Characterisation and Advanced Applications of the Steam Oxidation of AlGaAs

Gareth J. Michell

Ph.D. Thesis

School of Physics and Astronomy

Cardiff University

July 2010

UMI Number: U585372

All rights reserved

INFORMATION TO ALL USERS
The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript
and there are missing pages, these will be noted. Also, if material had to be removed,
a note will indicate the deletion.

UMI U585372
Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.
All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Form: PGR_Submission_200701

NOTICE OF SUBMISSION OF THESIS FORM:
POSTGRADUATE RESEARCH

APPENDIX 1:
Specimen layout for Thesis Summary and Declaration/Statements page to be included in a Thesis

DECLARATION

This work has not previously been accepted in substance for any degree and is not concurrently submitted in candidature for any degree.

Signed ...Gareth Mitchell......... (candidate)   Date ...22.7.10.........

STATEMENT 1

This thesis is being submitted in partial fulfillment of the requirements for the degree of ....Ph.D...........(insert MCh, MD, MPhil, PhD etc, as appropriate)

Signed ...Gareth Mitchell......... (candidate)   Date ...22.7.10........

STATEMENT 2

This thesis is the result of my own independent work/investigation, except where otherwise stated. Other sources are acknowledged by explicit references.

Signed ...Gareth Mitchell....... (candidate)   Date ...22.7.10.........

STATEMENT 3

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loan, and for the title and summary to be made available to outside organisations.

Signed ...Gareth Mitchell....... (candidate)   Date ...22.7.10........

STATEMENT 4: PREVIOUSLY APPROVED BAR ON ACCESS

I hereby give consent for my thesis, if accepted, to be available for photocopying and for inter-library loans **after expiry of a bar on access previously approved by the Graduate Development Committee.**

Signed ...Gareth Mitchell....... (candidate)   Date ...22.7.10.......

Acknowledgements

For their guidance, support and input over the years, I would like to thank my supervisors Peter Smowton and Huw Summers. Similarly, I would like to thank all of the members of the Photons and Matter group that have provided help, practical advice and invaluable discussions in times of confusion. In this area, I would particularly like to thank Stella Elliott for putting up with my stream of interruptions and questions and her constant willingness to help answer them. My thanks also go specifically to James Ferguson for giving his own time to help out practically in my final weeks, and to Deepal Naidu for providing a sympathetic ear and always being just that little bit closer to the deadline than me.

Special thanks must go the cleanroom staff of Cardiff University, Karen Barnett and Angela Sobiesierski and to Gareth Edwards for their continued efforts above and beyond to solve my problems and to get samples made even under seemingly impossible time restrictions. I also thank my family and friends for their general support throughout, and for the provision of so many enjoyable distractions from work.

Finally, my kind thanks go Des Hub who has been, over the course of my time in Cardiff, a tutor, a housemate and a friend. Without his endless generosity, inspiration and free accommodation this would certainly have not been possible.

Abstract

The high temperature, wet oxidation of buried, high aluminium content AlGaAs layers has seen a lot of attention in recent years thanks to its ability to form good quality, deep and uniform oxide layers; a unique attribute among the III-V materials. It is the unusual combination of relatively high density and good porosity of the oxide material, that allows the oxidation to reach much greater depths than that of other materials, and this is what makes the process so widely employable and versatile.

This project is based on two systems of oxidisable layers, incorporated into two wafers, both of which were designed as part of the work. One wafer, referred to as Wafer A, contains 30nm 98% aluminium layers in both of the cladding regions, and the second, Wafer B, has spaced stacks of 1nm 95% aluminium layers placed directly into the active region.

The oxidation depths of both wafers are mapped over a range of temperatures and times, and maximum measured rates of 2μm/min and 1μm/min are reported for Wafers A and B respectively. Further to this, a non-trivial dependence of the oxidation rate on dopant species is reported. It is found that although p-type material initially achieves greater oxidation rates, after a depth of approximately 8μm has been reached, the n-type material over takes. This result is contrary to the popular belief that p-type material oxidises faster, regardless of depth.

The two wafers are then made into working devices. In the case of Wafer A, longitudinal oxidation of broad-area devices are used to create unpumped saturable absorber sections. These sections are responsible for various effects including the lowering of laser threshold currents by up to 30%, the narrowing of laser near-fields by up to 50%, and the production of repetitive wavelength shifting behaviour, which is investigated in great detail. Wafer B is oxidised laterally to produce current confinement directly in the active region. Threshold current reductions of around 20% are recorded. Lastly, cross-sectional electron microscope images reveal a simple and versatile method of fabricating buried microlenses from the structure of Wafer B.

Chapter 1 : Introduction and Laser Background 1

1.1 Introduction 1

1.2 Background principles of lasers 2

1.2.1 Laser diode basics 2

1.2.2 Band diagrams and density of states 3

1.2.3 Laser cavity modes and index guiding 6

1.2.4 Modal gain and the threshold condition 8

1.2.5 Gain and loss measurements by the multisection technique 9

1.2.6 Saturable absorption and self pulsation 11

1.3 Thesis structure and motivation 14

1.4 References 16

Chapter 2 : Background of Oxidation of AlGaAs 18

2.1 Introduction 18

2.2 Development 18

2.3 Chemistry of AlGaAs oxidation 20

2.4 Physical and electrical properties of oxides 22

2.5 Factors affecting oxidation rates 26

2.5.1 Sample temperature and $Al_xGa_{1-x}As$ composition 26

2.5.2 Layer thickness 28

2.5.3 Doping 29

2.5.4 Carrier gas flow and water bubbler temperature 30

2.6 Summary of applications 31

2.6.1 Current apertures 31

2.6.2 Waveguiding 32

2.6.3 Reflectors and wavelength control 34

2.6.4 Surface roughness reduction 34

2.6.5 Phosphide-based materials 35

2.7 References 36

Chapter 3 : Experimental Techniques 39

3.1 Introduction 39

3.2 Steam Oxidation 39

3.3 IVLT measurements 41

3.4 Atomic Force Microscopy (AFM) 44

3.5 Scanning Electron Microscopy (SEM) 47

3.6 Streak camera 48

3.7 The near-field camera 51

3.8 The multisection technique 53

3.9 The integrating sphere 56

3.10 References 58

Chapter 4 : Design and Analysis of Oxidisable Material Systems 59

4.1 Introduction 59

4.2 Wafer structures 59

4.2.1 Simulation results 64

4.3 Measuring the oxidation rate 68

4.4 Oxidation rates in 30nm layers 71

4.4.1 Dependence on temperature 71

4.4.2 Dependence on time and dopant species 72

4.4.3 Dependence on etched sidewall profile 75

4.5 Oxidation rates in stacked 1nm layers 78

4.5.1 Dependence on temperature 80

4.5.2 Dependence on time 81

4.6 Device reliability 82

4.7 Summary 85

4.8 References 86

Chapter 5 : Effects from Lateral Oxidation of Active Region Layers 88

5.1 Introduction 88

5.2 Electrical performance 88

5.3 Sub-threshold near-fields 89

5.4 Laser performance 92

5.5 Fabrication of buried microlenses 96

5.6 Summary 100

5.7 References 102

Chapter 6 : Effects from Facet Oxidation of Cladding Region Layers..... 103

6.1 Introduction 103

6.2 The first result – Threshold dependence on cavity length 103

6.3 The proposed mechanism 107

6.4 Experimental evidence 109

6.4.1 The saturable absorber 109

6.4.1.1 Turn-on delay 109

6.4.1.2 LI characteristics 112

6.4.2 Near-field narrowing and mechanism 114

6.5 The second result – Repetitive wavelength shifting 120

6.6 Extending the proposed mechanism 122

6.7 Further experimental evidence 123

6.7.1 Facet heating 123

6.7.1.1 Catastrophic optical damage 124

6.7.1.2 RWS analysis to show facet heating 126

6.7.2 Mode superposition 129

6.8 Further analysis 131

6.9 Summary 135

6.10 References 137

Chapter 7 : Conclusions and Future Work 138

7.1 Conclusions 138

7.1.1 Modifications to oxidation equipment 138

7.1.2 Wafer designs and oxidation rate characterisation 138

7.1.3 Dopant dependence of oxidation rate 138

7.1.4 Sidewall profile dependence of oxidation rate 139

7.1.5 Current confinement by active region oxidation 139

7.1.6 Microlens fabrication 139

7.1.7 Near-field and threshold current changes by facet oxidation 139

7.1.8 Repetitive Wavelength Shifting 140

7.1.9 Losses added by oxidation 140

7.2 Future work 140

7.2.1 Oxidation rate dependence on doping 140

7.2.2 Reduction of delamination 141

7.2.3 Near-fields and thresholds of Wafer B 141

7.2.4 Microlenses 141

7.2.5 Reducing COD in Wafer A 142

7.2.6 Laser turn-on delay 142

Presentations and publications 143

Chapter 1 : Introduction and Laser Background

1.1 Introduction

Laser diodes have found many applications since their conception [1] and have benefitted from several leaps in design from the original homojunction devices, [2] through quantum wells [3] and, most recently, into quantum dots. [4] The devices discussed in this work are all quantum dots emitting at wavelengths around 1.3µm, one of the key wavelength ranges for modern fibre-based telecommunications.

Traditionally, the presence of oxides in or on III-V semiconductor devices, such as laser diodes, would have often been a sign of degradation and a cause of diminishing performance. In silicon technologies, however, the high quality, stable thermal oxide of silicon has long been used as an integral part of device design. For example, silicon dioxide is key to the function of standard silicon MOSFETs. It was not until the late 1970s, however, that an oxide of a reasonable quality was reported for III-Vs. [5]

Since that first report of steam oxidation of high aluminium-content AlGaAs, the technology has been investigated widely and incorporated into many areas of III-V device design. Most notably, in recent years, has been the fabrication of very low threshold current VCSELs, which incorporate partially oxidised AlGaAs layers as current-confining apertures. [6]

Despite this success, the technology is still relatively young and certainly worthy of further study, both in the oxidation process itself and in its eventual applications. It is to this body of knowledge that this work aims to add, by studying new oxidisable layer systems and new methods of employing them.

1.2 Background principles of lasers

This section aims to provide some, but not all, of the background physics behind the results in later chapters. Any further detail required can be found in any of the major relevant text books.[7] [8] [9] [10]

1.2.1 Laser diode basics

A laser diode is a semiconductor device, very similar to a normal p-n diode, but with two main differences; the bandgap is such that recombination of electron-hole pairs across it has a high probability of producing a photon and the material is located between two parallel mirrors. In the presence of sufficient numbers of excited electron-hole pairs, this situation leads to the emission, over a very narrow range of wavelengths, of intense and coherent beams of light.

Figure 1.1 shows a diagram of a standard, broad-area laser device. It depicts the oxide-defined stripe of the top contact (usually 50μm wide), the light-emitting active region, and defines the axes of the device as they will be referred to throughout this work.

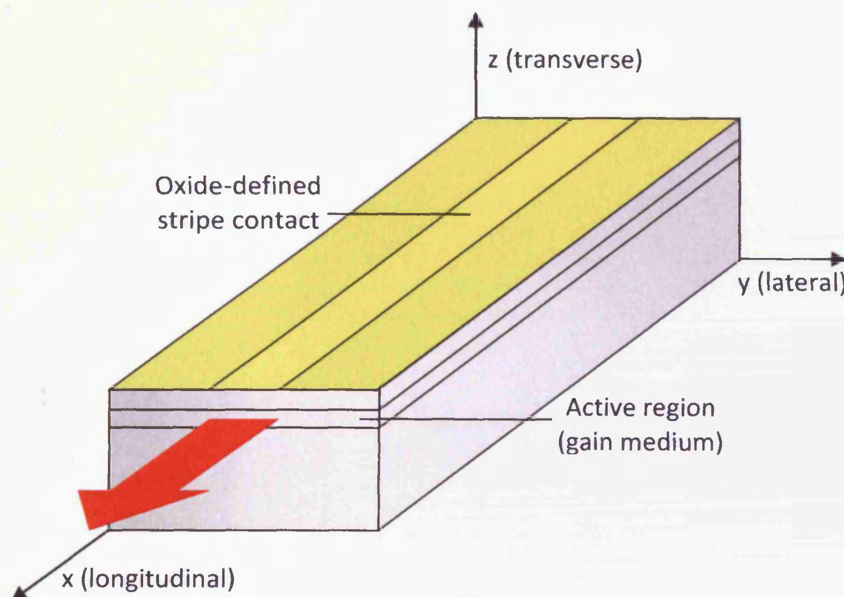

Figure 1.1: Definitions of axes as used in this work

The basic condition for laser action in any type of laser is that optical gain in the device matches the sum of all optical losses. In the case of laser diodes, this gain comes from electrical pumping. Pumping a semiconductor laser, be it electrically or optically, means exciting electrons from the valence band into the conduction band. Once excited, electrons can recombine with holes of the same momentum and

release photons, either by spontaneous or stimulated emission. Whilst still within the lasing medium, there is also a chance that emitted photons may be reabsorbed, creating another electron-hole pair. These three processes are represented below in Figure 1.2.

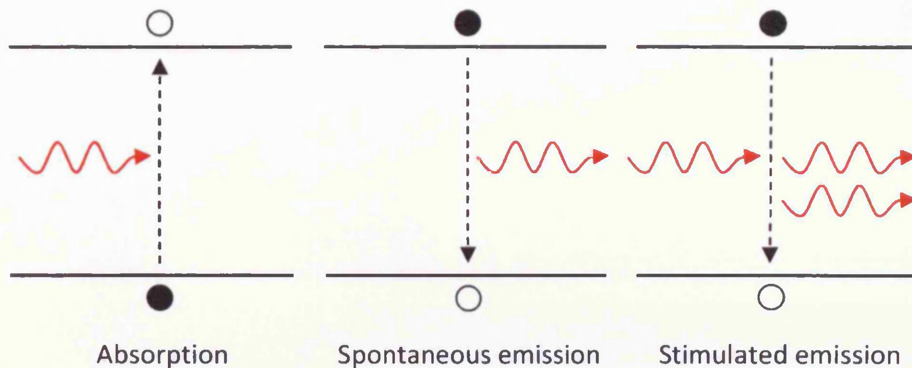

Figure 1.2: Absorption and emission events in a simple, two-level system

There are also non-radiative methods of relaxation possible, Auger and defect-centre recombination for example, but these are not illustrated. In lasers, the vast majority of the outputted light, once the laser has fully turned on, comes from stimulated emission. The consequence of this, and the defining property of a laser, is that the light is coherent, highly directional, and emitted over a very short range of wavelengths.

Both of the radiative processes require excited electron-hole pairs, the availability for recombination of which, depends both on the presence of electrons and holes in the conduction and valence bands respectively, and on the wave vector, or momentums of both particles. Information on both of these factors is derived from the density of states and band structure of the material.

1.2.2 Band diagrams and density of states

As mentioned earlier, laser diodes, unlike most standard diodes, are made from material with a band structure that allows photons to be produced from recombination events. The specific property required for this is a 'direct' bandgap. An energy against wave vector graph, such as the one in Figure 1.3, shows the configuration of the bands of a material and whether it has a direct or indirect bandgap.

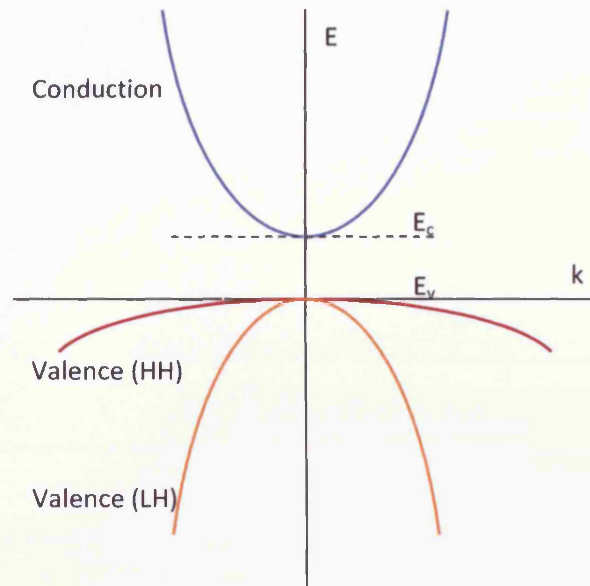

Figure 1.3: E-k diagram for a typical direct-gap semiconductor. Valence band is split into light- and heavy-hole bands

It is because of the law of k-selection that a direct gap is necessary for emission of photons. Only electron-hole pairs with the same momentum (vertical transitions on the E-k diagram) are able to recombine via spontaneous or stimulated emission. Non-vertical transitions are possible, although, in these cases, a second electron or a phonon must be involved in order to conserve momentum. In such transitions, energy is almost always ultimately lost to heating the semiconductor lattice, and photons are not produced.

To fully understand the band structure of a material, the density of states (DoS) function must also be considered. The DoS is the number of available states per unit energy per unit volume and its form depends strongly on the dimensions of the material, or any confining potential barriers. Figure 1.4 shows DoS functions for the conduction and valence bands of ideal bulk, quantum well and quantum dot samples.

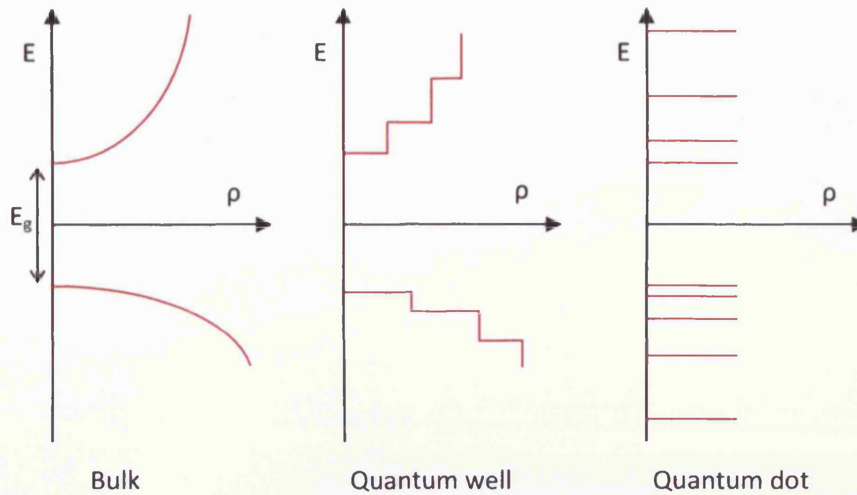

Figure 1.4: Density of states for ideal bulk, quantum well and quantum dot material

By multiplying the DoS function by the Fermi function, or probability of occupation function, shown in Equation 1.1, the number of occupied states as a function of energy can be determined.

$$f(E) = \frac{1}{1 + e^{\left(\frac{E-E_f}{kT}\right)}} \tag{1.1}$$

In terms of full laser structures, all of the above information can be used to draw a band diagram of any particular layer system. Figure 1.5 shows a sketch of a band diagram of a simple quantum well (QW) device.

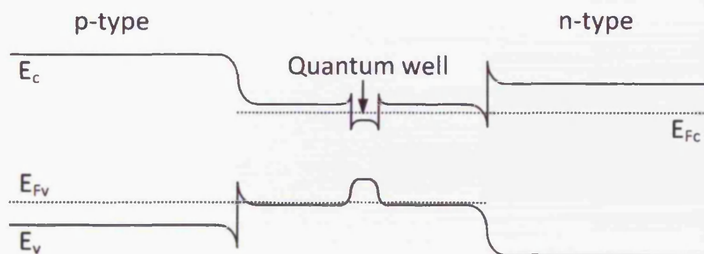

Figure 1.5: Band diagram of a simple quantum well device under forward bias

In such a diagram the x axis displays distance through the device in the direction of layer growth, and the y axis shows energy. Taken into account in this diagram is the band bending at the interfaces between different materials, caused by a combination of the band offsets, which are fixed quantities for any pair of materials, and by the equalisation of the different Fermi energies as two materials are brought into contact. Shown in the diagram are the 'quasi' Fermi energies, or those applying to

electrons in the conduction band, E_{Fc}, and to holes in the valence band, E_{Fv}. It is these quantities that will later be related to the ability of a material system to exhibit a net gain.

The factors considered in this section provide information on how many electron-hole pairs are available for radiative recombination and at what energy separation. Although this relates directly to the energy (or wavelength) of the emitted photon from any single recombination event, further factors are involved in the determination of the eventual lasing wavelength. These are the refractive index of the material and cavity length.

1.2.3 Laser cavity modes and index guiding

By definition, a semiconductor laser requires the gain medium to be located between two parallel mirrors to provide optical feedback. Consequently, a resonant cavity is formed, within which only certain wavelengths are allowed to propagate by the assertion that a standing wave is formed and nodes (points of zero intensity) must be located at the mirror surfaces. These allowed wavelengths are the longitudinal cavity modes and are referred to by their order number. Figure 1.6 shows the first, second and fourth order modes of a simple cavity.

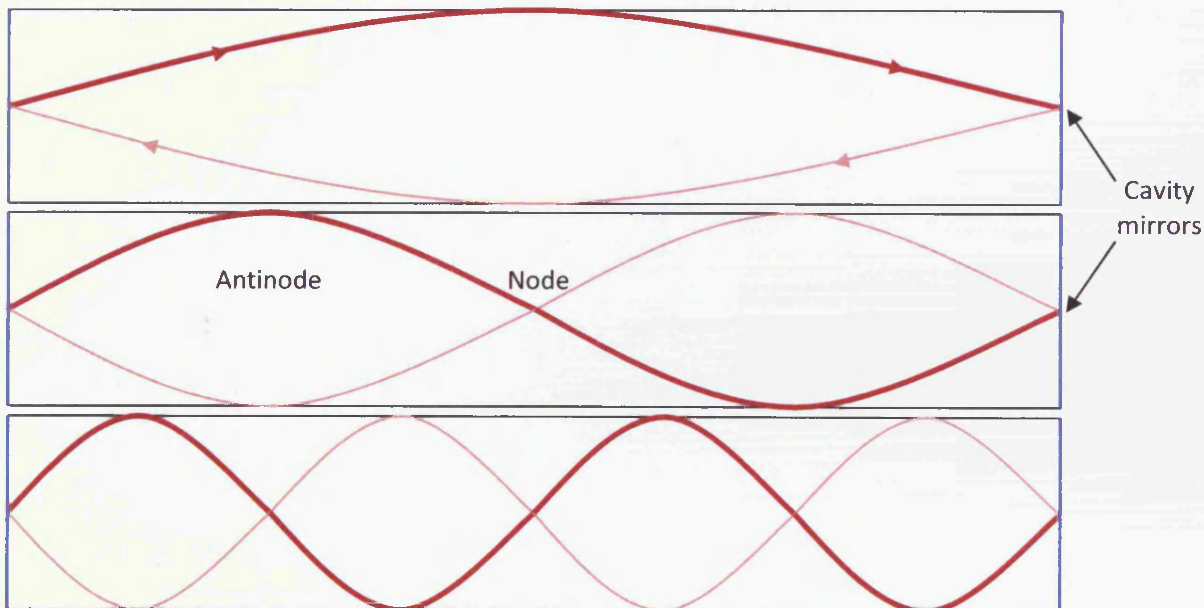

Figure 1.6: The 1st, 2nd and 4th longitudinal modes in a simple single cavity

An important and calculable figure is the wavelength spacing of these modes. In the figure above, the difference in wavelength between each mode depicted is clearly significant, when compared to the

absolute wavelength value. In fact, each successive picture above represents a doubling of the allowed wavelength. However, in a typical semiconductor laser, the cavity length is far greater than the wavelength of the light within it, so the wavelength spacing between adjacent modes is very small, in comparison to the actual wavelength, and can be said to be approximately constant over the small number of modes that are likely to be involved. Using this approximation, the mode wavelength spacing, $\Delta\lambda$, within a simple cavity can be calculated using Equation 1.2.

$$\Delta\lambda = \frac{\lambda^2}{2nL} \tag{1.2}$$

In this equation, n is the refractive index of the laser material and L is the length of the cavity. For the devices used in this work, emitting around 1.3µm with cavities lengths of a few millimetres, mode wavelength spacings of around 0.5nm are typical. For more complex cavities, with multiple sections, this equation can be extended to calculate the overall mode spacing by considering the length and refractive index of each section individually. This is shown for a two-section device in Equation 1.3. [7]

$$\Delta\lambda = \frac{\lambda^2}{2(n_1L_1 + n_2L_2)} \tag{1.3}$$

The same mode structure also exists in the transverse and lateral directions, although, in the case of the transverse, the confinement of the light is almost always sufficient to restrict operation to a single mode. Confinement of the light in the transverse direction is achieved by refractive index guiding through careful selection of the growth layers. Figure 1.7 shows how this would be achieved in the same simple QW system for which the band diagram was drawn in Figure 1.5.

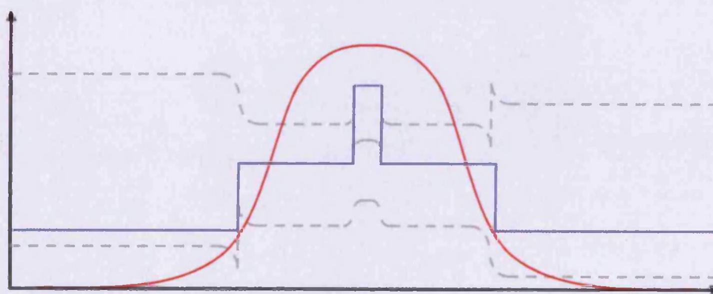

Figure 1.7: Refractive index profile (blue) and guided light intensity profile (red) on top of QW band diagram (grey)

The step changes in refractive index effectively provide partially reflective interfaces, whose strength of reflection depends on the size of the step. In the case of quantum wells, the wavelength of the light will always be larger than the width of the well, so confinement entirely within the well could never be

achieved. This refractive index guiding, along with all the other factors discussed so far, will all have an effect on the overall gain of the material.

1.2.4 Modal gain and the threshold condition

Modal gain, in general, is defined very simply as the rate of increase of optical intensity in a single mode, either per unit time or unit distance. Relating this to the processes shown in Figure 1.2, the modal gain per unit time is equal to the stimulated rate, minus the absorption rate. Multiplying this by the refractive index over the speed of light, c, gives the modal gain per unit distance.

The modal gain required to reach threshold in a semiconductor laser diode is that which completely and exactly overcomes all the loss inherent to the device. This is quantified by Equation 1.4 where G_{th} is the modal gain required for threshold, α_i is the intrinsic waveguide loss, L is the cavity length and R is the mirror reflectivity.

$$G_{th} = \alpha_i + \frac{1}{L}\ln\left(\frac{1}{R}\right) \tag{1.4}$$

Both G_{th} and α_i are measured per unit length. The parameter α_i is a fixed value for any particular device and, under most conditions, is irrespective of the level of population inversion. Mechanisms distributed along the whole length of devices, such as scattering, contribute to this value. The remaining term in the equation, referred to in its entirety as the mirror loss, is the only other inherent loss that must be overcome to reach threshold. For a standard, as-cleaved facet of a GaAs-based semiconductor in an atmosphere of air, the mirror reflectivity, R, has a typical value of about 0.3, and simply quantifies the proportion of light that escapes from the facets during every half round-trip of the cavity. The quantity R is defined in equation 1.5, where n_1 is the refractive index of the material and 1 is the index of air.

$$R = \left(\frac{n_1 - 1}{n_1 + 1}\right)^2 \tag{1.5}$$

The threshold requirement stated in Equation 1.4 can be more rigorously related to many of the preceding sections of this chapter through the Bernard-Duraffourg condition for achieving net gain, [11] which is given in Equation 1.6. It states that for net gain the quasi Fermi level separation must be greater than or equal to the photon energy, ħω, which must be greater than or equal to the bandgap.

$$E_{Fc} - E_{Fv} \geq \hbar\omega \geq E_c - E_v \tag{1.6}$$

The argument behind this statement is as follows. For net modal gain (net stimulated rate) to be achieved, the rate of downward transitions of electrons must be greater than the total upward rate. This is given by Equation 1.7, where f_1 and f_2 are the probability of occupancy by electrons of the lower and upper states respectively, and C is a constant relating to the probability of a transition.

$$f_2(1-f_1)C > f_1(1-f_2)C \tag{1.7}$$

From the Fermi-Dirac distribution:

$$f_2 = \frac{1}{1+e^{\left(\frac{E_c-E_{Fc}}{kT}\right)}} \tag{1.8}$$

$$f_1 = \frac{1}{1+e^{\left(\frac{E_v-E_{Fv}}{kT}\right)}} \tag{1.9}$$

which, when put back into Equation 1.7, simplifies to:

$$\Delta E_F > hv \tag{1.10}$$

where hv is the photon energy. Equation 1.10 is equivalent to the Bernard-Duraffourg condition given in Equation 1.6.

In this way, once Equation 1.4 has been met, and threshold reached, the gain must pin at a fixed level even if the pumping level is increased further. If this was not the case, and the net gain exceeded the total losses, the optical intensity would increase without limit, which is clearly not possible. From threshold upwards, any carriers injected above threshold level are assumed to go directly into the stimulated emission process, and are converted into photons at a 1:1 ratio.

1.2.5 Gain and loss measurements by the multisection technique

The net modal gain is separate, but related, to the intrinsic material gain, g, by Equation 1.11, where Γ is the optical confinement factor (dependent on waveguiding properties such as those already mentioned) and α_i is the waveguide loss.

$$G = \Gamma g - \alpha_i \tag{1.11}$$

Experimental determination of the modal gain has historically been performed by a number of methods, but the most practical and direct is the multisection device method, developed in Cardiff University. [12] [13] This technique requires only a single device with its top contact split into individually pumped, identical sections. The principle of the experiment is that identical sections, pumped to the same level,

will produce exactly the same light output, so, by measuring the output from all combinations of the front two sections, the gain and loss spectra of the material can be derived simultaneously.

To determine the gain the measured spectra from pumping both of the front two sections together is divided by that of just the front section alone. The intensity produced from one section of length l is:

$$I_l = \int_0^l \beta I_{spon} e^{(G-\alpha_i)l} dx \qquad (1.12)$$

$$I_l = \beta I_{spon} \frac{e^{(G-\alpha_i)l} - 1}{G - \alpha_i} \qquad (1.13)$$

Dividing the same equation for a length of $2l$ by Equation 1.13 then, gives:

$$\frac{I_{2l}}{I_l} = \frac{e^{(G-\alpha_i)2l} - 1}{e^{(G-\alpha_i)l} - 1} \qquad (1.14)$$

By expansion of the top line, this reduces to:

$$\frac{I_{2l}}{I_l} = e^{(G-\alpha_i)l} + 1 \qquad (1.15)$$

Rearranging then gives the final result:

$$(G - \alpha_i) = \frac{1}{l} \ln\left(\frac{I_{2l}}{I_l} - 1\right) \qquad (1.16)$$

I_{spon} is the uniform intensity of spontaneous emission given out along the length of the device, and β is a factor governing how much of that intensity is coupled to the lasing mode. Similarly, the absorption, A, is found by dividing the measured intensity from the front section alone by that of the second section alone. In these equations, C is a collection efficiency factor and R is the reflectivity of the device mirrors. The measured intensities when pumping the two sections individually are given by:

$$I_{meas1} = C(1 - R)I_1 \qquad (1.17)$$

$$I_{meas1} = C(1 - R)(e^{(-A-\alpha_i)l})I_2 \qquad (1.18)$$

Dividing I_{meas1} by I_{meas2} and making the assumption that $I_1=I_2$ gives:

$$\frac{I_{meas1}}{I_{meas2}} = e^{(A+\alpha_i)l} \qquad (1.19)$$

Rearranging gives the result:

$$(A + \alpha_i) = \frac{1}{l}\ln\left(\frac{I_{meas1}}{I_{meas2}}\right) \tag{1.20}$$

Data yielded by the experiment is representative of Equations 1.16 and 1.20 and will give accurate gain and loss spectra and provide a value of α_i with an error or around ±1cm^{-1}, provided that the sections are truly identical and there is no round-trip gain in the device. Practical considerations for achieving this are discussed in chapter three.

1.2.6 Saturable absorption and self pulsation

A saturable absorber is an element of an optical device that has a certain level of loss, which can be dramatically, and often suddenly, reduced by sufficient intensity of incident light. Such elements have been known about and used in lasers for many years, and can take several physical forms. Individual epitaxial layers, externally coupled material, and separately contacted but monolithic sections in laser diodes can all be designed to have saturable absorption properties.

The most common use for saturable absorption is to produce self-pulsation, i.e. a bistable optical output from a constant pumping source. In the case of semiconductor laser diodes, the easiest way of creating a saturable absorption is in the 'passive' regime, where the absorber can be made simply by leaving some amount of the device material unpumped. By this method, it is the active region (quantum dots or quantum wells) that provide the majority of the absorption, as they inherently have states available at and around the lasing wavelength.

The principle behind pulse generation from a saturable absorber is relatively straightforward. As the optical intensity within a laser cavity grows, excited electrons will be generated in the upper states of any absorbing material in the system, thus slightly reducing its ability to absorb further. With sufficient intensity, a situation is approached where all the available absorbing states are occupied, and the loss caused by the absorber is significantly and suddenly reduced. When this happens, the net gain is increased, and so the optical intensity is allowed to grow very rapidly. This is shown in Figure 1.8 below.

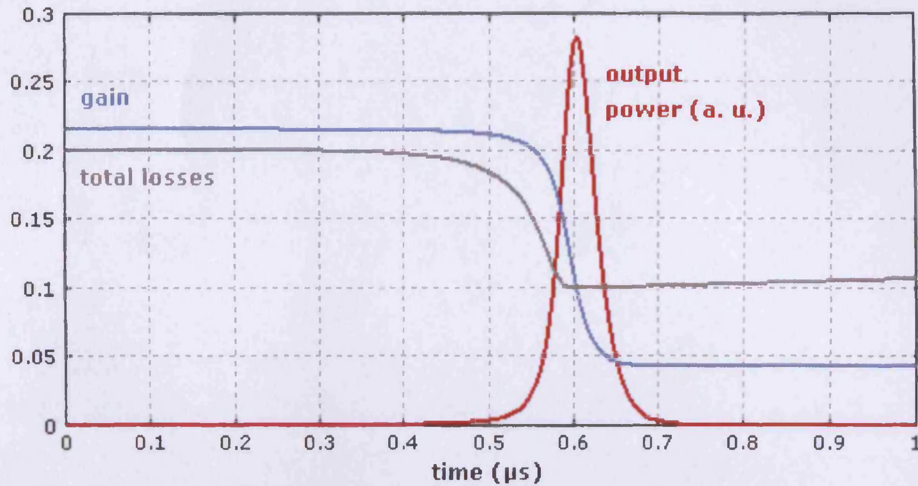

Figure 1.8: Pulsation generation through saturable absorption [14]

This sudden increase in optical intensity is entirely generated by stimulated emission, and so is a large drain on the available excited electrons in the gain medium. After a short period, conduction band electron concentrations become depleted and the overall gain begins to drop, shortly followed by the optical intensity. This signifies the end of the pulse and allows electrons in the absorber to begin recombining, so raising the losses of the absorber again. The amount of time this takes can be varied by the design of the absorber and is referred to as the absorber recovery time.

The pulsation mechanism just described is that of 'Q-switching'; so-called because the sudden drop in losses effectively raises the Q factor of the cavity momentarily to form a pulse. Typical pulse durations from Q-switching are in the low nanosecond range, however, there is another common form of self-pulsation – mode locking – that can also be achieved through saturable absorption, and with which, pulse lengths as low as five femtoseconds can be produced. [15]

In the case of Q-switching, the saturable absorber can be physically located anywhere within the optical mode space; indeed, it can be along the whole length of the device if it comes from an epitaxial layer. With mode locking, however, the absorption must be localised and carefully placed in one or more locations along the length of the device. This is because the principle behind mode locking lies in the optical wavefront passing through areas of saturable absorption at regularly spaced intervals in time, and it is the spacing of those intervals that determines the repetition rate.

The most common method of including a saturable absorber in laser diodes, for the purpose of producing mode locking, is to fabricate a separately contacted section on one end of the device, which

can be grounded or reverse-biased whilst the rest is pumped as normal. This provides an absorption that the wavefront experiences exactly once every round-trip, which, after several passes, will establish a 'steady state' of one pulse per round-trip. The variation of gain, loss and optical power in the cavity of such a device is shown in Figure 1.9.

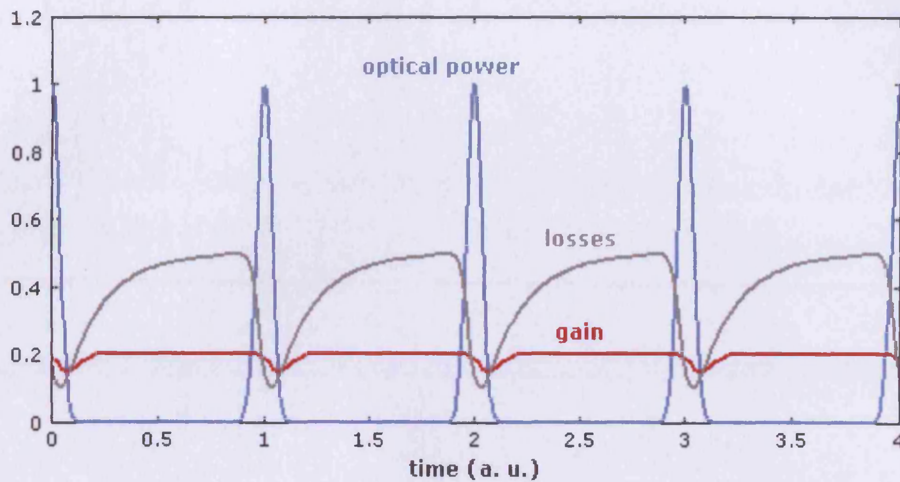

Figure 1.9: Mode locking behaviour [14]

In the example above, the absorber is showing very slow recovery. The only difference that would be seen here with a fast recovering absorber would be that the losses would return to their maximum value almost immediately after the end of the optical pulse.

It is clear that both Q-switching and mode locking are produced by the inclusion of very similar structures in laser diodes; indeed it is not uncommon for both processes to be generated in a single device, sometimes even at the same time.

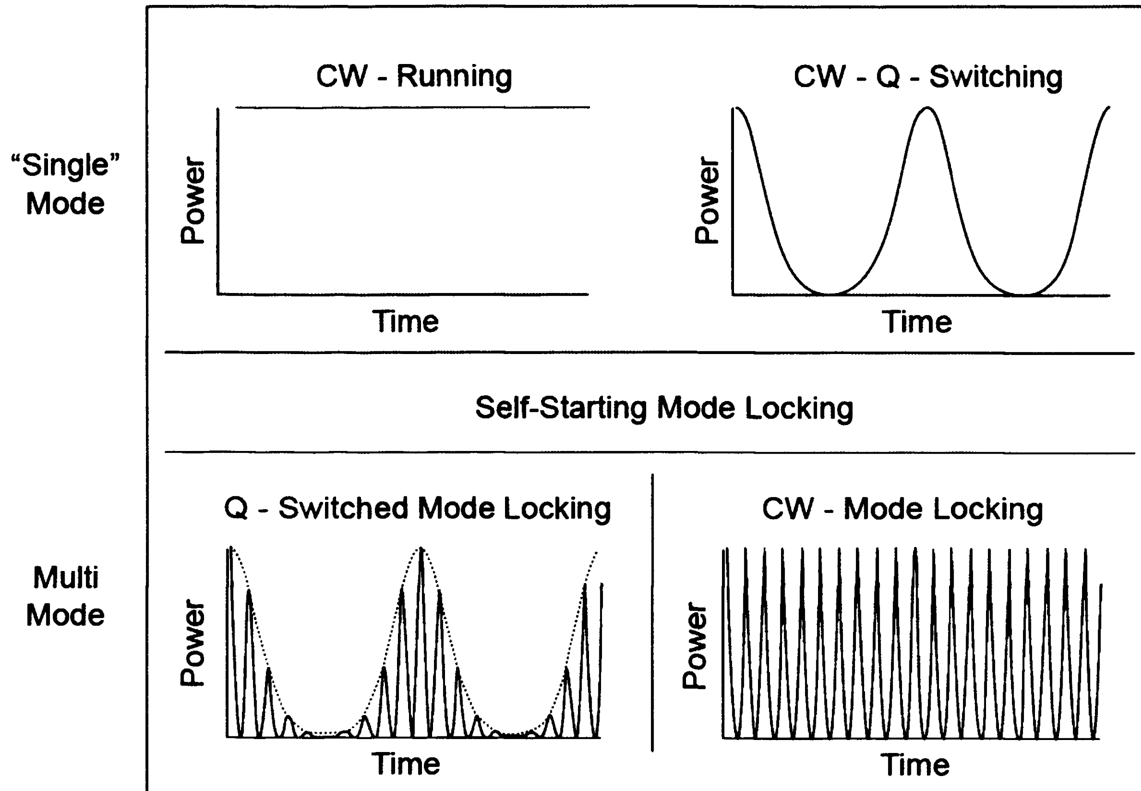

Figure 1.10: Single and simultaneous modes of pulsation [16]

Often, all four possible modes of performance depicted in Figure 1.10 can be achieved in one device, simply by varying the drive current and temperature. With careful design of the absorber parameters, such as the saturation intensity and the recovery time, it is possible to control which of the four regimes appear under what conditions.

1.3 Thesis structure and motivation

This thesis comprises seven distinct chapters. The first two, including this one, give largely background information; firstly, on the basic physical principles behind lasers of all kinds and, secondly, on what is known of the mechanism, chemistry and current applications of steam oxidation. Chapter 3 is given to the consideration of the main experimental techniques, stating details of their setup, use and, where applicable, their limitations. The great majority of the results displayed in this thesis are taken from two wafer designs; one with oxidisable layers in the active region, and one with layers in the cladding regions. Chapter 4 relates results based on the oxidation rates of both of these wafers, and their dependence on various parameters. Chapters 5 and 6 discuss the effects on working devices of the

oxidation of the active and cladding region layers, respectively. To close, Chapter 7 contains conclusions and suggestions for future work.

The steam oxidation of AlGaAs has been developed considerably in recent years, although its main applications remain reasonably limited, with the vast majority of reports on its use being based on current aperture formation in VCSELs. The motivations of this work are, therefore, as follows. Firstly, to continue from the developmental work performed in previous years in Cardiff University [17] in refining and characterising the oxidation equipment and process. This will provide the obvious benefit of making the oxidation process both easier and more repeatable. Secondly, this work will be based entirely on edge-emitting devices, which have only recently received relatively minor consideration in the literature with regard to AlGaAs oxidation. Within edge emitting structures, both thick, cladding region oxidisable layers and thin active region ones will be investigated practically. Both of these layer systems will be novel in some way, and the latter of the two is, as yet, unreported in its entirety.

Furthermore, this work will attempt to branch into different modes of use of the oxide layers. For example, all oxidation in edge emitting devices to date has been in the lateral direction, in pre-cleave ridge devices. If oxidation is employed in the longitudinal direction in broad-area devices, sections of unpumped material and saturable absorbers could be produced. If this were the case, it would provide a much simpler and easily variable method of creating such structures than has previously been used. Following on from that, it may be possible to create self-pulsation in the same devices. Altogether, this project has the potential to greatly expand the general base of knowledge in the area of AlGaAs oxidation, as well as to produce new devices and novel optical effects from what is, essentially, a very simple and accessible technology within the realm of III-V materials.

1.4 References

1. *Stimulated optical radiation in ruby.* **Maiman, T. H.** 187, Nature, 1960.

2. *Coherent ligh emission from GaAs junctions.* **Hall, R. N.; Fenner, G. E.; Kingsley, J. D.; Soltys, T. J.; Carlson, R. O.** 9, Physical Review Letters, 1962, Vol. 9.

3. *Quantum states of confined carriers in very thin AlGaAs-GaAs-AlGaAs heterostructures.* **Dingle, R.; Wiegmann, W.; Henry, C. H.** 14, Physical Review Letters, 1974, Vol. 33.

4. **Dingle, R. and Henry, C.H.** *Quantum effects in heterostructure lasers. 3982207* U.S., 21 September 1976.

5. *Self-terminating thermal oxidation of AlAs epilayers grown on GaAs by molecular beam epitaxy.* **Tsang, W. T.** 5, Applied Physics Letters, 1978, Vol. 33.

6. *Microstructure of laterally oxidized AlGaAs layers in vertical-cavity lasers.* **Twesten, R. D.; Follstaedt, D. M.; Choquette, K. D.; Schneider, R. P.** 1, Applied Physics Letters, 1996, Vol. 69.

7. **Coldren, L. A.; Corzine, S. W.** *Diode Lasers and Photonic Integrated Circuits.* John Wiley & Sons, Inc., 1995.

8. **Yariv, A.** *Optical Electronics in Modern Communications.* Oxford University Press, 1997.

9. **Wilmsen, C. W.** *Physics and Chemistry of III-V Compound Semiconductor Interfaces.* Plenum Press, 1985.

10. **Sze S. M.** *Physics of Semiconductor Devices.* John Wiley & Sons, 1981.

11. *Laser conditions in semiconductors.* **Bernard, M. G. A.; Duraffourg, G.** 7, Physica Status Solidi (b), 1961, Vol. 1.

12. *Determination of single-pass optical gain and internal loss using a multisection device.* **Thomson, J. D.; Summers, H. D.; Hulyer, P. J.; Smowton, P. M.; Blood, P.** 17, Applied Physics Letters, 1999, Vol. 75.

13. *Characterization of semiconductor laser gain media by the segmented contact method.* **Blood, P.; Lewis, G. M.; Smowton, P. M.; Summers, H. D.; Thomson, J. D.; Lutti, J.** 5, IEEE Journal of Selected Topics in Quantum Electronics, 2003, Vol. 9.

14. **RP Photonics.** Encyclopedia of Laser Physics and Technology. [Online] [Cited: 24 March 2010.] http://www.rp-photonics.com/encyclopedia.html.

15. *Semiconductor saturable-absorber mirror-assisted Kerr-lens mode-locked Ti:sapphire laser producing pulses in the two-cycle regime.* **Sutter, D. H.; Steinmeyer, G.; Gallmann, L.; Matuschek, N.; Morier-Genoud, F.; Keller, U.; Scheuer, V.; Angelow, G.; Tschudi, T.** 9, Optics Letters, 1999, Vol. 24.

16. *Semiconductor saturable absorber mirrors (SESAMs) for femtosecond to nanosecond pulse generation in solid-state lasers.* **Keller, U.; Weingarten, K. J.; Kärtner, F. X.; Kopf, D.; Bruan, B.; Jung, I. D.; Fluck, R.; Hönninger, C.; Matuschek, N.; Aus der Au, J.** 3, IEEE Journal of Selected Topics in Quantum Electroncs, 1996, Vol. 2.

17. **Pope, I.** *The construction of a steam oxidation chamber.* Cardiff University, 1999.

Chapter 2 : Background of Oxidation of AlGaAs

2.1 Introduction

This chapter aims to give a detailed background of all aspects of the oxidation of aluminium-based III-V semiconductor compounds. A description of the development of the oxidation process from its original conception in the late 1970s will be given along with details of the chemical reactions involved, characterisation of the oxide films and some of the major applications employing the technique to date.

2.2 Development

The silicon semiconductor industry has long employed the controlled growth of silicon dioxide (SiO_2) in the fabrication of integrated circuits (ICs). In fact, it is the high quality of the oxide, its high density, mechanical stability and insulating properties that has enabled the production of silicon ICs to develop into the global industry that it is today. [1] Until 1990 however, this was far from the case for III-V semiconductors.

Despite extensive research into developing a similarly employable oxide in a III-V system, all oxides grown were found to be specifically detrimental to device performance. Various authors as late as mid-1990 reported on the environmental degradation of III-V devices caused by natural oxidation of AlGaAs, citing cleaved edges, cracks and fissures at layer interfaces. [2]

The first report of steam oxidation being used in a III-V system came in 1978 when Tsang reported on the self-termination of the oxidation of fully exposed AlAs layers on top of GaAs. [3] The paper reports wet oxidation being performed at temperatures of 100°C on AlAs layers of varying thickness (Figure 2.1).

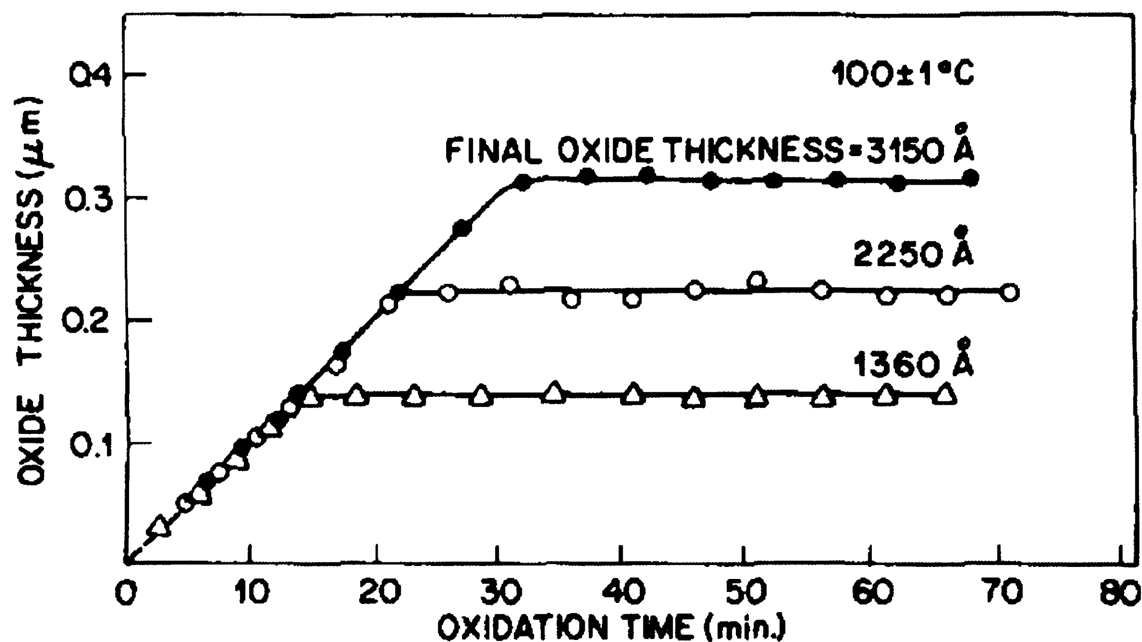

Figure 2.1: Increase in layer thickness with oxidation time. Initial layer thicknesses were 850, 1250 and 1550Å. [3]

It is important to note that with a growth temperature of only 100°C, the oxidation process here resulted in an increase in the layer thickness by a factor of up to 1.8. Tsang goes on to compare these results to dry oxidation at higher temperatures where this is not the case, and attributes the differences to the chemical make-up of the oxide produced.

