Total_Process_Time) * STD.T2_cost_rate;


% Level 2
T2_Total_Should_be_cost = T2_Setup_Costs...
                        + T2_Turning_Cost;
T2_Profits = T2_Total_Should_be_cost * STD.T2_Profit_Margin;
```

```
% Level 1
OUTPUT.Total_Turning_Op2_Cost = T2_Total_Should_be_cost + T2_Profits;
```

```
function OUTPUT = Total_Ultrasonic_NDE_Cost(INPUT,STD,OUTPUT)

% Level 6
U_NDE_Inspection_Surface_Area = INPUT.COS_SF_Area * STD.U_NDE_Number_of_scanning_angles;

% Level 5
U_NDE_Scan_Time = U_NDE_Inspection_Surface_Area / STD.U_NDE_Inspection_Speed;
U_NDE_Probe_Adjustment_Time         =         STD.U_NDE_Number_of_scanning_angles         *
               INPUT.Number_of_COS_Edges * STD.Time_To_Adjust_Probe;

% Level 4
U_NDE_Op_Time = U_NDE_Scan_Time...
                + U_NDE_Probe_Adjustment_Time;

% Level 3
U_NDE_Setup_Costs = STD.U_NDE_Setup_Time * STD.U_NDE_Setup_cost_rate;
U_NDE_Inspection_Costs = U_NDE_Op_Time * STD.U_NDE_Cost_Rate;

% Level 2
U_NDE_Total_should_be_costs = U_NDE_Setup_Costs...
                              + U_NDE_Inspection_Costs;
U_NDE1_Profits = U_NDE_Total_should_be_costs * STD.U_NDE_Profit_Margin;

% Level 1

OUTPUT.Total_Ultrasonic_NDE_Cost = U_NDE_Total_should_be_costs...
                            + U_NDE1_Profits;
```

```
function OUTPUT = Total_Turning_Op3_Cost(INPUT,STD,OUTPUT)

% Level 6
T3_Rough_Volume_to_Remove = (INPUT.COS_Vol - INPUT.Disc_2D_Vol) * INPUT.T3_Per_Rough;
T3_Semi_Volume_to_Remove = (INPUT.COS_Vol - INPUT.Disc_2D_Vol) * INPUT.T3_Per_Semi;
T3_Finish_Volume_to_Remove = (INPUT.COS_Vol - INPUT.Disc_2D_Vol) * INPUT.T3_Per_Finish;
T3_Dif_Rough_Volume_to_Remove     =     (INPUT.COS_Vol     -     INPUT.Disc_2D_Vol)     *
               INPUT.T3_Per_Dif_Rough;
T3_Dif_Semi_Volume_to_Remove      =     (INPUT.COS_Vol     -     INPUT.Disc_2D_Vol)     *
               INPUT.T3_Per_Dif_Semi;
T3_Dif_Finish_Volume_to_Remove    =     (INPUT.COS_Vol     -     INPUT.Disc_2D_Vol)     *
               INPUT.T3_Per_Dif_Finish;

% Level 5
T3_Total_Disc_Flipping_Time = STD.T3_Flipping_Time * STD.T3_Number_of_Flips;

T3_Rough_Machinging_Time             =             T3_Rough_Volume_to_Remove             /
               STD.T3_Rough_Rate_of_Material_removal;
T3_Semi_Machining_Time = T3_Semi_Volume_to_Remove / STD.T3_Semi_Rate_of_Material_removal;
T3_Finish_Machining_Time            =            T3_Finish_Volume_to_Remove            /
               STD.T3_Finish_Rate_of_Material_removal;
T3_Dif_Rough_Machinging_Time          =          T3_Dif_Rough_Volume_to_Remove          /
               STD.T3_Dif_Rough_Rate_of_Material_removal;
T3_Dif_Semi_Machining_Time           =           T3_Dif_Semi_Volume_to_Remove           /
               STD.T3_Dif_Semi_Rate_of_Material_removal;
T3_Dif_Finish_Machining_Time          =          T3_Dif_Finish_Volume_to_Remove          /
               STD.T3_Dif_Finish_Rate_of_Material_removal;

