

The University of Nottingham

UNITED KINGDOM · CHINA · MALAYSIA

Dynamics and Performance of Surface Finish Measurement Styli

By

Greg A. Duckworth MSci

Thesis submitted to the University of Nottingham for the degree of Doctor of Engineering

November 2019

Abstract

The main aim of this project was to characterise measurement stylus dynamic performance prior to manufacture and application. This work investigated each area of stylus production, and its influence of the dynamic performance of the stylus. The literature surrounding stylus dynamics and the wider performance of measurement instrumentation is sparse, with much of the development being undertaken in private.

At present, much of stylus design and production is based upon traditional attitudes and practices in; stylus component design, manufacture and stylus assembly, the use and application methods of adhesive, assembly fixturing technique and stylus materials. Investigations were performed into each of these specific areas to understand their impact on overall stylus production.

A novel system of stylus fixturing was developed. This adjustable fixturing is significantly more cost effective than current fixture implementations, whilst covering a wider range of applicable styli and allowing the in-process measurement of geometric tolerances. Stylus assembly and measurement methods were developed and implemented using the newly developed fixturing approach. This resulted in an increased assembly consistency, and an improvement in the repeatability and reliability of the dimensional measurement of styli.

The observation and characterisation of stylus dynamics has long been a challenging task. A system for the measurement of stylus dynamics was developed, using finite element analysis to inform design decisions, resulting in a realised measurement gantry. The gantry was characterised using accelerometers, allowing the Finite Element Analysis predictions to be tested. A piezoelectric actuator was used as a programmable measurement surface, this allowed automated testing of stylus performance on a range of simulated surface waveforms, with controllable amplitude and frequency.

The capabilities afforded by the fixturing system allowed the creation of custom styli, in conventional materials (Carbon Fibre, Aluminium and Ceramics), as well as unconventional materials (Resin, additively manufactured). These styli were tested using the gantry test system and the conventional materials were recommended for stylus applications in the future. The dynamic testing methodology allowed the observation of the primary resonant frequency of the stylus, even in circumstances where the stylus would undergo flight at frequencies below the stylus primary resonant frequency.

Acknowledgements

I would like to thank the multitude of people who have enabled me to complete this journey:

My partner, Nicola Hale, for her years of unwavering support, friendship and patience, and for putting up with my many interesting ideas.

To my parents, Paula and Leonard Duckworth, who graciously gave me a home and space to work, long after I moved away to university.

To my partner, Hannah Coop, for helping me through difficult times, and for her encouragement when I needed a gentle push.

To my grandparents, Paul and Olga Newstead, who continue to mentally challenge me and encourage me to always better myself.

To my brother, Ben Duckworth, for helping me to maintain perspective on my situation, in true sarcastic style.

To my academic supervisor, Simon Lawes, for supporting me academically and for dealing with my many quirks.

To my industrial supervisor, Mike Mills, for keeping me honest and patiently educating me in far too many ways to list.

To my friends and colleagues at Taylor Hobson and the University of Nottingham for freely their sharing of knowledge and expertise.

Without the support of everyone listed here, and many more besides, it is doubtful that I would have made it this far.

Thank you.

Table of Contents

1. Introduction 1
1.1. Background 1
1.2. Research Methodology 5
1.3. Aims and Objectives 11
2. Literature Review 13
2.1. Measurement Stylus Overview 13
2.2. Overview of Fixturing Theory 17
2.3. Range of Conventional Stylus Materials 19
2.4. Additively Manufactured Materials 25
2.5. Additive Manufacturing Processes 28
2.6. Ceramic Injection Moulding 36
2.7. Design for Manufacture and Assembly 38
2.8. Design for Additive Manufacture 41
2.9. Stylus Dynamics Theory 48
2.10. Stylus Dynamics in the Literature 56
2.11. Dynamics in Micro-CMM 61
3. Stylus Tolerancing, Manufacture and Production 63
3.1. Stylus Tolerancing 63
3.2. Summary of Current Stylus Fixturing and Development of Adjustable Fixturing 73
3.3. Use of Adhesives in Stylus Assembly 110
4. Materials Investigation - Stylus Dynamics 130
4.1. Stylus Dynamics Testing System 130
4.2. Initial Experimentation to Determine Parameter values for the Stylus Dynamics Measurement System 138
4.3. Realised Stylus Dynamics Measurement System 140

5. Materials Testing 179
5.1. Failed Styli 180
5.2. Realised Styli 180
5.3. Summary of Testing Methodologies 186
5.4. Range of Experiments Conducted on the Gantry System 189
6. Conclusions 229
6.1. Objectives 229
7. Further Work 232
7.1. Stylus Dynamics on Real Surfaces. 232
7.2. Repeatability of Stylus Mounting 232
7.3. The Development of a Suite of Standard Stylus Literature 232
8. Appendices 234
8.1. Investigation into the requirements of a stylus dynamics testing system 234
8.2. Coolant Absorption Data 239
9. References 240

1. Introduction

1.1. Background

1.1.1. Contact Metrology

Contact metrology is a vital part of mechanical engineering and production; it is used in a wide range of industries including; automotive, aerospace, lens making and optics, bearing manufacture, injection moulding, and even the concrete flooring industry. Simultaneous measurement of both surface finish and form, have allowed the development of now-ubiquitous technologies such as aspheric lenses, gothic arch bearings, turbine blades and fuel injection nozzles; even giving rise to an effort to redefine the standard of mass itself (Azuma, et al., 2015).

Modern manufacturing and production require that components must be standardised, and therefore interchangeable. This means that the components must be manufactured within tolerance, regardless of which specific machine tool or batch. These requirements have increased the demands placed on modern metrology, both from a dimensional and a “form and finish” perspective. The need for process control was the initial driving factor behind the development of commercial surface measurement instrumentation (Reason, Hopkins, & Garrod, 1944).

The performance of a mechanical system can be highly dependent on the surface finish of the individual components. The retention of oil on an engineered surface is related to the profile and finish of that surface. This has a profound effect on the performance of internal combustion engines; requiring cylinder liners to be made of a tough material and lapped to specific surface pattern and finish. (Willn, 1972), (Willis, 1986). The characteristics of the surface of the underside of downhill skis has been shown to be proportional to the sliding friction (Moldestad, 1999).

There are two general metrology paradigms: contact and non-contact. Each type of measurement system has advantages and specific applications that

would be difficult or impossible for the other measurement type. The measurement of 2D surface patches can be done using contact or non-contact methods but is considerably quicker when using 2D optical instruments. The measurement of the roughness of the inside wall of a 2 mm diameter hole is extremely difficult for optical systems and relatively straightforward for contact systems.

A schematic representation of a typical contact-based surface measurement instrument is shown in Figure 1–1. The measurement loop[1] for this instrument is shown in blue.

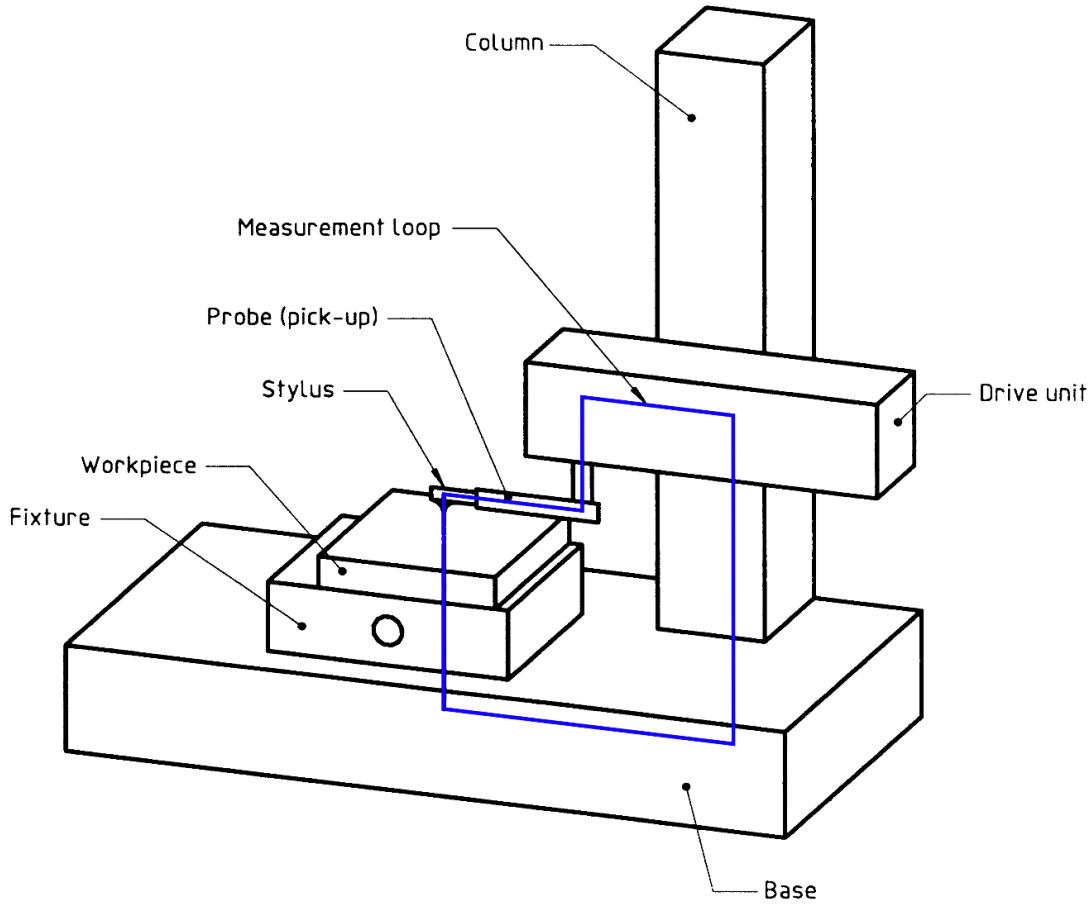


Figure 1–1: An outline of the typical surface finish measurement instrument, specifically highlighting the measurement loop (shown in Blue). Image taken from (ISO 3274, 1996).

[1] The measurement loop is "A closed chain which comprises all mechanical components connecting workpiece and the stylus tip, e.g. positioning means, workholding fixture, measuring stand, drive unit, probe (pick-up)" - (ISO 3274, 1996)

1.1.2. Taylor Hobson

Taylor Hobson has a rich history at the forefront of precision measurement and manufacturing. It was founded as a lens manufacturing company and initially developed the metrological instruments and techniques for internal use only. The company realised the value in releasing measurement instruments to the wider world and switch focus from lens manufacture to instrumentation and analysis.

In the present day, Taylor Hobson offers a wide range of contact and non-contact instrumentation, designed for specific markets including bearings, lenses, automotive, aerospace and defence. It offers the highest dynamic range surface instrumentation presently available and the company is developing instrumentation and software to support Metrology 4.0 and Industry 4.0.

1.1.3. Measurement Styli

Contact metrology systems typically use some form of stylus as a replaceable interface between the measurement gauge and the workpiece. A schematic of a common surface finish measurement stylus is shown in Figure 1–2. The stylus must perform several key roles within the general measurement system:

Rigid mounting of a measurement tip

The stylus must rigidly retain the reference surface; typically, a diamond conisphere or a sapphire sphere. This stylus tip must be rigidly retained such that the stylus dynamics do not interfere with the surface measurement. Stylus stiffness is becoming more important as surface roughness is improving with the increasing ubiquity of super-finishing technologies.

Repeatable and replaceable

Styli are typically a wearing component, and as such, each new stylus must be a drop-in replacement for the worn stylus. Geometric tolerances are necessarily small to ensure that automated measurement routines do not need adjustment after each stylus change.

Highly configurable geometry

The range and variety of surfaces and geometries that require measurement is so vast that it would be impossible to make a stylus capable of measuring them all. To that end, new geometries of styli must be able to fit into current measurement gauges to allow measurement of new geometries of component.

From a business perspective, styli should be affordable, as the cost of stylus replacement is an ongoing cost for the user. This means that the stylus must not be made from prohibitively expensive or difficult to machine materials.

The development of measurement styli was originally done in the 1940's. Since that time, many projects within academia and privately within Taylor Hobson have been concerned with the development of the measurement stylus instrument in accordance with the modern technologies. (Whitehouse, Enhancement of instrument and machine capabilities, 1996). Many materials have been tested (Jiang, Scott, Whitehouse, & Blunt, 2007) and new technologies adopted (Jiang, Scott, Whitehouse, & Blunt, 2007). There has been a concerted effort to simulate a measurement instrument, from simple rigid body simulations of styli (Damir, 1973), through to complex numerical methods that include stylus resonances and damping (Song & Vorburger, 1996).

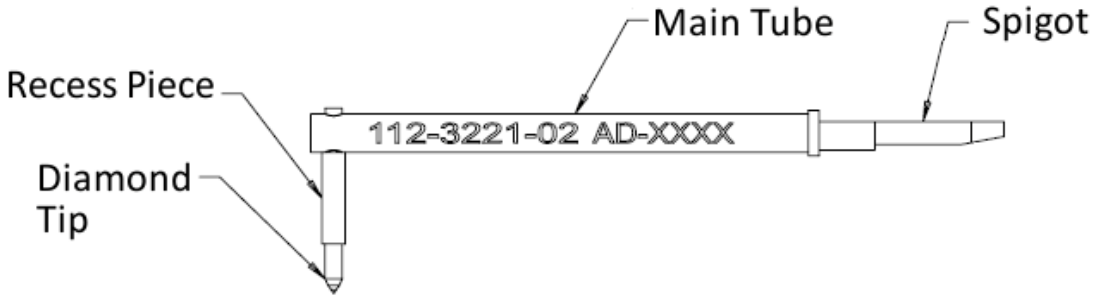


Figure 1–2: A common Taylor Hobson surface finish measurement stylus, including labels of the typical components.

1.1.4. Thesis

This thesis is borne of the latest effort of Taylor Hobson to innovate and improve their stylus production, from design, manufacture, assembly, testing and quality control. This thesis includes, either wholly, or in part, the commercial reports that have been generated for Taylor Hobson on; Tolerancing, Adhesives, Fixturing, Additive manufacture and Stylus Materials. It includes the development of a stylus vibration measurement gantry that was investigated and whose performance was specified in a prior MSc project. The automated testing routines and mass-data analysis techniques that were developed for this thesis are also included.

1.2. Research Methodology

The problems faced by Taylor Hobson are complex and unique, given their prominence within the metrology industry. The advent of new manufacturing methods and modern production technologies has enabled the mass production of surfaces that would not have been possible using traditional technologies.

An example of the challenges faced within modern stylus metrology is the measurement of fine-ground surfaces within a fuel injection port. A schematic of a fuel injection system is shown in Figure 1–3, for reference. The tolerances on the geometry and surface finish of a fuel injection nozzle are very exacting.

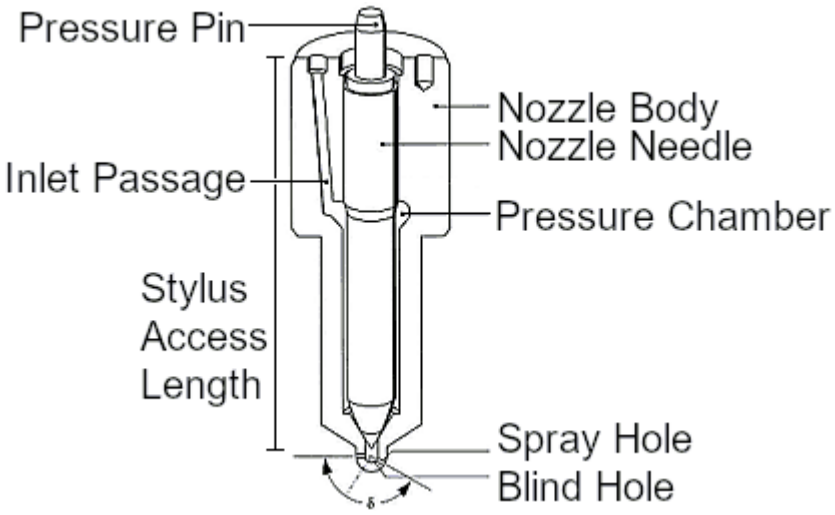


Figure 1–3: : A labelled diagram of a fuel injection nozzle (Jääskeläinen & Khair, 2013). Note the aspect ratio of the stylus access length to the diameter of the nozzle hole.

Grinding is typically used to achieve this tolerance reliably and the process must be finely controlled. Proper measurement of an injector nozzle should return both the surface finish of the main diameter and the angled land at the end of the bore; it should also be able to give the angular relationship, including the blend radius between the two features. To complicate the measurement further, these features can be as much as 100mm away from the opening of the bore.

An off-the-shelf stylus would not be suitable for measuring such a complex geometry, so a special range of injector-styli were developed (Figure 1–5). These can be made from as many as 20 separate components, each requiring complex manufacturing and assembly to produce a finished stylus. The produced stylus must be assembled to geometrical tolerances of ±50 µm between the diamond tip and spigot. The complexity of such a stylus means that it is possible to produce two styli which are visually identical, have the same mass and calibrate very similarly; and still give varying surface finish results.

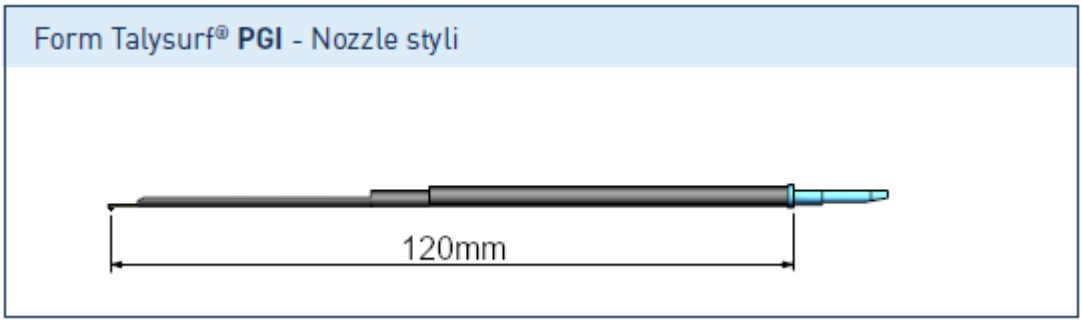


Figure 1–5: A render of a Nozzle stylus, used on PGI systems in fuel injector measurement applications. This stylus is 120mm long with a maximum diameter of 5mm. This design is necessary to measure fuel injector nozzles.

The long time period over which the stylus metrology field has developed has led to some present-day design choices being based in traditional thinking. This is a sensible approach for several reasons; backwards compatibility, support for instruments in the field (This can be 30+ years for some instruments), continuity between families of instruments and high production standards. Traditional thinking can also be a limitation as it makes it difficult

to integrate new production technologies and develop manufacturing and assembly processes.

1.2.1. The problems faced by Taylor Hobson

The challenges and problems faced within Taylor Hobson, regarding stylus production, can be categorised under three headings: process reliability, academic understanding and dynamic testing. These three areas of investigation provide a distinct and useful view of the challenges, each leading to specific aims and associated tasks to gain the understanding necessary for this project.

1.2.2. Process Reliability

A complete understanding of the production processes involved in stylus manufacture is necessary before improvements can be made. The current manufacturing and assembly processes are complex and have few overall guiding principles; often having several undocumented changes to mitigate manufacturing or assembly problems.

The production challenge faced by Taylor Hobson in manufacturing styli is large. Upwards of 5000 styli are manufactured each year, yet only 12 of the most popular models of styli make up 50% of production by volume. This is in stark contrast to the special styli, most of which are only ever produced in small quantity. There were more than 500 models of stylus made in 2017 which demonstrates the flexibility required of the design, manufacture, assembly and testing teams.

Most styli are made of 3 or 4 components, with geometrical tolerances of ±50 µm placed on the displacement between the tip and the interfacing spigot. (Figure 3–4). These assembly tolerances are readily achievable by a skilled technician. The same geometrical tolerances are placed on styli which have several tips. There may be as many as 4 tips, not necessarily of the same type, on styli with considerably more components and assembly steps.

The lack of guiding stylus design principles frustrates the communication from the designers to the production team and the lack of a quick and simple feedback mechanism does not allow assembly difficulties to be mitigated at design time. When styli are reviewed, only the higher-volume styli are considered. Specific design changes are made but these improvements are generally not transferred to lower volume styli. To further frustrate the feedback process, the CAD package used to design the styli has changed several times since moving to CAD in the 1980s.

The manufacturing and measurement capabilities have increased significantly in recent years with the acquisition of a HAAS 5-axis vertical milling machine and a Starratt AV300 optical measurement system. This has allowed a significant increase in the number of styli manufactured each year and has improved the tolerance of manufactured components and produced styli. The HAAS VMC has also made it possible to machine certain styli from solid carbon fibre, such that a separate spigot is not necessary.

The increased measuring capabilities have sped up the stylus inspection process and reduced the reliance on human reckoning. There has also been a push for more standard measurement routines, but this has been hampered by lack of an overall guiding document. The specification and tolerances placed on the stylus in the drawings is either difficult or impossible to measure using the current inspection fixturing. See section 3.2.6 for more details.

These issues, as observed in the first few weeks of the project, lead to a series of aims:

1. Investigate the design, manufacturing and assembly processes of the stylus cell.
2. Investigate the tolerancing, both theoretical and realised, for the produced styli.
3. Document and collate the information on the engineering choices that have been made regarding: materials, adhesives, geometry, fixtures and tip characteristics

4. Document the current processes to enable Taylor Hobson to make informed choices about the production processes going forward.
5. Collect data that would allow the creation of a set of standard documents for the production of styli in the future. (A set of stylus standards is beyond the scope of this work)

1.2.3. Academic Understanding

Taylor Hobson has a long history of collaboration with academia. This has allowed Taylor Hobson to keep abreast with the state-of-the-art in metrology. The solutions to the problems faced in this project should be informed by current academic thinking. The problems faced by Taylor Hobson suggest the following areas of research would be useful:

- Stylus design and tolerancing
- Adhesive selection and usage
- Proper cutting and subsequent cleaning of pultruded carbon fibre
- Fixturing techniques
- Suitable materials for stylus production
- Measurement using modern microscopes
- Gauge dynamics and damping
- Stylus flight and its effects on surface finish measurement

These areas of academic understanding cover a wide range of topics: requiring specific understanding of the nuances of current stylus production within Taylor Hobson and a wider appreciation of the state of the art in each area.

Individual studies into each area resulted in reports being generated and presented. These presentations typically included recommendations for changes to the current manufacturing and assembly processes.

A number of aims were generated from these requirements for increased academic understanding and as such, reports were written on:

- An evaluation of the tolerances in styli, both theoretical and what is achieved in practice.
- Adhesive types and application methods, including recommendations for future use
- Materials that can be used for stylus manufacture, including additive manufacture
- Modern fixturing and measurement as applied to stylus production
- A vibration testing gantry system, allowing automatic dynamic characterisation of styli, for use in production.

1.2.4. Stylus Dynamics Measurement System

The design of surface finish styli must fulfil several competing requirements. They must be low mass and low inertia, they should be very stiff, dimensionally stable and geometrically accurate to 50 µm to ensure comparability between stylus measurements. They need to be measurable and traceable, clearly display the model number and be cosmetically appropriate.

The aforementioned requirements placed upon a stylus are concerned with the materials properties and the geometric tolerance of a realised stylus. These requirements are incomplete as the dynamic performance of the stylus is not specified. For the dynamic performance of a stylus to be monitored, it must first be reliably measured. To allow the measurement of stylus dynamics, a measurement system must be devised.

1.3. Aims and Objectives

From the previously mentioned areas of investigation (1.2.2, 1.2.3 and 1.2.4), an overarching project aim was created:

1.3.1. Aim

To investigate the relationship between the dynamic performance of surface finish measurement styli and the design and production realities within a commercial metrology setting.

1.3.2. Objectives

From the stated aim, and with regard to the previously mentioned areas of investigation (1.2.2, 1.2.3 and 1.2.4), a series of objectives were chosen:

1. Understand how material properties affect stylus performance.
    a. Testing of a range of materials for performance in surface measurement applications
    b. Characterisation of gantry and its use for dynamics testing
2. Characterise the influence of current production processes on stylus performance
    a. Tolerancing, geometrical requirements and the approach to fixturing and assembly.
    b. Adhesive choice, adhesive application, non-fixtured features, component tolerances and in-process measurement
3. Investigate the applicability of Polymer AM, CIM to stylus production
    a. Design and produce styli using AM.
    b. Investigate AM and alternatives in context of styli.
4. Develop methods of predicting stylus performance, prior to application.
    a. Design, manufacture, assemble and characterise a stylus dynamics measurement system
    b. Develop data collection, analysis and mass-analysis techniques
    c. Investigate the dynamic performance of styli

    d. Gather and analyse enough data on the dynamic performance of styli to allow a useful specification and tolerance to be stipulated at design time.
5. Develop, improve or verify the current production strategies to enable the rapid and economical production of styli
    a. Understand production flow, from drawing stage to delivery
    b. Identify areas that can be improved, such as fixturing and adhesive approach
    c. Document current approach and its effect on business.

1.3.3. Deliverables

These aims and objectives, coupled with the original project brief from Taylor Hobson, lead to a set of deliverables:

1. Report on Tolerance and Fixtures
2. Report on Materials
3. Report on Adhesives
4. Report on AM
5. Development of Fixture techniques including proof of concept fixtures

2. Literature Review

2.1. Measurement Stylus Overview

2.1.1. Anatomy of a measurement stylus

The stylus in a contact-based metrology machine is vital to the proper performance of the overall system. The design of a measurement stylus has several requirements to fulfil. They must interface with the measurement gauge in a reliable and repeatable way, they need to constrain one or more 'contact geometries' that will actually contact the part to be measured. The stylus must be stiff such that any measured surface finish or form is transmitted to the gauge with good fidelity, this is especially important in high precision surface measurement or roundness measurement when measuring components with a high UPR (Undulations Per Revolution).

Individual stylus components are labelled in Figure 2–1 and discussed further in Section 3.1.1.

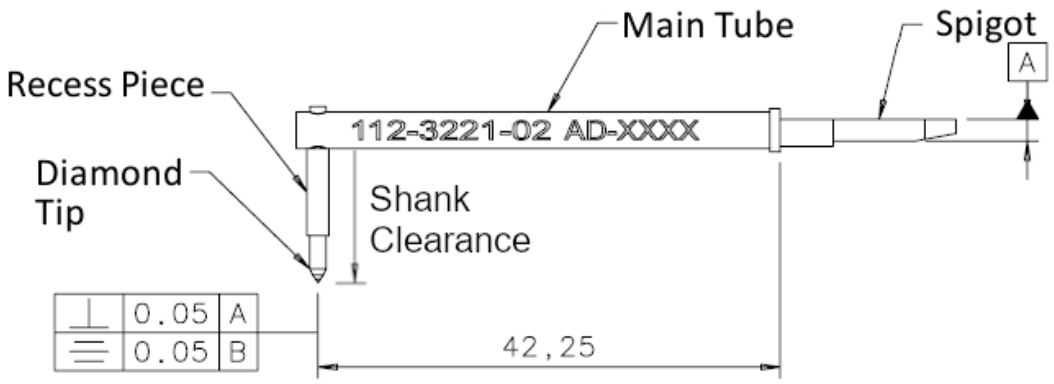


Figure 2–1: A schematic of a common surface measurement stylus. The basic components are labelled, and the geometric tolerances are shown in the table (bottom left). The shank clearance is also highlighted.

Surface styli also typically have several further requirements. The mass and moment of inertia must be kept low to keep the dynamic performance of the stylus-gauge system high. The geometry of surface styli is also typically more complex than roundness styli; this is because the surfaces to be measured are geometrically complex. Styli are expected to be drop-in replacements and so surface styli need a symmetry and perpendicularity tolerance which roundness

styli typically do not have. If a stylus is used to measure the form of a component such as a lens, then there is also an additional specification: shank clearance. Shank clearance is the distance from the underside of the stylus main tube to the measurement tip, as shown in Figure 2–1. Shank clearance allows the stylus to be used to measure behind the peak of a component such as a lens, or into the race of a bearing ring, without interference.

2.1.2. Mounting Considerations

Measurement styli are typically replaceable and must mount to the measurement gauge in a repeatable manner. The form of the stylus dictates the specific repeatability requirements. For example, styli which are rotationally symmetric, such as standard CMM and roundness styli, need to repeatably mount the stylus tip about the neutral axis of the gauge. This is very important for CMM applications as the gauge must accept contacts from the positive and negative X and Y axes, as well as from the negative Z axis. By contrast, a roundness measurement instrument will typically contact the component with the stylus in a single orientation.

CMM styli use a thread and a flat to establish a datum and to retain the stylus in the measurement probe. The roundness stylus uses the carbon fibre body, retained within a 'V' in the gauge, and a pin in the stylus provides the rotational constraint. Examples of these styli can be seen in Figure 2–2,. The manufacturing and assembly of the roundness stylus is more complex than

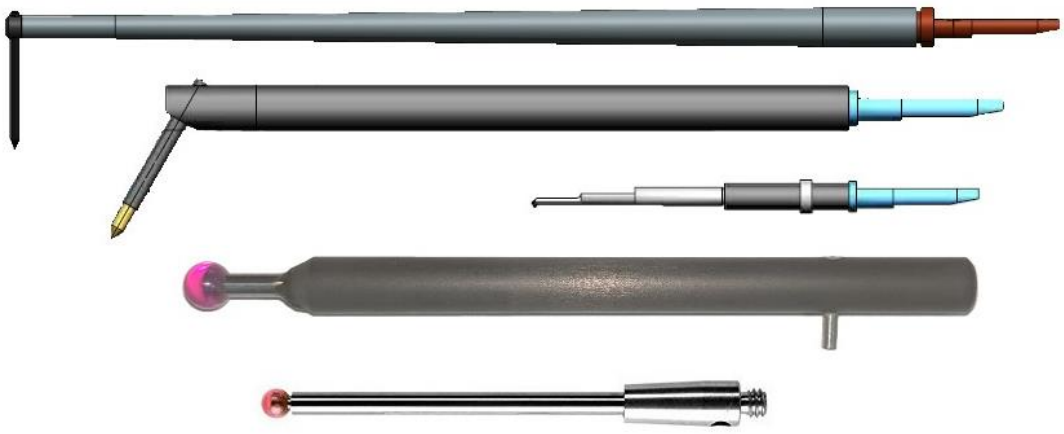


Figure 2–2: A range of measurement styli. From top to bottom: A long-reach surface finish recess stylus, a forward pointing surface finish stylus, a miniature-bore stylus, a standard roundness stylus and a typical CMM stylus.

the CMM stylus, although the mounting system used by the roundness stylus should have a better positional repeatability. The CMM probe is typically adjusted to ensure the probe tip runs true with the gauge axis.

Surface styli are constrained in two ways: a precision fit between a diameter and a bore, and a flat retained by a spring. The extension of the stylus from the stylus tube is controlled by a shoulder on the spigot contacting the front of the tube (B in Figure 2–3).

The spigot and stylus tube (shown in Figure 2–3) are held in contact by a retention spring. This spring sets the rotation of the stylus within the tube by pressing on the flat (Figure 2–3, labelled A). The spring also forces the spigot to contact the inside of the tube. The contact will be in the form of four line-contacts, as shown in green in Figure 2–3 (labelled C). The forces at the stylus tip are minimal (~200mg maximum), as such the stylus is unlikely to move within the stylus tube.

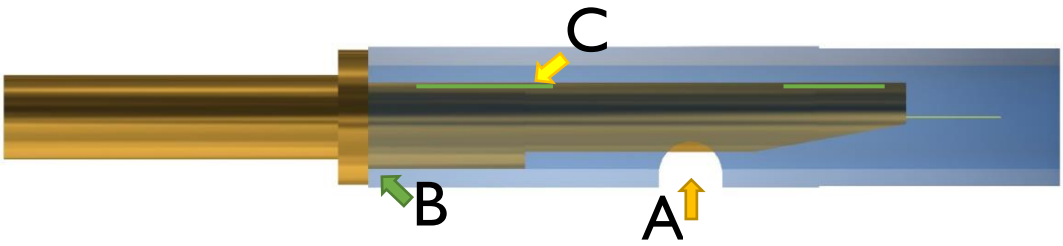


Figure 2–3: A render of a surface measurement stylus spigot (orange) inside a gauge tube (blue). The notch (labelled A) is clearance for the retention spring to press on the spigot. The shoulder of the spigot sits against the end of the gauge tube to set the extension of the stylus.

The surface finish stylus mounting system requires the highest precision of the three designs, in both component manufacture and assembly. The three highlighted mounting systems are not kinematically constrained, nor are they resistant to improper contact between the stylus and a component.

If a stylus is driven into a component or fixture, even at low speed, the results can be the destruction of the gauge and/or the stylus. This contact causes the indentation of the bearing balls onto the pivot pins. This results in noise and sporadic measurement issues, typically when the stylus passes through a certain angle, allowing the pin to become slightly offset as a ball bearing drops

into the indent. Mahr have demonstrated contour measurements with a stylus affixed to the gauge by way of a kinematic 3-ball mount, retained by magnets (Mahr, 2016). This system will protect the gauge in the event of most crashes, although it will not protect the stylus. A closeup of the kinematic stylus mount is shown in Figure 2–4.

The magnetically mounted, kinematically located stylus has several advantages over the other mounting approaches. The repeatable locating of the stylus removes the requirement for recalibration after each stylus change as calibration data can be recalled electronically for each stylus. The magnetic retention system limits the force applied to the gauge and allows the stylus to breakaway if it is improperly contacted.

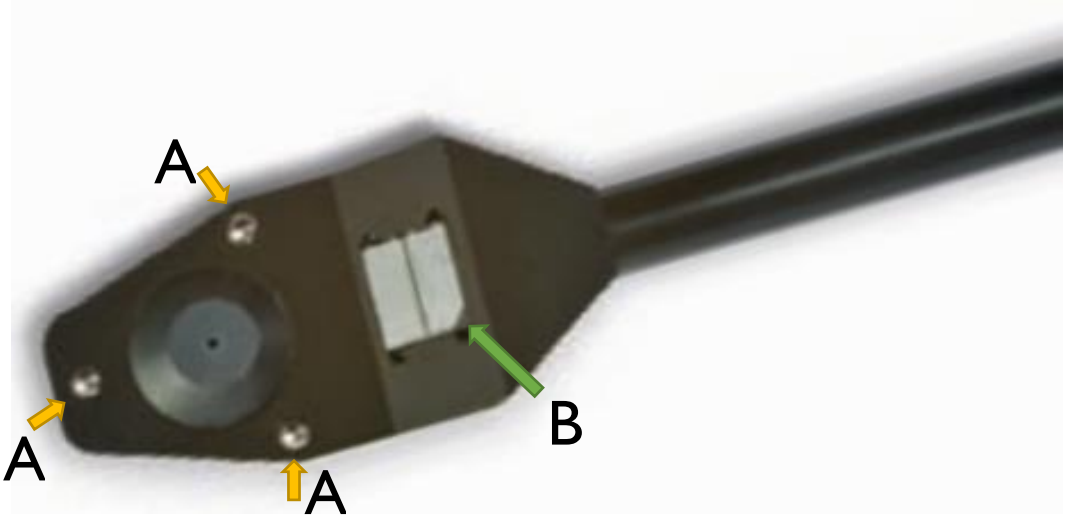


Figure 2–4: A closeup of the Mahr contour measurement stylus, focussing specifically on the kinematic mounting points (Labelled A) and the retention magnet (Labelled B). The measurement tip is not shown. (Mahr, 2019)

However, this mounting approach also greatly increases the moment of inertia of the stylus, limiting the system to low-frequency contour applications as opposed to high-frequency surface finish applications. Further, the magnet is attracted to any ferrous component, either in the measurement gauge or as a part to be measured. The variance of the magnetic force leads to an unpredictable stylus tip force, and therefore poor-quality surface finish measurements. Contour measurement systems typically use a longer stylus;

therefore, the magnet is further from the component to be measured. Contour systems are not as sensitive to changes in tip force, mainly because the contour measurement stylus tip radius is much larger than a surface finish measurement stylus tip.

Also, the Mahr kinematically mounted stylus is different to the other stylus mounting methods in that the stylus is mounted on the side of the gauge, necessarily offset from the bearings within the gauge. In contrast, surface finish measurement gauges typically mount the stylus between the bearings, minimising any cantilever effects that the offset stylus would have on the gauge.

2.2. Overview of Fixturing Theory

Fixturing is the process of work-holding and component registration within manufacture and assembly. This is done by selectively constraining the number of degrees of freedom of the component(s). Part location and part retention are usually considered separately, although at the small forces associated with stylus assembly make this separation less important than for larger components.

The fixture design should be robust against an acceptable range of component tolerances. The component tolerances are dictated by the manufacturing process. This dependence may limit the range of available manufacturing methods or will impact on the achievable tolerances of the final assembly.

There are many competing factors in fixture design. The older (more traditional) designs of stylus fixture place value in the longevity of the fixture and are not typically designed to utilise other tools during the assembly process. The choice of material and the geometrical design of the fixture indicate that it was designed to be produced using manual processes. The benefits of longevity are offset by the manufacturing cost and lead time for this fixturing approach.

An important consideration when designing fixtures is to identify the key physical relationships that should be embodied within the finished part. For measurement styli, the most important relationship is the displacement from the reference surfaces of the spigot to the stylus tip. In most typical measurement scenarios, where the stylus shank clearance is not vital, the only important relationship is this offset between the spigot and the diamond. Although the visual presentation of such a stylus plays an important commercial role.

A fixture which directly constrains the diamond tip as well as the spigot would satisfy the basic assembly criteria most simply. This may allow the fixture to be ‘dialled in’, after which the fixture can be used without need for further adjustment. This approach does not easily allow a range of shank lengths to be made in the same fixture. This fixture, once set, would not rely on measurement, but would not benefit from it either. This type of fixture is susceptible to tolerance changes and may be difficult to implement.

For further information on fixturing see (Nee, Senthil Kumar, & Tao, 2004)

2.3. Range of Conventional Stylus Materials

Measurement styli have been traditionally manufactured from a range of materials, including aluminium, beryllium-copper, carbon fibre, silicon nitride, alumina and a range of other ceramics (Renishaw, 2020), (Taylor Hobson, 2019). The size and shape of measurement styli has been dictated by the usage of the instrument and the engineering trade-offs that the component design entails.

Simply increasing the length of the stylus will effectively increase the range of the measurement gauge, although the resolution is degraded proportionally. The increase in length of the stylus will also change the dynamics of the stylus-gauge system. Longer styli will have a larger moment of inertia; this reduces the ability of the stylus to respond to higher surface frequencies. A stylus with a lower moment of inertia will track a certain amplitude of sine wave to higher frequencies than a stylus with a high moment of inertia, even if the static surface force is identical. (Tian, Liu, Chetwynd, Shirinzadeh, & Zhang, 2010)

The choice of material will influence everything from the initial design to the manufacturing methods, choice of adhesive, appropriate fixturing techniques and overall dynamic performance. The perfect material would be light, stiff, machinable, affordable, bondable, dimensionally stable[2] and have a low Q-factor[3]. No single material has all these properties and so compromises must be made. The volume of production (1000's per annum) and the range of designs (500+) requires that components be manufacturable in small to medium quantities. This generally excludes high volume production techniques, such as injection moulding, as the large variety of components would make production expensive.

[2] The dimensional stability of a material is a measure of the linear dimensional change from exposure to heat or other environmental conditions. (Surberg, Stratton, & Lingenhöle, 2007)
[3] The Q-factor or quality-factor is a dimensionless parameter describing the damping properties of an oscillator. A high Q-factor (>0.5) corresponds to a physical system that will resonate for a long period of time before coming to rest. A low Q-factor (<0.5) corresponds to a system that will not resonate but instead exponentially decay to rest. (Harlow, 2004).

The influence of moment of inertia on measurement performance means that the material stiffness per unit mass (specific modulus) is a more useful material property than plain material stiffness. The specific modulus of a material is a useful comparator, typically resulting in a design with the lowest mass and therefore lowest moment of inertia, for a given stiffness (manufacturingterms.com, 2020). The pivot of a typical measurement gauge is ~20mm behind the front extent of the stylus tube. This pivot is the point around which the moment of inertia is measured. This means that mass at the spigot-end of the stylus will not have as great an effect on the moment of inertia as mass at the tip.

A conceptually straightforward approach to maximise the stiffness of the stylus, while keeping the moment of inertia low, is to make the stylus tapered. This concentrates the mass near the spigot, where it will have less of an effect on the moment of inertia, whilst maintaining the overall stiffness. See (Silverman & Farrah, 2018) for further details on Other approaches, such as additive manufacturing, can allow for more complex designs, potentially resulting in a better stiffness-inertia trade-off.

2.3.1. Aluminium

The spigot and stylus tube are made from aluminium to minimise the mass of the moving parts in the gauge. Aluminium alloys were historically used as a stylus material in form measurement systems, although it has been superseded by other materials such as carbon fibre. The stiffness of the aluminium stylus is comparatively low and therefore the surface measurement performance is limited. Aluminium has a low damping coefficient, meaning that the stylus is more susceptible to external vibration, and that induced vibration will take longer to dissipate than for other materials. Aluminium can be anodised to improve the surface hardness and is dimensionally stable if the stock is heat-treated, it also has a comparatively low density. The common 6000-series aluminium alloys are readily machinable especially if heat treat and aged (Commonly T6). Materials properties for aluminium can be found here: (AZOM, 2005)

2.3.2. Beryllium-Copper

Beryllium Copper is a copper alloy with a considerably higher density than aluminium or carbon fibre (~8.25 mg/mm3). It is typically used in repeated stress applications such as springs or load cells. It is readily machinable (including stamping and acid-etching) and is good at dampening vibration. It is currently used for components within measurement instruments such as flexures or springs. The cost of the material is higher, and it is less readily available, than aluminium. The good vibration dampening properties would make Beryllium copper a good candidate for stylus designs but there are some considerable limitations.

The density of Beryllium Copper is much higher than carbon fibre (8.25 kg/m3 vs 2.5 kg/m3), this would require considerable design and manufacturing to keep the moment of inertia and stiffness comparable to carbon fibre. Further, the Beryllium in the alloy is toxic and the material is not ROHS compliant. This would necessitate a costly waste management process for manufacturing and distribution. Materials properties for Beryllium-Copper can be found here: (MatWeb, 2020)

2.3.3. Carbon Fibre

Carbon fibre is a composite material consisting of a matrix of individual carbon fibres which are impregnated with resin. This can be in the form of thin woven sheets, stacked into layers and then impregnated with resin; or individual fibres coated with resin, then drawn through a die, in a process called pultrusion. Since the components that make up a stylus are small, and the variety of styli is large, machined pultruded carbon fibre tube makes the most practical sense and so is used for the main tube and the recess piece. The carbon fibre tube is pre-ground to ensure the OD is consistent and within the engineered tolerances. The ID of the carbon fibre tube is governed by the mandrel used in the pultrusion process.

This material used in stylus manufacture is available in pre-ground lengths which simplifies the production process. It has a relatively low density (~1.55 mg/mm3) and a very high specific stiffness, although the directional

nature of the carbon fibres does make the material properties non-homogeneous. The fibrous nature of the material results in delamination and fibre pull-out during manufacture, meaning special machining strategies and tooling is recommended. The material is porous and should be thoroughly cleaned before adhesive is applied, the porosity means that adhesives will bond well with the surface. The abrasive nature of carbon fibre will result in fast tool wear during manufacture, meaning frequent checking of achieved tolerances is necessary. Materials properties for Carbon Fibre can be found here: (Tap Plastics, 2019)

2.3.4. Silicon Nitride

Silicon nitride is a hard, high strength ceramic compound that is typically used for high temperature applications where mechanical properties must be preserved. It can be powder processed and is used extensively in bearings. It is relatively low density (2.8 – 3.2 mg/mm3) and very high strength and stiffness. Current, large-radius stylus tips are made from silicon nitride, owing to its wear resistance.

This material is difficult to machine and has a limited range of appropriate adhesives, but powder processing may allow the high-volume production of small parts such as the spigot or Novus end-piece, by ceramic injection moulding. The small grain size and the high strength of Silicon Nitride means that thin-walled components, such as the Novus end-piece, would be reliably producible, in bulk. The machinability of this compound limits its use for general stylus applications, although simple styli, such as roundness styli, may be producible. Materials properties for Silicon Nitride can be found here: (Azom, 2001), (Riley, 2004)

2.3.5. Macor™ and other Ceramics

Ceramic is a term given to non-metallic materials with crystalline or partly crystalline structure. They are typically very hard, high stiffness materials with applications in aerospace, automotive, bearings, turbine and cutting tool technologies. They can be shaped using a variety of methods, including machining (typically using special tools), injection moulding and casting.

Contact coordinate measurement machines have used ceramic styli (typically alumina or similar) for a long time, relying on the ceramic stylus breaking to save the touch probe in the event of a crash.

The properties of ceramics vary depending on the exact ceramic conditions. Many ceramics can be readily formed and machined in a 'green' state and then transformed, typically using heat, into the harder form of the material. This transition is typically also accompanied by a predictable shrinkage as the ceramic is consolidated. Once hardened, ceramics are typically considerably more difficult to machine as they are as hard as the tools used to machine them.

Macor™ is a glass-ceramic which was developed by Corning Inc. It was developed to be much more machinable than other ceramics, without loss of stiffness by comparison to other ceramics. This material is readily available from ceramics suppliers and has good dimensional tolerances from the extrusion process.

Machining features within Macor, such as drilling cross holes, is certainly possible, although is very difficult and time consuming. The machining process creates a mildly irritating dust which should not be inhaled, but this may be readily managed using coolant. This is similar to the dust released from dry-machining carbon fibre. Materials properties for Macor can be found here: (Precipart, 2018), (Morgan Technical Ceramics, 2019)

2.3.6. Comparison between Pultruded Carbon Fibre and Alumina

Both carbon fibre and ceramics are commonly used for measurement styli in surface finish and CMM applications, respectively. Alumina, Macor, silicon nitride and silicon carbide are common ceramics used in wider manufacturing. Alumina and Macor are used in stylus manufacture for CMM purposes because of their high strength and brittleness (machine tool touch probes rely on the self-sacrificing nature of the stylus to protect the gauge in the event of a crash (Renishaw, 2019)). Pultruded carbon fibre, in the form of hollow

cylinders and solid rods, are currently used in the production of Taylor Hobson styli.

Comparing pultruded carbon fibre (PCF) against alumina; their specific stiffnesses are comparable although few of their other properties are as similar. PCF is a non-homogenous material with a different set of materials properties along the fibre orientation axis than transverse to this axis. The cross section of pultruded carbon fibre can be seen in Figure 2–5.

For stylus applications, the fibres are oriented along the axis of the main tube and along the axis of the recess piece. The fibres being aligned along the axis of the stylus is beneficial as any bending of the stylus will result in increased fibre tension; the individual fibres have a high tensile strength which in turn resists the bending forces. The positions of the fibres within pultruded carbon fibre are typically disordered, with each fibre having a varying number of neighbouring fibres and epoxy. This non-uniformity does not affect the gross performance of the carbon fibre as the fibre diameter is much smaller than the diameter of the stylus components. The distribution of fibres (grey) within the epoxy (black) can be seen in Figure 2–5.

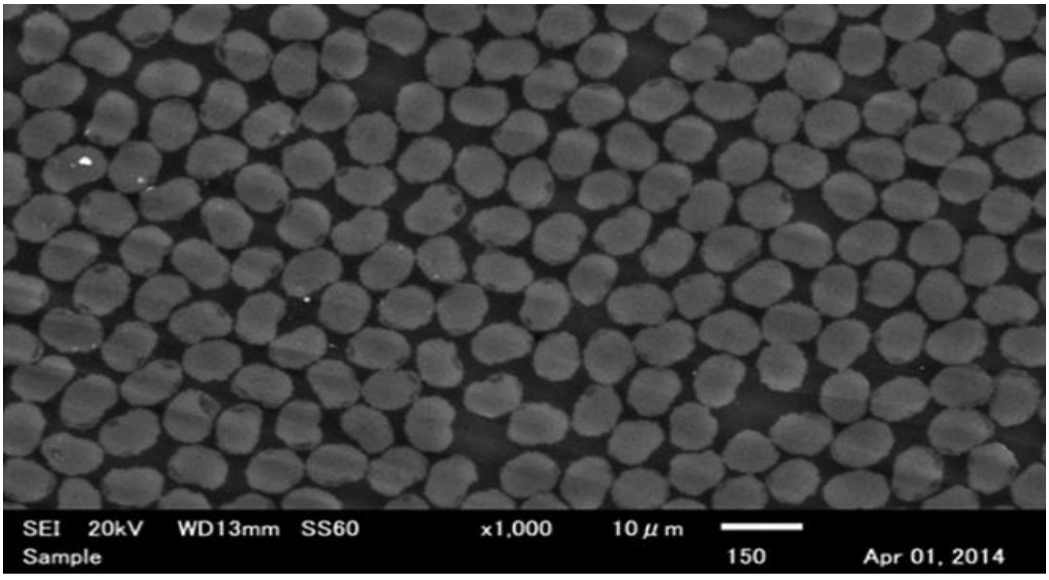


Figure 2–5: A microscope image of the cross section of pultruded carbon fibre. The fibres (grey) are surrounded by epoxy (Black). Note the fibre density within the resin medium (Bing, Anchang, & Limin, 2015)

Alumina is brittle, with a very low deformation before failure, whereas carbon fibre does not have a brittle failure mode. The density of alumina is higher than PCF, meaning that the same geometry cannot be used interchangeably with each material, if the same moment of inertia is desirable, as is the case with stylus design. The table of properties (Table 2-1) gives PCF a distinct advantage over Alumina in terms of weight, stiffness, and thermal expansion. However, Alumina can be ceramic injection moulded, allowing for a higher production volume, without sacrificing production tolerances.

Property	Alumina	PCF
Stiffness	360 GPa	228 GPa
Density	3.8 mg/mm^3	1.75 mg/mm^3
Specific Stiffness	95 GPa mm^3/mg	130 GPa mm^3/mg
Thermal expansion coefficient	7.2 µm/mK	±0.1 µm/mK
Failure Mode	Brittle failure	Fibre pull-out, splitting.
Machinability	Low	Medium
Cost per stylus	~£5	~£15
Adhesives	Epoxies, acrylates, ceramic	Epoxies, acrylates
Notes:	Ceramic Injection Moulding	

Table 2-1: Material property comparison of Alumina and Pultruded Carbon Fibre. (Tap Plastics, 2019) (Morgan Technical Ceramics, 2019) (Precipart, 2018)

2.4. Additively Manufactured Materials

There are a number of additive manufacturing technologies that are both capable of the precision and repeatability required for stylus manufacture and cost effective enough to compete against CNC manufacturing. Although other technologies exist, the additive technologies that process materials appropriate for stylus manufacture are: sintering, fused-filament fabrication and resin-based printing. A brief description of each of these processes is available in Section 2.5. Examples of a single component produced using these processes can be seen in Figure 2–6.

The range of materials that can be used in fused-filament fabrication include plastics such as; ABS, PLA, Nylon, ASA, Polycarbonate, PC-ISO, polyphenylsulfone, PET and TPU. (3D Hubs, 2015), (Computer Aided Technology (CATI), 2019)

Sintering can produce components made from; polypropylene, polyamide and polyamide blends (including blends with metal), Nylon and Nylon blends. Direct Metal Laser Sintering (DMLS) can produce components from; Aluminium alloys, Copper alloys, Nickel alloys, Steel, Stainless Steel, Titanium and Titanium alloys. (GE Additive, 2019), (Stratasys, 2019)

Resin printing is restricted to UV curable resins although limited success has been had with resin with suspended particles of a different material. Resin printing can produce components in materials such as Silicone or UV curable resins containing carbon fibres for added stiffness and thermal resistance. (3D Hubs, 2019), (Formlabs, 2019)



Figure 2–6: A demonstration of the 3 previously mentioned additive manufacturing technologies. The layer artefacts, surface finish and minimum feature size can be seen on each piece. (Res3D, 2018)

The range of materials listed in Section 2.4 is large and contains materials that are not suitable for stylus manufacture or even fixture manufacture. Many engineering materials in the additive field are developed for biocompatibility or to cleanly burn-out from a mould for casting purposes. (Alifui-Segbaya, Varma, Lieschke, & George, 2017), (Maker Juice Labs, 2019). The production of fixtures using additive techniques could use any of the common materials found in these technologies.

Table 2-2 details some of the mechanical properties of the specific resin used in the production of the additively manufactured stylus (Formlabs™ Tough Resin), along with a comparison material: Nylon. Generally, the cured resin is less mechanically capable than Nylon, with the exception of the Young's modulus (suggesting that the resin will retain its shape better under load than the Nylon).

The Formlabs™ resin should theoretically perform mechanically worse than Nylon, for a given geometry; although the geometric freedom associated with additive manufacture will help to improve the specific stiffness of properly designed AM components.

Property	Formlabs™ Tough Resin	Nylon
Tensile Strength	55.7 MPa	70 MPa
Young's Modulus	2.7 GPa	1.8 GPa
Flexural Strength	60.6 MPa	117 MPa
Flexural Modulus	1.6 GPa	1.8 GPa
Impact Strength	38 J/m	64 J/m
Adhesives	Epoxies, acrylates	Some acrylates
Notes:		Adhesive selection is limited

Table 2-2: Material property comparison between Formlabs™ Tough Resin (SLS) and Nylon (Conventionally Machined). (Formlabs, 2019)

2.5. Additive Manufacturing Processes

Additive manufacturing is an umbrella term that incorporates a large range of technologies and materials; from plastic and resin processes, to powder bed laser processes which can produce parts in metals or ceramics, even processes which use concrete to manufacture complete houses. (All 3DP, 2019). Figure 2–7 shows the large range of additive technologies, each having distinct advantages and application areas.

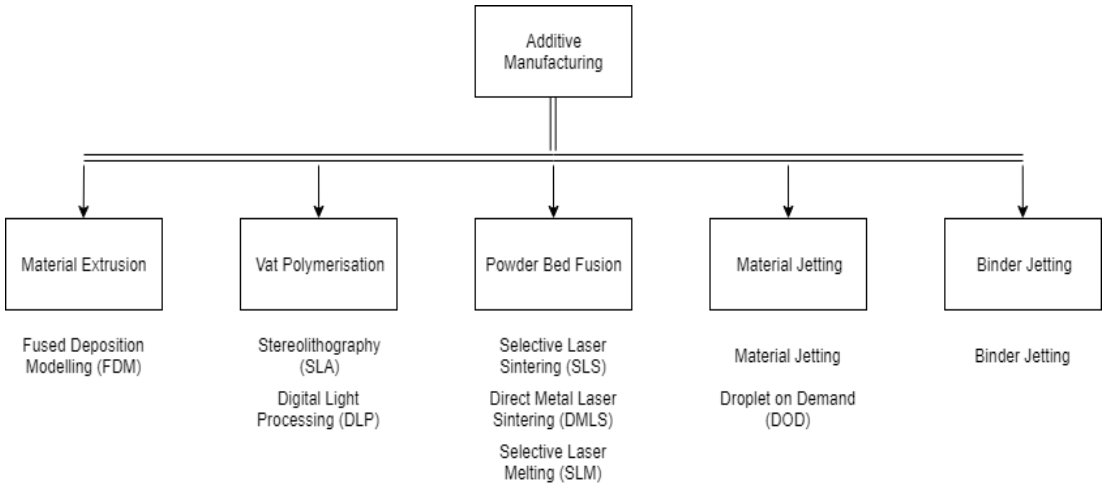


Figure 2–7: An overview of a range of Additive Manufacturing processes. Image adapted from (3D Hubs, 2018)

2.5.1. Powder Bed Fusion

Selective laser sintering (SLS) and selective laser melting (SLM) are examples of powder bed fusion processes. They use a bed of powdered material which is then selectively heated using a laser/ion beam. This creates a small volume of bonded material. After each layer is complete, the bed is lowered, and fresh powder is applied. These technologies can use a wide range of materials from metals such as, aluminium, copper alloys, steels, stainless steel, tool steel, titanium, nickel super-alloys and gold and plastics such as nylon (neat or glass-filled), polystyrene and green-sand for casting. (Loughborough University, 2019), (Song, Li, Deng, & Hu, 2007)

Sintering can be used to produce large and complex components from materials such as nylon, titanium, stainless steel, aluminium and other rarer

metals. The resultant model is fragile and requires consolidation which is typically achieved in an oven. The post-sintering heating causes the component to shrink and to significantly reduce the internal porosity. Material densities of 99% (sintered vs bulk) are possible, with strength and stiffness characteristics to match. (Mower & Long, 2016) The surface finish is usually rough in texture, although it varies with the specifics of the manufacturing process. The finished part can be machined to improve the finish or to achieve the dimensional tolerances required.

The surface finish of sintered components is rough compared to conventional manufacturing and other additive technologies, although this can be useful for intentionally textured surfaces. The post-sintering consolidation step reduces the layer artefacts, although it can slightly alter the geometry of small features.

The localised heating that is done as part of the sintering process introduces some stresses into the final component. This can result in warping during and after the printing process, especially on high aspect ratio parts. These effects can be mitigated by careful process control. A schematic of a laser-based powder bed fusion machine can be seen in Figure 2–8.

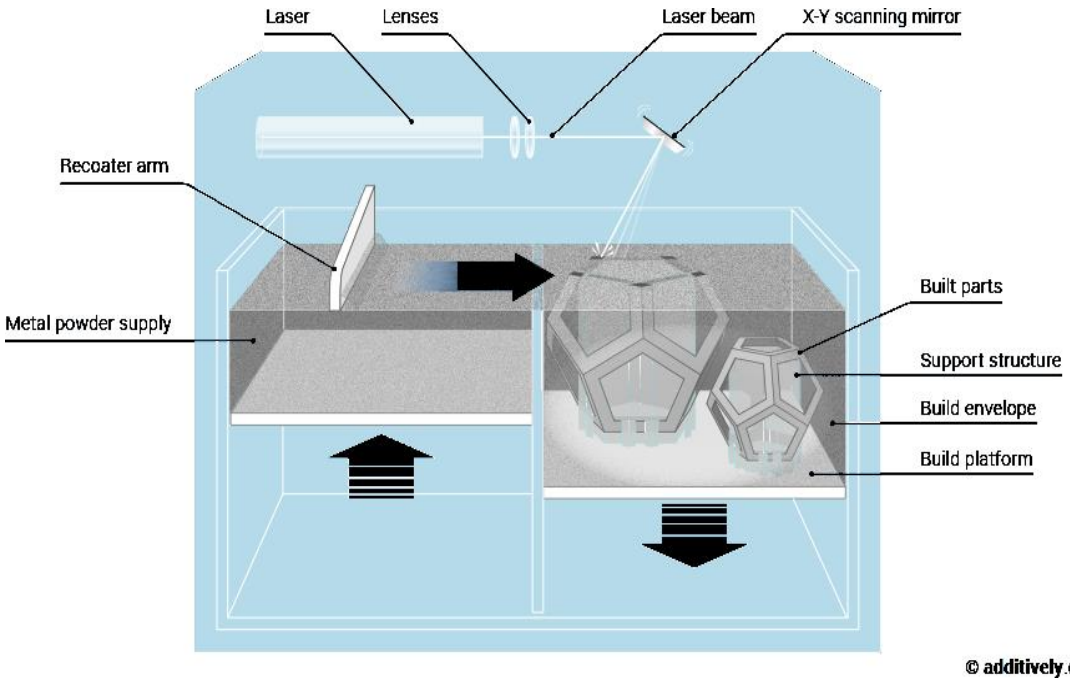


Figure 2–8: A labelled figure of a typical SLS machine, showing the powder supply (left) and the build chamber (right). Image courtesy of (Additively, 2018)

Capabilities

SLS and SLM both have a minimum layer height of 0.06 mm, with a minimum wall thickness of ~1 mm and a minimum feature size of 0.5 mm. Most features of a typical printed component will not need support, thereby reducing post-processing time. The surface produced by the machine is typically a 'matte-like grainy surface'

2.5.2. Material Extrusion

Fused Filament Fabrication (FFF), also referred to as Fused Deposition Modelling (FDM), is a thermal extrusion process that is typically limited to plastics, although a range of engineering plastics and plastics with reinforcing materials are becoming readily available. A plastic filament is extruded through a heated nozzle which is moved to build up a component in layers. The first layer is printed onto a heated bed which can be problematic for parts without a flat face. Implementations of this technology allow for dissolvable support material, multiple extruders, carbon fibre reinforced plastics, power-failure protection, customisable support features, automatic bed levelling, filament sensing and print queuing. A schematic of a typical FDM machine is shown in Figure 2–9.

This technology is typically used to produce proof-of-concept designs, rapid prototyping and fast iteration of components, custom jig manufacture and small quantity production of cases or other non-structural components. The print time is proportional to the component volume, layer thickness and the nozzle diameter. Support material is typically simple to remove, and post processing is straightforward, although time-consuming. (Tofail, et al., 2017)

The layer artefacts are typically more evident for material extrusion processes than for sintering, typically requiring some post processing to remove completely. Holes and small features will also require tuning of the printing process to be printed reliably. Components manufactured using this technology would typically require conventional machining to ensure the dimensional tolerances are achieved.

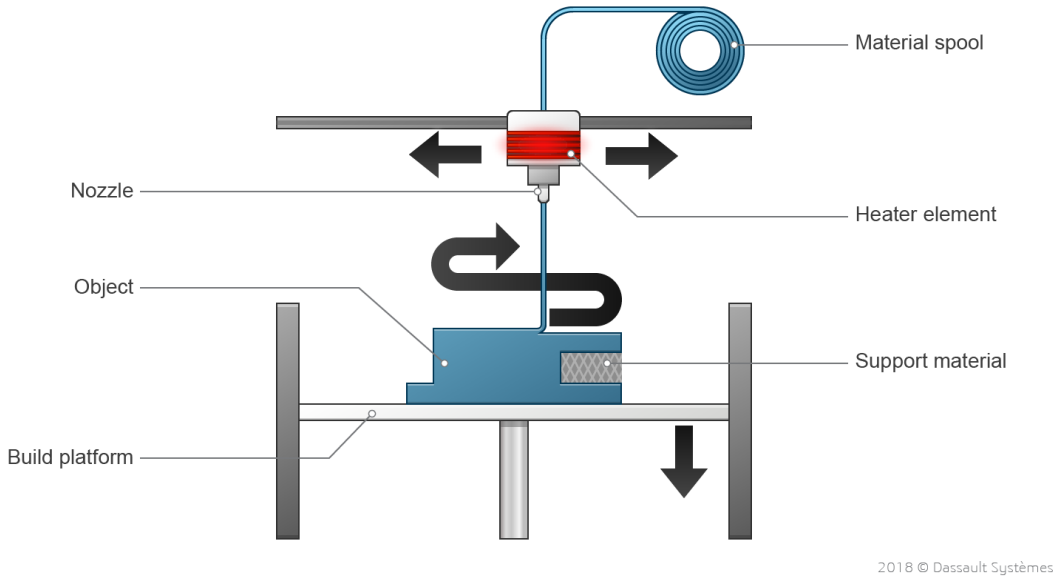


Figure 2–9: A Labelled figure of the typical Material Extrusion machine. (3D Experience Marketplace, 2019)

Capabilities

The smallest common nozzle diameter is 0.15 mm and the smallest layer height is 0.05 mm. Features under 0.5 mm usually require some optimisation before they will print reliably. As with sintering, long aspect ratio components will typically be more susceptible to warping, although this has become much less of an issue with more recent printing bed developments, material selection and designer awareness.

Features such as threads can be made using this technology, although they must be oriented with the spiral-axis of the threads aligned with the z-axis of the printer. The minimum thread that can be reliably printed is M6 x 1, although this will require careful tuning of the machine settings if the hole cannot be hand-tapped after printing. Overhangs of 45 degrees from normal, are typically manageable, and can be up to 60 degrees if the surface quality of the overhang is unimportant.

Printing of enclosed volumes is normal, although thought must be given to internal support material if the volume is required to be hollow. Wall

thicknesses for thin walls must be an integer of the nozzle diameter. For a 0.4 mm nozzle, the wall thicknesses could be 0.4 x N, up to a thickness of ~4 mm. With wall thicknesses greater than 4 mm, an infill is generated so the wall thickness is now continuous.

Vertical pillars should be greater than 3 mm in diameter for any pillar taller than 5 mm. Thinner pillars may work where there is adequate part cooling but are prone to being disturbed by the nozzle as the pillar height increases.

