15: Light and Geometrical Optics

15.1 Optical Behavior of Light

Most of Optics can be understood based on a description of light which would be largely familiar to Newton, who believed light to be composed of streams of some sort of particle rather than a wave. Later physicists would demonstrate that light was in fact a wave, while even more recently we discovered that it was somehow both (quantum mechanics is weird! just wait!). But despite this complex back and forth about the nature of light, many of its most important properties can be described in the old particle stream picture. The fundamental assumption we make about light while doing optics which maps to a particle interpretation is the idea that light is a *ray*. A ray is a concept of pure geometry and refers to something which starts in direction and then never stops. A ray has infinite extent in one direction. An optical ray of light has this properties, and we assign a light ray certain rules to determine how it acts when it encounters something different on its path.

15.1.1 Absorption, Reflection, and Transmission

Optical arguments are generally based on the assumption that little or no light is absorbed when interacting with a material. We use ray arguments that assume those rays continue with infinite extent rather than being absorbed and stopping, and then afterward can correct for any loss of brightness or intensity or whatnot based on how far from ideal the surfaces and materials involved really are.

Absorption of light means that somehow the energy of the light is getting sucked up by a material rather than continuing to travel with the light beam. Any detection of light generally involves its absorption: cameras of all sorts and our eyes, for instance, both function on recording the deposited energy of light on some receptor. But while the detection must assume absorption of light, its propagation must assume an absence of absorption.

When a ray of light encounters a surface (or interface between materials) it can either be absorbed, reflected back into the material it started in, or transmitted through into the other material. The study of basic optics boils down to understanding under what circumstances a ray of light is likely to do each of these things, and what the relevant angles are.

Diffuse vs. Specular Reflection

There are two different ways in which light can reflect off of objects, which are readily apparent in our everyday lives. Both are the same on a microscopic level, but appear quite different on a macroscopic level due to the nature of material surfaces rather than any property of the light itself. *Diffuse* reflection is what we typically see and refers to a situation in which light striking a surface will reflect off in all directions more or less uniformly. Diffuse reflection masks the simply rule governing how a light ray reflects off of a surface. In order to discuss how this works, we have to consider *specular* reflection. Specular reflection is what happens if we have a very smooth surface and all light rays incident upon it are reflected in the same direction. If we were to look closely enough at the rays impacting irregular surfaces we would find this is true there as well, but because the rough surfaces typical of our surroundings has small regions of surface facing all sorts of different directions, the rays end up going all over the place.

Figure 15.1: 32-4

Both forms of reflection are essential to our understanding of our visual experience. Specular reflection gives us our first ideal system from which to start investigating the underlying optical laws, while diffuse reflection explains our ability to see. Our eyes (which we will discuss more later) fairly obviously function by collecting light rays. They can only see rays that find their way into our pupils. If everything were sufficiently smooth for specular reflection, all the light from a given light source (like a light-bulb in a room) would reflect off in the same direction. If your eye happens to be in some other direction from that object, you wouldn't see anything. This is what is happening if someone were to shine a flashlight at a mirror from the side while you stand directly in front of it. You may know the flashlight is being shone at the mirror, but you don't see the mirror being illuminated by it. In such a world we would be surrounded by darkness, punctuated by brief bright flashes of light. Luckily, most objects reflect diffusively, in every direction, and we can see an entire room illuminated by a single source of light.

15.1.2 Specular Reflection and Flat (Plane) Mirrors

While most objects reflect diffusively, it is still quite common to find a flat reflective surface. Bathroom mirrors are an obvious example, but a still pond or lake can serve the same function. Any surface which is sufficiently smooth can act as a mirror, tho some materials will absorb rather than reflect a greater portion of the light incident upon them.

The striking feature of mirrors as opposed to other objects is that they produce an *image*. We look into the mirror to see an image of ourselves looking back. You can do this precisely because each ray of light which reflects off of your face and strikes the mirror reflects based on the same unique rules, and as we will show these rules serve to produce what we call an image.

