

University of New Hampshire

University of New Hampshire Scholars' Repository

Honors Theses and Capstones

Student Scholarship

Spring 2014

Forcing Mutual Coherence in Diode Laser Stacks

Jonathan R. Wurtz

University of New Hampshire - Main Campus

Follow this and additional works at: <https://scholars.unh.edu/honors>



Part of the [Atomic, Molecular and Optical Physics Commons](#), and the [Optics Commons](#)

Recommended Citation

Wurtz, Jonathan R., "Forcing Mutual Coherence in Diode Laser Stacks" (2014). *Honors Theses and Capstones*. 189.

<https://scholars.unh.edu/honors/189>

This Senior Honors Thesis is brought to you for free and open access by the Student Scholarship at University of New Hampshire Scholars' Repository. It has been accepted for inclusion in Honors Theses and Capstones by an authorized administrator of University of New Hampshire Scholars' Repository. For more information, please contact Scholarly.Communication@unh.edu.

Forcing Mutual Coherence in Diode Laser Stacks

UNH Honors Senior Thesis Project

Advisor: Dr. Bill Hersman

Jonathan Wurtz

Abstract

This paper will discuss both theoretical and experimental attempts to improve the spatial beam quality of diode laser stacks using an external optical system. An overview and derivation of the mathematics of both the optical system and diode lasers will be discussed. The experimental setup will be presented, as well as the fundamental theoretical and experimental results that suggest the external optical system used for this thesis fails to improve the beam quality of a diode laser stack.

Contents

Introduction	3
<u>Theory</u>	
Population Inversion	4
Diode Lasers and Silicon Junctions.....	5
Diode Laser Rate Equations	6
Simulating Diode Lasers	7
Laser Resonators	10
Beam Propagation Methods	12
Matrix Propagations.....	14
<u>Experiment</u>	
Background	17
Assembly	19
Phase One: Calibration.....	20
Phase Two: Experimentation	23
Results and Conclusions.....	25
Theoretical	25
Experimental	27
Future Work.....	29
Conclusion.....	29
References	30
Acknowledgements.....	30

Introduction

This document describes the work done for senior thesis work during the spring of 2014, investigating effects of mutual coherence in diode laser stacks. Diode lasers are a common type of solid-state laser which are used in applications ranging from scanners to laser welding and cutting. However, high-power diode laser systems have an inherent trade-off: A single diode laser is difficult to scale to high power, so instead arrays of many single low-power emitters are assembled together to create a diode laser stack, which now outputs a high-intensity beam. However, each diode lases independently, which means that the spatial beam quality, which allows for tighter control of the beam, is poor.

The goal of this project, as is outlined in the paper below, was to improve the spatial beam characteristics of a diode laser stack by correlating the light from each diode with other diodes, which means the entire assembly acts like was one large emitter. This was done using an external optical system, which reflects light from the laser back on itself and thus mixes light from different diodes together, causing a correlation.

In this project, this goal was pursued in two ways. The first was a pursuit from a theoretical aspect, which is introduced in the first half of this paper. Here, the goal was to characterize the theoretical light output from a physical system, by understanding both how the laser system works as well as the external cavity. This paper will reflect those pursuits, by first explaining how a diode laser works, as well as discussing rate equations and attempts to solve them. Next, the mathematics of laser resonators and optics propagation will be developed, as well as an explanation of applications and a demonstration of some of the mathematics.

The second goal was a pursuit from an experimental aspect, which is introduced in the second half of this paper. Here, a physical system was built utilizing a 100W diode laser stack to see if the theoretical results could be confirmed. This paper will first explain the physical principles of the system, and then continue by showing the physical system and listing the calibration procedures to make it work. Once the explanation is complete, the paper will conclude with a discussion of some of the results of the project, with an emphasis on what did and did not work.

Population Inversion

A laser uses the principle of population inversion, where there are more electrons in the excited state of a material than the ground, or non-excited, state. These electrons can do one of two things in order to become de-excited: they either spontaneously emit photons, or they be stimulated to emit photons. When there is population inversion in the material, there is positive gain in the material, which means that there are more photons that come out of the material than go into the material.

Spontaneous emission is the same process that occurs in neon lights. In this process, an excited electron moves from one energy state to a lower energy state, where the energy difference is released as a photon of that energy with a characteristic average time delay at which the electron switches states, normally measured in picoseconds to nanoseconds. In this way, a gas light works by electrically exciting a gas via high voltage, which then glows at specific wavelengths of light based on the energy differences between excited and non-excited states.

Stimulated emission is completely different. When an electron decays to the ground state spontaneously (spontaneous emission), it does so by interacting with a virtual photon—in other words something that “pops in and out of the vacuum”. Stimulated emission is when the excited electron instead interacts with a real photon of approximately the same energy as its energy change. Here, the photon emitted has the exact same frequency, direction, and phase as the stimulating photon, effectively “replicating” the incident photon. Thus, the incoming light is amplified (doubled) by the excited atom.

Similarly, a photon can be absorbed by an atom in the ground state if its energy is close to that of the change in energy of an electron to an excited state. In this case, incident light upon this material will be damped out and attenuated, and the energy in the photon flux will be changed to electronic energy in the excited atoms.

There is then a balancing act between stimulated emission and absorption: suppose there is a material with a gas of photons at a specific energy corresponding to an electronic energy level change (By gas it is meant that photons are uniformly distributed in space and angle). There are two competing processes of stimulated emission, which increases the number of photons but decreases the electronic excitation, and absorption, which decreases the number of photons but increases the electronic excitation. The absorption rate (photons per atom per second) is based on the number of atoms in that material that are in the ground state, $\frac{d\theta}{dt} = -\theta B_G N_G$, where θ is the photon density (photons/volume), N_G is the density of electrons in the ground state (electrons/volume), and B is the Einstein coefficient describing the rate at which photons are absorbed (1/time). The emission rate is based on the number of atoms in that material in the excited state, $\frac{d\theta}{dt} = \theta B_E N_E$. There are a finite number of atoms in the material, so thus there are a finite number of electrons, so $N_g + N_E = N_{tot}$. By an argument of time reversal, $B_G = B = B_E$, so the total rate of photon change is given by $\frac{d\theta}{dt} = \theta B(N_E - N_G)$. If there are more electrons in the ground state, then the amount of photons in the material will decrease. If there are more electrons in the excited state, then the amount of photons will

increase. This is population inversion—with more of the material in the excited state than the ground state, then it is possible to get more photons out of a material than you put in. This is how a laser works—by exciting a material there is positive gain in that material. Gain, in this sense, is the increase (+ gain) or decrease (- gain) in the number of photons as a ray of light goes through the material. Thus, the name of the game for any laser system is to get this effect.

Diode Lasers and Silicon Junctions

This section will introduce the laser system used in this thesis: the diode laser. The diode laser is a biased PN semiconductor junction that is doped to have a direct bandgap with an energy width at the photon lasing energy. In this section, the basics of semiconductors and doping for lasers with PN junctions will be discussed.

A PN junction is where a positively doped semiconductor is put next to a negatively doped semiconductor. Silicon naturally has 4 electrons and 4 holes (holes here mean the lack of a filled electron orbital, and in this context are considered to be particles in and of themselves), so there is no surplus of either. However, silicon can be “doped”, meaning extra materials can be added, such that there is a surplus of holes (P-type, or Positive-type), or a surplus of electrons (N-type, or Negative-type). The doping elements can be determined by looking at the periodic table—elements to the left of silicon (for example, Gallium) have 3 electrons and 5 holes, making the doping P-type. Equivalently, Phosphorus has 5 electrons and 3 holes, making the doping N-type. This doping is worthy of lifetimes of work; other elements can be used, in different concentrations, with different deposition processes, generating a massive number of different kinds of materials, each with different properties.