Following this, it was not until late 1990 that the process was improved and a mechanically stable, high quality oxide was produced. Dallesasse et al. at the University of Illinois were the first to report this, when they performed a wet oxidation process at temperatures around 400°C. [4] Not only did they develop a successful oxidation process, but were also the first to demonstrate the controllable progression of the oxidation laterally into layers buried in an epitaxial structure, and the monitoring of that progress. The paper is also the first to describe quantitatively the rate of oxidation as proportional to the fractional aluminium content of the material, and suggests very clearly the importance of the role of hydrogen in producing good quality oxides. It should also be noted that no delamination is reported despite oxidised layers being buried in the wafer structure. This suggests that the expansion factor of 1.8 reported previously by Tsang is either reduced or removed in this case.

This seminal paper led to many others over the next 15 years refining and expanding the process. Details of the most important of these can be found in the following sections.

2.3 Chemistry of AlGaAs oxidation

It is clear from the very early work on the development of the wet oxidation process that both understanding and controlling the details of the chemistry is critical to producing repeatedly high quality oxides. In another early paper from Dallesasse et al. [5] the complexity of the issue is made clear. The paper lists Al_2O_3, $Al_2O_3{\cdot}H_2O$, $Al_2O_3{\cdot}3H_2O$, $AlAs{\cdot}8H_2O$, $AlO(OH)$, and $Al(OH)_3$ as some of the possible products of oxidation by water vapour at room temperature only.

As has now been fairly extensively investigated, the chemistry that takes place at low temperatures (100°C or below) is quite different from that during higher temperature oxidation (300°C and above). The crucial differences between the two regimes are the greater availability of ionised hydrogen at high temperatures, and the likelihood of certain key reactions, as quantified by the Gibbs free energy, being greater at high temperatures.

Clearly, under ideal conditions all of the AlAs in the oxidisable material will be fully converted to Al_2O_3. Any GaAs present may well be left unchanged, or at least will oxidise at a much lower rate, as will be discussed later. To ensure this, all other unwanted by-products of oxidation must either be removed from the sample or undergo further reaction. The primary oxidation reaction, as shown below in Equation 2.1, [6] produces As_2O_3 and H_2 as by-products. Further reaction may then occur between the two by either of the two mechanisms described by Equations 2.2 and 2.3.

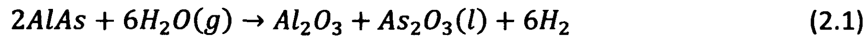

$$2AlAs + 6H_2O(g) \rightarrow Al_2O_3 + As_2O_3(l) + 6H_2 \quad (2.1)$$

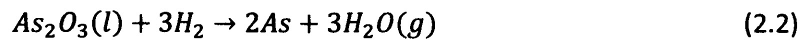

$$As_2O_3(l) + 3H_2 \rightarrow 2As + 3H_2O(g) \quad (2.2)$$

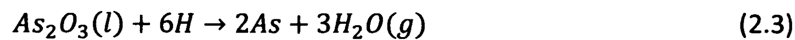

$$As_2O_3(l) + 6H \rightarrow 2As + 3H_2O(g) \quad (2.3)$$

The Gibbs free energies of these latter two equations at 425°C are given as ΔG^{698}=-131 kJ/mol and -1226 kJ/mol respectively [7], where a larger negative number implies a more energetically favourable reaction, and a positive value would imply an unfavourable one. If the temperature were lowered to 100°C, Equation 2.2 would become energetically unfavourable [8] and Equation 2.3 would not be able to proceed, as the availability of atomic hydrogen would be far less. For these reasons, at low temperatures, As_2O_3 remains within the oxidised material and contributes to the poor quality of the final layer by way of weakening and increasing it in volume, which causes cracking and delamination.

As a by-product, atomic arsenic may react further to form AsH_3, both of which are volatile enough within the relatively porous lattice of the oxide to escape from the sample and be removed with exhaust gases. [9] This transport of products out of, and reactants into, the sample can, under some circumstances, affect the oxidation rate of the material. This will be discussed in a later section.

Another factor that is important in the control of the wet oxidation of $Al_xGa_{1-x}As$ ($x \geq 0.8$) is the presence of molecular oxygen, O_2. Early tests showed conclusively that the use of oxygen as a carrier gas for the steam instead of nitrogen, far from enhancing the oxidation rate as might have been expected, actually inhibits the oxidation completely. Ashby et al. show this very clearly in their 1997 paper that compares four different carrier gases used in an oxidation process of $Al_{0.98}Ga_{0.02}As$ at 440°C. All the gases show the same oxidation rate of just under 1µm/min apart from oxygen, which displays no oxidation at all even after 900 minutes (Figure 2.2).

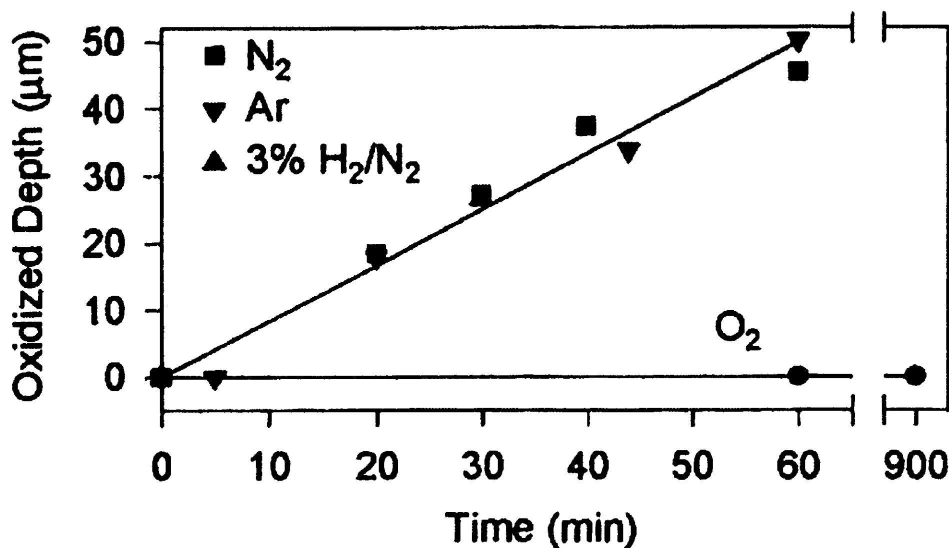

Figure 2.2: Suppression of oxidation by use of O_2 carrier gas [6]

The reason for this apparently total suppression of oxidation is given as the rapid consumption of atomic hydrogen, [6] which has already been shown to be pivotal in the reduction of As_2O_3 to As. With the suppression of the main wet oxidation mechanism, oxidation can only proceed by a mechanism associated with normal dry conditions. This has been shown to proceed at rates several orders of magnitude lower than those of wet oxidation, [4] which accounts for the results shown here.

Despite this clear result in the case of high aluminium content AlGaAs, later studies into the use of O_2 as a carrier gas revealed employable effects with material compositions approaching pure GaAs. In 2005,

Luo and Hall presented work showing that the controlled addition of trace amounts (≤ 1%) of O_2 into the normal N_2 carrier gas promoted the rate of oxidation of GaAs without retarding that of AlAs, thereby reducing the selectivity of the oxidation process in terms of material composition. [10] They report a tenfold increase in the oxidation rate of $Al_{0.3}Ga_{0.7}As$ and the formation of a dense and smooth oxide material, as characterised by atomic force microscopy and scanning electron microscopy images.

2.4 Physical and electrical properties of oxides

Strongly linked with the chemistry of the oxidation process are the physical properties of the resulting oxide layers. Characteristics such as the morphology, density and strain of the layers will greatly affect the performance of the devices containing them. Many of these characteristics have been extensively studied, and have already revealed some of the more intricate details of the oxidation process.

As early as 1996, high resolution cross-sectional TEM images were being produced of the lateral oxidation front as it proceeded through high aluminium-content AlGaAs layers. From these images, information on both the final phase of the oxide and any intermediate chemical steps can be inferred. Twesten et al. presented some of the earliest work of this type. [11] Their data, given in Figure 2.3, shows a 17nm amorphous interface between the unoxidised AlGaAs and the oxide.

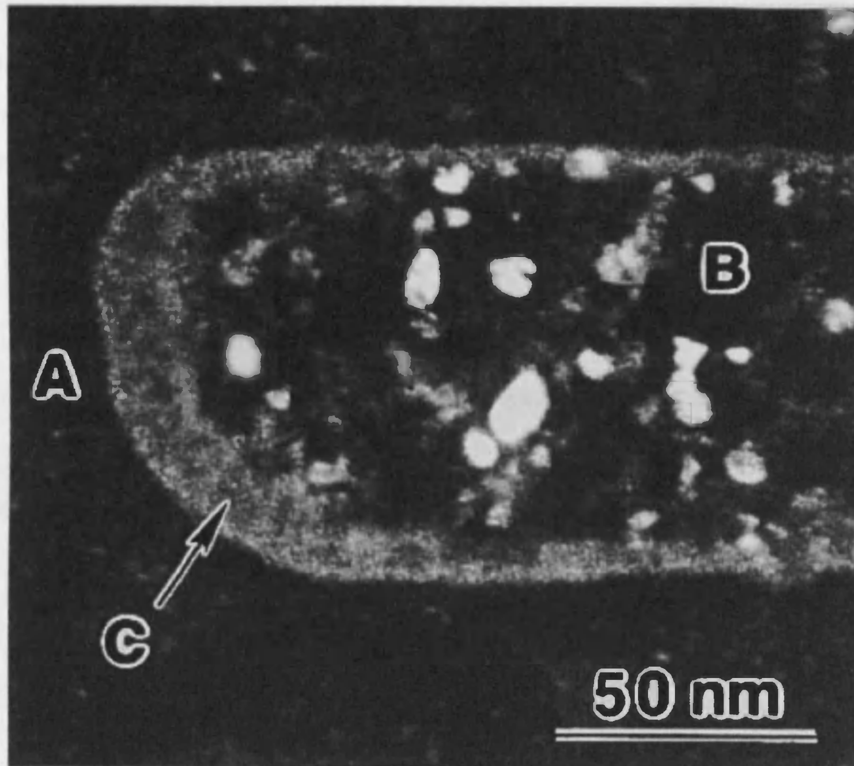

Figure 2.3: TEM darkfield image showing A) unoxidised AlGaAs B) polycrystalline Al_2O_3 and C) amorphous interface area [11]

Furthermore, diffraction patterns taken within area B show specifically that the oxide formed is polycrystalline gamma-phase Al_2O_3 (γ-Al_2O_3), rather any other crystalline structure of the same compound. Grain size is reported to be an average of 4nm with no directional growth preference relative to the unoxidised lattice.

An important comparison to make at this point is between the volume of the final oxide layers. As discussed previously, the early data presented on wet oxidation, performed at relatively low temperatures (≈100°C), showed an increase in the total volume of the material after oxidation. Conversely however, these more recent studies of oxidation at higher temperatures show a net shrinkage of layers during oxidation. Twesten et al. quantify this as a 6.7 ±0.2% linear reduction in layer thickness for 92% aluminium AlGaAs. This serves to confirm the presence of different chemical mechanisms in the two temperature regimes. When compared to the reduction in layer thickness that would be expected for a complete conversion to fully dense γ-Al_2O_3, given as 20% by Twesten, it is clear that the final oxide is not completely pure and is more porous than might be expected. However it is these factors that, in part, allow the oxide to be employable in working devices. A full shrinkage of 20%

would likely result in catastrophic delamination and fracturing, and a more dense layer would inhibit transport of reactants and products in and out of the layer, thereby hindering the oxidation process. Literature gives no evidence that gallium is able to precipitate out during the oxidation process, but various sources report different final compositions of the oxide layer. The main discrepancies seem to be between an amorphous solid solution of $(Al_xGa_{1-x})_2O_3$ and polycrystalline γ-$(Al_xGa_{1-x})_2O_3$ as described above. The reasons for these differences seem to be related to post-oxidation annealing and imaging by various electron beam methods, both of which tend to cause crystallisation into the gamma phase. [9]

It is worth noting that, although a 6.7% shrinkage in layer thickness will not usually be enough to physically damage any device employing it, it will often induce a significant amount of strain into the surrounding layers and may therefore affect the electrical and optical performance of the device.

Further properties that have now been comprehensively investigated are the resistivity and refractive index of Al_2O_3 as grown by wet oxidation of AlGaAs. Early measurements reported the room temperature resistivity of oxides grown from pure AlAs as being at least 5×10^{11} Ωcm. [12] Later, more detailed studies have shown comparable numbers and went on to show the relationship between the fractional composition of AlGaAs and the resistivity of the oxide that resulted from it. This data, from Ashby et al. is shown below.

Table 2.1: Resistivity values of various AlGaAs compositions. 'As level' is residual post-oxidation arsenic left at the semiconductor-oxide interface and leakage current is through capacitive devices of the same material [13]

Starting Al mole fraction x	As level (@5 V)	Leakage (A/cm^2)	Resistivity (Ω cm)
1	Low	7.5×10^{-8}	2.7×10^{12}
	High	7.5×10^{-7}	2.7×10^{11}
0.98	Low	2.9×10^{-9}	6.9×10^{13}
	High	2.5×10^{-7}	8.0×10^{11}
0.90	Low	3.7×10^{-9}	5.4×10^{13}
	High	1.0×10^{-7}	1.2×10^{11}

Resistivities of this order of magnitude potentially allow the oxide layers to act as effective insulators amongst doped AlGaAs layers, and as such, form effective current apertures, as will be described in a later section.

Refractive index has been documented in a similar way. As early as 1991, the index of native oxide layers had been measured as being between 1.5 and 1.6, compared to the 3.1 of the unoxidised material. [14] Again, more detailed studies followed, resulting in highly accurate measurements of the refractive index of oxides grown from various compositions of AlGaAs over a range of temperatures. In 1999, Hall et al. measured indexes ranging from 1.490 (x = 0.9, grown at 400°C) to 1.707 (x = 0.3, grown at 500°C) at 632.8nm by the prism coupling method. [15] Their data is shown in Figure 2.4. The maximum in the data clearly suggests that two processes are present in determining the refractive index. Hall et al. give these processes as the replacement of gallium with lighter aluminium atoms, leading to a reduction in index at high values of x, and the presence of light gallium hydroxide species at low values of x.

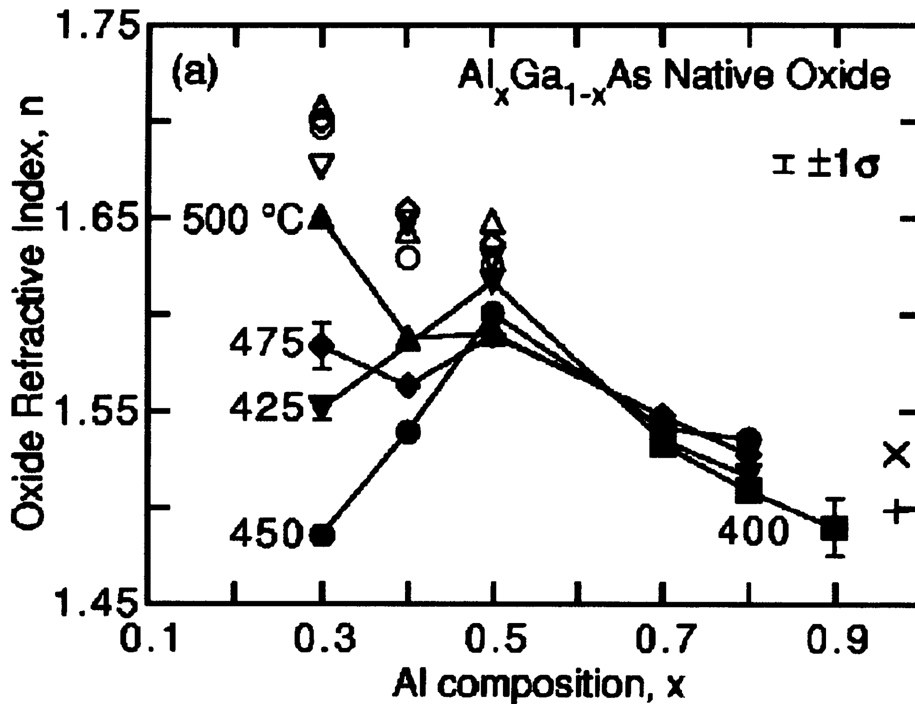

Figure 2.4: Refractive index for a variety of compositions and temperatures [15]

Again, this property has enabled the development of many new applications. The relatively low refractive index of the oxide compared to the AlGaAs it is grown from gives the potential to fabricate lateral waveguiding effects, thus controlling the optical modes in a device. This does however rely on the accurate control of the depth of lateral oxidation. This is a complicated process, which relies on a large number of variables, and is discussed in the next section.

2.5 Factors affecting oxidation rates

Many of the technologies where wet oxidation of AlGaAs has found application rely on lateral oxidation of layers to a controlled depth - for example, the fabrication of current apertures in VCSELs. To achieve this, the rate of oxidation must be made highly predictable by the investigation and control of all the variables that affect it. Amongst these variables are temperature, gas flow, gas composition, fractional AlGaAs composition and layer thickness. This section aims to detail all of the controlling parameters and to outline typical oxidation rates achievable across their range.

2.5.1 Sample temperature and $Al_xGa_{1-x}As$ composition

As has already been shown, the temperature of the sample during oxidation has a profound effect on the chemistry of the process. The depth of lateral oxidation has been shown by various sources to increase in either a linear, [16] [17] or parabolic fashion [18] [19] depending on whether it is limited by the reaction rate or the diffusion rate respectively. Generally, at higher temperatures reaction rates are greatly increased, so the oxidation is limited by diffusion of the reactants into the sample and of the products out. Clearly, with increased oxidation depth, this diffusion will take longer, and so oxidation depth will tend to progress in a parabolic fashion to a certain limit. At lower temperatures, reaction rates are much lower, and so determine the overall oxidation rate. As reaction rates do not drop with increased oxidation depth, a linear progression results.

Under these conditions, it can be logically inferred that any lateral oxidation process, if given enough time and depth of material, will eventually reach the point where diffusion becomes slower than the reaction and a parabolic limit is reached. Equally, any oxidation can be kept within the linear regime by lowering the temperature, as long as the reaction chemistry does not suffer. It is therefore solely temperature that determines when the transition from the linear regime to the parabolic one occurs, and thus, how great an oxidation depth can ultimately be achieved.

Numerous sources have reported data showing exactly how oxidation rate varies with temperature. The vast majority demonstrate an exponential increase with increasing temperature. Typical data is given by Choquette et al. [9]

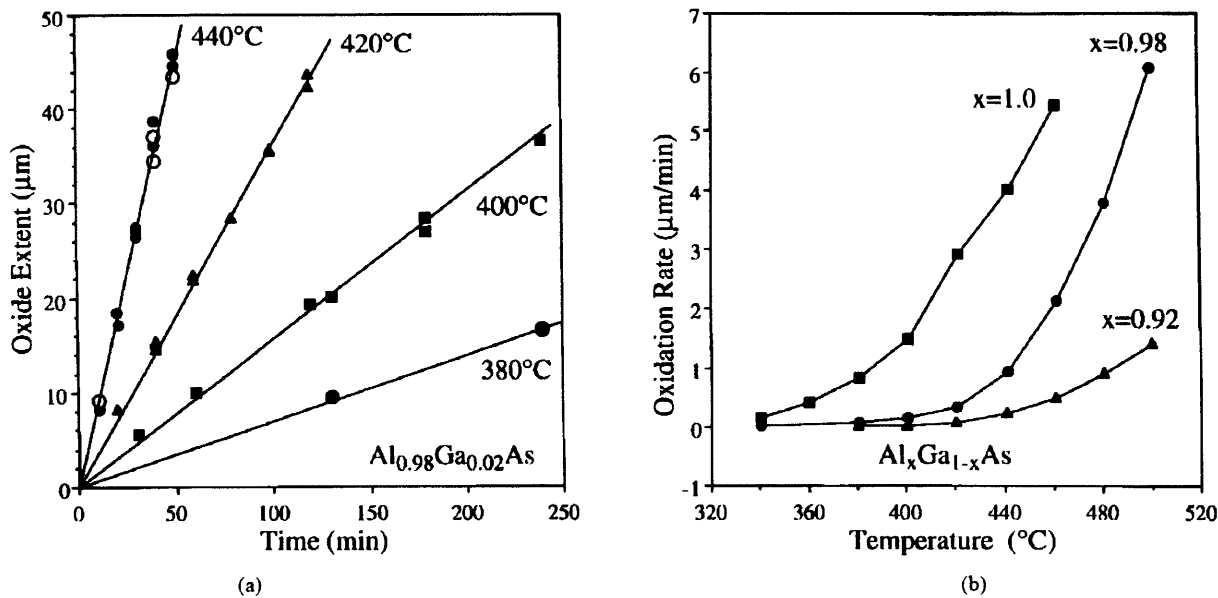

Figure 2.5: Variation of oxidation rate with (a) time and (b) sample temperature [9]

From this data it can be seen that even under moderate conditions ($x = 0.98$ and $T \leq 500°C$) rates of approximately 5μm/min can be achieved. This is quite acceptable for most applications where depths of only a few microns are likely to be desired. Comparing this data to the older results of oxidation at 100°C in Figure 2.1, which shows an oxidation rate of only 0.01μm/min, confirms both the exponential relationship and the benefits of oxidising at higher temperatures. Furthermore, comparing to data presented in literature on the rate of oxidation of pure GaAs serves to highlight the importance of the presence of aluminium and the selectivity achievable when oxidising layered structures. Oh et al. gave the rate of oxidation of GaAs at 500°C as 0.005μm/min, [20] more than three orders of magnitude smaller than that of AlAs at the same temperature.

This selectivity between different compositions of $Al_xGa_{1-x}As$ has become the basis of many of the applications to employ wet oxidation. Figure 2.6 shows typical data displaying the relationship between composition and oxidation rate. Given the logarithmic y-axis, the drop in oxidation rate with decreasing aluminium content is very steep, with an order of magnitude drop between $x = 1.00$ and $x = 0.95$. Although the drop becomes less severe at lower values of x, the aluminium fraction remains a critical factor in achieving a useful oxidation rate.

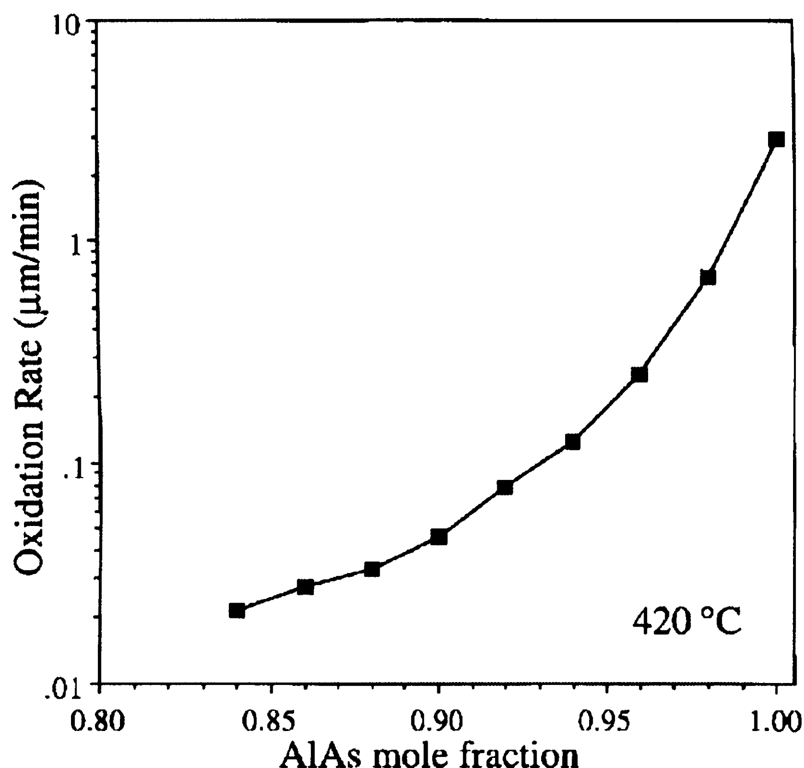

Figure 2.6: Variation of oxidation rate with $Al_xGa_{1-x}As$ composition [9]

Further to this, the compositions of the layers surrounding the layer in question also have an influence over oxidation rate. It has been shown that burying a thin 98% aluminium layer in the middle of a thick 92% layer will enhance the oxidation rate of the 92% material by an amount proportional to the proximity to the 98% layer by way the cross-diffusion of oxidant species, thus producing a tapered oxidation front rather than a square one as would otherwise result. [21] Similarly, multiple AlGaAs layers spaced with GaAs barrier layers will oxidise faster as the width of the barrier layers is decreased. [22] Both of these effects are attributed to vertical diffusion of oxidant species between AlGaAs layers contributing to the total volume of oxidised material, thereby increasing the lateral rate.

2.5.2 Layer thickness

In the case of individual Al(Ga)As layers, completely isolated from any other oxidisable layers, a definite dependence on the layer thickness has been reported by numerous sources over recent years. Data presented by multiple authors [23] [24] show good agreement, and suggest that oxidation rate increases only marginally for layers above 200nm in thickness, but a sharp decrease is observed for those below

80nm. Choquette presents this data for oxidation of AlAs at 400°C and shows a decrease of two orders of magnitude between 70nm and 10nm thick layers.

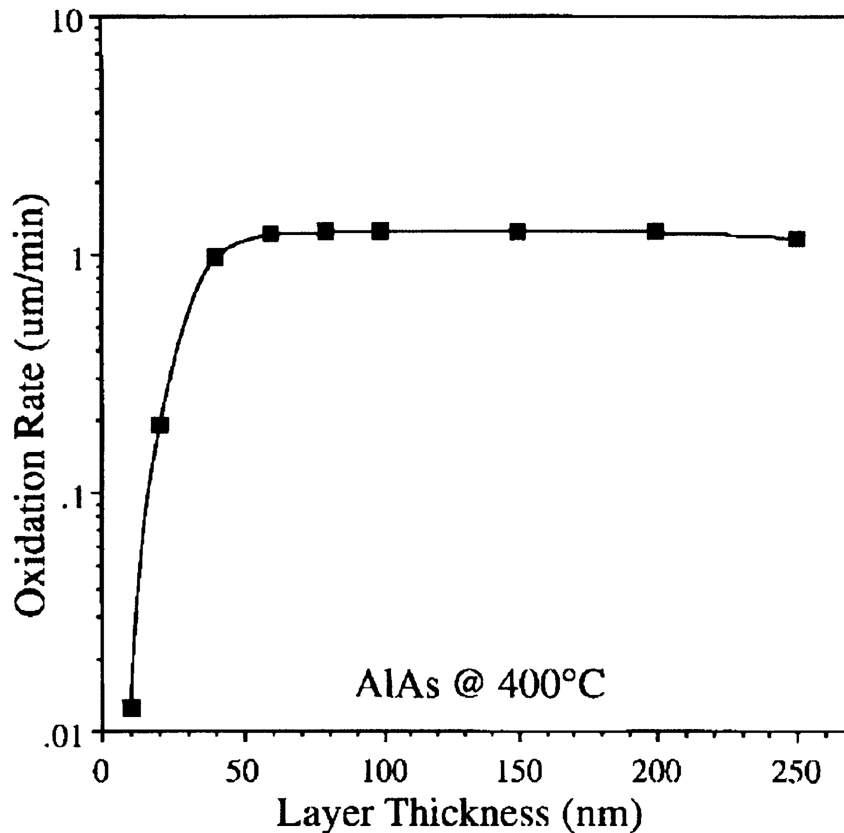

Figure 2.7: Variation of oxidation rate with layer thickness [9]

By way of explanation of this effect, Naone and Coldren present a detail theoretical model and corroborating experimental results based on the increased curvature of the oxidation front in very thin layers. [25] Based on the Gibbs-Thomson effect of increased vapour pressure across highly curved surfaces, they predict greatly increased activation energies of the intermediate, rate-determining reaction steps occurring at the oxide/AlGaAs interface.

2.5.3 Doping

Although it seems little research has been performed to determine the effect of dopant species on oxidation, it is nevertheless an important factor in a system where the speed of oxidation can be governed purely by the chemical reaction rate. The chemical definition of oxidation (increasing the oxidation state) is 'the loss of electrons', and the rate of any chemical reaction in a closed environment is strongly affected by the concentrations of reactants and products present. Given the multiple stages

involved in the chemical process of oxidation (Equations 2.1 to 2.3), it is hard to determine whether the overall effect would be a tendency to generate or use up electrons. However, if the concentration of either electrons or holes in the immediate environment is already high, it is quite conceivable that the reaction rate could be altered, one way or the other.

A doping dependence was observed in 1992 by Kish et al. during high temperature surface oxidation of p- ($\approx 9x10^{18}cm^{-3}$) and n-doped ($\approx 4x10^{17}cm^{-3}$) samples of $Al_{0.6}Ga_{0.4}As$, where the p-type samples oxidised approximately four times faster than the n-type. [18] They found no compositional differences between the two oxides after growth, but attributed their results to diffusion rates, as will be discussed again later. To date, no specific data has been presented on the lateral oxidation rate of buried AlGaAs layers.

2.5.4 Carrier gas flow and water bubbler temperature

As has already been discussed as part of the reaction chemistry section, the carrier gas chosen can affect the oxidation process. The use of pure oxygen will completely suppress any oxidation, whereas the use of a very small percentage of oxygen in a nitrogen carrier will greatly reduce the AlGaAs composition selectivity by increasing the rate of oxidation of pure GaAs.

Equally as important is the rate of flow of the chosen carrier gas. This, coupled with the temperature of the water the gas is bubbled through, can be optimised to achieve the maximum oxidation rate. As Geib at al. reported in 1997, [26] keeping the bubbler temperature just below 100°C is beneficial. Increasing the flow rate is reported to increase the oxidation rate up to a certain point, where the oxidation rate appears to saturate. This can be explained by considering the consumption of reactants by the oxidation reaction. At any given temperature, assuming the process is reaction rate limited, the oxidant species will be consumed at a fixed rate. Once the rate of supply of this species has equalled the rate of consumption, no further gain can be made by supplying the oxidant at a faster rate, so the oxidation rate saturates. This data is shown in Figure 2.8. The increased 'saturated' rate at the higher bubbler temperature, as compared to the lower, can be explained by the higher partial pressure of the steam.

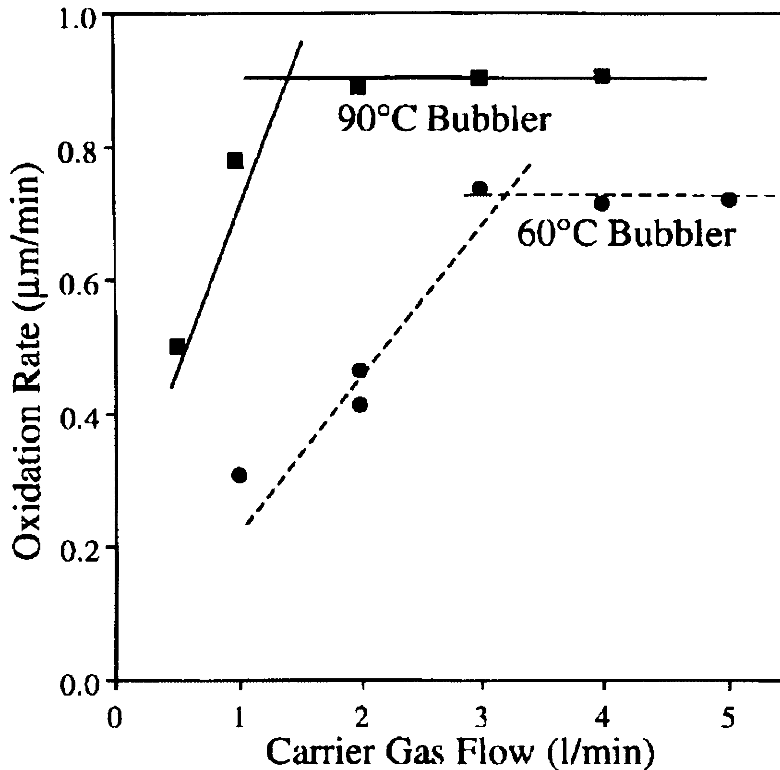

Figure 2.8: Variation of oxidation rate with carrier gas flow at different bubbler temperatures [26]

2.6 Summary of applications

Since the development of the process forming mechanically stable, high quality oxides from III-V material, many different researchers have employed the technique in a wide variety of applications. Although reports of these developments are now numbering in the hundreds, most can be categorised into one of a few more general areas. This sections aims to present the more important of these areas, and to illustrate them briefly with a few examples.

2.6.1 Current apertures

Probably the most immediately obvious and most common use of lateral oxidation is the formation of buried current apertures. This technique is both very simple to realise and highly beneficial in terms of threshold currents of laser devices. The basic premise is that, by careful wafer design and controlled, selective oxidation of high aluminium content layers, current injected into a broad-area laser can be funnelled into the centre of the device by the insulating properties of the oxide formed at the devices' edges. With the correct placement of the oxidisable layers within a laser structure, this can lead to an

effective increase in the current density in the active region and thus, a reduction in the threshold current of the device.

This technique has been used extensively in VCSELs, where oxidising inwards in all directions simultaneously is ideal due to the often cylindrical structure of the devices. [27] Recently, however, current apertures have begun to be investigated in planar devices as well. In 2000, Park et al. demonstrated successful current apertures in very wide ridge devices by oxidising prior to cleaving the laser facets, thereby keeping the facets free of oxide and confining the current in the lateral direction only. [28] Figure 2.9 is a schematic of the device they created including a sketch of the band diagram, which demonstrates that the inclusion of the oxidisable layer produces only a small barrier in the central, unoxidised region, and so will only affect the electrical properties of the device very minimally.

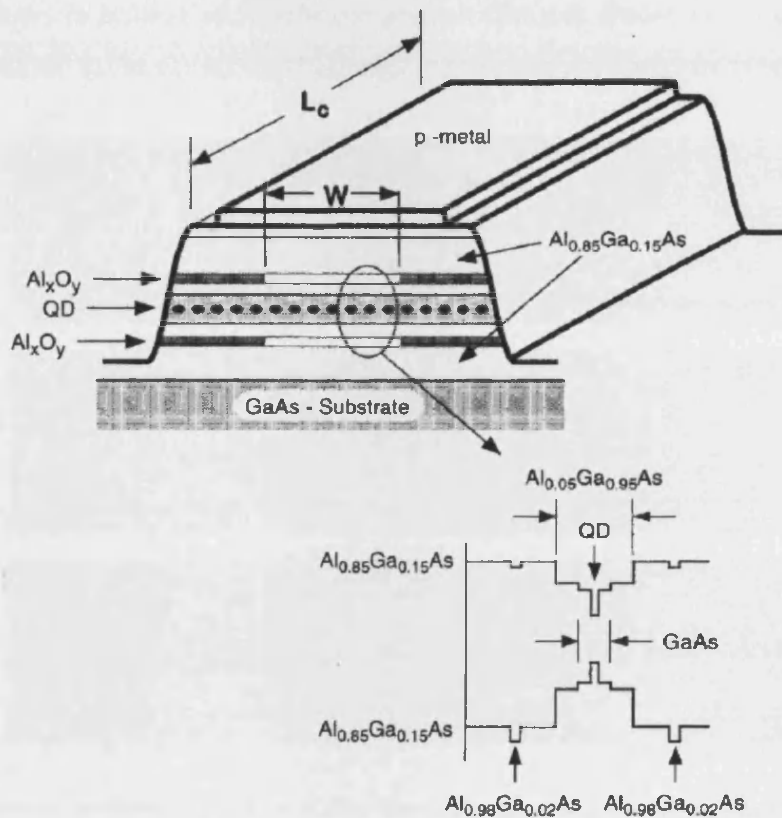

Figure 2.9: Lateral current confinement in a planar laser device [28]

2.6.2 Waveguiding

The difference in refractive index between Al_2O_3 and its unoxidised parent material is another widely employed property of the wet oxidation system. With oxide indices reported between 1.4 and 1.6, there

is great potential for creating strong lateral waveguiding systems in GaAs-based devices with typical indices of 3.0 to 3.4.

Again, this has been widely employed in VCSEL systems where optical modes can be confined in much the same way as the current confinement described in the previous section. Several more exotic applications have been developed over recent years, however, that make use of the same property. In 1993, Caracci et al. used the surface-grown wet oxide of $Al_{0.85}Ga_{0.15}As$ to fabricate waveguide S-bends with significantly lower loss than the same shaped waveguides fabricated by impurity-induced disordering. [29] This was a large step towards optical integration of multiple devices, as the waveguides could now be made on the same substrate as laser chips and other passive devices.

More recently, in 2006, Chang et al. used the wet oxidation of a chirped short-period superlattice of AlAs and GaAs layers to achieve an AlGaAs composition that was graded in the direction of growth. When this was oxidised to the correct depth, circular microlenses were produced (Figure 2.10).

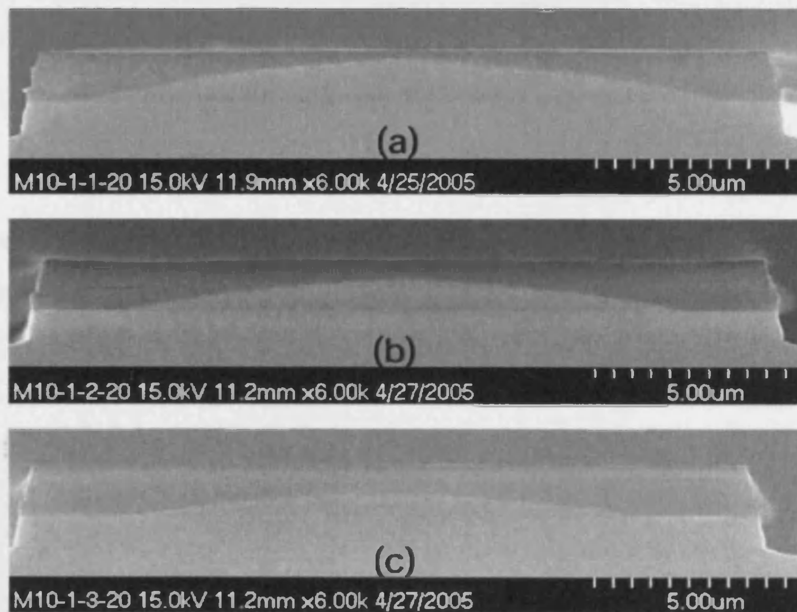

Figure 2.10: Microlenses fabricated by oxidation of AlAs-GaAs superlattice (a) 20 minute oxidation (b) 30 minutes (c) 45 minutes [30]

In this case, oxidation provides a cheap, repeatable and single-step method of producing very useful passive components, which were previously fabricated using much more complex and expensive methods.

2.6.3 Reflectors and wavelength control

The relatively low refractive index of Al_2O_3 also allows the material to be used as a reflector in the correct situation. In 1995, Huffaker et al. published work demonstrating the use of quarter-wavelength thick, fully oxidised AlAs layers as reflectors underlying standard resonant cavity LEDs. They report an increase in differential light output efficiency to ≈5.5% as compared to ≈2.5% without the oxide reflector. [31]

A similar setup was used in 2002 by Macaluso et al. to tune the resonant wavelength of a vertical microcavity. [32] A thick $Al_{0.7}Ga_{0.3}As$ layer was placed under the normal microcavity structure, with an AlAs 'supply layer' below that. As the AlAs oxidised much faster than the 70% aluminium layer but remained porous to the oxidant materials, the $Al_{0.7}Ga_{0.3}As$ layer above oxidised vertically as well as laterally from the edges of the device. This effectively gave a moveable mirror at the bottom of the resonant cavity, which could be used to tune the wavelength over a range of approximately 52nm. Again, this would have previously been achieved by the much more complex and expensive method of micro-electromechanical moveable external mirrors, but the use of oxidation allows a simple, repeatable, on-chip solution.

2.6.4 Surface roughness reduction

The presence of high surface roughness can be seriously detrimental to optical device performance. A very rough surface will reflect any incident light in an unpredictable manner and will therefore increase both waveguide bend losses and general scattering loss. In the case of narrow-etched ridge lasers, it is the semiconductor/air interface itself that is designed to act as the laterally confining waveguide and, as such, interacts strongly with the propagating optical mode. Furthermore, the etching method that is most often used to fabricate these narrow structures is a dry, physical one, and so will tend to produce rougher surfaces than a chemical etch. These two factors together mean that surface roughness becomes one of the primary sources of optical loss in such devices.

Whilst it is true that the parameters governing the dry etching process can be optimised to reduce surface roughness, this is only ever going to afford minor improvements. In 2007, Liang and Hall used the oxygen-enhanced non-selective wet oxidation method described in Section 2.3 to solve this problem. [33] By subjecting their AlGaAs ridges to the non-selective oxidation process, they created an oxide barrier of approximately 200nm between the semiconductor and the air. The resulting

semiconductor/oxide interface, when viewed in cross-section in an SEM (Figure 2.11), was seen to be far smoother than the etched surface of the same sample (σ ≈ 1-2nm rather than ≈ 100nm).

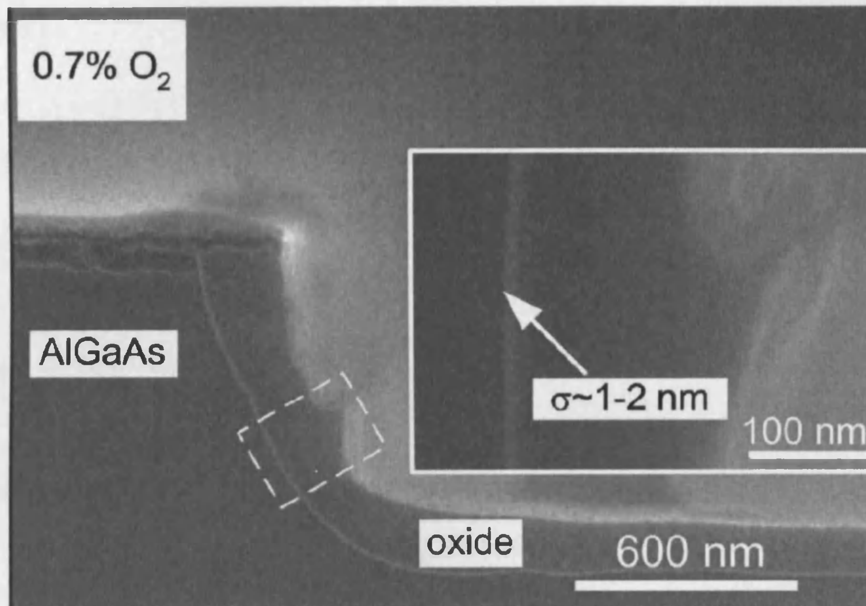

Figure 2.11: Decrease in surface roughness by non-selective wet oxidation [33]

This reduction of effective surface roughness of around 100 times enables a high index contrast lateral waveguide with low bend and scattering loss to be formed simply and cheaply, which will in turn, lead to high-quality, single-mode laser emission with a lower threshold current density.

2.6.5 Phosphide-based materials

Although the AlGaAs material system is the one that has been most extensively and productively researched because of its high oxidation rates and good quality, porous oxides, there has also been a certain degree of success reported in work on the oxidation of phosphide-based materials such AlGaInP. It is generally noted that oxidation rates in this material system are much lower and much less selective with respect to the aluminium content, but with the correct oxidation parameters, good quality oxides have been achieved.

In 1998, Sun and Treat reported on the successful manufacture of a native-oxide-confined ridge waveguide laser emitting at 670nm. [34] The oxide was grown from p-doped $Al_{0.5}In_{0.5}P$ 3.5µm ridge in much the same manner as described in the previous section, and resulted in threshold currents below 13.5mA and differential quantum efficiencies of 35%/facet.

2.7 References

1. *The silicon insulated-gate field-effect transistor.* **Hofstein, S. R.; Heiman, F. P.** 9 : Proceedings of the IEEE, 1963, Vol. 51.

2. *Environmental degradation of AlGaAs-GaAs quantum-well heterostructures.* **Dallesasse, J. M.; El-Zein, N.; Holonyak, N.; Hsieh, K. C.; Burnham, R. D.; Dupuis, R. D.** 5 : Journal of Applied Physics, 1990, Vol. 68.

3. *Self-terminating thermal oxidation of AlAs epilayers grown on GaAs by molecular beam epitaxy.* **Tsang, W. T.** 5 : Applied Physics Letters, 1978, Vol. 33.

4. *Hydrolyzation oxidation of AlGaAs-GaAs quantum well heterostructures and superlattices.* **Dallesasse, J. M.; Holonyak, N.; Sugg, A. R.; Richard, T.A.; El-Zein, N.** 26 : Applied Physics Letters, 1990, Vol. 57.

5. *Stability of AlAs in AlGaAs-AlAs-GaAs quantum well heterostructures.* **Dallesasse, J. M.; Gavrilovic, P.; Holonyak, N.; Kaliski, R. W.; Nam, D. W.; Vesely, E. J.; Burnham, R. D.** 24 : Applied Physics Letters, 1990, Vol. 56.

6. *Wet oxidation of AlGaAs: The role of hydrogen.* **Ashby, C. I. H.; Sullivan, J. P.; Choquette, K. D.; Geib, K. M.; Hou, H. Q.** 6 : Journal of Applied Physics, 1997, Vol. 82.

7. **Kubaschewski, O.; Alcock, C. B.; Spencer, P. J.** *Materials Thermochemistry.* UK : Pergamon, 1993.

8. *Wet oxidation of high-Al-content III-V semiconductors: Important materials considerations for device applications.* **Ashby, C. I. H.** Materials Research Meeting, 1999.

9. *Advances in selective wet oxidation of AlGaAs alloys.* **Choquette, K. D.; Geib, K. M.; Ashby, C. I.; Twesten, R. D.; Blum, O.; Hou, H. Q.; Follstaedt, D. M.; Hammons, B. E.; Mathes, D.; Hull, R.** 3 : Journal of Selected Topics in Quantum Electronics, 1997, Vol. 3.

10. *Nonselective wet oxidation of AlGaAa heterostructure waveguides through controlled addition of oxygen.* **Luo, Y.; Hall, D. C.** 6 : Journal of Selected Topics in Quantum Electronics, 2005, Vol. 11.

11. *Microstructure of laterally oxidized AlGaAs layers in vertical-cavity lasers.* **Twesten, R. D.; Follstaedt, D. M.; Choquette, K. D.; Schneider, R. P.** 1 : Applied Physics Letters, 1996, Vol. 69.

12. *Properties of Al2O3 optical coatings on GaAs produced by oxidation of epitaxial AlAs/GaAs films.* **Schubert, E. F.; Passlack, M.; Hong, M.; Mannerts, J.; Opila, R. L.; Pfeiffer, L. N.; West, K. W.; Bethea, C. G.; Zydzik, G. J.** 22 : Applied Physics Letters, 1994, Vol. 64.

13. *Wet oxidation of AlGaAs: Temporal evolution of compostion ond microstructure and the implications for metal-insulator-semiconductor applications.* **Ashby, C. I.; Sullivan, J. P.; Newcomer, P. P.; Missert, N. A.; Hou, H. Q.; Hammons, B. E.; Hafich, M. J.; Baca, A. G.** 18 : Applied Physics Letters, 1997, Vol. 70.

14. *Planar native-oxide index-guided AlGaAs-GaAs quantum well heterostructures.* **Kish, F. A.; Caracci, S. J.; Holonyak, N.; Dallesasse, J. M.; Hsieh, K. C.; Ries, M. J.; Smith, S. C.; Burnham, R. D.** 14 : Applied Physics Letters, 1991, Vol. 59.

15. *Refractive index and hygroscopic stability of AlGaAs native oxides.* **Hall, D. C.; Wu, H.; Kou, L.; Epstein, R. J.; Blum, O.; Hou, H.** 8 : Applied Physics Letters, 1999, Vol. 75.

16. *Wet thermal oxidation of AlGaAs compounds.* **Burton, R. S.; Schlesinger, T. E.** Journal of Applied Physics, 1995, Vol. 76.

17. *A detailed experimental study of the wet oxidation kinetics of AlGaAs layers.* **Nickel, H.** Journal of Applied Physics, 1995, Vol. 36.

18. *Dependence on doping type (p/n) of the water vapour oxidation of high gap AlGaAs.* **Kish, F. A.; Maranowski, S. A.; Höfler, G. E.; Holonyak, N.; Caracci, S. J.; Dallesasse, J. M.; Hsieh, K. C.** 25 : Applied Physics Letters, 1996, Vol. 60.

