Chapter 1

Fourier Optics

1.1 Introduction: image formation

In the figure below a mesh object is illuminated with collimated coherent radiation (produced by a laser) and a magnified image of the mesh is formed by the second lens, referred to as the transform lens. The magnified image is located in the plane a distance s_i away from the lens (at position $z = B$) which is itself a distance s_0 from the object. In the thin lens approximation, the magnification is equal to the ratio

$$m = \frac{s_i}{s_o}, \tag{1.1}$$

and the object and image distances (s_o, s_i) are related to the focal length f by the image equation

$$\frac{1}{s_i} + \frac{1}{s_0} = \frac{1}{f} \tag{1.2}$$

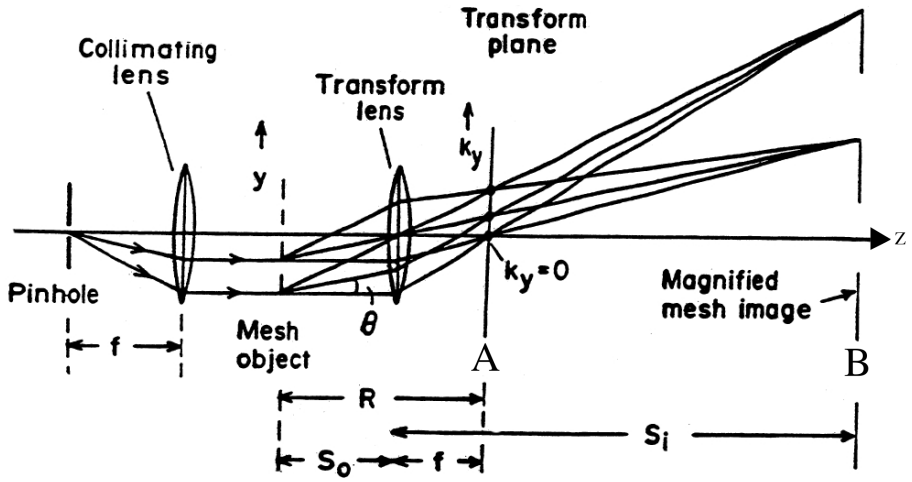


Figure 1.1: Magnified image formation with parallel coherent incident radiation.

In contrast, if the screen is placed at $z = A$, something else is produced. In particular, as will be shown below, the spatial 2-D Fourier transform [1] of the object image will appear in the plane at $z = A$

1.1.1 Optical Fourier Transform Produced by a Lens

In order to understand how a lens generates the Fourier Transform of the light wave emanating from the object, let's imagine that we break up this object wave into a superposition of plane waves (this can always be done for linear optical systems). Now, the lens simply focuses each of these plane waves to a different spot in the focal plane. So, the optical wave at the focal plane behind the lens is the 2-D Fourier transform of the optical wave leaving the object.

Let's see how this works. Consider a plane wave of complex amplitude $U(x, y, z) = A \exp[i(k_x x + k_y y + k_z z)]$ with wavevector $\mathbf{k}= (k_x, k_y, k_z)$, wavelength λ, wavenumber $k = (k_x^2 + k_y^2 + k_z^2)^{1/2} = 2\pi/\lambda$, and complex envelope A. The vector $\mathbf{k}$ makes angles $\theta_x = \sin^{-1}(k_x/k)$ and $\theta_y = \sin^{-1}(k_y/k)$ with the $y - z$ and $x - z$ planes, respectively, as illustrated in Fig. 1.2. The complex amplitude in the $z = 0$ plane, $U(x, y, 0)$, is a spatial harmonic (i.e. sinusoidal) function $f(x, y) = A \exp[2i\pi(v_x x + v_y y)]$ with spatial frequencies $v_x = k_x/2\pi$ and $v_y = k_y/2\pi$ (cycles/mm). The angles of the wavevector are therefore related to the spatial frequencies of the harmonic function by

$$\begin{aligned} \theta_x &= \sin^{-1} \lambda v_x \\ \theta_y &= \sin^{-1} \lambda v_y \end{aligned} \tag{1.3}$$

Recognizing $\Lambda_x = 1/v_x$ and $\Lambda_y = 1/v_y$ as the spatial periods of the harmonic function in the x and y directions, we see that the angles $\theta_x = \sin^{-1}(\lambda/\Lambda_x)$ and $\theta_y = \sin^{-1}(\lambda/\Lambda_y)$ are governed by the ratios of the wavelength of light to the period of the harmonic function in each direction. These geometrical relations follow from matching the wavefronts of the wave to the periodic pattern of the harmonic function in the $z = 0$ plane, as illustrated in Fig. 1.2.

If $k_x \ll k$ and $k_y \ll k$, so that the wavevector $\mathbf{k}$ is paraxial, the angles θ_x and θ_y are small (therefore $\sin\theta_x \approx \theta_x$ and $\sin\theta_y \approx \theta_y$) and

$$\begin{aligned} \theta_x &\approx \lambda v_x \\ \theta_y &\approx \lambda v_y \end{aligned} \tag{1.4}$$

Thus the angles of inclination of the wavevector are directly proportional to the spatial frequencies of the corresponding harmonic function.

Now, the different plane-wave components that constitute a wave may be spatially separated (i.e. isolated) by the use of a lens. Recall that a thin spherical lens transforms a plane wave into a paraboloidal wave focused to a point in the lens focal plane. If the plane wave arrives at small angles θ_x and θ_y, the paraboloidal wave is centered about the point$(\theta_x f, \theta_y f)$, where f is the focal length (see Fig. 1.3). The lens therefore focuses (maps) each plane wave propagating in the direction (θ_x, θ_y) onto a single point $(\theta_x f, \theta_y f)$ in the focal plane and thus spatially separates the contributions of the different harmonic functions.

[1] If you don't remember what a Fourier transform is, now is a good time to review it. Also see $http://www.falstad.com/fourier/$ for an instructional java applet.

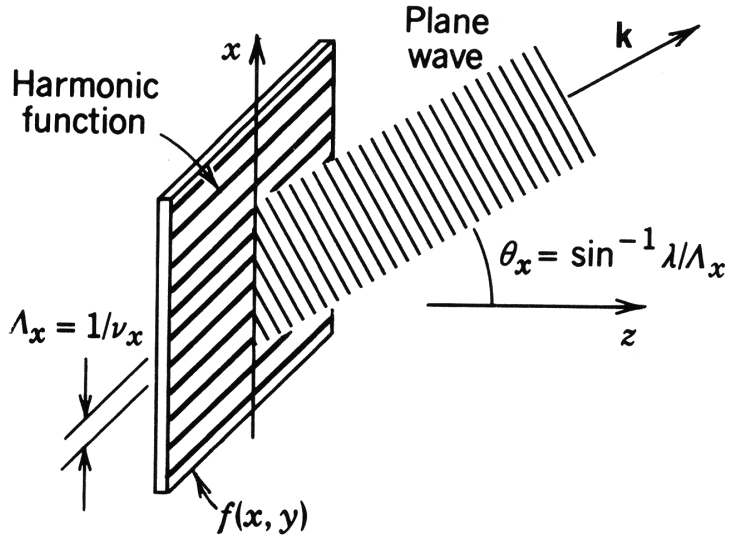


Figure 1.2: A harmonic function of spatial frequencies v_x and v_y at the plane $z = 0$ is consistent with a plane wave traveling at angles $\theta_x = \sin^{-1} \lambda v_x$ and $\theta_y = \sin^{-1} \lambda v_y$.