First, the *law of reflection* states that

The angle of reflection equals the angle of incidence, $\theta_r = \theta_i$ *Law of Reflection*

The angle of reflection is the angle made by a ray with the surface from which the ray has reflected, while the angle of incidence is the angle made by the same ray with said surface before contact. (Drawing these ray diagrams over again is a pointless waste of time, so I'm just going to steal Giancoli's)

Figure 15.2: 32-3

At this time we also define the *normal* as the line extending out from the flat surface perpendicular to the place of the surface, and on the side of the departing ray. Later when we discuss transmission, the normal will still be in the direction of the departing rather than incident ray.

Since rays travel in straight lines, we can use the law of reflection to make a number of arguments graphically. First, we can show what we mean by an "image" and demonstrate how it is created by a mirror.

Figure 15.3: 32-7

For any random object, there will be a large (essentially infinite) number of rays leaving each point on that object, all going in different directions. We use a few sample rays to characterize the whole system. Here, we are considering just 2 rays leaving each from 2 points on the same object. The slight difference in direction leaving the initial object means that when the ray encounters the mirror surface, they will reflect at a slightly different angle from slightly offset locations (B and B' for the bottom point A). This results in *diverging rays*. The rays we have chosen diverge such that they enter the eye on extreme sides of the pupil. The same is true for rays from the top of the object.

Once these rays are drawn, we need to consider how our brain will *interpret* the light rays entering our eye. Our brains do not have a built in conceptualization of reflection: we interpret what we see as tho all rays entering our eyes have traveled on straight, un-reflected, un-bent paths. So, we extend the rays entering the eye with dashed rays *passed the mirror surface* until they converge. Rays which originated at the same point on our object will now converge together at a single point on the opposite side of the mirror surface. Combining all or the rays together, this will produce a *virtual image* at the same distance from the mirror as the actual object, but on the other side. This is an image because to our eye it looks *just like* an object sitting at that location. The rays entering our eye from the reflection are the same as tho there were no mirror but rather the object were at the image location. It is a *virtual* image because the rays of light never actually pass through the imagine location: it is purely a result of our brain's interpretation. If one were to place photographic film at the image location, it would not record an image of the object. We will discuss *real images* later.

Note from this that we are assuming our eye can interpret rays which enter it on *diverging* paths. In fact, it can only properly interpret such rays.

Example: Full-Length Mirror

How tall does a mirror really have to be in order to show the whole of the person looking into it? How close to the floor does it really need to go?

Figure 15.4: 32-8

15.2 Mirrors and Reflection

15.2.1 Spherical Mirrors

Optics would be pretty boring if all we had were flat mirrors. Luckily we can curve reflective surfaces in all sorts of shapes. The next simplest (or at least most common) is the spherical mirror. A spherical mirror needn't be a sphere, the term refers to any reflective surface which takes the shape of a section of a sphere. In fact, we will generally assume that we are dealing with small sections of a sphere.

There are 2 types of spherical mirror, *concave* and *convex*. Concave mirrors are shaped like caves (this is seriously how I remember this) in that the reflective surface is on the inside of the curvature (like the inside of a cave vs. the outside of a ball). Convex mirrors are the other way around, the reflective surface is on the outside as tho we've covered a basketball with a reflective surface. Concave mirrors, as we will see, tend to be used to magnify objects by converging light rays, while convex mirrors cause rays to diverge and take in a wider field of view.

Any curved optical element, starting with spherical mirrors, will have a focal point and focal length. The focal point of a convex mirror is most easily defined, so lets start with that.

Figure 15.5: 32-14

For any section of a sphere we cut out (yes, we could cut out a bizarre shape, but we won't), there is going to be an axis of symmetry called the *principal axis*. This axis passes through the point which would be the center of the full sphere of which we have only a piece, and intersects our mirror at the center. Once we have identified this axis, we can define the focal point and from that the focal length.

