The term `f-stop` is most important in photography. It is endlessly discussed in the many books, lecture notes, and videos that teach and explain digital cameras and photography. However, the f-stop is also one of the least understood and most confusing terms in photography. Beginners tend to get confused by its precise definition and by the inverse relation between f-stop numbers and aperture sizes. This short document (prepared as auxiliary material for *The Computer Graphics Manual*) employs basic mathematics and graphics to explain f-stops.

In English we say "let's go take a picture," but in the past people were more likely to use the phrase "let's have our picture made." The word `take` implies that the picture somehow exists in space and time and only has to be taken (grabbed) by the camera. In contrast, `make` implies that the photographer has to *do* something in order to produce a good picture. Specifically, he has to think and make decisions. With old cameras, the photographer had to make all the decisions, but with modern, fully automatic digital cameras, we mostly have to decide whether we want the camera to make *all* the decisions, some of them, or none. Clearly, the more decisions we make, the more original and creative a picture we may produce. This is why even with a fully automatic camera, it still makes sense for the serious photographer to understand the basic decisions required to produce the perfect shot.

There is a vast amount of literature on photography, which is why this short document discusses only the f-stop, its definition, and its applications.

An f-stop (or f-number) is defined as the ratio of the focal length f of a lens to its aperture d (the diameter of the opening of the diaphragm, pupil, or iris placed inside or next to the lens). Thus, f-stop $= f/d$ is a pure (dimensionless) number. It is written as f/xxx, where xxx is a number. For example, $f/2.8$.

Recall that the film or image sensor of a camera is placed in the focus plane, which is normally located behind the focal plane of the lens. (This is explained on page 1233 of the book and is illustrated in figure 26.18b.) As the object to be photographed is moved away from the camera, the camera has to be refocused. This is done by moving the sensor closer to the focal plane (in practice, the lens and its focal plane are moved closer to the sensor). When the object is at infinity, the sensor has to be located at the focal plane in order for the image to be fully focused. Thus, the focal length of a lens is the distance from the lens to its focal plane, not the distance from the lens to the sensor (which is located at the focus plane).

If we denote the focal length of a lens by F, the distance between the object and lens by B, and the distance between the lens and the focus plane by S, then the thin-lens equation (an approximation for any lens that is much thinner than its focal length) relates the three quantities by

$$\frac{1}{B}+\frac{1}{S}=\frac{1}{F}.$$

If B is infinite, then $S = F$. If $B = F$, then S is infinite. For B values greater than F, S is greater than F. The derivation of this equation is shown later in this document.

We digress for a minute in order to explain the focusing scale found on lenses, which is always puzzling to beginners (and even after it becomes familiar to experienced photographers, it still remains mysterious). The figure shows a typical focusing scale and it is easy to see that the distance on this scale from 1.2 m to 1.5 m is about the same as the distance from 5 m to infinity; the scale is very nonlinear.

This behavior results from the nonlinearity of the thin-lens equation. When S is isolated in this equation it becomes $S = B \cdot F/(B - F)$, which is nonlinear in B. The *Mathematica* code listed here computes and prints S for several B values.

```
F = 0.1;
Do[Print[B F/(B - F)], {B, {0.2, 0.3, 0.4, 0.5, 1, 10, 100, 1000}}]
```

Running this code results in the sequence 0.2, 0.15, 0.1333, 0.125, 0.1111, 0.10101, 0.1001, and 0.10001. It is easy to see that small B values result in very different S values, while large and very large B values simply move S closer to F in smaller and smaller steps. (End of digression.)

For simplicity, we assume that the diaphragm opening of a lens is circular (it is often a pentagon or a hexagon). The area of a circle is $\pi r^2 = \pi(d/2)^2 = (\pi/4)d^2$; it is proportional to the square of the diameter d. Given a circle of diameter d and area a, it is easy to see that the circle of area $2a$ must have a diameter of $d\sqrt{2} \approx 1.414d$. Similarly, a circle of area $4a$ must have a diameter of $d\sqrt{4} = (\sqrt{2^2})d = 2d$, and a circle of area $8a$ must have a diameter of $d\sqrt{8} = (\sqrt{2^3})d \approx 2.8d$. The conclusion is that circles of doubling areas a, $2a$, $4a$, $8a$, … must have diameters that grow as $\sqrt{2^i}$ (or, equivalently $2^{i/2}$). Successive circles have diameters $\sqrt{2^0}d = d$, $\sqrt{2^1}d \approx 1.4d$, $\sqrt{2^2}d = 2d$, $\sqrt{2^3}d \approx 2.8d$, $\sqrt{2^4}d = 4d$, $\sqrt{2^5}d \approx 5.6d$, $\sqrt{2^6}d = 8d$, $\sqrt{2^7}d \approx 11d$, and so on.