In reference to the optical system shown in Fig. 1.1, let $f(x, y, z)$ be the complex amplitude of the optical wave in the x–y plane at some position z. The optical wave emanating from the object in the $z = 0$ plane located to the left had side of the lens is then $f(x, y, 0)$. Think of the light coming from the object as a superposition of plane amplitude waves, with each wave component traveling at a small angle $\theta_x = \lambda v_x$ and $\theta_y = \lambda v_y$ having a complex amplitude proportional to the Fourier transform of $f(x, y, 0)$, namely $F(v_x, v_y)$. Each of these plane wave component is then focused by the lens to a unique point (x', y') in the focal plane where $x' = \theta_x f = \lambda f v_x$, $y' = \theta_y f = \lambda f v_y$. The complex amplitude of the wave at point (x', y') in the output or focal plane (located at $z = A$) is therefore proportional to the Fourier transform of $f(x, y, 0)$ evaluated at $v_x = x/\lambda f$ and $v_y = y/\lambda f$. Thus we have

$$f(x, y, z = A) \propto F\big(\frac{x}{\lambda f}, \frac{y}{\lambda f}\big) \tag{1.5}$$

Short range irregularities in $f(x, y, 0)$ (sharp edges) contribute primarily to the diffraction field away from the optic axis, *i.e.* for large (x', y') or (k_x, k_y). Long range irregularities (soft edges) contribute primarily to the small spatial frequencies (k_x, k_y) near the optic axis in the diffraction pattern. Obviously, spatial filtering of the beam is performed in the transform plane. This filtering can be of simple low-pass or high-pass type, or it can involve selective phase-change of any particular spatial frequency component. Later we will investigate the filtering properties.

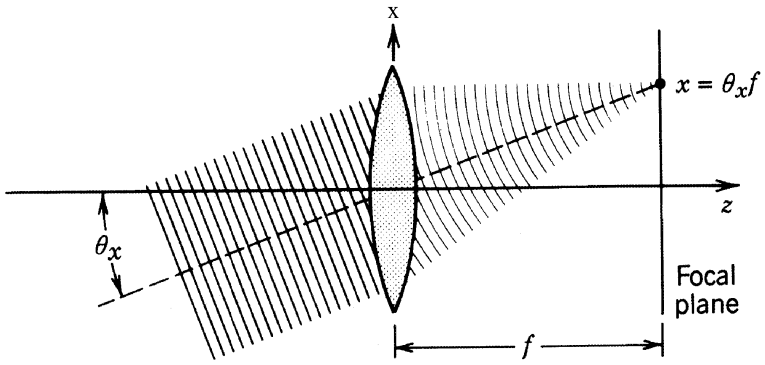


Figure 1.3: Focusing of a plane wave into a point. A direction (θ_x, θ_y) is mapped into a point $(x, y) = (\theta_x f, \theta_y f)$ at $z = f$.

1.1.2 Inverse Fourier Transform

So if the Fourier transform appears at $z = A$, how does an image appear at $z = B$? Very simply, propagation of the wave through free space from $z = A$ to $z = B$ generates the inverse Fourier transform! In this way, the object wave is modified by the lens so that propagation of the wave after the lens produces the Fourier transformed at the transform plane $z = A$, and then, by additional propagation through free space, the inverse Fourier transform is created at $z = B$. This process is what produces an image of the object at $z = B$.

To see how propagation generates the inverse transform, consider this: the electromagnetic wave which crosses the plane at $z = A$ can be though of as originating from a distribution of point emitters in that plane, and each of those point emitters illuminates the image plane (at $z = B$) after propagation by the distance d between A and B. If you write the image wave at $z = B$ as the sum of spherical waves generated by all these point emitters, you get something which, after some approximations, is exactly the inverse Fourier Transform!

Here's the same argument but with the math too. First consider just a single point source of light located on the z-axis in the Transform plane at $z = A$. This point source generates a spherical wave propagating outward (i.e. $E = \frac{1}{r}e^{i(kr-\omega t)}$ where r is the length of the vector pointing from the point source to the observation point). If we examine this spherical wave far from the source but close to the z axis, we can safely approximate it as a paraboloidial wave (i.e. $E \simeq \frac{1}{z}e^{ikz}e^{ik(x^2+y^2)/2z}$). If the point source emitter weren't located exactly *on* the z-axis (i.e. at $x = 0, y = 0$) but rather just off the z-axis, say at (x_i, y_i), then we simply replace $x \to x - x_i$ and $y \to y - y_i$ since the position variables come from the vector pointing from the source at $\vec{r_i}$ to the observation point at $\vec{r}$. In this case, the electric field produced in the plane at $z = B$, a distance d from the point emitter at $z = A$ is then

$$E(x, y, z)|_{z=B} \simeq \frac{1}{d}e^{ikd}e^{ik((x-x_i)^2+(y-y_i)^2)/2d} \tag{1.6}$$

Okay, now if we have a whole distribution of point sources of different amplitudes (brightnesses) in the plane $z = A$ described by the function $E(x, y, z = A)$, then

the electric field observed in the plane $z = B$ far from $z = A$ will just be the sum of all these point sources each located at some initial position (x_i, y_i) and with some amplitude $E(x_i, y_i, z = A)$. This sum is just the integral

$$E(x,y,z)|_{z=B} \simeq \frac{1}{d} e^{ikd} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E(x_i,y_i,z)|_{z=A}\, e^{ik((x-x_i)^2+(y-y_i)^2)/2d}\, dx_i\, dy_i. \tag{1.7}$$

This equation expresses the form of the diffracted field after propagation by a distance d. In the far field limit this reduces to the Fraunhofer diffraction expression

$$E(x,y,z)|_{z=B} \simeq \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} E(x_i,y_i,z)|_{z=A}\, e^{-i2\pi(xx_i+yy_i)/\lambda d}\, dx_i\, dy_i. \tag{1.8}$$

where k has been replaced by it's value $2\pi/\lambda$, we have dropped the unimportant phase prefactor ($\frac{1}{d} e^{ik(d+[x^2+y^2]/2d)}$) which doesn't matter if we only look at the field amplitude. Also, we have dropped the quadratic terms inside the exponent ($kx_i^2/2d$). This last step is okay so long as we're only interested in the field generated by point emitters lying close to the z-axis such that $kx_i^2/2d \ll \pi$. After these simplifications, we are left with just the integral which is the inverse Fourier Transform of the wave at $z = A$. The result is that if we modify the wave at $z = A$ by placing there a mask, the image which results at $z = B$ will be the convolution of the original image and the Fourier transform of the mask.