2.5.3. Vat Polymerisation

Vat polymerisation is a family of printing technologies that use a UV-curable liquid resin and a selective exposure process to create components layer-by-layer. It is the most accurate, and has the highest detail, of the 3 mentioned technologies. The available materials for this technology are typically hard UV-curable plastics, although the range of resins is expanding to include high toughness and high strength varieties. A schematic of a vat polymerisation machine is shown in Figure 2–10.

Printing time is generally longer when compared to FDM printing. The component must be cleaned and cured fully in a UV oven to allow the model to reach full strength. Therefore, vat polymerisation is typically better suited to small component prototyping or one-off production of intricate components.

In general, print time is related to the height of a component and the layer thickness; it is generally less affected by the overall volume than for FDM printing.

The resins that are available at the time of writing are not able to compete in terms of strength or stiffness when compared to the plastics used in FFF and the metals used in sintering. This limitation can be somewhat mitigated by a shift in the design to use more complex structures, such as lattices and internal geometry, to stiffen the geometry of a component. (Hussein, Hao, Yan, Everson, & Young, 2013)

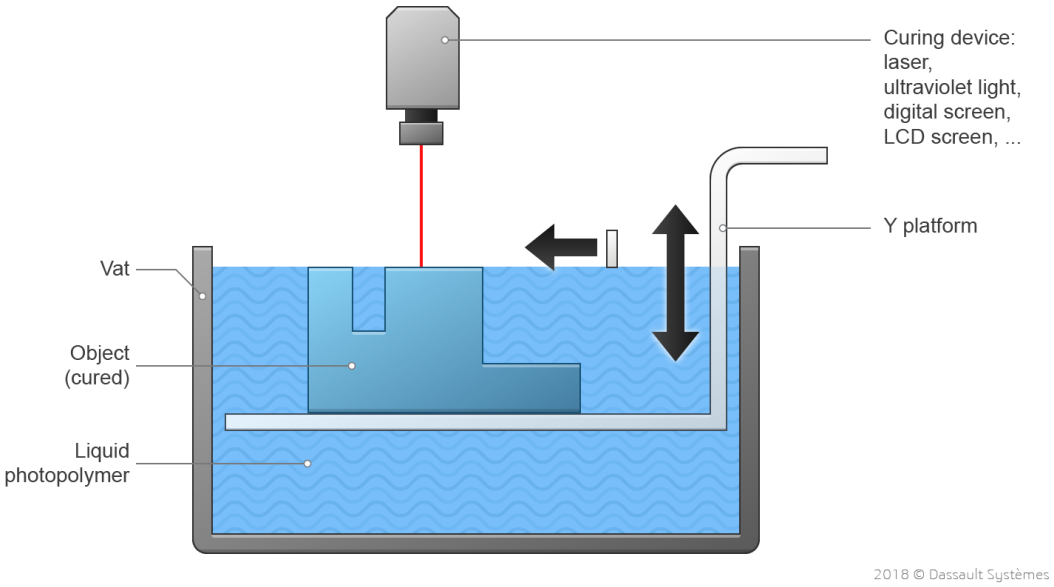


Figure 2–10: A schematic of a Vat Polymerisation Machine. (3D Experience Marketplace, 2019)

Capabilities

The resolution of this technology is typically limited by the minimum spot size, which can be as small as 0.03 mm for a laser diode-based system. The smallest practical layer height for resin printing is 0.05 mm. These process parameters result in a minimum feature size in the 0.1 mm range. Components can have internal features and hollow volumes, although care must be taken to ensure that trapped uncured resin can be drained. The minimum layer thickness is ~0.02 mm depending on the component geometry and specific printing technology.

The increased resolution allows for smaller printed threads, and for threads and similar features to be formed off-axis. Support structures for resin printers can be smaller in diameter (1 mm is typical), as the lateral forces on the cured components are minimal.

The surface finish for a resin part will generally be better than for the same part printed using FFF. The improved surface finish of resin printing is attributed to the very low inter-layer force. FFF achieves good layer adhesion

by ‘squishing’ the plastic to the preceding layer; creating a component with better structural properties but a worse surface finish.

The minimum wall thickness will be related to the spot size for scanning laser-based printers and related to the pixel density for projector/LCD based printers. Practically this means that scanning laser printing technology can print walls as thin as 50 µm and up, whereas projector and LCD based printers are limited to ~100 µm minimum wall thickness. Projector and LCD printers also have quantised exposure in the X and Y axes, although this is not normally noticeable and does not reduce part strength or surface finish quality.

Curing of thick-walled components can take considerable time and may result in the over-curing of thin sections on the same component. This can be mitigated by using internal lattice structures to retain the strength without the high resin density.

Summary of Additive Processes

The following table (Table 2-3) compares the previously mentioned additive processes including comparable numbers to help quantify the resolution of each process, appropriate materials, notes to summarise the relative strengths and weaknesses, and the cost per part.

Process	Min. Feature Size (mm)	Layer Height (mm)	Material	Cost of part	Notes
FDM	1.0	0.05	ABS, Nylon, PLA, acrylic, PEK.	Low	Good for rapid prototyping, jig manufacture.
DLP	0.2	0.025	Resin (range of hardness)	Med	High resolution, high detail, slow. Price is coming down.
SLS	1.0	0.07	Wide range of metals and plastics, casting sand.	Med	Suitable for production, layer artefacts are minimal. Good for batch work.
Material Jetting	0.05	0.02-0.04	UV cure resin, coloured resin, rubber-like material	Med-High	Highest accuracy, lowest warping, multi-colour, multi-material. Can make moulds for production.
Binder Jetting	1.0	0.1-0.4	Huge range of metals and plastics, sandstone, PMMA	Med	Full colour PMMA. Metals can be sintered or infiltrated. Large build volume, suitable for production.

Table 2-3: A table summarising the previously mentioned range of additive processes, along with the comparable properties of each process.

2.6. Ceramic Injection Moulding

Ceramic injection moulding (CIM) is a production process that is similar to injection moulding of plastics. It has several distinguishing features which make it much more appropriate for use in metrology stylus metrology manufacture, namely the excellent dimensional tolerance and surface finish of components. It can also be used to produce extremely small components, such as the gears inside a planetary gear system (Kryachek, 2004) (Stanimirovi & Stanimirovi, 2012). Some examples of the components that can be manufactured using CIM are shown in Figure 2–11.



Figure 2–11: A picture showing a selection of components produced using Ceramic Injection Moulding. Note the matchstick for scale. Image courtesy of (Ceramit., 2018)

In plastic injection moulding, the plastic is injected at an elevated temperature. This results in shrinkage and deformation in the final component, as well as a limit to the dimensional tolerance that can be achieved. Ceramic injection moulding is usually done at, or slightly above room temperature, which reduces the thermal deformation and improves the dimensional consistency of the component, by comparison to PIM.

The ‘green’ ceramic must be de-bindered before it is usable. This takes place in an oven, shortly after demoulding. This process shrinks the component, in turn, increasing the density. This shrinkage is uniform and well understood,

although it would need to be considered during the mould design process. The geometric deformation of the component pre to post cure is minimal, predictable and highly repeatable between components.

Dimensional tolerances of ~1 µm and surface finishes of 0.1 µm are possible with CIM. This fact, along with the high production rates which can be achieved, make it ideal for production of small, complex components such as gears. The CIM process can incorporate features such as threads or other locating features, which require no post machining to be functional. The ceramics typically have very high wear resistance, potentially increasing the lifetime of a product. (Precipart, 2018)

2.6.1. Pitfalls of CIM

The increased density of ceramics mean that the fundamental geometry of a component must be changed to produce components with the same mass and/or inertia. This may be straightforward, such as making a section hollow, or it may require a total redesign of the component.

The ease of design iteration is somewhat lost when using injection moulding technologies. Updating the geometry of a component requires a new tool to be made which may cost £10k or more. (Rex Plastics, 2019) The expected tool life of a mould may be as high as 10,000 units; making this a cost-effective proposition.

2.7. Design for Manufacture and Assembly

Design for manufacture and assembly is a design approach that ensures that any designed component is producible given the tools and technologies which are available. By analysing the component with a focus on its suitability for manufacture, extraneous features can be removed, or tolerances relaxed as necessary. (Stienstra, 2005).

Although design for manufacture and design for assembly can be treated separately, they are typically undertaken at the same time. These techniques can be characterised by a number of 'rules' (taken from (Chang, Wysk, & Wang, 2005)), several of which will be highlighted here:

1. Reduce the total number of parts.
2. Develop a modular design
3. Use of standard components
4. Design parts to be multi-functional
5. Design parts for multi-use

Reducing the total number of components in an assembly may reduce assembly time and increase product reliability, if done correctly. (Bayoumi, 2000). This is complimented by the use of standard components; a minor design change may be possible that would, for example, reduce the range of required fasteners in an assembly. This in turn reduces the number of opportunities for assembly error.

A modular design approach allows the design to be modified or extended to suit a different application. This may require the development of a standard coupling or interface that may then be used in a wide range of other designs. A modular design allows for easier maintenance and repair of a product in the field.

A custom component may be multi-functional if it incorporates two or more features such as structural support, electrical conductivity or heatsinking. This approach to component design will typically result in a reduction in parts count

and overall complexity, although care must be taken to understand the impact of these design decisions on the finished product.

A component may be multi-use when one component can be used across a range of products, by the inclusion of redundant features. For example, a component may have an extra set of mounting features so that it may be used in two separate products. In this example, the redundant feature is whichever mounting feature that is not used in that product. This approach generally reduces the number of custom components that must be manufactures.

2.7.1. Design for Machineability

Design for machineability is a conceptually similar to design for manufacture. Design for machineability is a specific set of guidelines that allow components to be evaluated for their machineability at design time. The guidelines for manual machining are different from those of CNC machining and different again from 5-axis machining. The capabilities of the available machine tools and operators must be considered at design time. (Khayal, 2019).

3D Hubs lists a set of design rules that should be followed to ensure that a design in readily machinable on a CNC mill (3D Hubs, 2020). The topics include:

1. Cavities & Pockets
2. Internal Edges
3. Thin walls
4. Holes
5. Threads
6. Small Features
7. Machine setups & Part orientation
8. 5-axis CNC machining

Cavities should be no deeper than 4 x the cavity width. This limits the length to diameter ratio of the cutting tool used to form the cavity. Internal vertical

edges cannot be sharp, instead, a fillet with the radius of the largest tool that can be used to machine the vertical edge should be used.

Thin walls and cross sections will quickly lose rigidity when machining, leading to chatter or a reduction in material removal rate, and therefore increased cost. Holes should be a standard diameter where possible (5 mm diameter drills are common; 5.37 mm drills are not). The recommended maximum hole depth is 4 x the drill diameter, although the technology exists to drill much deeper than this (with added expense).

As with holes, internal threads should be less than 4 x diameter deep, with 2 x diameter being preferable. Larger threads are preferred as they can be formed with larger stiffer cutting tools. Care must be taken to ensure the of a pre-drill hole is deeper than the thread length, to allow for chip evacuation.

Small features are readily achievable on modern CNC machines, although the smaller an internal feature, the smaller the cutting tool and therefore the longer the machining process will take.

Multiple part setups should be avoided where possible. There is a repeatability associated with each reorientation of a component so tight tolerances may be difficult to achieve over a high number of operations. Parts that cannot be machined in 1 or 2 operations may be better suited to production on a 5-axis milling machine.

5-axis CNC machines allow 5 sides of a component to be accessed in 1 setup. The associated setup repeatability is reduced to the kinematic error of the machine. Features may be readily machinable using a 5-axis CNC machine that would require special tooling on a 3-axis machine (chamfers or angled features are a good example.

2.8. Design for Additive Manufacture

Additive manufacturing (AM) has become increasingly accessible in recent years, allowing faster iterations on the design of components and assemblies, as well as the manufacture of near-net shape components in a single operation. Additive manufacture can significantly reduce the component count, especially on highly complex components. The penalty for complexity is much lower than for traditional CNC machining, even allowing the production of components that could not be manufactured any other way. The final component is not limited to the billet size, and the typical ratio of purchased material to utilised material is very high. Sensors and other internal components can be embedded in the component, without compromising the strength of the produced part.

These benefits do require a different paradigm in terms of design and approach to prototyping. CNC machined components can have a minimum feature size on the order of 0.1 mm, whereas most additive technologies require a feature to be at least 0.5 mm. The strength of the produced component is also dependent on the orientation in which it was manufactured as the inter-layer bond is typically weaker than the bonding between material in the same layer. Complex features can be produced using additive methods, although the orientation of the layers will have an effect on the produced feature.

2.8.1. Design Philosophy Differences

To properly attain the benefits of additive manufacturing, almost all of the conventional design philosophy must change. Traditionally un-machinable components are no longer limited by the tool access or the component geometry. Components can be embedded inside printed parts at print-time, meaning that integral brackets and support structures can be used instead of a dedicated separate component. Internal channels and ducts can be printed into a component, allowing airflow or cable management that is simply not possible with other manufacturing techniques.

This section details some of the key benefits of additive manufacturing and gives some guidance as to how to take advantage of them. The following concepts will be discussed in detail:

- Simplification of Design
- Complexity-cost trade-off
- Topology Optimisation and Weight Reduction

A detailed comparison between traditional manufacturing and additive manufacturing can be found here: (Pereira, Kennedy, & Potgieter, 2019)

2.8.2. Simplification of Design

During the conventional design of an assembly, the number of components is governed by several factors. A monolithic design is often preferable, but a component must be split when it becomes impossible to machine from a single billet. At the point that a component must be split, fasteners must be introduced, as well as other interfacing components such as washers or O-rings. Both components must now be given tighter tolerances such that they will fit properly and hold an external tolerance between them.

Components must also be split when some internal features are not accessible. Cavities and internal channels may also require access to machine, therefore requiring a separate component as a cover. The mounting considerations for a part may also necessitate a separate component or machining operation.

Conventional machining typically results in components which are prismatic or cylindrical, although this does not necessarily reflect the ‘natural’ shape of the component to be made. Ducts and tubing are conventionally formed in straight sections by drilling. The open end of the drilled hole must be plugged in a separate operation. This sharp corner will also drastically increase resistance to fluid flow, potentially requiring a larger duct and a larger component overall.

The Additive Way

The formation of internal geometry for an additive component means that the conventional limitations (accessibility in machining) are mostly irrelevant. A duct or channel can be formed in almost any shape (no longer limited to a series of straight sections), without requiring hole-plugs. This may avoid the need for the component to be split, thereby reducing the overall assembly complexity and the bill-of-materials (BOM).

The shape of the ducting is also much less constrained, with mounting or locating features that can be formed additively, therefore reducing the assembly time and complexity. The mounting and interfacing features can be machined to achieve the dimensional tolerances required.

2.8.3. Complexity-cost trade-off

A conventionally machined component has a cost associated with volume of material to be removed. As more material is removed from the billet, the work holding becomes more complex and the number of setups increase, thus the cost increases. In an additively manufactured component, the cost is proportional to the volume of the final component and the absolute number of layers required. A component with complex features, holes, undercuts and internal features, weighing 20 g would cost a similar amount to a simple 20 g solid block.

An example of the benefits of the increased manufacturing complexity offered by additive manufacturing would be an internal nozzle inside a pump housing. The nozzle would typically be a separately manufactured component, assembled after manufacture and requiring fixings or fasteners for retention. Using AM to produce the pump housing would allow the nozzle to be integrated directly at manufacture-time. A further benefit of the additive paradigm is that the nozzle geometry does not need to be conventional. A nozzle with a wide, flat outlet port would be as additively manufacturable as one with multiple outlet holes or another difficult-to-conventionally-machine geometry.

Bosses, counterbores, chamfers, fillets, O-ring grooves, model numbers and warning labels can all be produced in a single operation, often without requiring post-machining. The increased manufacturing capability extends to allow the creation of features that are not conventionally machinable, such as sharp internal corners, closed cavities, integrated cooling channels, captive components within components and thin walls. (Attaran, 2017)

The costs of the increased geometric freedom afforded by AM are varied. Surface finishes on AM components are typically much rougher than their subtractively manufactured counterparts and the dimensional tolerances are typically worse for AM components. These limitations may require secondary operations such as machining or coating in order to mitigate the impact of the manufacturing process on the performance of the component. The additive process will introduce design constraints such as the minimum wall thickness.

The drive to integrate as many components and features into a single additively manufactured component must be carefully considered. Although it is possible to manufacture a highly complex component using AM, the failure rate of the manufacturing method is likely to be related to the component complexity, aspect ratio and minimum feature size, meaning the assembly should be considered as a whole before optimising away all individual components. Further, if a component that typically requires frequent replacement during the lifetime of the machine is combined with other less-frequently replaced components, then the maintenance costs will be increased, even if the initial production costs are reduced.

2.8.4. Topology Optimisation and Weight Reduction

Topology Optimisation is the removal of the unnecessary material within a component, reducing the overall mass without significantly changing the stiffness. This ties in closely with the weight reduction benefits that come from the additive paradigm.

Weight Reduction

Conventional components are machined from billet, thus conventional design is typically concerned with the creation of features from the billet. This results in components with an optimised machining time and not a minimised mass. Additive components should be designed to minimise the mass (within sensible limits) which will in-turn reduce the cost of producing the component.

This approach naturally reduces the overall mass of the component but does require human intervention. It is often simpler to take a billet in cad and build-in the features as needed. Topology optimisation was created to aide with the transformation between a conventional design and an additive design.

Topology Optimisation

Topology optimisation is a tool which analyses the structure of a component and suggests volumes that can be removed, according to a set of loading criteria and a weight reduction target. First a component is created using conventional design practices. Next, features on the model are marked as important to preserve them during the optimisation process. Loading criteria are added to features on the model, the stresses generated by these loads are used to decide which bits of material can be removed. A mass reduction target is specified (say 50%).

Once these details have been entered, the shape is split into small finite elements to allow the shape to be deformed in similar way to conventional finite element modelling. The stresses of each volume element are calculated, with the lowest stress areas being removed. This process repeats, using the generated mesh from the iteration before, until the mass reduction target is reached, or the component would fail with the specified loads.

Topology optimisation can be repeated with several loading cases to give the designer an understanding of the relationship between load and important geometry. This technique can also be used to compare the effects of different design decisions on the final part (an M10 vs 2 M6 bolts, for instance).

Once the process is completed, the mesh is presented to the user. This mesh is typically considered a helpful guide as it may end up with voids, unconnected edges or difficult to model shapes. These systems do not typically account for the manufacturing process and as such may need to be checked for compliance (minimum feature size etc.). The 'organic' shapes that are generated by these techniques are often unintuitive to designers from a conventional background.

The topology optimisation method demonstrates good practice for additive design; important features are preserved, and all other connecting elements can be carefully optimised to avoid over-specification and therefore higher costs. The relationships between the components and those connected to it in the assembly, are used to drive the design choices, typically resulting in lighter and stiffer components which can be iterated more quickly. Indeed, additively manufacturing several versions of a component at a time is common.

An Example of Topology Optimisation with Specific focus on Design for AM

MIDAS nfx (Midas NFX, 2014) have conducted several case studies on topology optimisation, specifically focussing on DfAM. They looked at a typical mounting bracket, as shown in Figure 2–12. This approach helps to identify the volume of material which is critical to the performance of the component.

Several regions of this component are removed or reduced, resulting in a component with a 40% lower mass. The surfaces generated by the optimisation routine are typically not uniform and may not be aesthetically pleasing. To mitigate these problems, the optimised geometry may be used as a reference during subsequent design iterations. The topology optimisation process may be repeated after each design iteration, although the overall weight reduction is unlikely to improve significantly after the first run.

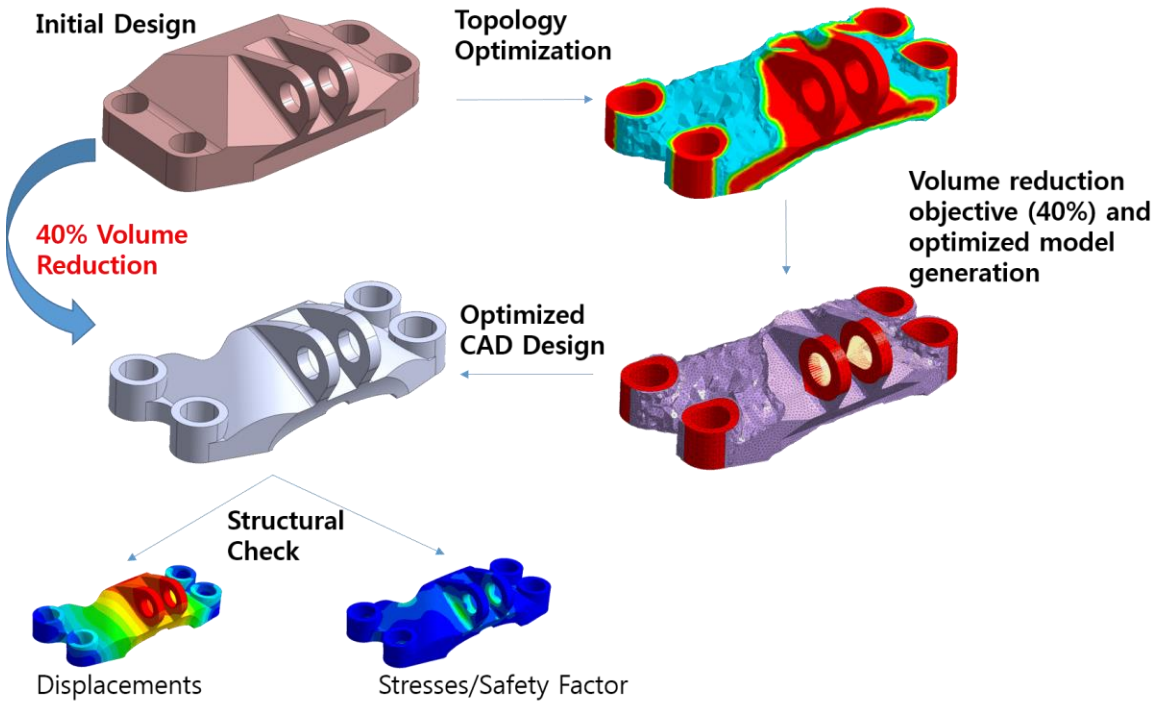


Figure 2–12: A typical topology optimisation flow chart, showing the CAD model at each step including the post-optimisation stress and displacement checking.

Note that some features are preserved in the model: the hole locations and the strut locations are necessary for interfacing with other components, so these must be preserved throughout the optimisation process.

2.9. Stylus Dynamics Theory

When measuring a component to the accuracies that are necessary for bearing or lens production, the entire measurement loop must be considered. The physical portion of the measurement loop for a typical surface finish instrument is shown in Figure 1–1. Every component in the loop must be more than capable of the level of static and dynamic performance required by the measurement that is being undertaken. The dimensional stability and excellent damping characteristics of granite and cast iron make them good materials for the construction of metrology systems (Slocum, 1994). The damping factor is related to the Q-factor of a material by the following relation:

$$\zeta = \frac{1}{2Q}$$

Where ζ is the damping factor and Q is the Q-factor. The damping factor of synthetic granite composite, Aluminium and Grey cast iron are shown in Table 2-4.

Material	ζ	Std. Deviation
Synthetic granite	0.0181	0.0019
Aluminium (6061)	0.0033	0.0005
Grey Cast Iron	0.0029	0.0004

Table 2-4: A Table of material damping properties for a synthetic epoxy-granite, 6061 Aluminium and Grey cast Iron. (Lovo, Pedroso, Erbereli, Purquerio, & Fortulan, 2018)

Active and passive anti-vibration systems are used to minimise the effect of ground-borne vibration on a measurement system. Environmental cabinets are also used to reduce the effect of air-borne vibration on the stylus, and to increase the temperature stability around the metrology system.

The measurement of a surface using a stylus instrument can be broken into several distinct stages: preparation, physical measurement, data collection and analysis. The dynamic performance will directly affect the physical measurement stage, requiring filtering and correction at the analysis stage.

During physical measurement, the drive unit (traverse unit) is lowered to bring the stylus into contact with the surface. The traverse mechanism moves the gauge and stylus in a well-defined and calibrated path, at a well-defined speed. The displacement of the stylus will be recorded, along with the horizontal position of the traverse unit.

From the perspective of the stylus, the forces it experiences are directly related to the surface which is being measured and the speed at which the measurement is being performed. Given that the measurement is performed at a constant traverse speed, the frequencies which the stylus experiences will be directly related to the spacing and amplitude of surface features. This assumes a constant surface gradient. The particular dynamics of the stylus-gauge system will influence the measured profile, with the influence being larger in frequency regions that are above the stylus flight frequency or around the stylus resonance frequency.

To avoid the measured profile from being affected by the stylus dynamics, the resonant frequency of the stylus must be above the pre-filter frequency. The practical result of this requirement is the careful selection of the measurement traverse speed to ensure that a particular surface can be measured reliably.

2.9.1. Stylus Flight

The stylus contact force is maintained by a spring under tension. The ISO specification (ISO 3274, 1996) for this force is: ”The nominal value of the static measuring force at [the] mean position of the stylus is 0,00075 N”. Surface instruments at Taylor Hobson are set to apply 800 µN of tip force. This is done to ensure that the stylus forces overcome the friction of the pivots and remain within specification over the entire gauge range. This force is kept small to minimise the surface pressure of a 2 µm radius stylus, thus avoiding scratching the surface under test.

A downside to this low force is the tendency for the stylus to leave the surface, typically called stylus flight. Stylus flight occurs in 2 scenarios: when a stylus ‘falls off’ a ledge (A in Figure 2–15) and when a stylus is moved upwards by

the surface and overshoots the corresponding downward slope (Figure 2–14 and B in Figure 2–15). These can be considered as separate phenomena as the root cause of each is independent. Overhangs and sharp edges are typically created intentionally as part of a component design, and thus should be avoided during measurement. In contrast, peaks and valleys are created at machining time, albeit at length scales much smaller than deliberate machined features. In order to mitigate the stylus flight as a result of machining marks, the measurement speed must be reduced.

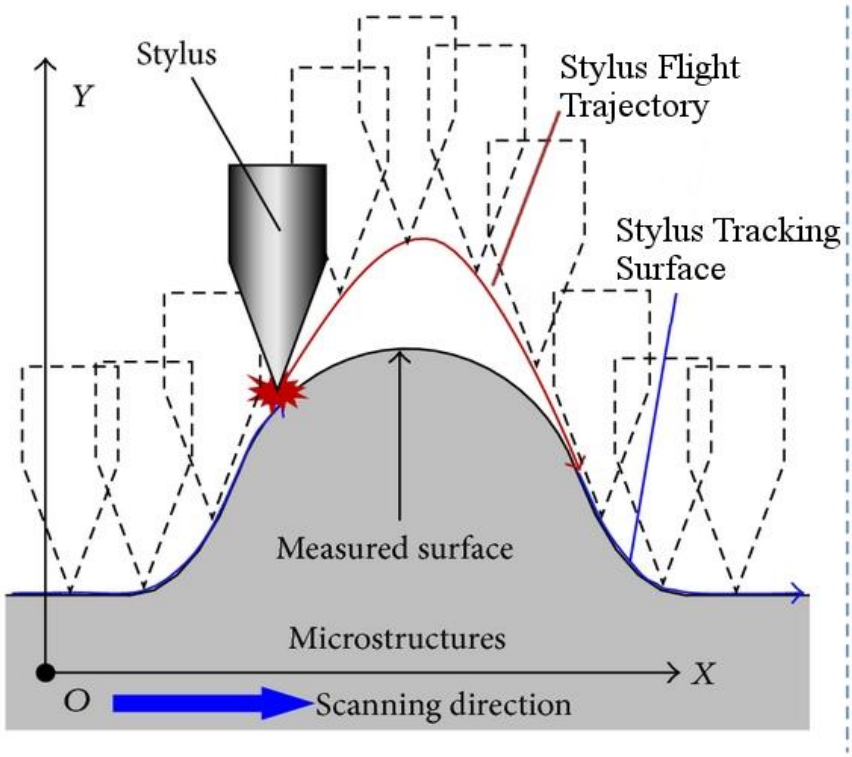


Figure 2–14: The motion of a stylus tip during stylus flight.

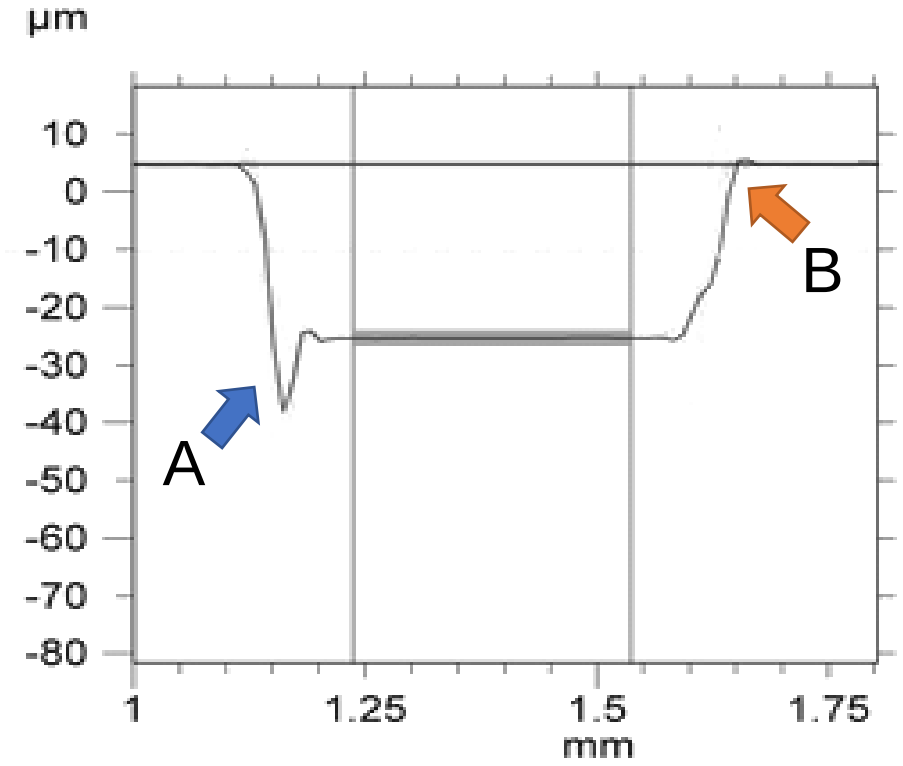


Figure 2–15: A profile of a step height measurement. (A): The downward transition and corresponding artefact. (B): The overshoot of the stylus after moving up a steep feature.

2.9.2. Step Height Changes

Drastic changes in the surface are naturally difficult to measure using conventional profilometry. The ability to 'catch' a falling stylus has been demonstrated but this should not be used during a measurement that must continue after falling from the step. If step changes are to be properly measured, a suitably slow measurement speed must be chosen, and the step-change artefacts must be excluded from the analysed data (Shown as A in Figure 2–15).

This step-change artefact is caused by the impact of the stylus tip after falling into the hole. The stylus tip contacts the bottom of the step and then the stylus undergoes bending. Once the stylus has absorbed the energy imparted by the fall, it releases it and the stylus tip will once again leave the surface, albeit to a much lower amplitude as energy is lost to the system. The frequency of this motion is related to the natural frequency of the stylus and the damping properties of the stylus-gauge system.

The path of overshoot of the stylus tip is characterised more by the energy in the stylus system at the point the stylus leaves the surface, and the dynamics of the stylus-gauge system as a whole, than the resonant frequency of the stylus itself. The magnitude of both phenomena can be reduced by measuring the surface with a slower traverse speed.

Steep sided components can be properly measured, at least in one direction, by tilting the component to eliminate the overhangs and avoid flanking (shown in Figure 2–16). The tilting of a component to allow favourable measurement conditions is normal and necessary when measuring the form and finish of steep-sided optics or the surface finish on gear teeth.

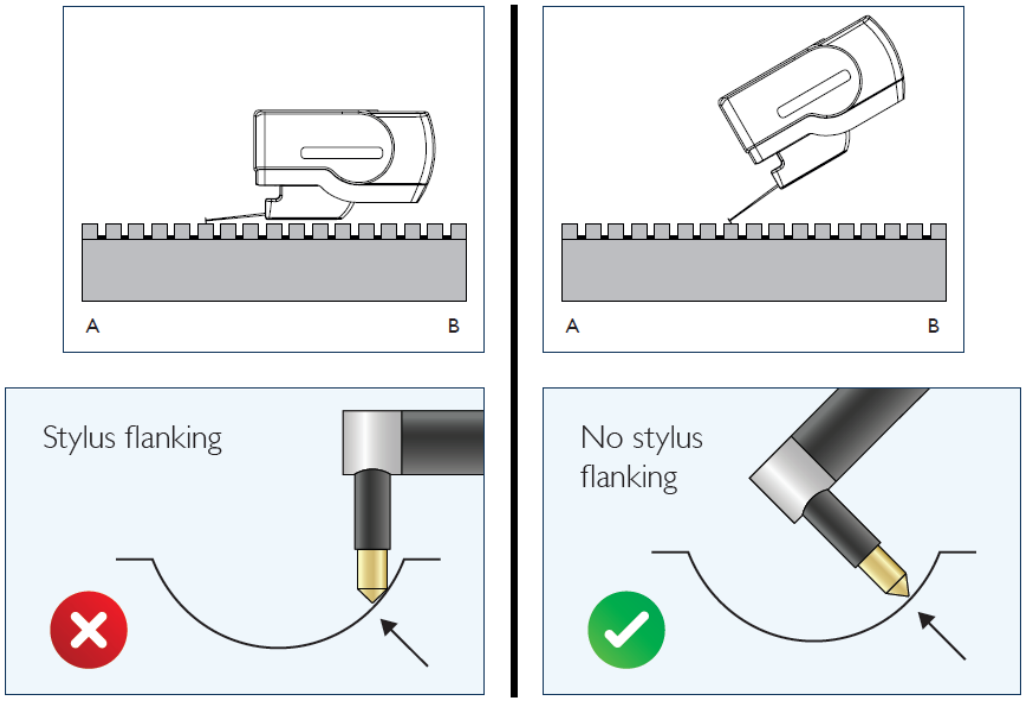


Figure 2–16: A schematic representation showing the effect that rotating a component (or in this instance, the traverse unit) can have on a measurement. (Taylor Hobson Ltd., 2019)

2.9.3. Surface Reaction Effects

Most surface finish measurements will be affected by surface reaction effects in some way. Measuring turned surfaces, perpendicular to the machining

direction, will give a highly regular machining pattern (Figure 2–18), related to the cutting tool geometry (tip radius) and the feed-per-revolution of the cutter. These sharp-peaks and wide valleys are very prone to inducing stylus flight, which will effectively broaden the peaks and change the shape of the measured valleys.

2.9.4. Surface Frequencies

Two surfaces can be used to demonstrate the concept of surface frequency components; one produced by turning and one perfect sinusoid with the same peak-to-peak length as the turned surface. Representations of these surfaces can be seen in Figure 2–17. The sinusoidal surface can be decomposed into only a single wavelength with a corresponding amplitude. The turned surface is decomposed into a large number of wavelengths and amplitudes as a result of the sharp transition between machining grooves.

The gauge will experience a surface which is close to these exact profiles with only a few modifications for the finite tip radius and the stiffness of the stylus-gauge system.

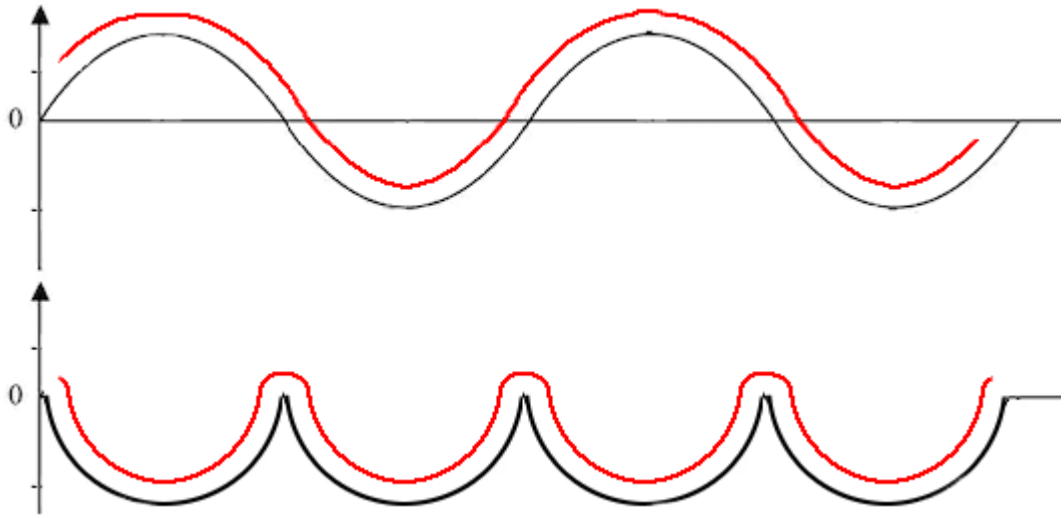


Figure 2–17: Simplified examples of sinusoidal (Top) and lathe-turned (Bottom) surfaces. The path of a locus of a stylus tip is shown in red for each surface, although the scale of the offset between the surface and the locus is exaggerated.

Since most surfaces are not perfectly sinusoidal, the frequency decomposition will be result in a large set of surface frequencies. For a given amplitude, the

dynamic force from a surface on a stylus is related only to the frequency. As a result, surfaces with many sharp discontinuities and other high-frequency features will cause stylus flight more readily than for lower frequency surfaces. This also means that low-amplitude surfaces are more susceptible to stylus flight at lower frequencies, resulting in difficulties in measuring super-fine surfaces at 'regular surface finish' speeds.

The finish of a turned surface is broadly related to the cutting tip geometry, the rotational speed of the workpiece and the federate of the tool across the workpiece. A theoretically perfect turned surface has the form shown at the top right of Figure 2–18. The lathe tool will wear over time and so the measured surface finish will change as a function of the cutter wear. The measured surface as a function of cutter wear is shown on the right-hand side of Figure 2–18.

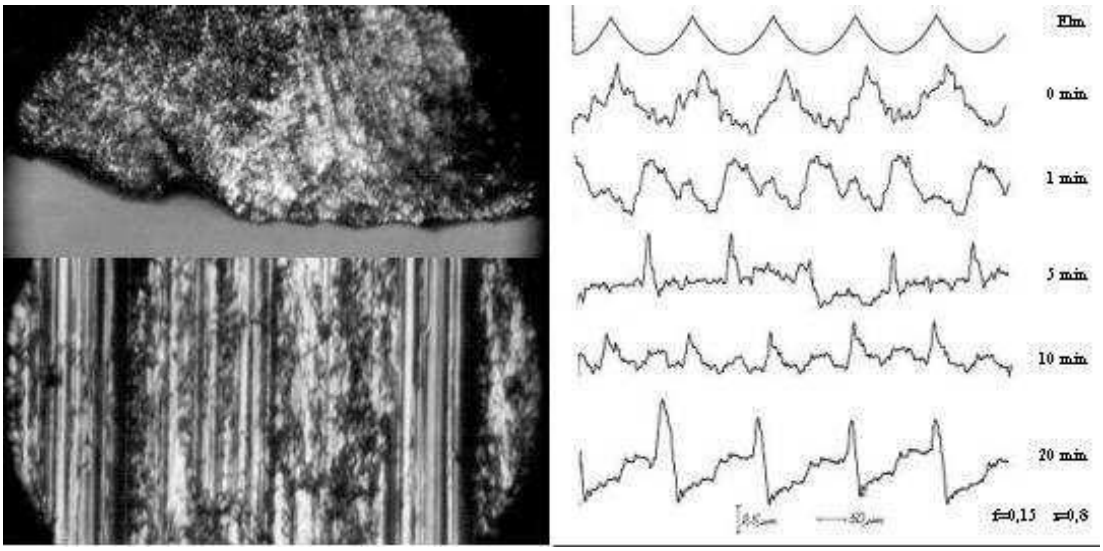


Figure 2–18: (Left-Bottom) A microscope image of a turned surface, with characteristic turning marks. (Left-Top) A microscope image of the cutting tool after use. (Right) Measurements of the surface finish of the workpiece as a function of time (and therefore wear of the cutting tool). (Palásti-Kovács, Sipos, & Czifra, 2012)

Note on Stylus Tip Radius Correction

If the wavelength of the sinusoidal surface is long when compared to the tip radius, then the effect of the tip correction on the surface will be minimal and will only result in a small 'spreading' of the power from a single wavelength to a small distribution of wavelengths with a peak at the original wavelength. Tip

correction will have the opposite effect on the higher frequency components of a surface, typical of sharp edges; these frequency components will be moved to lower frequencies as noise.

2.10. Stylus Dynamics in the Literature

2.10.1. Error in Measurement due to Stylus Kinematics

As far back as 1973, academics were making efforts to understand the errors in measurement caused by stylus dynamics. (Damir, 1973). The stylus was simulated as a mass on a spring, that would only be deflected in the vertical axis. The force of the stylus tip (ma) is equal to the reaction force, minus the spring force.

$$ma = R_z - K\{\delta + z\}$$
$$where: a = \ddot{z}$$

Stylus flight occurs when the reaction force $R_z = 0$. The acceleration term (a) is related to the local rate of change of the angle of the surface, as well as the traverse speed of the stylus. This problem is also not symmetrical; a stylus will be driven into an upward sloping surface, even if it would not track that same surface downward.

This model allows the prediction of the point at which the stylus leaves the surface but does not account for the inertia or damping of the stylus and gauge system, or the stiffness and dynamic characteristics of the stylus.

2.10.2. Stylus Flight in Surface Profiling

By 1994, the modelling techniques had become more sophisticated. (Song & Vorburger, 1996). The analytical model includes damping and spring force terms, although the model still lacks the terms for moment of inertia of the stylus-gauge system. The stylus is still modelled as moving only vertically, as opposed to in an arc, as in most modern surface finish gauges. This disparity will be more pronounced for form measurement as the range is typically larger than for surface finish measurement.

This paper goes on to highlight that stylus damping actually allows the stylus to track the surface, even when conditions would result in stylus flight for an undamped system. They also highlight a method of calculating the maximum stylus traverse speed for a given surface. It also gives an estimate of the

frequency response of the system and suggests that some of the observed disparities are caused by this limitation.

2.10.3. Enhancement of Instrument and Machine Capabilities

In 1996, DJ Whitehouse released a paper (Whitehouse, Enhancement of instrument and machine capabilities, 1996) that makes efforts to express the dynamics of a modern stylus instrument. It uses a spring-compensated, pivoting beam system, considering both damping and stylus stiffness. The model used in this paper is as follows:

$$-R(t) = M\ddot{z}(t) + T\dot{z}(t) + kz(t) + F$$

Where R is the vertical reaction force, M is the effective mass, T is the damping, k is the stiffness and F is the spring force.

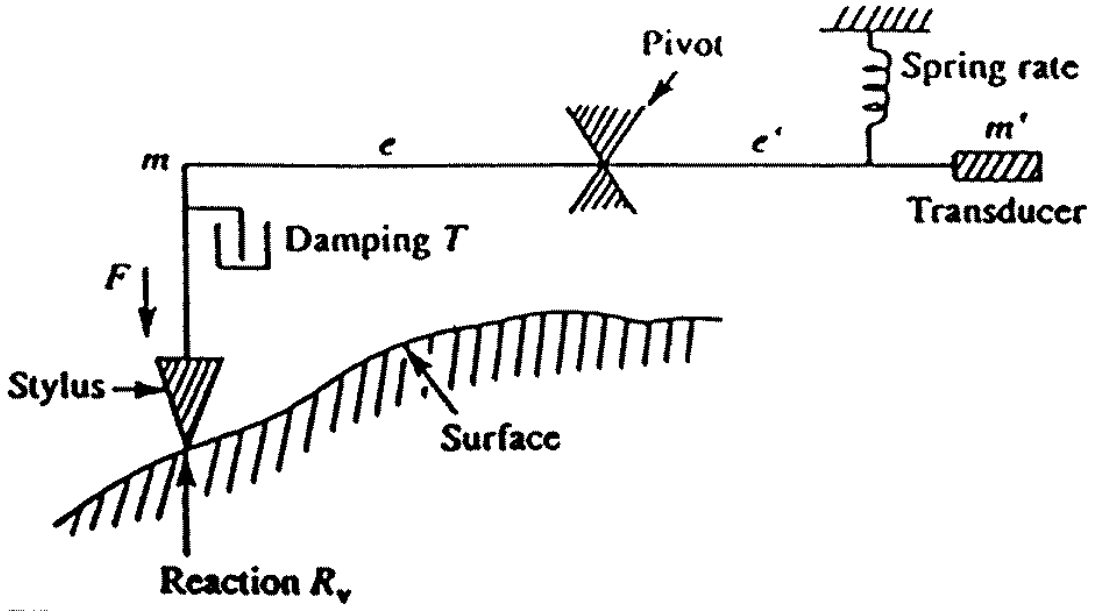


Figure 2–19: A schematic of the stylus-gauge system and the corresponding forces acting on that system. Taken from (Whitehouse, Enhancement of instrument and machine capabilities, 1996).

The complexity of this model means it is not generally analytically solvable. It is solvable in specific circumstances, for a given surface. The parameters of the function are then re-expressed in terms of average wavelength of the surface: (ω_a) and natural frequency of the stylus-gauge system (ω_n). The optimum damping factor is then calculated for a range of surfaces (single sine waveform through to totally random), with a frequency cut-off set to (ω_n). An

optimal damping factor is found to be $\xi = 0.59$, which is appropriate for all surfaces.

The paper concludes with an experiment looking at active damping, using a modified Talysurf™ 5 gauge. This demonstrated that active damping is achievable and practical for use when measuring surface finish.

2.10.4. Vibration analysis of stylus instrument for random surface measurement

This recent paper (Tian, Liu, Chetwynd, Shirinzadeh, & Zhang, 2010) implements a new dynamic model of the stylus system. It suggests methods for simulating random surfaces that are suitable for numerical analysis. In this case a surface of 140 points is generated and then a spline is fitted. This method increases the point density and ensures the curve has a finite second derivative. The analysis of the stylus dynamics is done using a finite error Runge-Kutta method. This is an iterative method of solving differential equations which allows the underlying physical equation to be defined at each point. In this case the simulation uses two distinct equations of motion, one for when the stylus is in contact with the surface and one for when the stylus leaves the surface (stylus flight). This paper also details the important role of the tip radius in the dynamics of the system and the accuracy of the measurement.

The topics discussed in this paper provide a useful overview of the current state of stylus system simulation. This work could be used as the basis of a project in simulating surfaces. Ultimately a project to incorporate the non-linear stylus dynamics (such as flexion) will begin by first understanding the basic models in the literature.

2.10.5. Dynamic modelling of the fidelity of random surface measurement by the stylus method

This paper (Tian, Liu, Zhanga, & Chetwynd, 2008) is written by the same group as the previous paper (2.10.4) and so uses a similar selection of techniques to generate the surfaces under test. It does however use a different dynamic model for the stylus system. This model is simpler than the one implemented in (2.10.4) as it only uses a single spring and a single damping component. It does however more clearly detail each of the terms of the dynamics equation. From this dynamic model, the tip flight criteria is substituted and then the resulting stylus tip trajectory is derived analytically. This paper goes on to evaluate the effects of scanning speed on some normalised parameters Ra, Rq, Δq and Sk to show that the parameters will deviate if the scanning speed is above the critical scanning speed. The critical scanning speed is a function of the surface to be measured and the stylus system. The effect of increasing equivalent mass and tip radius (tip truncation length) on the critical scanning speed is also shown.

This paper helps to detail the efforts into investigating the dynamic models of the measurement system as a whole. Some things are more explicitly discussed such as the technical implementation of the methods of surface generation and the analytical solution to the dynamic model under tip flight.

2.10.6. Notes on the General State of the Literature

The investigations into sources of error in contact metrology are often limited and disjoint. As computer power increases, people are more willing to investigate physical models but comparisons between theoretical prediction to actual measurement could not be found. Further, as the majority of measurement instruments are produced in a commercial context, the likelihood of sharing the results of investigations is greatly reduced.

The lack of generality in the investigations that have happened limit their usefulness. This comes with a whole series of very reasonable causes but represents just how difficult it is to simulate even a basic metrology system. The problem of understanding instrumentation dynamics has been talked about for many years (Whitehouse, Handbook of Surface Metrology, 1994), and yet the practical effect of stylus dynamics has not been demonstrated with comparison to theoretical models.

This lack of specific focus, coupled with general commercial sensitivity, stops the topic of sources of measurement errors and the resultant effects from being discussed. The general lack of standardised method for estimating error contributions also increases the difficulty in comparing distinct systems.

2.11. Dynamics in Micro-CMM

The perpetual drive for miniaturisation of components and assemblies leads to an ever-increasing demand on micro and nano-scale machining and measurement. (Dornfeld, Mina, & Takeuchi, 2006)

Dimensional metrology must rise to this challenge, exploring novel measurement technologies and approaches to ensure that the next generation of measurement instruments are able to keep up with modern manufacturing techniques. (Thalmann, Meli, & Küng, 2016)

Dimensional measurement tools have also undergone significant miniaturisation, with the accurate measurement of micrometre scale features using tungsten probe tips as small as 42 µm in diameter. (Xtreme Precision Systems, 2020). Reliable measurement at this scale relies on a good understanding of the surface forces that are present in the interaction between the measurement surface and the component. The dynamics of the measurement stylus (and the measurement system as a whole) significantly influence the measured data and must be clearly understood. (Bos, 2011).

An example of a micro-CMM is shown in Figure 2–20. (Küng, Meli, & Thalmann, 2007)

2.11.1. Comparison of micro-CMM and surface Profilometry

Scanning mode micro-CMMs are conceptually most similar to surface profilometers, in both cases, the probe is kept in continuous contact while the measurement is taking place. A significant difference between profilometry and scanning CMMs is the typical orientation of the contact geometry with respect to the workpiece. CMM probes may make contact with the component in almost any direction, while surface profilometry typically requires the measurement tip to contact the component in the vertical axis.

The contact forces for the two techniques are also somewhat different. For fine surface finish measurement, a tip force of 750 µN is recommended (ISO

3274, 1996), whereas surface forces in scanning micro-CMMs may be as low as 109 μN (Fan, Cheng, Wang, Chen, & Lin, 2010).

Efforts have been made to calculate and measure the effects of the finite stiffness of the micro-CMM probe on measurement (Bos, 2011), however, the difference in geometric complexity and the scale of a surface finish measurement stylus and a micro-CMM stylus limit the transferability of these conclusions.

The dynamics of micro-CMM systems may be more relevant if surface profilometry gauges are miniaturised or significantly altered in the future. It is certainly feasible that surface measurement instruments could be improved using similar concepts to those in micro-CMM systems. Given the differences between the current measurement technologies, the dynamics of micro-CMM systems will not be investigated further in this report.

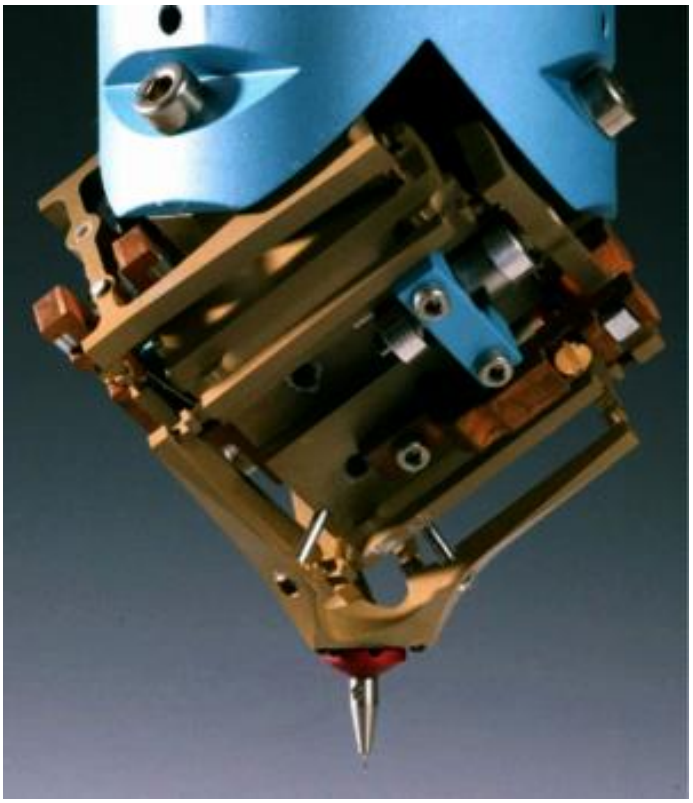


Figure 2–20: An ultra-precision micro-CMM using a low force touch probe

3. Stylus Tolerancing, Manufacture and Production

3.1. Stylus Tolerancing

This section presents research on the tolerances of common Taylor Hobson surface finish measurement styli, as expressed in their mechanical drawings, and highlights the sources of error in stylus production. The common stylus models are included (122-3221, 122-2009 etc.). A schematic of a common measurement stylus can be seen in Figure 3–1.

Data gathered from the stylus cell, surface cell and the roundness cell was used to identify the most numerous, and therefore important, styli. To accompany this classification is a note on the number of styli produced with a view to increasing the consistency of the higher volume styli. Finally, several future investigations have been suggested

3.1.1. Stylus Components

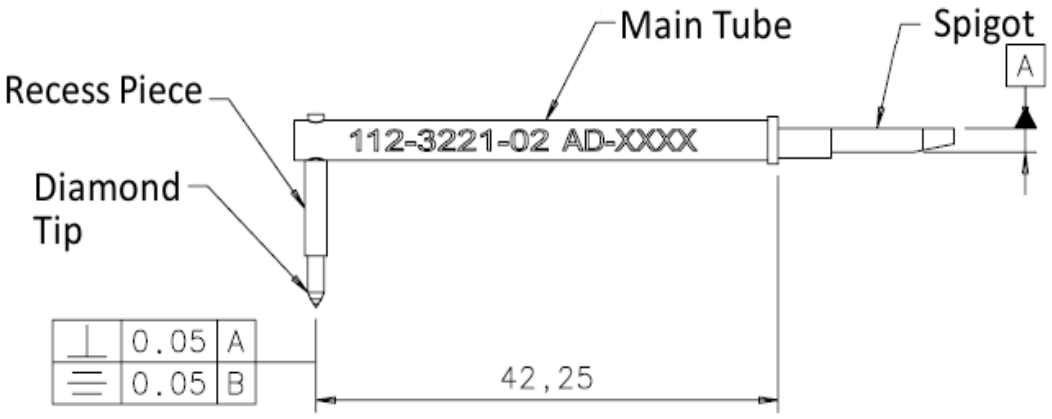


Figure 3–1: A common surface measurement stylus. The basic components are labelled, and the geometric tolerances are shown in the table (bottom left). This image is recreated here to simplify reference.

Spigot:

The spigot is a fundamental link in the metrology loop and so consistency and repeatability are paramount, as a result, the spigot has some of the tightest tolerances of any of the stylus components.

Main Tube

The main tube of the stylus sets the overall length of the stylus and as a result also the vertical resolution. The PGI gauge is fundamentally an angular measurement device and so a longer stylus will translate to a wider vertical range (for the same angular displacement), and a correspondingly reduced vertical resolution.

Recess Piece

The recess piece is required to extend the distance between the stylus tip and the body of the main tube. This length is commonly referred to as shank clearance.

Diamond Tip

The diamond tip if the precision contact surface of the measurement stylus. The tip must have a well-known form to allow the stylus correction algorithm to predict the actual surface from the path of the locus of the stylus tip.

3.1.2. Current tolerancing Specification

The drawings of the complete styli do not explicitly specify the positional tolerance of the diamond tip. The default drawing tolerance of (±100 µm) applies to the horizontal displacement of the tip from a feature on the spigot. Vertical displacement (shank clearance) is specified only where the tolerance is smaller than the default drawing tolerance. The shank clearance is referenced against the underside of the main tube.

Across the many years that these drawings have been authored, there are several different styles of tolerance and constraint. The perpendicularity constraint is not currently checked on the AV300 (Section 3.2.6) as the present fixture obscures the reference surface. The actual value of the

perpendicularity constraint changes between drawings. The other complication in this regard is that the reference surface changes between styli. These issues demonstrate a lack of consistency in the production of styli, maybe related to the CAD technology available at the time.

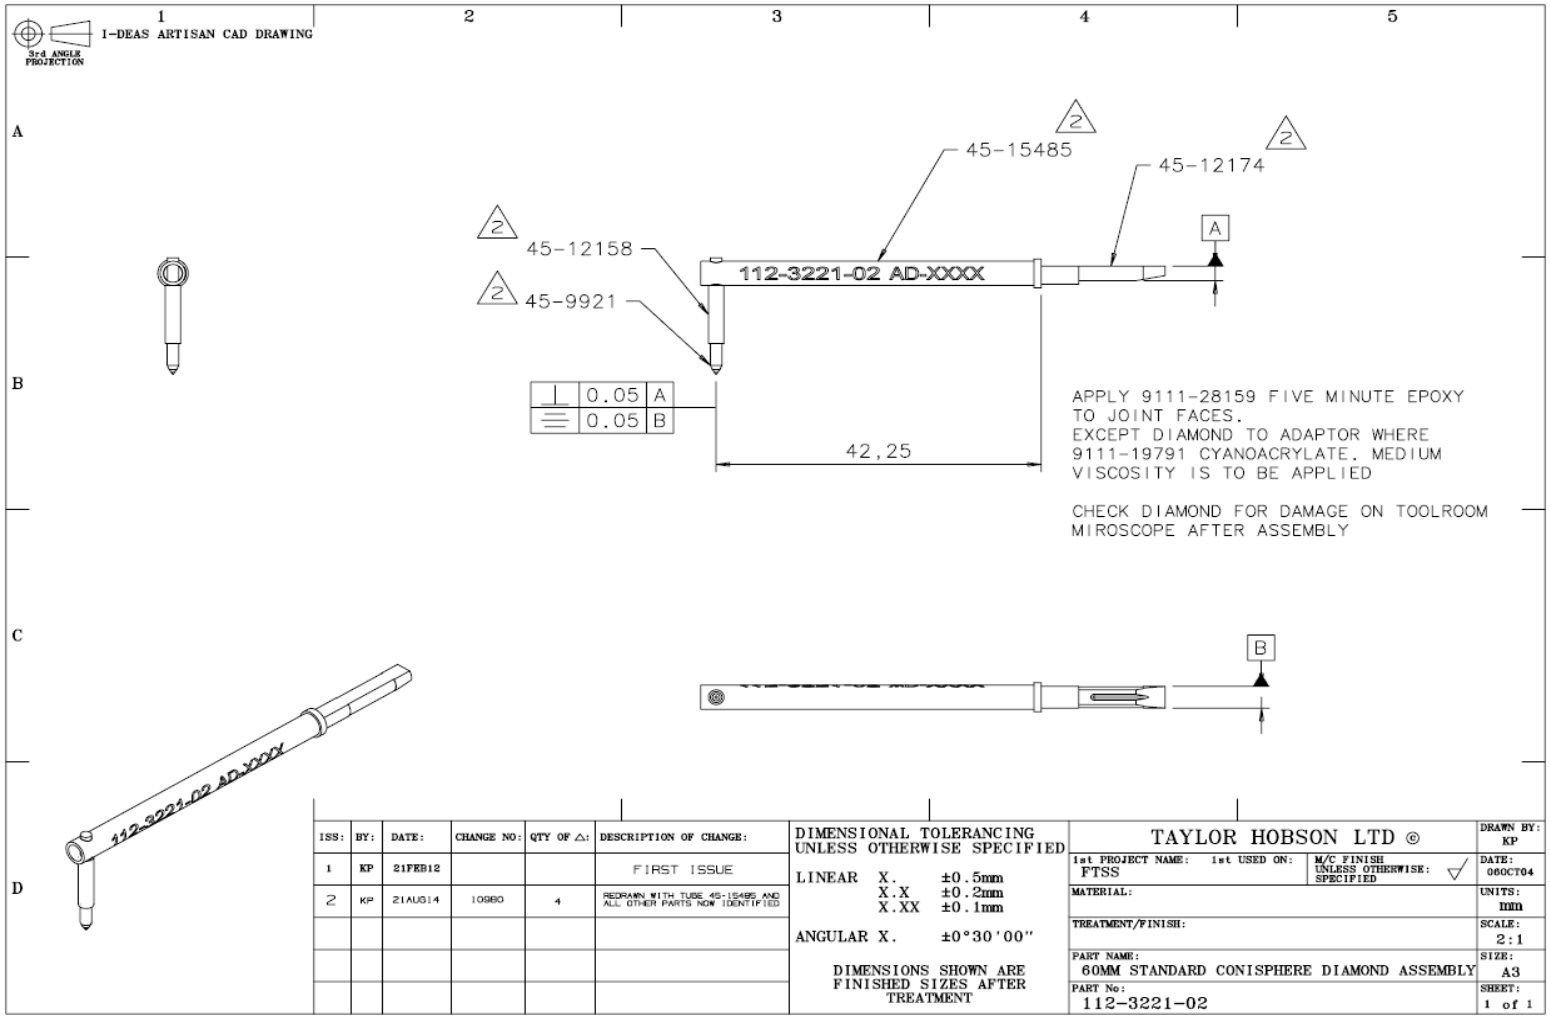


Figure 3–3: Inductive Stylus datum (from (Taylor Hobson, 2019))

Tolerance – Diamond mount symmetry and Perpendicularity

The diamond tip, and therefore the extension piece, of the stylus must be perpendicular to the flat on the spigot. There is a perpendicularity tolerance on the tip (of the 112-3221, for instance) of 0.05 mm. This means that the central axis of the diamond tip must be within ±50 µm of the ideal axis along its entire length. A typical stylus assembly drawing can be seen in Figure 3–4.

COMMERCIAL-IN-CONFIDENCE



COMMERCIAL-IN-CONFIDENCE

Figure 3–4: The assembly document for the 112-3221 stylus is shown. The perpendicularity datum (A) and the symmetry datum (B) are labelled, along with the corresponding tolerances.

3.1.3. Tolerance Stack Up

The positional and rotational tolerance of the measurement tip is the most critical dimensional tolerance of the stylus. This datum is measured offset from spigot and is different for the PGI and inductive styli, as shown in Figure 3–2 and Figure 3–3.

With each joint in the stylus, there can be several distinct modes of tolerance build up. An example of this is the fit of the spigot into the main tube. The sources of error for this joint are as follows:

1. The seating of the end of the tube against the spigot flat
2. The angular misalignment of the main tube onto the spigot
3. The fit of the spigot OD into the tube ID

This represents the fundamental degrees of freedom of the tube on the spigot and so some of the magnitudes of these will be constrained by the geometry. As the stylus must be a rigid component, the remaining degrees of freedom must be constrained by the adhesive. Some of these errors result in angular errors which therefore get worse as stylus length increases and others are fixed errors (concentricity errors, for example).

The overall specification of stylus tolerances is not universal. This means that some specifications that have been placed on the common designs may not be used in all cases and that two similar stylus designs may have different tolerances on the same dimension.

3.1.4. Spigot and Main Tube – Off Axis Rotation

The fit of the spigot into the stylus tube will give an angular error that is the same for a given length of stylus. The specification for the “Type A” tube is with an ID of 2.45mm, as per the internal drawing S54202. The tolerances on these dimensions are ±0.1 mm giving and ID in the range of 2.35mm to 2.55mm. The standard PGI spigot has an OD of 2.43mm to 2.45mm, which means that for almost half of the PCF tube ID tolerance the spigot would be

an interference fit. This demonstrates that the tolerance on the ID of the PCF tube is not appropriate for the tight tolerances of the spigot.

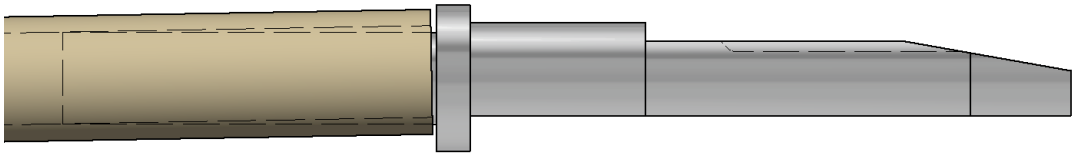


Figure 3–5: An exaggerated misalignment between the spigot (Right) and the main tube (Left) is shown.

With this tolerance range the best case is an identical ID and OD and therefore no angular error. In the worst case the OD of the spigot is 2.43mm and the ID of the tube is 2.55mm. The length of the inserted portion of the spigot is 9.9mm to 10.0mm. These numbers result in a potential angular error of 0.695°. This error would be very noticeable and may mean the stylus is investigated further. If this was the only error present in a typical 3221 stylus, then the diamond tip would be off by ~0.15mm along the axis of the main tube and ~0.50 mm vertically. This is well beyond the specified tolerances and would be rejected or re-worked upon inspection.

