Modelling High Peak-Power Solid-State Lasers

Adam Sroka

Doctorate of Engineering (EngD)

Institute of Photonics, University of Strathclyde

July 14, 2020

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

Signed:

Date:

Abstract

This thesis aims to improve existing design processes, increasing efficiency and reducing development costs through the use of numerical modelling techniques. This is achieved through leverage of existing modelling software, MATLAB and Zemax, and development of custom integrations, greatly enhancing their current capabilities.

This thesis is comprised of three distinct research projects the first of which details the development of a tool that enables automated analysis and interrogation of optical modelling systems. This tool was used to make recommendations in the design stages for a laser product developed at Thales UK Ltd., the sponsor company for this thesis.

The second major project concerned the modelling of a non-line-of-sight Light Detection and Ranging (LIDAR) system, using ray-tracing techniques to simulate time-of-flight responses from obstructed scenes. These models were used to analyse and understand some of the underlying relationships that governed this novel technology, and the author of this thesis presented these results at the conference *Advanced Optics for Defense Applications: UV through LWIR - Baltimore, United States* in 2016 [1].

The final research chapter details the development of a temporal and spatially discrete, numerical model, for investigating the evolution of gain within an active medium, during the pumping period. Amplified Spontaneous Emission (ASE) limits the energy available for useful gain and can severely impact the performance of high-power lasers, especially Q-switched systems. This model confrimed expectations around a novel rod geometry for suppression in high-power solid-state gain medium that could see benefits for Q-switched systems. These recommendations are the first steps to securing further investment to progress these designs into the physical domain and can be used to explore further, at significantly reduced costs.

This project has produced a number of tools to benefit those developing solid-state laser systems.

Contents

Acknowledgements

This thesis is the product of a long and arduous journey, one that would not have been possible without a great deal of help from some truly wonderful people. Here I'd like to thank a few of them.

Firstly, many thanks to all those from the Centre for Doctoral Training, Thales UK Ltd., and the University of Strathclyde. The vast network of support from behind the scenes hasn't gone unnoticed.

From Thales, I'd like the thank both Marc Smillie and Mark Silver for your support and friendship throughout. To my industrial supervisor Stephen Lee, for making the project a reality in the first place and imparting a wealth of lasers and engineering knowledge to me.

I'd especially like to thank my academic supervisor Alan Kemp - you're finally rid of me! Thank you so much for your eternal patience, support, and guidance. I'm lucky to have had the chance to work with you and will aspire to lead like you as I grow in my career.

I'd also like to thank Iain Rodger, for understanding exactly how it felt when things were tough and sharing a dark sense of humour as some kind of strange coping mechanism.

To Perry Pearcey and Simon Stevens, who have helped me as I grew and changed throughout this. Another day!

Finally, to my wife Joanne Martindale, words will never do. We've seen so much throughout the life of this project. You mean everything to me and without you none of this would have been possible, the last piece to my puzzle.

CHAPTER 1

Introduction

This chapter outlines and summarises the contents of the remainder of this thesis, aiming to introduce the tools developed within this body of work, and the motivations behind them. As the optical systems being designed were, in the primary use case, intended for production, there are many challenges faced when optimising the design of a system. Each of the following chapters is a self contained piece of work, with its own introductions and conclusions, making it possible to read and understand each of them in isolation. This introductory chapter is therefore more concerned with summarising the thesis as a whole, and discussing the common theme that binds the work herein. Chapter 2 contains a literature review on many of the topics related to the computational modelling of laser systems.

The overarching theme that runs throughout this body of work then, is the improvement of existing design processes, to increase efficiency and reduce development costs, through the use of numerical modelling techniques. This is achieved, primarily, through leverage of existing, powerful, modelling software by development of custom integrations, greatly enhancing their current capabilities. The main software used throughout this thesis, is a ray-tracing and stray light analysis program called Zemax. By interfacing this with MATLAB, a numerical computational suite, it becomes possible to apply the power of Zemax to simulate the ray-paths of solid-state laser systems in a number of different scenarios, and automate their analysis. A key theme that features throughout the entire thesis was minimal duplication of effort; there is little merit in designing and implementing a ray-tracing tool when one already exists. Another common theme that drives many of the decisions made within this thesis, is that of fitting into existing workflows; the tools developed here must consider ease of use by the primary target user, both of these design philosophies are discussed further in section 1.1. Building on the strong foundations of existing software and integrating them with other computational modelling tools enabled new capabilities to be investigated.

The thesis is separated into five chapters following this one. The next chapter is an overview of computation and numerical modelling techniques. Then the three research chapters follow, the first of which details the development of a tool that allows for automated analysis and interrogation of optical modelling systems, to aid in laser design. The second major project concerned the modelling of a non-line-of-sight Light Detection and Ranging (LIDAR) system, using ray-tracing techniques to simulate time-of-flight responses from obstructed scenes. Analysing these responses allows for the tracking of moving targets blocked from line-of-sight. Being able to accurately model these responses for complex scenes enables the investigation of dominant and limiting factors for these systems, with a much lower cost than physical experimentation. The final research chapter details the development of a temporal and spatially discrete, numerical model, for investigating the evolution of gain within an active medium, during the pumping period. Amplified Spontaneous Emission (ASE) limits the energy available for useful gain and can severely impact the performance of high-power lasers, especially Q-switched systems. Accurate models that enable the investigation and mitigation of ASE are beneficial when developing these systems and enable the user to trial new designs to improve performance with a lower cost and greater speed. Although each of the research chapters contain individual conclusions about the projects they discuss, a final chapter summarises the major research and concludes the thesis as a whole.

What follows is a brief discussion of the motivations behind the research carried out within each of the projects in this thesis, in section 1.1. The major findings of each of the research projects are then presented in section 1.2, including how each of the tools developed during this project were utilised and provided value to the end user. Finally, the remainder of the thesis is outlined in section 1.3, detailing the structure of the remaining chapters.

1.1 Motivations

The primary goal of this project was to develop tools that would enhance and improve existing laser design processes, for a number of different situations. These tools are not limited to the internal, optically active, processes that occur within laser gain media, but extend to the entire system.

It was decided early on that, should a tool already exists that accomplishes a necessary step in implementing a required model, then efforts should be made to utilise and integrate that with the existing framework. As nearly all of the software developed during this project was produced by a single developer, the level of sophistication and capability achievable in developing a new tool, from the ground up, would be severely lacking when compared to existing, established software that specialised in a similar job. More could be achieved by leverage of the larger budgets, bigger teams, more time and specialist knowledge that an established software company had invested in such tools. Further to this, the tools developed during this project each have real business applications within Thales UK Ltd., the host company for this project. This means that they will need to be maintained going into the future and fixed when extended routine use, potentially, exposes errors in their behaviour. Utilising existing software has the added benefit that this maintenance and error correction is, most likely, conducted by those that originally developed it. By focusing on integrating these tools with one another to produce new capabilities, the complexity of the software developed herein is reduced, limiting room for error, difficulty of debugging in the future, and the time cost associated with that development, which is at a premium in any project such as this.

Another key consideration, throughout the research, was that all of the tools produced in this project were easy to integrate into the existing design processes of the intended users, such as optical design and laser engineers. Any tool will face adoption issues if it is not user friendly and easily understood; developing a new tool that required a lot of time to understand may result in wasted effort as the intended users base abandons it and returns to their familiar processes. With this in mind, it was decided that focusing on enhancing the tools that were already in use during these design processes would result in minimal overhead when introduced to the users. By hiding the complexities of the models behind the familiar interfaces of these existing tools it made integrating and adoption into the design process straightforward and simple.

Chapter 3, the first of the technical chapters, strongly reinforces both of these themes. In developing a tool for assistance in laser cavity and optical system design, it quickly became apparent that the intended users could not be expected to learn a new interface to incorporate something new and complex. The tool was intended to improve and shorten the optical component design process.

As the optical systems being designed were, in the primary use case, intended for production, each of the measurements and parameters in the specification are subject to manufacturing tolerances. These are acceptable and anticipated deviations from the specification, within which the component can still fulfil its purpose. In isolation, under perfect conditions, each optical component can be straightforward to design and its behaviour well understood. The interplay between many components makes it much more difficult to accurately predict the overall performance, especially when extrinsic factors such as vibration or shock can interfere.

Strict tolerances on a component's specification have a higher production cost associated with them, as measurements must become more exact, more care has to be exhibited in producing the components. Alignment tolerances are one of the major focuses of chapter 3, a component that must be aligned to within a few hundred micrometres will take more time to align than one that only needs to be within a millimetre. When many components are placed together in a system, it is hard to anticipate how these tolerances can compound to affect performance. This leads to strict tolerances being specified as a rule-of-thumb. If the overall system can be interrogated and tested, quickly and in an automated fashion, it becomes possible to investigate what parts of the specification, if any, can have their tolerances relaxed, which in turn reduces costs.

The second project within this thesis, chapter 4, involved the development of a model for a non-line-of-sight LIDAR system, for tracking and positioning targets outside the field-of-view. The work outlined in this chapter was presented at the conference *Advanced Optics for Defense Applications: UV through LWIR - Baltimore, United States* in 2016 [1]. This model was based on the experimental work conducted by Gariepy et. al [2, 3, 4] and only the computational work was carried out as part of this thesis. This system had been developed and tested under laboratory conditions and gave promising results. In these early developmental stages, the limitations of the system were still not well understood.

The system had only been used on small scenes; understanding if the approach could be applied to scenes and targets on a human scale, in real time, would require much larger testing spaces and more sophisticated optical systems. A computational model, that allowed the researcher to arbitrarily alter the scene and optical system, affords the opportunity to understand the limitations of the technique for little monetary and time cost. By iterating through many systems at low cost using the model, research can then be focused on the more promising results, limiting room for error and misspent developmental time.

Chapter 5 is the final technical chapter in this thesis and presents the development of a numerical model for investigating laser gain medium activity for complex pumping geometries, often used in the production of high energy lasers by Thales UK Ltd. Of particular interest in the genesis of this project were gain limiting phenomena, such as Amplified Spontaneous Emission (ASE) and parasitic effects, that occur at high pump levels within solid-state lasers. Understanding the limits of pumping power and time, at which these effects begin to surface, or how certain components and pumping geometries can alter their prevalence, would allow for far better understanding of how to optimise a laser system to achieve maximal performance.

Many analytical solutions exist for modelling the behaviour of laser systems; many of these, however, are difficult or impossible to use when the cavity geometry is not easily represented in an analytical form. Fundamental object shapes, such as cylindrical rods and simple mirrors, are well understood and their interactions can be accurately modelled by analytical representations. Many systems in production, however, will require bespoke components and configurations that are not well understood and thus no accurate analytical model yet exists for them.

Computer Aided Design (CAD) programs allow the user to develop computational representations of complex, three-dimensional components and these programs are often used in the design of production lasers systems. As these tools were already in use throughout the design teams, enhancing the capabilities of the existing tools, as opposed to developing entirely new ones, would deliver the most benefit for the research investment and keep with the overall ethos of the thesis. The goal of this project was therefore to develop the capability to accurately model the behaviour of solid-state laser systems, which could incorporate any object or geometry that these CAD program could produce. This again would be achieved through interfacing ray-tracing software with a numerical modelling suite to build a temporally and spatially discrete model of the optical systems, using time-series and finite-element approaches.

Overall, this thesis aimed to improve the design process for a number of use cases by enhancing the existing techniques and tools already in use. Developing tools that are easy to adopt and use was a key consideration throughout and one which steered many of the decisions. By making use of advanced and established software throughout the project, combining the results of a ray-tracing package with a numerical computing environment, it was possible to achieve more in shorter time-frames than attempting to develop these capabilities from scratch.

1.2 Outcomes

Each of the three research projects that comprise this thesis produced many positive results and delivered tools that were put to use in their respective domains. This section will briefly discuss the major findings from each of the technical chapters, however chapters 3 to 5 each contain far more detailed discussions of the respective work.

The optimisation and tolerancing work in chapter 3 delivered a set of tools to quickly assess key points of interest, and focus for optimisation or cost saving efforts, within a design, in an automated and time-efficient fashion. This enables searching large combinations of complex optical parameters, across multiple components, with minimal input from the user, highlighting those areas that most contribute to changes in a system's performance. The model was successfully tested, experimentally, for a highly polarisation sensitive optical system giving increased confidence in its predictions for far more complex systems, that would otherwise take too long to reasonably test.

Three automated analyses were developed during this project to aid in the design process and allow the user to quickly identify where their efforts are best placed to improve performance. The first of the analyses, as presented in section 3.3.2.1, is concerned with finding optimal combinations of component parameters to maximise performance, by considering not only the peak performance, but also the possible deviations when tolerance ranges are taken into consideration. The second analysis, presented in section 3.3.2.2, seeks to highlight which parameters, both individually and in combination, occur most frequently in the highest or lowest performing subsets of possible combinations of all parameters within the system, allowing the user to single out those which may need tolerances adjusting for better mean performance. The final analysis, as detailed in section 3.3.2.3, considers each of the parameters individually to determine how they contribute to the overall variance in performance, considering all other parameters to be in pseudo-random states within their respective tolerance ranges.

The second project, concerned the development of a tool that could closely simulate the scene response observed by a non-line-of-sight LIDAR positioning system. The model was verified against a simple experimental case and the signal reconstruction for a stationary target was successful. It was possible to design large, complex scenarios that would have otherwise been difficult to test experimentally due to space limitations. This model made it possible to investigate the relationships between various factors for consideration when using the LIDAR system, such as distance of the target from the point of observation, in a much greater range of potential scenarios. Utilising ray-tracing approaches made it possible to investigate the paths of individual rays and determine the limits at which the response from a target fell below the background signal. Understanding these relationships and limitations allows the user to better understand the overall performance requirements when further development of the system is undergone.

The gain modelling project, in chapter 5, delivered a fully functional model that could accurately predict the gain evolution of arbitrarily shaped solid-state gain media and complex pumping geometries, as modelled in Zemax.

The results of an experimental investigation into the gain limiting effects of ASE were reproduced using the model and the simulated results closely matched the experiment. The model was then used to investigate a number of other pumping geometries and was put to use in making recommendations for designs which mitigated these effects.

Overall the projects in this thesis delivered on their original goals. In some cases these goals progressed and expanded as the research proceeded and what was achievable with the available resources became better understood. The outputs from each of the three research projects have been put to use in their respective domains and have been used to make design recommendations for further developments to other projects. Furthermore, these tools will remain in operation for future projects, affording the users an expanded repertoire of tools with which to answer potential questions, without the costs associated with experiment.

1.3 Outline

This section will briefly summarise the content in the remainder of the thesis. Each of the chapters will again be briefly introduced and their general structure presented.

The next chapter contains the overall literature review for the thesis. This is a broad and general overview of several themes key to this thesis, modelling of laser systems, computational and numerical modelling approaches, and how they can be applied to laser systems. Although some of the topics presented in the literature review are not prevalent throughout the remainder of the thesis, they have been included here either to better establish the context within which this project was conducted, or because they were considered a viable approach for accomplishing the goals of a particular project at one point. Each of the technical chapters also includes more specific reviews of the literature relevant to the project at hand. Specifics to any given chapter, therefore, have been excluded from the literature review.

The following three chapters present the record of original work carried out during this project. The first of these, chapter 3, details the development of a tool for improving the design process of complex optical systems, with many parameters, whilst considering the impact that deviations from specification and manufacturing tolerances have on performance. The chapter presents three automated analyses that aid the user to focus design efforts for improvements and optimisation. The tool is tested experimentally and finally applied to a system with many degrees of freedom and complex, polarisation sensitive components.

Chapter 4 presents the work on modelling the photon time-of-flight response from an obscured scene to simulate the results of a non-line-of-sight LIDAR system. A simple experiment illustrates that this model can replicate the response seen by the physical system, enabling computational estimation of a target position using the same process. The tool is then put to use determining some of the relationships and limitations that govern the final performance of such a system.

Chapter 5 documents the work conducted to produce a temporal and spatially discrete, finite-element and ray-tracing model to investigate the temporal evolution of gain in laser media, throughout the pumping period. The modelling process is first discussed and the algorithms presented in detail. The model outputs are then verified by reproducing the results of a paper investigating laser systems to mitigate the onset of high-power gain-limiting phenomena, as pumping power and time are increase. Finally the model is used to investigate other laser systems and their ability to mitigate these high-gain effects.

The last chapter concludes the work throughout the thesis. Each of the chapters are summarised, including their major components and the key findings from them. They are then discussed and key lessons learned during each of the project are presented. Future work for expanding on the tools developed herein is then discussed and recommendations made for where best to extract more value. Finally the thesis as a whole is discussed and critiqued in light of the key considerations set out at the start and the results of each of the research projects.

1.4 Hardware and Software

All of the computational modelling conducted throughout this thesis was performed on a PC workstation running Windows 7 SP1 version number 6.1.7601, with an Intel Core i7-3940XM 3 GHz quad-core processor, and 64 GB of RAM.

The numerical computing environment MATLAB [5] was used throughout. This software was updated twice-yearly following the standard MATLAB release-cycle, with version 20xxa being released in March and 20XXb in September, where 20XX is the release year. This work began in using MATLAB 2012a and the modelling tested and completed using MATLAB 2017b.

The ray-tracing software Zemax[6] is at the core of the modelling conducted in this thesis. Several versions were used during the project, the earliest of which being Zemax 12 Release 2 and the latest being Zemax OpticStudio V15 SP1. In Zemax there are several modes of operation. The models in this thesis use the non-sequential ray-tracing mode exclusively, although other modes were explored for some of the investigations.

CHAPTER 2

Literature Review

2.1 Introduction

Modelling is an important process in developing physical systems; it allows for critical decisions to be made on the basis of accurate predictions. Analytical models often provide a good first step in this process but can quickly become too difficult to implement, or inaccurate, when a system becomes complex. Fortunately, the speed of modern computation means numerical modelling techniques are easily achievable and often favourable in these circumstances. One such technique, known as ray-tracing, has become an invaluable, industry standard tool during the design process for complex optical systems.

2.1.1 Zemax, MATLAB, and Dynamic Data Exchange

Ray-tracing allows for accurate, detailed modelling of the performance behaviour of optical systems. Its usage and exploitation are at the core of this thesis and in particular a piece of software called Zemax and enhancement to its core functionalities [6]. The work presented in the later chapters expands upon the capabilities of Zemax, primarily by interfacing it with the numerical computing package MATLAB [5]. This allows for manipulation and representation of the data produced by Zemax, using all of the tools available in the MATLAB domain.

Dynamic Data Exchange (DDE) is a method of computational interprocess communication that allows one programme to access data made available by another programme. Zemax has a rich macro scripting language that allows some automation of the program for advanced modelling tasks. With these two technologies it is possible to control and manipulate Zemax, and its data from any other DDE capable program. Utilising and issuing DDE commands is simple in the Matlab environment making it straightforward to abstract any complexity into wrapper scripts and functions.

The design process utilising Zemax is simple, allowing complex systems to be quickly realised. Key performance characteristics can be interrogated to improve the user's understanding of a system's behaviour in a range of circumstances. Furthermore, this approach often affords a deeper understanding than is gathered from early experimental processes and can be accomplished in significantly shorter time-scales, for very little cost. These benefits make this approach to modelling an attractive strategy for lowering the risk of potential projects.

2.1.2 Models of light

The properties and behaviours of light have been studied at many levels of detail and abstraction. The appropriate model for any given task is often the simplest, or most efficient to implement, that considers all relevant phenomena. Whilst full quantum-electrodynamic simulations provide the most detailed view of a system, they are far more complex than geometric optics approaches [7, 8, 9]. Furthermore, they are unnecessary when, for example, the macroscopic properties of a lens are being determined. The long history of optics and photonics affords the use of many different tools, which can be best put to use dependant on the situation. From a computational perspective, different models determine which algorithmic approaches can be used in their optimisation. Often as models grow more detailed and complex they inhibit computational techniques that may otherwise improve their performance.

2.1.3 Outline

This chapter introduces the optical effects most important to the work carried out in this project. How these phenomena are modelled is then briefly summarised. Section 2.2 discusses some of the analytical and computational models that are foundational to the work in this thesis. The remaining sections discuss several phenomena that are of particular interest when optimising the designs of laser systems for defence applicaitons. Section 2.3 then introduces Amplified Spontaneous Emission (ASE) and its performance limiting effect in high power lasers. Parasitic Lasing is then discussed in section 2.4 and section 2.5 presents some of the models used to understand Excited State Absoprtion (ESA). Finally, in section 2.6 some methods of modelling the thermal phenomena in lasers are discussed.

2.2 Analytical models and computational methods

Analysing optical systems for design purposes becomes increasingly difficult with complexity. Several techniques have developed that allow for accurate predictions to be made about the output of such optical systems. These techniques traditionally involved solving geometric ray equations, but as systems become more complex these can become incredibly time consuming and difficult, sometimes even inaccurate [10]. Manipulating these techniques for use by computers allowed more complex systems to be rigorously analysed in much shorter spaces of time [11] and computational methods now continue to develop in order to keep up with modern requirements [10, 12].

This section will introduce some of these methods and the theory that they are founded upon.

2.2.1 Monte Carlo methods

The Monte Carlo method was introduced in the late 1940's, as a way of understanding physical systems, for which deterministic models were unsatisfactory [13, 14]. As a system grows in complexity, the mathematics required to produce solutions becomes ever more difficult to evaluate, and in some cases may not be known at all. A practical approach then, to better understand the system, would be to measure the output of multiple samples, taken with randomly varying inputs. Although certainty can never be reached without sampling all possible states, due to the law of large numbers [15, 16, 17], the model's accuracy improves as more samples are taken, tending towards the mean. This is how the Monte Carlo method works, an understanding of a solution is built up by examining the results from a large number of random samples.

Random sampling had seen used in earlier instances to solve numerical problems. One of the most well known examples dates back to 1777, known as Buffon's Needle [18, 19]. It was found that by repeatedly dropping a needle onto a surface, marked with equidistant parallel lines. The probability of any needle intersecting a line is given by,

$$\rho = \frac{2L}{\pi d} \tag{2.1}$$

where d is the needle length and L, being greater than d, is the line separation. As the number of needles dropped increases this method can be used to calculate the value of π, with increasing accuracy [20].

Neutron transport problems were one of the earliest applications of this approach where it saw extensive use and development in more recent times [21]. This was made possible by the advent of computers, which enabled individuals to simulate large numbers of random samples both quickly and more easily than was previously possibly. It is such a powerful and general technique, that it has used across most disciplines of science and engineering [22, 23, 24, 25].

In modelling light propagation, deterministic methods are often limited to restricted geometries, such as simple polygons, and uniform or ideal reflection and transmission properties. This is because the functions describing realistic objects can be difficult or impossible to define.

Monte Carlo methods can be used to interrogate any geometry or material property in this manner. The required series of large, monolithic functions, describing the system in an analytical model can be readily broken down into much simpler forms and then each part sampled randomly. In this case by randomly traced rays interacting with each part of the system. As more rays are traced, and thus more samples taken, the macroscopic behaviour of the system becomes apparent. The combined outputs of all the rays emulates the output from the optical system.

2.2.2 Maxwell's equations and scattering media

Consider an electromagnetic field $\vec{E}_{\text{inc}}$, $\vec{H}_{\text{inc}}$ incident upon a scattering object occupying a bounded region Γ_-, in three dimensional space $\mathbb{R}^3$, where the surrounding medium (Γ_+) equals

$$\Gamma_+ = \frac{\mathbb{R}^3}{\Gamma_-} \tag{2.2}$$

The total fields $\vec{E}_{\text{tot}}$,$\vec{H}_{\text{tot}}$ present in the surrounding medium (Γ_+) are equal to the sum of the scattered and incident fields.

$$\vec{E}_{tot} = \vec{E}_{inc} + \vec{E}_{sca} \tag{2.3}$$

$$\vec{H}_{tot} = \vec{H}_{inc} + \vec{H}_{sca} \tag{2.4}$$

and, finally, the fields present in the object (Γ_-) are known as the internal fields $\vec{E}_{\text{int}}$,$\vec{H}_{\text{int}}$. By using Maxwell's equations fields and material equations for the object it is possible to determine the internal and scattered fields from the known incident fields. Making the assumption that there are no field sources within the space and the material equations can be assumed to be linear (a safe assumption for low enough field strengths), Maxwell's equations simplify to

$$\nabla \times \vec{H}(\vec{r},t) = \frac{\varepsilon(\vec{r})}{c_0}\frac{\partial \vec{E}(\vec{r},t)}{\partial t} \tag{2.5}$$

$$\nabla \times \vec{E}(\vec{r},t) = -\frac{\mu(\vec{r})}{c_0}\frac{\partial \vec{H}(\vec{r},t)}{\partial t} \tag{2.6}$$

$$\nabla \cdot \vec{E}(\vec{r},t) = 0 \tag{2.7}$$

$$\nabla \cdot \vec{H}(\vec{r},t) = 0 \tag{2.8}$$

These eqs. (2.5) to (2.8) represent a system of eight equations for the six unknown field components [12]. To find these unknowns, solutions to these equations must be found. The wave equation can be found from manipulation of eq. (2.5) and eq. (2.6).

$$\nabla^2 \times \vec{E}(\vec{r},t) + \frac{\varepsilon(\vec{r})\mu(\vec{r})}{c_0^2}\frac{\partial^2 \vec{E}(\vec{r},t)}{\partial t^2} = 0 \tag{2.9}$$

From this the vector Helmholtz equation can be found by assuming the electric field takes the form $\vec{E}(\vec{r},t) = \exp(-\mathrm{i}\omega t)\vec{E}(\vec{r})$ i.e. is time harmonic [12]. The vector Helmholtz equation is given by

$$\nabla^2 \vec{E}(\vec{r}) + k^2(\vec{r})\vec{E}(\vec{r}) = 0 \tag{2.10}$$

$$\nabla^2 \vec{H}(\vec{r}) + k^2(\vec{r})\vec{H}(\vec{r}) = 0 \tag{2.11}$$

where 2.11 is the analogous vector Helmholtz equation for the magnetic field. These are differential equations that require boundary conditions to be solved. A requirement that the total fields are continuous across the surface of the object is one such condition and the radiation condition is another. The radiation condition states that the electric and magnetic fields approach zero as $\vec{r}$ approaches infinity [12].

A more complex model, encompassing the wave nature of light enables the user to observe effects such as diffraction and the modal behaviour of laser systems. This is especially useful when beam quality or cavity stability is important to the system being modelled. For many systems modelling the wave nature of light is not essential to accurately predict other key performance metrics, such as output power.

2.2.3 Differential equation methods

The techniques that solve the above equations are classified into various types. Those that do so numerically are called differential methods, such as the separation of variables method (2.2.3.1), the finite-element method (2.2.3.3), and the finite-difference time-domain method (2.2.3.2)).

2.2.3.1 Separation of variables method

The basic principle behind the separation of variables method (SVM) involves introducing a trial solution for separating the scalar Helmholtz equation. This solution is then used to produce differential equations for each of the components which are then solved [26, 27, 28]. With these solutions it is then possible to create wave functions that solve the vector Helmholtz equation. The incident, scattered and internal fields can then be expanded analogously to these wave functions as long as the wave functions satisfy the Helmholtz equation and are divergence free. Then by utilising the surface boundary conditions, a set of linear equations are formed that express the unknown expansion coefficients of the internal and scattered fields in terms of the known incident field [12, 26, 27, 28]. This method can be use in any coordinate system within which the Helmholtz equation is separable. While numerically this method is highly accurate, it can take a lot of computational power and time.

2.2.3.2 Finite-difference time domain

In this method time and space are no longer considered continuous and are represented by discrete intervals. This means that the spatial and temporal derivatives in the above equations are instead finite difference quotients. This is often done by choosing a grid of regular cells in three dimensions to represent the spatial domain and then iterating the calculations through a defined time step. This allows for the computation of an unknown value to be found in terms of the fields at earlier time steps. This is the equivalent of moving a plane wave or pulse across the spatial domain in discrete time steps and solving the equations numerically at every point in the grid for each step. Any object that may act to scatter the field is represented by a change in permittivity ε at its boundaries [29, 12, 30]. Grids with cell lengths of less than $\lambda/20$ are used and time steps should satisfy the Courant-Friedrichs-Levy condition [31] given by eq. (2.12).

$$c_0 \Delta t \leqslant \left[\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2} + \frac{1}{\Delta z^2} \right]^{-1/2} \tag{2.12}$$

In finite difference time domain (FDTD) calculations a defined domain containing the object is considered. It is beneficial from a speed of computation perspective to consider only the smallest possible domain. This however can lead to other problems, artificial reflections from the computational domain may occur. This will not only act to slow down processing speed by continuing to consider radiation that should have been considered to have left the system, but may also impair the accuracy of the results[10, 29, 12, 30]. Another problem this method has concerns curved surfaces, using a regular square or rectangular grid, curved surfaces are represented much in the same way they are in pixelated images. This can lead to numerical artefacts causing inaccuracies. One of the main strengths of this method is that discretising the spatial and temporal domains mean numerical solutions to large numbers of simultaneous linear equations like other techniques [12, 10].

2.2.3.3 Finite element method

The finite element method (FEM) involves solving the Helmholtz equation in the frequency domain by discretising space into defined intervals. Solutions can then be found numerically utilising the object boundary conditions. This method is similar to the finite difference time domain method detailed in section 2.2.3.2 in that a grid must be selected for the spatial intervals [32, 33, 34]. By enforcing continuity conditions across adjacent grid cells , a system of linear equations can be formed. Using techniques such as Gaussian elimination or the conjugate gradient method for example, these equations can be solved giving values for the field components at the node points [32, 33, 34]. This method forms coefficient matrices that are well understood and easily solvable by a computer making it particularly efficient.

Similar issues arise due to the finite computational domain as in the finite difference time domain method discussed in section 2.2.3.2 Issues also arise from enforcing the radiation condition, due to the nature of the FEM operating in the frequency domain.

2.2.4 Integral equation methods

2.2.4.1 Null field method

By making use of the Helmholtz equations and the vector-Green's identity, it is possible to derive the surface-integral equation.

$$\oint_{\partial\Gamma} \mathrm{d}\sigma(\vec{r}\,') \left[\frac{\mathrm{i}\omega\mu}{c_0} (\hat{n}_+ \times \vec{H}(\vec{r})) \cdot \vec{G}_0(\vec{r}\,', \vec{r}) + (\hat{n}_+ \times \vec{E}(\vec{r})) \cdot (\nabla \times \vec{G}_0(\vec{r}\,', \vec{r})) \right]$$
$$= \begin{cases} \vec{E}^{\mathrm{sca}}(\vec{r}) & : \vec{r} \in \Gamma_+, \\ -\vec{E}^{\mathrm{inc}}(\vec{r}) & : \vec{r} \in \Gamma_-. \end{cases} \tag{2.13}$$

Here $\hat{n}_+$ is a unit vector denoting the normal that points out of $\partial\Gamma$. The second case for this equation corresponding to when $\vec{r} \in \Gamma_-$ is the extended boundary condition and sometimes referred to as the null-field equation. It represents the continuation of the incident and scattered fields such that they cancel one another [12, 35, 36, 37].

The null-field method (NFM) first finds the boundary-surface fields using the second branch of the surfact-integral equation in terms of the known incident field. These are then substituted into the first branch of the equation making it possible to find the unknown scattered field. Due to the expansions utilised in this method, it is sometimes prone to convergence issues and can also take a long time to compute. It is however very good at giving highly accurate results in the far-field [12, 35, 36, 37].

2.2.4.2 Method of Moments

The vector Helmholtz equation can be written as

$$(\nabla^2 \vec{\vec{1}} - k_0^2)\vec{E}(\vec{r}) = \mathrm{i}\omega\mu_0 \vec{J}(\vec{r}) \tag{2.14}$$

where

$$\vec{J}(\vec{r}) = \begin{cases} \mathrm{i}\omega\varepsilon_0 \left[\varepsilon_r(\vec{r}) - 1\right] \vec{E}(\vec{r}) & : \vec{r} \in \Gamma_-, \\ 0 & : \vec{r} \in \Gamma_+. \end{cases} \tag{2.15}$$

By introducing $\vec{J}(\vec{r})$ as a volume current density. This enables the Helmholtz equation, a homogeneous differential equation with non-constant coefficients to be written as an inhomogeneous differential equation with constant coefficients, making it far easier to solve. A typical method for solving such an equation is to find the Green's function, a solution with a Dirac-delta inhomogeneity. In this specific case this is known as the free-space dyadic Green's function and is well understood. Using this Green's function and the volume current density defined in eq. (2.15) it is possible to obtain the following

$$\vec{E}(\vec{r}) = \vec{E}^{\mathrm{inc}} + k_0^2 \int_{\Gamma-} \mathrm{d}^3 r' \left[\varepsilon_r(\vec{r}\,') - 1\right] \vec{G}_0(\vec{r}, \vec{r}\,') \cdot \vec{E}(\vec{r}\,'). \tag{2.16}$$

Solving eq. (2.16) for the unknown $\vec{E}^{\text{sca}}$ field can be done with many methods. The method of moments (MoM) is a straightforward numerical approach that replaces the integral with a discrete summation over a finite number of volume elements [12, 38]. The volume is partitioned into sufficiently small volume elements, allowing simplifying assumptions to be made. This then allows a system of linear equations to be found that can subsequently be solved using Gaussian elimination or the conjugate gradient method [12, 38].

One of the key advantages of this method is that it can be applied to a wide range of objects including those that are optically active and , anisotropic and arbitrarily shaped. Also the computation only considers the volume of the scattering object removing some of the disadvantages differential methods experience.

2.3 Amplified Spontaneous Emission

Following the initial discussions by Dicke [39], further investigations into the effects of superradiance on laser output were conducted. Some of these investigations were carried out on what is now known to be Amplified Spontaneous Emission (ASE) , a different phenomenon that will be discussed more thoroughly in section 2.3.2. If spontaneous emission is being amplified within a gain medium then it will be extracting energy from the population inversion, this will in turn reduce the maximum gain experienced by stimulated emission that is available for the laser.

2.3.1 Superradiance

The term superradiance came about due to discussion of the classical treatment of spontaneous radiation in gases [39]. Previous treatments calculated excited molecules within a gas as radiating independently of one another. These were proposed on the assumption that the relatively large separation between molecules would prevent their states from affecting spontaneous emission probabilities [39]. This treatment could not therefore adequately explain coherent spontaneous radiation effects. A quantum mechanical treatment of the radiation process considering the whole gas as a single system, developed from research into nuclear magnetic resonance techniques, delivered a more general model for radiators in a magnetic or electric dipole transition [39, 40].

When an atomic system is coherently excited a macroscopic dipole moment is experienced by the atoms. Spontaneous emission can lead to a radiation field which then couples the remaining atoms in the system, resulting in synchronisation of their radiation fields and collective emission. This phenomenon gathered a great deal of interest and research was carried out to better understand and experimentally verify the underlying theory [40, 41]. It was found that if the conditions are sufficient for superradiance to occur then the rate at which the molecules emit is far greater than it would be for them to emit via incoherent routes and that the maximum intensity scales with the square of the number of atoms in the system [39, 41, 40].

2.3.2 Theory of ASE

After this research began to gather interest, it was outlined that a phenomena separate from the one outlined in section 2.3.1 was being confused with superradiance [39, 42]. Although work had been done previously to measure the effects of ASE, the phenomena was not properly recognised [43, 44]. Investigations into the theory of mirror-less lasers showed that, for a medium with a randomly prepared population inversion incapable of dephasing relaxation paths, the linewidth of the pulse would be broadened for a superradiant emission, whereas narrowing had been observed [42, 45, 46]. This narrowing however was evidence of spontaneous emission being amplified by the gain medium and it was proposed that these effects be described as Amplified Spontaneous Emission [47, 42, 48]. The first theoretical approach considered a very narrow signal growing to saturation within a gain medium [49].

Further work was then carried out to understand the theoretical conditions necessary for ASE to occur. Threshold conditions were then derived and related to the active medium length and the inversion density [50, 51]. Consider a length of material L and cross-sectional area a with a population inversion density of n. As N_p photons pass through this volume then the number of atoms that are stimulated to emit is given by

$$N_a = \frac{\sigma_r n N_p}{a} \tag{2.17}$$

Using the relationships between these values with the Einstein coefficient B and the radiation density

$$\rho(v) = \frac{N_p h\nu}{ac\Delta\nu_D} \tag{2.18}$$

Considering a Doppler line shape, we find

$$\sigma = \frac{c^2}{8\pi\nu^2\tau_{21}}\left(\frac{2\,(\ln 2)^{1/2}}{\pi^{1/2}\Delta\nu_D}\right) \tag{2.19}$$

It can then be determined that the threshold inversion required for ASE [50], when one spontaneously emitted photon can induce the emission of another photon after length L, by the equation

$$n_c = \frac{8\pi\tau_{21}}{L\lambda^2\phi}\left(\frac{\pi^{1/2}\Delta\nu_D}{2\,(\ln 2)^{1/2}}\right) \tag{2.20}$$

2.3.2.1 Lasing and ASE

With these relationships ASE effects could now be predicted and experimentally verified. Investigating relationships with other effects, in particular the threshold conditions required for lasing to occur, was also made possible. For a population inversion density $n = n_2 - n_1$, a relationship can be derived between the length of gain media needed to produce ASE for a given population inversion, L_c, and the threshold length of that material at which lasing will begin L_T [51, 52, 53].

For a population inversion density of n in the excited region of the material, the following two equations hold

$$nL_T = \frac{1}{3}\frac{8\pi\tau_{21}}{\lambda^2\phi}\left(\frac{\pi^{1/2}\Delta\nu_D}{2\left(\ln 2\right)^{1/2}}\delta_l\right) \tag{2.21}$$

$$nL_c = \frac{8\pi\tau_{21}}{\lambda^2\phi}\left(\frac{\pi^{1/2}\Delta\nu_D}{2\left(\ln 2\right)^{1/2}}\right) \tag{2.22}$$

where δ_l is the fractional losses after each pass through the laser resonator. Equation (2.22) is a rearrangement of eq. (2.20) to find the relationship in terms of length for a given inversion density [52, 51]. A direct relationship between the two values for lasing and ASE can now be made

$$\frac{L_T}{L_c} = \frac{1}{3}\delta_l \tag{2.23}$$

This tells us that the threshold length of gain material required for ASE to occur is greater than that required for lasing for the same atomic transition [43]. This can be explained by observing the fact that in a laser resonator, multiple reflections from the cavity mirrors effectively extend the length of the gain material. Only emission in a direction that will experience adequate reflections to reach a path length equal to L_c will contribute to ASE [51, 52, 50].

2.3.2.2 Intensity and saturation of ASE

Now consider a cylinder with cross-sectional area a and length L of gain medium, with population inversion density $n = n_2 - n_1$ between two energy levels (denoted by the subscript), at a time t and position x. If R_i is the rate at which the ith level is populated from a source level with population n_0 and A_i is the radiative decay probability, rate equations for the time dependent population densities can be written [51, 52]

$$\frac{\delta n_2}{\delta t} = -\left(n_2 - n_1\right)\frac{\sigma}{a}\left[N_p(x) + M_p(x)\right] - An_2 - A_2 n_2 + R_2 n_0 \tag{2.24}$$

and

$$\frac{\delta n_1}{\delta t} = \left(n_2 - n_1\right)\frac{\sigma}{a}\left[N_p(x) + M_p(x)\right] + An_2 + R_1 n_0 - A_1 n_1 \tag{2.25}$$

where $N_p(x)$ and $M_p(x)$ are the number of photons passing through the area a per second at position x in the positive and negative directions respectively. In the steady-steady state it can be shown that [51, 43]

$$n_2 = \frac{B + C\left[N_p(x) + M_p(x)\right]}{1 + E\left[N_p(x) + M_p(x)\right]} \tag{2.26}$$

and

$$n_2 - n_1 = \frac{F}{1 + E\left[N_p(x) + M_p(x)\right]} \tag{2.27}$$

The coefficients C, E and F characterise the specific system to be investigated. With these equations the rate of increase in the number of photons due to ASE over a distance in the steady-state can be caclulated by

$$\frac{\delta N_p}{\delta x} = \frac{F\sigma N_p}{1 + E\left[N_p(x) + M_p(x)\right]} + \frac{An_2 a\Delta\Omega}{4\pi} \tag{2.28}$$

where

$$\Delta\Omega = 2\Pi\left(1 - \frac{L - x}{\left[(L - x)^2 + r^2\right]^{1/2}}\right) \tag{2.29}$$

Here $\Delta\Omega$ is the solid angle into which the photons can be emitted for a medium of radius r. In this equation the population n_2 is approximated to constant in time as it varies slowly compared to the other terms. For medium lengths above the threshold for ASE ($x \geq L_c$), when a mode is excited continuously, the two counter-propagating waves grow over the entire length of the medium and their interaction with one another is gain dependent [51, 43].

Equation (2.28) can be rewritten in terms of measurable quantities such as the detected photocurrent of a detecting photomultiplier by replacing $N_p(x)$ by $I_p(x)$ times a contstant. Figure 2.1 illustrates these values for a rod of length L and radius r. Where I_p is the detected photocurrent for the wave travelling in the positive direction and J_p corresponds to the wave travelling in the negative direction [43, 51]. Including a third term I_q, representing an amplifier signal, it becomes possible with these three equations for proper calculation of the of affect ASE in an amplifying system.

$$\frac{\delta I_p(x)}{\delta x} = \frac{K_2 I_p}{1 + K_3\left[I_p(x) + J_p(x) + I_q(x)\right]} + \frac{K_1\Delta\Omega'}{\pi r^2} \tag{2.30}$$

$$\frac{\delta J_p(x)}{\delta x} = -\frac{K_2 J_p}{1 + K_3\left[I_p(x) + J_p(x) + I_q(x)\right]} - \frac{K_1\Delta\Omega'}{\pi r^2} \tag{2.31}$$

$$\frac{\delta I_q(x)}{\delta x} = \frac{K_2 I_q}{1 + K_3\left[I_p(x) + J_p(x) + I_q(x)\right]} \tag{2.32}$$

where

$$\Delta\Omega' = 2\pi\left(1 - \frac{x}{\left(x^2 + r^2\right)^{1/2}}\right) \tag{2.33}$$

Here the coefficient K_1 represents a loss constant for the medium through non-radiative decay, K_2 is the small signal gain and K_3 is a saturation effect coefficient [43]. These equations have no analytical solution but can be solved computationally by the Runge-Kutta method with the following boundary conditions $I_p(t_0) = 0$, $J_p(L) = 0$ and $I_q(t_0)$, the amplifier signal seed term, is a predetermined constant value [51].

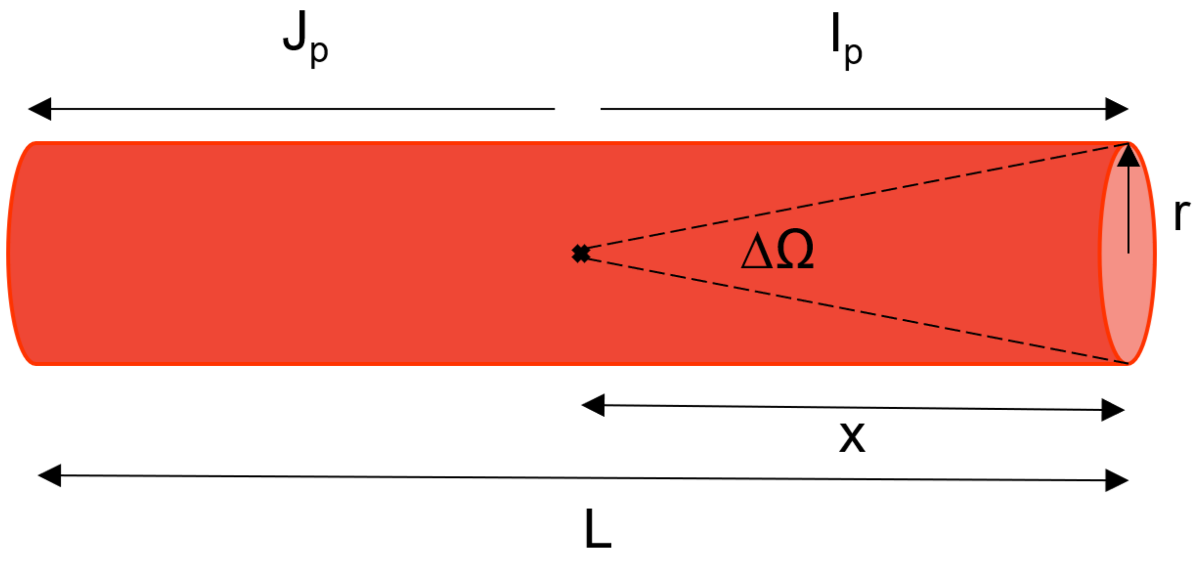

Figure 2.1: Illustration of the waves travelling in the positive I_p and negative J_p directions, emitting photons into the solid angle $\Delta\Omega$ from a point x, for a rod with length L and radius r.

2.3.3 High-power

For some applications achieving high output power is a key consideration. When pumping a gain medium, whilst below the threshold for parasitic lasing (introduced in detail in section 2.4), there are three main processes that will affect the population inversion density, the pump action, the spontaneous emission decay rate and ASE [54]. The spontaneous emission decay rate will not change for a given gain medium and thus the output power from spontaneous emission will increase linearly with pump power until the point of saturation or the threshold for another effect is reached. After its onset, as pump power increases, the output power contribution from ASE will increase non-linearly and begin to dominate, this can be a considerable problem for optical fibre amplifiers discussed in section 2.3.5 [55].

If spontaneous emission is being amplified within a gain medium then it will be extracting energy from the population inversion, this will in turn reduce the maximum gain experienced by stimulated emission that is available for the laser. A great deal of research has been conducted to try and understand this phenomenon and novel approaches have been made in attempts to overcome it [56, 57, 58, 59, 54].

Some attempts have been made to model these effects using three-dimensional ray tracing programs [56]. These methods however can be computationally expensive and often assumptions must be made to make them more manageable. These models make it possible to predict the Gain Length Product (GLP) achievable in the presence of ASE based on system parameters such as shape and size of the medium, pumping length, spectral line profile and Total Internal Reflection (TIR) angle [56]. The gain length product is simply how much gain a photon experiences along the length of its path through the system.

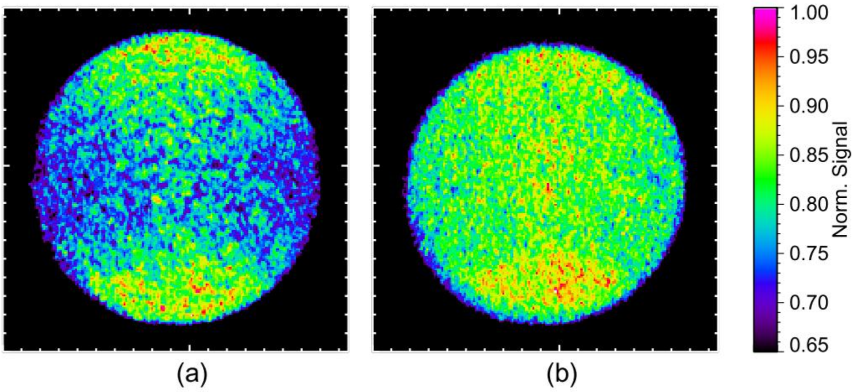

Figure 2.2: Normalised fluorescence intensity for a Ti:S laser rod without (a) and with (b) an edge cladding material. Images are normalised to maximum signal, showing that the signal is far more evenly dispersed across the medium when the cladding is applied [61].

2.3.4 Solid-state lasers

One method that has been utilised to suppress ASE involves bonding another material to the edge of the gain medium [58, 57]. This is known as an edge cladding. Edge claddings have to be designed to tight specifications if they are to work effectively. The refractive index must be closely matched to that of the gain medium to prevent surface reflections and allow ASE to couple to the cladding material [57, 58, 60].

These materials are then doped with an element that absorbs at the ASE wavelength [57, 58]. Doping a host material will cause a change in the refractive index dependent on the level of doping and the dopant itself. This means that a host material that will bond well to the gain medium can be used even if the refractive index does not match. By adjusting the concentration of the dopant, this discrepancy can be reduced to a minimum [57]. Room temperature-vulcanised silicone rubber edge cladding materials have been demonstrated that have refractive indices between 1.42 and 1.54, covering many common laser glasses [58]. Another large benefit of these techniques is that they are also easy to produce and can be applied to large gain media geometries. Figure 2.2 shows how the addition of an edge cladding material can reduce the intensity of fluorescence from a laser gain medium [61]. The normalisation of the signal shows that the 0.80 - 0.95 signal regions in (a) are more evenly spread throughout the medium after the cladding is applied, as shown in (b). This image also highlights the hourglass shape that is common for gain depletion by ASE in cylindrical rods [61].

Another method utilised in many cases is grinding the surfaces of the gain medium. This acts to reduce the surface reflections. Surfaces are treated with abrasive powders such as boron carbide. When grain sizes lower than 30μm are used they can have a polishing effect on parts of the material, having the opposite effect to that desired [62].

Larger grain sizes however have little effect on improving the extraction efficiency. Larger grains can however lead to more profound defects such as micro-cracks, which can increase the risk of stress fracture [62]. A drawback of this technique is the amount of pump radiation that stays within the gain medium due to TIR is reduced, resulting in a low pump confinement efficiency [63, 64]. Hybrid approaches to this problem have been investigated and demonstrated [63].

2.3.5 Fibres

ASE can be a considerable problem for fibre systems. Phase and frequency modulated signals are degraded by noise from ASE because of the potentially long transit lengths although intensity modulated signals do not suffer from this effect as badly [65, 66, 67, 68]. The presence of ASE in the fibre can lead to optical beating with the signal. This causes a random variation in intensity and therefore gain variations that will be experienced by the signal pulse [69]. ASE also contributes a considerable amount to the overall noise in these systems and accurate measurements of the optical signal to noise ratio is required to estimate the signal's quality [70]. Noise from ASE in signals applications can also make it difficult to distinguish between zeros and ones, this places a minimum energy limit on the system to account for this noise otherwise signal quality is lost.

ASE can also experience Kerr effects when mixing with the signal radiation [65, 66]. This can lead to large spectral broadening and result in timing jitter for the pulse arrival time [65, 66]. These non-linear interactions can be reduced by use of in-line filters which limit the degradation and reducing what is known as the Gordon-Haus jitter [65, 66, 71]. Polarisation can be used to filter ASE as long as there is a significant polarisation difference between the ASE and the signal. Photonic bandgap fibres consist of a low index core with a cladding formed by air and silica and offer another solution to the ASE problem. This cladding can be doped to tune the refractive index, allowing for wavelength specific confinement and can be tuned in such a way that ASE radiation is not efficiently confined within the fibre [72]. Careful selection of the core and cladding doping elements must be made in this situation.

ASE can also be of use in some applications where broadband emission is desired [55]. This will be discussed further in section 2.3.7.

2.3.6 Thin-disk lasers

Thin-disk lasers offer a novel approach to thermal management within solid state laser media. By mounting a gain medium with a large aperture to length ratio to a heat sink, and using multipass pumping geometries higher powers are achievable despite increases in waste heat generation [73]. Further discussion of the thermal issues in laser gain media can be found in section 2.6.

Increasing the power of these lasers can then be achieved by increasing the pump area of the gain medium. Increasing the aperture of the gain medium however leads to problems with ASE. The increased aperture means increased path length for spontaneously emitted photons travelling in the transverse direction. This results in an increased gain length product in the transverse direction, causing greater ASE [74]. ASE therefore infers a limit on power scaling by increasing the aperture size of thin-disk gain media, and estimations of these limits have been made [74, 75, 76].

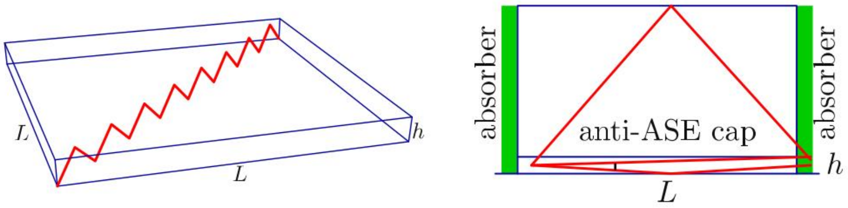

Figure 2.3: Path lengths of up to 2L are achievable by spontaneously emitted photons following the path illustrated on the left. With an undoped cap, these paths are suppressed and $\sqrt{2L}$ is the longest available path, as shown on the right [76].

Several methods have been proposed to mitigate ASE effects in thin-disk lasers. One involves adding an cap of undoped material to the thin-disk [76]. Due to the exponential growth of ASE in a gain medium, the longest path lengths experienced by spontaneously emitted photons are of greatest importance [76]. Path lengths of up to twice the longest dimension of a gain medium can be achieved by some photons, this is illustrated in fig. 2.3. By adding a cap of undoped material possessing a closely matched refractive index, this limits long path lengths experienced due to TIR at the pumping face. Now only photons emitted into a solid angle of $2\pi\frac{2h}{L}$, where L is the lateral dimension of the medium and h is its height, experience strong amplification in the medium [76].

Should the ASE within a gain medium experience Fresnel reflections at the boundaries then parasitic lasing can occur. Bevelling the edges of the thin-disk can prevent this from occurring although due to the machining process it can limit the usable area of the gain medium [77]. Further discussion of parasitic oscillations and suppression techniques are discussed in section 2.4.

2.3.7 Sources and other uses

There are many applications for broadband light sources with stringent criteria that ASE sources can match. Fields such as spectroscopy, medical imaging and fibre sensors, require high beam quality, brightness and intensity. ASE can provide a solution to these requirements [55, 78].

2.4 Parasitic Lasing

Another effect that can occur within a gain medium is the phenomena of parasitic lasing. This is distinguishable from ASE detailed in section 2.3. ASE can occur in all directions and at all frequencies in the fluorescent spectrum of the gain medium although it will be more evident in directions that maximise path length through the gain medium [79]. Parasitic lasing however forms modes, as some of the radiation will experience reflections and return to the point of origin with enough coherence to interfere [79]. Only ASE experiencing closed paths and sufficient feedback will result in parasitic lasing [54].

Parasitic oscillations occur due to a number of processes and are therefore difficult to manage [80, 81]. Any stimulated emission can form an undesirable laser mode if it experiences sufficient reflections. Fresnel reflections at the gain medium boundaries, as well as specular and diffuse reflections from surfaces around a gain medium can cause parasitic oscillations leading to undesirable modes. This results in a reduction of gain for the desirable modes and in high power systems can damage components of the laser [81].

2.4.1 Liquid claddings

In systems with a large gain spectrum, such as those required for ultrashort pulses, parasitics can considreably deplete available gain. Ti:Sapphire lasers have a large gain bandwidth making it well suited for generation of ultrashort (femtosecond) pulses as well as tuneability over a wide range due to the conjugate relationship between pulse length and frequency components. This also means that transverse lasing can occur at any wavelength within that range [60]. In section 2.3.4, methods of bonding an absorbant material to a gain medium to suppress ASE are discussed. As parasitic lasing comes about due to ASE, these techniques will also suppress parasitic lasing. For broad gain bandwidth media like Ti:Sapphire however selecting a material capable of absorbing the entire emission spectrum can prove difficult. One novel technique involves surrounding the gain medium with a broadband absorbing dye [82, 60, 61]. This method has the added advantage of allowing the liquid ot be utilised to cool the gain medium, improving thermal management, a topic further discussed in section 2.6.

2.4.2 Cant edges

Another approach to this problem is the use of cant edges for the laser gain medium [63]. Using bevelled edges prevents light being reflected by TIR and Fresnel reflections from oscillating between symmetrical surfaces, thus forming a cavity, an example of which is illustrated in fig. 2.4. Choosing the angles used carefully can prevent light from being confined in such a way after successive reflections [63, 83, 57, 59]. Using bevelled edges and carefully chosen angles can also assist with thermal effects experienced in lasers, again, this is discussed in more detail in section 2.6.

