
OPTICAL SCATTERING FOR SECURITY APPLICATIONS

Peter R. Seem

**Department of Physics
Imperial College London**

A dissertation submitted for the degree of Doctor of Philosophy.

September, 2009

Dedicated to my family, with love

Abstract

Laser Surface Authentication (LSA) has emerged in recent years as a potentially disruptive tracking and authentication technology. A strong need for such a solution in a variety of industries drove the implementation of the technology faster than the scientific understanding could keep up. The drive to miniaturise and simplify, the need to be robust against real-world problems like damage and misuse, and not least, intellectual curiosity, make it clear that a firmer scientific footing is important as the technology matures.

Existing scattering and biometric work are reviewed, and LSA is introduced as a technology. The results of field-work highlight the restrictions which are encountered when the technology is applied. Analysis of the datasets collected in the trial provide, first, an indication of the performance of LSA under real-world conditions and, second, insight into the potential shortcomings of the technique.

Using the particulars of the current sensor's geometry, the LSA signal is characterised. Measurements are made of the decorrelation of the signature with linear and rotational offsets, and it is concluded that while surface microstructure has a strong impact on the rate of decorrelation, this dependency is not driven by the surface's feature size. A new series of experiments examine that same decorrelation for interference effects with different illumination conditions, and conclude that laser speckle is not an adequate explanation for the phenomenon.

The results of this experimental work inform a mathematical description of LSA based on a combination of existing bi-static scattering models used in physics and ray-tracing, which is implemented numerically. The results of the model are found to be a good fit to experimental work, and new predictions are made about LSA.

Declaration

I affirm that the work contained herein is my own, except where the work of other authors is appropriately acknowledged.

A handwritten signature in black ink, reading "Peter P. Kim". The signature is written in a cursive style with a large, stylized 'P'.

Acknowledgements

First and foremost, I must gratefully acknowledge the guidance and support of my supervisor, Professor Russell Cowburn. Enormous credit is due him for shaping my ideas and energy into something vaguely resembling science, both through his active supervision and through providing a role model for what a researcher can achieve.

Similarly, Dr. James Buchanan provided invaluable guidance throughout my period of study, in all things from physics to proper English vocabulary. This work owes much to his patience and instruction. Great credit and thanks must also be offered to Julie Kite, who provided more administrative and cake-related support than a Ph.D. candidate has any right to have.

My academic colleagues at Imperial provided a dynamic environment, a sounding board for any number of ideas, and a multitude of reasons not to be working when I should have been. For two of those three reasons, I thank Dr. Dorothée Petit, Dr. Ana-Vanessa Jausovec, Dr. Dan Read, Dr. Oleg Petravic, Dr. Chee Lim, Marc Leat, James Berg, Dr. Laura Thevenard, Liam O'Brien, Huang Zeng, Emma Lewis, Dr. Simon Bühlmann, Dr. Fiona Turner and João Sampaio.

I was further aided throughout my experimental work by the skills of my engineering colleagues at Imperial, and must thank in turn: Shahid Hanif, Susan Parker, Ryan Burrows, Stephen Cussell, and Dave Bowler.

The various employees of Ingenia Technology must all be thanked for their support over the last few years, be it personal, intellectual, financial, or professional. In particular I am indebted to Mark McGlade and Kevin Graham for their varied and valued guidance.

I would also like to personally thank Dr. Yousef Jameel for his support, and more importantly for his vision and foresight in recognising the untapped value of research and researchers.

Several important pieces of experimental work performed in this study hinged on the supremely competent and professional support received from Sam Cox at PRP Optoelectronics and Kevin Amos and Craig Ashmore at Kuka Robotics, UK, and I can offer them only my thanks in return.

I have met many friends and colleagues with Ingenia Technology's various customers and partners, all of whom have been gracious and supportive, and whose cooperation aided in the development of this work. I must beg their understanding that I am prohibited from naming names, but I am grateful to each of them in turn.

My father has been a tremendous source of inspiration and confidence, no less during these years of study than in the 23 years which preceded them. He has taught me what it is to be a scientist and taught me what it is to be a man. At least I have a forum to appropriately thank him for one of those things.

Finally, I have been told that there are only two thank you's that are truly worth making, and neither can ever hope to be sufficiently articulated. In light of that wisdom, I can only conclude simply and wholeheartedly: thank you God, and thank you Momma.

ABSTRACT.....	3
DECLARATION	4
ACKNOWLEDGEMENTS.....	5
1 INTRODUCTION.....	12
1.1 Background	12
1.2 Motivation	15
1.3 Summary of Key Findings	15
1.3.1 Positional Tolerance Independent of Feature Size.....	15
1.3.2 LSA is not Spatially-Averaged Speckle	16
1.3.3 Additional Results.....	16
1.3.4 Resources Produced	16
1.4 Structure of the Document.....	17
1.5 Survey of the Literature.....	18
1.5.1 Overview	18
1.5.2 Previous Work in Scattering from Rough Surfaces.....	19
1.5.2.1 Theory of Scattering from Random Rough Surfaces.....	19
1.5.2.2 Numerical Simulations of Scattering	21
1.5.3 Previous Work in Laser Speckle.....	22
1.5.3.1 Statistics of Laser Speckle	22
1.5.3.2 Applications of Laser Speckle	26
1.5.4 Previous Work in Biometrics and PUFs	27
2 BACKGROUND ON LASER SURFACE AUTHENTICATION.....	34
2.1 Introduction	34
2.2 Principles of Operation of LSA	34
2.3 Illumination and Observation Geometry	36
2.4 Hardware	37
2.5 Positioning of the Target Item	39
2.6 Scans – Acquiring Data	42
2.7 Processing Raw Intensity Data	44
2.7.1 Encoding Velocity Information	44
2.7.2 Application of the Band-pass Filter	45
2.7.3 Developing the Binary Signature	46
2.8 Comparing LSA Signatures.....	47
2.9 Metrics for Evaluating Comparisons	49
2.10 Conclusions	52
3 PERFORMANCE OF LSA IN A FIELD TRIAL	54
3.1 Introduction	54
3.2 The Operational Environment	54
3.3 Four Products for LSA Identification	56
3.3.1 Overview of Products	56
3.3.2 Board 1	58
3.3.3 Board 2	60
3.3.4 Board 3	63
3.3.5 Board 4	65
3.4 Accurate Positioning in the Operational Environment.....	67
3.5 Results from Sampling Outtakes	69
3.6 Results from a Large Enrolment Data Set	72
3.7 Conclusions	74
4 TOLERANCE OF THE LSA SIGNAL TO TARGET MISPLACEMENT.....	78
4.1 Introduction	78
4.2 Tolerance to Misplacement of Real Products	79
4.2.1 Introduction.....	79
4.2.2 5-Axis Motion Control with Precision Stages	79
4.2.3 A Survey of Product Samples.....	85
4.2.3.1 Standard Samples	87

4.2.3.2	Plastic Samples	87
4.2.3.3	Sample Paper Products.....	87
4.2.3.4	Paperboard/Cardboard Samples	88
4.2.3.5	Miscellaneous Samples.....	89
4.2.4	Results of Tolerance Measurements	89
4.2.5	Analysis of the Dependencies of Sample Tolerances	92
4.2.6	Conclusions.....	96
4.3	Tolerance to Misplacement of a Model System	97
4.3.1	Introduction.....	97
4.3.2	7-Axis Motion Control with Robotic Arm	97
4.3.2.1	Motivation for Robotic Arm.....	97
4.3.2.2	Terminology Common to Robotic Manipulators	99
4.3.2.3	Vendor-Provided Control of the Arm	100
4.3.2.4	Customised Control of the Arm	101
4.3.2.5	Synchronising Data Acquisition with Motion Control.....	102
4.3.2.6	Integration with Existing Tools	103
4.3.3	Well-Characterised Ground Glass Samples	103
4.3.4	Additional Two Axes	106
4.3.5	Results of Tolerance Measurements II	108
4.3.6	Conclusions.....	112
4.4	Sources of LSA's Sensitivity to Positioning	113
4.5	Conclusions	115
5	EXAMINATIONS OF THE OPTICAL PHENOMENA UNDERPINNING LSA	118
5.1	Introduction	118
5.2	Decorrelation of Laser Speckle with Target Misplacement	119
5.2.1	Introduction.....	119
5.2.2	Foundations.....	120
5.2.3	Apparatus & Methodology for Speckle Decorrelation Measurements	122
5.2.4	Predicted Decorrelation from the Scattering Vector	125
5.2.5	Decorrelation Results with Translation and Roll	127
5.2.6	Conclusions.....	133
5.3	Prediction of a Huygens Wavelet Simulation	134
5.3.1	Introduction.....	134
5.3.2	Overview of the Model.....	134
5.3.3	Formulation of the Numerical Simulation	135
5.3.4	Results of the Calculations	140
5.3.5	Conclusions.....	144
5.4	Fractional Intensity from Spatially Averaged Speckle	145
5.4.1	Introduction.....	145
5.4.2	Generating Pseudo-Scans.....	145
5.4.3	Varying the Integrating Area	146
5.4.4	Conclusions.....	147
5.5	LSA Scans with Incoherent Sources of Illumination	148
5.5.1	Introduction.....	148
5.5.2	Sources and Optical Geometry.....	149
5.5.3	Results from Three Different Sources	152
5.5.4	Supersensitivity as a Manifestation of Speckle	154
5.5.5	Conclusion	156
5.6	LSA Scans with Different Wavelengths of Illumination	157
5.6.1	Introduction.....	157
5.6.2	Theoretical Independence of Speckle at Different Wavelengths	157
5.6.3	Apparatus & Methodology for Scans with Different Wavelengths	158
5.6.4	Results from Fibre-Coupled Illumination	160
5.6.5	Conclusions of the Fibre-Coupled Study	162
5.7	Conclusions	163
6	TWO-SCALE BISTATIC SCATTERING MODEL	166

6.1	Introduction	166
6.2	Foundations	167
6.2.1	Foundations of Surface Generation	167
6.2.2	Foundations of the Scattering Model	168
6.3	Formulation of the Two-Scale Simulation	169
6.3.1	Overview of the Two-Scale Model	169
6.3.2	Generation of Random Two-Dimensional Rough Surfaces	170
6.3.3	Defining the Illumination Envelope	172
6.3.4	Defining the Observation Plane.....	175
6.3.5	Diffuse Scattering from Planar Facets	176
6.3.6	Simulation of Complete LSA Scans	179
6.4	Numerical Results of the Two-Scale Simulation.....	181
6.4.1	A Single Simulated Scan	181
6.4.2	Fractional Intensity and Correlation Length of Simulated Scans.....	182
6.4.3	Detector vs. Detector Correlation	185
6.4.4	Simulated Tolerances to Misplacement	188
6.5	Outlook for the Numerical Simulation.....	192
6.6	Predictions of the Simulation for LSA	193
6.7	Conclusions	195
7	CONCLUSIONS AND OUTLOOK	198
7.1	Summary of Key Findings	198
7.2	Outstanding Questions.....	200
7.3	Suggested Improvements to the LSA System	201
	GLOSSARY OF TERMS.....	203
	REFERENCES.....	205
	APPENDIX A – CRADLE STAGE MOUNT	210
	APPENDIX B – FULL TOLERANCE DATA.....	211
	APPENDIX C – RBM, RBS, AND CROSSCOMM	215
	APPENDIX D – KRL PROGRAM	217
	APPENDIX E – KUKA CONNECTOR WIRING	221

1

INTRODUCTION

“But ye are forgers of lies, ye are all physicians of no value. O that ye would altogether hold your peace!”

—Job 13:4-5

1 Introduction

1.1 Background

Counterfeiting goods of value is hardly a new phenomenon. Archaeologists have found counterfeit coins nearly as old as the oldest genuine coins[1]. It is, however, on the rise. The most commonly-cited sources for the value of pirated and counterfeit goods are the Organisation for Economic Co-operation and Development (OECD) and the Counterfeiting Intelligence Bureau of the International Chamber of Commerce. A survey of those reports from recent years found that the estimated financial losses due to counterfeiting globally rose from 100 billion USD in 1993 to 610 billion USD in 2006. Counterfeits as a percentage of global trade increased from 3% to between 7% and 9% in the same period[2].

The cost of counterfeit goods is more than just the loss in sales revenue. Poor quality imitations can also damage a brand's equity, defined as the value carried by the awareness, attributes, benefits, images, thoughts, feelings, attitudes and experiences associate with a brand[3]. When consumers are unaware that the product they have bought is counterfeit, any dissatisfaction with the product will reflect on the original brand owner.

Beyond the economic cost, pirated and counterfeit goods are a growing source of revenue for organised crime and international terrorism. In testimony before the 108th US Congress' House Committee on International Relations, the Secretary General of Interpol testified that "Intellectual Property Crime is becoming the preferred method of funding for a number of terrorist groups." [4] In explaining the criminal motivation behind the shift, Secretary Noble told the committee that:

One estimate is that the profits from counterfeiting are similar to drugs trafficking; there is a return of €10 for each €1 invested.

Other estimates are that counterfeiting is more profitable than drugs trafficking... a computer game costs €0.20 to produce and sells at €45 while cannabis costs €1.52 a gram and sells at €12.

In terms of the levels of risk involved, the penalties are low, for example, in France selling counterfeit products is punishable by a two-year prison term and a €150,000 fine, while selling drugs is punishable by a ten-year prison term and a €7,500,000 fine.

Generally speaking, piracy refers to wilful violations of copyright, while counterfeiting refers to wilful violations of trademark[2]. The numbers quoted above are combined figures for piracy and counterfeiting, and modern discussions of the subject tend to focus on digital media, and therefore

piracy. Much ink has been spilled in discussing the introduction of digital rights management in music, movies and games, and lawsuits and prosecutions under the Digital Millennium Copyright Act[5], all of which relate to piracy. However, counterfeit goods can have more tangible and less easily quantified costs to health and human safety. Airplane parts, automobile parts, and pharmaceutical products all undergo stringent safety testing and regulatory oversight, all of which is absent in their counterfeit counterparts.

Laser Surface Authentication (LSA) is a new technique[6], developed in 2004, that may be well-placed to address the problems of counterfeiting. It is a method of illuminating an item's surface with a laser, and measuring the fine variations in the diffusely scattered light. The variations in intensity over a region of the surface are used to create an LSA signature that, like a human fingerprint, contains enough identifying information about the surface's topography to uniquely identify that surface against any other.

When speaking of anti-counterfeiting and authentication technology, two common metrics are used. Cost to Break (CtB) seeks to quantify the total resources which would need to be committed to circumvent the technology. It is generally assumed that a technology is useful to secure something of less value than its CtB[7]. The alternate measure of the security provided by the system is the probability that a particular counterfeit will be detected, which is defined as a function of the sampling method and efficiency, the quality of the counterfeit and the reliability of the feature[8]. Only the reliability, which is expressed as the probability that a fake feature will fail a check, is a function of the security feature itself. The rest are related to either manufacturer or counterfeiter practices.

These two metrics, taken together, help to define the qualities of a good authentication technology: it should be difficult (expensive) to produce a feature which would pass a check. The vast majority of the effort being dedicated to designing security and authentication features is expended in products and processes that all rely on the same asymmetry for their security. Security inks, chemical markers, radio frequency identification (RFID) tags, holograms and specialised printing all rely on the fact that the original manufacturer can perform a process that the counterfeiter cannot. Whether it is because the process is kept secret (e.g. RFID cryptographic keys) or takes specialised knowledge (e.g. chemical markers) or is simply too expensive for counterfeiters (e.g. laser engraved holograms), each of the security features mentioned above presume that the attacker will not be able to reproduce a process performed by the manufacturer.

In the 'arms race' of security technologies, this places the burden of innovation on the manufacturer, while the attacker need only replicate. Laser Surface Authentication is based on the premise that

security would be better served by flipping that around. Rather than creating a special feature, LSA uses the microscopic randomness that is already present as a result of the manufacturing process to identify and authenticate a product. Implementing LSA requires no special manufacturing technique, while replicating the security feature requires control far beyond that exercised by the original manufacturer. The burden of innovation, of performing new and complicated production techniques, is thus shifted onto the attacker.

LSA works by probing the microscopic randomness present on all but the most mirror-like of surfaces. The tiny imperfections in the surface are virtually unique and can be used to identify one surface from another, just as the variations in ridges and whorls allow fingerprints to identify one individual from another. When individual products, packaging or documents are authenticated in this way they are identified not just as genuine or fake but as one specific item amongst the millions produced. This further differentiates LSA from many existing and proposed security features, which can say only whether the feature is present or not. LSA combines the authentication of the security feature with item-level identification offered by per-item printing or RFID tags storing unique identification (UID) codes.

This opens up a second potential application of the technology: item-level track and trace. Track and trace refers to the manufacturer's ability to individually identify each item at each step in the supply chain, and record this information. With proper track and trace, a quality assurance office could pick up a box from a store shelf and know when and where that box was produced, to what distributor it was sold, on what truck it was shipped, through which reseller it reached the retailer, etc. The methods described so far protect against counterfeit goods—goods produced by someone other than the authorised manufacturer. But manufactures of many products face three additional challenges. When manufacture is outsourced, 'third shift' goods or overruns are produced when otherwise authorised personnel and equipment is used to produced additional goods, which are then sold through non-authorised channels. Similarly, authorised products may be intercepted at some point in the supply chain between production and the authorised retailer. These stolen goods—referred to as shrink— end up being sold through unlicensed channels. The final challenge is diversion, when authorised goods are initially sold to an authorised distributor but are then passed on and resold in another market. This is a particular problem with pharmaceuticals (where regulations and prices vary steeply from country to country), and in tobacco (where diverted cigarettes typically bypass customs controls and duties). In all three cases the item was produced on authorised machinery and thus will possess any manufactured security features. However, an accurate item-level record will show the step at which the product moved outside of authorised channels.

Attention to these problems is growing. In 2007 a law was passed in the United States requiring the Food and Drug Administration to develop a standard for identifying, authentication and tracking prescription drugs to secure the supply chain against counterfeit or diverted drugs[9]. Similarly, the European commission reached an agreement with Philip Morris for the tobacco company to implement improved track and trace capabilities for their packaging[10] to combat contraband and counterfeit cigarettes.

1.2 Motivation

By 2005 LSA was established as a laboratory technique, and a spin-out company had been formed to try and bring the technology to market. Prototype scanners had been through several stages of development and revision, and were just starting to be produced. Basic control and analysis software existed, aimed primarily at serving as a demonstration of the technology. These tools¹ had been tested in the laboratory on sheets of office paper, but in no other environments and on no other targets in significant numbers.

The approach to LSA up to this point was primarily phenomenological, and understanding had stopped at the generic term ‘diffuse scattering.’ The Nature paper introducing LSA[6] suggested laser speckle as the most obvious source of intensity variations, but this was merely a hypothesis with no supporting evidence, and had never been the focus of any investigation.

From that starting point, this thesis aims to achieve two important and related goals: firstly, to evaluate, characterise and understand the performance of LSA as a system; secondly to discover the source and nature of the optical effect. A number of unanswered questions surround the performance of the system, including: Under what conditions can it operate? On what materials is it suitable? What different factors will degrade performance, and how severely? The questions driving an investigation into the nature of LSA are: Is it a result of laser speckle? If so, why is it not more sensitive? If not, what alternative mechanisms could produce LSA?

1.3 Summary of Key Findings

1.3.1 Positional Tolerance Independent of Feature Size

Laboratory testing using ground glass as a model system will show that the feature size is not a determining factor in the sensitivity of the target item to misplacement. Match quality degraded by the same amount per mm of offset or per degree of offset for surfaces with three different levels of roughness. The impact of misplacement of the target will be quantified for surfaces with a Gaussian distribution of heights.

¹ Extant tools will be described in Chapter 2 as background.

However, experimental investigations using the LSA system on eighteen different real-world samples will show the interdependency of target material and positional tolerance. LSA signatures survive over as much as 10 mm of linear misplacement or 10° of rotational misplacement, or as little as 0.4 mm of linear and 0.25° of rotational misplacements, depending upon the target material. Rather than feature size, the material dependency will be linked to the higher-order statistical description of the surface.

1.3.2 LSA is not Spatially-Averaged Speckle

Experimental studies together with simple numerical simulation will show that under the LSA geometry laser speckle decorrelates considerably faster than does the LSA signal. Speckle will be shown to be similarly inadequate to explain the strength of the intensity fluctuations measured by LSA.

A study with sources of different degrees of coherence will demonstrate that signals of the nature and strength of LSA can be measured regardless of the nature of the light source. Illumination from a red LED and from a tungsten thermal source both produce strong repeatable signatures in the absence of speckle. It will be shown that changing the source from 671 nm to 532 nm or 808 nm illuminations does not significantly change the LSA signature, while it causes the speckle pattern to fully decorrelate.

1.3.3 Additional Results

This study will also demonstrate that LSA is not limited to a laboratory technique, and that it can perform adequately even in harsh operational environments. A field trial will demonstrate that four different categories of product, all scanned during live production, can be identified with terrific confidence. The trial will also expose some shortcomings in the technology, most importantly twin dependencies on target position and target material.

A model will be developed to describe LSA as the result of geometric optics and scattering from different scales of surface roughness. A numeric implementation of the model will show that it correctly demonstrates both the strength and the tolerance of the LSA signal.

Finally, numerical simulation will uncover the cause of the observed—but unexplained—correlation between different parts of the LSA signature.

1.3.4 Resources Produced

As key building blocks for the work to follow, this dissertation leaves behind two valuable tools: first, an experimental tool capable of automatically measuring LSA scans along a series of arbitrary paths; and second, a fully-developed numerical simulation of the technique.

The work presented in this thesis has been published in two patent families (US 2007165208 and US 2008044096) pending in the US, Europe, China, Japan, Argentina and with the World Intellectual

Property Organization. It has also been filed in one new patent application (US application serial number 12/501,303), and accepted for publication in *Optics Letters* ("The Impact of Surface Roughness on the Correlation of Diffusely Scattered Light under Linear and Rotational Displacements," submitted June 2009).

1.4 Structure of the Document

The remainder of this first chapter will establish the framework within which LSA must be discussed. A brief survey of the foundational work in optical scattering, laser speckle, and biometrics and security devices will be undertaken. This serves to introduce many concepts of use throughout the rest of the study. Where appropriate, additional work by other groups will also be covered in the specific chapter in which it is used, in sections entitled "Foundations."

Chapter 2 will provide a thorough look at the foundational work in LSA specifically. All of this work was conducted by the author's research group at Imperial College. However it was either highly collaborative or predates the author—in either case falling outside the scope of the dissertation work—and is presented here only as background material not readily available elsewhere. The principles of operation of the technique and the particular optical geometry are explained, along with the tools developed to carry out LSA measurements. Finally the analytical techniques are explained, including the signal processing routines, comparison algorithms, and metrics for performance.

Following these introductory and background chapters, the author's work is broken into four chapters, which compose the body of this study. In Chapter 3, the previously-developed scanning technique is the focus of a field-trial in a factory in Germany. This chapter examines the restrictions on real world deployments, and the performance of the technology outside the lab.

Chapter 4 returns to laboratory investigations, in order to quantify these dependencies. This chapter is broken into two parts, the first of which seeks to categorise the behaviour of a wide variety of surfaces to different potential misplacements. In the second part the same misplacements are applied to three artificial surfaces using a more flexible experimental apparatus.

Chapter 5 considers the questions of observation geometry and illumination source separately, and addresses the nature of the scattering phenomena. The behaviour of the classical picture of diffraction from a rough surface under misplacement is considered against that of LSA experimentally and numerically. A final pair of experiments investigates the significance of phase to the LSA signal.

Chapter 6 modifies the classical scattering picture to remove the contributions of phase and introduces a new geometric factor to represent the contribution of multiple scales of roughness. This model is

implemented as a full numerical simulation. Results are calculated to emulate LSA scans and target misplacement.

Chapter 7 ends this study with a summary of important results and a discussion of areas which remain for future investigation and improvement. A glossary of terms is provided at the end of the document.

The structure used here is the reverse of what might be considered traditional in a physics thesis, where the expectation is a chain of reasoning from theory to experimentation to application. However, the reverse structure is a better fit to the work presented here, both chronologically and logically. First, it more closely mirrors the path taken in conducting this research, and second, each chapter is best understood in light of the chapter which precedes it. The field work was performed early in 2006, and the wealth of questions raised therein set the direction of much of the work that followed. It was performance in the operational environment (Chapter 3) which demonstrated that positional tolerance needed to be understood more fully (Chapter 4). Similarly, it was the results of these tolerance studies which suggested that the assumed mechanism, laser speckle, might not be an adequate explanation. The experimental investigations of the underlying mechanism (Chapter 5) are what direct the decisions made in developing a functioning model (Chapter 6).

1.5 Survey of the Literature

1.5.1 Overview

Adequate analytic expressions for the scattering of electromagnetic radiation from rough surfaces were developed by the middle of the last century[11]. Although refinements were made[12, 13], the basic approach—using the Helmholtz-Kirchhoff integral over a surface approximated by plane tangents—remained unchanged. For the particular case of scattering from random rough media, the Kirchhoff or physical optics approach provided the basis for a statistical examination[14-16]. More recent works expand the statistical model into numerous specific cases, some of which begin to approach the case of LSA[17, 18]. With increased computing power more attention was paid to numeric models, both implementations of physical optics[13] and Monte Carlo simulations[19].

A comprehensive review of modern scattering models in use in physics and radar engineering was undertaken by Elfouhaily[20], with a complementary survey of numerical methods by Warnick[21]. Ragheb[22] provides an entry into work approaching the same problem from the perspective of computer graphics.

Of the many applications of speckle, the most relevant is its use in monitoring for small-scale changes to a surface. Work on small deformations and transient processes in general[17, 23] illustrate one of the

obstacles to LSA. These studies make use of the fact that speckle decoheres quickly, even for small changes of the surface being probed. This same sensitivity would seem to be a detriment to the robustness of LSA.

The final component of LSA is the metric used to compare scans. Using randomly occurring fine structure as an identifying feature is the standard approach in human biometrics. The same approach has been used both with magnetic[24] and optical[25, 26] probes to uniquely identify items, and has become a field unto itself. Physical unclonable functions (PUFs) are physical tokens whose response to a query or challenge is the result of a complex and difficult-to-model physical process[27]. LSA mirrors human biometrics and deviates from many PUFs in that it makes use of the surface as it occurs naturally. There is no need to create or embed artificial features with the specific aim of engineering a PUF as has been suggested with fibres[28] or additional coatings[29]. Much work has been done on the strength[30] and statistics[31] of biometric matching.

1.5.2 Previous Work in Scattering from Rough Surfaces

1.5.2.1 Theory of Scattering from Random Rough Surfaces

Almost any work on the scattering of waves, acoustic or electromagnetic, from a rough surface begins with an acknowledgment of the initial monograph of Beckmann and Spizzichino[11]. They begin with the Helmholtz-Kirchhoff integral, whose derivation can be found in many standard textbooks (see, for example [32]):

$$E_{scattered}(P) = \frac{1}{4\pi} \iint_S \left(E \frac{\partial \psi}{\partial n} - \psi \frac{\partial E}{\partial n} \right) dS \quad (1.1)$$

The field at a point P is calculated from the field and its derivative on a closed surface S surrounding that point, with the green's function ψ chosen so that $\psi = \frac{e^{ikr}}{r}$. With proper selection of the enclosing surface, the integral can be reduced to one over the scattering surface.

Applying the "Physical Optics" method, Beckmann[11] approximates the field on the surface as the field from a plane tangent to the surface at that point, thus:

$$E_{surface} = (1 + R)E_{incident} \quad (1.2)$$

and further assumes that the incident wave is monochromatic and planar, therefore:

$$\left(\frac{\partial E}{\partial n} \right)_{surface} = (1 - R)E_{incident} \mathbf{k}_{incident} \cdot \mathbf{n} \quad (1.3)$$

Here R is the reflection coefficient, $\mathbf{k}$ is the wave vector, and $\mathbf{n}$ is the surface normal. In order to decompose the resulting equations, the Fresnel coefficient is taken as constant along the surface. This means either assuming a perfect conductor, where this is exactly true, or using an average value over the surface. By making several other approximations (namely that the surface is constrained such that it begins and ends at the same height, that the area being integrated is large compared to the wavelength of incident light, and neglecting some edge effects) Beckmann reaches a straightforward solution for scattering from a two-dimensional rough surface. Taking E_{ideal} as the field reflected specularly for the same angle of incidence from a perfectly flat plane, where A is the total area of integration and F_3 is a constant determined by incidence and observation geometry, he finds:

$$E_{scattered} = E_{ideal} \frac{F_3}{A} \iint_A e^{i(\mathbf{k}_{incident} - \mathbf{k}_{scattered}) \cdot \mathbf{r}} dx dy \quad (1.4)$$

Where F_3 , which is the geometric constant for three-dimensional scattering, is defined

$$F_3 = \frac{1 + \cos \theta_1 \cos \theta_2 - \sin \theta_1 \sin \theta_2 \cos \theta_3}{\cos \theta_1 (\cos \theta_1 + \cos \theta_2)} \quad (1.5)$$

with the angles as shown in Figure 1.1.




Figure 1.1 In Beckmann and Spizzichino's geometric constant, θ_1 and θ_2 are the angle between the z-axis and the incident and scattered wavevectors, respectively. θ_3 is the azimuthal angle of the scattered vector, measured from the x-axis.

In an effort to account for shadowing and secondary (or higher) scattering effects, Liszka[12] follows a model laid down by Meecham[33] that uses an iterative approach to take into account multiple scatterings. Both authors develop a series solution for the field and its derivative on the surface, where the first order solution is the Kirchhoff approximation. Each successive term introduces an additional order of multiple scattering or shadowing to the previous term. Later work[34] suggests that higher order scattering terms are increasingly unimportant in the speckle regime.

From the fairly unified starting place of Beckmann's scattering model, the field quickly fragments as researchers adjust the theory to suit a particular regime of interest. The breadth of applications and degree of fine-tuning is captured in a pair of review articles and does not bear repeating here except in broad strokes. Warnick and Chew[21] reference some 275 works in the field of numerical simulation for rough surfaces. They cover seven different approaches to simulation based on approximations to scattering theories (such as the Kirchoff approximation covered above), 15 "exact" solutions based on integral equation constructions and a further three from differential equation constructions of the scattering problem. The variation is not limited to the numerical approach; Elfouhaily and Guérin[20] cite another 260 works covering 31 different approximate scattering wave *theories* which operate either at the high-frequency or low-frequency limit, or are what the authors call "unifying" techniques which attempt to bridge the two.

Finally, it is worth looking at one last set of modifications to the Beckmann integral which were not even included in the two reviews above. Ragheb and Hancock test a model which sits at the top of a chain of modifications made by introducing components of physics into the scattering models used in computer graphics, and *vice versa*. Depending on one's perspective, the key contributions building up to that model were either adding geometric optics[35] and Fresnel reflection coefficients[36] to the Lambertian model (used to calculate diffuse reflection from matte surfaces) or introducing Lambertian scattering[37] into the geometric term F_3 from equation (1.4) above.

1.5.2.2 Numerical Simulations of Scattering

In the late 1980's, researchers started to bring numerical simulations to bear on the problem of scattering from random rough surfaces. One of the early approaches using a Monte Carlo simulation[19] is illustrative of much of the work that followed it. Nieto-Vesperinas and Soto-Crespo simulated numerous one-dimensional perfectly conducting surfaces with various well-defined statistics and a linearly-polarised plane wave with varying angles of incidence. By evaluating numerically the integral expression for scattered field over an ensemble of surfaces, they developed numerous plots of intensity against observation angle for a variety of surface parameters. From these, and comparisons with analytical results predicted by Kirchhoff theory, they develop a test for the validity of those approximations:

$$\frac{\sigma}{T \cos(\theta)} < 0.2 \quad (1.6)$$

Here σ is the RMS deviation from a mean height of the surface, T the correlation length of the surface and θ the angle of incident radiation, measured from the surface normal. The authors hold that the

validity of the physical optics approach is independent of the wavelength of incident light, provided that σ is not many times larger than the wavelength of light.

As computing power increased, further complicating factors were added[13], but the approach remains similar—an ensemble of surface realisations were averaged to find peaks in the scattered intensity distribution. In all of these cases the motivation was to check the validity of the Kirchhoff, or other, approximations. Insofar as these simulations will be of benefit to LSA, attention must be turned back to the individual surface results, where intensity variations from speckle[19] dominated any of the specular or back-scattering peaks of the ensembles. Similarly, experimental results have only gone to confirm statistical models[38], while this work will focus on the behaviour of individual instances of a rough surface.

1.5.3 Previous Work in Laser Speckle

Scattering theory, as it has been discussed up to this point, has not depended on the specifics of the surface. As long as certain assumptions are valid the scattering formula apply equally well to fully specified, periodic, or random surfaces. In the particular case of a fully random surface it is often more informative to examine the statistical properties of both the surface and the resultant scattered intensity, as is the tradition in laser speckle. The Kirchhoff formulation from the previous section provides the basis for this analysis. The first and second moments of the scattered intensity show dependences on both the surface and the incident light, and these dependencies have been exploited for various applications of speckle metrology.

1.5.3.1 Statistics of Laser Speckle

As in physical optics, statistical analyses of speckle (such as that by May[15]) begin by assuming that the intensity at a point in space as a result of scattering from a random rough surface is the superposition of intensities from many small elements of the surface. When the surface is sufficiently rough (as most macroscopic surfaces are, compared to the wavelength of visible light) it can be said that each small component is statistically independent of its neighbour, and that their phase variation covers a complete 2π radians. Goodman[39] points out that when the amplitude and phase are random and independent and the phase distributed uniformly across the primary interval then the problem becomes identical to that of a random walk in the complex plane, where the final phasor amplitude $A(x, y, z)$ is the sum of N individual phasors with amplitude a_k and phase ϕ_k , such that:

$$A(x, y, z) = \frac{1}{\sqrt{N}} \sum_{k=1}^N a_k e^{i\phi_k} \quad (1.7)$$




Figure 1.2 When the intensity of light at a point is the sum of many independent intensities, the problem can be thought of as the vector sum of a random walk in the complex plane.

From this point, May and Goodman proceed along similar lines to determine the probability density function (pdf) from the real and imaginary amplitudes, and from that the pdf p and variance σ_I^2 of the intensity:

$$p(I) = \frac{1}{4\pi\sigma_I^2} e^{-I/2\sigma^2} \quad (1.8)$$

where

$$\sigma_I^2 = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N \frac{\langle |a_k|^2 \rangle}{2} \quad (1.9)$$

At this point a useful quantity, the contrast C , is often defined as $C \equiv \sigma_I / \langle I \rangle$. For fully developed, polarised speckle $\sigma_I^2 = \langle I \rangle^2$ and the contrast is always unity.

May uses the pdf to develop a host of useful properties, including the autocorrelation function and the power spectral density of the intensity in both free space and image space. The autocorrelation functions yields the average grain size for free space:

$$\sigma_g = \frac{1.22 \lambda z}{d} \quad (1.10)$$

where λ is the wavelength of light, z the distance between the observation and target planes and d the diameter of the illuminated spot. In image space z is the distance between the lens and the observation plane, and d is replaced with half the diameter of the of the lens aperture.

In a more recent work[18], Goodman extends his statistical analysis to include spatially or temporally averaged speckle signal, which is a case of particular interest to LSA. Assuming a photodetector with area A_D , he develops a parameter M depending on the covariance function of the speckle intensity and the covariance function of the detector photosensitivity, such that:

$$C = 1/\sqrt{M} \quad (1.11)$$

and

$$\left(\frac{\text{Signal}}{\text{noise}}\right)_{RMS} = \sqrt{M} \quad (1.12)$$

In the case of a photodetector with uniform responsiveness across its surface, M reduces to simply

$$M = \frac{A_D}{A_C} \quad (1.13)$$

where A_C is the correlation area of the speckle. M , then, can be taken to represent the number of grains influencing an intensity measurement. Goodman further verifies this by evaluating the detailed expression for the case where the photosensitive area is much smaller than a speckle grain, in which case M goes to unity.

All of the speckle statistics examined to this point depend only on properties of the light and the system geometry. To appreciate the avenues available for speckle metrology and the steps which lead to LSA, one can turn to Goodman's[39] analysis of the impact of surface.

Building on the simplified scattering model discussed in 1.5.2.1 and its approximations, a geometric treatment can show that the phase ϕ (see eq. (1.7)) of the reflected wave depends on the height h of the reflecting surface (here taken to be deviations from the x-y plane) and the angle of incidence θ measured from surface normal.

$$\phi(x, y) = \frac{2\pi}{\lambda} (\cos \theta_1 + \cos \theta_2) h(x, y) \quad (1.14)$$

It follows that for a height variance σ_h the phase variance σ_ϕ is as given below:

$$\sigma_\phi^2 = \left(\frac{2\pi}{\lambda} (\cos \theta_1 + \cos \theta_2) \right)^2 \sigma_h^2 \quad (1.15)$$




Figure 1.3 The angle of incidence, θ_1 , and the angle of observation, θ_2 , are measured from the average normal of the surface, $h(x)$.