The Abbé theory of image formation arose from Abbé's consideration of these effects when trying to view small objects through a microscope. It became clear to him that a lens (the objective lens of the microscope) always acts as a low-pass filter, since its finite size will admit only spatial frequencies up to an obvious geometrical limit. The bigger the lens and the larger the range of spatial frequencies admitted, the greater the image fidelity. In the limit of admitting only the undeflected central order ($\theta = 0$), no detail of the object at all will be present in the image. Only a uniform illumination (the so-called d.c. level) will be obtained in the image plane. This, therefore determines the size of the smallest objects which can be viewed by a particular microscope. This relationship is exploited to "clean" up the laser beam coming from the HeNe laser. By focussing the beam through a pinhole, we spatially filter out the high spatial frequencies (i.e. the transverse spatial structures) of the laser to give a clean uniform beam without structure.

1.2 Prelab

Consider the setup shown in figure 1.4. This is part of the basic optical setup for this lab. The beam from the laser is first expanded and collimated by the spatial filter and collimating lens. The beam then passes through the mesh object (one of the objects you will use in this lab). The optical wave transmitted by the mesh passes through the Fourier Transform lens and, due to the magic of wave propagation, at the focal distance (f) after that lens, the Fourier transform of the transmitted beam is generated.

The optical Fourier transform phenomenon arising from wave propagation through a lens plays an important role in pattern recognition, image restoration,

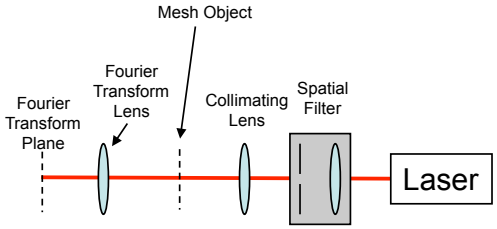


Figure 1.4: Part of the imaging setup for Fourier Optics lab. The spatial filter and collimating lens expand and collimate the laser beam which then illuminates the mesh object. The optical wave transmitted by the mesh passes through the Fourier Transform lens and at the focal distance (f) after that lens, the Fourier transform of the transmitted beam appears. A picture of the mesh Fourier transform intensity pattern is shown in Fig. 1.5

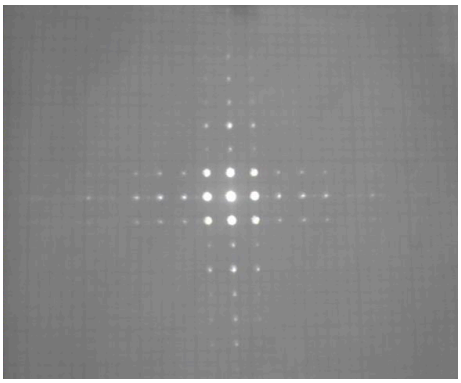


Figure 1.5: Picture of the intensity profile at the Fourier Transform plane produced by a wire mesh object. See Fig. 1.4 for the optical setup.

and optical image processing. A picture of this mesh Fourier transform intensity pattern is shown in Fig. 1.5.

Given the intensity pattern shown in Fig. 1.5, the Fourier Transform Lens focal length is $f = 40$ cm, and the graph paper squares are 0.5 mm on a side, **compute the distance between the wires in the mesh.**

To make imaging the intensity pattern at the Fourier transform plane, you will normally (in this lab) use a second lens to magnify the image before taking a picture of it.

If the wave is allowed to pass the Fourier transform plane, after some more propagation through free space, the optical wave undergoes an inverse Fourier transform, and the image of the transmitted optical wave appears some distance away from the lens. In general, wave propagation by a large distance through free space leads to a Fourier transform of the input wave. The lens simply shapes the input wave so that the Fourier transform of the input wave is complete at a distance f from the lens. Wave propagation after this point produces a second Fourier transform which is complete at a very specific location in z at the image

plane (s_i is the distance from the Image screen to the lens, s_o is the distance of the object to the lens, and $1/f = 1/s_i + 1/s_o$). This second transform is the same as an inverse transform but with the axes inverted, that's why the object image is inverted in this single lens setup. With an appropriate two lens imaging setup, there are 4 transforms and the inversion is made twice producing a non-inverted image.

1.3 The Spatial Filter

The purpose of the spatial filter is to remove any higher order spatial frequencies from the output of the laser to produce a "clean" Gaussian beam with which to illuminate different objects in this lab.

The spatial filter is comprised of a microscope objective and a pinhole mounted on translation stages and should be placed just after the laser. The microscope objective lens focusses the laser through the tiny ($\sim$ 25 to 50 μm) pinhole, and this produces a divergent, coherent beam with a clean spatial profile (i.e. no transverse spatial structure).

Remove everything from the rail except the laser. Align the laser so that it produces a beam which is centered over the optical rail and parallel to it. To check the alignment of the laser, place the pointed metal rod on the optical rail in a holder which allows no horizontal motion. Adjust the laser mount so that the beam remains on the point of the rod as it is moved along the entire rail, approximately 20 cm above the rail. If you find that there is a place where the rail rolls a bit and the beam misses the point (so to speak), try to avoid placing lenses or objects at this position.

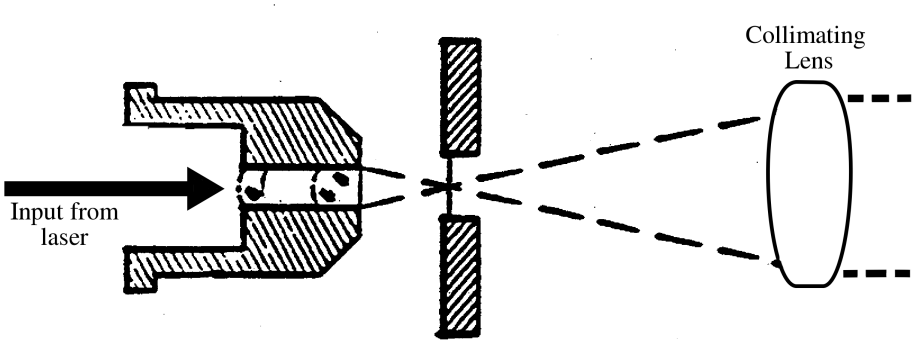


Figure 1.6: A spatial filter consists of a microscope objective, which can be adjusted to focus a laser beam through a small pinhole.

Next, insert and align the spatial filter. With the white screen on the rail, you want to see a large uniform beam without diffraction rings. There are six adjustments which you can make on the spatial filter assembly. It can be rotated about a vertical axis by turning the clamped rod in the holder. This is not a critical adjustment. The assembly as a whole can be raised or lowered, and moved side to side by turning the black knobs on the optical mount. The pinhole itself can be moved up or down, and side to side using the silver knobs attached to it. These are critical adjustments. Finally, the microscope objective lens can be focused on the pinhole using the knob near it parallel to the optic axis.