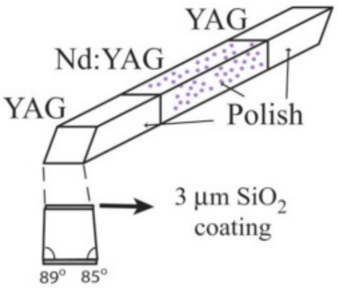

Figure 2.4: Example of a slab constructed with bevelled edges to prevent light being reflected by total internal reflection [63].

2.5 Excited State Absorption

In some solid-state gain media it is possible for absorption to occur from levels higher than the ground state. This is known as Excited State Absorption (ESA). It is often the case that the energy of the fundamental laser wavelength is sufficient to excite electrons in the meta-stable upper lasing level to higher energy states. Thus as laser radiation builds within the cavity, more ESA occurs. Gain media with broad absorption cross-sections are particularly susceptible to this issue, as are many materials used as saturable absorbers. Q-switched systems also suffer from a large contribution to total losses from ESA. This is due to that high peak population in the upper lasing level [84].

ESA acts as another loss mechanism in the cavity as the radiation is being used for unwanted absorption processes and is therefore not able to contribute to gain. If these losses are significantly large it can lead to an increase in the pump threshold power of the laser and a reduction in its efficiency. Pump radiation can also lead to ESA for some media which means special attention must be made to the optimisation of the pump wavelength to reduce this effect, sometimes resulting in utilising pump wavelengths that are not the same as the highest absorption wavelength for the gain media. Modelling the affects of ESA is possible and in some cases can be simplified with the inclusion of an additional absorption term. Should the higher-lying level have a significantly long lifetime in comparison to the upper lasing level, relaxation rate equations may be necessary to accurately determine ESA's effects.

Section 2.6.1.3 discusses upconversion, which produces laser radiation with a shorter wavelength than the pump source, this process commonly utilises ESA.

2.5.1 In Nd^{3+} doped gain media

Measurements of the slope efficiencies for Nd^{3+} gain media often differ from those theoretically predicted. ESA is a considerable factor in these discrepancies and has great effect on overall laser performance [84]. The losses due to ESA in Nd^{3+} lead to an infrared emission with a shorter wavelength than the laser radiation, this higher energy transmission has a non-exponential decay that shortens with higher pump-powers [85, 86].

When ESA occurs at the laser wavelength it remains relatively simple to accurately calculate the overall gain cross-section for the gain media. The absorption in the excited state competes with the stimulated emission process, acting to reduce the overall gain cross section (σ) [84, 85, 86],

$$\sigma = \sigma_e - \sigma_{ESA} \tag{2.34}$$

the stimulated emission cross-section (σ_e) can be taken from the fluorescence data allowing for straightforward calculation of the ESA cross-section (σ_{ESA}). For Nd^{3+} ions, ESA at the laser wavelength does not occur in YAG host media, and in YLF and LMA the polarisation of the radiation is important [85]. In most cases within these media it does not play significant enough a role in the gain dynamics to considerably alter laser performance [86].

Calculating the effects of ESA occurring due to the pump source becomes more difficult. In this case the reduction in efficiency becomes dependent upon the inversion density and pump light penetration into the crystal [85]. In host media such as YAG and YLF crystals ESA occurring at the pump wavelength is not as much of an issue as they often have lower doping concentrations than other media and the infrared fluorescence decays exponentially. However, if the ESA cross-section is of the same order of magnitude as the stimulated emission cross-section then it can seriously reduce laser performance [86]. ESA occurring from the upper lasing level due to the pump radiation in the visible region, between 400 and 900 nm, is calculated to have a cross-section no greater than 10% of the Ground State Absorption (GSA) cross-section in the same region [86]. This means that ESA only becomes significant at when the excitation density is large, which is often the case for flash-lamp pumped solid-state systems. Thermal effects that arise from ESA can account for nearly a third of the total thermal issues, as discussed in section 2.6 [86]. Around the 800nm mark however, very little ESA occurs from the upper lasing level, meaning high pump intensities at these wavelengths, typical of diode-pumped Nd^3:YAG systems, will not be significantly influenced by ESA.

2.5.2 Calculating ESA

Consider a longitudinally pumped Nd^3 solid-state laser emitting around 1.06μm. Making the assumptions that all of the photons that are absorbed in the excited state do not contribute to laser action and no GSA occurs at the lasing wavelength (as it is a four level system), it becomes possible to solve equations for the gain evolution.

$$\frac{\mathrm{d}I_p}{\mathrm{d}z} = -\left(\sigma_0^p N_0 + \sigma_1^p N_1\right) I_p \tag{2.35}$$

$$\frac{\mathrm{d}I_L^\pm}{\mathrm{d}z} = \pm\left(\sigma_e - \sigma_1^L\right) N_1 I_L^\pm \tag{2.36}$$

$$\frac{\mathrm{d}N_1}{\mathrm{d}t} = \sigma_0^p N_0 I_p - \left(\sigma_e + \sigma_1^L\right) N_1 I_L - \sigma_1^p N_1 I_p - W N_1 \tag{2.37}$$

Where $\sigma_0^p,\ \sigma_1^p,\ \sigma_e$and σ_1^L are the cross-sections for the GSA at the pump wavelength, ESA at the pump wavelength, the stimulated emission and ESA at the laser wavelength respectively. I_p is the intensity of pump radiation in the gain medium and $I_L^{\pm}$ is the laser radiation intensity in either the plus or minus direction. Lastly N is the ion density in either the excited (N_1) or ground (N_0) state and W is the spontaneous emission rate.
The pump threshold and laser output power can then be found by

$$P_{abs}^{th} \cong \frac{G^{2r}-1}{2\sigma_1^p\tau}\frac{1-\Theta_0}{1-\Theta G^r}h\nu_p S \tag{2.38}$$

$$P_{\text{laser}} = \rho_{\text{diff}}\left(P_{\text{abs}} - P_{\text{abs}}^{\text{th}}\right) \tag{2.39}$$

$$\rho_{\text{diff}} \cong \frac{h\nu_L}{h\nu_p}\frac{T}{2}\frac{2\sigma_1^p}{\sigma_e+\sigma_1^L}\frac{1}{G^{2r}-1}\frac{1-\Theta G^{2r}}{1-\Theta_0} \tag{2.40}$$

eq. (2.40) is for the slope efficiency of the laser. Here $h\nu_p,\ h\nu_L$ are the energies of the pump and laser photons, G is the single pass gain, S is the cross-section area of the pump and cavity modes inside the gain medium and Θ and Θ_0 represent the transparency of the gain medium above and below the lasing threshold respectively.

It then follows that,

$$\Theta = \Theta_0 G^{\sigma_0^P/(\sigma_e-\sigma_1^L)}G^{-2r} \quad \text{and} \tag{2.41}$$

$$\Theta = e^{-\sigma_0^p N_0 d}$$

$$P_{\text{abs}} \cong (1-\Theta_0)P_{\text{inc}} \tag{2.42}$$

$$G = L^{-1} = \left(1-\frac{T}{2}-\delta\right)^{-1} \tag{2.43}$$

where P_{inc} and P_{abs} are the power incident upon and absorbed by the gain medium respectively. L is the single pass losses and δ accounts for all fo the intracavity losses that are not due to the transmission of the output coupler.

In the above equations, r is a value that compares the relative magnitude of the ESA to the stimulated emission cross-section and is defined in eq. (2.44).

$$r = \frac{\sigma_1^p}{\sigma_e-\sigma_1^L} \tag{2.44}$$

In the continuous wave regime, $T,\ L$ and δ are generally small compared to 1. This allows eqs. (2.38) to (2.40) to simplify to,

$$P_{abs}^{th} = \frac{Sh\nu_p}{2(\sigma_e-\sigma_1^L)\tau}(T+2\delta)(1-F)^{-1} \tag{2.45}$$

$$\rho_{\text{diff}} = \frac{h\nu_L}{h\nu_p}\frac{\sigma_e-\sigma_1^L}{\sigma_e+\sigma_1^L}\frac{T}{T+\delta}(1-F) \tag{2.46}$$

where

$$S = \frac{\pi(\omega_p^2 + \omega_L^2)}{2}$$

and F is assumed to be very small [86, 85].

2.6 Thermal Considerations

Thermal damage and stress fracture limits the achievable average output power in solid-state laser gain media [87]. Damage and fracture of the gain element is not the only problem imposed on laser systems by increased thermal energy. Thermal issues can affect resonator stability, wall plug efficiency, beam quality and cause birefringence, wavefront aberrations and lensing effects [87]. The following sections will outline these various issues along with their specific causes.

2.6.1 Heating in solid-state lasers

Background optical absorption of both the pump and laser outputs and nonradiative relaxation are the primary causes of heat generation in solid-state lasers [87]. In high quality gain media nonradiative decay dominates in this regard and in four level systems such as Nd^{3+} gain media, this decay is necessary from the pump bands to the upper lasing level and from the lower lasing level back to the ground state. This necessary decay is due to what is known as the quantum defect, the difference between the photon energy required to pump a laser and the photon energy of the desired output [87]. This excess energy is often the dominant source of heat generation in such systems. Other contributors to increased thermal effects include nonradiative sites, concentration quenching, upconversion and excited state absorption.

2.6.1.1 Quantum defect

The quantum defect of a laser material relates to the difference in energy between the pump photons and emitted photons. As electrons must be excited to an energy level that is higher than the metastable, upper lasing level, the emitted photon energies will be less that the energy of the photons required to pump the laser transition [88]. The quantum defect efficiency can be calculated as the ratio of the lasing and pump energies ($h\nu_L$ and $h\nu_P$ respectively) given by

$$\eta_S = \frac{h\nu_L}{h\nu_P} = \frac{\lambda_P}{\lambda_L} \tag{2.47}$$

The difference between these two energies is often lost by fast nonradiative processes [87, 89]. The closer eq. (2.47) gets to unity, the more inherrently efficient a system becomes.

2.6.1.2 Concentration quenching

In high doping concentration gain media energy transfer processes between laser ions can occur. This can also happen in gain media with lower concentration when laser ions cluster together, particularly common when a host material is used that has a low solubility of the dopant. These energy transfer processes can reduce the metastable state lifetime and if the transferred energy is then lost through nonradiative processes an increase in thermal energy occurs. It also causes gain reduction, as energy available for stimulated emission is lost to other processes, resulting in an increase in the pump power required to meet lasing threshold [90, 91].

2.6.1.3 Upconversion

When an ion emits photons with more energy than those it absorbs this is known as upconversion. This can occur through various excitation mechanisms that allow the excited electron to move to increasingly higher energy levels. One such mechanism is sequential absorption by means of ESA as discussed in section 2.5. After an ion absorbs a photon, exciting an electron to a metastable state, for some species such as Tm^3, higher energy states exist that are metastable and can be reached through absorption of the pump photon energy. This can be more of a problem for systems with high pump intensities or broadband pumping sources, permitting a wider range of available transitions [92, 93, 86].

Cooperative upconversion occurs when two excited ions undergo an energy transfer process, exciting one ion into a higher energy state and relaxing the other. This is more common in systems with high doping concentrations and can become a particular problem if clustering as discussed in the section on concentration quenching (section 2.6.1.2). Thermal issues arise from upconversion when the multiply excited electrons undergo nonradiative decay to reach the higher metastable states. This can also cause a considerable waste of pump energy producing high energy photons that are not of the desired wavelength [92, 93, 86].

2.6.1.4 Nonradiative sites

Nonradiative sites and the issues they can cause have been a matter of interest for some time, before a thorough understanding was developed they caused a great deal of discussion about the true quantum efficiency of some solid state laser systems. In Nd^3 lasers, measurement of the quantum efficiency was made difficult because of the affect nonradiative sites had on the radiative rate [94].

Laser ions when excited will undergo various energy transfer processes with their surrounding environment. If there are anionic impurities in a gain medium that are sufficiently close to excited laser ions, often highly efficient, energy transfer will take place between them. These impurities may then undergo nonradiative decay and are thus referred to as nonradiative sites [95]. This is disadvantageous on two accounts, reducing the energy available for stimulated emission and increasing the thermal energy in the laser system [89, 87].

2.6.2 Thermally induced effects

As mentioned in the previous section, not only does the generation of thermal energy reduce the energy available for gain, it can also have other detrimental effects on laser performance. Thermo-optical phenomena such as alterations to the refractive index can arise from the thermo-mechanically induced stress-strain dynamics. This occurs in the form of birefringence, brought on by thermal stress in the crystal [96, 97]. Inhomogeneous temperature distributions can even cause crystal fracture in solid-state media. More subtle effects however will be discussed here as these are far more commonly seen in normal laser operating conditions. A summary of the effects caused by thermal energy in solid-state laser systems is given in fig. 2.5.

2.6.2.1 Polarisation

Materials that are normally isotropic can, when under stress, exhibit birefringence [97]. It is possible for the thermo-mechanical stress experienced by a solid-state laser material to be sufficient to induce such birefringence [96, 97]. In systems whereby polarised output is not required then this is not an issue, however polarised output can be important for many applications. Figure 2.6 shows how a polarised beam is partially depolarised after passing through a crystal exhibiting birefringence.

This can lead to increased losses in the system. As every time a depolarised beam comes into contact with a polarising element within the cavity those rays that are not of the correct polarisation will be lost [96, 97]. Figure 2.6 also highlights how the beam polarisation is not altered along the x and y crystal axes. This can lead to the beam spatial profile altering, if sufficient depolarisation losses are experienced by the depolarised part of the beam then a cross shaped intensity profile will be observed.

Another effect, known as bifocussing, arises due to the polarisation dependence of the optical path a ray takes through a gain medium. Whilst pumped the gain medium will exhibit different refractive indices for the radially polarised parts of the beam and the tangentially polarised parts, illustrated in the depolarised beam in fig. 2.6. This will lead to two thermal lens focal lengths [96, 97]. Further discussion of this effect is made in the derivation in section 2.6.2.2.

2.6.2.2 Phase

Due to the temperature dependence of a material's refractive index, thermal gradients can lead to what is known as thermal lensing. Solid-state media will often exhibit a thermal gradient in the radial direction, whereby the outer edges of the medium will be cooler than at the beam propagation axis. As the refractive index of a medium is temperature dependent this causes a refractive index gradient in the same direction. The result of this is that the optical path length appears longer through the core of the medium than it does at the edges resulting in a lensing affect. This phenomena can be further exaggerated by thermo-mechanical affects on the refractive index, known as the thermo-elastic effect and mechanical stress causing the end faces of the gain medium to bulge so

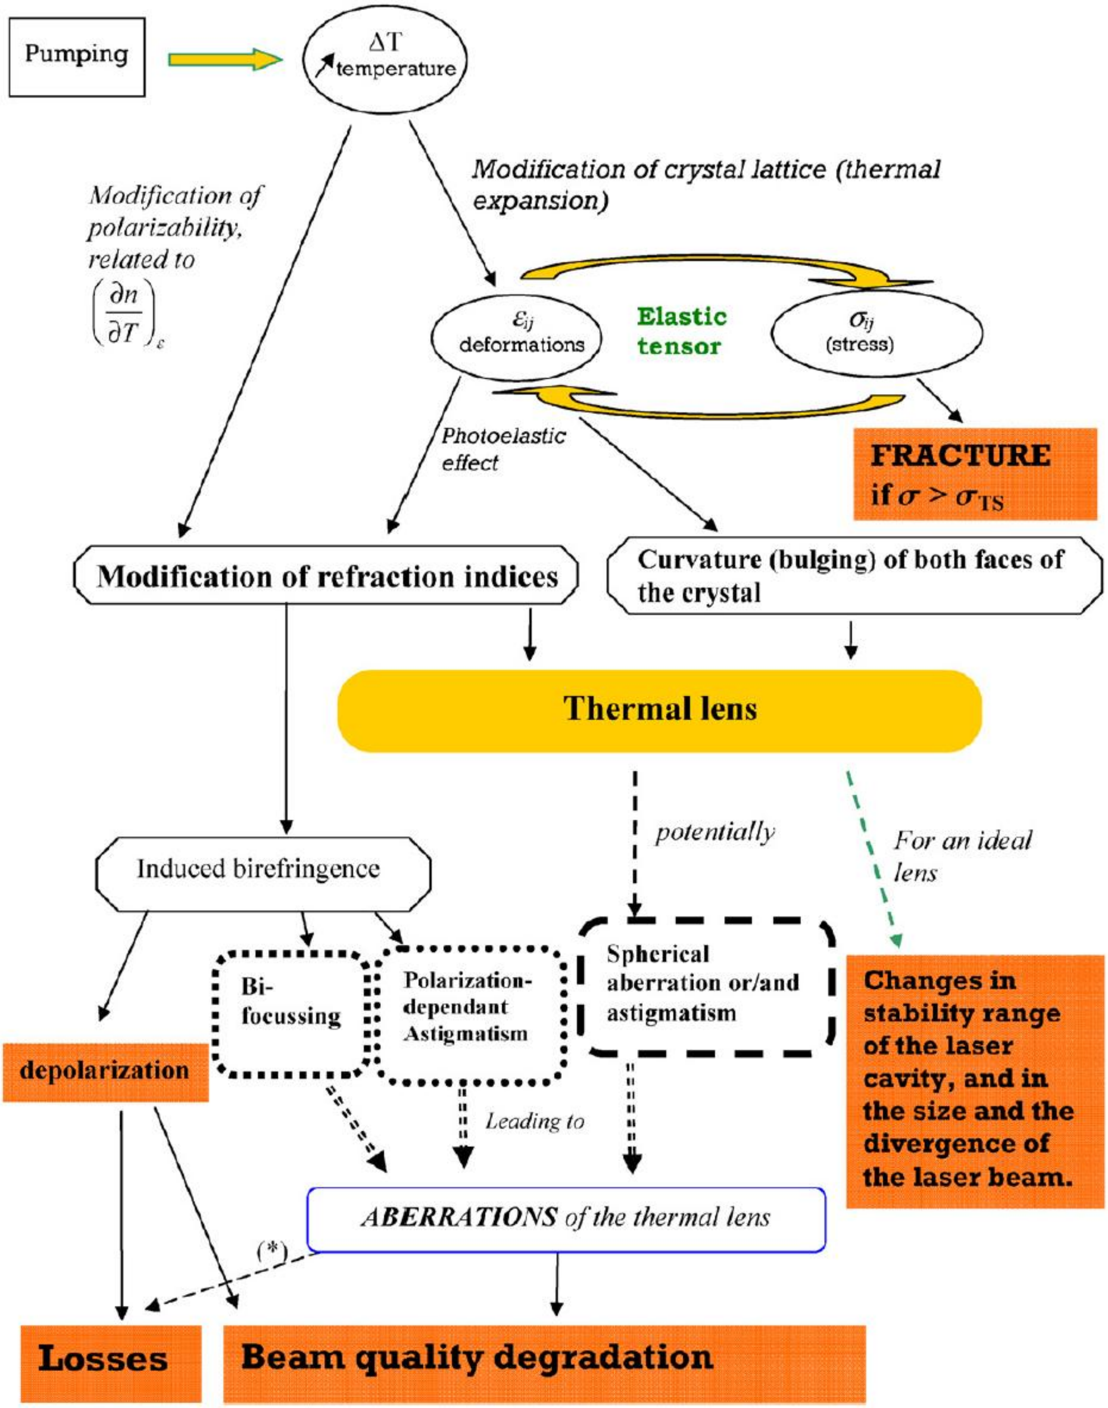

Figure 2.5: Summary of the thermal effects in solid-state lasers [96, 97].

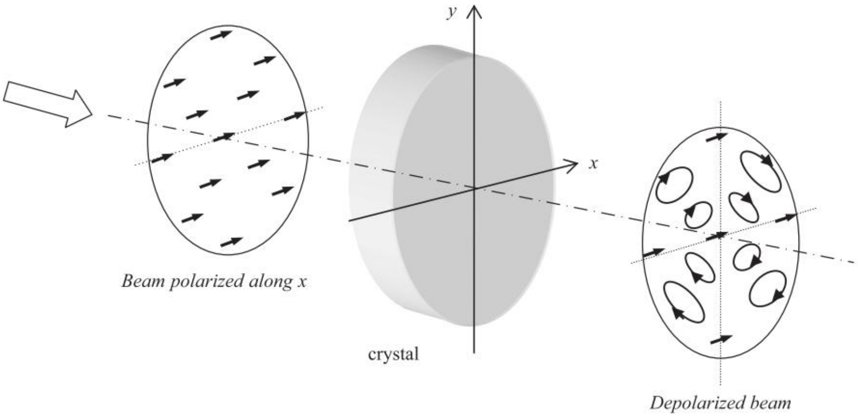

Figure 2.6: A polarised beam experiencing birefringence passing through a crystal under thermo-mechanical stress resulting in partial depolarisation [96, 97].

that it is in fact longer at the core [96, 97, 98].

Consider a cylindrical crystal of length L with refractive index n_0 at a temperature T_c. A ray travels along an optical path from the crystal face at $z = 0$ to a distance of $z = L + d$ where d is an arbitrary distance from the crystal. In the absence of external temperature or stress fields the optical path is

$$\delta^{\text{off}} = n_0(T_c)L + d \tag{2.48}$$

Turning the pump on, the optical path becomes

$$\delta^{\text{off}}_{r,\theta}(r) = \int_0^{L+\Delta L} n_{r,\theta}(T,\varepsilon)\,\mathrm{d}z\,\mathrm{d}r\,\mathrm{d}\theta + d - \Delta L(r) \tag{2.49}$$

and is now dependent on the lateral shift of the ray with respect to the crystal axis r, and on the direction of polarisation (which can lead to bifocussing). Here $\Delta L(r)$ is the expansion of the gain medium due to thermo-mechanical stress.

Expansion of the $n_{r,\theta}(T,\varepsilon)$ term leads to

$$\begin{aligned}\delta^{\text{off}}_{r,\theta}(r) = \int_0^{L+\Delta L} \Bigg[& n_0(T_c) + \left(\frac{\partial n_{r,\theta}}{\partial T}\right)_{\varepsilon}(T(r,z) - T_c) \\ & + \sum_{j=r,\theta,z}\left(\frac{\partial n_{r,\theta}}{\partial \varepsilon_j}\right)_T \varepsilon(r,z)\Bigg]\mathrm{d}z + d - \Delta L(r)\end{aligned} \tag{2.50}$$

In this equation the temperature derivative of the refractive index is only partial as calculated with constant strain. Changes in the rod length can be related to the axial strain ε_z by the equation

$$\Delta L(r) = \int_0^L \varepsilon_z(r,z)\,\mathrm{d}z = \langle \varepsilon_z(r) \rangle \tag{2.51}$$

Here angled brackets denote values intetgrated over the whole length (z-axis).
It can then be shown

$$\langle \varepsilon_z(r) \rangle = -\alpha_T(1+\nu)\langle T(r=0) - T(r) \rangle \tag{2.52}$$

Where ν is Poisson's ratio and α is the thermal expansion coefficient for the material.
Given that the rod length difference is small compared to the entire rod length ($\Delta L \ll L$) and the first-order terms in eq. (2.49) are much smaller than $n_0(T_c)$ the relative optical path length can then be written as

$$\begin{aligned} \delta^{\text{rel}}_{r,\theta}(r) &= \delta^{\text{on}}_{r,\theta}(r) - \delta^{\text{off}} \\ &= \left(\frac{\partial n_{r,\theta}}{\partial T}\right)_{\varepsilon} \langle T(r,z) - T_c \rangle + \sum_{j=r,\theta,z} \left(\frac{\partial n_{r,\theta}}{\partial \varepsilon_j}\right)_T \langle \varepsilon_z(r) \rangle \\ &\quad + (n_0 - 1)(1+\nu)\alpha_T \langle T(r) - T(0) \rangle \end{aligned} \tag{2.53}$$

This enables calculation of the path difference experienced by a ray travelling through a pumped gain medium relative to an unpumped medium. The optical path difference for two parallel rays at radial postions r and 0 is then

$$\Delta_{r,\theta}(r) = \delta^{\text{rel}}(0) - \delta^{\text{rel}}_{r,\theta}(r) = \chi_{r,\theta} \langle T(r) - T(0) \rangle \tag{2.54}$$

where

$$\chi_{r,\theta} = \left[\left(\frac{\partial n_{r,\theta}}{\partial T}\right)_{\varepsilon} + 2n_0^3 \alpha_T C'_{r,\theta} + (n_0 - 1)\alpha_T(1+\nu) \right]$$

Here $C'_{r,\theta}$ is a photoelastic constant, related to the refractive index shift in isotropic crystals under thermal stress and is a property of the gain material. This term explains the bifocussing, depolarisation and polarisation dependent astigmatism effects that occur due to thermal stresses. The $(n_0 - 1)\alpha_T(1 + \nu)$ term, refers to the bulging of the end faces. This occurs due to the compression of the inner part of the crystal as a consequence of thermal expansion of the lattice structure. It has been shown however that this term can predict too high a value if the diameter to length ratio of the rod is greater than 1.5 and therefore should be taken as a maximal value [98]. The $\chi_{r,\theta}$ term is referred to as the thermo-optic coefficient.

The temperature difference between the core and edges of the cylinder, integrated over the entire length of the gain medium is given by

$$\langle T(r) - T(0) \rangle = \int_0^L (T(0,z) - T(r,z))\,\mathrm{d}z = \frac{\eta_h P_{\text{abs}}}{4\pi K_c}\frac{r^2}{w_p^2} \tag{2.55}$$

Where η_h is the fractional thermal load, P_{abs} is the absorbed pump power, K_c is thermal conductivity of the material and w_p represents the radius of the pump beam waist. This equation highlights the quadratic dependence of the optical path difference on the radial position r, meaning the crystal acts similarly to a thin lens.

Using the paraxial approximation, the focal length of this thin lens can be calculated by

$$f_{\text{th}(r,\theta)} = \frac{r^2}{2\Delta_{r,\theta}(r)} \tag{2.56}$$

with a dioptric power of

$$D_{\text{th}} = \frac{1}{f_{\text{th}}} = \frac{\eta_h P_{\text{abs}} \chi}{2\pi K_c w_p^2} \tag{2.57}$$

eq. (2.57), however, only holds if the beam is not polarised or the photoelastic effect is negligible. Should this not be the case then the focal lengths will differ dependent on the radial and tangential refractive indices, in which case bifocussing will occur [96].

2.7 Conclusion

Designing high power laser systems is a complicated, multi-faceted process. This chapter presented some of the key considerations that must taken into account and introduce some of the methods and approaces used to model these phenomena. Building robust and accurate models allows the user to quickly investigate and iterate during the design process, faster and cheaper than physical experimentation allows. The following chapters each present some of the most relevant work specific to solutions being presented.

CHAPTER 3

Optimisation and Tolerancing

3.1 Introduction

in the design of complex optical systems, it is common to use purpose-built optical design software packages such was Zemax [6]. Zemax is an optical design programme used in the design and analysis of imaging and illumintation systems. Zemax uses ray-tracing to model the propagation of light as rays through an optical system. Aside from being developed for optical design workflows and optimised to perform large numbers of complex calculations in parallel, one of the main advantages of utilising Zemax is its extensive object, materials, and coatings libraries. These enable optical engineers to quickly model and test specific components, without having to manually encode their properties into the software.

Some systems are reliant on electro-optical processes that are not easily modelled using standard optical design software approaches, such as Zemax [6], for example those optical systems utilising polarisation phase retardation to control power transmitted through optomechanically Q-switched cavities. Introducing waveplates to manipulate phase retardation becomes essential to deliver this control. Precise alignment and calculation of waveplate retardation become key parameters when optimising these systems, the theory for which is discussed in section 3.2.

Zemax has several in-built tools for optimisation and tolerancing, but none of these adequately handle the retardation values of polarisation waveplates. Furthermore, the tools presented here are designed to be used in conjuction with those within Zemax to give a broader understanding of the optical systems being modelled.

This chapter presents a set of tools and analyses to enhance the design process of optical systems with many degrees of freedom. These tools help to find configurations to achieve optimal performance, but also highlight problematic components in a design and carefully consider the sesnitivity to misalignment of any configuration. An important aspect of any system that is to be manufactured are the permissible tolerance ranges for each of the components. In the case of optical systems even very small perturbations of alignment can have critical effects on output performance. Understanding which parameters can afford to have less stringent tolerance ranges, enables designers to better focus their efforts and looser tolerance ranges often lead to reductions in manufacturing costs.

By utilising the ray-tracing capabilities of Zemax [6] and automation, through an interface with MATLAB [5], it is possible to search large numbers of configurations to find these optimal designs. These tools are capable of utilising custom merit functions to determine optima, by querying multiple outputs from the model in Zemax, for each configuration. The process of how this automation is achieved is presented in section 3.3.1.

The analyses presented in section 3.3.2 offer the user a quick method for comparing the effects of multiple parameters. Reducing the n-dimensional output data into a graphical format enables direct comparisons between every degree of freedom being considered. The algorithms for how these analyses are carried out are discussed in section 3.3.2.

To verify these tools are applicable to real systems, and to give examples of the analyses introduced, a test case is presented in section 3.4. Here an example of a typical military, polarisation sensitive, laser is modelled and then reconstructed in the lab. The results are compared to demonstrate the accuracy of this approach. They are then analysed using the methods presented in section 3.3.2 and the results discussed. Another major advantage of this computational approach is the time saved compared to the laboratory testing. The test case in section 3.4 is limited to only five parameters due to reasonable time constraints, as large numbers of configurations are required to be tested and each new parameter increases the number of test cases exponentially. Even in the simplest of cases, where each parameter had only two settings, each additional parameter would double the effort required to test all cases, growing at a rate of 2^p, where p is the number of parameters. The computational tools do not suffer from this limitation as they can be left to run without supervision. This means it is possible to investigate far more complex systems this was, an example of which is given in section 3.5. Here 16 parameters are analysed and the findings discussed to highlight how this tool set enables the exploration of complex optical design problems.

It will be shown that the methods presented in the following sections will allow the users to test the effects alignment, configuration and polarisation have on complex optical systems. With many systems, such as the one presented by Lee et. al [99] there are many degrees of freedom that can significantly impact performance. Tightly controlling the alingment of these components during the manufacturing process ensures performance matrics are met, but this is costly and in many cases these constraints are perscribed from the experience of experts. These models inform the user which combinations of misalignments for the many components have the most significant impact on performance metrics. By identifying these combinations it enables the user to adjust designs to mitigate them accordingly. The ultimate goal being, to relax the constraints on tolerance ranges which should, in turn, reduce the cost to manufacture each of the individual components.

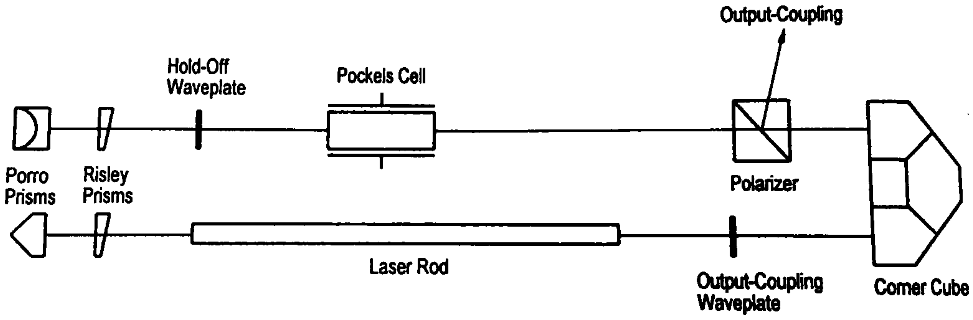

Figure 3.1: The expanded laser setup that the verification design was based on, adapted from [101].

3.1.1 Laser Cavity

This section presents the optical system used for experimental verification of the modelling approach, discussed further in section 3.4. The laser cavity is presented here to illustrate the types of optical systems this modelling approach was initially designed for. For the experimentl verification, sufficient complexity required to demonstrate the effects outlined in section 3.2 had to be balanced against the difficulties of highly-dimensional problems.

Verifying the model using a laser that was well understood and had general application outside the scope of this experiment was also an important consideration. This showcases the general application of the modelling technique, one that is not bound to specific systems in restricted environments. The setup shown in fig. 3.2 is based on the cross-porro prism resonator fig. 3.1, adapted from [100, 101].

The cross-porro prism laser setup in fig. 3.1 is highly polarisation dependent. The corner cube and porro prisms all induce phase-retardation, which will affect the output coupling for the laser. This means the two waveplates must have carefully calculated retardation values to correctly control output-coupling [100, 101].

Typically used in military applications, such as range finders and target designators, designs such as these are often used when environmental conditions are less than ideal and reliability is more important than optimal system performance. Systems such as these have been developed in response to extensive testing and application in the field. The self-aligning behaviour of retro-reflectors maintains performance over large temperature ranges and makes them insensitive to shock and vibration induced misalignments [100, 101]. The introduction of the corner cube in this case allows for folding of the resonator. Decreasing the total length is highly desirable when such systems are to be used in the field.

To achieve this improved reliability, however, precise alignment of the components must be achieved in the construction of the cavity. To demonstrate how sensitive these lasers can be to the alignment of individual components, and how the modelling techniques presented in this chapter can be applied to the design process, a simplified version of the output-coupling optics have been tested. Figure 3.2 shows the experimental setup used in verification.

A non-polarised 6 mW probe laser will emit a beam incident upon a polarisation beamsplitter. This will cause

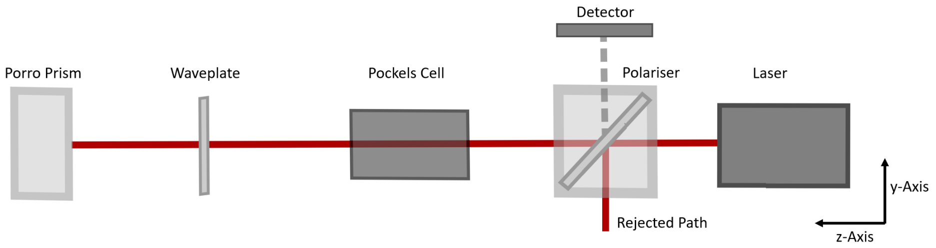

Figure 3.2: Diagram of the example system used for demonstration of this analysis based on the cross-porro prism resonator fig. 3.1, adapted from [100, 101].

approximately half of the emitted light to be reflected along the rejected path. The other half will transmit through the beamsplitter where it will next be incident upon a waveplate. This waveplate will cause a quarter wave retardation of the beam's polarisation state. The beam will then be reflected by a porro prism and make a second pass through the waveplate, causing it to be reflected by the beamsplitter along the output coupling path. Finally a detector will be placed to measure the transmission through the system.

This setup could readily incorporate a gain medium between the porro prism and the waveplate, allowing it to act as an amplifier stage of an optical system as seen in [100, 101]. This has been omitted in this case as it is not required to demonstrate polarisation retardation effects.

3.2 Modelling Polarisation in Optical Systems

Ray-tracing programs usually treat rays geometrically, considering position, orientation and phase. This means, that at a surface, each ray can be completely described by the intercept coordinates, the direction of travel (often expressed as direction cosines) and the phase [6]. As ray-tracing is primarily concerned with determining ray paths through a system, Snell's Law [102, 103] is enough to calculate a ray's refraction. Amplitude and phase need not be considered as they do not affect beam direction. For laser design however, both of these properties are very important.

Polarisation analysis extends ray-tracing programs, enabling them to consider the effects of optical coatings and interface reflection and absorption. This section will discuss how polarisation analysis is used in the context of ray-tracing and incorporated into the initial model.

3.2.1 Waveplates

Calculating the effect waveplates have on the polarisation state of an optical system can be achieved using Jones Matrices. The elements of this matrix can be calculated using eq. (3.1), where Φ_x is phase offset of the electric field.

$$M_R^P = \begin{pmatrix} 1 & 0 \\ 0 & \exp[i\delta] \end{pmatrix} \qquad \text{with} \quad \delta = \Phi_2 - \Phi_1 \tag{3.1}$$

Using eq. (3.1) it is possible to calculate the effect a waveplate has on a propagating wave using Jones Calculus [104, 105, 106, 107, 108]. Equation (3.2) gives the Jones vector for the electric field $\vec{E}$, describing the field components in the x and y axes:

$$\vec{E} = \begin{pmatrix} E_x \\ E_y \end{pmatrix} = \begin{pmatrix} E_{0x}e^{i\delta_x} \\ E_{0y}e^{i\delta_y} \end{pmatrix} \tag{3.2}$$

where E_0x and E_0y are the electric field amplitudes, and δ_x and δ_y are the phases in each axis. Jones Calculus is a powerful tool for analytical modelling of optical systems as it allows the calculation of the electric field after an element by multiplying the electric field vector with the relevant Jones Matrix for that element [104, 105, 106, 107, 108].

3.2.2 Initial Model

Retardation inducing waveplates modelled in Zemax by placing a Jones Matrix that behaves equivalently to the desired waveplate into the model [6]. Calculating the elements of this Jones Matrix is done externally to Zemax using eq. (3.1), which must then be manually input prior to each ray-trace. This process is time consuming for the user and is more prone to input errors. Once input, however, this enables accurate modelling of how waveplates, with differing retardations, affect the performance of optical systems. It must be noted that the Jone Matrix element is a two-dimensional construct in Zemax and will apply the effects of the waveplate element to all rays that pass through it, irrespective of any differences in path they might take through a three-dimensional waveplate. Despite this, the Jones Matrix treatment of waveplates in Zemax gives good agreement with experimental results, as show in figs. 3.3 to 3.4.

The waveplate's ultimate effect on the polarisation state is also dependent on its rotational alignment. Therefore, to test rotational behaviour, each angle must be adjusted individually for each axis of rotation and tightly controlled rotational tolerances must be enforced. This is both a computational and experimental issue as these adjustments take time and can be prone to error if done manually. Examining multiple axes of rotation, for each component, quickly escalates the dimensionality of the problem, introducing associated complexities.

Figure 3.3 shows the results of modelling the rotation of a waveplate in Zemax for one retardation value and how it compares to experiment. As can be seen, the results closely match that of experimental values and as such, this type of modelling can be considered a worthwhile method of prediction for certain component behaviours.

To generate computational results presented in fig. 3.3 the waveplate object in Zemax was rotated through 180° in 5° increments and the transmission through the system was measured on a detector object at the end of the

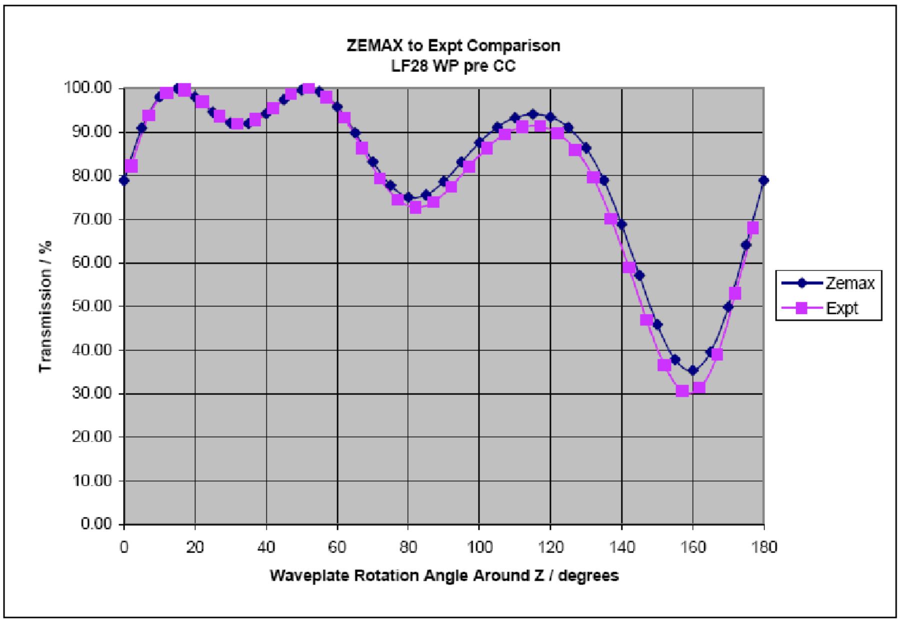

Figure 3.3: Single retardation value modelled for full range of rotations and comparison to experimental values.

optical path. Using this approach is was necessary to manually copy transmission values from Zemax and input them into a graphing program (Excel in the case of fig. 3.3) for every ray trace. Due to significant manual effort required, testing large numbers of rotation and retardation combinations becomes infeasible. The resource bottleneck becomes dependent on the user's interaction speed instead of computational speed or model efficiency.

Trying to remedy this bottleneck was the initial consideration for this project. The solution to was gradually developed and expanded upon to comprise this chapter.

3.2.3 Initial Findings

Whereas the modelling data represented in fig. 3.3 took a considerable amount of time to generate for a single retardation value, fig. 3.4 is a surface map of every combination of retardation and rotation values. Furthermore, it was generated automatically, with less manual interaction required by the user than the single slice view in fig. 3.3. The blue slice depicted in fig. 3.4 is the point at which the retardation, and therefore data, matches fig. 3.3 (although phase shifted). Not only is far more data produced using this method, presented in detail in section 3.3, it is quicker to do so than the previous work flow. This whole process is automated, meaning the user can simply input the ranges and step sizes they are interested in once, and leave the computer to generate the data. The plotting and data manipulation functions available in MATLAB also make it possible to visualise in the manner presented in fig. 3.4, which Zemax would otherwise be incapable of.

It is from this initial case that other analysis methods have been developed. Their details and outputs will be discussed in section 3.3.2. How the tool set generates data for analysis is the topic of section 3.3, the algorithm will be presented and discussed alongside some of the technical details of its implementation. In section 3.6 the tool set will be discussed in a general manner, with special attention given to its strengths and weaknesses as well as how it could be extended and developed further.

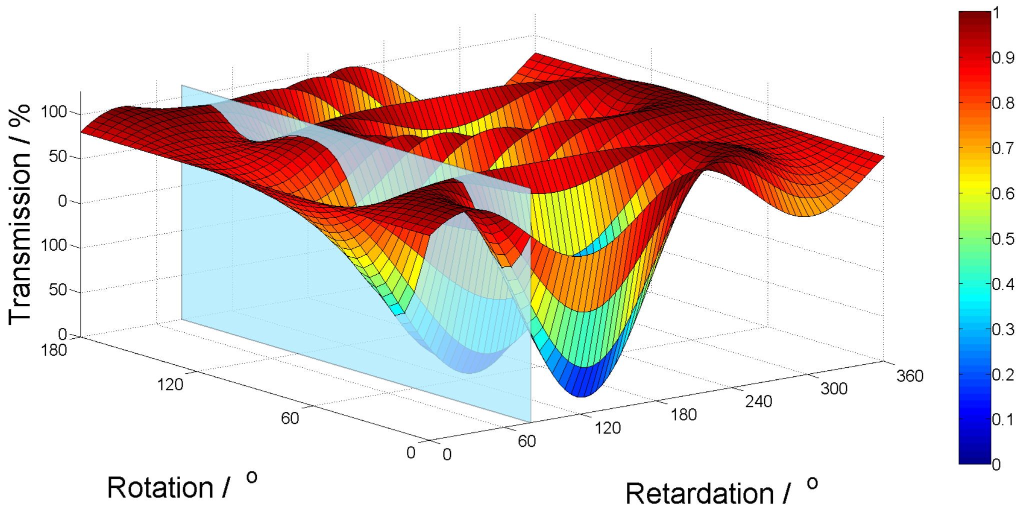

Figure 3.4: Transmission percentage through waveplate for full ranges of both retardation and rotation modelled, with slice drawn at point where rotation matched with single rotation value in fig. 3.3.

3.3 Methods

This section discusses how the data is generated for the analyses presented in section 3.3.2. The algorithm used to automate Zemax for data collection is presented, alongside brief consideration of design decisions made during its implementation. Finally the experimental design and setup is discussed. Due to the nature of the analyses, and dealing with large numbers of parameter combinations, a large number of measurements were required to be made to verify the approach. This meant a lot of effort put into experimental design was essential to reap adequate results, as discussed in section 3.4.1.de

3.3.1 Data Collection

All of the analysis done utilising this tool kit begins with the same algorithm for generating data. A description of the algorithm will be the subject of this section. Figure 3.5 depicts a flow chart of the process.

The first and most crucial step in the process of modelling any optical system is to generate the model within the Zemax environment. Any analysis done using these tools will be dependent on the accuracy of this model. In fig. 3.5 this has been compressed to a single step as it may often be the case that the model is generated by another user and generic Zemax modelling is not the topic of this thesis. Here it is assumed that accurate models can be generated prior to the invocation of these methods.

Once the model has been generated the components of interest must then be identified. Each component will have a number of parameters relating to a degree of freedom in the manufacture and alignment of that component, for example tilt in each of the cardinal axes. Each object to test and the properties of those objects must be clearly defined. It is then important to define the range over which to vary these parameters. For the purposes of verifying this model experimentally, the potential configurations of the system will be generated using simple

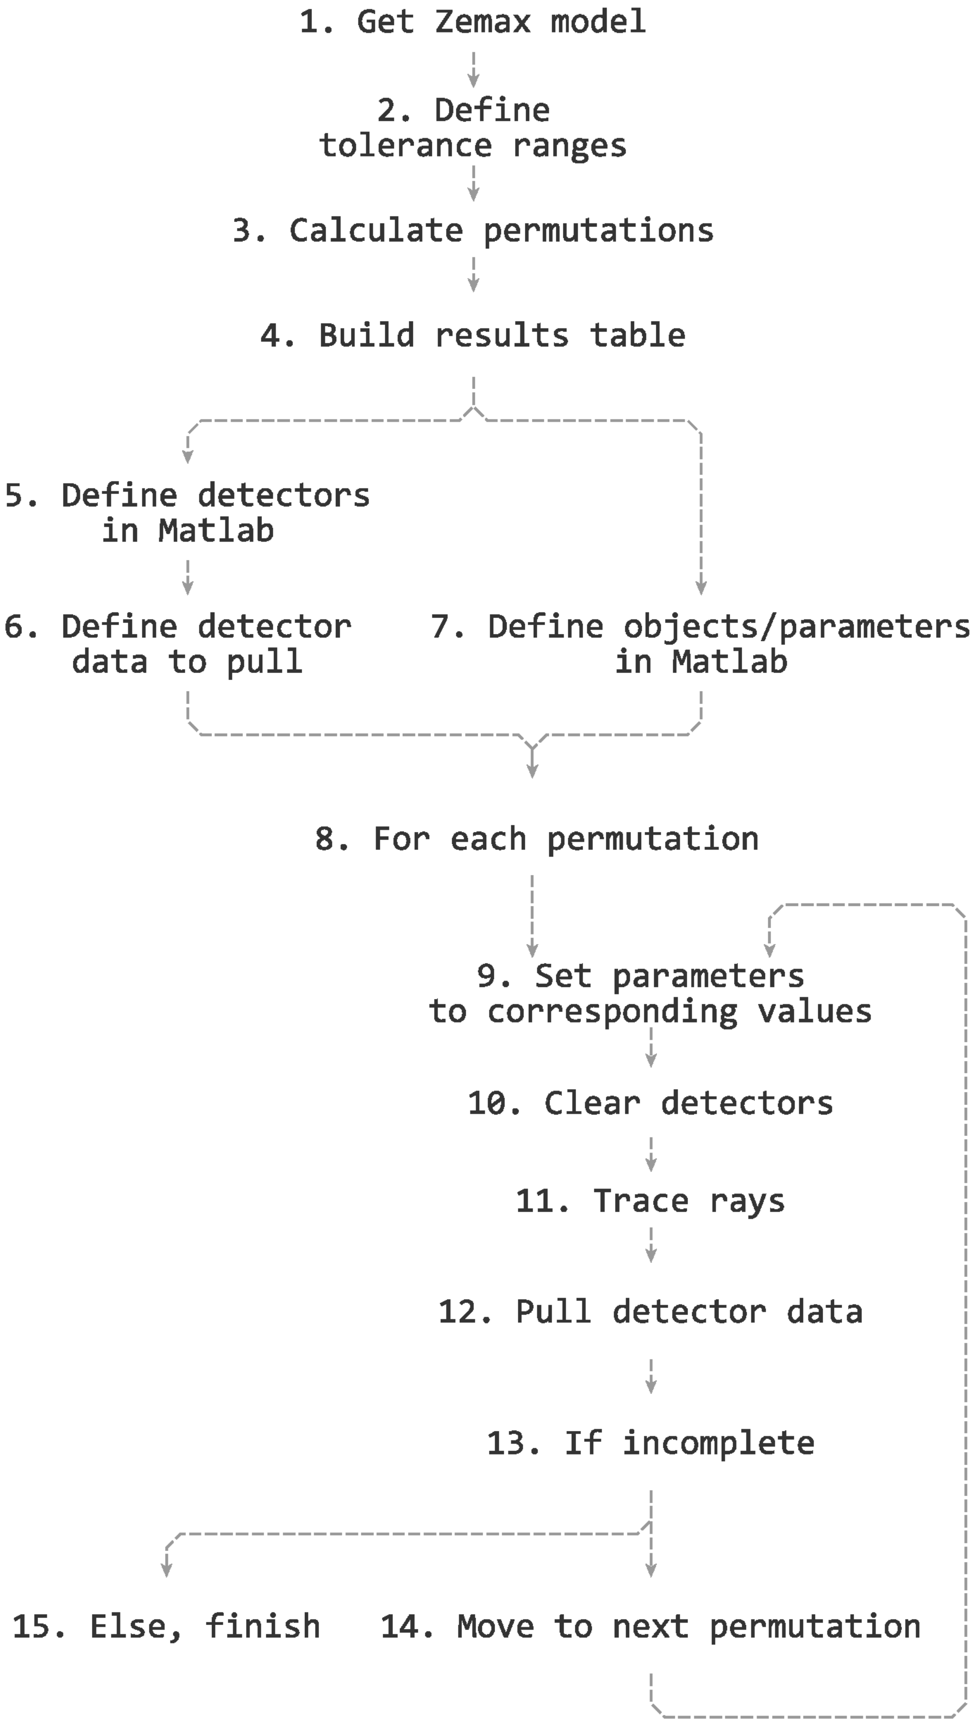

Figure 3.5: Algorithm for gathering data from Zemax.

Table 3.1: Example Zemax inputs table.

Input	Mirror X-tilt	Mirror Y-tilt	Waveplate X-tilt	Polariser Z-tilt
1	-2.0	-2.0	-1.0	-1.0
2	-2.0	-2.0	-1.0	0.0
3	-2.0	-2.0	-1.0	1.0
4	-2.0	-2.0	0.0	-1.0
5	-2.0	-2.0	0.0	0.0
…	…	…	…	…
225	2.0	2.0	1.0	1.0

interval sampling for each component. A set number of samples will be taken, at regular intervals across the entire tolerance ranges for the selected components. For example it may be required to test the system tilting a lens in the X-axis by ±2 degrees from the normal and taking sample points every $0.5°$ degrees. Once these definitions are made they must be translated into the corresponding object and parameter numbers within the Zemax environment.

Each unique object-parameter can be viewed as a single dimension in the test space. The total number of permutations (P) of the system to be tested is given by eq. (3.3)

$$P = \prod_{1}^{N} \frac{t_{\text{max}} - t_{\text{min}}}{s} \tag{3.3}$$

where N is the total number of object-parameters, t_{max} and t_{min} are the maximum and minimum values in the tolerance range and s is the step size for each dimension. An inputs table is then formed in MATLAB where each column corresponds to a degree of freedom in the system, in this case a parameter-object combination. Each row is then a unique combination of values for each of the degrees of freedom to be tested. Table 3.1 is an example of what this inputs table might look like for four degrees of freedom totalling 225 unique parameter combinations. With this inputs table then a results table can be made where each column represents a metric or detector to be measured in Zemax and each row is a result from the corresponding row in the inputs table allowing for easy reference of data values to the object-parameter values that produced them.

With the parameter values determined it is now required to define which columns in the inputs table correspond to which object-parameter numbers in Zemax. It is also necessary to define which objects in the Zemax environment MATLAB will query as a detector. Once this is done, the Zemax object can be inspected to ascertain appropriate data types to query from it, which can then be done using the associated command from the MATLAB-Zemax automation toolkit, developed as part of this thesis and underpins this work.

The remainder of the algorithm in fig. 3.5 depicts the process of generating all of the data. For every point in the results grid, firstly MATLAB assigns the appropriate value to each object-parameter. It then ensures the detectors are clear and runs a ray trace, the detector data is pulled from Zemax into MATLAB and placed into the correct point in the results table. This process is then repeated for every point in the results table until all P permutations have been tested and stored.

It is clear to see that this process is very general and can be used to produce large amounts of data with minimal input. The sampling scheme can easily be adapted to whatever best suits the current model and data requirements.

3.3.2 Analysis

This section will describe in detail some of the analyses performed using the data generated by the algorithm in section 3.3. Aiming to maximise transmission, it is crucial to understand how the tolerance ranges on key components affect the polarisation state of the light through the system. This allows for identification of the components and parameters that lead to the greatest change in transmission and therefore indicates which tolerance ranges need to be most tightly controlled. Each of the methods described in the next subsections start from the finishing point of the data generation algorithm as stated above.

3.3.2.1 Finding Maxima

Determining the optimum state in a real optical system can be incredibly difficult. In the base example presented here, maximising transmission through the system is the primary goal. Finding the highest value within the results is trivial using MATLAB but the maxima may not be the optimum solution. Some permutations may have high transmissions at the given point, but even a slight variation from that set of parameters could cause large variances. To manufacture a system with parameters so close to a maxima as this would require very tight tolerance ranges, increasing the time and cost to required to build the system with the necessary precision. The purpose of this analysis is to find a result within the set, that meets a set of conditions and continues to meet those conditions for a given number of steps away from that point in each dimension. Figure 3.6 shows the algorithm used here to perform this task.

Firstly criteria must be set to search for results within the set. In the example given here the highest point above a threshold value of 97.50 % was defined as the search criteria. Any logic can be used here though to determine the subset, a figure of merit could be used to define which results where interesting depending on how they were generated initially. Once this is done MATLAB can seek matching results and test further for stability.

How far to step away from the point must then be defined for each dimension. This can be done individually for each, or a single value can be set for all. It is important to note here that the greater the number of steps taken in each dimension, the longer this analysis will take. Large numbers of steps not only increase the analysis time, they also reduce the likelihood of finding a stable point. A balance must be found for highly sensitive systems, ultimately it is the distance from the test point that is important for manufacturing purposes but steps sizes that are too large might miss variances that have significant impact on performance.

