


THE LAMBDA PROJECT



EXPERIMENTAL GUIDE

May 21, 2004

Mike Maloney, Laboratory Director

Keith Nelson, Project Director

Written by Rebecca Slayton

Table of Contents

Guide to Presenting Experimental Work	3
I. Purpose	4
I.A. Introduction to very large scale integrated circuits (VLSI)	4
I.B. How are VLSI circuits manufactured?	6
I.C. How and why do manufacturers monitor the production of circuits?	7
I.D. How and why are manufacturing processes changing?	8
II. Theory	12
II.A. Introduction to acoustic waves	12
II.B. “Stiffness” in Matter: the Modulus	14
II.C. Longitudinal and Transverse Waves	17
II.D. Using surface acoustic waves to measure thin film properties	19
III. Methods	20
III.A. Introduction to Lasers	21
III.B. Generating acoustic waves with lasers: Interference	25
III.C. Detecting Acoustic Waves with Lasers: Diffraction	28
III.C.1. Amplitude Diffraction Gratings	29
III.C.2. Phase Diffraction Gratings: Transmission	33
III.C.3. Phase Diffraction Gratings: Reflection	35
III.D. Summary: Using lasers to generate and detect acoustic waves in thin films	36
IV. Analysis	39
IV.A. Typical Data and Interpretation	39
IV.B. Fourier Transforms	40
IV.C. The Acoustic Velocity	42
Summary of Appendices	43
Appendix I. Metrology for fabricating faster circuits	43
Appendix II. Using acoustic waves to depth profile film systems	43
Appendix III. Interference between crossed laser beams	43
Appendix IV. Physical properties of metals	44

Guide to Presenting Experimental Work

Any presentation of a scientific experiment should include the following elements:

1. Purpose: What do we want to learn from the experiment? Why does it matter?
2. Theory: How, physically, does the experiment give us access to what we want to learn?
3. Methods: What tools do we use to make the experiment work?
4. Analysis: How do we make sense of the data we've gathered?
5. Conclusions and future directions: What can we conclude on the basis of the data we've gathered? What future experiments would expand the scope of the conclusions?

The remainder of this manual will provide the information you need to answer these questions for the surface acoustic wave experiments in thin films.

I. Purpose


In this section, you should explain why semiconductor manufacturers are interested in studying thin metallic films, and what they're interested in knowing about them. Your presentation should address the following questions:

1. What is the structure of an integrated circuit? This includes drawing the structure, labeling the following parts, and explaining their function:
 - a. Conducting layers
 - b. Dielectric layers
2. What is the size scale of a modern circuit? Compare it to some common objects.
3. How do manufacturers typically make such circuits? This includes:
 - a. What is photolithography?
 - b. What steps need to be repeated in order to build up layers?
4. What material properties do manufacturers need to monitor, and why?
5. How is circuit fabrication changing, and why? This includes:
 - a. What materials are changing?
 - b. How are designs changing?
 - c. What problems are introduced by these changes?

I.A. Introduction to very large scale integrated circuits (VLSI)

Very large scale integrated (VLSI) circuits are a central component of modern computer chips. They speed up computer operations by making overall electrical circuits smaller and reducing the number of connections between them. Because current doesn't need to travel as far through the circuits, operations are repeated more quickly. In VLSI circuits, many small (local) circuits are connected upward into increasingly large (global) circuits. The very smallest feature we talk about in a computer chip is the transistor – this allows current to flow (or not), and is used to create small circuits. Today transistors are going as small as 20 nanometers! To give you a sense of scale, 1 nanometer (1×10^{-9} meters) is about 50,000 times smaller than the width of a human hair! In a silicon chip, 1


cubic nanometer contains about 50 atoms. To take an example from the life sciences, 2 nanometers is about the diameter of a strand of molecular DNA.¹



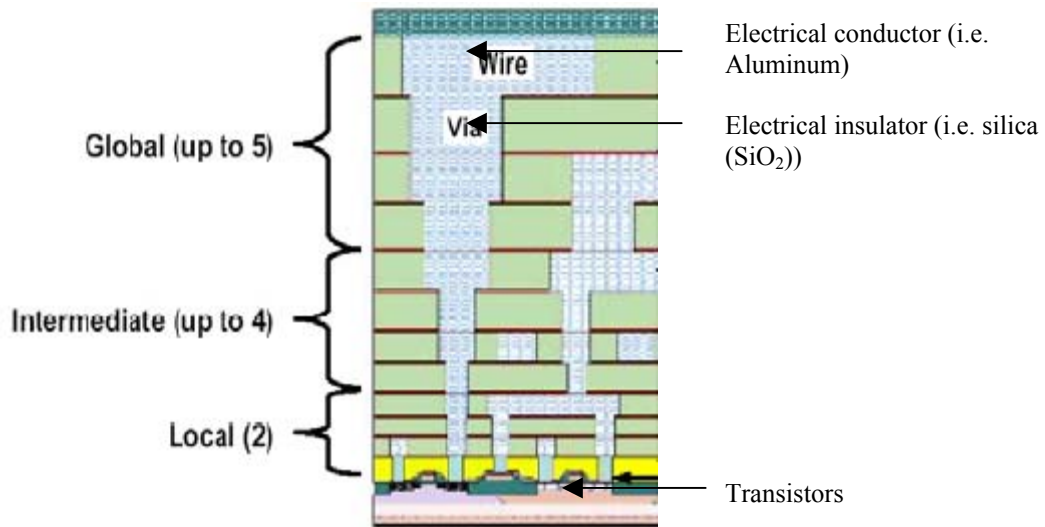
In a semiconductor chip, many transistors at the bottom of a stack make up small circuits that perform very small scale operations. These circuits are integrated into successively more complex circuits, performing more complex operations, by layering circuits on top of one another. Electricity flows between the different levels through interconnects made of a conducting metal such as aluminum. The dimensions of the interconnects vary, but are typically on the order of hundreds of nanometers. For example, the interconnects shown below in an Intel microprocessor range from 280 nanometers thick at the bottom of the stack, to 1,200 nanometers thick at the top.²

¹ Diagram and photo courtesy of http://www.chem.wisc.edu/courses/801/Spring00/Ch1_2.html

² This picture is borrowed from Intel: <ftp://download.intel.com/research/silicon/013microniedm2000.pdf>



An electrical insulator, or dielectric material, is used to electrically insulate the metallic interconnects. This is shown schematically in the figure below.