Each time the area of the diaphragm is doubled, the amount of light entering the lens is also doubled. This is why the sequence of numbers 1, 1.4, 2, 2.8, 4, 5.6, 8, 11, 16, 22, 32, 45, 64, 90, 128, 180, 256, … (i.e., the powers of $\sqrt{2^i}$) is a convenient measure of consecutive steps of exposure. We refer to these steps as stops, where the term `stop` means doubling or halving the amount of light that arrives at the sensor. However, because the aperture is in the denominator, small f-stops correspond to more light, while large f-stops correspond to less light. This is the confusing aspect of f-stops. Perhaps the best way to think of an f-stop is “the larger the f-stop, the more light stoppage there is.”

Exposure: The amount of light that is recorded by the image sensor.

The numbers in the sequence above seem rather arbitrary, but are actually easy to derive and memorize. If we denote them by f_i, then $f_1 = 1$, $f_2 = 1.4$, and $f_{i+2} = \sqrt{2^{i+2}} = \sqrt{4 \cdot 2^i} = 2\sqrt{2^i} = 2f_i$, for $i = 3, 4, \ldots$.

So far we have considered only the aperture of a lens, but the f-stop also depends on the focal length f of the lens. This dependence implies that the f-stop is a pure (dimensionless) number, one which has no units. We now explain why the f-stop is defined as f/d. Two explanations are provided. The first one considers the drop in the intensity of the light as it travels through the lens and the second one relates the f-stop to the exposure time in an intuitive way.

1. The focal length of a lens affects the amount of light reaching the sensor because the intensity of light falls off as the square of the distance it travels. As a simple thought experiment, consider a lens with a large aperture, say one meter. A huge amount of light enters this lens, but what if the lens is 100 km long? Looking through the end of the lens, we would see nothing. The light would have lost virtually all of its energy in its 100 km journey.

Here is a practical example. We consider two lenses with f/8 and with focal lengths of 24 mm and 80 mm. Their apertures are $24/8 = 3$ mm and $80/8 = 10$ mm. The latter lens admits more light, but we have to consider the well-known fact that the intensity of light is inversely proportional to the square of the distance it travels. When we consider this fact, the results are surprising.

The first lens has a radius of 1.5 mm and an area of $\pi 1.5^2$. The light travels 24 mm inside this lens, so its intensity drops by a factor of 24^2. At the end of the lens we measure a light intensity proportional to

$$\frac{\pi 1.5^2}{24^2} \approx 0.012271484375.$$

The second lens has a radius of 5 mm and an area of $\pi 5^2$. The light travels the much longer distance of 80 mm, so its intensity drops off by a factor of 80^2. At the end of this lens we measure a light intensity proportional to

$$\frac{\pi 5^2}{80^2} \approx 0.012271484375.$$

The conclusion is obvious. In order for the f-stop of a lens to be a true measure of the light intensity reaching the sensor, its definition should include the focal length f of the lens.

2. This explanation relates the f-stop to the exposure time in a simple way. Figure 1 shows a lens and the relation between the subject size h_0 and image size h_i. The figure makes it clear that the magnification M of a lens is given by

$$M = \frac{h_i}{h_0} = \frac{d_i}{d_0}.$$

The thin-lens equation above provides the relation

$$\frac{1}{f} = \frac{1}{d_0} + \frac{1}{d_i},$$

from which we obtain

$$M = \frac{d_i}{d_0} = \frac{f}{d_0 - f} = \frac{d_i - f}{f}.$$

Thus,

$$h_i = h_0 M = \frac{h_0 f}{d_0 - f} = \frac{h_0(d_i - f)}{f}.$$

The image is a circle of radius h_i and we have shown that this radius depends on the focal length f in a complex way. However, a typical focal length of a camera lens is 50 mm, whereas d_0, the distance of the object from the lens, is normally measured in meters. Thus, we can write $d_0 - f \approx d_0$ and obtain

$$h_i = \frac{h_0 f}{d_0 - f} \approx \frac{h_0 f}{d_0} \propto f.$$

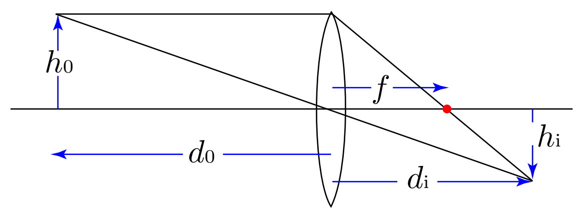

Figure 1: Lens Magnification.