19. *Kinetics of thermal oxidation of AlAs in water vapour.* **Ochiai, M.; Giudice, G. E.; Temkin, H.; Scott, J. W.; Cockerill, T. M.** 14 : Applied Physics Letters, 1996, Vol. 68.

20. *Steam oxidation of GaAs.* **Oh, T. -H.; Huffaker, D. L.; Graham, L. A.; Deng, H.; Deppe, D. G.** 21 : Electronics Letters, 1996, Vol. 32.

21. *Oxidation of AlGaAs layers for tapered apertures in vertical- cavity lasers.* **Naone, R. L.; Hegblom, E. R.; Thibeault, B. J.; Coldren, L. A.** 4 : Electronics Letters, 1997, Vol. 33.

22. *Barrier-layer-thickness control of selective wet oxidation of AlGaAs for embedded optical elements.* **Blum, O.; Ashby, C. I.; Hou, H. Q.** 21 : Applied Physics Letters, 1997, Vol. 70.

23. *Lateral wet oxidation of AlGaAs-GaAs depending on its structures.* **Kim, J. -H.; Lim, D. H.; Kim, K. S.; Yang, G. M.; Lim, K. Y.** 22 : Applied Physics Letters, 1996, Vol. 69.

24. *Dependence of lateral oxidation rate on thickness of AlAs layer of interest as a current aperture in vertical-cavity surface-emitting laser structures.* **Koley, B.; Dagenais, M.; Jin, R.; Simonis, G.; Pham, J.; McLane, G.; Johnson, F.; Jr., R. Whaley.** 1 : Journal of Applied Physics, 1998, Vol. 84.

25. *Surface energy model for the thickness dependence of the lateral oxidation of AlAs.* **Naone, R. L.; Coldren, L. A.** 5 : Journal of Applied Physics, 1997, Vol. 82.

26. *Fabrication issues of oxide-confined VCSELs.* **Geib, K. M.; Choquette, K. D.; Hou, H. Q.; Hammons, B. E.** SPIE, 1997.

27. *Low threshold voltage vertical-cavity lasers fabricated by selective oxidation.* **Choquette, K. D.; Jr, R. P. Schneider; Lear, K. L.; Geib, K. M.** 24 : Electronics Letters, 1994, Vol. 30.

28. *Low-threshold oxide-confined 1.3-μm quantum-dot laser.* **Park, G.; Shchekin, O. B.; Huffaker, D. L.; Deppe, D. G.** 3 : IEEE Photonics Technology Letters, 2000, Vol. 13.

29. *Native-oxide-defined low-loss AlGaAs-GaAs planar waveguide bends.* **Caracci, S. J.; Krames, M. R.; Holonyak, N.; Herzinger, C. M.; Crook, A. C.; DeTemple, T. A.; Besse, P. -A.** 16 : Applied Physics Letters, 1993, Vol. 63.

30. *Microlens fabrication by selective oxidation of composition-graded digital alloy AlGaAs.* **Chang, K. S.; Song, Y. M.; Lee, Y. T.** 1 : IEEE Photonics Technology Letters, 2006, Vol. 18.

31. *Resonant cavity light emitting diode with an AlxOy/GaAs reflector.* **Huffaker, D. L.; Lin, C. C.; Shin, J.; Deppe, D. G.** 23 : Applied Physics Letters, 1995, Vol. 66.

32. *Resonant wavelength control of a 1.3 μm microcavity by intracavity steam oxidation.* **Macaluso, R.; Robert, F.; Bryce, A. C.; Calvez, S.; Dawson, M. D.** Semiconductor Science and Technolgy, 2002, Vol. 18.

33. *Reduction of etched AlGaAs sidewall roughness by oxygen-enhanced wet thermal oxidation.* **Liang, D.; Hall, D. C.** Applied Physics Letters, 2007, Vol. 91.

34. *Characteristics of native-oxide-confined InGaP-AlGaInP quantum-well ridge waveguide lasers.* **Sun, D.; Treat, D. W.** 4 : IEEE Photonics Technology Letters, 1998, Vol. 10.

Chapter 3 : Experimental Techniques

3.1 Introduction

This chapter outlines the main experimental techniques used throughout the project. This includes the experimental setup of the steam oxidation process as well as details of structures used for the various experiments and the equipment used to measure them. The results presented in this work were taken using a wide variety of equipment and techniques. Atomic Force Microscopy (AFM), Scanning Electron Microscopy (SEM), IVLT (current, voltage, light, temperature) kits, a streak camera and the multisection technique, are the majority of the means used, and all of these are detailed here. Consideration is given in this chapter not only to the methods used but to the accuracy and suitability of the results obtained.

As a brief overview, all the samples used throughout this work have been through the same basic lifecycle. Samples are prepared on an area of material roughly one centimetre squared. This is then cleaved up into between 20 and 100 devices, depending on the length of device required. In the case of devices to be oxidised, the oxidation process may take place before or after cleaving, depending on the particular structures used and the desired location of the oxide. Once cleaved, devices are mounted onto copper heat-sink blocks on standard 'transistor-style' headers, which allow for single or multiple electrical connections to be made on each chip. Once in this state, devices can be interfaced with any of a number of measurement systems, some of which are detailed below.

3.2 Steam Oxidation

As has been outlined in Chapter 1, the technique of steam oxidation involves the control of several variables to achieve repeatable oxidation depths. The process involves bubbling pressurised nitrogen gas, which in this case comes from a standalone gas cylinder, through a flask of water, which is heated to boiling point. In the setup used in this project, shown in Figure 3.1, a glass diffuser is used to create smaller bubbles. This maximises the surface area of nitrogen in contact with the water and thus allows more steam to be carried through the system.

The steam is carried with the nitrogen into the main chamber, which is itself heated above 100°C to prevent the steam recondensing. The chamber comprises a long glass tube with the devices to be

oxidised sitting on a separate heating block at the furthest end. The chamber is purposefully made long and narrow to promote more uniform, laminar flow of gases and therefore a more controllable oxidation process.

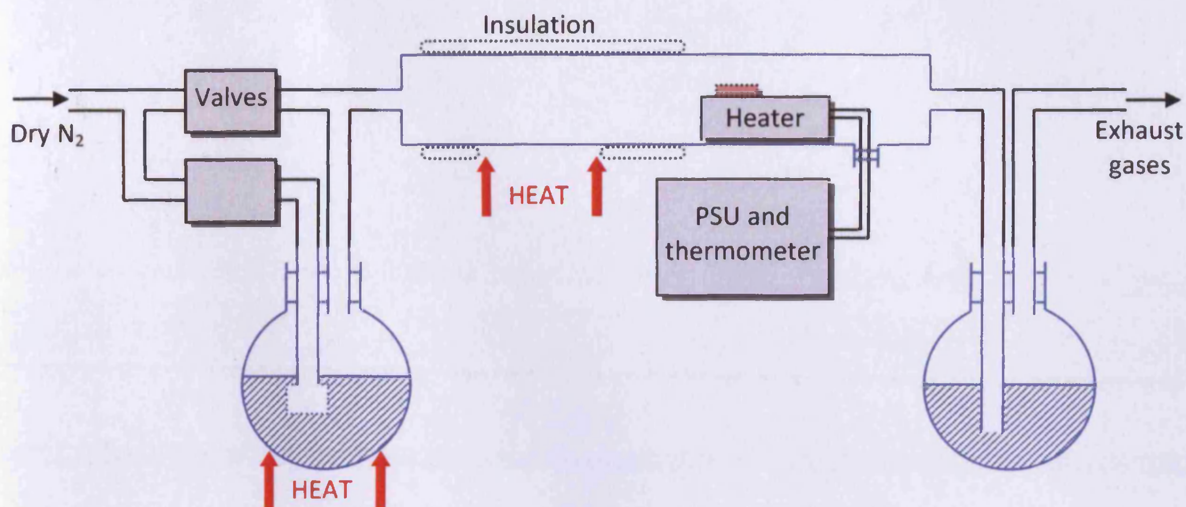

Figure 3.1: Steam oxidation experimental setup

The devices themselves rest on a heating block, which is powered with a dc power supply and monitored with an embedded thermocouple. During this project devices were elevated to temperatures ranging from 350 to 480°C. Devices were always orientated parallel to the direction of steam flow during oxidation, again to promote laminar flow and a more controlled oxidation process.

Finally, the exhaust gases, a combination of steam, nitrogen and expelled arsenide products are passed out of the main chamber and through a water bubbler. This time the water in the flask is not heated and serves to dissolve the arsenic, preventing it from being passed into the atmosphere.

Throughout the whole process the gas flow is kept constant by way of monitoring a flow meter at the input to the main chamber. As part of this project the valves shown in Figure 3.1 were added to allow rapid switching in and out of the steam on the nitrogen carrier gas. This enabled much more accurate control of the oxidation time, which was previously controlled only by the heating and cooling of the device heater block. In the new system, devices are brought both up to oxidation temperature and back down to room temperature in an inert atmosphere of dry nitrogen, thus minimising the effect of atmospheric oxygen on the oxidation process.

Current limitations of the kit are firstly, a tendency for steam to condense in the tubes before the main chamber causing turbulent flow as the water periodically partially blocks the tubes. This can be manually overcome by monitoring the process and ensuring any water is tipped into the main chamber where it will evaporate again. Secondly, oxidation cannot reliably be performed for times under one minute as it takes approximately this length of time for the steam to fully propagate through the system. This problem is due partly to the total distance the steam has to travel and partly due to the tubes in front of the main chamber being cold at the start of the process. For this project this has not been a significant problem as most devices have been oxidised for times of five minutes or longer, however rebuilding the kit to enable all of the tubes to be heated above 100°C would significantly reduce the effects of both of the problems discussed here.

3.3 IVLT measurements

Determining the threshold current of devices will be pivotal throughout this work and, in the vast majority of cases, it will be found using a custom-built apparatus, hereon referred to as the IVLT kit. This kit was developed as an earlier M.Phil project in Cardiff University. [1]

The kit allows IV (current-voltage) and LI (light-current) curves to be recorded electronically, either individually at room temperature, or in an automated fashion over a flexible range of temperatures. Figure 3.2 shows a schematic diagram of the constituent parts of the kit, starting with the signal generator at the top. The generator provides both the pulsed power to the device under test and the trigger signal with which the whole kit is synchronised. For the results taken within this work the signal generator is operated with a repetition rate of 1kHz and with pulse lengths varying from 500ns to 3µs, depending on the optical detector being employed.

The output amplitude of the pulse generator is left at maximum and fed through a voltage stepper box (a potentiometer is turned by a stepper motor), which allows both manual and automatic computer control of the signal amplitude fed to the laser under test. This in turn passes through a signal routing box that applies a voltage to either of the two possible connections on the standard TO5 header that all devices are mounted on, enabling two devices to be tested without changing samples if desired. This same routing box also directs information relating to measurements of the voltage and the current passing through the device (current is converted into a voltage signal via a transformer) onwards to the boxcar.

Also fed into the boxcar is the light signal, which comes directly from the detector placed on top on the device chamber. Depending on the specific purpose of each measurement the detector will either be an InGaAs or germanium photodiode operated in the photocurrent regime. The InGaAs detector has a shorter response time and so can operate at the lower end of the pulse length range but only has an active area of 1mm diameter. The germanium detector has a 5mm diameter active area and so can be aligned more repeatably and will collect a greater fraction of light emitted but it is much slower, requiring pulses up to 3μs where device self-heating must be taken into consideration.

The boxcar comprises of three nominally identical channels, which deal with the light, current and voltage signals separately. Each channel is triggered off the signal generator and has a variable 'gate', allowing the user to fine tune from exactly which part of the pulse the final data is taken. The boxcar effectively performs a 'sample and hold' task, averaged over many pulses, on each of the three signals and passes this dc voltage to an analogue to digital converter (ADC) in the PC, where the information is collated, stored and displayed as IV and LI graphs.

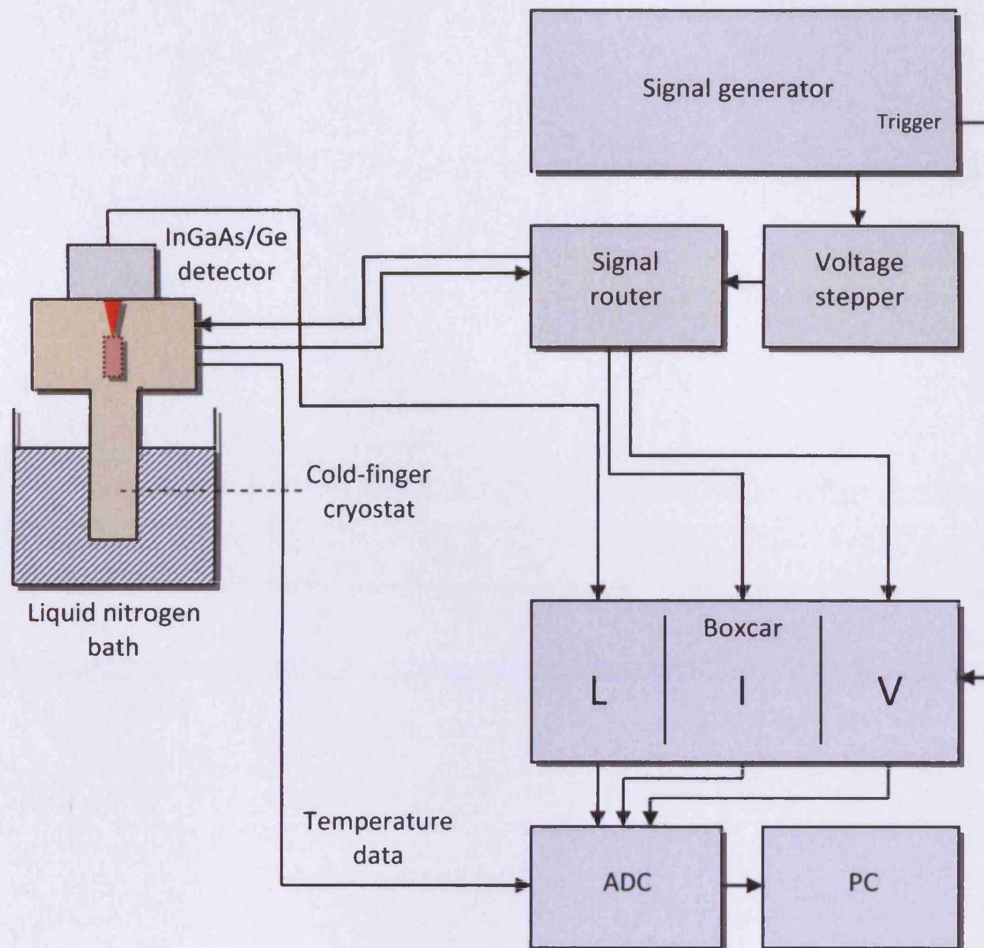

Figure 3.2: IVLT kit schematic diagram

This whole system can be automated to take results over a range of temperatures by the use of the cryostat device chamber. The device is kept under vacuum and is cooled down by way of a cold finger submerged in liquid nitrogen bath. Once below the lowest temperature desired, the liquid nitrogen is removed and the device is warmed slowly with the aid of a computer-controlled heater within the chamber. The chamber temperature is monitored by way of measuring the voltage across a known standard diode under a constant current within the chamber. Temperature intervals between measurements, current and light intensity limits can all be set as desired before a temperature run.

3.4 Atomic Force Microscopy (AFM)

One of the principal methods of determining oxidation depth used in this work will be cross-sectional AFM imaging of freshly cleaved, oxidised samples. AFM microscopes can be set up in a variety of ways to acquire a number of different sorts of data but, in the case of this work, all the results taken will be topographical facet images. General information on the workings of an AFM can be found in the original patent. [2]

Simple cross-sectional Scanning Electron Microscope (SEM) images could also be taken to measure the oxidation depth, and indeed some such images are included in this work, but the maximum resolution of conductivity-based SEM images of this nature is approximately 30nm, which is the size of the thickest oxidisable layers used here. AFM imaging has the advantage of being able to achieve resolutions below 1nm under optimum conditions.

The effect being exploited by the AFM to measure the oxidation depth is that, if exposed to the atmosphere, high aluminium-content AlGaAs will oxidise naturally and grow proud of the surrounding material, even at room temperature. If, however, the same AlGaAs has already been forcibly oxidised prior to exposing it to air, it can oxidise no further and so does not grow at all. Because of this difference in behaviour, a measurable step-change in height will be created on freshly cleaved facets at the point reached by the oxidation front during steam oxidation.

The AFM system used in this work is based on the Veeco NanoScope IIIa SPM Controller unit and the fundamental elements of the system are shown in Figure 3.3.

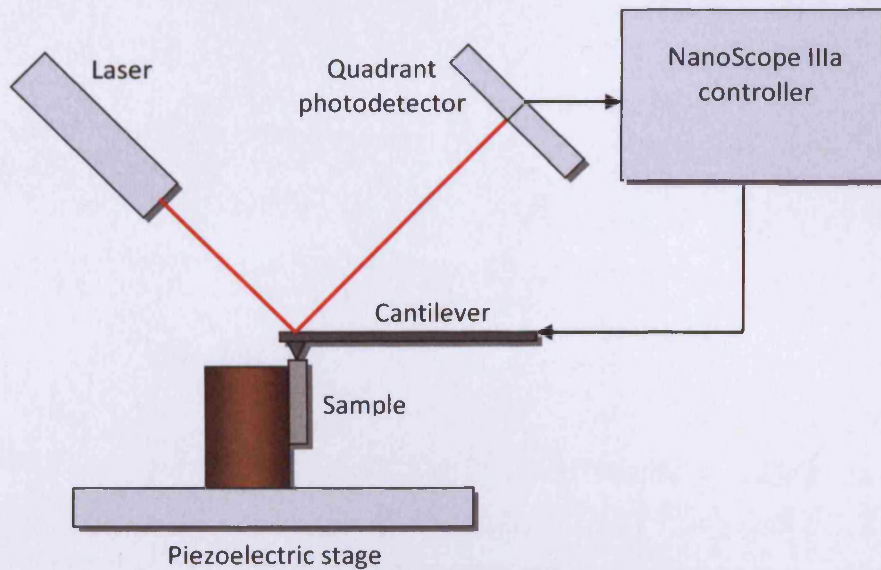

Figure 3.3: Diagrammatic representation of a basic AFM system

The functional component in any AFM system is microscopic needle or 'tip' formed on the end of a silicon cantilever. In the mode of imaging used here, called 'tapping mode', the cantilever is driven to oscillate at a frequency near to, but not at its resonant frequency. This causes the cantilever to oscillate with typical deflections of 100 – 200nm as it is gradually lowered towards the sample. When the sample surface comes into close enough proximity (not necessarily 'touching') with the oscillating tip, electrostatic forces (van der Waals, dipole-dipole etc.) cause the amplitude of the oscillations to decrease. As the tip is then rastered across the sample to create an image, a feedback loop uses information relating to these forces to maintain a fixed proximity between the sample and the tip by way of changing the voltage applied to the piezoelectric stage. As the tip is raised and lowered to match the contours of the sample surface, its deflection is measured very accurately by the deflection of the laser beam that is reflected off the back of the cantilever into a quadrant photodetector. Using this data a topographical image of the sample surface can be built up.

One limitation of the AFM system is the size of image it produces. Whilst scans can cover a fairly wide range of physical sizes, a maximum of 512 data points can be taken per scanned line. This limits the effective resolution for very large scan areas. The typical size of the oxide features being looked at in this work is between 10 and 30nm in thickness, which corresponds, by this principle, to a maximum scan size of about 3µm squared. Oxidation depths, however, will certainly be much larger than this. For this

reason the method employed to measure those depths will be to take one image at the oxidation front, track along the layer in steps of known size, and then take a final image at the sample edge. In this way, the total oxidation depth can be calculated, although it will also introduce what will be the greatest source of error in the results.

The scanning location on the AFM can be altered from where the tip is originally set down by way of an x and y offset value. These offsets are what will be used to track along the oxidisable layers during the measurements. This would not be a significant source of error if the samples were perfectly stable on their mounts, however, due to their nature, an ideal mount is not possible.

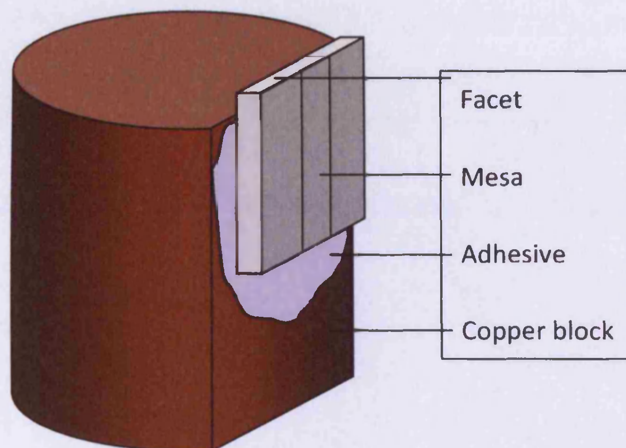

Figure 3.4: Method of sample mounting for AFM

Figure 3.4 shows how the samples for AFM and, indeed, for SEM are mounted. Ideally the adhesive shown in the diagram would be a thermally set, conductive silver epoxy as is used in normal device mounting. However, as the elevated temperatures or very long times required to set the epoxy would almost certainly alter the oxides visible at the facets, an alternative adhesive has to be used. This adhesive, although fairly firm, does not set completely hard and, therefore, leaves the sample very slightly mobile. When taking AFM images in this way, this results in a 'tracking error' of approximately 10%, which is indicated on the relevant graphs in later chapters. The AFM itself has a maximum calibration error of 10% before this tracking error is introduced.

The only other drawback of this method is that it inherently requires driving the AFM tip off the edge of the samples to take the final image of each measurement. Almost invariably, this damages the very delicate and expensive tips. Although this is clearly undesirable, it does not affect the final values

obtained as it is usually the case that the images just become slightly more blurred or 'shadowed', as shown in Figure 3.5.

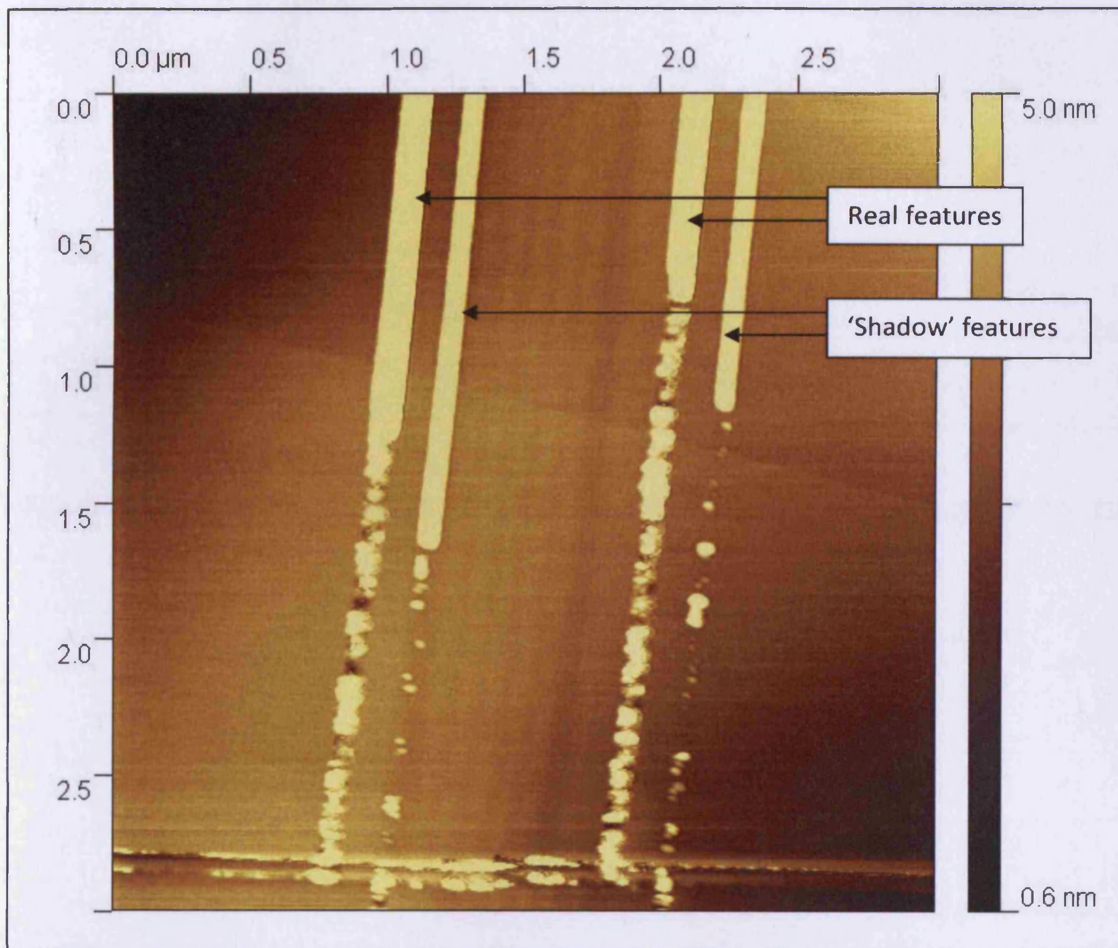

Figure 3.5: 'Shadowing' defects in AFM images caused by tip damage

This effect will exhibit when the damage to the tip is such that there are, effectively, two points to the tips, both scanning the same area but with some fixed offset. Clearly, accurate depth data can still be taken from this image if the shadow features are simply ignored.

3.5 Scanning Electron Microscopy (SEM)

SEM imaging, despite having lower resolution than AFM, is also used during the course of this work. Its advantages are its relative ease and speed of use and the ability to take images of much larger areas whilst still retaining fine detail.

The sample preparation is much the same for SEM as for AFM, although, for SEM, the samples must conduct all the way through to the stage in the microscope to avoid charging. For this reason, and because heating the samples after cleaving is less critical, silver epoxy is used to fix samples to the copper blocks.

SEM works by firing a focussed beam of high energy electrons downwards onto the surface of the sample and detecting the backscattered or secondary electrons. Backscattered electrons are those directly from the original beam that have been reflected or elastically scattered backwards off the sample surface. They are characteristically high energy and emitted mostly at very large angles relative to the sample surface. Secondary electrons, however, are emitted from sample atoms by inelastic scattering and have much lower energy levels. Such electrons are ejected over a far wider spread of angles. Both types of electron can be detected by scintillator units, which essentially turn ionising radiation into luminescence. Detectors, however, are placed differently inside the sample vacuum chamber to best suit the different angles of ejection.

Most commonly used of the two, and the type used in the results chapters of this work, are the secondary electrons. Whilst backscattered primary electrons enable the distinction of different chemical elements within a sample, they do not provide as good a 3D image due to their relatively indiscriminate nature with regard to sample surface angle. Secondary electrons give very good 3D images and also allow for good contrast between materials of different conductivity, particularly when the primary beam energy is kept low so as to minimise the penetration depth into the sample. This is clearly the most applicable method of imaging oxides within doped semiconductor material.

Two SEM systems are used during this work. These are a Jeol 6300 SEM and the SEM function of a Raith 50 electron-beam lithography unit.

3.6 Streak camera

A streak camera is a tool that can provide very high resolution, time-resolved pictures of short light pulses. Depending on the orientation of the devices being measured and the optional presence of a monochromator, these pictures can be time-resolved near-fields, in either the lateral or transverse directions, or time-resolved spectra, above or below threshold. Because of this versatility and the high level of time resolution available (the best commercial systems can achieve time resolutions in the order

of 200 femtoseconds [3]), the streak camera is the perfect tool to investigate the pulse dynamics that will be discussed in later parts of this work.

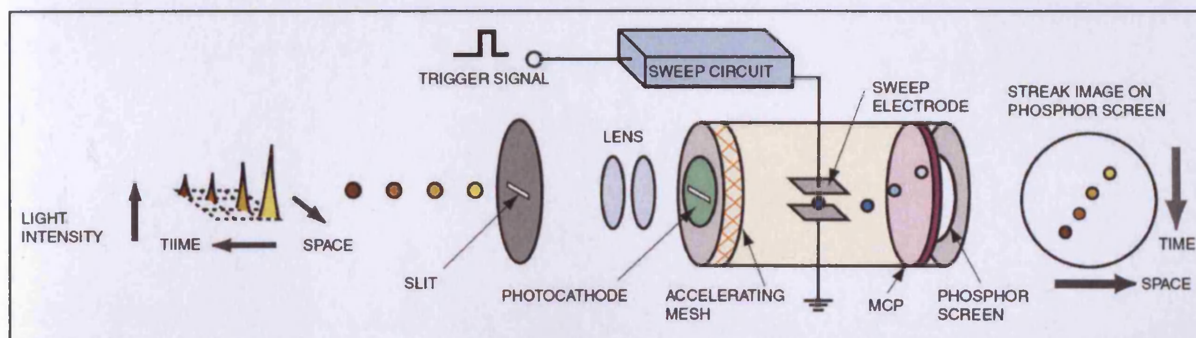

Figure 3.6: Operating principles of a streak camera [3]

Figure 3.6 gives a diagrammatic layout of the basic streak camera system. A signal with both temporal and spatial features is passed through a horizontal slit and focussed down onto a photocathode, where the incident photons are converted into electrons by the photoelectric effect, in numbers proportional to the intensity of the light signal. The resulting stream of electrons is accelerated away from the photocathode and through a pair of electrode plates. A swept voltage is applied across these plates as the electrons pass through them, which has the effect of spatially separating the electrons in the vertical direction according to their position in time. Finally, the electrons pass through a Micro Channel Plate (MCP) that gives off secondary electrons and so multiplies them in number by up 10,000 times. These electrons hit a phosphor plate and produce a light image with time resolved onto the vertical axis and the original spatial dimension of the source preserved on the horizontal. This image can be read and recorded by a CCD camera and computer. The simple addition of a monochromator in front of the first slit will change the final horizontal axis to wavelength, rather than distance, enabling time-resolved spectra to be taken.

In this work a streak camera system from Hamamatsu, incorporating a C5680 Streak Unit, C4742-95 CCD camera and associated control boxes, was used. In the case of any spectral data recorded, a Bentham M300 monochromator was added into the system. Devices under test were mounted horizontally on a custom-built, temperature-controlled stage, with three degrees of freedom (x, y and z axes), to aid alignment. The temperature of the mount was monitored with an embedded thermistor and controlled with a Peltier module. The Hamamatsu Streak Unit has provision for an external trigger source, which was utilised here in triggering the unit off the pulse generator used to power the laser devices. With this setup completed, device temperatures could be set anywhere between five and 80°C with an accuracy

of better than ±1°C, and the position in time of the electrical pulse, relative to the trigger signal, could be controlled for easy synchronisation with the streak camera.

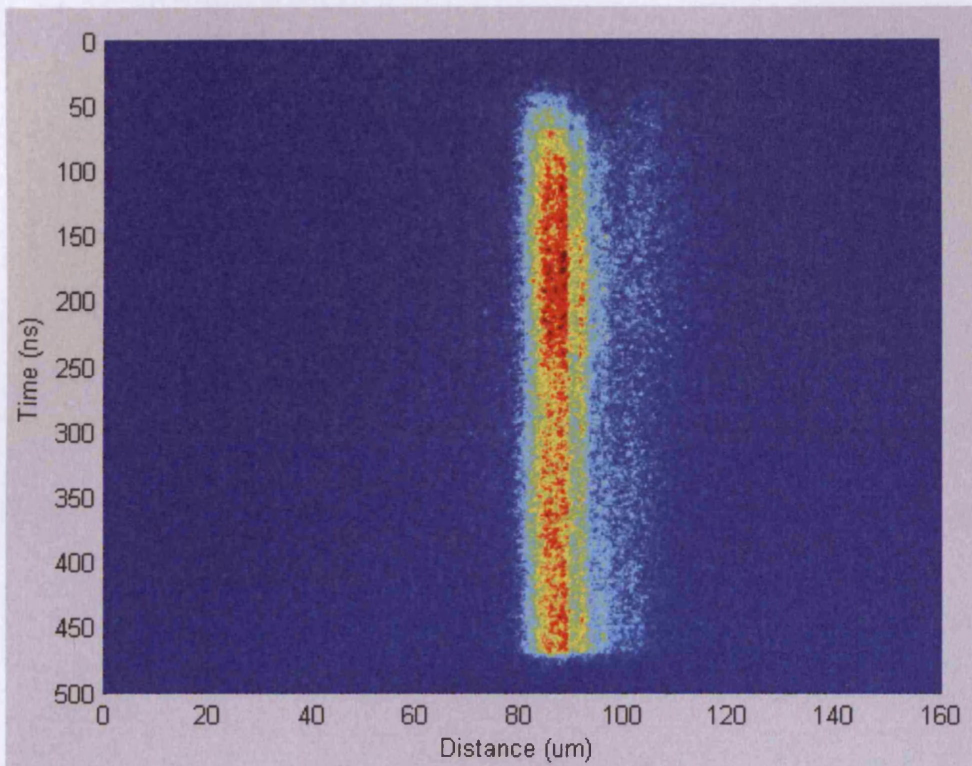

Figure 3.7: Typical streak camera data (no monochromator). Colours show intensity with blue being low, yellow is medium and red is high

An example of the time-resolved near-field images produced by the streak camera is shown in Figure 3.7. In the vast majority of cases throughout this work, the top of the image (t=0) is aligned to the start of the electrical drive pulse and the bottom of the image to the end; any absence of light seen up and down the picture is due to real effects. In the case above, the start of the pulse is aligned correctly, and the delay before the onset of a signal is due to the device behaviour. The end of the pulse, however, does actually occur at 475ns.

Figure 3.7 shows a simple capture of the real-time data displayed in the streak camera software, and is consequently visibly noisy. There are options available in the software, however, which are used later in this work, to integrate and/or average many identical frames in order to increase the signal to noise.

Due to the nature of the streak camera, errors in the time axis are very small. As the horizontal axis (distance or wavelength) is effectively unknown to the streak camera system, however, there are some

errors to be considered in its manual calibration. In the near-field and spectral measurements, the full range covered by the 'window' of the streak camera is measured simply tracking a near-field or spectrum of known width (and peak value in the case of the spectral data) across the screen. In both cases, this error in the absolute values can be conservatively estimated at around 10%, although this calibration need only be done once after the experimental components are fixed in position, errors in comparative data between devices will be much smaller.

3.7 The near-field camera

With the exception of time-resolved near-fields, which were taken with the streak camera system, all of the near-field data presented in this work was taken on the dedicated near-field camera system detailed here. The camera itself is a XenICs Xeva near-infrared device, with an InGaAs array as the active component. The array is sensitive to wavelengths from 0.9µm to 1.7µm and is 320 x 256 pixels in area.

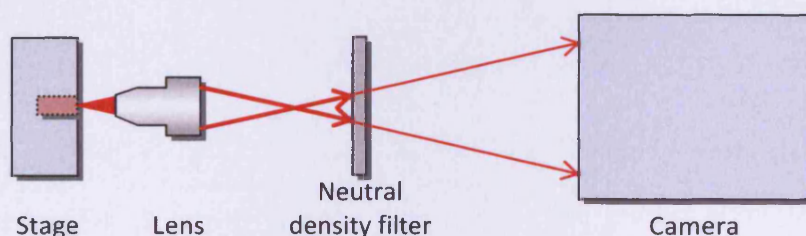

Figure 3.8: Near-field camera system as viewed from above

Figure 3.8 shows how the camera is used on an optical bench within the near-field kit as a whole. The devices to be imaged are mounted on a stage with the laser beam pointing directly towards the camera array. In order to align the devices accurately, the stage itself has four degrees of freedom, allowing motion in the x, y, and z axes as well as rotation around the axis of the light beam. An objective lens (commonly 5x or 10x) is placed close to the laser to collect the diverging beam and focus it on the camera array, which is approximately 50cm away. The camera's own lens is removed as this configuration allows greater magnification. An optional neutral density filter can be placed in the path of the beam to reduce the intensity of light incident on the array. This is usually required when looking at near-fields above threshold as the dynamic range of the array is relatively small.

The laser devices are electrically pumped, most commonly with a pulsed power supply, and the current is monitored with an oscilloscope via a closed-loop current probe. The camera is not triggered off this

power supply and so, simply integrates over a period of time, which is fixed by the user to a level that achieves the best signal to noise ratio whilst not saturating the array. The camera is controlled via a computer, where a real-time image of the near-field can be viewed. From this image a line scan is taken across the width of the whole device. It is this data that is saved and analysed as the near-fields seen in this report. As the physical width of the devices being measured is known (usually 300μm), the position data (x axis) can be calibrated. The intensity however, cannot be calibrated in absolute units and so can only be compared relatively where the integration time and attenuation of any neutral density filters are known.

To obtain 'good' near-field data, care must be taken in two areas. Firstly, as has already been mentioned, the array must not be saturated or artificially flat-topped near-fields will result. This can be avoided by using the correct level of filtering. Secondly, the camera software employs a variable 'auto-zeroing' mechanism. In some conditions, this can lead to 'zero' being set too high and the near-field be artificially cropped from the bottom. The addition of a controlled low level of background light in the form of a diffused desk lamp ensures that the zero is set sufficiently low as to include the full range of intensities present in the near-field. Figure 3.9 shows examples of data where both of these errors have been made alongside a set of good near-field data.

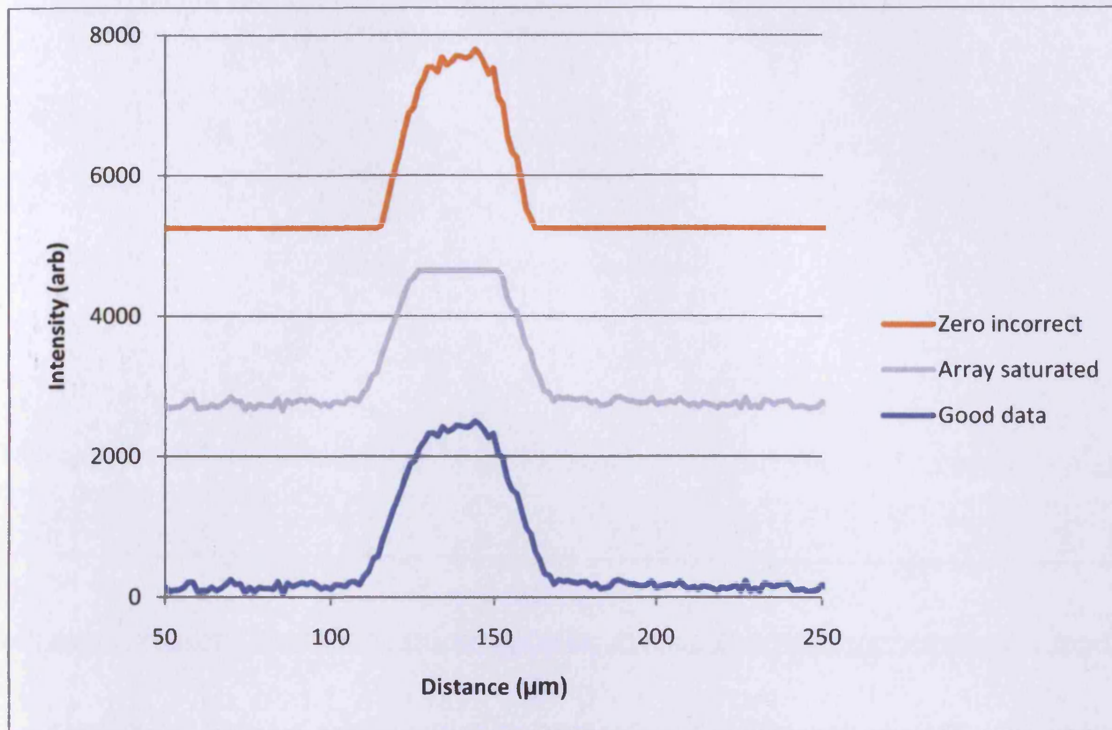

Figure 3.9: Examples of good and bad near-field data. Traces are offset on the y axis for clarity

If these experimental conditions are adhered to, and 'good' data is achieved, a measure of the near-field width can be obtained. Where the data can be approximated to a Gaussian shape, or even where all the data to be compared is of a similar, but non-Gaussian shape, the measure of the Full-Width-Half-Maximum (FWHM) can be used. This simply takes the width of the near-field at half of its maximum intensity. Where the near-field profile is heavy filamentary or otherwise irregular, this measure may not be appropriate.

3.8 The multisection technique

This section outlines the practical implementation of the mathematical principles of the multisection technique that were detailed in the first chapter of this work. The multisection technique is a conceptually simple, single-device method of determining both the modal gain and loss of a material. An example of the data that can be obtained is given in Figure 3.10.

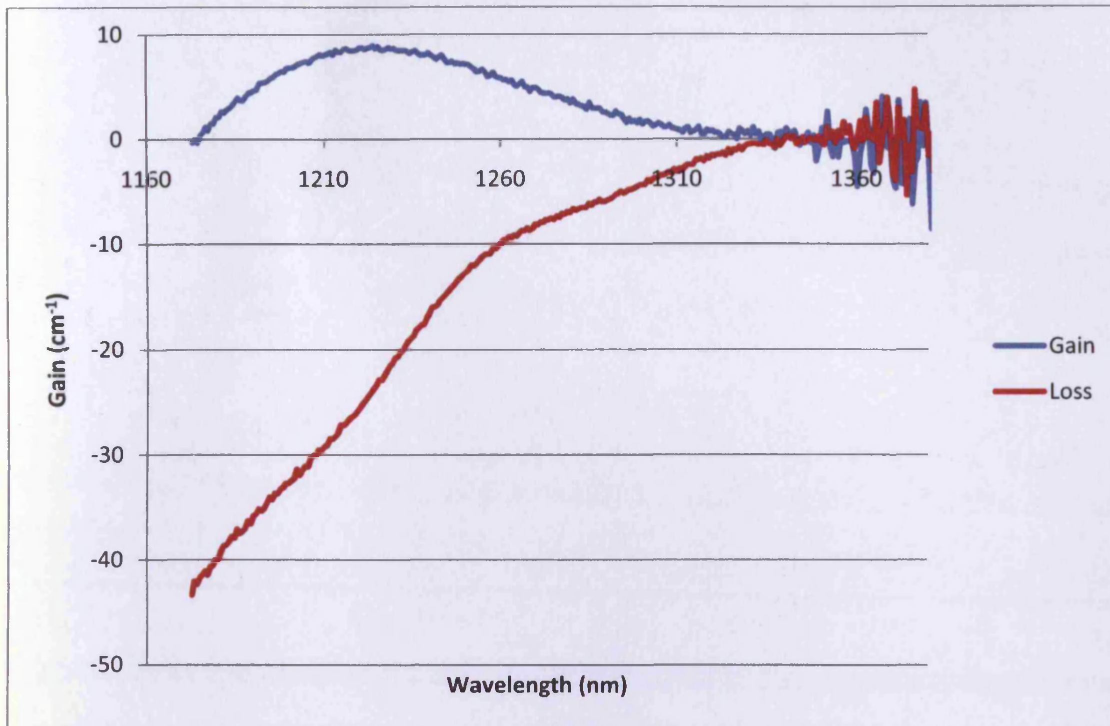

Figure 3.10: Typical gain and loss data from the multisection method

The devices used in this experiment are specially fabricated for the purpose. They are most commonly broad-area devices, but can be made from any other laser structure and comprise several identical sections, separated by shallow-etch 4μm trenches, termed section breaks. The sections in devices used in this work are all 300μm in length, although this can be varied as long as all sections are the same. The section breaks provide the crucial electrical isolation of the sections, although this isolation can never be total, as would be ideal. The section break trenches are etched to the bottom of the GaAs cap, and provide inter-contact resistances of between 100 and 1,000Ω. As a consequence, some current will always leak between sections. To ensure valid data, compensation for this can be made by supplying and monitoring the currents to each section individually. Figure 3.11 shows a diagrammatic picture of a multisection device. A real device would have more sections than the four depicted here (10 or more), in order to provide the most amount of absorbing material possible, and reduce any round-trip gain to an absolute minimum. For the same reason, multisection devices are often made with off-angle facets, or, as in the picture below, with the back facet purposefully damaged. The front two sections, those actually used in the experiment, are conventionally referred to as section 1 and section 2 in the order illustrated here.

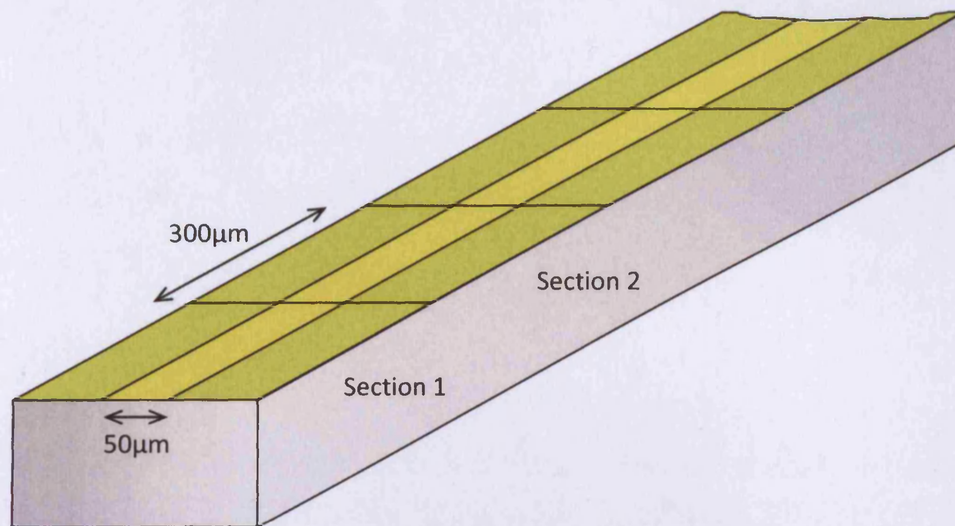

Figure 3.11: Multisection device with back facet damaged

Once fabricated, the devices are mounted p-side up on standard laser headers and the sections are all individually electrically contacted. Depending on the number of sections available, one or more of the rear sections may be actively grounded, or even reverse biased, in order to remove any carrier pairs generated by absorption.

For the experiment to yield accurate results, not only must round-trip gain be eliminated and the sections be physically identical, but they must also have matching IV characteristics and near-fields that show no filamentation. If all this is true, then the devices can be aligned to any appropriate equipment, and the requisite spectra can be taken. In the case of data taken in this work, the chosen equipment is a monochromator coupled to an InGaAs photodiode array, very similar to the one described in the previous section.

As a final note, the acquisition of 'good' data in this experiment relies heavily on correct alignment of the devices. Focussing exactly on the facet (not necessarily on the sharpest image of the mode) and limiting the collection angle to approximately 5° are the most crucial factors in this. An indication that the acquired data is reliable is the converging and flattening of the gain and loss traces at the long wavelength end of the graph, which equates to the energy of the bandgap, and indicates the value of the loss parameter, α_i. Details and justification of these criteria can be found in previous work done at Cardiff University. [4]

3.9 The integrating sphere

The final major piece of equipment used in this work is an integrating sphere. The sphere is a means of directly measuring the total optical output power of a laser in real units and is used here to investigate the high-power performance of devices.

When using an integrating sphere, the laser device is inserted, via a small opening, into the sphere, which is coated inside with highly reflective material, and is sealed so that no light can enter or escape. The reflective coating ensures that whatever light is emitted by the device is, after multiple reflections, distributed evenly across the whole internal surface of the sphere. At some point on the sphere's surface is an optical detector of small, but accurately know surface area. This detector can be positioned anywhere on the surface of the sphere as long as it is not directly in the primary beam path of the laser. Figure 3.12 shows the detector as a blue square on the top of the sphere.

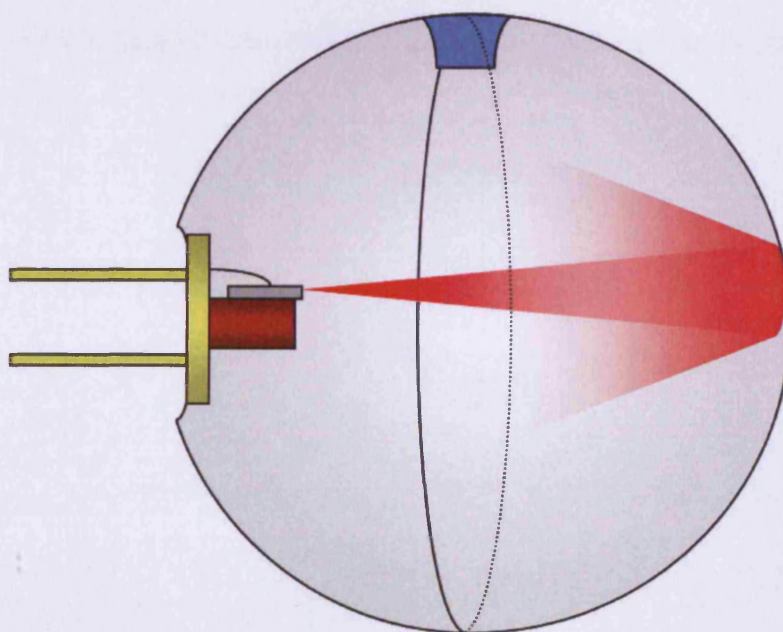

Figure 3.12: Diagrammatic view of an integrating sphere

By knowing the exact areas of both the sphere and the detector, the power incident on the detector can be integrated over the whole area of the sphere, and the total power can be calculated. The only real source of error in this method arises from the fact that the laser emits from both facets. Whilst a lot of the light emitted from the back facet will be reflected off the metal laser mount, some may well be

absorbed, which will reduce the total power reading. It is estimated that the final numbers given by the system will equate roughly to one and a half facets-worth, rather than the full two. Whatever this fraction, however, it will remain more or less constant for all devices mounted in the same way and so results can still be compared.