3.1.5. Spigot and Main Tube – Translation

The main tube is meant to be seated against the spigot. Any space between the spigot and main tube will result in a longer stylus than the specification. The extra displacement is not a predictable error, but it has a physical limit of 10.0 mm. The displacement error of shipped styli is limited as the stylus would be rejected at the visual inspection stage. This error is shown in Figure 3–6.



Figure 3–6: Gap between spigot and main tube is visible, taken from Previous work conducted at Taylor Hobson

3.1.6. Stylus and Main Tube – On Axis Rotation

The rotation of the stylus tube is not positively geometrically constrained. This means that any range of rotation of the main tube on the spigot is geometrically possible. There is a symmetry tolerance on the position of the stylus tip in relation to the spigot (datum B) of ±50 µm, shown in Figure 3–4. This tolerance is used to constraint the rotation of the main tube in relation to the spigot. The symmetry tolerance of the tip is small but typically achievable, and so is appropriate; however, there are a few specifics of assembly that must be highlighted. The lack of geometrical constraint means that the adhesive is the only mechanism used to set the orientation of the main tube on the spigot. When using current hard fixturing techniques: the adhesive is applied and fully cured before the stylus tolerances are measured. This results in styli that do not achieve the desired tolerance during assembly needing to be reworked.

The actual geometrical offset caused by a ±50 µm asymmetry error is small but the need for consistent stylus geometry is very important, especially in scenarios where the stylus is crested[4] on a component. This error may be alleviated by considering new manufacturing and assembly paradigms. This symmetry tolerance is measured on the AV300 but is not consistent across all

[4] Cresting a stylus refers to moving a component perpendicular to the measurement direction to ensure that the stylus will properly traverse over the highest point on the component during measurement. Cresting is particularly important when measuring spherical components such as lenses. Any departure of the stylus to either side of the crest of the component will result in a geometrical error in the measured profile.

stylus types. The 112-2009 has a symmetry tolerance of ±80 µm, as opposed to the standard ±50 µm of other styli.

3.1.7. Cross Hole Tolerance and Consistency

The general hole tolerances are defined by the internal document S54202, an excerpt of which is shown in Figure 3–7. These holes are currently drilled, as opposed to milled to dimension, which is likely to contribute to a lack of consistency, especially given the abrasive properties of the carbon fibre and the subsequent wear on the cutting tools.

The positional tolerance of the cross hole is ±20 µm. This seems appropriate considering the hole diameter tolerance is ±15 µm. There may be a need in the future for boring or milling of the cross holes to improve the consistency.

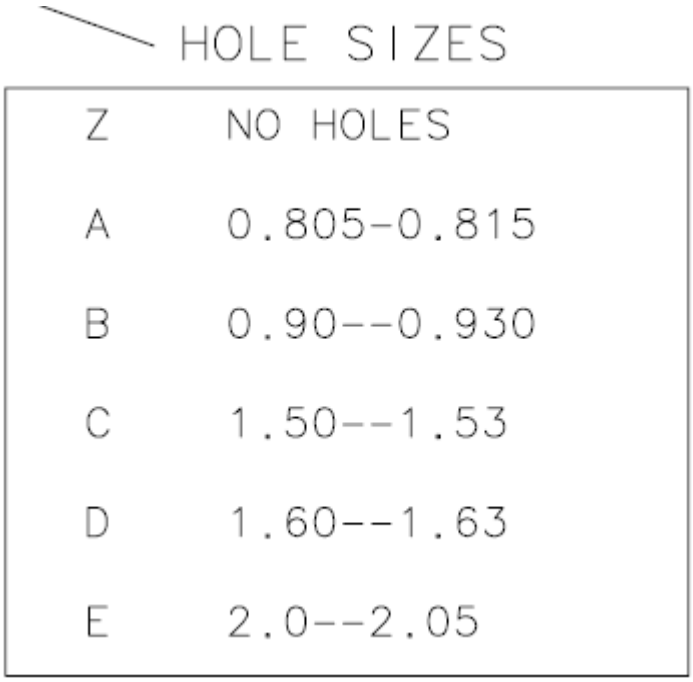


Figure 3–7: Cross hole size chart from internal document S54202

There may be a requirement to increase the nominal hole size if new assembly techniques can be employed successfully. The proper constraint of the stylus components would allow the fits between stylus components to be relaxed, easing the manufacturing burden.

3.1.8. Cross hole and Recess Piece

The cross-hole tolerance and the OD tolerance of the diamond or the recess piece is the source of the misalignment between these components. For the 112-3221 stylus, the cross hole is specified as (1.52mm – 1.55mm). This is not a standard size of cross hole (according to S54202) but the difference is not significant enough to warrant another drill size. The OD tolerance on the recess piece is 1.49mm-1.46mm. At best the size difference is 30 µm and at worst, 90 µm. This tolerance stack-up results in an angular misalignment between the main tube and the recess piece in the range of 0.5° and 1.55°. Practically this means the diamond tip of the 112-3221 can deviate by between 0.14 mm and 0.42 mm from the nominal point. The vertical change is minimal.

The potential deviation of the stylus tip is problematic, as even the tightest tolerance does not constrain the tip to better than the 0.1mm positional tolerance. This means that the constraint of the main tube to the extension piece is currently only achieved with adhesive.

As shown in Figure 3–8, the space between the main tube and the extension piece will be filled with adhesive if the gap is large (Blue Arrows). If the gap is reduced, then the adhesive will sit in a fillet alongside both components (Green arrows). This reduction in clearance between the components increases the geometrical constraint of the recess piece in the main tube. This is both advantageous and disadvantageous as the tighter tolerance will increase the difficulty of manufacture but will also increase the reliability and consistency of the final stylus when using the current ‘hard’ fixturing techniques

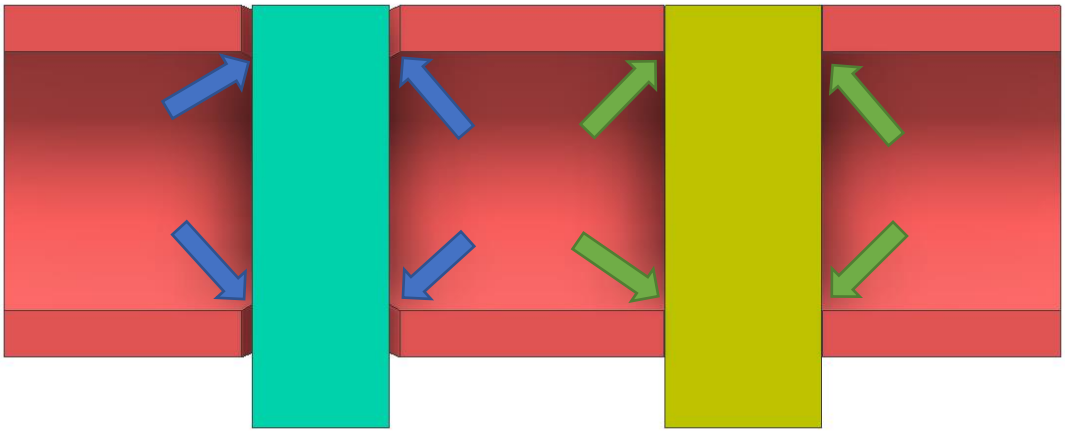


Figure 3–8: Cross section of two cross holes. These represent the joint between the main tube and the recess piece. On the left the gap between the main tube and the extension piece is large (Blue arrows) and on the right the gap is much smaller (Green arrows).

3.1.9. Diamond tip to Recess Piece

As with the previous error, there will be angular issues between the diamond tip and the recess piece. The OD of the mating spigot on the diamond tip used in the 112-3221 is specified as 0.86mm to 0.89mm. The ID of the hole in the extension piece is 0.90mm to 0.93mm, giving a range of clearances from 10 µm to 60 µm. The resulting angular range is 0.38° to 2.29° leading to a tip deviation from the nominal position of 0.04 mm to 0.24 mm.

3.1.10. Conclusions

The tolerancing of stylus components is presently disjoint, which may increase the complexity of component manufacture and storage, frustrate the stylus assembly process and increase the difficulty of achieving tight tolerances between batches of styli.

A top-down approach in creating a standardised suite of stylus drawings may be beneficial. An in-depth review of each production step within Taylor Hobson would allow the specification for all styli to be determined once, and only requiring updating as the available manufacturing and assembly technology improves.

The handling of the stylus component tolerance stack-up by the 'hard' fixturing must be investigated. Other approaches to stylus fixturing may prove fruitful, potentially allowing the stylus to be successfully assembled with tolerances that would result in an out-of-specification stylus if the hard fixturing were used.

During this research, it was observed that the datum face of the spigot is obscured by the measurement fixture during normal measurement on the AV300. The inability to reference the actual spigot may contribute to mismeasurement of the stylus length and symmetry, despite the quality of the measurement fixture. The stylus measurement process should be investigated further.

Given that the recess piece is geometrically constrained only by adhesive (according to the theoretical tolerances), the adhesive application process should also be investigated.

3.2. Summary of Current Stylus Fixturing and Development of Adjustable Fixturing

This section details the current approach that is taken to the fixturing of stylus components during the stylus assembly process. The deficiencies of the current assembly approach are be clearly documented as this will help to define the aims for a new fixturing solution.

The development of an adjustable fixture, along with accompanying design decisions, are also detailed. The proper documentation of the fixture design process will enable the fixturing approach to be easily reviewed as the capabilities of modern production technologies increase. The most obvious example of this is the growing applicability of additive manufacturing to fixture design.

3.2.1. Starrett AV300 and HAAS

Starrett AV300

The AV300 is a measurement microscope designed for automatic measurement and validation of components (See Figure 3–9). It has 3-axis motion capabilities, using high precision linear rails and magnetic linear motors on the X and Y axes. The motion capabilities are within single microns and the encoder resolution in 0.1um.

The table has back illumination and the object can be lit from within the microscope lens arrangement. Also included is a LED ring which helps to provide contrast when measuring a feature not on the silhouette of a component.

The AV300 measurement software uses contrast-based edge detection. From the contrast data it can create artificial axes and datum references. The ability to derive a measurement axis from a component is fundamental to the new style of fixturing. This effectively mitigates the requirements for the tight tolerance on existing measurement fixtures.



Figure 3–9: The Starrett AV300 measurement microscope.

HAAS

The HAAS machining centre is a machine tool that is used to produce almost all the carbon fibre stylus components. It is a standard 3-axis machining centre with the addition of a 4th and 5th axis on the bed of the machine. The spindle is capable of 10,000 RPM, the machine is equipped with a 24-pocket tool changer and has a comprehensive flood coolant system.

An advantage of the multi-axis capabilities of the HAAS is the ability to machine certain styli from solid. TalySeries™ styli are produced this way to improve the production rate. Other complex features can be machined such as compound angles or steps.

The advent of the HAAS has resulted in tighter and more consistent tolerances on holes and stylus lengths. The tighter hole tolerance has resulted in an increase in the amount of broken stylus main tubes. Process development is still on-going.

3.2.2. Background to Styli and Fixturing

Styli are designed to interface a standard measurement gauge with a range of different measurement tasks. Surface profile styli have tight assembly tolerances which ensure that a new stylus will be immediately compatible when it replaces an old stylus in a measurement system.

Styli are currently assembled using ‘hard’ fixtures as shown in Figure 3–15. This assembly method relies heavily on tight manufacturing tolerances and may induce unnecessary internal stresses within the finished stylus. This technique also requires rework in some instances. Any stresses introduced to the stylus at the assembly stage may not manifest as dimensional changes immediately. The large temperature and humidity changes associated with cargo shipment may mean an in-specification stylus may leave the UK and an out-of-specification may arrive at the destination. The replacement process is difficult and expensive and so internal stresses must be low from the start.

3.2.3. Motivation for Investigation

Taylor Hobson assemble upwards of 5000 styli per year, spanning more than 500 models. Roughly half of these styli are the most common ‘standard’ styli with the remainder being mostly special styli, produced in small quantities per month. The special styli are made up of over 450 different models and most (80% by number) of the special styli are made in batches of 2.

There are ~100 stylus types that are considered ‘standard’ or have the same geometry as a standard stylus. These styli may have dedicated assembly fixtures, dependent on the quantity of styli and the availability of time to design and produce fixtures. Fixturing of the styli allows for a considerably higher throughput and better consistency between batches of styli.

The fixtures that are currently used for assembly are different from those used for inspection; these dedicated assembly fixtures must be made within tight tolerances and are not adjustable. The lack of adjustability in the fixturing can result in a stylus failing to achieve the geometric tolerances, despite each individual component being within specification. This is caused by the tolerance stack-up of multiple components. The magnitude of the tolerance stack-up error increases as the number of components increases.

Most approaches to fixturing would help keep the tolerances and production standards high, although not all fixturing approaches reduce the need for rework. There are several benefits to using an adjustable fixture:

- Low-force component retention reduces the stresses imparted to the stylus while maintaining the geometric accuracy.
- Adjustability at cure-time reduces the need for post-cure rework.
- The AV300 is used to measure the stylus geometry during assembly, therefore the geometric tolerances on the fixture are much less critical.
- The use of standard components reduces the development time and machining complexity for new fixtures.

The new stylus fixture approach uses the AV300 for measurement during assembly. The use of the AV300 is advantageous because the tolerances of the fixture and some stylus components can be relaxed, while the accuracy of the finished stylus maintained by in-process measurement. However, this approach has the disadvantage of requiring the AV300 for assembly, where there was no such requirement beforehand. The amount of time per stylus on the AV300, during assembly, is short (~5 minutes). The fixture can be removed from the AV300 after the measurement process which helps to reduce the utilisation of the AV300s.

In the case that a stylus can now be assembled in a new fixture and was previously un-fixtured; the utilisation of the AV300 is significantly reduced. This reduction in utilisation is because non-fixtured styli are held in the measurement fixture with the individual components supported using gauge blocks and similar, while the adhesive cures. This process may take up to 30 minutes, during which time the AV300 cannot be used for other measurement.

3.2.4. Specification of Stylus Geometry

The specifications on stylus assembly and production have improved with each upgrade in manufacturing or measurement technology. The use of the HAAS machining centre has greatly increased the repeatability, and reduced the rejection rate, of stylus components. This is in contrast to the manual processes that were time consuming and lacked the repeatability of the modern machining centre.

The AV300 has helped remove the operator judgement and standardise the measurement procedure. This has resulted in a more traceable and reliable measurement process. The AV300 has allowed the automatic generation of a document for recording the measured tolerances of the stylus and can measure highly complex styli with little penalty.

The assembly drawings of styli use features on the spigot as datums, whereas the measurement process measures relative to the fixture or the stylus tube. There is also an implicit tolerance on the effective length (tip to spigot). See

Figure 3–4. A modified summary of the stylus dimensional specification is shown here: Figure 3–10

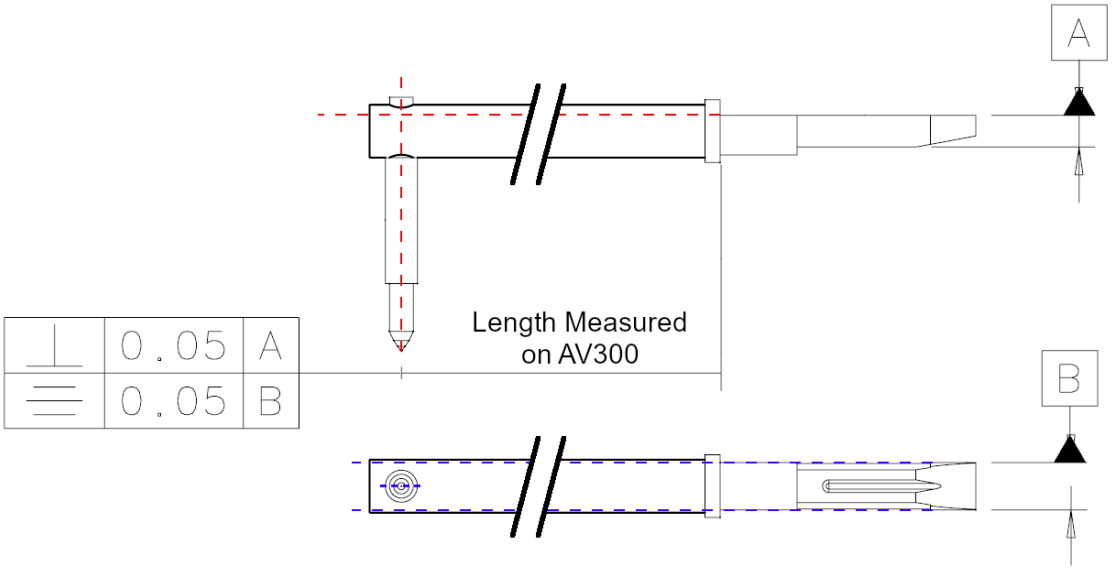


Figure 3–10: A diagram showing the datums of the perpendicularity tolerance (Red) and the symmetry tolerance (Blue). The stylus length that is measured by the AV300 is also labelled.

3.2.5. Build Fixtures

There are many different 'hard' fixtures that are used within the stylus assembly cell. An example of a typical hard fixture is shown in Figure 3–15. These fixtures are made from ground tool steel and are based on traditional fixturing techniques. The designs of these fixtures limit the measurements that can be taken before the stylus is removed from the fixture, as the spigot is obscured by the retention clamp.

These fixtures allow for fast assembly of a specific stylus but require tight manufacturing tolerances on stylus components and are not resilient to the compound stack-up of component tolerances. The high retaining force experienced by the stylus components may also impart stress into the finished stylus.

As previously mentioned, these fixtures allow for increased production rates compared to styli without designated fixtures. The limited adjustability of these

fixtures is both beneficial (simple: set once; never adjust) and limiting (small range of appropriate styli, tight fixture and component tolerances necessary)

The fixtures are made from tool steel and as such are difficult and expensive to make. The non-adjustable features in the tool steel are required to be within ±10µm so that the assembled stylus meets the symmetry and length requirements. This tight tolerance and material choice require the fixtures to be ground, which is typically an expensive and time-consuming process. This process is slow, and so the fixture development and procurement time is long.

However, the tool steel is very wear resistant, and as such, will withstand the assembly of thousands of styli. If the range of manufactured styli were limited, then this style of hard fixturing would be appropriate.

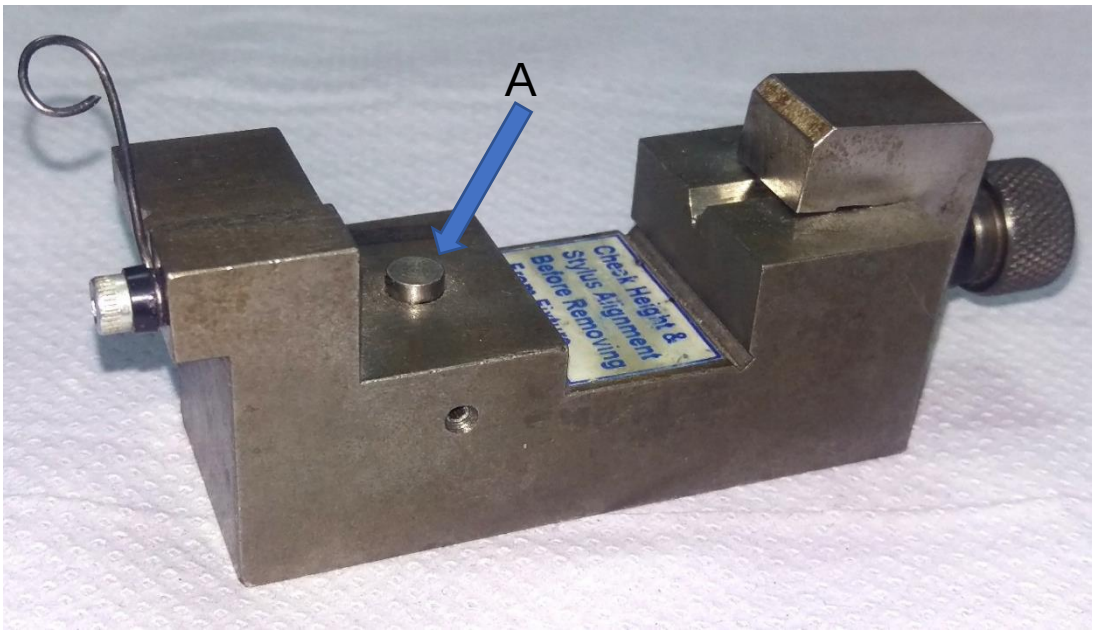


Figure 3–11: A picture of the rigid fixture used to assemble standard Inductive recess styli. The effective length of the stylus is fixed by-design. The vertical attitude of the main tube can be set by adjusting the lower stop, labelled A.

Some hard fixtures were designed to accommodate a range of stylus lengths by mounting the spigot support block and the recess piece support block onto a ground rail. This design can be seen in Figure 3–16. Each length option has a set of corresponding alignment holes to allow a good repeatability between length changes. The downside to this approach is the lack of geometric constraint on the non-spigot end of the stylus main tube. There is

also no provision for in-process measurement of the stylus during assembly, thereby potentially increasing the number of styli that must be reworked.

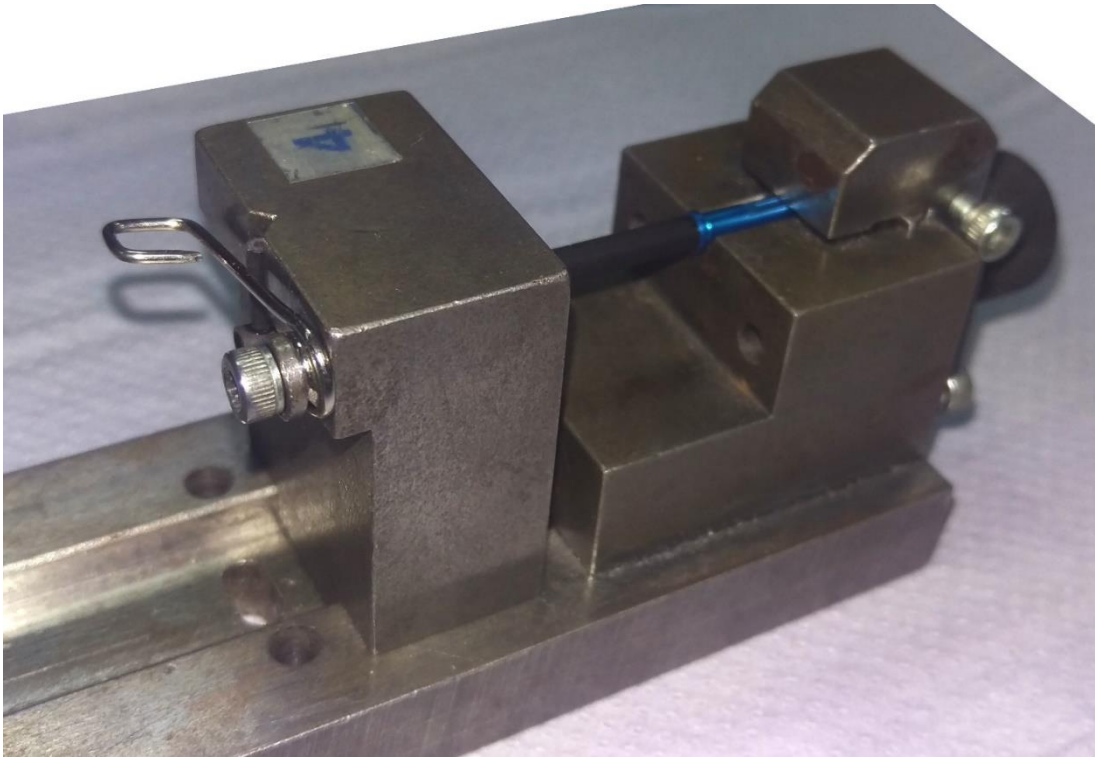


Figure 3–12: This fixture is similar to Figure 3–15, although this fixture is adjustable in length. The attitude is not controlled by a hard stop as in Fig. 3, but instead by the 'roof' of a machined hole in the movable portion of this fixture.

3.2.6. Current Stylus Geometry Measurement Process

After the stylus has been assembled, and the adhesive has cured, it is removed from the assembly fixtures and measured using a dedicated measurement fixture. The current stylus measurement fixture (shown in Figure 3–13) is designed to clamp the spigot of the stylus onto a vee that is ground along the centre of the fixture. The spigot rotation is set by a retaining clamp and thumb screw. This completely encloses the spigot and therefore measurements of stylus geometry must be made relative to the fixture.

The measurement process happens in 2 steps, Part 1 and Part 2, with the pertinent data being transferred to an excel spreadsheet after each step.

Stylus Measurement - Part 1

The stylus is secured into the fixture by hand, which is then placed onto the measurement bed of the AV300, oriented such that the tip points upwards and the spigot is to the left, according to the microscope view. The orientation of the measurement fixture for part 1 of the measurement process, from the point of view of the AV300, is shown in Figure 3–13

Part 1 of the measurement uses the outer edges of the fixture to define the Y-axis datum. The ground surface at the front of the fixture is used to set the X-axis datum. The operator positions the measuring crosshair over the stylus tip and the offset information is recorded. Specifically, the displacement from the front ground surface of the fixture to the stylus tip (effective length) and the displacement from the centre line of the fixture to the stylus tip (symmetry).

This process does not account for skew of the measurement fixture and so the absolute position and rotation of the fixture fences on the AV300 influences the measurement process. This is mitigated by a regular calibration schedule, as opposed to correcting the skew for each measurement.

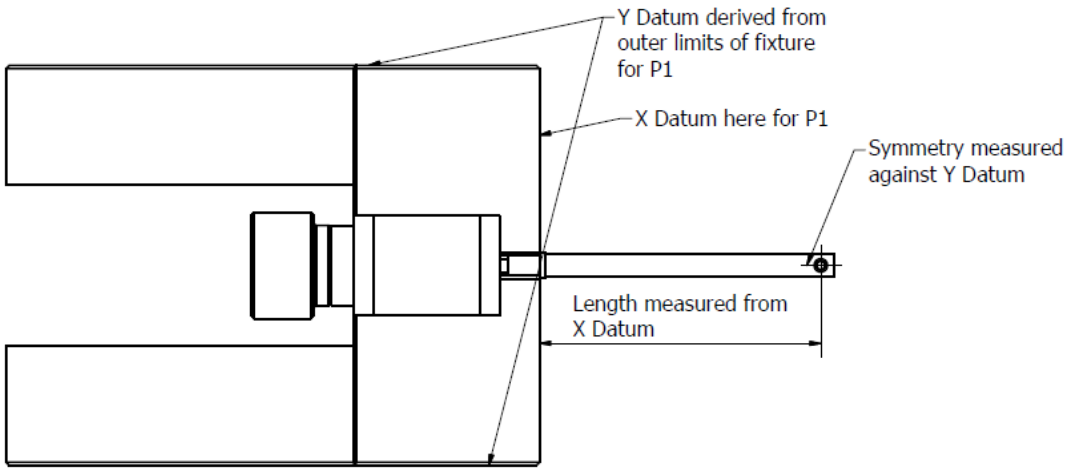


Figure 3–13: Diagram of the current rigid stylus measurement fixture, as seen from the point-of-view of the Starrett AV300 measurement microscope during part 1 of the stylus measurement process. The fixture is made from tool steel.

Stylus Measurement – Part 2

Part 2 of the measurement process is used to measure the shank clearance, so the stylus fixture is reoriented such that the stylus is in profile. The orientation of the measurement fixture for part 2 of the measurement process, from the point of view of the AV300, is shown in Figure 3–14. The Y-axis datum is taken from the (now) top of the fixture and the X-axis datum is preserved from Part 1 of the measurement process. The attitude[5] of the main tube is manually inspected but not recorded as there is no explicit tolerance placed on this angle at design time. A measurement is made from the bottom of the stylus main tube to the stylus tip to determine the shank clearance. The shank clearance is specified explicitly only for certain styli but is recorded for every stylus.

The geometric measurements from Part 1 and Part 2 are then transferred to an Excel™ document along with an image of the stylus tip. This document is included with the stylus and a copy is held within Taylor Hobson for process monitoring purposes.

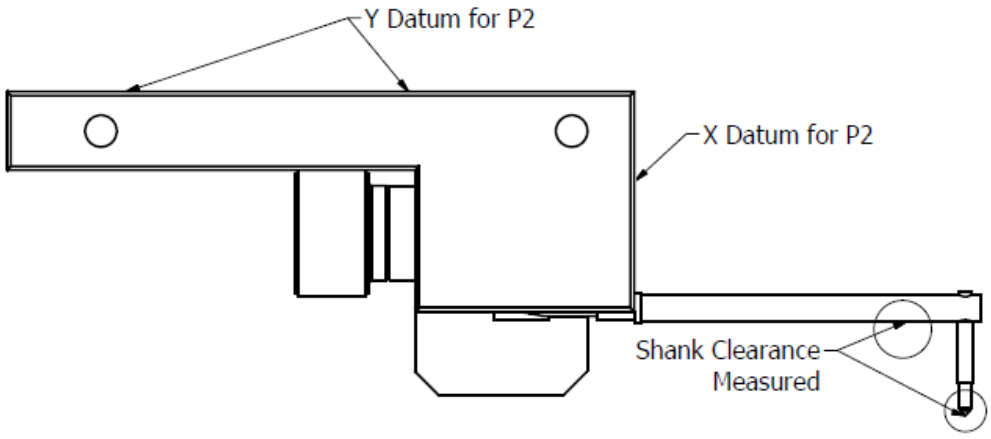


Figure 3–14: Diagram of the current rigid stylus measurement fixture, as seen from the point-of-view of the Starrett AV300 measurement microscope during part 2 of the stylus measurement process.

[5] The attitude of the main tube refers to the angular difference between the axis of main tube and the spigot, in the vertical plane of the stylus.

3.2.7. Alternative approach overview

This section details the development of a new stylus assembly fixture. This fixture was designed using a new approach to take advantage of readily available standard components and measurement tools. The fixture is easily adjustable, made from readily machinable materials and use standard components where possible. The fundamentals of conventional fixturing theory are still valid and have been summarised in Section 2.2.

As previously mentioned, the current fixturing approach has several shortcomings. These shortcomings will form the basis for evaluating new fixturing approaches.

In the current assembly process, the inspection of the stylus geometry is done only after curing has taken place. This means that any errors in the assembly process can only be corrected with the application of heat to the adhesive to allow for manual adjustment.

The styli that are assembled using the hard fixturing are prone to stress build-up, even when components are within drawing tolerances, because of the compound effects of the component tolerances on the final assembly. This happens because the fixture relies on the fit between mating components for alignment.

The hard fixturing is typically expensive and is limited in the range of styli that can be assembled. The cost of the fixture and the long development time mean that hard fixtures are available for only the most frequently produced styli. The expense usually means that only a single instance of a particular fixture is available, effectively restricting batch production.

These issues naturally result in a set of aims for the new fixturing approach:

1. The new fixture must be adjustable.
2. The new fixture should allow measurement during the cure process.

3. The new fixture should be cost effective when compared to the current hard fixtures.

The fixture being adjustable ensures that the tolerance stack-up will not have a detrimental effect on the produced stylus. This inherent adjustability allows some tolerances to be relaxed (cross hole diameter) and therefore the fixture must properly constrain all components of the stylus.

The current post-cure adjustment process uses a soldering iron to heat the 2-part epoxy and then the components are adjusted by hand. The epoxy is detrimentally affected by the heat which may have knock-on effects on the performance of the stylus. This rework process must be avoided where possible and is only necessary because they styli cannot be measured in the stylus fixtures during assembly. The stylus must be measurable and adjustable during curing.

The prohibitive cost of the current stylus fixtures, along with the long production time, means that very few models of styli have dedicated fixtures. The higher utilisation of the new fixture design during assembly means that multiple fixtures would help increase the throughput. The cost of the new fixtures should be less than 50% of the old fixtures, although lower is preferable.

A render of the proposed stylus fixture is shown in Figure 3–17. This design allows for in-process measurement and subsequent adjustment of the stylus during curing, ultimately reducing the need for stylus rework. The utilisation of the AV300 during stylus assembly is reduced as the fixture can be easily removed from the measurement platen. However, the fixture utilisation is generally increased to allow the adhesive to cure. This new assembly paradigm helps to decouple the AV300 utilisation from the fixture utilisation. This leads to an overall increase in assembly throughput. A further reduction in the number of styli requiring rework would also contribute to a more efficient assembly process.

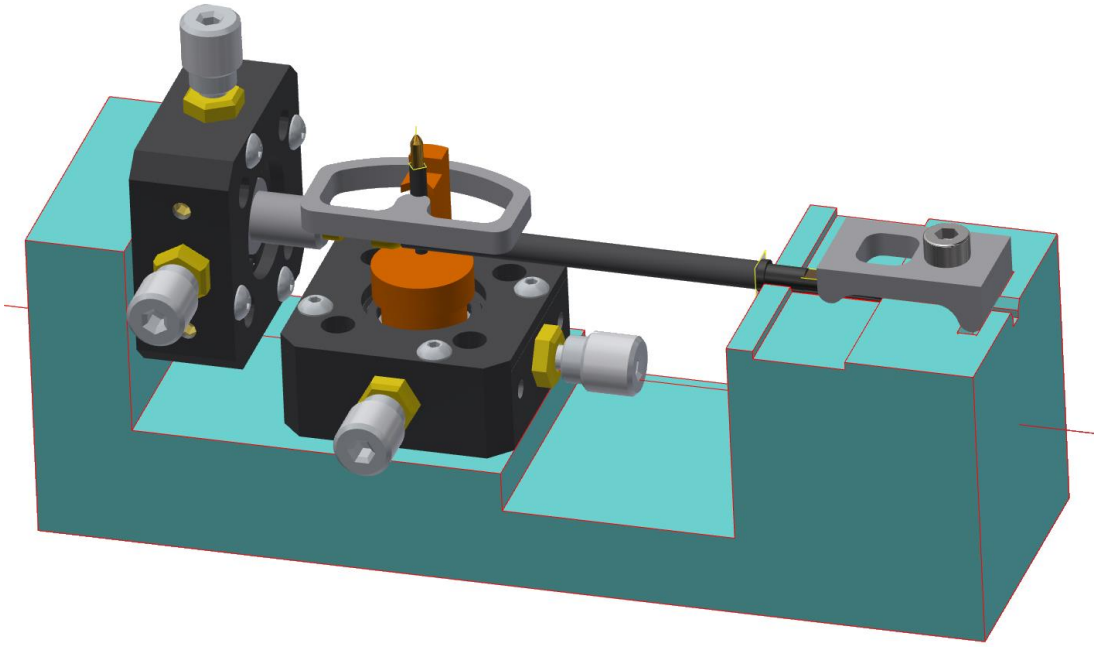


Figure 3–17: A render of the new stylus fixture, including a stylus and retention clip. Obscured by the retention clip is a spring plunger which retains the main tube against the spigot.

The performance of the stylus fixture relies on the precise adjustment of the stylus components during assembly. This adjustment must be quick, simple and backlash free; a task that typically takes a lot of engineering and finesse to achieve. This fixture uses standard components from the optics industry to achieve the fine motions necessary for stylus component adjustment. These standard components are available off-the-shelf, from several reputable companies, with a short lead time. For example, ThorLabs offers a 2-axis, flexure-based translation mount for lens systems. See Figure 3–18. (ThorLabs, 2014)

This translation stage has a 1mm range in each axis, eliminating the need for high-precision assembly of the assembly fixture itself. The central bore is threaded using a ‘standard’ optical thread (Ø0.535 in - 40) but accessories are available that allow for interfacing without the need to create this thread.



Figure 3–18: ThorLabs translation mount used in the adjustable stylus fixture.

3.2.8. Fixture Development Approach

A fixture for stylus assembly has been developed with several key aims that differ from the conventional design approach.

1. The fixture should be readily adjustable and remove the need for precision adjustment by hand.
2. The fixture should be cheap. This allows many fixtures to be produced and used at any one time.
3. The fixture should be easily manufactured on conventional machine tools, preferably from free-cutting materials such as aluminium or brass.
4. Any externally manufactured components should be as universal as possible to facilitate re-use in other fixtures.
5. The number of machining operations should be minimised. If the machining can be done in one operation then the fixture can be produced from bar-stock, greatly reducing the operator involvement and overall machining time.
6. The stylus fixture must allow in-process measurement as this is now a requirement of the assembly process.

Fixture Development Constraints

1. The spigot must be properly constrained in a repeatable and reliable way.
2. The spigot, main tube and diamond, must be visible from the AV300 in the 2 inspection planes.
3. The main tube must be constrained at the non-spigot end. It must also be retained on the spigot and allow for rotation of the main tube on-axis.
4. Recess piece must be retained in the main tube
5. The position of the diamond should be adjustable in effective length and symmetry.

New fixturing approach – achievement of proper constraint

The stylus spigot must be reliably positioned and retained as it is both a datum in the stylus assembly process and the geometric interface between the measurement gauge and the stylus tip. The body of the stylus is cylindrical, allowing a vee groove to be used for location. The rotation of the spigot is set by the retention piece. This spigot locating features are shown in Figure 3–19. The spigot retention piece is shown in Figure 3–20.

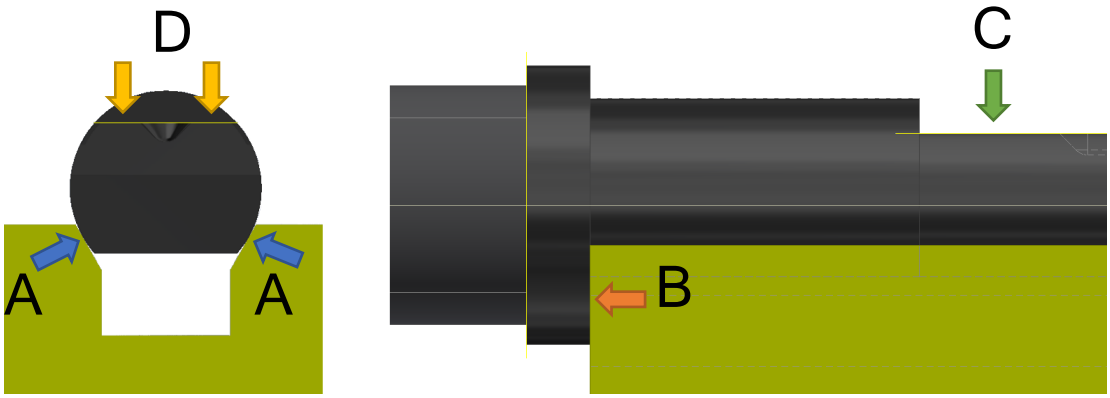


Figure 3–19: Schematic of the spigot retained in a well-constrained fixture. 4 DOF are constrained by the 'V' groove at A. One by the shoulder at B and finally the on-axis rotation is set at D. The spigot is retained in the 'V' by a retaining clamp at C which also serves to set the rotation.

The retention of the spigot must not obscure the view of the datum features from the AV300. Figure 3–20 details the features that are used for to derive the spigot datum.

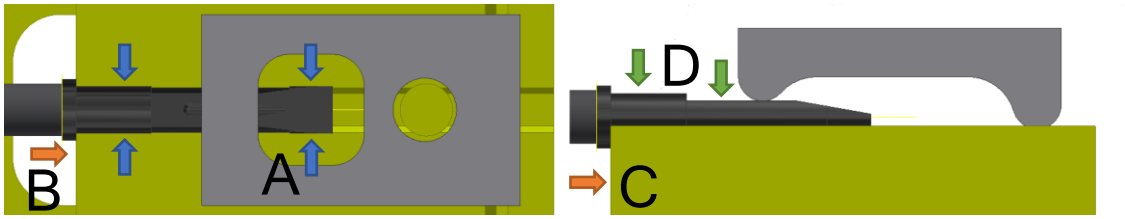


Figure 3–20: Schematic showing the spigot from the point of view of the AV300 during measurement. In P1 (Left), the X-axis datum (and skew) is taken from edges A, the Y datum is taken from B. In P2 (Right), the X-axis skew and X-datum are taken from D, the Y-axis datum is taken from C.

The stylus tip must be adjustable in symmetry and length, without the adjustment mechanism getting in the way of the measurement of the stylus tip. The stylus fixture should be as universal as possible, so the recess length of the stylus should not be limited or enclosed by the fixture. The main tube is adjusted on-axis, from the non-spigot end, which limits the length of the stylus in each fixture. Since the range of stylus lengths is small, only a small variety of fixtures are necessary.

The recess piece is adjusted from the back of the stylus. The adjustment thumbwheels are readily accessible, and the length of the recess piece is not limited by the fixture design.

The recess piece must be retained against the 'V' of the recess actuator (Labelled A in Figure 3–21). This may be achieved with a small elastic band or similar. This must not be a rigid coupling as this imparts stress into the stylus during assembly. This stress lead to the stylus changing shape once it is removed from the assembly fixture.

After testing the fixture in a real-world assembly scenario, a 3D printed retainer clip was designed (Figure 3–25). The clip takes less than 10 minutes to print and can be used without clean-up, straight from the bed of the printer. The use of the 3D printer allowed a very fast development time and the printed parts are suitable for production use.

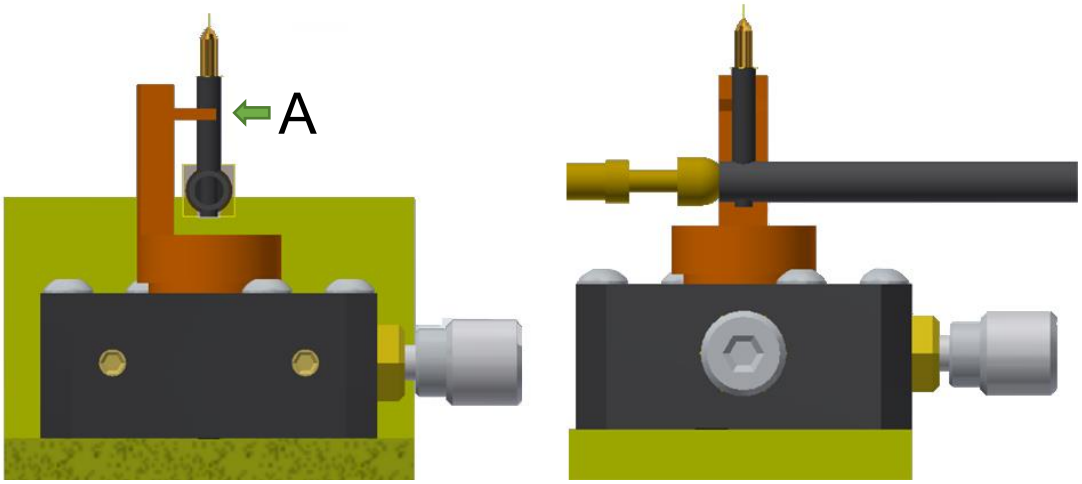


Figure 3–21: The recess adjustment mechanism is shown. The recess piece orientation is adjustable so that the position of the diamond can be set in terms of effective length and symmetry.

3.2.9. Realised Prototype

The fixture design was developed and improved several times before a physical prototype was realised (Figure 3–22). This prototype was produced for several reasons: to analyse the ergonomics of the fixture, to test the interaction between the new fixture and the Starrett AV300 and to produce some test styli for the adhesives testing investigation. A PTFE boss was used to locate the end of the stylus main tube; later being replaced with a pin-plunger.

The stylus main tube and spigot are physically constrained at A in Figure 3–22, above the length adjustment knob (labelled C in Figure 3–22). A through hole (labelled B in Figure 3–22) is provided to allow access to the adjusting knob by means of a hex-key. This was found to be unnecessary, as the length adjustment knob is not difficult to actuate, despite the presence of a stylus. This hole was removed from subsequent revisions of the fixture.

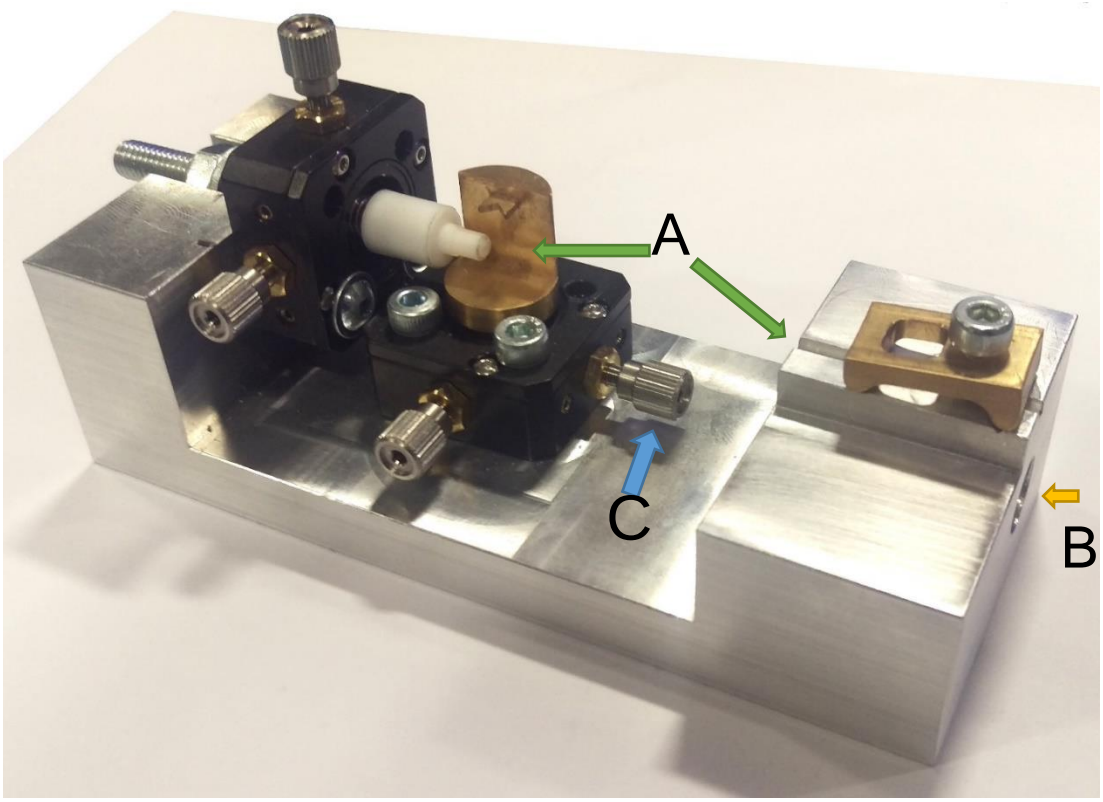


Figure 3–22: The first realised prototype of the adjustable stylus assembly fixture.

The AV300 was able to measure the spigot position in both fixture orientations. This is an important milestone in proving this new fixturing approach. The repeatability of AV300 measurement is documented in Section 3.2.16. This shows that the spigot is reliably locatable to ±4 µm. The magnitude of this errors influence on the stylus geometry is calculated in minimal, and well below the measurement error.

This prototype fixture also allowed the testing of other adhesive application and dispensing procedures.

3.2.10. Adhesive Application Testing

The prototype fixture allowed several novel adhesive application methods to be tested. The principal benefit of the developed fixturing approach that enabled the adhesive testing is that the stylus components are retained by the fixture, without need of adhesive.

Hypothesis

The assembly fixture can be used to hold the stylus components in place, without adhesive, such that the adhesive may be injected into the joints.

Testing

To reduce the complexity of the injection testing, only the joint between the spigot and the main tube, and the main tube and recess piece were tested. The diamond can be affixed into the recess piece beforehand as the assembly process is quite straightforward and the adjustable nature of the fixture allows small misalignments between the diamond tip and the recess piece to be mitigated.

The joints being tested are small and, in the case of the spigot to main tube joint, geometrically not suited adhesive injection. This necessitated the use of a low viscosity adhesive so that the adhesive could flow around the spigot. Most epoxy-based adhesives were too viscous to be readily injected, leaving only ‘thin’ (i.e. low viscosity) cyanoacrylate (CA).

Two 1 mm diameter holes were drilled in the side of the main tube, offset 90 degrees from the cross hole already present. The holes were positioned in the centre of the spigot boss (5 mm from the spigot shoulder) and slightly offset from the back of the cross hole (4 mm from the stylus tip end). These positions were chosen to allow the best access when using a syringe to dispense the adhesive.

Assembled styli were left to cure and then inspected to determine the bonding performance of this approach.

A stylus spigot is modified to allow adhesive to flow between the spigot and the ID of the main tube, whilst also locating the spigot within the main tube. A schematic of the modified stylus spigot is shown in Figure 3–23.

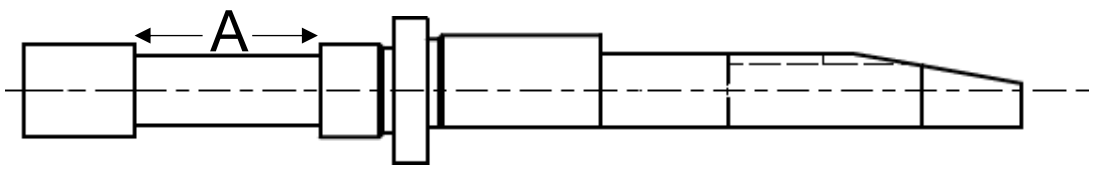


Figure 3–23: A modified version of the standard spigot. Space is left for the adhesive on the mating stem of the spigot (A).

Results

For the joint between the spigot and the main tube:

Despite using the modified spigots to allow space for the adhesive, the CA adhesive would typically cure before completely flowing around the joint. Cyanoacrylate adhesive cure time is inversely proportional to the thickness of the adhesive joint and is accelerated in response to moisture. The small adhesive joint thickness likely reduced the cure time to ~1 second; not enough time to allow the adhesive to flow. Efforts were made to remove moisture from the stylus components, but the testing did not give improved results.

For the joint between the main tube and the recess piece:

The space for adhesive in this joint is larger, and therefore the injection process itself was straightforward. The low viscosity of the CA adhesive meant that the adhesive tended to coat the inside of the main tube and would not fill the gap between the recess piece and the main tube. A medium viscosity CA was tested, giving much better results and a more predictable application, however the application process itself was more difficult as the medium viscosity CA adhesive would not flow easily down the 0.95 mm diameter needle.

The injection process may be developed in the future to give consistent results. However, the requirement of 2 different viscosities of adhesive per stylus, and the inconsistent bonding of the spigot to the main tube, has led to the rejection of this as an assembly strategy and is not recommended for further investigation at this time.

3.2.11. Stylus Assembly Summary

The stylus assembly procedure for the newly developed fixture is shown here:

1. The stylus components are dry fitted into the assembly fixture.
2. The fixture is adjusted to give a correct stylus configuration without the adhesive being applied.
3. The carbon fibre components are removed, and adhesive is mixed and applied to the appropriate joints.
4. The components are reassembled, ensuring good coverage of the adhesive, and put back into the fixture.
5. The stylus is immediately measured, and small adjustments are made to bring the stylus within geometric tolerance.
6. The stylus is left to fully cure before being removed from the fixture.

This process can be comfortably completed in 1 operation with adhesives that have a cure time of 10+ minutes. With quicker curing adhesives, the assembly process may need to be broken up to cure a single joint at a time. Splitting the assembly process does not significantly affect the total assembly time.

3.2.12. Recess Piece Retention Clip Development

The method of retention for the recess piece against the brass 'V' was observed to influence the final geometry of the stylus. Several retention methods were tested, from soft soldier, wire, elastic bands and blu-tac but these solutions were ultimately rejected.

This led to the design of a 3D printed retention spring that increases the consistency of the assembly process. See Figure 3–24 for a CAD model of the developed retention clip and Figure 3–25 for a photograph of the printed part.

Failed Retention Methods

The intention at design time was to use a short elastic band to retain the recess piece against the brass 'V'. The intent was to provide a retaining force without geometrically influencing or over-constraining the recess piece. The elastic

band provided the retention force but would not allow the main tube attitude to be adjusted without separating the joint between the recess piece and the main tube. To avoid this separation, the elastic band would have to be adjusted after initial measurement of the stylus, requiring a second run through the measurement process to ensure the stylus was still in tolerance.

The wire and the soft solder were able to slide easily along the recess piece, therefore reducing the magnitude of the separation phenomenon see when using the elastic band. Neither the wire nor the solder would provide an active retention force, instead acting as an adjustable limit to the separation between the recess piece and the 'V'. The wire and solder would also need to be unwound and rewound for each stylus, quickly fatiguing the retention piece.

Blu-tac was used in an effort to avoid the design and production of a retention component. The retention force was very unpredictable, and the blu-tac would cause the joint to separate as the elastic band had done.

Developing a 3D Printable Clip

Once the previous methods had been ruled out, a new design for a retention clip was sought. The measurement process needed to be able to measure the symmetry of the main tube so the clip must not obscure the view of the main tube during Part 1 of the measurement process. For this reason, the clip was designed to be larger than necessary, with the retention geometry inset from the edge (Labelled A in Figure 3–25).

The curved portion of the clip was designed to allow the clip to flex, providing the retention force, and for the clip to be twisted into place over the recess piece and the brass 'V'.

The clip thickness is 2 mm. This was chosen to ensure the printed component was robust enough to be useful, but to not be significantly taller than the brass V geometry (1 mm).

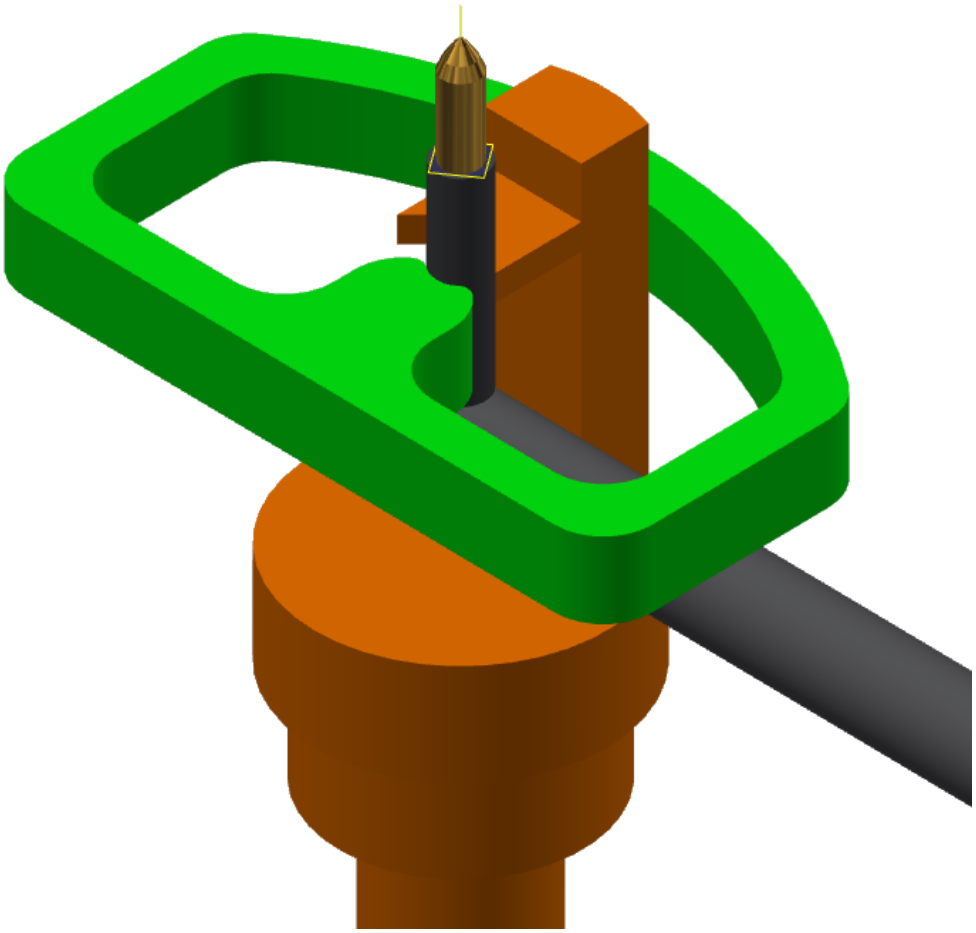


Figure 3–24: A render (Right) of the 3D printed retention clip (Green), with the stylus (Black) and the brass recess adjustment piece (Orange)

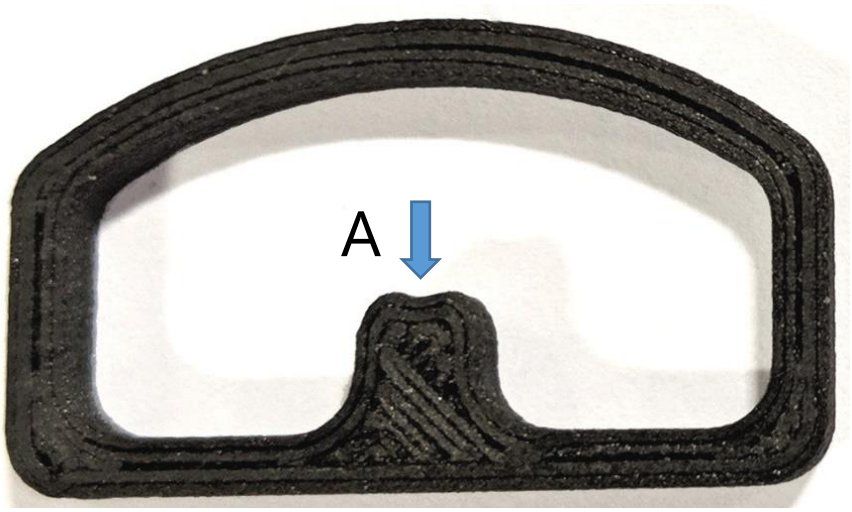


Figure 3–25: A picture showing the 3D printed clip developed to retain the recess piece against the brass 'V' of the new assembly fixture

The print settings were adjusted to change the retention force and the final print settings recorded for future prints. The component required a small amount of finishing work to remove the first layer burrs to ensure reliable retention.

The thickness of the retention clip did have a small effect on the recess piece orientation. The positioning of the retention clip above or below the brass V would change the angle of the recess piece around the axis of the main tube. This effect was well understood within a few assembly cycles and was not difficult to mitigate.

3.2.13. Design for Manufacture Analysis

The machine shop within Taylor Hobson uses 2 standard sizes of bar to simplify the purchasing process. The fixtures are designed to fit within the larger bar stock which allows for simple batch production. The current design of the fixture is machinable in a single setup if it is produced on the Mazak.

Advantages when producing fixtures from bar stock, using a 5-axis machining centre include:

- A reduced overall fixture cost.
- Much lower lead-time.
- Shorter design feedback time.
- Simple and fast adaption to new stylus designs.

The use of aluminium instead of tool steel means that it is possible to machine the fixture in-house. Aluminium is cheaper than tool steel, although it will wear more easily and need replacing eventually. The use of the AV300 in the measurement and assembly process allows the manufacturing tolerances of the fixture to be larger without impacting the performance of the fixturing approach, helping to mitigate the impact of wear to the fixture.

The fixture was designed to be machined on the Mazak Integrex i200 Mill-turn machine. This machine is highly accurate and repeatable, as well as allowing

for in-cycle probing and tool wear measurement. These factors will increase the reliability of the produced fixtures. The Mazak is able to index the workpiece in the lathe spindle and can tilt the tool-head, as well as the usual XYZ movement. The features present on the fixture are designed to be machined using these 5-axis capabilities.

The features of the fixture are designed to allow conventional milling tools to be used in the 5-axis machining centre (as shown in Figure 3–26), instead of requiring custom form tools. The internal pockets have large corner radii to allow the pocket to be machined with sturdy end mills. All features to be machined on the part have lateral clearance (no thin deep pockets) and all threaded holes have a depth of less than 3X their diameter.

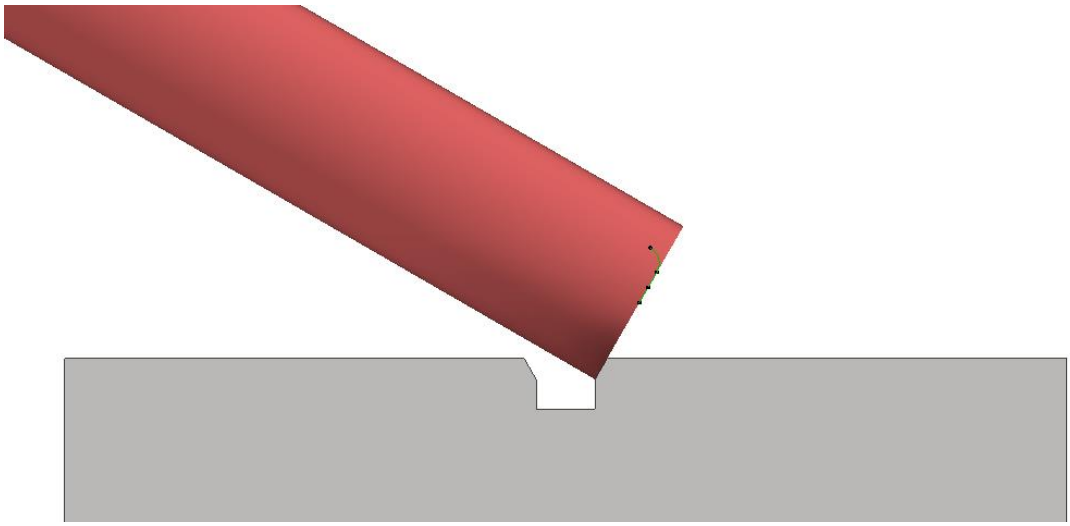


Figure 3–26: A simulation of the machining process for creating the 'V' that the stylus spigot locates into. The included angle is 60°. The dimensions of the 'V' are chosen such that a conventional endmill can be used (Shown in red)

The fixture is used in 2 orientations during the measurement phase. In phase 1 the fixture is oriented with the diamond tip pointing upwards and the spigot towards the X-positive end of the measurement table (Figure 3–27). In phase 2 the fixture is flipped such that the diamond is pointing towards the operator and the spigot is at the X-negative end of the measurement table (Figure 3–27). This change of orientation results in the best visibility during the adjustment and measurement phases of assembly.

The order of operations during the manufacturing process influences the realised tolerances of the fixture itself. For this reason, it is important to consider the manufacturing in relation to the datums of the fixture.

The faces highlighted in green in Figure 3–27 and Figure 3–27 are readily accessible during manufacture, although care must be taken when machining the end of the fixture that locates against the fence in phase 2 (Figure 3–27). This end is still connected to the bar stock and so this feature should be separately machined before the component is parted from the bar (as parting operation is less repeatable and reliable than milling, in terms of finish and dimensional tolerances.)

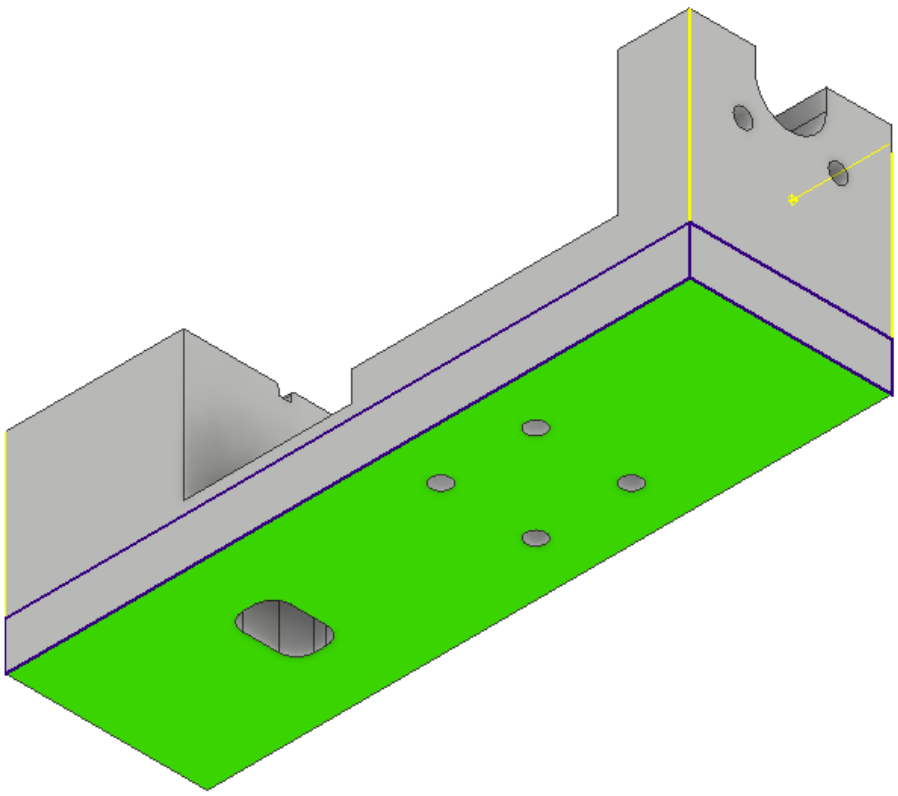


Figure 3–27: A diagram showing the orientation of the fixture during phase 1 of the measurement process. The green face is in contact with the glass and the blue lines represent the contact with the locating fence.