The area of the circular image is $A = \pi h_i^2$, so it is (approximately) proportional to f^2. The amount of light reaching each unit area of the image equals the total amount L of light (which is proportional to the aperture area or to $d^2 = (2h_0)^2$, the square of the diameter of the lens) divided by the area A of the image (which is proportional to f^2. Thus,

$$\frac{L}{A} \propto \frac{d^2}{f^2}.$$

We now relate this result to the exposure time (or shutter speed T). Experiments, as well as our intuition, indicate that the exposure time is inversely proportional to the light per unit area of the image (L/A). Thus,

$$\frac{1}{T} \propto \frac{L}{A} \quad \text{or} \quad T \propto \frac{A}{L} \propto \frac{f^2}{d^2}.$$

By defining the f-stop as f/d, we obtain the important result

$$\text{Exposure time } T \propto \text{f-stop}^2.$$

Doubling T has the same effect as increasing the f-stop by one stop and halving T is equivalent to decreasing the f-stop by one stop.

Shutter speed: Exposure time; the length of time the shutter is open (notice the use of the word speed to indicate time).

The discussion of magnification also leads to a simple derivation of the thin-lens equation. The two colored triangles of Figure 2 are similar, which implies the following:

$$\frac{d_i}{d_0} = \frac{f}{d_0 - f} \Rightarrow d_i(d_0 - f) = d_0 f \Rightarrow d_i d_0 = d_0 f + d_i f \Rightarrow \frac{1}{f} = \frac{1}{d_i} + \frac{1}{d_0}.$$

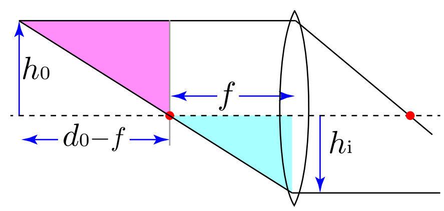

Figure 2: Derivation of the Thin-Lens Equation.

Circle and rectangle (a digression). An important point that has been ignored so far is that the lens is circular but the sensor is rectangular. Thus, the lens generates a circular image, parts of which are cropped when it ends up on the rectangular sensor. This means that some of the light sent by the lens, perhaps a significant part (gray in figure 3), never falls on the image sensor and is lost. The analysis here shows that the light loss is at its minimum (but is still a significant 36%) when the aspect ratio of the sensor is 1 (i.e., when it is a square). Rectangular sensors cause more light to be lost.

Figure 3 shows a rectangle of height h and width w (and an aspect ratio $a = w/h$) inscribed in a circle of radius R. The geometry is simple and yields

$$\frac{1}{a} = \frac{h}{w} = \frac{h/2}{w/2} = \tan\theta, \quad \text{so} \quad h = 2R\sin\theta \quad \text{and} \quad w = 2R\cos\theta.$$

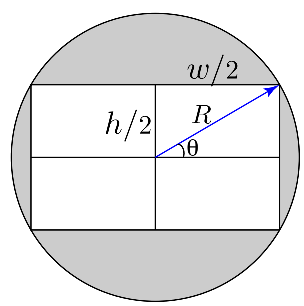

Figure 3: Rectangle in a Circle.

The ratio of the areas of the rectangle and the circle is therefore

$$\frac{h \cdot w}{\pi R^2} = \frac{4}{\pi}\sin\theta\cos\theta.$$

(Notice that it does not depend on the radius R.) The *Mathematica* code listed here computes this ratio for various values of the aspect ratio a and shows how the maximum ratio of areas (about 64%) is obtained for $a = 1$ (a square sensor) and goes down to 38% for $a = 3$.

```
w[a_] := Module[{t}, t = ArcTan[1./a]; 4 Sin[t] Cos[t]/Pi]
Do[Print[w[a]], {a, .25, 3, .25}]
```

Running this code produces the sequence of ratios 0.299586, 0.509296, 0.611155, 0.63662, 0.621092, 0.587649, 0.548472, 0.509296, 0.472543, 0.439048, 0.408924, and 0.381972. Thus, even for aspect ratios close to 1, much of the light collected and focused by the lens at so much cost and effort is lost.

Incomplete lenses. Because of practical considerations, lenses are not always manufactured at their maximum diameter. Figure 4 shows four lenses of the same focal length, but with smaller and small diameters. Only one of these lenses looks "complete," and the others are missing their outer regions. Such incomplete lenses have the advantage of being lighter and also easier to manufacture, but they have smaller effective diameters and therefore small maximum apertures. Such a lens should be specified by both its focal length and maximum aperture.