With this many knobs, alignment is clearly going to be an iterative process. First, align the system horizontally and vertically to get a clear beam along the optic axis. Then focus the system. Then fix the transverse position once again to return the beam to the axis, and so on.

It is best to make the two horizontal adjustments simultaneously. Move the assembly a little, then the pinhole, then the assembly, etc. Treat the vertical adjustments in the same manner. Watch the beam to see if you are making an improvement each time you refocus.

When adjusting the spatial filter, or any other device, assume that the previous user has left it close to the proper alignment. So remember how you found it, start with small adjustments in both directions, and do not make changes without observing their effect.

Check that the beam still runs parallel to the rail. Before you move on, ask a TA to check your spatial filter. If you can't get a uniform circle in ten minutes, ask a TA to help. It is often easier to absorb a demonstration than written instructions. Useful tip in case of difficulty: If you don't see any light through the pinhole, increase the distance between the lens and the pinhole until you can see where the beam hits the area around the pinhole. You can now centre the beam at the pinhole and see light transmitted. You can then readjust the distance until the diffraction pattern disappears.

1.4 The Lenses

Converging lenses, which are generally thicker at the center than at their margins, focus collimated light (a parallel beam) to a point in their rear, or secondary focal plane. Diverging lenses are wider at the perimeter and "focus" collimated light in their rear focal plane (see 1.7). This is the definition of a lens' focus.

Images that are produced on the side of the lens opposite to the object are termed real because they can be seen on a screen. Images on the same side of the lens as the object are called virtual, and are seen only by looking through the refractive surface (from the opposite side, of course). Diverging lenses always produce virtual images, but a converging lens will produce either depending on the position of the object. In this experiment, we concern ourselves only with real images. From a simple lens, real images are usually inverted, whereas virtual images are upright (see 1.7).

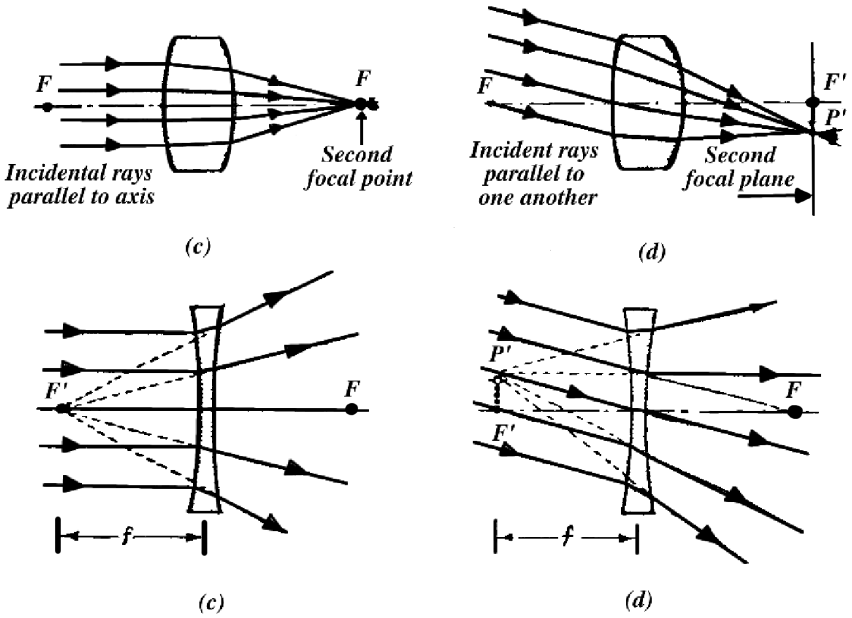


Figure 1.7: Focal planes of converging and diverging lenses.

1.5 Alignment

Figure 1.8 illustrates the *approximate* positions of the major optical components on the 4m optical bench. First check the alignment of the spatial filter pinhole and ensure that output beam is round and symmetrical. Then make fine adjustments to the transverse position of the spatial filter system with respect to the laser beam (i.e. the laser beam propagates along the z axis and you are to adjust the x and y position of the entire spatial filter system) to make the output beam propagate down and along the optical rail at a constant height.

Then place the collimating lens as shown in fig. 1.8 so that the laser light diverging from the spatial filter is collimated along the entire length of the rail. Make sure the collimating lens transverse (x and y) position is correct so that the beam propagates along the rail (along z) at a constant height. Position the end mirror so that the beam runs parallel to the secondary rail and check the laser beam spot size on the wall to see that it remains collimated up to that point.

Next, place the transform lens on the bench about 70cm from the collimating lens, and lock it. Avoid the junction of the two rail components. Center the transform lens, checking that the reflections return to the laser. With the screen or your hand, find the place along the rail where the light is brought back to a focus. Place the screen on the rail beyond the focus, so that the beam has re-expanded to a disc of 10 or 15 cm.

Finding the object plane. In the following sections, you will have to place various apertures and objects in the object plane. The best way to handle this is to first find the plane with an easy aperture/object and then change

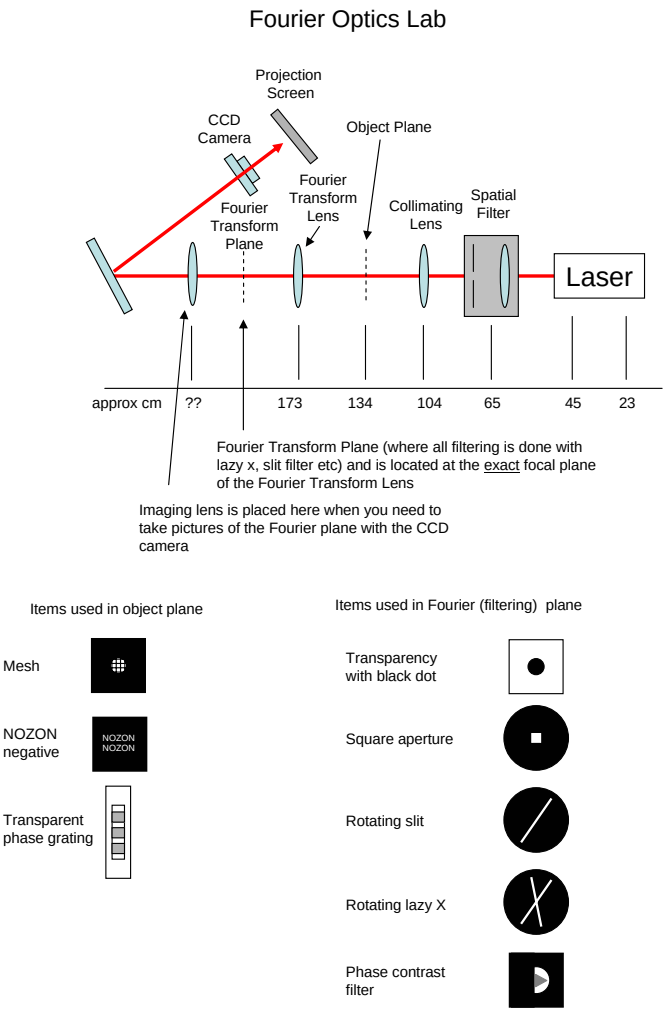


Figure 1.8: Experimental arrangement showing the *approximate* positions of the major optical components on the 4m optical bench as well as the various object and filter components used in this lab. The laser is mounted on one end, the plane mirror on the other. Position the mirror so that the beam illuminates a screen on the secondary rail.

apertures while leaving the mount locked down in that position. To find the object plane, use either the NOZON negative aperture or the aperture which consists of a razor blade with a wedge of clear glass glued to it. Place this in between the collimating lens and the transform lens and move it along the rail (i.e. along z) until a clear and sharp image appears on the screen on the secondary rail. You can have your lab partner (or CCD camera) watch carefully for the exact position at which the image of the razor blade becomes sharp and in focus without any additional structure arising from diffraction. Lock down the position of this mount and replace the razor blade aperture with your aperture of choice.