To gain qualitative information, Goodman assumes Gaussian statistics for the height distribution of the surface and a Gaussian form for its normalised correlation function. Specifically, plots of this situation show that there will be a strong specular component to the reflected light only for situations where $\sigma_\phi \leq 1$. This, taken together with the scalar difference between σ_h and σ_ϕ , should be recognised as roughly² equivalent to the Rayleigh criterion[13] dividing rough surfaces from smooth for the case of specular ($\theta_1 = \theta_2$) reflection:

$$\frac{2\pi}{\lambda} \sigma_h \cos \theta < \frac{\pi}{4} \quad (1.16)$$

Goodman concludes his analysis by considering a particular physical set up (that of Fujii and Asakura[40]) and applying its parameters to his calculations. He ends with an expression for the contrast which depends on the surface statistics, the wavelength of light and the imaging geometry. He shows that for values of σ_h/λ greater than approximately 1/2, the value of C saturates at unity. This is the working definition of “fully-developed” speckle.

The extreme sensitivity to translational and axial displacement exhibited by speckle patterns forms the basis of two other implementations of speckle metrology which will be examined in the next section. May[15] lays the groundwork for these applications by examining the effects of small translations on the statistical model developed above. Her work is concerned with applications involving doubly-exposing a photographic plate and the development of Young’s fringes, and she develops statistically two conditions for translation to result in significant changes in the speckle intensity-distribution. In the case

² the R.H.S. of inequality (1.16) becomes 1/2 instead of $\pi/4$

of lateral or in-plane translation of the illuminated target, where ζ is the displacement, d the illuminating aperture diameter and z the distance from the illuminated target to the observation plane, the test is:

$$\zeta \geq \frac{\lambda z}{d} \quad (1.17)$$

Similarly, for an axial translation, ε , along the surface normal, the speckle pattern will not decohere when:

$$\varepsilon \ll \lambda \frac{4z^2}{d^2} \quad (1.18)$$

1.5.3.2 Applications of Laser Speckle

Working in the regime of partially developed speckle (less than 0.5 μm average feature size for 632 nm incident light), Leonard[41] measured experimentally the contrast of speckle images scattered from metal surfaces of varying roughness. Within the limitations of a small data set, she found agreement with the expected behaviour (a linear relationship between contrast and surface roughness). In particular, she found that varying the angle of incidence could extend the range of partially-developed speckle, since the effective surface roughness is reduced as the angle of incidence approaches glancing. The author's results represent the most straightforward application of speckle as a metric for surface roughness.

A wider range of surface roughness (1 μm to 30 μm) is tested by the speckle-pattern correlation method. While the metrology itself is too involved[42] to be of use to LSA, the method has provided data on the decorrelation of speckle patterns with small movements of the target plane. Yamaguchi[17] undertook just such a study, calculating numerically the decorrelation between a free-space speckle and image-space speckle under rotations, strain and in-plane translations. Of particular note amongst the results was the fact that free-space speckle patterns were one to two orders of magnitude more sensitive to small displacement than were their image-space counterparts.

The work of Sjødahl[23, 43] provides a good, and particularly relevant, example of using speckle to monitor for small changes to a surface. Rather than looking for translations or rotations of the surface, Sjødahl instead seeks to minimise bulk movement of the target and instead monitor localised microstructural changes to the surface itself. To gain information about localised areas of change, the speckle pattern is captured in image space, where decorrelation in a sub-section of the speckle pattern corresponds to localised changes in just one area of the surface. In all other regards, this approach is similar to the previous metrologies where the rapid decorrelation of the speckle pattern provides an

easily-detected indication of hard-to-detect changes in the target. His chosen target—paper—is a volume scatterer so his speckle measurements are sensitive to microstructural changes both at and below the surface. His chosen method of inducing these microstructural changes is to wet the paper slightly and watch the progress of structural changes as individual paper fibres absorb water, saturate, and then dry.

Paper was the most common surface analysed by LSA, and the one for which the robustness of LSA to damage or changes to the surface was most heavily tested. It is interesting to note the apparent conflict between the two methods of analysis: it has already been demonstrated that the LSA signature remains after a paper target has been soaked and dried again[44]. Sjodahl finds that saturating fibres with even a single μl of water causes a 100% decorrelation of the speckle patterns. Similarly, drying causes total decorrelation of the speckle pattern with time. Unfortunately, the reference was updated for each correlation measurement so the most valuable correlation to LSA—that of initially dry paper speckle to speckle after a wet-dry cycle—was not reported in the article.

1.5.4 Previous Work in Biometrics and PUFs

While material security technologies—from Coca-Cola’s trademarked bottle shape to the dozens of current or proposed inks, markings and tags embedded in a banknote[45-49]—have typically relied on specialised manufacturing processes, human security technologies or biometrics have relied on pre-existing, random, virtually unique and intrinsic properties of the subject. Until recently there had been only limited attempts at cross-over, such as a method described by Indeck of identifying magnetic storage media[24], or a system for recording the scattering from specially-fabricated rough surfaces[25].

As it is reading magnetic microstructure rather than topological microstructure, both the target and the probe of Indeck’s technique are quite different from LSA. He uses a magnetic head to read fine variations in the base magnetisation of magnetic storage medium. Thus Indeck’s trace is of magnetisation against position, while LSA’s trace is of intensity against position. His methodology is quite similar though, and is one of the earliest scholarly works to take this approach. A single trace corresponding to a line on the target medium is initially recorded during a registration phase and this data is stored. At a later date, a second trace is recorded and the signal verified by means of a sliding cross-correlation. When the signals are from the same physical area of the same medium there is a corresponding spike on the cross-correlation. Their signal processing, while elementary, is entirely capable of good discrimination of one-dimensional time-varying signals, and provides a starting point for LSA signature matching.

The method for generating a unique code out of scattered light described by Pappu[25], conversely, relates to LSA in its particulars but not in methodology. A laser probes fine surface structure, which is

read out through the pattern of scattered light. Their work suggests that even minimal changes to the incident geometry (a rotation of the beam by 1.7 mrad or a translation of 60 μm) causes a complete decorrelation in the pattern captured.

Pappu's work is further relevant because it formed the nucleus around which PUFs would develop as a field. The basic concept behind a PUF, using intrinsic randomness to identify a man-made item, has been investigated by the US Department of Energy's Office of Arms Control since at least the late 80's[50] and several sources cite an even earlier work from 1983[51] which the author could not obtain. Modern examples are easy to come by [29, 52, 53], particularly for valuable integrated circuits such as processors and field-programmable gate arrays.

The PUF community quickly reached a consensus[54-56] on the definition and goals of physically unclonable functions. The physical feature chosen for the PUF should be:

- 1) Simple to read
- 2) Difficult to model or characterise
- 3) Beyond modern capabilities to manufacture

The first requirement ensures that the system is deployable practically; the early PUF-like work in nuclear arms monitoring proposed using some kind of portable scanning electron microscope to inspect and verify the microstructure of each treaty-limited item[57]. Such a technically-complex solution was acceptable when the stakes were thermonuclear war, but modern applications require much simpler hardware to be feasible.

The second principle is generally interpreted as requiring a large challenge space. Modern PUFs should mirror Pappu's first system in being multi-variant, so that there is not just one possible measurement but many. And for each variation of measurement (such as angle of illumination or frequency of interrogation) the PUF should generate a different response. A good PUF will have many potential challenges, and the responses will not correlate. As a caveat, the physics of the property should be complex; an arbitrary response must not be easily calculated from a subset of measurements of the property.

Finally, the third principle means that PUFs must live up to the "unclonable" in their name. The vast majority of PUFs realise this either through using nanoscale features which are too small to be controlled during manufacture, or through a large number of three-dimensional features such that duplicating one is technically feasible but vastly more complicated than manufacturing the original, and not cost-effective.

It is tempting to think of LSA as just such a system. Indeed, Skoric—one of the leaders in the field—included LSA as an optical PUF in motivating his statistical analysis of the security and repeatability of the information contained in laser speckle[58]. However, LSA deviates from classical PUFs in two important respects.

First, LSA makes no claim to a varied challenge space. One of the principle differences between LSA and the first optical physical one-way function[25] is that LSA’s geometry sacrifices the ability to interrogate the target in a near-infinite number of ways in favour of simple hardware and ease of deployment. As a consequence LSA is vulnerable to replay attacks[59] and can only operate in a trusted-user/trusted-terminal mode. This makes it inappropriate to the embedded-device authentication that drives so much of current PUF research. But it is in keeping with the needs of many manufacturing industries, where counterfeit detection is typically managed by a few quality assurance officers and is not in the hands of consumers or retailers.

Second, LSA only ever operates in one of the two possible modes open to PUFs. Skoric himself defined the two potential modes of operation[60] as “identification” and “authentication.” Identification as he describes it is comparing two different responses. Directly comparing a pair of responses determines whether those responses are the same, and is akin to a forensic scientist comparing a signature lifted from a crime scene with one stored in a database. Authentication, on the other hand, uses a response to generate a cryptographic key. Authentication is exclusively of concern to cryptographic systems, where the random and irreproducible nature of the PUF is used to make a secure physical token. In this mode, the PUF plays the same role as a password. Although it may be feasible, there has been no move in LSA research so far to apply it to key generation. As such, the rest of this document will use “identify” and “authenticate” interchangeably, always in reference to what Skoric would describe as identification.

Where PUFs and LSA do overlap strongly is in their measures of entropy. While modern PUFs use a wide variety of physical properties, probes and processing techniques, their methodologies all converge once a single binary sequence is generated. This sequence becomes the signature or key, and can be analysed using techniques developed in human biometrics[31].

These tools, explained succinctly by Daugman[61], can also provide a measure of the security of LSA. Whether the underlying data is generated from a human iris, a two-dimensional speckle pattern or an LSA trace is unimportant. Once it is in the form of a binary sequence, the same basic comparison routines and statistics apply. Each bit is compared to the corresponding bit in the signal to be matched. Each comparison, broadly speaking, is taken to represent an independent two-outcome (match or not-

match) test, or a Bernoulli trial. An ensemble of Bernoulli trials generates a binomial probability distribution, where the expected probability, p , of m matches out of N bits takes the form[62]:

$$p(m, N) = \frac{N!}{m!(N-m)!} p_0^m (1-p_0)^{N-m} \quad (1.19)$$

Here, p_0 is the probability of a single bit matching. The values of N and p_0 are calculated from the mean, μ , and standard deviation, σ , of the data by:

$$p_0 = \mu \quad (1.20)$$

$$N = \frac{\mu(1-\mu)}{\sigma^2} \quad (1.21)$$

Having established the functional form of the data, information can be extracted from the fitted values for N and p_0 . In his analysis, Daugman makes the important point that the effective N calculated from the mean and standard deviation will not be the same as the number of data bits compared. Rather, the fitted N is the number of independent bits, and will be referred to as N_{eff} . The ratio of independent bits to tested bits tells us the degree of correlation between sampling points.

Daugman goes on to describe two more related statistical tools, Hamming distance and d' . Hamming distance is defined simply as the ratio of different bits to total bits, and can be calculated for each comparison. A large number of comparisons can be plotted as a histogram, where the x-axis is hamming distance. Since a comparison of two different scans of the same iris should match in many more bits than two scans of different irises, you would expect to see two clusters of data in such a plot. One peak will contain unrelated iris comparisons and be centred around a hamming distance of 0.5, while the other will be composed of like iris comparisons and be centred around a lower hamming distance. Figure 1.4 shows an example from Daugman's work[61].




Figure 1.4 When plotting a histogram of the hamming distance between different signatures, two clusters emerge: signatures of unrelated items have hamming distances distributed around 0.5, signatures from the same item will show a lower hamming distance. Reproduced from [61].

These hamming plots show diagrammatically the level of discrimination offered by the chosen metric, indicated by the degree of overlap of the two peaks. The decidability[63] index d' quantifies this separation, where the subscripts correspond to the two different peaks, as:

$$d' = \frac{|\mu_1 - \mu_2|}{\sqrt{\frac{\sigma_1^2 + \sigma_2^2}{2}}} \quad (1.22)$$

While d' is useful as a rough figure of merit, it assumes normally-distributed data in hamming space. Even in the case of binomial distributions, where it will approach Gaussian for a large data set, the most severe divergence from normal is in the tails which d' seeks to quantify. Its usefulness is further undermined by the fitting algorithms which are used to find matches, and which typically skew the data towards lower hamming distances. Both non-Gaussian 'same' data and the skewed 'different' data are visible in Figure 1.4.

2

BACKGROUND ON LASER SURFACE AUTHENTICATION

*“I remember the days of old; I meditate on all thy
works; I muse on the work of thy hands.”*

—Psalms 143:5

2 Background on Laser Surface Authentication

2.1 Introduction

This chapter will cover the current state of the art in Laser Surface Authentication, which is necessary as background for the rest of this study. It will include work done concurrent to, but outside the scope of, this dissertation by engineers at Ingenia Technology and researchers at Imperial College to develop the technique.

The second section will cover the general principles of operation behind the technology, including the framework in which it is typically deployed and the limitations of the system. Section 2.3 describes the optical and mechanical geometry of the core LSA sensor, while the following section describes the two primary hardware devices developed to use that sensor. The most critical factor in using the current hardware—accurate positioning—will be examined separately in Section 2.5.

The next two sections detail the creation of a signature from a target surface, beginning with an explanation, in Section 2.6, of how the target surface is probed. This is followed in Section 2.7 by a description of the signal processing techniques used to turn the raw data into the LSA signature.

Finally, Sections 2.8 and 2.9 explain the algorithm used to compare two LSA signatures and the analysis needed to determine if the match passes or fails, respectively.

2.2 Principles of Operation of LSA

LSA seeks to use the fine variations in light diffusely scattered from any optically rough surface to generate a signature or ‘fingerprint’ which is unique to that surface. All but the most mirror-like of surfaces are rough, and will produce such a signature. Carefully measuring the diffusely scattered light from the surface of, for example, a credit card allows a duplicate to be distinguished from the original based on the existing microscopic imperfections in each one.

Just like a human fingerprint, a single LSA signature in isolation is of no value. Two signatures must be compared to determine if they came from the same surface, which means any item that is to be authenticated must have already been scanned and the original signature must be available. In practice, this means there are two different modes of scanning. First, new items are enrolled; their signature is captured for the first time and stored. Most commonly this is done at some point during the production process, as the items are already moving on a production line.

Later, when the item is to be authenticated, there is a second scan—the verification scan. The original signature must be retrieved, and checked against the verification scan. If the storage is linked to the

item, like with an encrypted barcode printed directly on the product, or if there is a unique identity code (UID) which can be linked with the signature in the database, then this is simply a matter of retrieving the correct signature from storage and comparing it with the verification scan (one-to-one). For cases without a UID, the verification scan must be compared against each scan in the database individually to determine if the item being scanned matches one of the known items in the database (one-to-many).




Figure 2.1 As a closed-loop system, LSA signatures must be registered and stored for every item at the time of production. A later scan in the field can be compared to the original signature to verify the authenticity of the item.

LSA cannot inherently tell a genuine item from a counterfeit; in general all LSA signatures are equally different. Two repeated scans of a counterfeit item will match as well as two repeated scans of a genuine item, and scanning an item which was not enrolled will produce the same result (no match found) whether the item is a counterfeit or genuine. In practice this means that enrolment must be done in a trusted system, and must be secured. For barcodes or on-item memory, this means encrypting the signature with a private key during enrolment; for databases it means certifying the enrolling systems. ‘Third shift’ goods, described in Section 1.1, are therefore rejected as well. Unauthorised production cannot be entered into the database, so it will be rejected as counterfeit even though it is identical to a genuine product.

A single piece of software, IsaControl™, governs this process. It interfaces with all types of LSA hardware, processes scan data, sends scans to and reads scans from a database or barcode, and performs comparisons between scans. IsaControl is written to use a common driver definition for hardware. This is helpful for the research described in this work, because it allows a custom piece of laboratory hardware to be interfaced with the existing system. By writing a new driver for experimental kit, the intensity data can enter IsaControl exactly as if it came from a commercial piece of hardware. The end result is that experimental data can be processed and compared using these existing routines, exactly the same way data taken in the field trial is processed and compared.

2.3 Illumination and Observation Geometry

At its most simple, Laser Surface Authentication requires only four things: an illumination source, a detector, a target surface, and relative motion between them. Although a series of different detection geometries have been tried over the last few years, all LSA systems currently in use have settled on a single geometry, which has become the defined sensor block.

For the purposes of the LSA geometry, the illuminating optical axis is the z axis and the target plane is the x-y plane, so the illumination is normally incident. An inexpensive commercial 650 nm laser diode is used, which comes pre-packaged with an adjustable collimating lens behind a 5 mm aperture. For the sake of repeatability the collimating lens is positioned about 400 μm too close to the laser diode, so the beam is still very slightly diverging rather than perfectly collimated. The nearly-collimated beam passes through a cylindrical lens with a focal distance of 8.3 mm, which focuses the beam only in the x dimension. The focal plane of the cylindrical lens defines the target plane. In this plane, the illumination is a focused line 5 mm long in y and approximately 100 μm wide in x.

A photodetector is positioned such that it is gazing at the centre of the laser line, in the target plane. The optical centre path, from the laser diode to the laser spot to the detector, is parallel to the x-z plane. The detector is at an observation angle of 45° from the z axis, and is 10 mm back from the centre of the focused line. The photodetector is behind a 4.5 mm diameter collecting lens, with an acceptance angle (defined as the angle for which responsivity has fallen to half of its maximum value) of $\pm 10^\circ$




Figure 2.2 A laser diode (A) is normally incident on the target surface, focused to a 5 mm x 100 μm spot (C). Diffusely reflected light is measured by a photodetector (B), with an observation angle of 45° and at a distance of 10 mm.

This basic geometry is outlined in Figure 2.2 above. For working LSA sensors, a second photodetector is positioned opposite the first, at -45° . A single 'slice' of the LSA sensor is a mechanical mount which combines all of the optical elements above: the laser diode, collimating lens, cylindrical lens, and two photodetectors and collecting lenses.

2.4 Hardware

A complete LSA sensor block consists of three slices or channels, aligned in x and z, and stacked in y with an 11 mm pitch. The sensor block, along with a camera used for alignment, is illustrated in Figure 2.3.




Figure 2.3 The LSA sensor head. Reproduced with permission from the IsaControl User's Manual.

The majority of LSA deployments require only two types of scanners: one for use on the production line, and one for use in the field.

The production line scanner (Figure 2.4) is simple, consisting only of the sensor head from Figure 2.3 behind a glass window, in a metal enclosure. A shielded data cable carries power to the lasers, detectors, and camera, and returns the video signal and the analogue photodetector signals to a PC-based acquisition system.




*Figure 2.4 Two views of the ILS2100, the standard LSA hardware for use on production lines.
Reproduced with permission from Ingenia Technology.*

The sensor can be mounted on any system where the target item is already moving, like a production line's conveyor belt. The sensor must be oriented such that the item is moving along the sensor's x axis. The focal plane of the lasers is 3 mm beyond the front face of the glass window, so that the target item is in focus without making contact with the ILS2100. Data acquisition is performed by a PCI card, controlled directly by IsaControl. The hardware and electronics are designed to have a maximum rate of 100 kHz. This can easily cope with production speeds of 2 m/s, and higher speeds are possible with a reduced spatial resolution. The acquisition system has external inputs for encoder data and for a trigger signal. The system can begin acquiring scan data either on a digital trigger signal or by monitoring any one of the photodetector signals for a spike or step characteristic of a box passing beneath the lasers.

The portable scanner (see Figure 2.5) uses the same LSA sensor head, this time mounted to a linear motion stage and gazing up through a glass window. In this unit the focal plane of the sensor coincides with the outside face of the glass window, so a target item is placed directly onto the unit (see Figure 2.5, inset).



Figure 2.5 The IFS4000 is the standard portable LSA hardware. Inset shows how an item is positioned for scanning. Reproduced with permission from Ingenia Technology. Inset reproduced with permission from Ingenia Technology.

The IFS4000 has a maximum scan length of 70 mm. The stage travels at approximately 15 cm/s, and data is sampled at 15 kHz. An onboard microcontroller controls the motion of the scan head and the data acquisition, and communicates with its IsaControl driver through a USB interface.

2.5 Positioning of the Target Item

In order for two scans to match well they must be taken from the same position³ on the target item, which means the correct x, y, and z coordinates as defined in Figure 2.2. The x position is the position of the scan data within the scan. Variations in this dimension are handled through an encoder and through offsets introduced when comparing signatures, both which will be described in the following sections.

The z position is defined as described above: 3 mm from the surface of the ILS2100 or in contact with the surface of the IFS4000. With the field hardware this is automatically true provided the item sits flat on the scanner. Care must be taken when installing production line hardware, to ensure that the sensor is the correct distance from the item. In most cases the sensor is above the target item and the target item is sitting flat on the belt, so there is little variation in this dimension after installing the head.

Ensuring the correct y positioning requires the most care, and can be done in one of two ways. Correct y positioning for items with at least one hard edge is described first.

³ How accurately the positioning must be repeated will be addressed in Chapter 4.

On the field scanner, the L-shaped bracket (visible in Figure 2.5, illustrated in Figure 2.6) provides a hard reference for the positioning of the target item. With the item pressed into the corner of the bracket, the centre channel is always positioned at a fixed distance, y' , from the edge of the item. During assembly of the scanner, a calibration process ensures that y' is 26.5 mm




Figure 2.6 The red lines indicate the path swept out by each of the three lasers. Factory calibration ensures the two distances x' and y' . Reproduced with permission from the IsaControl User's Manual.

Correctly installing a production line sensor requires positioning the sensor such that its centre channel is also 26.5 mm from the edge of the target. This is generally achieved through a set of mechanical guides. For example, conveyor feeding systems can be arranged to ensure that the item is flat against a main guide rail as it passes the sensor. The sensor head can then be installed with the correct y' relative to this guide rail. This configuration is illustrated in Figure 2.7.




Figure 2.7 **(a)** In a production setup, the centre laser diode (A) is arranged to shine down on the target item (B) as it moves along the conveyor belt (C). A mechanical guide (D) ensures that the edge of the item is a fixed distance, y' , from the centre laser. **(b)** With the field unit, the central laser (A) shine up at the target (B) and sweeps beneath it. A mechanical bracket (C) on the face of the scanner ensures that the edge of the item is the same distance, y' , from the middle channel.

The method of alignment described above is faster and easier for the operator, but it is only suitable for rectangular items, or items that can otherwise be fit to the bracket. For irregularly shaped items the alignment is performed using the camera shown in Figure 2.3. Factory calibration ensures that the centre of the camera's field of view is always lined up with the centre of the channel 2 laser. If the camera is aligned to a predefined template using markings on the item, this will ensure that the LSA scan is taken in the same place for every scan. Figure 2.8 shows an example of this process.




Figure 2.8 Using the video feed to align markings on an Ingenia business card. Reproduced with permission from the IsaControl User's Manual.

2.6 Scans – Acquiring Data

To generate the raw data for an LSA scan, the sensor block must be moved relative to the target surface. In the field scanner model this means mounting the sensor block on a linear motion stage and passing the block under a stationary target. In the production line model the sensor block is stationary while the target surface passes beneath it. In both cases the results are the same; it is only the relative motion which matters.

The signal from the photodetectors is amplified electronically and sampled by a 16-bit analogue-digital converter. The rate of data acquisition is matched to the speed of the relative motion, such that data is sampled every 10-25 μm .

Each photodetector 'sees' the diffusely scattered light from its laser spot. As described in Section 1.5, the exact intensity of the scattered field at any point in space is dependent on the specifics of the surface being illuminated. Thus the detected intensity will be slightly different for each position of the surface which is illuminated. By measuring the intensity every 10 μm along a length of the surface, a profile is developed of the diffusely scattered light for thousands of illuminated positions.

The length of the surface scanned for an LSA signature varies by the material type and the desired statistics, but it typically ranges from as little as 15 mm to as much as 100 mm. An intensity profile is measured for each detector in each slice, so a scan with a standard head will have 6 different profiles, corresponding to left and right views from three different strips of the surface. The data from one illumination path is referred to as a 'channel,' and each channel contains data from two detectors.




Figure 2.9 The top part of the figure shows the raw intensity data, from both detectors of all three channels, for a 60 mm LSA scan of a roughened glass surface. The raw data is generated from the six photodetectors. Laser and photodetectors are represented as they were in Figure 2.2. Each point in the trace corresponds to the intensity of the diffusely scattered light registered at each detector for a particular illumination position.

The random variations in measured intensity are a result of the random roughness of the surface, and are reproducible. Provided the surface has not changed, illuminating the same region of the surface will generate the same integrated intensity at the photodetector. Figure 2.10 shows a 10 mm segment of two repeated scans of the same ground glass surface.




Figure 2.10 The red and blue traces are the raw intensity values of one photodetector from two different scans of the same surface.

2.7 Processing Raw Intensity Data

2.7.1 Encoding Velocity Information

The repeatability of the LSA data, evident in Figure 2.10, is the basis of its use as an authentication technology. The raw data above was taken on laboratory equipment with extremely well-controlled motion. In general the source of motion is not accurate down to 10's of micron over as much as 100 mm. For this reason an encoder signal is usually acquired along with the photodetector data.

The encoder is a digital signal where one period corresponds to a fixed distance travelled. The encoder data is analysed and the local velocity for each point is determined, assuming a constant velocity across each encoder period. Using this velocity data, the raw data is recast onto points of fixed spatial resolution.

The encoder signal is acquired in parallel with, and at the same rate as, the intensity data. Since the data is acquired at a fixed interval in time, the encoder data will be stable for a fixed velocity, while the period will appear to shorten for faster motion and elongate for slower motion. This is illustrated in Figure 2.11.




Figure 2.11 One period of the encoder signal corresponds to a fixed distance travelled. When acquired at a fixed rate in time, the encoder signal gives a measure of the velocity. Reproduced with permission from the IsaControl User's Manual.

In the field model the encoder is an additional optical sensor mounted with the LSA head which reads a pattern of lines printed at a known resolution and mounted to the inside of the case. In the production model, this usually takes the form of a magnetic, optical, or mechanical switch mounted on a rotating component somewhere in the belt system. If no driveshaft or pulley is accessible, a friction wheel can also be mounted to the underside of the conveyor belt. In either case, one rotation must always corresponds to the same amount of travel of the belt.

2.7.2 Application of the Band-pass Filter

The detection electronics used for LSA are DC-coupled, so the raw intensity data has a non-zero mean that is determined by the background lighting and the overall reflectivity of the surface. For non-rigid surfaces, wrinkles and creases in the surface will also create long-wavelength influences on the scattered light as the wrinkles bring the detector closer or further from the average specular direction.

Sampling at 10-20 μm steps is oversampling for a beam width of 100 μm . Correlation between neighbouring points is useful in the comparison routines, but it also subjects the system to higher-frequency noise than is necessary.

To eliminate both long-wavelength DC changes and high frequency noise the spatial data is passed through a Bessel band-pass filter. The specifics of the filter vary from material to material. The filter is typically first or second order, with a high frequency cut-off between 100 μm and 400 μm and a low

frequency cut-off between 800 μm and 1.5 mm. The data from Figure 2.10 is shown after AC-coupling in Figure 2.12.




Figure 2.12 The red and blue traces are the data from two different scans of the same surface, after conversion into spatial coordinates and filtering with a 2nd order 300 μm – 1100 μm band-pass.

After being AC-coupled, the data can be additionally smoothed by convolution in the spatial domain with a normalised Gaussian distribution. When used, the Gaussian kernel has a width of anything between 30 μm and 200 μm .

2.7.3 Developing the Binary Signature

Finally, the smoothed AC-coupled signal is used to generate two binary sequences: the LSA signature, S , and the mask, M . The binary sequence contains the same number of points as the AC-coupled signal, except each point is only a single bit rather than 16 bits. The bit stores only whether the point is above or below zero. Each point in the AC-coupled signal is mapped directly to the corresponding point in the binary signature; the binary sequence's bit is set to 1 if the AC value is zero or positive, and set to 0 if the AC value is negative.

In reducing the AC-coupled data to a simple 'above or below the mean,' the signature becomes particularly sensitive to values very close to zero. Small differences in the AC value of most points will not change the corresponding bit value, but when the original value is already close to zero then a small difference can flip the binary sequence's value. To combat this, a mask is generated. A minimum threshold value is set, as a fraction of the AC signal's standard deviation. Typically this value is between 5% and 50% of one standard deviation. For each point the mask is set to 0 if the absolute value of the

AC point is less than this threshold and 1 otherwise. In this way, the mask identifies which bits should be ignored because they are more prone to errors.

2.8 Comparing LSA Signatures

The heart of the comparison process is the calculation of the bit match ratio (BMR) between two scans, A and B. The scans are often referred to as the comparison scan and the reference scan. Mathematically speaking, S and M are n -dimensional binary vectors and the BMR is the inverse of the hamming distance between them, normalised by their length n . More simply, it is the ratio of bits which agree to total bits in the sequences. Taking the masks into account as well, the BMR between two scans, A and B, is given by

$$BMR = 1 - \frac{|(\overline{S_A} \otimes \overline{S_B}) \cap M_A \cap M_B|^2}{|M_A \cap M_B|^2} \quad (2.1)$$

Where $\cap$ and $\otimes$ are the bitwise AND and XOR operations, respectively, and $\overline{A}$ is the negation or one's complement of A. Note that for binary vectors, the squared magnitude is equivalent to the sum of the elements, or alternatively to the number of 1's in the sequence.

Non-rigid materials may have stretched or contracted between enrolment and verification, or the scan paths may not be exactly aligned between the two scans, or the encoder signal may simply not be perfect. In any of these cases, the two spatially-converted signals will not line up to the nearest 10 μm across the full scan length. In addition, most motion control systems are not accurate enough for the scan to start at *exactly* the same position on the item each time. The equation described above cannot accommodate any of these imperfections.

Instead, each scan is broken up into a number of segments or blocks, depending on scan length and material. These blocks must be long enough that there is sufficient data contained within each block for a reasonable comparison, but short enough that any errors in tracking over the block's length are not significant. Typically, blocks are 5 mm – 10 mm long.

The BMR between corresponding blocks is calculated as in equation (2.1), for a variety of offsets. Offsets are introduced by left- or right-shifting the comparison scan. The BMR is calculated for each block for offsets of as much as several millimetres forward and backward, and the offset which produces the best BMR is recorded. A trivial example is illustrated in Figure 2.13.




Figure 2.13 BMRs are calculated for a range of offsets, to determine the location of the best match for each block.

The offset at each block indicates the difference in x position between that section of the two scans. The linear offsets introduced to each block are not intended to correct for misalignments and stretches on their own, but rather to determine a rough-cut look at the stretch required to align the two scans. The offsets of each block between two scans are shown in green in Figure 2.14.




Figure 2.14 The best offset found for each of 6 blocks between two scans (green), and the 2nd order polynomial best fit to those offsets (blue).

A polynomial fit is calculated from the individual block offsets. Depending on the nature of the stretching, this can be anything from a zeroth order fit (indicating only an offset in start position) on up to nth order for complicated stretches (usually caused by jerk in a production line). First order (indicating a linear stretch or an incorrect velocity in one scan) and second order (indicating an incorrect velocity in both scans or an uncorrected acceleration in one scan) fits are the most common.

The fit data, which was calculated block-by-block from the binary sequences, is then applied as a smooth stretch to the AC-coupled data of the comparison scan. A new binary sequence and mask are calculated from the stretched data, and a final BMR computed for the full reference and stretched comparison scans.

This final BMR is the one reported by IsaControl when performing a comparison. It should also be noted that when comparisons are calculated using only one block, this is exactly equivalent to repeatedly calculating the BMR as per equation (2.1) for the range of offsets and reporting the best BMR.

2.9 Metrics for Evaluating Comparisons

Developing statistically significant results from, and thus confidence in, the LSA technique requires analysing large numbers of comparisons. For the purposes of this type of analysis it is useful to separate results into two categories: the BMR from comparisons between repeated scans of the same item, and the BMRs from comparisons between scans of different items. These results are generated from

collections of scans from a set of nominally identical items, for example a ream of office paper. First, a reference library is created by scanning each item once. Second, a comparison library is generated by taking a second scan of every item. The two data sets, the like-item distribution and the unlike-item distribution, are generated from comparing each pair of scans of the same item and each unique pair of scans of different items. In total, two libraries of n scans will produced n like-item comparisons and $\frac{n^2-n}{2}$ unlike-item comparisons. This is illustrated in Figure 2.15. If only the unlike-item data is required, the reference library and the comparison library can contain the same data.




Figure 2.15 For two collections of scans (the reference collection, in purple, and the comparison collection in blue) of the same 5 items, the like-item distribution is made up of pairs of scans of the same item (green). The unlike-item distribution is made up of unique pairings of different scans (solid red) while repeated pairings of unlike items (dotted red) are ignored. Reproduced with permission from the IsaControl Technical Supplement.

Since the metric used to compare LSA scans, BMR, is the complement of the hamming distance, the same tools described in 1.5.4 can be used to analyze BMR. For a collection of comparison results, inverse hamming plots can be created by plotting the two histograms of the like-item and the unlike-item BMRs on the same horizontal axis. An example is shown in Figure 2.16. The independent number of bits, N_{eff} , can be calculated from the unlike-item distribution using equations (1.20) and (1.21).

Because each bit has only one of two possible values, half of the bits would be expected to match between two unrelated strings, and thus the unlike-item distribution should be centred at 0.5. The offsets and fitting involved in the comparison routine, however, improve the average unlike-item match.

How far above 0.5 the mean drifts is an indication of how aggressively the comparison algorithm was allowed to search for the best match. In the example hamming plot below, the mean is closer to 0.54.




Figure 2.16 A hamming plot generated from 50 sheets of paper. The like-item ("Same") distribution is in green, and the unlike-item ("Different") distribution is in red.

The separation between the two distributions provides a measure of the discriminatory power of the system. The decidability index discussed in Section 1.5.4 is one measure of this (and is shown on the right in Figure 2.16) but it assumes a Gaussian distribution for both data sets. While the unlike-item distribution is the result of a random process and is generally a good fit to a Normal distribution, the like-item is influenced by human error (manifest in the accuracy of the second scan) and tends to more complicated, often asymmetric and heavy-tailed. An empirical false accept rate (FAR) can be calculated solely from the unlike-item distribution.

Match/no-match decisions are made by choosing a threshold value for BMR, above which the signatures are said to match and below which the comparison is said to fail. The FAR is the probability that two unrelated scans will have a BMR equal to or greater than the threshold value. This probability is calculated from the complementary⁴ cumulative distribution function (ccdf) of the unlike-item distribution, evaluated at the threshold value. Assuming a Normal distribution with mean μ and standard deviation σ , and where erf is the error function[64] and the independent variable x is the BMR, the probability of x being a part of the unlike-item distribution is:

⁴ Where the cdf is the integration of the pdf from negative infinity to x , the ccdf is the integration from x to positive infinity.

$$P(x) = \frac{1}{2} - \frac{1}{2} \operatorname{erf}\left(\frac{x - \mu}{\sigma\sqrt{2}}\right) \quad (2.2)$$

Because the probabilities encountered tend to be very small, it is easier to work in terms of the exponent, so Security Level (SL) is defined as:

$$SL(x) = -\log_{10}[P(x)] \quad (2.3)$$

Thus the FAR for a given threshold, x_t , is approximately $1:10^{SL(x_t)}$.

Determining the BMR threshold involves weighing three factors: the statistics of the material as determined by a hamming plot, the acceptable error rate, and the size of the search space. If, through local storage or the use of a UID, the correct signature is known exactly, then the search space is one. This is known as a one-to-one comparison. If a database must be searched, referred to as a one-to-many comparison, then the search space is the number of scans which must be checked, after any available meta data has narrowed the list of potential matches as far as possible.

Searching a database requires computing the BMR with each signature in the database, and each must be considered a separate comparison for the purposes of computing the error rate. For example, if an acceptable error rate is one false positive per million challenges, this would be achieved by a security level of 6 for a one-to-one match. However, if there are a ten million items in the database then a single verification request (one challenge) will result in ten million comparisons, and thus about ten false positives. If D is the number of scans in the search space, then an acceptable error rate of one per N challenges becomes a rate of one per ND comparisons. In the previous example, then, the threshold must be chosen to give a security level of at least 13.