I.B. How are VLSI circuits manufactured?

Although precise processes vary greatly with the type of chip actually being produced, we can outline some of the major steps involved with fabrication. A relatively detailed description of the process can be seen at:

http://www.sematech.org/corporate/news/mfgproc/mfgproc.htm#steps1_2

Typically the process begins with the deposition of an insulator, silica, on a highly polished silicon wafer (see steps 1-3 on the Sematech site). Manufacturers then use several steps of photolithography and etching to create a complex design. In photolithography, a layer of light-sensitive material (called photoresist) is coated onto a material, and a mask pattern is used to selectively expose some areas of the material to light. The areas of photoresist that are exposed to light are hardened, while the non-exposed areas can be removed with a chemical solution. This process creates a protective pattern of photo-resist. (See step 4.)

In subsequent etching steps, chemicals are used to remove material from the unprotected region (See step 5.) These areas can then be filled with other materials, and the process can be repeated creating a complex design of uniform thickness (See step 6.) The flow of current through the chips can be controlled by changing the conductivity of some parts, using doping processes (See step 7). Finally, the local circuits are interconnected with conducting layers such as aluminum, to create more complex global circuits (See step 8). Aluminum interconnects are fabricated by first depositing a layer of aluminum, and then etching away portions using photolithography (steps 3-5, again, but this time for aluminum). The etched away portions are filled in with an electrically insulating material, creating isolated interconnects as shown in the figure above. By repeatedly depositing aluminum layers, etching portions away, and filling in the gaps, multiple circuits can be interconnected.

I.C. How and why do manufacturers monitor the production of circuits?

Every time a new layer of metallic film is deposited on a semiconductor chip, the manufacturers need to ensure that it has the appropriate thickness, uniformity, and grain size. If the metal is too thick in a particular dimension, electricity flows more slowly in that direction. This would slow the operation of the circuit. Similarly, the existence of many microscopic “grains” and other non-uniformities in the film can slow the flow of electrons through the metal. Thus, after every step of film deposition, photolithography, and etching, manufacturers need to make measurements to ensure that they’ve done it right. Often they use a “monitor wafer” to test the resistance with a four-point electrical probe. Sometimes the monitor wafer is scratched and examined with other tools to ensure

that it has the appropriate structures. However, since monitor wafers must be thrown away after they have been touched or destroyed through testing, manufacturers would prefer to monitor their products without actually touching them. The experiment that we are using in the outreach lab can accomplish this by using lasers to generate and detect acoustic waves traveling in the films.

I.D. How and why are manufacturing processes changing?

While the basic concepts and technologies required to produce such large structures are established, large scale circuits are also constantly changing as manufacturers try to make their circuits faster. In the past, computers have been made faster largely by decreasing transistor and circuit sizes, reducing the amount of time required for electricity to flow through circuit. To give you a sense of how decreased sizes have translated into faster clock speeds, take a look at the evolution of microprocessors at Intel (see Table I below).


Table I. Evolution of Microprocessors from Intel³

Name	Date	Transistors	Microns	Clock speed
8080	1974	6,000	6	2 MHz
8088	1979	29,000	3	5 MHz
80286	1982	134,000	1.5	6 MHz
80386	1985	275,000	1.5	16 MHz
80486	1989	1,200,000	1	25 MHz
Pentium	1993	3,100,000	0.8	60 MHz
Pentium II	1997	7,500,000	0.35	233 MHz
Pentium III	1999	9,500,000	0.25	450 MHz
Pentium 4	2000	42,000,000	0.18	1.5 GHz

As this table shows, more transistors can be fit into smaller chips, and the shorter delay times in circuits allows the computer to run with a faster clock speed. In other

³ Data taken from: <http://electronics.howstuffworks.com/microprocessor2.htm> .

words, faster computers have been made possible in large part because of shrinking transistor size, labeled feature size in the figure below.⁴



For comparison, a piece of paper is approximately 100 microns (μm), or $1/10^{\text{th}}$ of a millimeter thick. This means that the transistors in the mid-1980's were roughly $1/10^{\text{th}}$ the thickness of a piece of paper...already very small...and they have only gotten smaller! For example, in 2002, Intel announced a new technology involving transistor sizes as small as 50 nanometers. In other words, the minimum size is now 200 times smaller than what it was in the mid-1980's.⁵

But there are limits to how much can be gained by decreasing the size of transistors. At this point, the time delays associated with interconnects are becoming the significant contributors to overall circuit speed. Hence, some companies have begun to make interconnects faster by replacing aluminum with copper. Since copper has a lower resistance than aluminum, electrical current travels more quickly through the circuits, speeding up their operation.