The focal point refers to a point at which incident rays are "focused" into a single point. In particular, it refers to the point at which a series of rays which are all perfectly parallel with the principal axis will be reflected into. The *focal length* f is the distance from the mirror to the focal point.

Unfortunately, a true focal point doesn't actually exist for a spherical mirror. Even perfectly parallel rays as described will be reflected into slightly different points along the principal axis. A *parabolic* mirror is needed for a true focus point. However, parabolic mirrors are far more expensive, so most of the time we make due with spherical shapes and deal with the "spherical aberration" as it comes. One of the ways we can minimize this problem is by keeping our mirror very small with respect to the radius of curvature. The closer our rays are to the principal axis (the closer to $\theta_i = \frac{\pi}{2}$) the smaller the effect of the spherical aberration. Such rays are called "paraxial" because they are close to the axis. In most of our discussions, we assume this to be the case and ignore spherical aberration.

For a spherical mirror, the normal extends from the surface of the mirror at which a ray is incident and extends out through the point C corresponding to the center of the sphere. The law of reflection tells us that $\theta_i = \theta_r$ about this normal. Geometry can then show that the angle between the principal axis and the normal is this same angle again. For an isosceles triangle, the two side corresponding to equal angles have equal lengths. We can use this in the $\theta = 0$ case to see that $f = \frac{r}{2}$ where r is the radius of curvature.

Images

Spherical mirrors produce images in a more complex fashion than flat mirrors, and can in fact produce two different kinds of images, real and virtual. To figure out what is going on, we imagine that we have an object sitting on the principal axis of a mirror and start drawing some sample rays and see what happens. We could just pick randomly, but the experience of others provides guidance on which rays are likely to be the most helpful.

Figure 15.6: 32-15

1. Draw a ray from the top of the object parallel to the principal axis until it reaches the mirror. We know these types of rays must pass through the focus, so draw it there.

2. Start again at the top of the object and draw a ray passing directly through the focal point. When this ray extends through to the mirror, we know that it will reflect back parallel to the principal axis (for the same reason as ray one passed through the focus, but in reverse).

3. Start at the center and draw a ray passing through the top of the object. This ray is traveling perpendicular to the surface of the spherical mirror (its coming from the center!) so it must reflect straight back, passing through the object and center point again.

If we've drawn our figure right, all 3 of these rays should converge on the same point. In fact, any ray leaving the same point on the object should at some point pass through this point of convergence: that's essentially how an image is defined and the goal for which we have defined/designed our mirrors the way we have.

So our converging rays define an image. How? Well, they tell us where all of the light from the top of the object will end up at. We could do the same exercise for each point on the object, and if we were careful would would find that the whole thing was meticulously reconstructed around the point we've already identified. Luckily we can get away with only doing it for the one point, and extrapolate from there. Since our image point I' is below the principal axis, our image will be inverted relative to the object. This is because the base of the object, which started on the principal axis, will always stay on said axis (try drawing all three rays above for the base of the object: each one is the same and stays on the principal axis). The image will then be between the image base and top I and I'. Since I' is below, we have an inverted image. Notice also that in our picture the image looks larger than the object! This is largely the point for which curved mirrors are used: creating images which are larger or smaller than the objects they come from.

Also notice that unlike the image we discussed from the flat mirror, the actual rays themselves pass through this image. This makes it real rather than virtual. Film placed here would capture an image, or the image could be seen projected on a screen. Also note that our eye can only see this image properly if it sits at or behind the imagine because we can only properly perceive diverging rays. This is a result of how our eyes are put together, as we will see later, but is in no way surprising. Nature has very few natural lenses, and thus an almost complete lack of converging rays, so we had no reason to evolve the capacity to interpret them.

Magnification

Having a diagram to guide us is all well and good, but trying to solve real problems that way is like trying to rely entirely on our ability to draw field lines correctly to calculate Coulomb's law for 2 point charges: you'll get the right general idea, but good luck getting any sort of precision or accuracy. Luckily, we can use geometry to develop an analytic equation.