2.10 Conclusions

This chapter has described the current state of LSA as a technology, which will be relied upon as background for the chapters which follow. It began with an overview of how LSA is used as an authentication technology. This was followed by a detailed explanation of the optical geometry and scanning method associated with LSA. The current LSA hardware, which will be used throughout this study, was described next, along with the operating procedures required to ensure correct positioning. The chapter finished with an explanation of the measure of similarity between two scans, their BMR, and the system used to transform BMR into a binary (pass or fail) decision.

3

PERFORMANCE OF LSA IN A FIELD TRIAL

*“And the rain descended, and the floods came, and the
winds blew, and beat upon that house; and it fell not:
for it was founded upon a rock.”*

—Matthew 7:25

3 Performance of LSA in a Field Trial

3.1 Introduction

The US Department of Defense (DoD) defines nine Technology Readiness Levels (TRL) between “basic principles observed” and “system proven.”[65] Most steps involve a progression in the technology: from principles to concept to experimentation, then on to components, subsystem, and full system prototypes. The next step, from TRL 6 to TRL 7, is to move the same full system prototype from testing in a simulated operational environment in the laboratory to testing in an actual operational environment. This “represents a major step up from TRL 6,” according to the DoD guidelines.

Certainly, moving from a laboratory environment to a true operational environment represents a major hurdle in technology transfer. Technologies which perform acceptably in a well-controlled environment with expert operators can stumble when introduced to the real world and non-expert users. Beyond success or failure, field trials are the only way to uncover the unpredictable problems which arise from a new technology. The most careful laboratory testing can only prepare for the expected problems.

For LSA this step was taken in cooperation with a potential customer, to allow a week long field trial on the floor of one of their factories. One of the reasons LSA is interesting as a security and authentication feature is that it is covert. Because it does not alter the original product or packaging in any way there is no way to tell when, or what part of, a surface has been scanned. This secrecy is of value to potential users and as a result the field trial is governed by non-disclosure agreements. As such, the brand owner, factory, and products involved must remain anonymous. As much description as possible will be given without compromising those identities. In a few cases this will include blacking out regions of photographs. Only identifying features are thus obstructed, and no data of interest to this scientific study is redacted.

3.2 The Operational Environment

This field trial took place in a factory in Germany, where boxes are produced for a number of different brand owners in consumer electronics and household goods. The trial was in conjunction with a parent company in household and personal goods, and so all of the boxes under investigation were from a handful of valuable brands deemed most at risk for diversion (see Section 1.1).

Paperboard packaging is originally produced in large (1 m^2) sheets of paperboard from which a number of boxes will be produced. Most of the details of the package, including background colours and the myriad of printing (brand identification, product descriptions, universal product code barcode, ingredients, standards compliances, etc.) are printed while the cartons are still together in the sheet.

After these print runs, the individual cartons are cut out of the sheet and creased. These individual cartons, still unfolded and flat, are then run through a machine which applies any item-level processes. For the products under investigation here, this meant printing a second barcode containing a unique ID (UID) code for each item. Finished cartons are packaged and shipped to another facility where they are glued, folded, and filled with the product.



Figure 3.1 An LSA sensor head (circled) mounted above the conveyor belt on the machine which processes individual cartons, before they are formed into boxes.

The individual-carton machine was the obvious point at which to capture the LSA signature for two reasons. First, each item passed through one at a time, which meant only a single sensor head needed to be mounted. Second, it was at this stage that the UID was determined. The end-customer for this deployment wanted to have a central database which stored the signatures of all the packages they produced. In order to interact with their existing track and trace solution, each signature in the database needed to be linked to that product's UID.

The LSA computer constantly monitors the DC level of one of the detectors. It determines when the intensity level rises above a certain threshold, which indicates that a box has moved under the sensor, and triggers the acquisition of scan data. Here, acquisition was at 100 kHz, and between 10 and 17

thousand points were acquired depending on the product. The conveyor belt moved at approximately 1.1 m/s, which gave a raw spatial resolution of approximately 11 μm and a total scan length of 110 mm – 187 mm. The number of points, and therefore the scan length, was adjusted for each product to ensure that data was captured for the entire length of the carton. The machine processed 2 – 3 boxes per second, and had an encoder with a 1 mm period.

Moving from the laboratory to a live environment brought with it three immediate disadvantages. First, the conveyor involved was older and less smooth than the one used in laboratory testing. Variations in speed meant encoder processing was even more important than normal, as there could be significant variations in the instantaneous velocity. The belt and motor also introduced vibrations both in- and out-of-plane. In-plane vibration manifested as a general degradation of the LSA signal, while out-of-plane vibration results in oscillations in the detected intensity, as the product moves slightly closer or slightly further from the sensor head.

Second, the production machine was not dedicated to LSA, meaning there were numerous other concerns which influenced its mechanics. Various rollers and guides were in place to prepare cartons for printing, to prevent jams and to feed cartons in the desired direction. In addition to being obstacles, some of the rollers—in an effort to keep one part of the box flat for printing—would warp the carton such that the edges were lifted off of the belt. This both changed the focal distance for the LSA sensor and made the edges of the carton much more susceptible to vibration.


Third, a production environment is electrically very noisy. Each production line involved several industrial motors along with high-current power cables, all of which generated radiative noise. In addition, several motors were connected to the same mains loop on the factory floor that powered the LSA hardware, and were grounded to the same metal chassis that supported the LSA sensor head. This meant that both the power and the ground to the LSA electronics carried considerably more noise than the equivalent system in the laboratory.

3.3 Four Products for LSA Identification

3.3.1 Overview of Products

Four different boxes were being produced during the trial week. Board 1 and Board 3 were moderate weight and moderately smooth paperboard, with a matte finish. Board 3 had a complicated background pattern, while Board 1 did not. Board 2 was a fairly standard black paperboard, but with a particularly smooth black coating on the outside. Board 4 is the coarsest of the paperboard samples. It is a natural brown material with faint corrugation and printed green on one side.


Board 1



Board 2



Board 3



Board 4



Figure 3.2 The normally-flat cartons were folded for these photographs, so both the inside and outside could be seen. In each case the larger tab shows the underside of the carton (which will become the inside, when formed into a box). The smaller tab shows the upper (outside) surface, which was the surface scanned with LSA. A few regions have been blacked out to hide the identity of the products.

Each product is discussed individually in the following four subsections (3.3.2 – 3.3.5). In each case the scan data are presented in the same way (Figure 3.3, Figure 3.6, Figure 3.7 and Figure 3.8). Plot (a) always shows the raw intensity data (on a 16-bit scale) as measured by one detector, for the full length of the scan. This is plotted in points, where a point is $10\ \mu\text{s}$ or approximately $11\ \mu\text{m}$. The encoder signal has not been processed yet at this point, so the spatial resolution is only approximate. DC intensity shows the influence of printing and the colour and reflectivity of the surface.

The region of the full scan which was selected for further processing is circled. The plot (b) always shows just that region after applying the encoder signal. It still shows raw DC intensity, but now against position in space, with a fixed spatial resolution. Comparing the (b) plots from each of the four figures gives a sense of the different fractional intensities for each product.

The third plot shows the same region, in spatial coordinates, after passing through a band-pass filter (see Section 2.7.2). The fourth plot shows the binary signature derived from the AC-coupled signal. These two plots give a sense of the density of effective bits in the scan. In each case the final LSA

signature used to compare scans was made of the data in plot (d) along with similar data from two additional channels. Each channel was between 37 mm and 43 mm long.

The following discussion will highlight what is unusual about each product. Board 1 will serve as the baseline, while Boards 2 – 4 will each be used to demonstrate some of the non-ideal features of real products.

3.3.2 Board 1

The raw intensity data from a scan of Board 1 makes clear the negative influence of printing. The large swings in reflected intensity from 1,000 – 3,000 and 10,000 – 12,000 in Figure 3.3 (a) both correspond to regions of dense printing on the surface of the box. While there will be LSA present in those regions, it will be swamped by the strong changes in intensity from fine print, which will be the same on every carton. Conversely, in the region selected the DC Intensity appears uniform, and the only variations in intensity are those of interest to LSA.

The AC-coupled signal (c) is generated with a 2nd order Bessel band-pass filter with upper and lower cut-offs at 400 μm and 900 μm , respectively⁵. The density of bits will provide a good standard against which to measure the other three surfaces.

⁵ Cut-off frequencies were chosen empirically to optimise the results for each surface. The method will be discussed in Section 3.5.




Figure 3.3 Board 1, full DC in time (a), cropped DC in space (b), cropped AC in space (c), binary signature in space (d).

3.3.3 Board 2

In test cases in the laboratory, boxes were anywhere from 2 cm to 4 cm tall, and a box under the sensor head produced a much larger DC signal than the surface of the belt so far out of focus. This produced a clear top hat in the intensity profile (see Figure 3.4), the rising edge of which was used to trigger the beginning of a scan acquisition.




Figure 3.4 The rise in DC intensity associated with a box moving under the LSA sensor. The red line indicates a potential threshold value for triggering. Reproduced with permission from the IsaControl User's Manual.

With the unfolded boxes used in this trial, the target surface is only the thickness of the paperboard (~1mm) above the belt's surface. This fact, combined with a smooth black material, meant that the DC level with Board 2 in focus was not significantly different from the background level. The step at 9,500 in Figure 3.6 (a) marks the edge of the carton, and was not large enough to reliably trigger from.

Instead of triggering from the DC step, as was standard, it was necessary to trigger off of a transient spike caused by the edge of the carton. The spike comes as the specular reflection from the corner of the cut passes over the detector. One detector sees this spike on the front edge of the box (Figure 3.5), while the other sees it on the back edge.




Figure 3.5 The detector on the left will see a spike in intensity as it catches the specular reflection from the corner of the package. The detector on the right will see a similar spike on the other end of the product. Reproduced with permission from the IsaControl User's Manual.

Since the spike is used as a trigger to begin the acquisition, it is not itself acquired and is not shown in any of the figures. Two similar spikes are visible in the DC intensity in Figure 3.6 (a), at about 2,000 and 7,500, caused by the same specular spike from the creases in the box.

The AC-coupled signal is generated with a 2nd order Bessel band-pass filter with upper and lower cut-offs at 300 μm and 2000 μm , respectively. This smooth black surface produces a fairly weak LSA signal, which manifests in two ways. First, the density of bits is noticeably lower than for Board 1 (Figure 3.3 (d)). Second, the standard deviation of the LSA signal is only about 10 points on a 16-bit scale, which is only 2-3 times the typical noise level.




Figure 3.6 Board 2, full DC in time (a), cropped DC in space (b), cropped AC in space (c), binary signature in space (d).

3.3.4 Board 3

Even the subtle stripes printed on this carton, visible on the tab pictured in Figure 3.2, come through as much stronger changes in intensity than the LSA. Those stripes cause the oscillations from 7,000 to 9,000 in Figure 3.7 (a). Like with the text on Board 1 it is simplest to choose a section of the scan which is free of that pattern.

There is also a region of the carton where the background colour slowly changes (from 2,500 to 5,000 in Figure 3.7 (a)) which could not be avoided without making the total scan length too short. This slow change in DC level is largely, but not completely, removed by the 2nd order band-pass filter with cut-offs at 115 μm and 2000 μm . The component which survives the filter can be seen most clearly in the binary signature in Figure 3.7 (d). The bit density is lower in the first half of the scan, where the large trends push the mean above or below zero, so neighbouring bits correlate more strongly.




Figure 3.7 Board 3, full DC in time (a), cropped DC in space (b), cropped AC in space (c), binary signature in space (d).

3.3.5 Board 4

Board 4 has underlying corrugation, which gave the LSA signal a strong periodic component. Periodic signals can produce matches between unrelated items when the two signatures are in phase, and need to be removed. Because this periodic component was quite broad, the best results were found with a narrow 200 μm – 350 μm pass band. The periodicity which is visible in the AC intensity in Figure 3.8 (c) is a result of this narrow bandwidth, rather than the original periodic component.

This periodicity also has the effect of making the bit density appear higher than it actually is. All of the signatures in this section have a 23 μm spatial resolution, thus all have the same number and density of actual bits. The important measure of statistical independence, the effective number of bits (N_{eff} , see equation (1.21), Section 1.5.4), takes into account the correlation between neighbouring bits and counts only the independent bits. This is roughly the property observed in the each of the (d) plots, because the eye picks out the density of transitions, which also decreases as correlation lengths increase. However, Figure 3.8 (d) is misleading because the periodicity makes neighbouring sections anti-correlate so that transitions are much denser than independent bits. Analysis of the comparisons in the next section will calculate a value for N_{eff} , which will show that this surface has roughly the same information density as Board 2.




Figure 3.8 Board 4, full DC in time (a), cropped DC in space (b), cropped AC in space (c), binary signature in space (d).

3.4 Accurate Positioning in the Operational Environment

As a retrofit of an existing production line, this trial required the LSA system to work within the confines of the mechanical system already in place. In the first instance, this limited the area of each product which could be scanned. The ideal region for LSA would be a light background colour, free of any fine printing (text or artwork), and long in the direction of travel (to maximise the scan's usable length). Setting up a new LSA system involves visually inspecting the target item and selecting the most appropriate area. Such spaces are rare on a lot of modern packaging and, insofar as they exist at all, are often along the edges of the item. In several cases, the machine itself was only marginally wider than the unfolded boxes, and the LSA head could not fit close enough to the edge of the machine to take advantage of these spaces. This particular machine also had a series of rollers and a central support, all of which further restricted the possible positions in which the LSA head could be mounted.

More importantly, the existing mechanics affected how the boxes were registered. Unfolded boxes were introduced into the machine as a stack, and feeder belt passed them individually onto the main conveyor. The chute which held the stack of boxes was the only place on the machine which set the horizontal (y) position of the cartons, and it was adjusted manually each time a different kind of box was loaded into the machine. Critically, the operators typically adjusted *both* sides of the chute to fit the new box. This meant there was no fixed reference from which the edge of the box could be reliably measured (see Figure 2.7).

Instead of relying on a fixed mechanical reference, the camera integrated into the production line head was used to capture a template image. Later rescans with the portable scanner were positioned by adjusting the product until the field scanner's integrated camera showed an image which matched the template. This further compromised the areas which could be used for LSA scans, since a completely blank section of the product would have no features on which alignment could be performed.

The IsaControl software allowed red markers to be overlaid onto the video feed, to aid alignment. The markers in each template in Figure 3.9 highlight which features were used for the visual alignment on each of the four products tested during this trial.




Figure 3.9 The templates used for visual alignment of each board. Several regions have been blocked in green because they show identifying brand markers. The black bars in Board 3 are from an additional calibration pattern attached the box, and are not a part of the product and therefore not available as alignment features.

The template used for Board 4 (lower right) is close to ideal: it contains several long, straight, high contrast features. Moreover, it has features which can be used separately for horizontal and vertical positioning, which makes the manual alignment easier on the user. The other three templates each demonstrate a different potential problem for visual alignment.

Board 3's template in Figure 3.9 (lower left) has many of the same characteristics, but the features are further from the image's centre. The camera used in both pieces of hardware has a pinhole lens, which distorts the image at the edges somewhat. This makes straight features near the edge of the image less suitable for comparison between different video feeds.

The template image capture from Board 1 is high contrast, but is also off centre. More importantly, this template has only a few curved features available, which are harder to use because different offsets cannot be separated. When making small corrections to the position of the box on the field scanner, it is much easier to first correct rotation, then the left/right position, then the vertical position. With

curved markings, like in Figure 3.9 (upper left), the features will only align with the template when all three axes are simultaneously correct.

Board 2's template demonstrates the importance of finding high-contrast features. It has several straight features with hard edges, and so should be relatively easy to use. However, with white illumination and a black and white camera, the features are hard to distinguish from the background material, as can be seen in Figure 3.9 (upper right).

Using manual visual alignment (especially for boards 1 and 2) was a slow and slightly finicky process. While this was acceptable for a first trial with personnel dedicated to the LSA system, it quickly became clear that it would not scale up as a solution, with operators who had less specialised training and a number of responsibilities competing for their time.

3.5 Results from Sampling Outtakes

With the machine running a production schedule, it was a simple matter to build up a large database of enrolment scans. In order to take a verification scan the carton had to be removed from production, so this was only done a handful of times for each carton type. Each carton removed, called an outtake, was then scanned on a field scanner using the alignment template above.

Inverse hamming plots are used as the final evaluation of the system's performance: how well it can discriminate between like and unlike comparisons, determined by the range and separation of the two distributions. As described in Section 2.9, hamming plots are created by plotting two histograms on the same scale. In the results from this trial (Figure 3.10), dark grey shows the resulting BMRs from comparing the outtake scans with the enrolment scan of the same box. The BMRs from comparing the verification scan with the enrolment scan of the rest of the cartons are plotted in light grey. The result for each carton is shown in Figure 3.10.




Figure 3.10 Inverse hamming plots from each of four different cartons. Enrollment scans were taken on a live production line, and verification scans on a portable scanner.

Because there is no overlap between the two peaks, all four results would provide 100% discrimination of the sampled outtakes for an appropriate threshold value. Nevertheless, none of the results are as well separated as results obtained in the laboratory (see for example Figure 2.16, where the like-item distribution sits between 0.95 and 0.99) as the various disadvantages of a live environment take their toll.

The number of independent bits N_{eff} and the decidability index d' are calculated from each hamming plot as per equations (1.20) – (1.22), and the values listed in Table 3.1.

	Board 1	Board 2	Board 3	Board 4
Different Comparisons	10035	2692	8385	8535
Same Comparisons	10	8	10	10
μ	0.538	0.593	0.543	0.533
σ	0.0136	0.0216	0.0108	0.0204
N_{eff}	1340	517	2130	598
d'	15.5	10.2	9.5	9.8

Table 3.1 Statistics calculated from each of the hamming plots in Figure 3.10.

Since d' is the number of standard deviations separating the two peaks, it implicitly assumes Gaussian distributions. While this is reasonably accurate for the unlike-item distributions, the like-item distribution is far too small to determine if it is Gaussian and calculate a reliable standard deviation. Nevertheless, it is valuable to be able to express the quality of the system in a single figure of merit, as d' does. In particular, this value is used to optimise the parameters used in processing the raw scans. For example, the cut-off frequencies in each of the band-pass filters described in Sections 3.3.2 – 3.3.5 were determined empirically by maximising d' . Each dataset was processed using different parameters, and a hamming plot was computed each time. The band pass which produced the best d' was used for the results presented in this chapter.

In order to evaluate the false accept rate (FAR) for each product, a pass/fail threshold value must be chosen. For simplicity here, the threshold is set to midway between the lowest 'same' BMR and the highest 'different' BMR. Finally, the security level (see equation (2.3)) is calculated for the best and worst of the outtakes. These results are reported together in Table 3.2

	Board 1	Board 2	Board 3	Board 4
Threshold BMR	0.661	0.784	0.616	0.682
FAR (one per 10^x)	19.0	12.4	11.1	12.7
Minimum outtake BMR	0.724	0.808	0.640	0.748
Minimum outtake SL	42.0	22.9	18.8	25.4
Maximum outtake BMR	0.766	0.982	0.676	0.862
Maximum outtake SL	62.5	43.1	34.3	57.8

Table 3.2 Security Level (SL) and False Accept Rate (FAR) for each product.

To put these values in real terms, it is helpful to consider an example. Searching a database of ten million (10^7) items, Board 3's FAR of one per 10^{11} comparisons means one false accept per 10,000 ($10^{(11-7)}$) challenges. Similarly, Board 1's FAR of one per 10^{19} comparisons becomes one per 10^{12} challenges.

3.6 Results from a Large Enrolment Data Set

All of the metrics shown in Table 3.2 assume that the unlike-item distribution is Gaussian. The largest of the sampling results in Figure 3.10 involve just over 10,000 comparisons, which is not large enough to expect ideal statistics. Even the largest published hamming plot[6] is based on only 500 items, and contains only 125,000 comparisons. During this field trial, the production line head was aligned for Board 4 and the data above collected by the end of the day. The product was then run for an entire 8-hour shift by local operators. Over the course of the shift, 21,230 boxes were registered, which was the largest database collected up to that point. Although no outtakes were taken during this period, cross-comparing all of the scans in the database allows for an extremely well resolved difference distribution, containing over 225 million comparisons (Figure 3.11, top). The mean and standard deviation of the dataset were calculated and used to overlay a Gaussian fit in red.

In order to evaluate the tails of the distribution, both the BMR data and the fit data from the top of the figure are reproduced as a probability plot in the bottom of Figure 3.11. The cumulative frequency at each BMR was calculated, and expressed as a percentage of the total count. This is plotted on a scale stretched such that a perfect Gaussian cdf forms a straight line.




Figure 3.11 An unlike-item distribution for Board 4 containing 225 million comparisons, and the calculated Gaussian fit, on a linear (top) and a probability (bottom) scale. The inset shows an error around a BMR of 0.5.

A binomial fit was also calculated, but at such large N it is indistinguishable from the Gaussian. The fit makes it apparent that the actual data is skewed towards higher BMR as a result of the optimisation inherent in the comparison routine, the periodicity of Board 4, or both. From the probability scale, the

heavy tail can be evaluated quantitatively. While the Gaussian fit predicts probability of $1:10^5$ at a BMR of ~ 0.62 , the tail of the experimental data does not reach the same probability until a BMR of 0.7

Such a large dataset also uncovered an error in the software which computes comparisons, which underestimates BMRs very close to 0.5 (Figure 3.11, inset). This error was too subtle to be noticed on less-developed distributions.

It must be emphasised that the distribution in Figure 3.11 is the result of comparing scans taken with the production line scanner to scans taken by the same scanner, while the sampling results are comparisons between a production line scanner and a field scanner. As such the distributions are slightly different, and this large set cannot be used to directly estimate more accurate SL and FAR values. However, two of the properties of the system which would cause BMR to skew high (comparison optimisation and spatial periodicity in the product) would be present in both the production line and field scans, so this data is a good indication that the asymmetry in boards 2 and 4 in Figure 3.10 are more than artefacts of small data sets and that the SLs in Table 3.2 are overestimates.

Ideally, a high-resolution distribution like the one above would be computed from comparing production line scans with portable field scans, and like-item distributions would be calculated for statistically-significant sample sizes. Both these analysis require a large number of verification scans, on a large number of outtakes. Despite the obvious benefits, this is unlikely to be achieved until either verification scans can be automated and do not require removing the sample from circulation, or a long-term trial is conducted where even a low sampling rate eventually builds up a large data set. In either eventuality, it remains as future work to be completed.

3.7 Conclusions

First and foremost, the results presented here are a strong endorsement of LSA's applicability even in 'hostile' environments outside of the laboratory. Despite the additional electrical noise, mechanical vibration, less-than-ideal movement and other limitations imposed by a live production environment LSA was able to recognise all of the outtakes presented to the system for verification. Moreover, a statistical analysis of the results indicate that searching a database of ten million items even in the worst case examined would yield a false accept rate of one per 10,000 challenges, and in the best case one false accept per 10^{12} challenges.

Although generally successful, this field trial also raised a number of questions and issues with the system as initially implemented. Much of what was learned in the trial has now been incorporated into advances in the system. The machine shown in Figure 3.1 has been replaced by a newer one, which incorporates a number of changes designed to fix the physical limitations discovered in this trial: a

vacuum belt keeps the boxes flat on the belt; an improved gantry allows the LSA head to be positioned across the full width of the machine; a fixed horizontal reference simplifies positioning during installation. The electrical isolation and shielding of the head and cable have also been improved since this early trial, to help cut down on electrical noise picked up by the sensitive detectors. Finally, the triggering signal now comes from an external sensor, which is less likely to be troubled by a low-reflectivity product.

Other questions do not admit such ready solutions. The positioning of the target item was exposed as a glaring weakness. The quality of the verification match was heavily dependent on how carefully the item was positioned to match the template, and the results presented were achieved through a slow and careful alignment which is unrealistic to expect of the typical operator. But positioning requirements did not just increase the difficulty for the end user; they also impacted the setup of the system. Ideally, the quality of the LSA scan would be the driving force behind selecting the scan area, but in this trial less-than-ideal scan positions had to be chosen to satisfy the need for visual markers to use for alignment. In the immediate future, fixing this will require developing a less arduous positioning system. A longer term goal should be a system which is inherently less sensitive to the exact position of the scans, which would open up a wealth of possible new applications. The first step in developing such a system is to understand exactly how the current system tolerates misplacement.

This trial also made clear the strong material dependence of LSA. Although all of the surfaces tested here exhibit the microscopic roughness which diffusely scatters light, they are clearly not equal from the perspective of LSA. All four of the results presented in Figure 3.10 were taken using the same production machine, the same production line scanner, and the same portable field scanner and yet each of the four boards produced very different statistics. The large variation in performance stems from the properties of the target item.

These two questions, material dependence and positional tolerance, are the focus of the investigations described in the next chapter.

4

TOLERANCE OF THE LSA SIGNAL TO TARGET MISPLACEMENT

*“If thy brother trespass against thee, rebuke him; and if
he repent, forgive him.”*

—Luke 17:3

4 Tolerance of the LSA Signal to Target Misplacement

4.1 Introduction

Using LSA to authenticate an item depends upon repeated measurements returning the same signature, which in turn depends upon scanning the target in the same location with the same orientation. LSA's tolerance to misplacement defines the precision and repeatability with which the mechanics must be designed at every step, from the quality of the linear bearings in the field scanner to the accuracy of the production line mountings. Positional tolerance is the first criterion which must be met in any successful deployment, and so understanding that tolerance will help to define the real-world applicability of the technology.

The experiments in this chapter seek to characterise the behaviour of LSA as a system. For all of the tests conducted here, the LSA components and the sensor geometry are left exactly as described in Chapter 2. Two different motion-control rigs are developed to perform the scan motion under a variety of conditions not achievable with the standard hardware.

This chapter is divided into two main sections. The first seeks to describe fully the behaviour of a representative sample of real-world materials. The second uses more carefully controlled samples to separate the impact of feature size (amplitude and in-plane correlation) from the general nature of the surface.

In Section 4.2 the target is mounted such that it can be rotated in all three dimensions, while the sensor block can be moved along the scan direction or offset in either of the perpendicular directions. The scan direction is fixed in space relative to the orientation of the sensor head, while the target is offset about this path. This effectively mirrors the impact of misplacing the target item on the field scanner.

Eighteen real-world samples are characterised using this apparatus. The LSA signature's robustness, fractional intensity, and characteristic length are all found to depend heavily on the target surface. Low fractional intensity is found to be an indicator of a signal which decorrelates quickly with misplacement.

In Section 4.3 an improved apparatus is developed using a robotic arm to provide more flexible motion control. In addition to the offsets mentioned above, this section tests rotations of the sensor head rather than the target, so that the scan path remains fixed relative to the target item. This mimics the effect of misplacing the sensor head on a production line.

Three different ground glass samples are tested with the robotic rig. The surfaces are generated with the same process and so share the same height distribution⁶, but cover a range of heights from 420 nm to 3.1 μm . The impact of misplacement is found to be independent of the surface height deviation and the in-plane correlation length.

4.2 Tolerance to Misplacement of Real Products

4.2.1 Introduction

Later sections and chapters will control individual aspects of the surface or the LSA sensor and draw more fundamental conclusions. This section is more holistic: it seeks to examine the final behaviour of real-world samples to misplacements or misalignments between the enrolment and verification scans, in order to establish the baseline for later investigations.

A special laboratory rig (the ‘5-axis’ rig) is constructed which can rotate the target about any axis, and move the scan head to any position in space. This allows the linear and angular correlation of LSA signatures to be measured for a variety of targets. Five different linear or rotational offsets were measured for each sample.

Samples were prepared from 15 different LSA trials and deployments, along with three standard samples. Except for the three standards, each sample is a piece of either the product or the packaging that a real customer hoped to authenticate with LSA. Once again, the results of the material trials represented here are governed by non-disclosure agreements. As such, very few of the samples can be explicitly identified. Each sample will be described as fully as possible without revealing any recognisable brands or products.

This survey of a wide variety of samples reveals a strong material dependency on the sensitivity of LSA to misplacement. The misplacement required for complete decorrelation of the LSA signal varies by an order of magnitude across different materials.

4.2.2 5-Axis Motion Control with Precision Stages

At its most basic, the experimental rig used here needs to be able to move either the sensor head or the sample through all six degrees of freedom. This was achieved using two separate stacks of Newport precision motorised stages. The first stack provided translation through the x-y-z space. The second stack provided rotation in all three dimensions about a single point in space.

⁶ Random rough surfaces are typically described by the standard deviation and correlation length of their height profile. To fully describe the surface, the distribution (i.e. Gaussian, exponential, etc.) of the heights must also be known. The source of the roughness generally determines the nature, and thus the distribution of the surface. As a trivial example, a bead-blasted surface will likely have a very different height distribution then one where the roughness is from fractures and fissures, although both could have the same RMS height and correlation length. See p.9 – 17 of Ogilvy[13] for a good discussion of the source and effect of different height distributions.

The M-ILS100CC linear stage has 100 mm of travel with 500 nm resolution and a repeatability of better than 1.5 μm . Three such stages can be mounted perpendicular to each other, as illustrated in Figure 4.1. A point attached to the plate of the third stage can be moved to any position in a 100 mm x 100 mm x 100 mm cube.




Figure 4.1 Three crossed M-ILS100CC linear stages provided high resolution positioning within a 100 mm cubed space.

Three different stages were assembled to provide control of rotation. The URS75CC is a rotation stage, which means it rotates the mounting plate continuously in-plane (about its normal) through 360°. It has a resolution of 0.0005° and a repeatability of 0.002°. The M-GBM80CC and M-BGM50CC are cradle stages, and provide rotation of the mounting plate about an axis parallel to the plate but raised above it. They have a range of at least $\pm 30^\circ$, a resolution of 0.0005° or better and a repeatability of 0.001° or better. The GB80 and GB50 are designed such that, with the two stages stacked, their axes of rotation are coincident. By crossing the two stages and mounting them centred above the rotation plate of the URS75, as shown in Figure 4.2, the three different axes of rotation all cross at a point 15 mm above the plate of the top stage.




Figure 4.2 Two crossed cradle stages stacked on a rotary stage provide accurate rotation in all three dimensions about a single point in space.

A mount was created (see Appendix A for technical drawing) for the plate of the top cradle stage that allowed samples to be mounted facing up, so that the point of rotation was located on the sample's surface. Samples could be either 14 mm x 40 mm x 1 mm or 20 mm x 50 mm with a negligible thickness.

A single-slice of LSA sensor was mounted on the third linear stage of the linear stack, gazing downward. The two stacks were aligned so that the centre of the laser spot in the focal plane coincided with the centre of rotation of the rotary stack. This is shown schematically in Figure 4.3, and pictured in Figure 4.4 and Figure 4.5.




Figure 4.3 Schematic of the 5-axis rig.




Figure 4.4 Photograph of the 5-axis rig.




Figure 4.5 The LSA sensor block mounted above a piece of a sample product box. The sensor block is brought into contact to ensure that it is level with respect to the sample. It will be raised to the focal distance, 5.5 mm, before data is taken.

Each stage is connected to one of two Newport ESP-300 controllers. The controllers can communicate with a PC over an RS232 link, exposing the motion stages to computer control. The linear stage aligned with the x-axis of the sensor head was used to generate the scan motion. Only one stage at a time could be moved, so the other five stages were sent to their positions before the scan and were then stationary during acquisition. A National Instruments 9215 USB analogue-digital converter (16-bit) was used to capture the raw intensity data from the two photodetectors at a rate of 2 kHz. These two pieces of hardware were integrated together into a custom driver for IsaControl. The movement and the data acquisition were synchronised within the driver, and the resulting data passed into the system to be handled as if it came from an authorised LSA scanner.

Offsets of the target item between two scans are described in terms of three linear and three rotational directions. The construction of the 5-axis rig mirrors the operation of an LSA field scanner; the orientations of the sensor block and of the scan path are invariant, and only the sample is rotated. Offsets describe the sample, so a negative offset in the z direction means the sample has been moved further from the sensor head, while a positive offset means it has been brought closer. The linear axes

align to the coordinate axes as defined in Section 2.1, such that offsets along y are called 'transverse,' offsets along z are called 'focus,' and offset along x are called 'scan.' This is shown in Figure 4.6.




Figure 4.6 The three translational offsets. 'Scan' is a movement along the scan direction, transverse is a movement along the length of the laser spot, and focus is a movement along the incident illumination path.

Rotational offsets describe the rotation of the target about the different axes as show in Figure 4.7. A rotation about the z axis is yaw, about the y axis is pitch, and about the x axis is roll. The origin of these rotations is the point on the target's surface which corresponds to the middle of the scan area.




Figure 4.7 The three rotational offsets, describing the rotations as if the target item were a plane, flying in the scan direction, with its wings in the plane of the target surface.

4.2.3 A Survey of Product Samples

Most of the 18 samples used in this experiment are pictured in Figure 4.8 below. They are divided into five categories and described in detail in the Sections 4.2.3.1 – 4.2.3.5.




Figure 4.8 (a) Paper (b) Label 1 (c) Business card (d) Card 1 (e) Card 2 (f) Plastic 1 (g) Plastic 2 (h) Label 3 (i) Cheque (j) Wrap (k) Board 4 (l) Board 5 (m) Board 6 (n) Board 7 (o) Sachet (p) Copper

4.2.3.1 Standard Samples

The first three samples examined here were simply surfaces which are often used internally as a benchmark for the LSA system, and were not associated with any particular customer. First, “Paper” is a standard quality office printer paper: Océ Black Label 80 gsm white A4. Second, “Label 1” is a peel-off label with a slightly waxy finish for spool-fed label printers: Z-Select 2000T matte coated thermal transfer paper with adhesive. Finally, “Business card” is an Ingenia Technology business card, which is on 400 gsm matte Artboard stock, laminated on both sides. These three samples, in that order, are used internally as an ‘easy,’ ‘moderate,’ and ‘hard’ surface to LSA, as determined by signal level and robustness.

In all three cases the sample region was plain white, and clear of any printing. The office paper was measured using a Dektak scanning tip profiler. Over a 4 mm length with 1 μm horizontal resolution, the paper has an RMS height deviation from the mean of 3.1 μm .

4.2.3.2 Plastic Samples

Plastic cards (ISO standards 7810 and 7813) are used primarily as financial instruments (credit, debit, and ATM) or as access control (identity cards, key cards). In both cases it is easy to see why there would be an interest in preventing forgeries. Two different plastic cards are tested here. “Card 1” is a full mock-up of an access card. The sample region is royal blue and free of other printing. “Card 2” is a blank card (plain white) but of lower quality, which makes it more prone to scratches and scuffs. Its sample region is plain white.

“Plastic 1” is a piece of the bottle from a household product, made of opaque orange plastic, injection moulded. “Plastic 2” is a piece of the casing of an electronic product. It is a cast part, and has a smooth beige finish.

4.2.3.3 Sample Paper Products

Four different paper-like products have been tested for material suitability for LSA. “Label 2” (not picture in Figure 4.8) is a high-security destructible label, with a metallic finish and a faint checkerboard pattern (5 mm x 5 mm squares) throughout. The sample region avoids all printing, but the background pattern was unavoidable. “Label 3” is a white peel-off label with some printing, although the sample region is clear save for a single black line. It is noticeably thinner than “Label 1.” “Cheque” is the back of a mock-up of a personal cheque, and is covered in dense printing. The sample region is clear of text, but does include the Guilloche⁷ pattern. “Wrap” is a very thin, very smooth paper product, used to wrap

⁷ Intricate patterns of fine lines, arranged in geometric patterns[66].

high-volume products for transport. The sample region is clear of any printing, but has a uniform dark red background.

“Wrap” was an unusual material. It felt paper-like in its flexibility and weight, but was extremely smooth and had no evident fibre structure. This was confirmed by imaging with an atomic force microscope. The image, after processing using WSxM[67], is shown in Figure 4.9. A 5 μm by 5 μm region of the surface has an RMS height deviation of only 21 nm.




Figure 4.9 An atomic force microscope image of the surface of “Wrap.”

4.2.3.4 Paperboard/Cardboard Samples

The most common materials for LSA suitability testing are paperboards. Any number of industries use paperboard boxes for their products. In those cases where the interest in LSA is as a track and trace technology, it is sufficient and far simpler to scan the packaging rather than the product itself. “Board 5” and “Board 6” are similar in appearance, but from different industries. They are similar in nature to Boards 1 and 3 (see Sections 3.3.2 and 3.3.4); both are smooth, matte paperboard. The areas chosen for the sample are a uniform colour: light blue and white, respectively. “Board 5” is from the pharmaceutical industry, and is slightly thicker than “Board 6,” which is a packaging used in the tobacco industry.