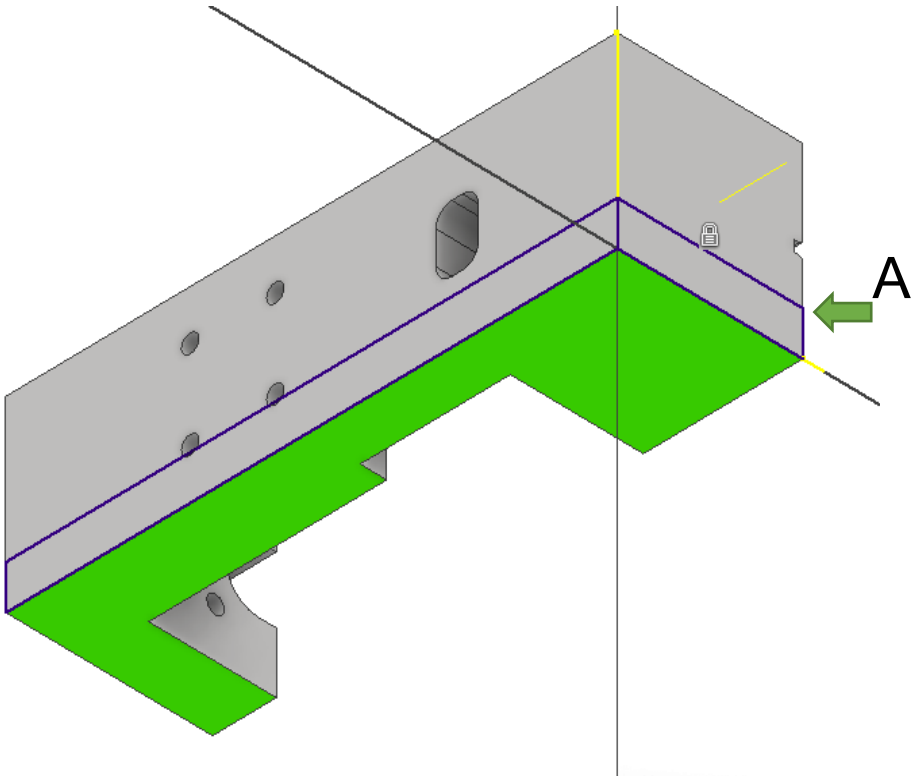


Figure 3–28: A diagram showing the orientation of the fixture during phase 2 of the measurement process. The green face is in contact with the glass and the blue lines represent the contact with the fence. Note: the fixture is flipped lengthwise between Part 1 and Part 2

3.2.14. Measurement Process of the Stylus during Assembly

The stylus is measured in 2 phases during assembly. Phase 1 is concerned with the symmetry of the stylus tip and the main tube (when compared to the position of the spigot). This is done by orienting the stylus with the diamond pointing up. With the stylus in this orientation it is possible to measure the position of the spigot, against which the position of the tube and tip can be measured. The orientation of the assembly fixture for phase 1 of the measurement process is shown in Figure 3–29.

In phase 1 the spigot orientation is measured through the window (labelled A in Figure 3–29). The centreline of the spigot is generated, and the back of the spigot boss is measured to use as a length datum. The floating end of the main tube is then visually inspected using projected lines on the AV300. The tip position is then focussed on and any adjustments in the symmetry and length can be made.

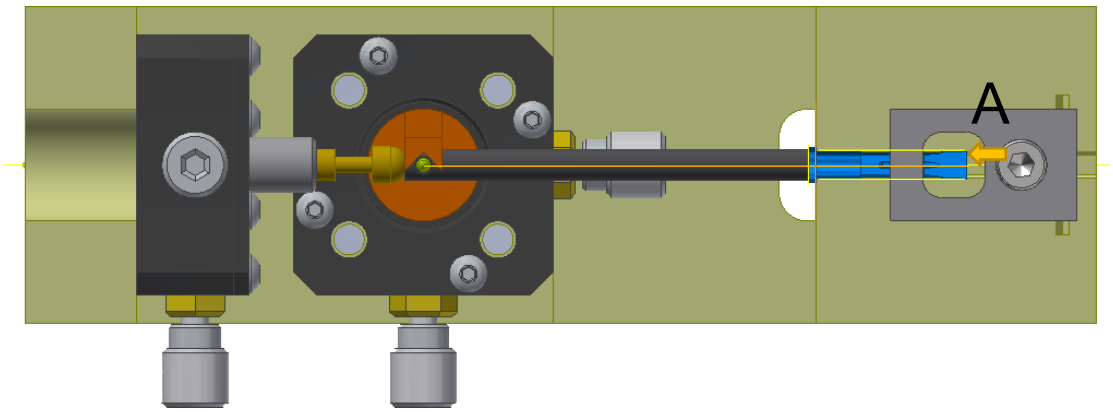


Figure 3–29: A view of the assembly fixture as seen by the AV300. This is the physical arrangement of the fixture during phase 1 of the measurement process.

In phase 2 the fixture is flipped lengthwise and the position of the spigot is measured. The orientation of the assembly fixture for phase 2 of the measurement process is shown in Figure 3–30. A construction line is generated from the spigot and the position of the spigot boss is used as a length datum. The floating end of the spigot is inspected to set the attitude of the main tube. The position of the tip is then measured again in this orientation.

This process is fast enough that the fixture can be readily adjusted within the cure time of the adhesive, and that each feature on the stylus can be visually inspected before the adhesive is no longer adjustable.

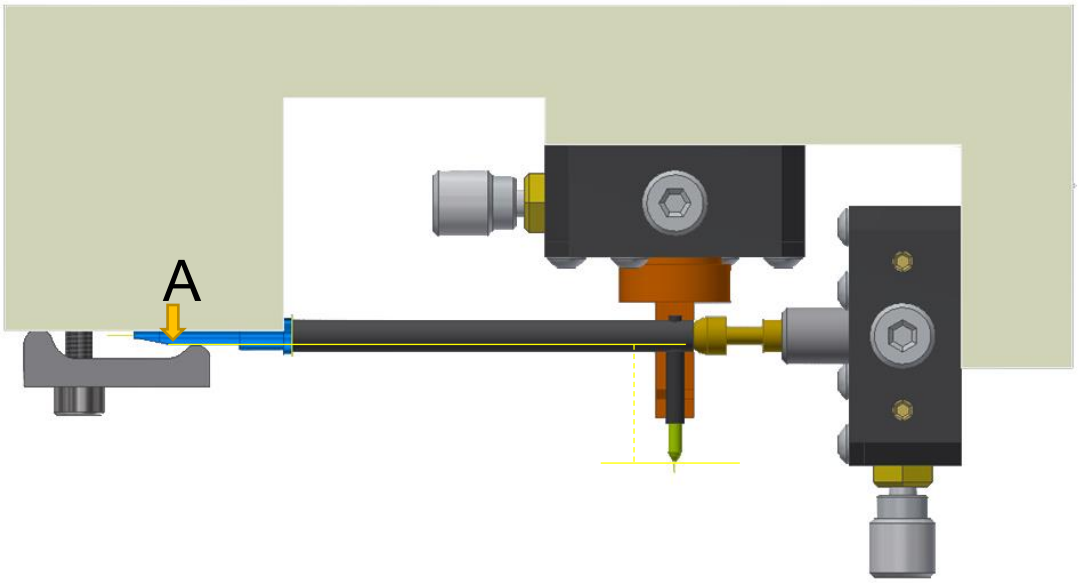


Figure 3–30: A view of the assembly fixture as seen by the AV300. This is the physical arrangement of the fixture during phase 2 of the measurement process.

3.2.15. Costing of the new Fixturing Approach

The fixtures were designed to be manufactured cheaply in-house and to use readily available standard components. Each machined component is designed to be manufactured in a single setup where possible, and the important tolerances can be achieved by using standard cutting tools such as drills or reamers. As a result of these efforts, the cost per fixture is kept low.

The aluminium stock for the fixture base costs ~£90 per metre and machining time on the Mazak is £91 per hour. The expected machining time is ~40 minutes and the stock length used per part is ~130mm. This means 7 parts can be made per metre (allowing for a chucking piece). The time cost is £60 and the material cost £13. The other ancillary components cost very little in terms of material (£10 estimate for the remaining pieces) and the machining time is estimated at 30 minutes.

The stages that are used from ThorLabs are £130 each in single quantities, although the price may improve if several are ordered at once. The stage blanking plates cost £11.

The total cost for a basic 60mm PGI fixture is:

Item	Number	Price per each (£)
ThorLabs XY Stage	2	130
Stage Blanking Plate	2	11
60mm PGI body	1	73
Hold Down	1	20
Recess Actuator	1	35
Pogo Pin	1	4
Printed Components	1	1
Total		410

Table 3-1: A breakdown of the costs of producing a single fixture.

The fixtures for the standard inductive stylus (Figure 3–15) cost upwards of £1500, although they are expected to last longer as they are made from tool steel.

Lifetime Analysis

It is difficult to estimate the lifetime of the new assembly fixture because the fixture has not been used in production and the wear characteristics have not been documented for either the new aluminium fixture or the old tool steel fixture.

The Thorlabs XY stages are flexure-based and therefore have very few individual wearing components. The Pogo pin contacts every stylus which is manufactured and is therefore more likely to wear. Pogo pins are used in the electronics industry for high volume testing processes and are typically specified for 100,000 cycles minimum (Harwin, 2017).

The hold down piece and the recess actuator (with the brass V that the recess piece is retained against) are both made of brass. The wear on these components is likely to be minimal as neither piece is retained with significant force and the motion between the contact surfaces is very small. The nature of the new fixture technique also reduces the effect of a worn component on the geometry of the assembled stylus.

The main body of the new fixture is aluminium and the cylindrical spigot mounts in the V groove in the main body. The contact between the spigot and the V groove is two lines. Given the tight geometric tolerance (therefore little change in the contact location between spigots) and hard anodising of the spigot the most likely location to wear on the fixture is the V groove.

The oldest stylus production fixtures that are currently in service are more than 15 years old. Roughly 500 styli per year are assembled using that fixture giving an estimated 7500 styli assembled in its service life. The cost to replace that fixture today would be ~£1500 (estimated by the head of manufacturing).

The new stylus fixture costs ~£400, leaving £1100 for replacement parts during the 15-year lifetime. This corresponds to a maximum of 7 replacement fixture bodies over the 15-year period, meaning each fixture body would need withstand the assembly of more than 1100 styli to be the same value proposition as the current fixture.

The aluminium body may be modified to reduce the wear in several ways. Most simple, the body could be hard-anodised or otherwise coated. A small design change would allow a standard replaceable steel V to be inserted instead of machining the V groove directly.

Applicability of new Fixture to Production

This type of fixture can be used with a wide range of styli. Of the top (by number) styli, 16 out of 28 can be assembled using these fixtures. Most of the remaining styli already have hard fixtures (Roundness and TalySeries™ styli). These 16 types of styli account for 1400 in number per year.

3.2.16. Starrett AV300 Measurement Repeatability

This section details the research into the performance of the AV300 regarding measurement of stylus features. It compares the current 'hard' fixturing to the adjustable fixture developed in Section 3.2.8. Also detailed is an absolute comparison between the three different AV300's that are in use in the stylus assembly cell.

Two experimental procedures were used to compare the performance of the newly developed fixture against the current stylus fixtures (in terms of measurement performance of the fixtures on the AV300s).

The first experiment used a single AV300 and compares the spread of measurements of stylus length and shank clearance in each fixture, the results are shown in Figure 3–32. The experiment was performed without resetting or re-homing the AV300. The same stylus was used in both fixtures.

1. Simple repeat measurement. Fixture and stylus are not moved between measurements
2. Fixture is re-seated on the AV300 between measurements.
3. Stylus is re-seated in the fixture and the fixture re-seated on the AV300 between measurements.

These three experiments were repeated for each fixture. The first experiment should give the smallest error distribution, representing the baseline repeatability of the AV300. The third experiment most closely simulates the real-world usage of the stylus fixtures.

The testing was expanded to compare the measurement performance of each AV300. The results of this are shown here: Figure 3–31. The stylus remained in the fixture throughout this testing process to eliminate reseating as a source of uncertainty. The stylus length and shank clearance were measured on all three AV300s in turn using both fixtures.

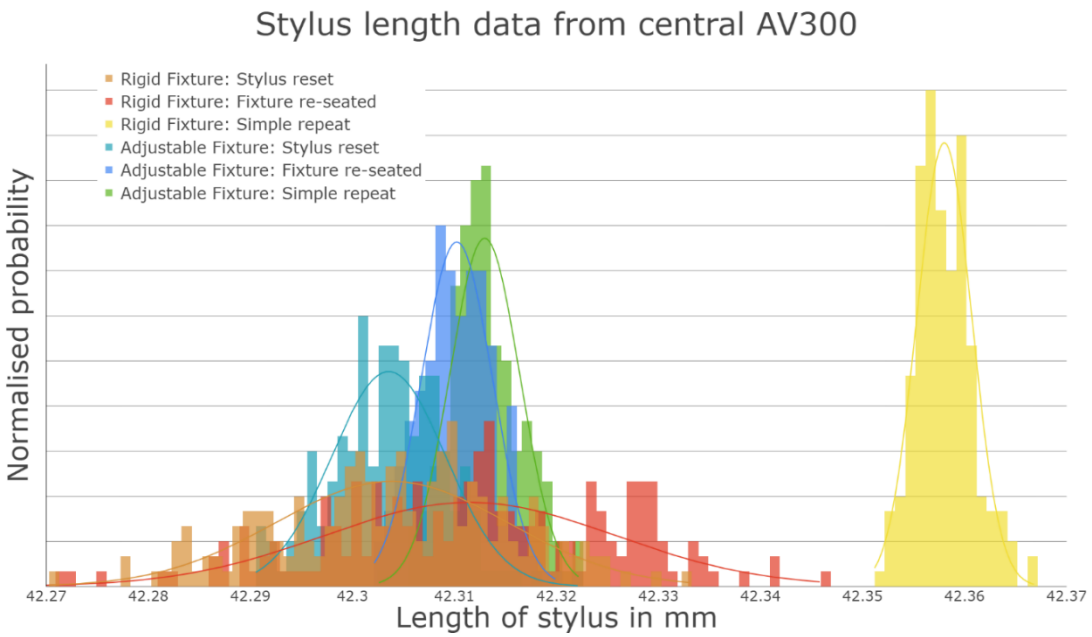


Figure 3–32: A graph of the stylus length as measured using a single AV300.

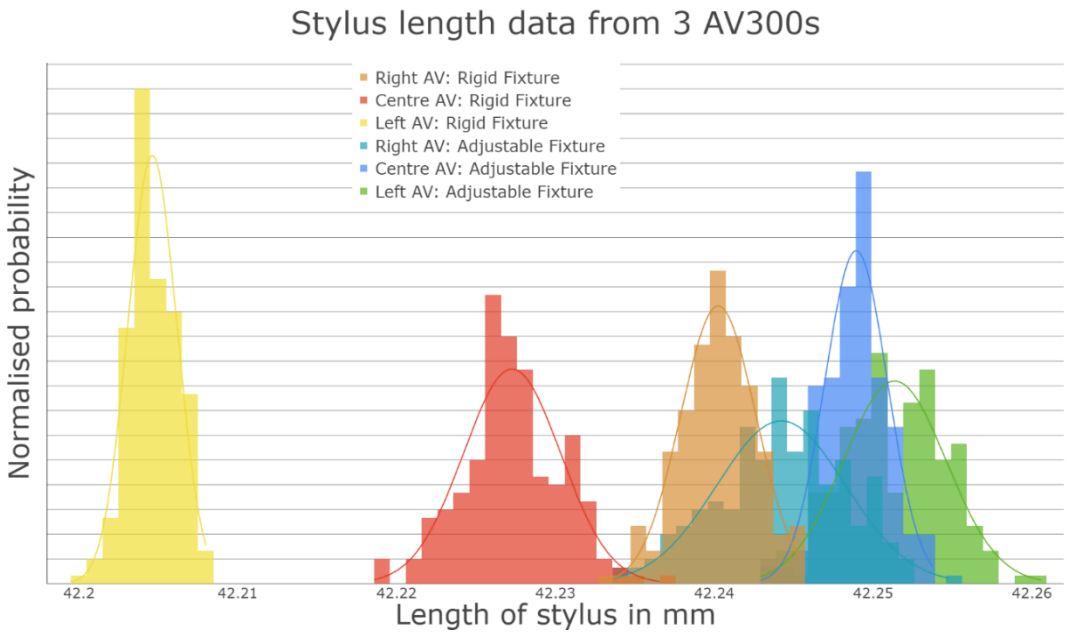


Figure 3–31: A graph of the measured stylus length using the 3 different AV300s.

Figure 3–32 shows the results from the measurement of a stylus on a single AV300. Two fixtures were compared using the same testing regime for each: simple repeat, fixture re-seated and stylus and fixture re-seated. The distributions in red, orange and yellow represent the current rigid fixture and the blue, cyan and green representing the adjustable fixture.

The widths of the distributions associated with the reconfigurable fixture are considerably smaller than that of the rigid fixtures (Red and orange VS. blue, cyan and green), although the thinnest distribution is the ‘simple repeat’ test using the rigid fixturing. The widths of these distributions should be compared to the length and symmetry tolerance as defined on the stylus drawings. Typically, the symmetry must be within ±50 µm of nominal so the measurement repeatability must be several times better than this. The standard deviations of the red and orange data are 13 µm and 12 µm, respectively. By comparison the green, blue and cyan are 3 µm, 3 µm and 5 µm respectively.

The grouping of the distributions is also important. The yellow distribution can be discounted as it is separate from the remaining distributions, likely caused by a setup error. The remaining distributions all overlap and have means which are within a few µm of each other (±5 µm from mean). This shows that there is good agreement between the fixturing techniques.

Figure 3–31 shows the results of measuring a stylus using 3 AV300s (Note: a different stylus was used from the previous test). The rigid and reconfigurable fixtures were used, the stylus was set once in each fixture, before being measured on each AV300. Each measurement set is a simple repeat measurement without the fixture or stylus being adjusted. This should give the best repeatability (thinnest distribution)

In this test the width of each distribution is very similar, although the rigid fixture performs better overall in terms of repeatability on a single machine. This is expected as the measured features on the rigid fixture are considerably larger than the measured features on the spigot. The largest standard deviation in this graph is 4 µm.

In contrast to the graph in Figure 3–32, the distributions in Figure 3–31 do not broadly agree on a stylus length. Using the rigid fixtures, the range of the means of the data is 36 µm, (42.240 mm – 42.204 mm). Using the reconfigurable fixture, the range of the means is 7 µm, (42.251 mm – 42.244

mm). This demonstrates that the method of measurement employed using the reconfigurable fixturing is more robust against the differences between the AV300s.

The difference between the actual measured lengths is problematic and so the stylus was measured using the shadowgraph. The length was determined to be 42.264 mm which does not lie within either measurement range. More investigation is needed to understand the measurement performance and consistency of the AV300s and methods for checking the calibration more regularly than once-per-year should be reviewed.

3.2.17. Conclusions and Recommendations

Fixture Usage

After testing the stylus fixture in a production scenario, several recommendations were made. First, this type of stylus fixture should be offered as an alternative to the presently used fixtures as it has been shown that the new fixtures can produce styli within the appropriate geometric tolerances, without the need for post-cure adjustment. The utilisation of the AV300's is also reduced.

Measuring a stylus by using the spigot as a reference has also been shown to work well. (3.2.16). The different AV300's correlate more closely when measuring with the spigot reference than using the fixture as a reference. (Styli were left in the fixtures between measurements. No adjustments were made between AV300's and each ran the same measurement program.) The work to develop the measurement process for the new fixture also demonstrates that the adoption of the new fixtures simply for the measurement of styli may be beneficial.

Further Fixture Development

The adjustable fixturing approach has proven to be a capable replacement for the current hard fixturing. It has been applied to many of the simple, high volume, styli. It is certainly possible to apply this approach to other styli with varying levels of effort.

Roundness styli are typically simpler to fixture but the utilisation of the AV300 is still high during assembly as the whole cure process happens on the measurement bed. Similar techniques could be applied to roundness styli which would increase the throughput of the assembly process.

It is also possible to fixture the small-bore styli using this technique, as shown in Figure 3–33. This fixture development will result in the same benefits being extended to another range of high-volume styli.

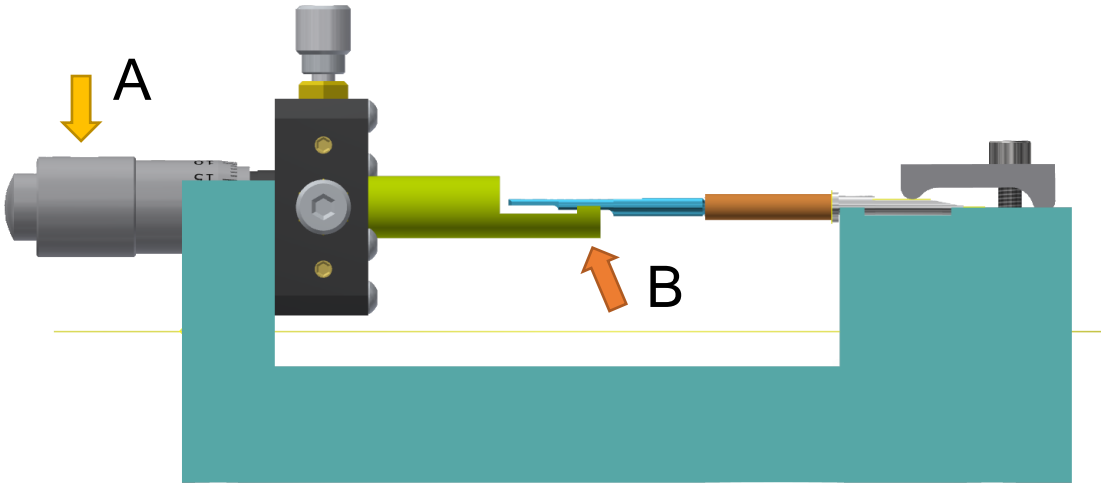


Figure 3–33: A render of a small-bore stylus fixture. The plunger has been replaced with a micrometer head (A) to allow for simple length adjustment. The stylus is held against the adapter (B) by a 3D printed component (not shown).

Standard Components - ThorLabs, Norelem etc.

Using standard components in fixtures greatly reduced the design time and complexity. ThorLabs and Norelem are examples of companies that provide a range of standard fixturing components. ThorLabs is primarily concerned with the optical components, where precise adjustability is paramount. Norelem is a supplier of standard products for fixturing and bespoke automation. These components are typically expensive but invaluable when needing to fixture or automate a low-volume product.

Another benefit of using standard components is the ability to replace one if it is broken or to easily replicate an existing assembly (fixture etc.) The downside of the use of standard components is the potential lack of range and the requirement to interface with the ‘standard’ components (The thread within the SCP05, for example)

3.3. Use of Adhesives in Stylus Assembly

This section highlights the topic of adhesives and their use in stylus production. A discussion of the current assembly procedure is used to highlight areas for future improvement. Recommendations are made on the appropriate surface preparation techniques as well as the recommended type of adhesive.

The stylus geometry is also evaluated, resulting in a suggested change to the spigot geometry and the carbon fibre preparation procedure. The amount of coolant uptake is measured and determined to be minimal; the use of IPA solvent is recommended to help with the surface preparation process.

3.3.1. Motivation

The current approach to adhesives in stylus assembly was developed more than 25 years ago. The range of available adhesives has increased greatly in that time, so it is worth exploring the choice of adhesive for stylus assembly. The joint geometry contributes to the quality of stylus and the ease of assembly, so is explored in this research.

The application method affects the resultant stylus in terms of both static and dynamic properties. The current mixing and application techniques vary greatly between people which makes monitoring the adhesive performance more difficult. The assembly technique will also be discussed with reference to the recently developed stylus fixturing systems.

The available surface preparation techniques are compared, and recommendations are made with a view to increasing the consistency of the adhesive bond.

3.3.2. Overview of Current Production Approach

This section will detail, in roughly chronological order, the stages of stylus assembly where the adhesive is an important component. It gives a robust idea of why this report was written and what problems must be dealt with.

A Typical Recess Stylus

Most of the currently produced surface styli have a geometry similar to the stylus detailed in Figure 3–34, although several other configurations are well represented (Small Bore etc.)



Figure 3–34: A drawing of a typical recess stylus. The main components are labelled, and the typical assembly tolerances are shown in the bottom left.

3.3.3. Production of Stylus Components

The styli at Taylor Hobson are made from pultruded carbon fibre tube and rod. The bundle of carbon fibres is dipped in epoxy and drawn through a die, with or without a mandrill, producing tube or rod, respectively. This rough tube is then cut to length and ground between centres at PMP grinding. These shorter lengths of ground tubes are delivered in batches to the stylus cell at TH. The communication between TH and PMP is very limited but the products that we receive are dimensionally consistent. The geometrical consistency of the carbon fibre is very good (ID and OD are both within 20µm between parts and batches).

These ground tubes are then cut to rough length using the Studer cylindrical grinder or the manual multi-axis grinder, using a cut-off wheel. The blank is then loaded into the HAAS 5-axis mill where it is milled to final length and the other features (cross holes etc.) are machined. This is done using the high-speed spindle and flood coolant. These finished components are then dried using an air blast then stored in the KANBAN system. Components are

retrieved from the KANBAN system and then etched on the laser engraving system.

The metal components used in the stylus are machined in-house and then finished externally. The spigots, for example, are machined from 6061-T6 aluminium that is subsequently anodised and dyed Blue. The anodising process does not have the required tolerance for coating thickness (±25 µm anodising tolerance vs ±5 µm required dimensional tolerance) and so the rejection ratio is high (Estimated at 30%). The dimensional tolerance ±5 µm on the spigot is imposed because the fit of the stylus into the retaining tube needs to be very consistent between different gauges. Nickel plating technology can reliably achieve the required coating tolerances (±2 µm), this technology is currently being investigated by Manufacturing Engineering to replace the anodising on the spigots. (Loto, 2016)

3.3.4. Assembly Process of the Stylus

The stylus components are collected in a plastic bag, along with the order details and the box, and then sent to the assembly bay. The actual assembly procedure may vary between configurations but for a standard recess stylus is:

1. Spigot is secured in a fixture and the main tube is levelled using a stack of gauge blocks.
2. Adhesive (Devcon 5-min) is dispensed onto paper and mixed by hand, using a small screwdriver or similar.
3. Diamond tip is fixed into the recess piece by hand, a small piece of wire is used to spread the adhesive.
4. Main tube is adhered onto the spigot, whose pitch is constrained by the gauge block stack. The yaw and roll of the main tube is set by hand, without verification.
5. Once the sub-assemblies are cured (~5 minutes), fresh adhesive is applied to the remaining joint (between the main tube and the recess piece) and the pieces are assembled

6. The assembly is placed under an AV300 and the measurement program is run.
    a. This is often repeated several times, with small changes being made by hand during each iteration
7. If the stylus is out of the specified tolerance, then the appropriate joint is heated with a soldering iron until it can be moved back into place.
8. The stylus tip is photographed, and the certificate is printed.

This assembly process is highly manual and relies on post cure correction methods. Adhesive manufacturers (including Devcon) recommend against the use of heat on a cured joint as it significantly weakens the adhesive bond. (Adams, Coppendale, Mallick, & Al-Hamdan, 1992). There is also little time given to a visual inspection after assembly

After the adhesive has cured the stylus is temporarily packaged and the labels affixed to the box. The stylus is checked on an appropriate instrument and given a certificate of calibration. This stylus is then held in stock (common styli only) or shipped immediately to a customer.

Some common styli already have non-adjustable fixtures for use in assembly. This approach can lead to difficulty when stylus components vary within their tolerance range. The result of this is more rework and likely less reliability, although the time taken is reduced.

3.3.5. Stylus Component Tolerancing

As shown in Figure 3–35, the spigot and the stylus tube are bonded together with a lap joint. If the tolerances of the drawing are the same as the manufacturing tolerances, then the gap between the spigot OD and the tube ID can be anywhere from -100µm to +120µm. This could result in an interference fit, which may explain some of the cracking that has happened after assembly.

After measuring several samples of the ground PCF tube using a set of digital callipers, it was concluded that the ID and OD are extremely consistent: typically measuring 2.52 mm and 3.35 mm respectively. The spigot diameter is consistently 2.44-2.45mm

The different types of adhesive need different bond thicknesses. This thickness will influence the strength of the stylus joint. There is not a requirement for maximum possible joint strength, but the space reserved for the adhesive should allow consistent application and bonding. The thickness will also influence the curing time, thicker layers taking longer to cure (in most cases).

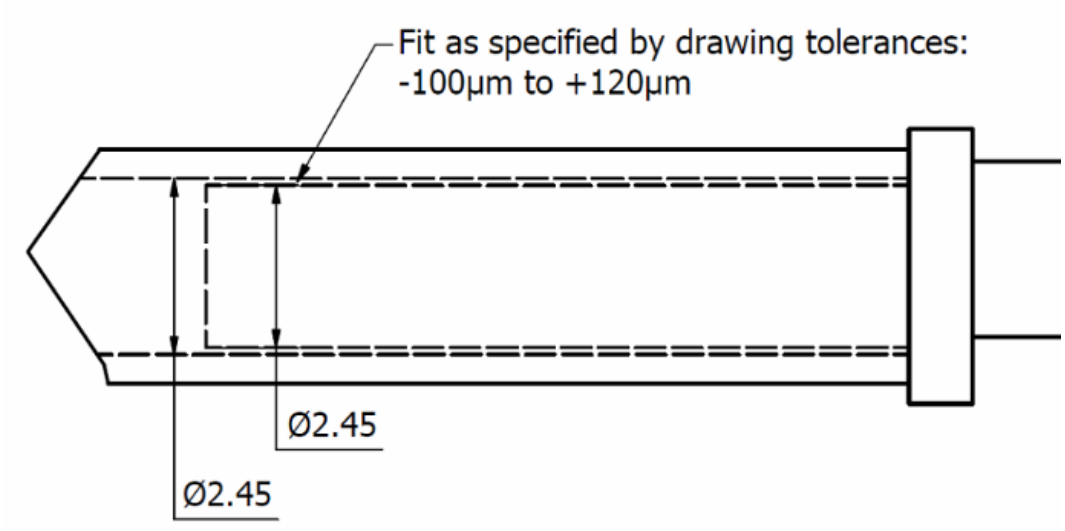


Figure 3–35: Exploded view of the mating surface between a spigot and a stylus tube. Dimensions and tolerances are shown.

3.3.6. Areas for Evaluation

Manufacturing

The manufacturing of the components may affect the bond strength of the adhesive joints in a couple of ways: the coolant may allow oil ingress into the PCF tube, and the drilling process may cause delamination of the inner layers in the PCF tube.

The flood coolant is an emulsified oil and the pultruded carbon fibre is a porous medium. (A material safety data sheet for the cutting oil is available here: (Lubricants, 2003)). The porosity of the carbon fibre suggests that the oil will penetrate the surface of the PCF tube and then may interfere with the bonding process of the 5-minute epoxy. The absorption of oil has been tested and has been shown to be very limited. The absorption is also limited by the 10-minute average duration of typical stylus component machining. The absorbed coolant layer is ~400nm thick. See page 116 for more information.

The drilling process also causes fibre tear out and layer delamination in the PCF tube. Delamination of the fibres can be seen in Figure 3–36 (Labelled A). This may weaken the tube and cause failure during, or after, the assembly process. The adhesive joint may also be affected if the wetted surface includes the delaminated fibres.

The curing process (in 2-part epoxies) is usually accompanied by a small shrinkage and the ‘movement’ of the adhesive. In the longer curing adhesives this can be as little as 0.5% but Devcon 5-minute epoxy can exhibit up to 7% shrinkage by volume (Devcon, 2012). The shrinking of the adhesive may cause cracks to propagate or further delaminate the fibres from the tube. This increase in internal stresses may also result in longer term failure, especially if the stylus is subjected to temperature or humidity fluctuations (soldering iron or plane flight).

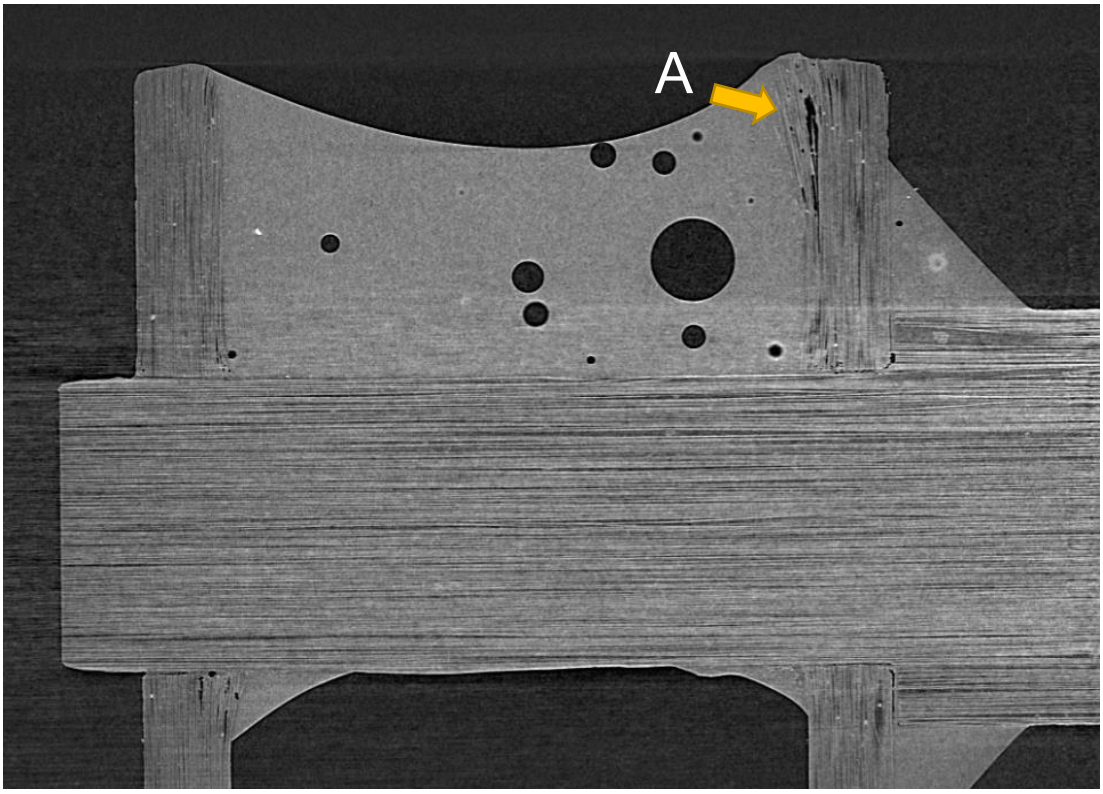


Figure 3–36: An X-Ray CT Image slice of a stylus adhesive joint. Tear-out of the PCF tube can be observed at the top right of the image. Delamination is also visible above and below the drilled cross-hole

Cleaning and Handling

At present, there is little to no process time given over to cleaning or surface preparation. Typically, an airgun is used to blow off the coolant after the machining of the stylus components. The surface of the PCF tubes is generally also dried using blue roll, although only on the OD of a component.

Cleaning Efficacy - Coolant Absorption Experiment

Carbon fibre components were tested for coolant uptake in a typical machining cycle. 2 standard stylus tubes were used as a proof of concept. The tubes were 87.1 x 4.78 x 3.00 mm. These tubes are ground between centres which results in an increased surface porosity vs unground stock.

One tube was used a reference (labelled tube 1), while the other was subjected to 10 minutes in the path of the coolant nozzles in the HAAS machining centre (tube 2)

Both PCF tubes were dried externally using blue-roll and blown out using an airgun. This procedure is currently the only cleaning that is routinely performed on machines PCF stylus components. After drying the components were weighed using a top-pan balance. After the stylus tubes were weighed, they were then cleaned with IPA solvent and weighed again.

Before exposure to coolant, tube 2 had a mass of 1.399 g ± 0.29 mg. After exposure to the coolant, drying with Blue roll, and an air gun, tube 2 had a mass of 1.402 g ±0.38 mg. After cleaning with IPA solvent, the mass of tube 2 was reduced to 1.400 g ± 0.25 mg; this is a reduction of 2.35 mg ±0.45 mg. This demonstrates the efficacy of cleaning the PCF tube with IPA solvent post machining.

The uncertainty (standard deviation) of the mass of Tube 2 is 0.29mg before testing and 0.25mg after cleaning. The difference in the average masses is 0.8mg which is ~3 standard deviations and therefore the amount of absorbed oil from the start to the finish is measurable using this equipment. The raw data and the corresponding statistical properties is in Appendix 8.2.

0.8mg cutting oil remained in the surface after tube 2 was cleaned with IPA. (~5% emulsified oil, 95% distilled water) is close to 0.8µl. This is spread over the surface area of the tube (inside and out)

$$\pi \times 4.78mm \times 87.1mm = 1308mm^2 \ (OD)$$
$$\pi \times 3mm \times 87.1mm = 821mm^2 \ (ID)$$
$$0.8\mu l \div (1308 + 821)mm^2 = 376nm$$

The oil is absorbed (assuming an even distribution of oil throughout the surface of the PCF) to a depth of 376 nm.

The surface roughness of the PCF tube is likely to be relevant in determining the adhesive bond strength. The measured Pt of the PCF tube is ~3.5 µm which still gives a good mechanical bonding surface even with the oil ingress.

Component Tolerance Checking

Stylus components, once manufactured, are not dimensionally checked before they are cached. There are two failure modes of the stylus components that could be mitigated at this stage, before the components are sent to engraving and assembly.

There are a wide range of ‘standard’ cross-hole sizes which can be used (0.8, 0.9, 1.2, 1.5, 1.6 and 2.0mm). These holes are drilled on the HAAS which leads to a good repeatability on the diameter. Although the tolerances that are achieved are very good, there is no checking of the hole diameter after machining. This can result in the recess piece being an interference fit with the main tube. This increases the difficulty of assembly greatly and often leads to chipping or splitting of the PCF tube, as shown in Figure 3–37.



Figure 3–37: Microscope image of a chipped cross-hole in a PCF tube. The carbon fibres run length-ways along the main tube, therefore the fibres do not contribute to the strength of the end section, resulting in shear failure.

The position of the cross hole (and therefore effective length) was measured and the variance was determined to be ±0.06 mm. This is close to the specified tolerance of ±0.05 mm that appears on most drawings. The actual position of the tip relative to the spigot is specified to a tolerance of 50 µm. The observed variability in the displacement of the cross-hole from the end of the stylus is already more than the specified tolerance. The routine measurement of this dimension may require a specially adapted measurement tool. It may instead be possible to use the AV300 although this would increase the demand for a machine which already has a high utilisation.

While these two failure modes (specifically, the variance in the diameter of the cross hole and the variance in the position of the cross hole.) are not directly related to the adhesive, the adhesives are currently used to help mitigate these problems. The dimensional discrepancies that occur naturally during stylus assembly are currently mitigated by re-heating of the adhesive. Solving the dimensional issues would reduce the re-work needed and therefore result in a better adhesive joint.

Adhesive Mixing

The hand mixing of the adhesive may be replaced by the nozzle mixing to increase the consistency of the adhesive before application, and to reduce the amount of air bubbles introduced during hand mixing. The adhesive should be applied to both components, following a specific technique, to increase the surface wetting and therefore improve the consistency of the bond.

Air bubbles have a detrimental effect on the bonding characteristics of the adhesive (Permabond, 2016). The trapped air will expand at a different rate to the epoxy, potentially causing stress cracking of the adhesive or the PCF tube, especially given the practice of reheating already cured adhesive joints. Trapped air bubbles will also reduce the effective surface area during wetting and make proper application of the adhesive more difficult.

3.3.7. Alternatives to Current Adhesive

Types of Adhesive

There are several other types of adhesive that would be suitable for the assembly of styli. The application method, bond strength, suggested adhesive layer thickness, cure time, shrinkage, hardness, colour, viscosity and range of mixing accessories all contribute to the choice of the correct adhesive for this application.

The range of appropriate adhesives fall into 3 categories: 1-part cyanoacrylate, 2-part epoxy and 2-part methacrylate. Each of these adhesives have certain advantages according to; the materials, surface preparation methods, fixturing techniques and type of stylus (completely fixtured and simple vs special and totally handmade). There is also the potential to use a 'potting compound'. This is traditionally used for sealing electronics or weatherproofing other small components, but it can also be used as an effective adhesive.

The current 2-part epoxy (Devcon 5-Minute (Devcon, 2012)) is well suited in terms of predictable bond strength, cure time and visual consistency, although it does require mixing and does weaken significantly when heated. The high viscosity of the adhesive reduces the need for fixturing of complex styli, as the adhesive can be used to tack components in place, although this does increase the requirements for clean bonding surfaces

The cyanoacrylate adhesives can have a broad range of physical properties but, because of their curing method, tend to cure very quickly. Cure speed is also inversely proportional to bond layer thickness which makes these adhesives difficult to apply pre-assembly. The cure time can be as little as 5 seconds with typical joints like the diamond tip to recess piece. (Henkel - Loctite, 2008)

Comparison of Adhesives

Referring to Table 3-2, there are several adhesives which would fulfil our requirements.

DP410, produced by 3M, has a working time of 12 minutes which would allow proper adjustment of the stylus fixture under the AV300, it also has a very high tensile strength and a low shrinkage. The fixture time (time before the stylus can be removed from the fixture) is not specified; the author estimates this to be 20-40 minutes from similar adhesives. (3M, 2013)

Penloc VT Plus (produced by Eurobond) specifically mentions applications with composites and fibre reinforced plastics in the datasheet, although the cure time may be too short for use with the reconfigurable fixture approach. It also states that the bond is less susceptible to surface contamination. (Eurobond, 2015)

Fastfix (produced by Eurobond) is similar to DP410, but cures in clear as opposed to off-white. The bond strength may be reduced than for the DP410 but still a considerable increase on the present Devcon 5-Minute Epoxy. (Eurobond, 2016)

The Devcon 5-Minute epoxy is currently used in stylus assembly and is noted for having poor bonding to carbon fibre and other composites. The viscosity of this adhesive (10 Pascal Seconds) can be used to compare to the other adhesives. (Devcon, 2012)

Hysol 934NA is a 2-part potting compound. At 20°C the cure time is 45 minutes, which is excessive when compared with other solutions. The adhesives cure time is reduced when the temperature is raised. This is advantageous as a several styli may be assembled using the same batch of adhesive, thereby reducing the number of used mixing nozzles and wasted adhesive. The styli can be cured in an oven which can take less than 10 minutes. This adhesive would not be appropriate for use on the un-fixtured styli. (Loctite - Hysol)

DP8810NS is a 2-part epoxy that is blue-green in colour. The advantage of this adhesive is the rapid structural cure time of 20-25 minutes. This adhesive contains 0.25mm diameter glass beads to help control the bond thickness. This will cause problems when used on the spigot as the glass beads may cause binding. The glass beads should have little effect on the adhesive joint between the main tube and the recess piece. (3M - Scotch Weld, 2019)

Loctite 4307, a UV cure cyanoacrylate, was used in an experiment with automated dispensing systems. This adhesive is similar to other cyanoacrylates but will also cure rapidly when exposed to UV. This adhesive would be suitable for stylus assembly if adhesive would not cure as a response to the small gap. The short cure time did not allow sufficient movement of the stylus components to make this approach useful for stylus construction. (Loctite, 2013)

Recommendation for adhesives

The author suggests that DP410, Penloc VT Plus, Fastfix and Hysol 934NA are tested in stylus assembly, using the new reconfigurable fixtures to ensure repeatability. Feedback from this testing should allow the appropriate choice of adhesive to be made. The Hysol 934NA has a much greater cure time than the other adhesives which may require a process change to be properly utilised. (Oven curing once the stylus has been verified under the AV300)

Name	Type	Working Time (min)	Fixture Time (min)	Cure Time (min)	Colour	Viscosity	Shear Strength (23°C)	Hardness	Notes
3M DP410	2-part Epoxy	12	-	180	Off-White	70 Pa.s	38MPa	-	High strength, Low shrinkage
Penloc GTi (Eurobond)	2-part Methacrylate	5-7	60	240-360	Transparent	5-6 Pa.s	7.67MPa (Tensile)	45-55 Shore D	
Penloc VT Plus (Eurobond)	2-part Methacrylate	4-6	10-12	-	Off-White	50-100 Pa.s	24-31MPa (Tensile)	125000 Psi	High temperature, minimal surface prep, resistant to hydrocarbons, specific applications with fibre reinforced plastic
Fastfix (Eurobond)	2-part Epoxy	5-10	30	120	Clear	20 Pa.s	20MPa 47Mpa (Tensile)	80 Shore D	Free Flowing, Chemical resistant.
Loctite 499	1-part Cyanoacrylate	10-60 s	Up to 60 s	24 h	Clear	18 – 40 Pa.s	5-15MPa	-	Far too fast for practical fixturing.

Name	Type	Working Time (min)	Fixture Time (min)	Cure Time (min)	Colour	Viscosity	Shear Strength (23°C)	Hardness	Notes
Devcon 5-Minute Loctite 408	2-part Epoxy 1-part Cyanoacrylate	3-6 10-60 s	10-15 30-45 s	12 h 72 h	Clear or Light Amber Clear	10 Pa.s 0.01 Pa.s	13MPa (Tensile)	85 Shore D -	Poor adhesive properties with PCF Cure speed and low viscosity make this impractical for stylus assembly
Penloc GTi (Eurobond)	2-part Methacrylate	5-7	60	240-360	Transparent	5-6 Pa.s	13.7MPa	45-55 Shore D	
Loctite 5-Minute	2-part Epoxy	5-7	20	24 h	Translucent Yellow	-	22MPa		Sand-able and low shrinkage. Limited long-term temp. range

Name	Type	Working Time (min)	Fixture Time (min)	Cure Time (min)	Colour	Viscosity	Shear Strength (23°C)	Hardness	Notes
Hysol 934NA (EA 934NA)	2-part potting compound	45	8 h	7 d	Grey	80 Pa.s	26MPa	-	The cure time greatly decreased by heating. 60°C = 20 min.
3M DP8810NS	2-part Acrylic	8-12	12-20	20-25	Blue-Green	45 Pa.s	25MPa	-	Quick-Curing, high strength. Contains glass beads
Loctite 4307	1-part Cyanoacrylate	2-30 s	10-30 s	24 h	Clear	0.6 – 1.2 Pa.s	6MPa	82 Shore D	UV Cure in addition to Moisture cure. Linear shrinkage is 15%

Table 3-2: A table detailing the range of available adhesives that would be suitable for stylus assembly. Several of the pertinent properties are included here to allow for simple comparison.

3.3.8. Alternative Adhesive Joint Geometry

The strength of an adhesive bond varies depending on many factors. The bond thickness (spigot OD to tube ID, for example) may reduce the bond strength by up to 90% if it is too thin. The adhesive joint being heated after cure will also harm the adhesives properties and may cause further stresses and leads to cracking of the PCF in the field. (Huntsman, 2007)

The spigot and other stylus components can be modified to give the optimal bond thickness and allow more repeatable assembly. A new spigot design is shown in Figure 3–38. One issue with the current assembly process is the adhesive being pushed ahead of the spigot as the stylus is assembled. This results in a limited amount of adhesive being present between the spigot and the main tube, likely leading to stylus failure when the spigot becomes loose in the main tube.

A simple way to improve the adhesive retention around the spigot is to put an undercut on the stem. This retains the locating diameters at either end, while provides a space for the adhesive to reside. It also stops the adhesive from being pushed ahead of the spigot.

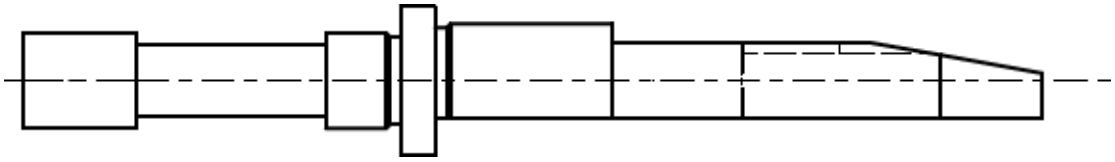


Figure 3–38: Diagram showing a potential new geometry for the stylus spigot. The depth of the groove in the stylus would be in the range 50-200µm.

The undercuts around the central boss allow the main tube to avoid being chamfered and reduce the likelihood of repeatability problems in the gauge.

3.3.9. Surface Preparation

The condition of the bonding surfaces is a very important factor in achieving a strong joint between components. The surfaces under question should be free from dust, grease, oil, degreaser and any other contaminant that would adversely affect the bonding process. (Huntsman, 2007)

As mentioned previously, the cutting oil used in the manufacturing process will ingress into the surface, potentially resulting an inconsistent bond. The solution to ingress of oil into the PCF is to use a mild solvent to clean the inside of the tube. See page 116 for further details.

The other type of adhesive problem faced in this scenario is the presence of carbon fibre dust or carbon fibres that have dislocated from the PCF tube. The adhesive may bond preferentially to the dust and the free fibres because of surface tension, and so the effective bonding area is greatly reduced.

The recommendation for minimising with these two issues is to:

1. Dry the components using Blue roll (Already happening)
2. Blast the ID of the tubes with air (Already happening)
3. Use a pipe-cleaner (or similar), saturated with mild solvent (IPA etc.) to clean the ID of the tube (and other adhesive surfaces)
4. Allow component to dry for 15 minutes in a well-ventilated area

The pipe cleaners should be changed after every few styli to avoid cross-contamination. Appropriate health and safety measures must be observed when using solvents. IPA is currently in use elsewhere, so the adoption of this solvent into a cleaning process should not be difficult.

3.3.10. Alternative Dispensing Technologies

The adhesives discussed can be dispensed in several different ways, each with a typical range of applications. The method that is currently used to apply the adhesives is by hand, with a wire or miniature screwdriver, after hand mixing. The 2-part epoxies typically have dispensing hardware which the tube of epoxy seats into, with the mixing nozzle fixed to the front.

This in-built mixing mechanism is beneficial in terms of part-to-part and batch-to-batch consistency, although there is a requirement to dispose of the used mixing nozzles and the amount of used adhesive increases. These costs are small when compared with the cost of a single returned stylus, although coordination between people and proper planning of the assembly stages would help minimise waste.

Low-viscosity adhesives can be dispensed from a typical medical syringe (thumb activated) or a pressure-time system (which increases the consistency of the dispensed volume). (GDM - Doug Dixon, 2005). The small size of the stylus means that the entry points for the adhesive are typically small (~1 mm). This limits the size of a dispensing nozzle and therefore the viscosity of the adhesive used.

There are some considerably more advanced dispensing technologies, although without proper fixturing techniques it is unlikely to give consistent results. (GEBE2, 2019)

The dispensing of adhesive by hand allows for a broad range of styli to be assembled. Given the large range of manufactured styli, this flexibility is important. The hand-assembly process can be made more consistent by using assembly tools with markings to aid the person assembling the stylus.

3.3.11. Conclusions

Surface Preparation

1. Formalise the cleaning process (Air blast, Blue roll etc.)
2. Implement pipe-cleaners and IPA solvent.
3. Designate ventilated space for drying.

Adhesive

1. Use a different adhesive (3M DP410 or Hysol934NA)
2. Purchase and implement dispensing equipment
3. Create guides on the amount of adhesive to be used as well as the application method etc.
4. Find application accessories (toothpicks with marks etc.)

Joint Geometry

1. Add undercut to standard stylus spigot
2. Remove the redundant stylus cross-holes
3. Create guidelines on the appropriate tolerances for stylus components
4. Create/buy measuring tools for stylus component measurement.

Misc.

1. Evaluate the current set of stylus assembly fixtures to investigate reports of hard fixture producing styli that are 150um out of specification
2. Reduce the reliance on the soldering iron by utilising proper fixturing techniques.

4. Materials Investigation - Stylus Dynamics

The investigations into stylus tolerancing and the subsequent work on the development of the new assembly fixture allowed for the consistent production of in-specification styli. The investigation into the proper stylus measurement technique allowed the produced styli to be dimensionally evaluated with a high degree of certainty on the required geometric tolerances. The work on the range of adhesives and application methods allows the proper selection of adhesive for a wide range or stylus materials. The new fixturing approach allows a wide range of stylus components to be properly fixtured for assembly.

In short, styli were produced that meet or exceeding the current tolerance requirements, without the need for rework, in a variety of materials beyond the normal pultruded carbon fibre. The ability to produce the wide range of styli allowed the testing of the dynamic properties of the stylus to begin.

4.1. Stylus Dynamics Testing System

This section is concerned with stylus-based contact metrology, specifically focussing on the measurement of surface finish as opposed to the form of a component. This focus helps to determine the limits of the dynamic environment that the stylus-gauge system will experience in normal operation. The spatial wavelengths and amplitudes of a surface translate into a dynamic frequency and amplitude that is experienced by the stylus during measurement.

The dynamics of the measurement system as a whole will have an effect on the measured profile. The measurement instruments themselves have had considerable development in improving their dynamic performance; a task made easier by the lack of physical limit to their mass and shape. There are requirements placed on the measurement stylus that cannot, for practical reasons, be solved by simply increasing the mass.

To develop a proper understanding of the dynamic performance of the instrument as a whole, a good understanding of each component or sub-system is required. The overall aim of the work performed into investigating stylus dynamics is concerned with the dynamic performance of the stylus-gauge sub-system. The lessons learned from this investigation will aid the development of metrological systems in the future.

4.1.1. Background

To give an accurate representation of a surface, a stylus must remain in contact with the surface throughout the measurement process. The stylus leaves the surface when the dynamic forces overcome the preloaded surface force (typically 800 µN for a standard PGI gauge[6]). The dynamic force exerted by a surface is proportional to the traverse speed, surface amplitude and surface type (ground, turned etc.). (As discussed in Section 2.9)

For a small amplitude sinusoid (200 nm peak-to-valley) it is possible to measure without leaving the surface at surface frequencies of up to 600 Hz. This frequency is the upper limit of the tracking capability of the PGI gauge and stylus system, for a purely sinusoidal surface.

If a stylus leaves a surface, it will not only give an overestimate of the height of the surface immediately following the flight event but will also disrupt measurement after it contacts the surface again by continuing to resonate. This will change surface measurement parameters, particularly affecting skew and peak parameters.

4.1.2. Motivation

The design of surface styli is often governed by the geometry to be measured. Most typical surface measurement applications (such as the measurement of lenses and bearings) dictate the geometrical requirements for the stylus and the tip and stylus. The radius of the lenses or the depth of the bearing race

[6] The PGI gauge is a surface finish and form measurement gauge developed by Taylor Hobson in ~1992. PGI stands for Phase Grating Interferometer, referring to the technology used to achieve the high resolution and large measurement range.

dictate the minimum shank clearance of the stylus, for instance. The range of contact angles between the tip and the measured surface dictate the tip angle and the requirement for surface texture measurement dictates the tip radius.

Complex measurement tasks, such as measuring the surface finish of a fuel injection port within an engine, necessitate a considerably more complex stylus geometry (shown in Figure 4–1). These styli are usually bought in small quantities. Each of these styli should not only be geometrically identical (within tolerances) but also should perform identically in measurement as well. The measurement performance of a complex stylus is governed by many factors including; geometric design, material properties, adhesive properties, distribution of adhesive, specific diamond geometry, effective tip force, gauge spring constant, moment of inertia and other factors.

Traditionally the performance of surface instruments could be improved by increasing the stiffness of the stylus and the rigidity of the overall system, with little understanding of the changing dynamics on the measurement. (Jiang, Scott, Whitehouse, & Blunt, 2007)

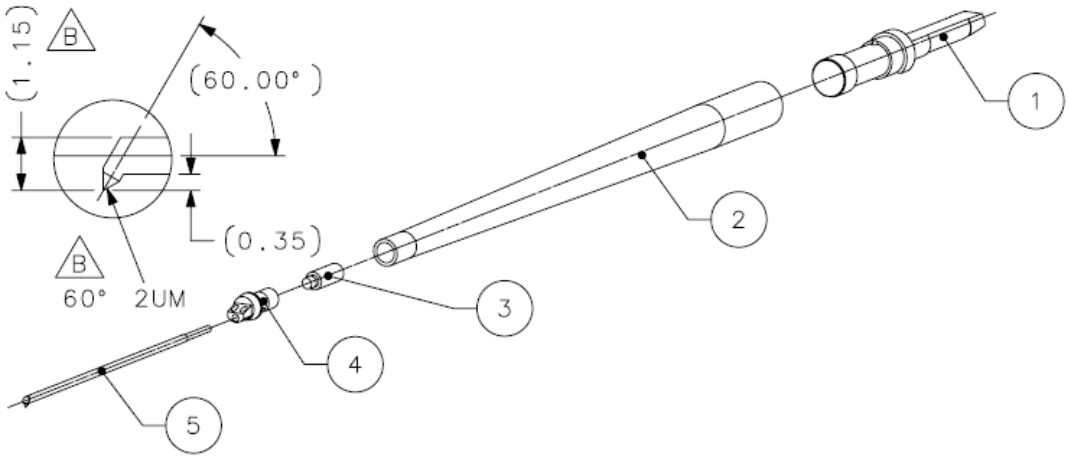


Figure 4–1: A diagram of a complex stylus assembly. The stylus used to measure the surface finish of deep internal holes.

The measurement ability of the AV300 optical-CMM and the tight tolerances of the manufactured components help to ensure that produced styli meet the required geometric tolerances. Extensive in-house testing of a stylus before

shipping would help to guarantee measurement performance, but this is costly and could result in damage to the stylus. Styli are currently tested by measuring a calibration artefact twice; the two traces must agree for the stylus to pass.

An alternative method of investigating the potential measurement performance of the stylus, before application, is the use of a dynamics testing system. The response of the stylus to specific excitation across the range of typical measurement surface frequencies and amplitudes is, conceptually at least, a good analogue for actual measurement. The excitation is generated using an actuator to act as a 'configurable surface'. The amplitude and frequencies can be adjusted to match a typical measurement application. A standard dynamic testing procedure allows for comparison between styli, potentially even between different styli of different geometric designs.

If the response of a stylus to a given frequency is repeatable, and the dynamic environment that the stylus will experience is directly relates to a surface measurement speed and amplitude, then a practical limit on the traverse speed of a measurement can be specified.

4.1.3. Requirements of a Stylus Dynamics Measurement System

Rigid Physical Mounting System

To allow the proper separation of the dynamic properties of the gauge-stylus system and the remaining measurement loop, the dynamics measurement system is required to be highly rigid. If the resulting design is more susceptible to noise than a typical measurement system (Like the PGI Martix Figure 4–4), then it will be difficult to distinguish the dynamic properties of the dynamics measurement system from the dynamics of the gauge-stylus system. The developed system must allow for styli with different shank clearances and lengths to be measured easily. The system needs to be readily manufactured, preferably with as few custom or non-standard components as possible. The vibration of the final system should be measurable and, if possible, the modes

of vibration that most affect the measurement system should be attenuated by design.

The measurement loop of the system must be kept as small as possible to give the best rigidity for a given mass/material/geometry. In current roundness and surface instruments it has become normal to use a single vertical column for the Z axis (cantilever design, Figure 4–2). The column design has a large Z-axis travel, which is an advantage for industrial applications but is unnecessary for testing the dynamic response of styli. A design which reduces the cantilever effect and therefore increases overall stiffness is a gantry configuration. This is typically 2 columns joined with a beam, from which the measurement system is suspended (gantry design). A comparison between these design approaches can be seen in Figure 4–2.

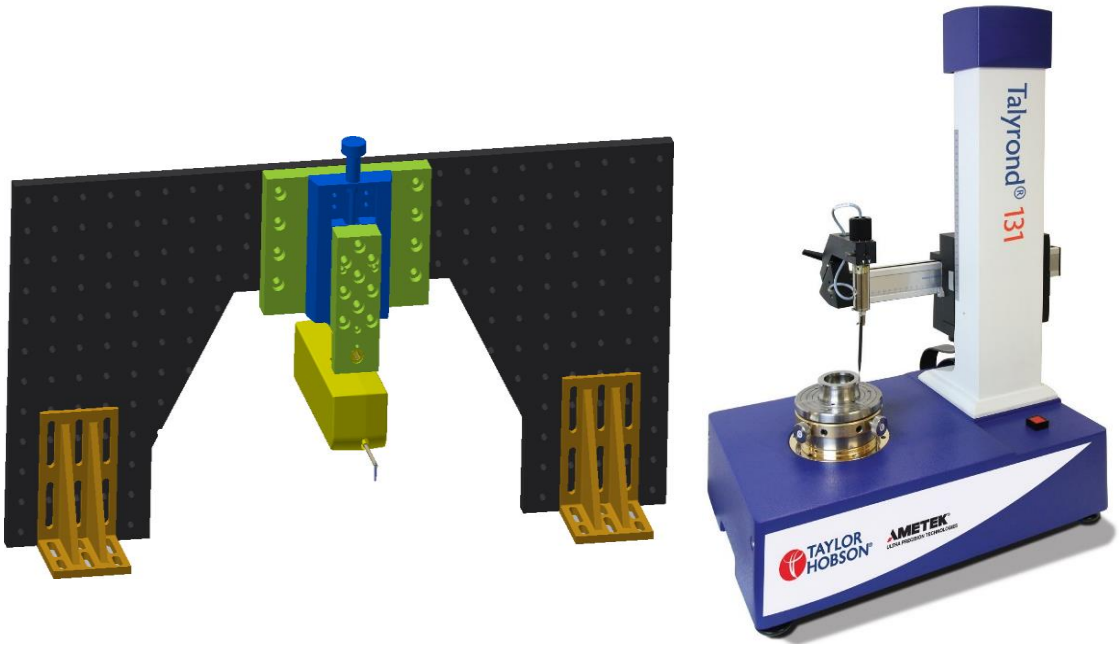


Figure 4–2: An example of a gantry based measurement system (Left), and a cantilever based measurement system (Right)

The weight of the PGI gauge should be supported in the same way as on commercial instruments to replicate measurement conditions. The PGI gauge itself is typically mounted underneath the traverse unit on a horizontal cylinder. The mounting point is chosen such that the centre of gravity of the gauge rests below the centre of the travelling portion of the traverse unit on a surface

instrument. The centre of mass of the PGI gauge in the vibration measurement system should be as close to the gantry as possible. The mounting interface for a PGI gauge is shown in Figure 4–3

The mounting system must allow for a range of stylus lengths and shank clearances to be measured. Some degree of adjustability in the vertical position of the gauge or the actuator is preferable as it allows the stylus position to be set to the centre of the gauge range before testing.

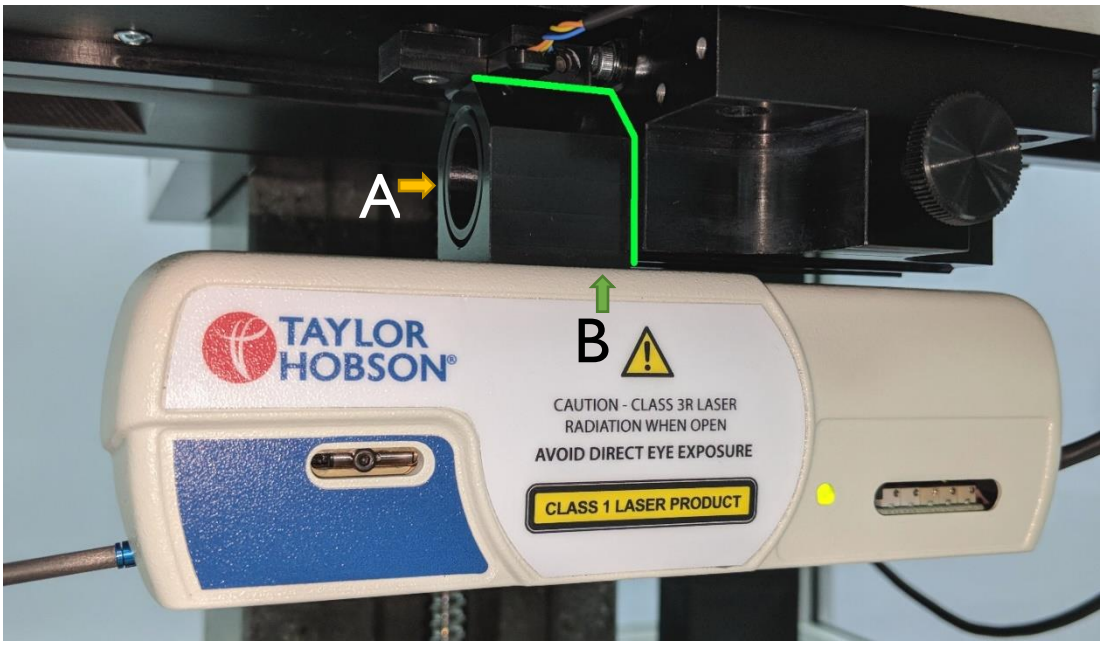


Figure 4–3: A photograph of a PGI gauge mounted under a traverse unit on a surface measurement system. The mounting boss of the PGI gauge (Labelled A) is highlighted by the green border for clarity. The tube on which the PGI gauge is hung is labelled B. The gauge is retained by a pair of grub screws.