Once these parameters are defined MATLAB can then perform the search for a stable point. This typically begins with the results set being searched in descending order when sorted by output of the merit function. As mentioned earlier, the subset can be ordered in any way deemed appropriate, or not at all. The algorithm then

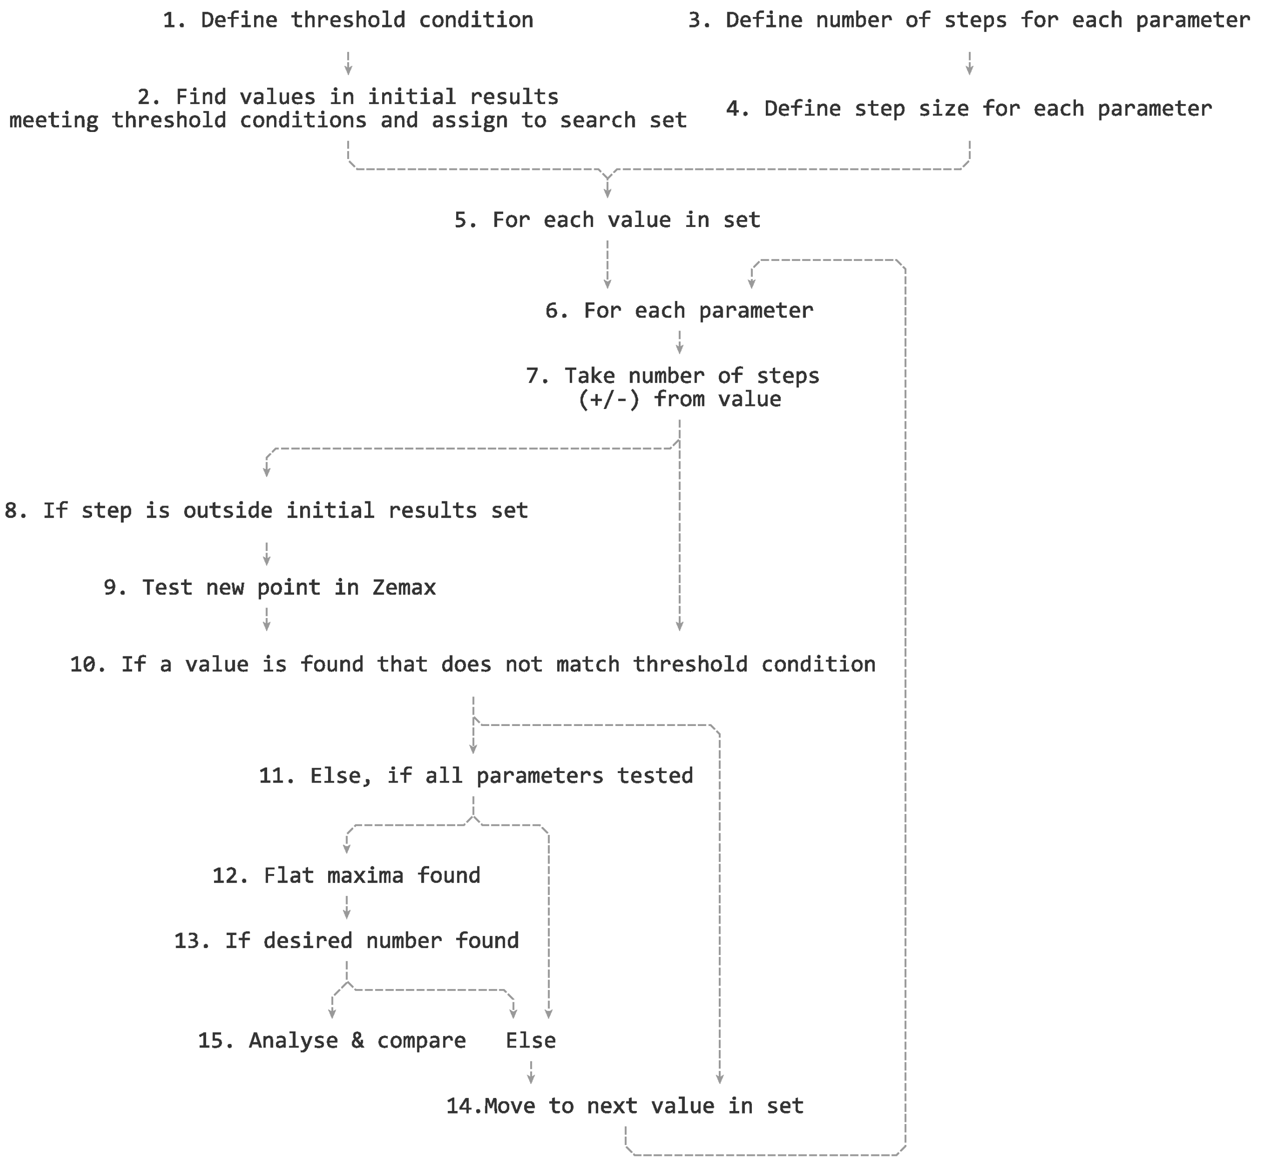

Figure 3.6: Algorithm used to find maxima that match threshold conditional for given number of steps in all dimensions.

starts at the first point, which will be referred to as the test point, in the subset and queries the input parameters. Then, for each dimension it queries every point plus or minus the given number of steps for that dimension. Should the algorithm encounter a point within the step boundary that doesn't meet the performance criteria, the test point does not meet the stability condition set earlier. The algorithm will then move onto the next test point.

It may be the case that a point within the given number of steps of a test point is not within the initial data set. If the user is interested in these extreme values outside of the initial data set, it is possible for MATLAB to quickly send a series of Dynamic Data Exchange (DDE) [109, 110] commands to Zemax to run a ray trace for that configuration. This allows for the ad hoc discovery of new stable points that may not have been initially considered.

If the algorithm finds that all points within the given number of steps from the test point for each dimension

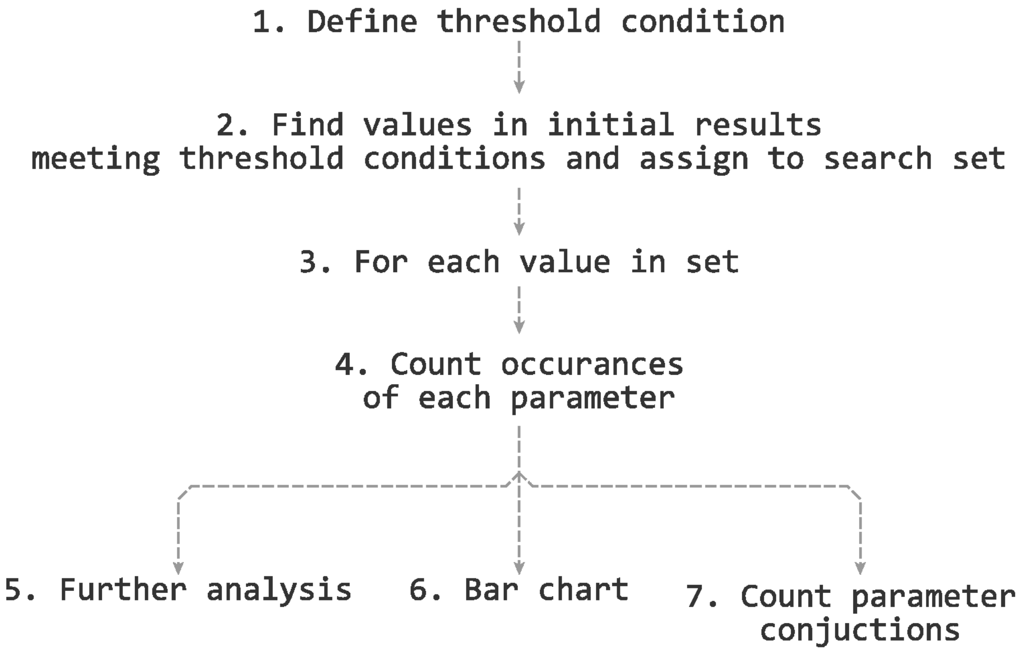

Figure 3.7: Algorithm for analysis comparing frequency of unique parameters to results of interest.

meet the threshold criteria, then a stable result has been found. At this stage, the algorithm can terminate and return the result, or continue until a predetermined stopping condition set by the user is met. If more than one result is found, further analysis can be performed to rank the results by user set criteria.

3.3.2.2 Parameter Frequency

This method aims to compare input variables amongst results of interest from the data set, to see if any particular inputs occur more frequently or are more strongly correlated to these results. This can then act as a first step in discovering why certain behaviours arise in complex systems and can assist engineers in isolation of causes. The algorithm is shown in fig. 3.7.

The user must then assert conditions that define the subset of results that are to be investigated further for commonalities. For example all results containing a transmission above or below a threshold value could be searched for to make up the subset of results. The wrapper function created in MATLAB has been written to allow for such searching to be done for one result output with ease. It is possible however to utilise any search criteria that satisfies boolean output simply by changing the search logic in the wrapper function. Realising far more complex search cases is therefore straightforward.

Once the search criteria have been determined, the subset of results that match are then collected. The algorithm then steps through each of the results in this subset and counts each occurrence of a parameter. These count values are stored for every parameter in every dimension. This allows for comparison between steps within a dimension and also between parameters.

This analysis is discussed alongside an example in section 3.4.3.

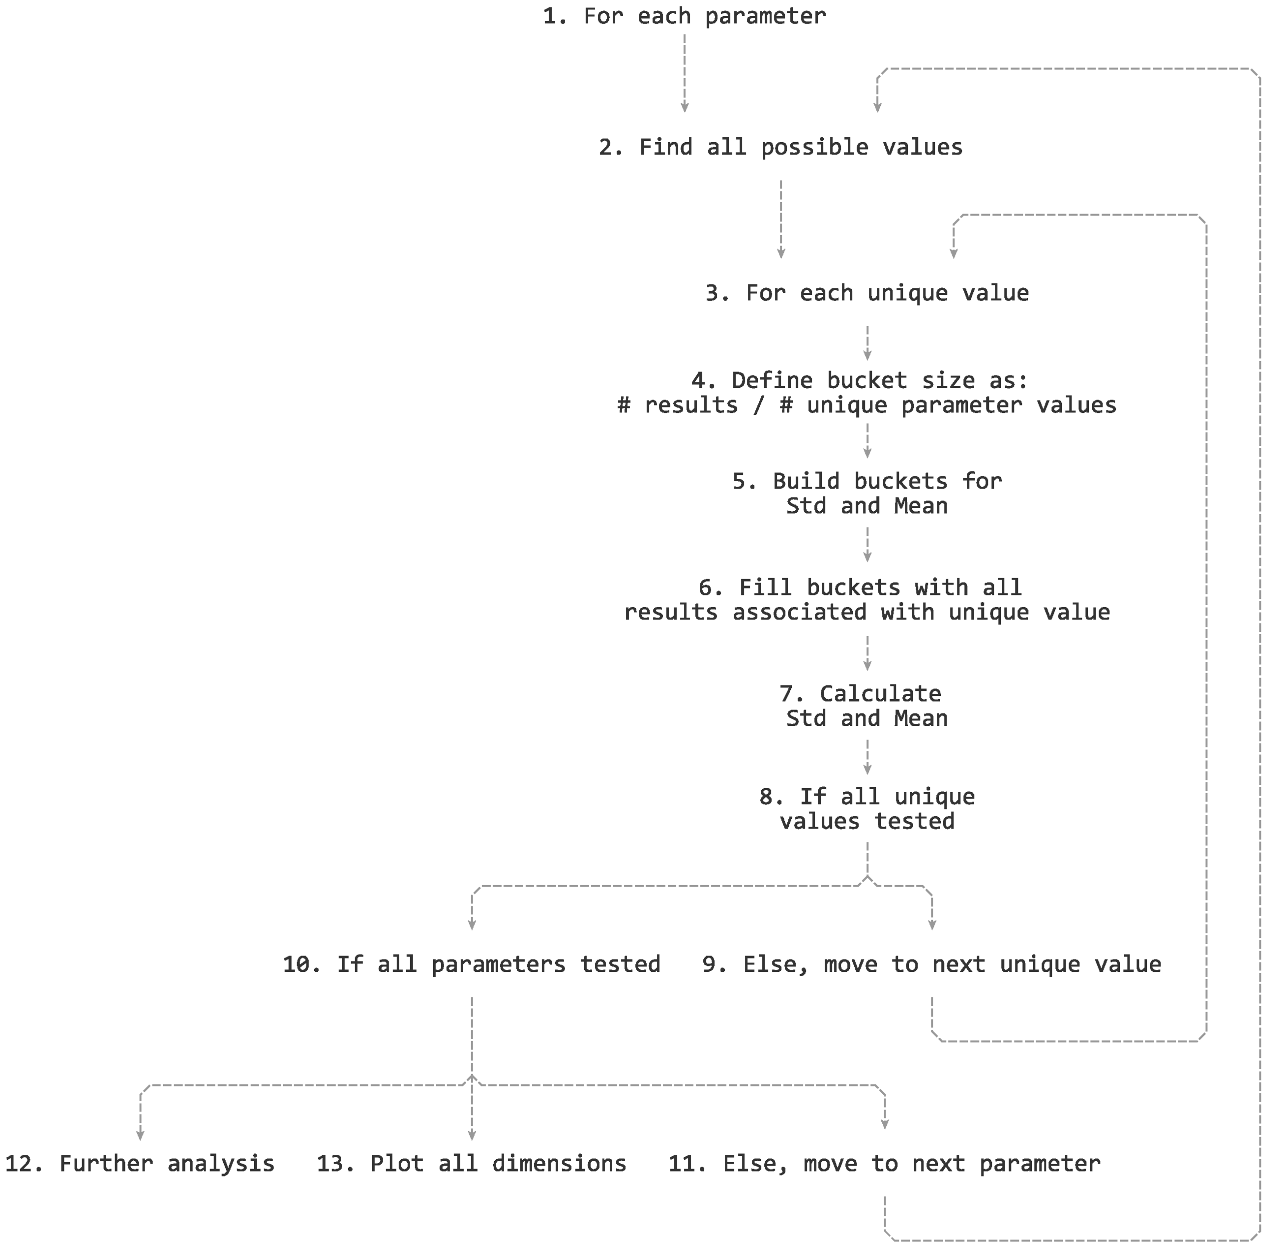

Figure 3.8: Algorithm for comparing variance of standard deviations and means of each parameter across unique results.

3.3.2.3 Variance Dominance

Discovering which variables lead to the highest variance in results can allow for a number of optimisations. That is the aim of this analysis. It is beneficial not only for understanding of the system but also to aid in further modelling efforts. The algorithm utilised to perform this is given in fig. 3.8.

Firstly the results and input data is taken and organised into a suitable format to allow for ease of processing. It is then required to determine the total number of dimensions and number of steps in each. Once this has been done the analysis is done on each dimension independently as they are looped through.

For a given dimension, the parameter values are stored. It is then required to build buckets to store the parameter specific result data in. For each parameter there will be a number of results equal to the total number of results divided by the number of steps taken in the given dimension. This subset of the results will therefore contain all permutations of the system whereby the given dimension was set to the current parameter value. In this respect, similarities can be drawn to Monte-Carlo analysis, in that for this subset of results one parameter is set and all others are in a variable state.

Further to this, smaller buckets must be created to hold the analysis results. Two of equal size for each dimension, large enough to contain one value per parameter. Once the subset has been determined the mean and standard deviation of all of the results are taken and stored in the smaller buckets. This is done for each parameter independently and the process is repeated through each dimension.

This analysis is discussed alongside and example in section 3.4.3.

3.4 Case 1: Experimental Verification

It was necessary to perform experimental verification of the modelling approaches presented in this chapter. An optical system that exhibited transmission dependant on polarisation was required. This system is described in detail in section 3.1.1.

One of the key objectives of this work was delivering a toolset that quickly and accurately predicts the output of complex, polarisation sensitive systems, with many parameters. This meant any verification would need to consider multiple parameters and many configurations. Section 3.4.1 discusses the decisions that were made in designing a sufficiently complex system, that could feasibly be tested.

3.4.1 Experimental Design

The purpose of this experiment was to demonstrate the accuracy and applicability this modelling approach has. Considering many degrees of freedom when optimising complex, polarisation sensitive, optical systems is one of the toolset's main advantages. Verifying the model's application to highly dimensional problems, in an experimental context, requires careful consideration of the experimental design.

Sufficient complexity must be shown to fully test the model, but considering many degrees of freedom leads to far greater numbers of configurations to be tested. Each new configuration introduces more potential for error. Misalignment of components during adjustment or data recording inaccuracies increase in likelihood as the experiment grows. Furthermore, the time required to perform the experiment increases polynomially with each new component-parameter introduced. Balancing the requirement for complexity with a reasonable set of input parameters was the focus of the experimental design.

The rotation of the waveplate about its z-axis has a significant effect on the transmission through the system. Due to this, the components were rotated about their z-axis, along the beam path, until maximum transmission

Table 3.2: Parameters used in the experimental verification of the model presented in chapter 3.

Component	**Parameter**	**Values (degrees)**	**Step Size (degrees)**
Polariser	x-Axis rotation	-2 to +2	1
Pockels cell	x-Axis rotation	-2 to +2	1
Pockels cell	z-Axis rotation	-2 to +2	1
Porro prism	y-Axis rotation	-1 to +1	1
Waveplate	Retardation	90, 180 and 289	-

through the system was observed. This alignment was then used as the basis for the rotations to be tested.

Table 3.2 lists the parameters that were varied during the experimental process. Five degrees of freedom were selected for the purposes of this experiment. Essentially the same experiment was performed three times, once for each waveplate retardation value. The order in which the parameters are listed in table 3.2 is also the order in which each of the parameters were adjusted through their full range. It should be noted that, for this system the rotations are not abelian. For each waveplate, the system was aligned to a defined base state as per the model. This state corresponded to a state whereby each of the components had zero degree rotation. These transmission values were in good agreement with the model.

Once aligned to the base state, each of the components was then rotated to the lower bound of its given range, for example -2 degrees in the case of the polariser. This was considered the first configuration and the transmission through the system was recorded. The polariser was then rotated in 1 degree increments to +2 degrees, with transmission recorded for each configuration. Once at the top of its range of values, the polariser was returned to the bottom of its range and the Pockels cell was rotated a single degree to its next value and the process repeated. Each component was adjusted in this order to minimise the number of adjustments reducing potential for introducing alignment errors. Each of the three waveplates, having retardation values of 90°, 180° and 289°, was tested in a total of 375 configurations.

The parameters in table 3.2 were chosen for having a significant effect on the transmission through the optical system and could easily be adjusted. These input parameters lead to a total of 1125 unique configurations.

3.4.2 Analysis and Discussion

All of figs. 3.9 and 3.11 are figures of the same data. Figure 3.9 presents the data for each of the waveplates separately, in the order that it was taken (as detailed in section 3.4.1). As can be seen from these three graphs, the data collected from the Zemax models (depicted as solid lines) is close to the data observed experimentally (shown here as coloured crosses). Some parameters affected the system transmission far more than others. The cyclic appearance of the graphs occurs when dominant parameters were adjusted from the higher, back to the lower end of their value ranges. This is most apparent for the 90° and 180° waveplates, where three large steps occur in the data at configuration 126 and 251. These steps correspond to the Pockels cell z-axis rotation being rotated from +2° to -2° , as the porro prism was rotated to its next state. The Pockels cell z-axis rotation did not exhibit the same dominant effect in the case of the 90° waveplate, as shown by the absence of this cyclic behaviour

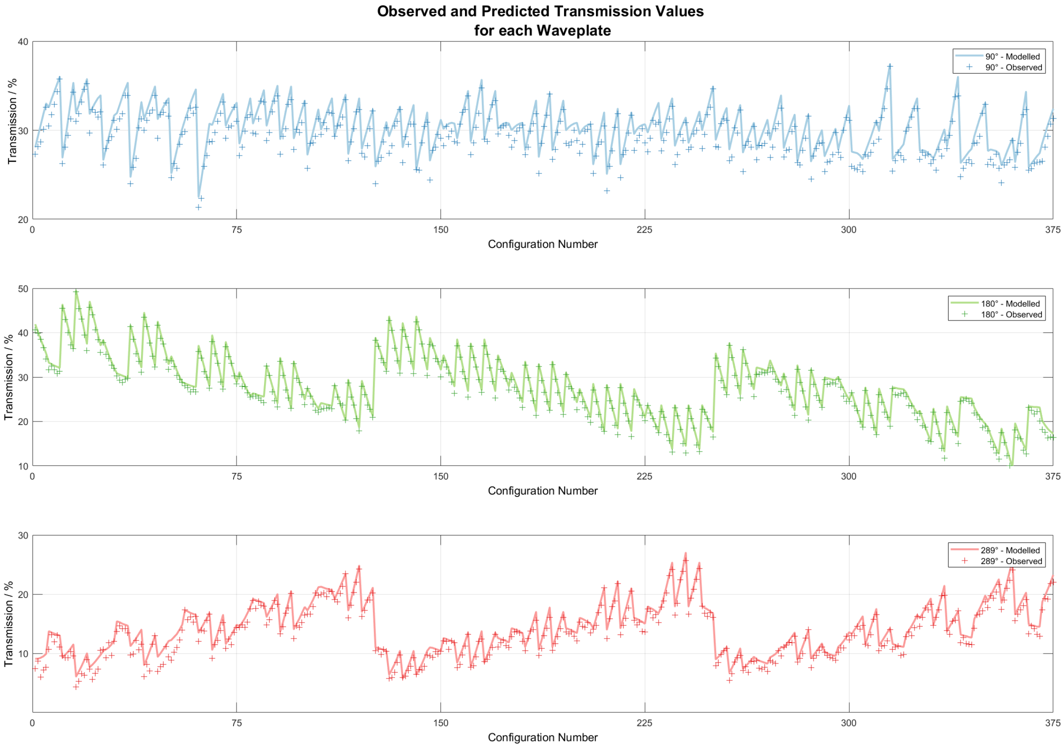

Figure 3.9: Transmission for each configuration of the test case, presented in the order of measuerment. The top (blue) graph corresponds to the results using the 90° waveplate, the middle (green) shows the results for the 180° waveplate, and the bottom (red) for the 289°. The crosses on the plots indicate the observed results from the experiment which are consistently lower than the modelled results, shown by the line.

in the first graph.

Figure 3.11 is a single figure, with the same data sorted in ascending order of transmission. The data from all three waveplates is plotted as model-observed pairs, in the same colours as fig. 3.9. This gives a clearer perspective on how well the modelled data matched that observed in the laboratory. For all three waveplates it can be seen that the model predicted higher transmission than that observed. This is expected, in this case absorption was not taken into account in the model, it could however be easily introduced. Also, in the model all components are perfectly aligned to the desired values and no shock or vibration is taken into account. Vibrations and misalignments can be introduced during the adjustment for the experimental process, both of which can lead to reductions in the transmission for a system this sensitive. It may also be possible that these discrepencies are caaused by a callibration error with the detector. The experimental results are approxiamtely 100 μW on average lower than the model, within the range that callibration errors could reasonably account for.

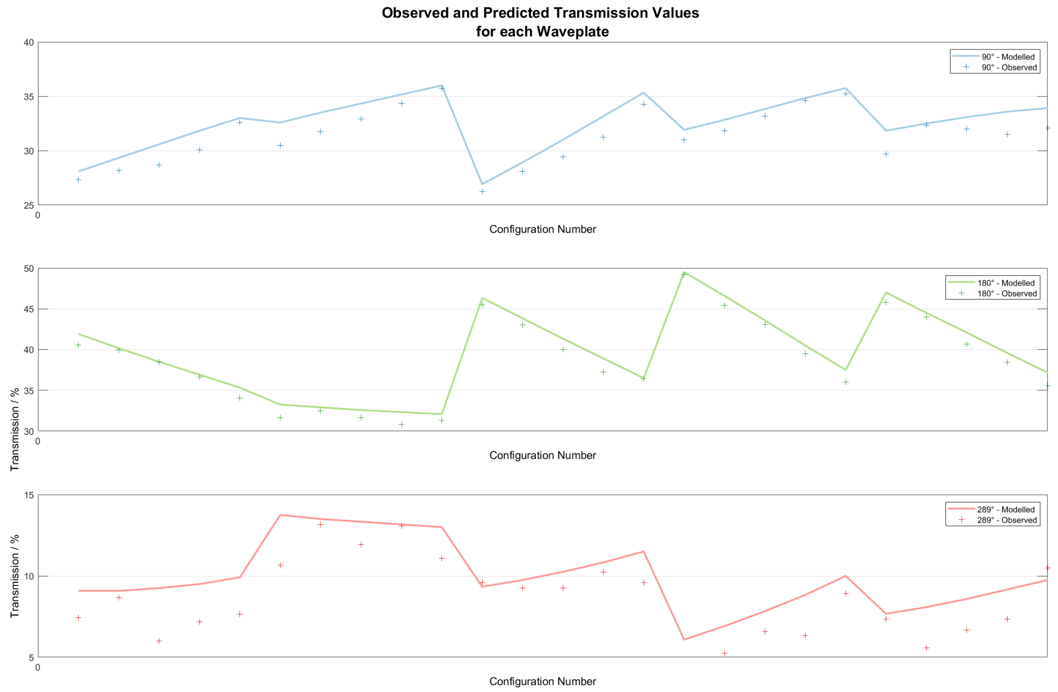

Figure 3.10: Close-up of the first 25 results from fig. 3.9 to highlight agreement between the modelled and experimental results. The top (blue) graph corresponds to the results using the 90° waveplate, the middle (green) shows the results for the 180° waveplate, and the bottom (red) for the 289°. The crosses on the plots indicate the observed results from the experiment.

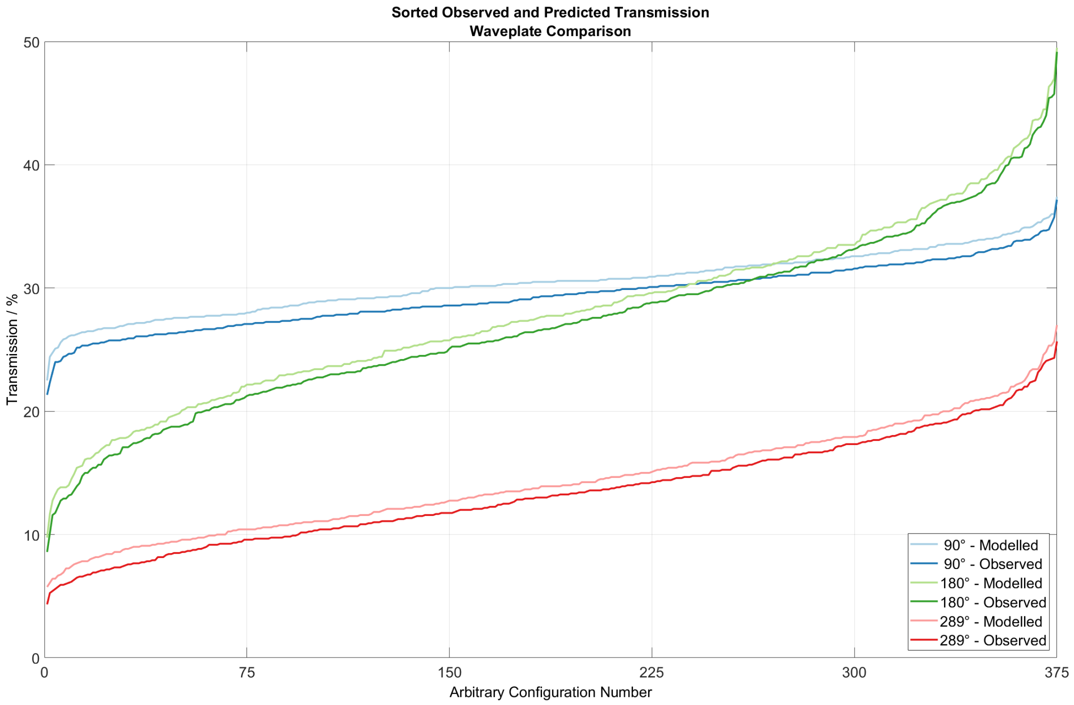

Figure 3.11: Transmission for each configuration of the test case. The data is exactly the same as that depicted in fig. 3.9, but for purposes of viewing these graphs have been presented in order of ascending power. The top (blue) graph corresponds to the results using the 90° waveplate, the middle (green) shows the results for the 180° waveplate, and the bottom (red) for the 289°. The darker lines of each colourindicate the observed results from the experiment which are consistenly lower than the modelled results.

3.4.3 Advanced Analysis

This section presents the application of the analyses from section 3.3.2 to the data produced for this cavity. These analyses were designed to assist the identification of combinations of components and parameters having the greatest impact on performance.

3.4.3.1 Parameter Frequency

Figure 3.12 shows the results of the analysis presented in section 3.3.2.2. In this case, a threshold of 50% transmission was set. Each row corresponds to a single degree of freedom and each column corresponds to a step from the aligned configuration, as detailed in table 3.2. As can be seen from the fourth row, there are only three columns with results because there are only three configurations for the Porro prism y-Axis rotation. In this case, for each grid square, the percentage is calculated from the 125 corresponding configurations. In all other cases, each grid square is representative of 75 possible combinations.

For each parameter configuration, the total number of data points that meet the condition is counted and divided by the total. Darker colours mean that a larger number of the configurations with that parameter at that value met the threshold condition, transmission lower than 50% in this case. This allows for quick comparison between configurations to see if any parameters in particular are more common in the set defined by the condition.

Similar frequency (colour) for all squares in a row indicate that the parameter has no configurations that are significantly more likely to meet the condition. An example of this is the x-Axis rotation of the Pockels cell in fig. 3.12, each configuration occurs with a similar frequency within the sub-50% transmission results. In contrast the z-Axis rotation of the Pockels cell has a larger range of frequencies for each configuration. As the Pockels cell is rotated towards +2° the frequency of results falling below 50% transmission increases. This indicates that this parameter in particular should be investigated further and more tightly controlled.

3.4.3.2 Variance Dominance

Figure 3.13 presents the output of the analysis discussed in section 3.3.2.3. These results are drawn from the subset for the 289° retardation waveplate. Each of the four graphs then represents one of the four remaining parameters. The mean and standard deviations for all of the results in the subset, where the graphed parameter is at the value on the x-axis, are then calculated. The mean for each parameter-value is plotted as a black cross and a band is then calculated that encompasses all values within one standard deviation of the mean at that parameter value.

This allows us to clearly see which parameters, at which values, dominate the output of the entire system. Bands that are thinner, having a smaller range, indicate parameter-values that greatly affect performance. Wider, broader range bands show that the configurations including this parameter-value were influenced more heavily by other degrees of freedom. Further to this, parameters with bands that remain flat and steady across all values

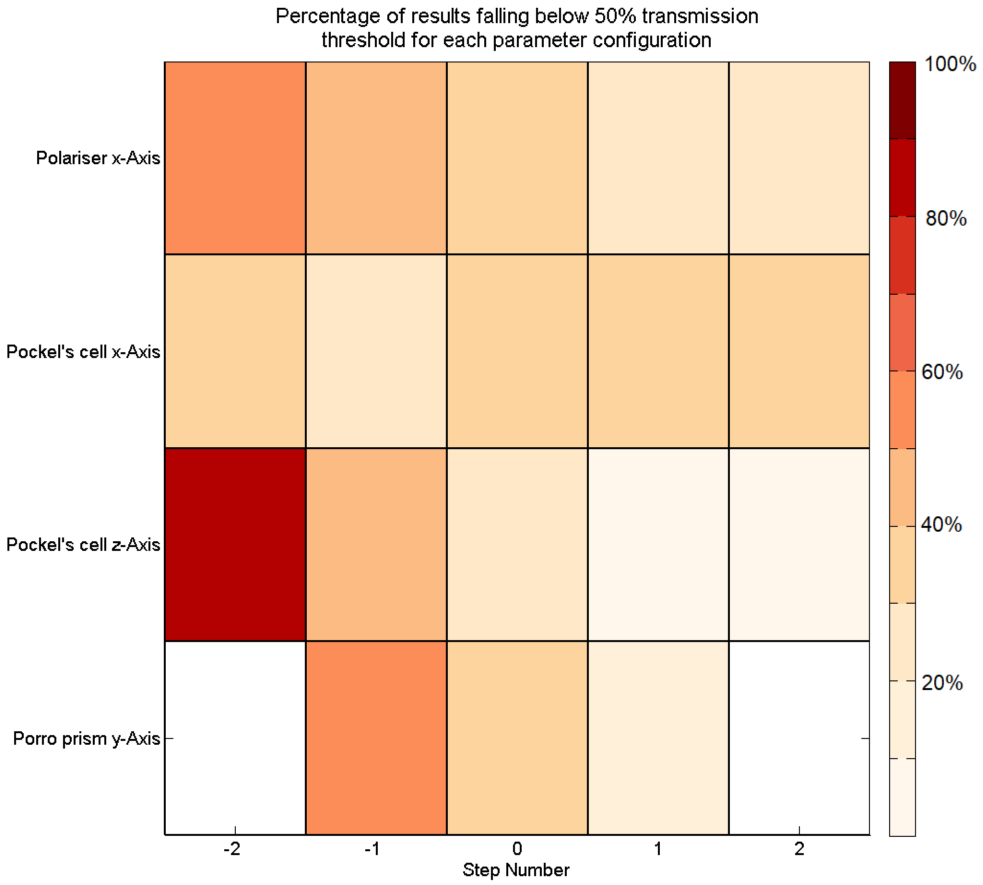

Figure 3.12: The output from the parameter frequency analysis outlined in section 3.4.3.1. Each cell in this grid corresponds to a parameter value that has been tested, note the blank cells at step number ±2 for the Porro prism y-Axis tilt as only three values were tested. The cells are shaded according to the percentage of results that fell below a 50% transmission threshold. The darker shading corresponds to a higher percentage of all results taken at that measurement falling below the threshold.

indicate that variance in this parameter doesn't greatly affect performance.

As can be seen from fig. 3.13 the Pockels Cell alignment has a strong effect on the transmission through the system. The z-Tilt exhibited a broad range of means, from 23% transmission (1.01 mW) in the -2° configuration to 45% transmission (2.32 mW) at +2° . The small tight band indicates that the performance varies relatively little with adjustments to the other parameters. In this case, aligning the Pockels Cell with a positive rotation about the z-axis may exhibit a more optimal output. The x-Tilt of the Pockels Cell shows a broader band and, thus, a less significant effect on the output, but a clear asymmetric distribution of the means can be observed about the 0° alignment. It would indicate that although a higher mean output may be seen, for all configurations, with a 1° negative rotation about the x-axis, slight deviations from this point in either direction would lead to a drop in transmission.

Both the Porro Prism and the Polariser x-Tilt have broad bands compared to the Pockels Cell z-Tilt. This is due to the aligment of the crystal axis along the z-Axis. Misalignment along this axis introduces birefringence leading to a larger phase shift and, therefore polarisation rotation, of the light passing through the crystal.

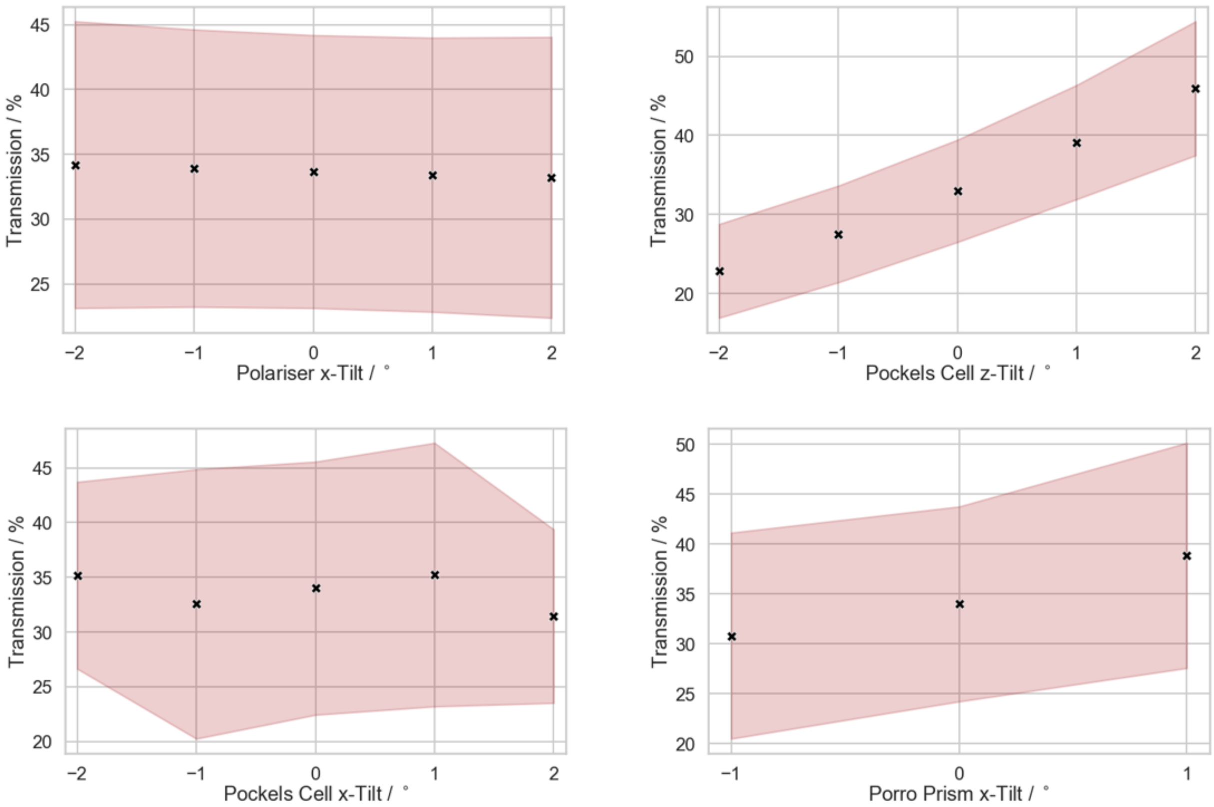

Figure 3.13: The output from the variance dominance analysis outlined in section 3.4.3.2. For each parameter at each value, all results were collected and the means calculate, displayed here as black crosses. The pink bands encompass one standard deviation of those results either side of the mean. Tighter bands indicate a parameter value that has a more significant impact.

The Polariser also shows very little variance in the means between the three possible values. This particular degree of freedom does therefore not need ot be controlled as strictly as others being considered in this system. Knowledge such as this should enable optical design engineers to make better decisions and relaxing alignment tolerances on components could ultimately lead to reductions in manufacturing cost.

3.5 Case 2: Complex Systems

This section is an example of how these tools can be used to model vastly more complex and general systems. One of the limitations in understanding tolerance ranges for multiple components is how they interact with one another. As the system becomes more complex, the number of degrees of freedom increase and testing all possible combinations becomes unmanageable. The system measured in section 3.4 had three parameters (Pockels Cell x-Tilt and z-Tilt, Polariser x-Tilt) with five possible discrete values, and two parameters (Waveplate retardation and Porro Prism y-Tilt) that had three values. Even this simple case led to 1125 combinations to be measured. Including more parameters increases the time taken exponentially and thus experimentally investigating many more parameters becomes unfeasable. The primary benefit of this approach is that, through computation and automation, vast numbers of parameters can be investigated thoroughly and at a lower time cost.

The remainder of this section will present the same analyses performed in section 3.4, on a more complex optical system, considering many more degrees of freedom.

3.5.1 Design

The system used in this section will again be based on the typical laser-rangefinder schematic as presented in fig. 3.1. As shown in fig. 3.14 the setup is almost identical to the one used in section 3.4 with the addition of a Corner Cube. The laser will first make a single pass through the corner cube before passing through the remainder of the system. This will add a further source of phase retardation which must be accounted for. It should be noted, the axes of rotation are defined in reference to the table.

In the perfectly aligned case it was found computationally that a waveplate with a 130° phase retardation, with a z-axis rotation of 60° , gave the maximum transmission. For this example a single value of phase retardation will be tested with ±2° tolerance, to highlight how deviations in this parameter affect performance.

In this example 16 parameters will be investigated in total, the x, y and z-axis rotations of five components and the phase retardation of the waveplate will be varied from their base state by ±2° . The z-axis is defined as parallel to the beam as it exits the laser, the x-axis is defined as being perpendicular to the page and the y-axis is defined as being perpendicular to detector face. Table 3.3 lists each of the individual parameters and the range of values to be tested.

The plane of the polariser is rotated by a few degrees around 45° about the x-axis in the base state and thus the range differs from the other components, as does the waveplate z-axis rotation and phase retardation, having values other than 0° in the aligned base configuration.

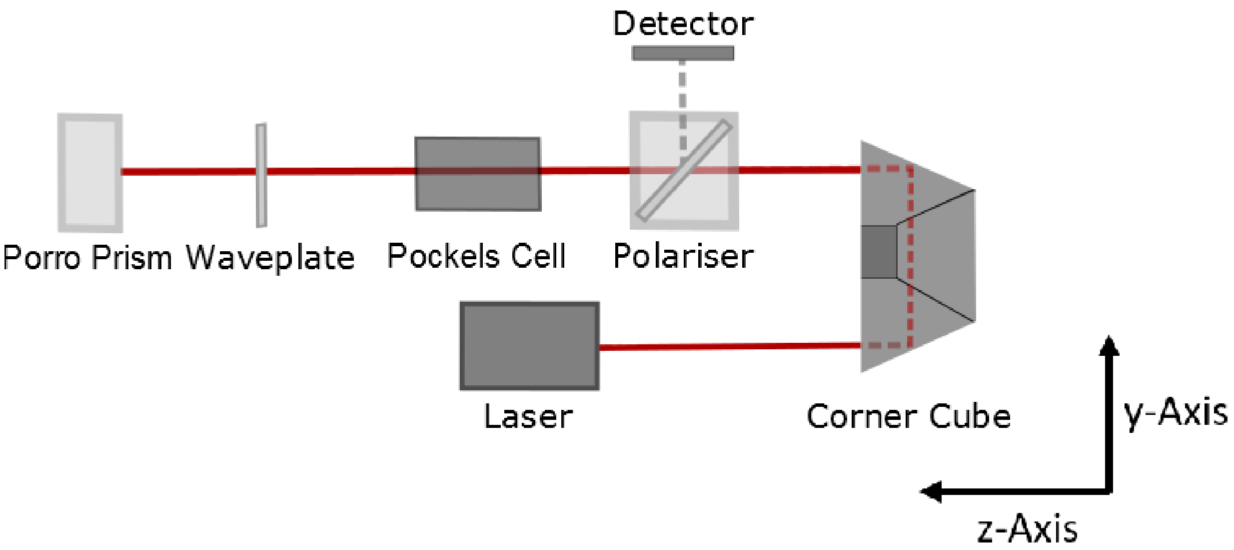

Figure 3.14: Diagram of the more complex system for demonstration of analyses. The axes of rotation are defined in reference to the table.

Table 3.3: Parameters used in the complex system demonstration analysis, where each axis of rotation is defined in reference to the table. Step-size has been ommitted from this table due to the random nature of the Latin Hypercube Sampling outlined in section 3.5.2.

Component	**Parameter**	**Values (degrees)**
Corner Cube	x-Axis rotation	-2 to +2
Corner Cube	y-Axis rotation	-2 to +2
Corner Cube	z-Axis rotation	-2 to +2
Polariser	x-Axis rotation	43 to 47
Polariser	y-Axis rotation	-2 to +2
Polariser	z-Axis rotation	-2 to +2
Pockels cell	x-Axis rotation	-2 to +2
Pockels cell	y-Axis rotation	-2 to +2
Pockels cell	z-Axis rotation	-2 to +2
Waveplate	x-Axis rotation	-2 to +2
Waveplate	y-Axis rotation	-2 to +2
Waveplate	z-Axis rotation	62 to 66
Porro prism	x-Axis rotation	-2 to +2
Porro prism	y-Axis rotation	-2 to +2
Porro prism	z-Axis rotation	-2 to +2
Waveplate	Retardation	125 to 135

3.5.2 Sampling Scheme

Although the computational approach to modelling these systems is quick, the high dimensionality of the data still precludes an exhaustive search of all possible configurations. If sampling were to be carried out in a similar fashion to section 3.4, a single point on each whole degree value for each parameter, then there would be a total of 5^{16}, more than 150 billion, combinations. For the system modelled here, it takes approximately 0.2 seconds per configuration to calculate the transmission in Zemax, were every combination to be investigated in this way then it would take over 950 years to calculate all transmission values.

Another advantage of a computational approach is the ease with which more complex sampling methods can be applied. Latin Hypercube Sampling (LHS) [111, 112, 113, 114, 115] is a method for sampling multidimensional distributions in a near random manner. It is an n-dimensional extension of a Latin Square. Placing a square grid over a two dimensional distribution and sampling from each row and column only once is the basis of a Latin Square. Extending this to an arbitrary number of dimensions is the premise of LHS. This method of sampling a dataset reduces the clumping and undersampling of certain parameter spaces that can occur from random selection.

This is applied to the data in this example by first selecting how many samples will be taken of the data set. The range of each parameter is then divided into many evenly sized segments. A random point is then chosen falling within each of these segments and this process is repeated for all parameters. The list of points for each parameter is then shuffled and combined with the others such that each segment for each parameter is only ever present in one solution. This method is useful for generating well distributed random samples and is often used when the measurement of those sample solutions is computationally costly [111, 112, 113, 114, 115]. It is well suited to the analyses presented here as it allows a more thorough exploration of the solution space than discrete, defined intervals.

In this example we have chosen to take 2×10^4 samples. This means for the 4° range of values each parameter can take, each segment within which a sample may be generated will be $2^\circ \times 10^{-4}$ wide.

3.5.3 Analysis and Discussion

With 2×10^4 measurements, each parameter having unique values for every sample, it is possible to observe the data with a varying degree of granularity. This is accomplished by regrouping the samples into larger segments for each of the parameters and aggregating the results. In the case of this example 20 grouping segments were chosen, each comprised of 1000 samples, and the analyses presented were performed on the mean results from these larger (0.25°) grouping segments. Controlling the granularity of the analyses is simply a case of altering the size of these grouping segments to include more or fewer samples in each.

3.5.3.1 Parameter Frequency

Figure 3.15 is the result of the analysis presented in section 3.3.2.2 performed on the grouped segment mean transmission values for this example. The threshold condition considered here is the same as in fig. 3.12; for each of the parameters at each value, the proportion of results falling below 50% transmission are compared.

Rows with little variation in colour indicate that they have less influence on the system across the measured range of inputs. Rows with larger variations in colour, such as the Pockels cell y-Axis in this case, indicate that they play a more dominant role in the variance of transmission across the sample. The Pockels cell y-Axis row shows that fewer than 10% of the measurements made when the parameter was between ±0.5° had a transmission lower than 50%. When this parameter was in the -0.5 to -1.75° and the 0.5 to 0.75° ranges, however, between 60 and 70% of the measurements made had transmissions lower than 50%. Less pronounced, but similar variance is shown for the Pockels cell x-axis, with a broader range about 0° within which fewer low transmission results occur.

Interestingly, the Porro prism and the Pockels cell z-Axes rows both show higher proportions of measurements falling below the threshold when close to the aligned case. This indicates that the system as a whole may be more stable with a deviation from orthogonal alignment for these two parameters.

The remainder of the graph being almost homogeneous, with nearly every cell falling in the 30 to 40% range, indicates that the Pockels cell alignment dominates the output of the system.

Conclusions could be drawn from this that tightening the tolerance range for the Pockels cell x-axis and y-axis to ±0.5° would increase the overall stability of transmission for the entire system. This might make it possible to relax other tolerance ranges, reducing the cost to manufacture.

3.5.3.2 Variance Dominance

Figure 3.16 shows the results of the analysis presented in section 3.3.2.3 for four of the parameters. These graphs were chosen as all of the others were very similar. Mean transmission for the other parameters remained steady across all values and each of them had broad, straight bands.

For each parameter the mean transmission is taken for each of the grouping segments and plotted as a black cross. The band is then plotted to encompass the all values within a standard deviation of the mean for that grouping.

Partitioning the results in this way treats all other parameters as in a random state within their tolerance range. This enables direct comparison between parameters and how transmission varies across their tolerance range.

Similar to the results presented in section 3.5.3.1 both the Pockels Cell x-Tilt and y-Tilt show the largest range of means. The Pockels cell y-Tilt graph is the only one to show transmission above 90% for any value, indicating that it is dominant in the system across such a large tolerance range. Any deviation from the aligned state greater

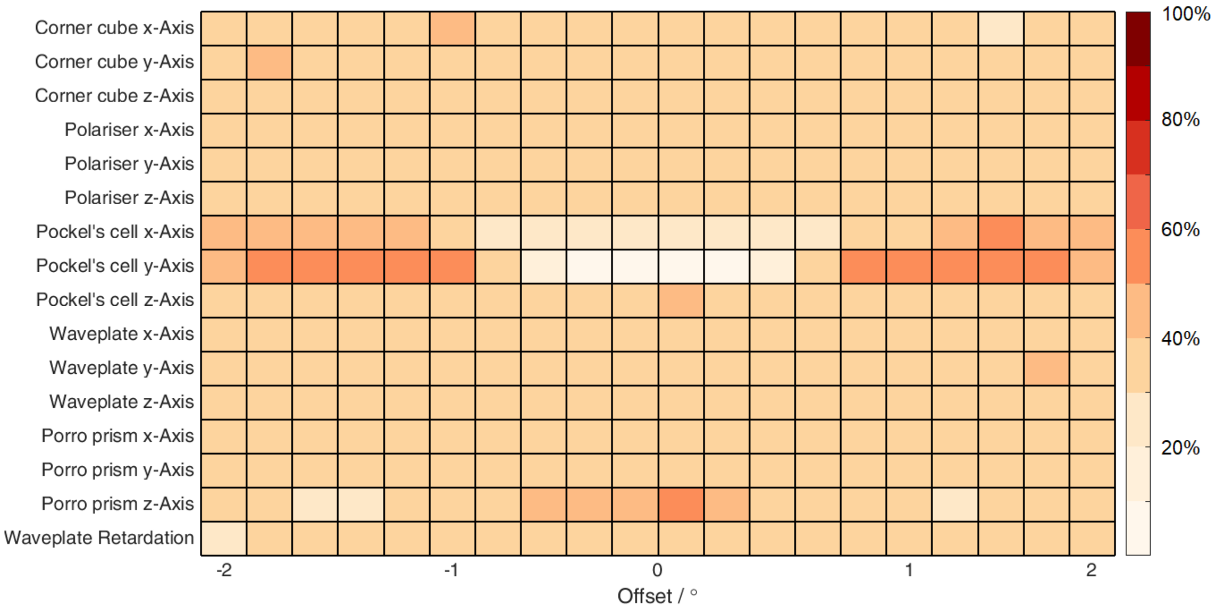

Figure 3.15: The output from the parameter frequency analysis, outlined in section 3.4.3.1, performed on the complex example. Each cell in this grid corresponds to a parameter value that has been tested. The cells are shaded according to the percentage of results that fell below a 50% transmission threshold. The darker shading corresponds to a higher percentage of all results taken at that measurement falling below the threshold.

than 0.75° in either direction leads to the mean transmission falling to below 80% for all other configurations.

The size of the bands also strongly indicate which parameters have the most influence on the transmission. The peak of the Pockels Cell y-Tilt graph has a very small standard deviation band despite all of the variations in the fifteen other degrees of freedom. As the deviation increases and the mean transmission drops, the band also becomes wider. This shows that once misalignment reaches a certain point, other parameters have a greater relative affect on the mean transmission than the Pockels Cell y-Tilt. This is due to the aligment of the crystal axis along theyz-Axis. Misalignment along this axis introduces phase shift and polarisation rotation, due to the birefringence of the crystal. Close to the aligned state, minor misaligments have a significant impact on the performance of the system.

Although shallower and broader, the peak on the Pockels Cell x-Tilt graph also shows a large decrease in mean transmission as misalignment increases. The band in this graph remains broad across all values. This is due to how dominant the y-Tilt is on the result set, where large variations in the y-Tilt alignment will greatly affect the mean for each of the x-Tilt values.

The Porro prism z-Tilt graph is interesting as there is a 16.3% range in the mean transmission values and the minimum occurs at no deviation from the aligned case. The maximum falls within the -1.25 to -1.5° grouping segment, but a rotation in the positive direction has a mean transmission only 1.7% lower. The band also remains fairly equal in width across all values, so none in particular have a stronger affect on the mean. The slope around the ±1.5° is also not as steep, but does decline, meaning that the mean transition would be more stable. This

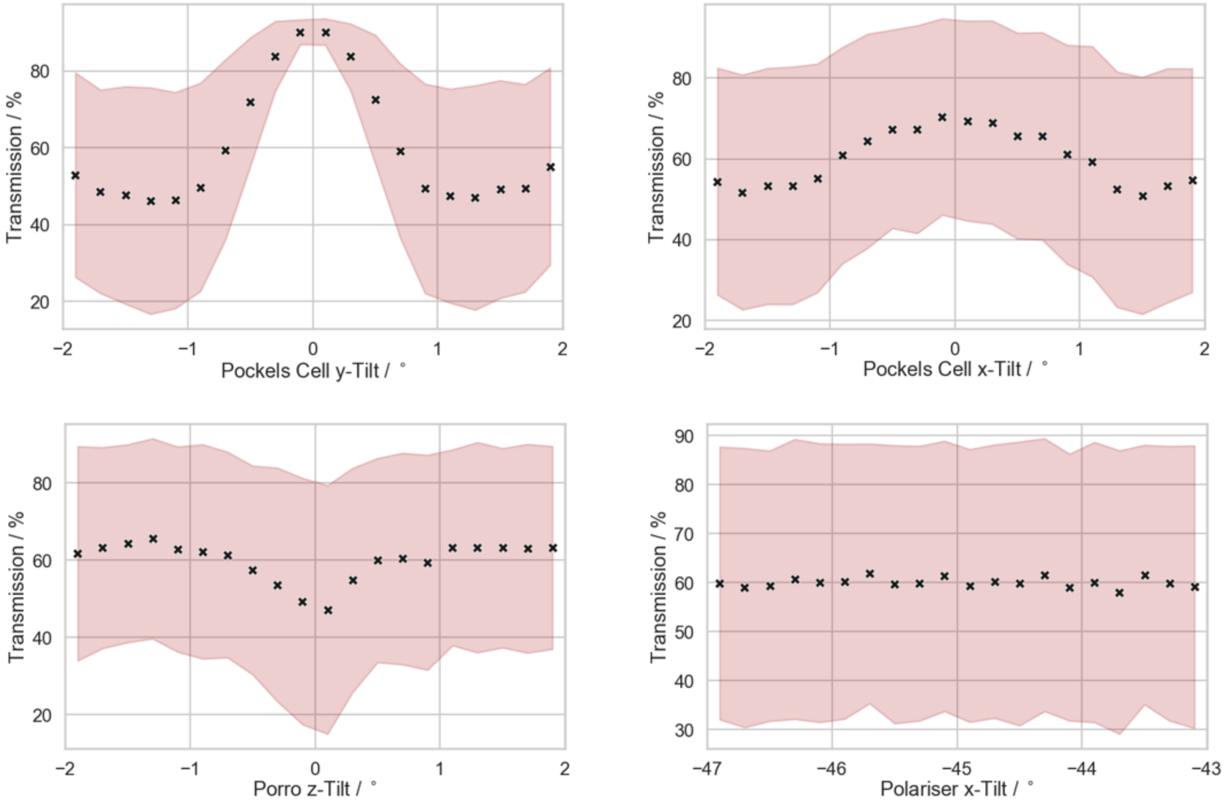

Figure 3.16: The output from the variance dominance analysis, outlined in section 3.4.3.2, performed on the complex example. For each parameter at each value, all results were collected and the means calculate, displayed here as black crosses. The pink bands encompass one standard deviation of those results either side of the mean. Tighter bands indicate a parameter value that has a more significant impact.

would make the ±1.5° more suitable for an optimum configuration for transmission stability, than the 0° state.

Other than the Polariser x-Tilt, all other graphs appeared flat with wide bands. This meant that the rotational alignment had a smaller effect on the transmission than those parameters selected here. Also the phase retardation of the waveplate also had less of an effect than the Pockels cell and Porro prism alignments.

The Polariser x-Tilt does not show large variance in the mean transition and the band remains broad across the entire range of values. This has been included to highlight what the results from a less significant parameter would look like using this analysis.

3.6 Conclusion

The techniques and analyses presented in sections 3.3 and 3.3.2, build on the capabilities of Zemax. They allow for answers to difficult questions concering the interplay of multiple parameters and thier effect on performance and enable the possibilities of further advanced analysis. Although the base example given in this report is that of a polarisation sensitive set of optics, all of these techniques are as general as the non-sequential capabilities of the Zemax software allow. Other example systems, with differing questions, could highlight worthwhile paths of development for the tool set.

Section 3.4 verifies how closely the results produced using this computational method imitate those gathered in the laboratory. Discrepancies arise due to the incomplete nature of the underlying models used within Zemax. Firstly, the ray-tracing model does not take into account diffraction effects which means the shape of the beam through each component would be different to that observed in reality. Secondly, for the model used in section 3.4, bulk absorption and surface coatings have not been considered, as these were not a core requirement from the users that led to this initial investigation. This accounts for the higher transmissions observed across almost all of the modelled values, when compared with the experimental. Despite this, the model gives a good understanding as to how the system behaviour changes as multiple parameters are altered.