Finding the Fourier transform plane. In the following sections, you will have to find the exact location of the FT plane and to place objects there either to completely block or simply phase shift certain spatial frequencies. The easiest way to do this is to start with the mesh object in the object plane since it gives a nice FT pattern. Then put the aperture of choice onto a mount with an x and y translation stage. Move the translation stage in the horizontal direction x far enough that the beam emerging from the transform lens lands to the side of the aperture on the black metal surface. Push the entire mount in the z direction until the Fourier Transform pattern is sharp and crisp. Then lock down the mount and remove or change out the aperture in the FT plane leaving the mount in place.

1.6 Mesh Filtering Experiment

1.6.1 Experimental Arrangement

Put in the wire mesh aperture and check that its image is focussed on the screen. It should be perfectly in focus assuming you found the object plane correctly from before. Using the mesh to find the object plane is very tricky because the Fresnel diffraction pattern from the mesh can be mistaken for a focused image. Be careful you have the correct location for the mesh since the following measurements depend critically on this step. The easiest way to verify you have found the correct location is to translate the mesh off-axis (say in the $x-$direction perpendicular to the optical axis) so that you can see in your image of the mesh part of the circular aperture to which the mesh is fastened. Ensure the image of the edge of the aperture is well focused and centre the mesh afterwards. Note the positions of the two lenses and the object (mesh), so that they may be easily replaced when moved. Also note the screen position and the angle between the two rails so that you can compute the object and image distances from the lens.

From the number of wires per cm given on the mesh, and the spacing of the image wires on the screen, estimate the magnification of this system. **Does this magnification agree with the thin lens formula (Equation 1.1)?**

1.6.2 The Transform Plane

In this arrangement, the Fourier transform plane is the back focal plane (a distance f from the lens) of the transform lens (See figure 1.1). A white screen placed here allows the Fourier transform of the mesh to be seen easily and

clearly, even with the room lights on. It is an extensive and intricate pattern. Each point in the pattern represents a unique spatial frequency, and is therefore the focus of all parallel rays in the object space making a certain angle with the optic axis. **What is the effect of moving the mesh?** Move it in the x, y, and z directions and rotate it. **Comment. How would this be useful for optical scanning systems, such as UPC symbols at the grocery check-out?**

The transform may be enlarged and viewed on the secondary rail by inserting a converging lens between the rear focal plane and the screen. Adjust the position of the converging lens until the diffraction pattern is focused on the screen. In this case the screen on the secondary rail and the plane of the spatial filter pinhole are conjugate, as may be seen by removing the mesh. Off-axis points in the pattern correspond to differences in optical path length from repeating parts of the mesh to the screen by integral multiples of the wavelength.

Attach a sheet of black graph paper to the screen. Position the CCD camera such that the entire screen will be imaged. Turn off the lights, but leave the door open so that the white lines of the graph paper are still visible in the image. Take a picture of the magnified diffraction pattern (directions can be found in Appendix A). Now measure the spatial frequencies present in the transform.

Are these frequencies what you expect given the wire spacing of the object? You can compensate for any distortions in the image due to the camera lens by using the grid on the graph paper. The information given in Appendix B might be useful in completing the required analysis.

1.6.3 Spatial Filtering of Vertical Lines

Return the lenses and the mesh to the positions for which a magnified image of the mesh is seen on the screen. Insert the slit in a holder and place it in the transform plane. The diffraction spots in the vertical direction are produced by the vertical spacings between the horizontal wires of the mesh. By aligning the slit to allow only the central vertical lines of the Fourier transform to pass, nearly all of the information about the horizontal separation of the vertical wires is blocked out. **What does the resulting image look like?** Record your observations and **explain them in reference to the mathematical expression for the Fourier Transform.**

It is possible to confirm the optical Fourier transform through image processing of the mesh grid image. This can be done using Matlab or Octave. Octave can be found on any of the Optics Lab computers by typing 'octave' in a terminal window, while Matlab can be run from any physics.ubc.ca xterminal by typing 'matlab&'. This confirmation will be considered bonus, and extra marks will be given accordingly.

1.6.4 Spacial Filtering to Produce 45° Lines

Strange as it may seem, an image pattern containing lines at 45° to the vertical may be produced by rotating the slit-filter to the symmetric 45° position.

Verify this and take a picture. **Explain why this is the case.**

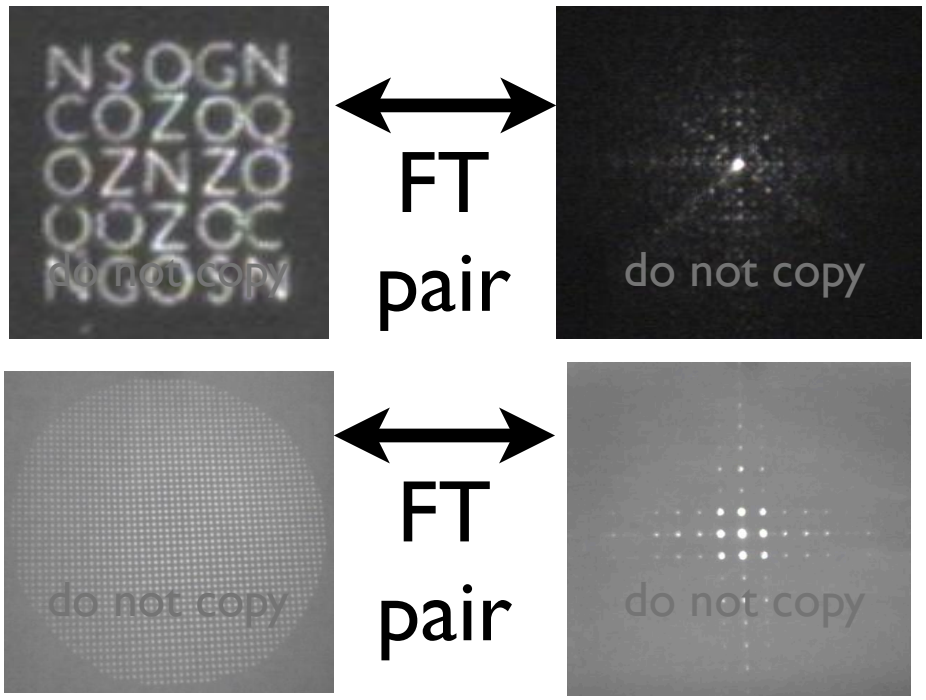


Figure 1.9: Images of the object-character array and mesh and their respective Fourier Transform images produced by a lens. These each form a "transform pair".