% Level 4
T3_Setup_Time = STD.T3_Initial_Setup_Time...
                + T3_Total_Disc_Flipping_Time...
                + STD.T3_Setting_Down_Time;
T3_Total_Turning_Op_Time = T3_Rough_Machinging_Time...
                           + T3_Semi_Machining_Time...
                           + T3_Finish_Machining_Time...
                           + T3_Dif_Rough_Machinging_Time...
                           + T3_Dif_Semi_Machining_Time...
                           + T3_Dif_Finish_Machining_Time;
T3_Total_Process_Time = T3_Total_Turning_Op_Time / STD.T3_Utilisation_Efficiency;

% Level 3
T3_Setup_Costs = T3_Setup_Time * STD.T3_Setup_cost_rate;
T3_Turning_Cost = (T3_Total_Process_Time) * STD.T3_cost_rate;
```

```
% Level 2
T3_Total_Should_be_cost = T3_Setup_Costs...
                          + T3_Turning_Cost;
T3_Profits = T3_Total_Should_be_cost * STD.T3_Profit_Margin;


% Level 1
OUTPUT.Total_Turning_Op3_Cost = T3_Total_Should_be_cost + T3_Profits;
```

```
function OUTPUT = Total_Machining_Cost(INPUT,STD,OUTPUT)