In this case only a single output parameter is being considered, transmission to the detector. Increasing the number of outputs recorded is trivial and does not significantly increase the computation time. This means the model will readily accommodate systems for which a performance metric, or merit function, is calculated through a combination of measurable output parameters.

The output from the analysis presented in section 3.3.2.2, as seen in figs. 3.12 and 3.15, gives the user a quick method for viewing which parameters, at which values, meet the desired condition. The examples presented here merely look at those transmission values falling below the 50% threshold, which is useful for identifying troublesome configurations and dominant parameters, as evidenced by fig. 3.15. It is possible to specify more complex logic however, enabling the user to investigate all degrees of freedom quickly and at a glance, for any criteria that was desired. Expanding on the logic makes this a very powerful tool for exploring highly dimensional data sets in a visual context, making it possible to quickly identify those parameters that are most crucial to the performance of the system overall. This would enable a production engineer to specify a set of conditions that are key to the performance of a system and the model would be able to quickly present which parameters most often breached those conditions. It would be simple to extend the process here to record each performance metric during the data collection stage and present the findings for each both separately and in a single view, if a proper merit function was defined. The production engineer could then quickly gather and assess essential performance data from a design and understand which parts of that design might need adjustment. The ultimate goal being, to enable tolerance ranges to be relaxed and reduction in design and manufacturing effort and, therefore, cost. It would also be simple to invert the threshold and identify which component parameters are seen most frequently in highly performant combinations. This would enable a quick method to refine an initial design, enabling the user to make sense of the interplay of parameters across many degrees of freedom.

Figures 3.13 and 3.16 present a single dimensional look at the highly dimensional solution space. They enable the user to understand how a system's performance varies according to the subject parameter of each graph. Clear, distinguishable differences between individual parameters enable conclusions to be drawn about not only how parameters affect the performance, but also how they are affected by others in the system. This analysis also enables quick identification of problematic parameters or value ranges. As is the case with fig. 3.15, the top two graphs show that the Pockels Cell x-Tilt and y-Tilt greatly reduce the overall system performance when they deviate more than 0.5° from the aligned state. This means that priority can be given to these parameters when determining manufacturing tolerance ranges. For other parameters, these small deviations have less of an effect and their tolerance ranges can therefore be reduced, ultimately reducing manufacturing costs. A production engineer could, therefore, use this analysis to identify which parameters are most sensitive to misalignment within the system. As was the case presented in section 3.5.3.1, a small subset of the parameters dominated the variance across all combinations. This allows the user to quickly identify where best to refine or refactor a design and provides a quantifiable measure against which to test tolerance ranges.

The algorithm presented in section 3.3.2.1 enables the user to explore the solution space in an automated manner. Although not presented here, it would be simple to extend the model so that is was capable of accepting custom merit functions. This would enable it to score each of the potential solutions found using this method, making it is possible to find configurations that meet many stringent criteria simultaneously. Furthermore, it is possible to query multiple outputs from Zemax. This means merit functions finding a balance between power, beam shape and polarisation state for example, could be used to optimise the system in any way the user sees fit. It also ensures that the configuration found meets the stability requirements, as set by the user, by searching neighbouring values in all parameters so that optima near steep declines in performance are not suggested.

Overall, this tool set enables engineers to quickly investigate systems with many degrees of freedom. It presents highly dimensional data in an intuitive, graphical format and quickly highlights areas to focus on when optimising designs.

CHAPTER 4

Non-Line-of-Sight LIDAR

4.1 Introduction

The models developed and presented in the other chapters of this thesis are concerned with modelling, through ray-tracing, the internal phenomena of laser systems. This chapter presents a model that is also reliant on ray-tracing, but of an application of laser technology. Many of the tools and techniques underlying the other models are applied here in subtly different ways. The non-line-of-sight LIDAR technique attempts to retrieve the position and track movements of target objects that are not directly in line-of-sight. This model allows the user to test the potential limits this technology, which would otherwise be difficult to ascertain due to current technological limitations or available laboratory space of sufficient size and shape. Later in the chapter relationships between the limitations imposed by target size, angle, distance, and position are investigated.

The ability to detect motion and to track a moving object that is hidden around a corner or behind a wall provides a crucial advantage when physically going around the obstacle is impossible or dangerous.

Laser illuminated detection and ranging (LIDAR) is a technique that allows for accurate mapping and positioning of objects over large distances. This is accomplished by sending a laser signal to the target and recording the time-of-flight of any reflected signal [116]. Figure 4.1 illustrates this process: the pulse is emitted from the device and is incident upon the target; a portion of this light is then scattered back to the detector and the time-of-flight is recorded. One clear limitation to this technique, however, is that it can only be used for objects in direct line-of-sight.

In recent years, technologies have emerged that seek to overcome the line-of-sight limitation, making it possible to observe objects that are obstructed from view. Although similar technologies have been demonstrated using radar [118, 119] and reflective surfaces to steer the light around obstructing objects, LIDAR based techniques offer many advantages over these other methods. It has been demonstrated, that imaging and positioning of an object can be achieved using a signal that has undergone multiple scattering events. Using LIDAR based

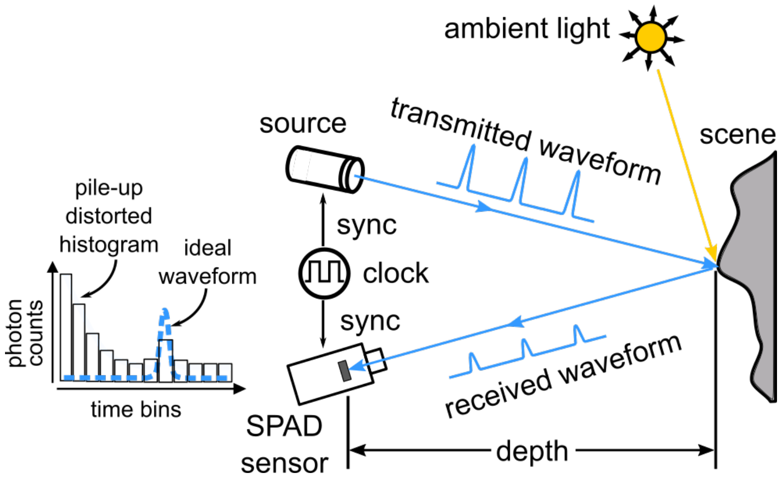

Figure 4.1: LIDAR measures the distance to a target by measuring the time-of-flight of a laser pulse illuminating the target and then scattered back to a detector. Here, a pulsed light source illuminates the scene with periodic short pulses. A SPAD sensor records the arrival times of returning photons with respect to the most recent light pulse, and uses those to build a timing histogram [117].

techniques [116, 120, 121, 122], this can be achieved in real-time and will work in any environment without need for reflective elements, making it far more suitable for handling unknown conditions and a wider variety of possible scenes.

Velten et. al have shown that is is possible to scatter light around an obstacle using flat vertical surfaces, such as walls or doors [120, 123]. Raster-scanning a laser beam enables illumination of the hidden scene from multiple points. Then, using a streak camera with picosecond temporal resolution, the backscattered light from the scene is detected. Using this approach, they could reconstruct full three-dimensional images of objects hidden from view [120, 124]. Acquisition times for enough data to build such an image are long, meaning this can not be done in real time for moving objects. And although the object reconstruction is accurate, the raster-scanning of the laser beam can introduce mechanical complications with regards to precision.

Based on a method of imaging through opaque barriers [125, 126, 127], another solution has been presented using light originating outside the field-of-view [128, 129]. When a hidden object lies between a light source and a scattering surface such as a wall, that is in line-of-sight, a speckle pattern is present in the reflection of that light from that surface. It is possible then to use post-processing to analyse that pattern and reconstruct the shape of the object. This approach therefore allows for use of surfaces, within field-of-view, to obtain information about hidden scenes. The post-processing stage is fast and only requires a single image to reconstruct information about the hidden scene. It is, however, reliant on a light source being placed behind the object of interest, directly illuminating it, limiting the applications to specific cases. Other methods have been explored for creating images through diffusing media using time-gated techniques [130].

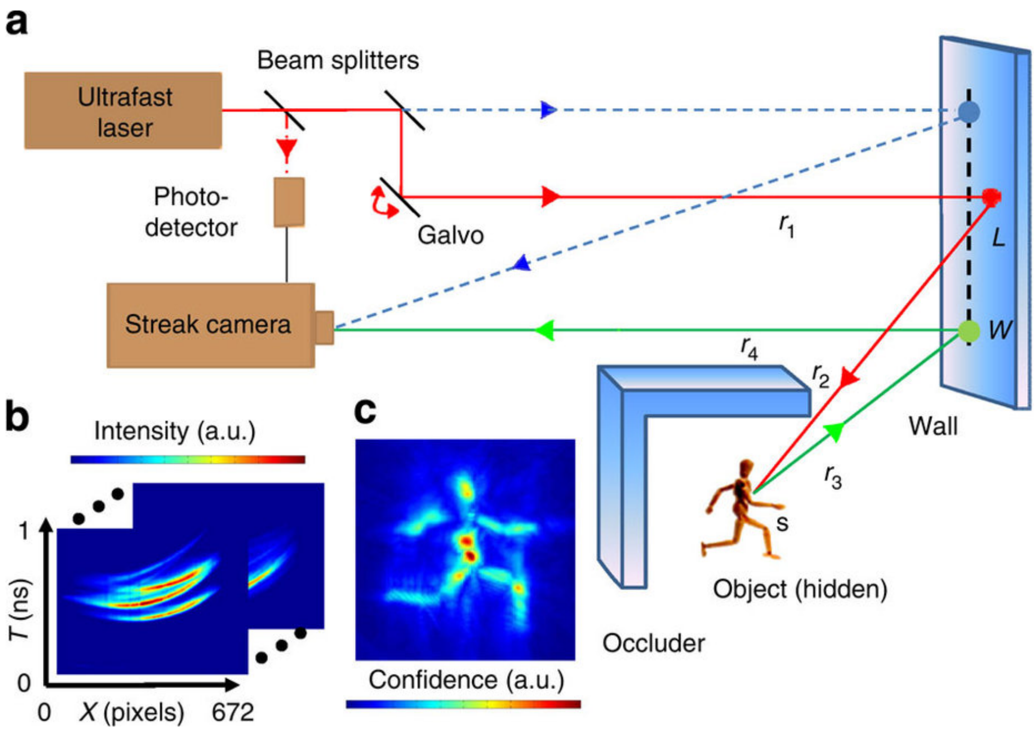

Figure 4.2: Reconstruction of images for hidden objects. (a) The capture process involves taking a series of images of the dashed line segment of the wall, with a streak camera. The laser pulse is raster-scanned to illuminate the scene from multiple points. Some of the light then scatters to the object and back to the wall where it is detected. (b) Examples of the normalised intensity images collected by the camera for a each pulse. (c) The reconstructed shape of the hidden object after the post-processing algorithm. Figure reproduced from Velten *et al.* [120]

4.1.1 Non-Line-of-Sight Detection and Tracking Model

Retrieving information about objects hidden from view and tracking their movements could see a number of real world applications. This chapter discusses a model of the method presented by Gariepy *et al.* [2]. This approach is capable of determining the position of a non-line-of-sight target and monitoring its movement, in real time, at the human scale . This technique eliminates the requirement for a close by wall, door or highly reflective scatterer, by utilising the floor as the imaging surface, relaxing some of the constraints on the real-world application of this technology. This was achieved using a two-dimensional array of single-photon avalanche diode (SPAD) detectors [2, 3, 4, 131, 132, 133] operating in time-correlated-single-photon-counting (TCSPC) mode [134, 135]. TCSPC is a method for detecting low light level signals by counting single photons from a periodical source and reconstructing the waveform from these counts over multiple repetitions.

SPAD detectors have extremely high sensitivity capable of detecting single photons with high temporal resolution. This allows for light from a laser pulse to be detected after multiple scattering events and the recording of precise time-of-flight for the returned photons. The high sensitivity of these cameras also means fast acquisition times and being an array of pixels gives a two-dimensional image, making it possible to track a moving target in sub-second time scales. The precise (approximately 50 ps) temporal resolution is achieved because of the detectors operating in TCSPC mode. TCSPC involves measuring arrival times of individual photons, emitted in a single pulse and scattered back from the scene. Multiple pulses are then fired each second and a histogram of the arrival times can be accumulated. These histograms are then analysed and comparisons made to isolate moving objects from static background [134, 135]. Section 4.2.2 further explains this technique, and how position is retrieved from the temporal data.

Being able to determine the target's position, with centimetre precision, illustrates how this technique could be used to track the movement of an object that is obstructed from line-of-sight. This approach is also able to track target movement in real time.

4.1.2 Time-Resolved Non-Sequential Ray-Tracing Model

The remainder of this chapter concerns the method proposed by Gariepy *et. al* [2] and developing an accurate approach to modelling experiments of this kind. These models provide a means to better understand the processes and, furthermore, the technological requirements to expand upon this work. Similarly to other chapters within this thesis, ray-tracing models will be developed, leveraging the complexity and power of the Zemax optical modelling software. These models will allow the user to examine each point in the path of each ray, from the laser, through all scattering events, to termination. This affords the user the ability to fully understand the distribution of rays as they traverse the scene and thoroughly interrogate how the size, shape, and positioning of objects within the scene, affect the target positioning calculations.

Demonstration and experimentation with this technique poses a number of difficulties, especially when regarding the spatial and temporal limitations of the approach. Being able to computationally model the response from a broad range of scenes mitigates some of the difficulty organising the large experimental scenarios required. This gives researchers a low-cost, low-risk method to priotitise further experimental work in development of non-line-of-sight target tracking approaches at human scales, without the need to secure scarce resources such as large laboratory spaces. It also enables a method for quickly evaluating what effect different sources or detector optics would have. Furthermore, Zemax is capable of modelling highly complex optical systems and sources. This enables users to test new components that they may not yet have access to in the laboratory. Reducing this barrier makes exploration of novel configurations and components quicker and easier, which may highlight the need to develop other enabling technologies. It also allows for the quick, automated, optimisation of the existing system through computational testing.

Section 4.2 presents an experiment involving non-line-of-sight targets, from which time-resolved photon counting measurements are taken. The method for retrieving the position of a target object is then introduced in section 4.2.2, giving a brief outline of the algorithm and the mathematical basis. Section 4.3 then discusses the modelling approach presented here, which makes use of non-sequential ray-tracing to determine how light, from a laser source, scatters within a scene with multiple scattering elements. The data collected from experiment is then used to verify results from the model in section 4.5. The model is then used to further investigate the technical requirements and understand the limitations of this approach in section 4.6. The model is then used to investigate some of the underlying relationships between objects within a scene, such as target position, target size, and the angle and distance of the detector from the initial scattering surface, in section 4.7. The model confirms a number of expected spatial relationships in this section which allow for quick estimation of expected signal power from a scene, the determining factor in distinguishing the detectability of a potential target.

4.2 Methods

This section covers how the data is collected and the process of determining the position of a non-line-of-sight target. Figure 4.3 illustrates a scene, encompassing a laser source, a target, a scattering surface and a detector.

To determine the position of a target at point $\mathbf{r_o}$, light must travel from it to the detector. Each detector in the array images a point in the field-of-view, point $\mathbf{r_i}$ on the wall in this case. The scene is illuminated with a laser pulse, emitted at a known time, incident upon a scattering surface, point $\mathbf{r_l}$, as shown in fig. 4.3. Some of this light is then scattered from point $\mathbf{r_l}$ onto the target at point $\mathbf{r_o}$. Finally, a portion of this light scatters to point $\mathbf{r_i}$ on the wall and to the detector.

As many scattering events must take place to interrogate the hidden scene, only a very small fraction of light from the laser pulse will scatter from the target to the detector. Described here is the simplest path the ray could take and return from the target; other scattering events will occur, dependant on both the positioning and geometry of other scattering surfaces. Furthermore, many paths exist that do not reach the target at all and these will in fact dominate the return signal. While it is possible to gate the detector to eliminate some of the background, it requires prior knowledge and understanding of the scene in order to do so effectively. Section 4.2.1 discusses how the background signal can be removed from the return, isolating the response from the target.

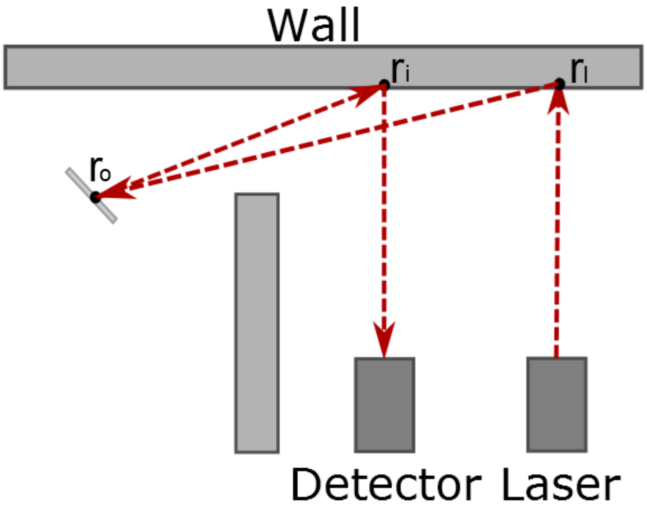

Figure 4.3: An example ray path scattered by a non-line-of-sight target. The light is first scattered by the wall, some toward the target. The target then scatters some of the light incident upon it back toward the wall. Finally a portion of that light is scattered onto the detector.

4.2.1 Background Subtraction

There are two methods for removing the background from the data set recorded by the detector. The first, and simplest, is only possible when a target free background measurement can be taken. The second method, involves taking multiple measurements from the scene and using the statistical median of response histograms to estimate the background signal. The median is used due to its sensitivity to changes in the scene [3].

When a target free background measurement can be taken, then a threshold after subtraction can be used to determine the target signal. For each set of data a response is considered to be foreground, scattered by the target in this case, if the condition in eq. (4.1) holds

$$P_{(t)} - B_{(t)} > \text{Threshold}, \tag{4.1}$$

where $P_{(t)}$ is the count in the histogram bin, with the target and $B_{(t)}$ is the count in the corresponding background bin, at time t. Typically, threshold values in the 10-15% range of the background signal, $B_{(t)}$, make for clear, easily distinguishable peaks with minimal loss of data.

If it is not possible to take a background measurement, which is a likely scenario outside of laboratory conditions, then the background must be estimated. This can be achieved with enough accuracy, by comparing the median response from multiple data sets, for the same scene, recorded at different times. If the target is moving, then the peak that corresponds to the signal scattered by it will move in time, as the target becomes closer or further away. It is expected that the speed of the target may potentially affect the size and shape of these peaks. Faster moving targets may cause the peaks to appear shallower and wider temporally. Objects within the scene that are not moving, such as the walls and ceiling, will give response peaks that vary on a much smaller scale and are thus easy to isolate.

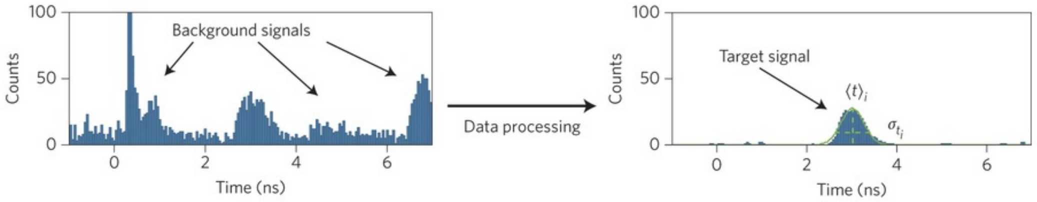

Figure 4.4: An example of how a typical signal looks before and after data processing. A histogram of arrival times is build up for each scene, shown left. Background subtraction allows for isolation of the target signal and fitting of a Gaussian to the signal's peak $\langle t_i \rangle$, with standard deviation σ_{t_i}, shown right. This figure is taken from Gariepy et. al [2]

.

This self-referenced background can be calculated with only a few acquisitions, in a matter of seconds. Should the target be stationary in this time, this method will not pick them up, as it only discriminates between moving and stationary objects in the scene. If the target begins to move however, locating it can begin as soon as the new background has been calculated.

For the case of verifying the model presented in this chapter, the simpler background subtraction technique has been used on the modelled data that relies on using only the known target rays in Zemax.

4.2.2 Position Retrieval

As has been stated above, the method for locating objects outside field-of-view relies on the recorded histograms of time-of-flight for the emitted photons. The response signal is gathered from all photons that land back onto the detector after multiple, random scattering events. As the histograms are built up over many pulses, they give the probability distribution for the arrival time of photons scattered back from the scene. After background subtraction, detailed in section 4.2.1, the remaining information is the probability of arrival times for photons scattered by the target.

It is then possible to map the time-of-flight information into a spatial probability density for the position of the target. If multiple detectors are used, the separate probability densities for each can be simply multiplied together: more detectors means a more accurate position can be retrieved for the target.

For a single detector, an arrival time will be recorded for each of the photons that is scattered from the scene. This time t_i is how long a photon takes, travelling at speed c, to scatter first from the centre of the laser pulse on the initial scattering surface to a point on the target $\vec{r_o} = (x_o, y_o)$ and then to the point on which the detector is centred $\vec{r_i}$. Possible locations for the target point $\vec{r_o}$ include all points on the surface of an ellipse defined by eq. (4.2) [2, 3, 4, 136].

$$|\vec{r_o} - \vec{r_l}| + |\vec{r_o} - \vec{r_i}| = t_i \times c \tag{4.2}$$

where $|\vec{r_o} - \vec{r_l}|$ is the distance from the centre of the laser point $\vec{r_l}$ on the initial scattering surface to the target and

$|\vec{r_o} - \vec{r_i}|$ is the distance from the target to the point at which light is collected, $\vec{r_i}$. Note that this method assumes only single scattering events occur within the scene, from the target. Secondary and tertiary scattering events add a source of error for this technique, but occur with low probabilities so as to be insignificant [2, 3, 4, 136].

Solving this equation gives an infinite number of solutions, lying on the surface of an ellipsoid with foci at $\vec{r_l}$ and $\vec{r_i}$. All possible scattering locations that can generate signals at the detector, at time t_i, lie on an ellipsoidal surface with evenly distributed probability.

In the absence of uncertainty in the measured signals, the data collected by the detector corresponds to a probability density $\mathrm{P}_{\text{ellipsoid}}$, for finding an object's location, given by eq. (4.3) [2, 3, 4, 136].

$$\mathrm{P}_{\text{ellipsoid}}(\vec{r_o}) \propto \begin{cases} 1, & \text{if} \quad |\vec{r_o} - \vec{r_l}| + |\vec{r_o} - \vec{r_i}| = t_i \times c \\ 0, & \text{otherwise} \end{cases} \tag{4.3}$$

It is possible to restrict this set of solutions by searching in a plane parallel to the floor. By making the assumption that a target's centre of mass will give the strongest return signal, a single plane can be selected at half the target's height. This also assumes that targets will not move significantly in the vertical direction. This is dependent on the types of target for the application, but can easily be changed and allows the ellipsoid to be collapsed into an ellipse. An error Δz, for estimating target height, can be introduced which will at worst be of the same order as the error $\Delta\vec{r_o}$, for determining the target coordinates. It will, however, be typically much smaller, decreasing as targets move further away.

A simpler form of eq. (4.3) can be written in elliptical coordinates, illustrated in fig. 4.5, as eq. (4.4), using $\vec{r_l}$ and $\vec{r_i}$ as foci

$$\mathrm{P}_{\text{ellipse}}(\vec{r_o}) \propto \delta(\epsilon - ct_i), \quad \text{where} \quad \epsilon = |\vec{r_o} - \vec{r_l}| + |\vec{r_o} - \vec{r_i}| \tag{4.4}$$

The proportionality sign here is as normalisation must occur over the surface of the ellipse, according to eq. (4.5), such that the probability of finding the target somewhere on the ellipse is equal to one.

$$\iint \mathrm{P}_{\text{ellipse}}(\vec{r_o})d\vec{r_o} = 1 \tag{4.5}$$

Uncertainties in the measured signals can be caused by a number of factors such as jitter in the laser or detector systems. Also the target, laser spot and detection point, all having a finite size, giving rise to uncertainty in the returned signal. From the experimentally recorded signals it has been found that these compound uncertainties have a Gaussian form [2, 3, 4]. As a result of this, the general expression for the probability density $\mathrm{P}_{\text{ellipse}}(\vec{r_o})$ becomes eq. (4.6)

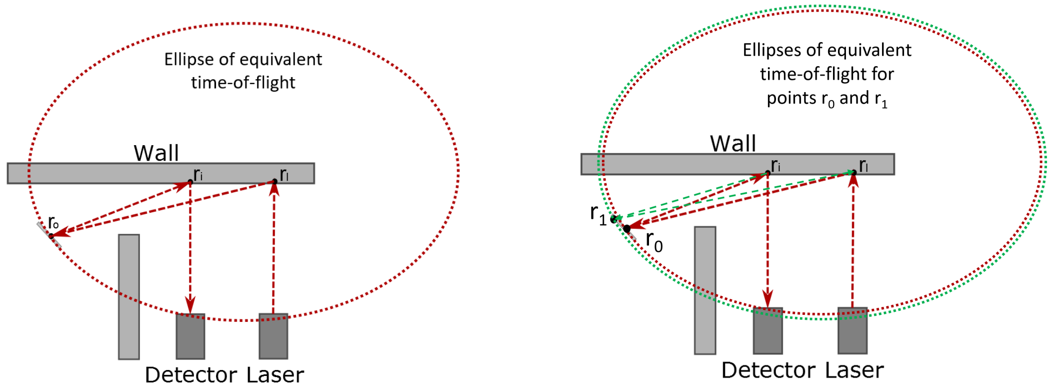

Figure 4.5: A simple illustration of the ellipse of equivalent time-of-flight for a ray incident on a target at point r_0 and another non-equivalent ellipse scattered from the same target, but from point r_1. For many rays scattered from the target, the point at whic these ellipses most-overlap gives the highest probability of location.

$$\mathrm{P}_{\text{ellipse}}(\vec{r_o}) \propto \exp\left[-\frac{(\epsilon/c - \langle t\rangle_i)^2}{2\sigma_{t_i}^2}\right] \tag{4.6}$$

where $\langle t\rangle_i$ is the mean arrival time for the isolated signal, and σ_{t_i} is the standard deviation.

As the probability densities for each detector are multiplied together, it becomes necessary to add a uniform probability density to $\mathrm{P}_{\text{ellipse}}$. If a single detector has a zero probability for a point, it will multiply all other probabilities by zero also. These occur because of the simplification of the ellipsoid to a two dimensional ellipse and the estimation of the target surface to a single point. The three dimensional ellipsoids will all overlap with the target surface, however, there is no guarantee that the ellipses will all overlap and indicate a single point. Adding the uniform probability ensures that no information is lost, should a single detector not fit the signal. Then the probability of finding the target in the search space $\mathrm{P}(\vec{r_o})$ considering the probability densities of all detectors $\mathrm{P}_i(\vec{r_o})$ is given by eq. (4.7)

$$\mathrm{P}(\vec{r_o}) = N\prod_{i=1}^{n}\mathrm{P}_i(\vec{r_o}) \tag{4.7}$$

where n is the number of detectors in the array, P_i is equivalent to P for each individual detector plus the added non-zeroing constant and N is a normalisation constant [2, 3, 4].

4.3 Model

The process for modelling experiments such as those outlined in sections 4.4 to 4.7, is presented here. Having accurate models allows users to quickly investigate a number of practical considerations about the experiment for little cost and without the difficulties associated with securing large laboratory spaces. These initial investigations can then be used to guide practical experimentation and develop better understanding. Figure 4.6 gives an

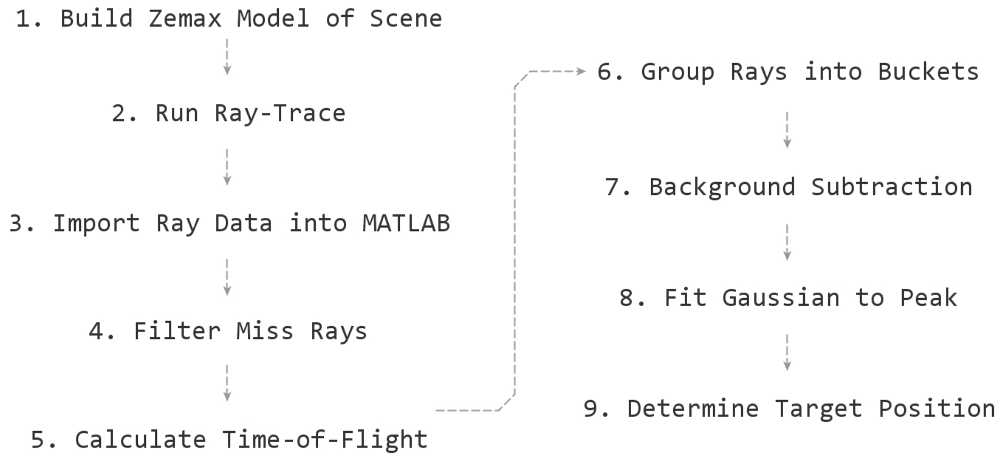

Figure 4.6: The process for building and analysing scenes, enabling the location calculation for a non-line-of-sight target.

overview of the algorithm used for building and analysing scenes to locate targets hidden from line-of-sight.

The first step required involves building a model of the scene in the ray-tracer, in this case Zemax. Utilising commercially available software allows for implementation of high accuracy models exploiting optimised and efficient ray-tracing algorithms. Non-sequential Monte-Carlo ray-tracing models simulate optical systems by casting rays and allowing for them to interact randomly with the objects defined in the scene. In most cases it is possible to model a scene using only the objects and tools within Zemax. However, if required Zemax allows for importing of more complex objects, such as those generated by three-dimensional computer aided design programs.

The scene construction is straightforward and can have as much or as little detail as the scene demands. Floors, walls and ceilings are constructed of simple rectangular objects with dimensions and orientations corresponding to the scene, as shown in fig. 4.7. The target is placed in four separate locations to demonstrate the differences in scene responses and the ability to distinguish between them.

4.4 Verification

This section describes the experiment used to verify the modelling approach presented in section 4.3. The entirety of this practical work was carried out by the team at Heriot-Watt University and these results were then processed by the author [1, 136].

The results from the model are compared to those gathered using a single pixel SPAD detector and laser source. Five configurations of the scene will be measured and modelled. The aim of this experiment is to demonstrate how the model can produce time-of-flight data that is similar to that collected by a physical system. Only a single detector will be used in this case, meaning position retrieval of the target will not be possible, as this is not the aim. Section 4.7.3 discusses retrieving the position of hidden targets using this model.

Table 4.1: The X and Y coordinates for each of the target positions to be used in the experimental verification of the model. Distances are measured from the laser spot on the wall, as illustrated in fig. 4.7.

Position	X (cm)	Y (cm)
1	-73.5	-98.0
2	-48.5	-93.0
3	-61.0	-110.5
4	-58.0	-47.0

4.4.1 Experimental Setup

Figure 4.7 illustrates the scene measured in the experiment. A 780nm wavelength laser source is placed 160 cm from a wall, to be used as the first and last scattering surface in the ray path. In this experiment a single pixel SPAD detector is used to collect the scattered signal. Light is coupled into the detector by a 1.27 cm, 20×, 0.40 NA lens and a 105µm, 0.22 NA fibre-optic cable. The lens collects light from a point 24.5 cm to the left of the laser spot. The target in this case, is a cardboard square, 14.9 cm on each side, oriented at 45° to the wall and then placed in one of four positions as given in fig. 4.7. The fifth configuration is the same scene in the absence of the target, giving a result for the background signal from the scene.

The laser illuminates a spot on the wall with 1 mW average power, running at a rate of 80 MHz, with a 10 fs pulse length. The SPAD detector is operating in TCSPC mode, this means that for each pulse only a binary result is recorded in each of the 64 ps bins, meaning that the detector only records the difference between observing any or no photons. If more than one photon is detected within that period, the detector simply counts one. The pulse length must be very small in comparison to the size of the detector bins.

As multiple pulses are fired, theses bins are accumulated over the acquisition time into histograms. Information about the scene is then gathered from the shape of these histograms. Should the objects within the scene move, these histograms will change significantly.

Observing the scene will generate unique histograms of response signals dependent on the target's location. After processing, each signal can then be used to predict a position for the target.

It should be reiterated here that once the model has produced time-of-flight data for the response signal, the post-processing to determine the position of the target can be treated in exactly the same manner as experimentally collected data.

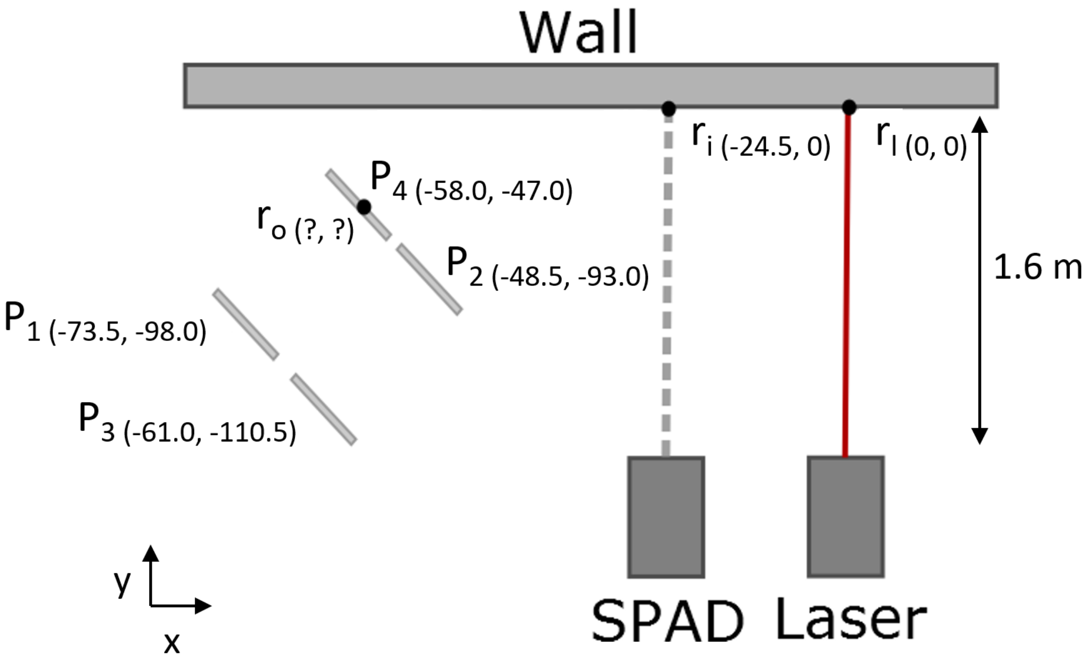

Figure 4.7: The experimental setup involves directing the centre of our laser spot onto a wall 1.6 m away. A 14.9 by 14.9 cm, square target is oriented 45° to the wall and placed at various positions. A fibre-coupled SPAD detector collects the backscattered light from a point on the wall 24.5 cm to the left of the laser spot [1, 136]

4.4.2 Non-Sequential Ray-Tracing

The albedo of the target and background scattering surfaces used in the experiment were approximated to be 20%, which is a number often used for targeting models [137] . Therefore, an 80-20 Lambertian scattering function was applied to each surface in the model, including the target, meaning the surfaces absorbed 80% of the incident light and scattered 20%. A Zemax source object was used to model the laser, this enables full control of the output to closely match that used in the experiment, which were chosen to replicate the parameters used in the original demonstration [2, 3, 4, 1, 136]. The laser's wavelength was set to 780 nm in the model.

Finally, a single SPAD detector was set to absorb all rays incident upon it, eliminating the potential for rays to be counted multiple times. The model differs from the experiment here in that the detector will count every ray that strikes it. The real SPAD detector pixels each record a binary flag, this flag is set to true if anything is detected within the same time bin, therefore, if the same detector pixel is struck multiple times in a single window, only one event is recorded. The model will only take a single run as opposed to the real system which builds up these counts over multiple pulses. It would be posisble to enhance the model to consider phenomena such as jitter, dark counts, and finite detection efficiency.

For each of the scenes, rays were traced in Zemax in simple ray-splitting mode, meaning that only one path is continued in the model and the overal intensity is reduced with each bounce. All of the analysis is then performed in MATLAB. As the ray database contains all of the spatial and intensity information about the ray at each interface, a great deal can be accomplished during the post-processing in MATLAB. The ray-tracing software has a limit of 4×10^9 rays, a number of 1.5×10^9 was chosen for the modelling here. Although increasing the number of rays does improve the model accuracy, there is a significant time increase for the model as the program approaches its limit,1.5×10^9 was determined to be a good balance between accuracy and speed. Firstly, once a ray-trace has completed and the resulting ray database has been saved, the binary output file is read into MATLAB. From each of the ray segments, total time-of-flight can then be calculated from the lengths and summed over the whole ray. A histogram is then constructed, counting how many rays fall within 64 ps bins from zero to the longest time-of-flight ray. This simulates the detector's 64 ps timing resolution and ensures the modelled data resembles that from the experiment. When rays are incident upon an interface they are labelled with the corresponding object number in Zemax, this makes it is easy to isolate which of them have scattered from the target by selecting only those rays with the appropriate label. This allows for perfect background subtraction, which can help to mitigate some of the computational limitations of tracing a very small number of rays, compared to the light used to interrogate a physical scene. Once the histograms have been built, both the experimental and simulated data can be compared for time-of-flight to target, as described in section 4.2.2.

4.5 Results

For the experimental data, a simple background subtraction scheme was utilised. For each set of experimental data a response was considered to be foreground, scattered by the target in this case, if the following condition held:

$$P(t) - B_{(t)} > \text{Threshold}$$

where $P_{(t)}$ is the count in the 64 ps bin for the scene interrogated with the target. $B_{(t)}$ is the count in the same bin, for the signal recorded of the background scene, at time t. The threshold value in this case was chosen to be 10 % of the background signal, B(t).

For the modelled data, those rays that had hit the target were extracted and placed into histograms with bins at the same positions. Figure 4.8 shows the two sets of data, corresponding to the target response at position four. The modelled data in blue, clearly overlaps with the experimental data in red. Both are measured in arbitrary units as the experimental data results from emitting a far greater number of photons when compared to how many rays can be traced. The model, in this case, also doesn't account for quantum efficiency or detector noise; these could both contribute to the differences in shape between both signals.

This difference in treatment of the background subtraction enables this modelling technique to yield results, comparable with experiment, and is the only point at which the post-processing differs from the real system. Computational limitations of the ray-tracing make multiple traces costly, each one taking several hours for the

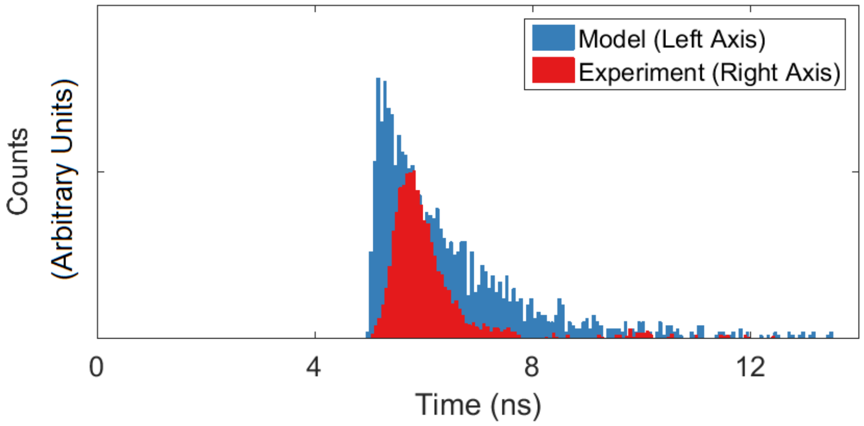

Figure 4.8: Histograms of response signal from target in position four. The data collected from experiment is displayed in red and measured against the right hand y-axis, and data from the model in blue and measured against the left hand y-axis.

scene described in section 4.3. The number of rays returned also means that the same background subtraction technique used on the experimental data leaves the modelled data unusable, as the small sample size leads the median to become less reliable as a proxy for the background signal. The background subtraction used in the model also leads to differences in the appearance of the histograms, as can be seen from section 4.5. The smoother rising edge of the experimental data histogram is caused by a combination of the threshold subtraction and the jitter both in the laser and the detector. The SPAD used in this experiment had an electronic jitter of 120 ps, leading to the delay of many of the counts, smoothing the appearance of the signal. It is also suspected that the median response background subtraction approach, also causes the long trailing end, observed in the target response histogram, to be lost and, therefore, not considered when fitting the Gaussian to determine flight time.

Table 4.2 compares the mean and standard deviation of the modelled time of flight results, to those gathered in the experiment. For the target in positions 2-4, the differences in mean time-of-flight between the model and experiment are all less than 0.4 ns. The first target position sees a larger difference in these means of 1.28 ns. This variance may be due to the lower number of rays incident on the target as it is moved further away from the wall, an affect discussed in more detail in section 4.7.3. The standard deviations of the results from positions 3 and 4 are also close at less than 0.2 ns, showing that the target position curves are have close peaks and similar width for these positions. The curve for position 1 is far broader when calculated from the experimental data, position 4, however, is the opposite, whereby the experimental data produced a narrower curve than that of the model. This variance, again, may be attributed to the relatively low number of target rays in the model, compared to the experiment. For positions 2-4, these results correspond to differences in estimated round-trip path length of less than 12 cm for a 5 cm target.

Table 4.2: Mean response time (and standard deviation) from target for both experimental and modelled data for each of the four target positions, given in table 4.1.

	Mean Target Response	
Position	Experiment (ns)	Model (ns)
1	8.53 (± 1.88)	9.81 (± 1.23)
2	8.42 (± 1.16)	8.49 (± 0.92)
3	10.06 (± 0.88)	9.93 (± 0.99)
4	6.10 (± 0.94)	6.48 (± 1.47)

4.6 Discussion

As can be seen from the results presented in section 4.5, this is an accurate approach to modelling LIDAR experiments involving scattering events. All of the modelled results have returned similar response times from the target to those calculated from the experiment. The results also show comparable standard deviation, corresponding to similar spatial precision when calculating range to target, for both the modelled and experimental data. As it is possible to perfectly isolate the target response using the model, small photon counts are not lost to the background, resulting in no loss of information.

As this model produces comparable results to those gathered from experiment, it provides a means to quickly model arbitrary scenes and apparatus for non-line-of-sight LIDAR experiments. Investigation of technical requirements and the limitations, for any given scene, can be conducted computationally, such as the maximum distance of the detector to the first scattering surface, that still allows for the target position to be resolved. By using this model, it is possible to investigate scenes that would be difficult to construct given physical constraints, such as acquiring very large lab spaces, or constructing complicated scenes. It also permits a deeper understanding of the affects, that different objects within the scene have on the detector signal and the resolvability of the target signal. This is investigated further in section 4.7.

Figure 4.9 gives two views of a more complex room simulated in Zemax. The dimensions of the room in this case are 1 by 1.2 m with a 0.6 m wall separating the 10 by 30 cm target from the emitter and detector. Again all surfaces have an 80-20 Lambertian scattering function applied to them and 1.5×10^9 rays were traced. The response from this room is shown in fig. 4.10, with the total response given in blue and only those rays that hit the target overlaid in red. As can be seen, the response from the target is small compared to the early reflection from the environment, The secondary peak at the trailing end of the data arises from the reflections caused by the walls at the far end of the room. The origin of these secondary peaks is not straightforward to prove from the experimental data, the model allows full exploration of the entirety of each ray path, making it easy to determine. it may be possible to isolate individual components by covering each of the background scattering surfaces with absorbing material. Although small compared to the size of the signal from the back wall, it is possible to isolate the target in the experimental case. This can be achieved by either taking a background signal, as done in section 4.4, or when that is not possible, using the median signal over several seconds serves as a good

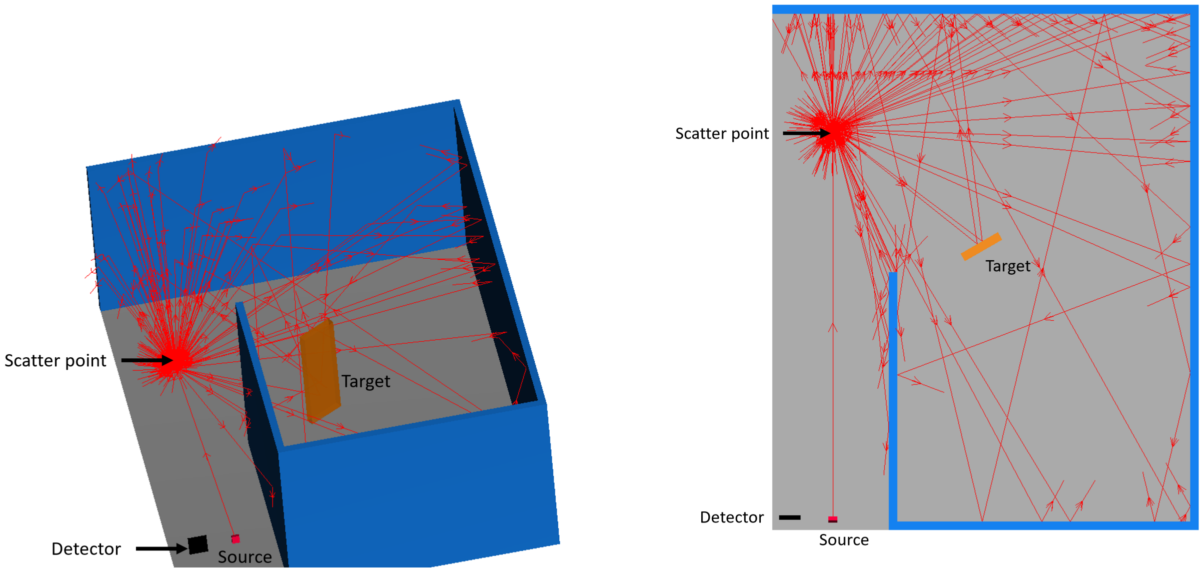

Figure 4.9: Top-down and rotated view of a more complex scene model in Zemax. Here the black square represents a detector surface, rectangular objects make up the walls and floor in blue and grey respectively. The orange rectangle in the centre of the room is the target. Finally the red cuboid, not required, has been added as a visual reference point for the laser source. Note that of the 200 random rays, drawn in red, none are reflected back to the detector region.

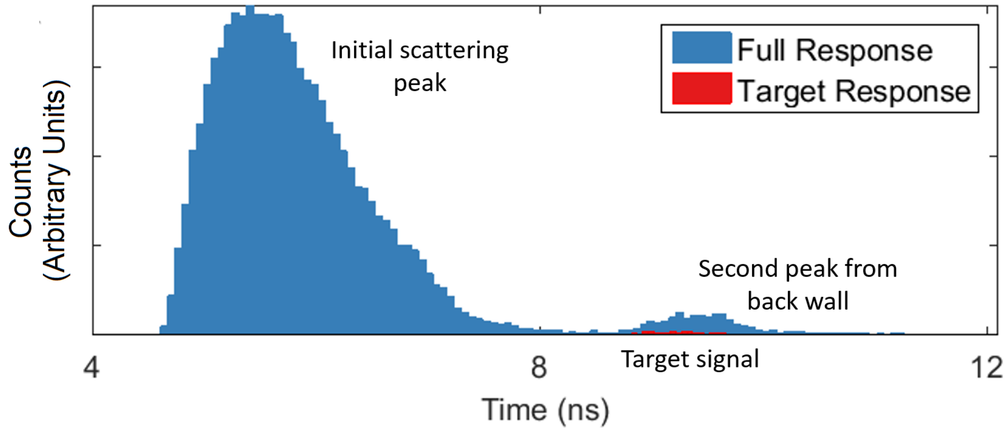

Figure 4.10: Graph depicting the response from the model presented in fig. 4.9. The full response signal in blue and only the rays that hit the target have been overlaid in red, highlighting how small a proportion of the rays returned from the scene can be used to determine the target position.

estimate for the background. This however only allows you to detecting moving targets, as anything stationary will be considered background using this median subtraction technique. [3, 4].

As the ray database contains all of the spatial and intensity information about the ray at each interface, a great deal can be accomplished during the post-processing in MATLAB. Once the rays for a scene have been traced, making quick comparisons between many different setups can be accomplished through processing. For example in the case of a camera based system, as proposed by Gariepy *et. al* [2], varying the camera's field of view through a range of potential values would be straightforward. Beginning with a flat surface as a detector, accepting light from any angle, searching for rays scattered only within a particular geometry would allow the user to simulate the response from a different field of view. This process could then be optimised using a merit function to determine the best field of view for experiment. Although not explored in this work, many parameters could be determined in this manner including pixel pitch, camera optics and source properties.

Techniques have been presented that allow for more precise location of a target, at the cost of requiring additional spatial information. This can be achieved by raster scanning [120] or use of a camera [2]. The modelling approach discussed in this chapter is compatible with these more complex techniques. For the sake of verification however, the simpler case of a single pixel and static source was used.

4.6.1 Eye-safety

Tracing large numbers of rays into a scene allows for a good estimation of the probability of return from a target. To track the motion of human-sized targets in large spaces, sub-second time-scale position retrieval must be accomplished, processing the experimental data at Hz rates or higher. As scenes become more complex and larger, more signal is required to ensure an adequate response from a target to distinguish it from the background. This translates to a requirement for more photons, and thus higher energy pulses, for physical systems. If this technology is to be deployed to track motion of uncooperative or unaware human targets, then eye-safety must also be considered, as accessible emission limits will define the limitations for maximum pulse energy.

The following calculations were made, using the maximum accessible emission limits for a Class 1 laser, from the safety standard documentation [138]. At a wavelength of 780 nm a single pulse must be limited to 38nJ, with a photon energy of 2.55×10^{-19}J this gives a total of 3.14×10^{11} photons per pulse. Were a source with a wavelength of 1530 nm used, a single pulse is limited to 8μW. Therefore for a 1ps pulse, 6.16×10^{13} photons of 1.30×10^{-19}J can be emitted and still be considered Class 1, a factor of 196 increase in the number of photons. To acquire the data necessary for positioning, however, many pulses must be emitted. The Class 1 limitations for repetitively pulsed lasers are more restrictive, reducing the average permissible power emitted to 0.57mW at 780 nm and 10.12 mW at 1530 nm. This still gives a factor of 34.8 increase in the number of photons that can be emitted for the longer wavelength system, offering proportionally shorter acquisition times. Tracing the scenes at both of these wavelengths gives no significant difference in the response, because surface roughness on typical objects is on a scale much larger than these wavelengths. One factor that may limit the improvement in acquisition time for a longer wavelength source, however, is detector efficiency and temporal resolution. It would also be important to investigate other efficiencies for using different wavelengths, as factors such as quantum efficiency, albedo, and scattering efficiencies could have a significant effect.

Determining if useful data can still be obtained within safety limitations is dependent on the scene. This model could be used to determine whether or not a target is resolvable by lasers, within safety regulations, for a given scene, With long enough acquisition time, any scene can be successfully interrogated. Due to the random scattering nature of surfaces typical in most scenes of interest, more power will correlate with faster acquisition times, eye-safety regulations limit the maximum permissible power for use in real world applications. Section 4.7 discusses some of the relationships determining how the specific make-up of a scene affects the response observed at the detector and, therefore, how these power limitations affect the ability to detect a target for a given scene.

Increasing the number of detectors used in the location of an obstructed target could drastically improve the acquisition time for a scene. This is realised in relation to other methods in the literature in the work by Gariepy et al. [2, 3, 4]. Modifying the collection optics could also have a significant impace on acquisition times and the minimum viable pulse energy for a scene. Although not conducted here, this model could be utilised to explore both of these approaches.

4.7 Scene Limitations

This section utilises the model to demonstrate some of the underlying relationships, between scenes and the apparatus parameters, for this positioning method. The outputs are discussed in the context of optimising performance and addressing the limitations of the technology. Many of these relationships are dominated by Lambertian scattering, which scatters the total incident power out into a partial sphere from the point of incidence.

4.7.1 Detector Range to Scattering Surface

This part of the investigation looks at the variation in power on the detector as it is moved further from the scattering surface. The model consists of the two main scattering surfaces, the wall and the floor. The laser source is placed 1 m from the wall, the detector is place 1.7 m from the wall, collecting light from a point 0.245 m to the left of the laser spot. A virtual square object with 2 mm sides was placed on the wall, centred in line with the centre of the detector. To simulate the requirement of collecting light from a single point, only light that's final scattering event was on this virtual object was recorded.

No collection optics were modelled in this instance, the detector was simply modelled as a plane surface. It is possible to mitigate the effects that detector range to the scattering surface with a sufficient detection scheme. Furthermore, Zemax allows the user to fully explore detection schemes to see how they would impact the making this a valuable tool in the design of these systems.

The target is a constant size 0.2 m square, kept in a constant position at 0.3 m from the detector point on the wall, at an angle 45° from the wall normal. The target centre, laser source, and detector are all 0.5 m above the floor and the wall is 1.0 m tall. The setup can be seen in fig. 4.11.

Figure 4.12 shows the detector range investigation results, where total power is indicated for each detector distance. The relationship between power and distance is the same for both the scene response and target incident rays, with power diminishing as distance increases. It is difficult to determine where the detector noise floor would be for a given scene using this model, as only a single light source (the laser) is used in the model, and it

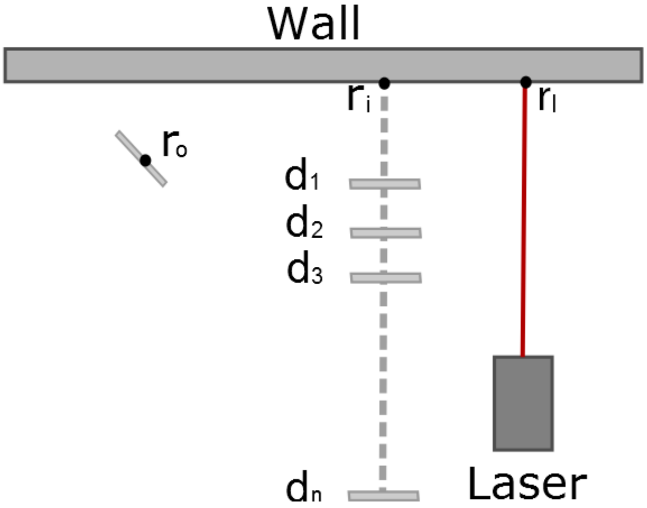

Figure 4.11: The model used to determine the relationship between response signal and distance of detector from initial scattering surface. Multiple detectors are placed at regular intervals from the surface. They do not interact with rays, allowing for each configuration to be tested simultaneously.

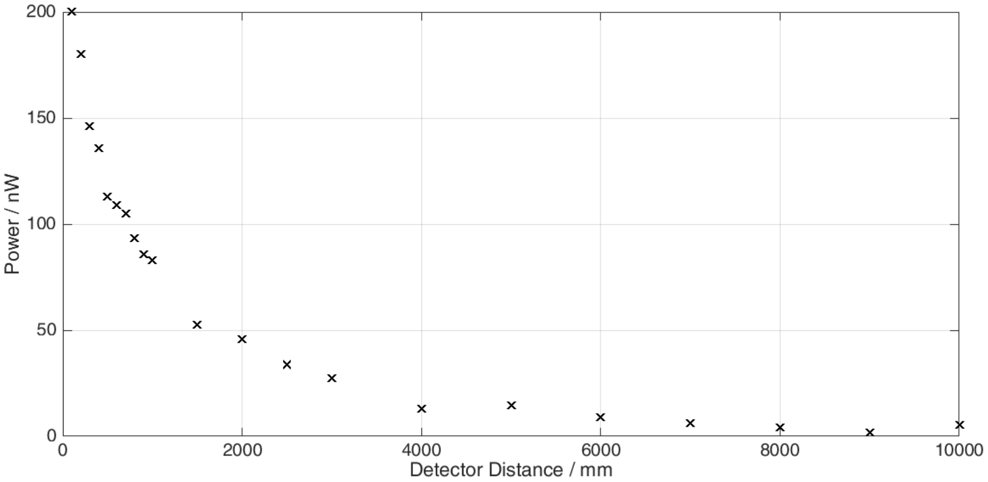

Figure 4.12: Total power measured on detector, against distance of detector from scattering surface.

would be computationally expensive and time consuming to simulate enough rays to use the same background subtraction method that is used with the real system. This is to be expected for a Lambertian scattering surface. The wall in all of these experiments is modelled as an 80-20 Lambertian scattering surface, which closely resembles the behaviour of a matte white wall, with an albedo of 20% [137]. Lambertian scattering scatters the total incident power out into a partial sphere from the point of incidence. As this is the case we know the intensity I at any point on the surface of that sphere must obey an inverse square law, namely $I = I_0/r^2$, where I_0 is the intensity at point on the scattering surface and r is the distance from that point ie the radius.