However, the introduction of copper also makes the fabrication process more complex. The primary reason for this is that copper cannot yet be etched like aluminum; we don't know of chemicals that can etch copper away without causing other problems in

⁴ Taken from: <http://www.ece.cmu.edu/~maly/maly/DAC94.pdf>

the fabrication process. So instead, copper is added to trenches already etched in a dielectric material. The upshot of this is that the deposition of copper depends upon the patterns that it is deposited into. This is not true of aluminum – because it is deposited first, and etched later, the pattern doesn't affect the quality of the film deposition in any way. This makes some previous methods of testing – such as four point probes – impossible to use, because they must be used with a metallic layer that has not yet been patterned. Copper never goes through an “unpatterned” stage because it is deposited into already etched areas of a dielectric material.

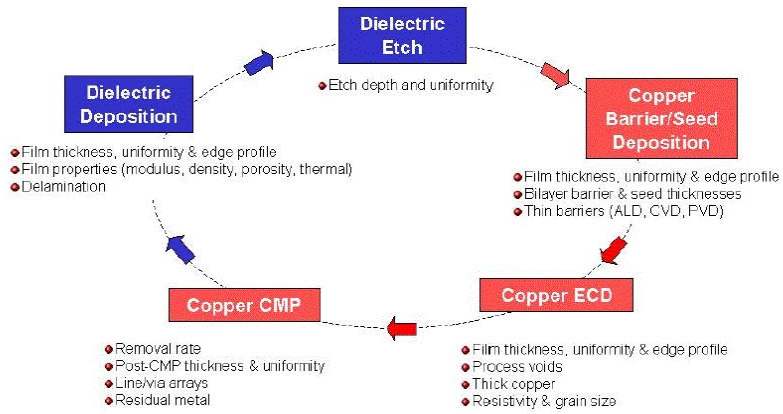
Copper interconnects are manufactured through at least five steps, known as the Damascene process:

1. The dielectric layer is deposited as a blanket layer.
2. The dielectric layer is etched using photolithography.
3. Copper barrier/seed deposition: a very thin barrier layer (to prevent the copper from diffusing into the surrounding material) and copper “seed” layer is deposited using a process known as physical vapor deposition.
4. Copper electrochemical deposition (ECD): This is a method of filling in the trenches through a chemical deposition process (done in a solution of copper ions).
5. Copper chemical-mechanical planarization (CMP): the copper layer is polished to a very high level.

The chart below shows some of the material properties that we want to be able to monitor during each of these steps. These properties include the thickness of the layers, the profile of the etched areas, their area, and the size of the grains grown in the metal. All of these properties can be monitored using surface acoustic waves in thin films.⁶ The next section will explain a little bit of why.

⁵ For more information on Intel's new technology, see their press release at: <http://support.intel.com/pressroom/archive/releases/20020813tech.htm>

⁶ Figure courtesy of Philips Advanced Metrology Systems.



II. Theory


Equation Section 2 In this section you should explain what an acoustic wave is, and why it can give us information about the thickness, density, and stiffness of materials. Your presentation should address the following questions:

1. What is an acoustic wave?
2. What does the acoustic frequency describe?
3. What is the mathematical relationship between the length of a traveling wave, its speed, and frequency?
4. How does the stiffness and density of a material affect the speed at which sound travels through it? This includes:
 - a. Explain what the modulus is (i.e. how do we quantify the “stiffness” of a material?)
 - b. Explain different types of material “stiffness”, and how they relate to different types of acoustic waves.
 - c. What is the mathematical relationship between stiffness, velocity, and density?
5. Why does the speed of a surface acoustic wave change when we vary the thickness of a metal film?

II.A. Introduction to acoustic waves

An acoustic wave is a periodic, or regularly recurring, displacement in matter. We quantify the speed at which the displacement recurs in terms of a frequency; the number of times that it recurs each second is the frequency in units of Hertz (Hz). For example when a musical instrument produces a single note, it is producing a regularly recurring pressure variation in air. To produce a musical note A, a string on a guitar must vibrate at 440 Hertz (Hz, or 440 times a second). The normal human ear can hear sounds as low as 20 Hz, and as high as 20 thousand Hz. It is also sensitive enough to tell the difference between 440 and 441 Hz – if the string vibrates a little too fast, the pitch “sounds” too sharp. The ear can detect sound that is only 2×10^{-9} times the pressure of air.

When waves travel, they carry energy, but not matter (unless we're talking about surfers, but that's a different story). In an idealized traveling wave, matter itself does not actually travel anywhere – it moves periodically, but inevitably comes back to where it began. To see this, consider the motion of a float sitting on top of a wavy pond. It makes a transverse motion, up and down, but usually does not travel much of anywhere.



In this case, the time for the float to complete one cycle, T , is equal to the time for a wave to advance one wavelength, λ , while traveling at a velocity, v . This is just another way of saying that the distance it travels is equal to its velocity multiplied by the amount of time it is moving, something that is true for all moving objects. In the case of a traveling wave we write this as $\lambda = vT$. If we substitute in the relationship between period and frequency, we have the wave relationship:

$$v = f\lambda \quad (2.1)$$


This relationship shows that as the frequency of the wave gets higher, it speeds up.

Again, it should be emphasized that the wave carries energy, but matter itself does not move anywhere. If you don't believe that the water molecules in this wave are not moving any more than the float, look at the individual particles in the animated waves at

<http://www.gmi.edu/~drussell/Demos/waves/wavemotion.html> How far do any of them actually travel?

II.B. “Stiffness” in Matter: the Modulus

All materials have a quantifiable “stiffness,” much like the stiffness of a spring. You may be familiar with the behavior of springs from classical mechanics. Classical mechanics tells us that if we pull a ball on a spring out of its resting, or equilibrium position, it oscillates with a natural frequency f . This is shown in the picture below.