Figure 15.7: 32-16

The Law of reflection tells us that the labeled θ are the same, which means that we have similar triangles $II'A$ and $OO'A$. Similar triangles mean the ratios of similar sides are the same,

$\frac{h_o}{h_i} = \frac{d_0}{d_i}$.

$OO'F$ and ABF are also similar (assuming $AB = h_i$ small because we are looking at small lenses relative to r), so

$\frac{h_o}{h_i} = \frac{OF}{AF} = \frac{d_o - f}{f}$

Combining these together, we find that

$\frac{d_o}{d_i} = \frac{d_o - f}{f}$

$\frac{1}{d_i} = \frac{1}{f} - \frac{1}{d_o}$

$\frac{1}{d_o} + \frac{1}{d_i} = \frac{1}{f}$ *The Mirror Equation*

This is a handy result and will keep popping up. At the moment, the straight up magnification is more interesting, however

$m = \frac{h_i}{h_o} = -\frac{d_i}{d_o}$

where we have introduced the first of a number of conventional signs to keep track of the fact that our image is inverted, so in a sense it has a negative high relative to before.

Sign conventions:

- Image height is positive if upright, negative if inverted
- d_i and d_o are positive if in front of the mirror, negative if behind (can be opposite signs if image and object are on opposite sides). Focal distance f also follows this convention, but so far we've only seen it being on the positive side.

Virtual Images

A concave mirror can also give virtual images. If we bring the object closer to the mirror than the focal point and redraw our 3 easy example rays, we find that they fail to converge on the front side of the mirror. When this happens, we trace them backwards across the mirror to simulate how our brain would interpret the diverging rays. We find a magnified but upright virtual image.

Figure 15.8: 32-17

Convex Mirrors

All of the above works, if properly applied, for convex mirrors as well. We can still trace out our rays and see where they go to figure out if we find a real or virtual image. Because the convex mirror causes rays incident from the front to diverge rather than converge, we get virtual (and upright) images for any object placement.

15.2.2 Snell's Law

When light passes from one transparent material to another, the velocity of light changes. We know this from our study of EM waves previously. The velocity of light is determined by a material's *index of refraction* which must be greater than or equal to 1. A larger index of refraction means light travels more slowly through that material. It also means that, due to a result called *Snell's Law*, that light will bend more passing from air to that material. Snell's law tells us what happens when light passes from one transparent material to another, and it isn't so simple as just carrying on at a different speed. It will in fact refract, or bend, at the interface.

Figure 15.9: 32-21

Unlike with reflection where $\theta_i = \theta_r$, Snell's law gives a more complex relationship between the incident and refracted angles:

$n_1 \sin\theta_1 = n_2 \sin\theta_2$ *Snell's Law*

In our figure, the index of refraction of water is greater than that of air. When light passes from one material to another, it will bend such that it travels closer to the normal while in the material with greater index of refraction. This means that when passing from small to large index, it will bend towards the normal, while wen passing from large to small index it will bend away.

Example: Passing through Glass

Figure 15.10: 32-24

When light passes through a region with a different index of refraction but returns to the original material, its direction overall will be unchanged, but all of its rays will be offset by some distance due to the refraction at both interfaces. As an interesting side note, the path the light takes turns out to always be the *fastest* way to get from its starting to end point, and can (with difficulty) be calculated based on this principal. This essential idea in one form or another is in fact possibly the most fundamental known physical principal and can be used to derive most of quantum mechanics, classical mechanics, quantum field theory, and much from solid state theories. Just a fun fact.

15.2.3 Dispersion

While all wavelengths/frequencies/colors of light travel at the same velocity in a vacuum, they can travel at different speeds inside of other materials and thus have different indicies of refraction. Because of this, red and blue are bent by different amounts when passing through the interface between two materials. This process is known as *dispersion* and is a general feature of wave propagation (tho we don't need to think of it that way yet).