It can be expected then that a plot of the inverse of the square root of the measured signal, $r = 1/I^{1/2}$, should form a straight line. This relationship is observed in fig. 4.13, where the power measurements of the target-incident rays have been transformed accordingly and the distance of the detector from the scattering surface is measured across the x-axis. A line of best fit has been plotted to further illustrate this relationship, which in this case has

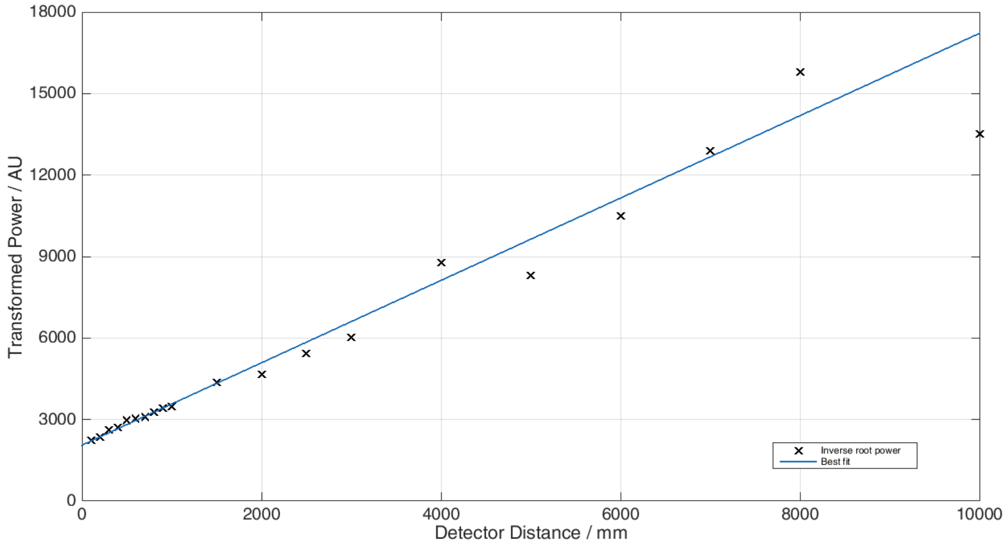

Figure 4.13: Inverse square root of power incident on detector against distance of detector from scattering surface. A straight line of best fit highlights the relationship, the intercept of which corresponds to the mean power incident at $r = 0$, on the scattering surface.

an R^2 value of 0.983. This line also has an intercept of 2057 which, when translated back to power, equates to 236.34 nW. A virtual detector of area 4 mm^2 was placed on the scattering surface at the detector point $\mathbf{r}_i$, where a mean power of 212.40 nW was observed from target incident rays. This demonstrates that the inverse square root transform of the power measurements holds well and the line of best fit can act as a means of estimating the total expected power from incident rays at point $\mathbf{r}_i$. The close fit of this $r = 1/I^{1/2}$ relationship highlights the substantial loss of available signal as the detector moves further from the scattering surface.

This inverse square relationship will mean that acquisition times for any scene will increase proportionally to the square of the distance to the scattering surface. In determining the maximum range that this method can be used at, the minimum response require from target-incident rays must be understood. The observed background signal will also fall off at the same rate, meaning that the background subtraction for typical scenes should not be significantly hindered until the target-incident power falls to within noise levels. The acquisition time for any scene will be specific to the scene observed, target and detector parameters and the power of the illuminating laser. Should acquisition time increase, the minimum resolvable spatial precision that the method could achieve for positioning the target suffers. If the acquisition time reaches the scales at which a potential target could travel significantly, within the scene, this minimum achievable spatial precision increases in line with the scale of that movement. Very short acquisition times are desired as the method assumes that no significant movement can occur in the acquisition time frame. As stated however, this can be accommodated for by improved detectors and signal collection or increased illumination power from the laser source. Acquisition times will scale inversely proportionally to any increase in the power detected. There are, however, practical limitations on these improvements and they are somewhat dependent on one another. For example, a certain level of laser power output will lead to bleaching of the detector, for a given scene, this would need to be remedied by improved detectors or signal processing. Complex signal processing methods come at a computational and therefore time cost, which can limit the achievable performance.

4.7.2 Detector Angle to Scattering Surface Normal

In this section we investigate how the angle of the detector, from the scattering surface normal, affects the response. The total power on the detector from all rays and the power from rays incident upon the target are compared.

Figure 4.14 illustrates the setup used. The detector is collecting light incident upon the wall at a spot 0.245 m to the left of the initial scattering event. A virtual square object with 2 mm sides was placed on the wall, centred in line with the centre of the detector. To simulate the requirement of collecting light from a single point, only light that's final scattering event is from this virtual object was recorded. The target is kept at a constant 0.2 m square size and in a constant position 0.3 m from the detector collection point at an angle of 45°. The detector is kept at a constant distance from the wall of 0.2 m and the angle is varied from 0-60 ° by steps of 1° The target centre, laser source, and detector are all 0.5m above the floor and the wall is 1.0 m tall.

Figure 4.16 illustrates the total power of the target incident rays, observed on the detector, for the various angles through 0-60deg. The variation in power throughout this range follows no trend and can be attributed to the relatively low ray sample numbers and random nature of the model. As the surfaces in this experiment have been modelled as Lambertian scatterers, the observed power at any given angle can be predicted using Lambert's Cosine Law [139]. Lambertian scatterers are often used in targetting models for defence applications. Normal to an ideally diffuse scattering surface, the intensity observed I_0 is given by eq. (4.8)

$$I_0 = \frac{I\cos(\theta)d\Omega dA}{d\Omega_0\cos(\theta)dA_0} \tag{4.8}$$

where $d\Omega$ is an angle into which rays are scattered from the surface area dA_0. At angle Θ to the normal, the scattering area appears as dA and the proportion of rays scattered in to the solid angle is given by $Icos(\Theta)d\Omega dA$ [139], which reduces to 0 at an angle of 90°, as the perceived area of the surface increases. Figure 4.15 illustrates this relationship. This relationship tells us that we expect to observe the same power regardless of angle between observing and scattering surface normals.

Figure 4.16 and section 4.7.3 also illustrate that the compound angle, between the detector and the target surface normals, has no significance on the power observed for lambertian scatterers. If there were to be significant variation in power with observing angle then the acquisition time for scenes would be directly impacted. The slight variation seen in fig. 4.16 is from the relatively low numbers of rays that could be traced using the approach presented here. The computational limits of the ray-tracing program mean that only 4×10^9 rays can be traced to sample each scene. As more rays are traced, the variance between these measurements decreases.

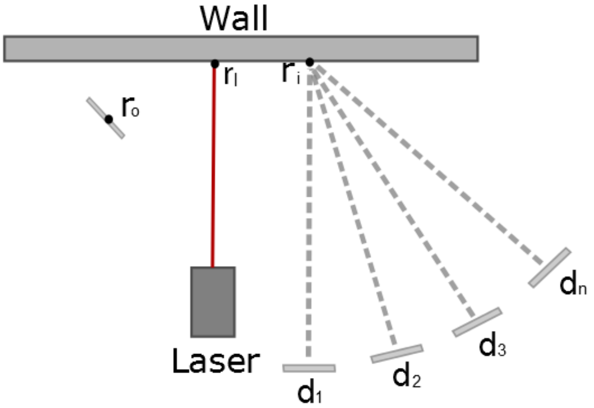

Figure 4.14: Several detectors placed equidistant from the detector point on the scattering wall, at 0-60° angles between the surface and detector normals.

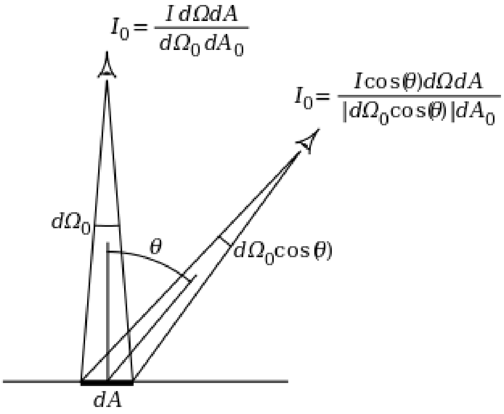

Figure 4.15: Observed intensity (from a Lambertian surface for a normal and off-normal observe, where dA0 is the area of the observing aperture and dΩ is the solid angle subtended by the viewpoint aperture [140].

The results presented throughout this chapter do not consider surfaces with a strongly specular scattering component. In reality, many surfaces exhibit significant specular reflection. The objective for testing each of the factors presented here was to highlight some of the core relationships between target and detector position for these systems. Modelling differing far more complex surface types is possible with Zemax and this model would, therefore, be capable of investigating these relationships.

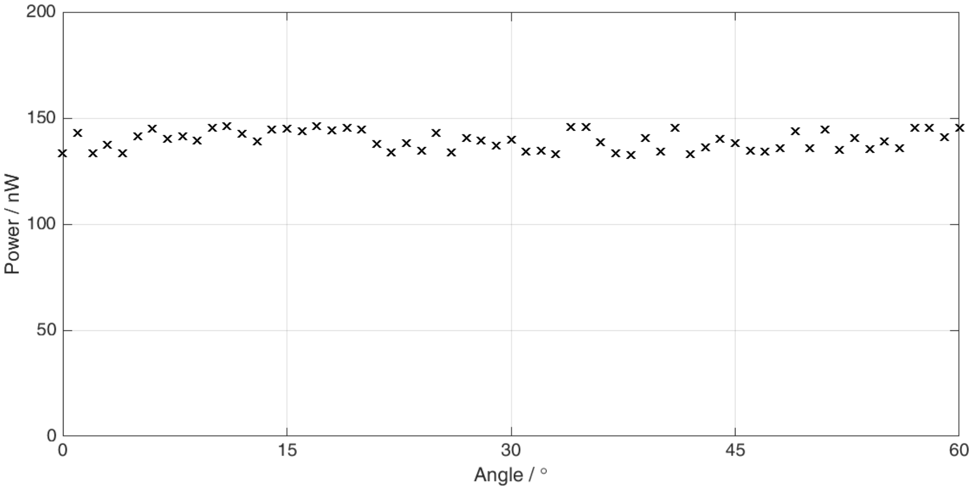

Figure 4.16: Total power on detector from target incident rays was measured for equidistant detector positions, with angle between the detector and scattering surface normals was varied from 0-60°.

4.7.3 Position of Target

In this section the model will be interrogated to look at how varying the position of the target, within the scene, affects the response on the detector. In this case three detectors were used to interrogate the scene. Each was placed at equal distance of 1 m from the scattering surface and separated by 100 mm from one another. A square target of with 50 mm sides was then placed within the scene and both its distance from the scattering surface and its angle to the surface normal were varied. The distance from the target's centre to the detector point on the scattering surface was increased from 100 to 1000 mm in 100 mm steps. The target's angle from the normal was increased from 0 to 60° in 5° increments. The target centre, laser source, and detector are all 0. 5m above the floor and the wall is 1.0m tall. Other than more target positions and the additional two detectors placed 100mm apart, the model here is identical to that presented in fig. 4.7 illustrating the experimental verification.

For each of the distances, the mean power from target-incident rays that were observed by each of the detectors can be seen in fig. 4.18. As the target moves further from the scattering wall, the power seen on the detector rapidly decreases. Figure 4.18 is shown on a logarithmic scale as the power change is so great, from 222.81 µW at 100 mm to 15.20 nW at 1000 mm.

Figure 4.19 shows an interesting relationship similar to that as seen in fig. 4.13. Section 4.7.1 demonstrated the inverse square relationship that the power observed had with distance from scattering surface. Following from that, as the light travels from the initial scattering point to the target and back to the detector point, two inverse square relationships should therefore be observed in conjunction. Good agreement is seen with the fourth root power dependency, on target distance from the detector point, as expected. This trend is illustrated in fig. 4.19.

It can be seen from fig. 4.19, that the fourth root of the incident power varies linearly with distance. Although there is a difference in distance from scattering point to target and target to detector point, making the assumption that they are equal provides a good estimate of expected power as distance of target changes. The R^2 value

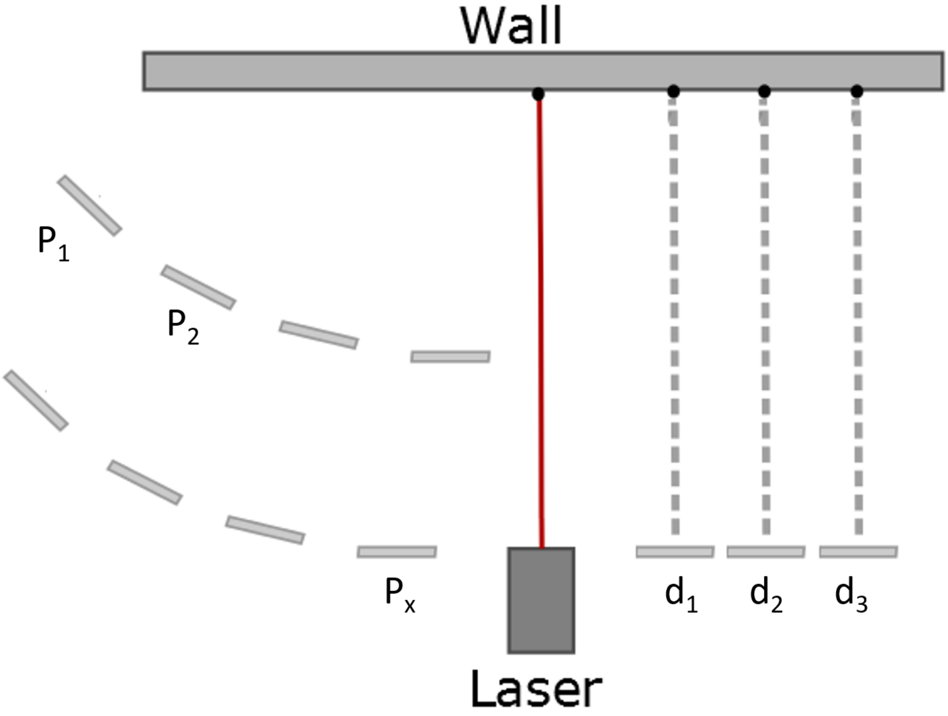

Figure 4.17: Illustration of the scene for used to determine relationship between target position and response. Here three detectors are used in positions d_1, d_2, d_3 and the overlapping probability densities calculated from their time-of-flight histrograms are used to estimate the position of the target. Both angle and distance from the scattering point will be varied as the target is moved through positions P_1 to P_x.

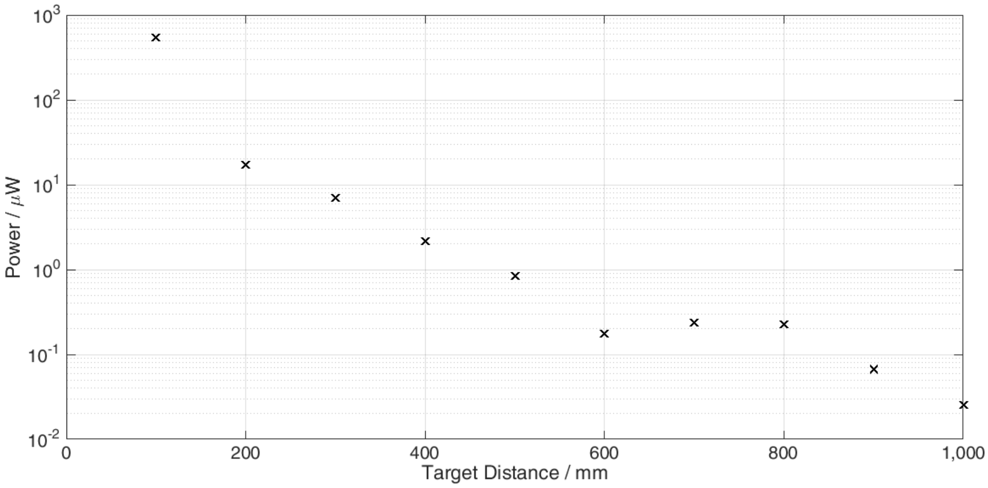

Figure 4.18: Mean power variation measured against target distance from detector point

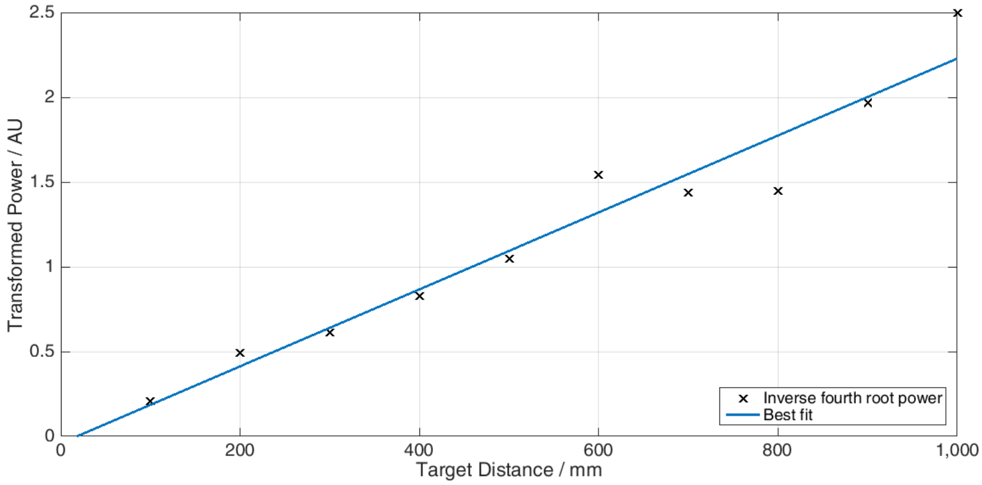

Figure 4.19: Translated inverse fourth root mean power variation measured against target distance from detector point

of the best fit line is 94.37 % showing that it is a strong linear fit. The intercept however falls below zero in this case, preventing meaningful comparison to the power on the wall as this would result in an non-physical value. This comes about, in part, due to the assumptions about path length made and the uncertainty in the data. The lower number of target-incident rays at larger distances make predictions more difficult as the sample size diminishes. At shorter distances, the path difference between scattering point to target and target to detector point also impact on the accuracy. As this path difference is larger, relative to the whole path length, the assumption made when plotting this graph and the inverse fourth root relationship has a greater impact. The method does illustrate the underlying relationships however and confirms the expected behaviour that the power would vary with the inverse fourth root of the target distance.

It can be seen then that as the target moves further from the scattering/detector points, on the initial surface, the acquisition time must increase to compensate for the loss of power observed back at the detector. A tenfold increase in target distance would translate to roughly a thousandfold increase in acquisition time to observe the same power.

Another consideration here is the total time-of-flight of the entirety of the laser pulse. For every 1m the target moves from the scattering point, the total time-of-flight increases by $2 \times (1/c) = 6.67 \times 10^{-9}$ seconds. The laser used in section 4.4.1 ran at 80 MHz, which means there is only a 12.5 ns gap between pulses. If the total-time-of flight for the target-incident rays exceeds the time between pulses, there will begin to be overlap. This overlap would make it impossible to take meaningful readings, as the light taking the shortest path, from the scattering point directly to the detector, cannot be gated. The repetition rate of the laser illuminating the scene would then need to be reduced to afford more time for each pulse, further impacting the acquisition time.

Similarly to section 4.7.2, the results here showed insignificant variation with the target's angle to the scattering surface normal. As eq. (4.8) states, the power observed is constant regardless of angle. Only should the target (or

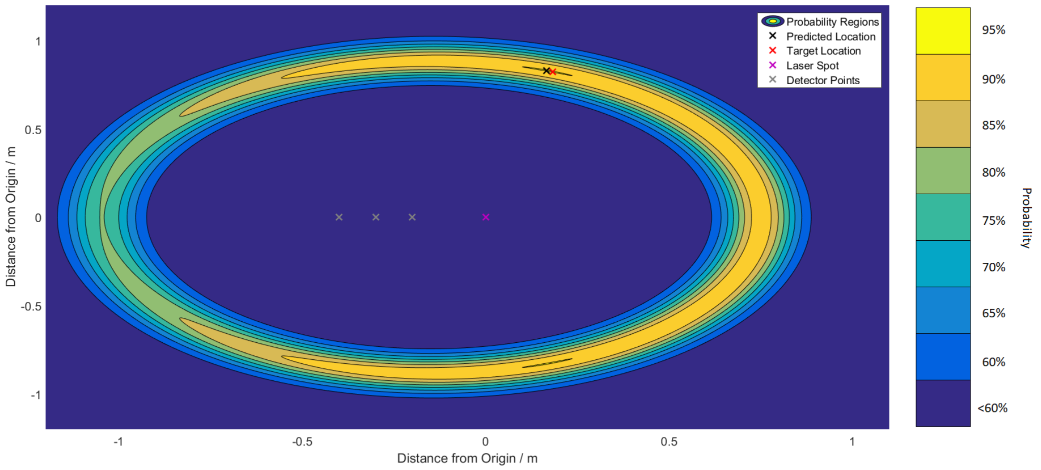

Figure 4.20: Map of results from the target position experiment. With a purple cross indicating the laser spot on the scattering surface and the origin, the grey crosses represent the points on the scattering surface from which the detectors gather light. A red cross indicates the actual position of the target's X-Y centroid. A black cross indicates whereabouts the model predicts the target to be. Finally each coloured area around the predicted position represents a confidence interval for the target's location, as predicted by the model.

detector) rotate about its own axis, would there be a significant impact on the observed power, as this effectively changes the aperture size through which these observations are made. This effect would be lessened in the real case due to the three dimensional nature of real targets. Whereas the change in aperture for an observed plane varies greatly as it rotates, this is not true at all for spherical or cylindrical targets. The use of planes in this experiment serves as a first order investigation into these effects, ensuring constant conditions. A more thorough examination into target shape could yield more accurate results in this context but is beyond the scope of this work.

Using eq. (4.5), predictions could be made for the target positions from the time-of-flight data gathered from the model. By combining the probability maps, from each of the three detectors, it becomes possible to construct probability regions for locating the target. Figure 4.20 illustrates how these look when mapped, with the highest likelihood region in yellow and the lowest in dark blue. Also the target location in red, the scattering point in purple and the three detector collection locations in grey are marked for reference. The target's predicted location, noted by the black cross here, is selected as the highest probability location within the searched space. Both the target's actual coordinates and the prediction are located within the small, highest probability region.

In fig. 4.20 the wall exists at $y = 0$, on which the laser spot and detector collection coordinates are located. It can also be seen that a second region, with equivalent confidence bounds as that in which the target and prediction are located, can be found beyond the wall. This is a product of the mathematics involved in making the prediction, using these methods either of these regions could likely hold the target. This method will always produced two regions of equal probability such as this, as these have equal distances from the foci of the ellipse. Often, these secondary regions occur within line of sight of the detector or, as in this case, beyond an obscuring surface which allows for them to be easily excluded when making predictions.

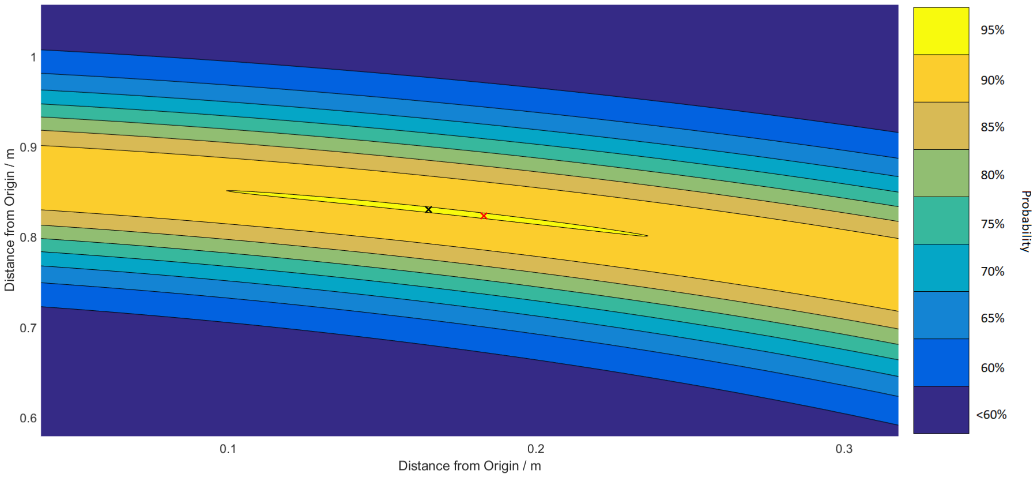

Figure 4.21: Zoomed map of positioning results.

Figure 4.21 is a closer view at the target location from fig. 4.20. In this case the distance from the prediction to the target coordinate was 18.4 mm. Over all of the predictions made using the model, a mean accuracy of 30.2 mm was achieved. Experimental results have seen precisions in the region of 25.4 mm [2, 3, 4] but these have used 1024 detectors in arrays and are not limited by computational power limiting ray numbers when interrogating the scene.

4.7.4 Target Size

It is expected that the final power on a detector, from target-incident rays, would increase as the size of the target increases. This section investigates how target size affects the response at the detector.

With the target and detectors kept at a single position, the model used here is similar to that depicted in fig. 4.13 considering only a single detector position. Both the target and the detector are kept at constant positions and angles. A square target was used of varying surface areas increasing from 10 to 100 mm^2 in 5 mm^2 increments. The target centre, laser source, and detector are all 0.5m above the floor and the wall is 1.0m tall.

Figure 4.22 shows the results from this investigation. As can be seen, the response on the detector from target-incident rays increases linearly with the target area in mm^2. This again can be predicted through use of eq. (4.8), as the observing aperture increases in size, so does the observed intensity. In this case, the target can be considered to observe the power scattered from the wall. Increasing this will, in turn, lead to a higher proportion of the illuminating laser pulse being incident back on the detector point. From this we can conclude that acquisition time should decrease linearly as target surface area increases.

Increasing the size of the target does also affect the background signal. As a higher proportion of the light is incident upon the target, less will illuminate the background scene. This effect was observed qualitatively from the model, however, the limiting nature of the computational approach prohibited quantifying this effect. The size of the background surfaces compared to the target, for a typical scene, means accurate investigation of this

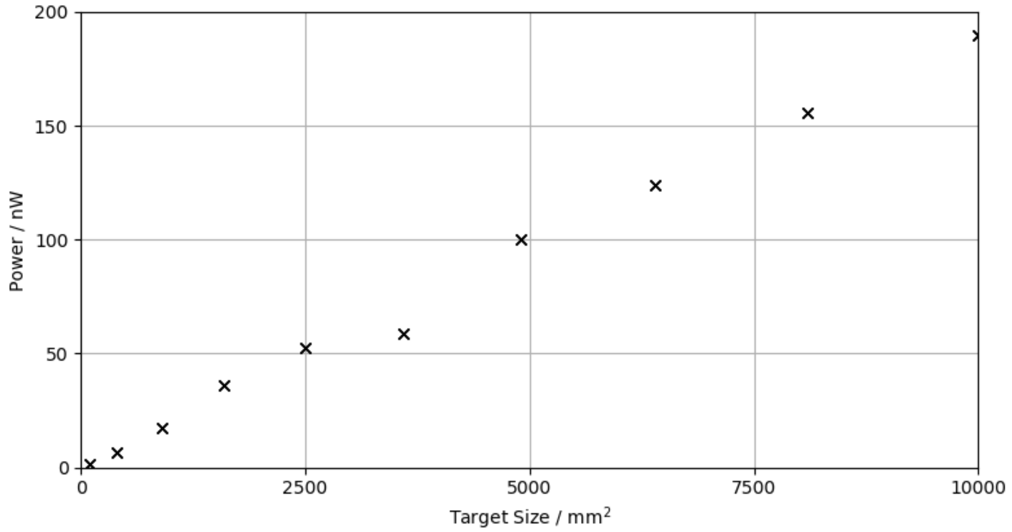

Figure 4.22: Total power on detector and power from rays incident upon target against size in mm^2.

relationship would require a much larger number of rays than the limits of the software. Lower intensity peaks from signal scattered by the walls behind the target were observed for larger targets, despite their relatively small size, compared to that of the walls. This is due to the projected size of the target, upon the wall, being much greater than that of target. This behaviour, of blocking signal from the background would be further affected by the orientation of the target within the room, in relation to the obscured walls and the target's distance from them.

Another consideration is the variation in precision when calculating the position of larger targets. The precision of this detection method is limited in part due to the finite size of both the target and the detector collection point. The mathematics presented in section 4.2.2 assume that the detector collection point and the target location, are singular points in space, which is not the case in reality. Increasing the size of either will increase the amount of light that successfully interrogates the scene at the cost of being less certain as to where their centre point is. As we are only concerned with finding a point in space that the target occupies, assumed to be the centre of mass, precision with which this point can be predicted decreases as target size increases. This is due to there being a greater number of locations from which target-incident rays can be reflected, increasing the potential number of possible ray paths, leading to a larger range of time-of-flights.

Figure 4.23 shows that minimum positioning accuracy does increase rapidly with target size for the smaller targets. Precision in this case has been measured as the distance travelled by light in one standard deviation of the time-of-flight measurements. As larger sizes are measured, the precision plateaus at approximately 330 mm. It is believed this is a side-effect of the limitations of the model. As the larger targets return many more rays to the detector, this increase in number of target-rays observed is balancing the increase in the position precision value due to size. The experimental results also report precisions in the range of 250 mm and the model plateaus at 330 mm. An increase in the number of detectors and a much larger sampling of rays, which in the model

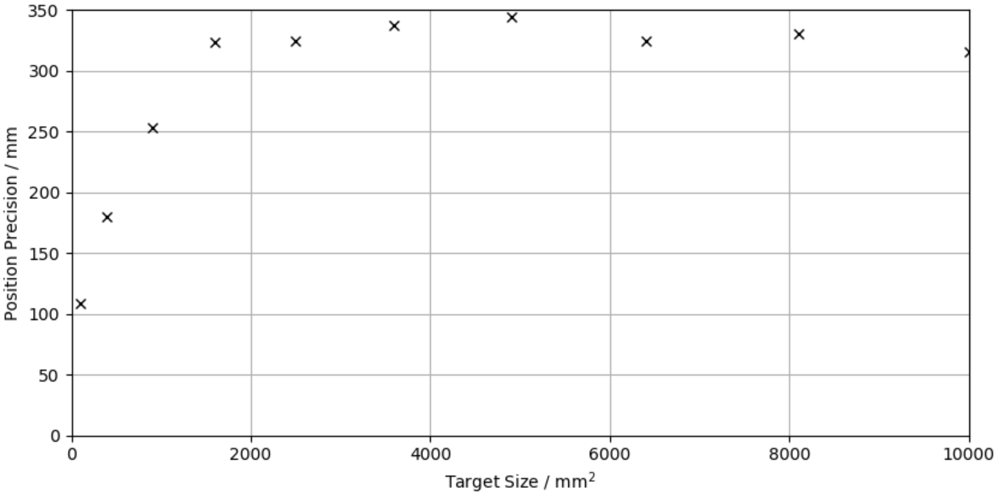

Figure 4.23: Variation of precision for time-of-flight calculation against target size.

are limited by computation time, would improve these minimum precision values. This may indicate that this model, due to its limitations, may not be best suited for precision investigations for this technique.

This model has shown that the size of the target greatly affects the response from the scene. A lower limit on detectable target size would be dictated by the background subtraction and noise in the detection method. Although larger targets will give better response and thus faster acquisition times, they do so at the cost of precision. Non-cooperative targets, of roughly human size and shape, are the most common use-cases intended for this technique and such should not vary too greatly, but with longer acquisition times and sufficient background subtraction, much smaller targets can also be tracked successfully.

4.8 Conclusion

A novel approach to modelling non-line-of-sight LIDAR ranging and positioning experiments has been presented. The model was successfully verified against experimental results. It has been shown that this model can be used to simulate results for this positioning methodology, providing new avenues for investigation into such systems. Through use of commercially available, optical modelling software, it is possible to quickly generate new scenes and investigate the effect varying components may have.

The model was used to investigate how a number of spatial parameters affect the application of this positioning approach. The inverse-square law forms the foundation for some of the more dominant factors which would be most likely to limit the accuracy and precision of this technique. Any change leading to a reduction in power on the detector, from rays that hit the target, will require a proportional increase in signal acquisition time to determine the target position with comparable precision. To reliably track moving targets, the target must move a distance no greater than its own size during the acquisition period. For human targets, this can limit the maximum acquisition to a few hundred ms [2].

The distance between detector and scattering surface was tested by placing a number of detectors interrogating the same scene at increasing distances. Power measured on each detector was then compared, a significant inverse-square relationship was observed as shown in fig. 4.13. For every doubling of distance between the detector and scattering surface, the power received would be quartered. This, therefore, limits the maximum distance at which the detector can be used and still recover a strong enough signal in a short enough acquisition time to track a moving target.

The angle between the detector and scattering surface normals was also tested. It was proposed that Lambert's Cosine Law [139] would apply here and that there would be no significant variance in power as the angle between these two normals were varied. The results in section 4.7.2 demonstrate this to be true. This means no special attention must be paid to optimise the angles between detector surface and that of the unknown target.

The effect that the target position itself had on the technique was also tested. Both the angle between the scattering surface normal and distance of the target centre from the initial scattering point. Similarly to section 4.7.2, the angle between the target and scattering surface normals has no effect on the power incident on the detector. As eq. (4.8) states, the power observed is constant regardless of angle. The target distance from the scattering surface, however, had the greatest effect on the observed power of all the parameters tested. Similarly to the model in section 4.7.1, since light from the initial scattering event is reflected to the target, an inverse square relationship is observed between power and distance of the target from that scattering point. For light to be observed at the detector, it must be reflected back to the scattering surface at the detector collection point. Figure 4.19 shows that an inverse fourth root relationship between the target distance and the scattering point gives a good approximation of the power expectation as the target distance varies. This would meant that, should the distance between a target and the scattering surface double, only a sixteenth of the power from target-incident rays would be observed, requiring a proportional increase in acquisition time.

Finally the effect that target size had on observed power was investigated. Figure 4.22 shows the proportional relationship between an increase in target size and the observed power on the detector. This would mean that, as target size grows, acquisition times will decrease. It would also indicate that the lower limit on target size, for tracking and positioning by this technique, would be limited by the noise and background subtraction. Variation in precision was also investigated using this model. Although it was anticipated that precision would decrease as the target size increased, this effect was not observed in the model for the larger target sizes. Figure 4.23 shows that the precision value plateaus at approximately 330 mm. It is believed that this is caused by the limitations of the computational approach. As so few rays are returned from the target for each scene, the increase due to size is leading to an improvement in precision as the sample increases which is partially balancing the expected decrease in precision. Improved ray-sampling of each target and a deeper look into the effect of size on precision would make for a worthwhile extension to this project.

For each of these factors, the acquisition time has been discussed as the parameter mostly affected. The laser used in section 4.4.1 ran at 80 MHz, which means there is only a 12.5 ns gap between pulses. If the size of a scene increases and the time for a photon to complete a round-trip of the scene breaches this 12.5 ns upper limit, overlap between pulses begins to occur. It would be difficult to distinguish which pulse a photon originated from and there would be significant risk of losing information. Therefore, the repetition of the laser source would have to decrease. This would have a knock-on effect of further increasing acquisition time, which also has an upper limit in the hundreds of microseconds before targets moving at a average human speed will begin to blur the positioning technique. One simple approach for reducing acquisition times, that has been demonstrated by Gariepy et al. [2, 3, 4], is to use an array of detectors. It would be possible to use this model the investigate how increasing the number of detectors affects the observed response power.

The limiting factor in using this modelling approach is the time it takes for the ray-tracing to complete. Using post-processing to vary parameters of interest such as field of view, or pixel pitch, in the case of a camera can mitigate some of this by reuse of a single ray database. Utilising post-processing and automation of the ray-tracing, it would be possible to develop optimisation routines to select components for improved experimental results. This would be a promising continuation of the work presented here.

CHAPTER 5

Amplified Spontaneous Emission in High-Gain Solid-State Lasers

5.1 Introduction

In the design of Q-switched lasers, it is often desirable to maximise the achievable energy storage in the gain medium prior to switching. Amplified Spontaneous Emission (ASE) can critically limit this achievable energy storage when gain reaches a threshold level. This limiting effect can be made worse by internal reflection in the gain medium, making the effective path length through the high gain region longer, increasing the likelihood of spontaneously emitted photons being amplified [56, 141]. Total internal reflection (TIR) can further cause ASE to strongly influence the performance of these systems [56]. Being able to identify how ASE develops within a specific gain medium and which paths through the system lead to the highest gain would enable designers to better understand this performance limiting phenomena. Gain media geometries and pumping schematics would be adjusted to limit those paths that most contribute to ASE. New and novel designs, that limit ASE, could be tested quickly and easily, without the need for costly manufacturing of bespoke components.

Many comprehensive theoretical and analytical models of ASE exist[80]. These analytical models, however, are difficult to implement for all but the most conventional of gain media geometries and pumping schemes. Modern high-gain solid-state lasers, like those used in defence applications, often have very complex designs that are optimised for a number of factors, including the overall size and weight of the system [99]. To fully understand the development of ASE within such laser systems a three-dimensional, time-dependent model must be developed.

A quasi-analytical technique was first used by Lowenthal et al. [141] to calculate the ASE effect in laser gain media with small aspect ratios. Due to the limitations of the existing technology, many simplifications were made to the model to make numerical computation feasible. A cylindrical coordinate system was used, allowing for some of the integration of the ASE over the volume to be computed analytically. Gain and flux were also only allowed to vary along the amplifier axis, being uniform across the cross-section. The authors state that performing these calculations in three-dimensions would be computationally expensive and an iterative time-step approach would be required to evaluate the effects of ASE[141].

Fully three-dimensional time dependent models had been developed previously [142, 56] but, again, were too costly to run with the technology available at the time. Lu et al [56] built upon the quasi-analytical approach used by Lowenthal et al. [141] to address this issue. This was attempted at the Lawrence Livermore National Laboratory to better understand the effects of ASE on face-pumped slab laser geometries [56]. Non-saturable absorption, ASE line-narrowing, and total internal reflection (TIR) were accounted for in this model, allowing for rapid determination of the gain-length product of rays travelling through high-powered slab amplifiers.

This model relied on two equations, one for the normalised ASE intensity and another that relates the ASE to a reduction in gain for any point within the laser medium. These were required to be solved self-consistently in order to determine the affects that ASE had on the laser amplifier [56]. The three-dimensional nature of both of these equations was a major limitation of the model at the time. Both would require significant computation to calculate the distribution of gain and ASE across the laser amplifier, meaning the model was essentially limited by the technology available.

To reduce computation time certain approximations were made. Firstly, the gain and intensity distributions were only allowed to vary along the z-axis of the amplifier. This meant that the ASE intensity was only computed at the centre of the gain medium and assumed to be uniform across the cross-section for each of the z-positions. The model also assumed that no reflection occurred at the end surfaces of the slab amplifiers, limiting the path of the spontaneous emission through the gain medium. The pumping model was also limited to only include those rays that would exhibit total internal reflection (TIR).

Despite these assumptions, when compared to experimental results, good qualitative agreement was found. This model, however, is optimised for specific use cases, slab amplifiers. The assumptions listed above may significantly impact the applicability of this model to more complex gain media and pumping schemes.

More recently Albach et. al [54] developed a model to investigate the effect ASE had on YB $^{3+}$ slab amplifiers. This model used a Monte-Carlo approach and was able to estimate stored energy density for a given point in three dimensions within a gain medium. The calculations relied on estimating the local depopulation due to ASE as a function of position, much in the same way as Lowenthal et al. [54].

This model also uses a number of approximations to simplify the calculation. Firstly, the pump power is taken to be uniform across the gain medium, much in the same was as Lownethal et. al [141], and all surfaces are considered to be perfectly absorbing for spontaneous emission. The model also ignores multiple internal reflections as the surfaces of the gain medium. The authors conclude that, although this assumption may work for gain media with aspect ratios of approximately one, such as cubes, thick slabs, and spheres [54], any geometries that exhibit significant internal reflection and TIR this could be problematic [54]. Gain media that are thin compared to the transverse show much stronger depopulation due to TIR when compared to thick crystals [54].

Developing a fully iterative, three-dimensional, numerical model for the effects of ASE on laser gain media would afford a greater understanding of the factors in cavity design that greatly impact it. In Q-switched lasers and laser amplifiers maximal energy storage is often a key design objective. It would be possible to test multiple configurations of gain media, pumping schemes, and optical designs to minimise the gain depleting effects of ASE. A fully iterative model would also allow the user to better understand at what point in the pumping period ASE begins to have a significant impact and enable them to adjust accordingly.

Although the existing models give good agreement with experimental results, they rely on artificially simple pumping schemes or assumptions around gain distribution. Many laser systems have highly complex pumping arrangements which can severely impact the gain distribution [99]. Diode stacks being one such example that give distinct striping of available gain within the gain medium. Utilising ray-tracing to first model these complex distributions before calculating the gain depleting effects of ASE may give a closer insight into the behaviour of ASE in these systems. This would allow the user to more easily explore and discover more optimal pumping schemes, cavity designs, and novel gain media to improve performance.

5.1.1 Numerical Gain Model

The model presented in this chapter will build upon the work and suggestions of previous quasi-analytical methods [56, 141, 142, 54]. One key difference in this approach, however, is the use of ray-tracing software to calculate three-dimensional, discrete models of the pump absorption within the gain media. Similar approaches have had success in several cases for designing lasers with complex pumping requirements [143, 144, 145, 146, 147, 148, 149, 150]. This approach uses the ray-tracing software Zemax [6], a commercially available optical design package. As an industry leading product, Zemax offers many benefits to this type of analysis, including large surface and source libraries, allowing users to quickly model systems with known components and materials, without having to define their properties each time. It is also possible to save and extract ray databases, holding information on each interaction a ray has as it passes through the system. This allows the exact ray trace results to be reused on multiple specifications of the system to compare performance. These benefits, along with the tool's high accuracy and ability optimise optical designs, make it commonly used in the design of systems for defence applications [99].

Zemax allows the user to measure pump absorption from any material or component in the system by defining it as a detector. This functionality is at the core of this model. By dividing the gain material into discrete three-dimensional pixels, known as voxels, it becomes possible to see the three-dimensional pump absoption throughout the gain medium.

Zemax does not, however, calculate gain. In this model, the pump absorption profile and the ray tracing database will be extracted and further analysed in MATLAB [5]. The evolution of gain, within the gain medium, will be calculated in discrete time-steps for each of the voxels. The rays passing through these voxels at each time step will then have their energy recalulated according to the gain they experience and the energy available in each voxel will be updated accordingly. This will result in an iterative, discrete, three-dimensional model of the gain distribution.

Despite advances in computational power, it is still necessary to make a number of approximations to make the model computationally feasible. It will, therefore, be vital to verify the predictions from the model. To do this, an experiment concerning ASE from the literature [59], has been replicated. The same set-up will be used, in this model, as was present in the physical experiment. Achieving comparable results should indicate the effectiveness of this model to the modelling of ASE in solid-state lasers. This will then give confidence in predictions of the model in other systems, enabling ASE reducing design decisions to be tested more quickly and easily.

5.1.2 Chapter Outline

Section 5.2 will discuss the model in detail. The algorithms for calculating the energy distribution at each time-step and for each voxel are presented. This section will also discuss how the ray-trace and pump absorption from Zemax are used to calculate spontaneous emission. Finally, some of the limitations of this model are presented.

In section 5.3 the approach to verifying the model is detailed. The experimental setup and results from the 2010 paper by Huß et al. were chosen to test the model against. Section 5.4 then discusses how the results from the model compare to those from the paper.

Section 5.5 will see the model used to predict the performance of other systems. The energy storage properties of a novel laser gain geometry is tested against a typical square rod and the results presented here.

Finally, section 5.6 summarises the chapter. Recommendations are made for improvements to the approach taken here and next steps in furhter developing this model.

5.2 Model

This section will discuss the process behind the numerical modelling of Amplified Spontaneous Emission (ASE) in solid-state lasers. An overview of the entire process will be presented here, then each of the component pieces of the model will be discussed in further detail.

Figures 5.1 to 5.2 presents the algorithm this model uses to calculate power output, from a laser, considering the effects of ASE. The process begins with determining the absorption of the pump, discussed in section 5.2.1. It is then necessary to calculate the gain across the medium and seed spontaneous emission rays, which is discussed in sections 5.2.5 to 5.2.3. These rays lead to energy lost from the system and gain depletion, the processes for which are considered in section 5.2.5. Section 5.2.6 then presents how the final power seen at the detectors is determined.

Figures 5.1 to 5.2 illustrate the procedure that this model follows. Each of the steps presented will be discussed in detail, including the mathematics underlying the physical phenomena that are being model. Although the

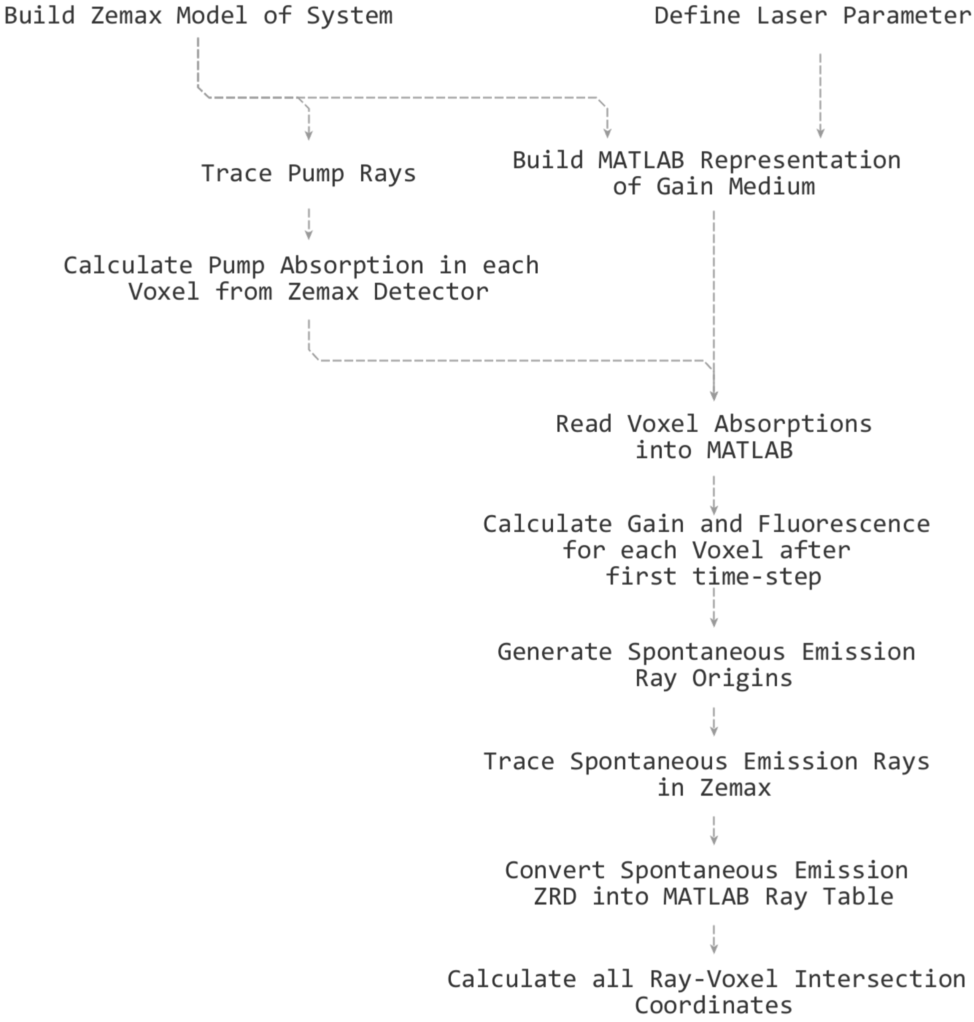

Figure 5.1: This flowchart summarises the initial steps in building the optical model in Zemax and MATLAB. Once pump absorption has been measured, the gain and flourescence can be calculated in MATLAB. This is used to define then build the spontaneous emission Zemax Ray Database (ZRD) which is converted back to MATLAB for the second, iterative phase of the model.

model is cyclic in some places, as many processes depend on dynamic values that are updated continually, this section will try to follow the figures from top to bottom, where possible.

Constraints relating to both available time and the current hardware and software limitations, lead to a trade-off that dictates the performance of this approach. Lowenthal and Eggleston first presented a model for small aspect ratio laser oscillators and amplifiers but, having far less computational power available to them, deemed it infeasible at the time [141]. Lu and Dong [56] later used this approach to model slab lasers but made simplifications, such as assuming transverse pump homogeneity and only considering rays that undergo Total Internal Reflection (TIR) [56]. Ideally, one would model a single ray for every photon emitted, in each time-step.

For photons of wavelength λ, for E Joules of emitted energy the number of rays, N_r required is calculated using eq. (5.1)

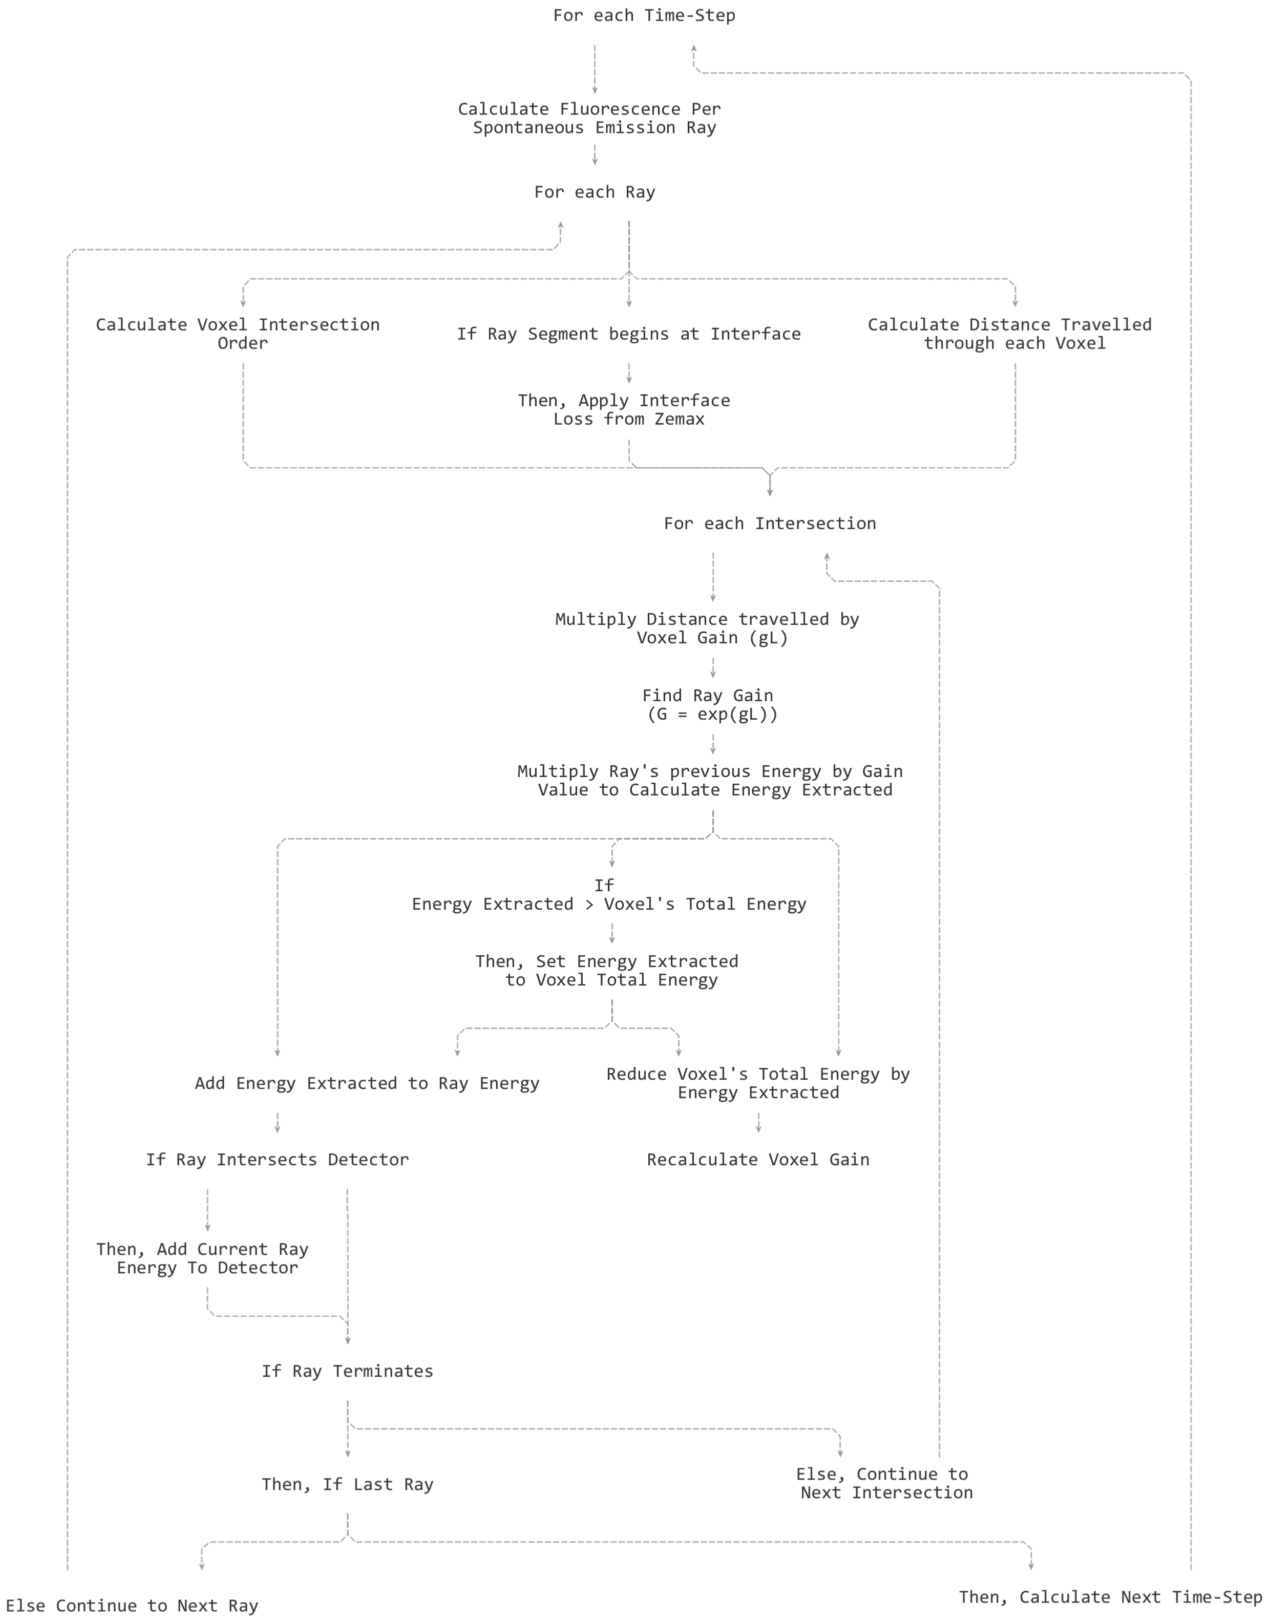

Figure 5.2: This flowchart summarises the calculations in MATLAB for each of the rays and voxels, at each time-step. Energy extraction and gain are calculated in an iterative manner for every ray and the voxels available energy is updated accordingly.

$$N_r = \frac{E\lambda}{hc} \tag{5.1}$$

where h is Planck's constant and c is the speed of light. Therefore for 1 μJ of energy emitted, at a wavelength of 1064 nm, 5.356×10^{14} rays would need to be traced to account for each photon. Zemax, however, is only capable of tracing 4×10^9 rays in a single trace. As all of the rays are calculated independently, it would be possible to run many ray-traces to get the numbers required, but for even relatively simple optical systems a trace with 4×10^9 rays can take an hour (for a typical desktop workstation). It would be possible to reduce the length of the time-step so that an even smaller number of spontaneously emitted photons are released, but this would take longer still, as the calculations in the model that aren't related to the ray-tracing would be repeated with the increased number of time steps. Also, as the system is pumped throughout the simulation, the time-step length would have to be adjusted to accommodate the dynamic changes in emitted power, meaning that the model's accuracy suffers. Using the approach taken here, as the number of rays traced increases and the time-step length decreases, the model more closely approximates the real system.