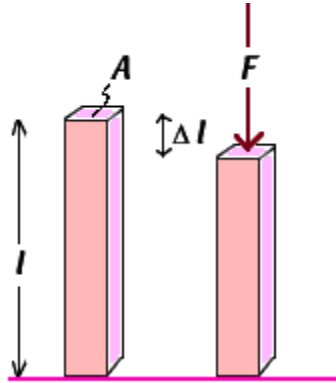
Classical mechanics also tells us that the natural frequency of the spring increases with its stiffness, which we quantify in terms of a spring constant, k (i.e. if k is larger, the spring is harder to move, or stiffer.) On the other hand, if the ball becomes heavier, the natural frequency decreases. The precise mathematical relationship between the stiffness of the spring, the mass of the ball, and the natural frequency is shown below.

$$f = \frac{1}{2\pi} \sqrt{\frac{k}{m}} \quad (2.2)$$

Solids are in many ways like many tiny atoms or molecules connected by springs; as a result, we'll see some direct analogies between the natural frequency of a ball and spring, and the frequency of acoustic waves in matter.

But when we want to specify how stiff a material is, we talk about its *modulus* rather than its spring constant. The modulus is simply a measure of how much stress we must apply to a material in order to strain it. The “stiffer” the material is, the higher the

modulus is. However, there are many different ways of compressing matter, and so there are many different types of moduli. For example, if we apply some pressure, or stress (force per unit area) to a rod, it will compress a little bit in the direction we're pushing.⁷



We call Δl the displacement, and strain is defined as the displacement per unit length, or $\Delta l/l_0$. Young's modulus is defined as the ratio between the stress we apply and the strain on the material:

$$Y = \frac{\text{compressive stress}}{\text{strain}} = \frac{F/A}{\Delta l/l_0} \quad (2.3)$$

Young's modulus for several materials is shown in Table II.1. We typically measure pressure in terms of Newtons per square meter, also defined as a Pascal.

Table II.1. Young's modulus for several common materials⁸

Material	Young's modulus Units of 10^9 Pa, or N/m^2
Rubber	7.0×10^{-4}
Bone	9
Wood	13
Concrete	30
Glass	65
Aluminum	70
Steel	200

⁷ <http://www.uvi.edu/Physics/SCI3xxWeb/Structure/ShearStress.html>

⁸ <http://hyperphysics.phy-astr.gsu.edu/hbase/hframe.html>

For example, if you weigh 132 pounds (60 kilograms), then standing on earth (where gravity constantly accelerates you downward at $a=9.8 \text{ m/s}^2$), then you exert a force of


$$F = ma = 130 \text{ kg} \times 9.8 \text{ m/s}^2 = 1300 \text{ Newtons}$$

If you distribute that force over an area of 10 cm^2 , then you exert a pressure of $1.3 \times 10^7 \text{ N/m}^2$, or Pascals. Now imagine standing on top of a concrete slab 10 cm thick, distributing your weight over only 10 cm^2 . Using the formula for Young's modulus, we know that the concrete will compress by an amount

$$\Delta l = \frac{F/A}{Y} l_0 = \frac{1.3 \times 10^7 \text{ Pa}}{30 \times 10^9 \text{ Pa}} \times 0.1 \text{ m} = 4.3 \times 10^{-5} \text{ meter}$$

or 43 microns.

Another way to deform matter is to apply a shear stress. Imagine putting a piece of tape on top of a block and pulling, so that it deforms as shown below.



The amount of force that it takes to move the top of the block relative to the bottom determines the shear modulus μ . This is defined by:

$$\mu = \frac{\text{shear stress}}{\text{strain}} = \frac{F/A}{\Delta l/h}$$

Table II.2 below shows the shear modulus for several common materials.

Table II.2 Shear modulus for several common materials⁹


Material	Shear modulus (Units of 10^9 Pa)
Rubber ¹⁰	1×10^{-3}

⁹ <http://hyperphysics.phy-astr.gsu.edu/hbase/hframe.html>

¹⁰ There is actually a range of values; here I use the low limit.


Glass	19-34
Aluminum	26
Steel	75-80

The bulk modulus turns out to be another very useful way of defining “stiffness” in matter. This is very similar to Young’s modulus, but we draw the picture slightly differently. The bulk modulus is defined as the amount of pressure required to produce a change in volume rather than a change in length.



II.C. Longitudinal and Transverse Waves

Most acoustic waves can be understood as combinations of two basic types: transverse and longitudinal. Transverse acoustic waves are slightly easier to visualize. These waves displace the material perpendicularly to the wavelength, or direction of travel.




The velocity of a shear wave traveling through a bulk material is determined by the shear modulus, μ , and the density of a material, ρ :

$$v_s = \sqrt{\frac{\mu}{\rho}} \quad (2.4)$$

In other words, the wave travels faster for stiff materials (high modulus) and slower for heavy ones (high density).


Longitudinal waves displace matter in the same direction that the sound is traveling. We can envision this by thinking of sending a similar wave along a slinky. Some regions of matter are compressed, while others are stretched.



The displacement caused by longitudinal waves is slightly more complicated, and must be described by a mixture of the bulk and shear modulus. The longitudinal modulus, $M = B + 4/3 \mu$ determines the velocity of the longitudinal wave by:


$$v_L = \sqrt{\frac{M}{\rho}} \quad (2.5)$$

Surface waves contain both longitudinal and shear components. You can see this by looking at the diagram below. The overall wave is moving left to right. The longitudinal component of the wave compresses some parts of the material in the same direction as the wave is traveling, while it is stretching other parts. But the transverse component of the wave makes some parts of the material go up, while other parts go down.