Because shorter (bluer) wavelengths are consistently slowed down more by light and thus bent more under dispersion, the concept of a rainbow of colors in a particular order is very familiar. This always results from some form of "white" light, which just means a collection of many different wavelengths of light traveling together, passing through one or more refractive interfaces. In the classic example of a rainbow in the sky, it is raindrops which are the refracting agents. Newton's famous experiments with color used a prism to accomplish the same feat.

15.2.4 Total Internal Reflection

An interesting result of Snell's Law is that below a certain angle θ_C, the solution for a transmitted ray would require that the sine function take a value greater than 1. Since this is mathematically impossible, we interpret this result to mean that below that critical angle, all rays are reflected rather than transmitted. This results in a number of fun results discussed in your text you should look at more carefully. In particular, it is this phenomena which allows for fiber optic cables and their myriad applications. It is total internal reflection which allows them to behave as near-lossless tubes of light.

15.3 Thin Lenses

Just as we were able to find more interesting behavior by taking the flat reflective surface and curving it, if we give a refractive material a curved surface we find an even more rich set of phenomena to investigate. Also as in the case of curved mirrors, however, we have to make some assumptions about our curved refractive objects (which I will henceforth call *lenses* because that's all they are). Like with curved mirrors, we would have parabolic surfaces in an ideal world. However, this is extremely hard to produce and so it is far more common to use spherical lenses. In the case of mirrors this meant we had to consider mirrors which were very small relative to their radius of curvature such that the rays were paraxial. In the case of lenses, we similarly deal with lenses which are much smaller than their radius of curvature. This results in lenses which are thin, and so we refer to lenses which satisfy this assumption of large radius of curvature as *thin lenses*. For our purposes, assume all lenses are thin unless otherwise specified.

There are six different ways to construct a thin lens: each of 2 sides can be either concave, flat, or convex. Two of these are the same by this criteria (I've counted having 1 side concave and 1 convex twice), but they behave differently depending on whether the concave or convex surface has the larger radius f curvature.

Figure 15.11: 33-2

We can use our same strategy of ray tracing to illustrate the behavior of thin lenses, and many of the same principals apply. Snell's law in particular tells us how rays will bend when passing into and out of a given lens. We also have focal lengths for lenses as well as mirrors. A (thin) lens has the property that any set of incident parallel rays will be focused into a single point on the focal plane, or alternatively caused to diverge such that they can be backtraced to a single point on the focal plane (real image vs. virtual image). We identify the focal length f of a lens as the distance from the center of the lens to the focal plane.

Figure 15.12: 33-3

Lenses which are thicker at the center will cause these parallel rays to converge to a real image (*converging lens*), while a lens with a thin center will cause them to diverge (*diverging lens*). The focal plane of a diverging lens is on the same side as the incident rays, while for a converging lens it is on the far side relative to incoming waves.

Figure 15.13: 33-5

Different uses call for different ways of describing the behavior of a lens, and so often the *Power* P is quoted instead.

$P = \frac{1}{f}$ $\qquad\qquad$ $[P] = \frac{1}{m} = D = diopter$

This is *not* the same as the magnification of a lens. A larger power corresponds to a larger magnification, but the relation isn't simple and direct.

15.3.1 Ray Diagrams for Thin Lenses

As for mirrors, we can characterize the behavior of a lens with a ray diagram. Also as with mirrors, we can use a few special rays to get the job done without drawing a bunch via trial and error until we figure out what we are doing.

Figure 15.14: 33-6

1. Draw a ray starting at the top of the object, parallel to the central axis (which passes through the 2 focal points) of the lens. Rays parallel to the central axis will be bent through the far focal point.