The two most important properties are the Fermi level of the material and the band structure of the material. The band structure of the material is the functional form of the energy relation to momentum for electrons. In free space, $E \sim p^2$, but because a semiconductor forms a crystalline lattice, there is a more complicated form. Furthermore, there are actually two bands for a semiconductor (for metals, too): the valence band, and the conduction band. In the conduction band, electrons are considered to be mobile (thus, the material is conductive) and can be considered to be the “excited state” of the material. The valence band can be considered the “ground state”. Nominally, there are two types of “bandgaps”, or minima-maxima points: direct, and indirect. For a direct bandgap, there is no momentum requirement for an electron to jump from the conduction band to the valence band. An indirect bandgap requires momentum transfer, normally from phonons (the quanta of sound waves). Diodes that output light are direct-bandgap, because photons transfer a negligible amount of momentum to the system. An example band structure is shown in figure 1.




Figure 1: Example band structure of a material.

Electrons are fermions, which mean they obey the Pauli Exclusion Principle—only one is allowed per state. As such, they all “fill up” the lowest energies until they hit a highest state: the Fermi energy. The Fermi energy of a semiconductor is in-between the conduction and valence band. Doping a semiconductor changes the Fermi energy of the material. When a P type and N type are put next to each other, there is a region in-between where the Fermi energy changes. It is then possible to bias the junction, which means that a voltage, and thus an electric field, is applied across the junction. This then causes the transitioning Fermi energy to split into two “quasi-Fermi” energies, because the electrons on one side are “higher energy” then the others, due to the voltage difference. These “quasi-Fermi” levels are defined now for both the conduction and valence bands, which means that there is some electron population in both bands—with more voltage comes more electrons in the conduction band, which means more current and thus more power through the junction. This is what is required for a laser: to have more electrons in the excited state then there are in the ground state—and this is exactly the case! By forward-biasing a specifically built PN junction, population inversion is achieved, yielding positive gain and thus lasing.

Diode Laser Rate Equations

Now that the basics of diode lasers are outlined, it is now relevant to outline rate equations. This is important because it allows for modeling of the laser output, one of the main theoretical goals of the project. Rate equations are a set of coupled differential equations which describe the rate of change in time of the amount of light in the laser and the density of electrons in the excited and ground states. Generally, these equations are quite difficult: they are normally first or second-order nonlinear differential functions of both space and time. However, they can be simplified in several ways: the steady-state solution, where nothing is changing in time (thus time derivatives are zero) is normally the solution that is pursued, and is relevant for times much larger than the nanosecond-scale rise times of the laser. Furthermore, the spatial components can be averaged out for spatially homogenous systems, yielding an equation of only two variables.

The diode laser rate equations are effectively impossible to solve analytically in their general form, and difficult to solve numerically without resorting to trivializing simplifications. The generalized rate equations are as follows, with characteristic values in the table in the appendix:

$$-\frac{dN}{dt} = -D\nabla^2 N - \frac{J}{qd} + \frac{N}{\tau_{nr}} + BN^2 + \frac{\Gamma a(N - N_0)}{\hbar\omega} |E|^2$$

This is the rate equation for the electron density, describing the amount of electrons in the excited state of a material. Here, $-D\nabla^2 N$ specifies the rate at which electrons diffuse around the material—this part “smooths” the distribution. $-\frac{J}{qd}$ is the rate at which electrons are added by current pumping. $\frac{N}{\tau_{nr}}$ is non-radiative recombination, where the electron falls back to the valence band via emitting a phonon, aka heat. BN^2 is the rate at which electrons are lost due to spontaneous emission. $\frac{\Gamma a(N - N_0)}{\hbar\omega} |E|^2$ is the rate at which electrons are lost due to stimulated emission, aka the lasing action. Γ

is the confinement factor, describing how much of the Gaussian-shaped mode in the current-flow direction (aka, 0.2microns thick) is actually in the active region. Below is the equation describing the light-field in its spatial extent and steady state.

$$\frac{dE_f}{dz} = \frac{i}{2k} \frac{d^2E}{dx^2} + E_f \left[\frac{\Gamma}{2} (1 - i\alpha) a(N - N_0) - \frac{\alpha_{int}}{2} + in_2k_0(|E_f|^2 + 2|E_r|^2) \right]$$

Here, $\frac{dE_f}{dz}$ is how the intensity increases along the propagation direction. $\frac{i}{2k} \frac{d^2E}{dx^2}$ is diffraction, using the Helmholtz approximation. The Helmholtz approximation is that in which the high-frequency spatial components in the electric field like that from oscillating light are split off from slowly-varying components like that from diffraction or index of refraction changes. Thus, $E_f = constant$ means that the light is still oscillating, but the field is not changing in any way. The diffraction term derives naturally from this approximation. $\frac{\Gamma}{2} (1 - i\alpha) a(N - N_0)$ is the gain or loss from there being population inversion. There is also a complex term in this equation, $i\alpha$, which is the linewidth enhancement factor. Basically this describes an index of refraction change that is due to there being a different electron density—this comes out as a “slowly varying” complex factor because now that there is a different index of refraction, the wavelength is changed from its central value and thus has a phase delay. $\frac{\alpha_{int}}{2}$ are internal losses, due to scattering and the like. $in_2k_0(|E_f|^2 + 2|E_r|^2)$ is the factor for self focusing/defocusing—a non-linear effect of higher-intensity light changing the index of refraction of the material. This part of the equation, similar to the linewidth enhancement factor, causes a slow phase advance or delay in the wave front.

Note that there is actually no time dependence in this equation: instead there is a differential with respect to Z, the propagation direction. Thus, this equation describes the gain of the laser. Time dependence, if required, can be introduced using an argument of conservation of photons (similar to conservation of charge) with unit analysis to include the speed of light:

$$\frac{c}{n} \frac{dE_f}{dt} = \frac{dE_f}{dz} + [-Gain Terms -]$$

Suppose there is a small box with some number of photons in it. There are photons flowing in, photons flowing out, and photons being generated. The spatial derivative tells how many are flowing in one side verses out the other, while the time derivative describes the “pile-up” in intensity. The factor of c, the speed of light, makes it so everything happens on very short timescales—for example, for a numeric simulation where the gain path was 300microns long, split up into 1000 bins (300nm to a bin), the natural timestep is $\frac{\Delta x}{c} = \frac{3 \cdot 10^{-7} m}{3 \cdot 10^8 m/sec} = 10^{-15} sec$, orders of magnitude below the natural rise time of the laser, which is on the order of $10^{-9} sec$.

Simulating Diode Lasers

The goal on the theoretical side of the thesis was to simulate and predict what the output of the diode laser system would be if the spatial coherence system was built. Part of that goal was to simulate the “action” of a diode laser—in other words, how much light is put out as a function of back reflection

and current. This section will explain some of the work done in order to calculate this function, outlining the methods used to try to calculate the diode laser action, and explain where they succeeded and why they ultimately failed.

Generally, the differential equations to be solved are 2-dimensional: one in the lateral direction perpendicular to beam propagation, and one in the direction of beam propagation. The third dimension is very thin, so it is assumed that the light forms a Gaussian of some width, which comes out of the equations as the factor Γ . One simplifying assumption made for simulations was each stripe in the beam propagation was independent of adjacent stripes; all X-derivatives were set to zero. This assumption works well in a thin, free-running laser, but does not work as well for this thesis—however, it simplified simulation considerably and could be reasonably easily added in, at the expense of extra required computing power (imagine: simulating a 1000x1 space vs. a 1000x100 space; the simulation speed is decreased by a factor of 100).