The approach taken here then, is to calculate the total power emitted at each time-step and divide that evenly amongst the total number of rays to be traced. This directly contributes to the distribution of where the spontaneously emitted rays originate, which is discussed in detail in section 5.2.3.

5.2.1 Absorption

Determining an accurate pump and absorption profile is the basis of this model and what differentiates it from analytical solutions. By utilising Zemax to design the optical system and trace the pump ray-paths, complex geometries can be easily simulated. Figure 5.3 compares the pump profile of an Nd:YAG rod pumped by a diode stack as seen by an Infra-Red (IR) camera, on the left, with the simulation in Zemax for the same pumping geometry on the right. As can be seen the model closely resembles the real system allowing for accurate modelling of complex pumping schemes, something that is not easily achieved with entirely analytical approaches [99].

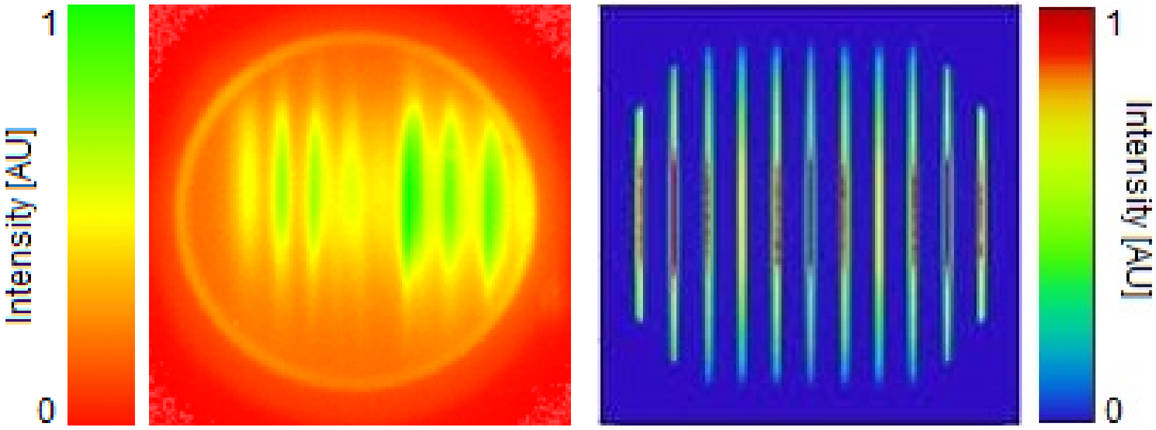

Figure 5.3: Nd:YAG rod pumped by a diode stack as seen by an IR camera (left), compared to the modelled pump profile in Zemax for the same pumping geometry (right).

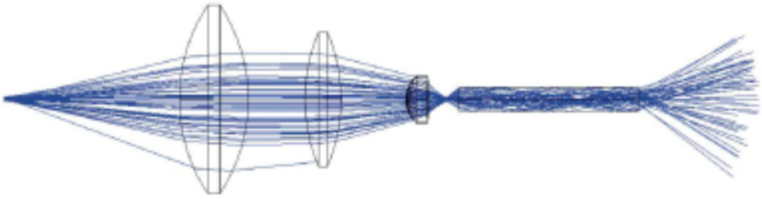

Figure 5.4: An example ray-trace from Zemax. Paths and material absorption are calculated for each ray as it passes through the system [59].

The pump profile is determined according to the rules underlying the Zemax ray-tracing software. Figure 5.4 illustrates a Zemax ray-trace, calculating paths and absorption as rays pass through the optical system [59]. Once the system has been designed and the correct materials and coatings placed on all of the components, a number of rays are traced randomly from the source. As the ray-tracing uses Monte Carlo [22] methods a greater number of rays gives a pump profile more closely resembling the real system. In the case of the model for ASE presented here 4×10^9 rays were used, the maximum number of rays permitted by Zemax. These rays pass through the system and Zemax recalculates their power, phase and direction at each interface. Should the rays pass through an interacting medium, such as the laser rod, it is possible to record how much energy is absorbed.

Absorption is recorded in Zemax through the use of detector objects and any object in Zemax can be made into a virtual detector. This allows the object to record information about the state of the system at each time-step, without it interfering with the simulation in the same way a real detector might. By assigning a number of pixels in each of the cardinal axes for a detector object, it is possible to control the resolution with which the object records this information. A detector then has $n_x \times n_y \times n_z \times = V$ number of cuboid voxels, where n_x, n_y, n_z are the number of pixels along each axis and a voxel being the three dimensional extension of a pixel. This ASE model will operate by considering each of the voxels as a separate unit of gain medium, in an array. The absorption for each of the voxels can then be taken to be the input power for the gain equations, discussed in section 5.2.5.

Once the pumping rays have been traced through the system, total absorption for each voxel can be calculated. For a single voxel, the energy absorbed from each ray is determined using eq. (5.2) and the path length that ray has taken through the voxel. One of the key benefits of using Zemax for the absorption calculation is that it supports a vast library of materials and is capable of handling custom absorption spectra, allowing most materials to be modelled easily. The intensity of a ray leaving a voxel, I_a, can be calculated using Beer-Lambert's Law [151], given by eq. (5.2)

$$I_a = I_0 e^{-\alpha l} \tag{5.2}$$

where I_0 is the intensity of the light entering the voxel, α is the absorbance of the material and l is the path length of the light through the voxel. Absorption for each ray can then be calculated as $I_0 - I_a$, assuming that the medium is far from saturated.

This must be done in the order with which the ray intersects with each voxel along its path, as the absorption in one voxel reduces the total energy available to subsequent rays that may intersect it. This is repeated for every ray to determine the total energy absorbed from the pump by each voxel. With use of three dimensional ray-tracing and proper treatment of surface coatings and material properties, it is possible to accurately simulate pumping profiles for complex cavities and pump sources.

One key aspect of this approach is the discretisation of the time component of the model, as discussed in section 5.2. This allows for the absorbed pump power to be converted into an energy value at each time-step This discrete method allows the energy for each voxel to be depleted, by spontaneous emission, before calculating the next time-step. Section 5.2.2 discusses the process of determining the gain seen by rays passing through each voxel. Once the absorbed pump energy has been calculated, for each voxel, the same value is used throughout the pumping phase. Although this has not been utilised here this could easily be altered, as the pumping power can be adjusted manually or by applying a time varying function to the pump values in MATLAB. This would allow more complex time-variant pumping to be considered with this model, if desired.

5.2.2 Pumping, Gain, and Spontaneous Emission

Once the absorbed pump power for each voxel has been retrieved from Zemax, population inversion and available gain can be calculated. With an accurate model of the population inversion, it is then possible to determine the distribution of spontaneous emission throughout the gain medium, which is discussed in section 5.2.3.

The fluorescence energy spontaneously emitted in the initial time-step must be calculated to allow for the individual ray energies to be set. This is also required to determine how much energy is available for gain at the end of the time step. This model breaks the calculations for ASE development over time into discrete steps. A time-step is defined and the gain and fluorescence is recalculated at each time-step. To get each ray's initial energy, we must calculate the pump energy absorbed in the initial time-step and the proportion of that which is lost to spontaneous emission.

Firstly, the energy within each voxel E_v must be calculated. For the first time-step, this is simply a case of multiplying the length of the time-step δt_1, by the absorbed power for each voxel A_v which is determined by the pumping model in Zemax, $E_v = \delta t_1 \times A_v$. The rate of electron excitation to the metastable level, for each voxel R_{pv}, can then be found by eq. (5.3)

$$R_{pv} = \frac{E_v Q}{V_v h \nu_p} \tag{5.3}$$

where Q is the quantum efficiency, ν_p is the frequency of the pumping light, h is Planck's Constant and V_v is the volume of the voxel [101]. Some of the electrons that are excited within the time-step will decay during that time, this decay must be accounted for when determining how much gain is available, which could then amplify spontaneously emitted photons, for each time-step. Assuming a square pump pulse, the maximum number of electrons raised to the metastable level n_v from the ground state n_0 can be given by eq. (5.4)

$$n_v(t_p) = n_0 R_{pv} \tau_f [1 - e^{(-t_p/\tau_f)}] \tag{5.4}$$

where τ_f is the fluorescence lifetime of the metastable level for the gain medium [101]. Since the total number of electrons raised to the upper level during the pump pulse is given by $n_0 R_{pv} t_p$ the fraction in the metastable level at the end of the pumping pulse, the storage efficiency η_{st}, is given by eq. (5.5)

$$\eta_{st} = \frac{1 - e^{-t_p/\tau_f}}{t_p/\tau_f} \tag{5.5}$$

From eqs. (5.3) to (5.5), it is then possible to calculate the gain for each voxel g_v at the end of each time-step using eq. (5.6)

$$g_v = R_{pv} t_p \eta_{st} \sigma_e \tag{5.6}$$

where σ_e is the spontaneous emission cross-section for the gain medium [101]. Equation (5.6) is used to calculate the gain for each of the individual voxels in the gain medium at the end of the first time-step.

For subsequent time-steps, the same process applies with regards to the absorption of additional pump energy; however other phenomena affect the available gain, such as gain depletion by spontaneously emitted rays. Gain depletion is included in this model and discussed in section 5.2.5. If the pump source is defined as constant and continuous, then the same amount of energy will be incident upon each voxel at every time-step. To determine the energy within each voxel after each time step, the amount of energy already present within the voxel must be known. Gain depletion is of particular importance when considering ASE, as large amounts of energy could be extracted by spontaneously emitted rays at each time-step, this is discussed in detail in section 5.2.5.

Using discrete time-steps makes working with gain values, expressed as a per centimetre value difficult, especially when considering changes during each time step. For this reason, the energy stored in each voxel made the calculations easier to compute. For each voxel, the population pumped into, and remaining in, the metastable level at the end of the time-step n_v is given by eq. (5.7)

$$n_v = R_{pv} t_p \eta_{st} \tag{5.7}$$

and the energy stored in a given voxel E_v, after the first time-step, can be determined by eq. (5.8)

$$E_v = n_v h \nu_l V_v \tag{5.8}$$

where ν_l is the frequency of the photons relaxing from the metastable level. From these equations we can then also calculate the amount of energy that is fluoresced in the first time-step E_f using eq. (5.9)

$$E_f = (1 - \eta_{st})\frac{E_v}{\eta_{st}} \tag{5.9}$$

Equation (5.9) gives the energy fluoresced from a single voxel, which is used to determine the origins of the rays to be traced. As discussed in section 5.2.5, each of the rays in this model share the same initial energy, equal to the total fluoresced in the time-step divided by the number of rays being traced.

Due to computational and algorithmic limitations, assigning individual energies to every ray at origin causes the modelling time to increase beyond reasonable limits. The compromise here, by giving rays equal energy, causes those parts of the gain medium that are furthest from the pumped region to see a greater proportion of the spontaneously emitted energy then they would in reality. Section 5.2.3) discusses the methods used to account for this somewhat through ray placement, by strongly weighting the origin for every ray towards that more intensely pumped region.

5.2.3 Seeding Rays

Once the population inversion has been calculated for a time-step it is then possible to determine the distribution of spontaneous emission rays throughout the gain medium. Zemax is only capable of tracing 4×10^9 rays in a single trace, a small number when compared to the number of spontaneously emitted photons. Each ray must then be given a proportion of the total spontaneously emitted energy.

Although the rays are traced randomly, the distribution of where they originate within the gain medium is not homogeneous for most pumping schemes. Considering the voxelised gain medium, voxels containing a greater absorbed energy are far more likely to spontaneously emit a photon than those with less energy. Figure 5.5 presents the pumping profile for an Nd:YAG rod, the darker blue end of the rod shows that very little of the pump is absorbed. This will mean fewer photons will be excited to the metastable level and this section of the rod will exhibit far less spontaneous emission than the brighter red and green section, where absorption is much higher.

Section 5.2.3.1 describes the method used to weight each voxel when determining the random distribution of emitted rays.

5.2.3.1 Roulette Wheel Selection

Roulette Wheel Selection is a method used in computer science, commonly in genetic algorithms [152], when weighted random selection of competing candidate solutions is required. In the case of this model, it has been used to randomly select voxels to emit rays, weighted by the proportion of the total absorbed pump energy in

Figure 5.5: Pump absorption profile for an Nd:YAG rod, as calculated by Zemax. False colour is used here to represent variation in absorption where dark blue corresponds to low and red to high absorption.

the system.

For any voxel i, the probability of it being selected to emit a ray ρ_i is equal to the absorbed energy in the voxel a_i divided by the sum of all absorbed energy for every N voxels, as shown in eq. (5.10).

$$\rho_i = \frac{a_i}{\sum_{j=1}^{N} a_j} \tag{5.10}$$

This is accomplished using the algorithm shown in fig. 5.6. Firstly the pump energy absorbed in each voxel is calculated using the ray-trace of the source in Zemax, as discussed in section 5.2.1. The absorbed energy in each voxel i is then scaled by a_i/a_{max} giving a ratio of absorbed energy relative to a_{max}, the maximum absorption by a single voxel in the gain medium.

Then to determine which voxel a ray j will be originating from, a random value r_j between 0 and 1 is generated. A voxel i is then chosen randomly, if that voxel has a relative absorbed energy $a_i/a_{max} > r_j$, the ray's origin will be within that voxel. Should the converse be true $a_i/a_{max} < r_j$, another voxel is chosen randomly until one is selected that meets the criterion. This process is then repeated until each of the required rays has been assigned a voxel. The origin coordinates are then generated randomly, within the boundaries of the voxel. Finally direction cosines are also generated randomly, with respect to each of the cardinal axes. These six data points, coordinates and directions in all three axes, are then saved as a source file readable by Zemax, for tracing the spontaneous ray paths.

There are many other methods for selecting origin points for each of the rays in the system. Completely random sampling was tested in the first instance and gave good results. This method, however, leads to homogeneous energy density throughout the gain medium. For end pumped systems, the majority of the pump energy will be absorbed at one end of the rod. Many rays will intersect the boundary of the gain medium at less than the critical angle, immediately transitioning to the surrounding medium. In end pumped systems the concentration of light escaping at first incidence, is much higher, closer to the pumped end of the rod and relatively much lower at the opposite end. Seeding ray origins truly randomly leads to an even spread of this escaped energy. In the lower pump power regime where ASE is not dominant, this can lead to great spatial inaccuracies in the model, as the

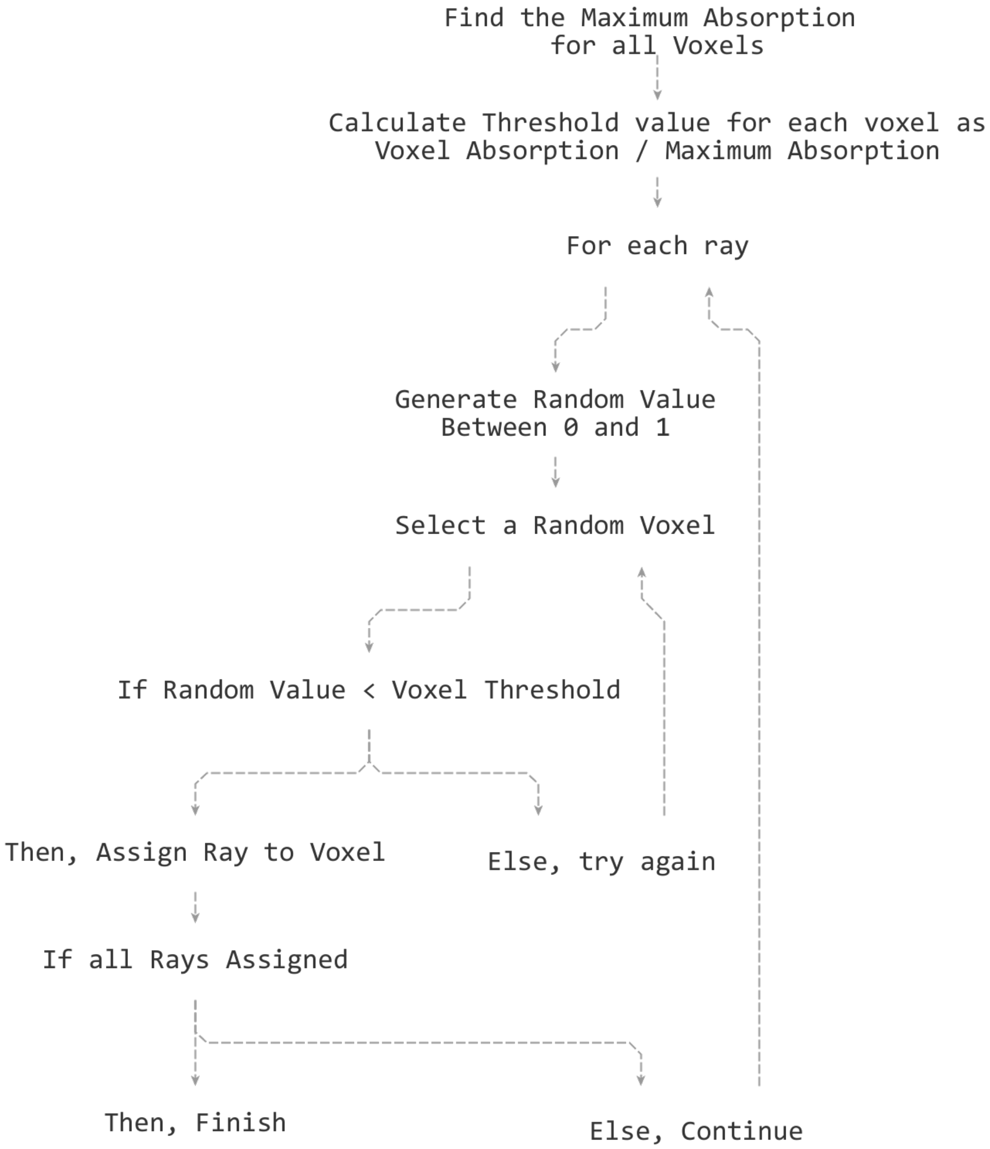

Figure 5.6: A flow-chart of the Roulette Wheel Selection (RWS) algorithm, used to place the ray start locations throughout the gain medium [152].

distribution of the escaped energy in the model does not reflect what is observed in reality.

It would also be possible to trace an equal number of rays originating in each of the voxels. The energy available to those rays would then need to be scaled, dependent on their origin voxel to allow for an accurate energy distribution across the model. This can lead to tracing a large number of rays that have low relative energies and never see any amplification, having very little effect on the model. The ray-tracing in Zemax (in the case of this model) would still incur a large computational cost for little improvement in accuracy. Furthermore, scaling individual rays would add another complicated step to the process.

Roulette Wheel Selection (RWS) enables sampling all of the voxels within the system whilst favouring those that have the largest absorbed energy, and rightfully eliminating those voxels that see no energy at all from the pump. Figure 5.7 illustrates the difference between the two sampling schemes. These two images are of exactly the same

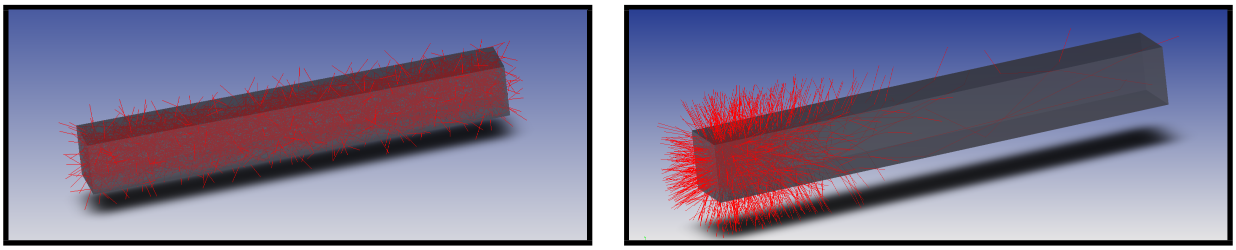

Figure 5.7: Zemax rendering 1000 rays of the fluorescence modelling for two ray placement schemes. The left image shows rays placed entirely at random, the right image shows same number of rays placed using the Roulette Wheel Selection algorithm [152].

rod, both being pumped from the left face. Within each of the rods are 1000 rays as traced by Zemax. The left hand image shows rays that have been generated completely randomly whereas the rays in the right image have been generated using RWS.

As can be seen, the RWS rays strongly favour the pumped end of the rod but there are still some rays generated in the less pumped region. The truly random rays are much more evenly distributed throughout the rod, having equal chance of being seeded in voxels that have seen no pump absorption or those that are strongly pumped.

5.2.4 Ray-tracing and Losses

Once the ray's origin points have been generated, as discussed in section 5.2.3.1, their paths must be traced through the system. Each of the rays is given the same power at origin and they are traced independently of one another. The loss each ray experiences as it travels around the system is calculated by Zemax at each of the interfaces.

Interactions at interfaces can be handled a number of different ways in Zemax. In this model the method used is called Simple Ray-Splitting [6]. For every interface event a portion of the ray may be transmitted or reflected, dependent on polarisation and the interface properties. Simple Ray-Splitting selects which path to follow, that of the transmitted or reflected light, at random weighted by the proportion of light reflected and transmitted. The energy that would have remained in the path that is not followed is simply subtracted from the ray's total energy. This leads to a single path for each of the traced rays that diminishes in energy at each of the interfaces.

Rays can terminate in one of three ways following this method; they exit the system and are not reflected back onto any of the modelled surfaces, they reach a maximum number of intersections (computational limit on the number of interface interactions a single ray can have) or they reach a minimum power threshold. In this model the minimum power threshold was removed. For inactive optical media, there is no phenomena by which a ray can gain power as it travels through the Zemax model and it therefore makes sense, from a computational perspective, to stop calculating the paths of rays that have become very weak. When gain is introduced to a system even very weak rays can encounter great amplification in their power and have significant influence on the system. This is of key interest when studying ASE. Unfortunately, the computational limit to the number of intersections a single ray can experience is limited to 4000. This may be inadequate for a ray that experiences significant gain throughout it's path so as to counteract the interface losses and it is not uncommon for rays to reach this limit. Investigating these rays in particular, by filtering the ray database for rays that have been terminated artificially due to this condition, highlights potential high-gain paths within the system.

It is possible to split every ray at each interface in Zemax, but this causes the size of the ray database output file to become very large and greatly increases the number of paths through the system. Following this splitting method would yield a model that more closely resembles the physical phenomena. The compromises required to run this model could further be eliminated with improvements in algorithm efficiency and computational power, which would be promising avenues for future development efforts.. Care would also need to be taken in calculating the order in which the ray branches are resolved, as they could potentially deplete the gain another branch could experience as it passes through the gain medium.

Once these rays are traced, the model can then calculate how much, if any, gain they experience. This is discussed in section 5.2.5.

The same ray paths are used for all time-steps in the model, which is another limitation due to computational resources.

The choice was made to use the same ray-trace throughout the mode. This is a slight compromise in the model accuracy for systems wherein the spontaneous emission distribution changes significantly, and can also impair the voxel sampling statistics depending on the initial rays. This is discussed in more detail in section 5.2.7.

5.2.5 Gain and Depletion

Once the ray paths have been traced in Zemax the next step is to calculate their interactions with the gain medium. In this model, each ray is treated independently and sequentially, meaning that gain is calculated for each ray, from its origin to termination, before the next ray is considered. This allows for the gain to be depleted, limiting the energy that can be extracted from each voxel. For each time-step the rays are selected in a random order so that one ray is not guaranteed to always deplete the gain of a voxel that another passes through just because it precedes it in the ray table.

Each ray is given equal starting energy, as calculated by eq. (5.9). The Zemax ray database is read into MATLAB and the voxel intersections are calculated. The voxels that the ray passes through and its segment length through them is then calculated. For the first voxel this will be the distance from the origin to a boundary edge (the voxels are cuboid, regularly sized and spaced). Then with the second voxel the ray passes through, for example, this length will be the distance between this first boundary intersection and an exit point, intersecting one of the voxel's five other faces. This is repeated for all subsequent voxels on the ray's path, until the ray leaves the gain medium, or meets its termination point calculated by Zemax.

These voxel segment lengths must then be multiplied by the gain values for each of the corresponding voxels to calculate how much gain each ray experiences. The gain coefficient for a voxel i, of volume V_v, is given by eq. (5.11)

$$g_i = \frac{E_i \sigma_e}{h\nu_l V_v} \tag{5.11}$$

where E_i is the energy in the voxel and σ_e is the spontaneous emission cross-section for the gain-medium [100]. The gain experienced G_{ij}, for a ray j, passing through a voxel i can then be calculated by eq. (5.12)

$$G_{ij} = e^{g_i l_{ij}} \tag{5.12}$$

where l_{ij} is the length that the j^{th} ray passes through the i^{th} voxel. For the first ray segment, the ray starting energy E_{fj} is then multiplied by G_{ij} to get the new ray energy E_{nj} upon leaving the voxel.

This treatment, however, allows for a ray to extract more energy than a voxel has. To ensure this does not occur, a maximum condition must be enforced on how much energy a ray can extract E_x, that is, all of the energy from the voxel. Therefore the new ray energy E_{nj} must be calculated using eq. (5.13)

$$E_{nj} = \min\{E_{fj}G_{ij}, E_{fj} + E_i\} \tag{5.13}$$

where E_i is the total energy in the i^{th} voxel. The new gain coefficient value for the voxel g_{in} can then be calculated using eq. (5.14)

$$g_{in} = \frac{(E_i - E_x)\sigma_e}{h\nu_l V_v} \tag{5.14}$$

where E_x is the energy extracted from the voxel. This process is then repeated for each of the voxels along the rays path until the ray hits a object boundary in Zemax.

Once all of a ray's segments have been traced, the results are stored and the next ray can be calculated. This process of sequential gain depletion ensures that the finite energy put into the gain medium, via pumping, is conserved as no thermal losses are considered here. Each ray following the first will see a unique gain array as it passes through the gain medium to its termination point. After every ray has been treated in this way, it is then necessary to move on to the next time step.

For time-steps after the first this process is repeated in a similar manner. The only significant difference being that the gain medium already has an amount of energy stored within each of the voxels. To determine the population at the end of the next time step, the electrons excited from the ground state and the remaining population must be considered. Equation (5.15) can be used to determine the remaining population in the metastable level n_{2i}, after gain depletion has occurred.

$$n_{2i} = \frac{(E_i - E_x)e^{(-\Delta t/\tau_f)}}{h\nu_l V_v} \tag{5.15}$$

where $e^{(-\Delta t/\tau_f)}$ accounts for the decay of those electrons to lower levels through spontaneous emission and Δt is the time difference from the previous time-step. The storage efficiency η_{st} of the electrons pumped from the lower level can be calculated using eq. (5.5), and the number of electrons $R_{pv}\Delta t$, given by eq. (5.7), wherein the pump time t_p becomes Δt. The stored energy value for the i^{th} voxel at the start of the next time-step E_{in} is then given by eq. (5.16).

$$E_{in} = h\nu_l V_v (n_{2i} + R_{pv}\eta_{st}\Delta t) \tag{5.16}$$

With the new energy values it is then possible to calculate the new gain in each of the voxels using the stimulated emission cross-section σ_e. The final value required is the fluorescence energy emitted in the next time-step E_{fn}. This must account for the spontaneous emission of the population already in the metastable level at the end of this time-step n_{2i} and the spontaneous emission lost during the next pumping window, as calculated using eq. (5.5). This can be done using eq. (5.17)

$$E_{fn} = V_v h\nu_l \sum_{i=1}^{n_{vox}} R_{pv}\Delta t(1 - \eta_{st}) + n_{2i}[1 - e^{(-\Delta t/\tau_f)}] \tag{5.17}$$

which takes the efficiency terms from eqs. (5.15) to (5.16) and attributes the electrons lost from the upper state population to spontaneous emission. The energy is then calculated by multiplying by the laser photon energy $h\nu_l$. This new fluorescence energy E_{fn} is then used to define the starting energy of the rays in the next time-step. E_{fn} is divided by the number of rays in the ray-trace and the gain calculations for the time-step can proceed.

Once this process has continued and all of the necessary time-steps have been calculated the model is complete.

5.2.6 Final Energy Detection

At the end of each time-step, once gain calculations have been completed, it is possible to determine the energy seen by any detectors. There are two types of detector that will be considered here, those that represent physical detectors such as photodiodes and those that are virtual.

Zemax is capable of recording the state of rays non-destructively, in a way real detectors can't, for any object they come into contact with in the model. This is a type of virtual detector, the object does not even need to have a physical counterpart in the real system, but Zemax can allow for the collection of useful information from these objects. The voxels in the gain medium are calculated in this way, a cuboid object is placed in the same space as the optical gain medium within the Zemax model but is not optically active, the rays ignore it during the ray trace. This object can then have a resolution defined in a number of pixel in each of the cardinal axes (xyz) which determines the size, spacing and number of voxels, distributed regularly and evenly throughout the object. These voxels can then record properties of the rays that pass through them, such as incident flux and polarisation. They are also capable of recording information about any interactions that occur within them, such as energy absorbed by a material, which is how the pump absorption is calculated.

The other type of detector used in this model, are those that represent real detectors in the physical system. These can be modelled simply, as flat objects that absorb rays that intersect with them and record their properties, or more accurately, behaving as the physical detectors do.

Once the detectors are defined and the model complete, it is then possible to review the energy incident upon the detectors. For each of the rays that interact with a detector object, at the end of each time-step, the total gain-adjusted energy is summed, giving the final energy.

This process is repeated for every time-step. The number of time-steps to consider is determined by the length of time the user wishes to model the system for and the length of the time-step to consider. A more representative and accurate model of the system is achieved by shortening the time-step, at the cost of increased computational time, as this remains the same for each time-step.

5.2.7 Limitations

This section will go into more detail over some of the compromises made in this model. Many of the alternatives presented here were tested before a final decision was made, to ensure they did not compromise the accuracy of the model too greatly. The driving factor in many of these decisions were the time taken for the model to complete. Increased computational power and algorithm efficiency would be promising avenues for further development.

In section 5.2.4 the decision to use a single ray-trace for the entirety of the model was made. A fuller treatment would be to recalculate the weights for every voxel according to eq. (5.10) at every time-step. As the metastable level is depleted each time-step through amplification, as discussed in section 5.2.5, the distribution of energy throughout the gain medium could drastically change. This would lead to the fluorescence for the next time-step to also differ. It would then be necessary to recalculate the origin voxels for each of the rays and retrace the new ray-paths. These steps, however, take a significant amount of time compared to the rest of the calculation and recalculating for every time-step causes the model to take an unacceptable amount of time. Furthermore, the process of converting the ray database to MATLAB requires saving a file, which is typically several gigabytes in size, to disk and is often one of the slowest computational operations. If all ray-traces are to be saved for further analysis, storage space quickly becomes an issue having to store hundreds of these files. Initial trials of this fuller approach were conducted and saw only marginal variations at much greater a computational cost.

In section 5.2.5 it was stated that gain depletion is calculated on a ray-by-ray basis and rays are selected in a random order. It would be possible to calculate the temporal order in which each of the ray-voxel intersections occur, for all rays, and then follow this ordering for depletion purposes. This would be a further area for development on this model and may give a more accurate representation of the gain depletion phenomena for higher numbers of rays. With the limited number of rays traced in this model this ordering has less of an effect, as multiple rays passing through the same voxels would be far more common with a greater number of rays.

5.3 Verification

To verify this model, an experiment concerning ASE from the literature [59], has been replicated. The same set-up will be used, in this model, as was present in the physical experiment. Achieving comparable results should indicate the effectiveness of this model to the modelling of ASE in solid-state lasers.

The experiment used here will be the one presented by Huß *et. al* 2010, in "Suppression of parasitic oscillations in a core-doped ceramic Nd:YAG laser by Sm:YAG cladding" [59]. This experiment demonstrates how Sm:YAG cladding is capable of suppressing parasitic oscillations in cylindrical Nd:YAG laser rods, reducing the role ASE has in gain depletion within the gain medium. By introducing another medium around the laser rod, that absorbs the 1064 nm wavelength light, any spontaneous emission that travels transverse to the barrel of the rod sees large losses. By increasing these losses, less energy can be extracted by some of the spontaneous emission rays that will not couple to the intended cavity mode.

This experiment was chosen as it had comparisons of multiple rod geometries and the measured effect of ASE on the system. It also illustrates the temporal evolution of fluorescence from gain media, both with and without ASE suppression. Replicating these results would be a good indication that this model can be used to assess the affects of ASE.

This section will first present the laser pumping scheme used in both the model and experiment. The absorption profile in the laser rod, as generated by Zemax, will be discussed for the Sm:YAG clad rod in comparison to a homogeneously doped Nd:YAG rod. The results from the literature will then be presented and discussed, for both the Sm:YAG clad and unclad rods. Then the figures produced by this model will be compared to those from the literature. Finally, other findings from the model, regarding the energy extraction in this pumping scheme, will be presented.

5.3.1 Pumping Scheme

In the experiment a DILAS GmbH, N7F-806.7-1000Q-H207 quasi-continuous wave diode laser was used as the pump source, with a peak power of 1kW and a pulse duration of 200 μs. The pump laser was temperature tuned to 806 nm wavelength, with a 2.5 nm full-width half-maximum (FWHM) spectral width [59]. This was then passed through an 800 μm diameter, 0.22 numerical aperture fibre. The pump light, output from the fibre, was modelled as an area emitter with the same size as the fibre core with equivalent properties in Zemax. This same set-up was modelled in Zemax for the purposes of the model.

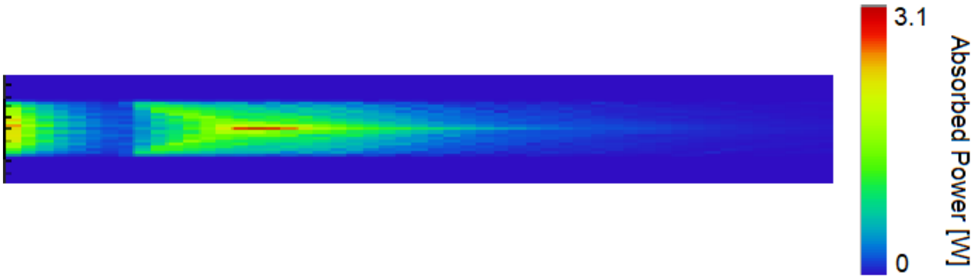

Figure 5.8: A two-dimensional aggregation of the Pump profile as calculated by Zemax for the pumping scheme presented by Huß et. al [59].

The dichroic mirror of the laser cavity was coated onto the laser rod, modelled as anti-reflective at 808 nm and high-reflectivity at 1064 nm wavelength, using Zemax's coating functionality. The other end of the laser rod had an anti-reflective coating, at 1064 nm. The properties of the coatings used in the experiment are not detailed in the Huß paper. Due to the computational limitations affecting the model, as discussed in section 5.2.7, the coatings were set to be perfect transmitters and reflectors at the specified wavelengths. That is, the anti-reflective coatings transmitted, and the high-reflectivity reflected, all of the power at the given wavelengths. Absorption by the coatings was set to zero.

Two 30 mm long and 4 mm diameter cylindrical, composite core-doped, Nd:YAG laser rods were used in the investigation. Both rods had YAG with 1.0 at. % Nd^{3+} cores of 2 mm diameter and different cladding materials, one with undoped YAG and the other doped with 4.0 at. % Sm^{3+}. Both rods also had polished barrel surfaces.

The light from the source point was then focused, a distance of 2.5 mmfrom the end of the rod, to achieve multiple internal reflections. This was achieved by using an 8 mm focal length lens.

The volumetric detectors allow for examination of the three-dimensional absorption profiles, of objects within Zemax. For detectors assigned 51 pixels in each of the three x, y and z axes, the absorption profile for both the core-doped rods are similar, fig. 5.8 shows the absorption for undoped YAG clad rod. These profiles match those of their comparative examples, presented in the Huß paper, affirming the accuracy of this pumping model. As can be seen a second maxima occurs approximately 6 mm along the rod due to internal reflection of the pump light by the rod's polished sides. There are little reflections at the Nd:YAG and cladding interfaces, as the refractive indices are very similar. It should be noted, the small black lines on the left edge of the pump absorption profile in fig. 5.8 are pixel markers from the Zemax software.

5.4 Results

This section will compare the results from the Huß paper, to those produced by the model. Figure 5.9 show the intensity on the photodiode over time, for both the of the core doped rods. A number of pumping powers are used to illustrate how the development of ASE limits the maximum achievable intensity in the YAG clad rod.

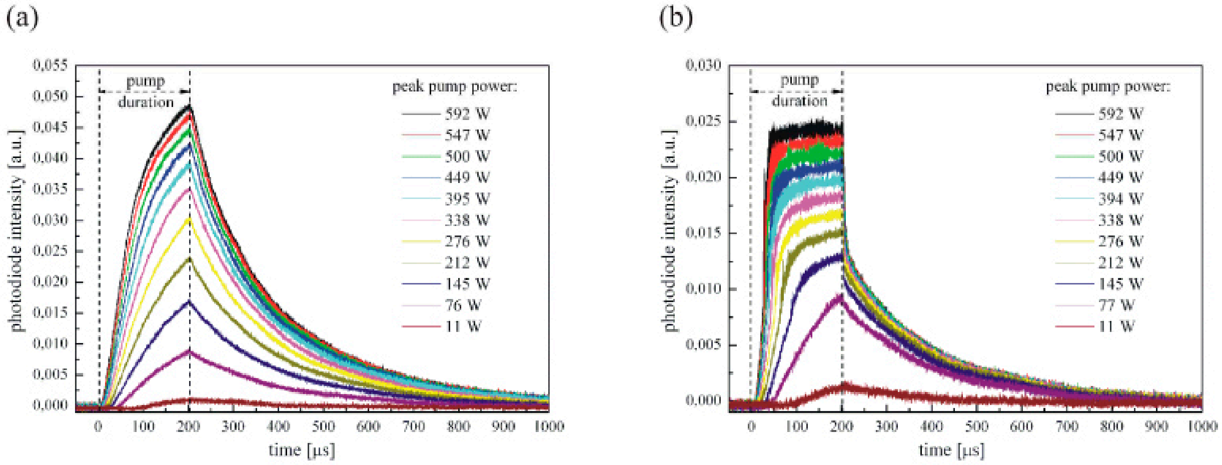

Figure 5.9: Comparison of the time-resolved fluorescence measurements of the Nd:YAG laser rod with both the Sm:YAG (a) and YAG (b) cladding, as presented by Huß et al. [59]. Note, the photodiode intensity is presented in arbitrary units [a.u.].

The rods are pumped for 200 μs and the decay of the fluorescence intensity is monitored on the detectors to a time of 1000 μs. The YAG clad rods plateau to a steady-state, above a pump power of 145 W. As shown in figs. 5.9 to 5.10, the onset of this plateau occurs earlier as the pump power is increased, indicating stronger parasitic effects as the potential gain increases. In the absence of parasitic effects, the fluorescence on the detector would plateau higher and at a later time, as population inversion would be able to build. In contrast to this, the Sm:YAG clad rods do not plateau as pump power increases. The cladding suppresses these parasitic effects by absorbing energy from rays that travel transversely across the rod. This reduction in energy means those rays are not able to deplete as much gain, as they are not able to stimulate as many photons, when travelling through the pumped regions of the rod. It should also be noted that higher intensities are recorded on the photodiode for the Sm:YAG clad rod for all pumping powers above 76 W, with the peak intensity reaching almost double that achieved by the YAG clad rod at its highest pumping power.

Figure 5.10 shows the time-resolved fluorescence signal from both rods, pumped at equivalent powers, normalised for direct comparison of signal shape. The YAG clad rod reaches a steady-state plateau before 50 μs whereas the Sm:YAG clad rod keeps increasing throughout the pumping period. The decay is also noticeably different for both rods, being much steeper for the YAG clad rod. This indicates the parasitic oscillations in the YAG clad rod deplete the gain much faster than in the Sm:YAG rod. As the rays oscillating within the rod can achieve higher intensities, due to ASE, they can stimulate more emission from the population inversion limiting the maximum steady state gain and leading to faster decay. This explains why an elbow is seen in the decay more prominently for the YAG clad rod.

These results, showing the clear differences in performance between the two rod types, were replicated using the model. A time-step of 2.5 μs was used for all of the models. Seven pump powers at; 11, 77, 145, 276, 394, 500 and 592 W were investigated for comparison with the Huß paper results.

Figures 5.11 to 5.12 show the same results as presented from the paper in fig. 5.9, for a number of different

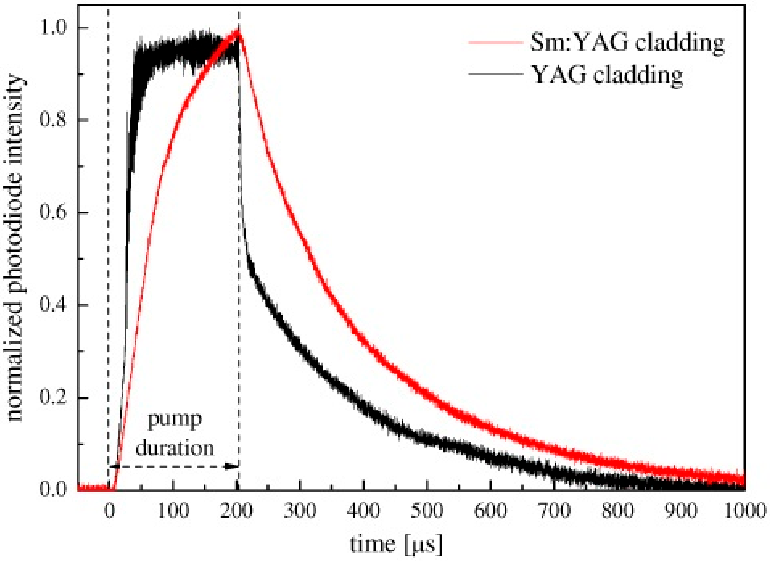

Figure 5.10: Comparison of the normalised time-resolved fluorescence for the Nd:YAG laser rod with both the Sm:YAG (red line) and YAG (black line) cladding, as presented by Huß et al. [59].

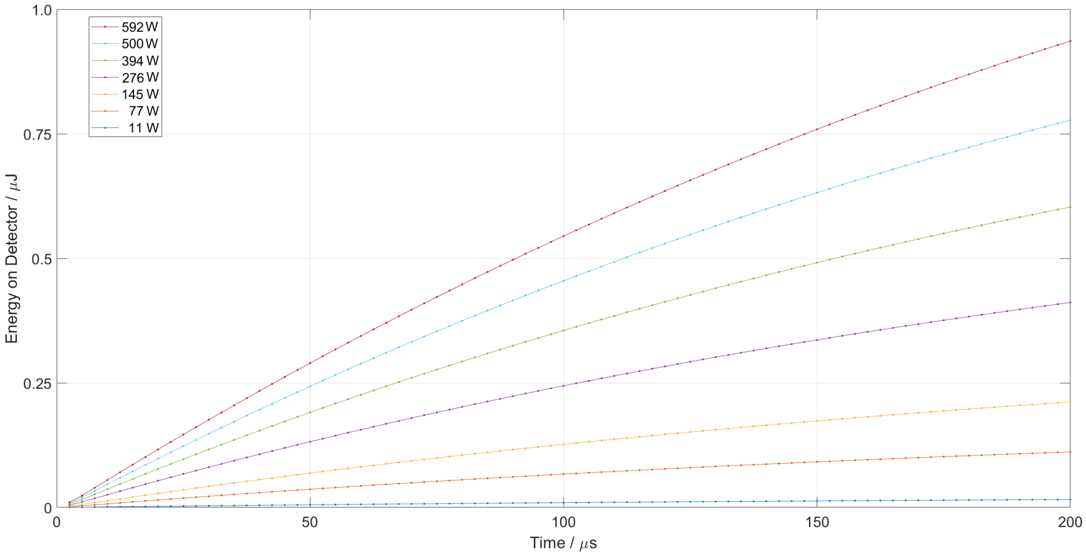

Figure 5.11: Modelled results of the energy observed on the detector in each time-step for the Sm:YAG clad rod for several pumping powers, giving good agreement to the experimental results shown in part (a) of fig. 5.9.

pumping powers, for the Sm:YAG and YAG clad rods respectively.

Table 5.1 higlights the difference between the two claddings, comparing the final energy stored in each at the end of the 200 μs pumping period.

As can be seen, in comparison to the results from the Huß paper, this method can closely simulate the behaviour of ASE and parasitic effects on laser performance. This allows for a better understanding of the ASE process during, and after, the pumping cycle by enabling the user to quantifiably test and measure the effects changes in the system have on this behaviour. This affords users the capability to investigate multiple pumping schemes to consider which would best reduce these effects, prior to spending time or money building them in the laboratory.

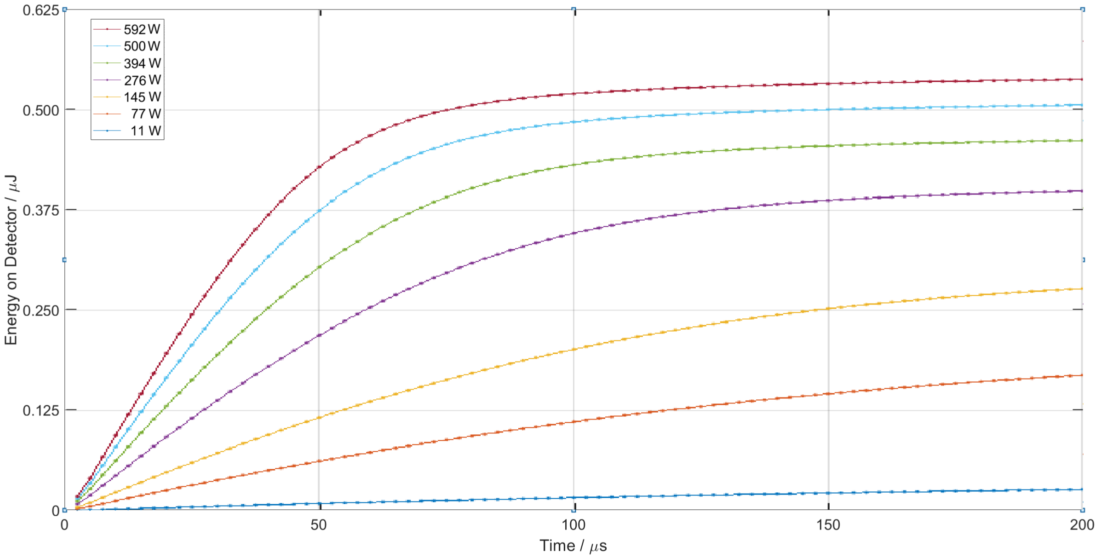

Figure 5.12: Modelled results of the energy observed on the detector in each time-step for the YAG clad rod for several pumping powers, giving good agreement to the experimental results shown in part (b) of fig. 5.9.

Table 5.1: Comparison of the energy stored at the end of the 200 μs pumping period for both the Sm and YAG clad rods, according to the model.

		Energy Stored		
Pumping Power (W)	Energy Input (mJ)	Sm (mJ)	YAG (mJ)	YAG/Sm (%)
11.00	2.20	1.28	1.27	99.69%
77.00	15.40	8.94	8.71	97.38%
145.00	29.00	16.83	15.91	94.51%
276.00	55.20	32.02	28.71	89.68%
394.00	78.80	45.67	39.72	86.98%
500.00	100.00	57.92	49.35	85.21%
592.00	118.40	68.53	57.63	84.09%

5.4.1 Energy Extraction

This model also allows for investigation of processes that can not be readily monitored, or directly measured, in a physical experiment. This is achieved using virtual volumetric detectors that can cohabit the same space as physical objects, that do not interfere with the optical phenomena. One such example of this are the detectors used to determine the three dimensional pump absorption profile, within the laser rods.

As the rods are discretised and each voxel is modelled separately, it is also possible to monitor all of the processes that occur within each of them. Knowing exactly how much energy is extracted from the gain medium at each time-step, enables pumping scheme optimisations to consider that as a metric. This is something that can not readily be achieved with experiment, as there is no direct method for differentiating the energy from fluorescence at a given time and that from gain.

Figures 5.13 to 5.14 shows the sum of the total energy extracted, from each of the voxels at every time-step, for the same two pumping schemes presented in figs. 5.11 to 5.12. This data is produced by totalling the new ray

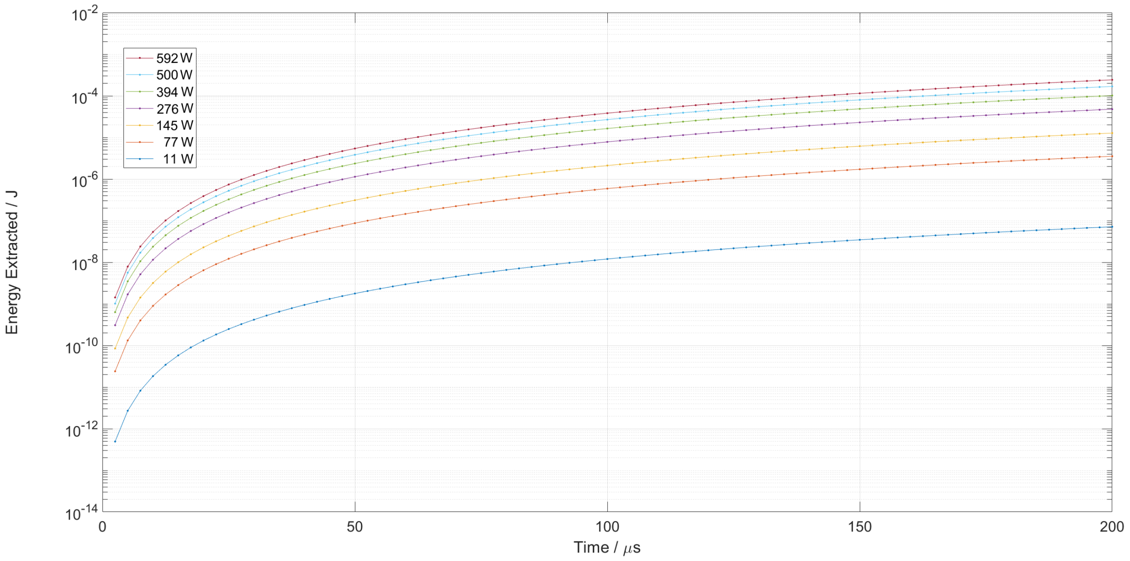

Figure 5.13: This figure shows the total gain energy extracted from all voxels at each time step for the Sm:YAG clad rod.

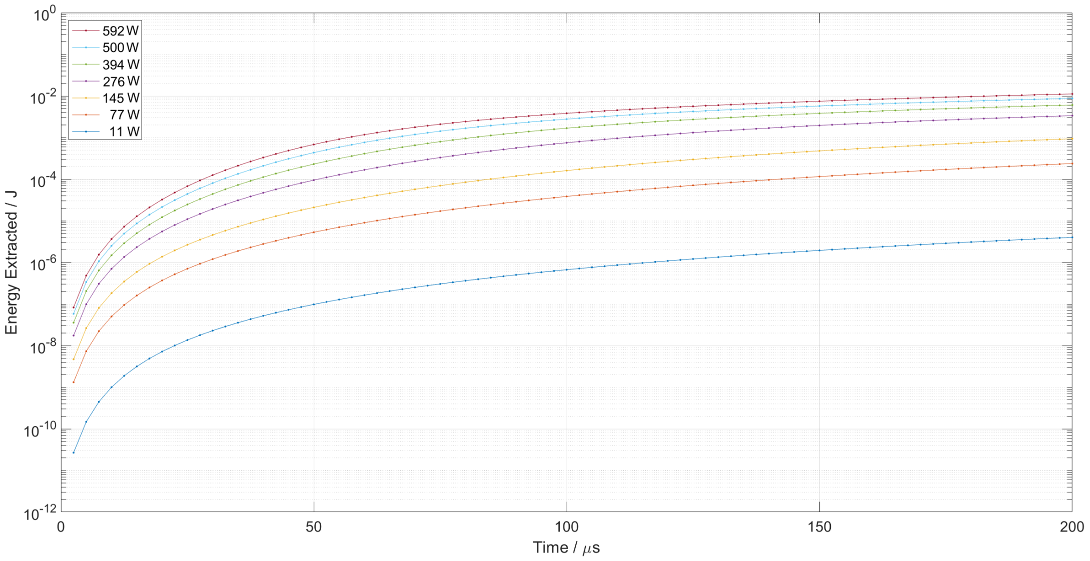

Figure 5.14: This figure illustrates the total gain energy extracted from all voxels at each time step for the YAG clad rod.

energy E_{in}, presented in eq. (5.13), for each of the voxels at each time-step. This is straightforward to record during the calculation in MATLAB and output as a separate result to the ray energy.

Similar to the YAG clad rod's total intensity graph in fig. 5.12, the total gain energy extracted reaches a plateau that occurs earlier as pump power is increased. The Sm:YAG rod results, however, increase throughout the pumping period and do not reach a limit. The scales of both figures should also be compared, with the YAG clad rod reaching a maximum at 72.1 W for the highest pump power, whereas the Sm:YAG rod only reaching 3.1 W. This highlights the considerable inefficiencies caused by ASE when parasitic oscillations are not suppressed. Comparing the 592 W pumping power results, after 100 μs the rays in the unclad rod are extracting 72 W of power and those rays in the Sm:YAG clad rod are extracting 1.1 W. Despite this large difference in energy extraction, a 2 mm square detector object, placed at a distance of 5 mm parallel to the longitudinal edge of the rods, in both models, indicates both rods are emitting roughly the same energy, approximately 0.44 mJ in the time-setp, or 2.2 W. This highlights how much energy is wasted by parasitic effects and that proper treatment of ASE can decrease wasted energy by almost 14 mJ each time-step, or 70 W, in the case of this system. Wasted energy will, ultimately, limit the maximum available energy stored within the gain medium, which can be a vital parameter in Q-switched systems.

5.4.2 Gain Depletion

This section will briefly compare the output from the model, both with and without gain depletion. This means that, when gain depletion is considered, after a ray have passed through a voxel and extracted some of the energy, the next ray that should pass through the same voxel will experience lower gain, until the gain is entirely depleted or the time-step is complete. When gain depletion is not considered each ray will experience the same gain passing through the voxel, regardless of any other rays that might have also passed through it. This should highlight the considerable difference gain depletion can make to discrete spatio-temporal laser gain medium modelling.

In this section a new model will be built to demonstrate these differences. In Zemax, a 3000 W circular source object of 2mm diameter, with a divergence of 20° was used to pump a 1%, homogeneously doped, square, Nd:YAG rod of 40 mm length and 5 mm width, at 808 nm wavelength. The model was run for 1000 μs to observe the fluorescence output on a 2 mm square detector placed parallel, 5 mm from one of the longitudinal edges of the rod. Using virtual volume detectors that contain the rod, it was possible to observe that 1200 W of power was absorbed throughout the entirety of the gain medium.