Stylus Excitation Method – Actuator

An actuator is used on the dynamics measurement system to simulate a controllable dynamic environment into the stylus, at its tip. This actuator must have a controllable amplitude and frequency and perform reliably across the entire range of expected frequencies. The amplitude of the vibration is limited to a few microns as most engineered surfaces have amplitudes within that range. The frequency of the vibration must be variable from DC to greater than 1000 Hz. The minimum vibration amplitude should be on the order of single nanometres, so the dynamic range of the signal generation must be

~4000:1 or 12-bits of resolution. Digital filtering would also exclude features at this frequency during the analysis phase.

Stylus Excitation Method – Electronics

The electronics system that is required to drive the actuator on the stylus dynamics measurement system has a few key requirements; governed by performance requirements or by usability and control features. The electronics system should:

- Be controllable and triggerable from the control PC
- Have a large dynamic range (at least 12-bit resolution) for the signal generation
- Have the ability to sweep a signal frequency over time, at frequencies from DC to >1 kHz
- Preferably have an Application Programming Interface (API) that would allow direct programmatic control, as opposed to being limited to GUI control.
- Preferably be capable of generating a variety of waveform shapes

Any electronics system that fulfils these requirements can be used to generate the drive signal for an amplifier or actuator controller.

System Control Software

The software used in the stylus dynamics measurement system must be able to coordinate the data capture, configuration and triggering of the actuator, exporting of the gathered data, analysis of the exported data, presentation of the analysis to the user and the archival of the data and the analysis for suture traceability.

This list of tasks is non-trivial and requires the cooperation and synchronisation of several distinct sub-systems. Further, the analysis process may take a considerable amount of time and processing power.

The analysis and the control software are not required to be implemented as a single monolithic solution, indeed, the development and debugging of the analysis process would be simpler if it can be run as a standalone solution. There are several ‘industry standard’ analysis packages which can be used in this application, such as MATLAB, R or Python with Numpy (Python, 2019), (Mathworks, 2019), (R-Project, 2019)

Further Considerations

The development of the dynamics measurement system should use as few custom manufactured components as possible. This will reduce the time to realise the design and likely reduce the cost of the system. Given the ubiquity and mechanical performance of standard optical components, such as those produced by (Newport or ThorLabs), the need for custom components is minimal.

The developed system should be reconfigurable and allow for future work which can be performed on the dynamics measurement system. The use of a linear stage to emulate a traverse unit would allow the gantry system to investigate stylus-gauge system performance in a way that is more like the typical measurement application. If possible, the dynamics measurement system should allow further addons to support future endeavours.

4.2. Initial Experimentation to Determine Parameter values for the Stylus Dynamics Measurement System

With the conceptual requirements set out clearly in Section 4.1.3, some practical investigation was required to determine the specifics and to investigate the current performance of a standard measurement system. This work was undertaken for the MSc project that preceded this thesis and so will only be summarised here. The relevant section of the MSc work is included in Appendix 8.1 for context.

4.2.1. Conclusions of the Initial Testing

The initial testing process was more than a simple proof of concept. The initial testing gave guidance on the requirements of the actuation method, the software requirements gave a snapshot of the dynamic performance that can be expected from the PGI Matrix system. Further, a baseline for the influence of environmental noise on a surface measurement system was established.

Further Requirements for Actuation Methods

The initial testing process demonstrated the feasibility of using an actuator to excite a stylus, however the fidelity and range of the chosen actuator would need to be considerably higher than either of the solutions that were tested. The poor static noise performance of the loudspeaker, along with the high coupling of the speaker cone and the air conditioning, makes it unsuitable for use in the realised system.

The Maplin piezo on the other hand is, by comparison to the speaker, a highly mechanically stable device. The range of deflection on the Maplin piezo is too small (of the order of 1 µm) to be suitable for the dynamics measurement system. There are several commercially available piezo actuation systems, such as the P-841 Preloaded Piezo Actuators from PI (Physik Instrumente,

2019). These are already utilised in other metrological and high precision instrumentation and have several advantages over the Maplin piezo system.

Requirements of the Software and Electronics System

The generation of an excitation signal was initially handled by a Philips PM5138 signal generator, which was incapable of interfacing to a host PC. This signal generator is highly capable in terms of output frequency and voltage level. The signal generator could generate a frequency sweep over a well-defined time-period, with a variety of wave shapes (Sine, Square, Triangle etc.). The DC offset was controllable, although the instrument was incapable outputting only DC, without AC.

The usage of this signal generator was exclusively manual and required configuration each time the sweep parameters were changed. The manual triggering of the signal generator did not allow for simple automatic trimming (trimming by timing) of the collected data. The PM5138 would also default to outputting a fixed frequency before the sweep was triggered, this would exaggerate the starting frequency when the data was analysed.

Despite the obvious capabilities of the PM5138, the inability to automate the data collection and to alter the signal properties means that another solution was necessary. PC control of waveform generators and similar 'benchtop' instrumentation has been available for some time. Instead the Analog Discovery 2 was identified. The Analog Discovery 2 is a 'PC based oscilloscope' that is already used within the electronics engineering department at Taylor Hobson.

The Analog Discovery 2 has an accompanying fully featured Python library, allowing the PC to configure the waveform at runtime. Ultra (the instrument control software) has a basic scripting interface to coordinate complex measurements and can call external programs when running a script. This interoperability allowed Ultra to control the measurement instrument, call Python programs to control the Analog Discovery 2, record data from the measurement gauge and export that data for analysis.

4.3. Realised Stylus Dynamics Measurement System

This section details the realisation of a stylus dynamics measurement system, based on the work conducted in section 4.2.

The choice of actuator, control electronics, software implementation, analysis methodology and the physical development of the gantry system are detailed here.

Finite element analysis was used during the design of the gantry to evaluate design choices with a view to supressing the modes of vibration that would be most disruptive to stylus measurement. The gantry design is optimised to minimise the deflection of the stylus tip, caused by environmental vibrations.

4.3.1. Realised Actuator

As was briefly mentioned in Section 4.2.1, a piezoelectric actuator stack fulfils the requirements set out in Section 4.1.3.

Physik Instrumente (Physik Instrumente, 2019) sell a wide range of positioning systems, typically using piezoelectric drives to provide the basic motion. In the simplest realisation of a piezoelectric motion system, a piezo crystal can be made to expand or contract proportionally with an applied voltage. Unconstrained piezo stacks are typically noisy and are likely to creep over time. To combat these issues, preloaded piezo stacks have become available; these are normal piezo stacks that have been compressed inside their metal casing. This preloading means that the elastic expansion of the casing helps to stabilise the piezo movement and increases the linearity of the response.

The range of motion of a basic PI piezo actuator is 15 µm, which is sufficient for the requirements of this application; however, the power dissipation requirement for the actuation of a sine wave with 5 µm amplitude and a frequency of 1000 Hz means that a physically larger actuator is necessary

The chosen model of actuator was the P-841.4 (Physik Instrumente, 2019). This has a range of 60 µm and an integrated strain gauge sensor. This sensor, coupled with the closed loop controller, is capable of positioning at the sub-nm level. The closed loop control increases the ability of the piezo to follow the input signal but is limited to ~250 Hz maximum frequency. This feature may be beneficial for checking the calibration of the PGI gauge but is not likely to be useful during actual measurement.

An advantage of the PI piezo actuator over the Maplin piezo is the greatly increased actuation range. The Maplin piezo has a limited range of motion well below 1 µm whereas the PI P-841.40 has an actuation range of 60 µm and is capable of sinusoidal motion at frequencies of up to 2 kHz with an amplitude of 1 µm. This makes the PI system a logical choice for the actuation system. A picture of the piezo actuator in place on the measurement system is show in Figure 4–8.

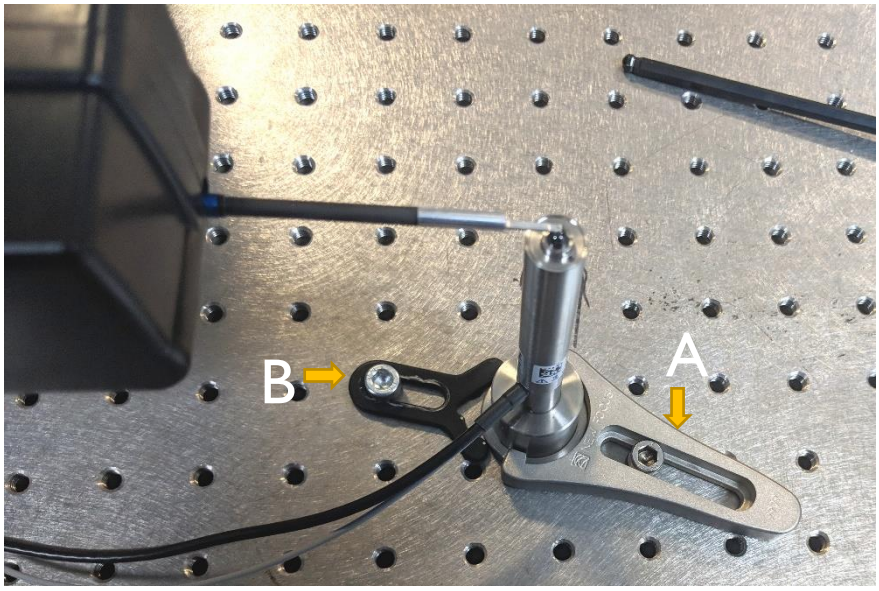


Figure 4–8: A picture of the piezo actuator, the clamp (A) and the repositioning-aide (B)

4.3.2. Realised Gantry Electronics

Digilent Inc. have created a highly configurable USB oscilloscope, logic analyser, signal generator and variable power supply. This device is controlled using a PC application and crucially has an Application Programming Interface (API). The Analog Discovery 2 has a 14-bit arbitrary waveform generator, capable of outputting 100 Msample/sec and a bandwidth of at least 12 MHz. Practically this means that the generation of high-fidelity sine waves, at frequencies up to 2 kHz, is easily attainable. (Digilent Inc., 2019)

Comparison between the Analog Discovery 2 and the PM5138

The Analog Discovery 2 does not have the same effective resolution as the Philips PM5138, although this was somewhat mitigated by using a series resistor between the AD2 output and the piezo amplifier. This means more of the 14-bit resolution can be used to generate the signal (higher ENOB), although the largest amplitude of the actuation is limited to a few microns.

When the frequency sweep was initiated on the PM5138, the signal transitions smoothly from the base frequency to the ramping frequency. This makes the automated clipping of the collected data quite difficult. The AD2, by comparison, outputs a fixed 0V until the frequency sweep is initiated, at which point the signal will rise smoothly and follow the appropriate waveform.

The AD2 is configured before each waveform is generated. This causes a period of undefined output from the waveform generator which must be accounted for during testing. The sharp transition between the DC and the waveform means the automated clipping of the data is straightforward.

The AD2 has many more capabilities that would be useful in future projects. There is a second waveform generator available, 16 digital IO that can be used to trigger analogue or digital events and 2 14-bit oscilloscope channels which can be used for collecting analogue data.

Drive Electronics - Piezo Amplifier

The amplifier that is supplied with the P-841.4 piezo stack is able to drive the piezo through its entire range of motion (60 µm) with a very high precision (~1 nm for open loop applications). In open loop mode, the amplifier effectively acts as a simple buffer with an amplification of 10 times. The input range for the amplifier is -2 V to +12 V which corresponds to a driving voltage of -20 V to 120 V. This applied voltage moves the piezo through its full range. (Physik Instrumente, 2019)

When using a time-varying signal, such as a sinewave, the behaviour of the piezo is a little different. The piezo and amplifier can handle step-changes in the drive voltage as the slew rate is very high. The small signal bandwidth is at least 15kHz for this configuration of drive electronics and piezo actuator.

Electronics System – Attenuator Overview

The signal that is generated by the AD2 is attenuated by a 10 kΩ and 50 Ohm resistor voltage divider. This is currently implemented using an in-line shielded project box to house the 10 kΩ resistor and a 50 Ω terminator on the input terminal of the piezo amplifier. This reduces the travel range of the piezo while increasing the vertical resolution. A schematic of the attenuator can be seen in Figure 4–9.

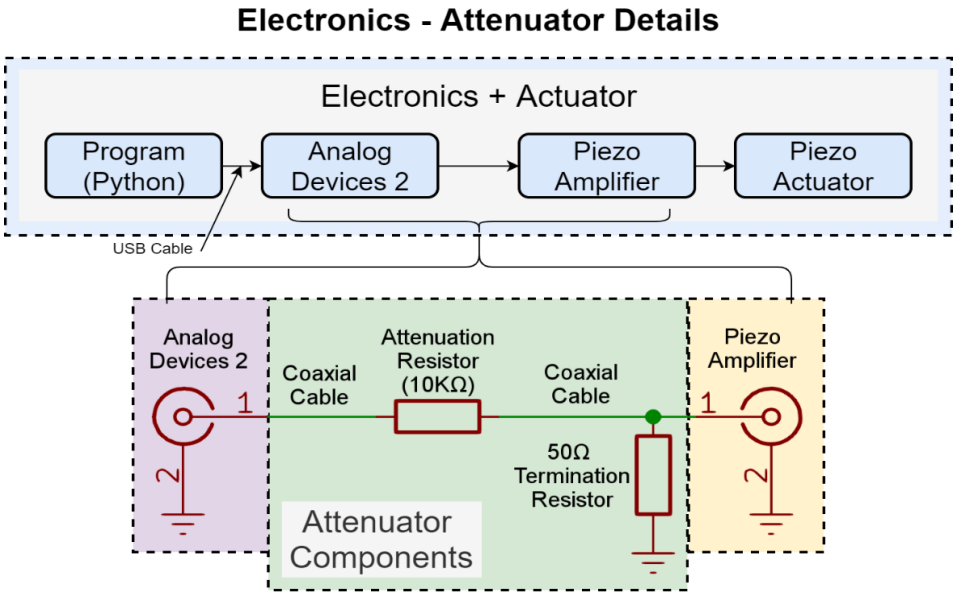


Figure 4–9: A schematic of the electronics system used in the gantry, with specific focus on the attenuation components. The voltage amplitude is defined at the point labelled A.

4.3.3. Realised Gantry System

The realised system that has been developed has a few basic components: A PGI gauge, an actuator, drive and control electronics, a physical support system and a PC for system control. The interaction of these components is shown schematically in Figure 4–10.

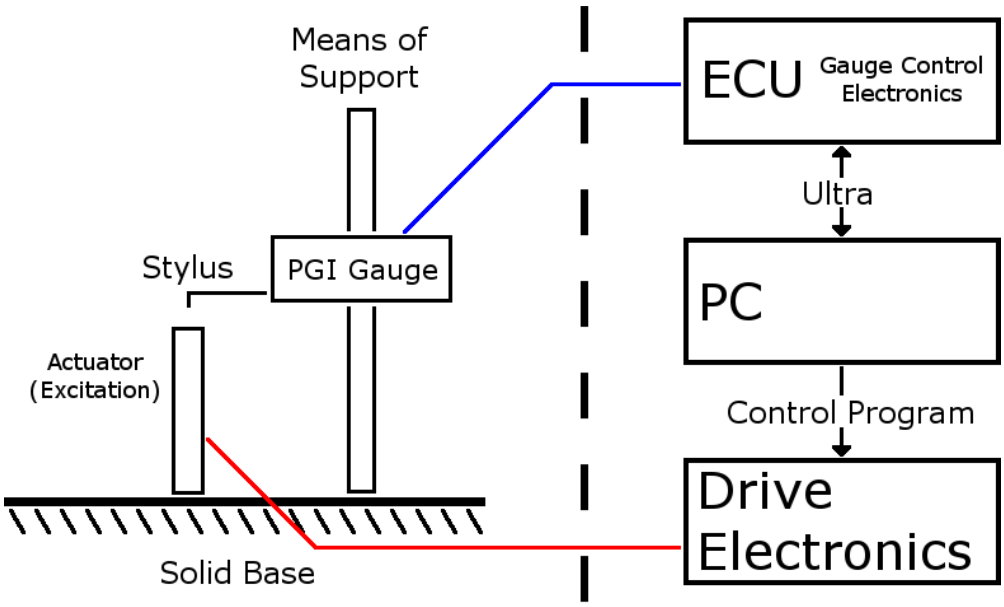


Figure 4–10: A schematic of the vibration measurement system, showing the physical implementation on the left of the dashed line and the electronic/control implementation on the right of the line.

The implementation allows the straightforward separation of the measurement system from the excitation system, thereby simplifying testing and development. This system can be easily adapted for other uses such as gauge testing. A standard PC and ECU were used to reduce the complexity of the implementation.

The chosen actuator is a PI P-841.4 preloaded piezo actuator with integrated strain gauge. This is driven by the E-501.00 Piezo controller with strain gauge feedback. The signal is generated by an Analog Discovery 2 (AD2); driven by the PC, using a custom piece of control software. The signal generated by the AD2 is attenuated by a 10k Ohm series resistor into a 50 Ohm terminator on the Piezo controller.

The physical support is provided by a gantry system, as shown in Figure 4–11. This design used the fewest custom components and gave the highest stiffness in the finite element modelling. This design allowed the PGI gauge to be mounted in the same way as on the commercial instruments. The gauge is adjustable in the Z-axis to allow for a range of shank clearances.

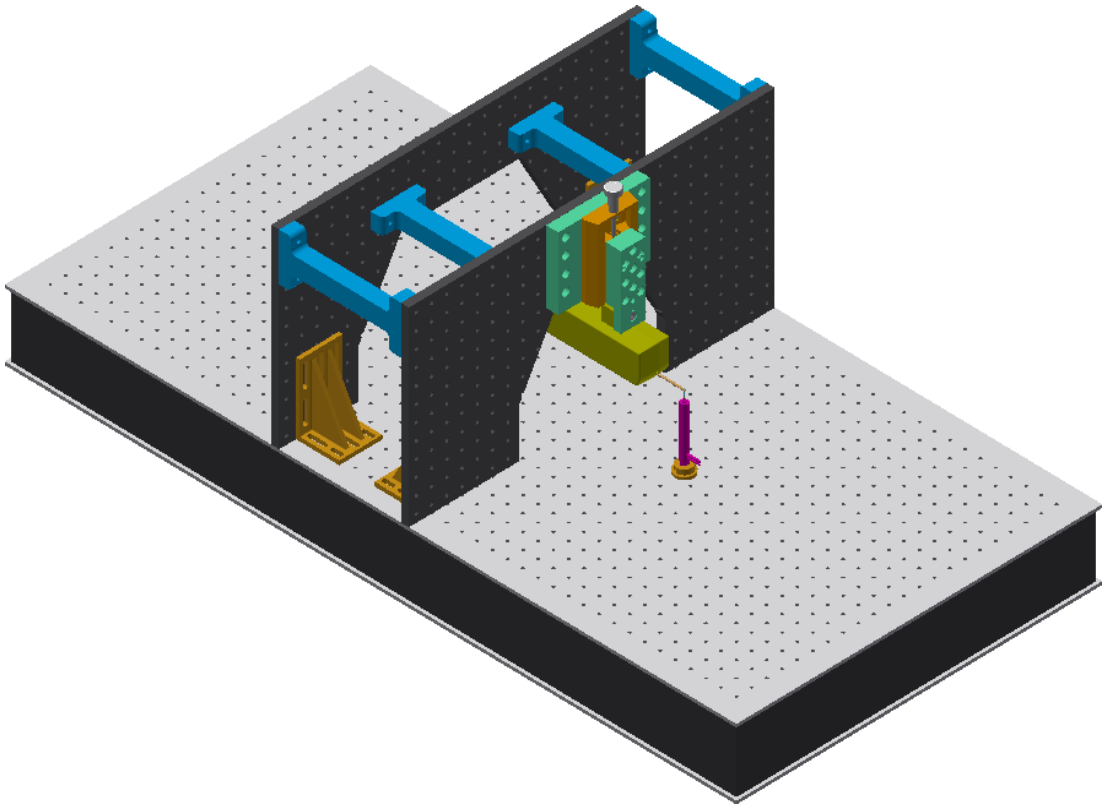


Figure 4–11: A render of the chosen gantry solution. It shows the honeycomb optical breadboard, on which the gantry is mounted. The gantry is mounted using 4 precision angle plates, with the z-axis stage mounted on the front. The PGI gauge (yellow) is mounted below this stage, above the piezo actuator (pink)

Development of the Gantry system using FEA

The optimal configuration of the plates and other elements within the gantry is not obvious and so several configurations were explored, in different stages, using finite element analysis (FEA). The number of potential configurations is large, so the different options were tested separately, and the most effective ideas were tested in the final product.

Optical breadboards are available in a variety of sizes, typically multiples of 150 mm, so initially 2 sizes of breadboard were tested using FEA (300 mm x 600 mm and 450 mm x 600 mm). The 600 mm x 300 mm breadboard can be seen in Figure 4–12. The increased height (450 mm vs 300 mm) did marginally reduce the amplitude of the primary and secondary modes within the plate, but the frequency of the primary mode was reduced significantly. This is problematic in this context as the pass band for the typical PGI gauge is up to 400 Hz so any vibration in this frequency range is directly transmitted to the measured results.

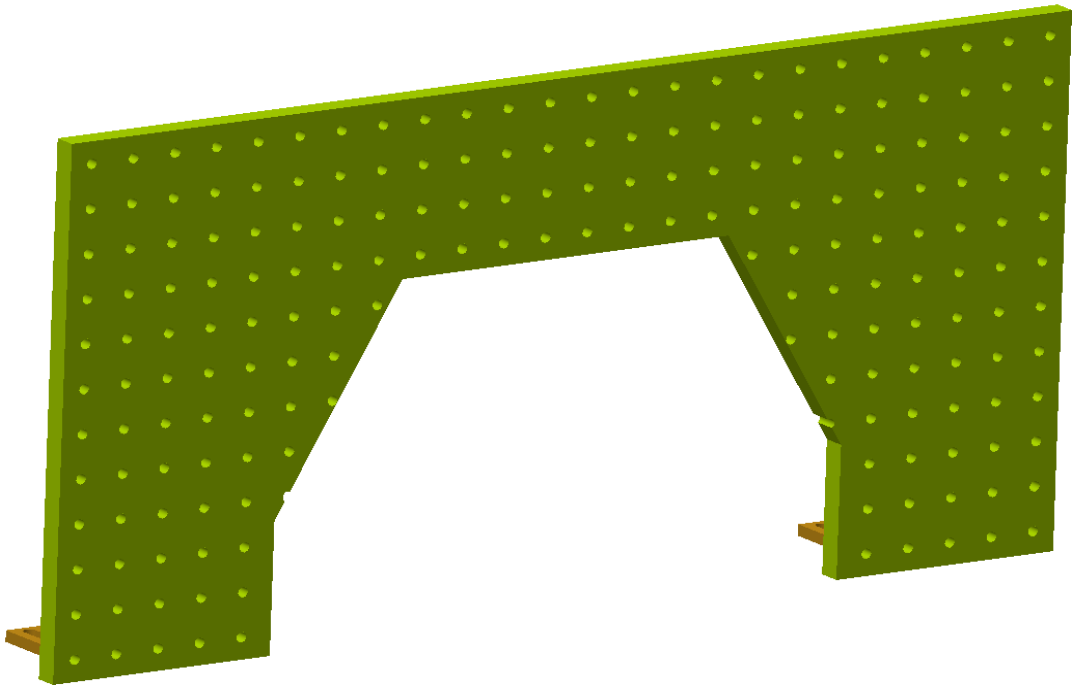


Figure 4–12: A render of the single plate configuration of the gantry. This plate is 300mm x 600mm with a cut-out to accommodate the PGI gauge which will hang below it

Several solutions were considered in an effort to reduce the amplitude of the primary vibration mode (Figure 4–17). The first of these was a set of bracing strips that would be secured across the front and back of the plate, shown in Figure 4–13. The fundamental geometry of the plate was very similar and so the primary mode shape was not significantly different. A small reduction in vibration amplitude was achieved with this method but the cost of the bracing was comparable to the optomechanical components themselves, so a different configuration was tested.

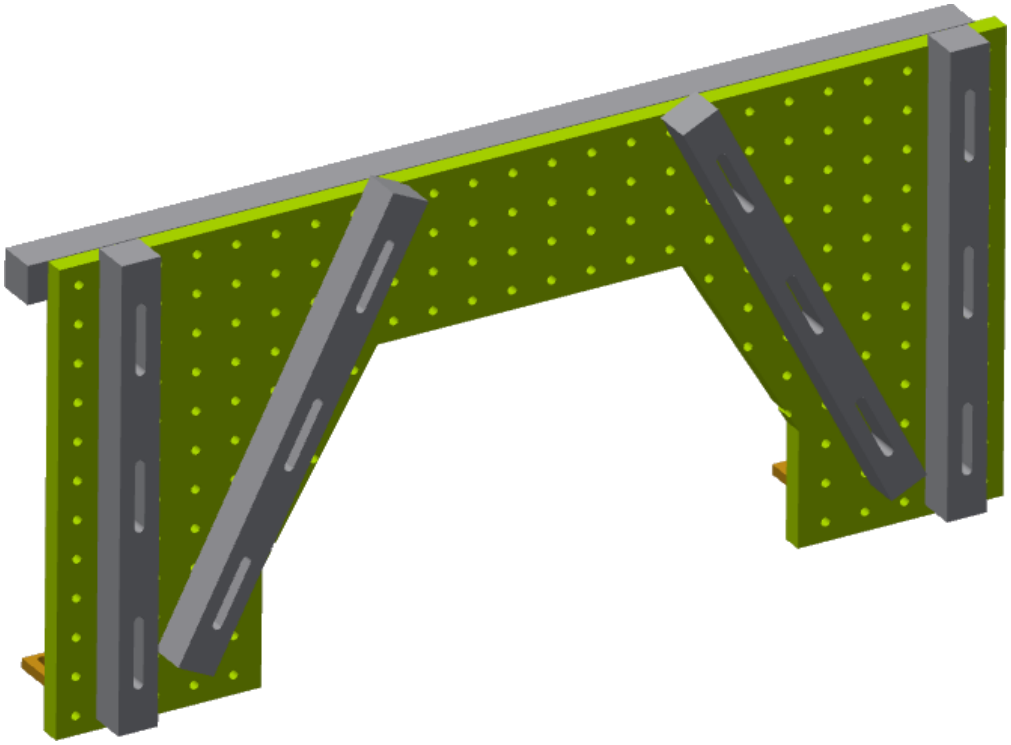


Figure 4–13: A render of the initial efforts to stiffen the single-plate gantry with steel braces.

A second plate was added behind the first. This second plate was added to allow the PGI gauge to be suspended above the Steinmeyer stage (Steinmeyer mechatronik GmbH, 2019) and to stiffen the gantry as a whole. The two plates could be fixed together using a variety of designs of braces or supports and so several different concepts were tested. The goal of the bracing was to attenuate the most problematic modes of vibration of the gantry system.

The 2-plate solution has some design constraints that were not present on the 1-plate solution. The plates must be spaced appropriately to allow 2 PGI gauges to be mounted concurrently, and the PGI gauge above the Steinmeyer stage must be raised up compared to the original vibration testing gauge. Another subtlety of the 2-plate design was the location of the angle mounting adapters (either inside or outside the plates). The location of the angle adapters was informed by FEA. Figure 4–14 shows the first double-plate gantry design, with further variations shown in Figure 4–15 and Figure 4–16.

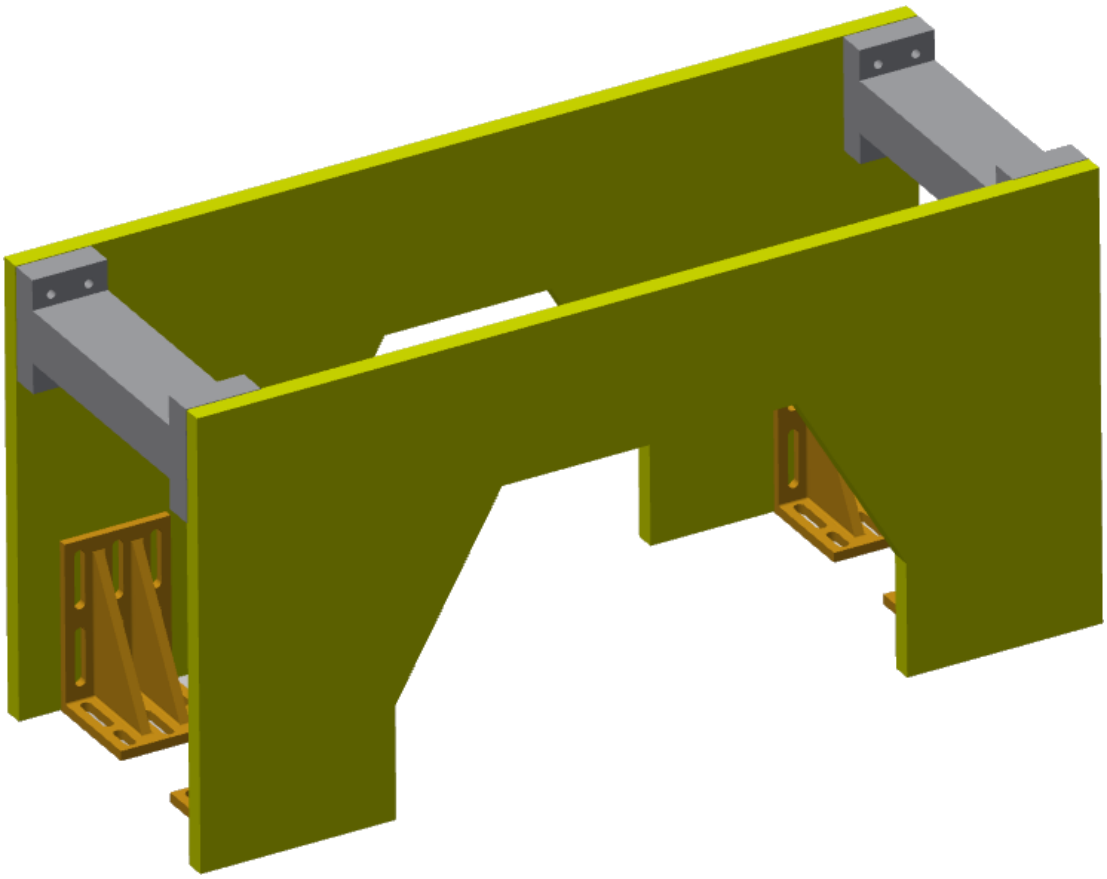


Figure 4–14: A render of the first double-plate design. This was as a baseline for evaluation of other configurations of gantry.

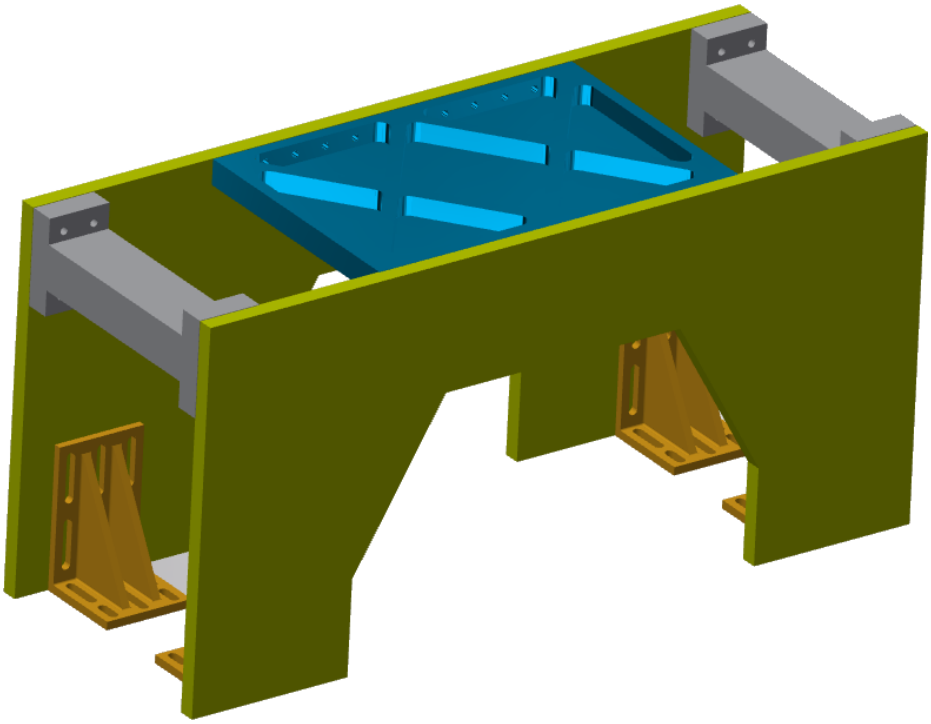


Figure 4–15: A render of the first improvement upon the initial 2-plate gantry. Note the use of 2 custom components (Bracing and ribbed-plate)

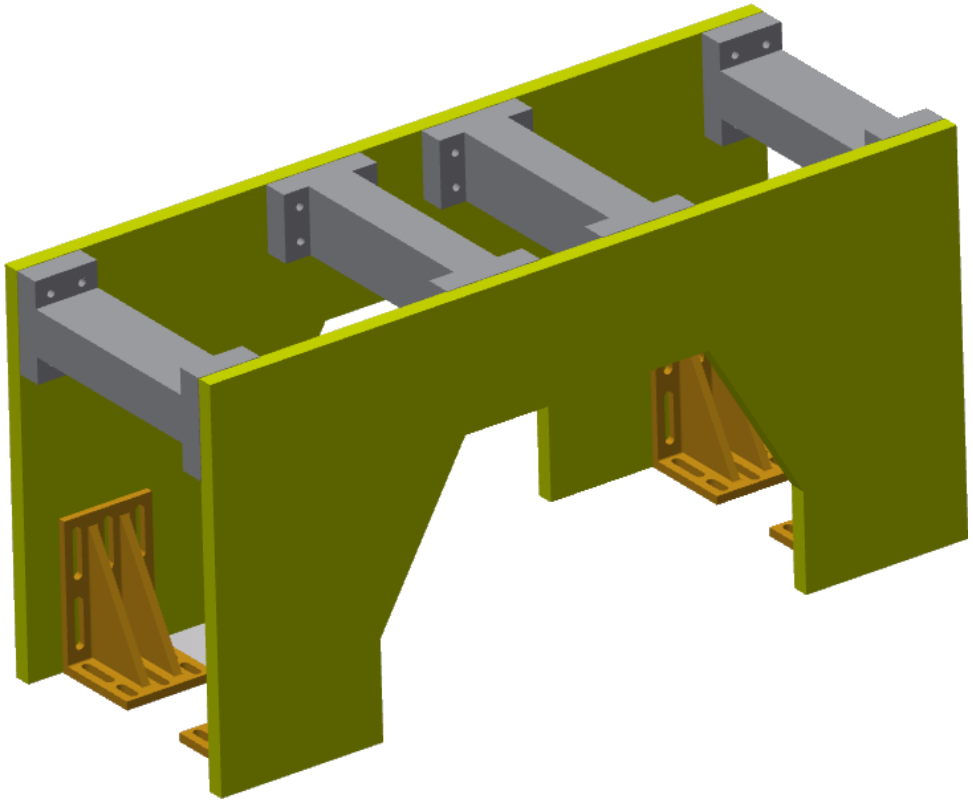


Figure 4–16: A render of an braces-only equivalent setup to Figure 4–15. This setup performed just as well and did not use a second custom component

First round – General form of a rigid gantry system

As mentioned previously, the gantry development was split into 2 rounds. This approach was chosen as it was not obvious which type of bracing would be most effective. Finite element modelling was performed on each gantry system, finding both the static and dynamic performance. The main parameter used for performance evaluation was the frequency of the primary mode of the vibration.

In each gantry configuration that was tested, the primary vibrational mode was always similar (see Figure 4–17). This is to be expected because of the fundamental geometry of the problem (including the mounting locations etc.) is consistent. The position of the applied force was chosen specifically to highlight the primary mode as it will cause the biggest displacement of the PGI gauge, compared with the other vibrational modes.

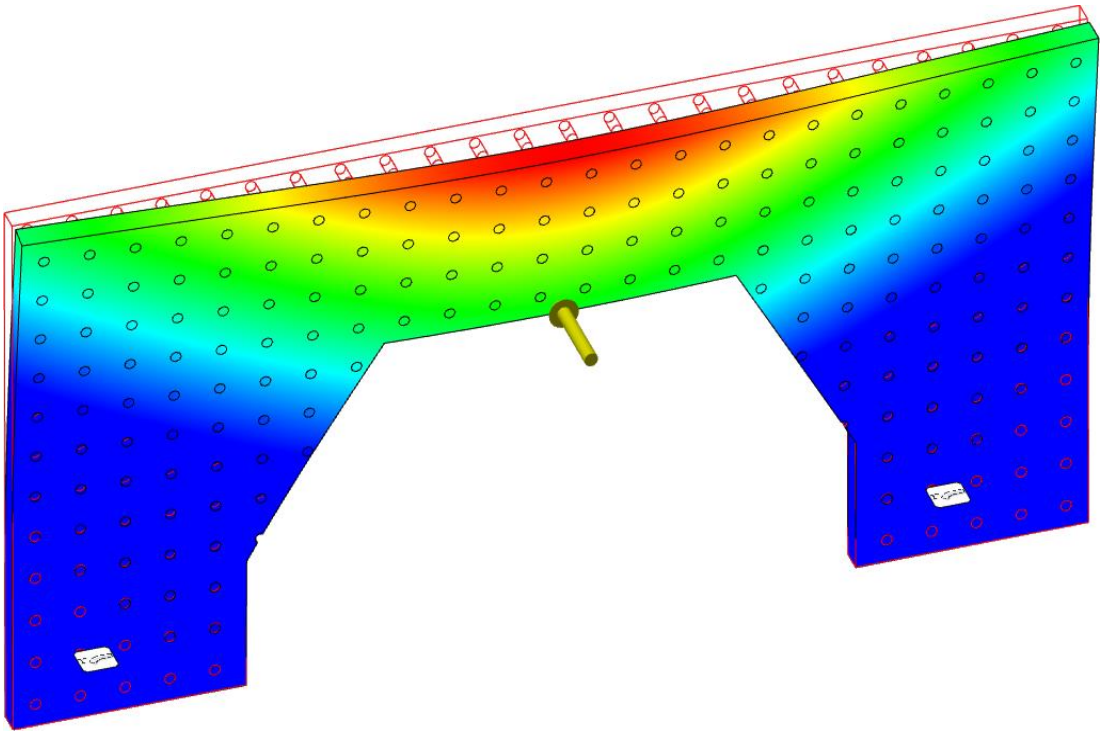


Figure 4–17: A basic finite element analysis of the dynamic performance of the single-plate gantry. Shown is the primary mode of displacement with a frequency of 136Hz.

The configuration shown in Figure 4–13 is an extension of the basic 1-plate gantry. It uses reinforcement beams to attempt to dampen the primary modes

of the plate, without interfering with the vertical slide (not shown, positioned front and centre). These beams did not supress the primary mode by any useful amount and would require 2 custom-machined steel components. This approach was almost as expensive as a 2-plate solution and was rejected.

Second round – 2-plate gantry systems

The configuration in Figure 4–14 was the first 2-plate method to be tested. The separation struts were designed to be very large so that they would not introduce their own dynamics into the situation. This configuration was not intended as a real solution but more as a reference for the other gantry configurations. For this reason, only 2 struts were used, and they were positioned furthest from the location most affected by the primary mode of the gantry. This was also the first configuration to use 2 different geometries of gantry plate (one side for the piezo and one side for the Steinmeyer stage). These geometries were similar, differing only in the height of the cut-out, but the asymmetry did bias vibration towards the side with less material (stage side)

As expected with this configuration, several of the lower-order modes were coupled between plates. The overall configuration was considerably stiffer than the 1-plate gantry, with a primary mode frequency of 214Hz. The dynamic modes that involve twisting of the whole gantry were more supressed in this configuration.

Several strut configurations were tested, with the aim of improving the overall stiffness of the structure. The use of a large, lightweight plate as a brace was investigated (Figure 4–15); the internal ribbing of the plate was designed to couple areas with low vibration amplitude on one plate with the areas of highest vibration amplitude on the other.

The FEA results for the plate-brace gantry configuration (Figure 4–15) are interesting for several reasons. Firstly, the primary mode frequency is lower than in the 2-plate configuration without the plate-brace (211 Hz vs 214 Hz). This represents a reduction in the specific stiffness of the gantry. Secondly, the coupling between the plates is greatly increased when compared to other configurations (this was the intended effect of the plate brace). Finally, the second and third vibrational modes were greatly supressed. The second mode (a twisting of the plates about the z-axis, at the top of the gantry plates) had a frequency of 438Hz. This contrasts with 290Hz for the same mode in the simple 2-plate gantry configuration. The typical third mode (a separation of the two gantry plates at their centres) seen in the other configuration is not present below 1050Hz in the plate-brace configuration.

The position of the plate-brace was lowered several times (in discreet steps as the bolt holes allowed) but this produced results that were not significantly different to this initial configuration. The plate-brace configuration was ultimately rejected because of the lack of increased stiffness in the primary mode and because of the need for 2 custom machined components. A full width plate was also designed but rejected because of the expense and difficulty of manufacture.

The next configurations to be tested used either 3 or 4 braces, arranged symmetrically between the gantry plates (Shown in Figure 4–16). The frequency of the primary modes was 194 Hz for the 3-brace configuration and 204 Hz for the 4-brace configuration. The dynamic performance of this group did not vary significantly across the range of brace configurations, nor did any of the brace arrangements supress any previously observed modes.

The reduction in the frequency of the first vibrational mode for the 4 brace configuration, when compared to the 2-brace solution, was likely caused by the increase in mass of the bracing. The coupling of the primary mode between the plates was similar to the plate-brace configuration, although the twisting modes were not as well supressed. As expected, the bracing greatly increased the coupling between the plates. The braces were moved around

the plates (as the grid-pattern allowed) but this produced no significant improvement.

Following the testing with the large braces (50 mm wide), the braces themselves were tested using FEA. It was found that their longitudinal stiffness did not increase significantly even when they were made 100 mm wide. The width was reduced to 25 mm wide (the hole spacing on the optical plates) and the gantry configuration testing was repeated.

As shown in Figure 4–22, the vibrational modes were very similar to previous configurations. Crucially the specific stiffness had been improved and the frequency of the primary mode had been increased from 194 Hz to 251 Hz. Several different locations and orientations of bracing were tested, including a 6-brace solution, although this exhibited similar coupling properties to the plate-brace solution and was rejected.

Different geometries of brace (thinner, weight removed from the central bar) were tested but they did not present a significant improvement over the implemented gantry configuration.

Implemented gantry configuration

The finite element analysis performed on the gantry systems suggested that a 4-brace solution would give the highest stiffness of all the configurations tested (FEA results shown in Figure 4–22). This configuration used only 1 type of brace which simplified manufacturing. The 2 central braces are offset from the top as this helps to reduce coupling in the second and third vibrational modes, which can be seen in Figure 4–23. A summary of the first 8 vibrational modes is shown in Figure 4–24. A render of the implemented gantry is shown in Figure 4–11.

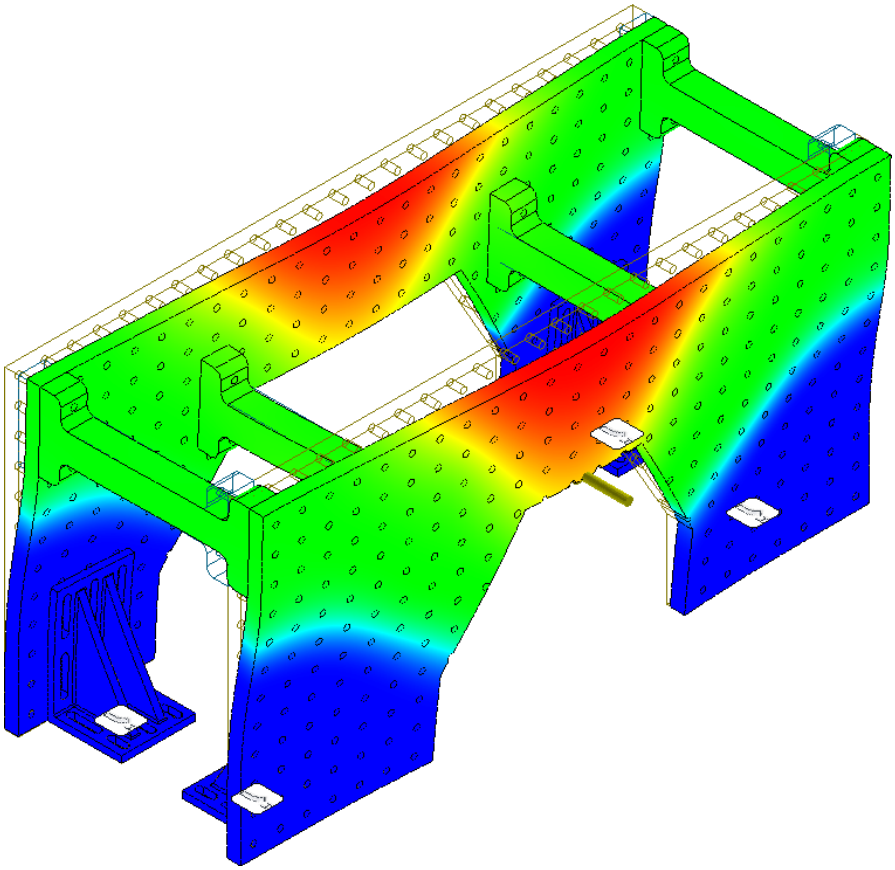


Figure 4–22: A render of the primary mode of the implemented gantry configuration.

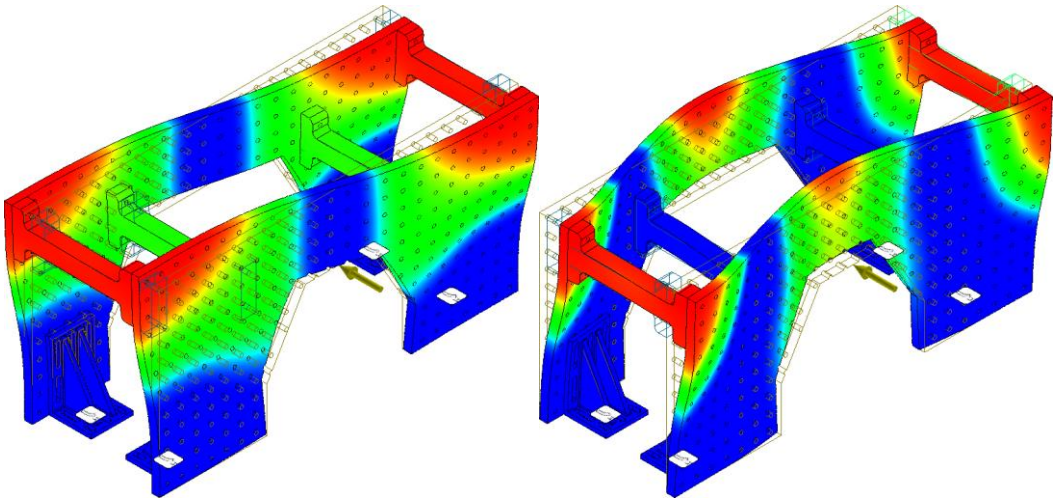


Figure 4–23: A Render of the secondary vibrational mode (Left) and the tertiary vibrational mode (Right) of the implemented gantry system.

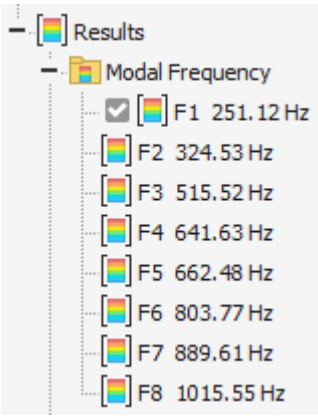


Figure 4–24: An extract from the FEA results showing the first 8 vibrational modes and their corresponding frequencies.

Notes on Finite Element Analysis

The finite element analysis system that was used was the built-in stress-analysis environment in Autodesk Inventor Professional (2015, 2016, 2017). The material properties used in the analysis were imported from the materials library within Autodesk Inventor; the closest of which was chosen for each component. The material used throughout the analysis was 'Aluminium, 6061', although the aluminium used for the brackets was 6062 Aluminium.

The contact between adjacent surfaces were automatically calculated by Inventor and were considered as an infinitely strong adhesive bond. This means that the faces in contact will translate together and experience the same forces. Autodesk inventor is not able to use bolted connections for dynamic simulations, only for static deflection simulations. The large number of bolts in the assembly (54 x M6 bolts, 25 mm on centres) and the large face-to-face contact between the angled brackets and the other components makes this an acceptable approximation to an adhesive joint.

The simulation used a single force of 10 newtons, applied at the top of the cut-out archway, as the perturbation (Shown as a yellow arrow in Figure 4–17). The position and direction of the applied force was chosen to maximise the displacement of the stylus tip. The actual forces experienced by the gantry system will be extremely small, although the prediction of vibrational behaviour

of the gantry has helped with the interpretation of the results of the stylus resonance testing.

Holes vs. no holes

The initial finite element analysis was performed on a gantry where the holes in the plate were modelled explicitly. This greatly increased the simulation time and the mesh complexity for the finite element analysis. Testing was performed using a fully modelled gantry and a gantry without the holes explicitly modelled. The mesh settings were increased for this test to give a mesh element every 0.005 of the bounding box length (2.8 million nodes). The reliability of the finite element analysis result is improved as the mesh element size is reduced, although this will give diminishing returns for each successive reduction in mesh element size. A choice of 0.05 as a relative mesh element length is a good trade-off between performance and accuracy.

The results of the FEA were very similar, although the plate with the explicitly modelled holes was slightly less stiff, as expected. They had similar vibrational performances (within 0.2 Hz) for all 8 vibrational modes that were tested. All intermediate FEA models were done with the hole-less plates to speed up testing, although the analysis of the implemented gantry was repeated with holes for inclusion here.

The mesh size was also varied to find the optimal mesh length. The stiffness of the model will be over-estimated if the mesh size is too large, but the computation time will be too long if the element size is too small.

The data in Figure 4–25 shows the positive correlation between mesh element sizes and the increase in frequency. From this data, the optimum mesh size of 0.05 was chosen to give a good balance between accuracy and simulation speed.

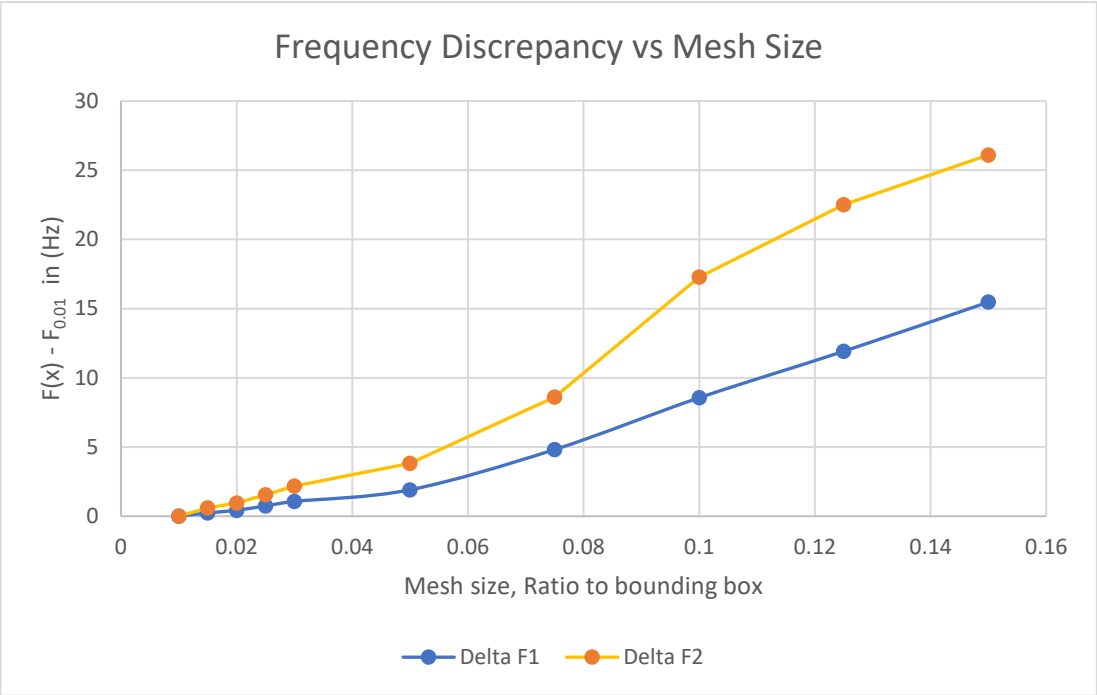


Figure 4–25: A graph showing the frequency discrepancy against the mesh size for the simplest gantry configuration. The modal frequency of the gantry with a mesh ratio of 0.01 is the frequency against which the other frequencies are compared.

4.3.4. Standard Optomechanical Components

As in previous design tasks, the use of standard components may greatly simplify the realisation of an engineering solution.

It is very common in optical system development to use a range of generic components that can be configured for a specific situation. Typical optical table setups require the stable and accurate positioning of many optical elements. The optomechanical components such as ‘breadboards’, angle plates, braces, linear stages and lens mounts are made to a very high standard to ensure compatibility with existing systems.

The breadboard (Left, Figure 4–26) is a large, flat aluminium plate, typically 12.7 mm thick, with tapped holes at regular intervals across its surface. The metric plates have M6 tapped holes on 25 mm centres across the surface. The array of holes provides a range of mounting options and will easily interface with the existing optical equipment that has been reserved for this project.

The right-angle mounting plate, (Right, Figure 4–26), is a precision 90° adapter made from aluminium. The mounting holes (tapped and clearance) are arranged to allow a wide variety of physical configurations. The angle plate has 2 reinforcement webs to increase the stiffness of the component.



Figure 4–26: (Left) A square optical breadboard, (Right) A 90° angle mounting adapter.

Some components were already available from previous projects within Taylor Hobson: This included a large (600 mm x 900mm) honeycomb breadboard, a granite surface plate (3ft x 4ft x 8in), a Steinmeyer linear stage, several manual linear stages and other assorted optics hardware and accessories.

The honeycomb breadboard has excellent vibration damping properties and is therefore used for the base of the vibration gantry system. The honeycomb breadboard rests on the granite surface plate to help isolate the table and reduce the transmitted floor vibration. The breadboard is not rigidly fixed to the granite to minimise the vibrational coupling.

4.3.5. Vibration in the gantry system

Measured Vibration

The gantry system was measured after assembly, to compare the real vibration against the simulated estimates. Three piezoelectric vibration probes were attached to the gantry, as shown in Figure 4–27. The locations of the accelerometers were chosen to be as close to the PGI gauge as possible, without requiring custom mounting components.

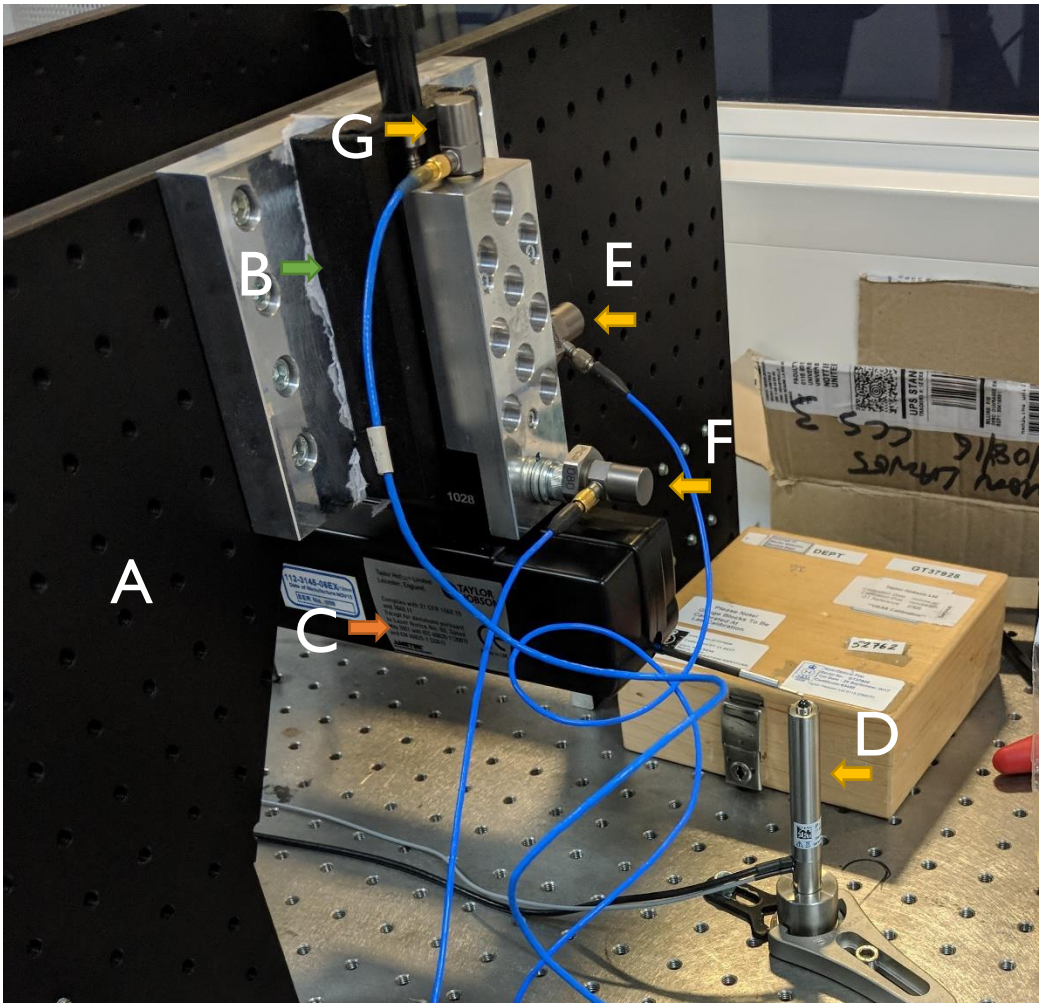


Figure 4–27: A photograph of the test setup used to measure the vibration characteristics of the gantry (A), the Z axis traverse and gauge mounting components (B), the PGI Gauge (C), the piezoelectric actuator (D) and the piezoelectric accelerometers for the X, Y and Z axes (E, F and G, respectively).

The ambient vibration experienced by the PGI gauge was recorded. The X axis and Y axis components are shown in Figure 4–28. The Z axis ambient vibration results are shown in Figure 4–29 for two different accelerometer locations, (left: on PGI gauge mount , right: on the floor below the machine)

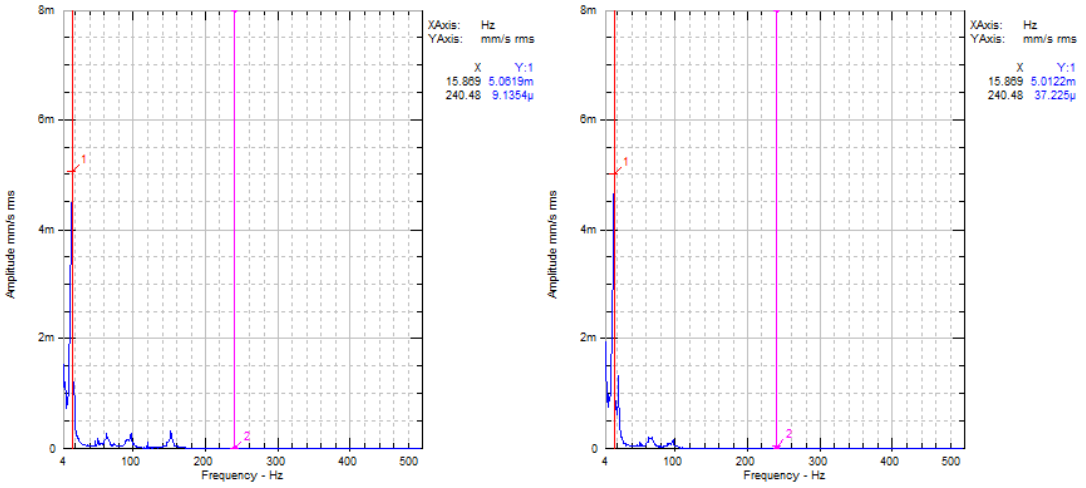


Figure 4–28: The amplitude vs frequency of the ambient vibration in the X axis (Left) and Y Axis (Right) of the gantry system, close to the PGI gauge.

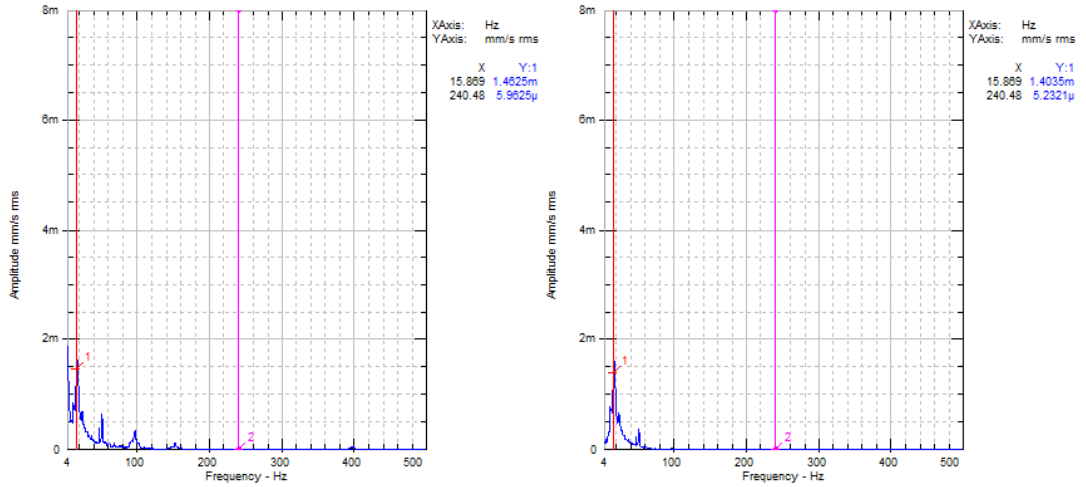


Figure 4–29: The amplitude vs frequency of the ambient vibration in the Z axis near the PGI gauge (Left) and the Z axis on the floor (Right)

The Z axis vibration amplitudes can be compared using a Frequency Response Function, as shown in Figure 4–30. This gives a measure of the susceptibility of the gantry to the floor-borne vibration (in the Z axis). The highest amplitude is 0.05 of the floor-borne vibration is transmitted, demonstrating the damping effects of the granite and the gantry system itself.