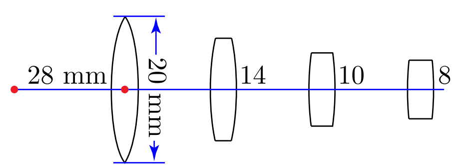

Figure 4: Four 28 mm Lenses.

Assume that these four lenses have a focal length of 28 mm and a full diameter of 20 mm, then the "complete" lens has a maximum f-stop of $28/20 = 1.4$, but the other three have bigger maximum f-stops because of their smaller sizes. If their diameters are, say, 14, 10, and 8 mm, respectively, then their maximum f-stops are $f/2.0$, $f/2.8$, and $f/3.5$ (because, for example, $28/8 = 3.5$). Note that these are the maximum f-stops. In addition, each lens can have bigger f-stops, corresponding to smaller diaphragm apertures. Thus, the last of these lenses, with a maximum f-stop of $f/3.5$, also has f-stops of $f/4.0$, $f/5.6$, and so on, up to a maximum which is determined by the smallest aperture its diaphragm can achieve (perhaps $f/16$ or $f/22$).

Clearly, the smaller the diameter of a lens, the less light it lets through into the camera. With diameters of 20, 14, 10 and 8 mm, the ratios of the areas of the lenses are $20^2/14^2 = 2.04$, $14^2/10^2 = 1.96$, and $10^2/8^2 \approx 1.56$.

The smallest f-stop of a lens depends on its maximum aperture, which is why it indicates the speed (slow or fast) of the lens. A fast lens is one that lets in more light because of a large aperture. It is fast because a large aperture generally allows a fast shutter speed. The 28 mm $f/2.0$ lens in our example is fast, but the same 28 mm lens with a maximum f-stop of $f/3.5$ is slow (a difference which is also reflected in the prices of the lenses).

Zoom lenses and f-stops. Currently (in 2012), even the simplest point-and-shoot cameras have zoom lenses, often with impressive zoom ranges. The range of f-stops in a zoom lens varies with the focal length. If the inscription on the lens says, for example, 25–125/4.5–22.5, then the smaller (wide angle) f-stops may be available only in the wide-angle range of focal lengths (around 25–50 mm), while the larger (telephoto) f-stops may be available only in the telephoto range of focal lengths (around 100–125 mm). Sometimes the inscription on the lens may be of the form `7.5-44.5 mm 1:2.8-4.8`. Here, f-stop 2.8 indicates the widest aperture that's available at the smallest focal length (7.5 mm) and f-stop 4.8 refers to the widest aperture at the longest focal length, 44.5 mm.

Wide angle: The widest aperture of a zoom lens (results in the widest field of view).

ND filters. Quite a few compact (point-and-shoot) cameras do not have an iris (diaphragm)! This surprising fact is hidden by camera makers, but is known to those exegetes who require the *complete* specifications of their cameras. In particular, the entire SD series of cameras made by Canon features an ND (neutral density) filter instead of a diaphragm. In conditions of much light, the computer in the camera sets it to a faster shutter speed. If this still results in overexposure, the computer swings the ND filter to cover the lens, which is usually equivalent to reducing exposure by three stops.

How can we tell whether a given camera has a diaphragm? Sometimes it is possible to actually see the iris and how it closes down when the camera is pointed at a source of strong light. If the lens is too small to see the iris, take several pictures with the same zoom setting and different lighting conditions, and then examine their EXIF metadata. In a camera with an ND filter, there will be just two apertures, for example, f/2.8 (ND filter out, not active) and f/8.0 (ND filter is active). Notice that the f-stop depends also on the focal length, so even without a diaphragm, a zoom lens features a range of f-stops when it is zoomed.

The advantage of ND filters is low cost. A diaphragm consists of several mechanical parts that have to be assembled and tested. The drawback is a constant depth-of-field.

Fractional stops. The stops discussed so far are numbers of the form $\sqrt{2^i}$ for $i = 0, 1, 2, \ldots$. They correspond to doubling or halving the area of the diaphragm. Many digital cameras also support fractional stops, which are interspersed between the traditional full stops, allowing for a fine scale of f-stops. Fractional stops are sequences of the form $\sqrt{2^{i \cdot j}}$ where j is a fraction, often 1/3, but also 1/4 or 1/2. In particular, the maximum aperture of a lens may be a fractional, not an integer, power of $\sqrt{2^i}$.