Figure 5.15 shows the output the model with gain depletion, in blue, and without in red. Without any gain in the system, we would expect 1200 W of 808 nm pump to produce a maximum fluorescence output energy of 0.911 mJ at 1064 nm wavelength in 1000 μs; this highlights that the model goes non-physical before 250 μs. For the model considering gain depletion the fluorescence slowly climbs reaching 700 mJ. The fluorescence climbs to an more than 30 times the available input energy. In this case, the power output only plateaus when the rate equations underlying the population inversion calculation reach an equilibrium steady-state and ground state bleaching begins to occur.

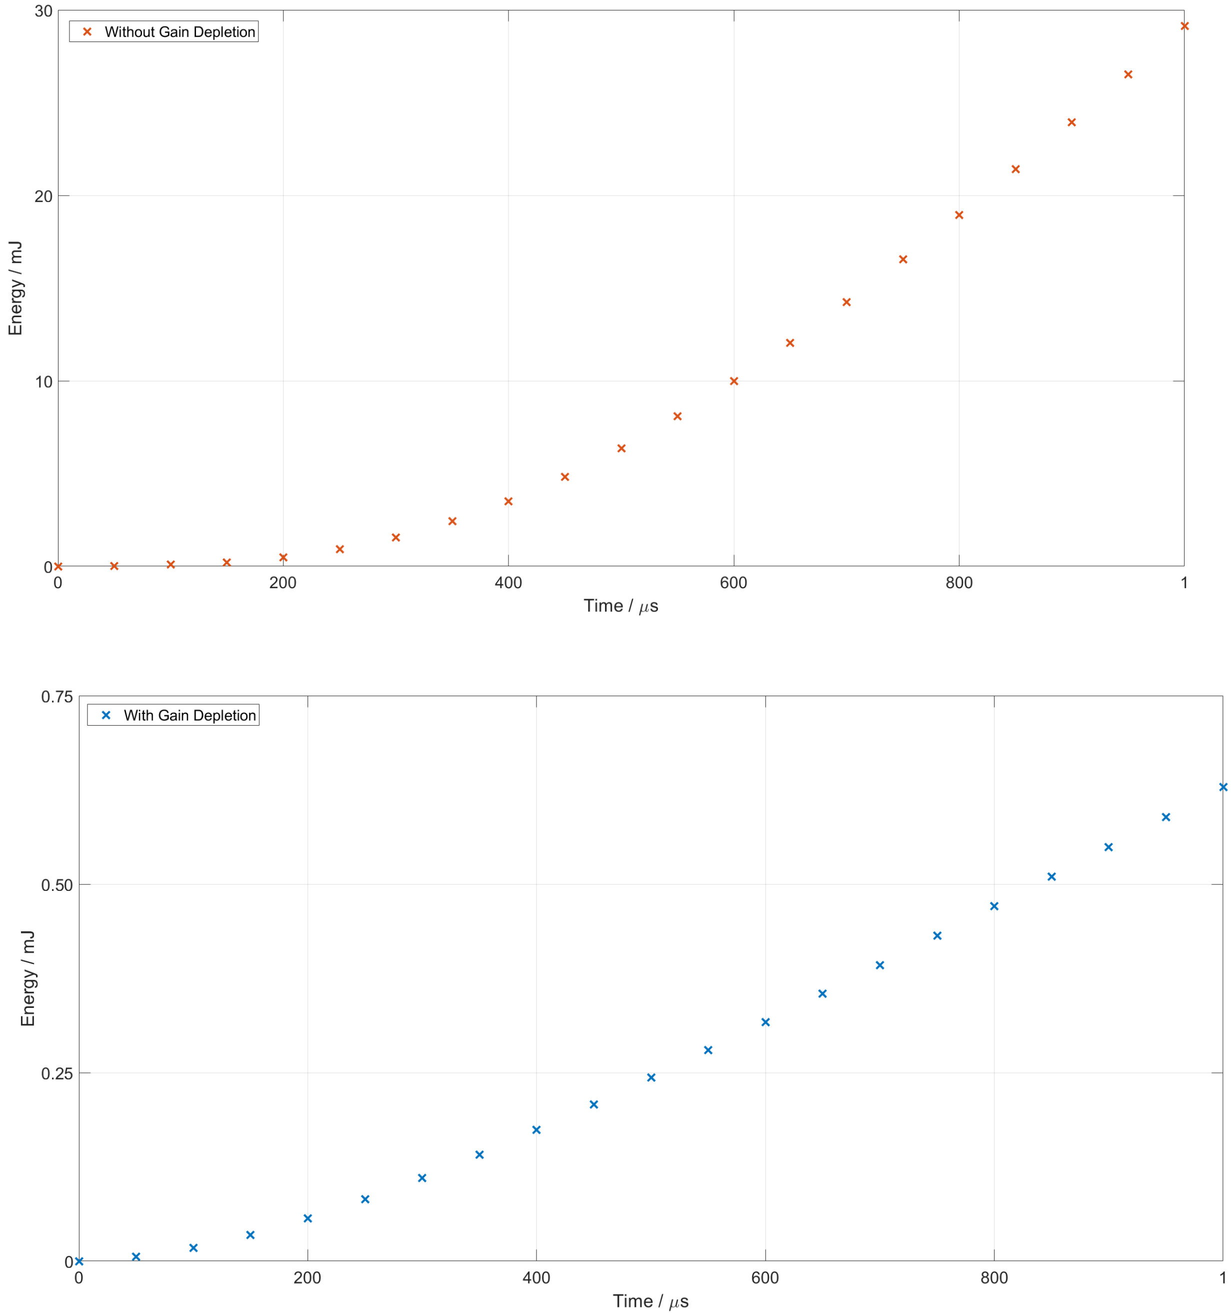

Figure 5.15: Comparison of two models for laser gain, both with and without gain depletion for a 1%, homogeneously doped, square, Nd:YAG rod of 40 mm length and 5 mm width, pumped at 808 nm wavelength.

5.5 Predictions

Having proven that this method of modelling solid-state lasers gives good agreement with experiment, it can now be used to predict the extent to which parasitic effects, such as ASE, limit energy-storage, a vital parameter in Q-switched laser systems. In this section different rod geometries have been investigated, using the model, to compare the effects ASE has on their performance. As can be seen from section 5.4, this model accurately predicts the time-resolved fluorescence output for solid-state laser rods experiencing ASE and parasitic effects, at a number of pumping powers. This allows the user to model and compare other interesting rod geometries prior to physical experiment, enabling them to test multiple designs and better understand their performance, whilst considering gain depleting phenomena. This, in turn, will enable them to test designs for lasers systems more quickly and at less cost, giving closer insight into how ASE may be affecting performance.

5.5.1 Gain Medium Geometries

One method of limiting parasitic effects and the impact ASE has on the maximum achievable performance of a laser gain medium, is to alter the geometry of the gain medium [54, 79]. Doing so can reduce the prevalence of total internal reflection of fluorescence within the medium, which can significantly increase possible path length. Two rod geometries have been compared in this section using the same pumping scheme as discussed in section 5.3.1. The first is a square rod typically used in defence applications such as target-designation or range-finding, optimising the performance of these types of laser are important because reducing the power demand to produce the same pulse strength can lead to smaller and lighter designs. The second rod is triangular, a design that will be tested to see if it exhibits any beneficial properties with regards to ASE suppression.

The 30 mm long square rod, with 2.63 mm sides was compared to a rod with an equilateral triangular faced rod, also 30 mm long and 4 mm on each side. This ensured that the pumping surface area, and therefore the pump intensity, of both rods were equal for the purposes of this comparison. The rods are homogeneously doped 2% Nd:YAG, typical for lasers used in defence applications [101], and had their central axis aligned with the centre of the pumping optics. These rods were modelled as having polished surfaces, again typical for defence lasers of this kind [99]. Both rods were pumped at nine different powers, 0.5, 2, 10, 50, 250, 500, 1000, 2000 and 4000 W. This was to investigate the performance of these rod geometries over a large range of pumping powers, from low to high gain, and examine how parasitic, gain depleting effects inhibit their performance. Time-steps of 5 μs were used for the model over a typical Q-switched pumping period of 200 μs [99, 101].

The square rod shape was chosen as it is very common in many laser systems for defence applications [101] and understanding its performance, in relation to other rods aids in a better understanding of how these gain depleting phenomena limit typical lasers.

The triangular rod shape offers a look at how an unconventional geometry might assist in the mitigation of some of these effects. Long paths within the laser rod and cavity, that aren't coupled with the laser output mode, contribute significantly to gain depletion. Total internal reflection leads to many long ray paths forming within the laser rod. Rays incident upon the surface of both cylindrical and square rods can experience significant internal reflection, at a wide range of angles. This leads is due to subsequent reflections occurring at angles very similar to the first reflected angle. For example should a ray, travelling perpendicular to the longitudinal axis of square rod, be incident upon the inner surface at an angle of 40° to the normal and be reflected, it will strike the next surface at 50° and will continue in this manner in a closed loop around the inside of the rod. Should this ray experience enough gain along this continuous path to counteract any losses, such as absorption or transmission at the surface, then the ray will continue to deplete potential gain from the system. Although this two-dimensional path, transverse to the length of the rod, does not make up the bulk of the random paths through the system, should it experience enough gain it can quickly grow to become a dominant effect on laser performance.

The triangular rod, however, has 60° angles between each of the surfaces. The same 40° incident ray, as described in the previous paragraph, will strike the second surface at 20° to the normal, meaning that much smaller incident angles are allowed after multiple reflections, which in turn are less likely to exhibit total internal reflection. The 60° between each surface means that rays incident at greater than 60° will not see the same pairs of angles through all of the reflections, as they do for square rods. Rays over 60° experience multiple changes to their incident angles along their path. For example, a ray with a 70° angle will strike the second surface at 20° to the normal, then be reflected back toward the first surface, where it will have an incident angle of 30° .

This limitation of the size of angles that can persist along a ray's path, within the rod, should decrease the number of rays experiencing long paths, and thus reduce parasitic oscillations. Comparing this rod shape to the more conventional square and cylindrical rods, should verify this hypothesis. It is for these types of comparisons that this model was originally intended, where developing potentially difficult or costly components would previously hinder experimentation.

5.5.2 Pump Profiles

Figure 5.16 shows the pumping absorption profiles for the triangular faced rod on the right and square rod on the left. The longitudinal profile shows that most of the pump is absorbed in a very short distance. In the triangular rod longitudinal image, the left side of the absorption profile corresponds to one of the corners of the triangular face, whereas the right is parallel to a flat side. Compared to the absorption for the square rod in fig. 5.16, it can be seen that less light is reflected from the left back into the centre of the rod.

Identical pumping was used for both rods, centred along the rods longitudinal axis, this leads to the similar high absorption circular pattern at the centre of both rods.

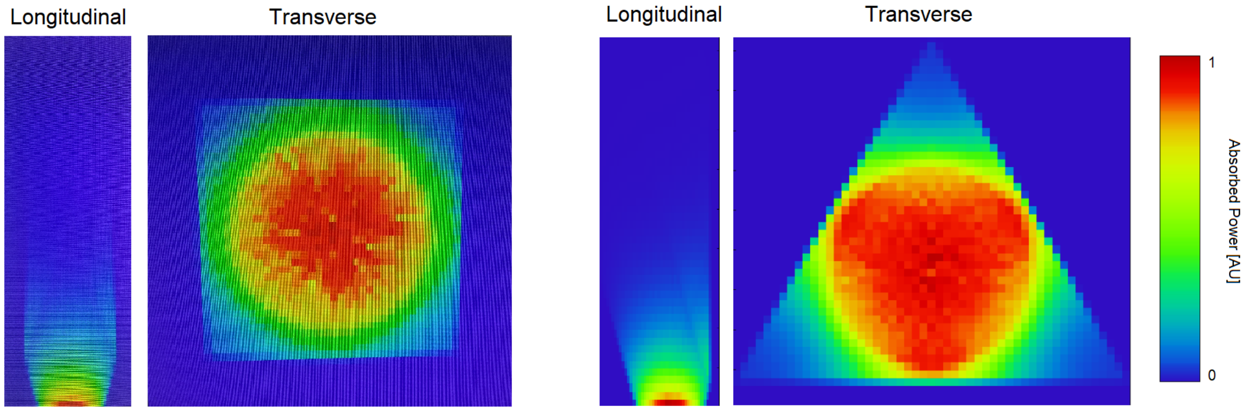

Figure 5.16: Both transverse and longitudinal pump absorption profiles for 2% Nd:YAG square rod (left) and triangular rod (right) as modelled in Zemax.

5.5.3 Results

Figures 5.17 to 5.18 presents the the energy measurements of the spontaneous emission for both rods. Energy was measured by a 2 cm square detector, parallel to the longitudinal edge of the rod, 2 cm away near the pumping edge. The results for multiple pumping powers are shown over the 200 μs pumping period. The square rod on the left shows that, as the higher pumping powers are used, the output energy on the detector begins to plateau earlier in the pumping period. The upper limit of energy detected does, however, increase through the entire range of pumping powers.

The triangular rod does not appear to reach this plateau in the same way that the square rod does, but steadily increases throughout the pumping period for all of the input powers. The absolute energy at the end of the pumping period is lower than that of the square rod, for all pumping powers except the highest of 4000 W. Some of this can be attributed to the reduction in total internal reflection of pumping power, meaning some pump light incident upon the rod is lost before being absorbed. This can be seen in the longitudinal pump absorption profiles presented in fig. 5.16, where the square rod clearly shows pump absorption further along the rod compared to the triangular rod, despite being the same material. Total internal reflection causes more of the pump light to remain in the rod for absorption, for the square rod. Most of the difference, however, occurs because the spontaneously emitted rays do not experience significant internal reflection. This results in the rays that are fluoresced, and hit the detector, have not depleted the gain as much, and thus the triangular rod storing more energy.

It should also be noted, however, that at the highest pump power of 4000 W, the triangular rod does exhibit more energy on the detector. This higher power also increases steadily towards the end of the pumping period, whereas the square rod plateaus around 100 μs.

It is also possible, using this model, to examine the total energy extracted from gain by all rays throughout the pumping period. Figures 5.19 to 5.20 shows this value for both rods at each of the pumping powers. For the square rod, the shape of the energy extracted graph follows a similar pattern to the detector energy graph in fig. 5.17. The amount of energy extracted increases as the pump period continues and reaches a plateau. For the triangular rod in fig. 5.20, however, it seems that the increase in extracted energy reaches a plateau within the first few time-steps of 5 μs. For the highest pump powers, it appears that the energy extracted also continues to increase throughout the pump period, in line with the continued increase in detector energy shown in figs. 5.17 to 5.18. Table 5.2 shows the comparison for energy stored at the end of the 200 mus pumping period for both rods. It is interesting to see that the triangular rod isn't as efficient in the middle pumping powers, but significantly outperforms the square rod for the highest pumping powers. This indicates that a lot more of the absorbed energy remains stored in the rod compared to the square rod and is not depleted as much by ASE and parasitic phenomena when gain is high.

The triangular rod geometry is effective at suppressing the gain depletion by ASE and parasitic oscillations for high pumping powers. As these simulations have been carried out in the absence of a cavity, the results are highly

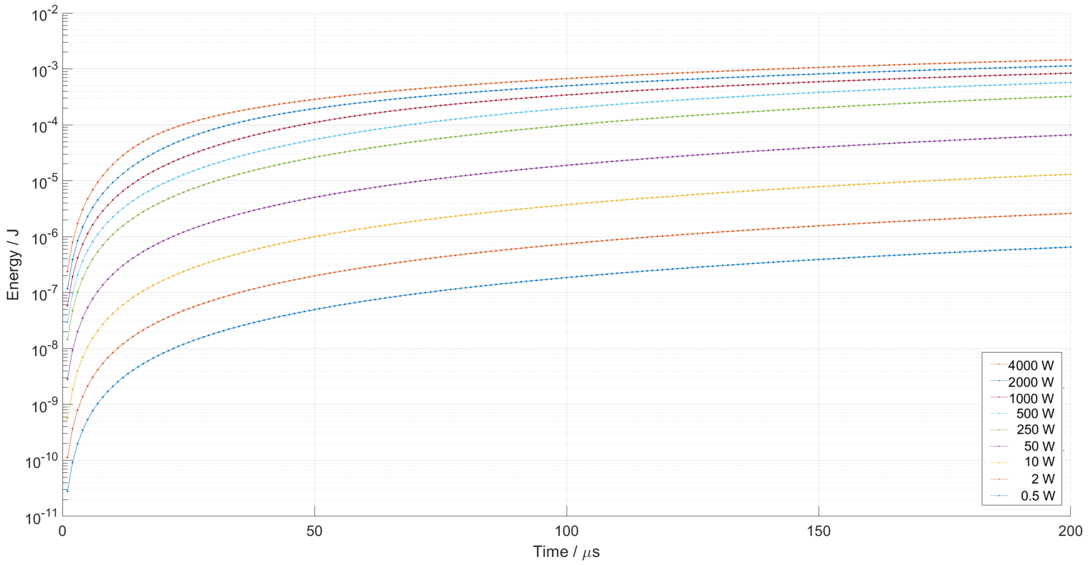

Figure 5.17: Energy observed on the 2 cm square detector over 200 μs for the square rod at pumping powers of 0.5 to 4000 W.

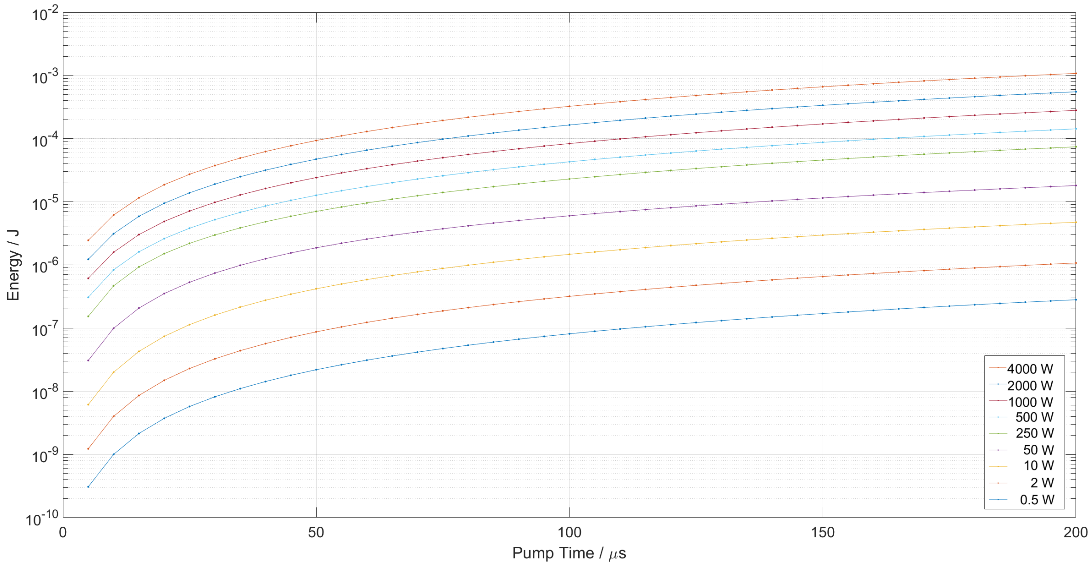

Figure 5.18: Energy observed on the 2 cm square detector over 200 μs for the triangular rod at pumping powers of 0.5 to 4000 W.

applicable to Q-switched lasers. The triangular rod geometry emits lower fluorescence at all but the highest pumping powers and time periods. This means, that for Q-switched operation, the triangular rod would hold more of the gain at the moment of switching when in the high power regime.

Table 5.2 higlights the difference between the two claddings, comparing the final energy stored in each at the end of the 200 μs pumping period.

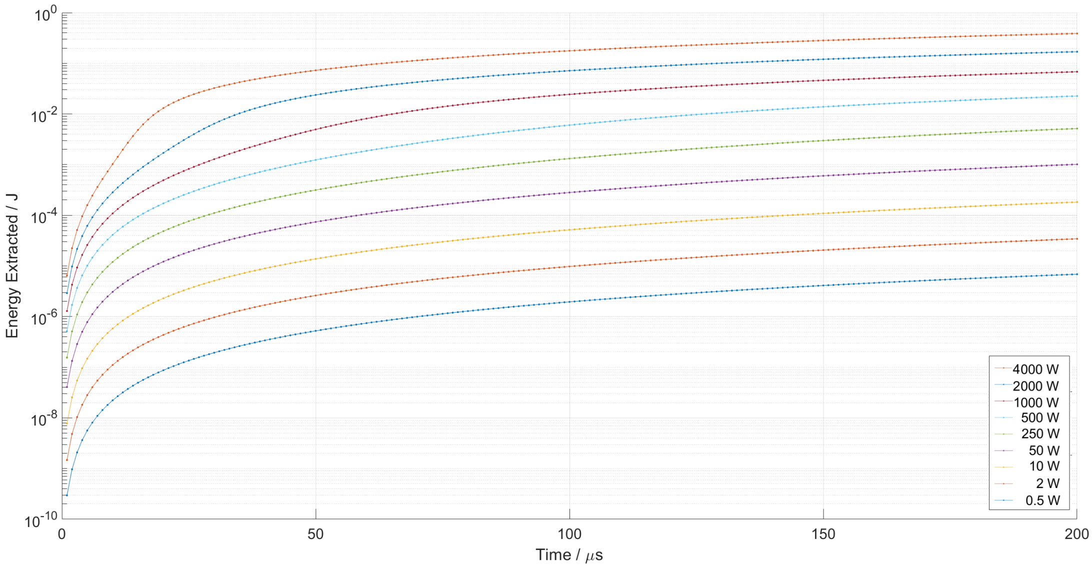

Figure 5.19: Total gain energy extract for all rays over 200 μs for the square rod at pumping powers of 0.5 to 4000 W.

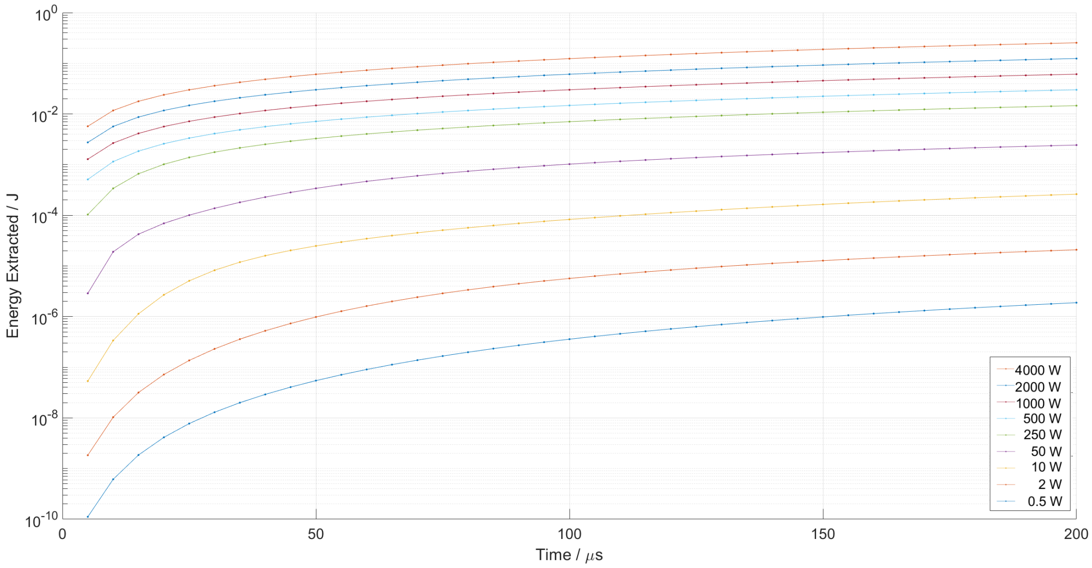

Figure 5.20: Total gain energy extract for all rays over 200 μs for the triangular rod at pumping powers of 0.5 to 4000 W.

Table 5.2: Comparison of the energy stored at the end of the 200 μs pumping period for both the square and triangular rods, according to the model.

		Energy Stored		
Pumping Power (W)	Energy Input (mJ)	Square (mJ)	Triangular (mJ)	Square/Triangular (%)
4000	800.00	76.27	210.14	36.30%
2000	400.00	62.75	108.11	58.05%
1000	200.00	47.97	55.17	86.96%
500	100.00	35.47	28.06	126.41%
250	50.00	23.88	14.43	165.43%
50	10.00	4.80	3.37	142.19%
10	2.00	0.98	0.90	108.79%
2	0.40	0.20	0.21	93.65%
0.5	0.10	0.05	0.06	91.13%

5.6 Conclusion

A method for modelling the effects of parasitic, gain depleting phenomena, has been discussed and presented. By utilising ray-tracing methods, accurate and detailed pumping and absorption profiles, for complex laser systems, can be analysed. This offers an advantage over analytical methods, whereby it becomes prohibitively difficult to model all but the simplest of geometries.

Making the spatial analysis discrete, using a finite-element approach, then enables the modelling of small, local behaviour within the gain medium. Treating each voxel as an individual set of laser equations allows for a three-dimension representation of the gain dynamic throughout the medium. Random rays are then emitted as fluorescence and traced again to see how they travel through the system, giving an accurate representation of the spatial output from the laser. Coupling this with a discrete temporal analysis then allows the development of gain, within the laser rod, to be examined throughout the pumping period, and not just in steady-state.

The model was verified using the results from a paper on the effects of ASE on a cylindrical rod geometry [59]. Modelled results closely match the shape of those from the paper for both rod types. This showed that significant reduction in parasitic effects could be achieved, by the use of an absorbing cladding around the active gain material in the laser rod. Matching these results so closely indicates that, the approach presented here, is appropriate for these phenomena.

Finally the model was used to compare the effects of ASE on two other rod geometries, a square rod and a rod with a triangular face. The model showed that the significantly less fluorescence was observed from the triangular rod and less of the available gain energy we extracted throughout the pumping period, when compared to the square rod. This indicates that the triangular rod shape may be well suited to suppressing gain depleting effects in high pump power Q-switched laser operation.

There were many aspects to this model and it heavily relied on the interconnection of multiple systems and techniques. It was designed to fit easily into the work-flow of laser engineers already using the Zemax ray-tracing tool to design laser systems. One drawback of the its complexity means little time was spent optimising the individual processes, resulting on some aspects of the model taking hours to complete a the required time-steps for a single laser. The transfer of data from Zemax to MATLAB also depends on the generation of multiple GB sized files. Which can be problematic for some systems. Improving these computational aspects of the model would be a worthwhile, first-step, in continuing development. Making the model faster and less resource dependent would allow for fast iterations within the design process of both the lasers being modelled and the model itself. A thorough review of the codebase developed for this model by an experienced computer scientist would no doubt find many areas for improvement and optimisation of the algorithms and data structures used. There are several large changes to the approach that could be adopted that would also dramatically improve run-times, such as introducing dynamicallly size time-steps that adjusted to shorter lengths at the periods of greatest change in the system.

Another development to consider for this model, would be retracing rays that terminate due to a minimum power threshold condition being met, that see enough gain to overcome that threshold. The default behaviour in Zemax is to terminate rays that reach 0.1% of their original power to save computational power tracing rays that have little effect on the system, this condition is switched off in this model. When gain is considered, some rays that would drop below this threshold might instead remain above it. Enabling the minimum relative power condition such as this, then retracing those rays terminated by Zemax that experience gain, would allow for more rays to be traced by discarding only those of little significance and investigating more ray-paths that Zemax might consider high loss.

CHAPTER 6

Conclusions

6.1 Introduction

This chapter will aim to summarise each of the three main projects comprising this thesis. The motivations and foundations for each of these projects will be discussed briefly, seeking to put into context, the work that has been carried out and the reasoning for doing so. Other research in similar areas will be discussed, placing the relevance of each project in the scope of other, current, work. The basics of the science behind each piece of research will then be reiterated, and the outline of the methods and techniques used will be presented. The main outcomes from each project will then be summarised again, along with the major findings, before finally assessing the success and shortcomings of each of the projects.

The chapter begins with discussion of the photon time-of-flight model, developed for the non-line-of-sight (non-LOS) LIDAR system. The new techniques developed utilising ray-tracing automation, for the optimisation and tolerancing methods, are discussed in the next following section. Finally, the gain evolution models for solid-state laser systems, which consider gain depletion effects such as amplified spontaneous emission(ASE), are then presented and summarised.

Once each of these summaries are concluded the possibilities for future development, of any of the work in this thesis, will be discussed. Recommendations for where the most value lies, in the next steps for each project, are then made. Finally, the thesis as a whole will be discussed and the author's perceptions of successes and any limitations relating to the work as a whole will be presented.

6.2 Optimisation and Tolerancing for Laser Manufacturing

This chapter detailed the development of a tool that could aid in the design of optical systems by assisting in the optimisation of multiple components simultaneously. Some optical design tools already have some optimisation capability, but when more complex optical components are introduced, such as polarisation retarding waveplates,

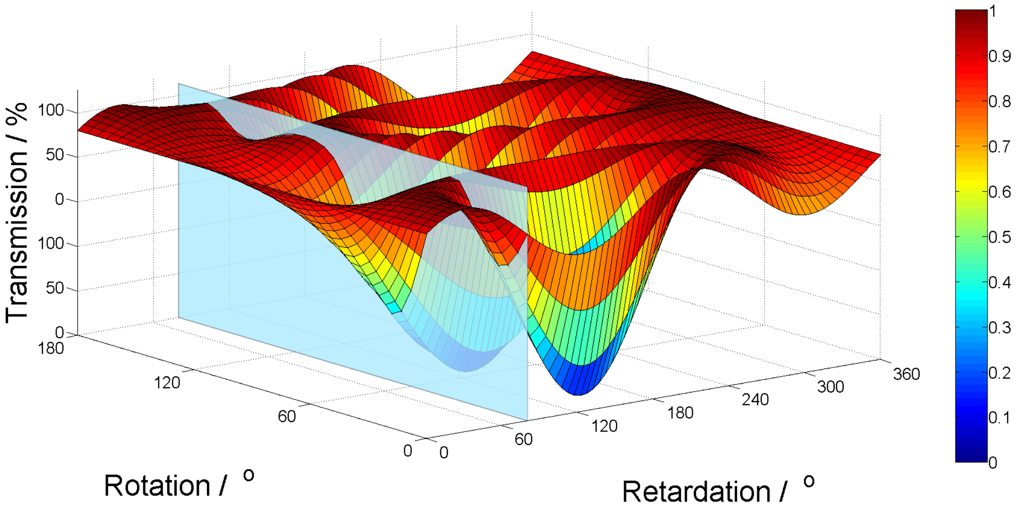

Figure 6.1: Full ranges of both retardation and rotation modelled, with slice drawn at point where rotation matched with single rotation value in fig. 3.3

the process becomes difficult or in some cases impossible to model using the conventional packages. The tool developed here also aids in the analysis of these systems with the use of three techniques which are detailed in section 3.3 and summarised in section 6.4.

The initial aim of this project was to automate a few measurements in Zemax for a system containing a polarisation retarding waveplate.

The effect a waveplate has on transmission is also dependent on the alignment of the waveplate with the polarisation of the wavefront. Varying the angle of a single waveplate can greatly affect the performance of an optical system, as shown in fig. 3.3. Calculating the appropriate Jones Matrix values for each of the retardation values had to be outside of Zemax, this made comparison of multiple waveplates over large ranges of angles difficult. This means that finding an optimal combination of waveplate and angle a costly process.

By automating control of Zemax, to allow for adjusting of objects within a model automatically, and the calculating of these Jones matrix values, it becomes possible to search the entire solution space quickly, and with minimal human interaction. Figure 6.1 shows the surface of transmission percentages for all combinations of both rotation and retardation for a waveplate in a polarisation sensitive optical system. A blue slice illustrates the subset of results that correspond to the single retardation value used, prior to automation.

With this new capability to calculate object properties outside of Zemax and to update and automate control using MATLAB, it became apparent that other, previously difficult, modelling tasks could now be achieved. The remainder of this project sought to further enhance the laser system designers work process, by enabling new automated analyses requiring minimal human effort to process. This affords the user new insights that otherwise might have been foregone due to time or monetary constraints, potentially improving the final design.

6.2.1 Analyses

Section 3.3.1 details the data collection process utilised by each of the following analyses, enabling all three to be performed after a single run. The system is first designed in Zemax, then the parameters of interest are specified in MATLAB, along with their range of potential values. The metric, by which performance is to be measured, is then also specified in MATLAB. In many systems this is not encompassed by a single metric, but many. For example, in many laser systems, peak output power and beam-shape can both be of importance and optimising solely for one and not both will lead to an inferior system design. As MATLAB is being used to automate and control the models in Zemax, it enables complex, bespoke, merit functions to be considered and evaluated for each of the potential solutions. The ranges for each of the degrees of freedom are then sampled, either exhaustively (if discrete values are specified) or via Latin Hypercube Sampling (see section 3.5.2, and the merit function is evaluated.

The first of the analyses, as presented in section 3.3.2.1, is concerned with finding optimal combinations of component parameters to maximise performance. Many optimisation algorithms will seek to merely minimise or maximise a given metric. Components in physical systems, however, are subject to manufacturing tolerances, acceptable deviations from the specified measurements on each component's parameters. The length of a laser rod, for example, may vary by several microns from one device, of the same design, to the next. Sometimes a solution may only be optimal very close to the specified measurements. Tighter restrictions placed on the tolerance for, any given component, will increase manufacturing costs.

The optimisation analysis presented here, seeks to address both of the above points. This tool allows the user to specify tolerance ranges and a performance threshold, and have the tool explore all potential solutions until a high-performing, stable result is found. Once the optimally performing combination of parameters has been found, MATLAB will then perturb their values by specified amounts and continue to evaluate the merit function, finding the highest performing solutions adhering to specified tolerance requirements. It is also possible to find solutions that offer the required performance over the largest range of parameters, increasing tolerance ranges will serve to reduce manufacturing costs, both monetary and time.

The second analysis, presented in section 3.3.2.2, seeks to highlight commonalities between parameters, both individual and combinations of, for high or low performing systems. After all of the results, have been recorded, this tool searches through subsets of those solutions to highlight parameters of interest. It is possible to take the worst performing 10% of solutions and then count how often each value, for each parameter, occurs in that subset. Values that occur most frequently within that group are likely to be related to the cause of the low performance.

It is also possible to search for other subsets, like the top 10% of performance values or for specific metrics relating to the system design. When confronted with many degrees of freedom, it is difficult to isolate the parameter, or combination of, that most influence a given metric. If the extreme ranges of a parameter's potential values occur most frequently in a subset containing the worst performing solutions in the search space, this parameter

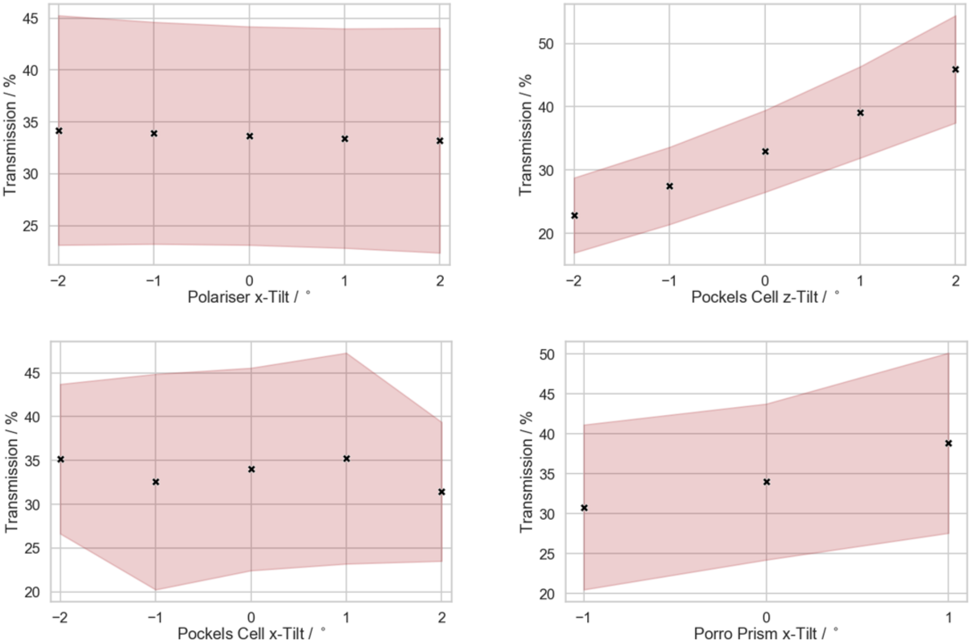

Figure 6.2: The output from the variance dominance analysis outlined in section 3.4.3.2. For each parameter at each value, all results were collected and the means calculate, displayed here as black crosses. The pink bands encompass one standard deviation of those results either side of the mean. Tighter bands indicate a parameter value that has a more significant impact.

is a good candidate for investigation, a smaller range or tighter tolerances might be beneficial to the overall system performance. This tool aims to guide the user in the process of refining a system design, by highlighting parameter values of potential interest, that might otherwise be lost amongst all of the other degrees of freedom.

The final analysis, as detailed in section 3.3.2.3, considers each of the parameters individually to determine how each of them contribute to the overall variance in performance. The analysis begin by taking a single parameter value and calculating the mean performance and confidence intervals, for all solutions containing that particular parameter value. This is repeated for each of the values of that parameter and graphs such as fig. 6.2 can be drawn. The mean of all results that include the parameter at a given value is plotted with a black cross, then a band is drawn which encompass results within one standard deviation of the mean.

6.2.2 Results

To verify the accuracy of this approach, and test out the analyses, an experiment was performed on a polarisation sensitive optical system, commonly used in defence lasers. Figure 3.1 illustrates the set-up used, a laser is passed through a polarisation beam splitter and a Pockel's Cell before interacting with a polarisation retarding waveplate, correcting for the retardation inducing components. After this the polarised beam is reflected by a Porro prism, making a second pass through the waveplate and Pockel's Cell and is finally reflected from the beamsplitter onto a detector. This set-up as it allows fine control of the output coupling onto the detector and folding of the long cavity dimension so that the system may fit into a smaller package; this set-up is also more resilient to misalignment from shock when compared to conventional systems.

Five degrees of freedom were chosen to investigate, three retardation values for the waveplate and four rotational alignments of three other, polarisation sensitive, components; the Pockel's Cell, the Porro prism and the beam-splitter. In total, 375 combinations of alignments were investigated for each of the retardation values, meaning 1175 solutions searched in total. The system was modelled in Zemax and the automation handled by MATLAB. In the laboratory, the components were mounted on rotational mounts and each of the measurements were taken at least three times to ensure good agreement. The results from the experimental work closely matched those of the model, which were slightly higher in also all cases. This was attributed to limited modelling of surface losses, material absorption of the laser light and imperfect alignment in the experiment.

One major advantage the model has over the phsyical system is the number of possible degrees of freedom that can be searched simultaneously. For merely five parameters, each with five possible values at most, 1175 results had to taken, three times. This took a week of lab time which, for a single system and only a few parameters, can be difficult to justify and very costly. Running this model, however, requires the user to only initiate the search and interpret the results. To demonstrate this, the optical system was extended with a Corner Cube and the number of degrees of freedom was greatly increased. In total, 16 parameters were each varied over a range of 5° . If just five discrete readings were taken for each of the parameters, there would be over 150 billion potential solutions. This analysis, instead, used Latin Hypercube Sampling to take 20,000 samples from the solution space and analyse the results.

The parameter frequency analysis highlighted that the tightening of the Pockel's Cell x and y-axes tolerance ranges would have a positive effect on the system's overall stability. A second analysis, incorporating those tighter tolerance ranges, saw that all remaining parameter values appeared evenly throughout the low performing results, indicating that no parameter value in particular could be solely responsible for limiting system performance. The variance dominance analysis indicated, again, that the Pockel's Cell alignment was the dominant factor in determining the overall transmission through the system. Mean performance reduced greatly as these parameters were perturbed greater than 0.5° from the aligned state. Many of the output figures from this analysis showed that the non-dominant parameters behaved similarly; their more subtle effect on transmission making little difference when compared to the Pockel's cell and to a lesser extent the Porro Prism.

Overall, this chapter presents a new set of tools to enable the optical system designer to very quickly assess key points of interest with little time-cost to themselves. Whereas previously, large search spaces, of complex optical parameters, were difficult or impossible, this set of tools automates the process and affords the user greater freedom to assess and optimise their design. The verification shows clearly that the model itself is very close to the results seen in experiment, for polarisation sensitive optical systems.

The key benefits of this approach though, are quickly realised in the extension of this work to a more complex system. Understanding the interplay between components, and which of them are of particular interest or importance for key performance measures, can be difficult to process when the system has many degrees of freedom. This tool allows the user to include as many parameters as they see fit and will quickly highlight those that may lead to important improvements when diagnosing or improving a design. Tools like this should not seek to replace the optical design engineer, but merely ease and enhance their workflow. This tool achieves that objective to support decision making in the design process and quickly limit the search space by exposing key parameters.

6.3 Non-Line-of-Sight LIDAR

This project sought to accurately model the processes involved in a novel, non-line-of-sight LIDAR, motion tracking and positioning system. The details of the method are found in section 4.2; semiconductor photon avalanche detector (SPAD), operating in time-correlated single-photon counting (TCSPC) mode are used to track the time of flight for photons emitted by a laser source. The optics that serve these SPAD detectors collect light from a small point on a scattering surface. The distance from both the laser source to the first scattering event, and the point at which the light is being collected from, is easily measured. Subtracting these from the total distance travelled leaves the maximum path length of the intermediary stage of the photon's flight. With the points of these two scattering events known it is possible to draw an ellipsoid, with those locations as the focii, using the distance travelled by a photon, calculated from the remaining time-of-flight.

If it is assumed that only a single other scattering event occurs, then that it must occur on the surface of the ellipsoid whereby the two known coordinates are the focii.

When multiple photons are scattered in this way, in a single laser pulse, it becomes possible to draw many different ellipsoids in the environment that satisfy these conditions, meaning many possible locations for the scattering surface. The point at which these many ellipsoids intersect is most likely to be the point from which the photons are being scattered. This allows the position of objects, even those out of line-of-sight, to be determined.

Many assumptions are made in the processing of this technique. These assumptions serve to simplify the calculation of a target's location but they also limit the precision with which the position of an object can be determined. Nevertheless, it is still possible to determine the location of a primary scattering event to within several centimetres. One simple assumption made using this technique serves to limit the dimensionality of the problem. As the primary application of this technique is to position a human sized target in scenes that are out of line-of-sight, it is assumed the vertical distance of the primary scattering event from the floor is 70cm. This is because the largest surface on a person is the torso and, on average, the middle of the torso is approximately 70cm from the ground. Another key assumption is that all photons are scattered from the same point on the target; this limits the number of possible paths making the problem solvable. This is necessary because it is computationally infeasible to know where on an object a scattering event occurs and, for targets that are out of line-of-sight, how large the bounds of the scattering surface are. Assumptions such as this directly limit the spatial precision with which objects can be positioned, but also limit the possible location of the primary scattering object to a single point on an ellipse in the horizontal axes, which greatly reduces the possible locations, making the problem far more easily solvable.

6.3.1 Time-of-Flight Analysis

The aim of this project was to produce a tool capable of modelling the process summarised above and in detail in section 4.2. The versatility of Zemax meant that producing representations of scenes, using built in geometric objects, was both quick and simple. It also made it possible to model the optical components used in a system such as this on vastly smaller scales than the dimensions of the scene, and calculate the ray paths accordingly. Taking the ray-database output from Zemax and interpreting it with MATLAB meant that any number of transformations could be performed on the initial data.

Zemax was therefore used to build very simple models of scenes, containing flat objects as walls, floors and ceilings, and smaller planes or cylinders to represent the target objects. Although not explored in this project, Zemax is capable of modelling much more complex geometries and importing objects designed in CAD programs. This would allow for realistic scenes with accurate models of furniture to be examined if required.

Once the scenes and optics had been designed, ray-paths were then traced. Many different approaches were used to overcome the computational limitations of this technique. Zemax can only, at most, trace $4e\times10^9$ rays, which for most scenes doesn't sample a large proportion of the available space. In the physical system, the target position is calculated from a number of laser pulses over a time-period. By making multiple ray-traces and concatenating the resulting ray databases, it was possible to build larger samples of the scene and more accurately model the process. It is also possible for the user to define which objects and scattering events they are interested in within the Zemax environment. This greatly reduces the number of rays required to model the scene, as the potential ray-paths are now limited and Zemax calculates what proportion of the would-be total rays are in the resulting data. This method, although best for developing understanding the positioning process, directly prohibits the modelling of one of the more difficult challenges of the technique; background subtraction.

In the physical system light can be scattered from any and all surfaces; there is no way to determine where the returned photons have been scattered from. This makes it difficult to find the position of stationary targets with no prior knowledge of the scene. The resulting time-of-flights histogram for each laser pulse will have varying intensities across the measured time window. To isolate those that were scattered from the target, the other scattering events must be subtracted from that histogram. This is either done by measurement of the scene without the target in it, or by comparing the median response for each histogram bucket over many pulses. One advantage of the model, however, is that perfect subtraction of the background is possible. Because the scene exists within the model, it is possible to filter out the paths that do not interact with the target. This allows for far fewer rays to be required before a position can be calculated.

With these concerns in mind, a combination of the above methods was used to determine accurate scene models.

Once the rays had been traced their time-of-flights could be calculated and analysed in MATLAB. Having the results in this format allowed for many interesting transformations to be performed. Automation of Zemax also allowed for many different scenes to be compared quickly, the results of which are summarised in section 6.3.3

6.3.2 Verification

To ensure that the model could accurately predict the process an experiment was performed to reconstruct the time-of-flight responses from a target at multiple positions.

Each of the scenes was modelled in Zemax, rays were traced and the resulting ray-database converted to MATLAB. Here the histograms from both the experimental data and the modelled data could be compared. Figure 6.3 shows the results from both the experiment, in red, and the scene in Zemax, in blue.

As can be seen form the y-axes, the total count for each of the time windows is far greater for the experimental results than those of the model. This is due to the computational limitations of the approach, it takes many hours to trace and analyse enough rays to get the requisite amount of data for this process. The experiment can collect enough data in seconds. This also accounts for the differences in the response shapes, as far fewer rays are used in the model the background has been removed by merely filtering it from the results in MATLAB, meaning that the long small tail of the response can be seen. As this is not possible with the experimental results a background response was taken, in the absence of the target. The difference in the two scene responses, with and without target, is then considered to be signal if the result is above a threshold count value. This leads to the more Gaussian shape of the response signal from the target in experiment.

Four target positions were calculated using the experimental and modelled results and were in good agreement. Very precise positioning of the target was not the goal of this comparison, as the technique is limited in the manner discussed above. This comparison illustrated that this modelling approach is applicable to reconstructing the time-of-flight response for targets in non-line-of-sight experiments. This means the approach can be used to test some of the limitations of the approach quickly and easily, which was done and is summarised in section 6.3.3.

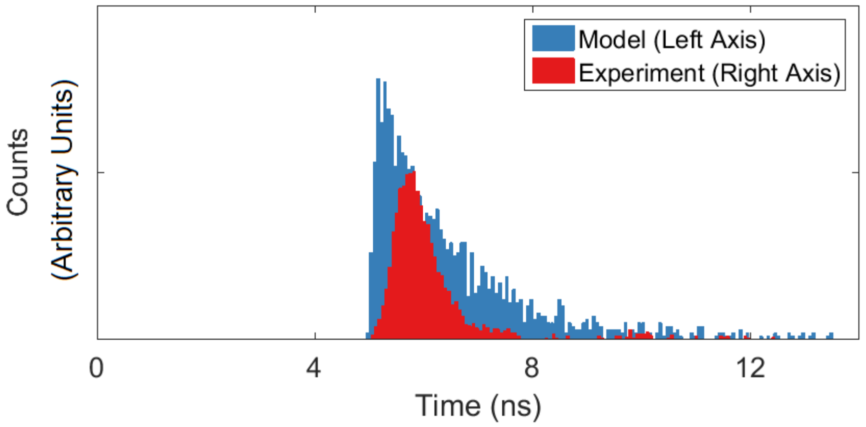

Figure 6.3: Histograms of response signal from target in position four. The data collected from experiment is displayed in red and measured against the right hand y-axis, and data from the model in blue and measured against the left hand y-axis.

6.3.3 Scene Limitations

The model was then used to initially explore some of the limitations this positioning technique might have in real applications. Four aspects of the physical setup were investigated in isolation for their potential affect on performance: range of the detector from the final scattering surface, angle made by the detector plane normal with the scattering surface normal, size of the target and finally position of the target within a scene.

Firstly, the variation in power on the detector was investigated as the distance from the final scattering event was varied. The more power that is incident on the detector, from each laser pulse, the faster the time-of-flight response from the scene can be built up. Even for scenes where the power on the detector is very low, it is possible to collect enough data to position the target, it just means the acquisition time is much longer. If the acquisition time becomes too long, more than a second, this can make positioning difficult or impossible as a human target could have moved significantly in that time. It was expected that an inverse square law would govern the relationship, between the detector distance from the final scattering event and the power collected, because of hemispherical nature of diffuse scattering events. As distance is doubled from the surface, the proportion of the scattered light collected is quartered. The results agreed with this hypothesis, meaning that acquisition times will increase proportionally to the square of the distance between the scattering surface and the detector.

Next, the angle of the detector to the scattering surface was investigated. The centre of the detector was placed at a range of angles, from 0 to 60° , equidistant from the scattering surface. The target and source were kept equal throughout. The results confirmed beliefs that compound angles between the detector and the scattering surface had no effect on the scattered power that was observed. This means that altering angle, when interrogating a scene, will not directly affect the acquisition time.

The position of the target within the scene was also of great interest. This would develop an understanding of the relationship between power observed on the detector and the target's distance and angle to the scattering points.

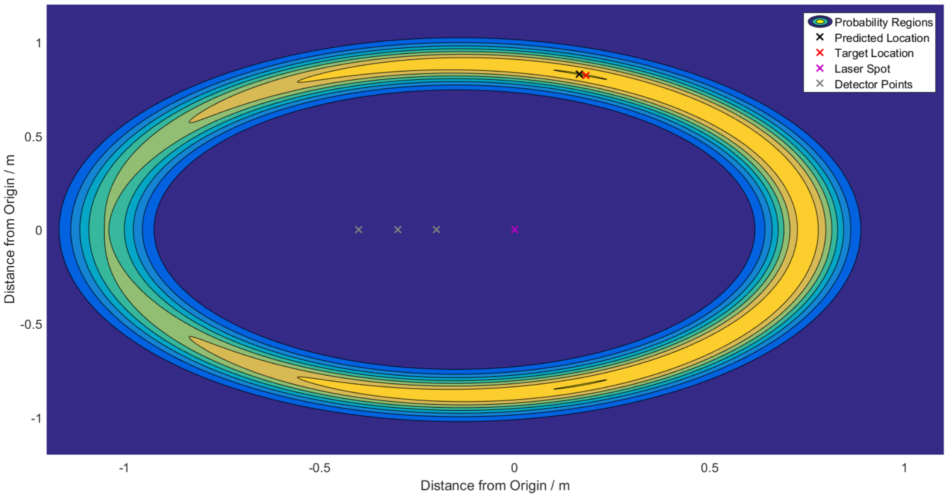

Figure 6.4: Map of results from the target position experiment. Black crosses indicate the actual position of the target's X-Y centroid. Red cross indicates whereabouts the model predicts the target to be. Finally the grey area surrounding the red cross indicates the 65% confidence interval for the target's location, as predicted by the model.

In this simulation the detector, scene and source were all kept constant and the target moved through the scene. Distances were increased from 100 to 1000 mm in 100 mm increments and angle from 0 to 60° , increasing by 5° . All permutations of these two ranges led to a total of 130 target positions being examined to determine the power-position relationship. The results for target angle to the scattering surface were similar to those of the detector angle investigation; there was no significant change in power across the range of angles observed. It must be noted that if the planar target were to rotate about its own vertical axis, this would change the effective aperture of the target and could lead to a reduction in power; this was not investigated, however as the three dimensional nature of real targets would negate this somewhat.

Testing the target's distance from the scattering surface did, however, confirm that power observed from the target decreases proportional to the fourth root of the increase in distance. For example, a factor of ten increase in the distance from the scattering surface would decrease the power reflected by roughly a thousand fold, greatly increasing acquisition time. It was concluded then, that this distance would be dominant in determining the possibility of positioning a target within a scene. It was also noted that maximum time-of-flight within the scene must be taken into consideration. Above a certain distance from the scattering surface, the light from the first laser pulse would still be in flight after the second fired. This would make it impossible to take meaningful readings, as the background response of the second pulse would obfuscate the distant target signal from the first.

From these results it was also possible to test the positioning calculation, as performed by MATLAB. Figure 6.4 shows the results of one such calculation, where the true position in red and the predicted in black are both within the area of highest confidence. Three detectors were used to calculate these positions, with a limited number of rays traced the confidence bands are quite broad but positioning can still be accomplished to centimetre precision.

Finally the effects of varying target size was considered. In the model, the planar objects used as targets were decreased in size geometrically and kept in the same position. The power response observed at the detector decreased proportionally to the size of the target. The target size also limits the precision with which its position can be determined. Larger targets reflect more of the laser pulse from a greater number of possible locations within the scene, this serves to give a stronger return signal at the cost of lower spatial precision. Overall, this investigation demonstrates the importance of target size for this technique and a means of calculating the variance in acquisition times as this parameter changes.

This project delivered a tool that could closely simulate the response observed by the positioning system proposed by Gariepy et al [153, 2, 3]. The benefits of this model allow the user to quickly test and investigate scenarios that are otherwise difficult, costly or impossible to replicate in the laboratory. It also allows the user to gain a good understanding of some of the limitations this approach faces. Combined with the more advanced optical modelling capabilities of Zemax, this would also afford the user a means of testing different optics configurations to maximise the system's performance. The work outlined in this chapter was also presented at the conference *Advanced Optics for Defense Applications: UV through LWIR - Baltimore, United States* in 2016 [1].

6.4 Amplified Spontaneous Emission

The goal of this project was to develop a modelling tool capable of predicting the effect of gain limiting phenomena in high energy, solid-state, laser systems. In high gain active media, spontaneous emission can lead to detrimental effects that serve to deplete the gain available for useful laser action. Spontaneously emitted rays, that do not couple to the laser cavity mode, can still experience gain when the gain-length product is above a certain threshold. This effect is known as amplified spontaneous emission, occuring when one spontaneously emitted photon travels far enough through the active medium to stimulate emission of another photon. This is a straightforward calculation when considering laser media with homogenous gain profiles and simple photon paths, where analytical solutions are simple. For the majority of real systems, however, this quickly becomes non-trivial and only coarse estimates of these effects can be made.

One particular problem with calculating the onset of ASE in real laser systems comes from the complexity of the pump absorption profile. For many pumping schemes, perfectly homogenous absorption profiles are not achiveable or desirable. Many absorption profiles can not be accurately represented by an analytical function. Figure 6.5 shows the transverse pump absorption profile for a typical solid-state laser system. Here the cylindrical rod is pumped by a 3kW stack of 810 nm diodes, leading to the striped peaks in the absorption profile.