The specific equipment used here is an Ophir 3A-IS-IRG integrating sphere head coupled with an Ophir Nova II optical power meter.

3.10 References

1. **Griffiths, K.** *Automated laser diode measurement, M.Phil.* UWCC, 1992.

2. **Okada, T., et al.** *Atomic Force Microscope. 5,260,824* USA, 9 November 1993.

3. **Hamamatsu.** Guide to Streak Cameras. [Online] [Cited: 5 March 2010.] http://sales.hamamatsu.com/.

4. **Lutti, J.** *Optical properties of InP/AlGaInP quantum dot laser heterostructures, PhD.* Cardiff University, 2005.

Chapter 4 : Design and Analysis of Oxidisable Material Systems

4.1 Introduction

Over the past couple of decades many authors have presented results on the oxidation rates of AlGaAs compounds, as has already been discussed in the background chapters. Whilst the data available on the topic is relatively abundant, parameters that affect the oxidation rate are so numerous and, in some cases, affect it so strongly that it can still be hard to predict the exact rate of oxidation that will be achieved under a given set of conditions.

This chapter therefore aims to characterise the specific oxidation equipment and range of parameter values that will be used to gather further results in later sections of this work. A set of test samples has been prepared from materials very similar to those which have been used to fabricate the working devices in the subsequent chapters. The majority of the results have been obtained by cross-sectional AFM imaging and, wherever possible, will be compared to literature values for oxidation under similar conditions.

The results presented in this chapter are taken from two wafer designs - one that is comparable, though not identical, to data already reported in the literature and that serves partly to characterise the oxidation equipment, and a second that is novel in its design and for which no oxidation rate information is available in the literature. This chapter also presents information on the dependency of the oxidation rate of the dopant species of buried material - a parameter that is also unrepresented in literature.

4.2 Wafer structures

The main aim of the first experiments detailed in this chapter is to map the oxidation rates achieved by the specific equipment used in the rest of this project on the relevant materials. To this end, very simple etched mesa structures were fabricated from the two wafers that were intended to be made into full working devices at a later stage.

Both of these wafers were designed specifically for this project using a computer simulation package called SimWindows.[1] SimWindows is a one-dimensional, finite-difference simulation package that uses, amongst other methods, Schrödinger's and Poisson's equations to simulate the performance of a variety of types of optoelectronic devices. It takes into account quantum confinement and tunnelling current, which is calculated using the WKB approximation. The software allows the designs made here to be optimised for as low a turn-on voltage and device resistance as possible, whilst adding in the necessary oxidisable layers. Both wafers were based on standard 5-stack quantum dot designs. The first wafer, in fact, has been grown twice, with some differences between the two iterations. Both iterations are based on the use of 30nm oxidisable layers, placed in the cladding regions, at a distance of 160nm from the active region. The earlier design, which existed prior to the commencement of this project, has two such layers, one in either cladding region. The second iteration, however, which was designed and simulated during this work, contains only a single oxidisable layer in the p-type cladding. Furthermore, in order to minimise the device resistance by way of reducing the potential barrier created by the high aluminium content layer, extra layers of sequentially lower aluminium content were added on either side of this layer. This has the effect of graduating the potential barrier created by adding in the high bandgap material and, therefore, reducing its overall effect when unoxidised. The oxidisable layers in both cases have been placed as close to the active region as possible, without being close enough to interact strongly with the optical mode. From this point onwards both iterations of this wafer will be referred to as 'Wafer A', however, it is exclusively the first (referred to as the 'test version') that is used in this chapter for characterisation of the oxidation, and exclusively the second (the 'device version') that is used in following chapters for working devices. The reason for this is that the first iteration of the design offers insight into more parameters affecting the oxidation rate, whereas the second provides better electrical performance and avoids concerns of differing oxidation depths between the p- and n-doped oxidisable layers.

The second wafer design incorporates oxidisable layers within the active region - a feature that has not previously been reported in literature. In this case, stacks of ten 1nm $Al_{0.98}Ga_{0.02}As$ layers interspersed with 1nm $Al_{0.75}Ga_{0.25}As$ layers were placed around each quantum dot layer. The purpose of this design was to avoid completely the effects of current spreading by having the current confining layers actually in the active region, and to provide much stronger lateral optical confinement, again by their proximity to the lasing mode. Once again, the design was optimised to give the minimum device resistance in the unoxidised state, whilst retaining enough oxidisable material in the active region to have a significantly different resistivity and refractive index when laterally oxidised in devices. Here, rather than reducing

the potential barrier to allow normal current flow, the very thin layers maintain a low resistance by allowing a significant amount of tunnelling current. This wafer will hereafter be referred to as 'Wafer B'.

Figure 4.1 below is not to scale but shows, for the purpose of illustration, the intended use of the oxidisable layers in Wafer A (test version). In broad-area devices, such as the one shown below, the lateral oxidation will have no effect on the performance of the laser as it will only penetrate to a fraction of the lateral distance that the stripe-forming oxide on top of the device covers. When such oxidation is performed after the cleaving of the device facets however, the resulting oxide will prevent the drive current from reaching the areas immediately behind the facets, and it is this feature that will be investigated in Chapter 6.

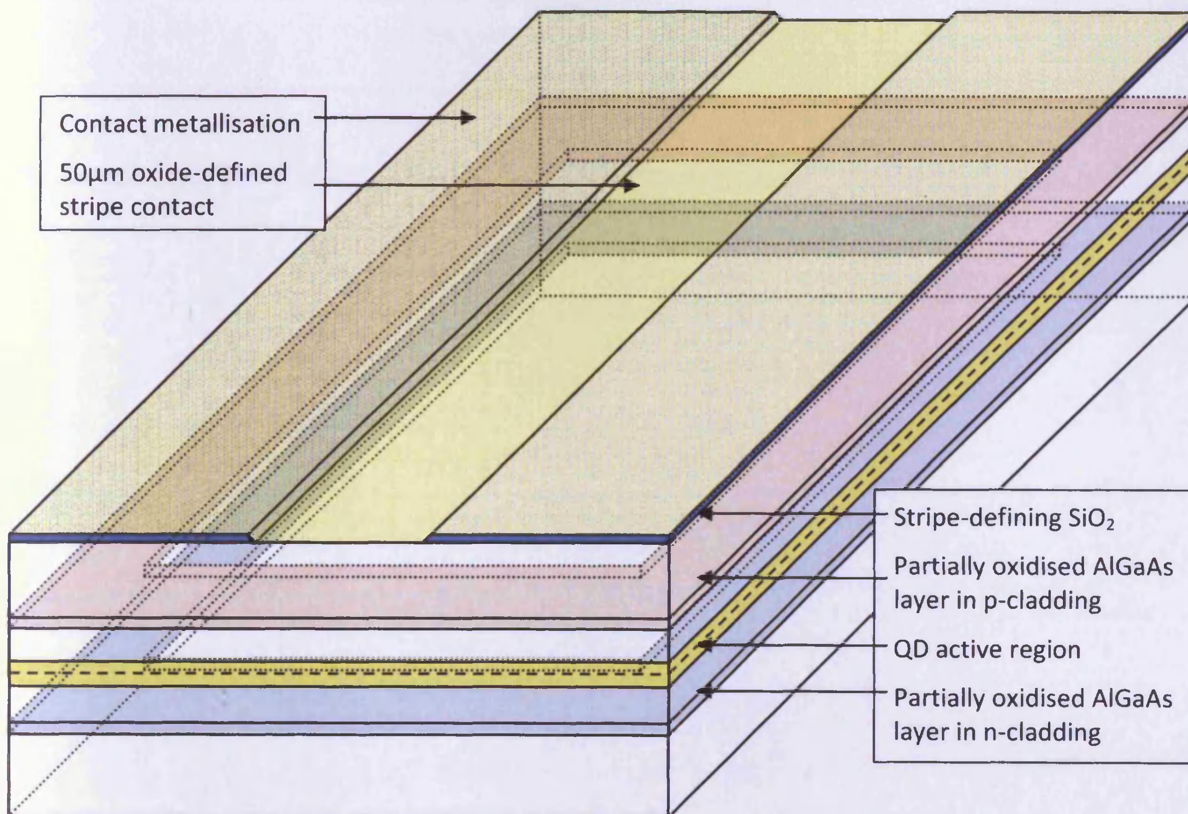

Figure 4.1: Schematic diagram of partially, post-cleave oxidised device of 'Wafer A'

Conversely, the eventual use of Wafer B is to oxidise ridge structures prior to cleaving the facets, so that it is only lateral oxidation that is present and that affects the device performance. However, for the purposes of this chapter, the same etched mesa structures have been fabricated from both wafer types. The detailed structures of both wafers as grown by Sheffield University by Molecular Beam Epitaxy are

given below in Tables 4.1 and 4.2. The specific tolerances for the mole fraction and layer thicknesses of these layers are not known.

Table 4.1: Growth structure of Wafer A, test version

Material	x	Thickness (nm)	Dopant type	Dopant conc. (cm^{-3})	Notes
GaAs	-	300	p	1e19	Capping layer
$Al_{(x)}GaAs$	0.45	1100	p	5e17	Cladding
$Al_{(x)}GaAs$	0.98	30	p	5e17	Oxidisable layer
$Al_{(x)}GaAs$	0.45	100	p	5e17	Cladding cont.
$Al_{(x)}GaAs$	0.45	0.5	p	5e17	Bandgap matching superlattice
GaAs	-	0.5	p	5e17	Pair repeated 25x
GaAs	-	50	-	-	Quantum dots in wells
$In_{(x)}GaAs$	0.15	6	-	-	
InAs	-	0.9	-	-	
$In_{(x)}GaAs$	0.15	2	-	-	
GaAs	-	50	-	-	Repeated 5x
$Al_{(x)}GaAs$	0.45	0.5	-	-	Bandgap matching superlattice
GaAs	-	0.5	-	-	Pair repeated 25x
$Al_{(x)}GaAs$	0.45	100	n	1e18	Cladding
$Al_{(x)}GaAs$	0.98	30	n	1e18	Oxidisable layer
$Al_{(x)}GaAs$	0.45	1000	n	1e18	Cladding cont.
GaAs	-	300	n	2e18	Capping layer

Table 4.2: Growth structure of Wafer B

Material	x	Thickness (nm)	Dopant type	Dopant conc. (cm^{-3})	Notes
GaAs	-	300	p	1e19	Capping layer
$Al_{(x)}GaAs$	0.80	1	p	5e17	Bandgap matching superlattice
GaAs	-	1	p	5e17	Pair repeated 10x
$Al_{(x)}GaAs$	0.80	1150	p	5e17	Cladding
$Al_{(x)}GaAs$	0.95	1	p	5e17	Oxidisable binary alloy stack
$Al_{(x)}GaAs$	0.75	1	p	5e17	Repeated 10x, last 3x undoped
$Al_{(x)}GaAs$	0.70	1.5	-	-	Graduated bandgap narrowing
$Al_{(x)}GaAs$	0.55	2.5	-	-	
$Al_{(x)}GaAs$	0.40	2.5	-	-	
$Al_{(x)}GaAs$	0.25	2.5	-	-	
$Al_{(x)}GaAs$	0.10	2.5	-	-	
GaAs	-	2.5	-	-	
In(x)GaAs	0.15	6	-	-	Quantum dot structure
InAs	-	0.9	-	-	
In(x)GaAs	0.15	2	-	-	
GaAs	-	2.5	-	-	
$Al_{(x)}GaAs$	0.1	2.5	-	-	Graduated bandgap widening
$Al_{(x)}GaAs$	0.25	2.5	-	-	
$Al_{(x)}GaAs$	0.4	2.5	-	-	
$Al_{(x)}GaAs$	0.55	2.5	p	5e17	
$Al_{(x)}GaAs$	0.70	2.5	p	5e17	
$Al_{(x)}GaAs$	0.95	1	p	5e17	Oxidisable binary alloy stack
$Al_{(x)}GaAs$	0.75	1	p	5e17	Repeated 10x, last 3x undoped
$Al_{(x)}GaAs$	0.70	1.5	-	-	Graduated bandgap narrowing
$Al_{(x)}GaAs$	0.55	2.5	-	-	

$Al_{(x)}GaAs$	0.40	2.5	-	-	
$Al_{(x)}GaAs$	0.25	2.5	-	-	
$Al_{(x)}GaAs$	0.10	2.5	-	-	
GaAs	-	2.5	-	-	Whole section repeated 4x
In(x)GaAs	0.15	6	-	-	Quantum dot structure
InAs	-	0.9	-	-	
In(x)GaAs	0.15	2	-	-	
GaAs	-	2.5	-	-	
$Al_{(x)}GaAs$	0.1	2.5	-	-	Graduated bandgap widening
$Al_{(x)}GaAs$	0.25	2.5	-	-	
$Al_{(x)}GaAs$	0.4	2.5	-	-	
$Al_{(x)}GaAs$	0.55	2.5	p	5e17	
$Al_{(x)}GaAs$	0.7	2.5	p	5e17	
$Al_{(x)}GaAs$	0.95	1	p	5e17	Oxidisable binary alloy stack
$Al_{(x)}GaAs$	0.75	1	p	5e17	Repeated 10x, all doped
$Al_{(x)}GaAs$	0.45	1150	n	1e18	Cladding
GaAs	-	1	n	1e18	Bandgap matching superlattice
$Al_{(x)}GaAs$	0.45	1	n	1e18	Pair repeated 10x
GaAs	-	300	n	2e18	Capping layer

4.2.1 Simulation results

The SimWindows program is able to take the layer structure of a material, such as those shown in the tables above, and calculate band diagrams, refractive index profiles, IV curves and more. Although the design process for the two wafers went through many iterations, this section discusses only the more important design points and shows the final performance results. Figures 4.2 and 4.3 display the equilibrium band diagrams of final designs of the wafers as obtained from the simulations. These are followed in Figures 4.4 and 4.5 by the forward bias versions, which also include the quasi Fermi levels.

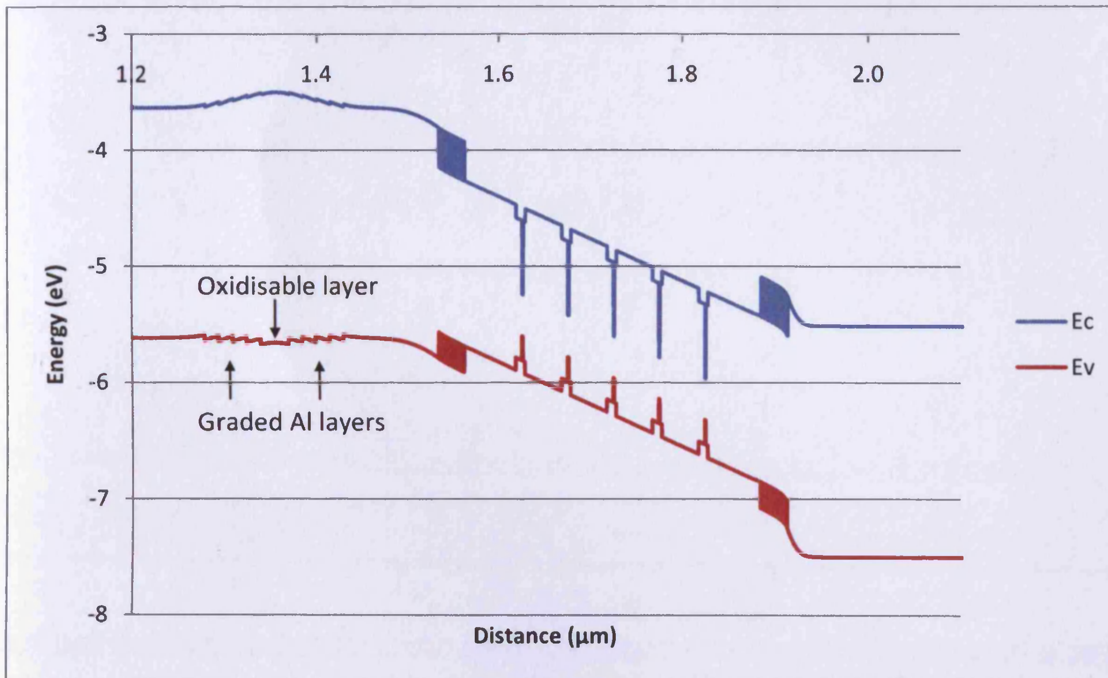

Figure 4.2: Simulated equilibrium band diagram of 'Wafer A' active region, device version

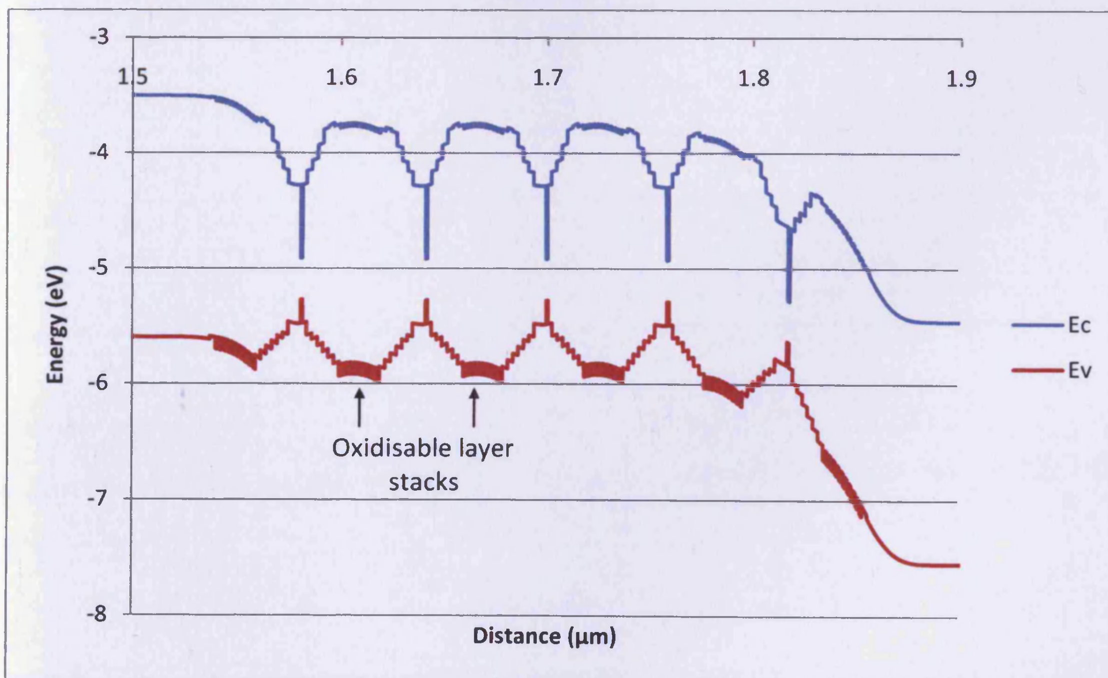

Figure 4.3: Simulated equilibrium band diagram of 'Wafer B' active region

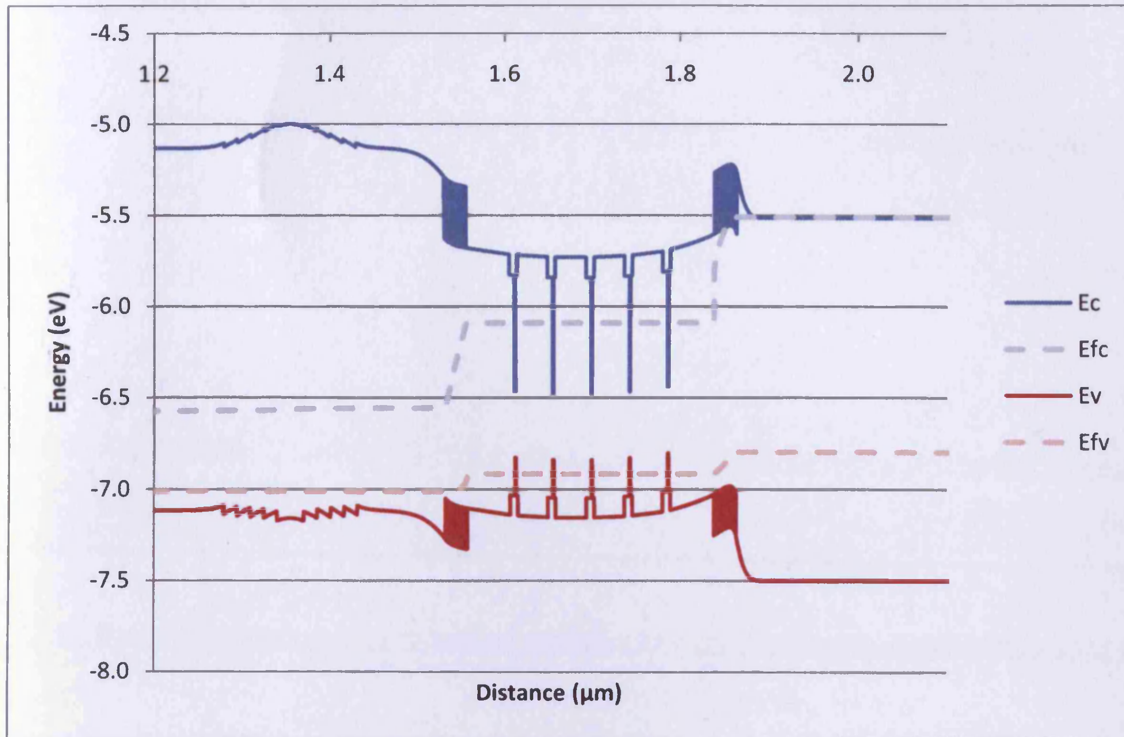

Figure 4.4: Simulated forward bias band diagram of 'Wafer A' active region, device version

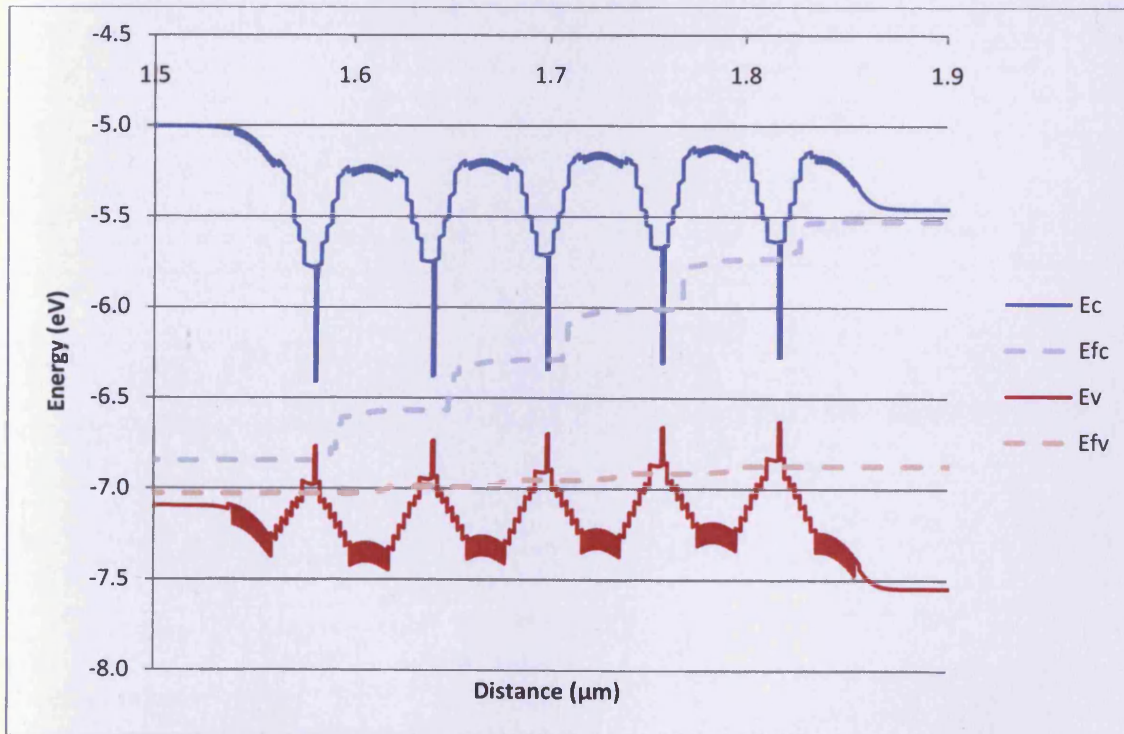

Figure 4.5: Simulated forward bias band diagram of 'Wafer B' active region

There are a few features that should be noted from the figures and tables above. First the smoothing effect of the bandgap graduation around the oxidisable layer in Figure 4.2. Figure 4.3 shows an unusually flat active region bandgap across the structure until the far right-hand end of the diagram where the bands bend suddenly downwards. This is due to the presence of p-doping in the active region, which aids recombination in the dots. The latter two figures show both wafers in overall gain. It should be noted that at the drive level shown in Figure 4.5, the two left-most quantum dot layers are not actually in gain at all. This may affect the threshold current of devices, but is a consequence again of the p-doping in the active region, and so may well be compensated for by increased dot performance. Lastly, it should be noted that Wafer B has 80% aluminium p-type cladding, rather than the usual 45%. This was necessary to offset the high bandgap oxidisable material that has been added to the active region and retain the transverse waveguiding effect of the structure. The n-type cladding has been left at 45% aluminium to minimise device 'on resistance'.

As has already been mentioned, the only design differences between the test and device versions of Wafer A are the removal of the n-type oxidisable layer, and the addition of graded aluminium layers on either side of the p-type layer. As can be seen in Figure 4.6, which shows the simulated IV characteristics of both wafer designs in comparison to a control wafer design with no oxidisable layers, the device version of Wafer A is predicted to have an 'on resistance' more than a factor of ten less than of the test version. Its IV curve as a whole, in fact, is almost identical to that of the control design.

The IVs of two versions of Wafer B are also plotted in the same figure. The trace labelled 'early design' incorporated undoped, solid, 10nm oxidisable layers in-between each successive dot layer. Clearly this produces a very poor IV characteristic, so the 10nm layers were broken up into stacks of 20 1nm layers, alternating 95% with 75% aluminium material in order to promote significant levels of quantum tunnelling current. It is also know that certain levels of p-doping around InAs quantum dots can be beneficial to the gain, [2] so this was added whilst not doping the actual dots themselves. Lastly, graded aluminium content layers were added to aid the step down in bandgap from the oxidisable material to the quantum dots. The final result of these changes is a nearly 250-fold increase in the electrical performance of the design, and an IV that, again, is very similar to that of the control wafer.

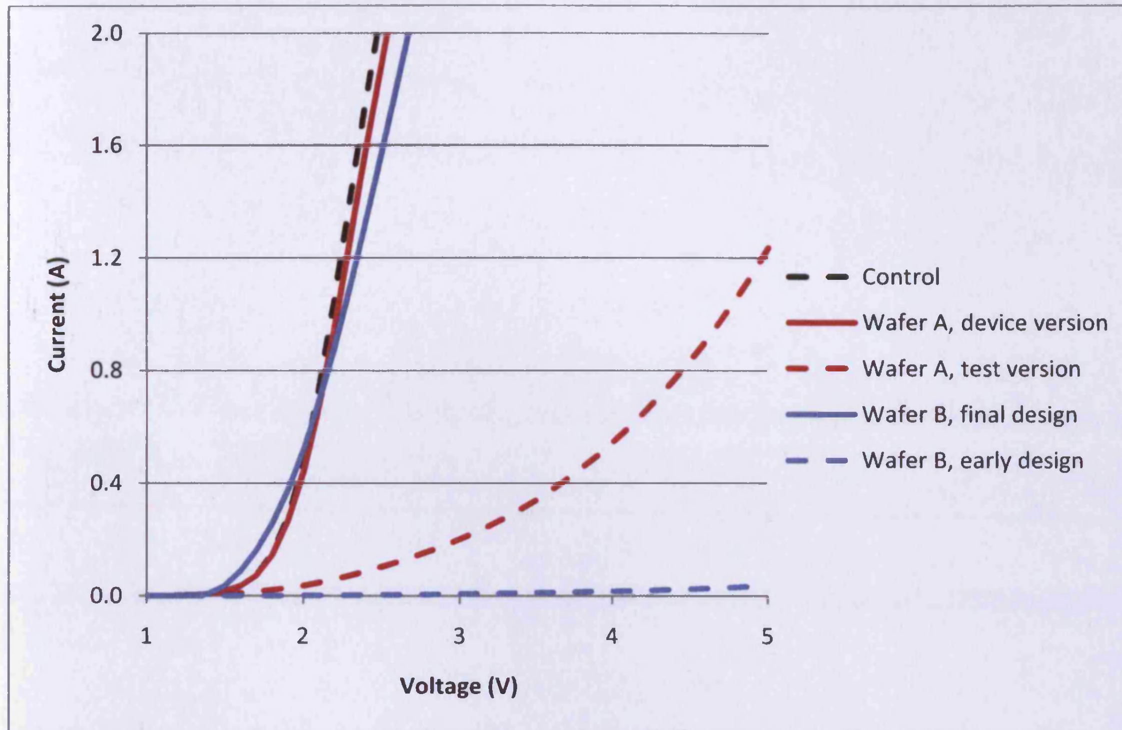

Figure 4.6: Simulated electrical performance of Wafers A and B

4.3 Measuring the oxidation rate

Due to the nature of the technology, measuring the rate of lateral oxidation can be problematic. In situ monitoring of the process in laser devices is all but impossible given the small scale and harsh environment in which the oxidation takes place. Several post-process methods have been used successfully by various authors, although most are destructive and all require careful sample preparation. The most commonly used methods are long-wavelength IR imaging, which allows the oxidation front to be viewed from above by way of its relatively low refractive index, and cross-sectional AFM, SEM or TEM imaging, all of which view the oxidation front directly. A combination of AFM and SEM imagining was chosen for this work due to the availability of the relevant equipment and the thickness of the layers to be viewed. The IR imaging technique works well with thick oxide layers, but at 30nm or less the layers used in this work are relatively thin, whilst still being within the resolution of both SEM and AFM imaging.

The sample preparation is the same in both of these cases. Both wet- and dry-etched stripe mesa samples, prepared by the cleanroom staff at Cardiff University, are used in this chapter. Typically the

mesas are 100µm wide, 2µm deep and run for the whole length of the sample. The etch serves to expose the buried oxidisable AlGaAs layers down the length of the samples, which are typically 2mm by 2mm with several stripe mesas on the surface at a pitch of 300µm. The samples are oxidised as described in Chapter 3 before being cleaved in half across the width of the mesa stripes. This leaves the cross-section of the partially oxidised layers visible as shown below in Figure 4.7. The diagram shows a typical sample with the remaining high aluminium content AlGaAs in hatched pink or blue (p- or n-type doping), the resultant oxide in solid pink or blue and the surrounding material in grey.

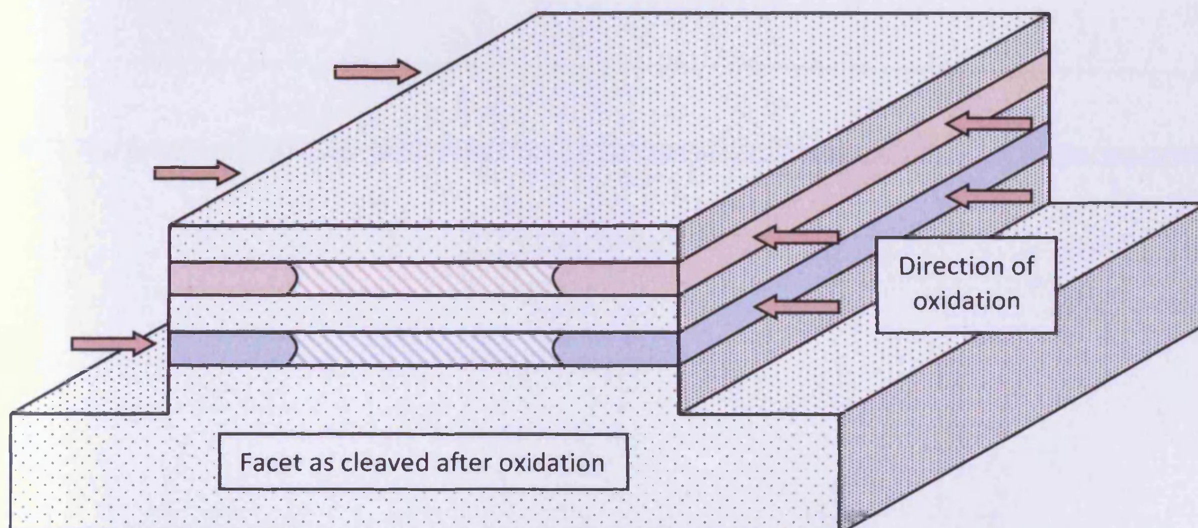

Figure 4.7: Mesa stripe structure as cleaved to measure lateral oxidation depth showing p-type (pink) and n-type (blue) doping

Important factors in the sample preparation for this method of measuring the oxidation depth are the quality of the cleave and the time allowed between cleaving and imaging the samples. In the case of either AFM or SEM imaging, a smooth facet surface is vital in terms of acquiring an accurate measurement of the oxidation depth. It is for this reason that the test samples are made wide enough to include several mesa structures, allowing enough distance across the sample for any damage caused by scribing to be left behind and the cleave to follow the natural lattice planes of the material cleanly. In the case of SEM imaging, the samples were put under vacuum as soon as possible, as it is the difference in conductivity between the oxide and the AlGaAs parent layer that is detected. With AFM imaging however, the cleaved samples must be exposed to air for a certain period in order to develop a rough, natural oxide on the previously unoxidised region of the AlGaAs layer, which is proud of the facet

surface. Previous work at Cardiff University has indicated that the optimum time for exposure to air in this situation is two hours. [3]

In both cases the samples were cleaved and mounted by hand onto the sides of small copper blocks, so that the freshly exposed facet faced upwards for ease of observation. Figure 4.8 shows an AFM cross-sectional image, typical of those used to take oxidation depth data. The image is topographical, with light colours corresponding to raised features, and shows the two oxidisable layers of Wafer A as well as the 5 quantum dot layers in-between. In this image, the oxidation has proceeded from the bottom of the picture upwards, and the oxidation front, where the smooth, raised native oxide begins, is clearly visible.

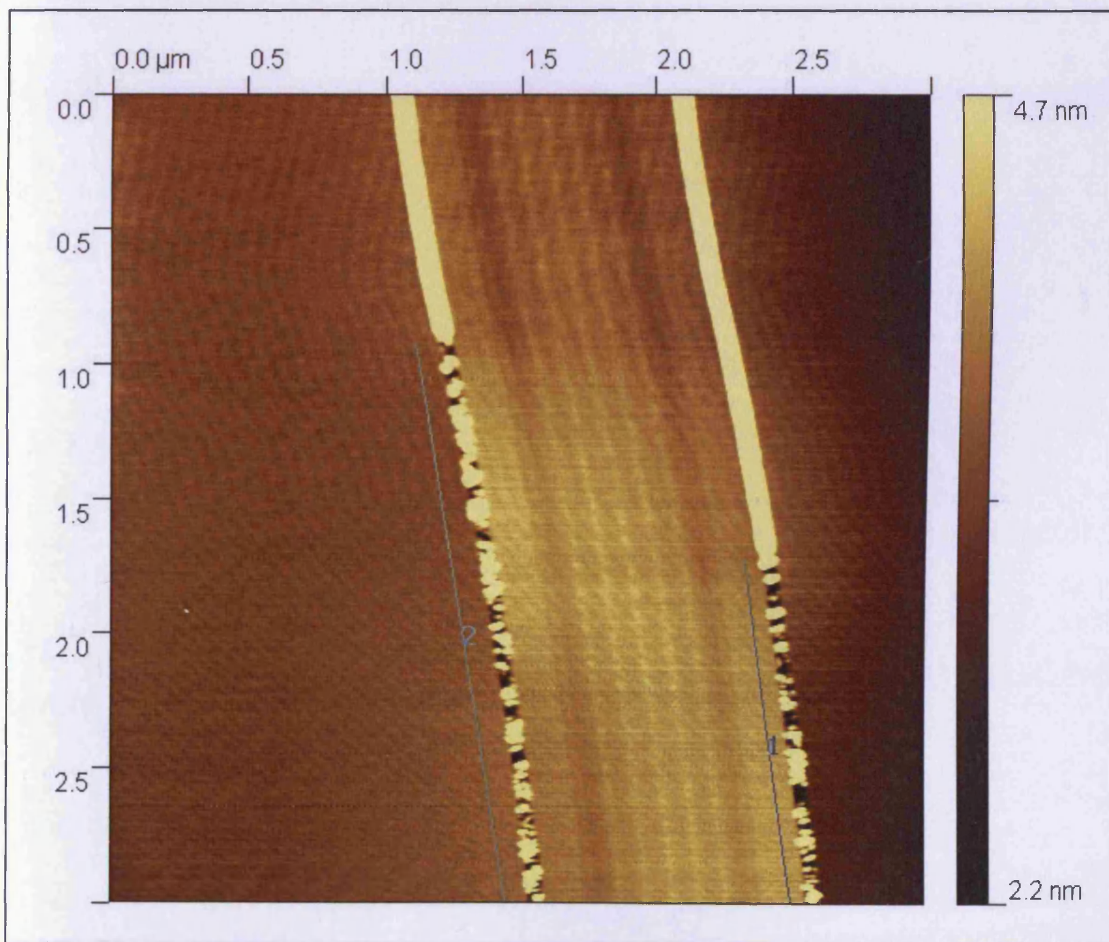

Figure 4.8: Atomic Force Microscopy image showing oxidation fronts (proceeding upwards) in p-type (1) and n-type (2) material

4.4 Oxidation rates in 30nm layers

This section reports oxidation rates of the layers in Wafer A and their dependence on various factors both in the wafer structure and in the physical structure of the samples.

4.4.1 Dependence on temperature

The first results taken for the purposes of checking and characterising the performance of both the material and the oxidation equipment were a series of 15 minute oxidations on samples of Wafer A held at a range of temperatures. The bubbler temperature was kept at approximately 90°C and the gas flow rate, although unknown in absolute figures, was monitored and kept constant and fairly high. As shown in the background chapters, work reported by Geib et al. [4] states that the oxidation rate remains constant with increasing gas flow rate above a value of 2l/min with a bubbler temperature of 90°C. Based on this, not knowing the exact rate of gas flow was deemed acceptable if kept high and approximately constant between oxidation runs.

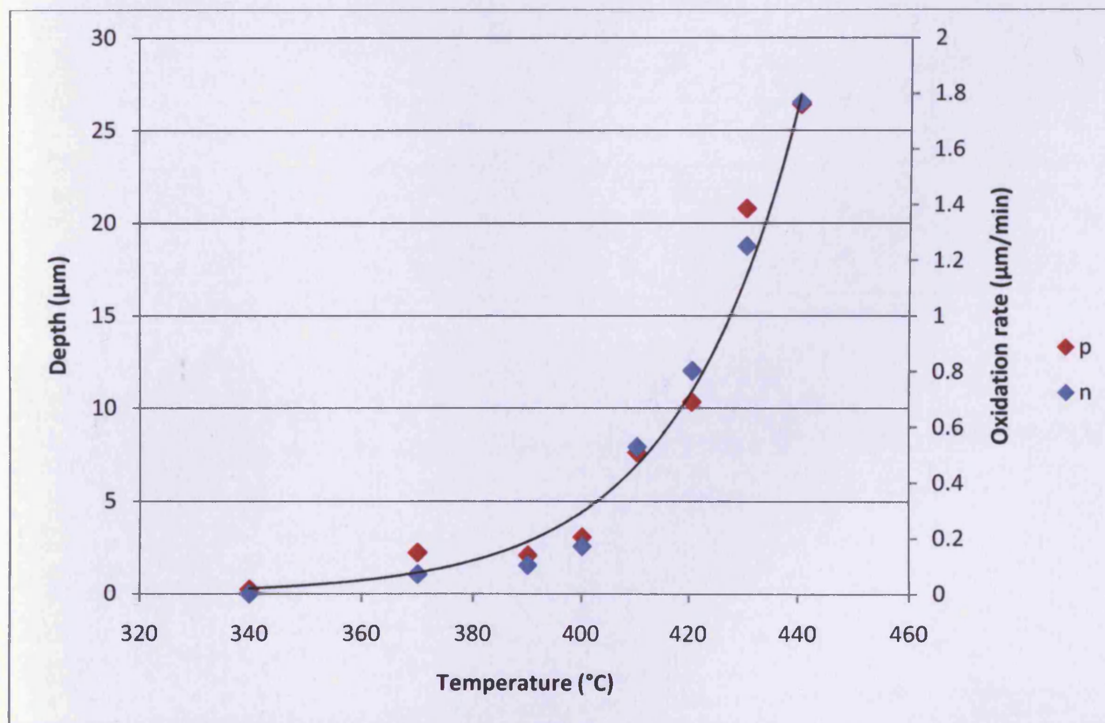

Figure 4.9: Oxidation depths and average rates achieved in Wafer A in 15 minutes

Figure 4.9 shows the depth and rate data from this experiment as taken by AFM cross-sectional topographical imaging. The rate data is a simple time-average and so assumes a constant rate during the

whole 15 minutes. This will later be shown to be a reasonable approximation for this mid-range time of 15 minutes, and is useful for comparison to literature values.

The data, in which p- and n-type doped material is distinguished, shows a very good fit to an exponential curve and, in that fact, agrees well with the data shown by Choquette et al. in their summary paper of 1997. [5] Choquette's data for $Al_{0.98}Ga_{0.05}As$ is for a layer of approximately twice the thickness and yet shows rates around 40% lower, which might seem contradictory, but the gas flow rate during oxidation and the neighbouring layers in the material are not stated, and both of these parameters would affect the result. The two sets of data are well within an order of magnitude difference of each other and, importantly, both follow the same exponential increase of oxidation rate over the same increasing temperature range. Furthermore, the activation energy, calculated by the Arrhenius relation, for the data in Figure 4.9 is 1.7±0.13eV, which agrees completely with Choquette's value of 1.75 at the relevant value of aluminium content.

A second important comparison to make to the literature is with Park et al., who used 30nm cladding region layers in edge-emitting wide-ridge lasers for the purposes of lateral current confinement in 2000. [6] The paper quotes a total oxidation depth of 14µm achieved in 4 minutes at 470°C, equating to an average rate of 3.5µm/min. This very relevant data point sits approximately 20% below the extrapolation of the exponential fit line in Figure 4.9. This discrepancy, although relatively small given the large number of variables involved, could certainly be explained by the turning at high temperatures of the rate-determining mechanism of the oxidation towards diffusion of reactants, rather than the reaction itself, as was discussed in the background chapters.

4.4.2 Dependence on time and dopant species

The next development of this experiment was to measure the progression of the oxidation front at a single temperature over a range of times, as this was the information that would be necessary to oxidise later devices to a controllable depth. Using the data above, the temperature of 410°C was chosen for a number of reasons. At 410°C the oxidation rate is moderate, allowing reasonable depths to be achieved in manageable times, but at the same time the temperature is low enough to expect a linear progression of the oxidation front to remain for fairly long oxidation times. By avoiding the non-linear progression, which has just been suggested to be present at 470°C, greater control and predictability over the total oxidation depth can be found.

Oxidation time was varied up to 30 minutes. It was suspected, given that the tubes leading to the oxidation chamber would be cold before every run began, that oxidation times below one minute may not yield accurate results due to steam condensing before reaching the sample.

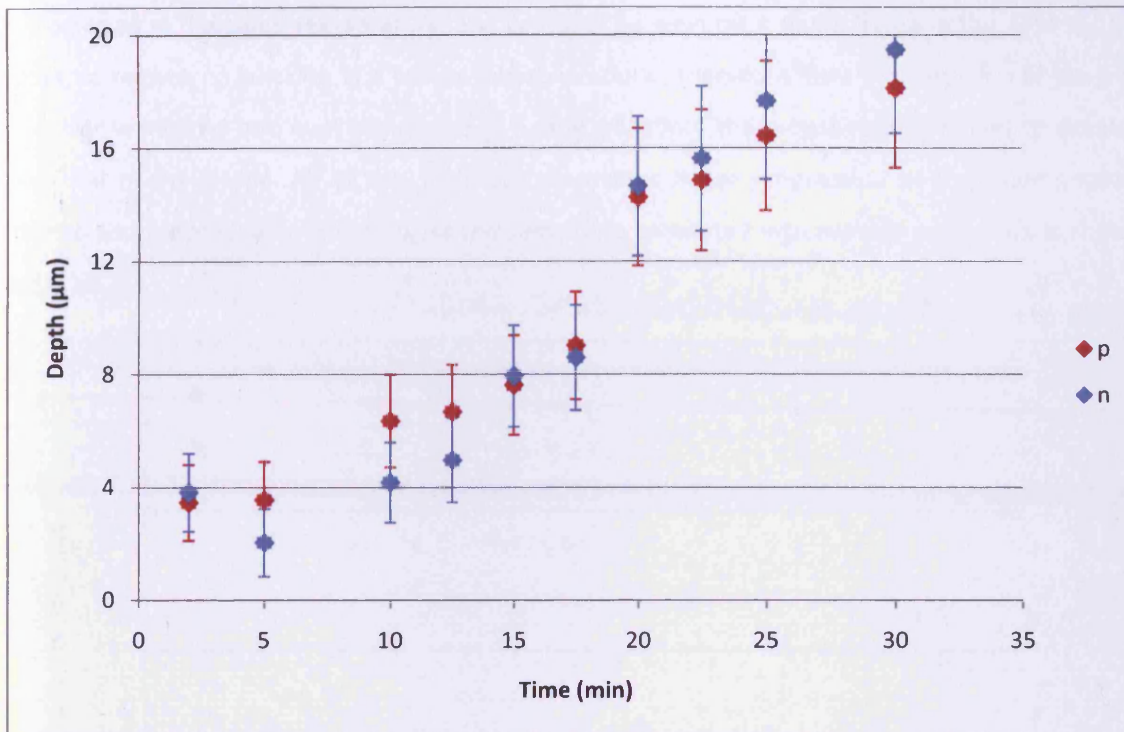

Figure 4.10: Oxidation depths achieved at 410°C in Wafer A

The data from this stage of the experiment is shown in Figure 4.10. Although there is an apparent break in the data at 20 minutes, there is no evidence for anything but a linear relationship between time and depth. What is represented, however, is definite difference of the oxidation rate between the two layers. Discounting the first pair of points, where the oxidation rate is unusually high, it appears that the p-type layer achieves a greater depth than the n-type during the first five minutes, giving it a 'head start'. However, across the graph, the difference between them seems to diminish steadily, and then grow in the opposite direction, leaving the n-type clearly ahead at 30 minutes. This could potentially be due to slight differences in aluminium composition or layer thickness, but as this material was grown by MBE, which is known to give the highly accurate layer interfaces, it assumed that this possibility is remote. Ideally though, a test structure with many differently doped layers of different thicknesses would be grown.

It could be said that the differences observed here between p- and n-type are within experimental error and are, therefore, not reliable. However, the errors bars displayed in this figure are derived from fluctuations in the oxidation temperature and from tracking errors in the AFM, as mentioned in the methods chapter, and relate only to the absolute depth values. The relative differences between each pair of points suffer neither of these sources of error as they are necessarily from the same sample, and so are oxidised at the same temperature, and can both be seen on a single frame in the AFM results in all cases, so require no tracking. It is reasonable to conclude, therefore, that the oxidation of the p-type layer is higher initially, but, over the course of a long oxidation, the n-type rate increases to eventually exceed that of the p-type. All of this, including the overall linear progression of the oxide growth, is confirmed and made clearer by looking at the same data converted into average rates. This is shown in Figure 4.11.

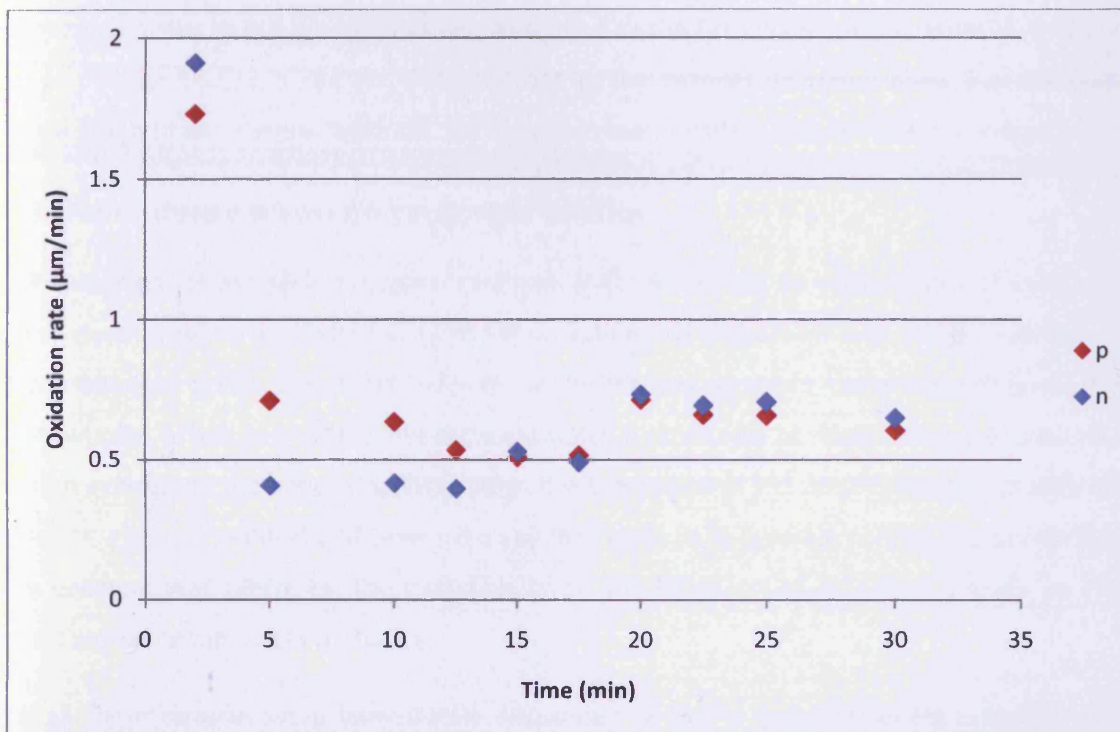

Figure 4.11: Oxidation rates achieved at 410°C in Wafer A

This dependence of the oxidation rate on the dopant species is a very significant one. The only previous major work on the subject is that presented by Kish et al. in 1992. [7] This work reports that it is the p-type, not the n-type, that has the faster rate, by up to a factor of four. The samples used in the paper, however, are all broad-area surface layers, unlike the thin, buried layers under study here. In a way,

these two sets of data agree. Looking again at Figure 4.11, it is clear that the oxidation rate of the p-type material is indeed greater than that of the n-type between times of two and 15 minutes. What the data suggests then, is that there is a mechanism affecting the overall oxidation rate that favours n-type material increasingly as the depth of the oxidation increases.