II.D. Using surface acoustic waves to measure thin film properties

As the picture above suggests, surface acoustic waves have some “depth” in the material they’re traveling through. When the wavelength is comparable to the thickness of the material it travels through, it begins to feel the material underneath it. For example, below you can see a diagram of a wave traveling through a thin film on a substrate.¹¹ The film on the left is not attached to the substrate, while the film on the right is. As you can see, the surface wave changes when it is tightly bound to the substrate.



This causes the velocity of the wave to change. As Equations (2.4) and (2.5) show, the velocity of an acoustic wave is higher for materials that are stiffer and “lighter” (less dense). Hence, when the wave is partly traveling in a stiffer and lighter material, the velocity rises. You can see this on the right hand side of the diagram above. This is why measuring the velocity of a surface wave can tell us about the depth of a film, how dense it is, and whether or not it is well attached to the layer beneath it.

¹¹ Figures taken from Dhar et. al., “Moduli determination in polyimide film bilayer systems: Prospects for depth profiling using impulsive stimulated thermal scattering.” *J. Appl. Phys.* 77 (9) 1 May 1995.

III. Methods


Equation Section 3

This section should describe how you generated and detected sound in thin metal films. You should be able to answer all of the following questions. In your presentation, you should briefly summarize answers to questions 2-4; you will not be able to go into full detail, but you should be prepared to field questions about details.

1. What is a laser? You should be able to explain:
 - a. What does each part of the word stand for?
 - b. What are the key differences between light produced by a laser and light produced by a light bulb or other ordinary light source?
 - c. You should be able to explain what each part of the acronym means.
2. How do we use lasers to generate an acoustic wave of a single wavelength? Explain:
 - a. Why does heating a metal film produce sound?
 - b. What is constructive and destructive interference?
 - c. How can we use constructive and destructive interference to create sound of a well-known wavelength?
 - d. What type of laser do we use to generate the acoustic wave? Specify the wavelength, whether it is continuous or pulsed, and why.
3. How do we detect surface acoustic waves with lasers? This includes:
 - a. What is diffraction?
 - b. What type of laser do we use to detect acoustic waves in our experiment? Specify the wavelength, whether it is continuous or pulsed, and why.
4. Draw a schematic diagram of the experiment and explain what each component does. You should trace the path of the laser beam through the following components:
 - a. The pump and probe focusing optics

- b. The transmission grating (phase mask) and spatial filter
- c. Imaging optics
- d. The sample
- e. The photodetector

III.A. Introduction to Lasers




As the diagram above shows, the acronym LASER stands for Light Amplification by Stimulated Emission of Radiation. As discussed below, there are essentially three parts to any laser:


1. A lasing medium (this would be between the mirrors in the diagram above) that can produce stimulated emission (more on this below). Stimulated emission is the primary reason why light produced by a laser is **coherent**.
2. Two mirrors that reflect light back and forth through the medium, as shown in grey above. This allows for us to use amplify the light produced by stimulated emission. This arrangement of mirrors tends to make light produced by a laser **collimated**.
3. A leaky mirror that allows some light out of cavity, as shown on the right.

First, a lasing medium is some substance that stores energy. It may be a gas, a liquid, a crystal. Any material contains electrons, and the electrons can be in different quantum energy states. When a material absorbs light, what is happening microscopically is that electrons are absorbing photons of light. The photon is destroyed as it gives its energy to the electron; the electron now has more energy than it had before, and we say

that it is in an excited energy state. After a short time the electron can lose the extra energy, sometimes by emitting a photon of the same energy as the one that it absorbed. You can see this as fluorescence, light that the material gives off. Usually the electrons just emit those photons by themselves, spontaneously. But if a photon of the right energy hits the excited electron, it can actually stimulate the electron to give up its energy in the form of a second photon. This is **stimulated emission**: one photon goes in, and two come out.




One important feature of stimulated emission is that the photons are released at the same time, coherently. This means that all of the crests and troughs in the light waves are lined up. This is very important, because if the waves and troughs do not line up, we get what's called destructive interference. Interference simply means that if you mix two waves they can either add up or cancel each other out. This is easy to see with pictures shown below.



On the left we have constructive interference – the wave crests and troughs line up, so the combined wave is larger. On the right we have destructive interference – meaning that the wave crests and troughs cancel each other, and we end up with no wave at all. The advantage of stimulated emission is that we get as much constructive interference as possible, because the light is released all at once.

In addition to giving us coherent photons, the stimulated emission process allows us to amplify light. This can only happen if the lasing medium contains a population inversion. This means that more than half of the electrons in the material have jumped up from a lower energy level to an excited level. The population inversion can be created in a number of ways – by applying a voltage to a diode, or by shining a bright light on the material. If the excited state contains more electrons than the lower energy level, then photons passing through the material are amplified significantly. Each photon that gets released can stimulate the release of another photon, so that a whole cascade of photons emerges.



Very well then, we have some amplification due to stimulated emission from a population inversion. But this is still not enough to make a really powerful laser – to do this, we repeat the amplification process many times. This is the purpose of the laser cavity – two mirrors make this cascade of photons bounce back and forth through the lasing medium, and they are amplified every time. The lasing medium is continually pumped up with energy so that every time the photons pass through the medium, they are additionally amplified. Finally, one of the mirrors must be a little bit leaky so that some of the energy can escape from the cavity.