The following *Mathematica* code can be used to compute and list any fractional f-stops.

```
s = Sqrt[2]; j = .333;
Do[Print[NumberForm[Sqrt[2^(i j)], {2, 2}]], {i, -2, 10}]
```

Setting variable `j` to 0.333 produces the sequence of 1/3 f-stops 0.79, 0.89, 1.0, 1.1, 1.3, 1.4, 1.6, 1.8, 2.0, 2.2, 2.5, 2.8, and 3.2.

Small f-stops. Can an f-stop be less than 1? In theory, yes, and this results in very fast lenses, but for practical reasons it is rare to find a lens with an f-stop smaller

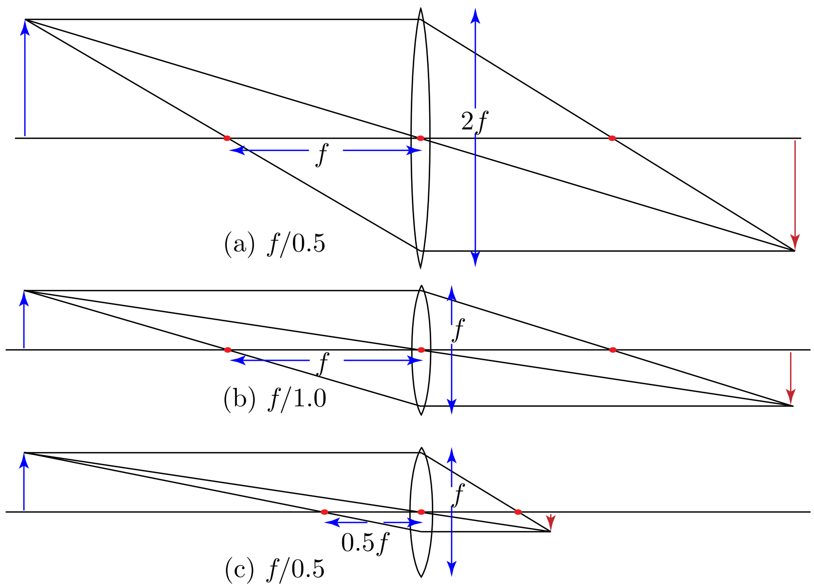

Figure 5: Small F-stops.

than 1. Figure 5 illustrates what happens when the f-stop of a lens is reduced from f/1.0 to f/0.5. In part (b) of the figure we see a lens with f/1.0. In part (a), an f/0.5 is obtained by doubling the diameter of the lens. The wider lens is heavier and results in a bigger, heavier, and more expensive camera. It is easy to see how the lens bends the peripheral light rays (those that hit the lens away from its center) more than in part (b). This makes it more difficult to design a compound lens (one that's made of several simple lenses) and also results in most of the light arriving at the image sensor at a sharp angle (obliquely). Image sensors lose sensitivity as the angle of incidence increases, resulting in an unevenly-exposed picture (the peripheral rays contribute to the total exposure less than the ones closer to the center).

The conclusion is simple but unexpected. The wider the lens, the more oblique light arrives at the sensor and is lost, a phenomenon also known as pixel vignetting.

The word vignette, from the same root as vine, originally referred to a decorative border in a book. Later, the word came to be used for a photographic portrait which is clear in the center, and fades off at the edges. Pixel vignetting only affects digital cameras and is caused by angle-dependence of the digital sensors. Light incident on the sensor at a right angle produces a stronger signal than light hitting it at an oblique angle.

`http://en.wikipedia.org/wiki/Vignetting`

Bending the light also increases chromatic aberration. The different wavelengths of light bend by slightly different amounts, which results in color bands or fringes in the final image. (Recall Isaac Newton's classic experiment where light was separated into the rainbow colors when bent by a prism.)

In part (c) of the figure, f/0.5 is achieved by halving the focal length of the lens. In addition to the sharp bending of the peripheral rays, the image plane is now located very close to the lens. In a DSLR camera, the mirror is located between the lens and the sensor, which is why enough space must be left for the mirror to swing up and down. In such cameras, the lens cannot be very close to the image sensor. (This problem does not exist in mirrorless cameras, such as the recent 4/3 and micro 4/3 types.)

Another limitation to small f-stops is the lens mount, as illustrated in Figure 6. Both Nikon and Canon lens mounts are shown and it is obvious that it requires ingenuity to get a wide beam of light past the lens mount and into the camera. (Notice especially the electrical contacts. They are used by the camera to control the lens, but they take space.)

Figure 6: Nikon and Canon Lens Mounts.