Unfortunately, it cannot be determined from these results whether it is the depth of oxidation or the fact that the layers are thin and buried that invokes this mechanism, as Kish's data goes only to a maximum depth of around 1µm, and the effect is not clearly observed here until a depth of approximately 8µm. The only other difference between the two sets of data is that Kish uses 60% aluminium material and so is strongly diffusion rate limited, whereas the results here are from 98% aluminium AlGaAs and are clearly reaction rate limited. Kish attributes the dependence of the oxidation rate on doping to the concentration of charged crystal defects, which control the rate of diffusion of oxidant species, and which are, themselves, controlled by the Fermi level of the material. It is possible then, as this mechanism is not the dominant one in the samples discussed here, that although the diffusion rate is faster in p-type material, the chemical reaction rate is actually faster in n-type material.

4.4.3 Dependence on etched sidewall profile

The final element of this series of experiments on Wafer A involves an earlier batch of mesa samples that are identical to the previous ones, with the exception that these were wet, rather than dry-etched. As a consequence of this, the etched sidewalls of the samples were not straight like those of the dry-etched samples. In fact, in all of the wet-etched samples, a very heavy horizontal V-shaped undercut was seen that centred on the upper oxidisable layer. It is the nature of the acid/hydrogen peroxide etch to oxidise the exposed material and then eat away that oxide. It is because of this mechanism that the heavy undercut was centre on the oxidisable layer, the edges of which would already be natively oxidised and more vulnerable to the etch.

Although these samples were immediately discounted as useful for determining accurate standard oxidation rates, the layer-centred undercut provided an opportunity to investigate the effects of restricting the access of the steam to the oxidisable layers. Figure 4.12 shows a cross-sectional SEM image (one minute oxidation) obtained from this experiment. Although close to the resolution limit of the SEM, the oxidisable layer is clearly visible as a darker line through the lighter grey GaAs material, and the oxidation front can be seen at the point marked. The oxidation was performed at the relatively high temperature of 480°C, as the presence of the undercut was hypothesised to reduce the oxidation rate

significantly, so higher temperatures were needed to compensate and make the oxidation deep enough to be visible.

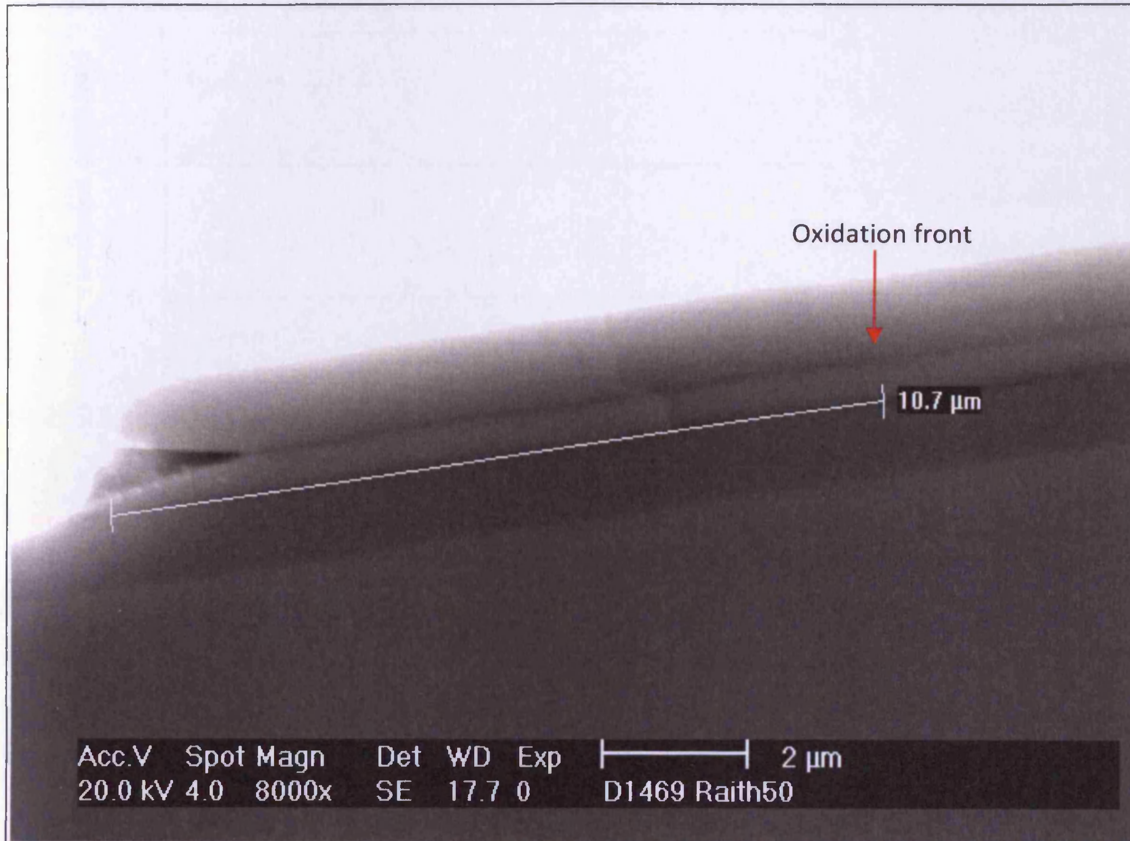

Figure 4.12: SEM image of oxidation depth after one minute

In all but the very shortest oxidation times, the n-type layer, which was not affected by the undercut, was seen to have oxidised completely through the 100µm structure. The p-type layer showed much slower oxidation, producing the data shown in Figure 4.13. The results show an initial depth of approximately 5µm, which remains unchanged until a time of one minute, when the rate of oxidation becomes very high. This rate then tails off rapidly to what appears to be a total oxidation depth limit of about 20µm.

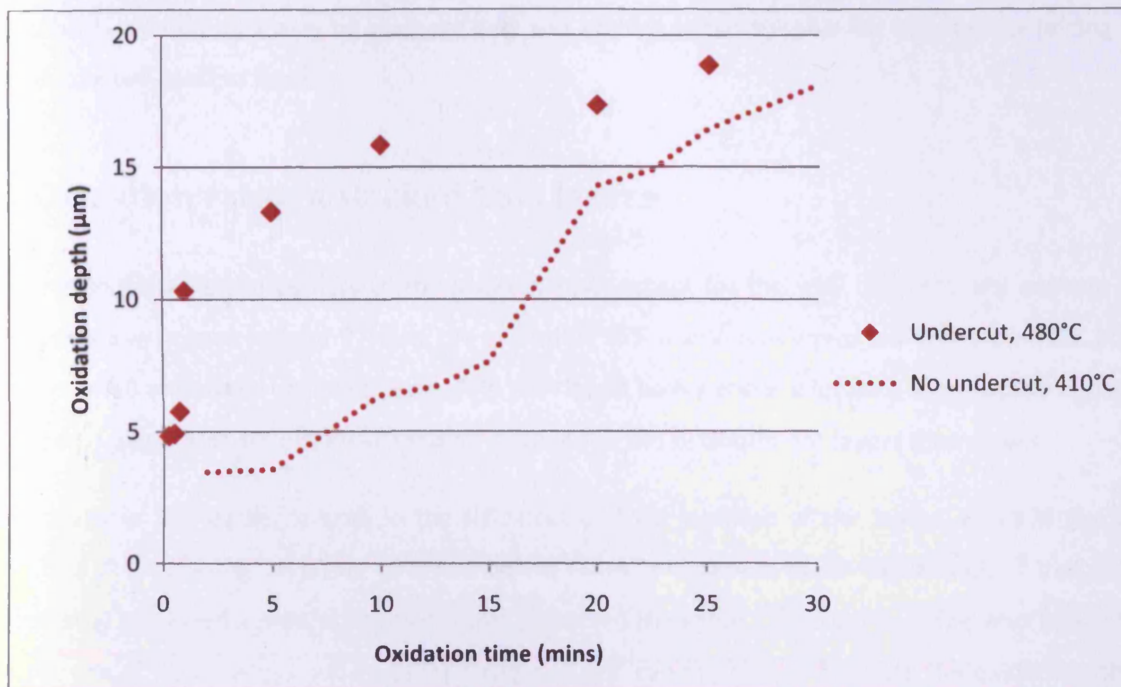

Figure 4.13: Oxidation depth over time for 30nm, p-doped layers. Undercut samples at 480°C, normal samples at 410°C

Included on the graph for comparison is the p-type data from Figure 4.10, taken at 410°C. This data, shown with a dotted line, is approximately linear, indicating that the rate of oxidation is governed by the rate of the chemical reaction, not by diffusion of the reactants. The new data from the undercut samples, however, is very different. Within the first minute, whilst the oxidation depth is small and diffusion is not a problem, the full, linear 480°C reaction rate is seen. After this point, the oxidation rate decreases rapidly until it is significantly less even than that of the 410°C data with straight sidewalls. The average rate of oxidation over the latter 15 minutes in the undercut samples is only 1.1% of the maximum rate in the samples, observed between 45 and 60 seconds, and the total oxidation depth is less than 20µm.

To compare these values to literature, Choquette et al. state that reactions between ≈350°C and 500°C will only become diffusion-limited after 50µm total depth (with 84nm thick layers), [5] and Ochiai et al., who show diffusion-limited oxidation curves comparable to those shown here, only observe a reduction in oxidation rate of around 50% under similar conditions. [8] Altogether this is good evidence that not only is oxidation in the undercut samples diffusion-limited, but it is limited to a far greater extent than would be caused by diffusion in samples with straight sidewalls. It is therefore the conclusion of this section of the experiment that the sidewall profile of any samples to make use of oxidation technology is highly important. Furthermore, given that the V-shaped undercut was very definitely centred on the

oxidisable p-type layer, it may be the case that wet etching is not suitable for samples containing such high aluminium-content layers.

4.5 Oxidation rates in stacked 1nm layers

This section presents similar data to the previous one, except for the very different and entirely novel arrangement of layers of Wafer B. Here, the oxidisable 95% aluminium layers are only 1nm thick, but are in stacks of 10 alternated pairs with 1nm 75% aluminium layers and are located in the active region. Six of these 10-pair stacks are placed alternately around the five quantum dot layers themselves.

The novelty in this design is both in the thickness and the position of the layers. In 1996 Kim et al. published data showing oxidation of layers below 11nm in thickness to be virtually nil, [9] and in 2001 Osiński et al. provided a theoretical model that predicted the same. [10] However, it has also been shown that the use of digital alloys and so-called 'feeder layers' can be used to enhance the oxidation rate. [11] For this reason, it was expected that this arrangement of layers would oxidise at rates similar to a 20nm layer of 85% aluminium material whilst retaining the benefit of allowing tunnelling current through each 1nm layer individually. Figure 4.14 shows an AFM of the oxidation fronts observed in the dry-etched mesa samples created from this wafer. The oxidation has proceeded from the bottom of the image upwards.

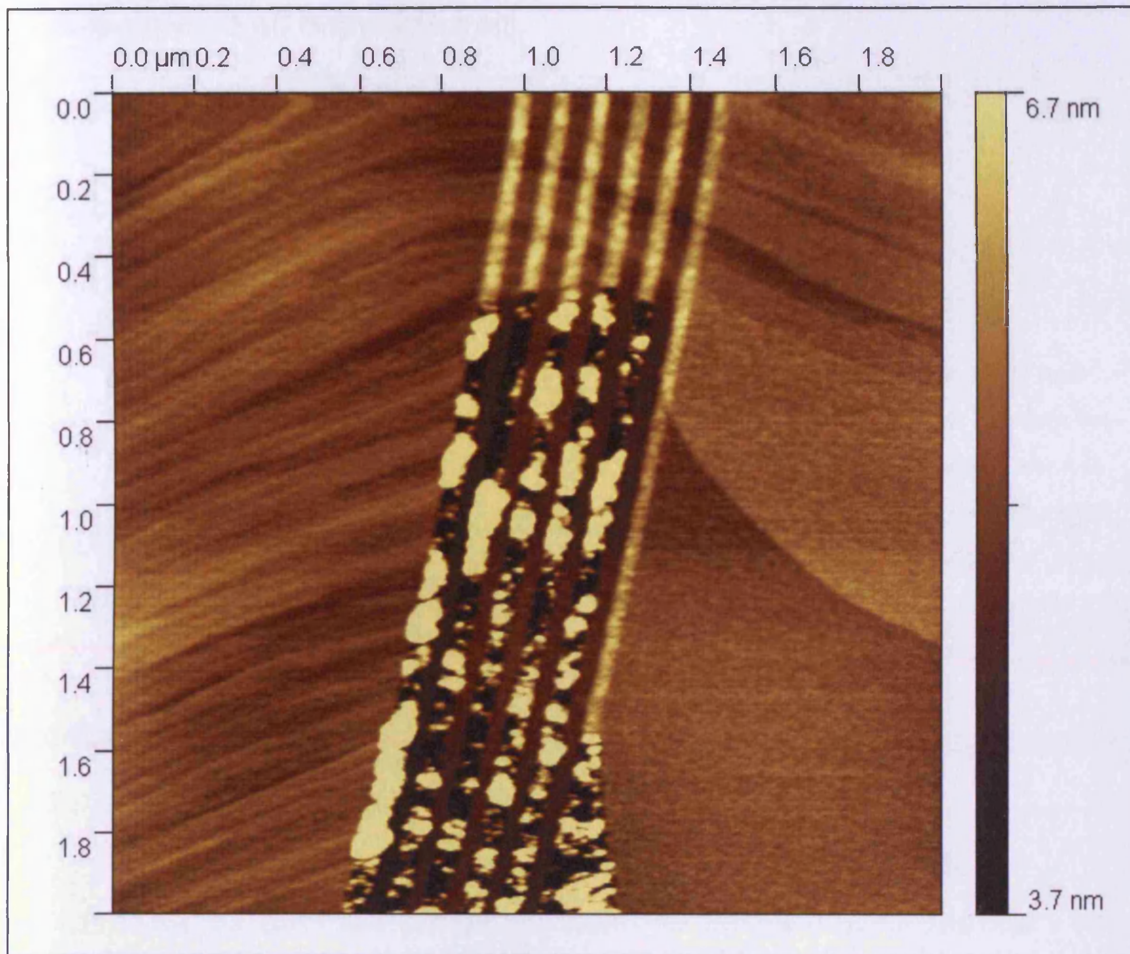

Figure 4.14: AFM image of oxidation fronts as seen in Wafer B (top of wafer is on right of image). Oxidation performed at 440°C for 30 minutes

The image quite clearly shows the six separate stacks of oxidisable layers and the slightly different depths they have all oxidised to. The central stacks have achieved a marginally greater depth than those on the outside by using the outer layers as 'feeder layers' in just the same way as a single layer of the same total thickness (≈300nm) would produce a curved oxidation front. The very uppermost stack (right-hand side of the image) has behaved differently for reasons that will be discussed later.

It should be noted that the x dimensions of the image above appear approximately 50% larger than they should (total active region thickness appears ≈500nm). This is due to the method of mounting the samples and the fact that the scanning direction is in the x axis. The y dimensions are accurate to the errors already discussed, as is confirmed in later chapters.

4.5.1 Dependence on temperature

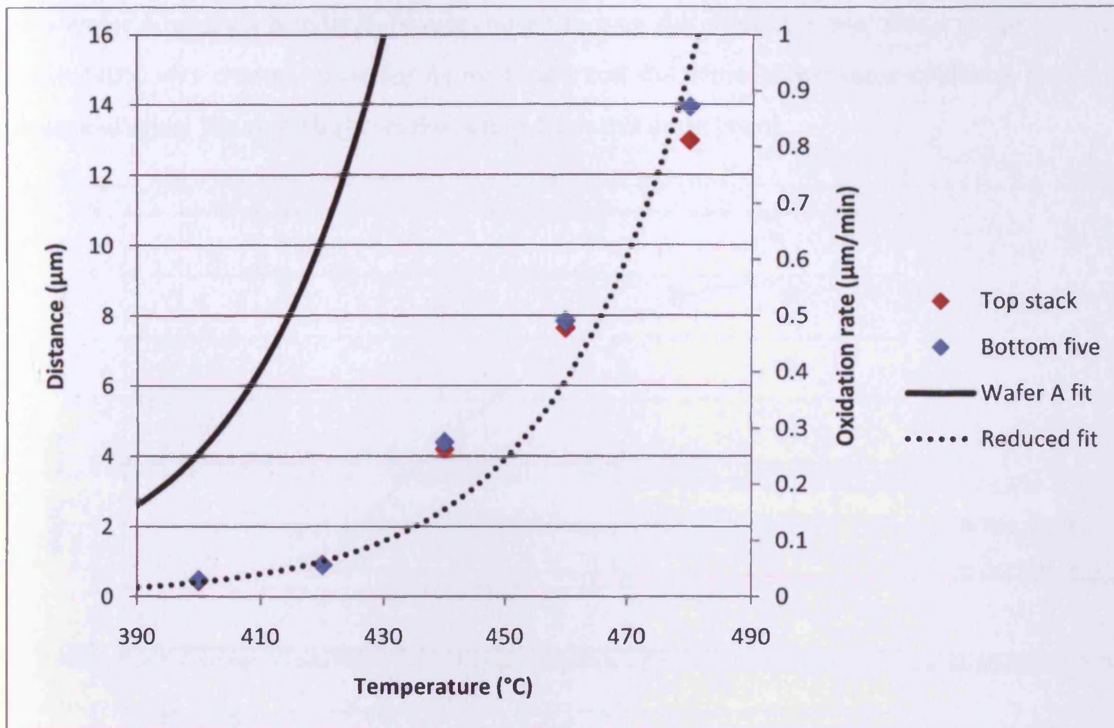

Figure 4.15: Oxidation depths and average rates achieved in Wafer B in 15 minutes

Figure 4.15 shows the depth and average rate results for oxidations performed over a range of temperatures for 15 minutes each. Also shown on the graph are two fit lines. The first (solid line) is the exponential fit line relating to the same data taken for Wafer A in Figure 4.9, and the second (dotted line) is the same exponential but reduced by one order of magnitude. The purposes of these lines are, in the first case, to compare the oxidation rates of the two wafers, and in the second, to compare to literature values. Taking information again from Choquette, [5] it can be estimated that in going from a 30nm layer of 98% aluminium (Wafer A) to what is effectively a 300nm layer with an average aluminium content of around 80% (Wafer B), the oxidation rates at all temperatures should drop by a total of approximately one order of a magnitude (aluminium percentage reduction ≈ -1.4 orders and thickness increase ≈ +0.4 orders). Looking at the Wafer B data and the 'reduced fit' line on the graph, although there is some scatter, the agreement is good. This demonstrates that the achievable oxidation rates have not suffered by the use of 1nm stacked layers.

4.5.2 Dependence on time

As with Wafer A, a single temperature was chosen to map the oxidation rate over a range of times. In this case 440°C was chosen, again for its moderate and therefore controllable oxidation rate and its relatively low value. Figure 4.16 shows the results from this experiment.

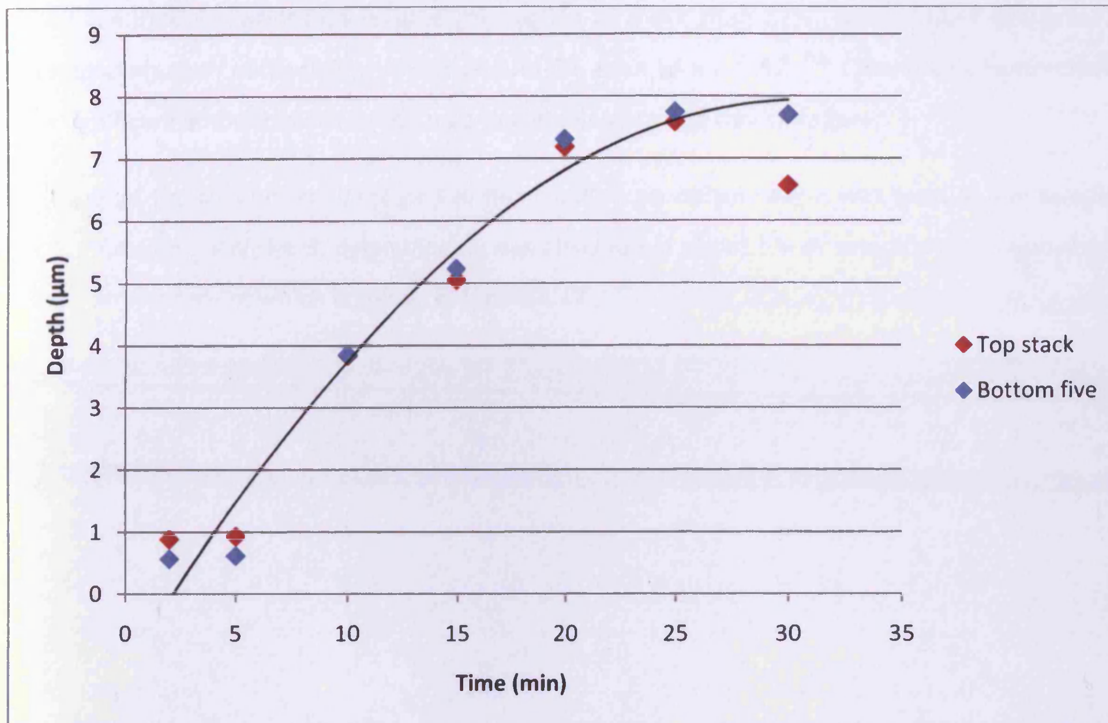

Figure 4.16: Oxidation depths achieved at 440°C in Wafer B

In this case, as distinct from the equivalent data shown for Wafer A, the data follows, fairly clearly, a parabolic curve that tends towards some eventual limit. This is fairly good evidence that at this temperature and with the layer system present in Wafer B, the oxidation rate is limited by diffusion mechanics rather than reaction rate.

In all of the oxidation depth results presented, the zero time data point has been omitted, however, by looking at the very short oxidation times, it would appear that there is some initial depth of oxidation already present (≈0.5 – 0.8µm). This is likely to be due to a certain depth of native oxide being present before the steam oxidation process and it would seem that its presence does not significantly affect the results at longer oxidation times.

4.6 Device reliability

Although the various oxidation rates have now been mapped and can be produced repeatably, one other major factor in the reliability of devices using this oxidation technology. One observation made here is that, if caution is not taken, there is a tendency towards delamination. The oxidation of pure AlAs has been reported to result in a volume contraction of more than 12%. This number decreases with decreasing aluminium content, but is still around 6% even at x = 0.92. [12] Clearly this contraction can cause a significant amount of strain and lead to delamination and device failure.

In the cases of the two wafers discussed in this chapter, no delamination was seen in any samples of Wafer A. However, in Wafer B, delamination was observed in about 5% of samples when cleaving after oxidation, an example of which is shown in Figure 4.17.

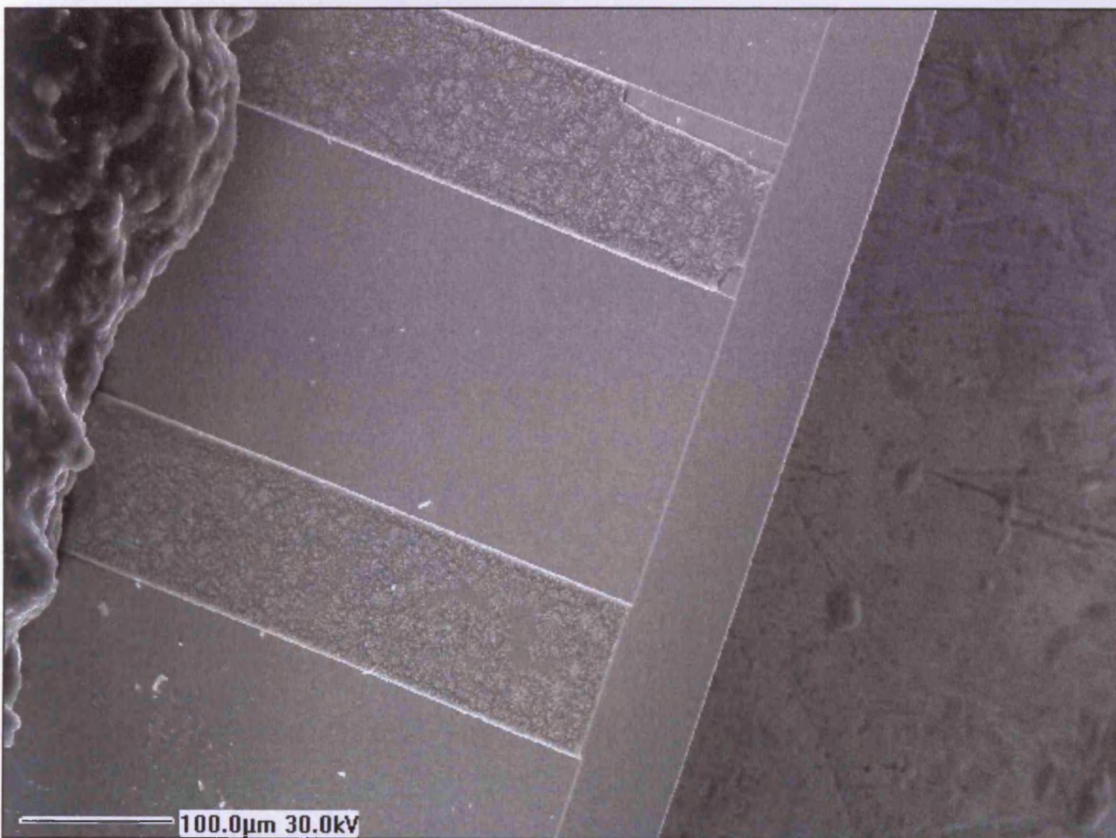

Figure 4.17: SEM image showing delamination in Wafer B (top mesa)

Digital alloys, such as the stacks used here in Wafer B, would normally serve to minimise the strain induced by the volume shrinkage of oxidised layers, and indeed this will be the case, even in the image

above. The explanation of the delamination lies with the p-type cladding layer and its proximity to the first oxidisable layer. In this wafer structure, the p-type cladding is a 1.15μm layer of $Al_{0.8}Ga_{0.2}As$, which would ordinarily oxidise at rates 100 times lower than AlAs. However, as it is placed immediately on top of the first 95% aluminium oxidisable layer, it has a very effective 'feeder layer' by which its oxidation rate is enhanced greatly. The 45% aluminium n-type cladding layer does not visibly oxidise in the same way, despite being partially exposed by the etch because its nominal oxidation rate is so much lower even than the 80% material that the feeder layer is insufficient to create the effect.

Returning to the AFM image in Figure 4.14, it can be seen that the uppermost oxidisable stack does not oxidise quite as far as the other five, as previously mentioned. Furthermore, in the lower half of the image, the oxide created by the steam oxidation process spreads much wider than the width of the stack, unlike the lower five stacks where the oxide follows the layers precisely. The latter condition is true for all samples and, by further SEM imaging, the full extent of this effect can be seen.

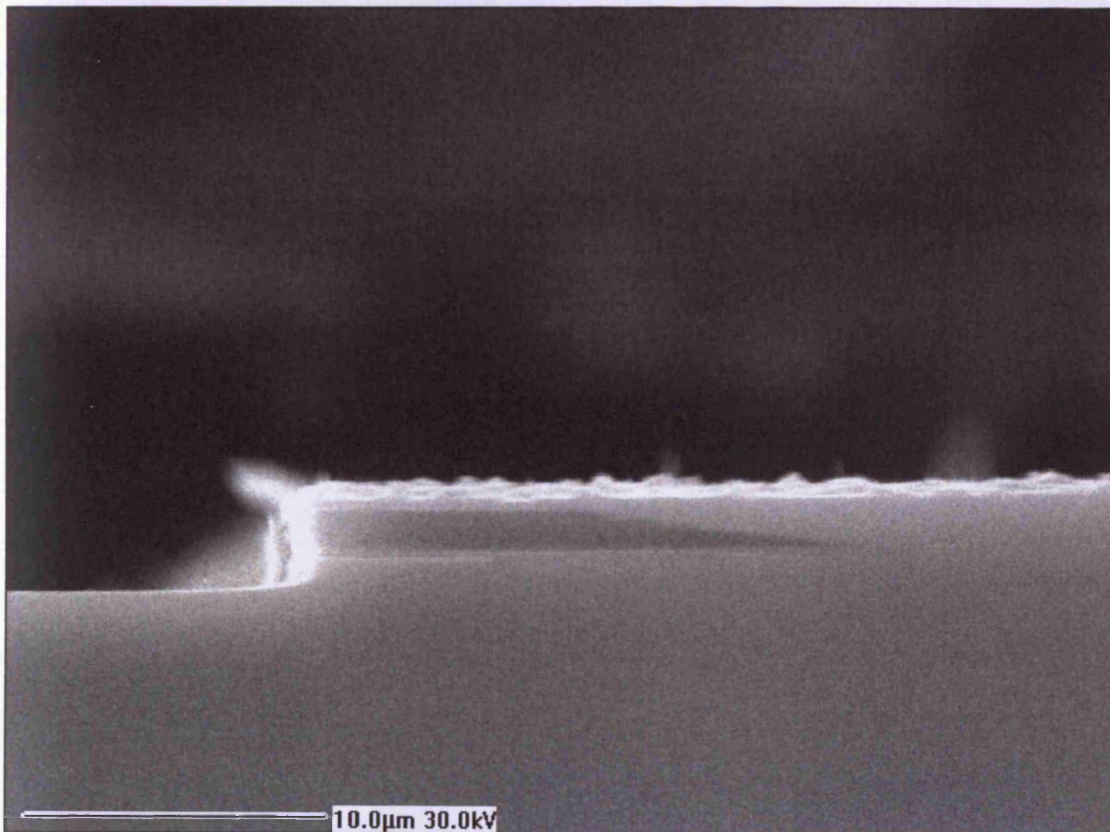

Figure 4.18: SEM image showing enhanced oxidation of p-type cladding

Figure 4.18 shows that the oxidation of the cladding layer is clearly seeded from the layer immediately below it, as evidenced in the image by its diagonal progression upwards and left through the layer until its full width is oxidised at a distance of about 8µm behind the oxidation front of the 95% aluminium layer.

Just as the higher aluminium-content layer enhances the oxidation rate of the cladding layer, so the reverse is simultaneously true, as the large cladding layer is a heavy diffusive 'sink' of oxidant species. This now explains the slightly reduced oxidation rate of the uppermost stack in comparison to the lower five. It also now becomes clear why increased incidents of delamination are observed; an approximate 10% decrease in volume of only a 100nm total material would create a fairly low degree of mechanical stress, whereas the same percentage decrease in more than one micron of material will produce considerably more. Indeed, where delamination occurred, its origin was found to be at the level of the oxidisable layers and the depth of the oxidation as is shown in Figure 4.19, where two mesas from the same oxidised sample are shown; one mesa has delaminated and the other has not.

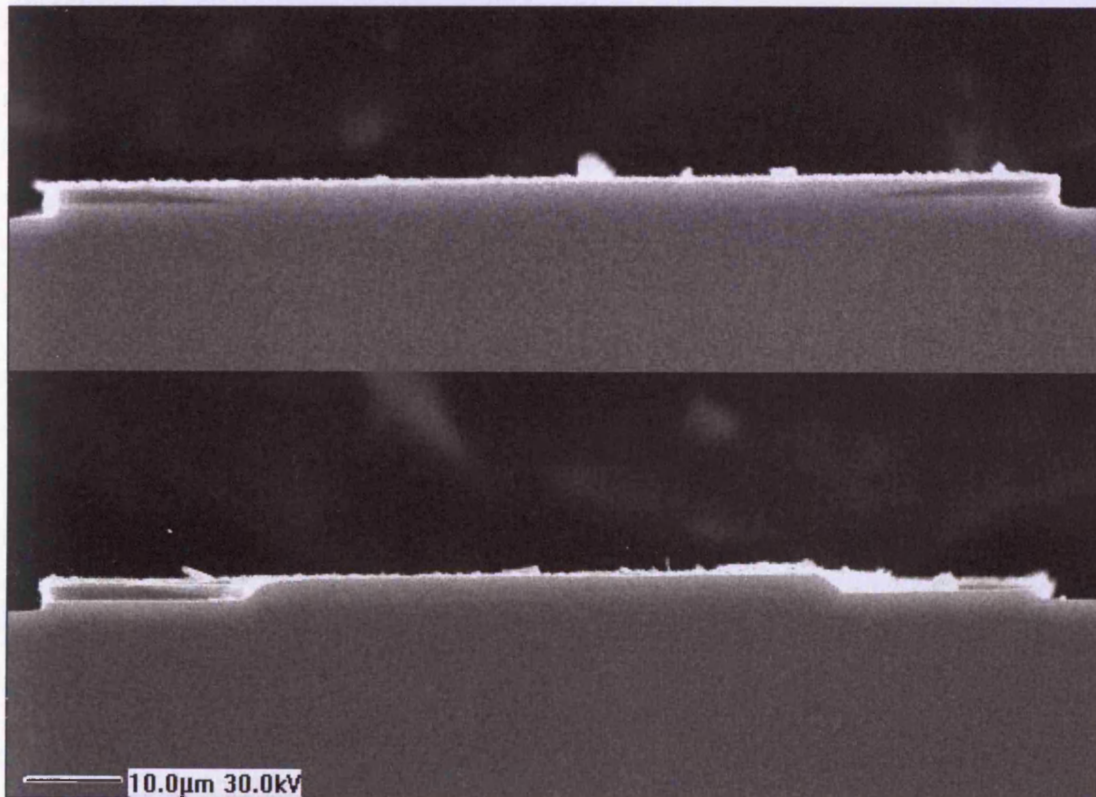

Figure 4.19: Delaminated (bottom) and surviving (top) mesas from the same sample

Clearly, this will remain an issue for devices fabricated from this material, although the investigation has revealed some of the mechanics of oxidation of neighbouring layers. Consequently, the delamination problem could be avoided in future wafer designs by adding a low aluminium-content barrier between the layers involved. In terms of the rest of this work, the frequency of incidents of delamination remains low, at around 5%. However, this figure could be reduced by keeping oxidation depths below 8μm, which, as this study has revealed, would reduce the total thickness of cladding material oxidised, and so minimise the strain induced.

4.7 Summary

As the first major set of results, this chapter has provided the basis upon which the remaining experiments will be performed. The mapping of the oxidation rates of both of the applicable wafers was a crucial preliminary step in fabrication of full, working devices and in the investigation of the effects created in those devices. Furthermore, the characterisation of the oxidation equipment and the comparison of standard layer systems to literature values have also been important in adding context and credibility to the results that followed.

Besides this introductory work, this chapter has also produced several novel, standalone conclusions. Two new wafers have been designed and simulated from scratch. This demonstrated that oxidisable layers can be incorporated into the cladding and, importantly, into the active regions without significantly degrading the electrical performance of the material.

Oxidation has been proven, for the first time, to take place in layers as thin as 1nm, when stacked in a binary alloy system, and rates have been shown to be repeatable and of high enough value to be practically employable. The effects of dopant species on lateral oxidation rate has been investigated along with that of heavy undercutting of etched sidewalls. In both of these cases, particularly in that of doping, considerable new insight has been gained. Information has been presented that is potentially contradictory to the generally accepted, although minimally reported dependence of oxidation rate on dopant species.

Finally, the limitations of the technology, with respect to the specific material used here, have been discussed; namely, the appearance of delamination in some samples. Reasons for this effect have been detailed, its likely impact on the rest of the work has been discussed, and methods to combat the problem in future works have been suggested.

4.8 References

1. **Winston, D.** *Physical Simulation of Optoelectronic Semiconductor Devices, PhD.* University of Colorado at Boulder, 1996.

2. *The effect of p doping in InAs quantum dot lasers.* **Sandall, I. C.; Smowton, P. M.; Walker, C. L.; Badcock, T.; Mowbray, D. J.; Liu, H. Y.; Hopkinson, M.** 11 : Applied Physics Letters, 2006, Vol. 88.

3. **Jenkins, C.** *Cross-Sectional Atomic Force Microscopy of III-V Semiconductor Device Structures.* Cardiff University, PhD, 2004.

4. *Fabrication issues of oxide-confined VCSELs.* **Geib, K. M.; Choquette, K. D.; Hou, H. Q.; Hammons, B. E.** SPIE, 1997.

5. *Advances in selective wet oxidation of AlGaAs alloys.* **Choquette, K. D.; Geib, K. M.; Ashby, C. I.; Twesten, R. D.; Blum, O.; Hou, H. Q.; Follstaedt, D. M.; Hammons, B. E.; Mathes, D.; Hull, R.** 3 : Journal of Selected Topics in Quantum Electronics, 1997, Vol. 3.

6. *Low-threshold oxide-confined 1.3-μm quantum-dot laser.* **Park, G.; Shchekin, O. B.; Huffaker, D. L.; Deppe, D. G.** 3 : IEEE Photonics Technology Letters, 2000, Vol. 13.

7. *Dependence on doping type (p/n) of the water vapour oxidation of high gap AlGaAs.* **Kish, F. A.; Maranowski, S. A.; Höfler, G. E.; Holonyak, N.; Caracci, S. J.; Dallesasse, J. M.; Hsieh, K. C.** 25 : Applied Physics Letters, 1996, Vol. 60.

8. *Kinetics of thermal oxidation of AlAs in water vapour.* **Ochiai, M.; Giudice, G. E.; Temkin, H.; Scott, J. W.; Cockerill, T. M.** 14 : Applied Physics Letters, 1996, Vol. 68.

9. *Lateral wet oxidation of AlGaAs-GaAs depending on its structures.* **Kim, J. -H.; Lim, D. H.; Kim, K. S.; Yang, G. M.; Lim, K. Y.** 22 : Applied Physics Letters, 1996, Vol. 69.

10. *Temperature and thickness depenence of steam oxidation of AlAs in cylindrical mesa structures.* **Osiński, M.; Svimonishvili, T.; Smolyakov, G. A.; Smagley, V. A.; Maćkowiak, P.; Nakwaski, W.** 7 : IEEE Photonics Technology Letters, 2001, Vol. 13.

11. *Barrier-layer-thickness control of selective wet oxidation of AlGaAs for embedded optical elements.* **Blum, O.; Ashby, C. I.; Hou, H. Q.** 21 : Applied Physics Letters, 1997, Vol. 70.

12. *Selective oxidation of buried AlGaAs versus AlAs layers.* **Choquette, K. D.; Geib, K. M.; Chui, H. C.; Hammons, B. E.; Hou, H. Q.; Drummond, T. J.; Hull, R.** 10 : Applied Physics Letters, 1996, Vol. 69.

Chapter 5 : Effects from Lateral Oxidation of Active Region Layers

5.1 Introduction

It has already been shown, by simulations earlier in this work, that by using very low layer thicknesses and stacking them alternately with lower bandgap material, high aluminium-content layers can be incorporated directly into the active region of devices with barely any detriment to their electrical performance. It has also been shown practically that these digital alloy stacks do, in fact, oxidise at appreciable rates, thanks to their proximity and to diffusion of oxidant species through neighbouring layers. In isolation, every indication leads to the conclusion that such thin layers would not oxidise at all.

These results are novel in themselves, however, if the material can be shown to make working lasers and if the partially oxidised material successfully confines the current or the optical mode, significant advances in laser fabrication could be possible. This chapter will, therefore, focus mainly on demonstrating the basic capabilities required of the wafer to make lowered threshold current lasers, and will go on to discuss some potential applications.

5.2 Electrical performance

The first step in assessing the performance of the wafer is to confirm the electrical properties predicted by the simulations. Although this chapter will deal otherwise exclusively with Wafer B, Wafer A is included in this section for ease of comparison. Figure 5.1 shows IV traces taken at 300K for 2mm devices of both wafer types, as well as a third control wafer, which is identical to the other two in all aspects other than the oxidisable layers.

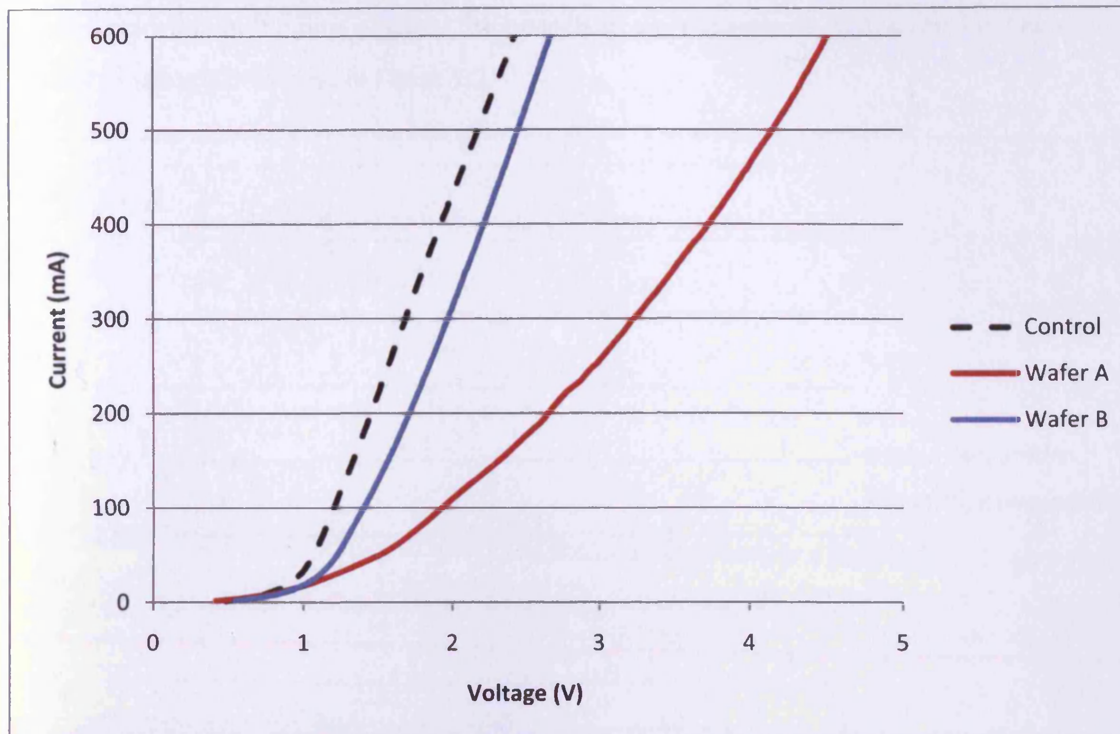

Figure 5.1: IV comparison between Wafers A, B and a control wafer using 2mm, 50μm stripe devices

The results clearly show that, whilst Wafer A does not perform quite as well as predicted, Wafer B has an on resistance very close to that of the control wafer, which contains no oxidisable layers, and a turn-on voltage only slightly higher. The marginal under-performance of Wafer A could be due to any one of a number of reasons; non-optimised growth temperatures of the oxidisable layers, overestimation of tunnelling current in the simulations, for example. Both wafers, however, are easily good enough to fabricate working laser devices, and with double the total thickness of oxidisable material over the whole structure in Wafer B, compared to Wafer A, it is reasonable to conclude that tunnelling current is playing a significant role in the performance of Wafer B.

5.3 Sub-threshold near-fields

The advantage of lateral oxidation within the active region is that the partially oxidised layers should have the capability of confining both the current and the optical mode simultaneously. A simple method of testing whether the first of these two effects is present is to measure the width of the near-field below threshold. For this purpose, a batch of 125μm-wide stripe mesas were etched and metallised. These structures were oxidised to various depths before their facet ends were cleaved, thereby allowing

only lateral oxidation in the final devices. The near-fields were measured as described in Chapter 3 and the FWHM values are presented in Figure 5.2.

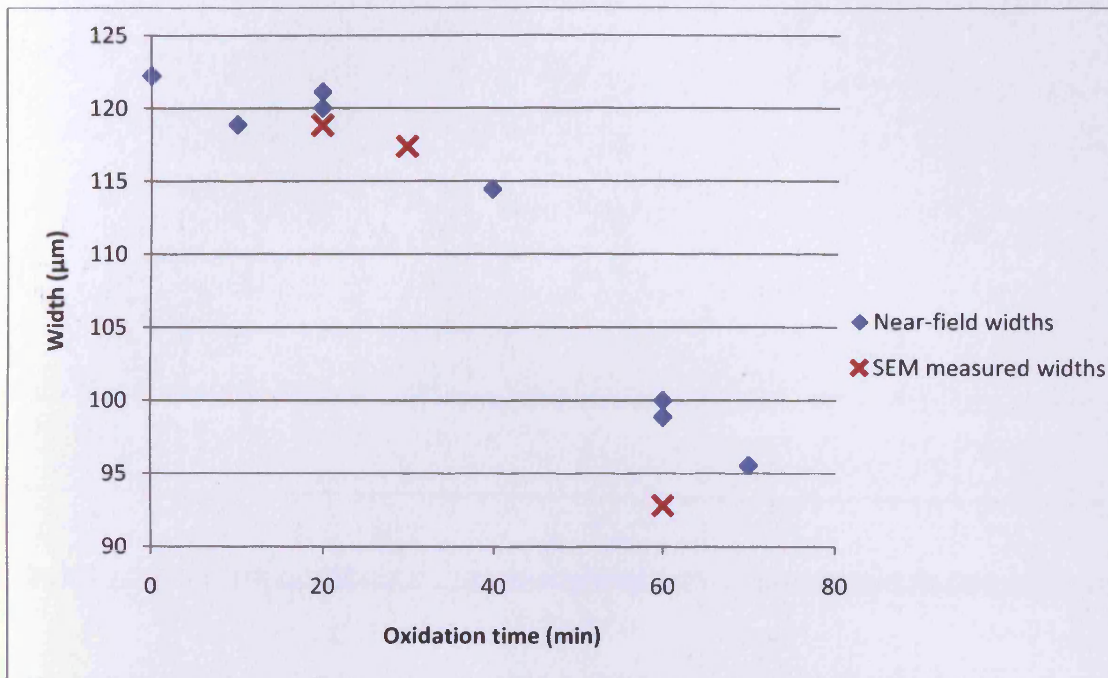

Figure 5.2: Reduction of sub-threshold near-field width by oxidation

All near-fields were measured at 300K and at a current of 200mA, well below the expected threshold of the devices. This ensures that the near-field width is due only to the path of the current (defined by the oxide), and not any modal behaviour, which would further narrow the width. This leads to fairly square near-field shapes, where the full-width, half-maximum parameter is very applicable and where errors are fairly small (±3µm). One such near-field is shown in Figure 5.3.

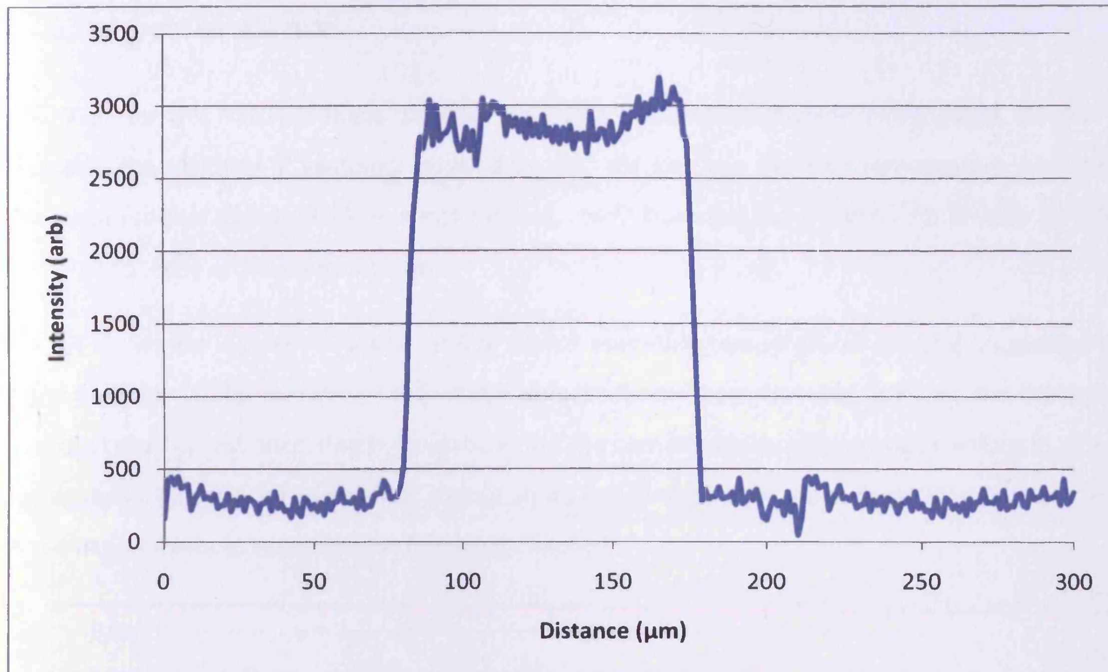

Figure 5.3: Typical sub-threshold near-field shape

At the time of fabrication, the oxidation rate in this wafer was not yet known more accurately than an order of magnitude estimation, hence, the size of the effect compared to the total width of the devices is relatively small. Ideally the mesas would have been etched much narrower. However, the results do still firmly demonstrate a reduction in the near-field width (FWHM) by current confinement with increasing oxidation time. The red crosses on the first graph represent a few samples from the batch where the depth of the oxidation front was imaged directly in an SEM, giving information on the total width of unoxidised material still left. The two sets of data show very good agreement, with the near-field widths being slightly larger than the material widths due to the relatively large proportion of light being emitted off-angle at these sub-threshold currents.