1.7 Dark-field Image

A "dark-field image" is obtained when the central (zeroth) order spot in the transform is blocked. The dark regions in the unfiltered image are a result of destructive interference of all the light coming from the object. If the DC (zero frequency) component is removed, complete destructive interference will not occur. Thus, previously dark regions of the image will appear bright. Since the high spatial frequency components required to define the edges of the apertures are still transmitted, the edges haven't lost any definition, so will appear more clearly. Here it is essential that the image be in focus. A difference in focussing of even a few millimetres can destroy the effect.

For this experiment, you should use a different aperture other than the wire mesh since the dark-field image of the mesh is somewhat difficult to interpret. Instead, you should use the aperture which consists of a razor blade with a wedge of clear glass glued to it. This aperture is used later in this lab to realize phase-contrast imaging. Place this aperture in the object plane (where the wire mesh was) and translate it until its image is clear and crisp. Now place the transparency with the black dot on it into the FT plane.

Describe your observations and take a picture for your lab book.

1.8 Character Recognition

1.8.1 Spatial Filtering to Transmit "N"s

Place the object-character array negative NOZON at the object position. Place a rotatable lazy-x spatial filter in the transform plane. This filter should transmit the N's if properly aligned. It should be in a holder with vertical and lateral adjustment. Orient the filter to transmit the N's but remove the O's and Z's. Most of the non-N letters will not be completely removed because of cross-correlation effects. If the filtering isn't perfect, your spatial filter might not be exactly at the Fourier Transform plane. The position of the filter can be optimized by observing the Fourier Transform pattern hitting the back side of the filter and adjusting the filter position to make the pattern as sharp as possible.

Take a picture. **How does this filtering work?**

1.8.2 Character Transform

The transform of these letters is a very detailed "snowflake" pattern and is shown in Fig. 1.9. As in Section 1.6.2, add a converging lens behind the transform lens and adjust its position until the transform is focused on the screen. Take a picture.

1.9 Phase-Contrast

The only spatial filtering discussed so far has involved altering the amplitude of the transmission function, by passing only selected spatial frequencies through the Fourier Transform plane. However, the phase of the light wave is also available at the Fourier Transform plane, and altering the phase of selected spatial frequencies using phase changing spatial filtering masks offers different possibilities for optical image manipulation and reconstruction.

Consider biological imaging applications where the objects you are trying to observe (say a protist swimming in water) are almost as transparent as the medium in which they are situated. In addition, different parts of the cell will typically have the same absorption coefficient and therefore result in the same output wave amplitude, but they have different indices of refraction. As a result, a substantial amount of information is encoded in the *phase* of the light. However, because photodetectors (like your eye) measure the intensity of the wave, this phase information is lost in a normal microscope. Phase contrast imaging is a technique by which spatial filtering is used to transform the phase variation across the wave into amplitude variations which can then be detected. The invention of this method was first applied in Phase-Contrast Microscopy, which earned Fritz Zernike a Nobel Prize.

Returning our attention to the protist, the index of refraction of this creature is close to that of water, so the transmission function will be approximately equal to that of water but with a small position dependent phase rotation (see figure 1.10A). A plane wave passing through the water will have the form e^{ikd} where d is the thickness of the water layer. This same plane wave will pick up an extra phase term from the modification of the transmission function due to the protist in the water. The complex amplitude of the plane wave passing through

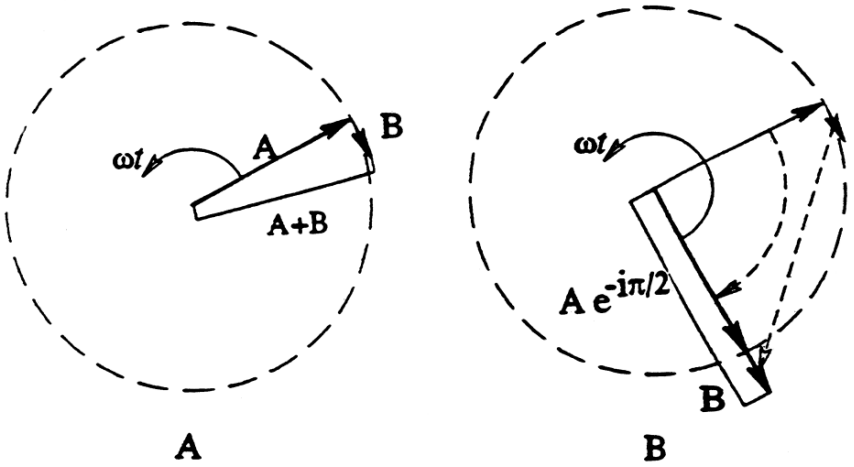


Figure 1.10: Phasor diagram illustrating the principle underlying phase contrast imaging. In the left image (A) we see the two parts of the field **A** and **B** adding. In the text, the field is $f(x,y) = 1 + i\theta(x,y)$, so here $\mathbf{A} = 1$ and $\mathbf{B} = i\theta$. Because one is real and the other imaginary, they are orthogonal and the length of the resulting vector is almost the same length as the original **A**. Therefore the contribution from **B** is hard to see. In the right image (B), The **A** component is first phase shifted by 90° and then added to **B**. In this case, they add in the same direction, and any changes in **B** produce a much larger change in the length of the resulting vector (the intensity of the field) than before when the two components were orthogonal. This increases the intensity contrast produced by the variation of the **B** component making it more visible. Phase contrast imaging is therefore especially useful for seeing clear objects, such as a fingerprint on glass, where the object encodes its presence into the *phase* of the light.

the protist will be $f(x,y) = e^{ikd}e^{i\theta(x,y)}$, where $\theta(x,y)$ is the position dependent additional phase shift induced by the protist body and its organelles. If we were to image this wave on a detector, the intensity profile, $I(x,y) \propto |f(x,y)|^2 = 1$, is completely uniform and we can't possibly see the creature like this. So, we need to implement phase contrast imaging to reveal the phase information. Let's see how this works.