Because of these reasons, constructing lenses with f-stops below 1 produces diminishing returns. We lose the advantage of a bigger aperture while increasing the cost of design and manufacture. An f/0.7 does not send twice the light to the sensor as f/1.0. Similarly, an f/0.5 (the next stop down) sends much less than twice the light of f/0.7.

In spite of this, some users demand fast lenses and lens makers respond to the demand. Two examples of very fast lenses are the Leica 50mm Noctilux-m lens, featuring f/0.95, and the GOI CV-catadioptric lens 20mm, with f/0.5.

The smallest possible f-stop. In this section we try to answer the question what is the smallest theoretical f-stop?

We start with the concept of curvature. Figure 7a shows a curve that starts (on the left) as a straight line (which we intuitively associate with zero curvature) and becomes highly curved as it propagates to the right. The curvature of a curve is a useful measure in differential geometry and it is defined at any point where the curve is smooth (it has a derivative) as follows. Given a point P_1 on the curve, we select a point P_0 on one side of P_1 and a point P_2 on the other side. We then construct the unique circle that passes through the three points (Figure 7b). We now slide both P_0 and P_2 toward P_1 while recomputing the circle. The circle that is obtained in the limit of this process is the osculating circle to the curve at P_1. The radius of this circle is the radius of curvature of the curve at that point.

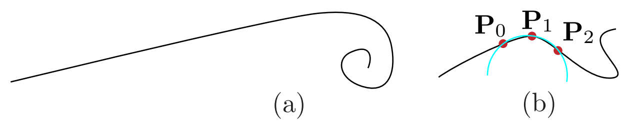

Figure 7: Radius of Curvature.

(The curvature itself is defined as the inverse of the radius and is a vector. It points from P_1 to the center of the osculating circle.)

A lens has two sides. Often, both are curved and have the same radii of curvature, but because they are curved in different directions, the radii have opposite signs (Figure 8).

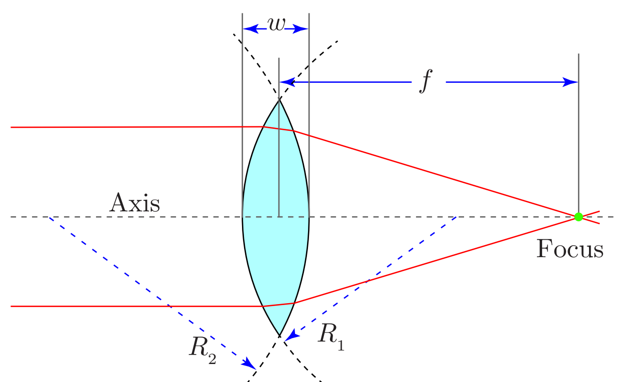

Figure 8: Radii of Curvature of a Lens.

Next, we need the concept of refractive index. The refractive index (or index of refraction) n of an optical medium is a number that describes how light is bent when entering the medium. The simplest way to understand refractive index is by means of Snell's law of refraction (Figure 9), $n_1 \sin\theta_1 = n_2 \sin\theta_2$, where θ_1 and θ_2 are the angles of incidence of a ray moving between two media with refractive indices n_1 and n_2. The refractive indexes of air, glass, quartz, sapphire, and diamond are 1, 1.5, 1.644, 1.77, and 2.42, respectively.

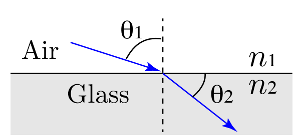

Figure 9: Snell's law of Refraction.

Armed with this background, we look at the lensmaker's equation. We denote the refractive index of air by $n_1 = 1$ and the refractive index of glass by $n_2 = 1.5$. Given a glass lens with radii of curvature R_1 and R_2, the lensmaker's equation expresses the focal length f of the lens in terms of these refractive indexes and radii

$$\frac{1}{f} = (n-1)\left[\frac{1}{R_1} - \frac{1}{R_2}\right],$$

where $n = n_2/n_1 = 1.5$. (The full equation also includes a term proportional to the width w of the lens, but we assume a thin lens and ignore this term.)

We are interested in a lens with the smallest possible f-stop, which is obtained by the largest possible focal length f. The lensmaker's equation implies that large radii of curvature yield large f. The largest radii of curvature are obtained when the lens is a sphere of diameter d. In that case, $R_1 = -R_2 = d/2$ and the lensmaker's equation becomes

$$\frac{1}{f} = 0.5\left[\frac{2}{d} + \frac{2}{d}\right] = \frac{2}{d}.$$

The f-stop of this lens is $f/d = 0.5$, so this is the smallest theoretical f-stop. It is clear that lenses made of materials denser than glass can have, at least in principle, smaller f-stops. Also, immersing a lens in a gas or liquid with $n_1 < 1$ may decrease its minimal f-stop even more.