The main simulation attempt was an analytic one. Instead of calculating the differential equation in 1 dimension, it is possible to average along that dimension and now have everything in terms of single discrete variables. However, for a high-gain laser, the light intensity on one side of the gain medium may be much less than the light intensity of the other, because the light is amplified many times over the region. Thus, a naïve single average may be incorrect for high gain, because it would not get the large-scale “non-linear” effects due to this light intensity discrepancy. Instead, a first-order average was taken, where the intensity field and excited electron density were taken to be:

$$E(z) = E_0 + Bz$$

$$N(z) = N_0 + \beta z$$

Here, E_0 and N_0 are the averages across the material, and B and β are the slowly varying linear components. It is then possible to plug these into our equations to calculate the solution. The intensity field equation becomes a non-linear differential equation of the form $\frac{dE}{dz} = E(a + bz)$, with the solution $E = \exp\left(az + \frac{bz^2}{2}\right)$. Assuming that the photon density is approximately two exponentials to zeroth order, the first order B can be found to be $B = \alpha^2 L E_0$. The electron density can be calculated via integration, canceling out higher-order terms of βB , B^2 and β^2 , to find that $\beta = \alpha^2 L (N_0 - N)$. Assembling everything it was found that the gain across the whole diode (one back-and-forth pass) was:

$$E_{out} = E_{in} \exp\left(\left(\frac{[\Gamma\alpha(N - N_0) - \alpha_{int}]}{2} + i(\alpha\alpha(N_0 - N) + n_2 k_0 E)\right)L\left(1 + \frac{[\Gamma\alpha(N - N_0) - \alpha_{int}]^2 L^2}{2}\right)\right)$$

Here, all quantities are given by the table in the diode rate equations section except L, which is the length of the cavity. The L^2 term is the error term, aka from there being a first-order linear function from B and β . The cool thing is, now this term can be used as a “check” for the zeroth order approximation. Thus, for numeric simulations, the zeroth order approximation can be used, but can be checked for validity by seeing that the following approximation is satisfied:

$$L^2[\Gamma a(N - N_0) - \alpha_{int}]^2 \ll 1$$

$[\Gamma a(N - N_0) - \alpha_{int}]$ is the real part of the gain term, so this is the criterion for the low gain regime, or that $\exp(La) \approx 1$. Coupled to unit analysis, this provides a good “consistency check”. It should also be noted that it is possible to calculate higher-order terms from this expression such as $E(z) = E_0 + Bz + Cz^2$, but that would be horrendously difficult (and for the purposes of this project, unnecessary). Similarly, the rate equation for electron states becomes, through integration:

$$D \frac{d^2 N}{dx^2} = -\frac{J}{qd} + \frac{N}{\tau} + BN^2 + \frac{\Gamma a}{\hbar\omega} (N - N_0)E_0 + \frac{\Gamma a L^2 E_0 [\Gamma a(N - N_0) - \alpha_{int}]^2}{\hbar\omega} (2N - N_0)$$

Here, the L^2 term is the error term from the linear factor. Similarly, the zeroth order approximation holds as long as this error term is much smaller than the rest of the expression.

With these two expressions in hand, it is now possible to do analytic calculations of diode output power as a function of current. The boundary conditions of the problem become that of the reflectance of the partially reflective mirror on one side:

$$\frac{E_{in}}{E_{out}} = \sqrt{R}$$

Here, R is the reflectance of the mirror (say, 30%). Now this problem becomes a question of solving 2 equations (gain and rate) and two unknowns (intensity out, and electron population). This becomes semi-trivial analytically using numerical solving methods, in this case using Matlab. Using values similar to those given by the example values in the diode laser rate equation section, a cavity length of $L = 200\text{micron}$, and a reflectance of $R = 75\%$, we get a plot such as that in figure 2. Here it can be easily seen that the laser does not lase when it is below threshold, then increases linearly with some slope efficiency past threshold, just as a natural lasing action does.

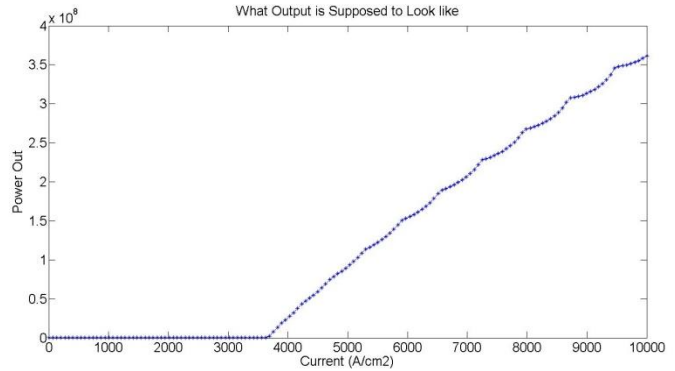


Figure 2: Theoretical output of a single diode laser. The laser hits a threshold at about 3,700 A/cm² then linearly increases in power.

At this point this method runs into a solid dead end. In order to get a simulation of the amount of light out of the system, it is necessary to state what the reflectance of the front facet is. For the system in question, there is no real “reflection amount” that can be pinned down—instead, there is just “some amount” of light reflected back on that particular section of the diode; that light may not necessarily come from that section so it is not possible at all to have such simple boundary conditions. Furthermore, it is important to look at the time dependence of how the modes lock together: there may otherwise be other eigenmodes which are competing, which would go unnoticed in the single-state

simulation such as this. Finally, it would be very difficult to actually calculate what the light output as a function of position along the laser would be—this is because of the inherent non-linearity in the exponential term. Because of this, it would be difficult, if not impossible, to apply the linear algebra methods which are otherwise used elsewhere in the system; instead, a calculation would have to “settle down” on a solution, meaning it starts with some initial value, then moves to a numerically steady-state solution.

A time-dependent solution had its own difficulties as well. As mentioned above, including time dependence in the intensity rate equations includes a factor of the speed of light—and light is very fast. Thus, the timestep required was now on the order of 10^{-15} seconds, which would require millions of cycles in order to replicate the 10^{-9} sec risetime of the laser-- numerically infeasible. A time dependent numeric simulation was still attempted once more with the 1D case (here including a dimension along the propagation direction), with a few tricks, including dynamic timestepping, where if any value was increasing or decreasing too quickly the program would slow down the timestep. However, it was not possible to get around the discontinuity that occurred at the lasing threshold; the simulation always diverged. This then meant that the more general calculation, where the laser now had some width and light was reflected by an optical system instead of simple mirrors, would also fail. This was one of the failings of this project; the inability to simulate the actual output of the laser diode stack. It was still possible to simply characterize the external optical system as a laser resonator to see if that part would work, whilst assuming one of the mirrors simply had a reflectance greater than one—analogue to a diode laser stack.

Laser Resonators

The requirement for a laser is to have a positive gain medium. How is it possible to construct a device that takes advantage of this effect to have light output? Imagine a device like that shown in figure 3. The region of positive gain material is surrounded on two sides by mirrors. One is partially reflective, meaning that a fraction of light is transmitted through the material whilst the other part is reflected




Figure 3: a generalized laser resonator.

back into the cavity. Now imagine a packet of light particles that rattles back and forth in this cavity. A single photon generated from spontaneous emission is emitted on one side of the gain region, and propagates backwards through the medium, increasing in intensity. It is then reflected by the fully reflective mirror and passes through the medium again, once more increasing in intensity. Finally, the packet of photons reflects off of the partially reflective mirror. Some are transmitted out of the cavity, but others are reflected back in, to be amplified by the gain region. The laser has reached a steady-state operation if the returning photon packet exactly matches the packet of the previous pass—in other words; it “resonates” back on itself, in a similar manner to a flute resonating at certain audio

frequencies. However, unlike an audio resonator which have a small ($<3\lambda$, normally) number of wavelengths within the length of a resonator, a laser may have tens of thousands of wavelengths.

This is an oversimplified view of a resonator, because it only acknowledges the picture in one dimension. A general resonator has curved mirrors, or optics that interrelate different spatial parts of the device with one another, and in this way, the resonator needs to be looked at in its cohesive whole.

Suppose there is a generalized laser resonator, with no gain material in the center. This is a simplifying assumption; everything gets more difficult to calculate if a localized gain region is included, because now beam propagation calculations also need to account for this region, and the amount of gain may respond to the light intensity in that region (a phenomenon called “bleaching”). It is possible to calculate, using methods explained in the next section, the propagation of the configuration of light on one mirror (or any surface in the resonator) through the entire system, back onto that surface. A steady-state of the system, which means the system is no longer evolving in time, is when the light pattern that leaves the surface is exactly the same as that which arrives back at that surface, modulo some real number between zero and one. This is a re-iteration of the eigenvalue problem:

$$\hat{P}I(x, y) = \lambda I(x, y)$$

Here, $I(x, y)$ is a complex function of the position on the surface describing the amplitude and phase of the light field at that point. $|I(x, y)|^2$ then describes the intensity of light (in units of Power per unit area). \hat{P} is some operator which describes how light propagates through the system—for example, a Gaussian beam of some width looks like some other function (say, two Gaussians next to each other) after it is sent through the optical system—this operator would then describe the shift in functions. The value λ , as mentioned above, describes the amount of power lost in the system— $\lambda=0$ is equivalent to complete loss, and $\lambda=1$ equivalent to no losses. There can also be some complex phase to λ as well, but this is relatively unimportant-- this means that the real wavelength of light that is propagated through the system differs by some very small amount; a phase of π simply means that λ needs to change by $\frac{\lambda}{2L}$, where L is the total length of the cavity. Generally, the wavelength can increase by $\frac{n\lambda}{L}$ and the equations would stay the same—this is because that is the same as adding a factor of $2\pi n$ to the phase. One can imagine this as picking $N=30,000$ cycles across the resonator (“the 30,000th harmonic”) being indistinguishable from picking $N=30,001$ cycles (“the 30,001th harmonic”). This differs from musical instruments, which operate on the fundamental or lowest harmonics ($N<10$, normally).