After all of this amplification, however, a relatively high power laser may only produce 20 Watts of power – whereas a typical light bulb will easily produce closer to 100 Watts of power! Most of the energy that goes into producing a population inversion gets wasted anyway; only a small fraction winds up in the laser beam. So why go through all this trouble? The light produced by a laser is unique in three ways. First, it is coherent, unlike the light produced by a light bulb or a neon sign. This is part of the reason that lasers can be so useful in surgery and other applications. Marching in lock-step, the energy of the light can cause matter to move in coordinated, specific ways at the atomic or molecular level. The energy from a light bulb is too disorganized for this – it can warm things up, but it can't exercise any coherent control over the matter.

Second, laser light is nearly monochromatic – it contains only one wavelength. (Strictly speaking this is not true; especially in pulsed lasers a small range of wavelengths is produced. However even this pulsed light is close to monochromatic.) Lasers are monochromatic because the photons are all produced by the same quantum transition in the lasing medium; thus they all have the same energy. Other sources of light, such as neon signs, are also monochromatic because they are created from a quantum transition.


However, they are not also coherent – the light is produced by spontaneous rather than stimulated emission, so the emitted photons don't all have the same phase.

Third, laser beams tend to be collimated, meaning that all of the light waves are traveling in a nearly straight line. The two mirrors that bounce it back and forth through the lasing medium create this straight-line path. If the mirrors weren't aligned properly, such that the reflected beam began to wander off of this path, it would stop being amplified. The collimated nature of lasers makes them easier to work with because they can be directed long distances and focused to smaller regions.


III.B. Generating acoustic waves with lasers: Interference

There are plenty of ways to generate acoustic waves in matter. The simplest way is to give the material a “push.” For example, striking a drum or a gong produces a set of distinctive waves that your ear can recognize. Similarly, if you drop a rock into a lake, you produce a ripple that is the sum of many simple sine waves. Another way to generate acoustic waves is to heat a material very suddenly. This is how lightning creates thunder; it heats a region of air that expands very quickly. We do something very similar when we heat a metal with a laser. The metal expands very rapidly, launching what we call a wave packet: many different frequencies bunched together in space. To generate an acoustic wave of a well-specified wavelength, or frequency, we use constructive or destructive interference between two laser beams.

As described above, light waves can interfere either destructively or constructively. A single beam of laser light consists of light waves that are in phase with one another – all the crests and troughs line up, so interference is only constructive. We can represent a single laser beam as a series of lines representing wave crests. The spacing between the crests corresponds to the wavelength of light, λ . Then a single laser beam traveling toward the bottom of this page would be represented as shown below:



If we cross two laser beams, at an angle, some of the crests and troughs will not line up – they will cancel. This means that crossing two beams of light creates an interference pattern, or regions of dark and light intensity. You can see this in the picture below.



The spacing of the interference pattern Λ depends on both the wavelength of the light, λ , and the angle at which the two beams cross, θ , according to:

$$\Lambda = \lambda / (2 \sin(\theta / 2)). \quad (3.1)$$

Take a look at the a set of attached transparencies with the same color of lines (see Appendix III. Interference between crossed laser beams.) The lines represent the wave crests of a laser beam; the arrows show the direction the beams are traveling, and the distance between the lines corresponds to the wavelength of light, λ . Now try crossing the arrows at an angle, θ . You should see a pattern of dark and light intensity; this corresponds to destructive and constructive interference of light with a spacing of Λ . What happens to the grating spacing as you increase the crossing angle? Now use two the transparencies with the red lines and try crossing them at the same angle as the green beams (a). Notice that the wavelength is larger. Is the resulting angle larger or smaller? Do your observations match what you would expect based upon equation (3.1)?


The interference pattern produced by two crossed laser pulses can be used to generate a sound wave of a well-defined wavelength in a thin metal film. This is because wherever the light interferes constructively, we have a region that is heated and quickly expands. Note that it's very important that the pulses that generate the acoustic wave – commonly called “pump” or “excitation” pulses – are shorter than the time it takes for the sound it generates to move through a full wavelength. If they are too long, they will relative to the frequency of sound. If they were longer, then the acoustic waves in the thin film would move

III.C. Detecting Acoustic Waves with Lasers: Diffraction

To detect the acoustic wave, we use a second, probe laser that “diffracts” off of the surface wave. To better understand what this means, the following sections will introduce you to the principles of diffraction.

III.C.1. Amplitude Diffraction Gratings

Like interference, diffraction comes from the wave nature of light. Whenever a wave hits an edge or a corner, it bends around it a little bit. Most of the time you don't notice the bending of light waves, because the wavelength of light is so short compared to sizes of edges that you usually look at. An edge that looks sharp to the naked eye, which has difficulty sizing up objects much smaller than a millimeter, looks soft and round relative to a tiny wavelength of light. However, if the edge is very sharp you will see more bending, or more diffraction. When light passes through a slit that is close in size to the wavelength of the light, we can observe bending. If we look at the intensity of the light far away from the slit, we will see that the waves that have been bent in different directions have traveled different distances. This results in regions of constructive interference (where the waves add up) and regions of destructive interference (where the waves cancel each other). This is shown in the picture below.¹²




The narrower the slit is, the more bending we see, and the broader the distribution becomes on the other side of the slit:

¹² Adapted from <http://hyperphysics.phy-astr.gsu.edu/hbase/hframe.html>




When we add additional slits, we see more complicated effects because light that gets bent as it travels through each slit can interfere with the light from other slits. For example, if there were no diffraction, or bending of light through each slit, the pattern from two slits would look like the picture below.¹³




¹³ From <http://hyperphysics.phy-astr.gsu.edu/hbase/hframe.html>

In other words, interference between the light waves that pass through different slits causes a regular pattern of peaks. In reality, since waves bend a little (diffract) when they pass through the slits, we see an “envelope” of intensity with the central peaks brightest. As we add more slits, the peaks get sharper and better separated, as shown below.