“Board 7” and “Board 4” are more unusual paperboard, both from the household products industry. “Board 7” is similar in nature to Board 2 from Section 3.3, but with a gold-coloured coating smoothed to a near-mirror finish. The sample is clear of all printing, and oriented so that the metallic coating is the

surface being tested. “Board 4” is one of the samples from the field trial, described in Section 3.3. Again the sample is free of any other printing, but focuses on the coloured surface. It is turned so that the scan runs along the lines of corrugation, rather than across them.

4.2.3.5 Miscellaneous Samples

“Sachet” is a thin plastic sachet used to individually wrap blister-packs in the pharmaceutical industry. It is foil-lined on one side and printed white and yellow on the other. A clear white portion of the printed side is prepared as the sample. The sample is from a sachet which was sealed shut and then cut open, and is smoothed out on the 5-axis’ mount. “Board 8” is the same material as “Board 6”, but it is then covered with a layer of clear cellophane wrapping, as is typical when the product is unopened. Finally, the sample “Copper” is a rectangular segment of a brushed copper sheet that had already been cut, fabricated, and stamped to form a metal cap.

4.2.4 Results of Tolerance Measurements

A routine was written into the IsaControl driver that automated the process of scanning at various offsets. Each axis of offset corresponds directly to a motion stage. The desired stage was selected and its minimum and maximum positions were defined, along with a step size. The software would then take a ‘home’ scan (the scan with zero offset on all axes), followed by a scan with each of the desired offsets.

The sliding comparison used to calculate BMR will undo any offset in the scan direction, so that axis is ignored. Focus and transverse offsets are each measured experimentally, focus from -10 mm to 3 mm in 250 μm steps, and transverse from -3 mm to 3 mm in 100 μm steps. Yaw is measured from -10° to 10° in 0.5° steps; pitch and roll are measured from -5° to 5° in 0.25° steps.

The BMR between the home scan and the scan at each offset was calculated, and recorded as a function of offset. Each of the 5 offset axes was measured individually for each of the 18 prepared samples. Each trace of BMR against offset produced a peak at 0 offset.

As an example, the BMR against each offset axis is reproduced in Figure 4.10 for the three samples in the ‘Standards’ set. Individual tolerance curves for the other 15 samples are shown in Appendix B.




Figure 4.10 Focus (a), transverse (b), yaw (c), pitch (d), and roll (e) offsets for the three 'Standards' samples.

Finally, for one sample (card 1) two dimensions of offset were calculated simultaneously. Transverse and yaw were mapped together, as were pitch and roll. The resulting falloffs, in two dimensions, are plotted in Figure 4.11 and Figure 4.12.




Figure 4.11 The impact of simultaneous transverse and yaw offsets on the BMR from "Card 1."




Figure 4.12 The impact of simultaneous pitch and roll offsets on the BMR from "Card 1."

The results from the two figures above are important only insofar as they show that the drop in BMR is roughly independent for each axis. These two-axes plots confirm that there are no complicating dependencies and justify the following treatment of each axis individually.

4.2.5 Analysis of the Dependencies of Sample Tolerances

For each sample, the home scan's AC signal was reprocessed using only a high-pass filter with a cut-off of 2.5 mm. The AC signal's standard deviation was calculated, along with the full width at half maximum (FWHM) of its autocorrelation function. This standard deviation over the mean intensity of the raw (DC) signal gives the fractional intensity dI/I for each sample's LSA signature, while the FWHM gives its correlation length.

For the intensity profile $h(x)$, the autocorrelation function $\mu(x)$ is defined[68] as

$$\mu(x) = \int_{-\infty}^{\infty} h(\tau)h(\tau - x)d\tau \quad (4.1)$$

A BMR of 70% is a typical threshold value, so the width of each tolerance curve is determined by the width of the peak at 0.7. As can be seen in Appendix B, these widths can vary widely; depending on the material, the BMR might fall from its peak value to below the threshold in a single step, or it might never fall below the threshold for the entire range investigated. Pitch and roll, particularly, often return 10° as the width, indicating simply that the BMR never fell below 0.70 over the range examined. The widths for every axis of each sample, along with the correlation length and fractional intensity of the LSA signature, are listed in Table 4.1.

Sample	$dI/\bar{I}$	Correlation Length (μm)	Focus Width (mm)	Transverse Width (mm)	Yaw Width (mm)	Pitch Width (mm)	Roll Width (mm)
Paper	0.0066	360	9.00	4.1	8.5	10.0	10.0
Label 1	0.0042	830	4.00	1.8	6.0	9.75	10.0
Business Card	0.0013	360	2.50	0.5	2.5	2.25	0.25
Card 1	0.0132	330	6.25	2.4	8.5	10.0	10.0
Card 2	0.0011	330	3.75	0.7	4.5	1.50	1.00
Plastic 1	0.0129	88	5.25	3.4	6.5	9.75	10.0
Plastic 2	0.0016	800	3.50	0.4	3.5	2.25	1.50
Label 2	0.0200	640	5.25	2.0	7.0	10.0	10.0
Label 3	0.0063	810	4.25	2.5	7.5	10.0	10.0
Cheque	0.0089	140	6.25	2.9	7.5	10.0	10.0
Wrap	0.0077	320	8.25	3.1	5.5	10.0	10.0
Board 5	0.0021	500	2.75	0.2	3.0	2.25	0.25
Board 6	0.0029	470	3.75	1.3	7.0	7.25	6.50
Board 7	0.0287	700	2.00	0.4	7.0	1.25	0.25
Board 4	0.0120	400	10.0	4.1	11	10.0	10.0
Sachet	0.0030	750	2.50	0.5	2.0	2.25	0.50
Copper	0.0239	400	9.00	4.4	8.5	10.0	10.0
Board 8	0.0071	1200	3.75	1.6	7.5	6.50	4.50

Table 4.1 A summary of results from all 18 material samples.

Several patterns emerge in the data above. First, in every case the LSA signature is less sensitive to changes in focus position than to changes in transverse position. Materials tend to be either quite sensitive or quite insensitive to pitch and roll, with little middle ground. Roll in particular has only three materials for which match quality has not already fallen away by 1° but does not survive for the full 10° range measured.

It is also interesting to note that there is as much variation within groups of samples as there is between groups: for example both the smallest and second-largest transverse tolerances are in the cardboard set, and the two plastic cards are at either end of the scale in terms of sensitivity to misplacement. Paper products are something of an exception, as none of the paper samples show the extreme sensitivity to misplacement present in some of the other samples. Finally, the addition of a transparent plastic wrapping (Board 6 versus Board 8) had little impact on the LSA's sensitivity to misplacement, as would be expected for a sensor system operating in the visible spectrum.

The data collected here also provide a good chance to observe the variation in fractional intensity and correlation length of the LSA signature. It is clear from Table 4.1 that even when the optical geometry is identical, the fractional intensity and autocorrelation width can vary widely. The fractional intensity runs from a minimum value of 0.11% on plain white plastic card to a maximum of 2.9% on a paperboard sample. Similarly, the autocorrelation width of the LSA signal runs from 88 μm (Plastic 1) to 1.2 mm (Board 8), depending on the sample material. The largest autocorrelation is associated with the paperboard covered with transparent cellophane, and is probably an indication that the wrap was not perfectly smooth against the backing, but had some large (on the scale of LSA) wrinkles. However, three other surfaces which were easily prepared flat have autocorrelation widths in excess of 800 μm , so the range covered by different material surfaces is still an order of magnitude.

Fractional intensity and correlation length do not appear to be related to one another. It might have been expected, for example, that both a larger correlation length and a larger fractional intensity would go along with larger surface features. However, Figure 4.13 shows no such relationship between the two.




Figure 4.13 No relationship is evident between the correlation length of the LSA signal and its fractional intensity for a variety of target materials.

Given that, it is worth looking for trends in tolerance widths against each signal parameter individually. The widths for each of the axes in Table 4.1 are plotted below as a function of the signature's correlation length in Figure 4.14 and its fractional intensity in Figure 4.15.




Figure 4.14 Scatter plots of the tolerance to misplacement in each axis, measured as the width of the displacement before the BMR falls below 0.7, against the correlation length of the LSA signature, for 18 different surfaces.

From the figure above, there is no clear relationship with correlation length. However, the near-binary nature of the roll tolerance is readily apparent, with most surfaces sitting either very near zero or at the maximum value of 10°. To a lesser extent this is true of pitch, while focus, yaw, and transverse are all fairly evenly distributed. Finally, a few of the samples exhibit extreme sensitivity in transverse and roll, where the LSA signal decorrelates in only a few hundredths of a degree or a few hundred microns.

For surfaces with an in-plane feature size greater than the beam width, the correlation length of the signal should relate to the correlation length of the surface. Given that the surfaces are generally isotropic, one might expect transverse tolerance to depend on the surface's feature length as indicated by the signal correlation length. This relationship does not appear in the experimental data because of the length of the laser line. Just as with the relationship between surface correlation width and signal correlation length, the relationship between surface correlation width and transverse tolerance would only manifest when the surface features are longer than the laser's 5mm length.




Figure 4.15 Scatter plots of the tolerance to misplacement in each axis against the fractional intensity of the LSA signature, for 18 different surfaces.

There is a more noticeable trend in Figure 4.15, with tolerance as a function of fractional intensity. While it is not a clear enough trend to extrapolate any sort of fit, there is a noticeable clustering in the lower left corner, and a general rise in widths on all five axis as fractional intensity increases. This indicates that there is a least a correlation between low dI/I and greater sensitivity to misplacement.

4.2.6 Conclusions

This section has examined the behaviour of the LSA signature from 18 different real-world samples under changes in the relative position and orientation of the target and the sensor head. Using a specially-constructed laboratory apparatus, it has been shown that the tolerance to misplacement can vary by as much as an order of magnitude, depending on the sample material.

Furthermore, it is shown that the decorrelation widths of the LSA signal along different axis of misplacement are largely independent. The fractional intensity and the characteristic length of the LSA signal are independent of each other, and strongly impacted by the material under investigation; both varied by over an order of magnitude amongst the 18 samples tested here.

Finally, it was concluded that samples with lower fractional intensity also tend to be less tolerant to misplacement between enrolment and verification scans.

4.3 Tolerance to Misplacement of a Model System

4.3.1 Introduction

The data presented in the previous section provided a useful description of surfaces already encountered in LSA. In order to develop a better understanding, and predict the behaviour of surfaces not yet encountered, it is helpful to look at a well characterised model system.

This section will examine the tolerance to misplacement of the LSA from three polished glass diffusers, each with different feature sizes. Importantly, examining these surfaces will show, first, that tolerance to misplacement is *not* determined by the size of the surface feature and, second, that maximum achievable match quality *is*.

In order to take tolerance measurements on a wider variety of targets—including the 2 mm thick glass samples used here—and to measure Yaw and Pitch in isolation, an entirely new experimental apparatus had to be developed. This robotic scanner is explained in detail first, followed by the experimental methods and results.

4.3.2 7-Axis Motion Control with Robotic Arm

4.3.2.1 Motivation for Robotic Arm

The apparatus described in the previous section generated important tolerance measurements for a number of real world samples. But it had two important shortcomings. First, the geometry of the mounting plate meant that only samples which were less than 1 mm thick could be positioned at the coincident point of rotation. Second, each of the Newport motion stages were controlled separately and stacked in such a way that rotational offsets were applied to the target rather than the sensor. As a consequence, yaw and pitch measurements included both a rotation of the incident and observation angle *and* a rotation of the scan path. To do otherwise would require precisely synchronised control of multiple stages, which was not possible.

Modern industrial 6-axis robotic arms are able to reach any point inside their working volume, with any orientation. They can provide extremely well-controlled, highly repeatability movement. As such, a robotic arm makes the ideal framework for a system able to scan an arbitrary path with an arbitrary orientation. A KR 5 sixx R850 from Kuka Robotics with a standard LSA Production Line head formed the basis of this LSA apparatus, pictured in Figure 4.16.



Figure 4.16 The LSA Robot Scanner, an ILS2100 integrated with an industrial pick-and-place robot and custom control software.

The tool was implemented within IsaControl as a piece of custom hardware. In this way, LSA scans could be acquired and processed by the system just as if they had come from one of the official LSA scanners,

but with far greater flexibility in position than could be afforded by either the IFS4000 or the Newport stages.

4.3.2.2 Terminology Common to Robotic Manipulators

6-axis robotic arms traditionally number their axis beginning with the base and moving outward, as shown in Figure 4.17. Axis 1, 4, and 6 are called waists or wrists, and spin through at least $\pm 180^\circ$. Axis 2, 3 and 5 are elbows which can bend through $\pm 170^\circ$.




Figure 4.17 Axis are numbered 1-6, starting from the base and moving out. Figure reproduced with permission from Kuka Robotics.

Although individual motors can be controlled, the robot's operating system will transform Cartesian coordinates and orientation into positions for each axis motor, meaning positional commands can be in a simple x-y-z a-b-c form. The origin of the fundamental coordinate system, called the "world coordinates" is at the base of the robot, oriented as show in Figure 4.18.




Figure 4.18 The home coordinate system, with its origin at the centre of the robot's base. Figure reproduced with permission from Kuka Robotics.

4.3.2.3 Vendor-Provided Control of the Arm

As part of their control package, Kuka provides a three-level control rack (see Figure 4.19). The drive electronics for all six motors are controlled by Kuka software running on a real-time operating system, VxWorks. A second PC running Windows XP provides a familiar user interface, and runs the top-level piece of Kuka's control suite, the Human Machine Interface (HMI).

The HMI serves two purposes. First, it translates user directions, in near-real time, to the VxWorks system. It is through the HMI that the robot can be controlled manually. Second, the HMI is the development environment for the Kuka Robotics Language (KRL). Instructions (including movement commands, logic, and memory access) are written as a program in KRL. The HMI allows the user to select a program to execute, and compiles the instruction set for transfer to VxWorks, where it is executed.

Positioning information in KRL has two additional layers beyond the world coordinates. The robot control software allows the operator to define a different origin, called the base coordinates, and the active location on the robot, called the tool control point.

First, the tool control point is designated. This teaches the robot a position and orientation, relative to the centre of the face plate after axis 6. An instruction to move to a particular point in space causes the robot to move this tool position to the designated location. For the purposes of this experiment, the

centre of the channel two laser, in its focal plane, is designated the tool position. The tool coordinates are oriented so that the laser's short axis (the scan axis from Figure 2.2) is aligned with the world coordinate x-axis, with the laser shining down. With this tool and orientation the scan motion is generated simply by instructing the robot to move from some point, (x, y, z) , to a new point $(x + \delta, y, z)$.

Second, the base coordinates are defined. The new origin can be at any position and any orientation with respect to the world coordinates. Once a base has been defined, movement commands are, by default, with respect to this base. This means that the same set of movements can be performed about different centres by redefining the base coordinates between each iteration.

Rather than designating a particular scan centre and direction, and then performing all the required vector transformations on each of the four positions, it is simpler to redefine the base coordinate system. A new origin is defined, aligned to the target surface and centred at the scan position. Kuka's software system then performs the transformations automatically, so that the scan motion is performed relative to this new origin.

4.3.2.4 Customised Control of the Arm

The HMI is designed primarily for industrial use, where the final application involves the robotic arm repeating a predefined series of motions. It is built on the assumption that the operator will 'teach' the correct path, by manually moving the robot through key positions in slow motion. Additional tools are provided to control important timing elements, and to fine-tune the path. Once the movement controls have been codified in KRL, there is only limited opportunity for further input, and that is generally restricted to flow control (for example, awaiting a digital trigger from a programmable logic controller). There is no mechanism for a more complicated operation, such as redefining the base coordinate system, on the fly. In order to achieve this, the additional architecture shown in Figure 4.19 had to be developed.




Figure 4.19 The control hierarchy for the robotic acquisition system. Connections are as follows: A) National Instruments proprietary drivers; B) Custom driver for IsaControl; C) Custom TCP/IP link; D) Kuka proprietary DLL; E) Internal PCI bus; F) Kuka power/control cable; G) Single 24V digital signal over BNC; H) LSA shielded data cable Z) VxWin, part of Kuka's proprietary software package. Link (Z) is only used during programming and setup.

A small program, Robot Bridge Slave (RBS), was written to run directly on the Windows environment on the Kuka controller. A companion program, Robot Bridge Master (RBM), ran on the IsaControl PC. RBM ran embedded within a custom driver written for IsaControl, and created a link through TCP/IP with RBS. A more detailed description of the communication between RBM, RBS, and the Robot Instructions can be found in Appendix C. Appendix D describes in more detail the commands written in KRL, and available to LSA control through this link.

4.3.2.5 Synchronising Data Acquisition with Motion Control

The optics and acquisition system used in this experiment were borrowed from the ILS2100 (see Section 2.4). The standard LSA sensor block was mounted to the end of the robotic arm, and the power and signal lines from the lasers and photodetectors were run through 3 m of shielded energy chain cable. The same National Instruments analogue acquisition PCI card was used to measure the photodetector intensity.

The command to begin a scan passed through a TCP/IP link between two non-real time operating systems, before being passed into VxWin. This introduced an unpredictable latency, such that simultaneously issuing the command to perform a scan and beginning data acquisition did not necessarily synchronise the movement of the scanner and the acquisition of intensity data. Instead, the custom IsaControl driver used a feature of the National Instruments system and prepared the

acquisition to begin with a rising edge on a digital control line. The Kuka controller also has digital control lines available, which are controlled directly by the KRL program running in VxWin. As a result, the controller's digital output is exactly synchronised with the movement of the robot. This synchronised output was used to begin the waiting data acquisition at exactly the same place in the robot's path each time. Appendix E shows the pin connections necessary to interface the Kuka controller's digital I/O with an ILS2100 trigger signal.

4.3.2.6 Integration with Existing Tools

Final control of the robot scanner was through a custom hardware driver for the IsaControl software. The user interface exposed through IsaControl let the operator specify a position and orientation for each scan. Any arbitrary point in space, within the 850 mm reach of the robotic arm, could be designated the origin. Rotations from the world coordinate could also be specified, such that the orientations of the x-, y- and z-axes were also arbitrary. The robot tool would then acquire data while moving the centre of the channel 2 laser, in the focal plane, through a linear path from (-25,0,0) to (25,0,0) in the new coordinate system. The tool tip was moved at 50 mm/s, and data sampled at 5 kHz. Thus a 50 mm LSA scan could be taken about any point in space with 10 μ m resolution.

The robot's movement system defined the orientation of the tool separately from the orientation of the origin. The scan motion, by default, aligned that x-axis of the tool with the x-axis of the coordinate system as the tool point was moved along the pre-defined path. The final facility added to the IsaControl driver allowed an offset to be introduced between those two orientations. Thus the *path* travelled by the LSA sensor head was decoupled from the *orientation* of the head.

4.3.3 Well-Characterised Ground Glass Samples

The 5-axis data presented earlier covered a wide range of real-world materials, but the manufacturing techniques for cardboards and plastics are not designed to create well-defined surfaces. Strong anisotropies, wide variations in feature size, and poorly-behaved surface height functions (e.g. near-discontinuities from fissures, multi-valued from overhangs) make these real world surfaces difficult to model and characterise.

To gain a better understanding of the impact of microstructure on LSA in general and tolerances in particular required well-characterised surfaces. Three samples of ground glass diffusers, each with different surface statistics, were used throughout this section.

The diffusers were produced by Thorlabs, polished with three different polishing aggregates measured by grit number. The grit describes the density of wires, in lines per inch, in the mesh used to sort the polishing aggregate's particles. Ignoring wire diameter, then, grit is roughly the inverse of the diameter

of the largest particles in the aggregate[69]. Thus, 120-grit glass was polished with particles as large as 210 μm ; 220-grit glass was polished with particles as large as 115 μm ; and 600-grit glass was polished with particles as large as 42 μm .

There is no simple relationship between grit and surface parameters, as feature size and height also depends on the method and duration of polish. In order to characterise the samples, measurements were taken with a Dektak stylus surface profiler. The profiles were 2 mm long with 1 μm linear resolution, and an example from each surface is shown below. The average RMS deviations from the mean height were 3.1 μm , 1.6 μm and 420 nm for the 120-grit, 220-grit and 600-grit glass, respectively.




Figure 4.20 Surface height profiles measured from ground glass diffusers for three different levels of polish: 600-, 220- and 120-grit. RMS height deviations are 3.1 μm , 1.6 μm and 420 nm, respectively.

Plotting the distribution of heights for each surface as a histogram shows that the ground glass surface profiles are approximately Gaussian, as shown by the fits in Figure 4.21. This is expected for a polished surface[70], where the final height of any particular feature is the result of a number of independent mechanical events.




Figure 4.21 All three surfaces have approximately Gaussian distributions of heights.

Surface height only describes the out-of-plane dimension of the rough surface; fully characterising a surface requires, broadly speaking, a description of both the out-of-plane and in-plane length scales.[71] When the surface is isotropic, a single length scale is sufficient to characterise the in-plane feature size. The in-plane characteristic length can be found by calculating the FWHM of the autocorrelation function (see equation (4.1)) of the height profile.

The autocorrelation functions of the three height profiles from Figure 4.20 are shown in Figure 4.22.




Figure 4.22 Autocorrelation functions for each of the surface height profiles shown in Figure 4.20. FWHMs are 420 nm, 1.6 μm and 3.1 μm respectively.

The correlation length for each surface is calculated from the FWHM of the autocorrelations above, which are 45 μm , 36 μm and 16 μm for the 120-, 220- and 600-grit glass, respectively. In all three cases this correlation width is narrower than the laser's focused width.

All three samples were square, 50.8 mm per side, and 2 mm thick. Only one side was left roughly ground as described above; the other was polished to optical quality. For these samples, surface height and surface wavelength cannot be treated separately and are referred to together as simply surface roughness. For simplicity, when surface roughness is the independent variable in a figure it is shown as surface height, but in each case this also corresponds to an increase in surface correlation length.

4.3.4 Additional Two Axes

It was generally most convenient to mount the samples with the target surface facing up, somewhere in the centre of the robot's reach. The ground glass samples discussed here were transparent to varying degrees, so the samples were raised off of the work surface by 100 mm, and supported only at their corners. This minimised the amount of light which reached the detector by reflecting off of the surface

behind the glass. The working origin was defined so that the x-y plane was coincident with the target's surface. All three glass surfaces were isotropic, so the orientation of the plane was arbitrary.

Seven different scanning routines were defined, each describing the axis under investigation: Transverse, Height, Yaw, Yaw tracking, Pitch, Pitch tracking, and Roll. In each case a minimum and maximum value were defined, along with a step size. Every routine began with a scan at the default position and orientation, called the home scan. A series of scans were then acquired beginning with the designated axis offset by the minimum value, with the offset incremented by the desired amount and the scan repeated until the maximum offset was reached. The series of scans all included a scan at 0 offset. The 0 offset scan and the home scan are in nominally the same position and orientation, with one taken before and one taken during the series.

Transverse, Height, Yaw, Pitch, and Roll offsets are exactly as described in Section 4.2.2. Unlike in that section, the path and orientation of the scan head were altered rather than the target itself, but the same relative offset was produced. 'Yaw tracking' and 'Pitch tracking' adjust the orientation of the scan head while retaining the original scan path. The difference between Yaw and Yaw tracking, and between Pitch and Pitch tracking, are illustrated in Figure 4.23. Offsetting Roll is a rotation about the scan direction, so there is never a change to the scan path.




Figure 4.23 Red indicates either the focused spot (a-b) or the incident path (c-d). The vector indicates the scan path. In Yaw (a) and Pitch (c), the scan path is rotated along with the sensor orientation. In Yaw tracking (b) and Pitch tracking (d) only the sensor is rotated, while the scan path remains true to the x-axis.

Yaw and Yaw tracking scans were taken from -5° to $+5^\circ$ offsets from the home position, in 0.25° steps. Roll also went through $\pm 5^\circ$, while Pitch and Pitch tracking went from -4° to 4° , all in 0.25° steps as well. Transverse data were taken through ± 5 mm, and Height data were taken from -2 mm to 8 mm, both in

0.25 mm steps. Each axis was run four times per sample, with the sample rotated one quarter turn between each run.

Only the central channel of data (both detectors) was used. Scans were processed using a 300 μm – 1.1 mm band pass filter, and no mask (0% standard deviation filter as described in Section 2.7.3) was applied. No encoder signal (Section 2.7.1) was available, but the robot’s motion was smooth enough and precise enough that encoding was not necessary. The data were cropped down to the central 15 mm, so that the origin (and thus the centre of rotation for Yaw) was exactly half way through the scan data.

At each scan position, the BMR was calculated between the home scan and the local scan and recorded as a function of the offset from the home position. Comparisons were calculated using a single block, which was allowed to search over 2 mm (see Section 2.8).

Each surface produced a different background DC level of light. In order to have accurate measurements of any intensity-dependent components of the noise, motionless scans were taken for each sample. With the scan head positioned at the correct focal distance, intensity data was acquired at the usual rate, but without moving the sensor head along the surface. By eliminating the variations in intensity due to changes in topography, these motionless scans isolated the laser, detector, and other noise sources from the desired signal. Only electrical noise introduced when the motors were engaged could not be measured in this way.

4.3.5 Results of Tolerance Measurements II

First, the data collected here provides an opportunity to examine the signal parameters in terms of RMS surface height. The mean DC level $\bar{I}$ and AC standard deviation dI were calculated for each home scan, and averaged to yield a $dI/\bar{I}$ value for each surface. While mean levels fell as expected (smoother surfaces scatter less light diffusely), the fine variations in the scattered light decreased even more, resulting in a steady decline in $dI/\bar{I}$. Next, the home scan for each axis was compared to the 0 offset scan for each axis, and these results were averaged together to give a maximum BMR value for each surface. Finally, the signal-to-noise ratio was calculated, by dividing the $dI/\bar{I}$ of each surface by the $dI/\bar{I}$ of that surface’s motionless scan. Noise levels were consistently 0.15% of $\bar{I}$, so S/N fell with $dI/\bar{I}$. This also suggests that the dominant source of noise was intensity dependent, like laser noise. Detector dark noise and other similar absolute sources are not significant for this system at these intensity levels. Maximum BMR levels tracked very closely to these two values, providing strong evidence that the quality of match achievable is governed in this case by the noise levels of the system; not a surprising result. All four values are plotted in Figure 4.24. Error bars are the standard deviations of the values.




Figure 4.24 Signal-to-noise, fractional intensity, and maximum achievable BMR all increase monotonically with increasing surface roughness.

The complete results of each run are plotted as the BMR between the home and local scan, against the offset of the local scan. Each axis is plotted separately in Figure 4.25, with the results for all three surfaces plotted together. Error bars indicate minimum and maximum values over the four runs.




Figure 4.25 Tolerance measurements for each axis taking using on three different ground glass targets.

As would be expected, most axes are symmetric about 0 offset. Height is not, showing something of a plateau between 0 mm and 2 mm of offset, with a similar falloff in both directions outside of this range.

In fact, the falloff with Height is approximately symmetric about +0.75 mm. This indicates that the actual focal plane of the lasers in an LSA head is slightly longer than the nominal focal plane, coming in to focus about 3.75 mm from the sensor head rather than 3 mm. This is a consequence of the assembly process. The nominal focal plane is defined by the focal length of the cylindrical lens (see Section 2.3), but that distance assumes a perfectly collimated beam. Since consistency in assembly is more important than perfect collimation, the collimating lens is always positioned against a hard stop rather than being individual adjusted, leaving the outgoing beam slightly diverging.

A focal distance of 3.75 mm would also explain the slight dip and rise seen between 1 mm and 2 mm: if the home scan is taken with a slightly out-of-focus—and therefore slightly larger—laser spot, then the laser spot will shrink and then grow as the sensor head moves back and the target surface passes through the beam waist. The spot will come closest to matching its original (home) size when it is out of focus by approximately the same amount on the other side of the waist.

The results from Section 4.2 clearly indicated a material dependency in the tolerance to misplacement. In light of that, the most surprising aspect of these results is that the rate of falloff appears to be independent of feature size. This can be quantified by calculating the slopes of each of the tolerance curves.

To get the most accurate values, the rising linear region of each curve was isolated. For Transverse, this was the region from -2.5 mm to 0 mm; Height was selected from -2 mm to 0 mm; Yaw and Yaw tracking from -5° to 0°; Pitch and Pitch tracking from -3° to 0°; and Roll from -5° to 0° offset. For each of these regions, a line of best fit was calculated, and the gradient of the line is report in Figure 4.26.

Over a 15 mm scan rotated about its centre, the errors in tracking from 5° of Yaw and 4° of Pitch are only 660 µm of transverse and 520 µm of height offset, respectively. Neither produces a significant amount of loss compared to what results from the rotation itself, so with and without tracking behaved almost identically. To make the graph clearer, only the result with correct tracking are reported.




Figure 4.26 The gradient was extracted from a linear region of each tolerance curve and plotted as a function of RMS surface height. This rate of decline in BMR gives a measure of the sensitivity of LSA to misplacement.

The graph above represents one of the most important findings of this study. Only Yaw appears to show any dependence on the size of the surface microstructure, and even that is within the error. In the other cases, this confirms the qualitative observation of Figure 4.25—that different surface feature sizes do not impact that rate of decay of the LSA signal. The only strong impact of surface feature size is to decrease the signal strength, and therefore the maximum achievable match.

4.3.6 Conclusions

The experiments conducted in this section are instructional because of the simple nature of the samples, compared with those of the previous section. First, polished glass surfaces had reasonably well-controlled surface profiles. Second, they were clear of any larger features or printing which would influence the behaviour of the LSA signal. With only the surface feature size being varied, two important conclusions could be drawn.

First, the parameters of a surface's height profile—its RMS deviation in height and its correlation length—do not have a significant impact on the sensitivity to misplacement of the surface's LSA signature. BMR falls by roughly 15% for every mm moved in the transverse direction, 10% for every mm in height, 5% for every degree of in-plane rotation, and 2-3% for every degree of out-of-plane rotation.

Second, this section has shown that not only does the detected light level drop for smoother surfaces, but the fractional variation in that scattered light falls as well. This means smoother surfaces produce a weaker LSA signature, which is more susceptible to noise and exhibits a corresponding poor match quality.

4.4 Sources of LSA's Sensitivity to Positioning

The results from the Newport 5-axis rig on product surfaces suggest a strong material dependence, while the results from the robotic 7-axis on ground glass samples do not. For even the smoothest surface examined in Section 4.3, the LSA signature survives many degrees of rotation. Given the different types of surfaces under investigation in the two sections, that suggests that the *form* of the surface height is more important than the size of the surface in determining how LSA behaves with misplacement.

The 'form' of the surface is used here to encompass both the distribution of the surface heights and the less-well-behaved aspects of the surface which elude simple statistical descriptions. Microscopic images of paper show that it is far from a smooth Gaussian surface. The nature of the layered fibres in paper mean that features can protrude in a way that makes the surface height profile multi-valued. Brittle materials like hardened plastic are likely to have cracks and fissures, while metallic surfaces can have grooves and tooling marks, both of which make near-discontinuities in the surface height profile. In contrast the polished glass samples were reasonably well controlled surfaces with a simple Gaussian height distribution. Since the study on these better-characterised surfaces has shown that surface height and correlation length do not impact LSA tolerances to misplacement, it can be concluded that the wide variation seen in Section 4.2 comes not from the different surface heights, but from the distributions which govern their nature. That is, it would seem that a Gaussian surface and an exponential surface with the same RMS height and correlation length would differ more from each other than two Gaussian surfaces with different RMS values.

Most of the tolerance data collected from real-world samples runs the range from very slow to very fast falloffs, and admits the explanation above. However, there are a few materials for which the signal's decay stands out as differing not just by a matter of degrees, but as of an entirely different nature. These supersensitive cases lose most of their match after a single step of offset; that is, their central peak is exactly one point wide. Section 4.3 shows no evidence at all of the supersensitivity observed in Section 4.2. As an example of this, Figure 4.27 contrasts the roll data from the plastic samples (from Appendix B) with the roll data from the ground glass (from Figure 4.25).




Figure 4.27 (LEFT) Tolerance data from the 5-axis rig on the four plastic samples. The falloff of Plastic 2 (pink) and Card 2 (red) are of a completely different order than the falloff of Plastic 1 (black) and Card 1 (blue). (RIGHT) Tolerance data from the 7-axis rig on three ground glass samples. None of grit 120 (blue) 220 (red) and 660 (black) falloffs show a sharp initial drop.

The supersensitivity on the left can be interpreted differently, however, in light of the 7-axis results. First, it is important to note that the central spike does not fall all the way to complete decorrelation (BMR of 0.5). In fact, if the central point is removed from the data for Plastic 2 and Card 2, the two plots would look quite similar: reasonably slow, similar falloffs each starting from a different maximum BMR.

This interpretation makes sense in light of an important difference between the two systems. The 5-axis rig is driven by Newport stages with a repeatability of <500 nm, while the 7-axis rig is driven by a robotic system with a repeatability of 10 μm – 30 μm . In those cases where the central peak is exactly one point wide, it is plausible that the central peak is so sensitive to misplacement that the robotic motion control misses it entirely, even when returning to nominally the same position. This suggests that there are actually two phenomena at work: one whose rate of decay depends on the nature of the surface and is reasonably slow, and another that decays almost immediately. In this sense the 5-axis results are misleading from the perspective of LSA, because they include an exceptionally good match that depends upon a signal not accessible to most LSA systems. Investigations in the next chapter will propose laser speckle as the source of this supersensitive peak.

Putting aside any considerations of supersensitivity, it is instructional to try to construct a purely phenomenological description of the quality of LSA match. The final BMR will be the simple combination of two factors: the reduction $\emptyset$ in BMR due to offsets in each axis, and the maximum BMR M .

$\emptyset$ takes the simple form:

$$\emptyset = c_{tran} \cdot x_{tran} + c_{focus} \cdot x_{focus} + c_{yaw} \cdot \theta_{yaw} + c_{pitch} \cdot \theta_{pitch} + c_{roll} \cdot \theta_{roll} \quad (4.2)$$

where x 's and θ 's are the misplacements between enrolment and verification scans in the various axes, and c 's are constants determined by the distribution of the surface. For Gaussian surfaces, the constants are the average values for each axes from Figure 4.26, so:

$$c_{tran} = 0.15 \text{ BMR/mm} \quad (4.3)$$

$$c_{focus} = 0.10 \text{ BMR/mm} \quad (4.4)$$

$$c_{yaw} = 0.052 \text{ BMR/deg.} \quad (4.5)$$

$$c_{pitch} = 0.027 \text{ BMR/deg.} \quad (4.6)$$

$$c_{roll} = 0.017 \text{ BMR/deg.} \quad (4.7)$$

The maximum BMR depends on the noise level N of the system and the fractional intensity of the LSA signal. The fractional intensity is itself determined by the roughness of the surface (RMS height σ_h and correlation length l_c). While not enough data are available to develop a functional form for M , several relationships and limiting behaviours can be inferred. From Figure 4.24, M is monotonically increasing with σ_h and l_c and monotonically decreasing with N . Furthermore, a perfectly noiseless system would yield a maximum BMR of 1.0, and complete noise would have a maximum BMR of 0.5, so

$$\lim_{N \rightarrow \infty} M(\sigma_h, l_c, N) = 0.5 \quad (4.8)$$

$$\lim_{N \rightarrow 0} M(\sigma_h, l_c, N) = 1.0 \quad (4.9)$$

The final BMR, then, can be expressed as

$$BMR = M(\sigma_h, l_c, N) - \emptyset(x_{tran}, x_{focus}, \theta_{yaw}, \theta_{pitch}, \theta_{roll}) \quad (4.10)$$

where $\emptyset \leq M - 0.5$, and $\emptyset$ has been explicitly defined for Gaussian surfaces.

4.5 Conclusions

The first two sections in this chapter have examined the behaviour of LSA from both real-world and well controlled surfaces. The decorrelation of the LSA signal with changes in orientation and position of the sensor or the target surface was examined for rotations about and translations in all three dimensions.

Section 4.2 found that the signal parameters (fractional intensity and characteristic length) and the width over which the signal decorrelated with each axis of offset were all strongly dependant on the surface being scanned. Over the 18 surface examined, the maximum offsets over which the LSA signature survived varied by at least a factor of 5, and as much as an order of magnitude. Signal parameters varied by over an order of magnitude.

Section 4.3 used a flexible motion-control system and a set of ground glass diffusers to isolate the impact of the surface's amplitude and in-plane feature size from the surface type. It found that the *rate* of decorrelation is unaffected by the surface's feature size. The tolerance width, as described in Section 4.2, will change based on feature size only because the maximum match quality will be different.

Changes in the rate of decay with respect to misplacement of the LSA signal between different materials are attributed to the nature (i.e. Gaussian distribution, exponential distribution, or even more complex multi-valued functions) of the surface. A step-like initial decay seen on some materials is considered a likely result of a secondary optical phenomenon. Finally, a simple phenomenological model reflecting the various dependencies of match quality was proposed.