So why is this important? Well, it is possible to solve the above equation for particular functions of $I(x, y)$ and λ if the form of \hat{P} is well known. These solutions, of which there may be an infinite amount, form the eigenstates and eigenvalues of a cavity. A well-known postulate of laser physics is then that any laser resonator will generally pick the eigenmode that has the lowest losses.

In the case of a perfect laser, only the lowest-loss eigenmode will be selected. However, for a general system, there is always spontaneous emission which acts as “noise” which probes all eigenmodes equally. Thus, there are always higher-loss eigenmodes, which can be imagined to be driven by the spontaneous emission in the system, at amplitudes inversely proportional to the loss in the

system. Thus, by calculating the eigenmodes of a resonator, it is reasonably possible to calculate what the output of the laser looks like.

Beam Propagation Methods

An important part of numerical optics is calculating how light travels and propagates through an optical system. Without any mathematical way of describing the system, it is not possible to simulate and predict anything about the system—which is, without getting philosophical, the basic goal of physics. Luckily, it is possible to simulate the propagation of light through a system. This section will derive the equations used for optical propagation, and then explain the assumptions made in order to simplify the mathematics to become solvable.

The equations that describe all of electricity and magnetism are Maxwell's equations. A well known fact is that by applying the equations together it is possible to derive a wave equation—in other words, self-sustaining, propagating disturbances in the electromagnetic field: Photons. The wave equation can be further simplified to cut out time dependence by using the Helmholtz equation:

$$(\nabla^2 + k^2)E = f(x)$$

Here, E is the electric field vector, k is the wave number, and $f(x)$ is some source function. This can be solved via way of Green's functions to solve for E in all space. Generally, a Green's function attempts to solve the equation of the form:

$$\hat{P}X = f(x)$$

Where \hat{P} is some differential operator, X is the function to be solved for over the space of the differential equation, and $f(x)$ is the source function. One can obviously draw the parallel between this and the Helmholtz equation, if $\hat{P} = (\nabla^2 + k^2)$. A Green's function is a solution of the form:

$$\hat{P}G(x) = \delta(x)$$

Where $G(x)$ is the Green's function, and $\delta(x)$ is the delta function. Thus, the solution of X becomes the convolution over all possible delta functions:

$$X(\vec{y}) = \int G(\vec{y} - \vec{x})f(\vec{x})d^3x$$

The Green's function solution to the Helmholtz equation in three dimensions is:

$$G(\vec{x}) = \frac{\exp(ik|\vec{x}|)}{|\vec{x}|}$$

So the solution to the intensity of light at some point due to some light input becomes:

$$X(\vec{y}) = \int \frac{\exp(ik|\vec{y} - \vec{x}|)}{|\vec{y} - \vec{x}|} f(\vec{x})d^3x$$

This expression can be further simplified by considering the source function along a 2D surface instead of a generalized 3D space. This is because light does not “come from nothing”; instead, it comes into the system from an external source (such as a laser) and thus the source can be described as a light source coming in from one side. Thus, our integral becomes:

$$I(\vec{x}') = \iint_{\text{surface at } z=0} I_{in}(x, y) \frac{\exp\left(ik \left|\vec{x}' - \vec{x}\right|\right)}{\left|\vec{x}' - \vec{x}\right|} dx dy$$

This is equivalent to the Huygens principle, which is that every point on the surface of a wave front is the source of a spherical wave. Here, imagine that the surface is a wavefront in space, so $I_{in}(x, y)$ is a constant phase, with some slowly-varying amplitude variation. If we look at only one spot on that wavefront, the Green’s function term now looks like the propagation of a spherical wave—we have a factor of $\left|\vec{x}' - \vec{x}\right|$, which is precisely the distance away from the spot. The exponential component then describes the phase of that spherical wave. Also, remember that the intensity out is $|I(\vec{x}')|^2$, so our spherical wave then drops off with $\left|\vec{x}' - \vec{x}\right|^{-2}$, as it should using an argument involving conservation of energy flux.

In practice, this integral is difficult to analytically and numerically calculate—there is a component of the Euclidian distance between two points, which now involves square roots, as well as a rapidly oscillating complex exponential term. In order to actually analytically or numerically solve this integral, a few simplifying assumptions are made.

The first assumption involves the reduction of dimension. In this integral, there are two dimensions to integrate over, corresponding to the surface integral over the source function. However, things become considerably simpler computationally if the integral is now over one dimension less: Imagine a sum over a 1:10,000 element vector (1D) versus a 10,000:10,000 element matrix (2D); everything becomes very computationally expensive, very quickly. This integration can be simplified as follows.

Suppose the intensity function now goes as $I_{in}(x, y) = I'_{in}(x)I''_{in}(y)$, so the intensity in the x and y directions are uncorrelated and separable. Next, suppose we want to calculate the output $I(x, y)$ only for when $y = 0$. Thus, it is now possible to do this integral in two separate parts: one for X, and one for Y. After integrating over Y for some function $I''_{in}(y)$, we obtain the integral of the form:

$$I(x', z') = F(z') \int I_{in}(x) \frac{\exp(ik|x' - x|)}{\left((x' - x)^2 + z'^2\right)^{\frac{n}{2}}} dx$$

Here, the factor of $F(z')$ is a focusing factor from the y direction, which comes from the phase modulation of $I''_{in}(y)$. The exponential factor of n ranges between ½ and 1 and comes from the amplitude modulation of $I''_{in}(y)$. There are two trivial functions of $I''_{in}(y)$: the first is $I''_{in}(y) = 1$, which corresponds to the light being infinite in extent in the Y direction. Doing this integral in Y yields $F(z') = 1$ and $n = \frac{1}{2}$. The other function is $I''_{in}(y) = \delta(y)$, which corresponds to a small localized line of

light, yielding $F(z') = 1$ and $n = 1$. For this project, the function was normally the former, so the general integral is of the form:

$$I(x', z') = \int I_{in}(x) \frac{\exp(ik|x' - x|)}{((x' - x)^2 + z'^2)^{\frac{1}{4}}} dx$$

This by inspection also has the same form as spherical wavefronts, except instead now they are cylinders, which can be motivated using an argument of conservation of energy with gauss's law.

The second assumption is one to make a step from analytic calculations to numeric calculations. Numeric calculations, in general, are much more soluble to simulation, because there does not necessarily need to be any sort of closed-form solution for any general system. Furthermore, there is a lot of existing architecture, mainly in the form of programs like MATLAB, which are happy to crunch through numbers to get an approximate result. Analytically, an integral can be approximated as a sum over finite elements, or for some specific x' value:

$$I(x', z') = \frac{1}{N} \sum_{i=1}^N I_{in}(x_i) \frac{\exp(ik|x' - x_i|)}{((x' - x_i)^2 + z'^2)^{\frac{1}{4}}}$$

It can now be recognized that this is equivalent to a dot product of 2 vectors of dimension N. Furthermore, the x' part can be pixilated along a j index, so we have:

$$I(x'_j) = \sum_{i=1}^N I_{in}(x_i) M_{ij}$$

Where:

$$M_{ij} = \frac{\exp(ik|x'_j - x_i|)}{((x'_j - x_i)^2 + z'^2)^{\frac{1}{4}}}$$

Is a matrix of dimension N by N. Thus, it is possible to calculate the intensity distribution of light after propagating through free space by a certain amount z' by simply multiplying the input light into the system by some matrix M_{ij} .