Unfortunately, due to the unforeseen enhanced oxidation rate of the p-type cladding layer, as detailed in Chapter 4, it cannot be categorically said that the thin, active region layers of Wafer B alone are sufficient to confine the device current. However, with a resistivity between 10^{12} and $10^{14}\Omega$cm and a reported breakdown voltage of between 4 and 11MVcm^{-1} [1] [2] (equating to 8 to 22V across each of the six 20nm stacks), it would seem highly likely that the current confinement would still occur to the same extent even without the extra oxide thickness added by the cladding layer.

5.4 Laser performance

Finally, then, for this batch of mesa devices, the lasing characteristics were investigated. Due to their large width, the majority of samples, oxidised or not, did not lase at room temperature. However, it should be noted that standard 50μm stripe devices, made from the same wafer, did so with thresholds comparable to those of the control wafer.

Figure 5.4 shows the LI traces at a selection of device operating temperatures for one unoxidised mesa and one oxidised for 60 minutes – one of the longest times used. Notable features are the reduced threshold at the highest operating temperature and the comparable maximum slope values in all cases, indicating both that there is no radical change in α_i, and that no saturable absorption is present. This latter point will relate to results in the following chapter.

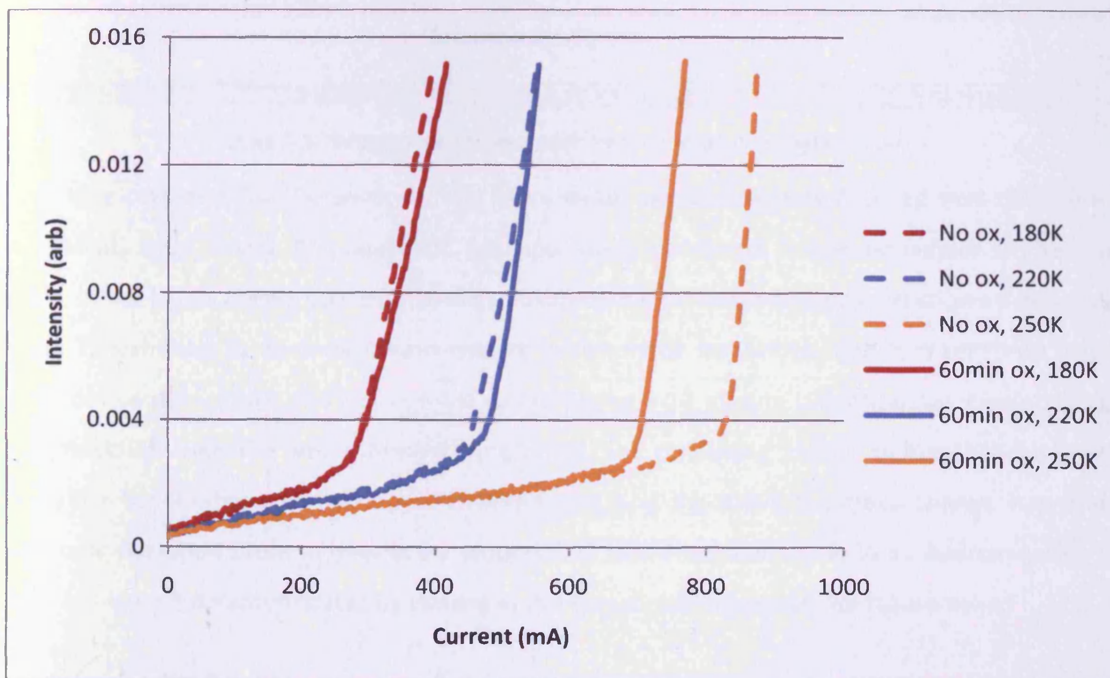

Figure 5.4: Comparison of oxidised and unoxidised mesa LI characteristics

As an extension of this, threshold values were recorded for several oxidation times over a wide range of temperatures. These are displayed in Figure 5.5, where a cursory inspection of the data reveals no obvious detriment caused by the oxidation process.

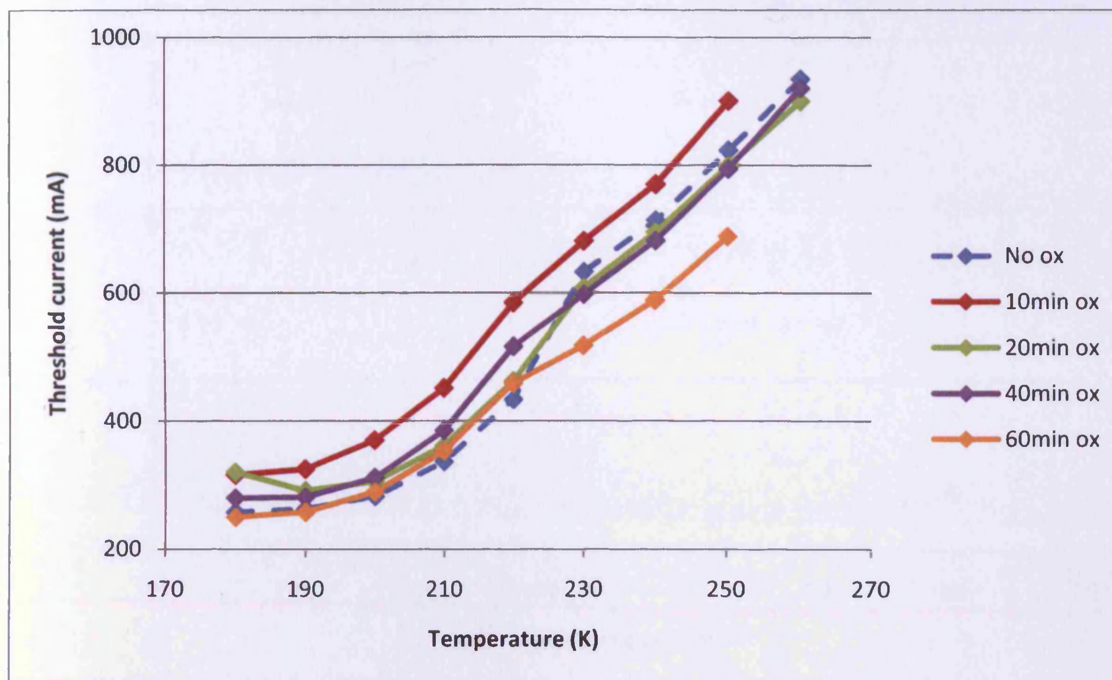

Figure 5.5: Threshold against temperature for a range of oxidation times

With further consideration, however, a little more detail can be suggested. Given that these are very wide devices, even where the near-field has now been narrowed, mode behaviour is likely to be complex. In the figure above, the data for the unoxidised device (blue dashed line) shows a discontinuity at 230K. This is likely to be a significant change in the mode behaviour, and it is mirrored in all the oxidised device data, albeit at 220K in most cases. Whilst a 10 minute oxidation has produced slightly raised thresholds over the whole temperature range, the remaining traces, below this discontinuity, show no significant difference from the unoxidised case. Looking above the mode change, however, the three longer oxidation times all produce a reduction in threshold, and do so in an approximately linear fashion, as Figure 5.6 demonstrates by looking at the thresholds for one temperature value.

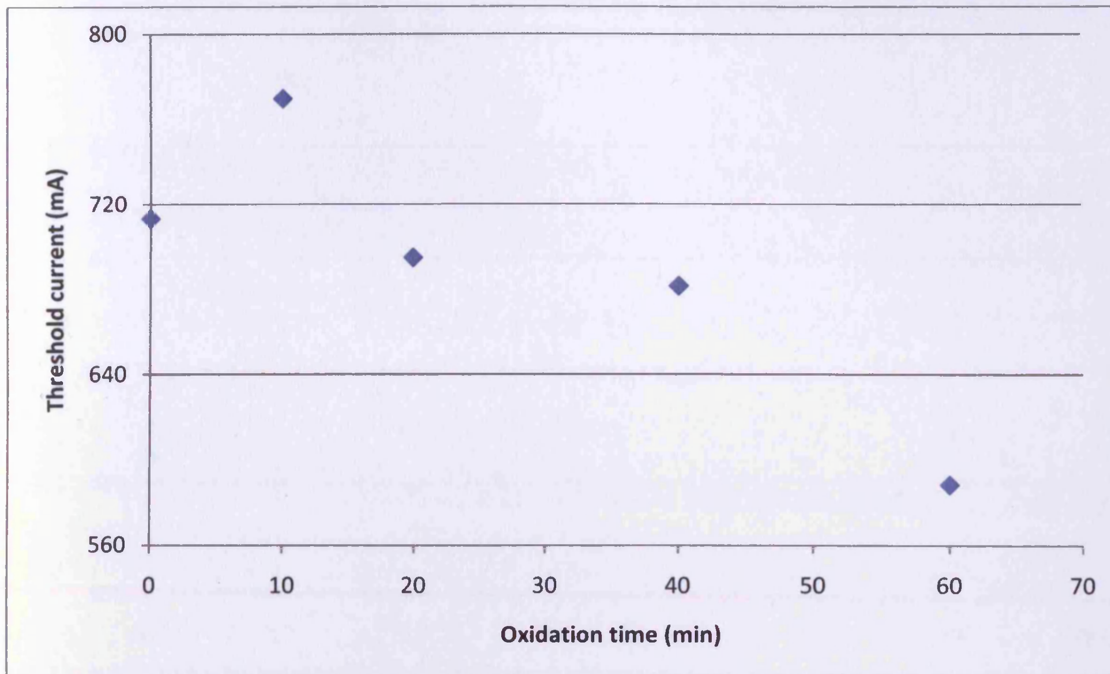

Figure 5.6: Threshold against oxidation time at one temperature, 240K

If viewed in threshold current density, taking into account the now different widths of the devices, as in Figure 5.7, the same data shows an approximately constant value for all oxidation times (apart from 10 minutes, where current density is unusually high, suggesting a poor device). This confirms that it is purely the current confinement that has lowered the threshold currents, and no significant extra loss or damage has been added into the devices by the oxidation process.

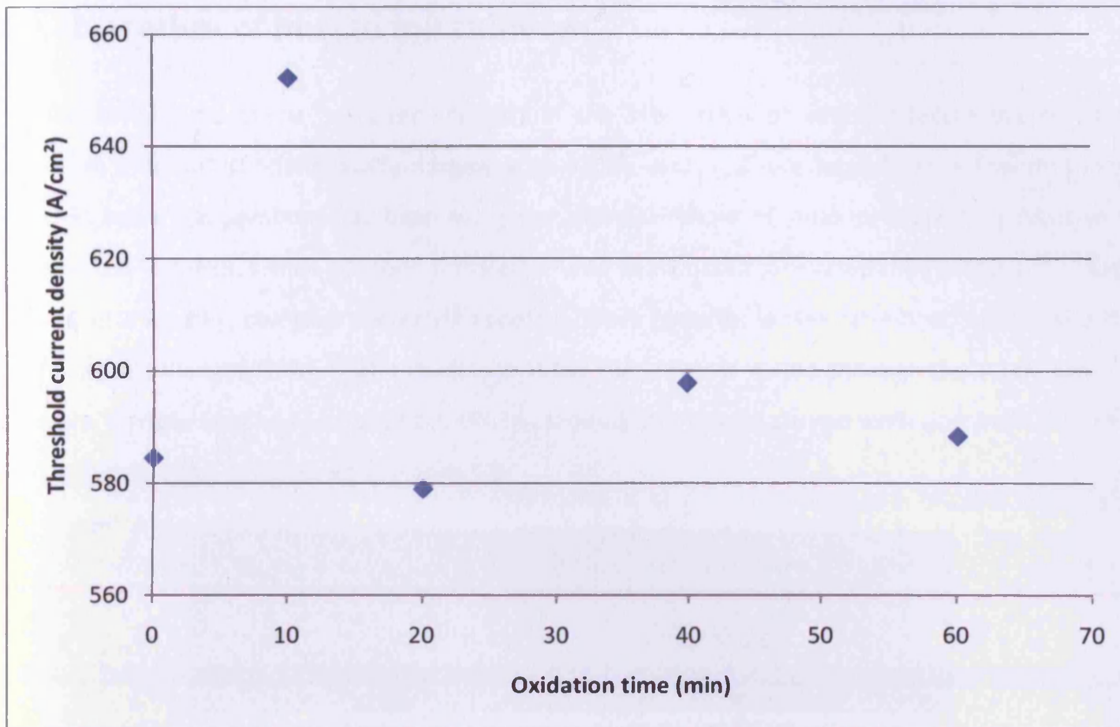

Figure 5.7: Constant threshold current density for all good oxidised devices

Generally, in etched laser structures such these, non-radiative recombination through sidewall surface roughness can be a significant source of loss. A plausible explanation for the trends seen in the all the data above is as follows. At low temperatures, below the mode change discontinuity, the selected mode of the devices is, by chance, already narrow enough not to be affected by the presence of the oxide even at its maximum depth. At higher temperatures, the mode changes to one of a higher order, with a wider lateral profile where it is now in range of the oxidation front, and so, benefits from the current confinement and threshold is reduced. It is possible that in the case of the 10 minute oxidation that the oxidation depth is not sufficient to outweigh the extra losses that may have been added by oxidising the sidewall as a whole, although, given the proximity of the oxidisable layer to the active region, it is more likely that the device generally has a higher threshold for some unrelated reason.

To test this theory, the near-fields of the devices were recorded over the whole range of relevant temperatures. Their behaviour was so heavily filamentary, due to their large width, that definite conclusions were impossible to draw. Ideally, this experiment would be repeated with much narrower devices, which would promote single mode behaviour and so, be easier to analyse.

CARDIFF UNIVERSITY ★ PRIFYSGOL CAERDYDD ★

5.5 Fabrication of buried microlenses

For quite some time there has been interest in the fabrication of semiconductor microlenses for integration, monolithic or otherwise, largely with VCSEL arrays. There have been a few methods for doing this, but most common has been using the thermal reflow of small photoresist islands to form hemispherical features, which are then transferred into semiconductor material by etching. [3] [4] Clearly, this method is lengthy, complex and hard to control. More recently, lenses have been made using digital alloy stacks of AlAs and GaAs. [5] The thickness of the AlAs layer is varied through the stack, producing, effectively, a graded AlGaAs composition. When oxidised, this gives a curved oxidation front that can act as very effective lens, as is shown in Figure 5.8.

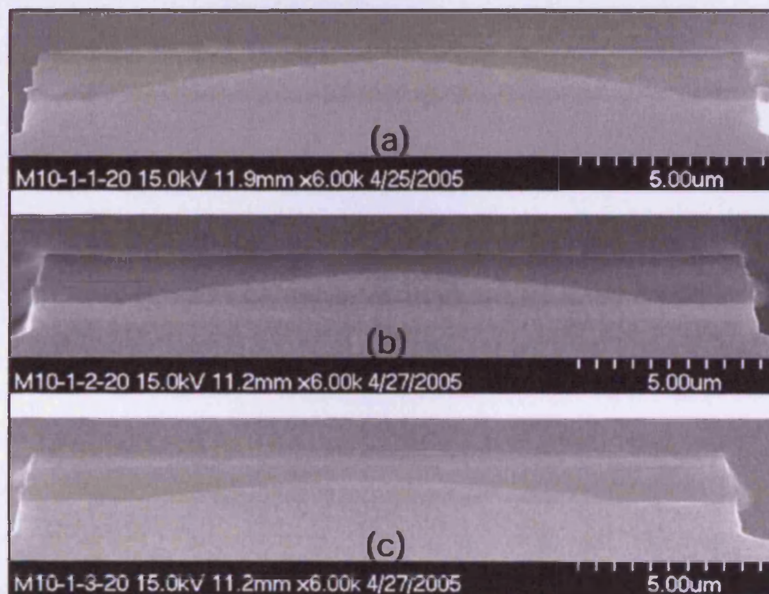

Figure 5.8: Microlenses made by oxidation for (a) 20, (b) 30, (c) 45 minutes of digital alloy stacks [5]

The creation of these lenses, however, has required the growth of 150 separate layers ranging in thickness from 90 down to two monolayers, with an accuracy of one monolayer. What has been discovered here, during the investigation into the performance of Wafer B, is a much simpler method of fabricating the same lenses, but involving only two layers.

The same proximity of the 80% aluminium p-type cladding layer to the first 95% oxidisable layer as has already been mentioned, creates, very effectively, a curved oxidation front in the thick cladding layer, as is displayed in Figure 5.9.

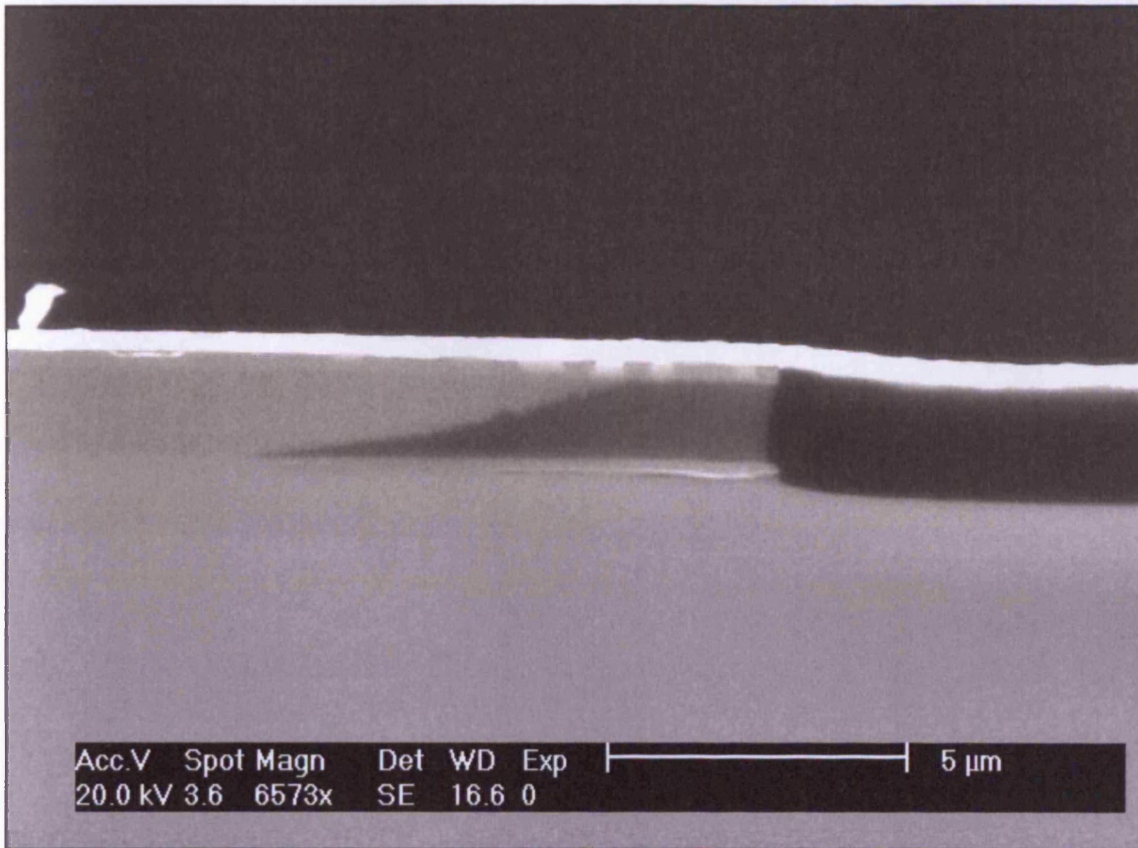

Figure 5.9: Detail of oxide front curvature

This curvature occurs because the oxidation rate of the 80% layer, which would be much lower in isolation, is enhanced by the neighbouring 95% layer. A simple qualitative simulation of the situation reveals more clearly the origin of the exact curvature shape observed. The two following figures show the results of this simulation. Figure 5.10 gives the shape of the oxidation front in the thick, 80% aluminium layer (where oxidation has proceeded inwards from right to left, as in Figure 5.9) where oxidant species and waste products may diffuse only through the thin, 95% aluminium 'feeder' layer. Figure 5.11 however, displays the same oxidation front, but assuming that once any of the 80% material has been oxidised, it becomes porous and oxidant species may diffuse through it directly from the sample surface (right to left) and aid the further transverse (bottom to top) oxidation of the 80% material. Note the large differences in the scales of the y axes, and the identical scales of the x axes. This demonstrates that the concave shape and the large transverse thickness of oxide achieved in the 80% material are only possible because of lateral diffusion of oxidant species through the 80% layer. The convex shape of Figure 5.10 is present because it is assumed in the simulation that the oxidation of the 95% layer is not linear, and tends towards some limit.

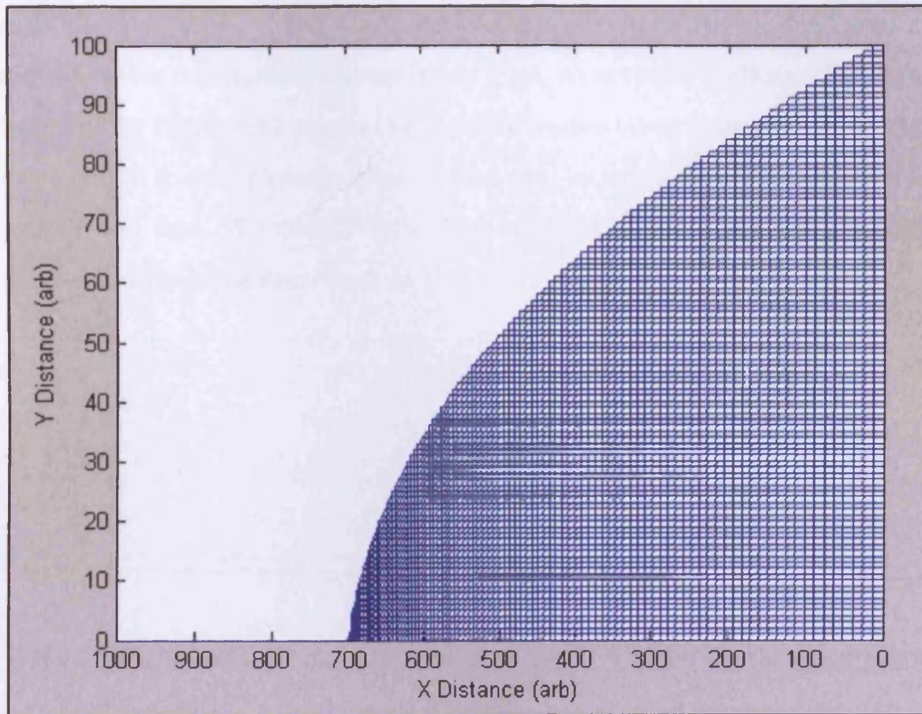

Figure 5.10: Oxidation front simulation results with no lateral diffusion

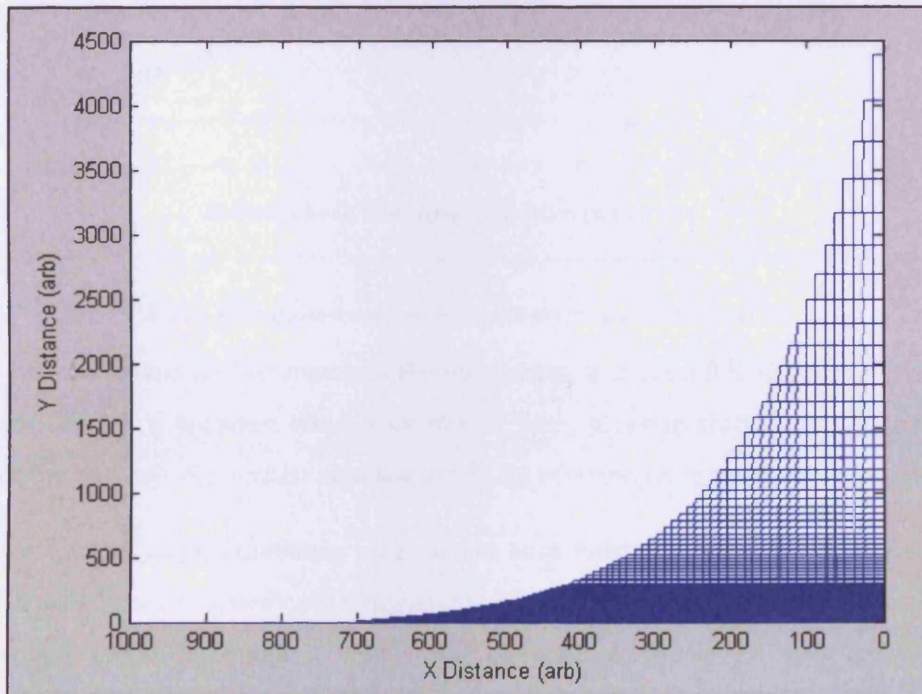

Figure 5.11: Oxidation front simulation results with lateral diffusion

As with Chang's paper of 2006, [5] the curvature of the oxide front closely matches a circular fit. By measuring and replotting points along the oxidation front, an accurate Radius of Curvature (RoC) could be found. The graph in Figure 5.12 shows the digitised values taken from the oxide front and two fit lines. The polynomial fit is an automatically generated one, included for the purposes of smoothing out the small errors in the data. The circular fit is then matched to the polynomial, thereby achieving a greater degree of accuracy in the final result.

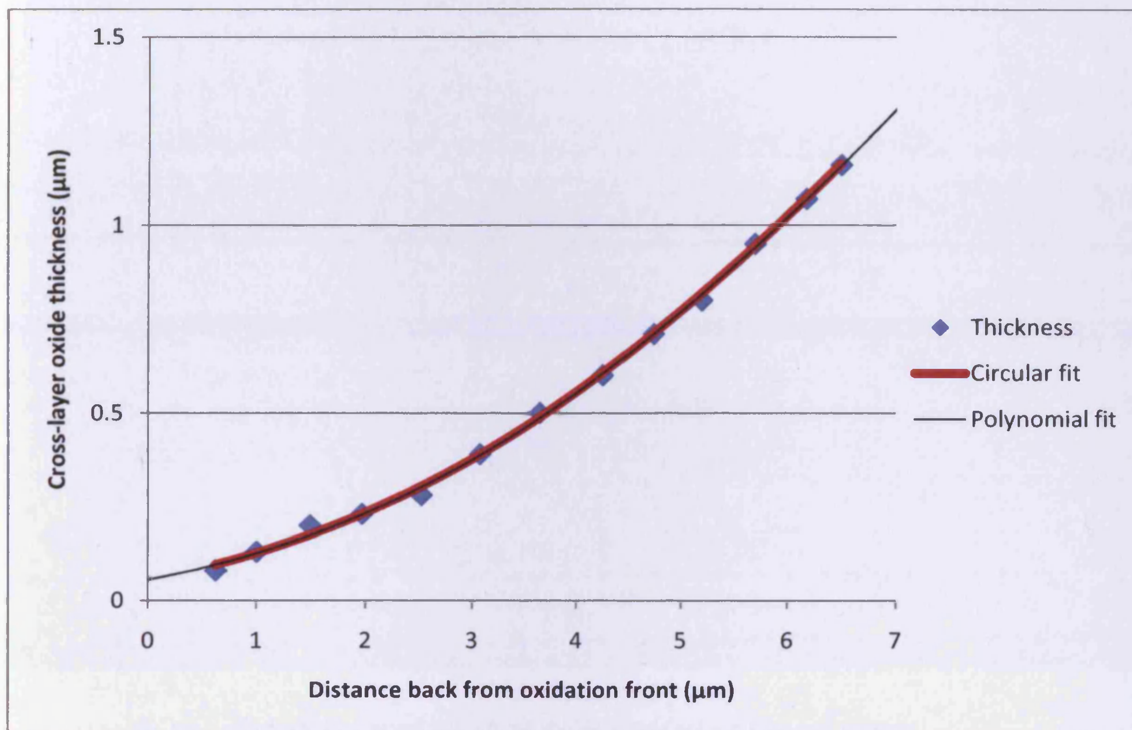

Figure 5.12: Oxide front curvature with polynomial and circular fits

The circular fit is an almost perfect match to the polynomial, and has a RoC of 29µm. The best result in Chang's paper reports a focussed beam spot size of 4µm, attained through a lens with an RoC of 33.4µm, implying that an even smaller spot size would be achieved using the layers discussed here.

A picture of a narrow ridge, containing what would be a functioning lens, is shown in Figure 5.13. Clearly, as the wafer was not specifically designed for this purpose, the lens is upside-down compared to those in Chang's paper (although it would also converge a beam this way around) and some considerable further work would be needed to determine the possible range of curvatures and lens sizes by varying the thickness and composition of the cladding layer, as well as the oxidation temperature. It should also be noted, as mentioned earlier, that this technique does currently lead to strain-related

delamination in approximately 5% of samples. As already mentioned, however, this can be reduced in a number of ways, and in the case of lens fabrication in narrow mesas, oxidising just deep enough for the feeder layer to oxidise completely through would help reduce the strain in the device considerably and would simultaneously produce the smallest lenses.

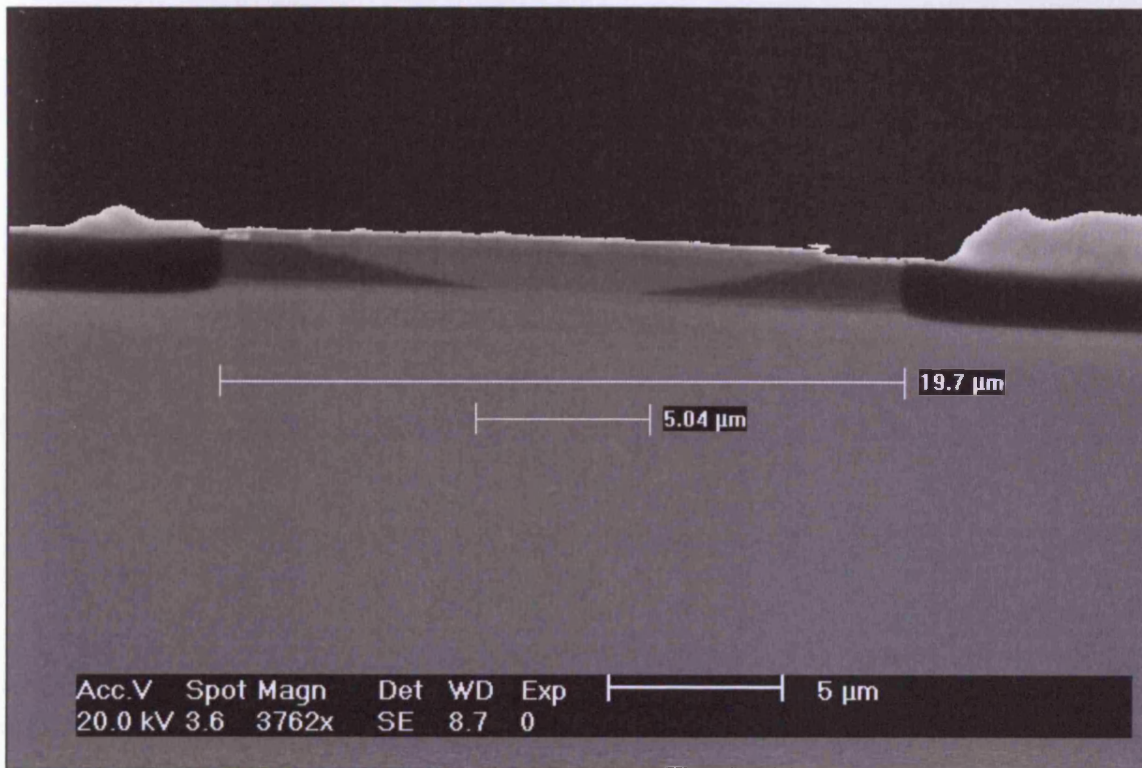

Figure 5.13: Curved oxide front shows potential for buried microlenses

5.6 Summary

This chapter has demonstrated some crucial points and has opened lines of investigation that might previously have been thought impossible. It has been shown that the incorporation of high bandgap oxidisable material into the active region can be achieved without appreciable detriment to the electrical performance of the wafer. Subsequently, such layers can be used to confine the device current directly in the active region, so negating the effects of current spreading entirely. As a direct consequence of this, broad-area devices have been fabricated with lower thresholds than their unoxidised equivalents, although the effect would be much more significant in narrow devices. By examination of the slopes of LI traces it has been determined qualitatively that no significant increases

in the value of α_i have been induced. Finally, a simple and novel method for the fabrication of monolithically integrated microlenses has been discovered. An RoC of 29μm has been recorded, which equates, theoretically, to a spot size smaller than 4μm; a value comparable to those reported in the literature from lenses fabricated by far more complex methods.

5.7 References

1. *Electrical conduction and dielectric breakdown in aluminium oxide insulators on silicon.* **Kolodzey, J.; Chowdhury, E. A.; Adam, T. N.; Qui, G.; Rau, I.; Olowolafe, J. O.; Suehle, J. S.; Chen, Y.** 1 : IEEE Transaction on Electron Devices, 2000, Vol. 47.

2. *Dielectric breakdown of ultrathin aluminium oxide films induced by scanning tunneling microscopy.* **Magtoto, N. P.; Niu, C.; Ekstrom, B. M.; Addepalli, S.; Kelber, J. A.** 14 : Applied Physics Letters, 2000, Vol. 77.

3. *Technique for monolithic fabrication of microlens arrays.* **Popovic, Z. D.; Sprague, R. A.; Neville Connell, G. A.** 7 : Journal of Applied Optics, 1988, Vol. 27.

4. *Highly uniform vertical-cavity surface emitting lasers integrated with microlens array.* **Eitel. S.; Fancey, S. J.; Gauggel, H. -P.; Gulden, K. -H.; Bächtold, W.; Taghizadeh, M. R.** 5 : IEEE Photonics Technology Letters, 2000, Vol. 12.

5. *Microlens fabrication by selective oxidation of composition-graded digital alloy AlGaAs.* **Chang, K. S.; Song, Y. M.; Lee, Y. T.** 1 : IEEE Photonics Technology Letters, 2006, Vol. 18.

Chapter 6 : Effects from Facet Oxidation of Cladding Region Layers

6.1 Introduction

One of the original motivations of this work was to investigate cheaper and easier methods of fabricating saturable absorbers in semiconductor diode lasers with a view to producing effects such as self-pulsation, as has previously been achieved with absorbers made using other methods. [1] Although conventional self-pulsation as such has not been achieved, several novel and beneficial effects have been observed, all of which can be related to saturable absorber-type effects and, in some cases, to self-pulsation. This chapter discusses exclusively Wafer A and deals with the two most significant groups of results from that wafer. In each case the result is described and quantified, a hypothesis for the physical mechanism behind it is proposed, and evidence backing that mechanism up is provided. The oxide geometry of the devices made and discussed in this chapter is very similar to that shown in Figure 4.1. The only difference being that it is the 'device version' of Wafer A being used here, so the oxidisable material is only to be found in the p-cladding. As demonstrated in Figure 4.1, the lateral oxidation will not affect device performance here and saturable absorbers will be formed at both ends of the lasers.

6.2 The first result – Threshold dependence on cavity length

Using the IVLT apparatus described in the experimental chapter, the threshold currents of a number of oxidised and unoxidised devices were measured. For this experiment, all the oxidised devices were oxidised for a period of 15 minutes at 480°C, which, by extrapolation of the results from the previous chapter, and from other literature, [2] is known to equate to an oxidation depth of approximately 150μm. This depth was not tested in the previous chapter as the mesa structures could not easily be made sufficiently wide due to the lack of the correct photolithography masks. It was chosen for this experiment, however, as absorbing sections fabricated by different methods are known to be required to be approximately this thickness to produce pulsation.

Figure 6.1 shows typical LI data for devices of a range of cavity lengths taken from a batch of around 20 lasers and compares the threshold currents at 300K of oxidised and non-oxidised devices of the same

length. The results show that oxidised devices with cavity lengths greater than 1.5mm have lower threshold currents than the equivalent non-oxidised device, whereas 1.5mm devices show a significant increase in threshold current.

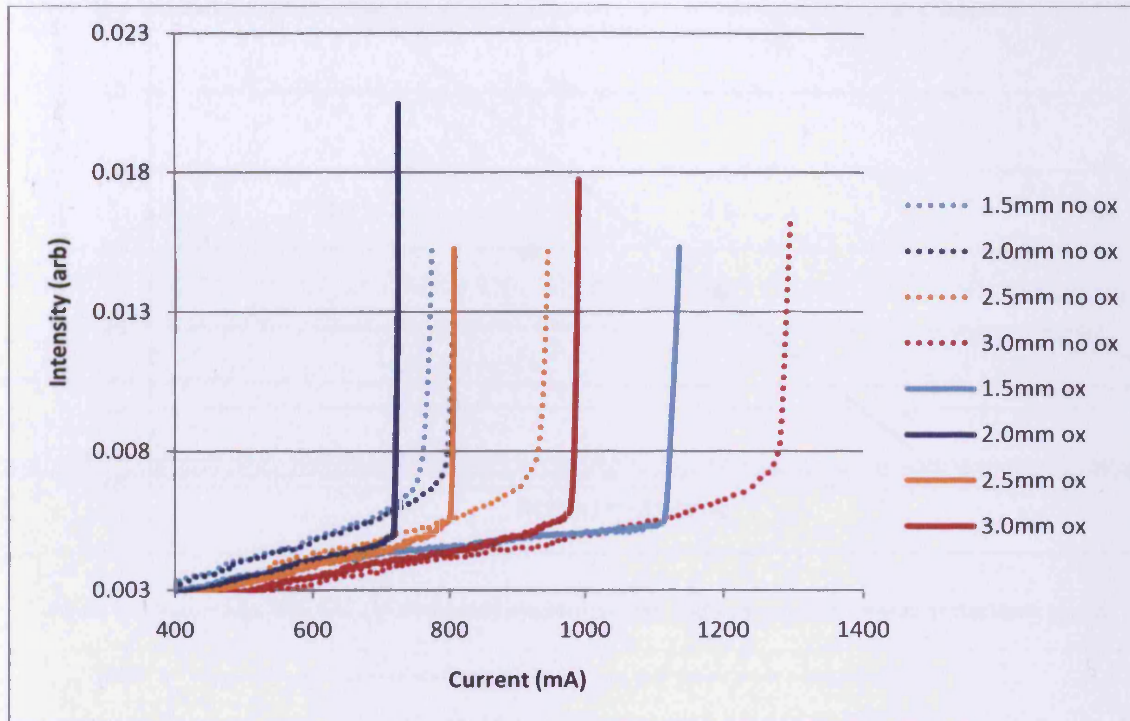

Figure 6.1: Comparison of LI traces for oxidised and standard devices of various lengths

Furthermore, the amount by which threshold is either increased or decreased appears to be proportional to the cavity length. This is quantified in Figure 6.2, which shows an approximately linear relationship where the percentage difference in threshold current ranges from approximately -30% to +30% with cavity lengths of 3mm and 1.5mm respectively. The data in this figure comes from only four individual devices and so the errors on these specific points are small. However, if more devices were fabricated in the same way, the statistical spread of threshold currents may lead to errors of around ±10 on the y axis. Errors in the x axis are very small. Figure 6.3 goes on to show that this change in threshold current is present over a wide range of temperatures, and even increases, as a percentage, at lower temperatures. The device measured for this figure is not included in the preceding data as it is from a different processing batch, and so is not quantitatively comparable. It does, however, demonstrate the stability of the effect with temperature.

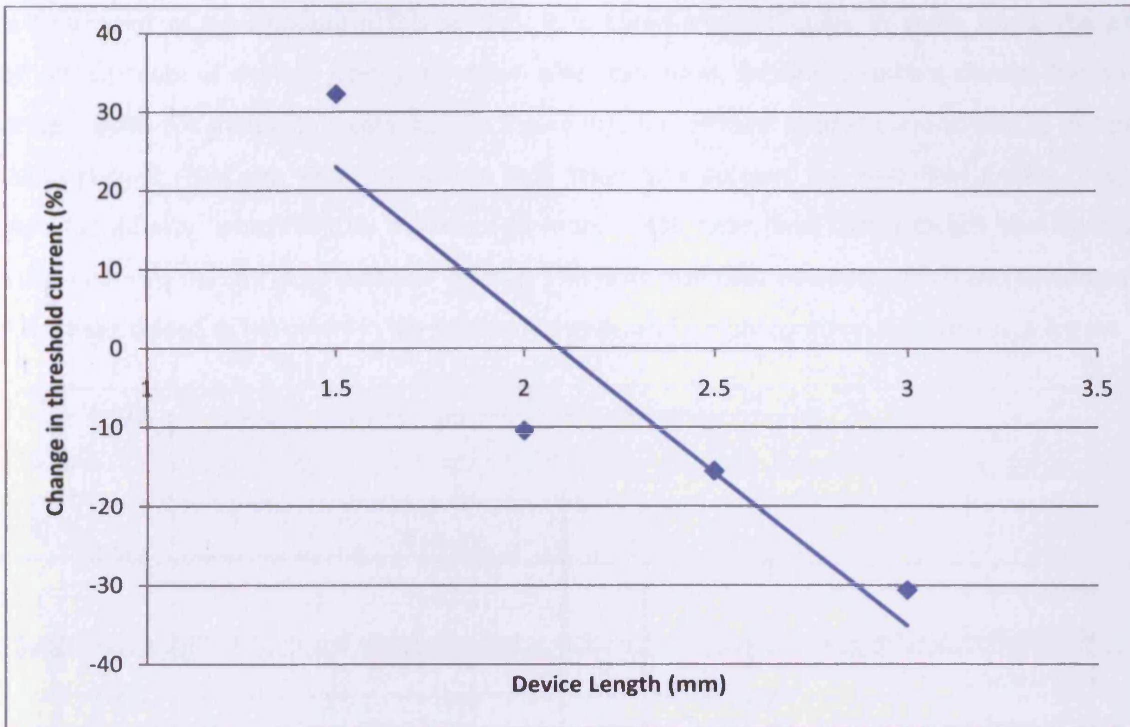

Figure 6.2: Percentage difference in threshold between oxidised and non-oxidised devices of the same length

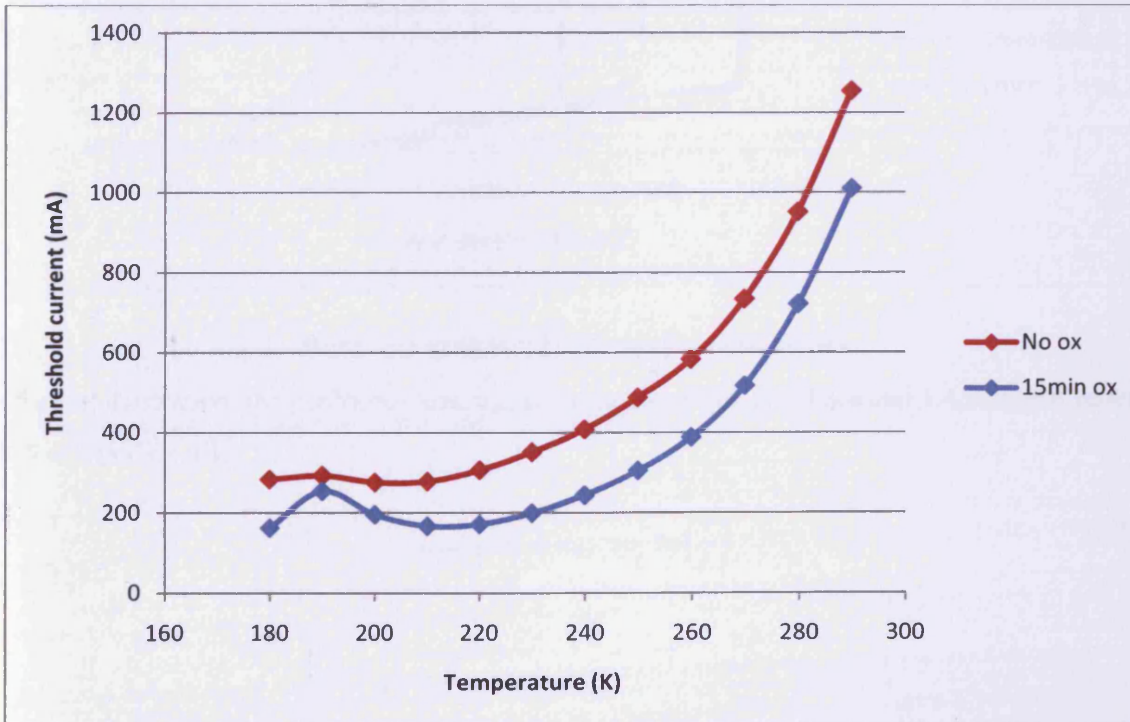

Figure 6.3: Temperature dependence of threshold of 4mm standard and oxidised devices

As a final point of information in this section, it is noted that although, in some cases, the actual threshold currents of devices have gone down after oxidation, threshold current density has always gone up. Figure 6.4 shows the same data as Figure 6.1, but plotted against current density instead of absolute current. This new depiction of the data takes into account the near-field widths of all the devices individually, which will be discussed in more detail later, and shows clearly the increase in threshold current density in all oxidised devices. This is as expected however, given the additional loss that has been added to the cavity in the oxidised devices, and simply confirms the process going on.

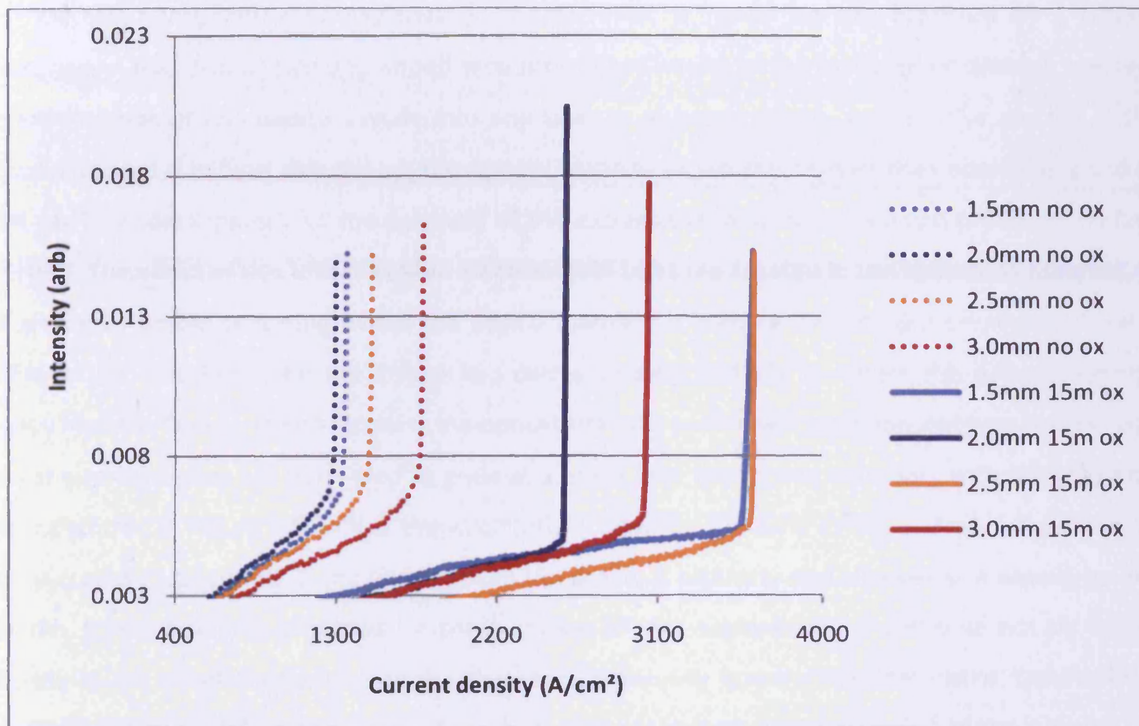

Figure 6.4: Comparison of LI traces using current density

For further illustration, the additional loss, α_{ox}, can now be added into Equation 1.4, which is rewritten below as Equation 6.1.

$$G_{th} = \alpha_i + \alpha_{ox} + \frac{1}{L}\ln\left(\frac{1}{R}\right) \tag{6.1}$$

6.3 The proposed mechanism

A hypothesis explaining this behaviour has been developed over the course of this project based on a combination of background knowledge and new experimental data. It should be noted whilst considering this hypothesis that the change in threshold is not numerically proportional to the change in the pumped length of the device after oxidation; nor is it due simply to the added optical loss, as both increases and decreases in threshold were observed.