2. Draw another ray starting at the top of the object, but this time passing through the near focal point. Once this ray reaches the lens, it will be bent parallel to the central axis (the reverse of ray 1 again)

3. Draw a third ray from the top of the object, but this time it should pass straight through the center of the lens. In the thin lens equation, this passes straight through because at the center the lens looks, to the

incident ray, like a uniformly thick sheet. There would be a small shift (as well light passes through a pane of glass), but this is proportional to the thickness and we are in the "thin" lens approximation

Once again, these 3 rays should converge on a single point and we can construct an image in the same fashion as we did with mirrors. Because the rays are passing through the image, it is in this case a real image.

The same process produces a virtual image for a diverging lens.

Figure 15.15: 33-8

15.3.2 Magnification

Using simple geometry and traced rays, we can develop relationships between the positions and sizes of our object and image as we did with mirrors. This time we will develop the lens equation, rather than the mirror equation (creative names, I know).

Figure 15.16: 33-9

Our variables have much the same meaning this time as last. Once again we use similar triangles to find our relationships. $II'F$ is similar to ABF and $AB = OO' = h_o$, so

$\frac{h_o}{h_i} = \frac{f}{d_i - f}$

$OO'A$ is similar to $II'A$ as well, giving us another equation:

$\frac{h_o}{h_i} = \frac{d_o}{d_i}$

Once again this gives us

$\frac{1}{d_o} + \frac{1}{d_i} = \frac{1}{f}$ *The Thin Lens Equation*

Hey! Its the same as the mirror equation! That's because the same geometry is at play in both cases, only one case uses reflection to accomplish it and the other uses refraction.

We have some signs to deal with when we do the same thing for diverging lenses (go ahead and work through the diagram)

1. Focal length is positive for converging, negative for diverging lenses

2. Objects on the side of the lens from which light is coming have positive distances, opposite side is negative
3. Image distance has the opposite sign convention as object distance. This should always work out to be positive for real images and negative for virtual images.
4. Height, as before, is positive for upright and negative for inverted images.

Once again, we define the magnification as

$m = \frac{h_i}{h_o} = -\frac{d_i}{d_o}$

If we look back at the definition of power and the lens equation, we find that

$P = \frac{1}{d_o} + \frac{1}{d_i} = \frac{d_o+d_i}{d_o d_i} = \frac{1+\frac{d_i}{d_o}}{d_i} = \frac{1+|m|}{d_i}$

So, as I stated before, a greater magnification (ignoring the issue of inversion) yields a greater power, but the relationship is nontrivial.

15.4 Applications

There are tons of applications for this stuff, and we have limited time. I am going to concentrate on two applications, but if you are interested at all I suggest you read the sections on cameras, telescopes, etc. that I'm skipping. I won't test on it them, but some of these bits are much more directly applicable to your everyday lives all on their own.

15.4.1 The Human Eye

Since we use our eyes to "see", which means to perceive images, it shouldn't be surprising that they function similarly to a camera. Our eyes are ridiculously good when measured against man-made cameras, but they do share many of the same structures.

Figure 15.17: 33-25

The eye consists primarily of a lens, the retina, the iris, and important support structures. Support barely begins to describe the functions of these other structures, however. The vitreous and aqueous humors have precise indicies of refraction which allow the lens and cornea to bend light just the right amount. Without this jump in the index of refraction, our eyes couldn't possibly focus on such a wide range of distances while maintaining their tiny physical size. Likewise the ciliary muscles are absolutely essential as without them we couldn't deform the lens and focus. It is interesting to note, incidentally, the huge difference in how the eye focuses relative to man-made cameras. All standard cameras focus by changing the position of lenses and the focal plane. The eye simply modifies the lens itself. This allows for an enormously more compact, faster system. Our eyes are able to change focus from near to far with an almost imperceptibly short delay (its easier to see when you are extremely tired) and almost never get confused. Even a high quality camera with a high quality lens cannot hope to accomplish the feat so well, without even considering the problem of poor light conditions. Anyone who has struggled to computer on film or flash card a dim scene can recognize how badly our eyes outperform our cameras.