As you can see from the picture below, the peak intensities become larger and larger as you add more slits.



To see how waves traveling in different directions interfere, go to <http://cat.sckans.edu/physics/ripple.htm> and observe the interference between two point sources of waves. Click on the “edit” button and add another source. Press the “calculate” button and watch the pattern that results. What’s changed? Try adding additional sources and changing the spacing between them. What do you notice about the relationship between the number of sources and the interference pattern? What about the spacing between them?

III.C.2. Phase Diffraction Gratings: Transmission


A diffraction grating changes the direction of the light that strikes it by “bending” light around a regularly spaced set of edges, in the manner illustrated above. The examples illustrated above show how to create a diffraction grating by blocking some regions of light. This is called an “amplitude” grating, because it changes (blocks) the amplitude of light traveling through some parts of it. But we can also diffract light by passing it through a “phase” grating, which changes the phase of light traveling through some parts of it. For example, a piece of glass with wells etched into it at a regular spacing will bend the light slightly as it passes through the wells. And because light travels more slowly through glass than through air, the light waves traveling through the

wells will come out ahead of those that travel through the full depth. In other words, the crests of light will be shifted – some will be ahead of others. We call this a phase shift.

The angle at which the beams diffract from a phase grating is the same as for the transmission grating described above.


$$\sin \theta = m\lambda/\Lambda . \quad (3.2)$$

Again, the angle increases as we increase the wavelength of light, and as we decrease the grating period Λ , or the distance from one well to the next. The picture below shows this working for a Helium Neon laser.¹⁴



As you can see, the beam comes in, and then splits into many parts because of the diffraction grating. The pattern produced from a square well amplitude grating is shown below.


¹⁴ Picture taken from <http://hyperphysics.phy-astr.gsu.edu/hbase/hframe.html>



III.C.3. Phase Diffraction Gratings: Reflection

Another type of phase grating reflects rather than transmits light. Again, the mechanism of diffraction is the same – light hits a small structure and bends slightly. Because the structure is regularly repeated, we get a well-defined diffraction pattern. A reflection grating is shown in the picture below.¹⁵

¹⁵ Taken from <http://www.gratinglab.com/library/handbook/chapter11.asp>



As this picture shows, light diffracting from the reflection grating diffracts at an angle β , meeting the Bragg condition: $\sin \beta = m\lambda/2d$. This formula is very similar to that for a transmission grating, with one exception: now the diffraction angle is inversely proportional to *twice* the grating spacing. As the picture above shows, light that comes in at the Bragg angle β diffracts along the same path that it traveled in on.


In our experiment, we diffract a probe laser beam off of a surface acoustic wave in a thin film. This looks very much like the picture shown above, except that acoustic waves don't have sharp edges. But the process is the same: the probe laser is aligned to the Bragg angle, so when it strikes the surface of the thin film, it diffracts along the same path that it came in on.

III.D. Summary: Using lasers to generate and detect acoustic waves in thin films

The optics we use to generate and detect acoustic waves in thin films are shown schematically in the diagram below. First, we use a transmission phase grating Λ_g to diffract excitation (pulsed) and probe (continuous) laser beams into various orders. We select the first order beams ($m=+1, -1$ in equation (3.2)), which are shown in step 1 below.

Next, the excitation pulses are crossed at the surface of a thin film. The constructive interference between them heats some parts of the film, which rapidly expand, launching a surface acoustic wave, or "ripple," of wavelength Λ_f . (See equation (3.1)). Note that if we use 1:1 imaging (i.e. the imaging lenses are identical, as shown below) the wavelength in the film is half the grating spacing (you should be able to prove this by using equations (3.1) and (3.2), and the picture below.)

Finally, the ripple in the surface of the thin film serves as a reflection diffraction grating, which we can detect with a constant probe beam. The probe beam that has been produced with the phase grating is automatically aligned with the Bragg angle, so that it reflects back on itself (you should be able to prove this using the equations below!) The diffracted probe beam is mixed with a reference beam (not shown below). The mixed signal is detected by a photo-detector, which produces an electrical current in response to the light. The current is sent into an oscilloscope (this converts the continuous current to a set of digitally stored numbers), and uploaded to a computer. The diffracted signal oscillates with the same frequency of the acoustic wave in the thin film, and this can be used to gain information about the metallic film.



1. Transmission grating of spacing Λ_g diffracts excitation and probe beams of wavelength λ_E, λ_P , respectively. The first diffracted orders have angles of θ_E and θ_P , respectively. (Section III.C.2.)

$$\sin \theta = \frac{\lambda}{\Lambda_g}$$

2. Excitation pulses of wavelength $\lambda_E=532$ nm, cross at an angle θ_C , in the thin film, generating an acoustic wave of wavelength, Λ_f (Section III.B.)

$$\Lambda_f = \frac{\lambda_E}{2 \sin(\theta_C/2)}$$

3. Probe beam of wavelength $\lambda_P=830$ nm diffracts from the acoustic wave in the thin film at an angle β (Section III.C.3).

$$\sin \beta = \frac{\lambda_P}{2\Lambda_f}$$

Acoustic Wave in thin film sample (much larger than actual size)

IV. Analysis

In this section, you should explain what the raw data you obtained from the experiment means, how you analyzed it quantitatively, and what your analysis tells you about the thin film sample or samples that you looked at. You should do the following:

1. Prepare a signal vs. time plot and explain the following
 - a. What the y-axis corresponds to, physically.
 - b. Compare the units of time that we are looking at to more familiar units of time.
 - c. What do the oscillations in the signal correspond to in the sample?
2. Prepare a plot showing the signal intensity vs. frequency, and explain how it relates to the signal vs. time plot.
3. Show how you can calculate the acoustic velocity.
4. Make a table and/or plot of the velocity vs. another variable. Variables may include:
 - a. The position of the thin film.
 - b. The metal the film is made out of, and its properties (modulus and density).
 - c. The structure of the film (does it have multiple layers, and if so, what are they?)
5. Explain why the velocity changes with the variable. You may want to make use of the appendices.