5

EXAMINATIONS OF THE OPTICAL PHENOMENA UNDERPINNING LSA

*“It is the glory of God to conceal a thing: but the
honour of kings is to search out a matter”*

—Proverbs 25:2

5 Examinations of the Optical Phenomena Underpinning LSA

5.1 Introduction

The fundamental studies undertaken in this section are all designed to understand what optical phenomena underlie LSA and account for the tolerances measured in the last chapter. An important initial question is whether the signal is or is not accounted for by spatially-averaged speckle. This chapter breaks down into two halves. The first examines justifications both for and against speckle as the principle component of the signal. The second examines the significance of coherence and phase to the measured signal.

The original Nature paper on LSA[6] attributes the measured intensity fluctuations to laser speckle. This is true insofar as speckle can be broadly defined to encompass all of the random fluctuations that result from coherent illumination striking a random rough surface. However, the term speckle is generally confined to the short-length, high-contrast pattern visible to the naked eye, which comes from the constructive and destructive interference resulting from different path lengths from source to scattering surface to observation point.

Many proposed applications of speckle explicitly rely on its sensitivity to the small changes to the surface structure[72] or translation[73], whereas the Nature paper described LSA's resilience in spite of crumpling, scribbling, light abrading, baking and wetting. On the face of it, then, a notoriously sensitive phenomenon like this classically-defined speckle is a poor fit for something as robust against small changes to the surface as LSA. The first two sections examine the accuracy of that position.

Section 5.2 conducts experimental measurements of the angular and translational decorrelation of free-space speckle images using the observation and illumination geometry of LSA. It finds that LSA shares more with a setup[42] known to suppress some of the extreme sensitivity to angular changes than it does with those setups that exploit it. With normal incidence, 45° detection, and a circular spot speckle images are found to survive over several degrees of offset, while other surfaces and illumination conditions reduce this tolerance.

Section 5.3 introduces a first-principles numerical simulation of diffraction, which broadly confirms experimental findings: phase interference resulting from different path lengths are sufficient to explain the sensitivity of the classical speckle image in the LSA geometry. Under ideal circumstances, a speckle pattern measured by an LSA detector could be expected to survive several degrees of rotation and roughly half the spot size in translation.

Section 5.4 examines another property of the LSA signal, the fractional intensity, and compares this to the fractional intensity that could be expected from spatially averaged speckle. These figures differ by an order of magnitude on office paper, showing that speckle is a particularly poor fit for the source of the LSA signal in this regard.

In the last two sections of this chapter, the question is approached from the other direction. Rather than examining the behaviour of speckle and comparing it to the behaviour of LSA, experiments are designed to completely remove speckle as a component of the signal.

Section 5.5 examines the LSA-like signals which result when the laser source is replaced with either a partially or totally incoherent source of illumination. It finds that fine variations in the diffusely scattered intensity persist even in conditions where there are no interference effects at all.

The experiment described by Section 5.6 uses careful positioning and three widely-spaced laser sources to quantify the fraction of the signal which depends on the speckle pattern in particular and the phase of the illumination source in general.

Taken together these experiments establish the second major finding of this study, that under most conditions speckle does not play any part in the measured LSA signal.

5.2 Decorrelation of Laser Speckle with Target Misplacement

5.2.1 Introduction

The experiments in this section all seek to quantify the angular or linear decorrelation of speckle. Studies of this type have been performed before for different geometries [34, 74, 75]. The goal in this section is not to supersede any of these studies, but to perform them for the geometry and illumination conditions of interest to this larger work. Large angles of observation, like the 45° detector position of LSA, invalidate the paraxial approximation Beckmann-Kirchoff theory requires, and have largely been ignored in existing experimental work. Similarly, a highly asymmetric illumination spot has never been of particular interest before LSA. The data in this section will fill those gaps.

The important measure here is how quickly the speckle pattern changes as the geometry is altered. Data were collected for in-plane translational offsets and for rotations of both incidence and observation angles together (the surface angle). Data for angular offsets were taken for a 2mm circular illumination spot and for a 2mm focused line, on both a 120-grit ground glass diffuser and office paper. Second, data for translational offsets were taken on one axis for a 2 mm circular spot on the ground glass diffusers and on both axes for a focused line on paper.

5.2.2 Foundations

The experiments in this section will be examining the decorrelation of speckle images. The normalised correlation coefficient (NCC) will be used as the metric to measure the degree of correlation. The NCC μ between two zero-mean two-dimensional functions (e.g. images), f and t , is defined as:

$$\mu(u, v) = \frac{\sum_{x,y} f(x, y) \cdot t(x - u, y - v)}{\sqrt{\sum_{x,y} f(x, y)^2 + \sum_{x,y} t(x - u, y - v)^2}} \quad (5.1)$$

and is calculated here using the fast algorithm described by Lewis[76].

It will also be useful in moving forward to have expressions like (1.17) and (1.18) which estimate the amount of movement it takes for a speckle image to decohere, but which do so for this particular geometry. Goodman[18] provides a framework to make these new calculations.

The method begins with the generalised geometry shown in Figure 5.1. The x-y plane (not shown) typically describes the observation plane, so the plane which defines the mean of the rough surface will be the α - β plane, with z pointing along the average surface normal. $\vec{k}_i$ and $\vec{k}_o$ represent the average incident and observed wavevectors, respectively, with θ_i and θ_o measuring their angle from the surface normal. β is the angle between the z -axis and the α - β plane. Grey lines indicate the in-plane projection of the vectors.




Figure 5.1 General scattering geometry for reflection from a rough surface, with arbitrary incident and outgoing wave vectors. Grey lines indicate the in-plane projection of the vectors.

The usual convention is adopted, where $|\vec{k}| = k = 2\pi/\lambda$, and the scattering vector is defined:

$$\vec{q} = \vec{k}_o - \vec{k}_i \quad (5.2)$$

The scattering vector contains all of the information (observation angle, incidence angle, wavelength) necessary to describe a particular scattering event. This is also the term of interest in the Beckmann-Kirchoff integral, eq. (1.4).

For clarity, the scattering vector will frequently be decomposed into the transverse (in the α - β plane) component, $\vec{q}_t$, and the normal component, $\vec{q}_z$, such that

$$q_z = k(\cos \theta_o + \cos \theta_i) \quad (5.3)$$

and

$$q_t = k(\sin \theta_o - \sin \theta_i) \quad (5.4)$$

The quantity of interest is how the speckle intensity changes as the conditions of the experiment change. This is expressed as the amplitude correlation function between two scattering systems, $\mu_A(\vec{q}_1, \vec{q}_2)$. Goodman[18] shows that the intensity correlation depends only on the difference between the scattering vectors

$$\Delta \vec{q} = \vec{q}_1 - \vec{q}_2 \quad (5.5)$$

through

$$|\mu_A(\vec{q}_1, \vec{q}_2)|^2 = |M_h(\Delta q_z)|^2 |\Psi(\Delta \vec{q}_t)|^2 \quad (5.6)$$

M_h describes the decorrelation that results because of any change to the ratio of the surface-height to the wavelength, including changes from the foreshortening of the surface height due to non-normal incidence or observation. Ψ describes both the translation of the speckle pattern with changing angles and the expansion or contraction of the pattern caused by a change in the wavelength[18].

While Goodman provides more general forms for both M_h and Ψ , this study will use only his values after two simplifying approximations. For a uniform circular spot of diameter D , Ψ reduces to

$$\Psi(\Delta \vec{q}_t) = 2 \frac{J_1\left(\frac{D \Delta q_t}{2}\right)}{\frac{D \Delta q_t}{2}} \quad (5.7)$$

where J_1 is a Bessel function of the first kind. M_h can be similarly reduced, by assuming a Gaussian surface height distribution with standard deviation σ_h , to:

$$M_h(\Delta q_z) = e^{(-\sigma_h^2 |\Delta q_z|^2)} \quad (5.8)$$

All that remains, then, is to calculate $\Delta \vec{q}$ for the particular cases of interest, which will be done in Section 5.2.4.

5.2.3 Apparatus & Methodology for Speckle Decorrelation Measurements

The illumination source for these experiments was a 633 nm Helium Neon laser (Melles Griot 25-LHP-151-230). For reasons of safety, a neutral density filter (optical density 0.5) was permanently attached, reducing the intensity from 15 mW down to just below 5 mW.

The CCD was a closed circuit TV camera, model VCB-3380P by Sanyo. The camera had a removable lens, which was left off for all of the experiments in this section, leaving the CCD element itself exposed to direct illumination. All measures of speckle patterns taken here are free-space, or objective⁸. The imaging device is a 1/3 inch CCD, with 752 x 582 pixels. Pixels are approximately 6.5 μm x 6.2 μm . The analogue video signal was digitised and captured by a Hauppauge WinTV-HVR video decoding card.

The camera had an Automatic Gain Control, which adjusted the gain setting in real time, to keep the image at a roughly constant mean intensity. This setting was left on, to remove any large-scale variations in the scattered intensity, and focus findings on the speckle.

Two kinds of Newport motorised stages, the M-ILS100CC linear stage and the URS75CC rotary stage (described in Section 4.2.2) were used for these experiments. Control of the stages was again computerised through their controller's serial communication driver.

Acquisition and processing for these experiments was performed by a new piece of software. The control software integrated the serial communication which controlled the Newport stages, access to Windows' DirectShow library for capturing still images from a generic video source, and an implementation of the two-dimensional cross correlation algorithm, (5.1). Two separate stages could be controlled. For each stage, a minimum and maximum position was defined, along with a home position and a step size. The control software moved all stages to their home position and captured a still image there. All stages were then moved to their minimum positions. From there, the stages were moved step by step to their maximum positions. With two stages, the first stage stepped through every position for each step of the second stage. At each position the software delayed for several seconds to allow the position and image intensity to stabilise, and then captured a new image. The local image was

⁸ Speckle can be detected either with or without a lens imaging the target surface onto the detector. When an imaging lens is in place, the measured intensity pattern depends upon the lens, and is referred to as subjective. The intensity pattern projected into free space, as is measured by an exposed CCD element, is called objective.[77]

experiments with angular decorrelation, d' was 59 mm. Translation of a circular spot on glass was taken with a d' of 65 mm, and d' for translation of a focused line on paper was 75 mm.




Figure 5.3 Generic components of the speckle decorrelation tests. The aperture and $f1$ could be changed depending on whether the desired illumination spot was a line or a circle. Mounting of the target and the CCD changed depending upon the dimension under investigation.

Rotations of the incidence and observation angles by the same amount were achieved by rotating the target surface itself, with the source and detector positions fixed. The target itself was mounted on a rotary stage (see Figure 5.4) such that the axis of rotation passed through the centre of the illumination spot, on the surface of the target. The stage was moved from -1.5° to 1.5° in 0.05° steps, sweeping the angle of incidence from -1.5° to 1.5° and the angle of observation from 43.5° to 46.5° .




Figure 5.4 Target angle offset

In-plane translational offsets were tested by mounting the target on a pair of crossed linear stages, as shown in Figure 5.5. For the circular illumination spot, the translational decorrelation was measured in just one dimension (y), while for the illumination line it was measured in both x and y . In both cases the y stage was moved from -1.5 mm to 1.5 mm in $60 \mu\text{m}$ steps. When used, the x stage was run from $-100 \mu\text{m}$ to $100 \mu\text{m}$ in $4 \mu\text{m}$ steps.




Figure 5.5 Translational offset

5.2.4 Predicted Decorrelation from the Scattering Vector

For the angular offset, the change in scattering vector needs to be calculated for a rotation of the incidence and observation angles by an equal amount. While changes to the speckle pattern from a translation of the surface are also measured, the primary driver of that decorrelation is the effective change to the surface profile, $h(x, y)$, which an analysis of the scattering vector will not elucidate.

Firstly, the bounds on changes to the transverse- and z-components of $\Delta \vec{q}$ must be calculated. Equation (5.7) assumes a circularly-symmetric spot, which is not always true here. However, the two dimensions of the illumination line will provide an upper and lower bounds for the expected decorrelation. Ψ is considered to decohere completely after the first zero of the Bessel function $J_1\left(\frac{D\Delta q_t}{2}\right)$. Taking the more conservative case of $D = 100 \mu\text{m}$ first, this occurs at

$$\Delta q_t > 0.077 \mu\text{m}^{-1} \quad (5.9)$$

Both the larger diameter of the line and the diameter of the circular spot are $D = 2 \text{ mm}$. In those cases the speckle will have completely decorrelated due to the Ψ component when

$$\Delta q_t > 3.8 \times 10^{-3} \mu\text{m}^{-1} \quad (5.10)$$

M_h can be said to have effectively decorrelated when it has fallen below $1/e^2$. Both the paper and the grit 120 glass have an RMS height deviation of approximately $3.1 \mu\text{m}$. From (5.8) with this σ_h , decorrelation occurs for either illumination shape at

$$\Delta q_z > 0.45 \mu\text{m}^{-1} \quad (5.11)$$

Next, the specific form of the scattering vector must be determined. The scattering vector for LSA geometry with a rotation of the target is shown in Figure 5.6. Rotating the target is equivalent to increasing both the angle of incidence and the angle of observation by some amount, δ .




Figure 5.6 (a) The original LSA scattering geometry. (b) The LSA scattering vector. (c) The LSA scattering vector after an equal change to the angles of incidence and observation.

In this geometry $\lambda_1 = \lambda_2 = \lambda$, $\theta_{i1} = 0^\circ$, $\theta_{i2} = \delta$, $\theta_{o1} = 45^\circ$, and $\theta_{o2} = 45^\circ + \delta$. Assuming δ is small and considering equations (5.3) - (5.5) gives

$$\Delta q_t = \left| \frac{\pi\sqrt{2}}{\lambda} \left(\frac{1}{2}\delta^2 + [\sqrt{2} - 1]\delta \right) \right| \quad (5.12)$$

and

$$\Delta q_z = \left| \frac{\pi\sqrt{2}}{\lambda} \left(\frac{1 + \sqrt{2}}{2}\delta^2 + \delta \right) \right| \quad (5.13)$$

Combining expressions (5.9) through (5.13) and using $\lambda = 633$ nm illumination shows that the speckle will decorrelate from the transverse component of the scattering vector from a 2 mm spot when $\delta > 0.074^\circ$, and from a 100 μm spot when $\delta > 1.5^\circ$, or in either case from the z-component when $\delta > 3.4^\circ$.

In all cases, it is the decorrelation caused by the transverse component of the scattering vector which would seem to dominate the decorrelation. However, the experimental evaluation of the correlation—the NCC—is calculated with an offset of up to 85% of the image size, or 3 mm. Taking the correlation coefficient at the best offset allows the observation point to effectively track across this distance, which is equivalent to 2.6° of arc. Absent any change in wavelength, Ψ only describes how the speckle pattern translates. Tracking the observation point through space nullifies this decorrelation. Ψ , then, will be ~ 1 until a change in surface angle of 1.3° (the observation angle changes by $+\delta$, while the speckle shifts by $-\delta$, creating a total shift of 2δ). After this point Ψ returns to the values predicted by the equations above, which are all effectively 0. Thus, for the measurement taken here Ψ will not be a steady falloff

but rather a tophat function with unit amplitude, so the measured correlation will reflect only the M_h portion until $\pm 1.3^\circ$.

5.2.5 Decorrelation Results with Translation and Roll

As expected (see equation (1.10)), the size of the speckle grains is inversely proportional to the diameter of the illumination spot, so the free-space speckle pattern is greatly elongated in the direction in which the spot is focused. The free-space speckle pattern recorded for the home position (normal incidence, 45° observation) on ground glass is shown in Figure 5.7 for a 2 mm circular spot and for a 2 mm focused line.




Figure 5.7 The objective speckle pattern from a 2 mm circular (LEFT) and a 2 mm focused line (RIGHT) laser spot normally incident on a ground glass diffuser, observed at 45° from the surface normal. The focused line is long horizontally and focused vertically.

As an example, Figure 5.8 shows a few of the series of images captured for rotation of the ground glass target when illuminated with a circular spot, along with the corresponding NCCs. The peak is visible as a black spot which tracks down through the NCCs, except in (h) and (n) where the offset is outside the range calculated by the NCC routine. The results that follow are presented as a plot of this peak NCC value against the offset of the local image.




Figure 5.8 Images (a) through (g) are the objective speckle patterns from a 2 mm circular spot incident on a ground glass diffuser. The images were taken with the targets surface's normal offset from the optical axis by the following amounts: -1.25° , -1.2° , -0.6° , 0° , 0.6° , 1.2° , 1.25° . Images (h) through (n) are the NCCs between the above speckle pattern and the 'home' speckle pattern, taken at the same position as (d). The NCC images have been inverted, to improve visibility. The sharp transitions between (h) and (i) and between (m) and (n) are from the peak value moving outside of the range of offsets calculated as part of the correlation function.

For rotations of the surface illuminated with a circular spot, Section 5.2.4 predicted a sharp transition at 1.3° , where the offset built into the NCC calculation could no longer compensate for the translation in the speckle and the suddenly reintroduced Ψ component forced the correlation to zero. This is clearly visible in the glass data in Figure 5.9. The theory also predicts that within -1.3° to 1.3° , the correlation is only M_h and should go like $e^{-\delta^4}$. The experimental data has sharp initial falloff from the zero offset position (the central spike in the glass data is exactly one point wide), but then appears to decorrelate at a more gradual rate in keeping with the predicted complete decorrelation at 3.4° . The paper target has roughly the same σ_h as the glass target, and so might be expected to show roughly the same behaviour. However, the expression used to predict the angular correlation, (5.8), assumed a Gaussian height distribution. While this might be an adequate match for a well controlled surface like the ground glass, it is not an accurate description of a complicated surface like the fibre structure in paper. The impact of a poorly-behaved surface—where overhangs and near-discontinuities mean the surface does not change smoothly with angle—are seen in the much stronger dependence of the paper’s speckle on angle of incidence and observation.




Figure 5.9 The correlation between speckle patterns after a rotation of the surface, taken on paper and grit 120 ground glass with a 2 mm circular illumination spot.

When illuminated with a focused line rather than a circular spot, the glass become more sensitive to rotations of incidence and observation angles, shown in Figure 5.10. While the theory outlined above predicts survival of the speckle pattern over larger angles with a smaller spot size, the measurements taken here show that the speckle from the 2 mm x 100 μm spot decorrelate much faster than should either a 2 mm spot or a 100 μm spot.




Figure 5.10 The correlation between speckle patterns after a rotation of the surface, taken on paper and grit 120 ground glass with a 2 mm focused illumination line.

The likely explanation is that the *focusing* of the incident illumination has a larger impact than the change in spot size. With a focused spot (particularly if it is at all out of focus, and therefore still converging or diverging) the assumption of an incident plane wave breaks down quite severely. Small changes in the incidence angle are compounded by changes in angle across the width of the focused spot.

For in-plane translations of the target, the region of the surface which is illuminated changes as the target is moved. The decorrelation is proportional to the amount of the illuminated surface which is different, and this effective change in $h(x, y)$ is the primary cause of the decorrelation observed. While there is an additional change in the effective observation angle as the spot tracks relative to the detector, this effect is small, particularly for the width of the focused spot. In addition, the decorrelation from a change in observation angle should have some dependence on feature size, but as can be seen in Figure 5.11, all three levels of polish behave identically. The angular correlation of speckle images from a 2 mm circular spot has a FWHM of 400 μm , and has fallen by $1/e^2$ after 500 μm , giving it a full width of 1 mm, or half the spot diameter.




Figure 5.11 A measure of the correlation between speckle patterns after a translation of the surface, taken on grit 120, 220, and 600 ground glass with a 2 mm circular illumination spot.

The final results are from the illumination of office paper with a 2 mm focused line. Figure 5.12 shows the peak correlation coefficient for combinations of offsets along both the long (y) and short (x) axis of the illumination line. The tilt of the offset map indicates that the long axis is actually off the y axis by 4.5° (note the x and y axes are scaled differently). The two-dimensional data also confirms that the correlation falls like the overlap between the original and offset illumination spots, without unusually severe or gentle decorrelation along the diagonals.




Figure 5.12 A measure of the correlation between speckle patterns after a translation of the surface along the x and y axis, taken on office paper illuminated with a 2 mm focused line. The spot is focused in the x-dimension and long in the y-dimension.

To calculate correlation widths, it is easiest to examine two line profiles, extracted at $x=0$ and $y=0$, respectively. Translation along the 2 mm axis (Figure 5.13) has much the same behaviour as translation of a 2 mm spot: a full width of $900\text{ }\mu\text{m}$ and a FWHM of $380\text{ }\mu\text{m}$. Along the short axis (Figure 5.14), translation of the line on paper has FWHM of $20\text{ }\mu\text{m}$ and decorrelates completely at $\pm 40\text{ }\mu\text{m}$. Again, both have a full width of roughly half the relevant spot dimension, in this case foreshortened by the small misalignment of the major axis.




Figure 5.13 A measure of the correlation between speckle patterns after a translation of the surface along the y axis, taken on office paper illuminated with a 2 mm focused line.




Figure 5.14 A measure of the correlation between speckle patterns after a translation of the surface along the x axis, taken on office paper illuminated with a 2 mm focused line.

5.2.6 Conclusions

This section has shown that the sensitivity of speckle can depend heavily on the form of the surface and not just its RMS height, which aligns well with the result of Section 4.2. Furthermore, the translational tolerance of speckle is a function of the dimensions of the illumination spot, just as it is with LSA.

Both of these results could be true of a number of scattering mechanisms, and show only that speckle is partially compatible as a source of LSA, which is far from showing that it is a primary component.

Conversely the decorrelation seen here with pitch/roll, especially for well-behaved surfaces like ground glass, is several orders of magnitude short of those found in Section 4.3. This argues strongly *against* speckle being the primary underlying source of LSA.

It has been demonstrated here that LSA does not operate in the regime where speckle is sensitive to μm of translation and minutes of arc. Although there are geometries which exhibit these sensitivities [78, 79], they generally stem from the Ψ component. Considering a region of the observation plane large compared to the grain size, as in LSA, is exactly the method required to eliminate this contribution to the decorrelation of the speckle pattern. The angular correlation lengths of several degrees, governed by M_h , can be achieved at the incident and observation angles used for LSA. However even these least sensitive geometries tested here showed speckle decorrelating more after 1° of offset than LSA does in 5° .

5.3 Prediction of a Huygens Wavelet Simulation

5.3.1 Introduction

Where the previous section demonstrated experimentally that speckle can survive macroscopic displacements in the LSA geometry, this section seeks to understand why that should be true theoretically. While the assumption that anything dependent on fractions of a wavelength could not survive a millimetre of displacement is obviously wrong, it is not immediately apparent how complex a model is required to reproduce this behaviour. As covered in the Section 1.5, there are a large number of complex surface-current models which have already been developed. Those models aim to include all the layers of complexity in scattering in order to accurately reproduce more complex phenomena like backscatter, localisation, and memory-effect[80]. For the purposes of this section, a detailed approach would actually muddy the waters. A first-principles simulation of diffraction from a rough surface will not include multiple-scattering or shadowing, and it will also present an easily understood physical interpretation. The results from this simulation directly demonstrate that the geometry of path length is sufficient explanation for the survival of speckle images beyond wavelength-sized displacements.

5.3.2 Overview of the Model

Figure 5.15 illustrates each of the pieces which must be developed as part of a Huygens' wavelet model for diffraction from a rough surface. A monochromatic plane wave with amplitude envelope $A(x)$ is incident on the random rough surface, $h(x)$. Each point on the surface emits as a Huygens point source, with its phase determined by its height above the surface. An interference pattern, $I(\beta)$, develops in the observation plane, β .




Figure 5.15 The basic components of Huygens' wavelet model of diffraction from a rough surface.

The model developed in this section will be limited to one-dimensional surfaces. A one dimensional simulation is sufficient to provide insight into both translation and roll. Since this simulation is primarily exploratory, the greatly increased computational cost of a two-dimensional simulation will be saved for the full model developed in Chapter 6.

The next section will explain in detail the formulation of the surface, the observation plane, the wave propagation, and the amplitude envelope.

5.3.3 Formulation of the Numerical Simulation

Typically, random rough surfaces are modelled as either Gaussian or exponential height distributions. There are obvious limitations to this description[13], particularly for samples like paper where the surface is a result of complex fibre structure. However, it is a reasonable fit for the ground glass samples, and simplifies the construction of the surface sufficiently to be worth implementing as a first pass.

The spectral method, described algorithmically for one dimension by Thorsos[81], is used here to generate a zero-mean, Gaussian stochastic surface. A spectral density weighting function, $W(K)$, is defined. For a Gaussian roughness spectrum with correlation length l and RMS height deviation, σ_h :

$$W(K) = \frac{l\sigma_h^2}{\sqrt{2\pi}} e^{-K^2 l^2/4} \quad (5.14)$$

The Fourier spectrum of the surface is generated by applying the above weighting function to a randomly generated amplitude, as:

$$F(K_j) = \sqrt{2\pi LW(K_j)} \left(\frac{R(0,1) + iR(0,1)}{\sqrt{2}} \right) \quad (5.15)$$

for $j \geq 0$, where L is the total length of the surface, $K_j = \frac{2\pi j}{L}$, and each appearance of $R(0,1)$ is a random number generated from a zero-mean Gaussian distribution with unit standard deviation. $F(K_j)$ must be real-valued at the maximum and minimum values of j , and at $j = 0$. For all other j , $F(K_j) = F(K_{-j})^*$. The surface height profile, $h(x, y)$, can then be recovered by the inverse Fourier transform of (5.15).

Using this method, one-dimensional surfaces of arbitrary length, correlation length, and RMS height can be created. The impact of the σ_h parameter is illustrated in Figure 5.16, and the impact of l is illustrated in Figure 5.17.




Figure 5.16 Three Gaussian surfaces generated with $L = 2$ mm, $l = 30$ μ m, and three different values of RMS height deviation: a) $\sigma_h = 1$ μ m; b) $\sigma_h = 3$ μ m; c) $\sigma_h = 9$ μ m.




Figure 5.17 Three Gaussian surfaces generated with identical $L = 2 \text{ mm}$, $\sigma_h = 3 \text{ }\mu\text{m}$, and three different values of correlation length: a) $l = 90 \text{ }\mu\text{m}$; b) $l = 30 \text{ }\mu\text{m}$; c) $l = 10 \text{ }\mu\text{m}$.

The method above always generates surfaces with a flat zero mean. In order to test changes in incidence and observation angle, it is necessary to rotate a particular surface about its midpoint, so that the centre line is offset from the x axis by the angle α . This is achieved simply by multiplying the surface height profile, $h(x)$, at each x by the rotation matrix:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos \alpha & -\sin \alpha \\ \sin \alpha & \cos \alpha \end{bmatrix} \begin{bmatrix} x \\ h(x) \end{bmatrix} \quad (5.16)$$

A new rotated height profile, $h'(x)$, is created by linearly interpolating a y value for each of the original x positions from the two nearest (x', y') positions.




Figure 5.18 Rotation of a surface profile about (0,0) by 1, 2, and 3 degrees.

It is helpful to define the observation plane, (β, n') , in terms of the surface coordinate system. There are two transformations required, illustrated in Figure 5.19. First, a rotation by the observation angle, θ , as in (5.16) transforms to the (m, n) system. Second, a translation along the m axis by the observation distance, z , transforms to the (β, n') .




Figure 5.19 A rotation by the observation angle (a), followed by a translation by the observation distance (b), transforms the surface coordinates (h, x) into the observation coordinates (n', β)

Thus for the position (x_0, h_0) in surface coordinates, its position in observation coordinates is:

$$n_0 = x \sin \theta + h \cos \theta - z \quad (5.17)$$

$$\beta_0 = x \cos \theta - h \sin \theta \quad (5.18)$$

And the distance from the position (x_0, h_0) to the position (n_1, β_1) is

$$r = \sqrt{(x_0 \sin \theta + h_0 \cos \theta - z - n_1)^2 + (x_0 \cos \theta - h_0 \sin \theta - \beta_1)^2} \quad (5.19)$$

The final phase of the contribution from each point source is determined by the two path lengths illustrated in Figure 5.20.




Figure 5.20 The two dotted lines indicate the two path lengths which together determine the phase detected at point β from a point source at x .

Ignoring the time dependence, the choice of initial phase is arbitrary, and for ease of calculations it is assumed to be zero at the mean ($h=0$) surface. The initial phase of the Huygens source at x , then, is $e^{-ikh(x)}$, and the phase change travelling to the detection point is e^{ikr} . Assuming a constant uniform reflectivity, and that our observation points are all on the β axis ($n = 0$), then the observed field at β is:

$$P(\beta) = \int_{-\infty}^{+\infty} A(x) e^{-ikh(x)} e^{ikr(\beta)} dx \quad (5.20)$$

where

$$r(\beta) = \sqrt{(x \sin \theta + h(x) \cos \theta - z)^2 + (x \cos \theta - h(x) \sin \theta - \beta)^2} \quad (5.21)$$

and the detected intensity is the squared magnitude of P .

When calculating the equation above, the actual bounds of integration are set by the amplitude envelope going to zero. This simulation assumes a normally incident plane wave with a wavelength of

633 nm and a Gaussian amplitude profile of 2 mm width ($\pm 2\sigma$). The observation plane is set at $z = 65$ mm with an observation angle of 45° . Observed intensity is calculated for a 5 mm length of the observation plane ($\beta = -2.5$ mm to $\beta = 2.5$ mm), with 500 nm resolution. A numerical implementation of (5.20) is shown in Figure 5.21, plotting, from top to bottom, the amplitude envelope, the surface height profile, and the detected intensity profile.




Figure 5.21 The Huygens' wavelet simulation of speckle. From top to bottom are plotted the amplitude envelope $A(x)$, the surface profile $h(x)$, and the intensity profile $I(\beta)$. All horizontal axes are in units of mm, as is $h(x)$. Intensity is in arbitrary units, such that the total integrated intensity of $A(x)$ is unity.

5.3.4 Results of the Calculations

The first measure of the simulation is to compare the calculated intensity profiles to experimental profiles. For the experimental data, a line profile is extracted from the reference image from the surface angle measurements of 120-grit ground glass illuminated with a 2 mm spot, from Section 5.2.5. The

simulated profile is generated using the simulation described above, where the surface is generated with an RMS height deviation of $3.1\text{ }\mu\text{m}$ and a correlation length of $45\text{ }\mu\text{m}$, to match the 120-grit surface.

The experimental data had a minimum intensity value of ~ 30 , on an 8-bit intensity scale. One expects speckle to include complete nulls in intensity, so it is reasonable to assume that there is a DC bias from ambient light. The data plotted in Figure 5.22 has had this background intensity subtracted out, and is normalised to 1.




Figure 5.22 Intensity profiles in the β plane calculated using the Huygens' wavelet simulation (a) and captured from a CCD (b) for a 2 mm spot shone on 120-grit ground glass. Both have been renormalised.

The experimental data has a mean intensity of 24.7 and a standard deviation of 13.4 on an 8-bit intensity scale, giving $C = 0.543$. The simulated data has a mean of 0.0195 and a standard deviation of 0.0106 on an arbitrarily normalised intensity scale, giving $C = 0.544$. The simulated and experimental contrast agree very well.

The other useful measure is the autocorrelation of the signal, which effectively measures the grain size of the speckle pattern. Equation (1.10) predicts a grain size of $25\text{ }\mu\text{m}$ for the geometry under consideration here. The autocorrelation of the two intensity profiles from Figure 5.22 are both plotted in Figure 5.23. The simulated autocorrelation, $28\text{ }\mu\text{m}$, and the experimental autocorrelation, $22\text{ }\mu\text{m}$, agree well with each other and with the predicted grain size, especially considering that precision of the experimental data is limited by the pixel size.




Figure 5.23 The autocorrelation function of the intensity profile for the simulated curve (solid line) and the experimental data (points).

Having thus developed some confidence in the simple Huygens' wavelet model described above, attention is now turned to the behaviour of the speckle pattern under changes to the illuminating geometry.

A series of intensity profiles, like the one above, were generated for the two different axes examined experimentally: surface angle and spot translation. The one-dimensional version of equation (5.1) was used to compute the normalised cross-correlation coefficient comparing a home intensity profile to each other profile. To simulate spot translation, a single 7 mm long surface was generated, and the scattered intensity profile was calculated for a series of amplitude envelopes. For each iteration, $A(x)$ was a 2 mm wide Gaussian, but the spot centre was moved from -2.5 mm to 2.5 mm in 100 μm steps. The home profile was computed with the spot centred on 0 mm. Figure 5.24 illustrates how the simulated intensity profile decorrelates as the spot centre translates across surfaces with statistics matching 120-, 220- and 600-grit glass. For comparison, experimental data for 120-grit glass from Figure 5.11 are included.




Figure 5.24 The decorrelation of simulated speckle intensities as the illumination spot is moved across the target surface, for simulated (solid lines) surfaces with statistics matching three different grit glasses. Experimental data (points) are included for reference.

The level of correlation in the wings, where the signal has completely decorrelated, is noticeably different between simulated and experimental data. This baseline level is set by the fraction of correlation between two unrelated signals, which is much higher for a one-dimensional signal. Put simply, a two-dimensional signal will have many more features in its area than a one-dimensional signal will have in its length, so the fraction of them that match by chance will be lower.

A similar process was used to simulate decorrelation with surface angle: a 2 mm surface was simulated, and the home intensity profile generated at $\alpha = 0^\circ$. Additional intensity profiles were calculated after rotating the surface profile from -2.5° to 2.5° in 0.01° steps, and their correlation with the home profile was evaluated. In order to establish error bounds on the simulation, the rotation simulations were repeated 9 times. The maximum and minimum values are plotted in dotted lines around the average. Again, the experimental data for 120-grit glass from Figure 5.9 are also included for comparison.




Figure 5.25 The decorrelation of simulated (solid line) speckle intensities as the target surface is rotated. Upper and lower bounds (dashed lines) after 9 iterations are included to establish the dependence on surface instance. Experimental data (points) are included for reference.

The step transition in the rotational data occurs at a different offset for the simulated data than it does for the experimental data. This happens simply because the experimental CCD element was 3.6 mm across, while the simulated observation plane ranged from -2.5 mm to 2.5 mm. An 85% correlation range (or 2.1 mm) covers 1.85° of arc at 65 mm, which is exactly where the simulated step occurs.

More significant, however, is the general disagreement in sensitivity between the wavelet model and experimental results seen in both Figure 5.24 and Figure 5.25. In both cases the experimental data exhibits greater sensitivity to changes in illumination geometry, decorrelating more at each offset position. There are two probable causes for this discrepancy. First, the model assumes a number of perfect conditions (for example, perfect alignment of the illumination spot, the target surface, and the axis of rotation; or a perfectly uniform plane wave) which would not be met in an experimental setup. Second, the model is of a one-dimensional system, or, to give it some physicality, a 'corrugated' surface, randomly varying with one dimension, but uniform and large with the other.

5.3.5 Conclusions

Ultimately the model's value is not in exactly predicting all aspects of speckle decorrelation, but in setting the upper bounds for such effects. The results of a Huygens' Wavelet simulation reinforce the conclusion from the previous chapter. Even in the best possible case, without any of the inevitable

degradations or tiny misalignments present in a real system, this tolerance still is not sufficient to explain the survival of the LSA signature itself.

5.4 Fractional Intensity from Spatially Averaged Speckle

5.4.1 Introduction

Tolerance to misplacement is not the only metric available to evaluate speckle against LSA. Examining the fractional signal strength, defined as the standard deviation over the mean of the signature, will show that the variations from speckle grains is not sufficient to account for the variations in intensity registered experimentally. Even fully developed speckle, which is likely an overestimation, would account for only one fifth of LSA's signal strength.

5.4.2 Generating Pseudo-Scans

The apparatus shown in Figure 5.3 and Figure 5.5 was used again for this experiment, but with the camera's Auto Gain Compensation turned off. A series of speckle images were captured for a target of office paper, illuminated with a 1 mm collimated spot. Images were captured at each x position from -10 mm to 10 mm, in 50 μm steps. Rather than the correlation values, in this run the average intensity of the image was calculated. This effectively turns the CCD into an integrating photodetector like the one used in LSA.

This mean intensity plotted against position (Figure 5.26) forms a trace like an LSA scan. The intensity traces had a small background slope, which was removed by calculating a line of best fit and subtracting this from the raw data. The background slope has already been removed in the figure below, and the calculation of standard deviation that follows.




Figure 5.26 An LSA-like scan captured by averaging all the pixels in an image, captured by a CCD for each position of the illumination spot.