OUTPUT.Total_Machining_Cost = INPUT.ThreeD_Feature_and_Finish_Cost;
```

This page intentionally left blank.

APPENDIX V – COST MODEL FIXED VALUES (CONTAINS PROPRIETARY DATA)

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

NOT FOR PUBLIC RELEASE

APPENDIX VI – CASE STUDY GEOMETRY AND GEOMETRY USED FOR NORMALISING DATA (CONTAINS PROPRIETARY DATA)

NOT FOR PUBLIC RELEASE

This page intentionally left blank.

APPENDIX VII – LOCATION OF DIGITAL CODE AND MODELS (CONTAINS PROPRIETARY DATA)

NOT FOR PUBLIC RELEASE

This page intentionally left blank.

REFERENCES, BIBLIOGRAPHY AND SOURCES

REFERENCES

Arrazola, P. J., T. Özel, et al. (2013). "Recent Advances in Modelling of Metal Machining Processes." CIRP Annals - Manufacturing Technology **62**(2): 695-718.

Azuma, T., Y. Tanaka, et al. (1997). Manufacturing and Properties of Newly Developed 9%CrMoVNiNbN High-pressure Low-pressure Rotor Shaft Forging. Steel Forgings: Second Volume. N. G. Nisbett and A. S. Melilli, American Society for Testing and Materials.

Bamberg, J. and M. Spies (2007). Opitimal Probe Arrangement for Ultrasonic Inspection of Spin Test Disks. The 34th Annual Review of Progress in Quantitative Nondestructive Evaluation (NDE). American Institute of Physics. **27**.

Bewlay, B. P., M. F. X. Gigliotti, et al. (2003). "Net-Shape Manufacturing of Aircraft Engine Disks by Roll Forming and Hot Die Forging." Journal of Materials Processing Technology **135**(2-3): 324-329.

Brujic, D., M. Ristic, et al. (2009). "CAD Based Shape Optimisation for Gas Turbine Component Design." Structural and Multidsciplinary Optimization **41**(4): 647-659.

Bunge, G. (2002). Method and apparatus of cooling heat-treated work pieces, US6394793 B1.

Bunge, G. and D. Furrer (2003). "Differential Cooling Technology Offers Jet Engine Design Engineers Performance Improvement Options." Aircraft Engineering and Aerospace Technology **75**(4).

Caputo, A. C., P. M. Pelagagge, et al. (2016). "Manufacturing Cost Model for Heat Exchangers Optimization." Applied Thermal Engineering 94: 513-533.

Castro, C. F., C. A. C. Antonio, et al. (2004). "Optimisation of Shape and Process Parameters in Metal Forging using Genetic Algorithms." Journal of Materials Processing Technology **146**(3): 356-364.

Chenot, J., E. Massoni, et al. (1996). "Inverse Problems in finte Element Simulation of Metal Forming Process." Engineering Computations **13**(2/3/4): 190-225.

Duverlie, P. and M. J. Castelain (1999). "Cost Estimation During Design Step: Parametric Method versus Case Based Reasoning Method." The International Journal of Advanced Manufacturing Technology 15(12): 895-906.

Farshi, B., H. Jahed, et al. (2004). "Optimum Design of Inhomogeneous Non-Uniform Rotating Discs." Computers & Structures **82**(9–10): 773-779.

Fisher, C. E., J. S. Gunasekera, et al. (1997). Process Model Development for Optimization of Forged Disk Manufacturing Processes. Steel Forgings 2nd Volume. E. G. Disbett and A. S. Melilli, American Society for Testing and Materials**:** 116-127.

Furrer, D. and H. Fecht (1999). "Ni-Based Superalloys for Turbine Discs." JOM **51**(1): 14-17.

Furrer, D. U., R. Shankar, et al. (2003). "Optimizing the Heat Treatment of Ni-Based Superalloy Turbine Discs." JOM **55**(3): 32-34.

He, B., P. J. Rohl, et al. (1998). CAD and CAE Integration With Aplication to the Forging Shape Optimization of Turbine Disks. 39th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference and Exhibit, American Institute of Aeronautics and Astronautics. **AIAA-98-2032:** 2742 - 2749.

Hebert, D., M. Friedman, et al. (2014). Cooling systems for heat-treated parts and methods of use. USA, Firth Rixson Limited. **US20140367898 A1**.

Huang, Z., C. Wang, et al. (2011). "Optimal Design of Aeroengine Turbine Disc Based on Kriging Surrogate Models." Computers & Structures **89**(1–2): 27-37.

JAR-E840 (2001). JAR-E Engines, Civil Aviation Authorities.

Jeniski Jr, R. A. and R. L. Kennedy (2006). "Nickel-Based Superalloy Designed for Aerospace." Advanced Materials and Processes **164**(12): 19-22.

Kesseler, E. and M. H. van Houten (2007). Multidisciplinary Optimisation of a Turbine Disc in a Virtual Engine Environment. 2nd European Conference for Aerospace Sciences (EUCASS), Brussels, also available as NLR-TP-2006-526.

Kim, D. and J. Park (1998). "Development of an Expert System for the Process Design of Axisymmetric Hot Steel Forging." Journal of Materials Processing Technology **101**(1-3): 223-230.

Klinger, C. and D. Bettge (2013). "Axle Fracture of an ICE3 High Speed Train." Engineering Failure Analysis 35: 66-81.

Kulon, J., P. Broomhead, et al. (2006). "Applying knowledge-based engineering to traditional manufacturing design." The International Journal of Advanced Manufacturing Technology **30**(9-10): 945-951.