Applying Matrix Propagations

Having a matrix representation of the propagation of light is very useful, because it is possible to easily simulate the propagation of light through free space. Furthermore, the beam propagation does not depend at all on what the input light pattern is—it works for any generalized input. Because of this, it becomes very useful for applications of calculating eigenmodes of a resonator, as discussed in the previous section. First this section will demonstrate some applications for this matrix method as well as a discussion of numeric artifacts and errors.

Because this method is numeric, it comes with inherent numeric artifacts. The largest numeric artifact is that of near-field non-physical diffraction due to under-sampling of the input light. Inherent to the calculation is that the source is not continuous—instead it is formed from a multitude of point sources, analogous to the continuous input beam going through a diffraction grating, with a slit width of $\frac{D}{N}$, where D is the total lateral distance in the simulation (normally about 5,000 to 10,000 microns) and N is the number of sample points (normally less than $\sim 2,000$ due to computational limitations). There is thus a first-order diffraction pattern that forms due to this numeric diffraction pattern, at $\sin(\theta) = \frac{\lambda N}{D}$, where λ is the wavelength, and $\frac{N}{D} = \frac{1}{\Delta D}$, the inverse of the slit width. One can see by inspection if $\frac{\lambda}{\Delta D} \geq 1$ then there is no diffraction pattern that forms; however, for physical systems with a reasonable width (say, 1mm wide) this would require over 10,000 elements—computationally expensive, at least for the machine that was used (a 4 year old Windows7 i3 processor with 4GB RAM). This then means that there are non-physical power losses within the system and nonphysical effects in the extreme near-field—however, these are easily corrected for. By varying the number of elements used in a simulation for propagating a square beam of light through one meter, it was discovered that the non-physical power loss was almost exactly proportional to $\left(\frac{\lambda}{\Delta D}\right)^2$ -- so all that was necessary was to divide all outputs by this amount; a great “fudge factor”! Of course, it is possible to derive this proportionality analytically, but after simulation it was completely unnecessary.

One good way to showcase how this method works is to make images of how an input beam propagates through free space. To do this, there would be two axes: the propagation direction, and the lateral direction. The lateral direction would be the direction of the “screen” that would be put in place to probe what the light intensity is at that point. The propagation direction, of course, is self-explanatory. An image can then be formed by using propagation matrices to move light from the source to some imaginary screen at distance Z , for multiple (200 or so) Z 's corresponding to different pixels in the image. The color at each point is then $|I(x_i)|^2$, the light intensity at that point in units of power per area. Below is an example of this process in action. Here, a widened Gaussian shape of the function $\exp(-x^4)$ has introduced random smoothed phase variations between 0 and 2π (smoothing generated in MATLAB using a low-pass Gaussian FFT technique) which can be equivalent to shining a laser light through a slightly bumpy piece of glass. The light pattern, then instead of being a nice smooth Gaussian shape, forms splotches as it propagates down 1m of length.





Figure 4: Propagation of randomized light through space. Left is the intensity of light along the beam path, and above is the shape of the beam showing the randomized phase front.

Another cool effect that can be succinctly demonstrated with this method is the Talbot effect. The Talbot effect is a near field phenomenon in which an infinitely repeating pattern—for example, like that of a diffraction grating—repeats itself indefinitely at certain lengths (mod some integer n). This matrix method applies well to demonstrate this effect, because it is very good at showing near-field phenomenon (of course, assuming that it is further away than that of the non-physical diffraction from input pixilation). However, things become remarkably simpler if the source is just delta functions at some repeating distance. With this calculated and applying MATLAB liberally, the following image can be found showing the light-intensity field with 30 elements, shown to the above in figure 5. Here, one can see that at the Talbot distance (which is $\frac{2a^2}{\lambda}$, where a is the distance between emitters and λ is the wavelength), the approximate delta-function light intensities are re-produced back on themselves, and is re-produced shifted up by $\frac{a}{2}$ at half-integers of the Talbot distance. Generally speaking, there is a cool, almost fractal pattern of light.

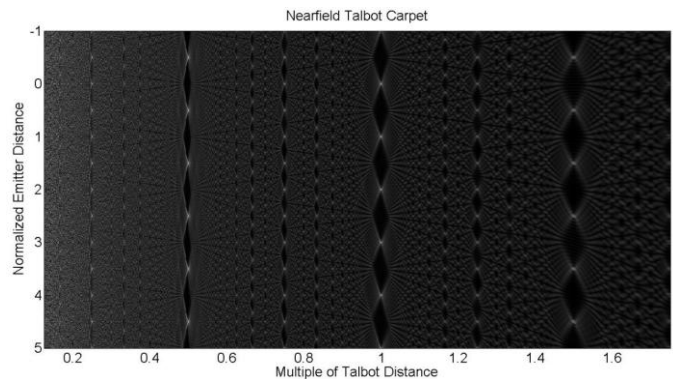


Figure 5: the Talbot Carpet. A pattern of delta-function light inputs is reproduced after the Talbot distance. Also note near-field effects due to under sampling on the left.

Another part of using matrices is the ability to append different parts of an optical system together, as mentioned previously. Imagine propagating light through some distance, then through a lens, then propagating down another distance, then through a piece of bumpy glass (a phase plate), and so on. The light on the front surface of the lens is exactly the input light times the propagation matrix through that distance. The light on the back surface of the lens is the same as the light on the front surface, but now with each point multiplied by some phase factor—equivalent to multiplying by a diagonal matrix with complex values whose amplitude is one, and its phase is varied slowly with the

functional form of $\theta(x) \sim x^2$. This can keep on going, so the light out at the other end of the system is exactly the light into the system, multiplied by a chain of matrices each describing a piece of propagation through the system, with the right-most matrix being the first propagation and the left-most matrix being the last propagation. This chain of matrices can be multiplied together to make a single compound matrix which describes the response of light through the entire system, which now can be some very complex entity if: for example, non-spherical mirrors, generalized lenses, and the like. An example propagation matrix is shown below in figure 6: in this case the system was actually described incorrectly (two mirrors facing each other, but instead of being spherical as intended the function describing them made them non-spherical). This matrix is complex-valued, so for the image the absolute value was taken so a color scale could be obtained.

An additional benefit of matrix calculations is the ability to leverage existing numeric infrastructure to analyze matrices, like that used in programs such as MATLAB. As is mentioned in the laser resonator section, the analysis of a resonator involves finding the eigenmodes and eigenvalues of the equation:

$$\hat{P}I(x, y) = \lambda I(x, y)$$

The entire derivation above is for describing the operator \hat{P} in terms of matrices, so now the function $I(x, y)$ becomes an N-dimensional vector and the eigenequation becomes the standard equation for finding the eigenvalues and eigenvectors of some N by N matrix!




Figure 6: an example compound optical propagation matrix. Light out is found by right-multiplying a vector of light in the normal way.

The Experimental Setup: Background

At this point the work shifts into the second part of the work done in this thesis—the experimental aspect. This section will explain the theoretical principle surrounding the experimental setup and explain what each of the components do.

The device, in its theoretical form, is rather simple. It is a long-focal-length lens, a spatial filter, and a spherical mirror which form an external resonator for a diode laser stack and beamsplitter, in a configuration as shown below.




Figure 7: Block diagram for the experimental optical setup.

First, the lens. The laser and spatial filter are both put at one focal length away from the lens; at this distance, the dependence on the lateral position does not matter, so the spatial filter “sees” every small diode in the laser stack at the same position. A focusing lens like the one used here has a useful property that is used here—the image formed at a distance of one focal length away from the lens is exactly the Fourier transform of the incident light pattern on the lens. This comes directly out of the calculations for beam propagation, with the correct approximations.

Next, the spatial filter. A spatial filter is a device that blocks light in certain places, and lets light through in others. A good example of a spatial filter is a single-slit diffraction plate, which can be as simple as a thin line (50microns, say) cut in a piece of paper. Light is allowed through the hole in the paper, but is blocked everywhere else. Another example would be a thin needle; except here, all light around a thin line is allowed, but light on the needle is blocked.