For a 1 mm diameter illumination spot with 633 nm illumination and an observation plane 64 mm away, equation (1.10) predicts a grain size of 49 μm diameter. Each original image as captured by the frame grabber is 320 pixels by 240 pixels, covering 4.8 mm x 3.6 mm of objective space, making each pixel approximately 15 μm on a side. For these sizes, a full CCD image could be expected to contain roughly 9,100 grains. Assuming the partially-developed speckle analysed in Section 5.3.4, such that $C = 0.5$, and assuming a $1/\sqrt{N}$ dependency on number of grains, the mean intensity of the series of images should have a $dI/\bar{I}$ of approximately 0.005, or 0.5%. The data in Figure 5.26 has a fractional intensity of 0.0186, nearly four times higher than would be expected from speckle grains alone.

5.4.3 Varying the Integrating Area

Having captured the full image allowed the data to be reprocessed for different effective observation areas. Each image is cropped down to a smaller central region of interest (ROI), and the new average intensity is calculated at each illumination point, generating pseudo-scans like the one above for different integrating areas.

The smallest ROI used is 240 μm by 180 μm , or 5% of the original length and width, which is large enough to include about 60 speckle grains. Intensity profiles and their fractional intensity are also calculated for ROIs which are 10%, 25%, 50%, 75%, and 100% of the original CCD image size. Taking the smallest image as the base unit, each larger ROI can be considered the sum of N units, where $N = 1, 4, 25, 100, 225$, and 400 for the different ROIs mentioned above. As the sum of independent components, the $dI/\bar{I}$ of each ROI should be $1/\sqrt{N}$ of the $dI/\bar{I}$ of the base unit. Figure 5.27 shows that expected fractional intensity as a function of N , which is linear on a log-log plot. The pseudo-scan was repeated three times, and error bars expressed below indicate min/max values. Extrapolating from the base unit to the full image size ($N=400$) gives an expected $dI/\bar{I}$ of 0.006, very close to the value predicted above. This is good confirmation that 60 grains is a sufficiently large number to fully realise the expected statistics for speckle (see Section 1.5.3.1), and that the base unit (5% ROI) is not too small.




Figure 5.27 The bottom trace shows the expected intensity for an image composed of N base units, assuming a $1/\sqrt{N}$ falloff from the base fractional intensity, while the upper trace shows the actual $dI/\bar{I}$ from the intensity profiles generated from regions N times larger than the base unit. Both fractional intensity and N are plotted on a log scale.

It is clear from the figure above that the $dI/\bar{I}$ calculated for each area of integration is larger than would be expected from a summation of base units. This deviation is particular severe for larger images, where it appears to be asymptotically approaching a value close to 0.02. The larger-than-expected signal strength suggests that there is an additional, longer-wavelength component that becomes more significant when the influence of the classical speckle grains is muted through averaging large numbers of them. Whether this longer length scale signal is a higher-order property of the interference pattern[82-84] or something simpler will be investigated further in the next two sections.

5.4.4 Conclusions

The image with an ROI of 100% in the dataset above includes approximately 24,000 speckle grains. It is useful to repeat the analysis performed in Section 5.4.3 for the LSA geometry, to see where that would fall on the Figure 5.27.

As described in Section 2.1, LSA is taken with a 650 nm laser, focused to a 100 μm by 5 mm line. The detector lens has a 4.5 mm diameter, and is positioned 11 mm back from the target surface. Running these numbers through equation (1.10) finds the expected grains will be 2.9 μm by 87 μm , for an approximate area of 250 μm^2 . The detector area is 64 mm^2 , and would therefore see about 250,000 speckle grains (this is equivalent to $N= 4160$ for the 60-grain base units used in Figure 5.27). Even assuming fully developed speckle with the highest possible contrast ($C = 1$), this means an expected $dI/\bar{I}$ of 0.002. With the contrast found in Section 5.3.4, this falls to 0.001. In fact, a typical LSA scan of

office paper (using data from Section 4.2.4) has a fractional intensity of 0.011, a full order of magnitude greater than can be attributed to speckle.

This section has shown that there is a longer-wavelength component to the intensity pattern in free space than the high-intensity fluctuations described by speckle. In the experimental geometry this signal appears to be 3-4 times larger than could be accounted for by spatially averaging speckle. In the LSA geometry, calculations show that intensity variations from spatially averaged speckle are a full order of magnitude less than the intensity variations typically recorded.

5.5 LSA Scans with Incoherent Sources of Illumination

5.5.1 Introduction

The previous sections have raised several problems with attributing LSA to speckle-like diffraction effects. Given these problems, a valuable line of inquiry would be to isolate the parameter required for interference effects: coherence. The following experiments examine side-by-side the performance when the illumination source is fully, partially, or weakly coherent as the standard laser diode source is replaced with an LED and a thermal source.

This section includes two rounds of experiments. The first compares the results of scans taken with a laser to scans taken using a structured LED as the illumination source. The second round uses more sources (red laser, red LED behind an aperture, tungsten filament behind an aperture) but only compares scans taken with the same source.

For a continuous wave system such as LSA, the illumination can be assumed infinite in time. Given that, the coherence length relates to the bandwidth of the source through a simple Schwarz inequality (see, for example, [32] p. 309):

$$\Delta\omega \cdot \Delta t \geq \frac{1}{2} \quad (5.22)$$

The equation above assumes Δt and $\Delta\omega$ are the standard deviations of the temporal waveform and the angular frequency spectrum, respectively. The coherence time, τ_c , is of the order of Δt , and the coherence length, L_c , is the distance the wave travels at the speed of light, c , during τ_c . Thus:

$$L_c \sim \frac{\lambda^2}{\Delta\lambda} \quad (5.23)$$

A laser diode has a typical bandwidth of a few 10's of picometres, while an LED might have a bandwidth of ~50 nm, and a thermal source could have a bandwidth greater than 600 nm. This gives us a coherence length of ~2 cm for the laser, ~10 μm for the LED, and less than 1 μm for the thermal source.

5.5.2 Sources and Optical Geometry

Lasers have a number of secondary properties that make them desirable as an experimental light source. First, they are easily collimated, and therefore focused. Second, they easily produce reasonably high intensities. In order to replace the laser source with an incoherent one, the new source must also be of reasonably high intensity focused to a sharp line.

Rather than attempting to collimate and then focus incoherent sources, it is simpler to create a source of light which already has the desired structure, and image that using a simple 1:1 relay onto the target surface. The structured source took the form of either a specially-constructed LED or a traditional source shone through an aperture.

The structured LED was manufactured by PRP Optoelectronics. It consists of a chip with 256 individual LED elements, each $15\ \mu\text{m} \times 22\ \mu\text{m}$, set at a $42\ \mu\text{m}$ pitch (see Figure 5.28). Although the elements are individually addressable, for the purposes of this experiment all anodes and cathodes are made common. When powered up together, the chip produces a red light source 10 mm long and $15\ \mu\text{m}$ wide. All but the central 5 mm of the chip's face were blacked out, to produce the desired illumination line.




Figure 5.28 A row of 256 emission elements, each $15\ \mu\text{m} \times 22\ \mu\text{m}$ and set at a $42\ \mu\text{m}$ pitch, was manufactured to provide a structured emission source $15\ \mu\text{m} \times 10\ \text{mm}$. Figure reproduced with permission from PRP Optoelectronics.

For the second set of experiments, using both an LED and a tungsten lamp source, each was shone through an adjustable rectangular aperture. The LED was a Farnell LRW5SN, drawing $\sim 500\text{mA}$. The LED was attached to a heat sink and mounted directly behind the aperture. The tungsten-halogen filament lamp (8V, 50W, Comar 8 LE 50) was focused into 3mm acrylic optical fibre. The fibre terminated with a 25mm focusing lens, shone onto the aperture.

Each of these three sources was placed in the “Structured Source” position in Figure 5.29, aligned such that either the aperture or the face of the chip was in the focal plane. The power supply to the LED and the filament lamp were adjusted until the detected DC level was approximately the same for each source. In each case the targets were one of three ground glass diffusers polished with 120-, 220-, and 600-grit aggregates, corresponding to RMS height deviations about the mean of 3.1 μm , 1.6 μm , and 420 nm (see Section 4.3.3).

The structured sources needed to be imaged onto the target’s surface, while the collimated laser required focusing in one dimension. For this reason, the collimated laser was not interchangeable with the other sources in the apparatus. Instead the setup was constructed such that the lenses, aperture, structured sources, and beam block (see Figure 5.29) were all removable, allowing either configuration (a), for the laser, or configuration (b) for any of the structured sources.

Unfortunately, neither the various structured sources nor the targets could be mounted in a way that allowed them to be removed and replaced without requiring a re-alignment. Instead, the process was to align setup (a), align setup (b), take data for (b), then remove (b) and collect data for (a). This process was repeated for each target using the laser and structured LED.




Figure 5.29 The incoherent illumination apparatus in each of its two possible configurations. Lens f_1 is a cylindrical lens with $f = 60$ mm. Lens f_2 is a circular lens with $f = 40$ mm. Both CCD and photodetector are positioned at 45° from the primary optical axis. The resulting illumination line on the target is oriented such that the long axis comes up out of the page.

Alignment used a dummy target constructed with fiducial markings that indicated where the 5 mm illumination line should strike. With only the cylindrical lens (f_1) in place, the collimated laser (5 mm diameter spot size) was focused to a line on the dummy target. The dummy target was then adjusted until the fiducial markings were correctly oriented to the laser line.

With setup (a) aligned, the lens f_1 was removed from its post holder without altering the post's position with respect to the rest of the optical setup. Similarly, a beam block was introduced in front of the laser source, such that the laser's position could remain unchanged. The circular lens, f_2 , and its aperture (introduced to eliminate edge effects from the strongly diverging sources) were introduced along with one of the structured sources, and these elements were then focused and aligned to the fiduciary markings.

Next, the dummy target was removed and replaced with the target of interest. A still image in objective space was captured with the CCD system. Several LSA scans were then performed by repeatedly moving

the linear motion stage through 30mm and sampling from the photodetector every 10 μm . Once all the measurements using configuration (a) were finished, the aperture, beam block, structured source and f2 were removed and f1 reintroduced. A new still image was taken, and new LSA scans were acquired in configuration (b).

5.5.3 Results from Three Different Sources

For the first round of testing, this process was repeated for each of the three different GRIT ground glass diffusers using the structured LED as the ‘structured source’ for each round.

It is accepted that the visible variations in intensity produced when a laser strikes a rough surface (i.e. speckle) are not produced by a less coherent source. For completeness, this was illustrated by capturing an image from the CCD under each experimental configuration. Figure 5.30 on the left shows the image captured by the CCD when the target was illuminated with a laser and on the right when the target was illuminated with the structured LED. It is clear that any partial or short-length-scale coherence present in the light coming from the structured LED is insufficient to generate the free-space speckle pattern, which is clearly evident under laser diode illumination and completely absent under LED illumination.




Figure 5.30 The image captured by the CCD of the light scattered from GRIT 120 ground glass in free space, under illumination from a laser (LEFT) and an LED (RIGHT). The light scattered from LED illumination is uniform, as is expected in the absence of interference effects.

This made the most surprising result from this first study the fact that the data taken in configuration (b) using the structured LED generated an LSA-like signal (see the red trace in Figure 5.31). What is more, the signals matched in their larger features. The strong degree of correlation between the scans taken in configurations (a) and (b) is readily apparent in Figure 5.31. When considered in the light of Figure 5.30, this is another significant blow to the original interpretation of LSA, as a spatially-averaged speckle.




Figure 5.31 A section of two scans over GRIT 120 ground glass, one using configuration (a) with a laser diode as the illumination source, and one using configuration (b) with a structured array of LEDs as the illumination source. A DC offset of ~1000 has been applied to the intensity value of the “LED Array” trace, so that the scans overlap, and similar features are more easily noted.

When processed (see Figure 5.32) and compared as LSA scans, the two signatures match in 75.3% of their bits. Even under illumination that produces no speckle grains, the scattered light is very similar. Certainly this says that at least a large fraction of the LSA signal is unrelated to the short-length high-contrast order of speckle, which dominates what the eye sees.




Figure 5.32 The same section of the same two scans shown in Figure 5.31, but after processing. When run through an LSA comparison, these two scans yield a BMR of 75.3%.

With a match of roughly 75%, it is difficult to draw any larger conclusions. Clearly the two signals have larger features in common, while finer features deviate. Given that the alignment process was based simply on a visual estimation of fiducial markings, it cannot be said conclusively whether the small deviations indicate that one quarter of the LSA signal is dependent on speckle, or whether they can be accounted for by small variations in pitch, roll, yaw, and transverse positioning between configurations (a) and (b). A better alignment procedure will be executed in Section 5.6, to address this concern. Furthermore, the autocorrelation widths of the two scans (140 μm for the laser scan, 210 μm for the

LED scan) indicate that the laser was more tightly focused than the LED. This would also account for some of the difference between the scans.

Results from the second round of testing serve to demonstrate that even the broadest thermal source can generate repeatable LSA-like variations in intensity. The second round differs from the first in that more sources are used, but no attempt is made to compare results across sources. Scans are taken with a laser, a red LED behind an aperture, and a tungsten source behind an aperture. Repeated scans of GRIT 220 ground glass with each source are shown in Figure 5.33. The important feature here is that the quality of match, both by eye and through software analysis (IsaControl reports each match as 98% or better) is unaffected by the source. Indeed, the variations in intensity of diffusely scattered light are present whether the illumination source has a bandwidth of 10 pm or 600 nm. Data from the tungsten-filament also undercuts any arguments about the weak or partial coherence of the LED generating localised interference effects across its correlation length, as no such argument could be made for a thermal source.




Figure 5.33 Blue and red traces are repeated scans over the same location on a GRIT 220 ground glass sample, taken with a) a laser diode, b) a red LED behind an aperture, and c) a tungsten filament lamp behind an aperture. Each source yields a BMR of at least 0.98.

5.5.4 Supersensitivity as a Manifestation of Speckle

It is informative to look at what happens to the structured LED scans, and to the correlation between LED and laser scans, for finer microstructure. The values in Table 5.1 are the BMR (expressed as a

percentage) from comparing the scan indicated by the row with the scan indicated by the column. All scans are taken from the same position on the target; scan number simply indicates the iteration.

Grit 120 – 3.1 μm						
	LED 1	LED 2	LED 3	Laser 1	Laser 2	Laser 3
LED 1				74.7	74.4	75.4
LED 2	96.5			74.8	74.8	76
LED 3	97.8	96.4		74.9	74.7	75.1
Laser 1						
Laser 2				99		
Laser 3				97.8	98.7	
						0.0054
						0.003

Grit 220 – 1.6 μm						
	LED 1	LED 2	LED 3	Laser 1	Laser 2	Laser 3
LED 1				70.6	68.4	68.7
LED 2	88.3			71.4	72.5	72.8
LED 3	87.6	92.9		69.3	71.6	71.6
Laser 1						
Laser 2				96.7		
Laser 3				97.4	98.5	
						0.0036
						0.002

Grit 600 – 420 nm						
	LED 1	LED 2	LED 3	Laser 1	Laser 2	Laser 3
LED 1				68.2	67.2	69.9
LED 2	78.3			63.8	65.5	66.5
LED 3	75.3	83.7		62.1	65.6	61.2
Laser 1						
Laser 2				85.7		
Laser 3				90.1	90.5	
						0.0013
						0.0012

Table 5.1 The quality of match (BMR as a percentage) from different sources, and between sources. The number after the source indicates the scan number. Green cells are comparisons between scans taken with the LED source. Blue cells are comparisons between scans taken with the laser source. Yellow cells are comparisons between the two sources.

Paying attention first to the green and blue cells: the values listed can be interpreted as the maximum achievable match quality, as they are from a high precision linear stage with nothing disturbed between scans. For both systems, the maximum match quality falls with finer microstructure, as expected (see Section 4.3.5 for a detailed discussion). Of interest here is the fact that it falls faster and sooner for the incoherent source than for the coherent one. This would seem to indicate that any coherence effects are of little importance with large surface features, but are an increasingly large portion of the signal as the surface features grow finer. Such an interpretation, however, runs directly counter to the results of Section 4.3.5, where the rate of degradation of LSA with linear and rotational offsets is independent of feature size. If the signal from GRIT 120 is dominated by geometric optics while the signal from GRIT 600 has a significant contribution from interference effects, the two would be expected to behave quite differently with misplacement.

One possible explanation is the extreme accuracy of the linear stage used here. As described, the results from Section 4.3.5 were from a motion control system with a repeatability of $\sim 10\text{ }\mu\text{m}$. The Newport precision stages used in this round of experiments were the same as in Section 4.2.2, with a positioning resolution of 500 nm. This would mean that the motion control in this round was accurate enough to preserve even the highly sensitive interference effects, while the tolerance data did not.

This can be understood in the light of how each signal changes with surface roughness. Section 4.3.5 showed that the fractional intensity of the LSA signal decreases for smoother surfaces. For fully developed speckle, neither the contrast nor the grain size depend upon the amplitude of the surface roughness. This means that the fractional intensity of the spatially averaged speckle will remain constant with decreasing feature size, while the fractional intensity of the LSA signal decreases. If the intensity fluctuations from the non-coherence part of LSA decrease sufficiently, then the intensity fluctuations from speckle will become more significant.

For particularly smooth surfaces, interference effects would become the larger component of the LSA. In those cases where speckle was significant, the LSA signal would degrade much more quickly with misplacement than for most materials. This would manifest exactly as the supersensitivity observed in Section 4.2.4. Furthermore, it would explain the absence of supersensitivity in Section 4.3.5, where the robotic motion is not repeatable enough to reproduce the speckle pattern 100%.

In fact, any contribution to the signal that decoheres between 500 nm and $30\text{ }\mu\text{m}$ out of position is effectively another source of noise to LSA, rather than a contributing factor. This is easily confirmed with an IFS4000. The accuracy with which a target can be placed up against the bracket of a field scanner is the working threshold for what constitutes repeatable LSA. A few repeated scans of each ground glass diffuser, removing and replacing the target in the brackets each time, shows the same loss of performance as the LED results in the table above. Scans on 120-grit glass match at 95% or better, scans on 220-grit glasses match at less than 95%, and repeated scans of 600-grit glass match at less than 85%.

5.5.5 Conclusion

The investigations present in this section have shown that an LSA-like signal can be generated with weakly coherent and fully incoherent illumination, just as with the standard coherent source. Furthermore, it has been shown that an indeterminately large fraction of the LSA signal, at least 75% but possibly more, matches between coherent and incoherent illumination. Finally, the interference effects become more significant for targets with finer surface structure, which is proposed as an explanation of the supersensitivity from Section 4.2.4.

5.6 LSA Scans with Different Wavelengths of Illumination

5.6.1 Introduction

The most serious limitation of the data collected in Section 5.5 was the experimental inability to strike precisely the same position on the target when changing sources. While visual alignment to fiducials was sufficient to qualitatively indicate correlation between different sources, it was insufficient to quantify the contribution from phase.

The ideal setup would allow for the wavelength of the source to be adjusted without changing the optical alignment in any way. While lasers that are tuneable over the visible spectrum exist[85], they are still prohibitively expensive and complicated to run. A more economical, if less dramatic, solution is to use fibre-coupled lasers. The fibre-optic cable would decouple the selection of the source from the aiming and focusing of the resultant illumination. This would, in turn, allow an analysis of the impact of source wavelength, freed from the uncertainties of variable positioning.

While all of the illumination used in such an experiment would be coherent, a wide enough selection of wavelengths can be chosen that the speckle patterns will be completely uncorrelated. In this way changing wavelengths will be sufficient to rule out any dependence on interference phenomena.

5.6.2 Theoretical Independence of Speckle at Different Wavelengths

Considering again the method used in 5.2.2, it is necessary to calculate the change in scattering vectors, $\Delta \vec{q}$, from a shift in wavelength with LSA's geometry, which is shown in Figure 5.34.




Figure 5.34 (a) The original LSA scattering geometry. (b) The LSA scattering vector. (c) The LSA scattering vector after an increase in wavelength.

Considering equations (5.3) - (5.5), and taking $\theta_{i1} = \theta_{i2} = 0^\circ$ and $\theta_{o1} = \theta_{o2} = 45^\circ$ gives

$$\Delta q_t = \left| \frac{2\pi}{\lambda_1} \left(\frac{\sqrt{2}}{2} \right) - \frac{2\pi}{\lambda_2} \left(\frac{\sqrt{2}}{2} \right) \right| \quad (5.24)$$

and

$$\Delta q_z = \left| \frac{2\pi}{\lambda_1} \left(\frac{\sqrt{2}}{2} + 1 \right) - \frac{2\pi}{\lambda_2} \left(\frac{\sqrt{2}}{2} + 1 \right) \right| \quad (5.25)$$

Equation (5.7) assumes a circularly-symmetric spot, which is certainly not true here. However, as correlation width varies inversely with spot size, the smaller dimension of the illumination ellipse provides an upper bound on the expected correlation. Ψ is considered to decohere completely after the first zero. Taking the most conservative case of $D = 50 \mu\text{m}$, this occurs at $\Delta q_t > 0.15/\mu\text{m}$. M_h can be said to have effectively decorrelated when it has fallen below $1/e^2$. From (5.8) with $\sigma_h = 3.1 \mu\text{m}$, this occurs at $\Delta q_z > 0.45/\mu\text{m}$.

The lasers used for this experiment will be green (532 nm), red (671 nm) and infrared (808 nm). Moving from the red to the IR laser, where $\lambda_1 = 671 \text{ nm}$ and $\lambda_2 = 808 \text{ nm}$, gives Δq_t of $1.22/\mu\text{m}$ and a Δq_z of $2.71/\mu\text{m}$. Moving from the red to the green laser, where $\lambda_1 = 671 \text{ nm}$ and $\lambda_2 = 808 \text{ nm}$, gives Δq_t of $1.22/\mu\text{m}$ and a Δq_z of $2.71/\mu\text{m}$. In all cases this is beyond the point where speckle will have decorrelated, both from the M_h and from the Ψ components. From the speckle point of view, each of the three laser tests should be unrelated.

5.6.3 Apparatus & Methodology for Scans with Different Wavelengths

The scan motion is generated using the same apparatus shown in Figure 5.29, with the sample mounted to a computer-controlled linear motion stage. The optical pathway is replaced with the one described in Figure 5.35 and pictured in Figure 5.36. All data was taken on GRIT 120 ground glass.




Figure 5.35 A diagram of the apparatus for measuring the LSA signal with different wavelengths of illumination. A camera mirrored the photodetector at 40 mm.




Figure 5.36 The apparatus for measuring the LSA signal with different wavelengths of illumination, with the printer label target loaded for focusing.

Three different fibre-coupled 10mW solid state lasers are used for this experiment: green (532 nm), red (671 nm) and infrared (808 nm). Each laser could be connected to a 1 m single-mode fibre optic cable,

terminating in a fibre-coupled collimator. The collimator was clamped to the optical breadboard, and the cable was strain-relieved, such that changing the source laser had minimal impact on the mechanics of the collimator.

A single photodetector was used throughout. The SD 5443's spectral response peaks at 808 nm, and is down to a relative value of ~ 0.3 at 532 nm. For this reason the value of the neutral density filter was adjusted for each laser, to bring the intensity level within the active region of the acquisition system.

The collimated light was passed through a x10 beam expander, and shone through a rectangular aperture. The aperture set the final length of the focused line to 3 mm. A cylindrical lens, f1, focused the collimated light in one dimension.

While moving the free end of the optical cable did not change the focus or the alignment of the collimator's output, it did introduce variations to the polarisation of the light. For this reason, a polarising element was introduced as well, which ensured the incident illumination was always p-polarised. Each time the source laser was changed, the optical cable was adjusted until the intensity coming through the polariser was maximised.

As in Section 5.5, five scans were acquired by moving the linear stage through 30 mm while sampling from the photodetector at 10 μm intervals, with a still image captured after each scan. This sequence was performed for the red laser, the green laser, the IR laser and finally repeated for the red laser.

The entire apparatus was inside a set of isolation curtains, which could be drawn shut to eliminate ambient light. The second round of red laser data controlled against any loss of BMR from decoupling and recoupling the laser sources and opening and closing the isolation curtains.

5.6.4 Results from Fibre-Coupled Illumination

The 2D normalised cross correlation described in 5.2.2 can evaluate the correlation of the speckle field for this particular observation point for each of the three different wavelengths. Evaluating the cross correlation between images captured by the CCD for different wavelengths of illumination will serve two functions. For the first and second pass with the red illumination, it will confirm that the experimental apparatus is precise enough to strike the same region of the surface after changing laser sources. Correlations between red and green or red and infrared speckle patterns will demonstrate that, despite originating from the same area of the target, the speckle patterns are unrelated.

In each case the first image is the same: the objective speckle field when the target was first illuminated with the red laser. Comparing two of the still images from the first red set gives a central peak in the

normalised cross correlation of 0.9139. This confirms that the Newport linear stage is repeatable enough to preserve speckle.

The normalised cross correlation of a red laser image with a green laser image, or a red laser image with an IR laser image, do not exhibit any central peak at all (see Figure 5.37). This is confirmation of the calculations from the previous section, showing that a change in the wavelength of incident light of $\pm 20\%$ is sufficient to cause the speckle to decohere completely for a fixed observation point at 45° .




Figure 5.37 The normalised cross correlation (right) between two images of speckle, one illuminated with the 671 nm light (left), and one with 532 nm light (centre). The faint vertical and horizontal lines on the cross correlation are artefacts of stitching the image together so that the zero-offset position is centred. A similar result is seen for infrared.

Returning to the red laser restores the central peak, although its maximum value has fallen to 0.319 (as shown in Figure 5.38).




Figure 5.38 The normalised cross correlation (right) between two images of speckle illuminated with the 671 nm laser, before (left) and after (centre) decoupling and re-coupling the fibre with a laser three times.

Working around the optical setup (decoupling the back end of the optical fibre several times, opening and closing the isolation curtain, the camera's constant re-normalising of intensity, and 20 passes of the Newport stage) has introduced only enough variation for the speckle to decohere half way, confirming

that the illumination spot is still in approximately the same position. Since the second red set comes at the end of the experiment, it is safe to assume that it is the worst positioned of the sets, and that the green and IR lasers are also striking the same position on the target.

Having established that the positioning is reasonably sound and that the speckle is fully decorrelated from one wavelength to the next, the LSA signatures themselves can now be evaluated. Scans from the first red set are taken as the reference scans. Scans within a set (that is, when the apparatus has not been touched between scans) have an average BMR of 0.999, so the drop to 0.90 (671 nm on Figure 5.39) can be interpreted as the loss from the fine variations in positioning and apparatus—the same variations which caused the speckle’s cross correlation to fall from 0.9 to 0.3. It is worth noting the respective ranges of BMR and NCC, when moving between the two units. BMR runs from 0 to 1, with 0.5 indicating no correlation. NCC runs from -1 to 1, with 0 indicating no correlation. Thus an NCC of 0.3 is still a significant degree of correlation.




Figure 5.39 The BMR of scans taken with the indicated wavelength of laser light, compared with a set of reference scans taken with 671 nm. When the apparatus is not readjusted in any way, scans match at 99.9%

Using a BMR of 0.9 as the control against minor changes in positioning, the impact of a change in wavelength is minor indeed. Decreasing the wavelength by 20% (to 532 nm) and increasing the wavelength by 20% (to 808 nm) causes the BMR to drop an additional 0.015 and 0.041, respectively.

5.6.5 Conclusions of the Fibre-Coupled Study

The experiment described here is an extremely important one in defining the optical phenomenon behind LSA. Section 5.5 established 25% as a very rough experimental upper bound for the component of the LSA match from interference effects. By carefully controlling the experimental apparatus, that

figure is now reduced to somewhere between 4.0% and 1.5%. This was done not by using incoherent illumination, but rather several sources of coherent illumination whose wavelengths were far enough apart that any phase-dependant signal had decorrelated completely.

As can be seen from the error bars on Figure 5.39, 2% – 4% BMR is well within the error for repeated LSA scans. This section has demonstrated that for the typical surface, represented here by ground glass, phase effects are not a significant factor in the LSA signal.

5.7 Conclusions

In this chapter it has been shown conclusively that the diffusely scattered light measured by LSA is not a result of phase-dependant interference effects in general, and in particular taking an LSA scan does *not* involve “making use of the optical phenomenon of laser speckle.”[6]

The first two sections examined the measured and expected tolerance of speckle to misplacement, using the LSA geometry. These results were considered in light of the data collected for LSA in Chapter 4, and it was concluded that classical speckle was not a sufficient explanation for the robustness of LSA on typical materials.

The third section concluded that the neither was classical speckle sufficient to explain the signal strength of LSA. Measurements taken in this section indicated the presence of longer-length components to the scattered intensity.

Whether these components were higher-order speckle correlation or another phenomenon was answered by the final two sections, the penultimate of which concluded that interference effects of any kind are not necessary to produce the kinds of intensity variations in diffusely scattered light that LSA measures. While this showed that an LSA-like signal could be generated from incoherent light, it could not determine whether the measured signal is a combination of speckle and other effects. The final section addressed this question, concluding that phase effects are a negligible part of the typical LSA signal.

6

TWO-SCALE BISTATIC SCATTERING MODEL

*“Every valley shall be exalted, and every mountain and
hill shall be made low: and the crooked shall be made
straight, and the rough places plain”*

—Isaiah 40:4

6 Two-scale Bistatic Scattering Model

6.1 Introduction

Using the findings of the previous two experimental chapters, the goal is to construct a scattering model which would accurately predict the behaviour of LSA and provide some confirmation of the conclusions about the primary mechanisms behind it.

The motivation behind this model is significantly different from the models which exist in the literature. Various scattering approximations are meant to help determine the *average* behaviour of an ensemble of surfaces with the same statistical definition. This is precisely the opposite of what is interesting to model for LSA: the *variations* between individual instances of surfaces with the same statistical definition. This difference in goals manifests, for example, in the choice of geometric factors based on local (variable) surface properties, rather than on average (constant) properties. Nevertheless, this model is built largely from components already justified and tested in existing models.

The model for the rough surface is simple Gaussian, described completely by the RMS surface height and correlation length. The incident illumination is considered a plane wave and wrapped in an amplitude envelope to define the spot size.

Diffraction-based models are completely dominated by short-wavelength variations[19]. The most rigorous option would be to follow the theory laid down by Parry in [86] and compute full diffraction pattern at a series of discrete wavelengths to cover the full bandwidth of the source, and sum results. First, for a wideband source like the tungsten filament in Section 5.5, this would effectively add another dimension to the already expensive calculations. Second, this would imply that the results should be different for a narrow, moderate, or wideband source, which experimentally is not the case. Essential, such a model would spend millions of processing cycles calculating fluctuations which would average each other out.

Instead, the scattering mechanism modelled here begins with a further simplification, justified by the findings of Sections 5.5 and 5.6: it is set at the infinite-frequency limit, to remove the influence of wavelength. At this limit, physical optics (e.g. equation (1.4)) collapses down to geometric optics. Onto this basic framework are added the angular dependence of Fresnel reflection coefficients to the local surface normal and a second scale of roughness, in the form of Lambertian scattering.

The next section describes the foundational work by other researchers. Section 6.3 describes the analytical form and the numerical implementation of each component of the simulation. The next

section describes the output of the simulation and compares those results to previous experimental data. In light of the simulation's output, Sections 6.5 and 6.6 discuss further improvements which could be made to the model and to the LSA system, respectively.

6.2 Foundations

6.2.1 Foundations of Surface Generation

The surface generation method described in Section 5.3.3 is expanded into two dimensions. Pardo-Iguzquiza and Chica-Olmo[87] break down the regions of symmetry in two-dimensional frequency space for a real surface, illustrated in Figure 6.1.




Figure 6.1 The symmetries in complex Fourier coefficients for a 2D surface with number of points N_1 and N_2 . * denotes a complex conjugate.

Two-scale models for surface roughness are fairly common, most often to describe radar interaction with the ocean surface where small capillary waves are superimposed on longer, wind-generated waves[13, 88]. Elton[89] extends this two-scale roughness to apply to optical interaction with a number of man-made surfaces, included coated paper, printed paper, and paints. Finally, Koenderink[90] pushes the model further, discussing four scales of roughness for manmade objects. More importantly, he extends the discussion beyond the two discrete scales found in paper or the ocean surface. In fact the RMS value used to describe rough surfaces characterises only the *largest* scale of roughness present on the surface. Except for fractured surfaces, the description of rough surfaces as many randomly-oriented and perfectly smooth facets is not a physical one. Surfaces are actually rough at every scale from the largest roughness described by the surface profile, down to the atomic-scale. This opens the discussion of two- or multi-scale roughness to surfaces, like ground glass, which are typically categorised by only one scale of roughness.

The pervasive nature of multiple scales of roughness motivate including them in the model in a different fashion. Two-scale models like Ogilvy's are applied to generate a more complex surface profile $h(x, y)$, which is then fed into a standard scattering model. This model will take a more recent approach (see, for example, [91]) and integrate the second scale of roughness into the scattering mechanism. That is, $h(x, y)$ will describe only a single scale of roughness, and the scattering mechanism will assume an additional level of roughness beyond that.

6.2.2 Foundations of the Scattering Model

The important contribution of the surface profile for speckle models is the local height, $h(x, y)$. This is what introduces the path-length difference, and therefore the phase shift, that leads to an interference pattern. For a geometric optics model, the important contribution is from the local gradient and the distribution of angles that compose the surface. Unfortunately, this also leads to the first problem encountered with geometric optics: specular reflection from planar facets can account for the observed intensity at large angles only for surface with an extremely wide distribution of angles. Generally, geometric optics is valid only for very rough surfaces or for illumination and observation where the paraxial condition is met. For surfaces with roughnesses and correlation lengths like those examined in the experimental chapters, the surface angle distribution rarely extends to 22.5° , the angle necessary to generate reflections at LSA's 45° observation angle using simple ray-tracing.

This is the difficulty Vernold and Narvey[37] highlighted. They had some success extending Beckmann-Kirchoff beyond the paraxial limitation by replacing the geometric factor, F_3 , with the simple Lambertian angular dependence, $\cos \theta$. They calculate this factor against the average normal for the whole surface, $\vec{m}$, rather than against the instantaneous or local mean, $\vec{n}$, at a particular point on the surface.

Using $\vec{m}$ is based on the assumption that the correlation length is much less than the illumination spot diameter, which is only true in one dimension for LSA. Moreover, eliminating the need to track phase has freed up considerable computational resources that can be dedicated to computing the angular dependence based on the local mean at each point on the surface. Effectively this model will move the geometric factor, F_3 in equation (1.4), inside the integral.

The final piece of the scattering model follows the step taken by Wolff, Nayar and Oren[36]. In an effort to update computer graphics scattering calculations beyond the simple Lambertian picture or the microfacet surface model, they introduced the Fresnel reflection coefficient into the geometric considerations for scattering. Their attenuation factor is an average of the s- and p-polarised reflection coefficients:

$$R_s = \frac{\sin^2(\theta_t - \theta_i)}{\sin^2(\theta_t + \theta_i)} \quad (6.1)$$

and

$$R_p = \frac{\tan^2(\theta_t - \theta_i)}{\tan^2(\theta_t + \theta_i)} \quad (6.2)$$

where θ_i is the incident angle and θ_t the transmitted angle, given by Snell's law for reflection at an air-medium interface, with index of refraction, n :

$$\theta_t = \sin^{-1}\left(\frac{\sin \theta_i}{n}\right) \quad (6.3)$$

6.3 Formulation of the Two-Scale Simulation

6.3.1 Overview of the Two-Scale Model

Figure 6.2 illustrates all of the pieces which must be developed as part of this scattering model. A plane wave with amplitude envelope $A(x,y)$ is incident on the random rough surface, $h(x,y)$. At each point on the surface the incident vector $\vec{k}_{inc}$ reflects specularly off the local tangent plane, about the local normal $\hat{n}$. Assuming each planar facet is itself rough on a smaller scale, diffuse light scatters about the reflected vector with an intensity determined by the local Fresnel reflection coefficients and a Lambertian envelope, $\rho(\theta)$. The resulting intensity $I(\alpha, \beta)$ is measured in the $\alpha - \beta$ observation plane.




Figure 6.2 The basic components of the two-scale scattering model for LSA.

6.3.2 Generation of Random Two-Dimensional Rough Surfaces

The surface is assumed to be isotropic and stationary. To fully describe the surface, the following parameters must be defined: the correlation length, L_c and the RMS height deviation, σ_h , desired for the final surface; the surface's extent in x and y (L_x and L_y , respectively); and the desired resolution, δ .

These parameters taken together determine the number of points in each dimension, $N_x = L_x/\delta$ and $N_y = L_y/\delta$.