Kulon, J., D. J. Mynors, et al. (2006). "A knowledge-Based Engineering Design Tool for Metal Forging." Journal of Materials Processing Technology **177**(1-3): 331-335.

Layer, A., E. T. Brinke, et al. (2002). "Recent and Future Trends in Cost Estimation." International Journal of Computer Integrated Manufacturing 15(6): 499-510.

Lonsdale, C. P. and R. D. Swarzell (2000). Recent Improvements in Wrought Railroad Wheel Production at Standard Steel. Railroad Conference, 2000. Proceedings of the 2000 ASME/IEEE Joint.

Manning, A. J., D. Knowles, et al. (2003). A Nickel Base Superalloy. **EP 1 193 321 B1**.

Masel, D. T., W. A. Y. II, et al. (2010). "A Rule-Based Approach to Predict Forging Volume for Cost Estimation during Product Design." International Journal of Advanced Manufacturing Technology **46**(1): 31-41.

Nagendra, S., J. B. Staubach, et al. (2005). "Optimal Rapid Multidisciplinary Responcse networks: RAPIDDISK." Structural Multidisciplinary Optimisation **29**(3): 213-231.

Niazi, A., J. S. Dai, et al. (2005). "Product Cost Estimation: Technique Classification and Methodology Review." Journal of Manufacturing Science and Engineering 128(2): 563-575.

Ou-Yang, C. and T. S. Lin (1997). "Developing an Integrated Framework for Feature-Based Early Manufacturing Cost Estimation." The International Journal of Advanced Manufacturing Technology 13(9): 618-629.

Poguet, J., A. Garcia, et al. (2002). Phased Array Technology Concepts, Probes and Applications. 8th European Congress on Non Destructive Testing. Barcelona, Spain**:** 1-6.

Pollock, T. M. and S. Tin (2006). "Nickel-Based Superalloys for Advanced Turbine Engines: Chemistry, Microstructure, and Properties." Journal of Propulsion and Power **22**(2): 361-374.

Rao, A. R., J.P.Scanlan, et al. (2007). "Applying Multiobjective Cost and Weight Optimization to the Initial Design of Turbine Disks." Journal of Mechanical Design **129**(12): 1303-1310.

Röhl, P., S. Srivatsa, et al. (1997). A Comprehensive Approach to Engine Disk IPPD. 38th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference and Exhibit and AIAA/ASME/AHS Adaptive Structures Forum, American Institute of Aeronautics and Astronautics. AIAA-1997-113: 05-39.

Rohl, P., B. He, et al. (1998). A Collaborative Optimisation Environment for Turbine Engine Development. 7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, American Institute of Aeronautics and Astronautics. AIAA-98-4734.

Schafrik, R. and R. Sprague (2008). "Superalloy Technology - A Perspeective on Critical Innovations for Turbine Engines." Key Engineering Materials **380**: 113-134.

Schreve, K., H. R. Schuster, et al. (1999). "Manufcturing Cost Estimation During Design of Fabricated Parts." Journal of Engineering Manufacture **213**(7): 731-735.

Shen, G. and D. Furrer (2000). "Manufacturing of Aerospace Forgings." Journal of Materials Processing Technology **98**: 189-195.

Szelazek, J. (1987). "Ultrasonic Detection of Cracks in Wheel Set Axles." NDT International 20(3): 177-180.

Tammieni, S. V., J. P. Scanlan, et al. (2007). A Hybrid Kowledge Based Swystem for Cost Modelling applied to Aircraft Gas Turbine Design. AIAA, ePrints Soton.

Urresti, I., S. Nikov, et al. (2009). Aerospace Gas Turbine Disc Distortion Modelling: Machining Sequence Optimisation. 12th CIRP Conference on Modeling Machining operations, . San Sebastian (Spain)**:** 13-20.

Winchell, W. (1989). Realistic Cost Estimating For Manufacturing: Second Edition, Society of Manufacturing Engineers.

Wise, S. (1968). "Ultrasonic Testing of Railway Components." Journal of the Institution of Locomotive Engineers 58(321): 77-110.

Wiseall, S., A. Phipps, et al. (2014). Method to Optimise Forgings for Ultrasonic Inspection and Cost. **GB1417762.0.**

Zerbst, U., S. Beretta, et al. (2013). "Safe Life and Damage Tolerance Aspects of Railway Axles – A Review." Engineering Fracture Mechanics 98: 214-271.

BIBLIOGRAPHY

Ausralian Transport Safety Bureau (2004). In-Flight Uncontained Engine Failure and Air Turn-Back, Boeing 767-219ER, ZK-NBC, 8 December 2002. Aviation Safety Investigation Report 200205780, Australian Government - Australian Transport Safety Bureau**:** 1-43.

Australian Transport Safety Bureau (2010). In-Flight Uncontained Engine Failure Overhead Batam Island, Indonesia, 4 November 2010 VH-OQA, Airbus A380-842. Aviation Occurrence Investigation - AO-2010-089 (Preliminary), Australian Government - Australian Transport Safety Bureau**:** 1-43.

Benz, M. G. (1999). Perparation of Clean Superalloys. Impurities in Engineering Materials, Impact, Reliability, and Control. C. L. Briant. New York, Marcel Dekker, Inc.**:** 31-48.

Davis, J. R. (1989). ASM Handbook, Volume 16 - Machining, ASM International.