Now, these two components—the lens and the spatial filter—can be combined to form a filter for incoming light. This filter would be analogous to highpass/lowpass filters in electronics, which respectively damp out low frequency signals and high frequency signals. However, here instead of having frequencies in time, there are instead frequencies in space—the space component being the lateral displacement on the surface of the lens. A lowpass filter (low frequencies) would be equivalent to a thin line in a piece of paper, because all high spatial frequencies are displaced from the center and are thus blocked by the paper. A highpass filter, equivalently, would be the needle, because all low spatial frequencies are around the center of the distribution, which would then be blocked off by the needle. For this project, a low-pass spatial filter was used, which physically was some small distance between two glass prisms.

Next, the spherical reflector acts as the back reflector for the resonator system, feeding back light back into the laser itself. It is a spherical mirror, with the radius of curvature exactly the same as the distance away from the spatial filter, which means that the light out of the other end of the spatial filter will exactly match the light reflected back into the spatial filter. However, this component required very fine accuracy in order for the resonator to work: any small angular deviation would cause the system to become unstable (photons could bounce out of the system or generally be lost).

The last component, the beam splitter, is what is used to output light from the system. A beamsplitter is generally a partially reflecting mirror at some wavelength. Thus, a significant fraction of the light power through the component ($\sim 50\%$) is reflected, while the rest is transmitted. Because the device inherently reflects from both sides, the returning light would also be reflected, yielding a feedback amplitude of at most $R(R - 1)$, where R is the reflection coefficient ($R = 0.5$ for 50% reflection).

The diode laser stack has been studied in detail in other parts of this paper. It should be mentioned that for the actual experimental setup, there was one fully-reflective mirror on the back side, but also a slightly reflective mirror on the front side (around $R = 0.1$, but the laser system is proprietary so the exact value was unknown). This reflection means that the external feedback has to compete with this internal reflection and requires the input light to be on the same order of magnitude (so at least 10% of light out of the diode comes back in).

The entire device as a whole acts like a mode selector—which, in the end, is the whole goal of this work: to have a laser select the mode in which all individual diodes are lasing together. Generally, there are two modes for the laser to act in: one in which each individual emitter is lasing on its own, and one in which everything is lasing together. Now, it is possible to use the general principle of laser resonators: a laser will select the lowest-loss mode. All that is then required is to make the “together mode” lower loss than the “individual mode”.

Suppose (incorrectly, but correct to zeroth order) that the output beam of one individual diode is a square beam, which means that the fourier transform is a sinc function of some width inversely proportional to the width of the diode. Now, the spatial filter can be set to a width such that the sinc function diffraction pattern from a single emitter loses almost all of its energy due to its small width, but almost all of the power from multi-emitter operation is fed through, because now the entire diode array acts as a single diode of that width. This effect is shown in figure 8, where instead of one emitter to 50 emitters (as would be in reality) is one emitter to two emitters.




Figure 8: Light in the red region goes through the filter; all other light is lost. More light is lost from a single emitter than a compound (wider) emitter.

The Experimental Setup: Assembly

Now that the general theoretical background behind the experimental device has been explained, it is now relevant to introduce the physical device used for this project. In this section, an overview of the experimental device will be given, with specifics on actual assembly, calibration, and observations for future work.

The experimental device was built in Dr. Olof Echt's lab on a large optical table. The lens had a 1-meter focal length, and the back reflector had a radius of 55cm, which meant that the entire setup had a length of just over 2.5 meters. The optical table only had a length of 2 meters, which meant a flat mirror was used to turn the beam by 90 degrees in between the lens and laser itself. The assembly of the system took place with two phases. In the first phase, all of the components were put in place except for the diode laser stack, which was instead proxied to be a small diode laser pointer to calibrate the system. In the second phase, the small diode laser was replaced with a 100W 50 element diode laser stack operating at 795nm. Below are some images of the system while in phase one, with the small diode laser pointer.

Phase One: Calibration with a Laser Pointer



Figure 9: The Full optical system. Here, all components are visible, as explained in the other sections. The laser diode has a piece of paper in front with a small hole (1cm by 1cm) cut in it. This is so it becomes easy to view the back-feedback coming through the system. Also seen is the two angular control knobs for the flat mirror, which are used as the proxy for aiming the laser diode.




Figure 10a: The Back Reflector. Here, a small spherical mirror is mounted on an angular stage. Using two knobs on the back of the stage allows control of the angle in two directions, with accuracies of about 10^{-3} radians. Also note the red laser dot from the laser pointer, which is the light that has gone through the spatial filter and is now reflecting off. Technically, there should be no side reflections, but the spherical mirror available had a lot of scratches on its surface, which cause losses in the system. The flat mirror is also on a similar angular stage.




Figure 10b: The Spatial Filter. Here, two prisms were put close together to create a low-pass spatial filter with a small, controllable-width slit. The right-hand prism is movable by a micrometer with an accuracy of about 5microns. The left prism is held stationary; shifting the entire filter to the left or right is instead done by changing the angle of the flat mirror to project the center of the beam more to the left or right. Also note the small red dot on the center of the slits; this is the light from the laser pointer which is calibrated to go through the slit.




Figure 11: Feedback back on the laser system. This is a close-up of the screen in front of the laser diode, with the lights off for better resolution and the system properly calibrated. The diode laser outputs its light through a small hole cut in a piece of paper. There is also some slight reflection from the 1m lens, which has been purposely angled as to not overlap the other effects. Also clearly seen is the feedback from the optical system, which has been shifted up slightly as to not overlap with the diode laser output. The light takes the form of single-slit diffraction. This is because the light reflected in the back reflector goes through the spatial filter again, which acts like a single slit, creating the same kind of diffraction pattern. The fringe width does change inversely proportional to the slit width, as is predicted.

Now there is the question of calibrating the system. In total, there were 5 dependent degrees of freedom to the system, and 4 independent, coupled degrees of freedom. Each degree of freedom corresponds to a calibration knob or dial within the system: the horizontal and vertical angles of both of the mirrors, and the slit width; however, it is best to imagine the degrees of freedom in the basis described below.

The first degree of freedom was the vertical angling of the system—changing this would shift the back-reflection up and down past the diode laser. This could be done by either changing the vertical angle of the flat mirror or the back reflector. This calibration was relatively easy to do; a quarter-turn of the vertical control knob would be enough to shift the reflection from a position like that of the figure above directly onto the diode laser.

The second degree of freedom was the width of the spatial filter. This quantity could be changed by turning the dial on the micrometer, which shifts the right-hand prism. The micrometer can be controlled to about 5 microns, which corresponds to the smallest possible nudge to the dial. However, this change is coupled to the lateral displacement of the filter (movement to the left or right) because widening the slit width by Δx shifts the center of the slit by $\frac{\Delta x}{2}$.

The third degree of freedom is, as just mentioned, the lateral displacement of the spatial filter. Shifting this degree of freedom now moves the back-reflection pattern slightly to the left and right. This, along with the first degree of freedom, can be used to aim the back reflection directly on the laser diode. Because the left-hand prism is fixed in place, it is possible to shift the system instead by changing

the horizontal angle of the flat mirror. This then shifts the laser dot left and right on the spatial filter, which is equivalent to lateral shifts. This calibration is somewhat difficult to do; an infinitesimal shift (<1deg rotation on the knob) on the corresponding angular knob for the flat mirror is all that is required to shift the beam from on-axis to off-axis. Shifting this degree of freedom also requires an adjustment in the horizontal angle of the back reflector.

The final degree of freedom is the horizontal angle in the back reflector. This is arguably the most important calibration; this component reflects back the light from the small slit directly into the slit again, so any mis-calibration would cause the reflected light to be absorbed by the spatial filter. As such, this adjustment is very prone to error. Like that of the lateral displacement of the spatial filter, an infinitesimal shift is all that is required to go from all light going through the filter, to no light going through the filter. The physical calibration process can be done by looking at the refracted and reflected light from the prisms. When the light is only going through one prism, then there is one dominant dot from reflection. Through the other prism, there is another dot shifted by about 5 degrees. When calibrated, there are instead two dots; one for each side of the “edge” light reflecting through each prism.