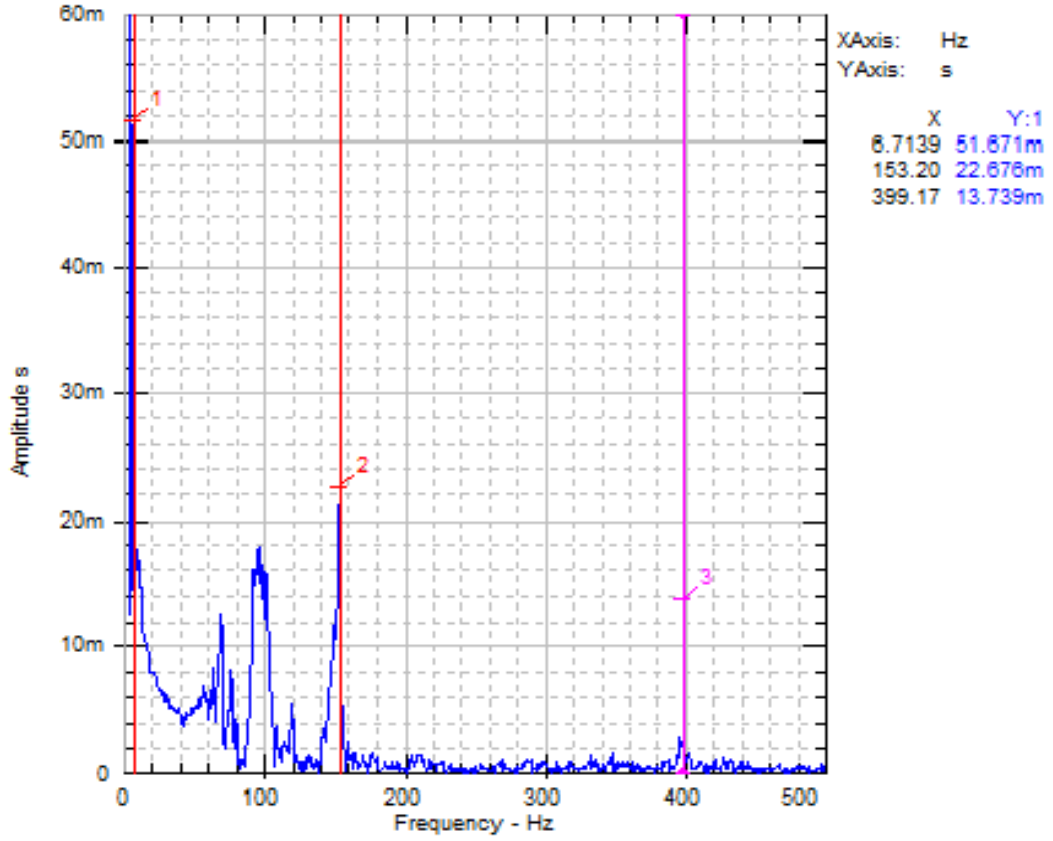


Figure 4–30: The frequency response function of the Z axis on the gantry when compared to the Z axis of the floor.

The ambient vibrations of the gantry system were minimal and almost entirely restricted to the region below 150Hz. This demonstrates the importance of measuring the environment when setting up measurement instrumentation.

An excitation was induced in the gantry in the form of a small, repeated impulse from a rubber-handled hex-key. This allowed a small and consistent excitation to be present throughout the data acquisition process. These impulses were light taps, using the rubber end of the hex key, in the direction of the axis being investigated.

Figure 4–31 shows the resonance peaks for the X and Y axes during excitation. The frequency of the X axis is lower as the gantry stiffness is lower in this direction. Although FEA is not directly comparable, the predicted frequency of the excitation in the Y axis is ~250 Hz and the experimentally measured frequency is 240 Hz.

The measured amplitudes in these axes are below the required noise floor for the environmental noise that would be appropriate for the installation of a Talysurf™ system.

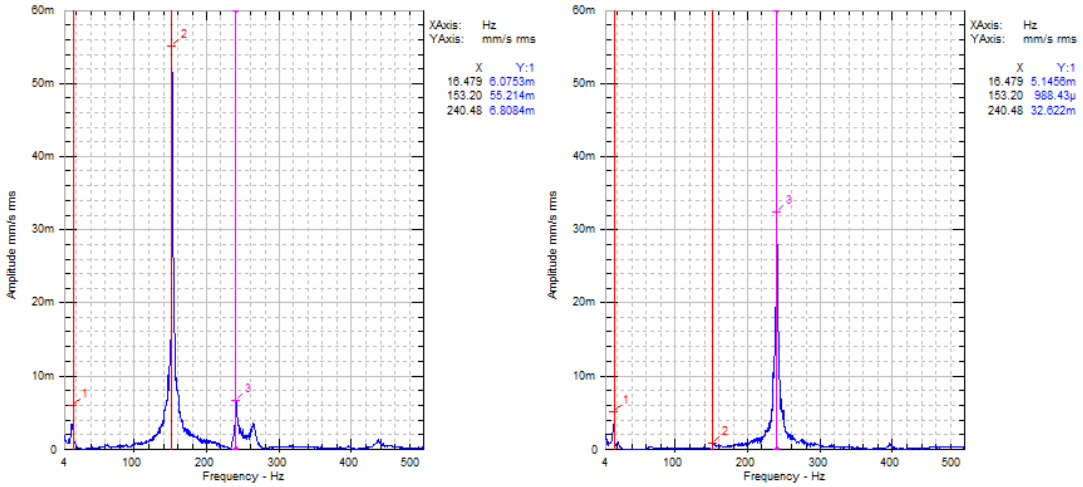


Figure 4–31: The amplitude vs frequency of the induced vibration in the gantry in the X axis (Left) and the Y axis (Right), while impulses are performed in the X Axis.

Figure 4–32 shows the greatly increased stiffness in the Z axis when compared to the X or Y axes. Although a resonance peak is still identified at 563Hz, the amplitude is only a few times the ambient amplitude for this frequency.

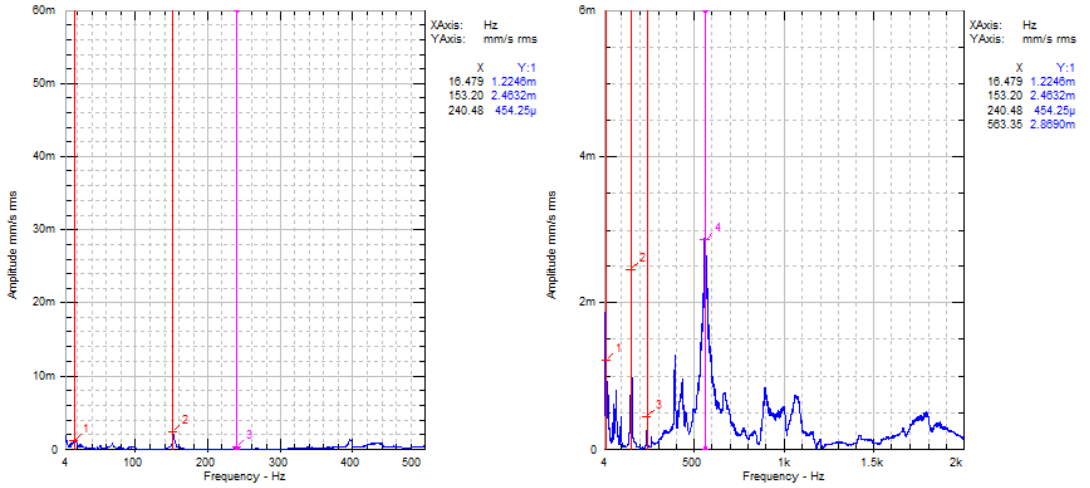


Figure 4–32: The induced vibration in the Z axis while impulses are performed in the Z axis. The figure on the left is the same scale as the X and Y axes graphs.

Sources of External Vibration

The gantry system must be resistant to vibrations from external sources. The design of the gantry itself will broadly determine the vibrational susceptibility of the overall system. The energy to excite the gantry system comes from a variety of sources, including ground vibration and air-borne vibration.

The gantry system rests on a large granite surface plate which weighs roughly 600kg and rests on a welded steel frame. This large mass will naturally resist low frequency vibration and provides a rigid base for the optical breadboard.

The bay in which the gantry is currently located is temperature stabilised by the air conditioning, varying by ±0.5 degrees per hour. The air conditioning fans cause vibrations in the air with a frequency of between 10 and 18 Hz. This can be clearly seen in Figure 4–28 and Figure 4–29. Air-borne vibrations will have a smaller effect on the gantry frame itself because of the low coupling between the air and the aluminium frame. The coupling between the air and the stylus is higher and so the air conditioning affects the stylus disproportionately by comparison to the rest of the system.

Floor-borne vibrations come from many sources, including, air compressors, machine tools, air conditioning, heating systems, road traffic, doors opening

and closing, people walking nearby, heavy lifting equipment and even iron-core transformers. By observation, the gantry in its present location is not susceptible to noise from the machine shop but is highly susceptible to the nearby door being opened or closed during measurement.

To reduce the transmitted noise, the supporting electronics for the gantry have been removed from the table. The PC and the piezo amplifier both contain cooling fans; the vibration of these fans is visible in a static noise measurement using the gantry.

4.3.6. Stylus Dynamics - Method of Analysis

The profiles that are gathered during the experimental phase are analysed using a purpose-written Python program to extract features from the gathered data for subsequent comparison. The collected data is typically a 15 second 'static noise' measurement[7]. This includes an initial 3 second delay to allow the analysis algorithm to establish the unperturbed gauge reading. This is followed by 10 seconds of actuation from the piezoelectric stack. The remaining 2 seconds is to ensure that all the data is collected, even with small delays introduced by the different control systems.

The profile data must be cut to exclude the data before and after the actuation. This is achieved by a thresholding process, using the initial noise to determine the threshold amplitude. The mean and the standard deviation of the first 2 seconds of the collected profile are calculated. As the actuation of the piezo starts 3 seconds after the data gathering, the first 2 seconds of gathered data is used to evaluate the environmental noise and the average amplitude of the gauge before actuation begins.

The first data point with an amplitude higher than 5 standard deviations above the mean is found. The algorithm backtracks from this point to the until the amplitude of the profile is within 1 standard deviation of mean and this point is used as the start of the actuation proper. This has a success rate of >98% and is neither computationally time consuming nor conceptually complex. Figure 4–33 shows a typical gathered profile, spanning from 1 second before to 1 second after actuation begins. The threshold that is used to discern the start of actuation is shown in red in Figure 4–33.

After the initial non actuation data is removed, the remaining data is then reduced from ~12s to 10.0 seconds by simply cutting the data to this defined length, starting at the actuation point. There are some features present at the end of the actuation, which would make it possible to use a similar approach

[7] A Static Noise measurement is a measurement which does not move any axes and simply records the gauge reading. This is typically used to assess the background or environmental noise of an instrument.

to corroborate the start and end times, but this proved to be difficult in practice, and lowered the success rate of the data cutting.

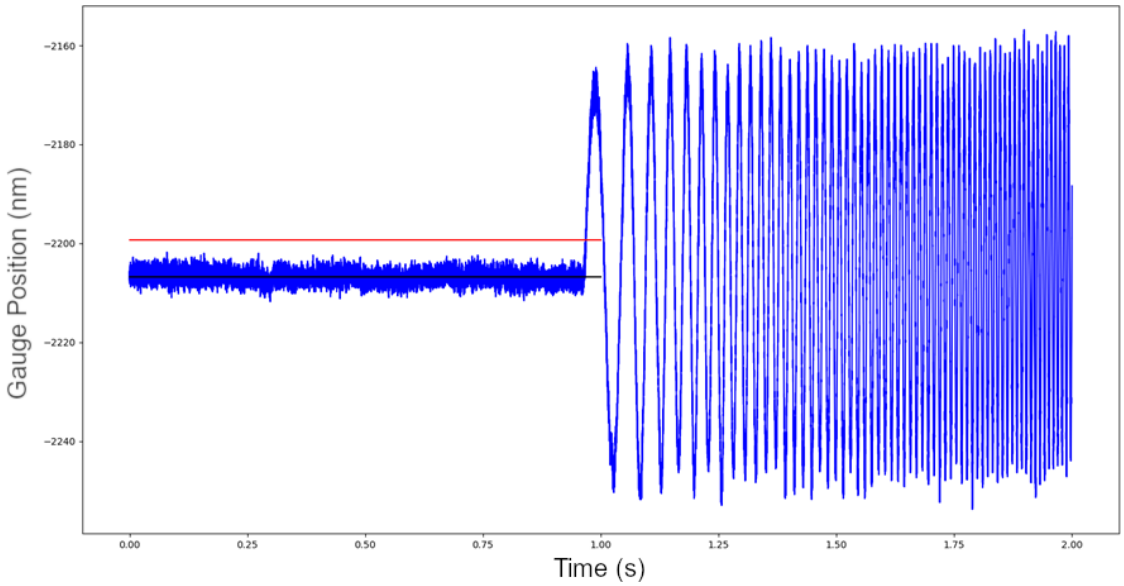


Figure 4–33: A plot showing the automatic cutting routine threshold levels (Red horizontal line) the no-actuation period mean value (Black Horizontal line) and the profile they were generated from (Blue). The image is focussed on the start of actuation for clarity, more lead-in data is used than is shown in the image.

The cut data typically consists of:

- A stable actuation phase, at a fixed amplitude (labelled A in Figure 4–34)
- A section of data which gradually rises to the resonance peak (labelled C in Figure 4–34)
- An abrupt change into stylus flight (labelled D in Figure 4–34)
- A steady amplitude reduction as the frequency becomes too high for the gauge to capture accurately. (labelled F in Figure 4–34)

There may also be a period of stability after the initial stylus flight (labelled E in Figure 4–34). This is normally observed when testing ceramic styli, although all other stylus materials may have this feature.

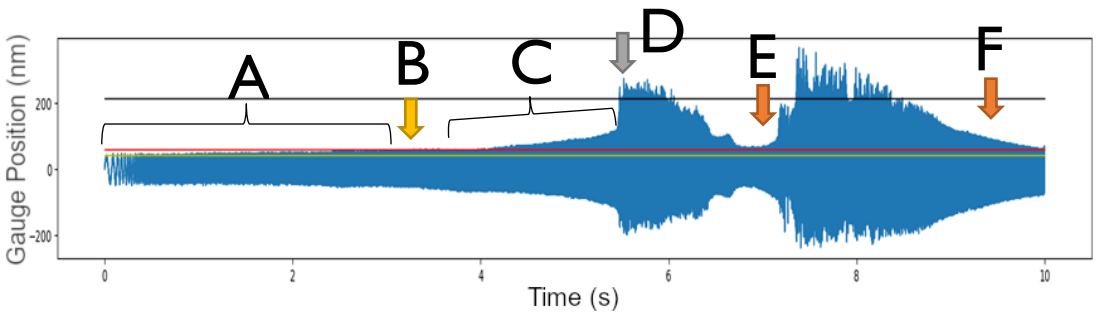


Figure 4–34: A typical 2V amplitude actuation for a ceramic stylus. The features are similar between most stylus materials. Labelled regions are; A: stable actuation region, B: small amplitude shift C: a region of changing amplitude, D: stylus flight, E: a region of stability, F: amplitude fall-off.

The initial region (labelled A in Figure 4–34) is used to characterise the actual actuation amplitude; this amplitude is presented in the summary table. The maximum height of region A and the maximum height of the profile overall are averaged to give the stylus flight threshold level (shown in black on Figure 4–34). The first time the profile exceeds this amplitude is used as the stylus flight frequency. This is correct in ~98% of cases, although this is subject to a few caveats.

A Fourier transform of the cut data, up to the stylus flight frequency, is performed. An example is shown in Figure 4–35. This gives a better understanding of the stylus frequency performance, as the Fourier transform helps to highlight noise, whilst removing the form that overwhelms the profile data. The subtleties of the Fourier transform are difficult for an algorithm to interpret, but a human who is even passingly familiar with the analysis procedures would be able to instantly tell if the automatic cutting and thresholding algorithms had worked correctly. See Page 177 for more details.

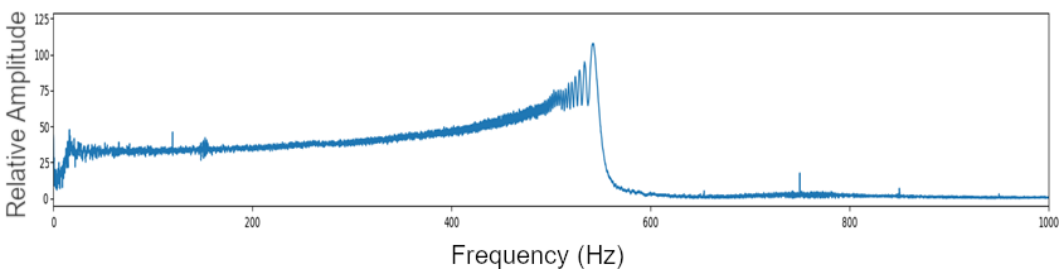


Figure 4–35: A Fourier transform of the profile data, using data up to the stylus flight frequency.

The profile gathering process is conceptually similar to a Fourier transform, insofar as the amplitude response of a system is related to the frequency sweep of the input, so certain features are visible on both the profile and the Fourier transform. Some higher frequency features are visible in the Fourier transform such as the spike at 750Hz and associated sidebands in Figure 4–35. The higher frequency information is present because the actuation signal is not a pure sine wave.

A spectrogram is made of the cut portion of the profile. A spectrogram is a series of short time Fourier transforms, each representing a short time period and often having some overlap with the previous segment. Spectrogram plots trade off frequency resolution for a more nuanced understanding of the progression of a function over time. An example spectrogram is shown in Figure 4–36.

The data contained within a spectrogram gives a more intuitive understanding of the dynamic behaviour of the stylus, including the stylus flight frequency. Several horizontal features are visible within the spectrogram, suggesting that there are some environmental or systematic influences on the recorded profile which are not changing over time. This may be the stylus beam resonance, although it is difficult to confirm this without alternative testing arrangements.

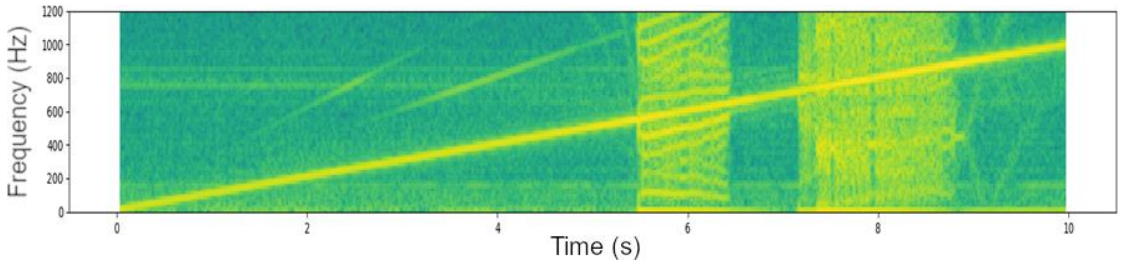


Figure 4–36: A spectrogram of the above profile, showing the linear progression from 10Hz to 1000Hz over 10 seconds. The regions of stylus flight are obvious when presented in this format. Several horizontal features can be seen in this representation (~800Hz, ~850Hz).

Finally, each peak and valley present in the profile data up to the stylus flight frequency is identified. This data is then used to determine the change in amplitude (dB) between the stable region and the build up to resonance region.

The profile, spectrogram, Fourier transform, and a table of results is presented for each individual stylus test; the render is saved as an image, allowing verification of the analysis at a later date. The date and time of the measurement, as well as the stylus flight frequency, are exported to an external file, allowing the gross changes over time to be observed. The exported image can be used to discriminate the profiles which have not automatically cut and so must be re-analysed using a manual cutting process. An example of the complete analysis is shown in Figure 4–38 and an example of an analysis performed on an incorrectly cut profile is shown in Figure 4–40

An example of stylus testing without stylus flight

The testing of low actuation amplitudes revealed a threshold, below which a given stylus may not undergo stylus flight at all. Although the testing algorithm still identifies a stylus flight frequency, manual inspection of the spectrogram shows that the characteristic vertical features typically associated with stylus flight are not present at these low amplitudes. Efforts were made to use the spectrogram to determine the stylus flight frequency, but the complexity of this task led to a considerably reduced reliability and so was not developed further.

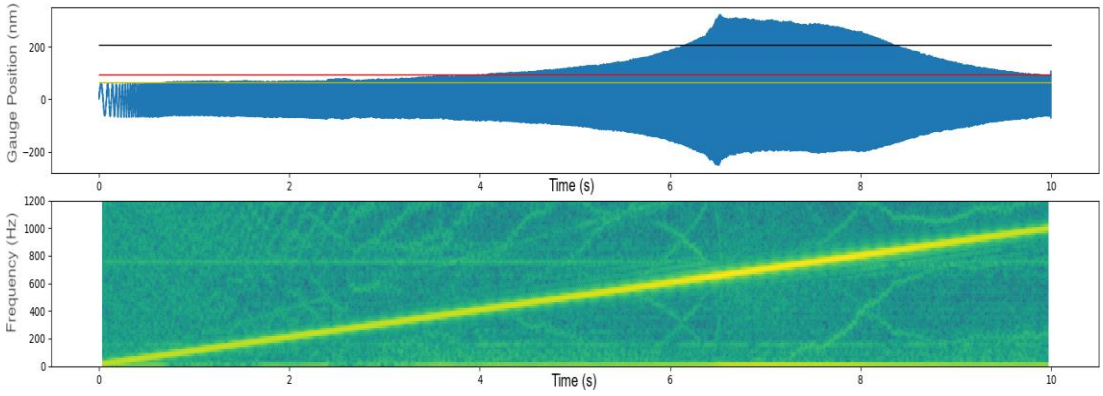
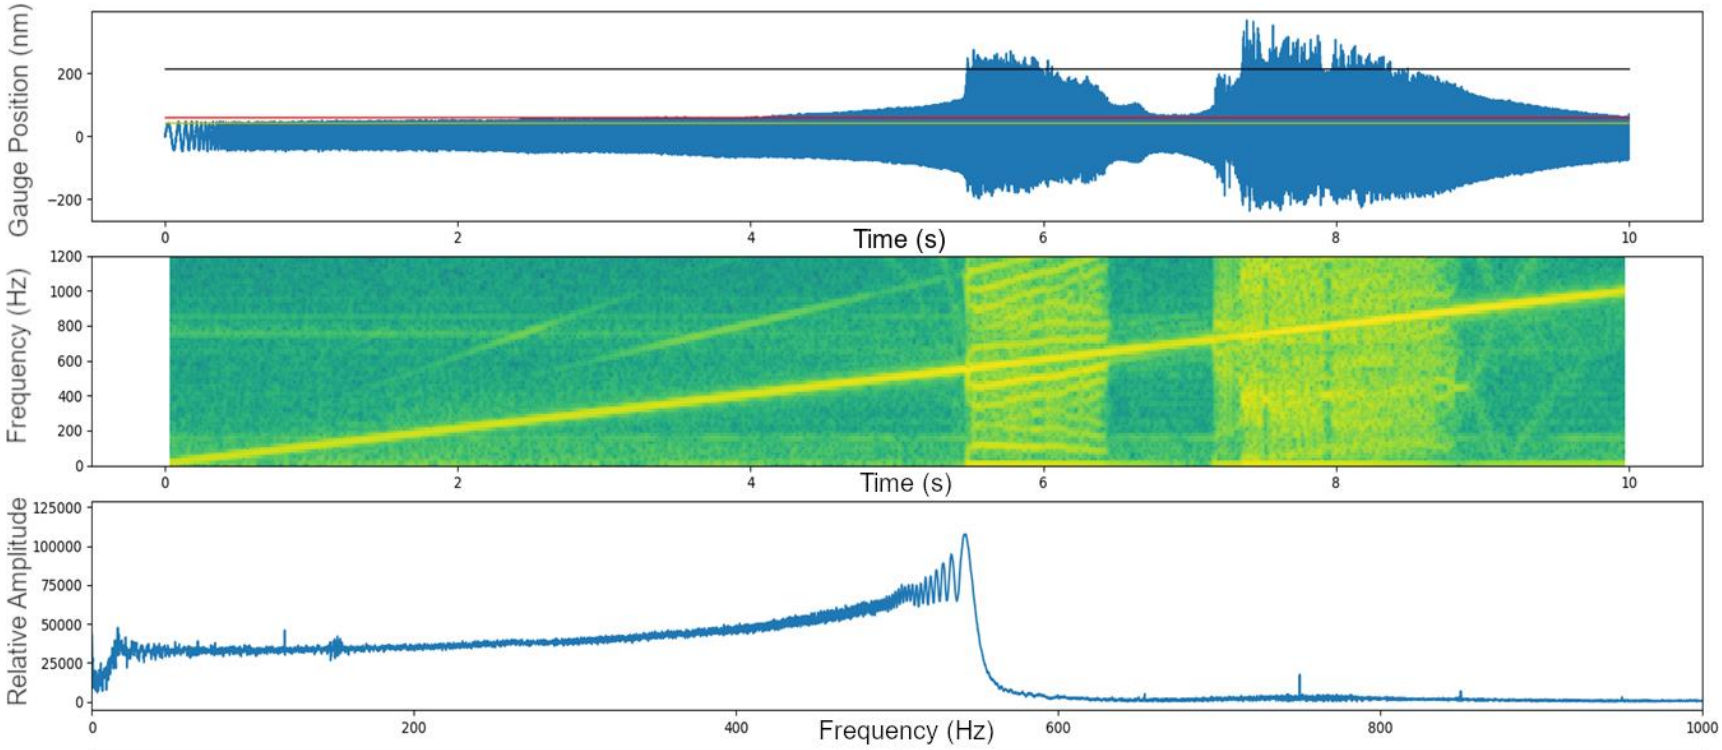


Figure 4–37: The dynamic testing analysis results for a ceramic stylus with a 1V actuation voltage. The spectrogram has no characteristic vertical features that would typically indicate stylus flight.

This approach may allow direct measurement of the stylus primary resonant frequency, although not all styli were found to have a lower limit on the amplitude for which flight occurs, limiting the usefulness of this technique.



Parameter	Value
Measurement Time	10.0 s
Initial Actuation Peak-to-Valley	121.00 nm
Maximum Stable Frequency	552.6 Hz
Minimum Stable Peak Height	36.70 nm
Minimum Peak Height in dB	-4.34 dB

Figure 4–38: A typical output picture of the automated analysis program.

4.3.7. Errors within the Automated Analysis Process

The automated analysis program is highly reliable, typically correctly identifying all features within ~98% of the collected profiles, but it can be susceptible to certain errors and near misses. These issues are plainly obvious to a human looking at the output data, but they can be extremely difficult to detect algorithmically.

Incorrectly Cut Data

The measured profile is subject to a small amount of gauge drift over the 15 seconds (typ.) of the measurement. This can sometimes manifest as a change in level of the initial pre-actuation period (Labelled A in Figure 4–39), resulting in a mean and standard deviation that do not give favourable cutting parameters for the automated cutting routine. Further, on some styli, there is a small glitch in the measured profile which will cause the cutting parameter to detect the actuation too early. (Labelled B in Figure 4–39)

The two issues shown in Figure 4–39 will cause the analysis program to give an overestimate of the stylus flight frequency, as this information is derived from the time between the start of the actuation and the occurrence of the stylus flight, linearly interpolated between 10 and 1000 Hz over 10 s. There are potentially several ways that the automatic cutting can be verified, but the analysis process reliability is high and so little would be gained by improving it further.

The simplest method of ensuring a particular profile is properly cut is to manually re-run the analysis process, thereby giving the user a chance to pick the start of the actuation. This manual re-analysis also gives the opportunity to explore the resultant graphs without the loss of resolution when the graph is exported as an image. The ability to zoom and explore the renders, without being limited to the image resolution, can be very useful to investigate other phenomena such as near misses.

The result of this premature actuation detection can be seen in Figure 4–40. Note the initial dead-time in the profile and the spectrogram. This dead-time

does not feature prominently in the FFT as the amplitude of the dead-time is small by comparison to the actual profile.

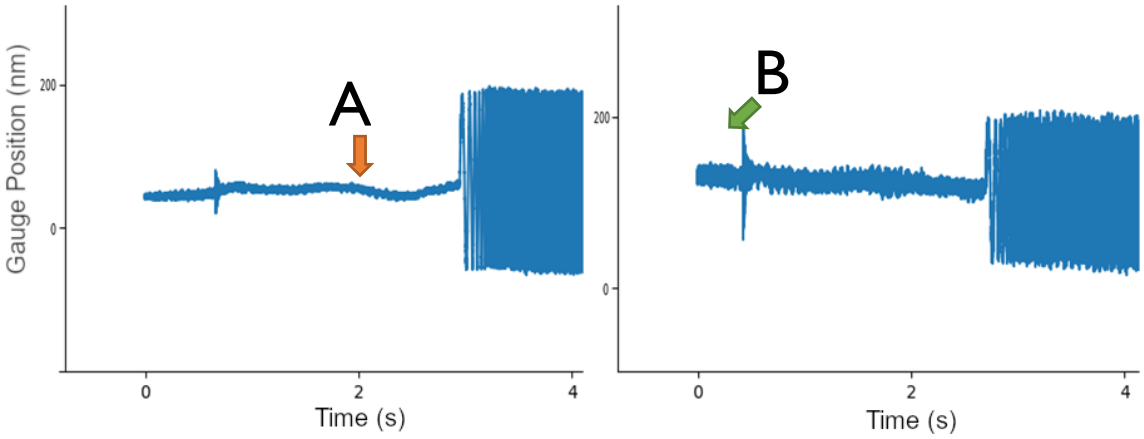
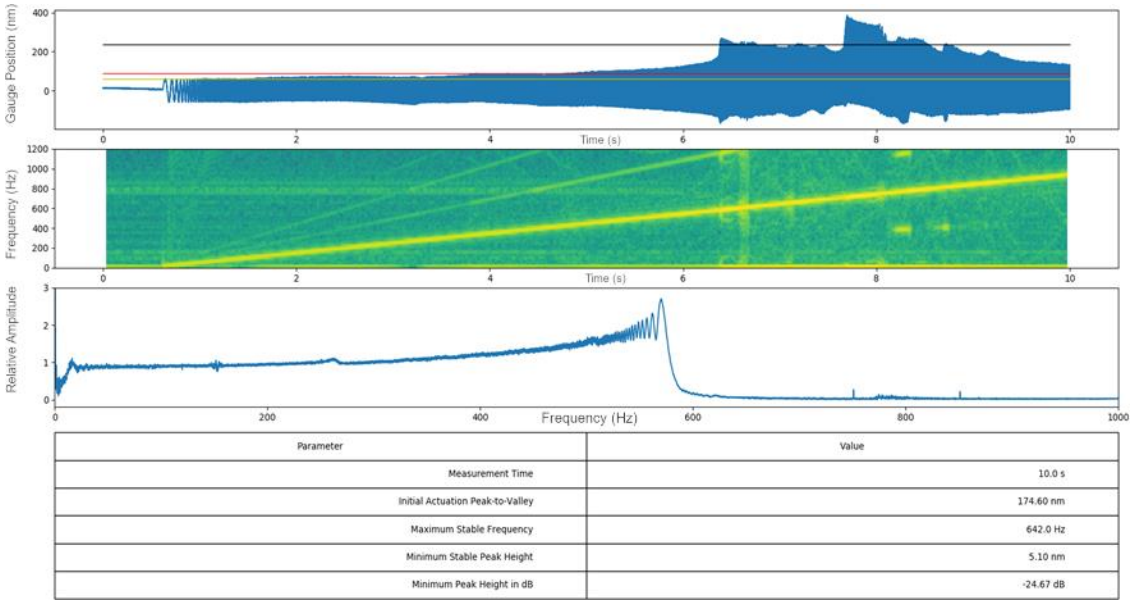


Figure 4–39: A graph showing the pre-actuation sections of a measurement. The build-up period is susceptible to drift (Labelled A) or to glitches (Labelled B)



Parameter	Value
Measurement Time	10.0 s
Initial Actuation Peak-to-Valley	174.60 nm
Maximum Stable Frequency	642.0 Hz
Minimum Stable Peak Height	5.10 nm
Minimum Peak Height in dB	-24.67 dB

Figure 4–40: A mis-cut analysis of a stylus actuation. This was caused by a small deviation in the profile before the actuation started.

Identifying Stylus Flight

The initial testing with the PCF styli resulted in a simple thresholding approach to determine the stylus flight frequency. This was simple and effective, avoiding the need for an algorithm of higher complexity. This thresholding level is set at the mid-level between the level of the initial actuation and the maximum point along the profile. The amplitude of the initial actuation period is typically very stable, although a glitch would cause issues if it occurred at this point. The absolute maximum height of the profile is typically less repeatable between measurements, although the threshold level derived from the absolute maximum value still works well in most cases.

The ceramic and the aluminium styli have slightly different characteristic actuation curves, examples of which are shown in Figure 4–41. The ceramic profile typically has a gradual rise to the maximum stylus flight amplitude, even if the stylus undergoes flight before the maximum amplitude is reached. The gradual change of amplitude causes the thresholding to vary between different measurements, as the transition from surface-tracking to stylus flight is no longer immediate. The actual point at which the stylus begins to leave the surface is not detected for several cycles of the piezo actuator, meaning that the Fourier transform results include stylus flight data and appear noisier.

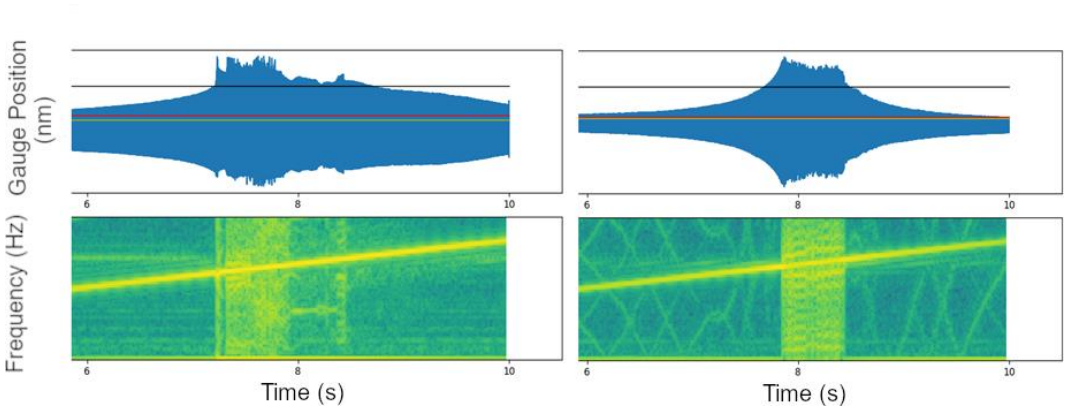


Figure 4–41: A comparison between the stylus flight build-up of a PCF stylus (Left) and a ceramic stylus (Right).

Figure 4–41 demonstrates the downside of using a thresholding approach. The PCF styli typically transition from resonance build-up to stylus flight in a sharp step, whereas the ceramic stylus amplitude reaches greater than 90% of the stylus flight amplitude before undergoing stylus flight.

Premature Thresholding of the Stylus Flight Frequency

If a stylus flight frequency is close to the stylus resonance frequency, then the build-up to the resonance peak will be visible before the stylus leaves the surface. The actual resonance peak is not typically visible, except for the additively manufactured stylus. The use of the triangular waveform shape also results in a resonance feature being visible, although it has not been conclusively proved that this corresponds to the beam resonance. This resonance feature is not properly mitigated by the thresholding algorithm and so the measured stylus flight frequency is low.

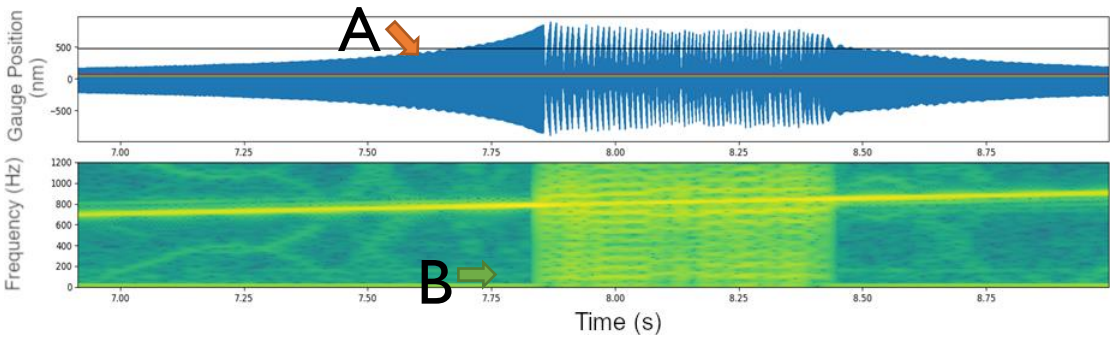


Figure 4–42: A section of the profile of a ceramic stylus, undergoing gantry testing. The horizontal black line (top) is the threshold value used to determine the stylus flight frequency. The spectrogram of the section is shown below.

As can be seen in Figure 4–42, the threshold level predicts the stylus flight position to be at 7.65 s (Labelled A), whereas the spectrogram shows the actual stylus flight to have taken place at 7.82 s (Labelled B). These correspond to frequencies of 767 Hz and 784 Hz respectively. The spectrogram could be used to more accurately predict the stylus flight condition, but the algorithm required would be significantly more complex than a simple threshold on the profile data. Premature thresholding can be accounted for at the analysis stage, or alternatively it can be considered reasonable to threshold early (albeit without labelling it the ‘stylus flight

frequency') as the stylus-gauge system is not able to capture the 'real' surface at these frequencies, even if the stylus hasn't undergone stylus flight.

Thresholding – Off by One Error

The resolution of the reporting of the stylus flight frequency is 0.1 Hz but there are factors that mean that this is not a realistic measurement accuracy. The thresholding approach works well for a sharp transition to stylus flight, typical of PCF styli, but can have introduce small errors for 'slow' amplitude transitions, such as for ceramic styli.

The simple thresholding technique will not exercise any judgement on the point at which a profile crosses the threshold, or if the profile does not continue to exceed the threshold after the first positive result. An off by one error in this context is the difference in frequency if the threshold had been just above point A in Figure 4–43. The next time profile can be above the threshold is not until (B), thus setting a limit to the effective resolution.

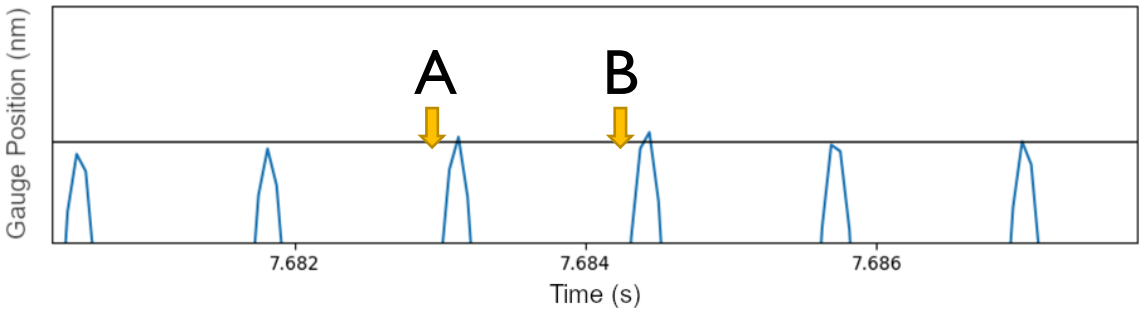


Figure 4–43: A highly magnified section of the profile shown in Figure 4–42, shown at the point at which the profile first crosses the stylus flight threshold.

In this instance A is at a frequency of 770.61 Hz and B is 770.81 Hz, although this limit is dependent on the absolute value of the frequency involved. In this instance, the threshold frequency would have an uncertainty of ±0.2 Hz simply from the off-by-one error.

Thresholding – Near Miss

Certain styli (usually the ceramic styli), display resonance features throughout their gathered profile. The threshold algorithm does not account for these

features which can result in a large disparity between the measured stylus flight frequency of two consecutive measurements. It should be noted that in both instances, the stylus flight frequency is incorrect, according to the spectrogram.

Figure 4–44 shows both a ‘near miss’ (top) and a ‘near hit’ (bottom). Since the flight threshold amplitude is related to the maximum amplitude of the profile, any period of large amplitude excitement can drastically change the threshold level. If the actual stylus flight amplitude observed in (Figure 4–44, bottom) were slightly higher, the threshold level would have missed the initial excitation (Labelled B) and the measured stylus flight frequency would have moved to a significantly higher frequency.

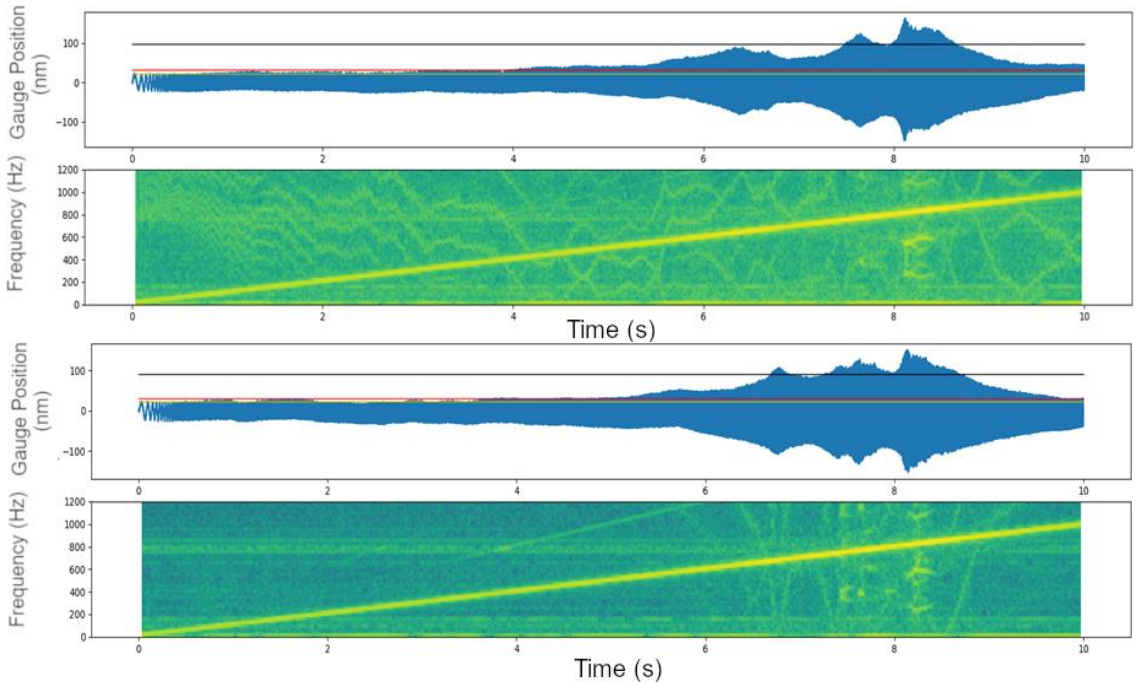


Figure 4–44: The analysis of a pair of consecutive measurements, giving vastly different measured flight frequencies. The amplitude used to drive the piezo for the measurements shown was 1Vpp.

Using the spectrograms, instead of the profile, the actual stylus flight frequencies can be more readily seen. In the top measurement of Figure 4–44, the stylus appears to leave the surface only at ~810 Hz, while in the bottom measurement, the stylus leaves the surface twice, at ~740 Hz and ~815 Hz. The thresholding technique fails to correctly identify stylus flight in either of the measurements shown in Figure 4–44. Fortunately, the higher actuator amplitude tests have a lower rate of misidentification and so this problem was not particularly prevalent in the materials testing experimental results.

Effect of late Thresholding on FFT

The accurate detection of the stylus flight frequency is complex, and the misidentification can have some effects on the conclusions that can be drawn about the performance of a particular stylus. The Fourier transform section of the analysis procedure is particularly susceptible to the inclusion of data beyond the stylus flight frequency.

Figure 4–45 demonstrated the effect that including a small amount of post-stylus flight data has on the Fourier transform analysis. The included data has energy distributed across the frequency spectrum (as shown in the spectrogram), resulting in a Fourier transform with many periodic features, making the analysis process considerably more difficult.

For example: The gantry system has a natural resonance at ~280 Hz, visible as a small decline in the Fourier transform amplitude in the correctly cut example (Top in Figure 4–45). This same gantry resonance feature is present in the incorrectly cut example (Bottom in Figure 4–45) but is considerably more difficult to identify consistently.

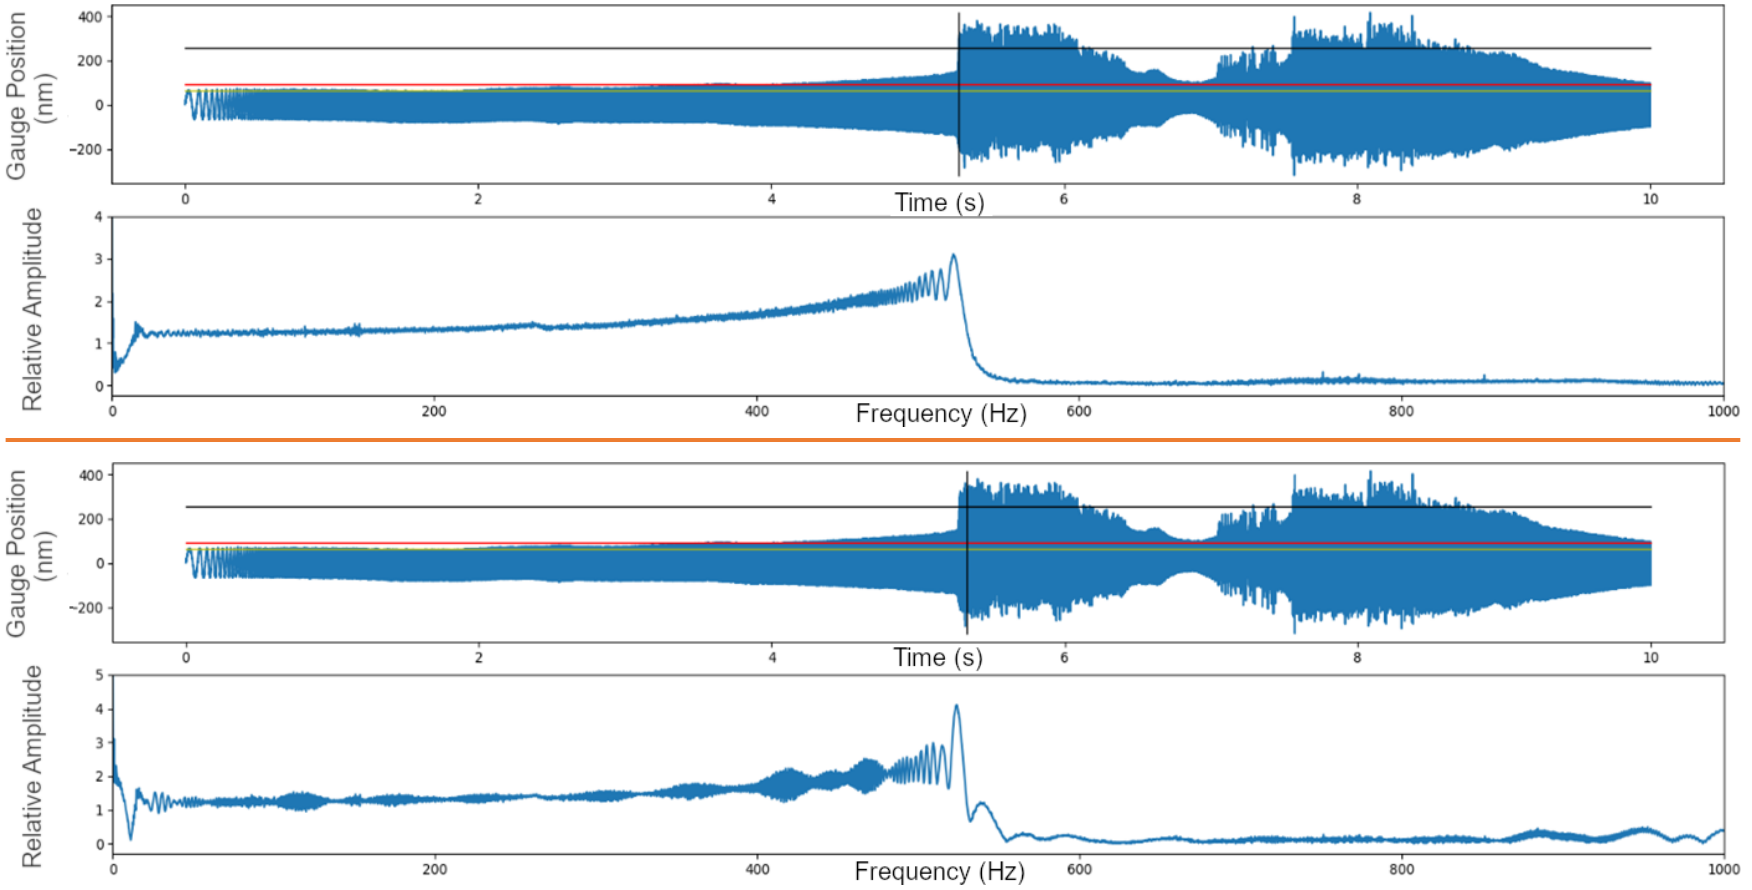


Figure 4–45: (Top) a standard analysis with a correctly detected stylus flight frequency (532.5 Hz) and the corresponding Fourier transform. (Bottom) The same profile with the stylus flight frequency 1% higher than the correct frequency (538 Hz).

5. Materials Testing

This section covers the range of potential and realised styli that were manufactured for testing purposes, the set of tests that these styli have been used for, the underlying reason for the testing and the results of the testing process. Factual information is presented, including the information required to repeat these tests in the future. Discussion of this information within the wider context of Taylor Hobson, and stylus metrology in general, is done in a future chapter.

The motivation for the investigation into the performance of a range of stylus materials is two-fold:
Firstly, there is a long history of using pultruded carbon fibre as the material of choice at Taylor Hobson. Significant advancements have been made in materials science and production techniques in the time since PCF was first used and so some traditional stylus materials must be investigated in a modern production context.

Secondly, Additive Manufacturing technology has become ubiquitous in recent years. Components with geometry which cannot be produced any other way have shifted the manufacturing and production paradigm in numerous engineering fields. Additive technology has the potential to change the stylus manufacturing process in the same way, and so AM must be investigated.

As previously mentioned in section 2.1, styli must be light, stiff, strong and made from readily machinable material, amongst other traits. A wide range of materials fit most of these criteria, and so several of the most promising materials were used to produce a set of specimen styli for testing purposes. These materials were compared against the standard carbon-fibre stylus and against the styli that were used to test the range of adhesives in section 3.2.10.

Two styli were printed in UV-curable resin using a Formlabs Form2 (Formlabs, 2019), 5 were made with a ceramic main tube and an aluminium nosepiece and 5 was made using an aluminium main tube and separate aluminium

nosepiece. This limited the geometrical differences between the ceramic and aluminium styli. Other styli were attempted but failed at various stages of the production and testing process. The realised styli can be seen in Figure 5–5.

5.1. Failed Styli

Several other materials had potential as stylus materials but ultimately failed before testing could be done. Main tubes made from PEEK and Nylon were manufactured, although the geometry made the process difficult and time consuming. The specific stiffness of these materials is low enough that the main tubes were trivially easy to flex and deform, so these materials were abandoned.

Other materials were unsuitable because of the form factor in which they are available. Silicon Nitride and other high-density ceramics would need to be available in the form of 1 mm diameter rods to keep the mass and moment of inertia of the stylus within specification. The stylus design would need to fundamentally change to allow these exotic materials to be used.

5.2. Realised Styli

5.2.1. Resin

Two styli, made from UV curable resin, were printed using a Form2 SLA printer (Formlabs, 2019). The geometry of these styli was chosen to emulate the standard 112-3221 stylus and the tapered oval-cross-section of the PGI Novus styli. These styli were printed using the ‘fine’ setting in the Formlabs software, giving a layer height of 0.1 mm. They were both designed to be hollow, but the components were not properly drained before curing so they are both almost solid, save a couple of bubbles. This will actually reduce their stiffness and increase their weight, meaning their specific stiffness is lower than predicted using FEA. A picture of the 3221-like stylus is shown in Figure 5–1.

The resin styli both shrank during curing. This can be accounted for during the design/printing phase, but the shrinkage is dependent on several factors including; printing orientation, model density, cure time, support structure and

wall thickness, so it was simplest to allow for shrinkage during assembly. The short main tube could not be butted up against the spigot during assembly, instead adhered with a ~1mm offset, shown on the right-hand side of Figure 5–1.

There is also a noticeable curvature in the main tube. This is partially as a result of the curing process, although the assembly fixture also contributed to the shape change. Despite the curved main tube, the key dimensions of the stylus were correct (when the stylus was used in the experiments). The spigot of this stylus had to be modified, reducing the mating boss in length and diameter, because the main tube wall was printed thicker than specified in the 3D model.



Figure 5–1: A closeup of one of the additively manufactured styli. This stylus is dimensionally similar to the 112-3221; while the second stylus has a main tube cross section that is oval. The wire wrapped around the main tube is used to increase the moment-of-inertia, allowing the stylus to be balanced in the PGI Bearings gauge.

5.2.2. Ceramic

A 40 mm long, 3 mm OD, 2 mm ID alumina tube was used as a substitute for the PCF main tube in this stylus. The alumina is extremely difficult to cross-drill so a nosepiece from the tapered Novus styli was used as an intermediary between the ceramic and the recess-piece. This was assembled using the 10-minute DP-410 adhesive from 3M, in the standard assembly fixture developed in 3.2.8. One of the realised ceramic styli is shown in Figure 5–2.

The ID of the ceramic required that the spigot was turned down to allow clearance. This would require a simple change to the spigot geometry in the event that ceramic is used for production styli.

The general straightness of the ceramic tubes is worse than the aluminium or PCF tubes, likely because the ceramic is formed in an extrusion process and is not machined further, whereas the aluminium and PCF are both machined as part of production. The stylus assembly required the use of the fixture developed in section 3.2



Figure 5–2: A closeup of the ceramic stylus. This specific stylus has been assembled using a proxy tip, without a diamond. The ceramic could not be cross-drilled and so a PGI Novus stylus nosepiece was used.

5.2.3. Aluminium

The main tube was manufactured from 6061-T6 aluminium to match the dimensions of the 112-3221 stylus main tube (3.35 mm OD and 2.5 mm ID). Although it would be possible to machine a cross hole in the stylus, the simpler method of using a Novus nosepiece was again chosen in preference, ensuring comparability with the ceramic styli, saving time and machining effort. One of the realised aluminium styli can be seen in Figure 5–3.

The main tube of the Aluminium styli was machined from solid stock; this presented a set of manufacturing challenges that may be avoided if Aluminium styli were made in production, where appropriate tube can be sourced to reduce the manufacturing burden.



Figure 5–3: A closeup of an example of the Aluminium stylus. A 2 µm radius, 90° conisphere diamond tip is used on this specific stylus.

5.2.4. Pultruded Carbon Fibre

Several styli were manufactured as part of the adhesive testing process. These styli are, for all intents and purposes, standard 112-3221 styli. The only difference is the applied adhesive is 3M DP410 and the assembly was done in the adjustable stylus fixture developed. These styli have been calibrated as standard production styli and have been shown to balance in a standard PGI gauge. An example of the PCF styli is shown in Figure 5–4.

If the pultruded carbon fibre stylus is considered as a simple beam system, fixed at the gauge pivot, the natural resonance would be of the ~800 Hz. The actual stylus should resonate at a lower frequency as it is made up of several components, some of which are made from aluminium. Further, the joint between the spigot and the spigot tube is not perfectly rigid and will therefore reduce the effective stiffness of the stylus.



Figure 5–4: A closeup of a pultruded carbon fibre stylus, made during the adhesive testing experiments. These styli have been calibrated, and proved to measure repeatably, on a Form Talysurf instrument.

The mechanical properties for pultruded carbon fibre are shown in Table 2-1.

A side-by-side comparison between the resin, ceramic, aluminium and PCF styli is shown in Figure 5–5.

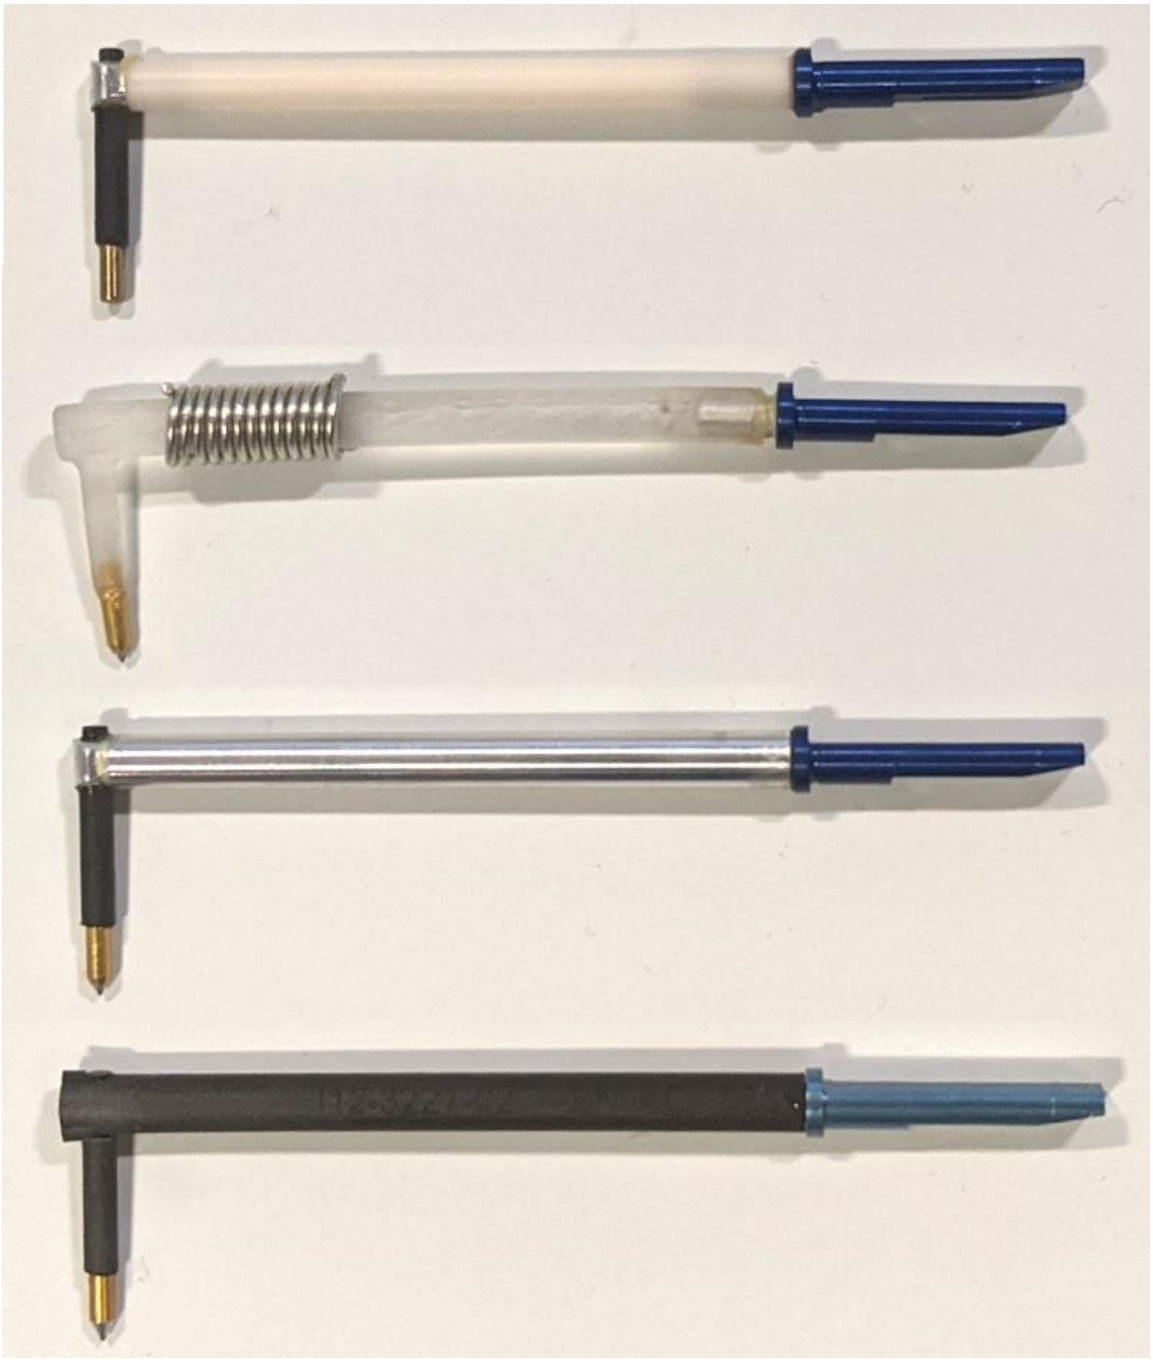


Figure 5–5: A comparison of the realised styli. From top to bottom: Ceramic, 3D printed resin, Aluminium and PCF.

5.3. Summary of Testing Methodologies

5.3.1. Mass and Inertia

Each stylus needs to be within a specified range of masses to ensure that it will be balanced properly when used in a standard PGI gauge. As the system is a rotating system, as opposed to a linearly traversing system, the moment of inertia is more important than the actual mass, but since the moment of inertia is directly related to the mass then the mass can be used as an indication of compliance.

Each stylus was weighed using a top-pan balance, and the tip force was recorded for each stylus when inserted into a standard PGI gauge that was set to give 80 mg force for a standard 112-3221.

5.3.2. Geometric

Each stylus is assembled using the adjustable fixtures, under the AV300 measurement microscope. This ensures that any discrepancy in individual component geometry does not have a knock-on effect on the final assembly geometry. After assembly, the adhesive is allowed to cure and then the styli are re-measured using the standard measurement fixtures and routines. This ensured that the styli do not warp during the assembly or curing processes. The assembly process for the styli requires the AV300 to verify the correct geometric relationship, within normal stylus tolerances.

Each stylus tip is 42.25 mm ±50 µm from the datum on the spigot, along the axis of the main tube (Length constraint). The tip is within 50 µm of the central axis of the spigot (Symmetry constraint) as shown in Figure 5–6. The false tips are made to be a close analogue to the real measurement tips in terms of mass and form; however, they are shorter than actual measurement tips by

~0.5 mm, meaning styli assembled using false tips are short. This does not appear to change the dynamics of the stylus in a significant way.

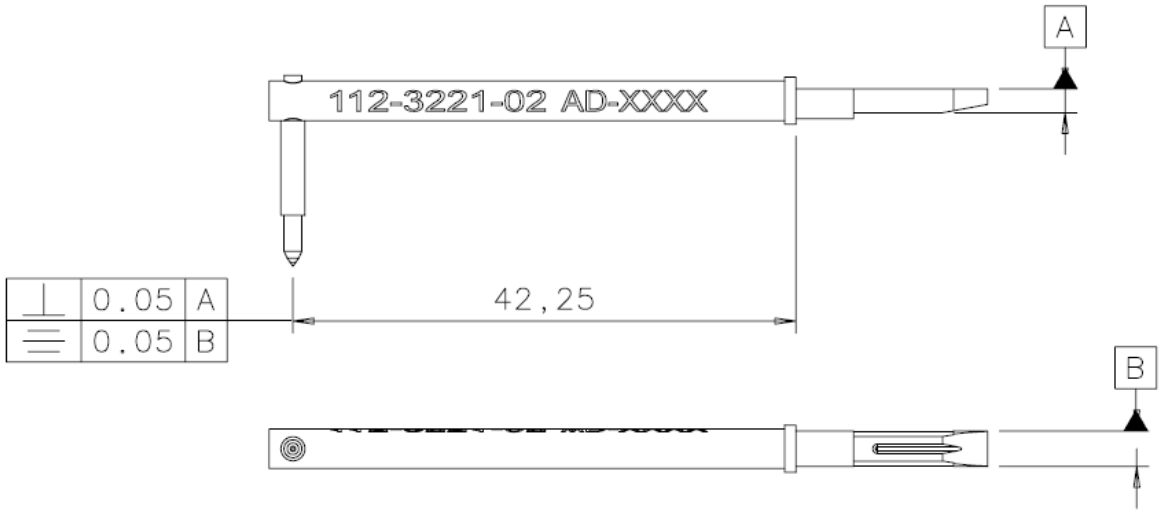


Figure 5–6: A diagram showing the constraints and datums used to define the geometry of a standard 112-3221 stylus.

5.3.3. Vibration

Each stylus was tested using either a standard PGI gauge or a PGI Bearings gauge[8], on the stylus dynamics gantry testing machine. The PGI Bearings gauge allows styli with a higher moment of inertia to be properly balanced. This process takes the stylus through a range of surface frequencies, beyond the point at which the stylus leaves the surface. If the natural resonant frequency of the stylus is below the stylus flight frequency, then this will be visible in the measured profile. If the stiffness of a stylus is low enough then the peak of the resonance can be seen before the stylus leaves the surface. Stiffer styli leave the surface before they reach their natural resonance peak and so estimates must be made instead. An example of the profile for a PCF stylus is shown in Figure 5–7, highlighting the inferred resonance peak and the actual stylus flight frequency.

[8] The PGI Bearings gauge is a modified version of the standard PGI gauge to allow longer styli to be used. The modifications make the gauge well suited to measuring form and roughness in bearing applications.