DeAntonio, D. A., D. Duhl, et al. (1991). Heat Treating of Superalloys. ASM Handbook, Volume 4 - Heat Treating. A. I. H. Committee, ASM International**:** 794-814.

Donachie, M. J. and S. J. Donachie (2002). Superalloys a Technical Guide, Second Edition. USA, ASM International.

Krautkramer, J. K. a. H. (1983). Ultrasonic Testing of Materials, Springer-Verlag.

National Transportation Safety Board (1990). Aircraft Accident Report NTSB/AAR-90/06 - United Airlines Flight 232, McDonnell Douglas DC-10-10, Sioux Gateway Airport, Sioux City, Iowa, july 19,1989. N. T. S. Board**:** 123.

National Transportation Safety Board (1996). Aircraft Accident Report NTSB/AAR-96/03 - Uncontained Engine Failure/Fire, Valujet Airlines Flight 597, Douglas DC-9-32, N908VJ, Atlanta, Georga, June 8, 1995, National Transportation safety Board**:** 131.

Raj, B., T. Jayakumar, et al. (2002). Practical Non-Destructive Testing. Cambridge, Woodhead Publishing Limited.

Reed, R. C. (2006). The Superalloys: Fundamentals and Applications, Cambridge University Press.

Rolls-Royce plc (2005). The Jet Engine, Rolls-Royce.

S.K. Srivatsa (2003). "Application of Multidisciplinary Optimization (MDO) Techniques to the Manufacturing of Aircraft Engine Components." Handbook of Workability and Process Design: 346 - 367.

Semiatin, S. L. (2005). ASM handbook, Volume 14A - Metalworking: bulk-forming, ASM International.

Smid, P. (2010). CNC Control Setup for Milling and Turning, Industrial Press, Inc.

Srivatsa, S. K. (2003) "Application of Multidisciplinary Optimization (MDO) Techniques to the Manufacturing of Aircraft Engine Components." Handbook of Workability and Process Design (ASM international), 346-367.

Valberg, H. S. (2010). Applied Metal Forming Including FEM Analysis, Cambridge University Press.

SOURCES

Blakey, M. C. (2006). Saftey Recommendation, National Transportation Saftey Board.

Brooks, J. (2010a). Discussion on the Application of Generic Forging Rules (Conversation). Glasgow (22 April 2010).

Brooks, J. (2010b). Discussion on Axi-Symmetric Forging (Conversation). Southampton (09 December 2010).

Bryant, D. (2010). SAMULET 5.6.3 Techinical Review (Conversation). Rolls-Royce, Derby (01 November 2010).

Cristinacce, N. S. A. (2010). Meeting with Nik Cristanacce (Conversation). Bristol (22 March 2010).

FSF Editorial Staff (1996) "Uncontained Disk Failure in Right Engine of DC-9 During Initial Takeoff Run Results in Rejected Takeoff and Aircraft Evacuation." Flight Safety Foundation - Accident Prevention **2011**, 1-14.

Gayda, J., T. P. Gabb, et al. (2006). Advanced Heat Treatment Technology for Superalloy Disks Verified, NASA Technical Reports Server. NASA, 23 April 2010**:** 135:136.

Hersman, D. A. P. (2010). Safety Recommendation A-10-98. National Transportation Safety Board. H. J. R. Babbitt, Administrator, F. A. Administration, Washington and D. C. 20591. Washington, National Transportation Safety Board: 1-8.

Institute for Metals Superplasticity Problems of RAS. (2011). "Isothermal Rolling." Retrieved 11 July, 2011, from http://www.imsp.ru/userfiles/SRD-800-in-working.jpg.

Kolstad, J. L. (1991). Safety Recommendation A-91-085. National Transportation Safety Board. H. J. S. Busey, Administrator, F. A. Administration, Washington and D. C. 20591. Washington, National Transportation Safety Board: 1-4.

National Geographic. (2011). "Crash Landing in Sioux City." Retrieved 11 July, 2011, from http://channel.nationalgeographic.com/series/seconds-from-disaster/2387/Photos#tab-Photos/0.

Olympus Industrial. (2011). "B-Scan." Retrieved 11 July, 2011, from http://www.olympus-ims.com/en/ultrasonics/intro-to-pa/].

Rolls-Royce (2008). RRP 58008 - Methodology for Consitruction of Ultrasonic Inspection Scan Plans and Ultrasonic Limitation Diagrams, Rolls-Royce - Aerospace Process Specification.

Rosenker, M. V. (2006). Safety Recommendation A-06-060. National Transportation Safety Board. H. M. C. Blakey, Administrator, F. A. Administration, Washington and D. C. 20591. Washington, National Transportation Safety Board: 12.

The Ohio State University. (2015). "Forging Technology." Retrieved 23 December, 2015, from https://ercnsm.osu.edu/forging-technology.

University of Strathclyde. (2015). "AFRC Research." Retrieved 23 December, 2015, from http://www.strath.ac.uk/afrc/research/.

Wagner, S. (2004). "Evolution of Trent HP Turbine Rotor discs - Based on the Trent 900 HPT Rotor." Operations Engineering Technical Forum.

Wilde Analysis Ltd. (2011). "DEFORM-3D Cogging Example." Retrieved 12 July, 2011, from http://wildeanalysis.co.uk/system/photos/367/preview/DEFORM-3D_Cogging_Example.jpg?1271245324.

Wiseall, S. (2009, 27/01/2010). SAMULET Task 5.6.3. Launch Meeting: Forging Shape COS Optimisation of a Gas Turbine Disc for Unit Cost.

Wright, D. (2010). Discussion on Ultrasonic Inspection Practices (Conversation). Derby (09 February 2010).