If m and n are the two indices, running from 0 to N_x and N_y , respectively, then the two-dimensional wave number $K_{m,n}$ can be defined as:

$$K_{m,n} = \sqrt{\left(\frac{2\pi m}{L_x}\right)^2 + \left(\frac{2\pi n}{L_y}\right)^2} \quad (6.4)$$

The spectral weighting function from Section 5.3.3 becomes the discrete spectral weight:

$$W_{m,n} = \frac{L_c \sigma_h^2}{2\sqrt{\pi} e^{\left(\frac{-K_{m,n}^2 L_c^2}{4}\right)}} \quad (6.5)$$

Finally, the two-dimensional Fourier spectrum can be calculated, piecewise, by:

$$F_{m,n} = \sqrt{\frac{2\pi L_x L_y W_{m,n}}{\delta^4}} \begin{cases} \frac{R(0,1) + iR(0,1)}{\sqrt{2}} & m, n \in \beta, \delta, \epsilon, \phi, \gamma, \kappa \\ R(0,1) & m, n \in \alpha, \chi, \eta, \psi \end{cases} \quad (6.6)$$

Once the regions $\beta, \delta, \epsilon, \phi, \gamma$ and κ (indicated in Figure 6.1) have been filled, their mirror regions are assigned their complex conjugate. With the full F matrix complete, an inverse FFT will recover the surface height.

$$h(x, y) = \mathcal{F}^{-1}\{F_{m,n}\} \quad (6.7)$$

It only remains, then, to calculate the local surface normal at each point. For each (x, y) position on the surface the point is taken as the centre of a rectangular plane of dimensions $\delta \times \delta$, with local normal:

$$\vec{n}_{x,y} = \begin{pmatrix} 2\delta \\ 0 \\ h(x+\delta, y) - h(x-\delta, y) \end{pmatrix} \times \begin{pmatrix} 0 \\ 2\delta \\ h(x, y+\delta) - h(x, y-\delta) \end{pmatrix} \quad (6.8)$$

To eliminate any potential edge effects in the transform, the surface is generated an extra $5L_c$ larger than desired, and then cropped back down to the desired size. A numeric simulation implementing this method is shown in Figure 6.3.




Figure 6.3 Random surface generated with desired RMS height and correlation length.

The index of refraction of the reflecting medium is designated at this point, although it is not used during surface generation. The resolution, δ , is determined by dividing the correlation length by the

“Calculations per correlation area” factor. The z axis on the 3D plot is in mm, while the x and y axes are in units of δ . In the example in Figure 6.3, $\delta = \frac{0.05 \text{ mm}}{7} = 7.14 \mu\text{m}$, and the surface is 1 mm on each side, hence the axes run from 0 to 140. The minimum and maximum extent of the surface in x and y determine both the surface dimensions, L_x and L_y , and the position of the surface relative to the origin, which is at 0,0. That is, a surface defined as 0 mm to 4 mm and a surface defined by -2 mm to 2 mm would both be 4 mm long, but in the first case the origin would be at one edge, while in the second case the origin would be in the centre.

Once the surface height function has been generated, the actual RMS height and correlation length (the FWHM of the autocorrelation of a line profile taken through the centre of the surface) are calculated, as a check. These values are displayed on the top right of Figure 6.3.

6.3.3 Defining the Illumination Envelope

The incident illumination is assumed to be a plane wave, fully described by a central vector and an amplitude envelope. The directional vector is applied identically across the whole extent of the spot.

Initially, two geometries were tried for the amplitude envelope: an ellipse with unit magnitude inside and zero magnitude outside, and an ellipse with a Gaussian profile in x and y. The first is clearly unphysical, and the second was a bad match to the LSA optics. This was particularly true in the y-dimension where the laser line was 3 mm long; a Gaussian profile has only a bright central region with long tails (see Figure 6.4, centre), whereas the aperture and cylindrical lens used in LSA produces a long central section of fairly uniform intensity, and sharper falloff at the edges. A hybrid, called “extended Gaussian,” was used to avoid both of these problems.




Figure 6.4 Three different amplitude envelopes applied to the illumination of the rough surface: unity ellipse (left); Gaussian ellipse (centre); and extended Gaussian ellipse (right).

For a spot with a full width of W , the extended Gaussian profile is defined piecewise so that the amplitude is unity to $\pm \frac{W}{4}$ and has fallen to 95% by $\pm \frac{W}{2}$. Defining $\sigma = \frac{W}{8}$ gives:

$$f(x) = \begin{cases} 1 & \sqrt{\frac{x^2}{2\sigma^2}} \leq 2 \\ \exp - \left(\sqrt{\frac{x^2}{2\sigma^2}} - 2 \right)^2 & \sqrt{\frac{x^2}{2\sigma^2}} > 2 \end{cases} \quad (6.9)$$

It must also be possible to rotate the illumination spot within the plane, so that its long axis is at an angle, ε , from the y-axis. Extending (6.9) into two dimensions and incorporating the rotating, the final amplitude envelope, $A(x, y)$, for an ellipse centred at x_0, y_0 becomes:

$$A(x, y) = \begin{cases} 1 & d(x, y) \leq 2 \\ e^{-(d(x, y) - 2)^2} & d(x, y) > 2 \end{cases} \quad (6.10)$$

where

$$x' = (x - x_0) \cos \varepsilon - (y - y_0) \sin \varepsilon \quad (6.11)$$

$$y' = (x - x_0) \sin \varepsilon + (y - y_0) \cos \varepsilon \quad (6.12)$$

so that

$$d(x, y) = \sqrt{\frac{x'^2}{2\sigma_x^2} + \frac{y'^2}{2\sigma_y^2}} \quad (6.13)$$

The numerical implementation of this process is shown in Figure 6.5.




Figure 6.5 Extended Gaussian amplitude profile, centred at 0,0 and rotated by -15°.

K_x , K_y , and K_z describe the incident vector, $\vec{k}_{inc}$. In addition to defining a centre, extent, and angle for the amplitude envelope, a threshold value (“Ignore below:” in Figure 6.5) was set to reduce computational time. Integrations over the surface were calculated only for points with an amplitude greater than this threshold value. This reduction becomes significant when large (15 mm x 8 mm) surfaces are defined with only a narrow region (100 μ m x 3 mm) illuminated at any given time.

6.3.4 Defining the Observation Plane

The x-y plane describes the scattering surface. An observation plane is also required to calculate scattered intensity; the observation plane is here defined as the detector. The vector $\vec{P}$ describes the position of the centre of the detector, relative to the centre of the illumination spot. The vector $\vec{N}$ describes the surface normal of the detector, or its gaze. In most cases $\vec{N} = -\vec{P}$, indicating that the detector is gazing at the centre of the laser line. If $\vec{s} = (x_0, y_0, 0)$ is the in-plane vector describing the position of the centre of the laser spot, then the centre of the detector is at $\vec{p} = \vec{s} + \vec{P}$.

The dimensions of the detection plane are described both in terms of a length (d_y) and width (d_x), and a number of pixels. The intensity at each pixel is calculated separately—this allows both the simulation of CCD-like detection of the free-space intensity pattern as well as an LSA-like simulation of an integrated intensity value.

The observation plane defines a second set of basis vectors, $\hat{\alpha}$ and $\hat{\beta}$, as:

$$\hat{\beta} = \frac{\vec{y}_{\parallel}}{|\vec{y}_{\parallel}|} \quad (6.14)$$

$$\hat{\alpha} = \frac{\hat{\beta} \times \hat{N}}{|\hat{\beta} \times \hat{N}|} \quad (6.15)$$

where

$$\vec{y}_{\perp} = (\hat{N} \cdot \hat{y})\hat{N} \quad (6.16)$$

$$\vec{y}_{\parallel} = \hat{y} - \vec{y}_{\perp} \quad (6.17)$$

Using these new basis vectors and the pixel size, $c_x \times c_y$, it is possible to describe the position vector, $\vec{R}$, to the centre of each pixel:

$$\vec{R}_{m,n} = \vec{p} + \left(m c_x - d_x/2\right) \hat{\alpha} + \left(n c_y - d_y/2\right) \hat{\beta} \quad (6.18)$$

where m and n index the pixels.

Detector sensitivity is not modelled explicitly. Instead it is combined with the illumination amplitude. For example, the photodetector used most commonly for LSA, the SD5443, has an acceptance angle of $\pm 10^\circ$. At a detector distance of 10 mm this corresponds to the central 3 mm of the 5 mm laser line. To model LSA with this detector, the length of the illumination spot is set to 3 mm.

6.3.5 Diffuse Scattering from Planar Facets

For the purpose of this scattering model, the position vector is defined as

$$\vec{r} = \begin{pmatrix} x \\ y \\ h(x, y) \end{pmatrix} \quad (6.19)$$




Figure 6.6 Specular reflection from a plane.

As illustrated in Figure 6.6, the specular reflection off of a plane with normal $\vec{n}_{x,y}$ reverses the component of the vector along the normal, so the reflected vector is:

$$\vec{k}_{ref} = \vec{k}_{inc} - 2(\vec{k}_{inc} \cdot \hat{n}_{x,y})\hat{n}_{x,y} \quad (6.20)$$

The intensity of the light reflected in this way is determined by the Fresnel reflection coefficient. This will be reflected in the quantity

$$\tau(x, y) = \frac{R_s + R_p}{2} \quad (6.21)$$

with R_s and R_p from equations (6.1) and (6.2), where

$$\theta_i = \cos^{-1} \left(\frac{|\vec{k}_{inc} \cdot -\hat{n}_{x,y}|}{|\vec{k}_{inc}|} \right) \quad (6.22)$$

and θ_t is then calculated from equation (6.3).

The detected intensity in the framework described so far will be $A \tau \vec{k}_{ref}$ if the reflected vector intersects with a pixel, and 0 if it does not.

The additional step is now taken of assuming that the facet under consideration is itself also an optically rough surface. In that case, light would reflect not just at the specular angle, but in a classic diffuse lobe around this angle. Some light will be detected from every position on the scattering surface that is ‘forward’ of the scattering vector, with the intensity of the light determined by the deviation from specular. ‘Forward’ here means contained in the hemisphere centred about the scattering vector.

Lambertian scattering is the simplest model for diffuse scattering, and predicts that the intensity of the scattered light will be scaled by the cosine of the angle between the surface normal and the observer. The assumption made here deviates from traditional Lambertian, insofar as it incorporates knowledge of the specular direction. Instead of the angle between the observer and the normal, the intensity is scaled based on the angle between the observer and the specular direction, as indicated in Figure 6.7.




Figure 6.7 A Lambertian-esque dependence will scale the intensity of the light reaching the detector element by the cosine of the angle θ .

With the position vector to the detector element $\vec{R}_{m,n}$ from (6.18) and the position vector to the scattering location $\vec{r}$ from (6.19) the new Lambertian scale factor γ takes account of the angle from specular and the solid angle subtended by the pixel. If ρ is defined as the dot product of the vector to the observer and the specular vector

$$\rho = (\vec{R}_{m,n} - \vec{r}) \cdot \vec{k}_{ref} \quad (6.23)$$

then the Lambertian scale factor is

$$\gamma_{m,n} = \begin{cases} 0 & \rho \leq 0 \\ \frac{\rho}{|\vec{R}_{m,n} - \vec{r}| |\vec{k}_{ref}|} \frac{c_x c_y}{|\vec{R}_{m,n} - \vec{r}|^2} & \rho > 0 \end{cases} \quad (6.24)$$

Combining equations (6.10),(6.20),(6.21) and (6.24) leads to the final expression for the intensity of light scattered from surface S to detector element m, n :

$$I_{m,n} = \iint_S A \tau \gamma_{m,n} \vec{k}_{ref} dx dy \quad (6.25)$$

Each element in the integral above has a dependency either on the local surface height or the local surface normal. The total integrated intensity registered by the detector is calculated by summing over all m and n .

Figure 6.8 shows the result of a numerical implementation of this scattering model. The $\vec{P}$ and $\vec{N}$ vectors are defined, along with the height, width, and pixel count of the detector. The resulting plot is of the intensity value at each pixel, and “integrated Lambertian intensity” reports the intensity summed across all pixels.

The “image” detected by this simulated CCD is uniform to within ~5%. This is a useful check against the CCD images shown in Figure 5.30, contrasting the free-space intensity pattern generated by a laser with the pattern generated by an LED. In this simulation, just as in reality, removing interference effects removes the local variations in scattered intensity, effectively ‘washing out’ the CCD image. If the conclusions of Chapter 5 are correct, removing the local intensity variations will not be sufficient to remove the variations in scattered intensity between different illumination positions.




Figure 6.8 The free-space scattered field calculated numerically. The x- and y-axis are in units of pixels, while the z-axis is an arbitrary scattered intensity.

6.3.6 Simulation of Complete LSA Scans

The simulation as described so far is able to calculate both the free-space intensity pattern and the total integrated intensity for an arbitrary detection element in space, from a surface with specific statistics when illuminated by an ellipse of arbitrary size, with an arbitrary incident angle. In order to simulate a full LSA scan, it must repeat the calculation of integrated scattered intensity for a series of illumination positions. To simulate tolerance measurements, like those taken in Chapter 4, whole scans must be repeated with small changes to illumination and observation conditions.




Figure 6.9 In order to simulate an entire LSA scan, the integrated scattering intensity is calculated for the same surface but for a series of different illumination centres, moving across the surface in x. The intensity at each position is recorded, just as in a real LSA scan.

In all cases, the first step is to generate an instance of a surface with the desired statistics. This is done only once for each scan or each axis, as the same surface must be used for all scattering measurements. The surface must be large enough for the desired scan or repeated scans.

To simulate a scan, the initial position of the illumination position must be described fully. The simulation is then given a step size and a maximum value, and the integrated intensity is calculated for an identically-sized illumination spot centred at each point between the initial and maximum values. Describing the position of the observation plane with the vector from the spot centre, rather than from the origin, ensures that the simulated detector tracks along with the illumination spot across the surface.

To simulate tolerance to misplacement along or about an axis, the above procedure for generating a simulated scan is repeated first for an initial set of conditions, and then with a series of small adjustments to those parameters between a user-defined minimum and maximum offsets, with an arbitrary step size. To simulate yaw, the angle of the spot's major axis and the vectors describing the position and gaze of the observation plane are all rotated about the z-axis from their initial position. To simulate transverse, the y-coordinate of the spot centred is stepped across a range of value, and the position and gaze of the detector track automatically. For roll and pitch, the incident illumination vector, the position vector of the observation plane, and the gazing vector of the observation plane are all rotated about the x- or y-axis.

For each offset position, the normalised cross correlation is computed between the scan taken with that offset and the original (home) scan. The maximum correlation coefficient as a function of offset is reported by the software.

6.4 Numerical Results of the Two-Scale Simulation

6.4.1 A Single Simulated Scan

As a first result, consider the similarity of the experimental and simulated scans shown in Figure 6.10. The experimental scan is taken from Section 5.5 for a red LED behind a 250 μm aperture, shone on 120-grit ground glass. The simulated scan was generated from a surface with a 3.0 μm RMS height deviation and a 40 μm in-plane correlation length, using a 220 μm by 3.7 mm illumination spot. Both sets of data have been renormalised to a mean of 1.




Figure 6.10 A real scan (TOP) taken from GRIT 120 ground glass, and a simulated scan (BOTTOM) using appropriate parameters.

The simulated scan is a slightly weaker signal and contains higher frequency components. These two properties are both functions of the illumination spot size, and will be examined in more detail in the next section.

6.4.2 Fractional Intensity and Correlation Length of Simulated Scans

The first two metrics against which the model can be quantitatively evaluated are properties of an individual intensity trace. The fractional intensity and the autocorrelation length can be checked, both for a match against a standard LSA scan of 120-grit ground glass, and against additional data showing how $dI/\bar{I}$ and autocorrelation width varied with aperture size.

Configuration (b) from Section 5.5 (see Figure 5.29) used an adjustable aperture to set the width of the illuminating LED line. Using that same rig, scan data was acquired for apertures from 100 μm to 1000

μm wide. The scans were analysed to extract the two parameters which depend on aperture width: the autocorrelation width of the signal, and the fractional intensity of the signal.

Autocorrelation width is taken as the FWHM of the central peak of the scan's autocorrelation function. Fractional intensity is calculated from the AC-coupled standard deviation over the DC mean of the signal. However, rather than the traditional band-pass filter, the AC-coupled signal is generated with a high-pass filter with a cut-off of 3000 μm . This ensures that the autocorrelation width is not inadvertently broadened by the application of the filter. Both autocorrelation width and fractional intensity were found to vary linearly with aperture width, and therefore the width of the illumination spot.

For comparison, six different scans were simulated, using parameters that match LSA on 120-grit ground glass. A 15.6 mm by 4.6 mm surface was generated at 5.7 μm resolution, with a 40 μm correlation length and a 3 μm standard deviation in. The illumination was normally incident and the spot was set to 3.7 mm long. The detector was positioned 10 mm from the focal point, at 45° (or $P = (7,0,7)$ and $N = (-7,0,-7)$ in mm) and each scan ran along the central 15 mm of the surface.

Only one surface was generated, and used for all six scans. The scans were identical except for the widths of the illumination line, which were set to 100 μm , 160 μm , 220 μm , 350 μm , 500 μm , and 700 μm to simulate different aperture widths. The FWHM and the $dI/\bar{I}$ of the simulated scans are calculated just as for the experimental data. The simulated and experimental results are plotted together in Figure 6.11.




Figure 6.11 The effect of aperture width on the autocorrelation and fractional intensity of experimental and simulated scans.

The simulated autocorrelation widths also show a linear relationship with aperture width, although with a steeper dependence. The difference in slope is likely caused by different definitions of spot width and intensity profile. As discussed, the model's width defines the 95% width of an extended Gaussian profile, while the experimental width defines the size of the aperture. If this corresponds to the 95% width of a Gaussian profile, the reported experimental width will be an overestimation compared to the simulated width. This would effectively elongate the experimental slope, which is in keeping with the figure above.

A more significant disagreement is in the *form* of the fractional intensity results in the top graph of Figure 6.11. A simple physical interpretation of the fractional intensity is that it is a function of the number of scattering centres being averaged by the illumination spot. Since the number of scattering centres goes linearly with the spot width, the fractional intensity would be expected to vary as the inverse of the square root of the width. This is reflected in the simulated results, but the experimental data shows a linear dependence. There are a number of complicating effects which could account for the deviation, for example both shadowing and multiple-scattering will change the effective number of scattering centres in a region of the surface.

Despite these shortcomings, the simulated fractional intensity results are encouraging. From basic physical considerations, the signal's autocorrelation width would be expected to grow along something

like $y=x$ with the aperture width and reproducing that result is only of limited utility. Conversely, there was no *a priori* reason to suspect that bistatic two-scale scattering would amount to 1% of the measured intensity, but these first pass results from the numerical simulation predict a fractional intensity of the same order as LSA. The error between simulated and experimental results is generally less than a factor of two, which suggests the model still needs some adjustment but is broadly suitable.

Finally, it is also worth noting that the experimental data used above was generated with an LED shone through an aperture. An LSA scan taken on 120-grit ground glass with the standard ILS2100 sensor head has an autocorrelation width of roughly 110 μm , which suggests that a simulated width of 220 μm is a good fit to the actual spot size. This is larger than the reported tightness of LSA's focus, but is again reasonable given the underestimation of the tail in the simulation's intensity profile.

6.4.3 Detector vs. Detector Correlation

One of the unexplained features of LSA is the degree of correlation between the two opposite detectors in the same channel. Performing a cross correlation between the signal from detector one and the signal from detector two tends to produce one of two results. For some surfaces there is a reasonably strong correlation, indicated by a positive spike at zero. For other surfaces, however, there is a strong anti-correlation, or a negative spike at zero. Explaining this behaviour is one of the triumphs of the model. It is particularly satisfying as a result because, unlike scan characteristics and positional tolerance, it was not one of the aims of the simulation. Rather, detector-to-detector correlation emerges as a result of the basic model.

Experimental examples of positive and negative correlations are shown in Figure 6.12. The two plotted correlation functions are the average of the detector-detector correlations from four different scans of each target surface. This helps to suppress the random fluctuations in the wings, and makes the spikes more noticeable.




Figure 6.12 The average normalised cross correlation between opposite detectors from the same channel, from a 15 mm LSA scan of paper (red) and 220-grit ground glass (blue).

Not only does the model correctly replicate this behaviour, it provides the explanation through the simulation of LSA from surfaces with different correlation lengths. Opposite detector signals can be generated from this model by simulating two scans over the same position on the same surface, one with the detector at $P = (7,0,7)$ mm and one with the detector at $P = (-7,0,7)$. In all cases the illumination geometry is the same as described above with a 200 μm spot width.

Four different instances of a rough surface are generated for each of four different surface correlation lengths: 10 μm , 24 μm , 37 μm , and 137 μm . All 16 surfaces have a 400 nm RMS height deviation. Left and right detector signals are calculated for each surface, and the cross correlations averaged for each of the surface types. The resulting correlations are plotted together in Figure 6.13.




Figure 6.13 The average normalised cross correlation between simulations of detectors 1 and 2 from surfaces with correlation lengths between $10\text{ }\mu\text{m}$ and $137\text{ }\mu\text{m}$. The illumination spot was $220\text{ }\mu\text{m}$ wide.

From the simulated data, it is apparent that opposite detectors correlate only when the surface's correlation length is much less than the laser spot width. As the correlation length approaches the spot width, an anti-correlation begins to develop. When the features are large enough that the laser spot only illuminates one feature (in the x dimension), then there is complete anti-correlation.

It is straight forward to see why anti-correlation occurs when the bulk of the scattered intensity is directed by just a single feature (in the x dimension): the specular reflection goes primarily in one direction, towards one detector and away from another. As more scattering centres are introduced and many different specular peaks average together, scattered intensity is directed in both directions and a single feature does not dominate the direction of the signal. This explains why the anti-correlation should be reduced, as the effects of different features will, in part, negate the effects of others.

To explain the positive correlation, it is necessary to look at the nature of the surface. As described in Section 6.3.2, the surfaces generated in this simulation are more than randomly-oriented planar facets. The facets together describe a smooth Gaussian surface with zero mean. Such a surface has some inherent symmetries: considering a "feature" as a hill or valley, then each feature includes both a rising and falling slope. Over the extent of a zero-mean surface, rises and falls (which correspond to increased intensity at one detector and decreased intensity at the other) do not simply average each other out, but tend to come in pairs. When an entire hill is illuminated, one detector sees an increase in intensity from one slope, while the other detector sees an increase in intensity from the opposite slope.

Averaged over many features, this produces partial correlation in the intensity signals. Thus it is the partial symmetry of the surface itself which generates the positive correlation.

This is confirmed in Figure 6.14, which shows the simulated detector-detector correlation for two different surfaces. The first is generated exactly as described above, such that the planar facets form a Gaussian surface with 10 μm correlation length and a 400 nm RMS height. The surface described as “Random Facets” is exactly that: a collection of randomly generated surface normal vectors, distributed over the approximately the same angular distribution ($\pm 4^\circ$ from vertical) and with the same facet size and illumination conditions as the Gaussian surface.




Figure 6.14 Positive correlation only occurs when a large number of scattering centres contribute to the final intensity, and the scattering centres are described by a smooth distribution (e.g. Gaussian). When the surface is described by randomly oriented facets without physical relationship to each other, no positive correlation develops.

Even when the facets are small enough to average out any noticeable anti-correlation, no positive correlation develops without the inherent symmetry of a real surface.

6.4.4 Simulated Tolerances to Misplacement

The final test of the model is against the tolerance data collected in Section 4.3.

In all of the following cases the home scan⁹ was generated from the central 15 mm, with the spot moved in 10 μm steps. The home scan was always taken with normally-incident illumination and a 45° observation angle. Based on the results from Section 6.4.2, the illumination width was set to 220 μm .

⁹ The initial scan taken with zero offset.

The exact spot length was also tuned empirically. The acceptance angle of the detector translates to approximately ± 1.5 mm of the laser line. The acceptance angle (as reported by the detector's datasheet) has heavier tails than the extended Gaussian amplitude envelope, which suggests a 3 mm long spot might be a slight underestimation. To first order, transverse and yaw tolerances can be considered a function of how much of the offset region of illumination overlaps with the original region of illumination. Considered in this way, it is clear that they will be directly impacted by the length of the laser line; transverse tolerance will improve as the laser line elongates while yaw tolerance will grow poorer.

To find the best laser length, yaw and transverse offsets were simulated (using the procedure described in 6.3.6) for different line lengths between 3 mm and 4.8 mm. To ensure that the illumination size was the only factor, the three yaw runs were simulated using the same surface realisation, as were the three transverse simulations. The dependence of transverse tolerance on laser line length is shown in Figure 6.15, and yaw dependence in Figure 6.16.




Figure 6.15 Three runs of transverse offset were calculated for a surface, using a $220\ \mu\text{m}$ by 3.0 mm, 3.7 mm, or 4.8 mm laser spot.




Figure 6.16 Three runs of yaw offset were calculated for a surface, using a $220\text{ }\mu\text{m}$ by 3.0 mm, 3.7 mm, or 4.8 mm laser spot.

The simulated data does not undergo the processing normally used for an LSA scan (see Section 2.7), and simulated scans are compared using a normalised cross correlation rather than with the optimised comparison routine leading to a BMR (see Section 2.8). To compare against experimental data, the data from Section 4.3 had to be reprocessed. The intensity data in space from the scans of 120-grit glass were extracted before AC-processing. The data were flattened by subtracting from each scan its linear best fit, and the normalised cross correlation was calculated between the flattened data from the offset scan and the home scan. Using this reprocessed data as the standard, a 3.7 mm line was found to be the best balance between matching to actual yaw and transverse data.

Yaw, transverse, pitch and roll offsets were repeated on five different surfaces for each axis, using a 3.7 mm x $200\text{ }\mu\text{m}$ illumination area. Each surface was a different realisation of a Gaussian random surface with a standard deviation in height of $3\text{ }\mu\text{m}$ and a correlation length of $40\text{ }\mu\text{m}$. For yaw, pitch and roll simulations, 15.6 mm (x) by 6.4 mm (y) surfaces were generated. For transverse simulations, the surfaces were 15.6 mm (x) by 20.4 mm (y). The transverse offset was simulated from -7 mm to 7 mm, pitch from -4° to 4° , and yaw and roll from -5° to 5° . Transverse was simulated in $250\text{ }\mu\text{m}$ steps, and rotational offsets were simulated in 0.250° steps.

The solid lines in Figure 6.17 show the average of the five simulated runs, while dotted lines mark the minimum and maximum envelope from the simulations of each axis. The points correspond to the experimental 120-grit data from Figure 4.25 processed in the simplified manner described above. Experimental data has a maximum NCC of about 0.98, while simulated data correlates perfectly (1.0) at zero offset. To simulate the small amount of noise present in the apparatus, simulated results were uniformly reduced by 0.02.




Figure 6.17 Solid lines show the average and dotted lines the minimum and maximum from five simulated runs. Data points are the four experimental runs from 120-grit ground glass from Section 4.3. The vertical scale is the same on the four main graphs. Insets show a closer view of the outermost two degrees of pitch and roll offsets.

Simulated yaw and transverse results both show excellent agreement with experimental results, with substantial overlap between the simulated envelope and the measured tolerance data. Pitch data are in broad agreement (within 5% NCC) with experimental values, but are not as closely matched. Viewed full scale, the pitch and roll results appear quite similar to each other. The insets make it clear that, in fact, roll is considerably less accurate. The simulation predicts a fall of around 3% from 4° of pitch, while the experimental loss is close to 8%. Experimental data shows a similar loss (~6%) from 5° of roll, while the simulation predicted just a 0.3% decline in correlation at this offset.

The disagreement is likely because the simulation assumes uniform normally-incident illumination. This simplifying approximation breaks down quickly above and below the focal plane of the laser. Changes in height, such as are introduced by pitch and roll, will produce both a change in the spot size and a non-uniform change in the incident wave vector across the width of the spot. Neither of these effects are

accounted for at present. It is significant that the least accurate axis—tolerance to roll—involves the largest changes in height.

6.5 Outlook for the Numerical Simulation

A number of avenues are open for improving the results for tolerances, fractional intensity and autocorrelation width (Figure 6.11). In several places, the inaccuracy of the amplitude envelope was a cause of error. This is because the envelop, in its current form, fills three different functions: it models the focused amplitude envelope of the laser spot (x); it models the aperture-based amplitude envelope of the laser spot (y); and it models the impact of the detector's acceptance angle. The extended Gaussian profile was used as the best compromise between these three components. The accuracy of each of these three profiles could be improved by implementing them separately. The amplitude envelope, $A(x, y)$, would become the product of two independent profiles, $A_x(x) \cdot A_y(y)$. The largest gain would probably come from implementing the detector's acceptance angle as an explicit function of ρ , although this would come at increased computational cost.

The Lambertian dependence used currently is only a specific, uniform form of the bidirectional reflectance distribution function (BRDF). It was implemented as one of the simplest approximations available and it is ripe for replacement. For a minimum of effort, the fixed cosine relationship contained in ρ in equation (6.24) could be replaced by the BRDF for the specific material from one of the publicly-available databases, such as the one provided by the Mitsubishi Electronic Research Laboratory[92].

The effect of noise is to reduce BMR or the NCC, which was done directly for the tolerance data presented earlier. This is acceptable here, with a single sample type and experimental data to show the magnitude of the drop. A more rigorous approach would be to introduce a couple of artificial noise terms, ideally a flat detector noise plus a second term that scales with the DC intensity. This would correctly show the increased impact noise has on both low S/N surfaces and on surfaces with weak diffuse scattering.

Like many early scattering models, both multiple scattering and shadowing[93] are neglected here, and could be introduced. Any improvement would have to be weighed against the additional computational complexity.

Finally, introducing non-uniform incident wave vectors would likely improve pitch and roll as explained previously. In addition this would allow the model to simulate the fifth and final axis, focus.

6.6 Predictions of the Simulation for LSA

Possible improvements notwithstanding, enough confidence has been developed in the model to begin looking at what predictions it might make for the behaviour of an LSA system. It is mechanically simplest to keep the illuminating optics normal to the target surface. The detector position would be easiest to change and so that is the focus of the following simulations.

Four detector positions are considered here: the standard position in the x-z plane (45°), near-normal and near-grazing positions (5° and 85°) in the x-z plane, and a right-angle detector at 45° in the y-z plane. Four scans are simulated from a surface with $3\text{ }\mu\text{m}$ RMS height and $40\text{ }\mu\text{m}$ correlation length, for detectors at each of these positions, and the results are shown in Figure 6.18.




Figure 6.18 Simulated scans from a surface like 120-grit ground glass, for four different detector positions. 45° is the usual LSA detector position in the x-z plane; 85° and 5° are near-grazing and near-normal in the x-z plane, and the Right-angle detector is at 45° in the y-z plane.

From each scan, the standard deviation and mean are calculated. The variations in intensity are all within 35% of the 45° detector, while the DC background levels change faster. The result is a significantly increased fractional intensity for near-grazing detection. These values are summarised in Table 6.1.

	85°	45°	5°	RA
Std.Dev.	7.48 E-6	1.02 E-5	1.38 E-5	7.17 E-6
Mean	5.60E-4	1.59 E-3	2.21 E-3	1.65 E-3
$dI/\bar{I}$	0.013	0.0064	0.0062	0.0043

Table 6.1 Standard deviation, DC mean, and fractional intensity for simulated scans with four different observation angles.

An increase in signal strength would benefit the LSA system, but only if it was not at the cost of increased sensitivity to misplacement. To predict the impact of different observation positions on tolerance, pitch offset simulations were performed with the detector in each of the three new positions. These results are plotted in Figure 6.19, along with the simulated pitch tolerance envelope for a detector at 45° (from Figure 6.17). This model predicts that near-grazing detection would be no more sensitive to rotation about the y-axis than is traditional 45° detection, and that near-normal and right-angle detection would both improve tolerance to pitch.




Figure 6.19 Simulated tolerance to misplacements in pitch for a surface like 120-grit ground glass, with different detector positions. The 45° position is the usual LSA configuration, and was simulated five times. Maximum and minimum values are indicated with dotted lines.

The final consideration for system performance is the impact of noise, as described in the previous section. Any constant noise terms become more significant as the overall intensity of the signal decreases. Smoother surfaces already have a lower DC level and it may be that the additional reduction from near-grazing detection is too much, and brings already weak singles below the noise level.

Accurate noise terms would need to be introduced into the simulation to make a final determination, along with simulations of the other three axes.

6.7 Conclusions

This chapter has developed a physical model for Laser Surface Authentication, based on diffuse scattering from planar elements of a surface with no phase contribution. The model built on existing work in scattering and computer vision to develop an expression for the bidirectional scattering from a stochastic rough surface.

A numerical implementation of this model was constructed, which allowed the scattered intensity to be calculated in a geometry which mirrored LSA: normal incidence, an elongated focal spot and observation at 45° . By introducing a relative offset in x between the target surface and the illumination and detection vectors entire LSA scans could be simulated, and by introducing linear and rotational offsets about other axes the tolerance data measured in a previous chapter could be simulated.

By adjusting the dimensions of the illuminating spot, the fractional intensity and correlation length of simulated scans could be brought into good agreement with their experimental counterparts. Detector-to-detector correlation emerged as an unexpected consequence of the model, and was found to depend upon the ratio between the in-plane roughness length scale and the width of the illuminating spot. Predicted tolerance to yaw and transverse offsets matched measured values very well. Predicted pitch and roll degradation matched experimental values to within 10% NCC.

Section 6.5 discussed several avenues for improving the current model. Most important was modelling the focal condition of the incident illumination. Finally, the last section began to look at what the model could tell future LSA development about the position of the detectors in space.

7

CONCLUSIONS AND OUTLOOK

*“And there are seven kings: five are fallen, and one is,
and the other is not yet come; and when he cometh, he
must continue a short space.”*

—Revelation 17:10

7 Conclusions and Outlook

7.1 Summary of Key Findings

LSA is well placed as both an anti-forgery technology and a track and trace solution. It can be a disruptive technology able to address these growing problems in shrink, diversion, black market trade, quality assurance, supply-chain management, brand protection, and counterfeit detection. An improved understanding of the technique and its underlying mechanism is essential in developing the system, planning deployments, and preventing spoofing. This thesis aimed to accomplish two things to further that understanding: to evaluate and characterise the performance of LSA as a system; and to discover the source and nature of the optical effect.

Chapter 3 has clearly demonstrated that LSA can perform in operational environments, outside the protective confines of the lab. The results of a field trial indicated that four different products, enrolled as part of live production, could be authenticated with a false accept rate of at least $1:10^{11}$ and as much as $1:10^{19}$. The field trial also exposed several weaknesses in the technology, namely a strong positional dependency in the optical geometry and a slight weakness in the statistical description. The trial also uncovered a strong material dependency.

Laboratory testing in Section 4.2 of controlled misplacement of the LSA sensor or the target item confirmed the material dependency uncovered in the field. Eighteen different real-world samples showed a wide range of tolerances. Depending upon the material, LSA signals survived over as much as 10 mm or as little as 2 mm of height offset, as much as 4.4 mm or as little as 0.4 mm of transverse offset, between 11° and 2° of yaw, 10° and 1.25° of pitch, and 10° and 0.25° of roll.

In Section 4.3, ground glass samples were used in conjunction with a new experimental rig to provide samples with different correlation lengths and RMS heights but the same statistical description. Testing on these samples uncovered one of the key findings of this dissertation: the feature size is not a determining factor in the sensitivity of the surface, meaning that that material dependency is attributable only to the higher-order statistical description of the surface. For surfaces with a Gaussian distribution of heights, BMR fell by 15% per mm of transverse misplacement, 10% per mm of height, 5% per degree of yaw, 3% per degree of pitch and 2% per degree of roll.

A phenomenological description of the match between scans was proposed in Section 4.4, of the form

$$BMR = M(\sigma_h, l_c, N) - \Phi(x_{tran}, x_{focus}, \theta_{yaw}, \theta_{pitch}, \theta_{roll}) \quad (7.1)$$

The basic LSA geometry was simplified to a laser and CCD detector in Section 5.2, and linear and rotational offsets were once again applied. An examination of the resulting speckle patterns, confirmed by a simple Huygens Wavelet simulation (Section 5.3), showed that under the LSA geometry laser speckle decorrelates about an order of magnitude more quickly than does LSA to out-of-plane rotations. Even under ideal conditions speckle would only be expected to survive one fifth the offset of LSA.

An analysis of the expected and measured signal from integrated laser speckle in Section 5.4 found that the signal strength is an order of magnitude weaker than that measured by LSA.

These findings motivated two final experiments to investigate the effects of eliminating interference phenomena from the LSA signal. First, a study with sources of different degrees of coherence in Section 5.5 found that signals of the nature and strength of LSA can be measured regardless of the nature of the light source. Illumination from a red LED and from a tungsten thermal source both produced matches of 99.9% BMR in laboratory conditions. Scans taken with different sources agreed with each other by as much as 75% BMR, despite poor alignment.