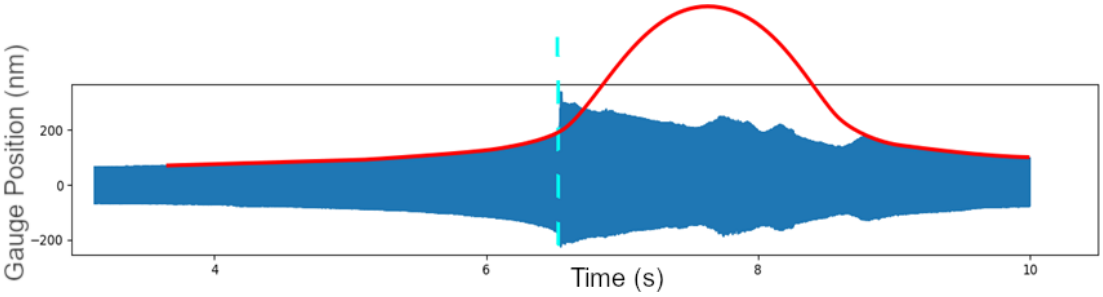


Figure 5–7: A closeup of a profile for a PCF stylus, showing a potential resonance peak (Red) and the stylus flight frequency (Light Blue).

The frequency at which the stylus leaves the surface is proportional to the applied stylus force and the amplitude of the actuation, amongst other things. The correlation between tip force and stylus flight frequency could be used to suggest that the tip force should be increased; this would only result in more damage being caused by the stylus tip. The stylus flight frequency can be used as a performance indicator if the stylus force is correctly set at 80 mg but becomes invalid if the stylus force higher than recommended.

The gantry system was also qualified using a large set of vibration tests, run over a period of 10+ hours. This allowed the actual performance of the gantry to be evaluated and helped to adjust the testing parameters of future experiments. A detailed description of the gantry testing is described in section 5.4.3

5.4. Range of Experiments Conducted on the Gantry System

The experiments undertaken on the gantry system are to test several hypotheses regarding the performance of the stylus, with relation to the material of the stylus main tube. To achieve this, the equipment used to test these hypotheses must be characterised such that meaningful conclusions can be drawn from the results.

The design and assembly of the gantry system has been covered, including the measurement of the dynamic properties of the realised gantry and the comparison to FEA simulation. The low noise-floor and excellent stability of the gantry provides a good starting point for the investigation into the dynamic performance of the selected range of materials.

Figure 5–8 shows a schematic overview of the gantry system, including the analysis steps that are takes after the data is collected. The shaded regions denote areas of the gantry which have been characterised in other sections of this thesis, and therefore do not need to be re-investigated. For the gantry results to have known uncertainty values, the remaining areas must be investigated and as such, a number of hypotheses were put forward to drive the testing methodology.

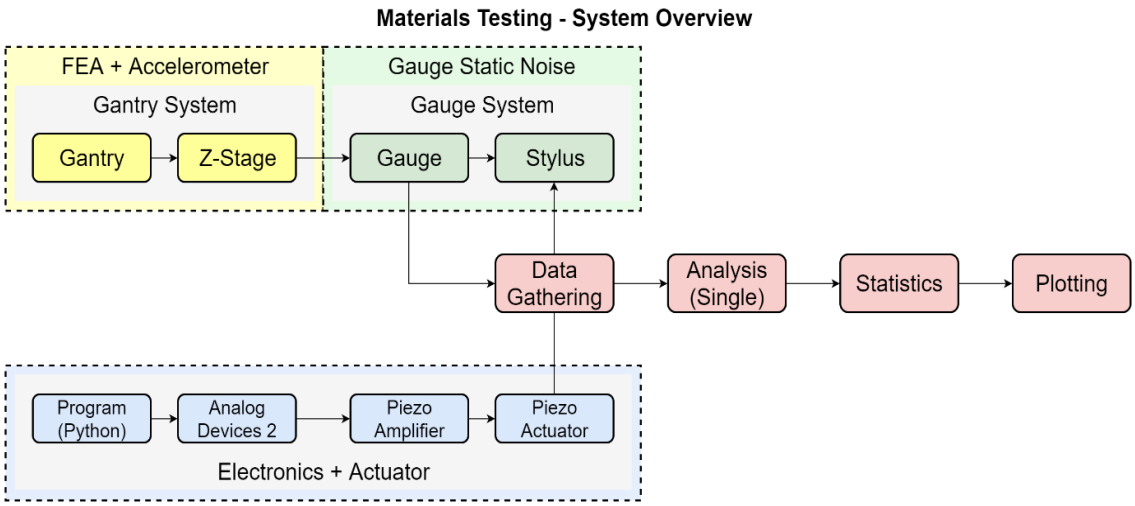


Figure 5–8: A schematic overview of the testing system used in this report. Shaded boxes denote the areas of the system that have been characterised previously.

5.4.1. Hypotheses

1. The physical, electronic and gauging elements of the gantry system are suitable for high frequency dynamic testing of styli.
2. The dynamics of a stylus can be usefully characterised on the gantry system.
    a. The measurement and analysis can be performed automatically after the stylus is properly set into the gauge.

5.4.2. Experiments

These hypotheses helped to define the large experimental parameter space. The initial experimental work was done to reduce or characterise each parameter space individually, to reduce the number of actual experiments needed for a robust under-standing of the stylus dynamic performance.

Experimental group 1

These experiments aimed to establish an understanding of the performance of the gantry system, over a range of experimental parameters. Where possible, each parameter was decoupled and evaluated separately, such that any experimental un-certainty can be properly apportioned.

1. Standard PGI Gauge, PCF stylus – Hypothesis: The gantry and analysis system will have a statistically characterizable performance, potentially related to the settling delay between subsequent measurements.
2. Standard PGI Gauge, PCF stylus – Hypothesis: The amplitude of the actuation input signal will affect the resulting analysis; that there is an optimal amplitude that can be used for future testing.
3. PGI Bearings Gauge, Ceramic Stylus – Hypothesis: The performance of the ceramic stylus, in the PGI Bearings gauge, will be broadly similar to the PCF stylus in the standard PGI gauge (as experiment 1). The effect of the range of amplitudes will be similar for the ceramic stylus in the PGI Bearings gauge, to the PCF stylus in the standard PGI gauge (as experiment 2).

4. Standard PGI Gauge, PCF stylus – Hypothesis: The waveform shape of the actuation input signal will have a significant effect on the measured stylus dynamic spectrum; that there is an optimal choice for actuation waveform for future experiments.

Experimental Group 2

These experiments aimed to establish the effect of several practical experimental factors on the repeatability of the gantry system. This helps to account for the un-certainty introduced when a new stylus is inserted into the gauge.

5. Standard PGI Gauge, PCF stylus – Hypothesis: The effect of the simple re-insertion of the stylus into the gauge is measurable.
6. Standard PGI Gauge, PCF stylus – Hypothesis: The effect of the re-insertion and spring force adjustment of the stylus into the gauge is measurable and expected to have a larger influence than found in experiment 5.

Experimental Group 3

These experiments are the actual materials testing experiments, using optimal parameters from experiment group 1 and with the context of the comparability between different stylus measurements, gained in experimental group 2.

7. Standard PGI Gauge, PCF stylus – Hypothesis: The dynamic performance of the PCF stylus is observable using the gantry system and analysis techniques.
8. PGI Bearings Gauge, Ceramic stylus – Hypothesis: The dynamic performance of the ceramic stylus is observable using the gantry system and analysis techniques. The results will be comparable with experiment 7.
9. PGI Bearings Gauge, Aluminium stylus – Hypothesis: The dynamic performance of the Aluminium stylus is observable using the gantry system and analysis techniques. The results will be comparable with experiment 7 and experiment 8.

5.4.3. Discussion of Experiments

As mentioned previously, the range of available tests that can be performed on the gantry system is large and multidimensional. The vibrational performance of the gantry has been previously established in Section 4.3.5. This section describes each experiment and discusses some of the key findings to give context and structure to the experimental results detailed in section 5.4.4.

Experiment 1.1

This test is to determine the best possible performance that can be expected form the gantry, therefore, the gauge, stylus, gantry, actuator amplitude, analysis technique and inter-measurement delay were all maintained, while a large number of repeats were performed (Experiment 1.1).

Experiment 1.2 and 1.3

The large variance of the results over time lead to two further experiments, testing if a delay between consecutive measurements would allow the gantry system to recover (Experiment 1.2 and Experiment 1.3). A delay of 30 seconds was chosen for future work as it gave the best improvement without such a serious impact on the consecutive measurement rate. A 30 second delay between consecutive measurements gives a measurement period of ~57 seconds, meaning 100 measurement results took just over 90 minutes to obtain. This data was also used to determine the optimal number of repeats per stylus test.

Experiment 2 and 3

The amplitude of the actuation is suspected to have a significant effect on the stylus flight frequency and the ease of the subsequent data analysis. This relationship was characterised by changing the amplitude of the actuation for each subsequent measurement, whilst performing a large number of repeats such that the general trend of the data can be mapped over time. Five actuation levels were chosen, across the range of suitable actuation amplitudes, 1 V, 2 V, 3 V, 4 V, and 5 Volts peak-to-valley.

Two experiments were performed to characterise the two, gauge-stylus systems that would be used for the subsequent experiments, (Experiment 2 and Experiment 3). Experiment 2 used the standard PGI gauge and a commercially produced PCF stylus (112-3221), with 100 repeats of the 5 actuation amplitudes.

Experiment3 used the Bearings PGI gauge and one of the styli with the ceramic main tube, produced for materials testing; this used the same 5 actuation amplitudes as Experiment 2 for a total of 1250 measurements. Experiment 3 deliberately collected more data so as to give comparability with experiments 1.1, 1.2 and 1.3, and a view on the overnight performance of the gauge.

Experiment 4

Discussion of the comparability of the gantry results to real-world performance lead to the consideration of the waveform that is used in the gantry system. For all previous testing, a sine wave had been used as the underlying form of the actuation signal. For experiment 4, a triangle waveform was used to emulate the available Halle standards. The higher range of spectral components in the actuation gave several benefits over the pure sine waveform, such as being able to investigate frequencies at higher values than the cut-off point in the FFT. These benefits were outweighed by the increased sensitivity of the analysis process and so all future experiments used the sinusoidal waveform shape.

Experiment 5 and 6 - Context

Testing of multiple styli on the gantry would require the repositioning of the gauge, using the Z stage on the gantry, the removal of the previously tested stylus, the insertion of the new stylus, resetting the stylus force and finally positioning the gauge to contact the piezo actuator for the next measurement cycle. The combined effect of these operations is likely to have a significant effect on the repeatability of the gantry system. Two experiments were performed to quantify the effect of the reinsertion (Experiment 5) and the

reinsertion combined with the resetting of the stylus force between tests (Experiment 6).

Experiment 5

Experiment 5 tested the variance introduced into the gantry system when the same stylus was removed and reinserted. This experiment consisted of a single stylus being inserted into the gauge, the tip force set, and the gauge correctly positioned over the piezo actuator. After this setup, the gauge and the gantry were not adjusted until all repeats had been completed. Five repeats of 100 measurements were per-formed, using the previously found optimal parameters, over the course of 2 days. This experiment helped to characterise the best-case performance for testing multiple styli on the gantry system.

Experiment 6

Experiment 6 was an extension of Experiment 5, modified to include the resetting of the stylus force which necessitates the adjustment of the Z stage between each re-peat. The variance observed in this experiment is as a result of the gauge adjustment and gantry adjustment and is much more in-line with the results observed in the other subsequent experiments.

Experiment 7, 8 and 9

With the characterisation of the gantry and gauge systems complete, the actual testing of the range of stylus materials could begin. This started with Experiment 7. Experiment 7 used the standard PGI gauge on the gantry to test 6 of the 7 standard PCF styli (112-3221). The excluded stylus was used of the previous testing (Experiments 1 through 6) and so was excluded from subsequent experiments to avoid cross-contamination of results.

Experiments 7, 8 and 9 were all conducted using the same experimental conditions; namely:

- Sinusoidal waveform shape, 3 Volts peak-to-valley
- Frequency from 10 Hz to 1000 Hz, over 10 seconds, in a linear sweep

- Gauge centred within range to better than 10 µm
- 100 repeats, taking ~23 seconds per measurement, with 30 seconds delay between consecutive measurements.
- Tip force set to between 750 and 850 µN

Experiment 8 used the PGI Bearings gauge and the ceramic styli, under the same conditions as Experiment 7.

Experiment 9 used the PGI Bearings gauge and the aluminium styli, under the same conditions as Experiment 7.

5.4.4. Summary of Experimental Results

Table 5-1 summarises the experiments, along with the key finding of each. A detailed discussion of the results, along with a presentation of the gathered data is in section 5.4.5.

Experiment	Investigation	Key findings
1.1, 1.2, 1.3	The best possible performance of the gantry and analysis, and the effect of including some settling time between measurements.	The analysis technique was improved using the large set of data, and optimised and automated for bulk analysis of said data. The delay between subsequent measurements was set for future experiments at 30 seconds. The repeat performance of the gantry was measured and the number of repeats per stylus was set at 100.
2	The effect of amplitude on the stylus flight point of a PCF stylus. The change of stylus flight amplitude over time, when including a range of actuator amplitudes.	Of the 5 amplitudes tested, the best compromise was found to be 3 Vp-v. This avoided the clipping and masking issues of the 5 Vp-v and, unlike the 1 V amplitude, reliably induced stylus flight into the measurement. The analysis technique worked more reliably on higher amplitude measurements.
3	The effect of amplitude on the stylus flight point of a ceramic stylus. The change of stylus flight amplitude over time, when including a range of actuator amplitudes. This also determined the baseline performance of the PGI Bearings gauge.	The performance of the gauge was shown to be acceptable for gantry measurements, and the stiffness of the ceramic stylus was confirmed to be within the range of stiffness of a standard PCF stylus. The change of flight amplitude vs actuator amplitude was similar to experiment 2 suggesting comparability between the standard PGI Gauge and the PGI Bearings gauge

4	The spectral response of the stylus, comparing the triangular and sine actuation waveform shapes.	The triangular waveform allowed the investigation of frequencies above the stylus flight frequency. The triangular waveform proved to be useful for manual analysis, although it cannot be used with the previously established analysis techniques without extensive modification.
5	The effect of reseating the stylus on the measured stylus flight frequency and the spectral performance of the stylus over time.	The repeatability of reseating the stylus was measured, producing 5 sets of measurements. Of these measurements, four were tightly grouped (within 2% of full scale) such that the individual means were within 1 standard deviation of the group mean. The fifth measurement was 1.4 standard deviations away from the group mean; while this deviation is small compared to the observed results for experiment 6, it still demonstrates that reseating a stylus is not a completely predictable process.
6	The effect of reseating the stylus and resetting the tip force between each set of repeats. The change in stylus flight frequency and the spectral performance of the stylus over time.	This experiment resulted in a similar set of results as experiment 5, with a much wider distribution between each repeat. No mean value, of any set of data, was within 1 standard deviation of another set of data. This suggests that one of the extra steps involved in resetting the stylus tip force is responsible for the increase in spread of the results, when

		compared to experiment 5. It also allows a large inter-stylus deviation to be explained for other experiments.
7.1 – 7.6	The performance of the PCF styli, under a set of standard experimental conditions, using the standard PGI gauge. The stylus flight frequency and the spectral performance were investigated.	The measured spectrum included some frequencies that were predicted from unrelated work, such as the air conditioning hum. Some previously unobserved frequency components of the PGI gauge were observed in the spectrum data. The results are discussed in greater detail in the text.
8.1 – 8.5	The performance of the ceramic styli, under standard experimental conditions, using the PGI Bearings gauge. The stylus flight frequency and the spectral performance were investigated.	The ceramic styli had a distinct frequency spectrum, compared to the PCF styli. The ceramic styli typically started stylus flight at a lower frequency, although often experienced a region of stability after an initial stylus flight. This was visible in the profile but made very obvious in the spectrogram of the experiment.
9.1 – 9.5	The performance of the aluminium styli, under standard experimental conditions, using the PGI Bearings gauge. The stylus flight frequency and the spectral performance were investigated.	The aluminium styli performed similarly to the ceramic styli, often having a period of stability between two regions of stylus flight. In general, the performance of the aluminium stylus was lower than the ceramic styli, as predicted by their relative stiffnesses.

Table 5-1: A table giving a summary of the range of experiments performed and the key findings of each.

5.4.5. Experimental Results

Gantry Background and Parameter Selection

The gantry system has been shown to have a very low noise floor, with only a single obvious vibrational mode below 1000 Hz. Gauges mounted on the gantry experience non-rigid-body vibration with an amplitude no higher than 0.5 nm, at frequencies between 20 Hz and 1000 Hz (See Section 4.3.5). Styli are susceptible to airborne noise, typically caused by the air conditioning system at roughly 12 Hz. For this reason, the lowest excitation frequency is set at 10 Hz.

The stylus is excited using a preloaded piezoelectric actuator, capable of sub-nanometre positioning at frequencies greater than 2000 Hz. The amplitude of the actuation is limited by the power dissipation of the piezo stack, although this amplitude limit is beyond the range necessary for the proper testing of the PGI systems. Practically, the actuation amplitude needs only to be ~200 nm peak to valley, as any larger amplitudes begin to introduce distortion. Unless multiple amplitudes were used in a single test, all experiments were conducted using an excitation 3.0 Volts peak-to-valley, corresponding to an amplitude of ~180 nm peak-to-valley. This voltage amplitude is output from the Analog Discovery 2 function generator (labelled point A in Figure 5–9).

At this amplitude, the stylus-gauge system appears to respond well up to frequencies of ~1000 Hz. This high frequency lies beyond the theoretical point at which the stylus would leave the surface for any typical surface amplitude. High surface frequencies are also supressed by a pre-filter which is applied to surface measurements before analysis (this filter has been disabled in this testing). Further wavelength filtering is typically applied during analysis $(\lambda_c \; and \; \lambda_s)$ also written as Lc and Ls, as well as a tip radius correction algorithm (morphological filter). The Lc filter is a long-wavelength reject filter (typically 0.8 mm) while the Ls is a short-wavelength reject filter (typically 2.5 µm). 1000 Hz was chosen as a reasonable upper limit to the actuation frequency, as this encompasses the filtering and pre-filtering for any sensible surface measurement.

Gantry Testing Schema – Experiment 1.1, 1.2 and 1.3

As mentioned in the discussion of experiments (Section 5.4.3), the gantry system as a whole must first be characterised. The testing of the suitability of the different stylus materials uses the stylus flight frequency, and the overall frequency spectrum, as a performance indicator and therefore the baseline stability and repeatability of the gantry must be established. This was done using a known-good measurement stylus (112-3221), at an output amplitude of 3 V, with a frequency sweep from 10 Hz to 1000 Hz over 10 seconds, with a sinusoidal waveform shape. This test was repeated every 22 seconds between 14:30 PM and 11:30 AM the following day, yielding approximately 3100 results. This is shown in Figure 5–9.

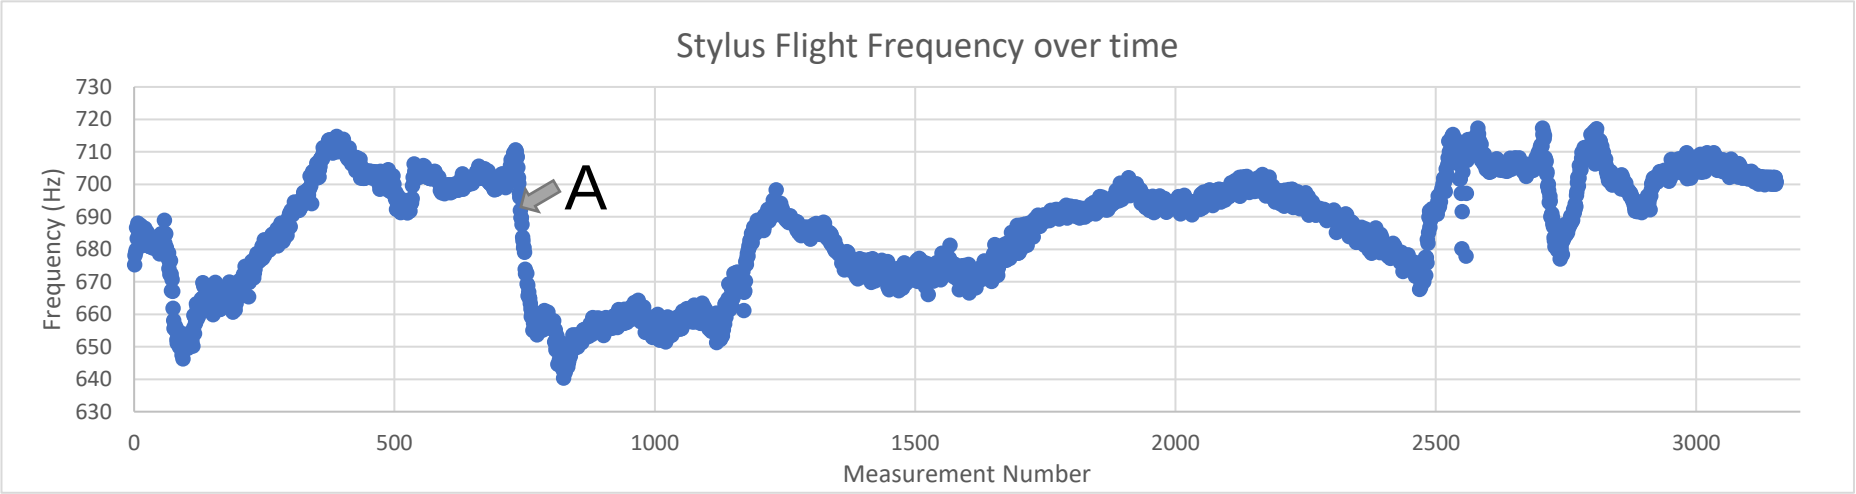


Figure 5–9: A graph of the stylus flight frequency over time.

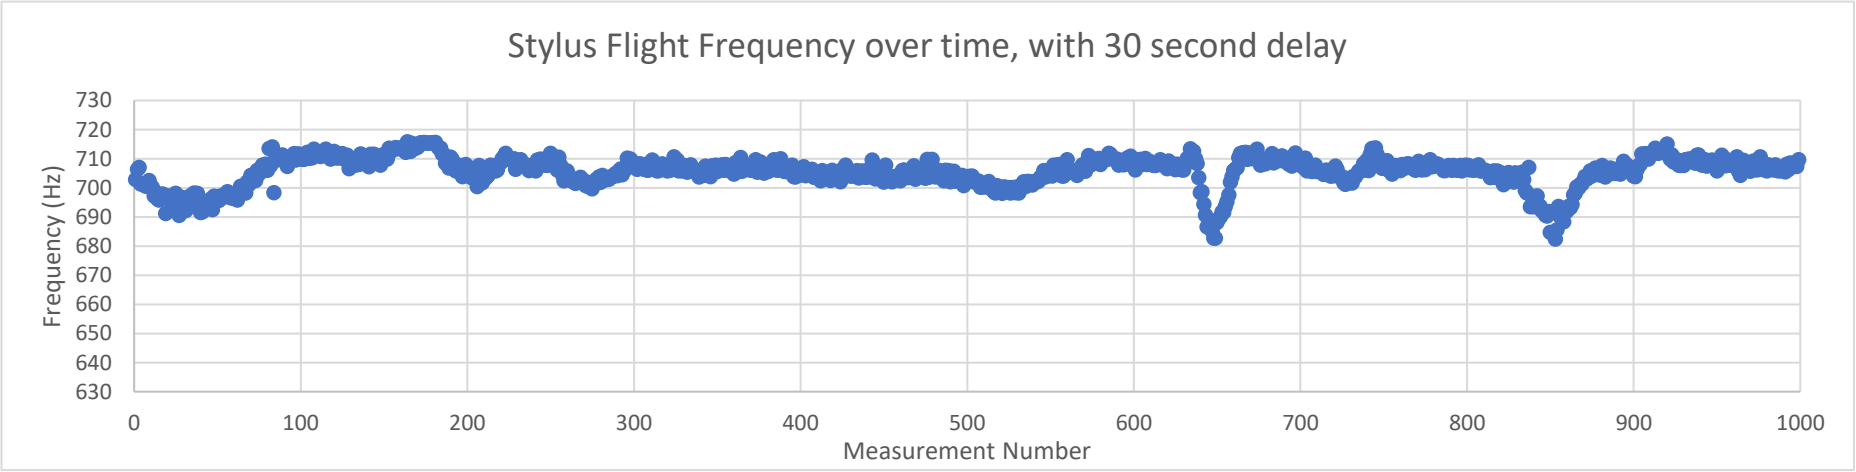


Figure 5–10: A graph showing the stylus flight frequency over time, with each measurement having a 30 second delay beforehand. Note the scale is identical to Figure 5–9 for comparison.

The results observed in the initial testing phase (Figure 5–9) have several interesting features. Firstly, the range of stylus flight frequencies is 76 Hz. This is far outside the expected deviations observed during the non-automatic testing of the gantry system. Secondly, the steep rise and fall of the measured flight frequency is suspicious. Gradual changes are generally observed to happen over the timescale of hours, whereas the steepest transition shown in this dataset is 710 Hz to 653 Hz in 14.5 minutes. This transition starts at 20:53PM, labelled A in Figure 5–9. This time in the evening does not correspond to any known event within the surrounding building.

A delay of 30 seconds was introduced between each consecutive measurement, increasing the time per measurement to 54 seconds. This resulted in the graph shown in Figure 5–10. This graph is presented at the same scale as Figure 5–9 to aide comparison. The increased relaxation time of the system seemed to drastically improve the stability. There are also significantly fewer deviations from the mean, and when deviations do occur, they are 'corrected' more quickly. The delay was increased to 60 seconds to see if further improvement would be seen. The system did not appear to be significantly more stable, and the further delay would greatly increase the testing time. The graph of the gantry performance with the 60 second delay is shown in Figure 5–11.

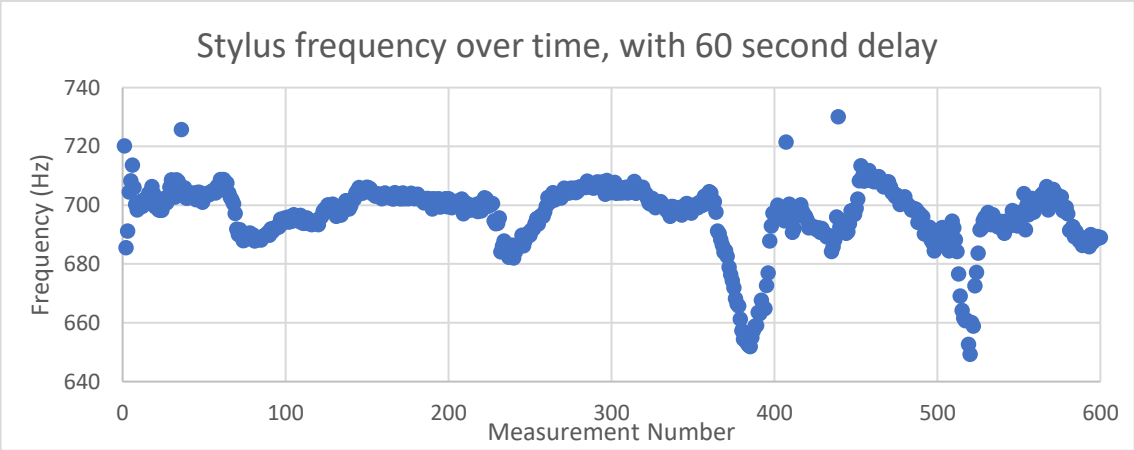


Figure 5–11: A graph showing the stylus flight frequency over time, with each measurement having a 60 second delay beforehand. Note the scale is identical to Figure 5–9 for comparison.

The mean and standard deviation for each delay time is shown in Table 5-2.

	Mean (Hz)	Std. Dev (Hz)	Delay (s)	Time between measurements (s)	Number of data points
Expt 1.1	686.5	17	0	22	3150
Expt 1.2	705.6	5.2	30	53	1000
Expt 1.3	696.4	10.8	60	85	600

Table 5-2: A Table summarising the Mean and Standard Deviations of experiments 1.1, 1.2 and 1.3.

The reason for the improved performance of the second experiment is not obvious, although it seems reasonable to conclude that a small delay would likely improve the stability of the overall experiment. The fact that a larger delay leads to a reduced stability may be explained by drift in the gantry system or an environmental influence, but the source of this increased instability is unknown.

Effects of surface amplitude of stylus flight – Experiment 2

The next test that was performed was to investigate the performance of the stylus-gauge system over a given range of amplitude, whilst observing the changes over time. This will help to categorise if a specific amplitude is more stable over time.

The experiment consisted of a single frequency sweep at each amplitude (1 V, 2 V, 3 V, 4 V and 5 V), with a standard delay of 30 seconds between each measurement. These 5 frequency sweeps are repeated 100 times over ~8 hours.

The collected and analysed data is graphed in Figure 5–12. There is a good degree of correlation between the different amplitudes and their flight frequencies. The correlation continues as the flight frequency changes, although the 1 V flight frequency appears to deviate the least from the mean.

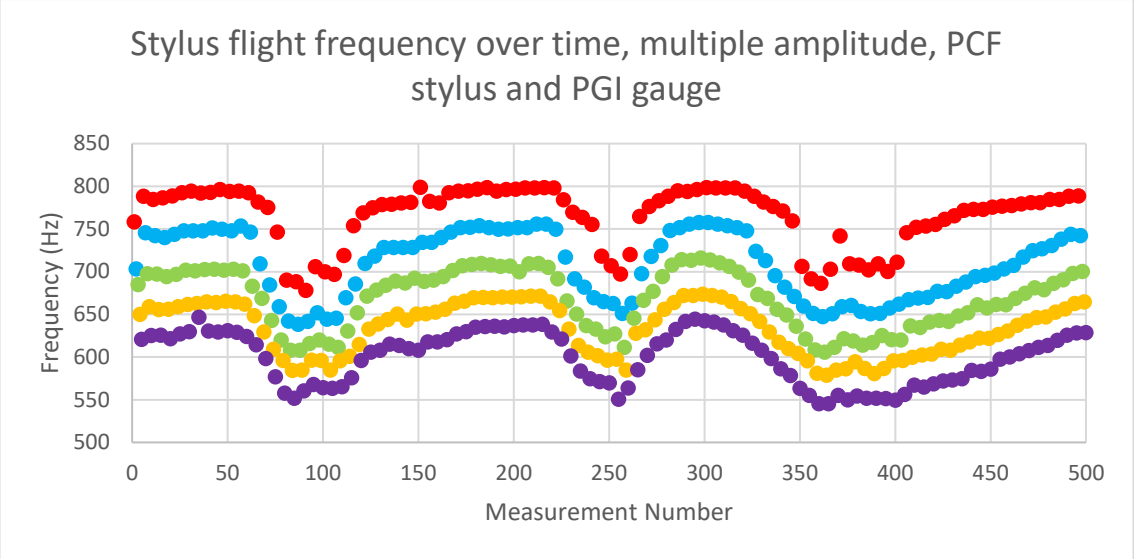


Figure 5–12: A graph of the stylus flight frequency over time, using 5 different amplitudes of actuation. 1V (Red), 2V (Blue), 3V (Green), 4V (Orange) and 5V (Purple)

In Figure 5–13, a line has been fitted to show the highest natural frequency that the stylus leaves the surface. The highest amplitude (5 V, purple) and the lowest amplitude (1 V, Red) are shown for comparison. The shaded areas represent the time and amplitude of the deviation from the modal performance. The 1V results return to the modal line one loop before the 5 V and leaves it 3 sets (of five measurement amplitudes) afterwards (labelled A and B, respectively).

The change in amplitude is more extreme for the 5 V amplitude than the 1 V amplitude.

These observations fit with the expectations of a dynamic system.

- Higher amplitudes of actuation are less stable than lower amplitudes. (Shorter period of stability)
- Higher amplitudes are affected more by changes in the system. (Faster decay, longer recovery)

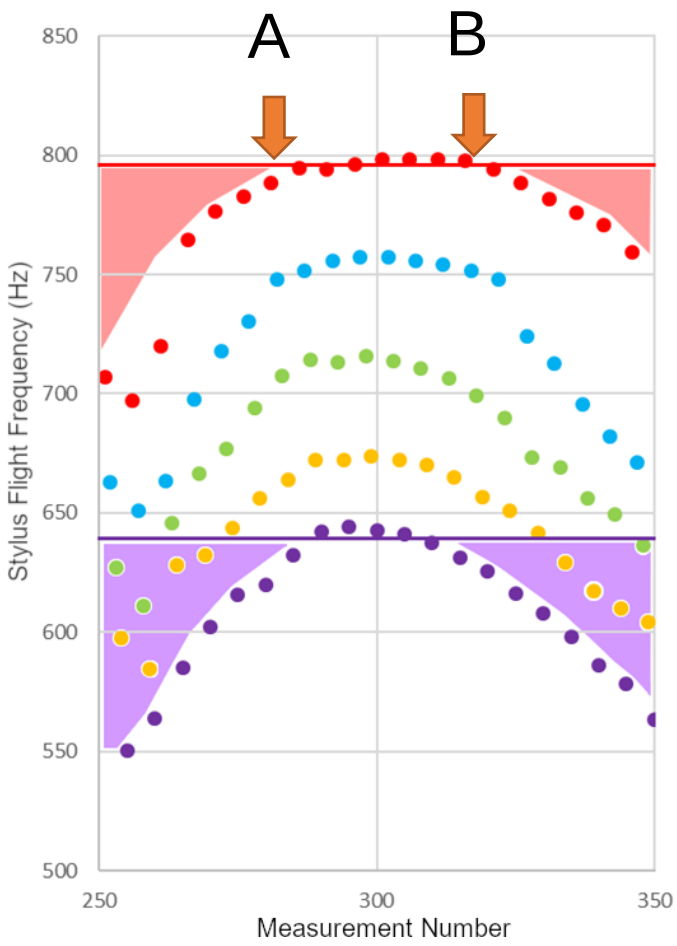


Figure 5–13: A close-up section of Figure 5–12 showing a fitted line for both 1V (Red) and 5V (Purple) amplitude.

A significant takeaway from the initial experiments is the choice of 3 V as a reasonable amplitude for actuation in future experiments. It gives a stylus flight frequency of ~700 Hz for a 60 mm PGI stylus like the 112-3221. This amplitude also corresponds to ~200 nm peak to valley triangle wave which is comparable to real surface finish measurements.

Note on the systematic error observed here

The deviation from the modal stylus flight frequency (as seen in Figure 5–12) is problematic and care must be taken to mitigate this error in other experiments. The magnitude of the observed effect is dependent on the absolute value of the stylus flight frequency. There was no obvious cause of this error, making mitigation ahead of time difficult. This error is an example

of a systematic error inherent to the measurement system and analysis procedure. The periodic nature of the error is unusual, although the magnitude of the effect appears to be similar during each epoch. There are several known errors and observed phenomenon, see Section 4.3.7 for further examples.

Effect of surface amplitude on stylus flight – Ceramic stylus – Experiment 3

Experiment 2 demonstrated the variance in stylus flight amplitude against the input actuation amplitude, for a PCF stylus on a standard PGI gauge. Experiment 3 is an extension of this experiment to include the ceramic stylus and the PGI bearings gauge. This is done for two reasons; first, to characterise the baseline performance of the PGI Bearings gauge, and second, to test the comparability between the ceramic and the PCF styli, under a range of actuation amplitudes. The results of this experiment can be seen in Figure 5–14.

The ceramic stylus and PGI Bearings gauge combination generally had a lower stylus flight frequency. The relationship between the amplitude of actuation and the stylus flight frequency is similar to the relationship found in Experiment 2, albeit with a few unique features.

Broadly, two phenomena were present in the ceramic stylus experiments that were not observed in Experiment 2 (Labelled A and B in Figure 5–14). These phenomena are related to the analysis procedure and to the dynamic performance of the ceramic stylus. The reduction of observed phenomena in the results for 3V actuation amplitude and above support the choice of 3 V peak-to-valley as appropriate for materials testing.

Profiles and spectrogram plots are shown for both the 1V (Red on Figure 5–14) and 2V (Light Blue on Figure 5–14) amplitudes, for both phenomenon A and B. These are shown in Figure 5–15, Figure 5–16, Figure 5–17 and Figure 5–18, respectively. The phenomena are discussed after the profiles are presented.

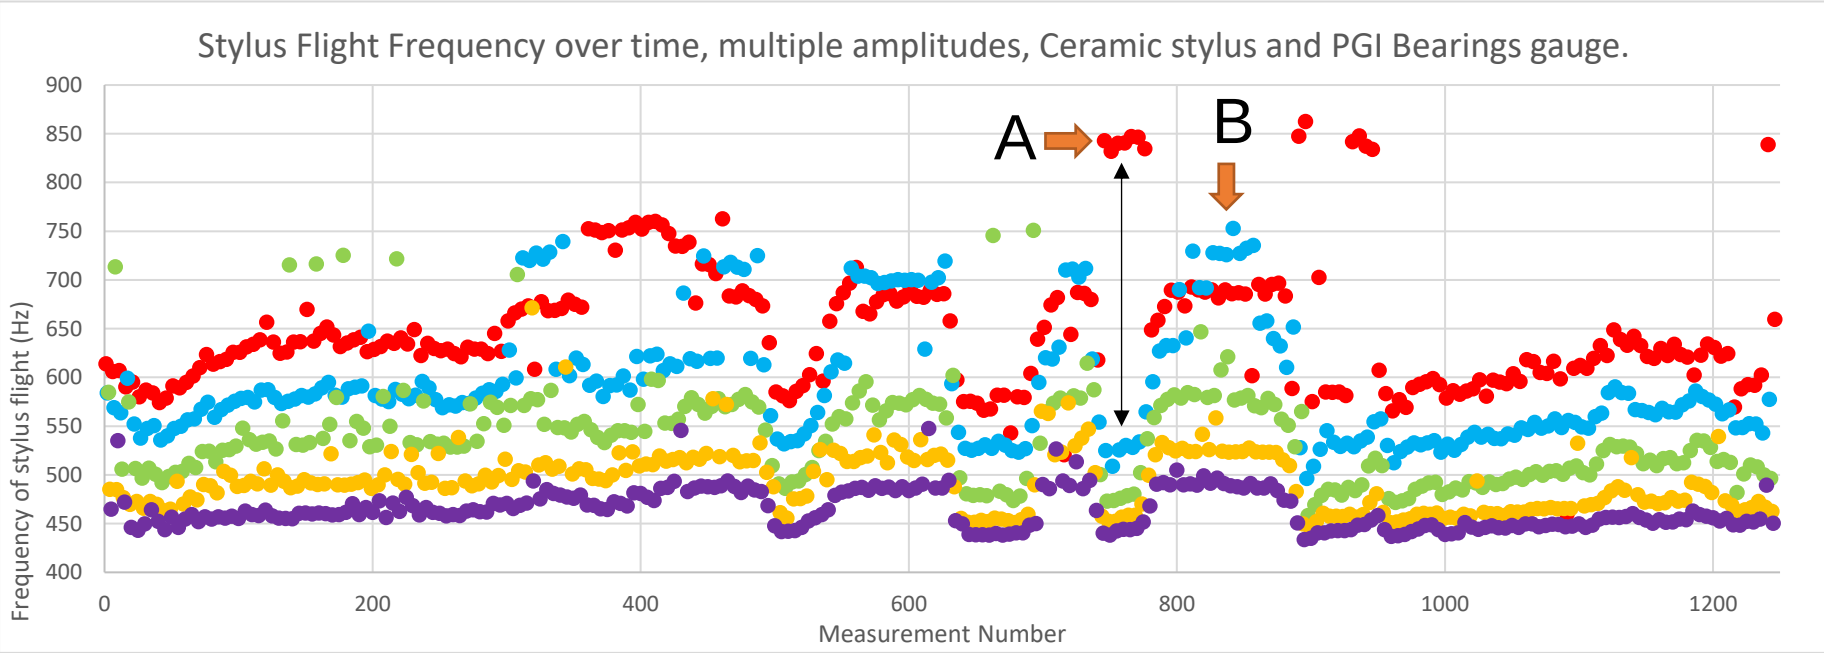


Figure 5–14: A graph of the stylus flight frequency over time, using 5 different driving amplitudes: 1V (Red), 2V (Blue), 3V (Green), 4V (Orange) and 5V (Purple). The ceramic stylus and PGI Bearings gauge were used for this experiment. Two phenomena specific to the ceramic styli are highlighted A and B, the cause of which will be discussed later.

The large deviation (Labelled A in Figure 5–14), is caused by small changes in the thresholding level and the resultant large shift in stylus flight frequency. This effect is discussed further on page 175.

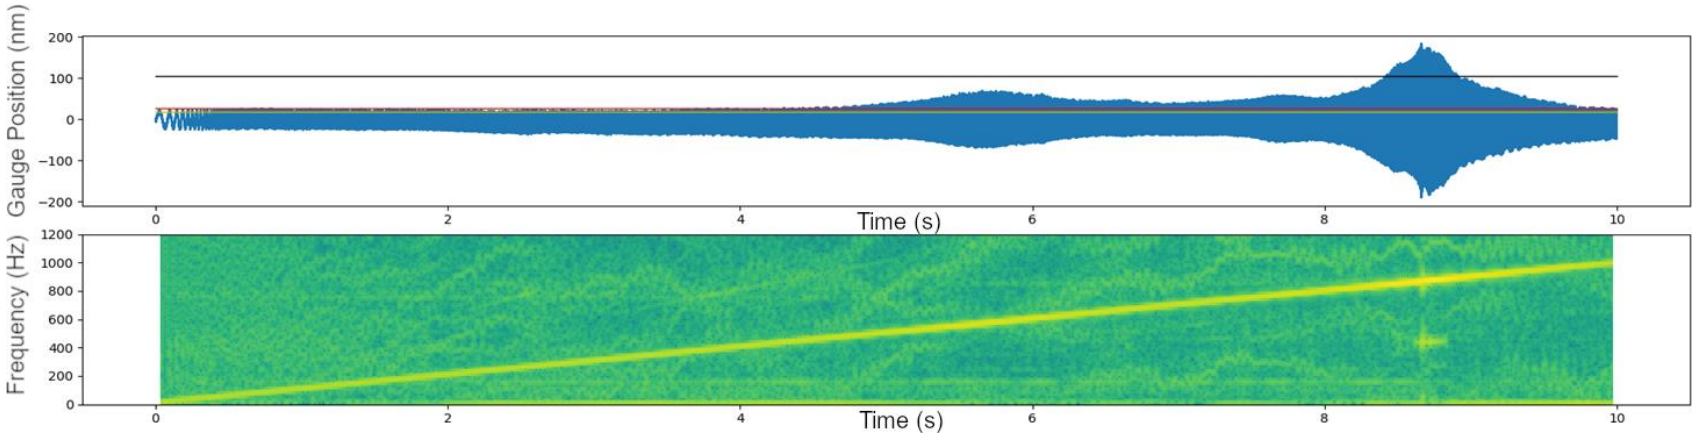


Figure 5–15: The profile and the spectrogram of the 1V amplitude data point highlighted as phenomenon A in Figure 5–14. The lack of wide vertical features in the spectrogram suggests that the stylus only leaves the surface for a short period at ~860 Hz.

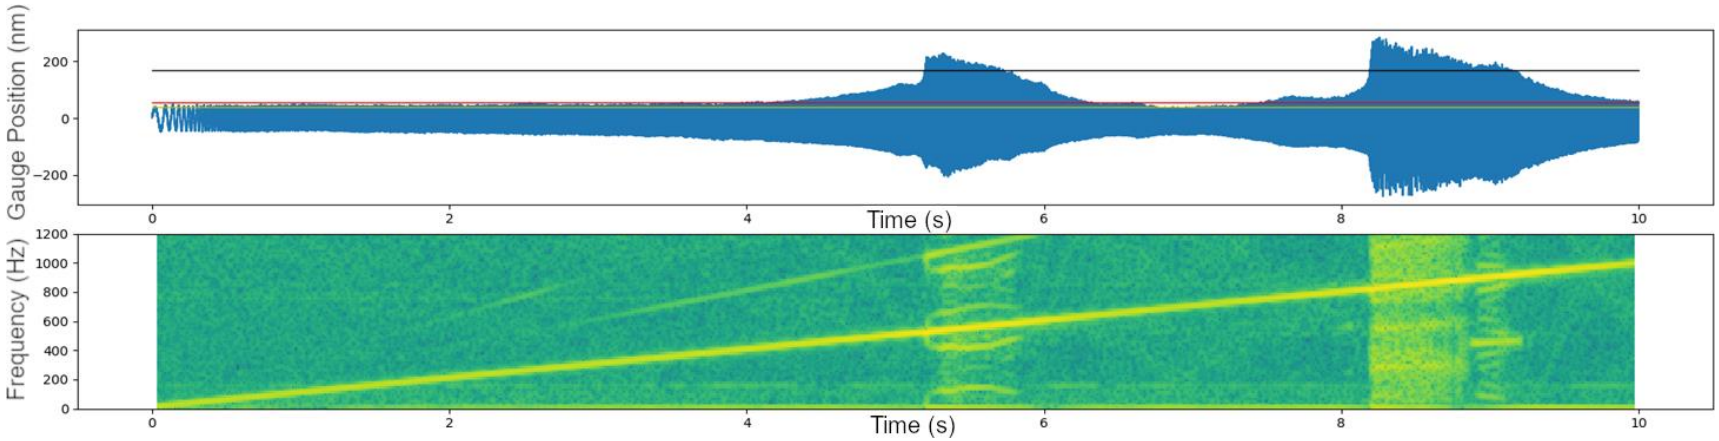


Figure 5–16: The profile and the spectrogram of the 2V amplitude data point highlighted as phenomenon A in Figure 5–14. The stylus leaves the surface at two separate points during this test (~525 Hz and ~825 Hz).

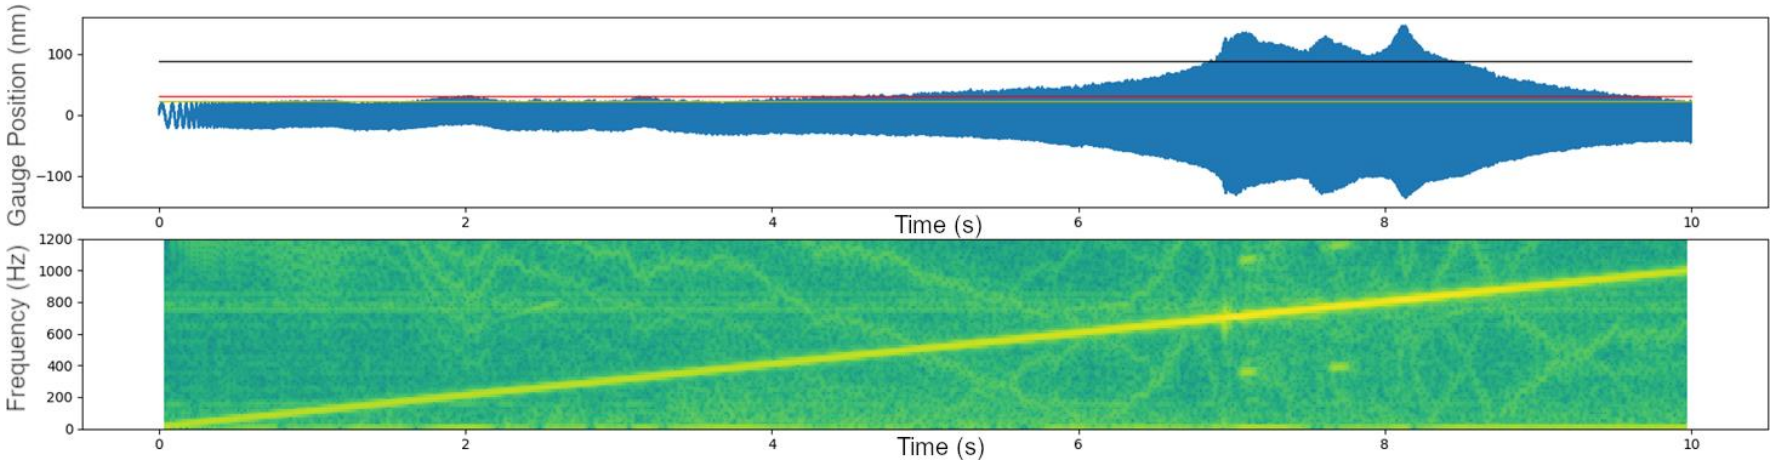


Figure 5–17: The profile and the spectrogram of the 1V amplitude data point highlighted as phenomenon B in Figure 5–14. The stylus appears to only leave the surface for a short period of time, at only a few frequencies.

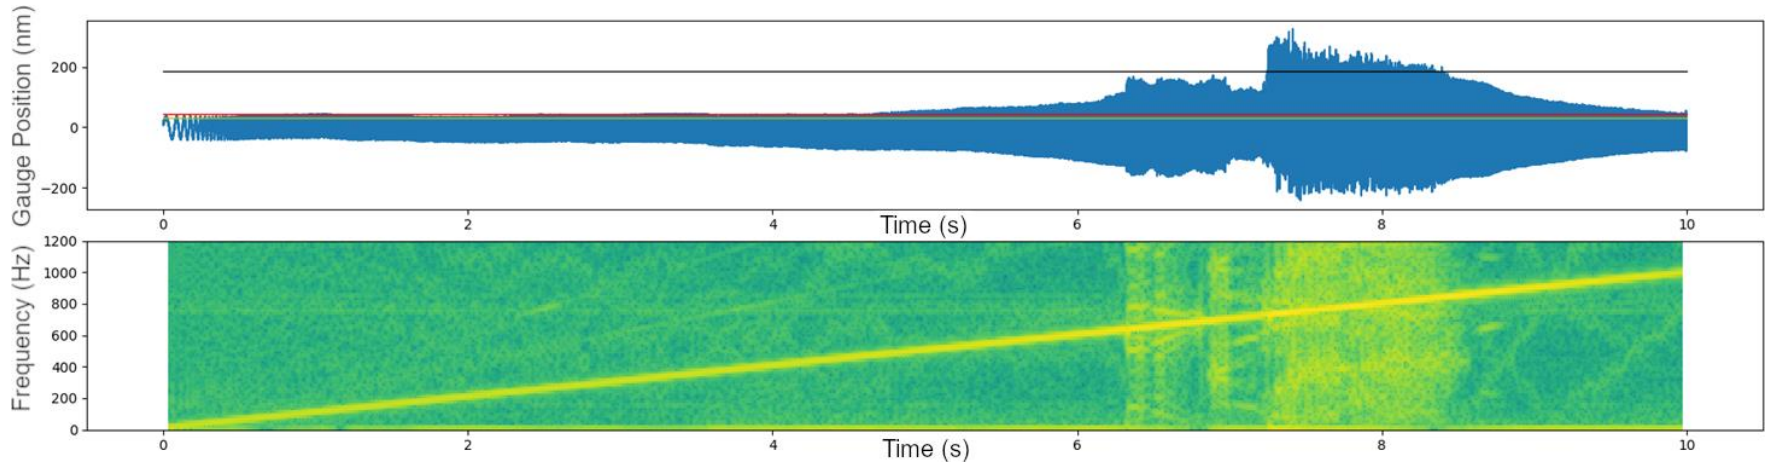


Figure 5–18: The profile and the spectrogram of the 2V amplitude data point highlighted as phenomenon B in Figure 5–14. The stylus leaves the surface at ~630 Hz but this is not detected by the thresholding analysis technique.

Experiment 3 - Phenomenon A

This phenomenon is characterised by the large discontinuity between some measurements at low amplitudes and the subsequent measurements at higher amplitudes. The profile in Figure 5–15 shows an increased amplitude, consistent with the expected primary mode of the ceramic stylus, at around 560 Hz. This resonance peak is not of high enough amplitude to be detected by the thresholding technique, when excited by the 1 Volt peak-to-valley actuator voltage.

This resonance peak is excited into stylus flight at a 2 Volt actuator voltage (Seen at 520 Hz in Figure 5–16). This flight is detected by the thresholding technique and therefore the large disparity between consecutive measurements is observed. This phenomenon should not be considered a failure of the analysis techniques, as it is actually representative of the real situation. This phenomenon does however highlight that the analysis of the gantry results is more complex than can be represented simply by a thresholding technique.

This phenomenon is not observed when using an actuation voltage of 3 Volts, meaning that the materials testing experiments should not need to mitigate this in their analysis.

Experiment 3 – Phenomenon B

This phenomenon is characterised by the flight frequency for a higher actuation amplitude being above the flight frequency of the previous lower amplitude. In the highlighted example (B in Figure 5–14), the 1 Volt stylus flight frequency is 690 Hz whereas the 2 Volt stylus flight frequency is 730 Hz. This is potentially explained by both a physical reality, and an analysis shortcoming.

The stylus leaves the surface at ~630 Hz in Figure 5–18 (2 Volt amplitude), corresponding with the predicted stylus bending frequency; this does not happen at a similar frequency in Figure 5–17. When the stylus does leave the surface in Figure 5–18, the increase in measured amplitude is small, by

comparison to the amplitude of the later stylus flight. The stylus in the 1 Volt actuation amplitude profile (Figure 5–17) does not leave the surface at ~630 Hz, rather the amplitude increases until the stylus leaves the surface at ~690 Hz, which is at a lower frequency than the second stylus flight event in Figure 5–18.

This suggests that bending of the stylus will bring on stylus flight at a lower frequency, fitting with the anecdotally expected behaviour. This flight amplitude is likely dependent on the amplitude of the actuation which explains why the higher amplitude measurements (3 Volts peak-to-valley and above) do not experience the same phenomena.

This small stylus flight amplitude is missed by the thresholding algorithm. The threshold amplitude is set halfway between the initial actuation amplitude (marked as a red line on the profile sub-plot, shown in Figure 5–20) and the highest peak of the profile. The threshold amplitude is marked as a black line on the profile sub-plot. The selection of the stylus flight amplitude is somewhat arbitrary, but this method was found to be sufficient in most cases for discerning the stylus flight frequency of PCF styli in a standard PGI gauge. More detail on this shortcoming of the thresholding algorithm are detailed in Section 4.3.7.

Actuation Waveform – Sinewave vs Triangle wave – Experiment 4

The initial choice of actuation waveform was a sine wave. This waveform shape would limit the actuation to a single frequency, hopefully allowing the measured profile amplitude to be directly related to the input frequency and amplitude.

Other waveforms are available on the control hardware, such as triangle, sawtooth and square wave shapes. The triangular waveform is the most similar to the (Halle Normale, 2019) standards that are commonly used for calibration and stylus testing within the UKAS laboratory.

Experiment 4 consisted of 200 measurements, alternating between a sinusoidal waveform and a triangular waveform. This interleaving approach would help to minimise the differential drift from environmental effects. Examples of the typical profile and subsequent analysis are shown in Figure 5–20 (triangle waveform) and Figure 5–19 (sinusoidal waveform).

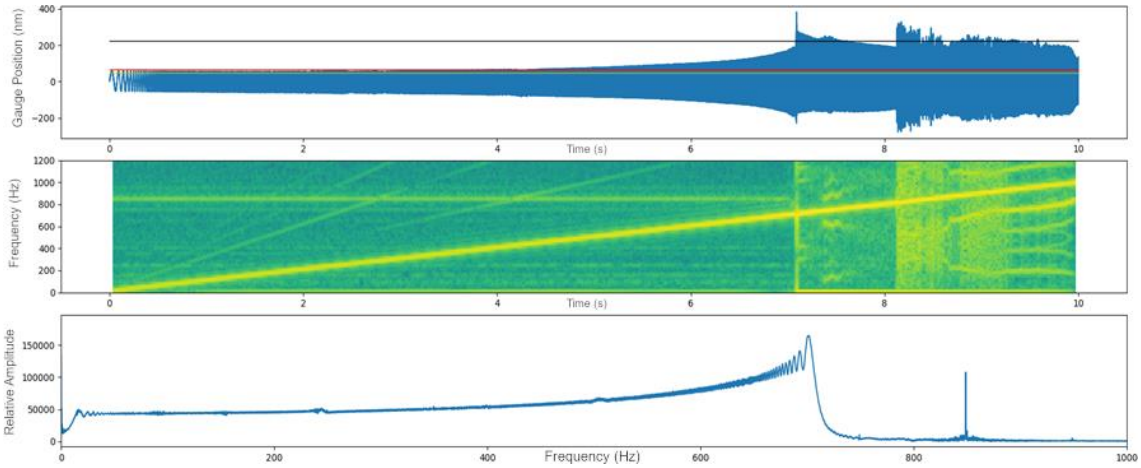


Figure 5–19: A recorded profile using a standard PGI gauge and 112-3221 stylus (PCF). A sinusoidal wave shape, with an amplitude of 3 Volts peak-to-valley, is used in this measurement.

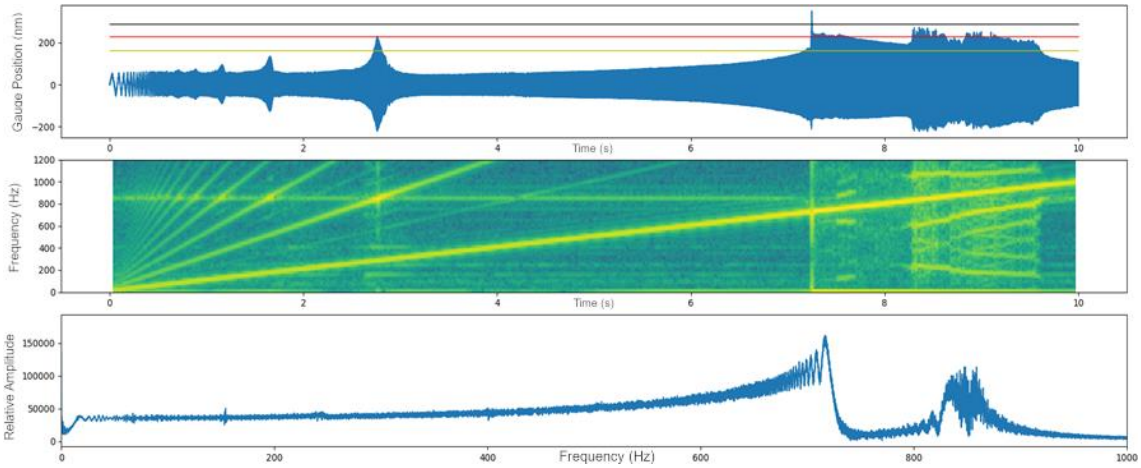


Figure 5–20: A recorded profile using a standard PGI gauge and 112-3221 stylus (PCF). A triangular wave shape, with an amplitude of 3 Volts peak-to-valley, is used in this measurement.

There are some subtle (and some not-so-subtle) differences between the collected measurements, the most obvious being the large spikes in the amplitude of the triangular wave profile at ~280 Hz, ~175 Hz, 124 Hz etc.

These amplitude spikes represent the 3rd, 5th and 7th harmonics of the now-excited resonance, that can be seen at 850 Hz, in the Fourier transform in Figure 5–20. These peaks are caused because a harmonic component of the triangle waveform couples with a natural resonance of the system. Triangle waveforms are composed of a series of odd harmonics, with diminishing amplitudes as the harmonic increases.

The peaks increase in amplitude, from left to right, as the left-most peak is a result of the coupling between the 11th harmonic of the triangle wave and the stylus bending resonance. The next peak is a result of the coupling between the 9th harmonic of the triangle wave, which has a higher amplitude than the 11th harmonic.

The triangle wave can be expressed mathematically as follows:

$$F(t) = \frac{\pi}{4}\sum_{i=0}^{N-1} -1^{i} \cdot n^{-2} \cdot (\sin[2\pi f_0 nt])$$

Where t is time, f_0 if the fundamental frequency and N is the number of harmonics to include in the approximation. The variable i is the harmonic label and is related to the mode number by the relationship:

$$Mode\ Number = 2i + 1$$

A more subtle difference between the two waveforms was the small increase in the stylus flight frequency when using a triangular waveform as opposed to the sinusoid. This is likely as a result of the energy of the sinusoidal waveform having only a single frequency. By contrast, the triangular waveform will distribute its energy over a range of frequencies. This variance is graphed in Figure 5–21.

The triangular waveform is a useful tool in investigating stylus frequency response above the stylus flight frequency. This tool is better suited for analysis of styli by-hand as the resonance peaks that interrupt the previously

flat regions of the profile that were used for analysis. Since the location of the peak features are dependent on the natural bending frequency of the stylus, their location may shift with different geometries and materials (because of their relationship to the actual stiffness of the stylus). This would further increase difficulty of the already complex analysis task.

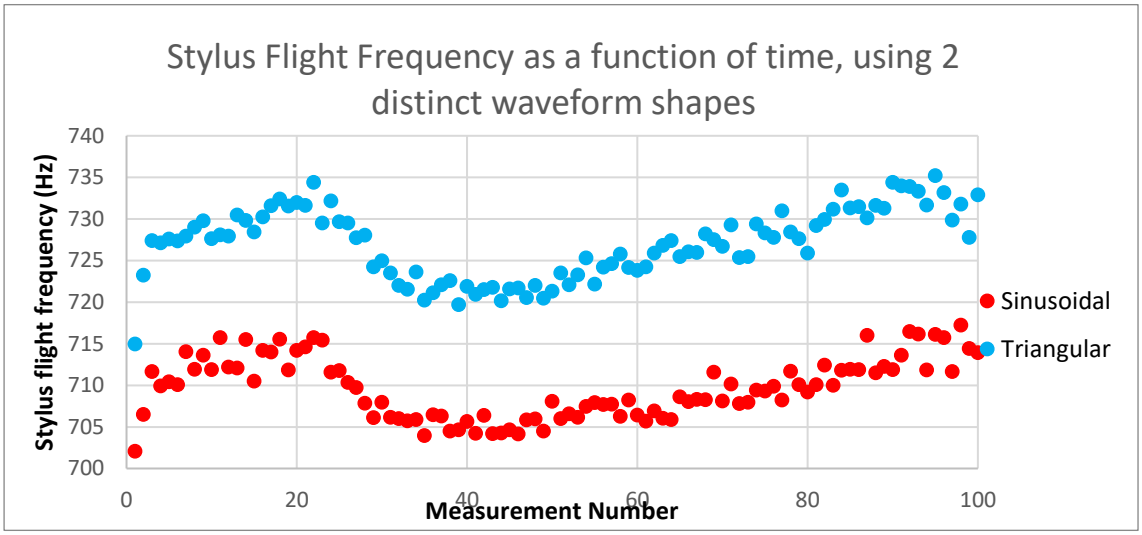


Figure 5–21: A chart showing the measured stylus flight frequency over time, using multiple waveform shapes.

Variance in stylus flight frequency caused by reseating – Experiment 5

The connection between the stylus and the gauge is designed to be highly rigid and repeatable, whilst allowing styli to be swapped easily. The retention spring in the spigot tube is stiff and is designed to fix the rotation of the stylus, as well as ensure that the spigot is firmly pressed into the spigot tube on 4 areas of contact (See Figure 2–3)

The stylus reseating arrangement is rigid and repeatable enough that both commercial partners and service technicians do not report measurement issues that have been solved by simply re-seating the stylus. This suggests that the stiffness of the retention system is sufficient to give a resonant frequency above the surface frequency of even super-fine surface measurements. The variance in measurement performance of the stylus as a function of the reseating process is not well understood; it has never featured

in an observed issue either at Taylor Hobson or reported by commercial partners.

As mentioned previously, the variance of the stylus dynamics caused by stylus insertion and the resetting of the stylus force must be investigated. Experiment 5 consisted of repeated insertion of the same stylus, without resetting or otherwise adjusting the gauge or the gantry. This allows the absolute best-case performance to be quantified.

The results of the testing are shown in Figure 5–22. The majority of the measurements fall in the region between 680 Hz and 720 Hz, although there are some significant, if temporary, departures from this region.

This graph suggests that, for a particular stylus and setup, the re-insertion repeatability is good. The maximum deviation from the nominal stylus flight frequency (700 Hz to 640Hz, in this instance) is still above the typical cut-off frequency for even fine sur-face finish measurement.