Phase Two: Experimentation with a 100W diode stack

Once the system was thoroughly characterized with the laser pointer by finding the general error tolerances, the laser pointer was replaced by a 100W, 795nm diode stack. This laser was a surplus laser from Dr. Hersman’s lab, consisting of 50 individual diode elements, each 1micron tall and 100 microns wide, with 100 microns spacing between them for a total element width of 1cm. the fast axis (vertical axis) was collimated using a small lens on the surface of the emitters, meaning that the rapidly diverging light (with a width of about 20 degrees) was focused down to become less divergent (about 2 to 5 degrees). The laser was stable on its own, and had an approximately 10% reflective front mirror and 100% reflective back mirror, although the exact reflectance is unknown because the information is proprietary. Because of the power of the laser, it needed an external cooling system, which was done via water cooling. The threshold of the laser was approximately 25A of current, and most experiments were done with less than 30A of current, so the total power output of the laser was less than 10W, for safety reasons. Figure 12 below is an image of the laser whilst lasing at about 30A of current.




Figure 12: The Diode Laser. The power supply and water circulator/cooler are external to this image. This laser emits light in the near infrared regime, which is invisible to the human eye. However, to a standard CCD camera, infrared light shows up as a purple color; thus, the purple glow from the front of the laser is the incoherent emission from the diodes which is scattered through the slow-axis collimating lens.

Because the system was pre-calibrated, or at least almost entirely calibrated, it was relatively simple to re-calibrate the system to the new laser. However, there were a few things that did make the calibration difficult to do.

The first was the problem of being able to see what the light output of the system looked like. As before mentioned, the light from the laser was emitted in the infrared regime, which is invisible to the human eye. A CCD camera was required to be a “looking glass” to see what the light pattern looked like, because as mentioned in the caption above, a standard digital camera is capable of seeing near-infrared light. This was not really that big of a deal, however, because a cheap camera with a screen could be quickly used.

The second problem was that of beam quality for the free-running laser. The beam diverged at an angle of about 2 degrees in both directions, which meant that when the light arrived at the 1m lens a significant fraction of the power was lost due to it being outside of the 2.5cm-diameter element. This now meant that it was literally impossible to line up the center of the beam with the spatial filter, because unlike the laser pointer there was no specific “center” to look at, just a general blob.

Because of this generally poor beam quality, the back-reflection on the laser also had a significant portion of the light not actually reflect back onto the laser itself. Instead, most of the light that was reflected back was spread in the vertical direction because of the divergence in the vertical direction, which went uncorrected. This is shown to the right in figure 13.

Results and Conclusions

Now, at long last it is time to present the results of the thesis. In this section, the final findings of the spatial filter optical system will be presented; first, in the theoretical sense, and also in the experimental sense. In this section will also be a discussion of what did and did not work with the project and how to improve for the future.




Figure 13: Diode Laser with Feedback. The purple light below and above the diode stack is the back-reflected light from the optical system, with the spatial filter width at around 1mm. Clearly seen is that a lot of the light does not make it back into the system, because of the vertical-axis divergence.

Theoretical Conclusions

The most important calculation for simulating the spatial filter system was twofold; first, if there were stable, low-loss optical modes within the system, and second what the tolerances for misalignment within the system were. However, conveniently these questions could be answered simultaneously through methods of finding the eigenvalues and eigenmodes of the optical system like that introduced in the theory section. A theoretical system was assembled by constructing matrices that describe exactly that of the spatial filter system: a lens, spatial filter, and back reflector, all coupled by propagation matrices. In order to improve numeric resolution, a few changes were made in comparison to the physical system. The focal length of the mirror and back reflector were both set to 20 cm, 5 times less than that of the physical system. The width of the laser was taken to be 10mm, the same width as the actual system, but was assumed to be a single continuous laser, instead of 50 discrete sub-lasers. The wavelength was taken to be 0.795micron, the same as the physical system. The spatial filter was taken to be of a width of 250microns. The compound matrix then looks like that of below, with the following eigenvalues:




Figure 14a: The Transfer Matrix of the Spatial Filter System: Clearly seen is what would be a diffraction pattern impulse response.

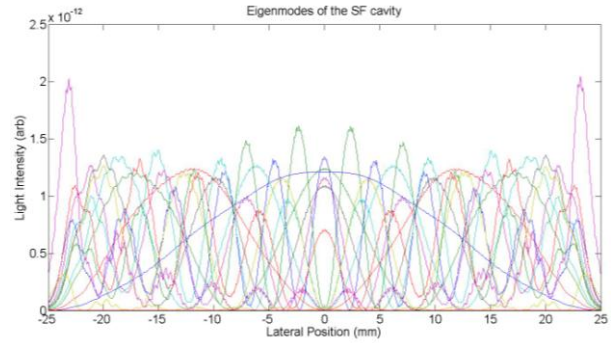


Figure 14b: Eigenvalues of the Spatial Filter Resonator. The intensity is each eigenmode times its respective eigenvalue; the exponent of 10^{-12} is an artifact of MATLAB's eigenvalues routine.

Thus, there are in fact stable modes which exist within the cavity. If there were not, then the eigenvalue routine would output what would look like simply random noise. This can be shown even with the same spatial filter cavity by looking at the logarithm of the intensity; There are the stable modes, which are at the top, then everything unstable is all around the same point at, in matlab's view, a power of about 10^{-53} .




Figure 15: Example of stable vs. unstable eigenmodes. Stable eigenmodes are at the top and have a specific pattern to them. The mass of lines in between 10^{-50} and 10^{-55} would be considered to be unstable eigenmodes, because their power is zero within numeric precision.

The second theoretical question to ask is what the region of stability of the resonator was in terms of how accurate each of the components has to be. For the spatial filter, there were four independent degrees of freedom, which were explained in the section above. However, two of those degrees of freedom, mainly the vertical displacement and spatial filter width, are irrelevant to inspection. The vertical displacement does not come into play with the numeric simulations in this project, because everything is strictly one-dimensional. The spatial filter width does not matter because all increasing the width does is let more or less power through the opening. Thus, there are two degrees of freedom to inspect: Lateral displacements of the spatial filter, and horizontal angular perturbations of the back reflector. The lateral displacements are relatively easy to simulate; this would correspond to shifting the diagonal row of ones to the left or right. The angular displacements would correspond to a row of mod-one complex phases along a diagonal that represent an angled shift in the beam from center (so a constantly increasing phase). Using the same parameters as the above simulation, the effects of mis-alignment are below in figure 16.



Figure 16a: Mis-Alignment of the back reflector: A slight mis-alignment can reduce the light through the system; a deflection of about 10^{-4} radians would cause a 20% loss through the system. For comparison, an angular deflection on this order would correspond to a shift of 1mm over 10 meters. This is doable with the experimental system.




Figure 16b: Mis-Alignment of the Spatial Filter. Here, the spatial filter is shifted off of center by some amount, while keeping all other values constant. Clearly seen is that once the side of the 250micron slit passes what was once the center the power through the system drops off almost immediately. This indicates that very fine accuracy is required for the flat mirror (10^{-4} radians)

These two theoretical results both suggest that very fine precision was required for the two most important degrees of freedom for the resonator. The back reflector, after these revelations, was improved via creating a long moment arm (10cm) coupled to a linear micrometer stage which had accuracy on the order of 10 microns. Thus, the angular resolution was then $10^{-5}/10^{-1}=10^{-4}$ radians, which was approximately what was required of the system. However, as mentioned above, it was very difficult to determine the center of the light distribution for the physical laser, which implies that hitting the center within 120 microns is effectively impossible; but not for lack of trying.