Further experiments with different wavelengths of illumination in Section 5.6 concluded that the influence of phase on the LSA signal was minimal. Using repeatable alignment and changing the wavelength of incident light from 671 nm to 532 nm or 808 nm reduced BMR by no more than 4%. Conversely it was shown that these changes in wavelength were sufficient to cause speckle to decorrelate completely.

The tolerance and fractional intensity of speckle, the role of phase in the LSA signal, and the impact of different wavelengths of illumination together form the second major result of this study: the variations in intensity measured in LSA are not the result of spatially-averaged laser speckle.

Finally, Chapter 6 saw a model proposed and developed numerically that describes the diffusely scattered light measured by LSA as a result of geometrical optics and the interaction of multiple scales of surface roughness. Simulations showed that the model accounts for both the signal strength and some aspects of the tolerance to misplacement for Gaussian surfaces. In evaluating the simulation against experimental results, tolerances to offsets in height and transverse agreed very well, while simulated tolerance to pitch and roll were underestimated by 5 – 10%.

Importantly, the model established that the previously unexplained detector-to-detector correlation depends upon the symmetry and correlation length of the target surface.

7.2 Outstanding Questions

From the chapter on field work, it is evident that larger databases are necessary. Larger enrolment sample sets are always useful, to further-define the statistics and add confidence when extrapolating out to the extreme edges of the distribution. In particular, though, larger verification data sets will be critical when moving forward. This has been difficult, firstly because enrolment scans vastly outnumber verification scans in most applications and secondly because the unlike-item dataset goes like the square of the number of items scanned, while the like-item dataset grows linearly with the number of items scanned. For these two reasons there are already suitably-large enrolment databases but only a handful of trials where verification scans number in the hundreds. Verification sets at least that large will be necessary to accurately calculate metrics that rely on the statistics of the like-item distribution: specifically d' and the False Reject Rate (FRR). FRR is the complement of the FAR, and requires extrapolating the tail of the like-item distribution. The next step would be to have a verification dataset large enough to create a high-resolution distribution like Figure 3.11 for the like-item distribution. This would allow a test of the proposed like-item statistics.

Chapter 4 was able to rule out feature size as a determining factor for the tolerance of surfaces with simply-described statistics. It could not, however, positively demonstrate that the nature of the surface was the determining factor. A dedicated microscopy study would be a useful addition to the data collected in the first half of that chapter. Simply being able to classify all 18 materials based on height, correlation length, and statistical description would add additional dimensions on which to sort the data, and could prove illuminating. Controlled surfaces, similar to what was undertaken in the second half of the chapter, could also be useful. If surfaces could be manufacture with common RMS height and correlation length but different height distribution (Gaussian, exponential, etc.) and different natures (with and without fissures, for example) then components of the complicated real-world surface descriptions could be tested individually.

This kind of work could also be paralleled in the modelling. In addition to the areas identified for improvement at the end of Chapter 6, surfaces with different statistical descriptions could be simulated. Surface generation is a separate process in the simulation, so it will be possible to introduce non-Gaussian surfaces to the model without substantial changes to the core algorithm.

Observation position is only one of a number of free parameters in the model that can now be explored with the aim of improving LSA. With more computation time the effects on signal strength and tolerance of beam width, incidence angle, observation distance and detector size could each be investigated. The results would point the way towards further improvements to LSA sensor geometry.

Finally, two more fundamental questions deserve further consideration. First, Chapter 6 posits that existing phase-dependent scattering models should—if averaged together over a series of bandwidths to produce a broad-band source—be able to generate an LSA signal. With sufficient computational time, this proposition could be tested. Second, more attention should be devoted to the suppositions about speckle as the source of “supersensitivity” to misplacement. If there is indeed a crossover point, where the fractional intensity due to LSA falls far enough that it is comparable to, or less than, the fractional intensity due to speckle, this should be confirmed experimentally. A study which took simultaneous measurements with an integrating photodetector and a CCD element with auto gain compensation from a variety of surfaces would allow a side-by-side comparison of the relative signal strengths. Alternately, reproducing the fibre-couple laser experiments on a few of the real-world samples could concretely establish a link between surfaces exhibiting “supersensitivity” and surfaces for which phase *does* play a role in the measured signal.

7.3 Suggested Improvements to the LSA System

In addition to the alternate detector positions simulated in Section 6.6, the work performed in this study suggests an important improvement to the base LSA hardware as described in Sections 2.3 and 2.4. Chapters 5 and 6, taken together, conclusively showed that a coherent illumination source is not necessary in an LSA sensor. But section 5.5.4 suggests this should go further: a coherent source is not simply unnecessary, it is actually a hindrance to measuring good LSA. Those materials from Section 4.2 which exhibited supersensitivity did so because the LSA was weak enough that the integrated speckle grains formed a significant portion of the measured intensity pattern. Taking the level to which the BMR has fallen after a single 0.250° step as a rough guide, speckle formed between 15% and 40% of the measured intensity profile for these surfaces.

Far from taking advantage of laser speckle, LSA has joined the myriad of metrologies[94-96] for which speckle is nothing but a source of noise. Figure 4.15 showed that the materials which exhibit supersensitivity also have some of the weakest fractional intensities. This is consistent with the LSA needing to be weak before speckle becomes significant, but it also means that these materials are already the most difficult to make work in a deployment, because they are already the most susceptible to environmental noise. Essentially, laser speckle is making a bad case worse, and eliminating it would go a long way towards improving the performance of LSA on these troublesome surfaces. Although many techniques have been suggested to minimise speckle[97-99], they are all appropriate for cases where the coherence is necessary for other reasons. Given that a coherence illumination system adds nothing but noise to an LSA device, the simplest approach would be to replace the laser diode source with an LED source.

In switching to an LED source it is possible to maintain comparable levels of intensity. An aperture severely restricts the available light, so the most intensity will be achieved with a structured element and that geometry will be the basis of this rough intensity calculation. Assuming that the system should maintain something like its current size, a reasonable lens would be a 12.5 mm diameter lens with a 10 mm focal distance. If this lens is positioned to directly image the structured LED (i.e. set at a working distance of 20 mm) and only the central 10 mm of the lens are used (to partially account for edge effects) then the lens will collect light from 0.19 sr. Output power for the PRP structured LED was not available, but a typical¹⁰ high-powered red LED has an advertised radiant intensity of 500 mW/sr from a semiconductor element approximately 1 mm². Rearranging the same semiconductor into a 5 mm long pattern of 125 elements as in Section 5.5.2 would use about 0.04 mm² of active area and theoretically produce 20 mW/sr, of which the lens could collect and focus ~4mW onto a 5 mm line. This is slightly more than would be available from a 3 mW laser, assuming all of the laser diode's light was collected and focused. The caveat is that a high powered LED array would take far more electrical power to generate the same level of optical power. The current drawn by PRP's structured LED would be unsustainable for a portable battery-powered unit.

From a technical standpoint, a superluminescent diode[100] (SLD) is the best alternative, as it would provide a similar intensity and power to a laser diode without the coherence. SLDs are not yet as mature a product as LEDs and costs are considerably higher at the moment: SLD modules cost hundreds of USD, versus tens of USD for laser diode modules and less than 1 USD for LEDs. Like all things, the added cost would have to be balanced against the improvement in speckle noise over a laser diode or the improvement in intensity over an LED source.

Finally, both SLD and LED sources tend to be more stable and less noisy than comparable laser diode sources[101]. Replacing the laser with an LED/SLD would not only remove speckle as a specific source of noise, it would also reduce the source noise of the system as a whole. These two effects combined could make a significant improvement on the performance of surface currently considered only marginal candidates for LSA.

¹⁰ Hewlett Packard HLMP-NG07

Glossary of Terms

5-axis Rig – The experimental apparatus which used precision linear, rotary and cradle stages to move the LSA sensor block through the scan motion and offset the sensor block from the target in the transverse, height, yaw, pitch and roll.

7-axis Rig – The experimental apparatus which used a robotic arm to move the LSA sensor block through the scan motion and offset the sensor block from the target in transverse, height, yaw, yaw tracking, pitch, pitch tracking and roll.

BMR – Bit Match Rate. The result of comparing two LSA signatures expressed as the ration of bits which agree to the total number of bits compared

ccdf – complementary cumulative distribution function. For random variable X with pdf $f(x)$ the ccdf $G(x)$ gives the probability that the random variable is greater than x , such that

$$P(X > x) \equiv G(x) = \int_x^{\infty} f(t)dt$$

cdf – cumulative distribution function. For random variable X with pdf $f(x)$ the cdf $F(x)$ gives the probability that the random variable is less than or equal to x , such that[102]

$$P(X \leq x) \equiv F(x) = \int_{-\infty}^x f(t)dt$$

Contrast – the ratio of the intensity's standard deviation to its mean for a speckle pattern; equivalent to the image's fractional intensity.

CrossComm – a dynamic link library supplied by Kuka Robotics which allows programmatic communication with the VxWorks layer of the KRC from the Windows XP layer, outside of the HMI.

Diversion – the practice of legally purchasing goods in one market for illegal resale in a different market. Colloquially referred to as grey market. When the resale in a different market is legal, this is referred to as parallel trade.

FAR – False Accept Rate. The rate at which false positives occur in a biometric system (ie the rate at which a challenge reports the item is genuine when it is not).

Fractional intensity – the ratio of the standard deviation to the mean for a 1D profile. For LSA this specifically refers to the standard deviation of the AC-coupled data over the mean of the DC-coupled data.

FWHM – full width at half maximum. The width of a central peak, measured from the two points where it is exactly half of its maximum value.

HMI – Human Machine Interface. The Windows XP program provided by Kuka Robotics which runs on the KRC and gives the user direct control of the robotic arm.

Home Image/Position/Scan – When collecting data across a range of positions, the home scan/image is the initial one taken, which defines zero offset and is used the reference against which subsequent data is compared. The home position is synonymous with zero offset on all axes.

KRC – Kuka Robotics Controller. The rack containing a VxWorks PC, a Windows XP PC, and the drive circuitry for the Kuka robot.

KRL – Kuka Robotics Language. The programming language used by Kuka Robotics to provide logic flow, digital I/O, and motion control for their robotic arms. It executes within the HMI, running on the KRC.

pdf – probability density function. For random variable X the probability density function $f(x)$ describes the probability that the random variable will fall within a given range by[102]

$$P(a < X \leq b) \equiv \int_a^b f(t)dt$$

Pseudo-scan – An LSA-like scan generated by calculating the scattered intensity by summing the pixels of a CCD image, rather than by measuring an integrating photodetector.

RBM – Robot Bridge Master. The IsaControl-side of the TCP/IP link between the IsaControl PC and the KRC's Windows XP PC. Runs embedded within a driver in IsaControl.

RBS – Robot Bridge Slave. The KRC-side of the TCP/IP link between the IsaControl PC and the KRC's Windows XP PC. Runs as a stand-alone program in the Windows XP environment, and uses CrossComm to communicate with the VxWorks environment and read and write KRL variables.

SL – Security Level. The negative exponent of the probability of a particular match occurring by chance.

Track and trace –The practise of creating an audit trail for pallets, cartons, or individual items from manufacture through to point-of-sale.

References

1. Kirovski, Darko, *Anti-Counterfeiting: Mixing the Physical and the Digital World*, in *Secure Component and System Identification*, Christof Paar, Editor. 2008: Berlin.
2. McGlade, Mark, *If You Can Make It, They Can Fake It*, M.B.A. Thesis, University of Surrey School of Management, 2008
3. Keller, K. L., *Brand synthesis: The multidimensionality of brand knowledge*. The Journal of consumer research, 2003. **29**(4): p. 595.
4. Noble, Ronald K., *The links between intellectual property crime and terrorist financing*, United States Congress, Committee on International Relations, July 16, Washington D.C., 2003
5. H.R. 2281, *Digital Millenium Copyright Act*, United States House of Representatives, 1998
6. Buchanan, James D. R., et al., *Forgery: 'Fingerprinting' documents and packaging*. Nature, 2005. **436**(7050): p. 475-475.
7. Schechter, Stuart, *Quantitatively Differentiating System Security*, in *Workshop on Economics and Information Security*, Hal Varian and Ross Anderson, Editors. 2002, University of California: Berkeley, CA.
8. Lehtonen, Mikko, *Anti-counterfeiting Business Case Report*, in *Building Radio Frequency Identification for the Global Environment*. 2007, ETH: Zurich.
9. H.R. 3580, *Food and Drug Administration Amendments Act of 2007*, United States House of Representatives, 2007
10. *Anti-contraband and Anti-counterfeit Agreement and General Release*, Philip Morris International and The European Commission, European Anti-Fraud Office, Appendix D, 2004
11. Beckmann, P. and A. Spizzichino, *The Scattering of Electromagnetic Waves from Rough Surfaces*. 1963, New York: Pergamon.
12. Liszka, E. G., *Scattering at a Rough Boundary - Extensions of the Kirchhoff Approximation*. The Journal of the Acoustical Society of America, 1982. **71**(5): p. 1093-1100.
13. Oglivy, J A, *Theory of Wave Scattering From Random Rough Surfaces*. 1991, Bristol: Adam Hilger.
14. Francon, M., *Laser Speckle and Applications in Optics*. 1979: Academic Press.
15. May, Marie, *Information inferred from the observation of speckles*. Journal of Physics E: Scientific Instruments, 1977. **10**: p. 849-864.
16. Weinrib, Abel and B. I. Halperin, *Distribution of maxima, minima, and saddle points of the intensity of laser speckle patterns*. Physical Review B, 1982. **26**(3): p. 1362-1368.
17. Yamaguchi, I. Ichirou, *Speckle Displacement and Decorrelation in the Diffraction and Image Fields for Small Object Deformation*. Journal of modern optics, 1981. **28**(10): p. 1359-1376.
18. Goodman, Joseph W., *Speckle Phenomena in Optics: Theory and Applications*. 2007, Greenwood Village, CO: Roberts and Company. 387.
19. Nieto-Vesperinas, M. and J. M. Sotoo-Crespo, *Monte Carlo simulations for scattering of electromagnetic waves from perfectly conductive random rough surfaces*. Optics Letters, 1987. **12**(12): p. 979-981.
20. Elfouhaily, T. M., *A critical survey of approximate scattering wave theories from random rough surfaces*. Waves in random and complex media, 2004. **14**(4): p. 1.
21. Warnick, K. Karl, *Numerical simulation methods for rough surface scattering*. Waves in random media, 2001. **11**(1): p. R1.
22. Ragheb, Hossein and Edwin R. Hancock, *The modified Beckmann-Kirchhoff scattering theory for rough surface analysis*. Pattern Recognition, 2007. **40**(7): p. 2004-2020.
23. Sjodahl, M. and L. Larsson, *Monitoring microstructural material change in paper through microscopic speckle correlation rate measurements*. Optics and Lasers in Engineering, 2004. **42**: p. 193-201.
24. Indeck, Ronald S. and Elias Galvinas, *Fingerprinting Magnetic Media*. IEEE Transactions on magnetics, 1993. **29**(6): p. 4095-4097.
25. Pappu, Ravikanth, et al., *Physical One-Way Functions*. Science, 2002. **297**: p. 2026 - 2030.

26. Haist, T., *Optical detection of random features for high security applications*. Optics communications, 1998. **147**(1-3): p. 173.
27. Tuyls, Pim, Boris Skoric, and Tom Kevenaar, *Security with Noisy Data*. 2007, London: Springer-Verlag.
28. Chen, Yuqun, M. Kivanc Mihcak, and Darko Kirovski, *Certifying authenticity via fiber-infused paper*. SIGecom Exch., 2005. **5**(3): p. 29-37.
29. Tuyls, Pim, et al., *Read-Proof Hardware from Protective Coatings*, in *Cryptographic Hardware and Embedded Systems - CHES 2006*. 2006. p. 369-383.
30. Schmid, Natalia A. and Joseph A. O'Sullivan, *Performance Prediction Methodology for Biometric Systems Using a Large Deviations Approach*. IEEE Transactions on Signal Processing, 2004. **52**(10): p. 3036-3045.
31. Daugman, John, *The importance of being random: statistical principles of iris recognition*. Pattern Recognition, 2003. **36**: p. 279-291.
32. Heald, Mark A. and Jerry B. Marion, *Classical Electromagnetic Radiation*. Third ed. 1995: Thomson Learning.
33. Meecham, W. C., *On the Use of the Kirchoff Approximation for the Solution of Reflection Problems*. J. Ration. Mech. Anal., 1956. **5**: p. 323-333.
34. Syrratt, RW, *Angular correlation of speckle patterns*, Ph.D. and D.I.C. Thesis, Imperial College, London Physics, 1994
35. Oren, Michael and Shree K. Nayar, *Generalization of the Lambertian model and implications for machine vision*. International Journal of Computer Vision, 1995. **14**(3): p. 227-251.
36. Wolff, Lawrence, Shree Nayar, and Michael Oren, *Improved diffuse reflection models for computer vision*. International Journal of Computer Vision, 1998. **30**(1): p. 55.
37. Vernold, Cynthia L. and James E. Harvey. *Modified Beckmann-Kirchoff scattering theory for nonparaxial angles*. in *Scattering and Surface Roughness II*. 1998. San Diego, CA, USA: SPIE.
38. Ruiz-Cortes, V. A., *Experimental light-scattering measurements from large-scale composite randomly rough surfaces*. Journal of the Optical Society of America. A, Optics and image science, 2002. **19**(10): p. 2043-2052.
39. Goodman, J. W., *Statistical Properties of Laser Speckle Patterns*, in *Laser Speckle and Related Phenomena*, J. C. Dainty, Editor. 1975, Springer-Verlag: Berlin Heidelberg New York. p. 9.
40. Fujii, H and T. Asakura, Opt. Commun., 1974. **12**: p. 32.
41. Leonard, Lisa C. and Vincent Toal, *Roughness measurement of metallic surfaces based on the laser speckle contrast method*. Optics and Lasers in Engineering, 1998. **30**: p. 433-440.
42. Léger, D., *Optical surface roughness determination using speckle correlation technique*. Applied Optics, 1975. **14**(4): p. 872.
43. Sjodahl, M. M., *Speckle correlation used for measuring microstructural changes in paper*. Proceedings of SPIE--the international society for optical engineering, 2003. **4933**: p. 33-8.
44. Cowburn, Russell, *Laser Surface Authentication - Biometrics for Brand Protection of Goods and Packaging*, in *Smart Packaging Technologies*, Joseph Kerry and Paul Butler, Editors. 2008, John Wiley & Sons: West Sussex. p. 281.
45. Buck, Roger D., US Patent US 2007/0108757 A1, *Security laminate intermediate using distinct infrared light spectrum*, 2005
46. van Renesse, R. L. *Paper based document security-a review*. in *European Conference on Security and Detection*. 1997.
47. Lin, Hai Hui, et al., US Patent US 2008/0087189 A1, *Security Pigments and the Process of Making Thereof*, 2007
48. Green, Alan, Ronald J. Reiss, and Douglas W. Wilson, Patent 5042842, *High Security Label*, 1991
49. Dittmann, J., L. C. Ferri, and C. Vielhauer. *Hologram watermarks for document authentications*. in *Information Technology: Coding and Computing*. 2001.
50. Zaluzec, N. J., et al., *Intrinsic TLI Surface Tag Directly Authenticated by a SEM*, in *General, Miscellaneous, and Progress Reports (UC-700)*. 1991, Argonne National Laboratory.

51. Bauder, D. W., *Systems Research Report no. PTK-11990*, in *Systems Research Report*. 1983, Sandia National Laboratories.
52. Kumar, S. S., et al. *The butterfly PUF protecting IP on every FPGA*. 2008. Anaheim, CA: Ieee.
53. Devadas, S., et al. *Design and implementation of PUF-based "Unclonable" RFID ICs for anti-counterfeiting and security applications*. 2008. Las Vegas, NV: IEEE.
54. Voloshynovskiy, Sviatoslav, et al., *Unclonable identification and authentication based on reference list decoding*, in *Secure Component and System Identification*, Christof Paar, Editor. 2008: Berlin.
55. Guajardo, Jorge, *Life Cycle of an FPGA Intrinsic PUF: From Start-Up to Butterfly*, in *Secure Component and System Identification*, Christof Paar, Editor. 2008: Berlin.
56. Kumar, Sandeep S., *Key Distribution for Wireless Sensor Networks and Physical Unclonable Functions*, in *Secure Component and System Identification*, Christof Paar, Editor. 2008: Berlin.
57. Fetter, Steven and Thomas Garwin, *Using Tags to Monitor Numerical Limits in Arms Control Agreements*, in *Technology and Limitation of Internatinoal Conflict*, Barry M. Blechman, Editor. 1989, The Johns Hopkins Foreign Policy Institute: Washington, D.C.
58. Skoric, B., *On the entropy of keys derived from laser speckle; statistical properties of Gabor-transformed speckle*. *Journal of Optics A: Pure and Applied Optics*, 2008. **10**(5): p. 055304.
59. Stinson, Douglas R., *Cryptography Theory and Practice*. 3rd ed. Discrete Mathematics and its Applications, ed. Kenneth H. Rosen. 2006, Boca Raton, FL: Chapman & Hall/CRC.
60. Skoric, B., et al., *Information-theoretic analysis of capacitive physical unclonable functions*. *Journal of Applied Physics*, 2006. **100**(2): p. 024902-11.
61. Daugman, John, *Results from 200 billion iris cross-comparisons*. University of Cambridge Computer Laboratory Technical Report, 2005. **635**.
62. Taylor, John Robert, *An introduction to error analysis*. 1997: University Science Books. 327.
63. Williams, Gerald O. *The Use of d' As a "Decidability" Index*. in *Security Technology, 1996. 30th Annual 1996 International Carnahan Conference*. 1996.
64. Zwillinger, Daniel, *CRC Standard Mathematical Tables and Formulae*. 31 ed. 2003, Boca Raton: Chapman & Hall. 910.
65. *Technology Readiness Assessment (TRA) Deskbook*, US Department of Defense, Deputy Under Secretary of Defense for Science and Technology, 2005
66. *Glossary of Security Documents, Security Features, and other related technical terms*, General Secretariat of the Council of the European Union, Working Party on Frontiers / False Documents, 2008
67. Horcas, I., et al., *WSXM: A software for scanning probe microscopy and a tool for nanotechnology*. *Review of Scientific Instruments*, 2007. **78**(1): p. 013705-8.
68. Kusse, Bruce and Erik Westwig, *Mathematical Physics*. 1998, New York: John Wiley & Sons. 668.
69. Griswold, Martin, Mark Harrison, and David Saltzberg, *Observation of light transmission through randomly rough glass surfaces beyond the critical angle*. *J. Opt. Soc. Am. A*, 2007. **24**(10): p. 3207-3210.
70. Griffiths, Brian, *Manufacturing Surface Technology*. Manufacturing Engineering Modular Series. 2001, New York: Taylor & Francis Books. 237.
71. Thomas, Tom R., ed. *Rough Surfaces*. 2nd ed. 1999, Imperial College Press: London.
72. Guelker, Gerd, Klaus D. Hinsch, and Claudia Hoelscher. *Detection of surface microstructure changes by electronic speckle pattern interferometry*. in *Holographics International '92*. 1993. London, United Kingdom: SPIE.
73. Briers, J. David, Glenn Richards, and Xiao Wei He, *Capillary Blood Flow Monitoring Using Laser Speckle Contrast Analysis (LASCA)*. *Journal of Biomedical Optics*, 1999. **4**(1): p. 164-175.
74. Lu, J. Q., *Angular correlation function of speckle patterns scattered from a one-dimensional rough dielectric film on a glass substrate*. *Applied optics*, 1997. **36**: p. 19.

75. Nieto-Vesperinas, M. and J. A. Sanchez-Gil, *Intensity angular correlations of light multiply scattered from random rough surfaces*. Journal of the Optical Society of America. A, Optics and image science, 1993. **10**(1): p. 150-157.
76. Lewis, J. P., *Fast Normalized Cross-Correlation*. 1995, Industrial Light & Magic.
77. Gabor, D., *The transmission of information by coherent light. III. Speckle noise*. Journal of Physics E: Scientific Instruments, 1975. **8**(4): p. 253-255.
78. Berkovits, Richard, *Sensitivity of the multiple-scattering speckle pattern to the motion of a single scatterer*. Physical Review B, 1991. **43**(10): p. 8638.
79. Yu, Francis T. S., et al., *Submicrometer displacement sensing using inner-product multimode fiber speckle fields*. Appl. Opt., 1993. **32**(25): p. 4685-4689.
80. Sheng, P., *Introduction to Wave Scattering, Localization and Mesoscopic Phenomena*. 2nd ed. Springer Series in Materials Science, ed. R. Hull. 2006, Heidelberg: Springer. 333.
81. Thorsos, E. I., *The Validity of the Kirchhoff Approximation for Rough-surface Scattering Using a Gaussian Roughness Spectrum*. The Journal of the Acoustical Society of America, 1988. **83**(1): p. 78-92.
82. Feng, Shechao, et al., *Correlations and Fluctuations of Coherent Wave Transmission through Disordered Media*. Physical Review Letters, 1988. **61**(7): p. 834.
83. Freund, I., *Memory Effects in Propagation of Optical Waves Through Disordered Media*. Physical review letters, 1988. **61**(20): p. 2328-2331.
84. Gu, Z. H., *High-order correlations from rough-surface scattering*. Applied optics, 2004. **43**(15): p. 3061-3065.
85. Moulton, P. F., *Spectroscopic and laser characteristics of Ti:Al₂O₃*. J. Opt. Soc. Am. B, 1986. **3**(1): p. 125-133.
86. Parry, G., *Speckle Patterns in Partially Coherent Light*, in *Laser Speckle and Related Phenomena*, J. C. Dainty, Editor. 1975, Springer-Verlag: Berlin Heidelberg New York. p. 78.
87. Pardo-Iguzquiza, E., *The Fourier integral method: an efficient spectral method for simulation of random fields*. Mathematical geology, 1993. **25**(2): p. 177.
88. Manninen, A.T., *Multiscalar Surface Roughness and Backscattering*. Progress In Electromagnetics Research, 1997. **16**: p. 175-203.
89. Elton, N. J., *A two-scale roughness model for the gloss of coated paper*. Journal of optics. A, Pure and applied optics, 2008. **10**(8): p. 85002.
90. Koenderink, Jan J., et al., *Bidirectional Reflection Distribution Function of Thoroughly Pitted Surfaces*. Int. J. Comput. Vision, 1999. **31**(2-3): p. 129-144.
91. Jafari, G. R., et al., *Two-scale Kirchhoff theory: comparison of experimental observations with theoretical prediction*. Journal of Statistical Mechanics: Theory and Experiment, 2005. **2005**(04): p. P04013.
92. Matusik, Wojciech, et al., *A Data-Driven Reflectance Model*. ACM Transactions on Graphics, 2003. **22**(3): p. 759-769.
93. Wagner, R. J., *Shadowing of Randomly Rough Surfaces*. The Journal of the Acoustical Society of America, 1967. **41**(1): p. 138-147.
94. Baribeau, R., *Influence of speckle on laser range finders*. Applied optics-OT, 1991. **30**(20): p. 2873.
95. Wagner, R. F., et al., *Statistics of Speckle in Ultrasound B-Scans*. IEEE Transactions on Sonics and Ultrasonics, 1983. **30**(3): p. 156-163.
96. López-Martínez, C., *Polarimetric SAR speckle noise model*. IEEE transactions on geoscience and remote sensing, 2003. **41**(10 Part 1): p. 2232.
97. Achim, A., *Novel Bayesian multiscale method for speckle removal in medical ultrasound images*. IEEE transactions on medical imaging, 2001. **20**(8): p. 772.
98. McKechnie, T. S., *Speckle Reduction*, in *Laser Speckle and Related Phenomena*, J. C. Dainty, Editor. 1975, Springer-Verlag: Berlin Heidelberg New York. p. 123.

99. Fukuda, S., *Suppression of speckle in synthetic aperture radar images using wavelet*. International journal of remote sensing, 1998. **19**(3): p. 507.
100. Alphonse, G. A., *High-power superluminescent diodes*. IEEE journal of quantum electronics, 1988. **24**(12): p. 2454.
101. Joindot, I., *Measurements of relative intensity noise (RIN) in semiconductor lasers*. Journal de physique. III, 1992. **2**: p. 1591.
102. Abramowitz, Milton and Irene A. Stegun, *Handbook of Mathematical Functions*. 9th ed. 1972, Mineola, NY: Dover Publications.

Appendix A – Cradle Stage Mount

This mount positioned samples of at most 1 mm thickness at exactly the point of rotation for the M-BGM50CC cradle stage. The mount attaches to M3 tapped holes on a 36 mm square.




Figure A.0.1 Custom faceplate for the BGM50 cradle stage, to position samples at the point of rotation.

Appendix B – Full Tolerance Data




Figure B.0.1 Plastics, see Section 4.2.3.2




Figure B.0.2 Papers, see Section 4.2.3.3




Figure B.0.3 Boards, see Section 4.2.3.4




Figure B.0.4 Miscellaneous, see Section 4.2.3.5

Appendix C – RBM, RBS, and CrossComm

When IsaControl requests a scan from the custom hardware driver, the RBM passes the request, along with the positional data, to RBS. RBS translates the positional data into a new base coordinate system, overwrites the existing coordinates, and set a command flag in the KRL program. The main KRL routine holds in a loop, waiting to receive a command.

First, the RBM receives a command from IsaControl through the custom driver interface. RBM's role is simply to translate commands from the standard IsaControl API into a simple text-based form suitable for transmission across a network. RBM and RBS establish a link between the two computers through a TCP/IP connection, and RBM passes any commands across this link.

RBS is a piece of custom software which was written to take advantage of a special dynamic link library (DLL) provided by Kuka support engineers. This CrossComm library is a set of functions which pass data between the VxWin environment and the Windows XP environment. Using this library it is possible to read and write to variables being used in the currently executing KRL program, from the Windows environment.

RBS runs on the KRC, in the background behind the HMI. It parses the text command from the RBM and extracts the variables needed by the KRL program. Position data, instructions, and timing and control flags are all declared as variables in the KRL program. Using the CrossComm DLL, RBS write the values from the text command into these variables, where the KRL program can use them. Appendix D details the KRL architecture more fully.

The following variables are declared in the KRL program and are available through CrossComm. WriteOnly properties are written on the IsaControl side, while ReadOnly properties are written by the KRL program.

```
Private WriteOnly Property itlEStop() As Boolean

Private WriteOnly Property itlProceed() As Boolean
Private ReadOnly Property itlAck() As Boolean
Private WriteOnly Property itlTX() As Boolean
Private ReadOnly Property itlRX() As Boolean
Private ReadOnly Property itlError() As Boolean
Private ReadOnly Property itlIdle() As Boolean

Private WriteOnly Property itlTask() As RobotTask

Private Property itlBase() As Pos
Private Property itlScanLength() As Double
Private Property itlAOffset() As Double
```

```

Private Property itlBOffset() As Double
Private Property itlCOffset() As Double

Private ReadOnly Property itlIDNum() As Long
Private Property itlCalib() As String

```

`itlEstop` is a Boolean flag linked to an interrupt in the KRL program, which will halt the robot's movement immediately if it is set to true. `itlProceed` through `itlIdle` are used to synchronise the execution of the two programs and the communication between them. `itlTask` stores an integer value which translates into one of the following specific commands:

```

Public Enum RobotTask As Integer
    Scan = 30
    SmallMoveScan = 35
    MoveCameraTo = 10
    SmallMoveCameraTo = 15
    MoveLSA = 20
    SmallMoveLSA = 25
    Calibrate = 40
    EndProgram = 250
    GoHome = 200
End Enum

```

`itlScanLength` through `itlCOffset` store values which control the position and orientation of movement and scan commands. `itlIDNum` stores a unique code which identifies each instance of the robot scanner, and `itlCalib` stores all of the robot-specific calibration data.

Appendix D – KRL Program

The KRL code is reproduced in full at the end of this appendix, following a brief overview of the program.

The main program executes as a continuous loop, which waits until a command is received and then executes the instruction. The command is written as a particular integer value to `itlTask` by the RBS.

The KRL program checks the command integer during each loop, and executes the appropriate subroutine based on that command (lines 31-66)

Most commands result in moving either a camera or the LSA tool to the specified coordinate. The scan command is slightly more complicated, and involves four distinct moves in order to be safe from collisions. First, the scan head is brought quickly to a position 75 mm above the target surface (line 100), centred over the desired scan location. The scan head is then lowered slowly (line 102) to 3 mm above the target surface and half the scan length back in x. Third, the head is moved forward in x by the desired scan length (line 106), at a fixed velocity while simultaneously the digital output line used for triggering is set high. At the end of the scan motion, the trigger line is set low again and the head is slowly raised (line 109) to the safe position, 75 mm above the surface, centred.

```

1  DEF Main()
2      INTERRUPT DECL 10 when itlEstop==TRUE do RemoteEStop()
3
4      base_data[3].X = 287.7
5      base_data[3].Y = 393
6      base_data[3].Z = 13
7      base_data[3].A = 0
8      base_data[3].B = 0
9      base_data[3].C = 0
10
11     tool_data[3].X = tool_data[4].X
12     tool_data[3].Y = tool_data[4].Y
13     tool_data[3].Z = tool_data[4].Z
14     tool_data[3].A = tool_data[4].A
15     tool_data[3].B = tool_data[4].B
16     tool_data[3].C = tool_data[4].C
17
18     itlACK = FALSE
19     itlRX = FALSE
20
21     PTP HOME
22
23     INTERRUPT ON 10
24     $advance = 0
25     $OUT[2]=FALSE
26
27     LOOP
28         itlIdle = TRUE
29         Listen()
30         itlIdle = FALSE
31         SWITCH (itlTask)
32             CASE 200
33                 Lift()
34                 PTP HOME
35             CASE 30
36                 itlError=TRUE
37                 Lift()
38                 Scan()
39             CASE 35
40                 itlError=TRUE
41                 Scan()
42             CASE 10
43                 itlError=TRUE
44                 Lift()
45                 BigPicture()
46             CASE 15
47                 itlError=TRUE
48                 SmallPicture()
49             CASE 20
50                 itlError=TRUE
51                 Lift()
52                 MoveLSA()
53             CASE 25
54                 itlError=TRUE
55                 MoveLSA()
56             CASE 40

```

```

57             itlError=TRUE
58             Lift()
59             itlCalibrate()
60         CASE 250
61             itlError=TRUE
62             itlIdle = TRUE
63             EXIT
64         DEFAULT
65             itlError=TRUE
66     ENDSWITCH
67 ENDLOOP
68
69     PTP HOME
70 END
71
72 DEF RemoteEStop()
73     INTERRUPT OFF 10
74     BRAKE
75     WAIT FOR (NOT itlEStop)
76     INTERRUPT ON 10
77 END
78
79 DEF MoveLSA()
80     $advance = 0
81
82     PTP XposHigh
83     LIN XPosCentered
84 END
85
86 DEF Scan()
87     DECL POS ScanStart
88     DECL POS ScanStop
89     DECL REAL Xstart
90     DECL Real Xstop
91
92     Xstart = -0.5 * itlscanlength
93     Xstop = itlscanlength
94     ScanStart.Z = - 75
95     ScanStart.X = Xstart
96     ScanStop.X = Xstop
97
98     WAIT SEC 0
99     $advance = 3
100     PTP posCentered
101
102     LIN_REL ScanStart
103     WAIT SEC 1
104     $OUT[2]=TRUE
105
106     LIN_REL ScanStop
107     $OUT[2]=FALSE
108
109     LIN XPosHigh
110     WAIT SEC 0
111     $advance = 0
112 END

```

```

113
114 DEF SmallPicture()
115     $advance = 0
116     PTP posCamCentered
117 END
118
119 DEF BigPicture()
120     $advance = 0
121     PTP XposCamHigh
122     LIN XPosCamCent
123 END
124
125
126 DEF Listen()
127     WAIT FOR (itlTX)
128     itlAck = FALSE
129     itlRX = TRUE
130     WAIT FOR (itlProceed)
131
132     tool_data[3].A = tool_data[4].A + itlAOffset
133     tool_data[3].B = tool_data[4].B + itlBOffset
134     tool_data[3].C = tool_data[4].C + itlCOffset
135
136     itlRX = FALSE
137     itlAck = TRUE
138 END
139
140 DEF itlCalibrate()
141     itlAck = TRUE
142     itlAck = FALSE
143 END
144
145 DEF Lift()
146     DECL POS Up
147     Up = {X 0, Y 0, Z 50, A 0, B 0, C 0}
148
149     IF $pos_act.Z < 70 THEN
150         LIN_REL Up
151     ENDIF
152 END

```

Appendix E - Kuka Connector Wiring




Figure E.0.1 Connector wiring for the X12 jumper and the X11 trigger signal for the Kuka Robotics Controller.