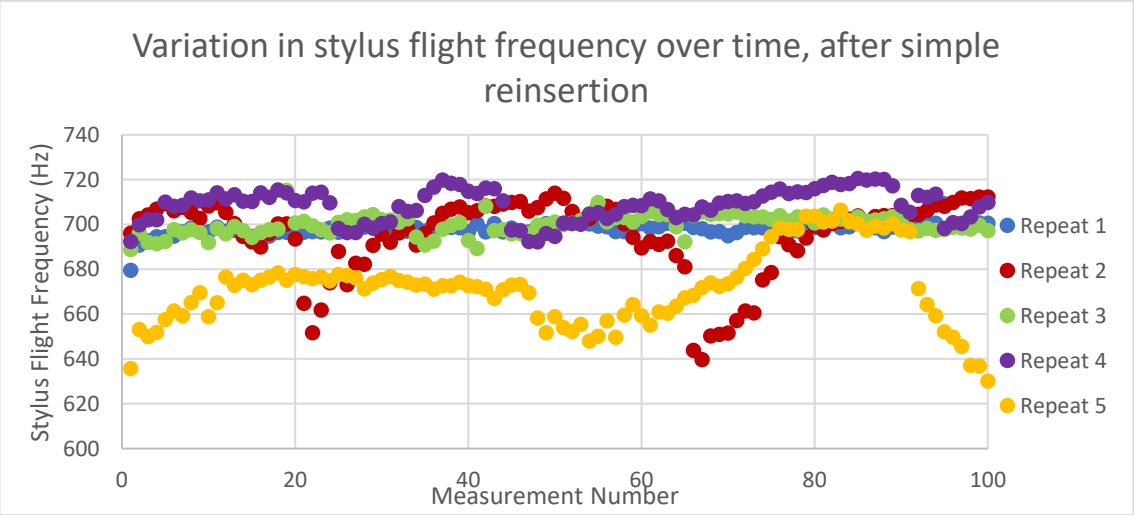


Figure 5–22: A graph of the stylus flight frequency over time. The stylus is re-seated between each repeat, without adjustment to the gauge or gantry system.

Further, the initial offset of repeat 5 is visible throughout the 100 measurements, except for the measurements from number 75 to 92. Between these points, the stylus conforms to the otherwise measured stylus flight frequency of 700 Hz, suggesting that there is a natural stylus flight frequency

(700 Hz, for this specific stylus and gauge configuration) and also a source of reduced stiffness that may reduce this stylus flight frequency.

The average stylus flight frequency of the collected data is 694.5 Hz, with a standard deviation of 16.85 Hz. If repeat 5 is removed from this analysis, the mean stylus flight frequency becomes 700.3 Hz with a standard deviation of 10.87 Hz. These results are both within a small margin of the results achieved in Experiment 1 in Section 5.4.5, suggesting that the effect of simply reseating the stylus is small.

This variability does not improve over the time frame of the testing process (~90 minutes) suggesting that the mounting conditions are stable during normal testing. It is hardly surprising that the stylus doesn't settle during testing, as in this application, the stylus tip only undergoes a maximum displacement of a few microns.

Note on the distribution of the data.

For all of the experiments conducted here, the data does not appear to be normally distributed. This is very obvious in Figure 5–22 as the range of data above the 700 Hz average is only 20 Hz, whereas below the 700 Hz average is 60 Hz. This does not stop the mean and standard deviation of the data from being calculated. It does how-ever mean that a confidence interval cannot be constructed from the mean and standard deviation parameters alone.

Variance in stylus flight frequency caused by resetting – Experiment 6

This experiment is an extension to Experiment 5 in that the influence of the stylus changeover process are characterised. This experiment focusses on the necessary ancillary tasks that must be performed when changing between styli, even when the styli are the same model (60mm PGI styli, for instance).

Before each run of 100 measurements:

1. The stylus is removed from the gauge
2. The same stylus is reinserted (only one stylus is used for experiment 6)

3. The gauge is reset (the stylus tube rests against an internal hard end-stop during this time)
4. The tip force adjustment screw is moved to the extremes of its range (to avoid hysteresis)
5. The stylus is contacted onto a jewel balance, using the z-axis on the gantry, until the gauge is centred in its range
6. The tip force is set at 75 - 85 mg reading on the balance, corresponding to 750 – 850 µN
7. The piezoelectric actuator is aligned under the stylus
8. The stylus is adjusted to the centre of its range
9. The normal automated testing process is performed (100 repeats, 30 second delay between subsequent measurements etc.)

The resulting measurements for experiment 6 are shown in Figure 5–23, with the mean and standard deviation of the stylus flight frequency, and the corresponding tip force being presented in Table 5-3.

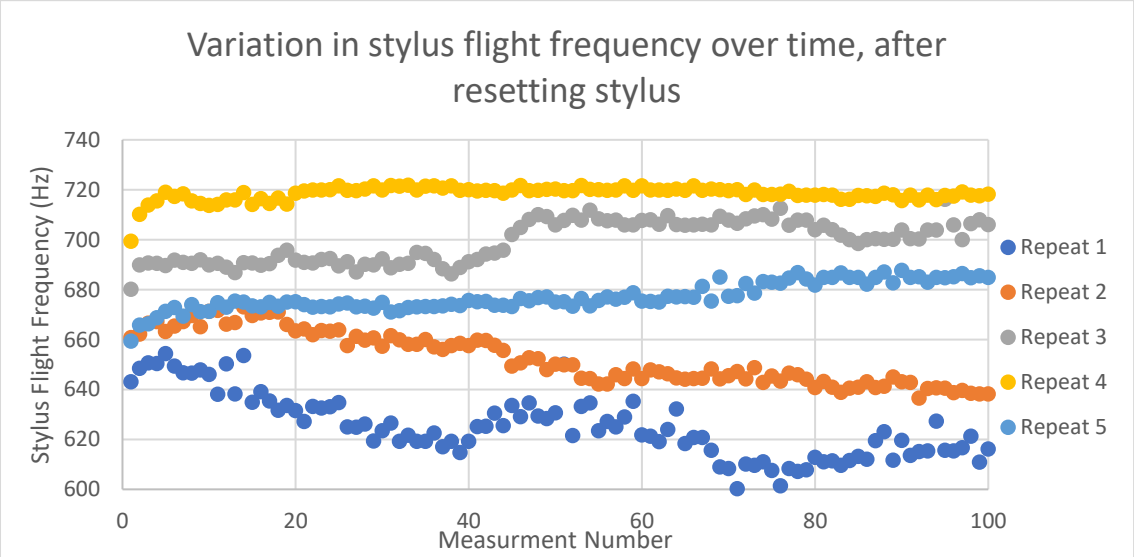


Figure 5–23: A graph of the stylus flight frequency over time. The stylus is re-seated, and the gauge reset, between each repeat.

Repeat #	1	2	3	4	5
Tip Force (mg)	84	74	84	84	80
Flight Freq. (mean) (Hz)	625.02	652.44	699.24	718.43	677.09
Flight Freq. (std dev.) (Hz)	12.62	10.32	8.23	2.91	5.50

Table 5-3: A table of the mean and standards deviations of each repeat in experiment 6

In contrast to the close distribution of results in experiment 5, the results for experiment 6 are much more widely spread. The values of the standard deviations shown in Table 5-3 are consistent with the expected performance of the gantry as observed in the previous experiments.

The comparatively low standard deviations, coupled with the large differences be-tween the means of adjacent (in terms of stylus flight frequency) measurements, mean that each set of results is unlikely to have been produced by chance, even from an adjacent measurement.

For example, assuming each set of 100 measurements is normally distributed:

Repeat 1 against repeat 2:

$$\frac{(652.44\ Hz - 625.02\ Hz)}{12.62\ Hz} = 2.173\ Standard\ Deviations$$

2.173 Standard deviations from the mean would suggest that 98.5% of all the data should lie within this limit, however, the vast majority of the data for repeats 2, 3, 4 and 5 are outside this limit. This suggests some external influence between each repeat, consistent with the stylus resetting process.

The lack of significant departures from the initial stylus flight frequency (such as those seen in Experiment 1.1, 1.2 and 1.3, Section 5.4.5) suggests that

each measurement was not impacted by an external environmental factor during the process of the data collection.

The repeats were performed at different times of the day, with different vibration sources and at different temperatures (stable at 21±2 degrees). The normal variation in temperature does not typically cause significant variation in stylus flight frequency. The environmental noise can have a significant effect on the performance of a sur-face measurement machine, but the gantry system appears to be significantly less sensitive to ground-borne noise, likely because of the 600 kg granite base and the type of analysis being performed (aggregate analysis will be less sensitive to short-time disturbances).

The drive equipment remained powered up throughout the duration of the experiment, thereby reducing the potential of drift as the equipment comes up to working temperature. The gauge was reset to within 10 µm of its zero point before each repeat, reducing the potential cosine errors in running the gauge without calibrated arcuate correction.

The repeatability of the stylus mounting system was shown in experiment 5 to be significantly better than the repeatability observed here. The largest remaining variable is the tip force and so is likely responsible for the differences in the stylus frequency performance. The literature predicts that a change in stylus force would result in a corresponding change to the stylus flight frequency.

This agrees with the experience of resetting the tip force during the experiment. The gauge force adjustment mechanism and the top pan balance, both have hysteresis. This makes the setting of the rather complex

The act of reseating the stylus and resetting the gauge seems to be the cause of the large variance between the individual stylus tests.

Stylus Material Testing – PCF – Experiment 7

Seven 112-3221 styli were assembled for the adhesive testing investigation. They were assembled using adjustable fixturing and tools to ensure the adhesives were properly mixed and evenly dispensed. Each stylus passed the standard calibration process. One of these styli was used in the initial gantry characterisation so the remaining 6 can be used to quantify the variance between nominally identical pultruded carbon fibre (PCF) styli.

Figure 5–24 shows the variance in the stylus flight frequency over time, for a set of standard 60 mm length, PCF styli (112-3221). There are several notable features of the graphs and some significant variance between each. Stylus 1 has the lowest absolute mean frequency but is highly stable. If the first 20 results are ignored, the standard deviation drops to just 5.4 Hz, which is comparable to the baseline gantry performance.

It can also be seen that there is some settling period, for the flight frequencies to become more consistent. This may be caused by mechanical, thermal, electrical or electrostatic phenomena; suggesting that the first few measurements after a stylus is inserted may be more susceptible to noise than measurements made after ~15 minutes.

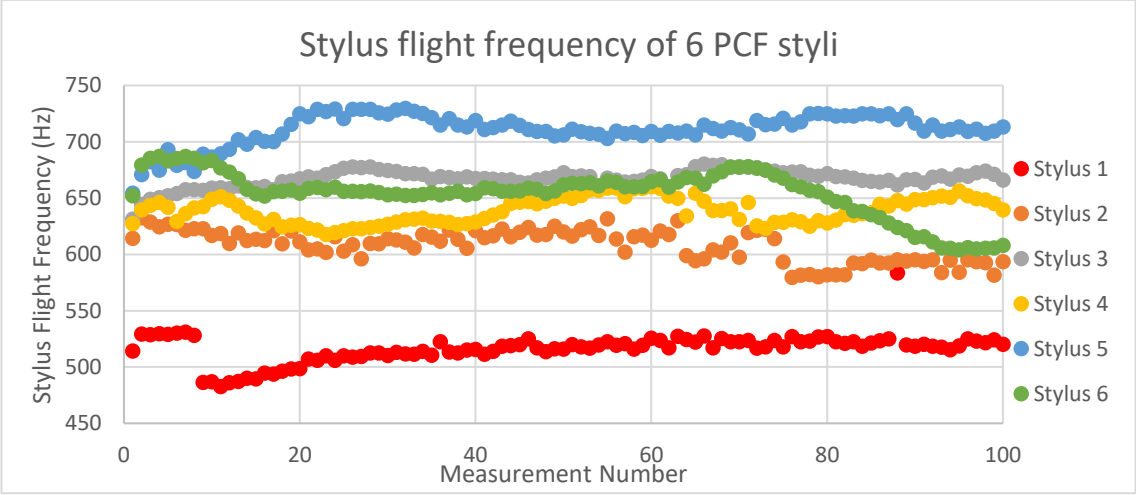


Figure 5–24: A graph of the stylus flight frequency over time of 6 nominally identical PCF styli. The gauge force was reset to 750 – 850 µN for each stylus.

The mean and standard deviation of the stylus flight frequency, and the corresponding tip force for experiment 7 is presented in Table 5-4.

Stylus #	**1**	**2**	**3**	**4**	**5**	**6**
Tip Force (mg)	76	100	78	82	84	84
Mean Flight Freq. (Hz)	515	608	667	638	710	654
Flight Freq. (std dev.) (Hz)	12.8	13.8	7.7	11.4	14.1	20.6

Table 5-4: A summary of the mean stylus flight frequencies and the standard deviations of the 6 standard PCF styli used in experiment 7.

A noteworthy physical explanation for the initial settling time and/or the longer-term drift, is the pressure exerted on the pivot bearings during stylus insertion or removal. The force exerted on the pivots is uncontrolled during stylus insertion, potentially causing plastic or elastic deformation. It is also likely that the oil in the pivot bearings is displaced during stylus insertion.

The large discrepancy between the mean stylus flight frequency of stylus 1 and the other 5 styli suggest that there has been some systematic error in the testing procedure which has resulted in an abnormally low stylus flight frequency. It may be explained by a discrepancy in the assembly procedure, such as the amount and location of the adhesive that was applied. It may also be a result of an incorrectly measured stylus tip force. If the force was low, then the stylus would undergo stylus flight at a lower frequency. The comparatively low standard deviation of this measurement (Stylus 1), coupled with only a single minor step change, suggest that the measurement was repeatable and therefore valid, despite the low mean stylus flight frequency.

If the results for stylus 1 are ignored, the remaining styli have characteristics similar to those gathered in experiment 6. The range of mean stylus flight frequencies for Experiment 6 is 93 Hz (625 – 718 Hz) and for Experiment 7 is 102 Hz (608 – 710 Hz). The results from both experiments 6 and 7 are free from sharp discontinuities and have low drift from start to finish.

The results for experiment 7 typically have a larger standard deviation than for experiment 6. The average standard deviation for experiment 6 is 7.7 Hz, whereas for experiment 7, it is 13.4 Hz. The increased values of the stylus flight standard deviation may be somewhat explained by the physical differences between the individual styli. It may also be as a result of the range of spring positions necessary to counterbalance the gauge with each new stylus, coupled with the slightly longer resetting process that is involved in experiment 7, by comparison to experiment 6.

The increase in the spread of mean stylus flight frequencies in experiment 7 (against experiment 6), suggest that the underlying dynamics of stylus flight are sensitive to minor changes in the stylus. The range of potential manufacturing and assembly tolerances will make it extremely difficult to understand the relationship between production quality and stylus performance. The general dynamic performance trends are still a useful indicator of stylus stiffness, but these experiments show that the stylus flight frequency results must be treated carefully. The stylus flight frequency results may become a more important indicator as the stylus length increases, or the geometry of the stylus becomes more complex.

Stylus Material Testing – Ceramic – Experiment 8

The heavier styli were tested using the PGI Bearings gauge, as the current designs are too heavy for the standard PGI gauge to balance. These styli have a main tube that is made from Rubalit©. Rubalit© is a tradename for a family of alumina-based ceramics, typically with an aluminium oxide (Al2O3) content above 96%. Some relevant mechanical properties for Rubalit© are presented in Table 5-5.

The high hardness of this material, coupled with its brittleness, make it extremely difficult to machine. As such, an aluminium end piece was used to avoid the requirement for machining. The density of this material is roughly twice that of carbon fibre, resulting in a slightly lower specific stiffness of the ceramic.

Property	Value
Density	3.85 kg/m3
Hardness	80 Rockwell C
Flexural Strength	360 MPa
Thermal Expansion	5.4 x 10-6 m/m°C

Table 5-5: A table of some of the pertinent mechanical properties for Rubalit©

The ceramic is available in a limited range of geometries, so the closest compatible geometry is chosen: a tube with OD of 3 mm, ID of 2.1 mm, and a length of 40.0 mm. This is compared to the 3.33 mm OD and 2.5 mm ID of the PCF tube. The geometry and material choices lead the ceramic tube to have a slightly higher absolute stiffness, implying that the ceramic stylus should perform better at higher frequencies, than the standard PCF styli.

The ceramic styli used in this experiment had a moment of inertia (when inserted into an appropriate measurement gauge) that was too high to balance using the standard PGI gauge, and so, the PGI Bearings gauge was used instead.

Figure 5–25 shows the variance in stylus flight frequency over time, for 5 custom ceramic styli, dimensionally similar to the 112-3221 type used in experiment 7. The settling period that was previously observed is still present in these results, although the settling period appears to be shorter in experiment 8 than for experiment 7. The range of stylus flight frequencies is considerably higher larger than observed in experiment 7. The mean and standard deviations of the stylus flight frequency, along with the corresponding tip force is presented in Table 5-6.

The range of mean stylus flight frequencies for the ceramic styli is 282 Hz, compared to 102 Hz in experiment 7. From Figure 5–25, there appears to be 2 distinct groups in terms of mean stylus flight frequency. The large range of mean stylus flight frequencies, along with the appearance of two groups in the data, suggest that there may be an external influence on the stylus flight.

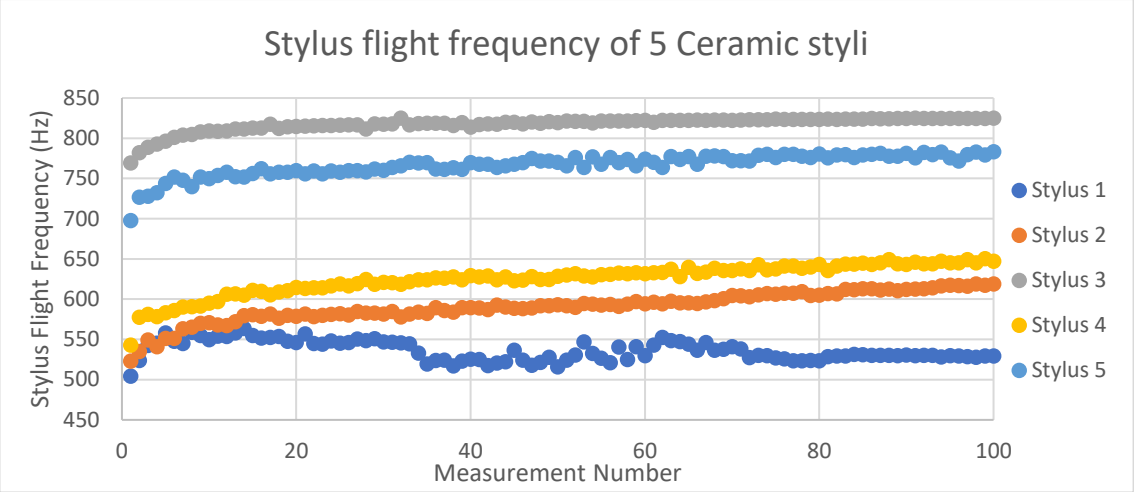


Figure 5–25: A graph of the stylus flight frequency over time of 5 nominally identical ceramic styli, in the PGI Bearings gauge. The gauge force was reset to 750 – 850 mN for each stylus.

Stylus #	**1**	**2**	**3**	**4**	**5**
Tip Force (mg)	76	84	84	80	86
Mean Flight Freq. (Hz)	535	590	817	625	766
Flight Freq. (std dev.) (Hz)	12.2	18.6	9.1	19.1	13.8

Table 5-6: A summary of the mean stylus flight frequencies and the standard deviations of the 5 ceramic styli used in experiment 8.

This external influence on the stylus flight frequency may be attributed to the different assembly methods, by comparison to the PCF styli; if so, the stylus flight frequency consistency may improve with time, as the assembly process becomes more practiced and well-understood. The influence may also be a from the PGI Bearings gauge itself. The gauge is designed to utilise significantly longer styli than a standard PGI gauge and therefore the tip force correction mechanism was at one extreme end of its range. If the spring that is used to correct the stylus tip force is only elongated a small amount, its behaviour may be different than for a spring at the middle of its extension range.

Stylus 1 had a moment of inertia so low that it required a weight to be applied to the main tube before it could be balanced properly in the PGI Bearings gauge. This increased moment of inertia may help to explain the lower flight

frequency. Stylus 1 appears to have some instability (from measurement 35 to measurement 70 in Figure 5–25). There is no obvious explanation for this phenomenon and further testing would be required to determine if the instability was a result of the environment or a property of the stylus.

Stylus 2, 3, 4 and 5 had stylus flight frequencies that followed a general positive trend over time. This may be a relaxation effect in the pivots, a drift effect from the drive electronics or an effect related to the gauge (although this is not seen in experiment 3). This suggests that the ceramic styli may need to be tested for longer periods to mitigate this effect.

The standard deviations for the styli in experiment 8 are generally larger than for experiment 7; the average standard deviation of this experiment is 14.6 Hz against 13.4 Hz for experiment 7. This suggests an increased instability in the stylus flight frequency, although this may not affect surface measurement performance.

The highest stylus flight frequencies were observed using the ceramic styli. The stiffness predicted during the design process is higher for the specific geometry of ceramic stylus, despite the lower specific stiffness of ceramic against carbon fibre. The increased actual stiffness should result in a delayed flight frequency; although this only seems to have happened in a single case. The actual stiffness of the styli has not been measured and so a variance in the achieved stiffness may also explain the large range of stylus flight frequencies across the group.

All ceramic styli were measured and were found to be within the standard geometric tolerances required for the production styli. The tip force for each stylus was set to ~850 µN, although, as previously mentioned, the adjustment mechanism was towards one extreme end of its range. The presence of a nosepiece in the design of the ceramic and aluminium styli introduces another adhesive bond and an extra component into the stylus assembly; the effect of which has not been separately characterised.

Stylus Material Testing – Aluminium – Experiment 9

To reduce the complexity of the comparison between the ceramic and the aluminium styli, and to reduce the manufacturing time, the aluminium styli have the same form as the ceramic styli. They are composed of a 6061-T6 Aluminium tube, 3.3 mm OD, 2.4 mm ID, 40 mm in length. This is a very similar form to both the ceramic and the PCF main tubes, although the Aluminium tube was made slightly larger to help increase the absolute stiffness of the styli. The moment of inertia of the Aluminium styli was too large to be balanced using a standard PGI gauge, so the PGI Bearings gauge was used instead.

Aluminium alloys are often a good choice for engineering situations that require a good strength to weight ratio, although they typically have a lower absolute and specific stiffness than steels or more exotic materials such as titanium. Aluminium has historically been used in the production of surface finish and contour styli and so, is included here. In line with its specific stiffness, the aluminium styli should have the lowest stylus flight frequency of the tested materials. Relevant materials properties of 6061-T6 Aluminium are presented in Table 5-7.

Property	Value
Density	2.7 g/cm^3
Hardness	95 Brinell
Flexural Strength	310 MPa
Thermal Expansion	23.6 x 10^{-6} m/m°K

Table 5-7: A table of some of the pertinent mechanical properties for Aluminium 6061-T6

Figure 5–26 shows the variance in stylus flight frequency over time, for 5 aluminium styli, machined and assembled to be comparable with the ceramic styli in experiment 8. The results presented in Figure 5–26 are distinct from the other experimental results in that they are evenly spaced over the range 400 – 700 Hz, having almost no overlap. Only stylus 2 showed the settling

characteristics observed in the other experiments (7 and 8). The mean and standard deviation of the stylus flight frequency, as well as the corresponding tip force, is presented in Table 5-8.

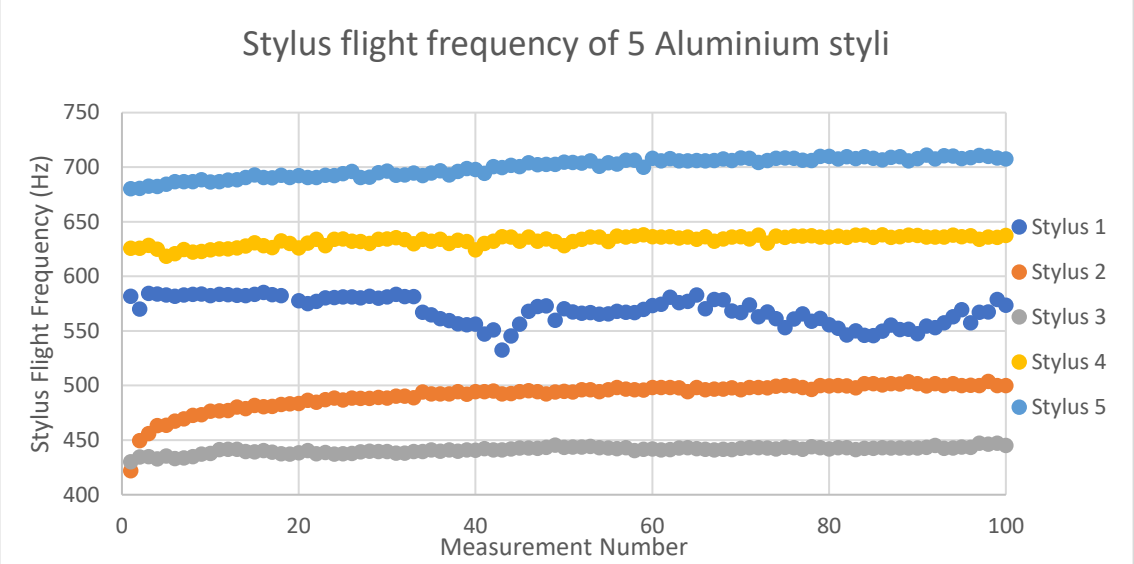


Figure 5–26: A graph of the stylus flight frequency over time of 5 nominally identical ceramic styli, in the PGI Bearings gauge. The gauge force was reset to 750 – 850 µN for each stylus.

Stylus #	**1**	**2**	**3**	**4**	**5**
Tip Force (mg)	80	78	80	78	84
Mean Flight Freq. (Hz)	568	490	441	632	699
Flight Freq. (std dev.) (Hz)	12.4	12.6	3.1	4.5	8.5

Table 5-8: A summary of the mean stylus flight frequencies and the standard deviations of the 5 Aluminium styli used in experiment 9.

The range of mean stylus flight frequencies for the Aluminium styli is 258 Hz. This is a considerably larger range than for the PCF styli (experiment 7) but is in line with the range for the ceramic styli (experiment 8). Both the ceramic and Aluminium styli were tested on the PGI Bearings gauge, but it is not possible to conclude that this is the sole reason for the larger ranges of stylus flight frequencies. The Aluminium styli had a similar moment of inertia to the ceramic styli, with two requiring nose-weights to al-low them to be balanced on the PGI Bearings gauge. The abnormally low (for the PGI bearings gauge)

stylus force spring position is also a common factor in the two sets of styli that were tested in the PGI Bearings gauge (Experiments 8 and 9). Further testing, along with some changes to the stylus geometry, would be necessary to determine which source of error to attribute the large range of stylus flight frequencies.

Stylus 1 seems to exhibit characteristics that are observed in experiment 3, in that the stylus flight seems to be stable at a characteristic frequency, and then departs from this frequency for a time. In this instance the initial characteristic stylus flight frequency (up to measurement 33 on stylus 1 in Figure 5–26) is very stable with a standard deviation of 2.94 Hz, whereas the results after this point have a standard deviation of 10.73 Hz.

The standard deviations of the stylus flight frequencies are lower than for both experiments 7 and 8, with an average standard deviation 8.22 Hz. This suggests that the Aluminium styli had a much lower tendency to deviate from their characteristic stylus flight frequency than the PCF of ceramic styli; even if (and maybe because) they have a lower actual stylus flight frequency.

The stylus flight frequency of the 2 worst performing styli is ~450 Hz. This frequency is low enough that it may not be completely filtered and will therefore be present in the results of the analysis. By contrast, the 2 styli with a flight frequency above 600 Hz would likely perform very well at measuring fine surfaces, demonstrating that Aluminium is suitable for measurement styli, if engineered properly.

As with the ceramic styli, all Aluminium styli were measured and were found to be within the standard geometric tolerances required for the production styli. The tip force for each stylus was set to ~850 µN, although, as previously mentioned, the adjustment mechanism was towards one extreme end of its range.

6. Conclusions

The research presented in this thesis explored a number of aspects of surface finish measurement stylus design, fabrication, validation and performance. The Objectives of this research are listed here for context:

6.1. Objectives

1. Understand how material properties affect stylus performance.
2. Characterise the influence of current production processes on stylus performance
3. Investigate the applicability of Polymer AM, CIM to stylus production
4. Develop methods of predicting stylus performance, prior to application.
5. Develop, improve or verify the current production strategies to enable the rapid and economical production of styli

This research resulted in a greater understanding of the concept of stylus performance, however, the reality of stylus performance has been shown to be more complex than originally expected. This experiment into the influence of the actuation waveform on stylus flight (Experiment 4, Page 211) demonstrates that the stylus flight frequency and the stylus resonance frequency are fundamentally separate. An implication of this separation is that if a stylus may be made stiff enough, such that its resonant frequency is above the flight frequency, then further increasing the stiffness may not further improve measurement performance.

The entire process of stylus production, from the initial design and drawing, through manufacturing, assembly, adhesive application and fixturing has been investigated; with recommendations on the best-practice in each of these areas being documented: Stylus components tolerancing should be standardised (Section 3.1.10). Coolant absorption into PCF is measurable and requires improved cleaning processes (Page 116). Proper adhesive joint geometry is important for achieving a consistent bond (Section 3.3.11). Stylus fixturing can be significantly improved (Section 3.2.17)

The application of Additive Manufacturing, insofar as it applies to stylus production has been demonstrated, with the production of a resin stylus (Section 5.2.1). Current AM technology is not yet capable of both the highly rigid materials properties and the high detail necessary for the production of surface finish styli.

The development and qualification of an adjustable stylus assembly fixture has been documented (Section 3.2.8), along with the modified measurement and adhesive application (Section 3.2.10) processes necessary for the new fixturing approach. A fast, simple to use, readily adjustable, low cost, adaptable fixture has been demonstrated (Section 3.2.8). This fixturing approach allows the detrimental effect of component tolerance stack-up to be mitigated; a feature which was simply not available with a traditional fixturing approach; ultimately allowing the creation of geometrically correct styli, using components that would have resulted in an out of specification stylus had hard fixtures been used (Section 3.2.6)

A stylus dynamics measurement system has been developed (Section 4.3), allowing the automated measurement of the dynamics of the stylus-gauge system. The software developed to accompany the dynamics measurement system, allows stylus flight to be readily measured. Stylus flight is typically discussed in the abstract; the reliable reproduction of this phenomenon allowed further investigation (Section 5.4.4) and demonstrates the efficacy of a spectrogram plot in the presentation of information regarding stylus performance. This measurement system also allowed the effect of surface frequency, amplitude and form to be investigated; an outcome of which was the explicit demonstration of the link between surface amplitude and the stylus flight frequency (Experiment 2, Page 203)

A range of styli have been produced using a set of suitable materials (Section 5). These styli have been characterised using the automated dynamics measurement process (Experiments 7, 8 and 9, Pages 220, 222 and 226 respectively). The materials have been demonstrated as appropriate for stylus construction, indicating the possibility of production of stylus components

using CIM. Ceramic injection moulding of stylus components would reduce the machining required for each stylus significantly.

The prediction of stylus measurement performance before manufacture has not been demonstrated in this project. Instead, the complexity of the task has been reduced with the investigations into dynamic performance of a stylus on the gantry. Areas for future work are suggested.

Of the recommendations made in this report, the updated fixturing approach has been implemented most readily within Taylor Hobson. The applicability of the fixturing technique to a wide range of stylus designs, along with the cost benefit against traditional hard fixturing, was obviously economically beneficial. The recommendation to adapt the spigot geometry to allow a more consistent adhesive bond thickness has been adopted, but the change to a new adhesive for stylus production is not yet commercially important.

7. Further Work

7.1. Stylus Dynamics on Real Surfaces.

The results obtained from the stylus dynamics measurement system in the materials investigation started to bridge to knowledge gap preventing the theoretical prediction of stylus measurement performance from a parametric description of the stylus.

The gantry system was designed with the idea that an artefact may be measured using a linear stage, using a wide range of programmable traverse speeds. This, coupled with the careful selection of surface waveform artefacts, would allow the measurement of surfaces similar to those that are simulated by the gantry and piezo at present.

7.2. Repeatability of Stylus Mounting

The large variance highlighted by the gantry experiments into the effect of reseating the stylus and resetting the gauge force are good indicators that this process is not as predictable as was hoped. The performance of surface measurement systems may improve if the repeatability and reliability of the stylus mounting mechanism can be properly characterised.

7.3. The Development of a Suite of Standard Stylus Literature

In the authors opinion, a complete set of documents should be produced which detail everything relating to styli, from the theoretical ‘best practice’ to a proper representation of what is currently achieved. This group of documents should also include information on:

- Proper tolerancing
- Standards for stylus drawings
- Evaluation procedures (temperature cycling, vibration etc.)
- Handling procedure

- Approach to fixturing
- Range of allowable stylus configurations (tip force, tip radius etc.)
- Adhesive best practice

These documents will serve to reduce the burden of repeated design as they specify much of the stylus. Further, these documents can be updated when the technology changes or improves, thus allowing only a small number of revisions as opposed to updating all stylus drawings.

Models of stylus flight should be further developed to allow for automatic measurement parameter selection for a given surface. The ability to assess the dynamic performance of a stylus has been demonstrated but this must be linked to real-world measurement scenarios.

8. Appendices

8.1. Investigation into the requirements of a stylus dynamics testing system

8.1.1. Initial Investigations

The initial work into the testing apparatus was predominantly concerned with the proper determination of the requirements for several of the key components to the stylus dynamics measurement system; namely:

- Gauge sensitivities
    - Mounting considerations
- Actuator selection
- Software selection
- Integration with current control software and hardware

The understanding of the areas highlighted in is key to properly informing any future design choices that will be made regarding the development of the resonance measurement system. The choice of each component of the system must be compared with alternatives and the most appropriate system assembled and tested to check that it meets the desired specification. The process of selection, testing and calibration of each of these components is detailed below and further development of the chosen implementations will be suggested.

The initial experiments took place on a complete measurement system equipped with a rotary stage and active anti-vibration mounts. The system was equipped with a high resolution PGI gauge with of a vertical resolution of 0.2 nm and a range of 12 mm, the traverse mechanism can be rotated to 45 degrees. The measurement instrument can be seen in Figure 8–1.

The large numbers of degrees of freedom in the PGI Martix system make it a highly versatile measurement system, but undoubtedly introduce extra noise when measuring the response of the gauge. For example, when testing the

dynamics of a stylus, there is no need for a rotary table or a traverse unit. The gauge does need to be able to be raised and lowered to allow for styli with a range of recess lengths to be measured. This is a requirement for the final testing system.

This system was mounted on an active anti-vibration mounting system which does change the dynamic response of the system. Figure 8–2 compares when the frequency domains when anti-vibration is on (Blue) and off (Red). This test was conducted using the modified loudspeaker (loudspeaker in series with a 1 kΩ resistor) and a Nexus 4 smartphone with a tone generation application. The amplitude was set by the phones volume which was at level 2 out of 7.

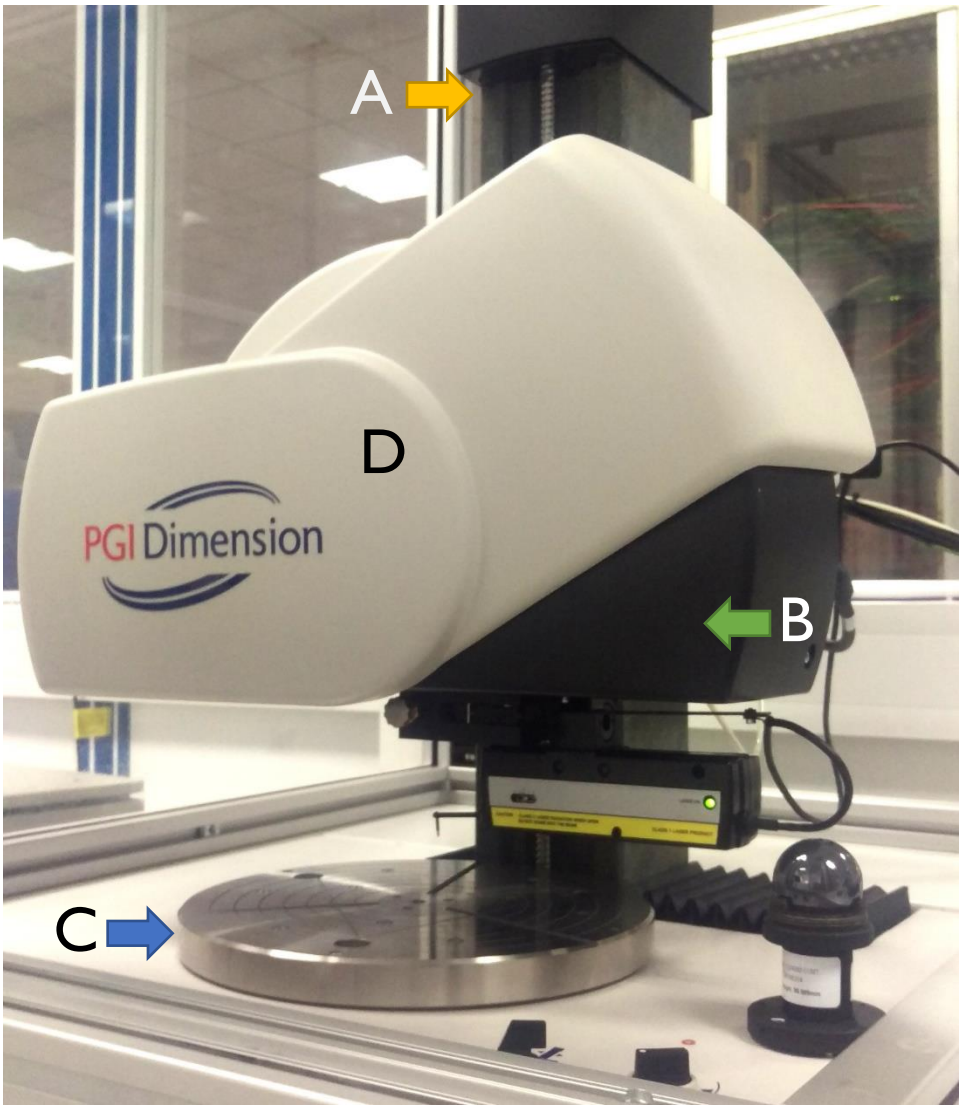


Figure 8–1: Photograph of the measurement system that was used to gather data from the loudspeaker and actuator. Visible are the column (A), traverse unit (B), rotary axis (C), the traverse unit tilt mechanism (D).

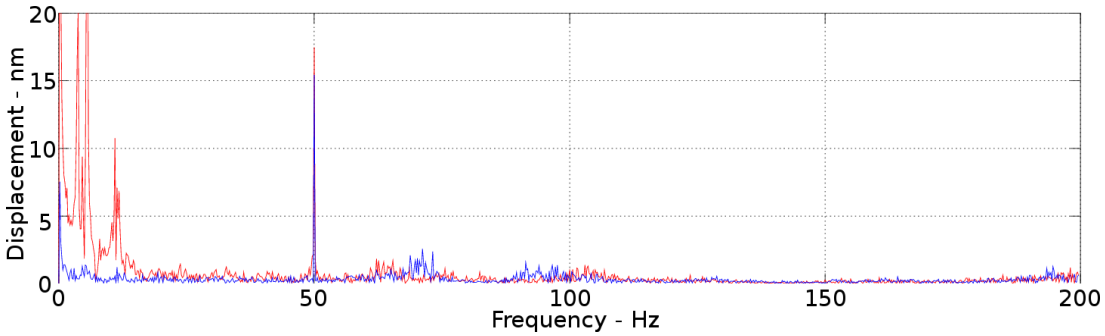


Figure 8–2: A Fourier transform plot of the data take by the measurement system in Figure 8–1. The actuation was provided by a loudspeaker and driven by a smartphone with a tone generation app set to 50 Hz sine wave. The Blue trace shows the gathered results with the active AV system on, and the Red trace is with the active AV system off.

Figure 8–2 has several interesting features. These have been highlighted in Figure 8–3. The plot on the left shows the difference the anti-vibration system makes at low frequencies. It is very clear that the active AV significantly reduces the ambient environmental effects. Air conditioning systems will typically induce a frequency of between 10 and 30 Hz into a system, this being almost completely mitigated by the active AV system.

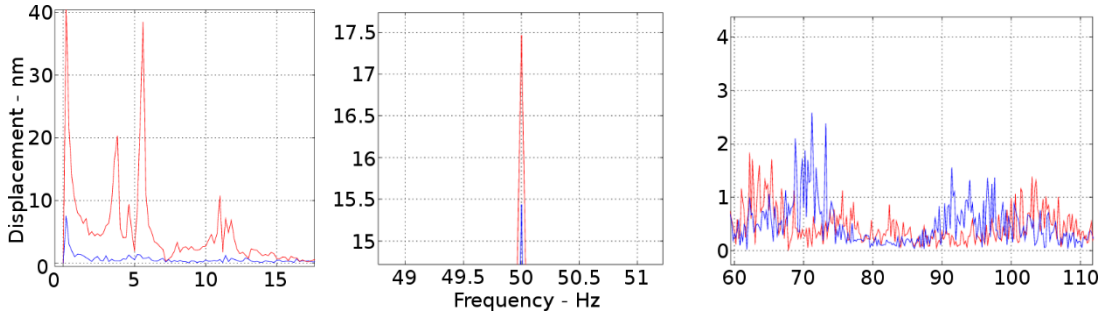


Figure 8–3: Several plots highlighting the important features in Figure 8–2. From left to right this shows the reduction in vibrational amplitude at low frequencies, the reduction of the amplitude of the driven signal at $50Hz$ and the nonlinear shift in noise at higher frequencies. (Blue: active AV on, Ref: active AV off)

The central image shows the deliberately induced signal at 50 Hz and the attenuation when the active AV is turned on. This was a poor choice of frequency as there may be contributions to the signal from the mains wiring. The image on the right shows the frequency shifting of the noise in the system at higher frequencies. The fact that the frequencies shift differently is unsurprising and is probably related to the attenuation characteristics of the

active AV system. The fact that this method can be used to characterise the performance of the active AV system shows promise.

8.1.2. Improvements to the Initial Test System

The promising results gathered using extremely basic tools in lead to an improved test system being implemented. This increased the range of tests that were available and allowed a basic piezo buzzer to be tested in place of the loudspeaker.

A Philips PM5138 (PHILIPS, 1994) signal generator was used to generate the drive signal for the actuator (the loudspeaker or the piezo buzzer). This signal generator lacks the external triggering or the ability to drive arbitrary waveforms but is a highly accurate and capable device.

The second actuator to be tested was a readily available piezo buzzer from Maplin. Since the piezo is essentially a solid-state device, the noise floor was much lower than the loudspeaker and the externally induced noise (signal not at 50 Hz) was considerably reduced. This can be seen in Figure 8–4. The noise measured when using the piezo buzzer as an actuator was more than an order of magnitude lower than for the loudspeaker. However, the simple piezo buzzer is not capable of the range of actuation necessary.

Some features in Figure 8–4 are similar to the ones in Figure 8–2. This suggests that those features are inherent in the measurement system itself as opposed to the method of actuation.

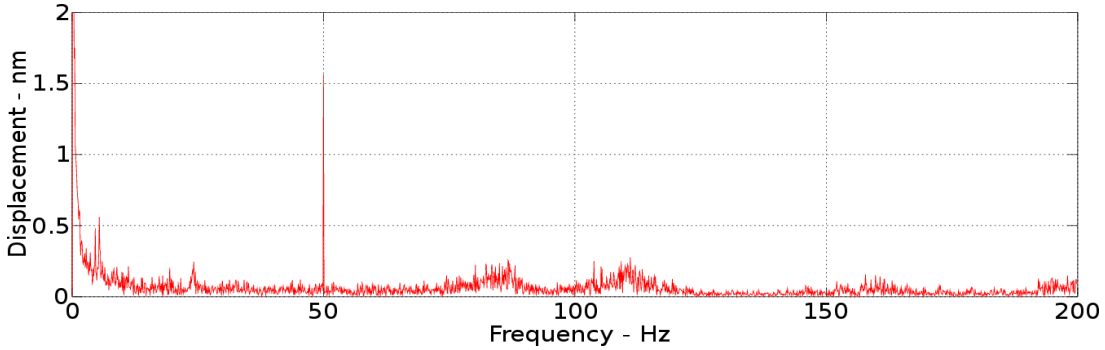


Figure 8–4: A plot of the frequency spectrum of the Maplin piezo. The piezo was being driven at 300 mVpp by the Philips PM5138. Note the vertical axis scale of 0.5 nm per division compared to the 5 nm per division in Figure 8–2

8.1.3. Software used in the Initial Test System

The data gathering and analysis in the initial testing was typically manual. The triggering of the signal generator was coordinated by-hand. This caused a considerable variation in the amount of data recorded before the actuation began, causing the subsequent analysis to be poor. The data was manually cut to improve the consistency of the analysis. This resulted in the requirement for the data gathering and analysis to be as automatic as possible.

The analysis of the data was performed in GNU Octave (GNU, 2019). Octave is a high-level interpreted language, most commonly used for numerical computation. Octave includes the ability to produce complex plots. Octave is designed to be ‘drop-in compatible’ with MATLAB (Mathworks, 2019).

8.2. Coolant Absorption Data

Repeat	Tube 1 (mg)	Tube 2 (mg)	Tube 3 (mg)	Tube 2 – Airgun (mg)	Tube 2 – IPA (mg)
1	1381.8	1399.4	1371.4	1403.0	1400.7
2	1381.7	1399.2	1371.6	1403.2	1400.8
3	1381.7	1399.6	1371.9	1403.4	1400.2
4	1381.6	1399.7	1371.4	1403.3	1400.6
5	1381.2	1399.5	1371.6	1403.2	1400.8
6	1381.1	1400.2	1371.4	1401.8	1400.8
7	1381.4	1399.7	1371.9	1402.2	1400.9
8	1381.2	1399.6	1371.1	1403.5	1400.8
9	1381.2	1400.0	1371.0	1403.1	1400.7
10	1381.6	1399.8	1371.0	1403.0	1400.8
11	1381.6	1399.7	1371.4	1403.0	1400.9
12	1381.5	1399.4	1371.2	1402.8	1400.4
13	1381.5	1399.8	1371.3	1402.9	1400.3
14	1381.0	1400.0	1371.5	1403.0	1400.5
15	1381.4	1399.8	1371.4	1402.9	1400.3
16	1381.3	1399.8	1371.4	1402.7	1400.3
17	1381.2	1399.8	1371.2	1402.8	1400.5
18	1381.2	1400.1	1371.1	1402.6	1400.4
19	1381.4	1399.8	1371.1	1402.9	1400.2
20	1381.5	1400.5	1371.4	1402.8	1400.2

Table 8-1: A table of the mass of the PCF tubes used to test the coolant absorption. Airgun and IPA refers to the cleaning method that was tested.

	Tube 1	Tube 2	Tube 3	Tube 2 – Airgun	Tube 2 – IPA
Average Mass (mg)	1381.40	1399.77	1371.36	1402.90	1400.55
Std. Dev of Mass (mg)	0.223	0.297	0.256	0.389	0.252

Table 8-2: A table of the statistical properties of the coolant absorption data.

9. References

3D Experience Marketplace. (2019). *3D Printing - Material Extrusion*. Retrieved 11 01, 2019, from https://make.3dexperience.3ds.com/processes/material-extrusion

3D Experience Marketplace. (2019). *3D Printing - Photopolymerisation*. Retrieved from https://make.3dexperience.3ds.com/processes/photopolymerization

3D Hubs. (2015). *FDM 3D Printing materials compared*. Retrieved 11 01, 2019, from https://www.3dhubs.com/knowledge-base/fdm-3d-printing-materials-compared/

3D Hubs. (2018). *Introduction to FDM 3D printing*. Retrieved 11 01, 2019, from https://www.3dhubs.com/knowledge-base/introduction-fdm-3d-printing/

3D Hubs. (2019). *SLA-DLP*. Retrieved 11 01, 2019, from https://www.3dhubs.com/3d-printing/processes/sla-dlp/

3D Hubs. (2020). *How to design parts for CNC machining*. Retrieved from https://www.3dhubs.com/knowledge-base/how-design-parts-cnc-machining/#rules

3M - Scotch Weld. (2019). *Scotch-Weld - DP8810NS Green.* Retrieved from https://multimedia.3m.com/mws/media/886363O/3m-scotch-weld-low-odor-acrylic-adh-dp8805ns-dp8810ns-tds.pdf

3M. (2013). *Scotch-WeldTM DP410 Off-White Product Data Sheet.* Retrieved 11 01, 2019, from https://www.rgr.ie/media/wysiwyg/Dp410.pdf

Adams, R. D., Coppendale, J., Mallick, V., & Al-Hamdan, H. (1992). The effect of temperature on the strength of adhesive joints. *International Journal of Adhesion and Adhesives*, 185-190.

Additively. (2018). *Laser Sintering*. Retrieved 11 01, 2019, from https://www.additively.com/en/learn-about/laser-sintering

Alifui-Segbaya, F., Varma, S., Lieschke, G. J., & George, R. (2017). Biocompatibility of Photopolymers in 3D Printing. *3D Printing and Additive Manufacturing*.

All 3DP. (2019, Feb). *How Much Does a 3D Printed House Cost in 2019?* Retrieved 11 01, 2019, from All3DP: https://all3dp.com/2/3d-printed-house-cost/

Attaran, M. (2017). The rise of 3-D printing: The advantages of additive manufacturing over traditional manufacturing. *Kelley School of Business, Indiana University*.

Azom. (2001). *Silicon Nitride (Si3N4) Properties and Applications*. Retrieved from Azom: https://www.azom.com/properties.aspx?ArticleID=53

AZOM. (2005). *Aluminium - Specifications, Properties, Classifications and Classes*. Retrieved from AZOM: https://www.azom.com/article.aspx?ArticleID=2863

Azuma, Y., Barat, P., Bartl, G., Bettin, H., Borys, M., Busch, I., . . . Fujimoto, H. (2015). Improved measurement results for the Avogadro constant using a 28Si-enriched crystal. *Metrologia*.

Bayoumi, A. M. (2000). DESIGN FOR MANUFACTURE AND ASSEMBLY (DFMA) CONCEPTS, BENEFITS AND APPLICATIONS. *Current Advances in Mechanical Design and Production VII.* Seventh Cairo University International MDP Conference.

Bing, L., Anchang, X., & Limin, B. (2015). Preparation of carbon fiber-reinforced thermoplastics with high fiber volume fraction and high heat-resistant properties. *Journal of Thermoplastic Composite Materials*.

Bos, E. J. (2011). Aspects of tactile probing on the micro scale. *Precision Engineering*.

Ceramit. (2018). *Ceramic Injection Moulding (CIM)*. Retrieved 11 01, 2019, from http://www.ceramit.pl/ceramic-injection,en

Chang, T., Wysk, R. A., & Wang, H. (2005). *Computer-Aided Manufacturing.* Prentice Hall International.

Computer Aided Technology (CATI). (2019). *3D Printing – FDM Materials*. Retrieved 11 01, 2019, from https://www.cati.com/3d-printing/materials/fdm/

Damir, M. N. (1973). Error in measurement due to stylus kinematics. *Wear*, 219-227.

Devcon. (2012). *Devcon 5-Minute Expoxy Technical Data sheet.* Retrieved 11 01, 2019, from https://krayden.com/technical-data-sheet/devcon_5_minute_epoxy_tds/

Digilent Inc. (2019). *Analog Discovery 2*. Retrieved from https://store.digilentinc.com/analog-discovery-2-100msps-usb-oscilloscope-logic-analyzer-and-variable-power-supply/

Dornfeld, D., Mina, S., & Takeuchi, Y. (2006). Recent Advances in Mechanical Micromachining. *CIRP Annals*.

Eurobond. (2015). *Penloc VT Plus – Technical data sheet.* Retrieved from https://eurobond-adhesives.co.uk/wp-content/uploads/2016/05/Penloc-VT-plus.pdf

Eurobond. (2016). *EUROBOND FASTFIX – TECHNICAL DATA SHEET.* Retrieved from https://eurobond-adhesives.co.uk/wp-content/uploads/2017/08/Fastfix-Technical-data-sheet.pdf

Fan, K.-C., Cheng, F., Wang, W., Chen, Y., & Lin, J.-Y. (2010). A scanning contact probe for a micro-coordinate measuring machine (CMM). *MEASUREMENTSCIENCE AND TECHNOLOGY*.

Formlabs. (2019). *Form 2*. Retrieved from https://formlabs.com/uk/3d-printers/form-2/

Formlabs. (2019). *Formlabs Resin Library*. Retrieved 11 01, 2019, from https://formlabs.com/uk/materials/

Formlabs. (2019). *How to Choose the Right 3D Printing Material*. Retrieved 11 01, 2019, from https://formlabs.com/blog/how-to-choose-the-right-3D-printing-material/

GDM - Doug Dixon. (2005). *Time Pressure Dispensing.* Retrieved from https://www3.uic.com/wcms/images2.nsf/(GraphicLib)/Time_Pressure_Dispensing.PDF/$File/Time_Pressure_Dispensing.PDF

GE Additive. (2019). *AP&C Powder*. Retrieved 11 01, 2019, from GE Additive: https://www.ge.com/additive/additive-manufacturing/materials/apc/materials-overview

GEBE2. (2019). *Robotic gluing*. Retrieved from https://gebe2-et.com/en/general-industry/robotic-gluing/

GNU. (2019). *GNU Octave - Scientific Programming Language*. Retrieved from https://www.gnu.org/software/octave/

Halle Normale. (2019). *HALLE Präzisions-Kalibriernormale GmbH*. Retrieved from http://halle-normale.de/framesets/englisch/startside/startside.html

Harlow, J. H. (2004). *Electric Power Transformer Engineering.* CRC Press.

Harwin. (2017). *P13-XX23 Pogo Pin.* Retrieved from https://cdn.harwin.com/pdfs/P13-XX23.pdf

Henkel - Loctite. (2008). *LOCTITE® 499 Technical Data Sheet.* Retrieved 11 01, 2019, from https://tdsna.henkel.com/americas/na/adhesives/hnauttds.nsf/web/5E36842C19441DB7882571870000D7A7/$File/499.EN.pdf

Huntsman. (2007). *Users Guide to Adhesives.* Retrieved from https://www.freemansupply.com/datasheets/adhesivesguide.pdf

Hussein, A., Hao, L., Yan, C., Everson, R., & Young, P. (2013). Advanced lattice support structures for metal additive manufacturing. *Journal of Materials Processing Technology*, 1019-1026.

ISO 3274. (1996). Geometrical Product Specifications (GPS) — Surface texture: Profile method — Nominal characteristics of contact (stylus) instruments. *ISO 3274:1996*.

Jääskeläinen, H., & Khair, M. K. (2013). Fuel Injection System Components. *DieselNet*.

Jiang, X., Scott, P. J., Whitehouse, D. J., & Blunt, L. (2007). Paradigm shifts in surface metrology. - Part I. Historical philosophy. *Proceedings of the Royal Society*, 2049-2070.

Jiang, X., Scott, P. J., Whitehouse, D. J., & Blunt, L. (2007). Paradigm shifts in surface metrology. Part II. - The Current Shift. *Proceedings of the Royal Society*, 2071-2099.

Khayal, O. (2019). *DESIGN FOR MACHINING.* Faculty of Engineering and Technology, Nile Valley University.

Kryachek, V. (2004). Injection Moulding (Review). *Powder Metallurgy and Metal Ceramics*, 336-348.

Küng, A., Meli, F., & Thalmann, R. (2007). Ultraprecision micro-CMM using a low force 3D touch probe. *Measurement Science and Technology*.

Loctite - Hysol. (n.d.). *LOCTITE EA 934NA AERO - Technical Datasheet.* Retrieved from https://webaps.ellsworth.com/edl/Actions/GetLibraryFile.aspx?document=13445&language=en

Loctite. (2013). *LOCTITE® 4307™ Technical Data Sheet.* Retrieved from https://tdsna.henkel.com/americas/na/adhesives/hnauttds.nsf/web/E15D8AE0A2512251882571870000D78B/$File/4307-EN.pdf

Loto, C. A. (2016). Electroless Nickel Plating – A Review. *Silicon*, 177-186.

Loughborough University. (2019). *About Additive Manufacturing - Powder Bed Fusion*. Retrieved 11 01, 2019, from https://www.lboro.ac.uk/research/amrg/about/the7categoriesofadditivemanufacturing/powderbedfusion/

Lovo, J. F., Pedroso, M. P., Erbereli, R., Purquerio, B., & Fortulan, C. A. (2018). Synthetic granite composite for precision equipment structures. *Matéria (Rio de Janeiro)*.

Lubricants, M. (2003). *MATERIAL SAFETY DATA SHEET - Cutting Oil.* Retrieved 11 01, 2019, from https://www.technologysupplies.co.uk/downloads/msds/SDS0106.pdf

Mahr. (2016, 11 01). *Mahrsurf XC2 Contour*. Retrieved from https://www.mahr.de/en/Services/Production-metrology/Products/MarSurf---Contour-Measuring-Units/XC-2/

Mahr. (2019, July). *MARSURF LD 130 / LD 260.* Retrieved from https://csisupport.rs/wp-content/uploads/2019/07/MarSurf_LD-130_LD-260.pdf

Maker Juice Labs. (2019). *WaxCast - Direct Investment Casting Resin*. Retrieved 11 01, 2019, from https://makerjuice.com/products/waxcast

manufacturingterms.com. (2020). *Specific modulus*. Retrieved from Manufacturing Terms: https://www.manufacturingterms.com/Specific-modulus-(specific-stiffness).html

Mathworks. (2019). *MATLAB*. Retrieved from https://www.mathworks.com/products/matlab.html

MatWeb. (2020). *Beryllium Copper, UNS C17000*. Retrieved from MatWeb: http://www.matweb.com/search/datasheet.aspx?matguid=9b159fd901454263b8a90c22a66f1988&ckck=1

Midas NFX. (2014, 10 4). *How to Reduce CAD Design Cost using simulation ?* Retrieved from http://feaforall.com/reduce-cad-design-cost-using-simulation/

Moldestad, D. A. (1999). *Some Aspects of Ski Base Sliding Friction and Ski Base Structure.* Trondheim: Norwegian University of Science and Technology.

Morgan Technical Ceramics. (2019). *MACOR® Properties*. Retrieved from Morgan Technical Ceramics: https://www.morgantechnicalceramics.com/en-gb/products/macor-machinable-glass-ceramic/macor-properties/

Mower, T., & Long, M. (2016). Mechanical behavior of additive manufactured, powder-bed laser-fused materials. *Materials Science and Engineering: A*, 198-213.

Nee, A., Senthil Kumar, A., & Tao, Z. J. (2004). An Advanced Treatise on Fixture Design and Planning.

Palásti-Kovács, B., Sipos, S., & Czifra, Á. (2012). INTERPRETATION OF "Rz = 4×Ra" AND OTHER ROUGHNESS PARAMETERS IN THE EVALUATION OF MACHINED SURFACES. *13th International Conference on Tools.* Miskolc, Hungary.

Pereira, T., Kennedy, J. V., & Potgieter, J. (2019). A comparison of traditional manufacturing vs additive manufacturing, the best method for the job. *Procedia Manufacturing*.

Permabond. (2016). *How to eliminate air from adhesive and remove air from a bond joint*. Retrieved 11 01, 2019, from https://www.permabond.com/2016/12/16/eliminate-air-adhesive-remove-air-bond-joint/

PHILIPS. (1994). *PHILIPS PM5138.* Retrieved from https://www.testequipmenthq.com/datasheets/PHILIPS-PM5138-Datasheet.pdf

Physik Instrumente. (2019). *P-841 Preloaded Piezo Actuators*. Retrieved from https://www.physikinstrumente.co.uk/en/products/piezoelectric-transducers-actuators/p-841-preloaded-piezo-actuators-412418452/

Physik Instrumente. (2019). *Preloaded Piezo Actuators.* Retrieved from https://www.physikinstrumente.co.uk/fileadmin/user_upload/physik_instrumente/files/datasheets/P-841-Datasheet.pdf

Precipart. (2018). *Technical Ceramics*. Retrieved 11 01, 2019, from https://www.precipart.com/products-capabilities/custom-mechanical-components/technical-ceramics/

Python. (2019). *Python.org*. Retrieved from https://www.python.org/

Reason, R. E., Hopkins, M. R., & Garrod, R. I. (1944). *Report on the Measurement of Surface Finish by Stylus Methods.* Leicester.

Renishaw. (2019). *Styli and accessories.* Retrieved from http://resources.renishaw.com/en/download/technical-specifications-styli-and-accessories--84826

Renishaw. (2020). *Styli materials*. Retrieved from https://www.renishaw.com/en/styli-materials--6423

Res3D. (2018). *Additive manufacturing - 3D printing*. Retrieved 11 01, 2019, from https://www.res3d.it/tecnologie

Rex Plastics. (2019). *The cost to make a plastic injection mold*. Retrieved 11 01, 2019, from https://rexplastics.com/plastic-injection-molds/how-much-do-plastic-injection-molds-cost

Riley, F. L. (2004). Silicon Nitride and Related Materials. *Journal of the American Ceramics Society*.

R-Project. (2019). *The R Project for Statistical Computing*. Retrieved from https://www.r-project.org/

Silverman, M. P., & Farrah, J. (2018). Bending of a Tapered Rod: Modern Application. *World Journal of Mechanics 8*, 272-300.

Slocum, A. H. (1994). *Precision Machine Design - Topic 9 - Design of the machine's structure.* MIT. Retrieved from http://web.mit.edu/2.70/Lecture%20Materials/Documents/Week%2003/PMD%20Topic%209a%20Structural%20design%201.pdf

Song, J. F., & Vorburger, T. V. (1996). Stylus Flight in Surface Profiling . *Journal of Manufacturing Science and Engineering*, 188-198.

Song, J. L., Li, Y. T., Deng, Q. L., & Hu, D. J. (2007). Rapid prototyping manufacturing of silica sand patterns based on selective laser sintering. *Journal of Materials Processing Technology*, 614-618.

Stanimirovi, Z., & Stanimirovi, I. (2012). 6 Ceramic Injection Molding. *Materials Science*.

Steinmeyer mechatronik GmbH. (2019). *Linear Stages*. Retrieved from https://www.steinmeyer-mechatronik.de/en/products/standardaxes/linear-stages/

Stienstra, D. (2005). *Introduction to Design for (Cost Effective) Assembly and Manufacturing.* Retrieved from http://me.gatech.edu/files/capstone/L071ME4182DFA

Stratasys. (2019). *Direct Metal Laser Sintering (DMLS)*. Retrieved 11 01, 2019, from https://www.stratasysdirect.com/technologies/direct-metal-laser-sintering#dmls-materials

Surberg, C. H., Stratton, P., & Lingenhöle, K. (2007). The effect of some heat treatment parameters on the dimensional stability of AISI D2. *Cryogenics*, 42-47.

Tap Plastics. (2019). *Pultruded Carbon Rods & Tubes.* Retrieved 11 01, 2019, from https://www.tapplastics.com/uploads/pdf/Pultruded_Carbon_Rods__Tubes.pdf

Taylor Hobson. (2019). *Styli & Accessories.* Retrieved from https://www.taylor-hobson.com/-/media/ametektaylorhobson/files/learning%20zone/th-styli-accessories_lowres_iss4_en.pdf?la=en

Taylor Hobson Ltd. (2019). *Fuel cell measurement.* Application Note A159: Fuel cell measurement, Leicester.

Thalmann, R., Meli, F., & Küng, A. (2016). State of the Art of Tactile Micro Coordinate Metrology. *Design and Applications of Coordinate Measuring Machines*.

ThorLabs. (2014). *SCP05 - 16 mm Cage System*. Retrieved 11 01, 2019, from https://www.thorlabs.com/thorproduct.cfm?partnumber=SCP05

Tian, Y., Liu, X., Chetwynd, D. G., Shirinzadeh, B., & Zhang, D. (2010). Vibration analysis of stylus instrument for random surface measurement. *Precision Engineering*, 586-591.

Tian, Y., Liu, X., Zhanga, D., & Chetwynd, D. G. (2008). Dynamic modelling of the fidelity of random surface measurement by the stylus method. *Wear*, 555-559.

Tofail, S., Koumoulos, E., Bandyopadhyay, A., Bose, S., O'Donoghue, L., & Charitidis, C. (2017). Additive manufacturing: scientific and technological challenges, market uptake and opportunities. *Materials Today*.

Whitehouse, D. J. (1994). *Handbook of Surface Metrology.* CRC Press.

Whitehouse, D. J. (1996). Enhancement of instrument and machine capabilities. *Nanotechnology*, 47-51.

Willis, E. (1986). Surface finish in relation to cylinder liners. *Wear*, 351-366.

Willn, J. E. (1972). Characterisation of cylinder bore surface finish— A review of profile analysis. *Wear*, 143-162.

Xtreme Precision Systems. (2020). *Xpress Precision Engineering*. Retrieved from xpresspe.com: http://www.xpresspe.com/