The relatively straightforward explanation of this result is based on the presence of a saturable absorption in the form of two unpumped sections of fixed length at the facets of all devices, created by the introduction of an insulating oxide into one layer to a certain depth. As the drive current in these oxidised devices is initially ramped up, the optical intensity within the cavities rises accordingly and with what can be approximated, for the purposes of this explanation only, to a Gaussian profile in the lateral direction. The effect of this is that excited electron-hole pairs are created in the absorbing material, with the greatest number occurring where the optical intensity is highest. This in turn creates a lateral loss profile in the absorbers with the lowest loss corresponding spatially to where the optical intensity is already highest. Clearly, from this point the optical intensity profile will gradually become sharper as the highest intensity areas are permitted to grow at a faster rate than those that start with lower intensity. This is depicted in Figure 6.5, which shows pictorially how the intensity profile narrows over time, with each successive loss profile being derived from the previous intensity and vice-versa. It should be noted that this figure is heavily idealised for the purposes of this explanation, and should not be taken as accurate in any other context. When the absorbers eventually become fully saturated, they will do so only in a relatively small central area, all of which will lead directly to a narrowing of the external near-field.

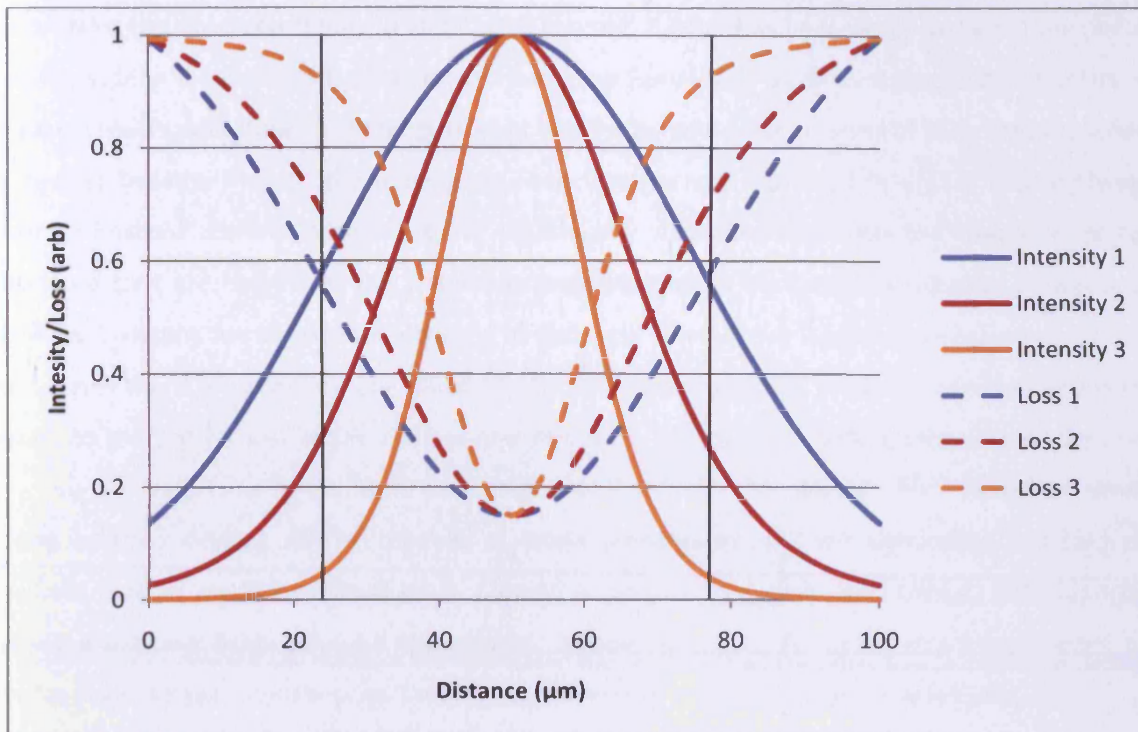

Figure 6.5: Normalised sub-threshold intra-cavity intensity and absorber loss profiles. Vertical black lines indicate approximate laser stripe width. Current spreading and off-axis emission of light spread the profiles wider than the nominal strip width.

On its own this would not lower the threshold as all that has changed from a standard device is that an extra loss has been added, so the threshold would, in fact, go up. However, when considering the oxidised devices at a point somewhere below threshold, it is clear that the added loss, which is localised at the facet regions, cannot significantly affect the levels of spontaneous emission but will reduce the build up of amplified spontaneous emission (ASE). Given also that the loss remains greatest at the sides of the devices, it follows that the ASE will be suppressed most in these areas, and least in the centre. These conditions will lead to a lengthening of the carrier lifetime and, therefore, carrier density in the side regions, relative to the centre. Consequently, there will be a diffusion gradient set up and carriers from both side regions will diffuse to the centre where they will add to the building ASE levels and, therefore, ultimately lower the threshold of the devices.

As a brief aside, this is good evidence that a significant portion of the emission well below threshold is ASE. Prior to this work, it has been debated as to whether the emission well below threshold was entirely spontaneous or if a portion of ASE was present but there was little experimental evidence either way. [3]

This explains the observed results in all but one respect. If what has been described were the picture in its entirety, then the thresholds of all devices would be lowered by a constant proportion, but this is not the case. Devices with large cavity lengths were seen to benefit most in terms of threshold reduction. As the cavities became shorter, the percentage reduction became less until finally, at a cavity length of 1.5mm, threshold actually increased quite significantly. This effect also fits the theory if the fact is considered that the lengths of the absorbing sections, and so the total extra losses introduced, are nominally constant for all oxidised devices. In the case then of the longer devices, proportionally to those losses there is more time per round trip for the suppressed ASE method described above to add carriers to the central axis of the devices and therefore the absolute optical intensity incident on the absorbers will be greater at the same current density than in shorter devices. For this reason, threshold in long oxidised devices will be reached at lower percentages of their equivalent standard device threshold than in shorter oxidised ones. Clearly, in devices of 1.5mm and below, the advantage of narrowing the near-field, although still present, is outweighed by the total extra losses added as the benefits available per round trip are simply not sufficient to maintain the same threshold.

6.4 Experimental evidence

The aim of this section is to provide experimental data to back up the various elements of the hypothesised mechanism. The presence of a saturable absorber and that of repeatable near-field narrowing are the two main premises that are proven.

6.4.1 The saturable absorber

A key premise of the whole hypothesis and crucial to the explanation of both of the two main results of this chapter is the presence of a saturable absorption within the oxidised devices that is not present in the control devices. The evidence supporting this is very strong.

6.4.1.1 Turn-on delay

The first evidence that was observed supporting the presence of saturable absorbers was the marked turn-on delay that was seen at currents just above threshold. Although this was first noted on the IVLT apparatus, the effect was investigated more thoroughly on the streak camera, as wavelength information could then be gathered as well. A number of oxidised and standard devices were compared, yielding results such as those shown in Figure 6.6.

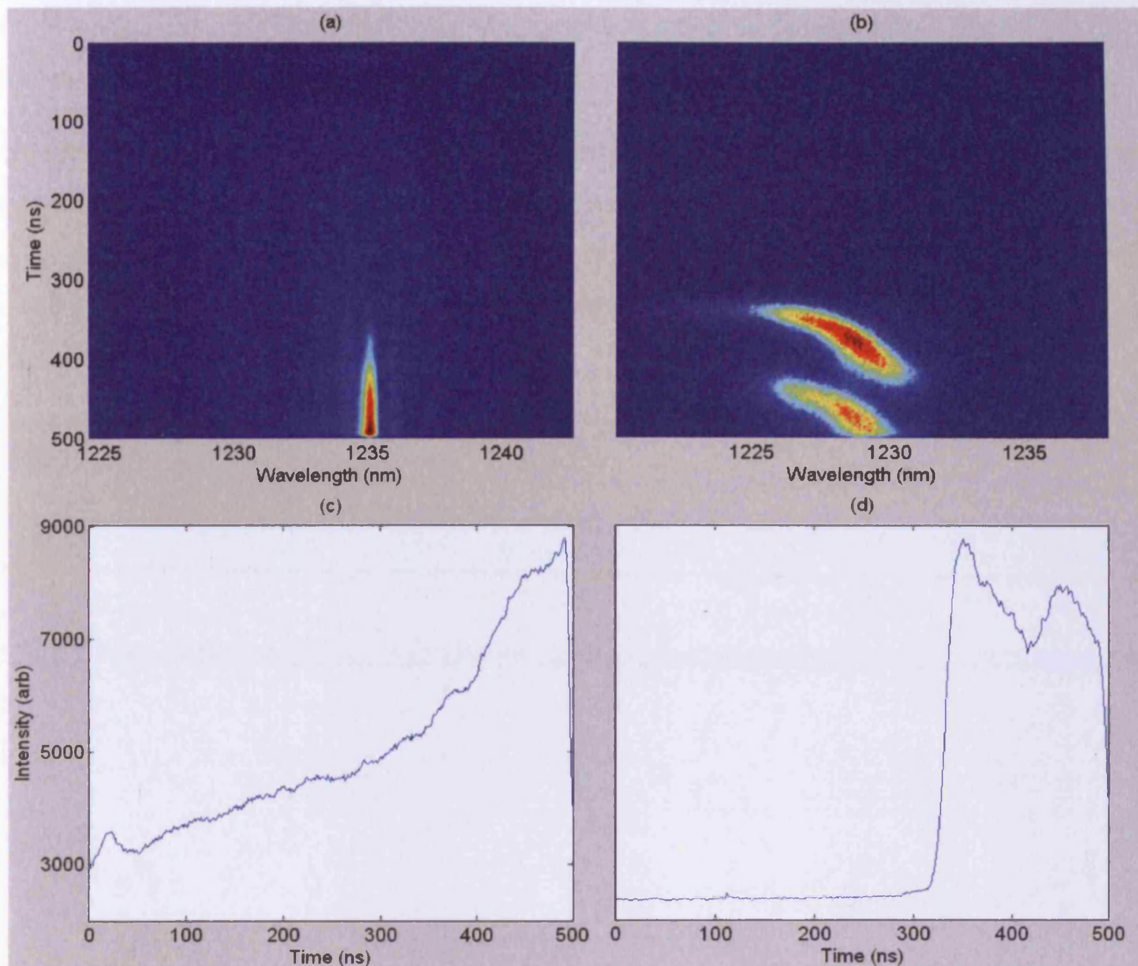

Figure 6.6: Turn-on delay data – (a) and (b) respectively show standard and oxidised device time-resolved spectra, (c) and (d) respectively show the corresponding time-intensity profiles integrated over all wavelengths

In both the above cases, the electrical drive pulse begins at t=0 and lasts for 500ns. The bottom two traces are the same data as shown above but integrated over all wavelengths. These are included for clarity, as the contrast on images (a) and (b) makes the differences in intensity before 400ns unclear. What is observed is that the standard devices actually also have a measurable lasing turn-on delay comparable to that of the oxidised devices. The presence of the saturable absorption has simply exaggerated the behaviour of the standard devices and made the delay more obvious in the oxidised devices. As can be seen from plots (c) and (d) of Figure 6.6, the saturable absorbers prevent almost any light at all escaping from the cavity until sufficient intensity has built up inside to bleach the absorbers and release the light in a very sudden burst. By comparison, what can be seen in the case of the standard device is a slow build up of amplified spontaneous emission (ASE) from t=0 to approximately 350ns when the build up notably accelerates and the device moves into full lasing action. The double

peaks in images (b) and (d) are related to more complex that will be discussed later, but are not relevant to the discussion here.

As an aside, the behaviour noted here from the standard devices is not well documented in literature. Plenty of work has been reported detailing turn-on delays relating to relaxation oscillations, which are in the region of 5ns and occur at all drive currents to some extent, but nothing of this order of magnitude. Figure 6.7 shows the delay in one such standard device progressing from around 400ns down to about 100ns as the drive current is increased.

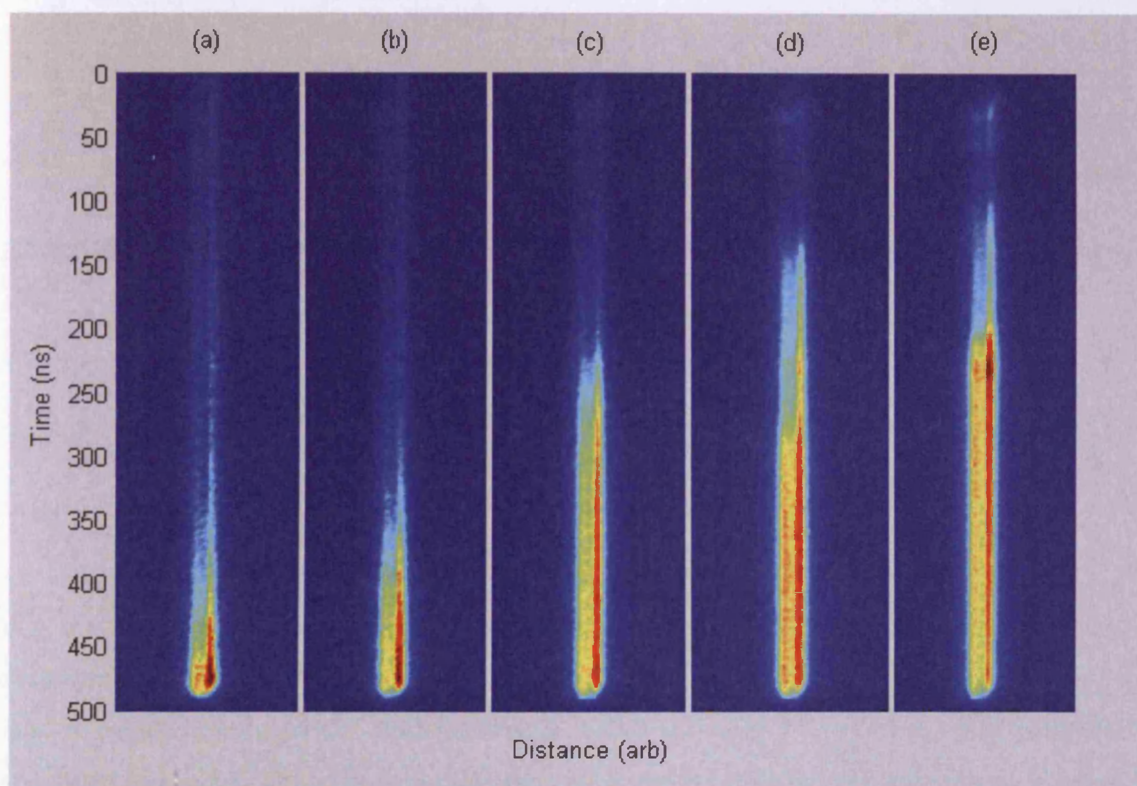

Figure 6.7: Progression of turn-on delay in a standard device shown at (a) 180mA (b) 182mA (c) 184mA (d) 192mA (e) 204mA

One possible explanation for this could be a weaker but similar effect to the ASE suppression described in the hypothesis section above. As the intensity builds up in the centre of the device carriers will be depleted at a faster rate than at the side regions and so carriers will diffuse from the sides to the centre, aiding the turn-on of the device. This effect would be most evident at currents only marginally above threshold, by being sufficiently slow as to be observed over the length of a single drive pulse, as it the case here. This was not one of the aims of this work however, and so the cause of this effect has not been thoroughly investigated; it is noted though that the effect does occur in standard devices, and that

oxidised devices merely exhibit an exaggerated version of exactly the same behaviour. This is confirmed by plotting the progression of the delay with increasing drive current of a standard and an oxidised device on the same axes, as is shown in Figure 6.8.

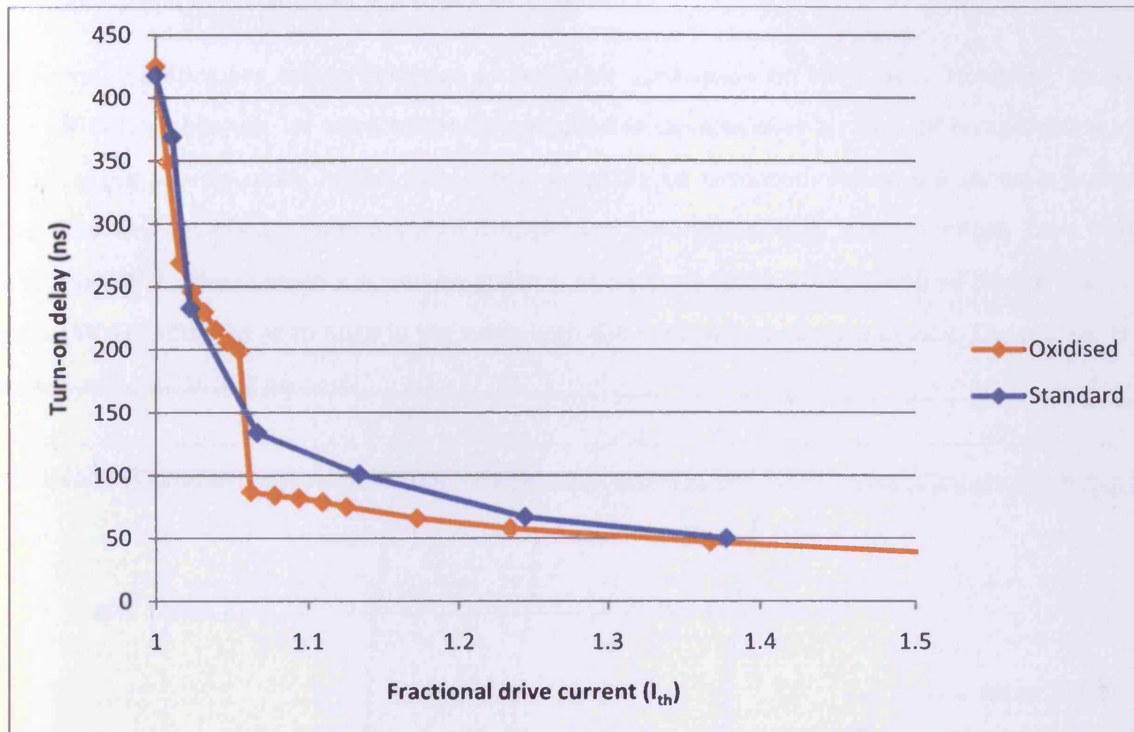

Figure 6.8: Comparison of turn-on delay progression with increasing drive current

Clearly, the two sets of data are in very good agreement when plotted against current as a fraction of each device's threshold current. The only discrepancy is the sudden jump observed with the oxidised device. This was common in the oxidised devices tested, although the exact size and position of the jumps varied somewhat. Although the explanation for these jumps is not yet understood, it would seem likely that it relates to further electronic levels in the absorbers becoming saturated and, the ASE being allowed to build up faster before turn-on. The evidence presented here, as a whole, is strongly supportive of the presence of saturable absorbers.

6.4.1.2 LI characteristics

Further supporting evidence came when comparing the LI characteristics of oxidised and normal devices. Oxidised ones showed a lower level of spontaneous emission below threshold, a much sharper turn-on at threshold, and a maximum LI slope above threshold approximately three times steeper than

those of the normal devices. These traits are all indicative of a saturable absorption, where again, the light is effectively trapped inside the cavity until sufficient intensity has built up to saturate the absorber. At this point, the losses within the cavity suddenly reduce, allowing the intensity to build further very rapidly, and giving rise to the characteristics described.

These characteristics are strong evidence of saturable absorption on their own. However, to acquire more definitive evidence, LIs were taken for comparable devices over a range of temperatures, up to optical output powers much higher than would normally be recorded. Figure 6.9 shows a number of traces from this set of data comparing an oxidised and a normal device, both of length 3mm between 260°C and 280°C. The current axis on this graph is in arbitrary units as the oxidised device traces have been artificial shifted so as to align in the x axis with the data for the control device. The gradients have not been affected in this process.

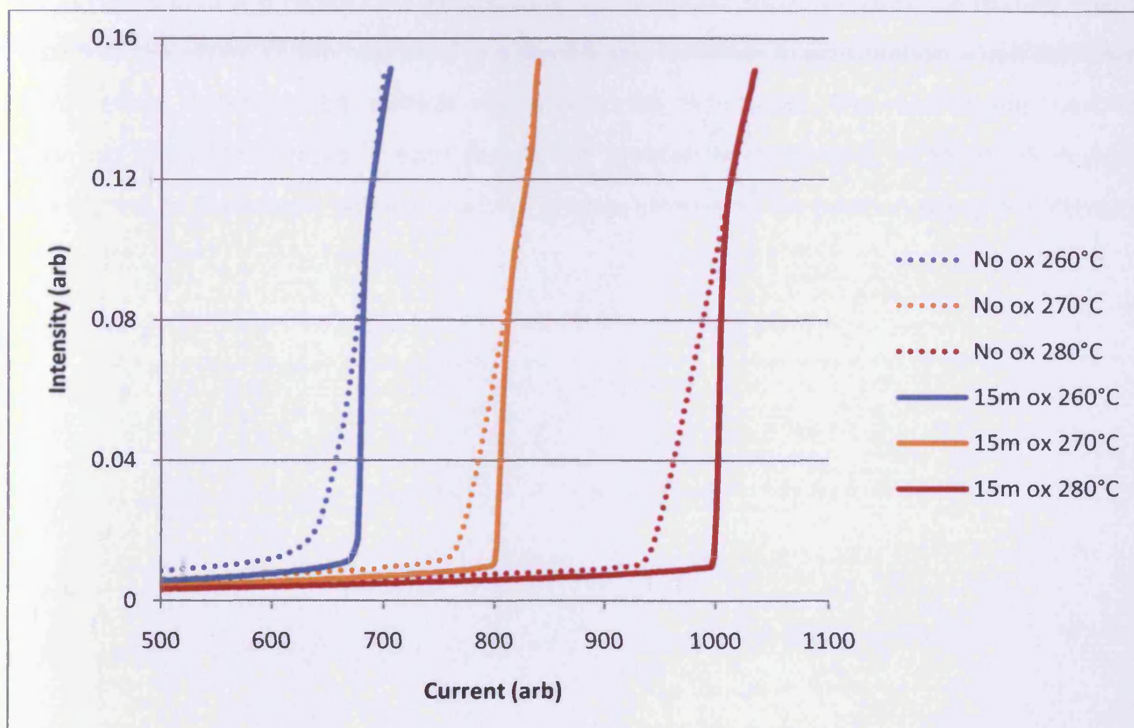

Figure 6.9: Changing LI slopes evidencing saturable absorption in 3mm device

As well as illustrating all the characteristics already described, the LI traces belonging to the oxidised device show the classic signs of saturable absorber behaviour. Namely, a slightly delayed turn-on (assuming the same theoretical threshold, which these traces have been shifted to simulate) and a large gradient immediately after turn-on that returns to a value matching the normal device after the point

where the effects of the absorber have been completely surpassed. The fact that the LI gradients of the oxidised device traces return to match those of normal devices in the upper part of this graph is the final piece of evidence needed to prove conclusively the presence of saturable absorption in the oxidised devices.

6.4.2 Near-field narrowing and mechanism

In this section, the narrowing of the near-fields cited earlier as the cause of the lower threshold currents is quantified and linked to the hypothesised mechanism. The results shown here are taken from devices from the same batch as those whose thresholds were documented in Figure 6.1, and wherever possible are the exact same devices.

Comparisons of the near-fields of 2.5mm and 3mm devices at relatively high drive currents of $1.5I_{th}$ are shown in Figure 6.10 and Figure 6.11 respectively. These results show not only the narrow near-field, but also that the shape of the near-field is a good single Gaussian approximation when compared to standard devices, which exhibit multiple and uneven near-field peaks. The relative alignment in the distance axis of the two traces in each case is not necessarily meaningful, although all devices are roughly aligned, as filamentary behaviour adds a random element to the position of highest intensity.

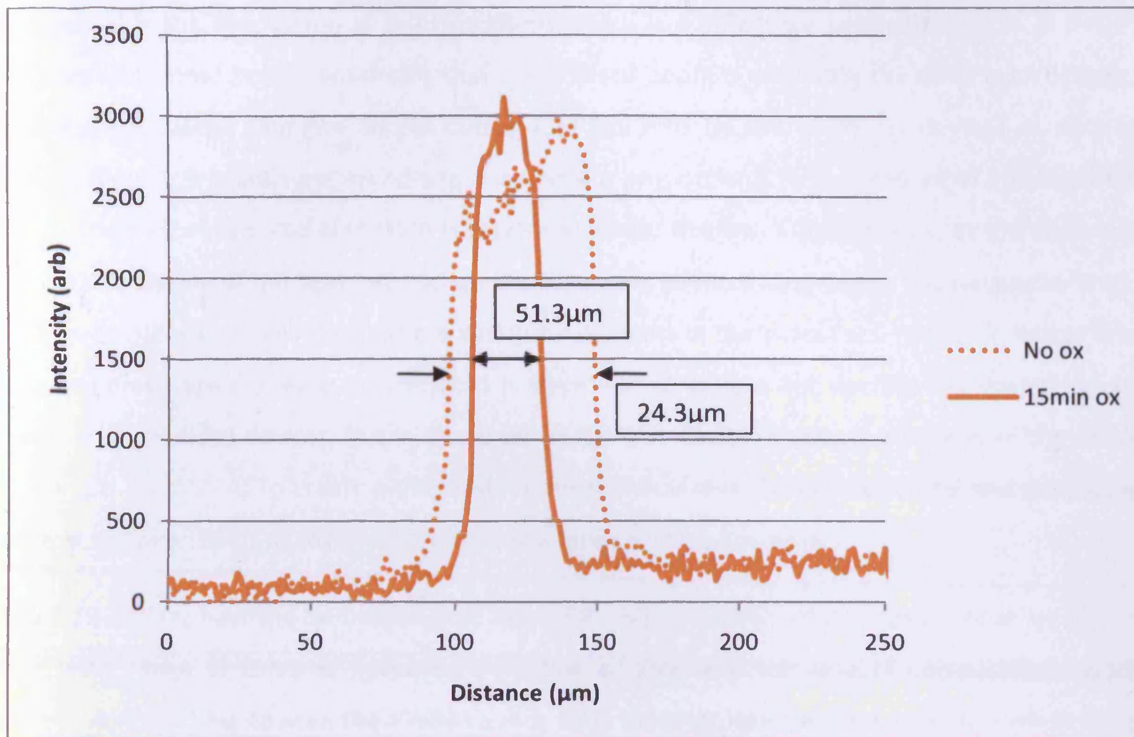

Figure 6.10: Near-field comparison of oxidised and standard devices with 2.5mm cavities at $1.5I_{th}$ drive currents

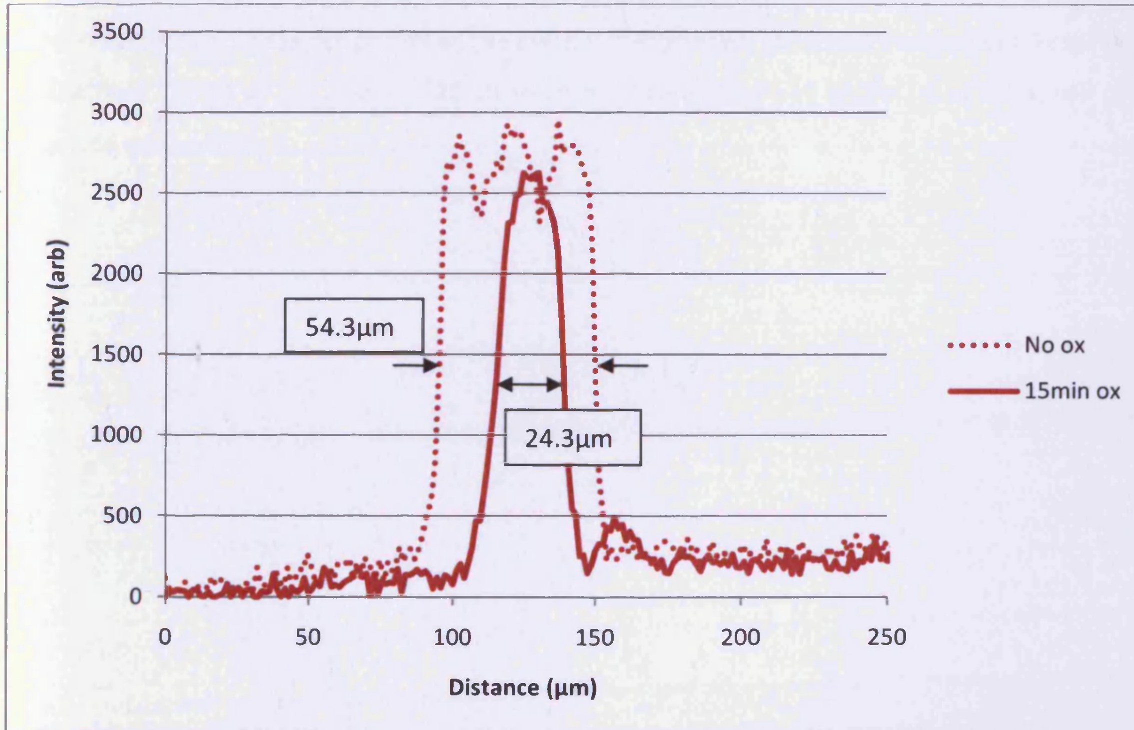

Figure 6.11: Near-field comparison of oxidised and standard devices with 3mm cavities at $1.5I_{th}$ drive currents

As suggested in the description of the hypothesis, there is a difference between devices of lengths at and above and below 2mm. Considering that the oxidised depth is nominally the same in all devices, it is reasonable to assume that the added optical loss will also be the same for devices of all lengths. However, the gain available per round-trip in a device is proportional to the amount of pumped material the light travels through, and therefore is greater in longer devices. Consequently, at the same current density, the intensity of the light incident on the absorbers within a long device will be greater than that in a short device and so, will create more electron-hole pairs in the absorbers. Hence, in longer devices, a greater percentage decrease in threshold is observed. Overall, a net decrease in round-trip gain is induced in all oxidised devices, but in those below 2mm in cavity length, the benefit of the near-field narrowing is insufficient to create a decrease in threshold current. This is confirmed and clarified when looking at the near-fields of several devices over a range of drive currents.

Figure 6.12 depicts how the 'full width, half max' (FWHM) parameter of devices of three lengths varies over a wide range of currents (given as a fraction of threshold for ease of comparison). In all the standard, non-oxidised devices the FWHM value stays more or less constant over the whole range of currents at values around 50μm, which reflects the width of the oxide stripe that forms the devices. At low currents, some values around 40μm are recorded, but this is not comparable to the narrowing seen in oxidised devices and does not persist as the current is increased. Long oxidised devices, however, are seen to achieve FWHM values as low as 15μm and remain consistently at widths of, at most, half of their non-oxidised equivalents.

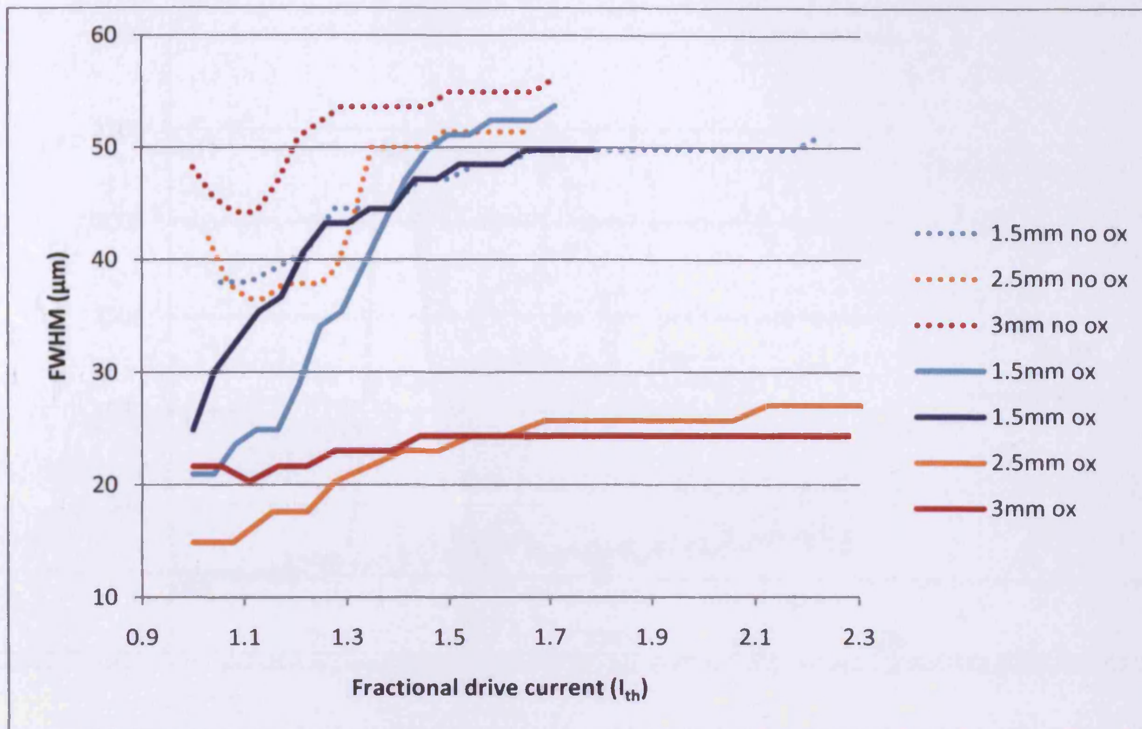

Figure 6.12: Comparison of near-field widths (FWHM) over a large range of drive currents

The most illuminating cases are those of the 1.5mm devices, of which two sets of data are presented to demonstrate repeatability. Despite their threshold currents having been raised by the oxidation process, they still show narrowed near-fields at currents up to approximately $1.5I_{th}$, at which point they widen again to match the FWHM values of their non-oxidised equivalents. Figure 6.13 and Figure 6.14 show this change in detail and demonstrate that near-field remains a good shape in both cases.

In all of the oxidised devices that exhibit narrowed near-fields, the remaining partially saturated areas of the absorbers slowly approach full saturation as the current density is increased. Due to the raised threshold current and short length of these devices, the threshold current density is very high indeed by the time the laser first switches on. This exact combination of device length and oxidation depth allows this rebroadening to be observed. From this it is predicted that the near-fields of even the very long devices would probably rebroaden if the current density was increased to produce the same photon flux outside the saturated regions, although this was not achievable in practice due to equipment limitations and potential damage to the devices.

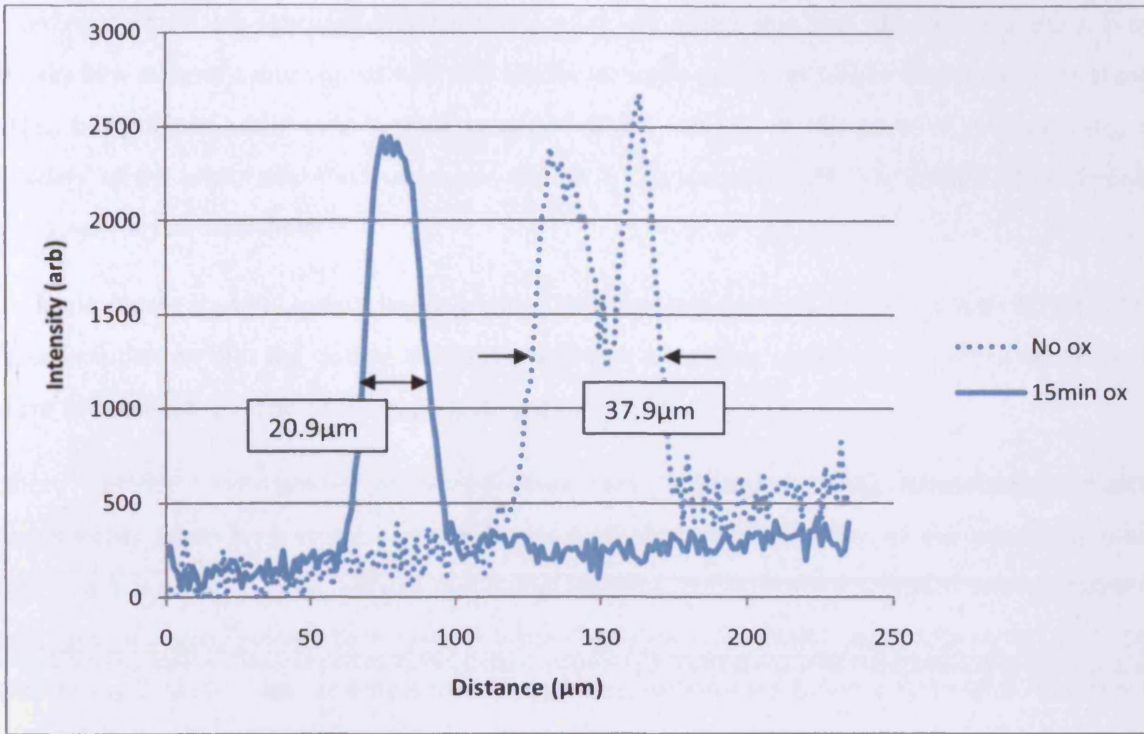

Figure 6.13: Near-field comparison of oxidised and standard 1.5mm devices at 1.05I_{th} drive currents

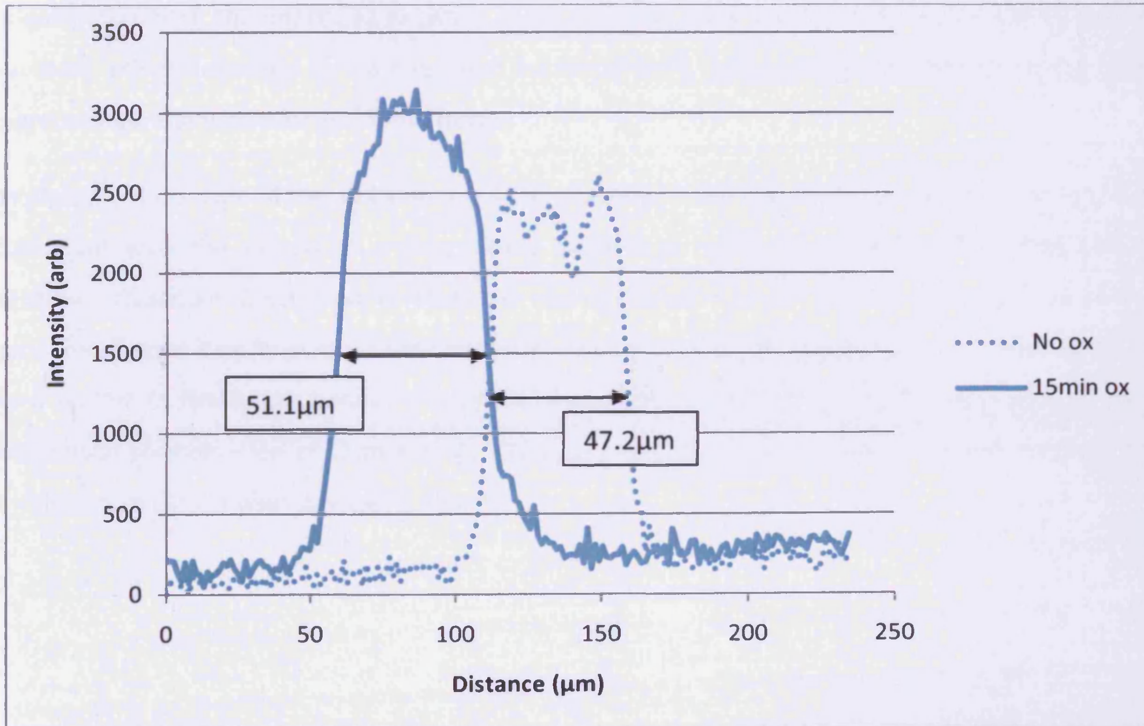

Figure 6.14: Near-field comparison of oxidised and standard 1.5mm devices at 1.5I_{th} drive currents

As confirmation of this rebroadening mechanism, it was noted that the rebroadening did not occur gradually in a manner synonymous with the steady increase in output power that is found in standard devices, but fairly abruptly over a small range of drive currents. At the point of rebroadening, both 'shoulders' of the wider near-field were seen to rise simultaneously until they matched the intensity of the original, narrow near-field.

Once again, this is a useful insight into a general and debated principle of laser diodes. It has not been conclusively proven that the optical intensity outside of the lasing near-field continues to rise as drive current is increased, but this data suggests strongly that it does.

A more detailed investigation of this process was undertaken using time-resolved near-field measurements taken on a streak camera. Figure 6.15 shows the variation of the near-field over the length of a 1μs pulse (a large, varying number of identical pulses were integrated and averaged in all streak camera measurements to achieve a better signal-to-noise ratio) taken exactly on the point of rebroadening. It shows that not only is the rebroadening very abrupt, but it is also slightly unstable for a short period of time. The labelled sections on the picture mark the three sharp changes from narrow (A) to broad (B), back to narrow (C) and finally settling with a broad near-field (D). As the current is gradually increased, the interfaces between these sections occur earlier in the pulse, giving rise to the 'shoulders' effect described in the near-field measurements, which are effectively the same data but integrated over the whole length of the pulse.

Although the time-scale of this behaviour is long compared with traditional Q-switching action, it is not inconsistent with the effects of a very slowly recovering saturable absorber, [1] in that the initial transition into section B suddenly releases the stored photons as the saturable loss goes to zero. As a result, the photon density in the cavity becomes depleted, and the loss temporarily rises again, giving rise to section C. Finally the behaviour stabilises in section D. This effect is very much akin to the well documented phenomenon of Q-switching, which can be purposefully induced by the employment of saturable absorbers in laser devices of all kinds.

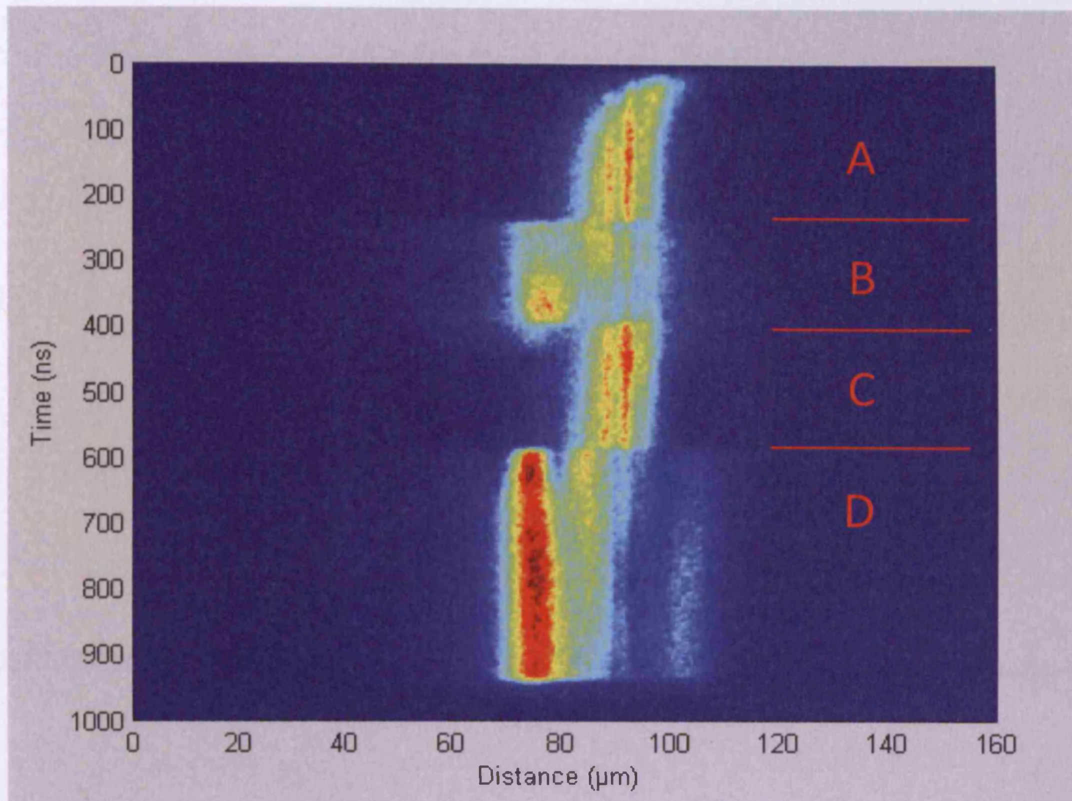

Figure 6.15: Time-resolved near-field of oxidised 1.5mm device taken on a streak camera showing the exact point of rebroadening

Overall, this section shows that the near-field narrowing is present in all oxidised devices discussed in this chapter, that it acts in accordance with the hypothesis in terms of controlling the threshold, and that it is controlled by saturable absorption, as evidenced by the rebroadening at high drive currents.

6.5 The second result – Repetitive wavelength shifting

The second novel effect created by oxidising laser devices in the longitudinal direction was a repetitive 'saw-tooth'-shaped wavelength shifting behaviour in the optical output over the duration of each electrical drive pulse. Figure 6.16 shows a typical example of this behaviour in a time- and wavelength-resolved streak camera image of the optical output of one device, integrated over several electrical pulses.

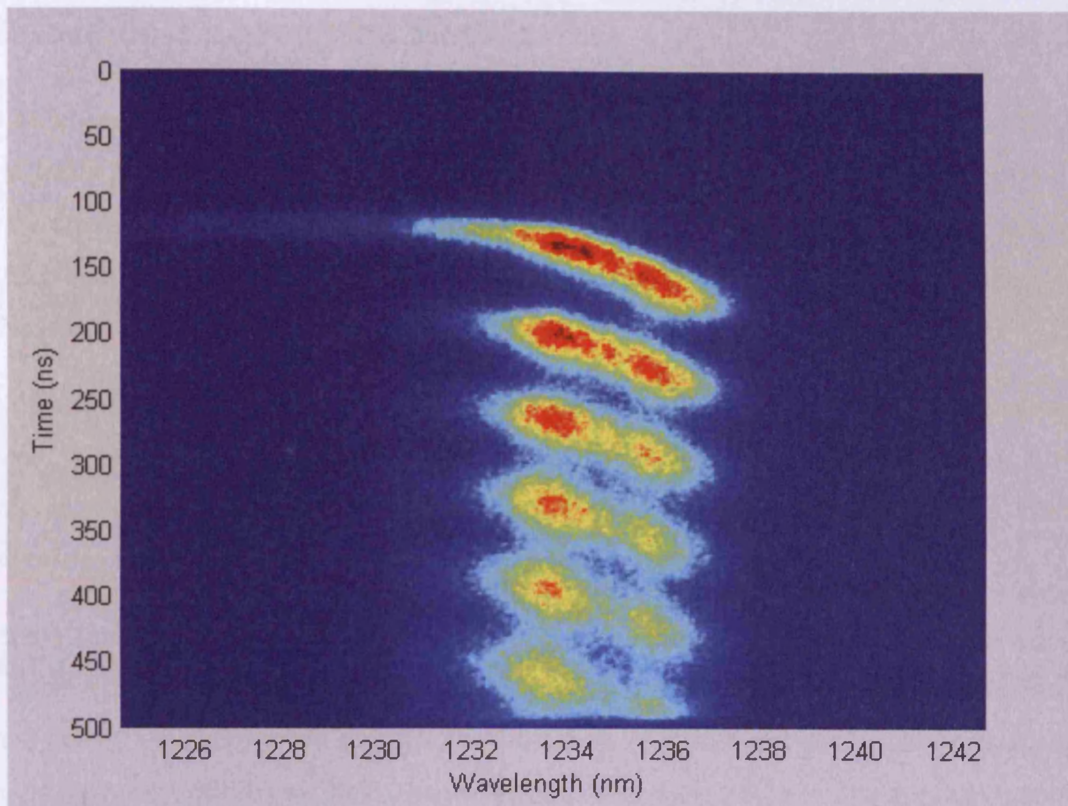

Figure 6.16: Typical time and wavelength-resolved behaviour showing RWS at a repetition rate of around 12MHz

The figure shows well defined shifts from short to long wavelength of approximately 4nm after a very brief initial slide in from much shorter wavelengths. At no point after its initial turn-on does the laser output switch off completely. Equally, there is no point in time when all wavelengths are being emitted simultaneously. For these reasons, the behaviour cannot be described as 'pulsation', although there are parallels to be drawn, as will be discussed later. Many oxidised devices were tested and the vast majority exhibited this behaviour between drive currents of 1 and around $1.5I_{th}$, and over all tested temperatures of between 10 and 60°C. Repetition rates for the behaviour started at around 10MHz at threshold (roughly as depicted in Figure 6.16) and rose with increasing drive current.

This behaviour is, to the best of the author's knowledge, novel both in its characteristics and in the method by which it was achieved. As such, the exact mechanism behind the effect is of significant interest, and it is this mechanism that is the focus of the remaining portion of this chapter.

The next two sections describe the mechanism by which the two results described here are proposed to be effected, and various experimental data supporting that hypothesis.

6.6 Extending the proposed mechanism

The physical explanation of the repetitive wavelength shifting (RWS) behaviour is simply an extension of the hypothesis given earlier. The mechanism is somewhat more involved, but remains centred around the same principle of selective bleaching of the absorbers. The main difference is that now not only must the spatial location of the saturation be considered, but also at what wavelength and what point in time the saturation is occurring.

Before detailing the mechanism behind RWS, some facts should be noted. Firstly, as has already been suggested, a narrowed near-field does not necessarily equate to a lower threshold. Conversely, however, all devices with narrowed near-fields were seen to exhibit RWS to some extent, even though in some cases heavy filamentation distorted the results.