Assuming the additional phase shift is small, $\theta(x,y) \ll 1$ (this assumption is not so important but makes the math less messy), the exponent can be approximated as $e^{i\theta} \simeq 1 + i\theta$. Dropping the overall phase of the plane wave in $f(x,y)$, we have

$$f(x,y) \simeq 1 + i\theta \tag{1.9}$$

$$|f(x,y)|^2 \simeq 1 + i\theta - i\theta - i^2\theta^2 = 1 + \theta^2 \tag{1.10}$$

Under this approximation, we can see that (again) there is no information showing up in the square of the field (intensity), since $\theta \ll 1$. So, what should we do? What would happen if we first changed the "1" in "$1 + i\theta$" to an "i"? This

would correspond to a 90° phase shift of the DC term (the "1") while leaving the spatially varying part (the θ) unchanged (See Figure 1.10B). In this case we have

$$|i + i\theta|^2 = -i^2 - i^2\theta - i^2\theta - i^2\theta^2 \quad (1.11)$$
$$= 1 + 2\theta + \theta^2. \quad (1.12)$$

Now the intensity has a 2θ term which is still smaller than the DC term but much larger than θ^2 and would give enough contrast against the flat illumination of the plane wave to see the protist. So, how do we affect this selective 90° phase shift of the DC term?

Since the DC term has no spatial variation, representing just a flat and featureless illumination from the plane wave, it will be focused to a point in the Fourier Transform plane and comprise the zeroth order of the diffraction pattern. This suggests a simple way to alter the phase of the uniform illumination, just insert a thin dielectric layer (film) at the central spot in the transform plane. A quarter wavelength phase change would produce the best results (turning the "1" into an "i"), which is what Zernike did, but clearly any phase change would help us distinguish between the object and the water. Commercial phase contrast microscopes often have an adjustable phase.

1.9.1 Experimental Arrangement

Place the transform lens back on the rail, and introduce, at the object position, a transparent phase grating consisting of parallel ridges and troughs ruled on a microscope slide. (See Figure 1.11.) The added path length in glass produces a small phase difference between the ridges and troughs. However, the image on the screen should be invisible, as this difference is very small. Ensure that the image is properly in focus at which point it should show no signs of the grating.

PLEASE DO NOT TOUCH THE GRATING SURFACE !

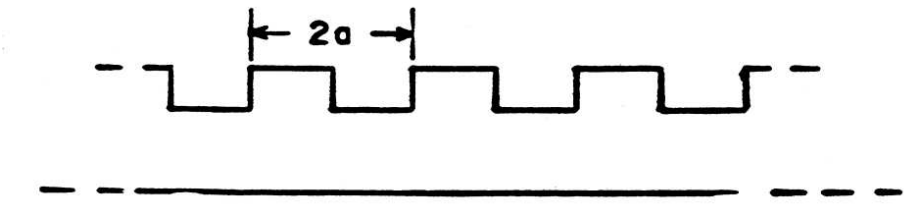


Figure 1.11: Cross-section of the transparent phase-grating.

Insert the vertical razor blade with a wedge of clear plastic taped to it. Adjust until the wedge covers just the central diffraction spot. The magnified image of the grating should be visible on the screen, with the ridges brighter than the troughs. Measure the grating spacing in the image. Compare this magnification with that found earlier in section 1.6.1.

Take a picture of the filtered image.

1.10 Diffraction

The term diffraction is generally understood to mean the propagation of light "around" the edges of obstacles or apertures. Diffraction phenomena are traditionally divided into two classes - Fresnel and Fraunhofer diffraction - although they are both manifestations of the same phenomenon. Fraunhofer diffraction refers to what happens to an optical pattern after propagation over a very large distance z in the paraxial approximation. Fraunhofer diffraction is the *far field* limiting case of Fresnel diffraction that arises when both source and observer are at very large distances from the diffracting screen such that the incident and diffracted waves are effectively plane. In this limit, the curvatures of the incident and diffracted waves can be safely neglected and the Fraunhofer diffraction pattern is just the Fourier transform of the initial pattern. If, however, either the source or receiving point is close to the diffracting aperture such that the wave front curvature is significant, then the result will be Fresnel diffraction [see Fig. 1.12 [2]].

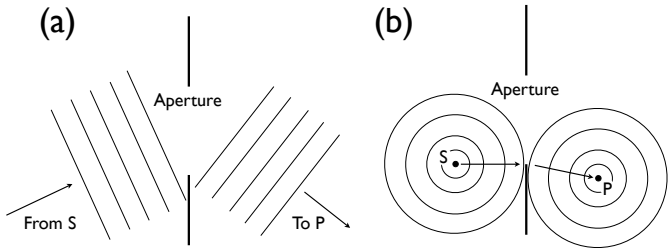


Figure 1.12: Diffraction by an aperture in the (a) Fraunhofer limit and (b) Fresnel limit

In order to be in the Fraunhofer limit, the distances between the source and diffracting aperture and between the diffracting aperture and receiving point should be large compared to a characteristic distance, R_0 given by;

$$R_0 = \frac{2\mathrm{w}^2}{\lambda} \tag{1.13}$$

where λ is the wavelength of the light, and w is a dimension of the aperture or obstacle. This relation for R_0 guarantees that the phase change across the width of the aperture is less than $\frac{\lambda}{4}$. Perfect Fraunhofer diffraction can only be achieved with collimated incident and diffracted beams.

1.10.1 Diffraction Pattern of a Slit

In this experiment, you will investigate the two (Fraunhofer and Fresnel) diffraction regimes by studying the diffraction pattern from a slit aperture of variable width. You should start by inserting the collimating lens after the spatial filter so that the beam wave fronts are parallel. Thus the source of illumination is effectively at infinity. Place the slit after the collimating lens and reflect the

[2] figure adapted from "Introduction to Modern Optics" by Grant R. Fowles

transmitted light with the secondary mirror onto the screen with the image visible by the CCD camera.

The diffraction pattern observed can be characterized by a dimensionless parameter, Δv, defined as;

$$\Delta v = \mathrm{w}\left(\frac{2}{R\lambda}\right)^{1/2} = \left(\frac{R_0}{R}\right)^{1/2} \tag{1.14}$$

where w is the width of the slit, and R is the slit-screen distance. Fresnel diffraction occurs for $\Delta v \geq 1$. Fraunhofer diffraction, the so called far-field approximation is characterized by $\Delta v \ll 1$. This condition can be established with either a large slit-screen distance as the name implies, or a narrow slit. In the lab, the slit width is the more convenient adjustment to vary the Δv value, since a large range of w is easily achieved with a turn of the micrometer. However, varying R is equally acceptable. For each measurement you make, you should compute your Δv value on the spot to insure that you are in the proper limit (recall that the wavelength of the He-Ne laser is 632.8 nm). Because the Fraunhofer diffraction pattern is essentially the Fourier Transform of the input wave, the Fraunhofer diffraction can also be observed at the focal plane of a converging lens (see Figure 1.13)

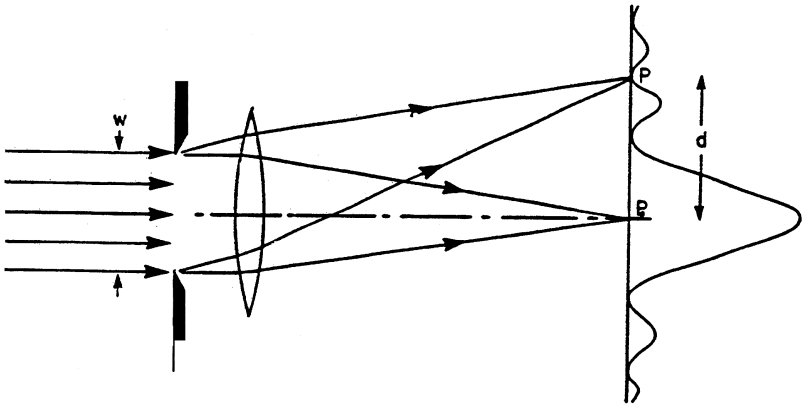


Figure 1.13: Fraunhofer diffraction pattern of a slit formed at the focal plane of a lens. In this part of the lab, you will generate your Fraunhofer patterns *without* the aid of a "Fourier Transform" lens placed after the aperture.