Determining the maximum photon path length within the gain medium also becomes non-trivial, due to the possibility of total internal reflection within the laser rod. Some spontaneously emitted rays will be reflected at the surface boundary with the rod, increasing their path length through the gain medium. A proportion of these reflected rays will end up in closed paths, continuously reflecting around the inside of the laser rod. If these

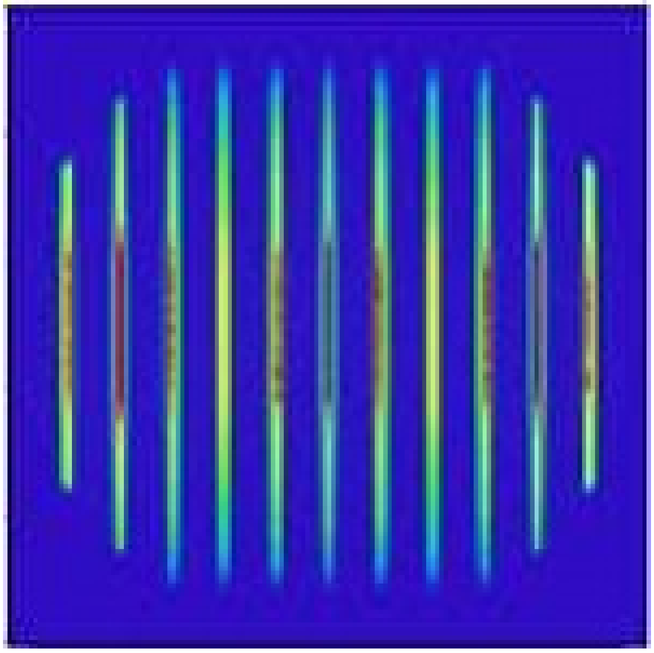

Figure 6.5: The pump absorption profile of a typical, solid-state, laser rod.

rays experience enough gain to counteract the losses they see at each reflection, their likelihood of exceeding the gain-length threshold for ASE is greatly increased. This leads to parasitic lasing wherein stimulated emission, not coupled to the useful laser output, limits laser performance by depleting the available energy. This is a concern for high power Q-switched lasers.

Predicting the threshold pump power at which these effects occur for any given system is difficult for all but the simplest of cases. This project sought to develop, accurate, time-resolved models of the gain evolution, that considered these gain depleting effects.

6.4.1 Approach

To accurately model the shape of complex pump profiles, such as that shown in fig. 6.5, a finite element method was used. The gain medium was divided into equally sized voxels, the three-dimensional equivalent of pixels. Each of these voxels were then treated as separate cells of active media. Their absorption was determined using the ray-tracing program Zemax. This allows for an accurate representation of the pump absorption profile, by first tracing pump-rays through a model of the system and then calculating the absoprtion of those rays as they pass through each of the voxels within the gain medium. As mentioned previously, ray-tracing has been used throughout this thesis, to numerically model the paths of rays in complex laser systems.

With an accurate absorption profile, it is then possible to calculate the gain distribution throughout the laser rod. This is achieved utilising the equations set out in section 5.2.5. The pumping period was also split into discrete time intervals, for each time-step, the pumping input power, cumulative stored energy, and fluoresccence power could then be calculated.

The fluoresccence output power from each of the voxels was then used to determine the spontaneous emission rays from the laser rod. A weighted pseudo-random method was used to determine the fluorescence ray distribution. Due to the limitations of current computational power, it would not be feasable to trace every spontaneously emitted photon.

For most of the work in this project only 1×10^4 rays were traced, a typical laser model, with 51^3 voxels and 1×10^4 rays over 200 time steps, would take approximately 2.5 hours to complete. These numbers were determined empirically to give a good balance between speed and accuracy.

This, far smaller, number of spontaneously emitted rays were then traced through the laser rod. The total fluorescence power, from every voxel in the laser rod, was then totalled and divided evenly between each of the rays. The rays directions at origin were determined randomly by generating random directional cosines in the three cardinal axes for the laser system model. Distribution of the ray origin points, though, was determined using a pseudo-random weighted process. The total fluoresccence power of each voxel was taken into consideration, for each ray voxels were selected at random and their normalised fluoresccence power output was used as a probability for selection for that rays origin point, selecting a new random voxel until successful. This lead to a three-dimensional fluoresccence distribution that more closely resembled that of an actual laser rod. The alternative approach, purely random voxel selection for each ray, resulted in a fluoresccence distribution that was more spread throughout the rod and disproportionately favoured voxels that saw little or no absorption.

The paths of these spontaneously emitted rays were then traced through the laser model, again using Zemax. This allows for accurate calculation of their losses as they travel through the system and interact with surface boundaries.

With these fluoresccence ray-paths it is then possible to caclulate their interactions with the gain in each of the voxels. Each ray was treated sequentially, the gain interactions along a single ray-path were caclulated from origin to termination, before calculating the next ray. Whenever a ray passed through a voxel, the gain experienced by the ray, according to the path length through that voxel, was calculated. The ray's energy entering the voxel, multiplied by the gain was then used to determine the ray's exit energy (as it left the voxel). The energy extracted was limited to the voxel's stored energy to ensure a ray could not exit with more energy than was available. For example, if a 1 μJ ray saw a gain coefficient of 1.5 passing through a voxel that only had 0.4 μJ of stored energy, then the ray would exit with only 1.4 μJ, extracting the maxiumum available energy.

After each ray's interaction with a voxel has occured, the energy remaining stored within the voxel was updated. This way, if another ray interacted with the same voxel, it would experience lower gain and less energy available for extraction.

When this process was complete for every ray, the next time-step would be calculated. Another time-step's worth of pump energy would be added to the voxels and the gain distribution updated. This time-resolved approach allows for analysis of the temporal evolution of the gain in the laser medium, highlighting the onset of gain limiting effects. One advantage of this approach meant that virtual detectors could be used, that recorded parameters not easily measured in a physical system, without interaction with the rays. Recording how much available power was extracted at each time-step, by spontaneously emitted rays, enabled direct comparison of the energy lost to parasitic effects.

6.4.2 Results

The accuracy of the model was verified by simulating results from the literature [59]. Models of the lasers used in the 2010 paper by Huß *et. al* were made, cylindrical core-doped 2% Nd:YAG rods, with different claddings, YAG and 1% Sm:YAG. The paper demonstrates, experimentally, how the Sm:YAG cladding suppresses the onset of ASE and parasitic effects, by limiting the path lengths of rays that travel transverse to the rod. This means less gain is lost to spontaneous emission as gain builds up within the laser rod, during the pumping period.

The pumping and absorption profiles for these two lasers were then simulated using Zemax and used in the model. The rods were pumped for 200 μs at seven pump powers, including; 11, 77, 145, 276, 394, 500 and 592 W. The model was then set to calculate the change in gain in time-steps of 2.5 μs.

As can be seen in section 5.4, the model results closely resembled those of the experiment, indicating the accuracy of this modelling approach. By keeping track of small, discrete volumes of the gain medium, and the processes by which they absorb or lose energy over time, it becomes possible to model the effect of high gain phenomena, in complicated laser systems.

This approach also allows the user to analyse these phenomena in any laser capable of being designed in Zemax. With the diverse modelling capabilities of Zemax, and the ability to import CAD objects into the software, novel laser pumping schemes can be tested and investigated easily. This was demonstrated by analysing the high-gain behaviour of two other rod geometries, in the same pumping setup as those tested in the Hußpaper.

A square and a triangular faced rod were pumped for 200 μs at a range of nine pump powers; 0.5, 2, 10, 50, 250, 500, 1000, 2000 and 4000 W. This was to see how they performed over a large range of pumping powers, all of which would have practical applications, and to investigate the pumping threshold at which high-gain phenomena, such as ASE, occur. Due to tighter angles between each of the rods edges, when compared to the square faced rod, the subset of angles that will see multiple reflections from to total internal reflection, is reduced. Both pump absorption profiles were compared and show significant differences, as seen in figs. 5.17 and 5.18.

The temporal evolution of fluorescence power on the detector and the gain power extracted by all rays, during the pump period, was compared for both rods. The fluorescence results in figs. 5.17 and 5.18 show that, at 500 W and above, the rate of fluorescence on the detector, for the square faced rod, begins to climb steeply and then plateau. As pump power increases this effect becomes more pronounced, the plateau begin reached at approximately 100 and 45 μs for the 500 W and 4000 W models respectively. The maximal fluorescence power achieved at this plateau does increase with pumping power, but not proportionally to the increased input, with a maximum value of approximately 8 W seen for the 4000 W pump model.

The triangular rod does not exhibit this plateau behaviour, at any pumping power. For each of the powers investigated, the fluorescence from the triangular rod increases steadily throughout the entire pumping period. The maximum power achieved on the detector is also lower at all pump powers except the 4000 W model, where the square rod reaches 8 W and the triangular 9 W. This would mean that for longer pumping periods at the highest pump intensities, anything greater than 160 μs, would mean more pump energy lost when compared to the square rod. For the 200 μs period examined, however, total fluorescence energy (area under the graph) for the 4000 W model, is still lower for the triangular rod compared to the square. The latter reaches a value of approximately 6 W in the first 25 μs, 75 μs earlier than the triangular rod.

It was also possible to track the total gain power extracted from all voxels at each time-step, using virtual detectors, something not directly measurable in the lab. Figures 5.19 and 5.20 show theses results for the 200 μs pumping period. The square rod, again, appears to reach a plateau as the pumping power increases. The triangular rod also plateaus and does so far more quickly, within the first 20 μs, at all pump powers. This shows that the rate of gain power extraction reaches a limit quickly and amplified spontaneous rays are not depleting increasingly larger proportions of the available gain as the pump period goes on. The triangular rod also plateaus at much lower maximal values than those achieved by the square, rod for the same pumping powers. This indicates that the rate of depletion is much lower for the triangular rod in all cases, except for the shortest pump periods.

These results highlight the potential advantages of using triangular faced rods in Q-switched systems, where high-gain is required to build in the laser rod before the onset of lasing action. Less energy is lost to the fluorescence and spontaneous emission during the pumping period, leaving more available for useful lasing. The innate geometry, limiting long internal path lengths, enables higher energy storage efficiency when compared to square rods. These results are promising and warrant further physical investigation in a laboratory.

Overall these results highlight the usefulness of the model presented here. A proposed pumping scheme can be quickly investigated and compared to existing systems without the monetary or time costs associated with fabricating the physical components. This enables a fail-fast and cheaply, agile, iterative approach to high-gain laser design that could not otherwise be realised.

6.5 Future Work

Each of the projects presented throughout this thesis present multiple opportunities for continued development. Outlined below are some recommendations for how best to extract additional value from the work presented herein. As each of the three project in this thesis exist within their own domains, their potential for extension will be discuss separately.

6.5.1 Non-Line-of-Sight LIDAR

The challenges associated with the non-line-of-sight LIDAR system required a tool that could replicate the time-of-flight response of complex scenes. The model, although able to produce results analogous to those seen in experiment, samples the scenes with many orders of magnitude fewer rays than would be necessary to thoroughly interrogate the potential ray-paths. This is mainly due to computational limits associated with the underlying ray-tracing engine, Zemax. Due to the nature of this model, results from multiple ray-traces of the same scene could be combined to produce a larger, more thorough data set. It would be possible to run multiple versions of Zemax, over a cluster of computers in parallel, or possibly automate many samples of the model on a single machine. Scaling up the sampling size would begin to capture far larger proportions of the potential paths and a more representative result, from any given scene, would be possible. This could easily be achieved using automation of both the sampling and processing, using MATLAB and the tools already presented in this thesis. This paper [145] demonstrates a similar approach used to gather sufficient ray-tracing samples of a complex optical system, using MATLAB and Zemax in server-client configuration. One of the key problems with collecting much larger samples would be in the large file sizes that are generated in the intermediate steps of processing, when data is transferred from MATLAB to Zemax. Careful management of the processing phase could eliminate the need for these large files to be saved, otherwise large data storage platforms may be necessary.

Another interesting avenue for development with this project would include tracking targets as they move through a scene. Reconstructing the positions of a target as it moves through a scene would require multiple ray-traces to be performed at time-intervals. This type of analysis would also allow for the investigation into stationary background subtraction techniques, isolating the moving target across the time separated ray-traces.

Either of these would serve to enhance the capabilities of this tool and allow the user to gain a more comprehensive understanding of the underlying processes in this positioning technique.

6.5.2 Optimisation and Tolerancing

The optimisation work presented in this thesis offers several opportunities for further development. Much in the same way that the non-line-of-sight LIDAR modelling could benefit from more computational power, so could this system. As the individual ray-traces for each of the configurations are entirely independent of one another, they could all be run in parallel. Again a system whereby the ray-traces to be performed are calculated and scheduled to run across numerous machines could greatly speed up this process.

Finding patterns in the complex, highly dimensional data can also be quite time-consuming, even for the most experienced users. Some of this work, searching over large result data sets, could be done automatically using machine learning techniques. These approaches allow computational models to be formed algorithmically based on the data at hand, without human intervention. Machine learning is capable of outperforming human experts in a number of fields. By utilising these techniques, complex patterns could be identified sooner and more nuanced questions could be answered about the underlying behaviour of the optical system.

This optimisation model could also greatly benefit from the inclusion of some diffraction modelling. All of the calculations regarding beam-shape and size are limited to the output of ray-tracing equations. Accurate diffraction and cavity modelling would give a more detailed understanding of the stable modes that exist within a laser cavity. This would greatly enhance those metrics and allow the user to better understand the effect that component changes had on the final model. Beam waists and stable modes play an important part in determining the extraction efficiency of a laser system. Some initial work was carried out using the software GLAD [154]. A working interface between the diffraction modelling software and MATLAB was created, enabling control and automation of the diffraction modelling, and transfer of data from Zemax. Although this was a success, there was not sufficient time to fully understand the program and how to make best use of it for these applications. Continuing this work could also benefit the ASE modelling, as any stable modes within a system could contribute to parasitic processes, limiting gain.

6.5.3 Amplified Spontaneous Emission

The ASE modelling project was very successful in developing a tool that could predict the onset of gain-limiting phenomena. One key area for improvement for this model would be the speed at which it runs, and the efficiency of the computations. Greater accuracy and precision can be achieved with shorter time intervals and more rays, but these both greatly increase the computational time. The existing code base was developed over a long period by a single person, who did not have a strong background in computer science or programming. Due to the serial nature of the gain depletion modelling, how the voxel interactions of one ray limit the gain available to those calculated after it, there are likely many algorithmic improvements that could be made to the code. These dependent interactions from ray-to-ray also mean the order in which rays gain calculations are done is significant. Improvements could be made to better handle the ordering of these rays, not just sampling them randomly at each time-interval. It would be possible to calculate the time at which each of ray-voxel interface collisions occur, for all rays in a time-step, and calculate each section in chronological order.

Large intermediate file sizes and slow hard-disk reading times also limited the number of rays that could be used in the model. Being able to trace more rays would greatly increase the likelihood that long, gain limiting, paths are found for any system. Improvements between the data transfer process between Zemax and MATLAB may improve this. Tracing a new set of rays for each time-step would also be a possible avenue for development. As these intermediate files are so large, only a single ray-trace is used for all time-steps. Retracing the rays, with their distribution recalculated to account for the previous time-step's depletion, might see a spatial change in the fluorescence as the pumping period goes on.

6.6 Final Thoughts

The tools developed throughout this project have had tangible industrial impact. The optimisation and tolerancing tools were used to make recommendations in the design stages for the system presented by Lee *et al.* [99]. The non-line-of-sigh LIDAR models were used to analyse and understand some of the underlying relationships that governed this novel technology, and the author of this thesis presented these results at the conference *Advanced Optics for Defense Applications: UV through LWIR - Baltimore, United States* in 2016 [1]. The amplified spontaneous emission models confrimed expectations around a novel rod geometry for suppression in high-power solid-state gain medium that could see benefits for passively Q-switched systems. These recommendations are the first steps to securing further investment to progress these designs into the physical domain and can be used to explore further, at significantly reduced costs.

Overall this project has produced a number of tools to improve and benefit those developing solid-state laser systems. It sought to produce software that aligned with existing work-flows and enhance, instead of replace, tools that were already in use. A key consideration throughout has been the reuse of existing software, so that minimal time is spent redeveloping capabilities that already exist. It has met these goals and the models have been utilised, with great success, since their inception.

Bibliography

[1] Adam Sroka et al. "Time-resolved non-sequential ray-tracing modelling of non-line-of-sight picosecond pulse LIDAR". In: *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 9822. 2016, p. 98220L. ISBN: 9781510600638. DOI: `10.1117/12.2223589`.

[2] Genevieve Gariepy et al. "Single-photon sensitive light-in-fight imaging". In: *Nature Communications* 6.1 (Dec. 2015), p. 6021. ISSN: 20411723. DOI: `10.1038/ncomms7021`.

[3] Genevieve Gariepy et al. "Detection and tracking of moving objects hidden from view". In: *Nature Photonics* 10.1 (Jan. 2016), pp. 23–26. ISSN: 17494893. DOI: `10.1038/nphoton.2015.234`.

[4] Genevieve Gariepy et al. "Picosecond time-resolved imaging using SPAD cameras". In: ed. by Keith L. Lewis and Richard C. Hollins. Oct. 2016, 99920N. ISBN: 9781510603882. DOI: `10.1117/12.2241184`.

[5] *MATLAB*. 2020. URL: `http://www.mathworks.com`.

[6] ZEMAX. *Leading Optical Product Design Software for Engineering Teams - Zemax*. URL: `https://www.zemax.com/` (visited on 02/09/2020).

[7] Benjamin Gallinet, Jérémy Butet, and Olivier J.F. Martin. "Numerical methods for nanophotonics: Standard problems and future challenges". In: *Laser and Photonics Reviews* 9.6 (Nov. 2015), pp. 577–603. ISSN: 18638899. DOI: `10.1002/lpor.201500122`.

[8] Herbert Walther et al. "Cavity quantum electrodynamics". In: *Reports on Progress in Physics* 69.5 (2006), pp. 1325–1382. ISSN: 00344885. DOI: `10.1088/0034-4885/69/5/R02`.

[9] Stefan Scheel and Stefan Yoshi Buhmann. "Macroscopic quantum electrodynamics - Concepts and applications". In: *Acta Physica Slovaca* 58.5 (2008), pp. 675–809. ISSN: 03230465. DOI: `10.2478/v10155-010-0092-x`.

[10] MR Foreman and Peter Török. "Computational methods in vectorial imaging". In: *Journal of Modern Optics* 58.5-6 (Mar. 2011), pp. 339–364. ISSN: 0950-0340. DOI: `10.1080/09500340.2010.525668`.

[11] GH Spencer and M Murty. "General ray-tracing procedure". In: *JOSA* 52.1951 (1962), pp. 672–678. DOI: `10.1364/JOSA.52.000672`.

[12] F. Michael Kahnert, Jakob J. Stamnes, and Knut Stamnes. “Surface-integral formulation for electromagnetic scattering in spheroidal coordinates”. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 77.1 (Feb. 2003), pp. 61–78. ISSN: 00224073. DOI: `10.1016/S0022-4073(02)00075-4`.

[13] Nicholas Metropolis and S. Ulam. “The Monte Carlo Method”. In: *Journal of the American Statistical Association* 44.247 (Sept. 1949), p. 335. ISSN: 01621459. DOI: `10.2307/2280232`.

[14] Nicholas Metropolis et al. “Equation of state calculations by fast computing machines”. In: *The Journal of Chemical Physics* 21.6 (June 1953), pp. 1087–1092. ISSN: 00219606. DOI: `10.1063/1.1699114`.

[15] Dirk P. Kroese et al. “Why the Monte Carlo method is so important today”. In: *Wiley Interdisciplinary Reviews: Computational Statistics* 6.6 (Nov. 2014), pp. 386–392. ISSN: 19390068. DOI: `10.1002/wics.1314`.

[16] N. Etemadi. “An elementary proof of the strong law of large numbers”. In: *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 55.1 (Feb. 1981), pp. 119–122. ISSN: 00443719. DOI: `10.1007/BF01013465`.

[17] Arthur B Yeh. “A Modern Introduction to Probability and Statistics”. In: *Technometrics* 49.3 (Aug. 2007), pp. 359–359. ISSN: 0040-1706. DOI: `10.1198/tech.2007.s502`.

[18] J. F. Ramaley. “Buffon’s Noodle Problem”. In: *The American Mathematical Monthly* 76.8 (Oct. 1969), p. 916. ISSN: 00029890. DOI: `10.2307/2317945`.

[19] Arakaparampil M. Mathai. *An Introduction to Geometrical Probability: Distributional Aspects with Applications*. Newark: Gordon & Breach, Science Pub, 1999, p. 576. ISBN: 9789056996819.

[20] Malvin H. Kalos and Paula A. Whitlock. *Monte Carlo Methods: Second Edition*. Wiley-VCH, July 2009, pp. 1–203. ISBN: 9783527407606. DOI: `10.1002/9783527626212`.

[21] H. KAHN. “Random sampling (Monte Carlo) techniques in neutron attenuation problems–I.” In: *Nucleonics* 6.5 (May 1950), 27, passim. ISSN: 0096-6207. URL: `http://www.ncbi.nlm.nih.gov/pubmed/15412767%20https://europepmc.org/abstract/med/15423762`.

[22] K. Binder et al. “Monte Carlo Simulation in Statistical Physics”. In: *Computers in Physics* 7.2 (1993), p. 156. ISSN: 08941866. DOI: `10.1063/1.4823159`.

[23] Gilles Celeux. “Bayesian Inference and Markov Chain Monte Carlo Methods”. In: *Construction Reliability: Safety, Variability and Sustainability* (2013), pp. 207–226. DOI: `10.1002/9781118601099.ch11`.

[24] Siddhartha Chib and Edward Greenberg. “Markov Chain Monte Carlo Simulation Methods in Econometrics”. In: *Econometric Theory* 12.3 (Aug. 1996), pp. 409–431. ISSN: 0266-4666. DOI: `10.1017/S0266466600006794`.

[25] Radford M Neal. “Probabilistic Inference Using Markov Chain Monte Carlo Methods Acknowledgements”. In: *Intelligence* 45.September (1993), p. 144. ISSN: 10959572. DOI: `10.1016/j.neuroimage.2009.01.023`.

[26] Nikolai V. Voshchinnikov. “Electromagnetic scattering by homogeneous and coated spheroids: Calculations using the separation of variables method”. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 55.5 (1996), pp. 627–636. ISSN: 00224073. DOI: `10.1016/0022-4073(96)00006-4`.

[27] Tom Rother. “General aspects of solving helmholtz’s equation underlying eigenvalue and scattering problems in electromagnetic wave theory”. In: *Journal of Electromagnetic Waves and Applications* 13.7 (1999), pp. 867–888. ISSN: 15693937. DOI: `10.1163/156939399X00330`.

[28] F. Michael Schulz, Knut Stamnes, and Jakob J. Stamnes. “Scattering of electromagnetic waves by spheroidal particles: a novel approach exploiting the T matrix computed in spheroidal coordinates”. In: *Applied Optics* 37.33 (Nov. 1998), p. 7875. ISSN: 0003-6935. DOI: `10.1364/ao.37.007875`.

[29] Stephen D. Gedney. “Introduction to the Finite-Difference Time-Domain (FDTD)) method for electromagnetics”. In: *Synthesis Lectures on Computational Electromagnetics* 27 (Jan. 2011), pp. 1–250. ISSN: 19321252. DOI: `10.2200/S00316ED1V01Y201012CEM027`.

[30] J. F. Lee. “Numerical solutions of TM scattering using an obliquely Cartesian finite difference time domain algorithm”. In: *IEE Proceedings H: Microwaves, Antennas and Propagation* 140.1 (1993), pp. 23–28. ISSN: 0950107X. DOI: `10.1049/ip-h-2.1993.0004`.

[31] A. Taflove and M.E. Brodwin. “Numerical Solution of Steady-State Electromagnetic Scattering Problems Using the Time-Dependent Maxwell’s Equations”. In: *IEEE Transactions on Microwave Theory and Techniques* 23.8 (Aug. 1975), pp. 623–630. ISSN: 0018-9480. DOI: `10.1109/TMTT.1975.1128640`.

[32] Michael A. Morgan and Kenneth K. Mei. “Finite-Element Computation of Scattering by Inhomogeneous Penetrable Bodies of Revolution”. In: *IEEE Transactions on Antennas and Propagation* 27.2 (1979), pp. 202–214. ISSN: 15582221. DOI: `10.1109/TAP.1979.1142065`.

[33] M. A. Morgan et al. “Finite element-boundary integral formulation for electromagnetic scattering”. In: *Wave Motion* 6.1 (1984), pp. 91–103. ISSN: 01652125. DOI: `10.1016/0165-2125(84)90025-8`.

[34] R. Coccioli et al. “Finite-element methods in microwaves: A selected bibliography”. In: *IEEE Antennas and Propagation Magazine* 38.6 (Dec. 1996), pp. 34–47. ISSN: 10459243. DOI: `10.1109/74.556518`.

[35] T. Wriedt and A. Doicu. “Formulations of the extended boundary condition method for three-dimensional scattering using the method of discrete sources”. In: *Journal of Modern Optics* 45.1 (Jan. 1998), pp. 199–213. ISSN: 13623044. DOI: `10.1080/09500349808231681`.

[36] P. C. Waterman. “Matrix Formulation of Electromagnetic Scattering”. In: *Proceedings of the IEEE* 53.8 (1965), pp. 805–812. ISSN: 15582256. DOI: `10.1109/PROC.1965.4058`.

[37] Michael I. Mishchenko, Larry D. Travis, and Daniel W. Mackowski. “T-matrix computations of light scattering by nonspherical particles: A review”. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 55.5 (May 1996), pp. 535–575. ISSN: 00224073. DOI: `10.1016/0022-4073(96)00002-7`.

[38] A. Lakhtakia and G.W. Mulholland. “On two numerical techniques for light scattering by dielectric agglomerated structures”. In: *Journal of Research of the National Institute of Standards and Technology* 98.6 (Nov. 1993), p. 699. ISSN: 1044-677X. DOI: `10.6028/jres.098.046`.

[39] R. H. Dicke. “Coherence in Spontaneous Radiation Processes”. In: *Physical Review* 93.1 (Jan. 1954), pp. 99–110. ISSN: 0031-899X. DOI: `10.1103/PhysRev.93.99`.

[40] Luc Vinet and Alexei Zhedanov. “A ’missing’ family of classical orthogonal polynomials”. In: *Journal of Physics A: Mathematical and Theoretical* 44.8 (Nov. 2011). Ed. by Intergovernmental Panel on Climate Change, pp. 1–30. ISSN: 17518113. DOI: `10.1088/1751-8113/44/8/085201`.

[41] R. Bonifacio, P. Schwendimann, and Fritz Haake. “Quantum Statistical Theory of Superradiance. I”. In: *Physical Review A* 4.1 (July 1971), pp. 302–313. ISSN: 0556-2791. DOI: `10.1103/PhysRevA.4.302`.

[42] L. Allen and G.I. Peters. “Superradiance, coherence brightening and amplified spontaneous emission”. In: *Physics Letters A* 31.3 (Feb. 1970), pp. 95–96. ISSN: 03759601. DOI: `10.1016/0375-9601(70)90168-4`.

[43] L. Allen and G. I. Peters. “Amplified spontaneous emission III. Intensity and saturation”. In: *Journal of Physics A: General Physics* 4.4 (July 1971), pp. 564–573. ISSN: 0022-3689. DOI: `10.1088/0305-4470/4/4/020`.

[44] H. Kogelnik and A. Yariv. “Considerations of Noise and Schemes for Its Reduction in Laser Amplifiers”. English. In: *Proceedings of the IEEE* 52.2 (Feb. 1964), pp. 165–172. ISSN: 15582256. DOI: `10.1109/PROC.1964.2805`.

[45] W. T. Silfvast and J. S. Deech. “Six dB/cm single-pass gain at7229 Å in lead vapor”. en. In: *Applied Physics Letters* 11.3 (Aug. 1967), pp. 97–99. ISSN: 0003-6951. DOI: `10.1063/1.1755052`.

[46] DA Leonard and WR Zinky. “Coherence Properties of the Superradiant 5401 Å Pulsed Neon Laser”. en. In: *Applied Physics Letters* 12.4 (Feb. 1968), p. 113. ISSN: 00036951. DOI: `10.1063/1.1651915`.

[47] Lewi Tonks. “Photon Avalanches from a Population Inversion”. en. In: *Journal of Applied Physics* 35.4 (Apr. 1964), p. 1134. ISSN: 00218979. DOI: `10.1063/1.1713580`.

[48] Amnon Yariv and RCC Leite. “Super radiant narrowing in fluorescence radiation of inverted populations”. en. In: *Journal of Applied Physics* 34.11 (Nov. 1963), p. 3410. ISSN: 00218979. DOI: `10.1063/1.1729206`.

[49] A. D. White, E. I. Gordon, and J. D. Rigden. “Output power of the 6328-Å gas maser”. In: *Applied Physics Letters* 2.5 (1963), pp. 91–93. ISSN: 00036951. DOI: `10.1063/1.1753793`.

[50] G. I. Peters and L. Allen. “Amplified spontaneous emission I. The threshold condition”. In: *Journal of Physics A: General Physics* 4.2 (Mar. 1971), pp. 238–243. ISSN: 0022-3689. DOI: `10.1088/0305-4470/4/2/009`.

[51] L. Allen and G. I. Peters. "Amplified Spontaneous Emission and External Signal Amplification in an Inverted Medium". In: *Physical Review A* 8.4 (Oct. 1973), pp. 2031–2047. ISSN: 0556-2791. DOI: `10.1103/PhysRevA.8.2031`.

[52] L. Allen and G. I. Peters. "Amplified spontaneous emission. II. The connection with laser theory". In: *Journal of Physics A: General Physics* 4.3 (May 1971), pp. 377–381. ISSN: 0022-3689. DOI: `10.1088/0305-4470/4/3/015`.

[53] Stig Stenholm and Willis E. Lamb. "Semiclassical Theory of a High-Intensity Laser". In: *Physical Review* 181.2 (May 1969), pp. 618–635. ISSN: 0031-899X. DOI: `10.1103/PhysRev.181.618`.

[54] D. Albach, J.-C. Chanteloup, and G. le. Touzé. "Influence of ASE on the gain distribution in large size, high gain Yb
3+:YAG slabs". In: *Optics Express* 17.5 (Mar. 2009), p. 3792. ISSN: 1094-4087. DOI: `10.1364/oe.17.003792`.

[55] Yuen H. Tsang, Ashraf F. El-Sherif, and Terence A. King. "Broadband amplified spontaneous emission fibre source near 2 μm using resonant in-band pumping". In: *Journal of Modern Optics* 52.1 (Jan. 2005), pp. 109–118. ISSN: 0950-0340. DOI: `10.1080/09500340410001703841`.

[56] Q Lü and S Dong. "Numerical and experimental investigation on ASE effects in high-power slab amplifiers". In: *Optics & Laser Technology* (1993). URL: `http://www.sciencedirect.com/science/article/pii/003039929390019C`.

[57] H. Yagi et al. "Y3Al5O12 ceramic absorbers for the suppression of parasitic oscillation in high-power Nd:YAG lasers". In: *Journal of Luminescence* 121.1 (Nov. 2006), pp. 88–94. ISSN: 00222313. DOI: `10.1016/j.jlumin.2005.10.006`.

[58] LA Hackel, TF Soules, and SN Fochs. "Gain media edge treatment to suppress amplified spontaneous emission in a high power laser". In: *US Patent 7,463,660* (2011). URL: `http://www.google.com/patents?hl=en%7B%5C&%7Dlr=%7B%5C&%7Dvid=USPAT7894496%7B%5C&%7Did=sxYpAQAAEBAJ%7B%5C&%7Doi=fnd%7B%5C&%7Ddq=Gain+media+edge+treatment+to+suppress+amplified+spontaneous+emission+in+a+high+power+laser%7B%5C&%7Dprintsec=abstract`.

[59] Rafael Huß et al. "Suppression of parasitic oscillations in a core-doped ceramic Nd:YAG laser by Sm:YAG cladding". In: *Optics Express* 18.12 (June 2010), p. 13094. ISSN: 1094-4087. DOI: `10.1364/OE.18.013094`.

[60] Sebastien Laux et al. "Suppression of parasitic lasing in high energy, high repetition rate Ti:sapphire laser amplifiers". In: *Optics Letters* 37.11 (June 2012), p. 1913. ISSN: 0146-9592. DOI: `10.1364/ol.37.001913`.

[61] Klaus Ertel et al. "ASE suppression in a high energy Titanium sapphire amplifier". In: *Optics Express* 16.11 (May 2008), p. 8039. ISSN: 1094-4087. DOI: `10.1364/OE.16.008039`.

[62] KV Galoyan and KM Khachatryan. “Influence of the surface treatment of a cylindrical yttrium aluminum garnet crystal on laser energy characteristics”. In: *Quantum Electronics* 4.2 (Feb. 1974), pp. 260–261. ISSN: 0049-1748. DOI: `10.1070/QE1974v004n02ABEH006718`.

[63] Arun Kumar Sridharan et al. “Zigzag slabs for solid-state laser amplifiers: batch fabrication and parasitic oscillation suppression”. In: *Applied Optics* 45.14 (May 2006), p. 3340. ISSN: 0003-6935. DOI: `10.1364/AO.45.003340`.

[64] W. F. Krupke et al. “Spectroscopic, optical, and thermomechanical properties of neodymium- and chromium-doped gadolinium scandium gallium garnet”. In: *Journal of the Optical Society of America B* 3.1 (1986), p. 102. ISSN: 0740-3224. DOI: `10.1364/josab.3.000102`.

[65] Amos Hardy and R. Oron. “Signal amplification in strongly pumped fiber amplifiers”. In: *IEEE Journal of Quantum Electronics* 33.3 (Mar. 1997), pp. 307–313. ISSN: 00189197. DOI: `10.1109/3.555997`.

[66] Francesco Matera and Marina Settembre. “Comparison of the performance of optically amplified transmission systems”. In: *Journal of Lightwave Technology* 14.1 (1996), pp. 1–11. ISSN: 07338724. DOI: `10.1109/50.476131`.

[67] J. P. Gordon and L. F. Mollenauer. “Phase noise in photonic communications systems using linear amplifiers”. In: *Optics Letters* 15.23 (Dec. 1990), p. 1351. ISSN: 0146-9592. DOI: `10.1364/ol.15.001351`.

[68] F Matera and M Settembre. “Nonlinear compensation of chromatic dispersion for phase- and intensity-modulated signals in the presence of amplified spontaneous emission noise”. In: *Optics Letters* 19.16 (Aug. 1994), p. 1198. ISSN: 0146-9592. DOI: `10.1364/ol.19.001198`.

[69] Per Kylemark et al. “Noise Characteristics of Fiber Optical Parametric Amplifiers”. In: *Journal of Lightwave Technology* 22.2 (Feb. 2004), pp. 409–416. ISSN: 0733-8724. DOI: `10.1109/JLT.2003.822152`.

[70] J. H. Lee et al. “A review of the polarization-nulling technique for monitoring optical-signal-to-noise ratio in dynamic WDM networks”. In: *Journal of Lightwave Technology* 24.11 (Nov. 2006), pp. 4162–4171. ISSN: 07338724. DOI: `10.1109/JLT.2006.883120`.

[71] A Mecozzi et al. “Soliton transmission control”. In: *Optics Letters* 16.23 (Dec. 1991), p. 1841. ISSN: 0146-9592. DOI: `10.1364/ol.16.001841`.

[72] C. B. Olausson et al. “Amplification and ASE suppression in a polarization-maintaining ytterbium-doped all-solid photonic bandgap fibre”. In: *Optics Express* 16.18 (2008), p. 13657. ISSN: 1094-4087. DOI: `10.1364/oe.16.013657`.

[73] Alan J. Kemp, Gareth J. Valentine, and David Burns. “Progress towards high-power, high-brightness neodymium-based thin-disk lasers”. In: *Progress in Quantum Electronics* 28.6 (2004), pp. 305–344. ISSN: 00796727. DOI: `10.1016/j.pquantelec.2004.11.001`.

[74] Jochen Speiser. “Scaling of thin-disk lasers–influence of amplified spontaneous emission”. In: *Journal of the Optical Society of America B* 26.1 (Dec. 2008), p. 26. ISSN: 0740-3224. DOI: `10.1364/JOSAB.26.000026`.

[75] Dmitrii Kouznetsov et al. “Surface loss limit of the power scaling of a thin-disk laser”. In: *JOSA B* 23.6 (2006), p. 1074. ISSN: 0740-3224. DOI: `10.1364/JOSAB.23.001074`.

[76] Dmitrii Kouznetsov, JF Bisson, and Kenichi Ueda. “Scaling laws of disk lasers”. In: *Optical materials* 31.5 (Mar. 2009), pp. 754–759. ISSN: 09253467. DOI: `10.1016/j.optmat.2008.03.017`.

[77] Aldo Antognini et al. “Thin-disk Yb:YAG oscillator-amplifier laser, ASE and effective Yb: YAG lifetime”. In: *IEEE Journal of Quantum Electronics* 45.8 (2009), pp. 993–1005. ISSN: 00189197. DOI: `10.1109/JQE.2009.2014881`.

[78] Martin Gorjan, Thibault North, and Martin Rochette. “Model of the amplified spontaneous emission generation in thulium-doped silica fibers”. In: *JOSA B* 29.10 (Sept. 2012), p. 2886. ISSN: 0740-3224. DOI: `10.1364/JOSAB.29.002886`.

[79] N.P. Barnes and B.M. Walsh. “Amplified spontaneous emission-application to Nd:YAG lasers”. In: *IEEE Journal of Quantum Electronics* 35.1 (1999), pp. 101–109. ISSN: 00189197. DOI: `10.1109/3.737626`.

[80] J. B. Trenholme. “Fluorescence Amplification and Parasitic Oscillation Limitations in Disc Lasers”. In: *Naval Research Laboratory Memorandum Rep. 2480* July (1972). URL: `https://apps.dtic.mil/dtic/tr/fulltext/u2/747234.pdf`.

[81] Arthur N. Chester. “Gain Thresholds for Diffuse Parasitic Laser Modes”. In: *Applied Optics* 12.9 (Sept. 1973), p. 2139. ISSN: 0003-6935. DOI: `10.1364/ao.12.002139`.

[82] S. Guch. “8.5 Enhanced Parasitic Suppression in Large Aperture Disk Lasers Using Liquid Edge Claddings”. In: *IEEE Journal of Quantum Electronics* 11.9 (June 1975), pp. 852–853. ISSN: 15581713. DOI: `10.1109/JQE.1975.1069001`.

[83] J. H. Campbell et al. “Large-Aperture, High Damage-Threshold Optics for Beamlet”. In: *NASA STI* 5.1 (1995), pp. 1–10. DOI: `10.2172/88479`.

[84] L. Fornasiero et al. “Excited state absorption and stimulated emission of Nd3+ in crystals III: LaSc3(BO3)4, CaWO4, and YLiF4”. In: *Applied Physics B: Lasers and Optics* 68.1 (1999), pp. 67–72. ISSN: 09462171. DOI: `10.1007/s003400050587`.

[85] Yannick Guyot and Richard Moncorge. “Excited state absorption in the infrared emission domain of Nd3+doped Y3Al5O12, YLiF4,, and LaMgAl11O19”. In: *Journal of Applied Physics* 73.12 (1993), p. 8526. ISSN: 00218979. DOI: `10.1063/1.353381`.

[86] Y. Guyot et al. “Excited-state-absorption and upconversion studies of Nd3+-doped single crystals Y3Al5O12, YLiF4, and LaMgAl11O19”. In: *Physical Review B* 51.2 (Jan. 1995), pp. 784–799. ISSN: 0163-1829. DOI: `10.1103/PhysRevB.51.784`.

[87] T.Y. Fan. “Heat generation in Nd: YAG and Yb: YAG”. In: *Quantum Electronics, IEEE Journal of* 29.6 (June 1993), pp. 1457–1459. ISSN: 00189197. DOI: `10.1109/3.234394`.

[88] Osama Helal et al. “Different efficiency mechanisms inside solid-state laser: Calculation and experimental work”. In: *... and Microelectronics (ICOM), ...* 4 (2012), pp. 23–25.

[89] David C. Brown. “Heat, fluorescence, and stimulated-emission power densities and fractions in Nd:YAG”. In: *IEEE Journal of Quantum Electronics* 34.3 (1998), pp. 560–571. ISSN: 00189197. DOI: `10.1109/3.661467`.

[90] Peizhi Yang, Peizhen Deng, and Zhiwen Yin. “Concentration quenching in Yb:YAG”. In: *Journal of Luminescence* 97.1 (Apr. 2002), pp. 51–54. ISSN: 00222313. DOI: `10.1016/S0022-2313(01)00426-4`.

[91] David C Brown. “Power Densities and Fractions in Nd:YAG”. In: *IEEE Journal of Quantum Electronics* 34.3 (1998), pp. 560–572.

[92] P. Yankov and N. Saliyski. “Upconversion fluorescence emission of Er3+/Yb3+ codoped transparent glass ceramics: kinetics of the excited states”. In: *Seventeenth ...* 8770 (Mar. 2013). Ed. by Tanja N. Dreischuh and Albena T. Daskalova, pp. 87701I–87701I–7. DOI: `10.1117/12.2014268`.

[93] M. Pollnau et al. “Upconversion-induced heat generation and thermal lensing in Nd:YLF and Nd:YAG”. In: *Physical Review B - Condensed Matter and Materials Physics* 58.24 (Dec. 1998), pp. 16076–16092. ISSN: 1550235X. DOI: `10.1103/PhysRevB.58.16076`.

[94] D.P. Devor. “Nd: YAG quantum efficiency and related radiative properties”. In: *Quantum Electronics, ...* 25.8 (1989), pp. 1863–1873. ISSN: 00189197. DOI: `10.1109/3.34046`.

[95] V Lupei. “Low-heat solid-state lasers”. In: *Romanian Reports in Physics* 63 (2011), pp. 1236–1253.

[96] Sébastien Chénais et al. “On thermal effects in solid-state lasers: The case of ytterbium-doped materials”. In: *Progress in quantum ...* 30.4 (Jan. 2006), pp. 89–153. ISSN: 00796727. DOI: `10.1016/j.pquantelec.2006.12.001`.

[97] W A Clarkson. “Thermal effects and their mitigation in end-pumped solid-state lasers”. en. In: *Journal of Physics D: Applied Physics* 34.16 (Aug. 2001), pp. 2381–2395. ISSN: 0022-3727. DOI: `10.1088/0022-3727/34/16/302`.

[98] A.K. Cousins. “Temperature and thermal stress scaling in finite-length end-pumped laser rods”. In: *IEEE Journal of Quantum Electronics* 28.4 (Apr. 1992), pp. 1057–1069. ISSN: 00189197. DOI: `10.1109/3.135228`.

[99] S. T. Lee et al. “A compact laser target designator”. In: *Laser Technology for Defense and Security XII*. Vol. 9834. SPIE, May 2016, 98340Q. ISBN: 9781510600751. DOI: `10.1117/12.2222989`.

[100] W Koechner and Michael Bass. *Solid-State Lasers: A Graduate Text*. 2003. ISBN: 0387955909.

[101] Walter Koechner. *Solid-State Laser Engineering (Springer Series in Optical Sciences)*. Springer, 2006, p. 750. ISBN: 038729094X.

[102] N. Bloembergen and P. S. Pershan. "Light waves at the boundary of nonlinear media". In: *Physical Review* 128.2 (1962), pp. 606–622. ISSN: 0031899X. DOI: `10.1103/PhysRev.128.606`.

[103] Andrew S Glassner. "Surface Physics for Ray Tracing". In: *An Introduction to Ray Tracing*. Ed. by Andrew S Glassner. London, UK: Academic Press Ltd., 1989. Chap. 4, pp. 121–160. ISBN: 0-12-286160-4.

[104] R. Clark Jones. "A New Calculus for the Treatment of Optical SystemsI Description and Discussion of the Calculus". In: *Journal of the Optical Society of America* 31.7 (July 1941), p. 488. ISSN: 0030-3941. DOI: `10.1364/josa.31.000488`.

[105] Henry Hurwitz and R. Clark Jones. "A New Calculus for the Treatment of Optical SystemsII Proof of Three General Equivalence Theorems". In: *Journal of the Optical Society of America* 31.7 (July 1941), p. 493. ISSN: 0030-3941. DOI: `10.1364/josa.31.000493`.

[106] R. Clark Jones. "A New Calculus for the Treatment of Optical SystemsIII The Sohncke Theory of Optical Activity". In: *Journal of the Optical Society of America* 31.7 (July 1941), p. 500. ISSN: 0030-3941. DOI: `10.1364/josa.31.000500`.

[107] R. Clark Jones. "A New Calculus for the Treatment of Optical Systems IV". In: *Journal of the Optical Society of America* 32.8 (Aug. 1942), p. 486. ISSN: 0030-3941. DOI: `10.1364/josa.32.000486`.

[108] Edward Collett. *Field Guide to Polarization*. SPIE, 2009, p. 134. ISBN: 0819458686. DOI: `10.1117/3.626141`.

[109] Raymond Chen. "Why are there broadcast-based mechanisms in Windows?" In: *The Old New Thing* (June 2005). URL: `https://devblogs.microsoft.com/oldnewthing/20050627-00/?p=35193`.

[110] Raymond Chen. "Please feel free to stop using DDE". In: *The Old New Thing* (Feb. 2007). URL: `https://devblogs.microsoft.com/oldnewthing/20070226-00/?p=27863`.

[111] M. D. McKay, R. J. Beckman, and W. J. Conover. "A Comparison of Three Methods for Selecting Values of Input Variables in the Analysis of Output from a Computer Code". In: *Technometrics* 21.2 (May 1979), p. 239. ISSN: 00401706. DOI: `10.2307/1268522`.

[112] Ronald L. Iman, Jon C. Helton, and James E. Campbell. "An Approach to Sensitivity Analysis of Computer Models: Part I—Introduction, Input Variable Selection and Preliminary Variable Assessment". In: *Journal of Quality Technology* 13.3 (July 1981), pp. 174–183. ISSN: 0022-4065. DOI: `10.1080/00224065.1981.11978748`.

[113] a B Owen. *Orthogonal Arrays for Computer Experiments, Integration and Visualization*. 1992. URL: `https://www.jstor.org/stable/24304869?seq=1`.

[114] Boxin Tang. "Orthogonal Array-Based Latin Hypercubes". In: *Journal of the American Statistical Association* 88.424 (Dec. 1993), p. 1392. ISSN: 01621459. DOI: `10.2307/2291282`.

[115] Kenny Q. Ye. “Orthogonal Column Latin Hypercubes and Their Application in Computer Experiments”. In: *Journal of the American Statistical Association* 93.444 (Dec. 1998), p. 1430. ISSN: 01621459. DOI: `10.2307/2670057`.

[116] Ulla Wandinger. “Introduction to Lidar”. In: *Lidar: Range-Resolved Optical Remote Sensing of the Atmosphere*. 2006, pp. 1–18. DOI: `10.1007/0-387-25101-4_1`.

[117] Anant Gupta et al. “Photon-flooded single-photon 3D cameras”. In: *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. Vol. 2019-June. 2019, pp. 6763–6772. ISBN: 9781728132938. DOI: `10.1109/CVPR.2019.00693`.

[118] A. Sume et al. “Radar detection of moving objects around corners”. In: *SPIE Defense, Security, and Sensing*. International Society for Optics and Photonics. 2009, p. 73080V. ISBN: 9780819475749. DOI: `10.1117/12.818123`.

[119] Bhavana Chakraborty et al. “Multipath exploitation with adaptive waveform design for tracking in urban terrain”. In: *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*. IEEE. 2010, pp. 3894–3897.

[120] Andreas Velten et al. “Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging”. In: *Nature Communications* 3 (2012), pp. 745–748. ISSN: 20411723. DOI: `10.1038/ncomms1747`.

[121] Otkrist Gupta et al. “Reconstruction of hidden 3D shapes using diffuse reflections”. In: *Optics Express* 20.17 (Aug. 2012), p. 19096. ISSN: 1094-4087. DOI: `10.1364/oe.20.019096`.

[122] Endre Repasi et al. “Advanced short-wavelength infrared range-gated imaging for ground applications in monostatic and bistatic configurations”. In: *Appl. Opt.* 48.31 (Nov. 2009), pp. 5956–5969. DOI: `10.1364/AO.48.005956`.

[123] Liang Gao et al. “Single-shot compressed ultrafast photography at one hundred billion frames per second”. In: *Nature* 516.7529 (Dec. 2014), pp. 74–77. ISSN: 0028-0836. DOI: `10.1038/nature14005`.

[124] Mauro Buttafava et al. “Non-line-of-sight imaging using a time-gated single photon avalanche diode”. In: *Optics Express* 23.16 (2015), p. 20997. ISSN: 1094-4087. DOI: `10.1364/OE.23.020997`.

[125] I M Vellekoop and A P Mosk. “Universal optimal transmission of light through disordered materials”. In: *Phys. Rev. Lett.* 101.12 (2008), p. 120601.

[126] A llard P. . Mosk et al. “Controlling waves in space and time for imaging and focusing in complex media”. In: *Nature Photonics* 6.5 (2012), pp. 283–292.

[127] Jacopo Bertolotti et al. “Non-invasive imaging through opaque scattering layers”. In: *Nature* 491.7423 (2012), pp. 232–234. ISSN: 00280836. DOI: `10.1038/nature11578`.

[128] Ori Katz et al. "Non-invasive single-shot imaging through scattering layers and around corners via speckle correlations". In: *Nature Photonics* 8.10 (Oct. 2014), pp. 784–790. ISSN: 1749-4885. DOI: `10.1038/nphoton.2014.189`.

[129] Ori Katz, Eran Small, and Yaron Silberberg. "Looking around corners and through thin turbid layers in real time with scattered incoherent light". In: *Nature Photonics* 6.8 (Aug. 2012), pp. 549–553. ISSN: 17494885. DOI: `10.1038/nphoton.2012.150`.

[130] H. Chen et al. "Two-dimensional imaging through diffusing media using 150-fs gated electronic holography techniques". In: *Optics Letters* 16.7 (Apr. 1991), p. 487. ISSN: 0146-9592. DOI: `10.1364/ol.16.000487`.

[131] Justin A. Richardson, Lindsay A. Grant, and Robert K. Henderson. "Low dark count single-photon avalanche diode structure compatible with standard nanometer scale CMOS technology". In: *IEEE Photonics Technology Letters* 21.14 (July 2009), pp. 1020–1022. ISSN: 10411135. DOI: `10.1109/LPT.2009.2022059`.

[132] Justin Richardson et al. "A 32×32 50ps resolution 10 bit time to digital converter array in 130nm CMOS for time correlated imaging". In: *Proceedings of the Custom Integrated Circuits Conference*. IEEE. 2009, pp. 77–80. ISBN: 9781424440726. DOI: `10.1109/CICC.2009.5280890`.

[133] Cristiano Niclass et al. "Design and characterization of a CMOS 3-D image sensor based on single photon avalanche diodes". In: *Solid-State Circuits, IEEE Journal of* 40.9 (2005), pp. 1847–1854.

[134] Desmond O'Connor. *Time-correlated single photon counting*. Academic Press, 2012.

[135] Wolfgang Becker. *Advanced time-correlated single photon counting techniques*. Vol. 81. Springer Science & Business Media, 2005.

[136] Susan Chan et al. "Non-line-of-sight tracking of people at long range". In: *Optics Express* 25.9 (May 2017), p. 10109. ISSN: 1094-4087. DOI: `10.1364/OE.25.010109`.

[137] J. L. Monteith and T. R. Oke. "Boundary Layer Climates." In: *The Journal of Applied Ecology*. Vol. 17. 2. 1980, p. 517. DOI: `10.2307/2402350`.

[138] *Safety of Laser Products - Part 1: Equipment Classification and Requirements*. BS EN 60825-1:2014 (E), Standard. British Standards Institution, Aug. 2014.

[139] Johann Heinrich Lambert. "Lamberts Photometrie: (Photometria, sive De mensura et gradibus luminis, colorum et umbrae) (1760)". In: *W. Engleman, Lambert's Photometrie. Leipzig* 127.128 (1892), p. 432.

[140] Inductiveload. *Lambert Cosine Law [Public Domain]*. 2008. URL: `https://en.wikipedia.org/wiki/Lambert%7B%5C%7D27s%7B%5C_%7Dcosine%7B%5C_%7Dlaw%7B%5C#%7D/media/File:Lambert%7B%5C_%7DCosine%7B%5C_%7DLaw%7B%5C_%7D2.svg` (visited on 09/05/2016).

[141] Dennis D. Lowenthal and John M. Eggleston. "ASE Effects in Small Aspect Ratio Laser Oscillators and Amplifiers with Nonsaturable Absorption". In: *IEEE Journal of Quantum Electronics* 22.8 (Aug. 1986), pp. 1165–1173. ISSN: 15581713. DOI: `10.1109/JQE.1986.1073122`.

[142] M. Lu et al. "Numerical simulation of a CW-pumped Cr: YAG passively Q-switched Yb: YAG pulsed laser". In: *Optics and Lasers in ...* 47.6 (June 2009), pp. 617–621. ISSN: 01438166. DOI: `10.1016/j.optlaseng.2008.12.008`.

[143] D. Liang, J. Almeida, and E. Guillot. "Side-pumped continuous-wave Cr:Nd:YAG ceramic solar laser". In: *Applied Physics B: Lasers and Optics* 111.2 (May 2013), pp. 305–311. ISSN: 09462171. DOI: `10.1007/s00340-013-5334-4`.

[144] Said Mehellou et al. "Stable solar-pumped TEM00-mode 1064 nm laser emission by a monolithic fused silica twisted light guide". In: *Solar Energy* 155 (2017), pp. 1059–1071. ISSN: 0038092X. DOI: `10.1016/j.solener.2017.07.048`.

[145] Christopher Reimer. "Monte Carlo tolerancing tool using nonsequential ray tracing on a computer cluster". In: *International Optical Design ...* 7652 (July 2010). Ed. by Julie Bentley, Anurag Gupta, and Richard N. Youngworth, pp. 76522M–76522M–12. DOI: `10.1117/12.871008`.

[146] K. Ertel et al. "Optimising the efficiency of pulsed diode pumped Yb:YAG laser amplifiers for ns pulse generation". In: *Optics Express* 19.27 (Dec. 2011), p. 26610. ISSN: 1094-4087. DOI: `10.1364/OE.19.026610`.

[147] Magdalena Sawicka et al. "Modeling of amplified spontaneous emission, heat deposition, and energy extraction in cryogenically cooled multislab Yb
3+:YAG laser amplifier for the HiLASE Project". In: *Journal of the Optical Society of America B* 29.6 (June 2012), p. 1270. ISSN: 0740-3224. DOI: `10.1364/JOSAB.29.001270`.

[148] Ulrike Fuchs, Uwe D. Zeitner, and Andreas Tünnermann. "Ultra-short pulse propagation in complex optical systems". In: *Optics Express* 13.10 (May 2005), p. 3852. ISSN: 1094-4087. DOI: `10.1364/opex.13.003852`.

[149] Nicola Coluccelli. "Nonsequential modeling of laser diode stacks using Zemax: simulation, optimization, and experimental validation". In: *Applied Optics* 49.22 (Aug. 2010), p. 4237. ISSN: 0003-6935. DOI: `10.1364/AO.49.004237`.

[150] Dawei Liang and Joana Almeida. "Design of ultrahigh brightness solar-pumped disk laser". In: *Applied Optics* 51.26 (Sept. 2012), pp. 6382–6388. ISSN: 15394522. DOI: `10.1364/AO.51.006382`.

[151] D. F. Swinehart. "The Beer-Lambert law". In: *Journal of Chemical Education* 39.7 (1962), pp. 333–335. ISSN: 00219584. DOI: `10.1021/ed039p333`.

[152] Tobias Blickle and Lothar Thiele. “A comparison of selection schemes used in evolutionary algorithms”. In: *Evolutionary Computation* 4.4 (Dec. 1996), pp. 361–394. ISSN: 10636560. DOI: `10.1162/evco.1996.4.4.361`.

[153] TJ Garrod et al. “Long wavelength surface-emitting distributed feedback (SE-DFB) laser for range finding applications”. In: *Society of Photo-Optical ...* 8241 (Feb. 2012). Ed. by W. Andrew Clarkson et al., pp. 824113–824113–4. DOI: `10.1117/12.909639`.

[154] *GLAD*. 2014. URL: `https://www.aor.com/html/glad.html`.