Considering first that the saturable absorption occurs in the main within the quantum dot active region in the unpumped end sections of the devices, it is clear that a number of electronic levels exist and therefore several wavelengths are available for saturation. As with the spatially selective saturation that causes the narrowing of the near-field, the wavelength that is emitted with the greatest intensity from the gain sections of the devices will be that which is selectively saturated in the first instance, and will therefore be the initial lasing wavelength. From this situation, with no other change in conditions, a stable state would be set up, and the lasing wavelength would remain constant indefinitely. However, what this hypothesis proposes is that an extra variable of enhanced localised heating of the facets caused by a combination of the absorption itself and various subsequent non-radiative recombination mechanisms is present.

If it is assumed that both facets heat at a rate much greater than the bulk of the device, an explanation of RWS becomes possible. When heated, the band gap of any lasing material, whether pumped or not, becomes smaller, and consequently the lasing or absorbing wavelengths available become longer. The differing heating rates of the absorbing and gain sections lead to the range of available absorbing wavelengths increasing at a much faster rate than those available for lasing. Importantly, the absorbing states that are fully or partially saturated remain so whilst shifting upwards in wavelength as the occupying electron-hole pairs are heated along with them. As a result of this, for a period of time, it is preferable in gain terms for the lasing wavelength to 'track' the saturated absorber states, giving rise to the sloped part of the RWS behaviour.

The final new element of the hypothesis that explains the sudden and repetitive switch-backs seen in RWS is to do with mode spacing. By oxidising inwards from the facets, what is created is a region that is totally electrically unpumped and therefore has a lower carrier density than the pumped sections even when the absorbers are fully saturated. This difference in carrier densities is initially quite an abrupt junction located just short of the total oxidation depth due to the small amount of current spreading that occurs, and is significant enough to produce a notable discrepancy between the refractive indices of the two regions. What results from this is a device with three distinct cavities (although two of them are nominally identical) rather than just one. The lasing modes allowed to propagate in such a device would then be only those that satisfy all three cavities simultaneously. Effectively, it is now a superposition of the three individual mode 'combs' that determines the spread of permitted modes in the device as a whole.

This, coupled with the facet heating, provides an explanation for the rest of the RWS behaviour. As the absorbing sections heat, the saturated levels within them increase in wavelength, and the lasing emission tracks those levels until such a point that it becomes preferable to switch to the next, now more widely spaced, allowable mode, which has itself now shifted up in wavelength to approximately the original lasing wavelength. This process then cycles and repeats, producing the full RWS behaviour as shown in Figure 6.16.

6.7 Further experimental evidence

Once again, this section will provide the practical evidence to support the extensions to the hypothesis as detailed above. The two features discussed here are increased localised heating of the facets and the link between RWS, the absorbing sections and near-field narrowing.

6.7.1 Facet heating

The final remaining element of the hypothesis to be confirmed with experimental evidence is the presence of greater facet heating than in standard devices. This section provides that evidence in the form of results relating to catastrophic optical damage (COD) and a more detailed analysis of the RWS behaviour.

6.7.1.1 Catastrophic optical damage

One of the secondary motivations behind this whole project was investigate whether or not buried oxidisable layers could be used to suppress or delay the onset of catastrophic optical damage in diode lasers. The central idea behind this was that by oxidising the buried layers inwards from the facet at both ends of the devices, carriers would be removed from small sections of material immediately behind each facet, thereby reducing the effect of non-radiative surface and defect recombination, and leading to a reduction in facet heating, which is known to contribute to the early onset of COD. [4] [5]

To test this theory, a batch of 24 nominally identical 2mm long devices were fabricated from the same wafer as discussed in the rest of this chapter. Half of the devices were mounted as cleaved and half were oxidised for 15 minutes at a temperature of 400°C and then mounted in the same manner. The power-current (PI) characteristics of all 24 devices were taken using the integrating sphere as described in the experimental methods chapter. During this experiment, the devices were pumped with an electrical pulse generator using pulse lengths of 1μs and repetition rates of 1kHz. Any references to 'half pulse rate' or 'half rep-rate' relate to these original values.

As a preliminary result, Figure 6.17 shows a selection of this data as well as some extra, taken with different drive pulse lengths and repetition rates. What it demonstrates is that the 'roll-over' feature seen in all of the traces is caused, in the vast majority at least, by bulk heating of the devices, not by heating of the facets alone. This is evidenced by different maximum optical powers achieved and different drive currents at which the roll-over occurs when both the pulse length and repetition rate are halved. The 'retrace' data was taken after all of the others for both devices to demonstrate that the differences seen are not due to device degradation as might be caused by driving them to such high currents.

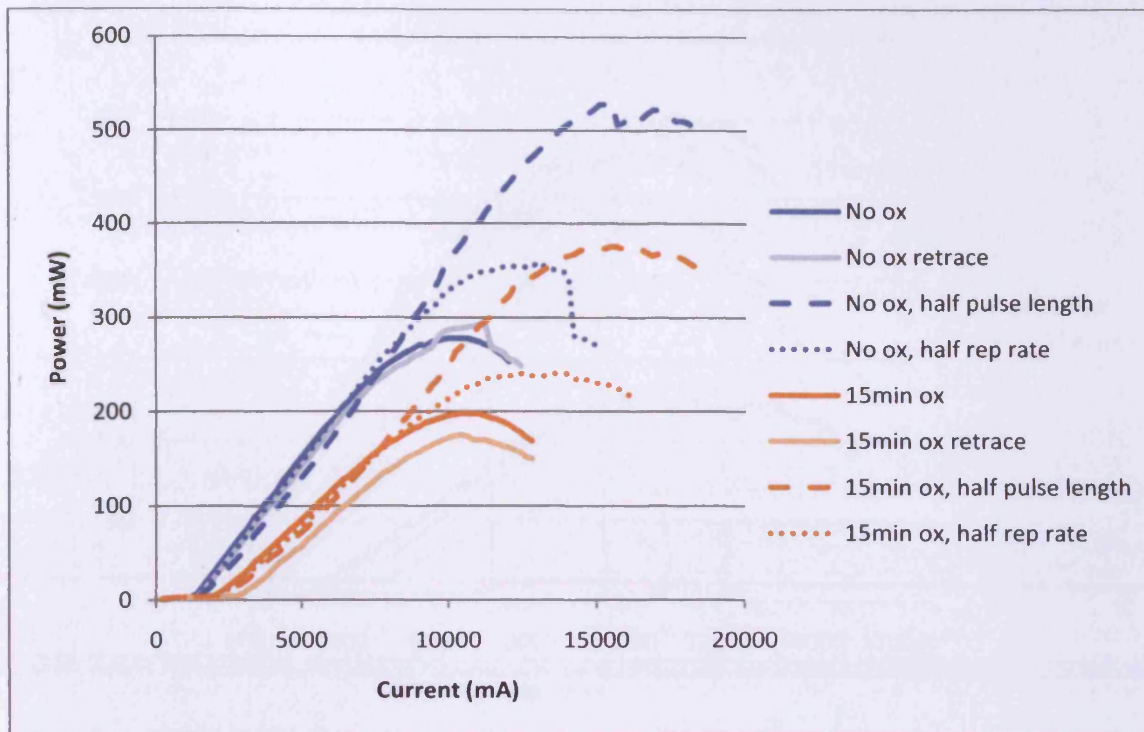

Figure 6.17: PI traces demonstrating thermal effects in oxidised and standard devices

All of the data from this experiment is now shown in Figure 6.18. The first observation to make is that thermal roll-over occurs at approximately the same drive current in all devices, whether oxidised or not. This demonstrates again that is it the heating caused by the electrical drive conditions that gives rise to the roll-over and, importantly, that the oxidised devices behave no differently to standard ones in this respect. It can therefore be concluded that the bulk heating of all devices is truly the same. There will, of course, be some differences caused by greater non-radiative recombination in devices with higher thresholds, but this experiment demonstrates that this is a much smaller effect than the simple 'I^2R' resistive heating caused by passing a current through the material.

With this in mind then, it should be considered that of the 12 oxidised devices tested, seven were seen to fail catastrophically, compared to only one of the 12 standard devices. If the bulk heating of the two device types is the same, as has just been proven, then the increased occurrence of high-power device failure can be attributed, with a great degree of certainty, to highly localised heating of the facet regions.

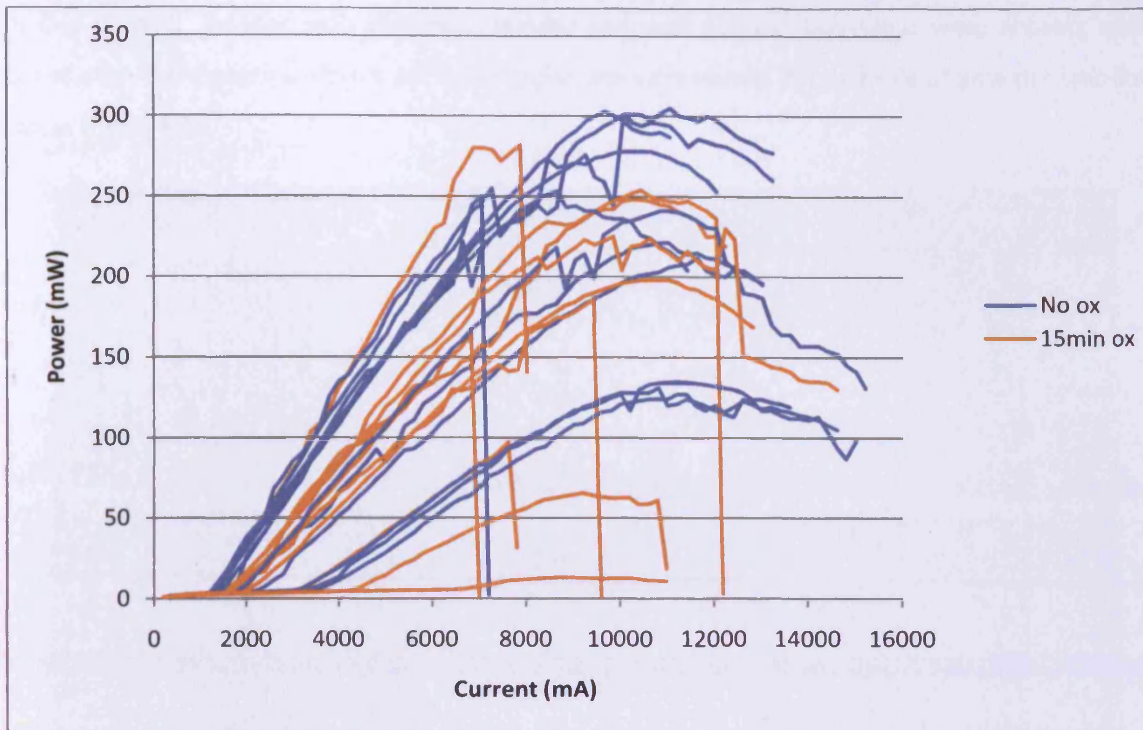

Figure 6.18: PI characteristics of oxidised and standard devices showing catastrophic device failures

An unfortunate secondary conclusion to be drawn from this experiment is that the oxidation performed here has, in fact, made the COD effect worse. However, this does not imply that oxidisable layers used in this manner could never be of use in this respect. With either a pre- or, more likely, post-oxidation treatment of the facet to remove any oxide from the area of the facet that lasing mode interacts with, this technique of removing carriers from small end sections may still prove beneficial in combating COD.

6.7.1.2 RWS analysis to show facet heating

Although the COD data provides good evidence of facet heating, extra weight can be added to the argument by more detailed analysis of the RWS behaviour. The hypothesis discussed in this chapter proposes that the wavelength shifting is caused by the preferential heating of the absorbing sections. If this is the case, then the rate at which that wavelength shifts should be proportional to the rate at which the temperature is increasing. Furthermore, while the rate of the heating will be constant over the whole duration of each drive pulse, the rate of heat loss from the same region will increase as the temperature rises, so overall, the temperature should rise at a decreasing rate towards some eventual limit that may or may not be met within the length of the pulse.

With this in mind, devices with the most regular and well defined behaviour were chosen, and the slopes of each wavelength shift in a particular pulse were measured. An example of how this was done is shown in Figure 6.19.

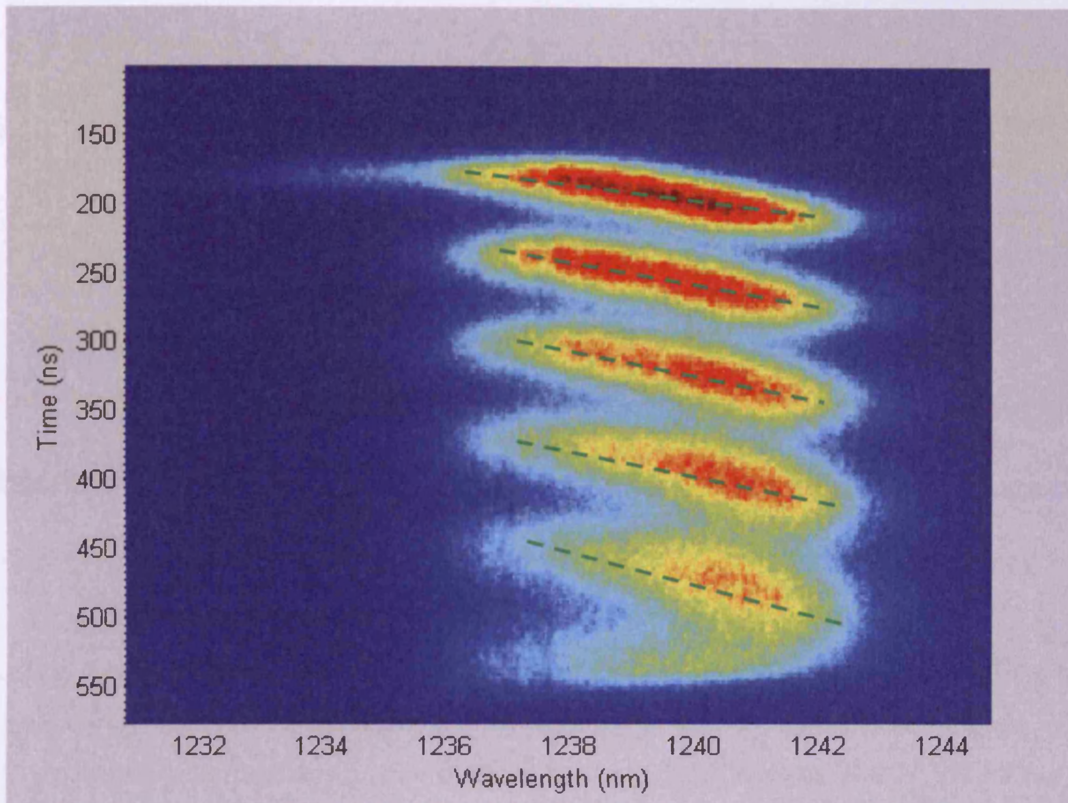

Figure 6.19: Method of measurement of shift rate

In the data shown above, each shift covers approximately the same wavelength, but each successive one takes slightly longer to achieve it than the previous one (note t=0 is just off the top of the image). This data was quantified and recorded for a range of ambient temperatures, and is shown plotted against the position in time through the drive pulse at which each shift occurred in Figure 6.20.

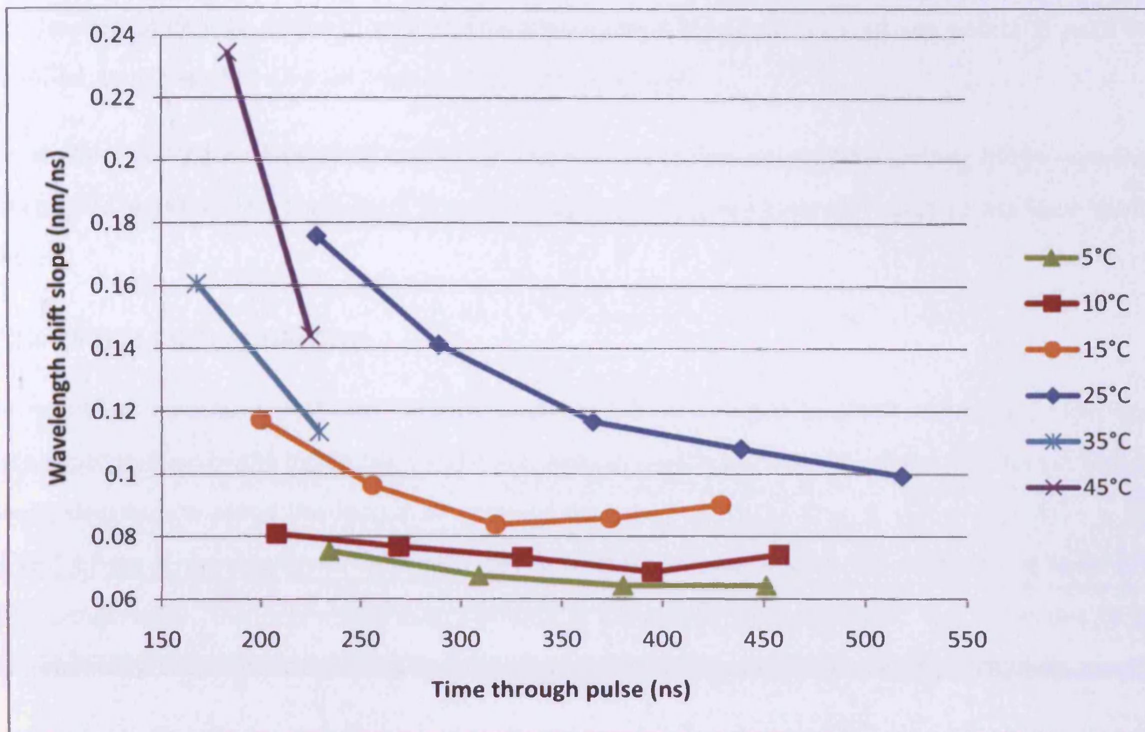

Figure 6.20: Shift rate data for one device over a range of environmental temperatures

Looking first at the data taken at 25°C, the pattern expected can clearly be seen. The initial slope value is high, equating to a high rate of change of temperature in the absorbing sections. The slope value then tails off and begins to approach some limit which is never quite reached, just as the rate of heating would be expected to.

The theory is afforded even further credence by examining how this pattern changes with different ambient or bulk temperatures. In these experiments, the devices were heated or cooled to the selected ambient temperature with a temperature-controlled mount employing a Peltier module. The temperatures stated on the graph reflect those of the bulk of the devices when under test. When held at a low temperature, the bulk of the devices will act as a much better heat sink to the absorbing sections, which means they heat up more slowly and exhibit slower wavelength shifting. Conversely, when heated, the bulk material becomes a poor heat sink and the wavelength shifting becomes faster. It may appear that the 35°C and 45°C data does not quite fit this pattern as they ought to have higher slope values. However, it should be considered that the device threshold is also changing over these measurements, and it is not only the absolute slope value that is a measure of the rate of heating but

also the rate of change of that slope. When the gradient between the first two points of each line is examined, both the 35°C and 45°C data sets fit the trend well.

Overall, this is very good evidence that it is a thermal mechanism causing the shifting of the wavelength, and coupled with the data from the COD experiments is very good indication that excess facet heating is present.

6.7.2 Mode superposition

The extended hypothesis proposes that the origin of RWS is in a combination of localised facet heating and a superposition of the mode combs of three distinct cavities, created by discontinuities in the carrier density distribution along the length of oxidised devices. Further to this, it states that RWS is always present where a narrowed near-field has been produced. Having proven the presence of facet heating fairly categorically, the only remaining element is the mode superposition. This is harder to prove experimentally, but some insight can be gained by revisiting the near-field rebroadening data.

Figure 6.21 shows data taken from the same device, under the same conditions as seen earlier in Figure 6.15, but this time with a monochromator added into the system to achieve what is effectively time-resolved spectral data during the rebroadening process. Although the data is perhaps not as clear, the same lettered sections relating to narrow and broad near-fields are present and labelled in the same manner.

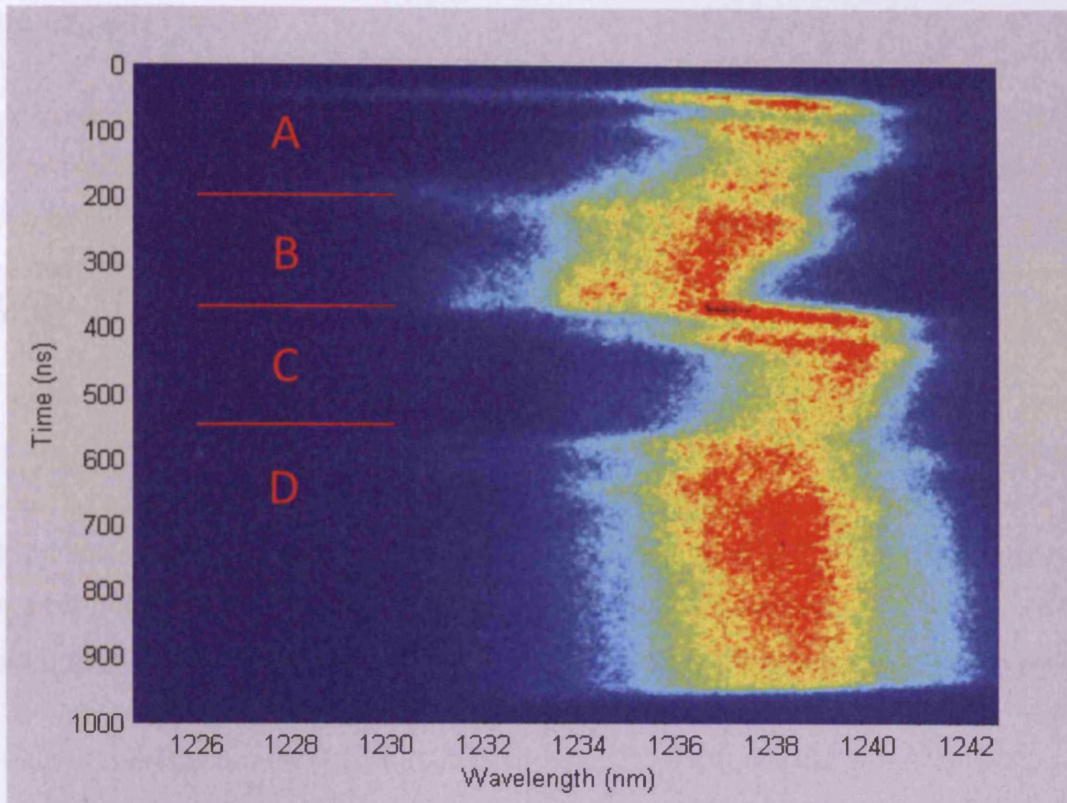

Figure 6.21: Time and wavelength-resolved streak camera data showing the same transitions in 1.5mm device at 1.4I_{th}

What can be seen is that in sections A and C, equating to the narrow near-field sections in the previous figure, there is definite RWS behaviour. Section D (broad near-field) shows some poorly defined shifting behaviour, and section B (broad) shows no discernable wavelength shifting at all. In all of the sections that show RWS, and indeed in all RWS-exhibiting devices measured during this entire project, a common feature is that the RWS behaviour becomes gradually less well defined over the length of the drive pulse, or in this case, the length of the section.

The discontinuities in carrier density cited as the origin of the two new 'cavities' and the superposition of modes will become less well defined over the duration of the pulse, as self-heating increases the ability of carriers to diffuse through the material.

Furthermore, there is a very definite difference in wavelength between the two section pairs. During sections B and D, the emission is slightly broader and reaches lower wavelengths than in section A and C. This corresponds in the hypothesis to the removing of the more selective mode control enforced by the partially saturated absorbers, and is a strong piece of evidence supporting the hypothesis as a whole.

6.8 Further analysis

Although all elements of the effects presented in this chapter have now been discussed in detail, there is still some insight to be gained from further analysis of the devices in general. As yet, no gain or loss data has been examined, and, given that the facets of these devices have been purposefully oxidised – a process that is normally avoided to prevent raising device losses – this information should be considered.

If this method of oxidation is to be used in devices to produce wavelength shifts, pulsation or any other output, the level of extra loss inherently introduced should be known. If the extra loss was evenly distributed along the whole length of devices, the multisection method, as described earlier, could be used to measure it directly. As the loss is localised at the facets, however, the front two sections of any such devices would not be identical, and so the experiment will not work, as was discussed in the Chapter 1. An alternative approach is to use gain and loss data from unoxidised devices of the same material, and derive the additional loss in oxidised lasers with the aid of additional information on the differences in threshold current density. It is this method that is documented in this section.

The gain and loss data presented in Figure 6.22 was obtained by the multisection method, using devices comprising ten identical, 300µm-long sections, with the back facet purposefully damaged to minimise reflections back into the cavity, and reduce round-trip amplification to an absolute minimum. The spectra were measured and recorded on a computer-interfaced monochromator and photodiode array system, as described in Chapter 3.

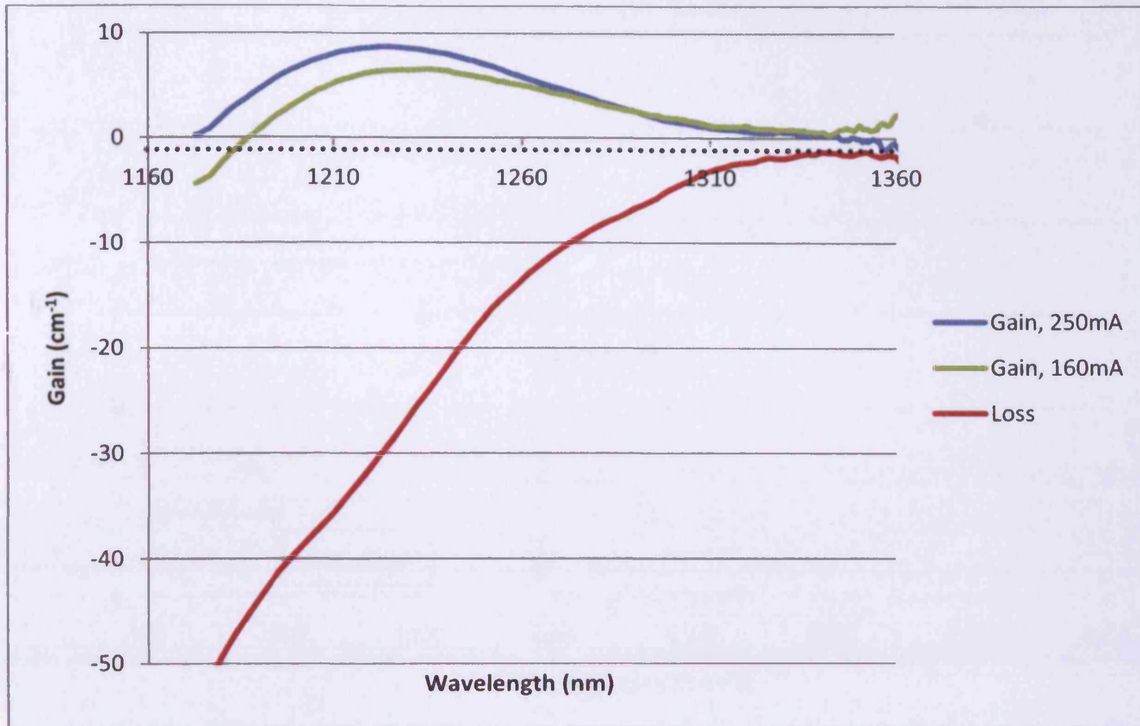

Figure 6.22: Gain and loss data for an unoxidised Wafer A device. Dotted line at -1cm^{-1} shows estimated α_i

The graph shows examples of gain spectra and a single loss spectrum, taken at 300K and smoothed for clarity. The gain spectra have been taken at carefully selected currents, corresponding, in a single 300µm section, to the known threshold current densities of the key batch of devices under discussion here. The loss data is independent of drive current.

Although the agreement in this data is not perfect, an estimate of around 1cm^{-1} can be made for the intrinsic loss parameter, α_i, by examining the convergence of all the traces at the long wavelength. This value of 1cm^{-1} was confirmed, with an error of ±2cm^{-1}, by similar measurements made on different apparatus by James Ferguson of Cardiff University. It should be noted that, despite this error, α_i can clearly never be positive, and also that the value of α_i will not affect the conclusions that will be drawn here, as all gain measurements are comparative. From this figure, a gain-current curve for the material can be drawn by plotting the peak gain value from each trace against its corresponding current density. This data is shown in Figure 6.23.

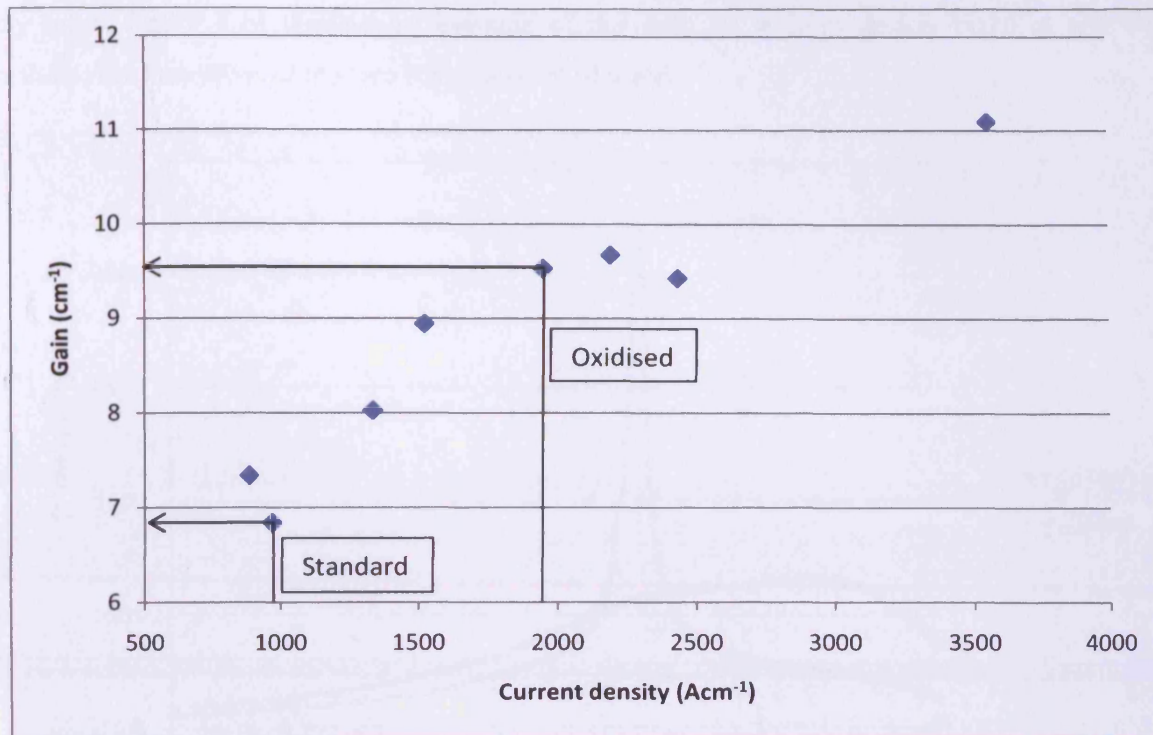

Figure 6.23: Peak gain-current curve for Wafer A. Arrows denote peak gain values for 2mm standard and oxidised devices

The gain-current density relation is an inherent property of a material, and can, therefore, be applied here to both standard and oxidised devices. Using the threshold values from the original data in Figure 6.1, the peak modal gain can be found for each device, and those values for standard and oxidised devices of the same length can be compared. It should be noted that in calculating the threshold current densities for this process, the device widths have been taken as the FWHM values of each laser individually. This is because the oxidised devices have significantly narrower near-fields, as has already been detailed.

Even in those devices that showed a reduction in threshold current in the initial results of this chapter, once the near-field width has been taken into account, all oxidised devices show an increase in threshold current density, as would be expected. On the graph above, this equates to an increase in peak gain required to reach threshold, which, in turn, is equal to the total extra losses added to the devices by oxidation. At this stage, however, this value is averaged over the whole length of each device, which, although useful, can be refined with a few more steps.

The final piece of information required is the lasing wavelength. This was obtained through a simple spectral analysis of the emission from the original laser devices, using a monochromator and photodiode

array setup. Figure 6.24 displays an example of this data for a 2mm device, taken at and below threshold. The intensities of the two traces are not to scale.

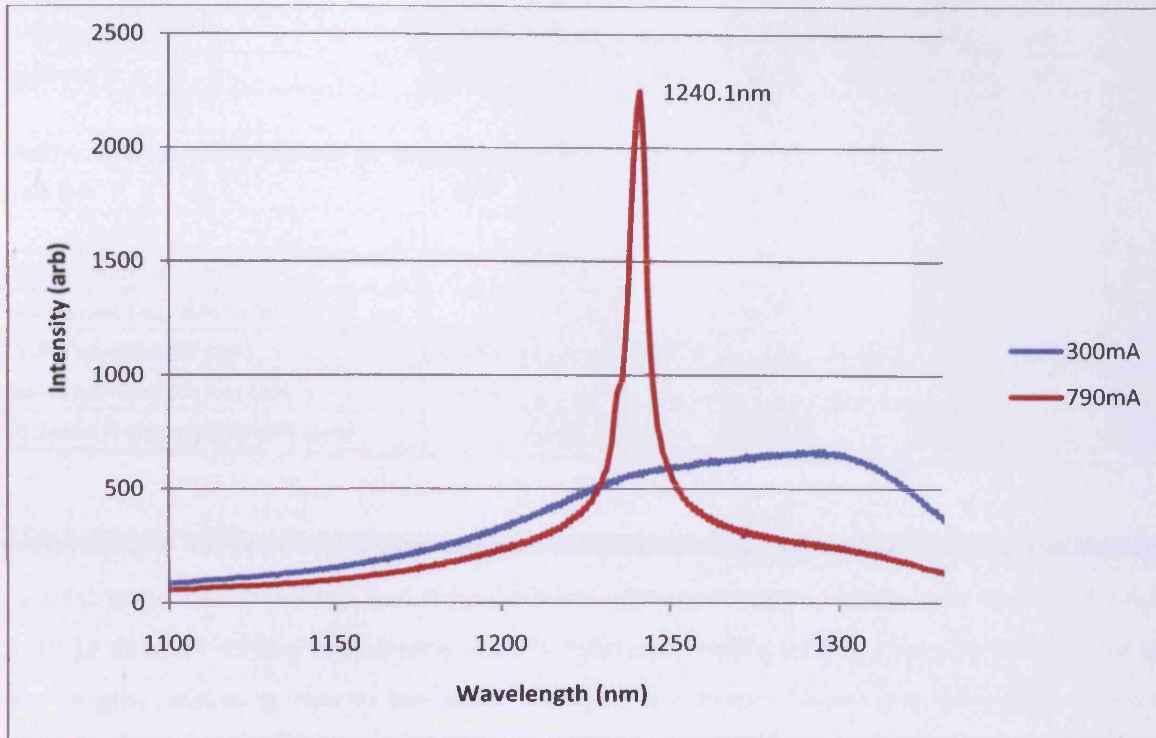

Figure 6.24: Emission spectra at and below threshold of 2mm, non-oxidised Wafer A device

Applying the lasing wavelengths to the loss (absorption) data in Figure 6.22, loss values per centimetre of standard material can be found. The loss per centimetre in the absorbers of oxidised material ought to be very much the same, as the oxide layer was purposefully placed far enough away from the active region to not affect the optical mode significantly. If this is the case, it ought to be possible to derive the original oxidation depth from the values of total extra loss added per device and the standard loss per centimetre by dividing the latter by the former. Table 6.1 is a summary of the data for this entire process for all device lengths.

Table 6.1: Calculation of added losses and absorber lengths. Oxidised devices shown in blue columns

Cavity Length (mm)	1.5	1.5	2	2	2.5	2.5	3	3
I_{th} (mA)	755	1115	795	720	930	805	1280	980
$FWHM_{th}$ (μm)	37.9	20.9	41.1	18.6	41.9	14.8	48.2	21.6
J_{th} (A cm^{-1})	1328	3557	967	1935	888	2176	885	1512
Multisection equivalent currents (mA)	219.1	586.8	159.6	319.4	146.5	359.0	146.1	249.5
Gain (cm^{-1})	8.0	11.1	6.8	9.5	5.0	9.7	7.3	8.9
Total added loss (cm^{-1})		3.1		2.7		2.3		1.6
Added loss per round trip		**0.92**		**1.08**		**1.17**		**0.96**
Lasing wavelength (nm)	1228.5		1239.9		1242.9		1249.7	
Lasing wavelength loss (cm^{-1})	-27.8		-22.3		-20.7		-17.0	
Equivalent absorber length (μm)		**82**		**121**		**140**		**141**

The key results in this table and the common sense checks that support the reliability of the method are the added losses per round trip and the equivalent absorber lengths, which have averaged values of 1.03 ±0.13 and 121 ±38μm respectively. Clearly both parameters should have the same value for all device lengths, assuming exactly the same oxidation conditions. Given the fairly large number of measurements involved in this calculation, each with their own errors, for all of the added loss per round trip values to lie within 13% of their average shows good agreement. Similarly the equivalent absorber lengths all lie within 37% of their average, which is itself, only 19% lower than the expected oxidation depth of 150μm.

The difference between the intended and derived values of the absorber length is within the scatter of the results and, therefore, it cannot be said whether the difference is real or not. This could only be confirmed and quantified by repeating the experiment of this section for a much larger batch of laser devices.

6.9 Summary

In summary then, this chapter has presented two main novel results, both obtained by a method that is itself novel in its application. The method of oxidation of selected layers inwards from the device facets introduces a variable level of saturable loss into cavities of AlGaAs-based laser diodes, but does so in a much cheaper and easier way than has previously been achieved with other methods. The oxidation

method also allows for p-side-down mounting, favoured in commercial devices for greater heat dissipation, unlike other methods. The results obtained from the employment of this method are, firstly, the narrowing of near-fields and consequent changes in thresholds brought about solely by saturable loss. As well as being greatly improved in terms of shape and reduction of filamentation, near-fields were narrowed typically by around 50%, and thresholds were raised or lowered by up to 30% depending on device length. Secondly, the production of repetitive wavelength shifting behaviour at repetition rates of around 10MHz and total shift lengths of 5nm was also achieved by the same simple method.

These effects have been studied extensively and, in doing so, detailed hypotheses of the physical mechanisms behind them have been proposed and, through the course of this chapter, substantiated with various experimental evidence. The hypotheses are based on a combination of saturable absorption, increased localised heating of the facets and increased mode spacing due to the superposition of modes from multiple reflective interfaces.

6.10 References

1. *Semiconductor saturable absorber mirrors (SESAMs) for femtosecond to nanosecond pulse generation in solid-state lasers.* **Keller, U.; Weingarten, K. J.; Kärtner, F. X.; Kopf, D.; Bruan, B.; Jung, I. D.; Fluck, R.; Hönninger, C.; Matuschek, N.; Aus der Au, J.** 3 : IEEE Journal of Selected Topics in Quantum Electroncs, 1996, Vol. 2.

2. *Advances in selective wet oxidation of AlGaAs alloys.* **Choquette, K. D.; Geib, K. M.; Ashby, C. I.; Twesten, R. D.; Blum, O.; Hou, H. Q.; Follstaedt, D. M.; Hammons, B. E.; Mathes, D.; Hull, R.** 3 : Journal of Selected Topics in Quantum Electronics, 1997, Vol. 3.

3. *Amplified spontaneous emission and carrier pinning in laser diodes.* **Chuang, S. L.; O'Gorman, J.; Levi, A. F. J.** 6 : IEEE Journal of Quantum Electronics, 1993, Vol. 29.

4. *Improvement of catastrophic optical damage level of AlGanInP visible laser diodes by sulphur treatment.* **Kamiyama, S.; Mori, Y.; Takahashi, Y.; Ohnaka, K.** 23 : Applied Physics Letter, 1991, Vol. 58.

5. *Direct measurement of facet temperature up to melting point and COD in high-power 980-nm semiconductor diode lasers.* **Sweeney, S. J.; Lyons, L. J.; Adams, A. R.; Lock, D. A.** 5 : IEEE Journal of Selected Topics in Quantum Electronics, 2003, Vol. 9.

Chapter 7 : Conclusions and Future Work

7.1 Conclusions

In conclusion of this work as whole, this section will go through the project, summarising all the work undertaken and state the main achievements and their relevance to the general body of knowledge in the field.

7.1.1 Modifications to oxidation equipment

Preliminary work completed before the main body of experiments was started involved practical improvements to the oxidation equipment. Specifically, the kit was modified to include valves allowing an easy switch between dry and steam-carrying nitrogen. This allowed more accurate timing of oxidation runs, the beginning and end of which were previously controlled only by the slow ramping up and down of the sample temperature.

7.1.2 Wafer designs and oxidation rate characterisation

The oxidation rate of two oxidisable layer systems have been characterised over a wide range of oxidation times and temperatures. One of these systems, Wafer A, was comparable to other systems reported in the literature and so served as an opportunity to compare the performance of the oxidation process. Results from this wafer showed good agreement with multiple literature sources. The second system was an entirely novel one. Wafer B was novel in two factors – the use of very thin, 1nm layers in alloyed stacks with lower aluminium composition material, and the placement of those stacks directly within the active region. The oxidation rates of this wafer have been measured directly and also show good agreement with values predicted by extrapolation of other data. The oxidation mechanics of the two systems have been revealed by demonstrating that under all conditions tested Wafer A showed a linear progression of the oxidation front, whereas Wafer B displayed a parabolic progression. This is indicative of a reaction-rate-limited mechanism in Wafer A, and a diffusion-rate-limited one in Wafer B. In both cases, this is the first practical confirmation of this information in these layer systems.

7.1.3 Dopant dependence of oxidation rate

A dependence of the oxidation rate on dopant species has been discovered. It was found that although p-type material oxidised fastest for short oxidation depths (<8μm), the n-type material then overtakes

and remains ahead for all deeper oxidations tested here (up to 20μm). This conclusion might be contrary to popular belief, which is based on early reports of p-type material oxidising faster in very large, surface-deposited layers.

7.1.4 Sidewall profile dependence of oxidation rate

The importance of straight sidewalls in structures to be oxidised has been demonstrated by showing the heavily diffusion-limited rate of oxidation in samples with a large undercut. The same layers, when etched with straight sidewalls, show a linear, reaction-limited rate, implying that the access of steam to the oxidisable layer has been severely restricted in the undercut sample.

7.1.5 Current confinement by active region oxidation

Lateral oxidation of active region layers has been successfully used to confine the current in broad-area laser devices. It has been shown for the first time that this is possible, not only without significantly detracting from the electrical performance of the wafer, but also without inducing a large increase in α_i. As a direct consequence of this current confinement, the threshold of the lasers has been reduced over a range of operating temperatures. It is predicted that this range and the size of the effect would be increased by using the technique on narrower devices.

7.1.6 Microlens fabrication

A simple and novel method of fabricating buried microlenses has been uncovered. The growth of a thin 'feeder' layer on top of a thick, mid-range aluminium content layer has been shown to result in a curved oxidation front, and consequently, a curved refractive index profile that could be employed as a lens structure. Initial results show an RoC of 29μm, corresponding to a focussed beam spot size of less than 4μm.

7.1.7 Near-field and threshold current changes by facet oxidation

By oxidation in the longitudinal direction from the device facets of lasers, saturable absorbers have been created. This has been confirmed by examination of LI characteristics and by the observation of unusually sharp turn-on of optical output. The saturable absorbers have been used to narrow and improve the uniformity of the near-fields of laser devices, and consequently, to change their threshold currents. Changes in threshold ranged linearly from +30% in 1.5mm devices to -30% in 3mm devices. A full mechanism for this behaviour has been hypothesised, and in doing so, evidence for significant sub-

threshold ASE has been acquired. This is a subject with much debate and little reported evidence in the literature, and is therefore a valuable conclusion to be drawn.

7.1.8 Repetitive Wavelength Shifting

The novel effect of RWS has been reported. Using the same saturable absorbers, repeated wavelength shifts and switchbacks were created. Typical the wavelength shifts occurred over a range of 5nm and with repetition rates of around 10MHz. The effect is not conventional self-pulsation, but is related and, again, has been explained with a full hypothesised mechanism. RWS is believed to be caused by a combination of the saturable absorbers, enhanced localised heating of the facets and a superposition of modes caused carrier density discontinuities in the cavity.

7.1.9 Losses added by oxidation

Finally, the losses added the longitudinally oxidised devices, irrespective of device length, have been calculated to be 1.03 ±0.13 per round trip. This was achieved using gain and loss data taken by the multisection method to plot a gain-current curve. From this, the differences in threshold current density between oxidised and standard device were used to determine the total added losses in each device. The loss per round trip was then used to derive the original, known oxidation depth as a means of confirming the loss value.

7.2 Future work

Through its course, this work has given an insight into several areas of investigation that warrant further attention. This section gives a list of these areas and details of how each of them might be begun.

7.2.1 Oxidation rate dependence on doping

Although the results presented here show a clear trend, they do not conclusively provide a reason for it. Given that the data seems, at first, to contradict other literature results, it would seem prudent to more thoroughly investigate the causes. It is suggested here that the doping dependence is different for thin, buried layers than for large amounts of surface-deposited material, but the question remains as to whether the cause is the thickness of the layers or the fact that they are buried, or indeed, a combination of both. A simple method of establishing the answer to this question would be to design a test wafer containing several well isolated layers of varying thickness and of a range of dopant levels of

both types. By oxidising basic mesa structures of this material to various depths the two possible causes could be isolated, and greater detail in the relationship between oxidation rate and doping could be established. This information could be invaluable in the manufacture of VCSELs, where oxide current apertures are often employed in both cladding regions of the same device.

7.2.2 Reduction of delamination

A worthwhile investigation would be to look into exactly what conditions of oxidation cause notable incidence of delamination and what methods can be used to reduce it. Specifically, looking into what thicknesses of single layers can be partially oxidised without delamination issues, and whether that figure can be increased, in the specific case of curved oxidation fronts, by having the deepest points of the curved front actually oxidised all the way through. This would be beneficial to the fabrication of microlenses, if it were the case. Further to this, other methods such as backfilling ridge structures after oxidation could be investigated as a method of strengthening. Again, a simple test wafer could be grown for this purpose.

7.2.3 Near-fields and thresholds of Wafer B

In order to substantiate the theory suggested in Chapter 5 regarding the temperature dependence of the lowering of threshold by oxidation, near-fields of the devices need were taken for the temperatures in question, although they proved inconclusive.

A step further would be to fabricate narrower ridge devices from the material and repeat the experiment. This may well provide information of the effect of the oxide on the lasing mode, if any. If the ridges were also made into multisection devices, the gain and loss of the material could also be measured. This would be an interesting insight as the oxide would be in contact with the lasing mode along the whole length of the devices, unlike in Wafer A, and so might affect α_i significantly.

7.2.4 Microlenses

This is an area that could be developed considerably. The setup of a thin, high aluminium feeder layer grown immediately on top of a thick, lower aluminium layer is a simple and versatile one that could produce a variety of shapes and sizes of microlens. Again, a preliminary test wafer would need to be grown to investigate the exact relationships between factors such as the thickness and composition of both layers and the temperature of oxidation. This could be fairly efficiently done, however, by growing a single wafer with multiple sets of these pairs of layers, and it would simultaneously provide a greater

insight into the cross-layer diffusion of oxidant species and waste products than is current available in literature.

Caution might need to be taken in a few factors, namely the problem of delamination and of the roughness of the interface between oxidised and unoxidised material. Ultimately though, with careful design and selection of aluminium compositions, it ought to be possible to fabricate VCSEL devices of various sizes with buried oxide current apertures and monolithically integrated oxide microlenses from a single etch and oxidation.

7.2.5 Reducing COD in Wafer A

A brief exploration of the effect of oxidation on the COD limit was undertaken as part of the work on Wafer A. The devices used in this test though, were not optimised to actually reduce COD, as this was not the main aim of the experiment. Despite the negative result with regard to COD reported here, the use of the oxide system in Wafer A still has the potential to be used in protecting against COD. It is highly likely that it was the oxide grown unintentionally on the active region of the facet that was the root of the detrimental effects seen in Chapter 6. A valuable investigation would now be to create a very similar batch of broad-area devices, oxidise them probably to a lesser depth than was used here, and then treat the facets to removed the surface oxide from the area that will be seen by the lasing mode. The oxide in these areas is likely to be relatively thin, in the order of 10nm, and so should be easily removed without affecting the intentional oxide layer growth. Suitable methods for doing this might be ion milling or a gentle wet oxide etch. Oxide layer depths of only a few microns ought to be sufficient to isolate the facet areas from carriers, thereby reducing non-radiative recombination.

7.2.6 Laser turn-on delay

One unexpected observation that came out of this work was the measurement of very large turn-on delays in both oxidised and standard devices. The effect was highlighted in the oxidised lasers by the exaggerating effects of the saturable absorption but was also clearly present in standard lasers. A plausible mechanism behind this effect has been suggested here, but as there seems to be no explicit mention of such behaviour in the literature further investigation would seem warranted. Perhaps the best initial step towards this would be to develop a model to test the theory suggested here. Some insight might then also be gained into the abrupt jumps seen in the turn-on delay in oxidised devices, which are currently more or less unexplained.

Presentations and publications

Characterisation and applications of AlGaAs oxidation in 1.3μm edge emitting quantum dot lasers. **Gareth J. Michell, Peter M. Smowton and Huw D. Summers,** SIOE, 2007

Manipulation of optical modes in quantum dot laser diodes by selective oxidation of high aluminium content AlGaAs layers. **Gareth J. Michell, Peter M. Smowton and Huw D. Summers,** SIOE, 2009

Manipulation of optical modes in quantum dot laser diodes by selective oxidation of high aluminium content AlGaAs layers. **Gareth J. Michell, Peter M. Smowton and Huw D. Summers,** Photonics West, 2010