Experimental Conclusions

Even with the dismal outlooks as suggested by theory, the experimental side of things moved forwards. However, before anything else, it would be necessary to see if there was any sort of feedback response from the system to begin with. The visible effect of this would be that the raw output of the laser would shift by some amount when there was feedback, because whatever the external resonator did would reflect a significant fraction of the light back, causing mixing and thus a difference in the distribution. This could be done with a very simple test: looking at the output light through a beamsplitter, with just the free-running laser (so the back reflector would be blocked off) and with the back reflector uncovered and the spatial filter removed from the system. The beam shape was imaged using a standard CCD camera focused on a piece of paper about 20cm away in the path of the exit beam from a beamsplitter in the main beam. The two images were then shifted to exactly match each other (at least for lateral shifts) for comparison, using an X on the piece of paper. The results, combined with the difference image, are below:

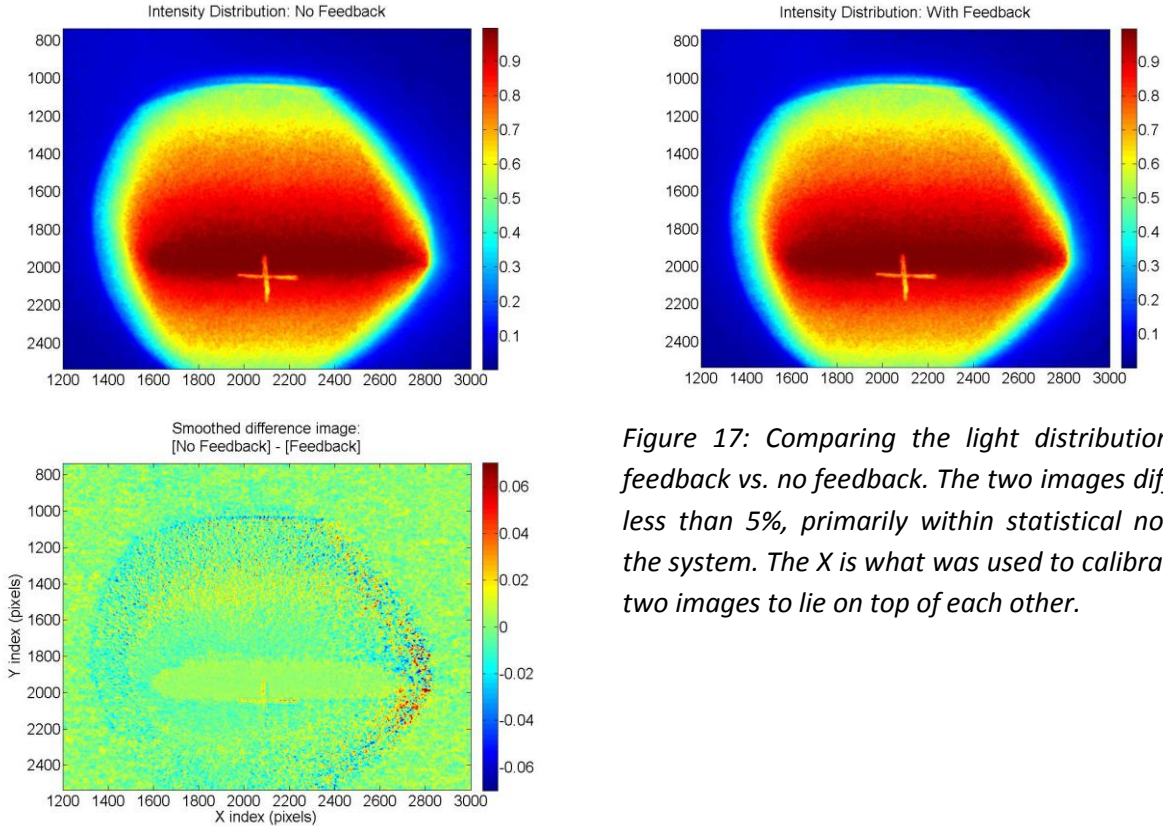


Figure 17: Comparing the light distributions for feedback vs. no feedback. The two images differ by less than 5%, primarily within statistical noise of the system. The X is what was used to calibrate the two images to lie on top of each other.

This result, even if quite preliminary, suggests strongly that there is simply not enough light getting back into the system to make any difference to the optical output, even when there is no spatial filter to diminish the feedback amount. This is primarily due to the before mentioned effect that the vertical axis is not collimated and thus a large fraction of the output power is not fed back into the diode itself. This, if anything, is the final indication that this effect does not work—at least, in this iteration. However, with future construction this might have the possibility of being fixed.

Future Work

Even with the dismal theoretical and experimental results found over the course of this project, there is still much room for improvement. The vertical divergence, for example, could be corrected by inserting a cylindrical lens within the system such that, compatible with other components, focuses all vertically diverging light back on the same line back on the laser. The angular accuracy of the components can also be improved by introducing long levers coupled to micrometer stages, as was mentioned in the sections above.

As another improvement, there is always the possibility for other types of external resonators. One such possibility, which was examined over the course of this work, is shown to the right. Here, the light from one diode reflects off both surfaces of a material, causing light to reflect to adjacent diodes, which creates mixing and thus preferential selection of the mode in which all diodes lase together.




Figure 18: Theoretical diagram for another type of external resonator.

Conclusion

In conclusion, much of the work done within this thesis suggests that the proposed system does not work. On the theoretical side, results show that the angular resolutions required for both the back spherical reflector and the flat mirror are below the physical tolerances of the system. Furthermore, the beam quality of the non-phase-locked diode laser stack made it impossible to determine the center of the light distribution, which is required in order to calibrate the flat mirror within good precision. As the last failure, the system did not compensate for the vertical divergence of the diode laser, which meant that much of the reflected power was lost due to it returning above or below the diode stack. However, many improvements can still be made, all of which move toward one of the eventual goals of high power laser physics: to have an efficient, high-quality, scalable, high-power laser system.

References

- [1] Jackson, D. *Classical Electrodynamics*. 3rd ed. New York: Wiley, 1962. Print.
- [2] Voelz, D. *Computational Fourier Optics: A MATLAB Tutorial*. Bellingham, WA: SPIE, 2011. Print.
- [3] Rutz, E. *Single laser beam of spatial coherence from an array of GaAs lasers: Freerunning mode*. *Applied Physics Letters*, Vol. 46, Issue 4552 (1975)
- [4] Rediker, R.H. *Operation of individual diode lasers as a coherent ensemble controlled by a spatial filter within an external cavity*. *Applied Physics Letters*, Vol. 46, Issue 133 (1985)
- [5] Yaeli, J. *Array Mode Selection utilizing an external cavity configuration*. *Applied Physics Letters*, Vol. 47, Issue 89 (1985)
- [6] Marciante, J. *Nonlinear Mechanisms of Filamentation in Broad Area Diode Lasers*. *IEEE Journal of Quantum Electronics*, Vol 32, Issue 4 (1996)
- [7] Partanen, H. *Spatial Coherence of Broad Area Diode Lasers*. *Applied Optics*, Vol. 52, Issue 14, pp. 3221-3228 (2013)

Acknowledgements

I would like to thank my advisor, Dr. Bill Hersman, for his intellectual support as well as supplying the components for the diode laser system. I would also like to thank Dr. Echt for allowing me to use his lab, as well as Dr. Jan Distelbrink for his assistance setting up components for the system.

Appendix

Table 1: Characteristic values for a diode laser		
Variable Name	Description	Normal Value
D	Diffusion constant	$33 \frac{cm^2}{sec}$
N	Conduction Band electron density	$10^{18} \frac{1}{cm^3}$
J	Current per unit area	$1000 \frac{A}{cm^2}$
q	Electron charge	$1.6 * 10^{-19} C$
d	Heterostructure thickness	200nm
τ_{nr}	Nonradiative recombination time	5nsec
B	Radiative recombination (spontaneous emission)	$1.4 * 10^{-10} \frac{cm^3}{sec}$
Γ	Confinement Factor	0.5
N_0	Critical electron density	$10^{18} \frac{1}{cm^3}$
$\hbar\omega$	Photon energy	$2.5 * 10^{-19} eV$
$ E ^2$	Light intensity	$10^8 \frac{W}{cm^2}$
E_f, E_r	Complex intensity of lasing	$10^4 \frac{\sqrt{W}}{cm}$
k	Wavenumber in the material	$2.2 * 10^6 cm^{-1}$
α	Linewidth Enhancement factor	0 to 7
α_{int}	Internal losses	$10 cm^{-1}$
n_2	Self-Focusing/Defocusing	$10^{-17} \frac{cm^2}{W}$