Start by adjusting the slit width to be very narrow until you see the characteristic Fraunhofer pattern. Record and reproduce in your lab write-up the image of the observed diffraction pattern. **Are the minima zero? Why or why not?**

Now slowly widen the slit width watching the diffraction pattern as you do. You should be able to see the transition from Fraunhofer to Fresnel diffraction, with the initial loss of the zeros, and subsequent increase of secondary maxima. Record the image of the Fresnel patterns at three interesting Δv's. For each measurement you make, you should compute your Δv value to insure that you are in the proper limit. Your images should look something like those in Fig. 1.14. Since the CCD has an auto gain circuit, it may help to turn on the

room lights so that the patterns are visible and not completely saturated. Your analysis will be done by taking a line out of the image, so as long as there is a part of the image that is not saturated, you can extract the data. Because the CCD has a limited dynamic range, the image may not show clearly the low intensity wings of the diffraction pattern that will be obvious to you by eye.

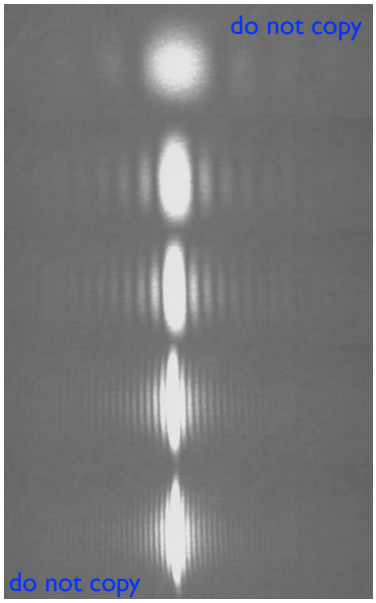


Figure 1.14: Diffraction patterns from a slit of various widths.

1.10.2 Analysis of the slit diffraction

Take your images and import them into an array in MATLAB, maple, mathematica, octave, or your favorite math tool. Plot the intensity as a function of position along an axis perpendicular to the slit axis (let's call that the x-axis, and this axis is horizontal in the images shown in Fig. 1.14). To improve the signal to noise of your data, you may need to integrate (add pixel intensities) the image along the y direction. Since there might be some distortion of the image along the vertical (y) direction (as seen in Fig. 1.14), you should probably only sum up the 10 adjacent pixels along a vertical line for each sample along x. Be sure that you choose an area of the image that is not saturated. In Fig. 1.14, this area would correspond to the top of the images above the point where the central spot is completely saturated. The problem of saturation is probably most apparent in the lowest image where the interference fringe spacing is the smallest. Calculate your values of Δv, and compare your diffraction patterns to theory by plotting your experimentally measured data against the predicted diffraction patterns at the corresponding Δv. For this you should numerically evaluate the form of the diffraction intensity using Matlab or Mathematica using the theory presented.

It is convenient to represent the intensity of the diffraction pattern at a

distance, d from the optic axis in terms of the dimensionless parameter,

$$z = \frac{d}{\mathrm{w}} \tag{1.15}$$

Thus the edges of the geometric shadow correspond to $z = \pm 0.5$.

In the Fraunhofer limit ($\Delta v \ll 1$), the intensity variation in the diffraction pattern is given by the Fourier transform of a square function,

$$I(z) \propto (\Delta v)^2 \sin^2 \beta/\beta^2 \tag{1.16}$$

where $\beta = (z\pi/2)(\Delta v)^2$. The pattern has a global maximum at $\beta = 0$, with subsidiary maxima on both sides. These maxima are separated by zeroes in the intensity for

$$\pm\beta = n\pi \text{ or } \pm z = \frac{2n}{(\Delta v)^2} \quad n = 1, 2, 3, \ldots$$

Notice that the intensity maximum is proportional to $(\Delta v)^2$ and therefore to w^2, while the width of the central maximum is proportional to $(\Delta v)^{-2}$ or to w^{-2}.

For Fresnel diffraction, ($\Delta v > 1$), the general expression for the intensity is given by;

$$I(z) \propto \left[C^2(v) + S^2(v)\right]_{v_1}^{v_2} \tag{1.17}$$

where $C(v)$ and $S(v)$ are the Fresnel cosine and sine integrals, respectively, given by

$$C(v) = \int_{v_1}^{v_2} \cos(\frac{\pi x^2}{2})\mathrm{d}x$$

and

$$S(v) = \int_{v_1}^{v_2} \sin(\frac{\pi x^2}{2})\mathrm{d}x$$

The limits on the variable of integration are

$$v_1 = -(z + 0.5)\Delta v \qquad \text{and} \qquad v_2 = -(z - 0.5)\Delta v$$

In the limit of small Δv, equation 1.17 reduces to equation 1.16.

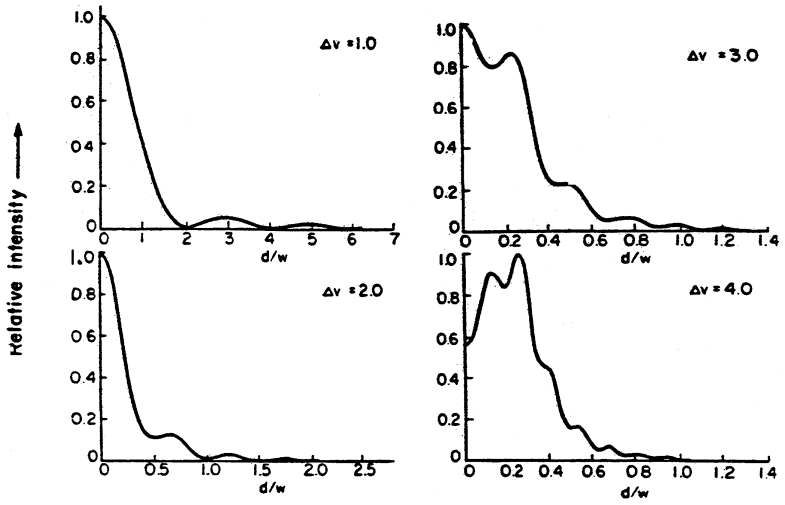


Figure 1.15: Line intensities of Fresnel diffraction patterns of a slit for various values of the parameter Δv. Only one half of the symmetrical pattern is shown.