Sharpness of a lens. A shallow depth-of-field is sometimes desirable, in order to emphasize certain objects in the photo. In general, however, we are interested in sharp pictures; those parts of the image that should be in focus should be as sharp as possible. This raises the question how does the sharpness of a lens (and with it, the quality of the images taken with the lens) depend on the f-stop?

Here is the theoretical background. Generally, decreasing the aperture increases image sharpness. This is especially true for out-of-focus objects, and for cheap lenses, because distortions caused by spherical aberration (a spherical lens is easy to manufacture, but it does not produce a sharp focus), coma (in an imperfect lens, light entering at an angle is not focused to a point), and astigmatism (the image of a circle is an ellipse) are less apparent at smaller apertures. The expensive, high-quality lenses available today suffer less from these distortions, and normally produce the sharpest images at about one or two f-stops below the maximum aperture. At smaller apertures, however, sharpness is reduced because of diffraction (spreading out of a light wave that is leaving a small opening).

That said, a search of the professional literature and of the many relevant articles found on the Internet yields many conflicting opinions. Some claim that a lens is at its sharpest (at its "sweet spot") in the middle of its f-stop range. Others claim that the sharpest f-stop is two stops above the middle of the range or is the third f-stop in the range. Still others claim that there is no direct relation between the f-stop and the sharpness of a lens and that each lens produces the sharpest images at a different f-stop. The only fact on which there seems to be a consensus is that a lens is not at its sharpest in its extreme f-stops (because of light dispersion and diffraction effects).

Lenses were already known to the ancients, but in modern times they became popular in the late 1200s with the invention of eyeglasses (spectacles). Grinding and polishing lenses for eyeglasses was the start of the huge optical industry we have today. In contrast, photography started much later, in 1839, with the photographic processes of Fox Talbot and Louis Daguerre. Thus, lens making is older than photography, and lenses made today, even cheap ones, are sharp enough for most practical purposes. When a blurred photograph is produced, the user should suspect camera vibrations or object movement rather than poor lens quality.

The following quotation, from June 3, 1937 (by Ansel Adams in a reply to Edward Weston's request for lens suggestions) shows that even as early as 1937 photographers had no need to worry about the sharpness of their lenses.

> Any good modern lens is corrected for maximum definition at the larger stops. Using a small stop only increases depth.
>
> —Ansel Adams, *An Autobiography*, 1985, page 244.

The obvious conclusion is that the sharpness of a given lens should best be determined by a simple test such as the following. Take a large sheet with a printed image that has lots of small, sharp details (a printed text, perhaps from a large book or a newspaper, may be ideal). Tape it to a wall in a quiet room free of vibrations. Place the camera on a tripod in front of the wall. Adjust the exposure and focus, and lock the mirror (if the camera is a DSLR). Take several pictures with a range of f-stops, using a shutter release cable or the camera's timer. Print the results, place them side by side on a large table, and compare them visually. If the test image contains small, sharp details, such as small text, it is generally easy to decide which test picture (or pictures) is the sharpest. Repeat this test for each lens that you own.

F-stop Rules. Once f-stops and their properties are understood, it is easy to come up with a few simple rules on how to select the f-stop (and shutter speed) of a camera for shooting various subjects. Casual photographers tend to use point-and-shoot cameras and most of those are automatic. The user has very little control over aperture and shutter speed, which is why the rules listed here do not apply to these cameras. On the other hand, DSLR, 4/3, and micro 4/3 cameras support non-automatic modes and offer the following important modes (the figure shows a generic mode dial with the P, S, A, M modes and several automatic modes):

- P (program) mode, where the camera selects both the aperture (f-stop) and shutter speed automatically. This mode is not considered automatic because the camera selects only the exposure (aperture, shutter speed, and perhaps also ISO), while other parameters—such as the light metering mode, exposure compensation, white balance, drive mode (single shot or continuous), self-timer, raw or JPEG output, and flash—can still be set manually by the user. Thus, the user is given a measure of creative control in this mode. In contrast, the auto mode leaves (almost) no control to the user and any camera set to this mode behaves as a a point-and-shoot camera.
- S or Tv (shutter priority or time value) mode, where the user selects the shutter

speed and the camera chooses the correct f-stop. A camera may offer many shutter speeds, but any given lens may have only a few f-stops. Thus, if the camera cannot find the appropriate f-stop for the shutter speed selected by the user, it indicates this by selecting the best f-stop it can and flashing it on the camera's LCD screen.