IV.A. Typical Data and Interpretation

A typical set of data might look like the main plot below:




FIG. 1. ISTS data from a nominal 5500 Å copper film electron-beam deposited on a silicon substrate. The Fourier transform of the data shows that the film supports an acoustic waveguide mode of frequency 270.02 MHz.




The big window shows how the intensity of diffracted light (labeled ISTS Signal Intensity on the y-axis) varies with time. Note that the entire window is contained within 200 nanoseconds, or $1/5^{\text{th}}$ of a millionth of a second! Physically, the oscillations show the frequency of an acoustic wave as it moves across the surface of the thin film sample. Each time the acoustic wave travels through a full wavelength, it causes a full oscillation in the signal.

IV.B. Fourier Transforms

The inset plot shows a Fourier Transform, or power spectrum of the data. The central frequency of the power spectrum corresponds to the most predominant frequency in the big window. In this case, it looks like the peak of the power spectrum is at about 270 MHz. This corresponds to the frequency of an acoustic wave, and means that the acoustic wave has a period of about $1/(270 \times 10^6 \text{ Hz}) = 3.7 \times 10^{-9} \text{ s}$, or a little less than 4 nanoseconds. This is about the amount of time required for light to travel 1 meter (it travels 30 cm/nanosecond).

In general, the Fourier Transform shows the relative power of different frequencies that go into making up the signal vs. time plot. Any signal in real time can be thought of as a sum of many different signals. The power spectrum decomposes a signal

in real time into a sum of sine waves, and then shows the amplitude and frequency of each wave. This is shown in the plots below. The top plot shows the sum of three waves, each with a different frequency and amplitude. If you add these up, you get the oddly shaped plot in the middle. The Fourier transform, or power spectrum, at the bottom, shows the amplitude and frequency of the original waves.



To try out different wave sums and look at the corresponding change in the position of peaks in the power spectrum, go to the website resources page and download the excel file on Fourier transforms.

IV.C. The Acoustic Velocity

We know the wavelength of the acoustic wave from the transmission grating we used to generate the two excitation beams, and the optics we used to cross them (see Section III.D. Summary: Using lasers to generate and detect acoustic waves in thin films). Hence the acoustic velocity can be determined straightforwardly from the wave relationship, $v = f\lambda$ (equation (2.1), discussed in Section II.A. Introduction to acoustic waves, p 12.)

The velocity of the surface acoustic wave in the thin metal film will *not* match either its longitudinal velocity or its transverse velocity in the bulk metal. Rather, it will vary to the extent that the wave is shared between the top film layer and the underlying layers. As the top film layer becomes thinner, we expect the acoustic wave to “feel” the effects of the underlying layer more substantially. Using this knowledge, and Appendix IV. Physical properties of metals., you should be able to qualitatively explain why you see the velocity change with the film structure or material.

Summary of Appendices

Appendix I. Metrology for fabricating faster circuits.

Philips Advanced Metrology Systems, “Advanced Metrology for Copper/Low-k Interconnects,” 2003.

This paper provides a rationale for using surface acoustic waves in the fabrication of new, faster circuits. Students need not understand all of the details; it is included to provide further information about how interconnect fabrication processes are changing, and what challenges these changes pose for quality assurance.

Appendix II. Using acoustic waves to depth profile film systems.

Dhar, Rogers, Nelson, Trusell, “Moduli determination in polyimide film bilayer systems: Prospects for depth profiling using impulsive stimulated thermal scattering.” *J. Appl. Phys.* 77(9), p 4431. 1 May 1995.

This is a highly technical paper – as a high school student, you should not expect to understand the math! However, you should be able to get a sense of what we can learn with our experiment. This paper also gives good illustrations of the ways that different types of waves travel through thin films.

Appendix III. Interference between crossed laser beams.

This appendix consists of two sets of two transparencies, which can be used to demonstrate how crossed laser beams interfere in a grating pattern, and hence generate acoustic waves. In particular, by changing the angle between the laser “beams” represented on the transparencies, you can convince yourself that increasing the angle decreases the wavelength of the grating pattern. Similarly, by comparing the grating pattern created by crossing beams of different wavelengths, you can convince yourself that increasing the laser wavelength increases the grating spacing. In other words, you can convince yourself of equation (3.1), qualitatively.

Appendix IV. Physical properties of metals.

Grigoriev, Meilikhov, eds. Handbook of Physical Quantities, (New York: CRC Press) 1997. pp 178-180.

Dwight E. Gray, ed. American Institute of Physics Handbook, Third Ed. (American Institute of Physics) 1972. Excerpts from Sections:

Section 2b. *Density of Solids*. Trent, Stone, Lindsay. Pages 2-19 – 2-22.

Section 3f. Acoustic Properties of Solids. W.P. Mason. Pages 3-98 – 3-104.

This provides information on the densities of different solids, and the velocity of different types of acoustic waves in solids. It should be useful for explaining why the surface acoustic wave velocity changes with changes in thin film structure and composition.