- A or Av (aperture priority or aperture value) mode, where the user selects the f-stop and the camera chooses the correct shutter speed.
- M (manual) mode, where the user is fully responsible for the exposure. The user selects the aperture, shutter speed, and ISO (the latter is not discussed here).

Some photographers prefer the program mode for general photographic work because of the following:

- You can set the camera to compute exposure on the whole scene or only a certain spot (such as a face or a flower).
- You can tell it if you need to focus continuously on a moving object (servo focus) or on a stationary one.
- You can set the camera to a burst (take many pictures in a row) or a single picture.
- You can manually control the use of the flash.

Other photographers may not care for the program mode. A typical claim is "it does not leave me enough creative control over the exposure."

If these modes are available on your camera, then the following simple rules might be useful as guidelines:

Rule 1. If the subject is moving, there are three choices, freezing it, blurring it, and panning the camera.

1.1. To freeze a moving object, use the S mode, set the shutter speed to 1/1000 s or faster (slow speeds may be used if the object moves slowly, such as an elephant), and let the camera select the correct f-stop. It has already been mentioned that a camera may offer a large selection of shutter speeds but any given lens may have only a few f-stops. If the lens does not offer the ideal f-stop, the camera will indicate a wrong exposure.

1.2. A blurry subject is sometimes considered artistically superior to a frozen subject. Use a slow shutter speed and let the camera select the correct f-stop. Because the subject will be blurry, the background should be sharp, so choose the M mode and select a small aperture (large f-stop).

1.3. An experienced photographer may be able to pan the camera and keep it steady on the moving object. The object will come up sharp and well focused, while the background will be blurry. Shutter speed must be slow, perhaps 1/15 or 1/30 of a second, and the f-stop should be set to an intermediate value, to achieve maximum sharpness. Depth of field is irrelevant in this case, because the background will be blurry anyway.

Rule 2. If the subject is stationary, select the Av mode and one of the following options:

2.1. If a shallow depth-of field is desired (such as when shooting a portrait or macro), select a small f-stop (i.e., large aperture) and let the camera choose the shutter speed.

2.2. If a large depth-of field is desired (such as when shooting a landscape), select a large f-stop (i.e., small aperture) and again let the camera choose the shutter speed. A tripod may be needed if the shutter speed is slower than about 1/60 s.

2.3. If depth-of field is unimportant (such as when shooting distant objects), select a medium f-stop, for best sharpness, and let the camera choose the shutter speed.

Photography is an art form; it does not exist merely to duplicate reality. Thus, there is more to photography than just exposure (aperture, shutter speed, and ISO). The rules above should be considered rough guidelines for beginners, illustrating how the f-stop and shutter speed are employed in general photographic work to obtain images that are technically correct. In order to "graduate" from a mere photographer to a photographic artist, a person must master composition, which is the artistic side of photography. An aspiring photographer should study the rules and techniques of image composition, but the best way to master it is to practice. For more information on the technical and artistic sides of photography, consult the many books and training material available, and then go out and shoot, shoot, shoot. This is the way to become a great photographer.

Disclaimer. Digital photography has released us from the toil, chemicals, red eyes, running nose, and darkness of the darkroom and so has started a revolution in photography. More people than ever buy cameras (mostly of the automatic, point-and-shoot type) and use them extensively. Because of this, a vast amount of literature—in the form of books, websites, and training videos—currently exists. Most of this literature concentrates on the more expensive cameras and discusses exposure (aperture, shutter speed, and ISO) and the use of lenses. However, the basic point-and-shoot cameras are improving all the time, and many an enthusiastic photographer has discovered that these cameras can do virtually everything their large, expensive cousins can. They set themselves to the right exposure and only sometimes need a little help from the user in the form of exposure correction (or exposure compensation).

Thus, experienced amateurs often feel that there is no need for expensive cameras with manual modes and no need to become intimate with the basics of exposure. Simply point your little camera at your subject and shoot. Satisfied? you are done. Don't try to improve on your picture. Otherwise, use the exposure compensation feature of your camera to lighten or darken the next picture, and then reshoot. Canon cameras use the symbol ± 0 for exposure compensation, while Sony uses `EV` (exposure value) for the same thing. If the picture is still unsatisfactory, the white balance feature can be used to vary colors in many ways, from cool (blue, green, and white) to warm (orange, yellow, and red) according to personal preference.

References. The relevant Wikipedia article (`en.wikipedia.org/wiki/F-number`) has more information and historical notes on f-stops.

Questions, error corrections, and comments should be sent to `dsalomon@csun.edu`.

Dec 2012