

Resolving the halftone resolution issue: how many dpi does it take to make an lpi?

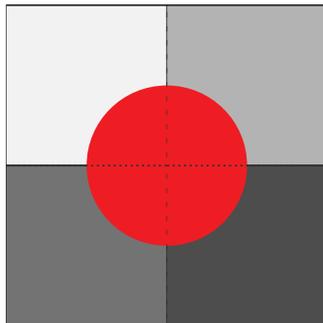
by Brian P. Lawler

When scanning an image, or working with images from digital cameras, one should understand the mathematics of halftone reproduction. It's easy math; it requires only multiplication or division to get the needed information.

As with any complex subject, *understanding the concept* is the most important factor in finding success. With resolution issues *the concept* is to scan an image with an *adequate* amount of information so that an excellent halftone image results. Too much – or too little – information in the scan can either ruin the image or cause a deadline to be missed and run-up unnecessary expense in production.

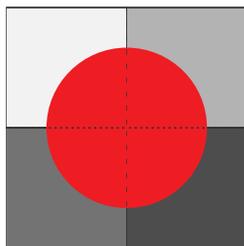
Digital images are elastic

A digital photograph, made up of pixels in rows (often called a *raster* image), can be reproduced large or small – these images are elastic – to a point. I often read or hear preposterous information about digital images, information that

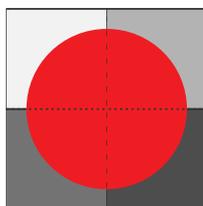


These images illustrate the relationship between halftone dots and the pixels in the file that created them.

We need to oversample pixels for conventional halftones – but to what degree?



Conventional wisdom dictates that the oversampling ratio should be 2:1. But many have disproved this theory, opting instead for lower ratios – values of 1.25:1 or 1.5:1 – which provide good images that print as excellent halftones, and are usually indistinguishable from halftones printed from higher resolution files.



is based on misunderstanding of the basic concepts of image elasticity and the relationship of images to their final halftone patterns.

Armed with the information in this handout, you will be able to work with images and turn

Oversample Rate	85 lpi	133 lpi	150 lpi
1.25:1	1.13 MB	2.74 MB	3.47 MB
1.5:1	1.61 MB	3.93 MB	4.98 MB
1.75:1	2.18 MB	5.34 MB	6.8 MB
2:1	2.84 MB	6.96 MB	8.85 MB

This chart demonstrates how an 8.5 x 11 inch full-bleed image changes in size according to its oversample ratio and halftone frequency. Curiously, the 8.85 MB and 4.98 MB files reproduce the same. Which would you rather store? Which would you rather transfer across a network?

them into excellent printed halftones. The resulting images will be effective for the desired purpose as long as you work within the practical limitations of image elasticity, and as long as you are careful to scan correctly and to work from good originals whenever possible.

The concept of oversampling

Oversampling is the process of scanning *more than one pixel* for each halftone dot in the final reproduction. The idea is sound: to prevent the duplication of a pixel value in halftone patterns (sometimes called *pixelization*), one must provide the halftone generating device with enough information as to provide a unique tonal value for each and every potential halftone dot in the final image.

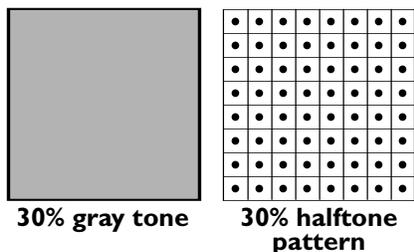
Certainly there are areas in a photograph where the pixels are of the same reflective value in the file, and the halftone dots in the the final reproduction will be the same. This is fine, as long as these similarly-toned pixels represent the tonality of the original image.

What we don't want to happen in an image is for a *single* pixel in an image to generate two or more halftone dots in the the final halftone. When this happens, you get jagged edges on diagonal lines (computer scientists call them *jaggies*). These are pixels that don't replicate the tones in the original image accurately.

Many theories exist on the issue of scanning, and most of them work. Whether you abide by the conventional wisdom of two pixels per halftone dot (2:1) or a lower relationship (1.5:1 or 1.25:1 for example), you will usually get a good image. Just

remember that there is no absolute right or wrong with regard to this process (despite what I say from time to time).

Another thing to remember is that there is such a thing as *too much image resolution*. Adobe PostScript RIPs, (and I assume those devices that mimic Adobe RIPs) set an image resolution limit of 2.5 times the halftone frequency, thus the upper limit is 2.5. Providing any more than that will result in useless processing in the RIP to reduce the oversampling ratio to 2.5 or less before processing the halftone.



Tonal values are represented by halftone dots as a percentage of printed area. The halftone dots on the right represent 30 percent of the possible printed area on the page, thus create the illusion of a 30 percent gray value. This ratio of inked-to-non-inked area is called “Dot Area.”

Conventional halftone dots

When we use halftone patterns, the printed image is creating a *simulation* of tonality. It’s really just an optical illusion that there are varying tones of gray or color on the printed page. In reality there is just one solid color of ink. To create the *illusion* of tonality, halftone dots are rendered with varying areas. Small black dots reproduce the light tones in the original, while larger black dots reproduce the mid-tone values in the original, and tiny white dots reversed-out of black reproduce the shadows in the original.

Our eyes and brain work together to average the reflective values of halftone dots into *perceived* tonality. We look at little dots of solid ink, and we *think* we see tonal variation, thus halftone patterns work. The halftone process was invented by lithographer Frederic Ives in 1878 (He was half of the famous *Currier & Ives* partnership).

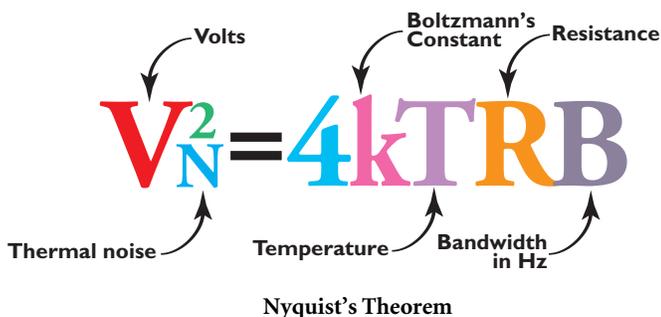
While most halftone patterns are based on variations of round or elliptical dots, it’s not necessary for halftone images to be made of dots at all!

Stochastic tonal patterns

Recent developments in screening technique have revived a system developed in the 1950s that uses

the same digital data to make dotless patterns for the description of tonal information.

Called *stochastic*, or *frequency-modulated*



screening, the technique is to disperse the spot components of the halftone pattern randomly to create the *illusion* of tonality. If an imagesetter makes a single halftone dot by drawing a cluster of spots in a circular pattern, what would happen if those same component spots were distributed randomly (or pseudorandomly) around the same cell area of the halftone dot – but not clustered together? Interestingly, these stochastic techniques result in the same image, but with no recognizable pattern of dots!

There are several advantages of using stochastic pattern images in print. First, these images have no screen angles, which makes it possible to print more than four colors to make a full-color image without creating moiré patterns.

Second, stochastic images appear to have more detail, and this is most likely a result of the information in an array of dots centering itself where the image had detail information, rather than in the center of the halftone cell as is the case with conventional electronic halftone images.

Nyquist's Theorem

Harry Nyquist was an electrical engineer who worked at AT&T Bell Labs from about the 1919 until his retirement in 1954. The theorem that carries his name was a watershed for engineers involved in signaling and signal-to-noise ratio calculations.

From 1912-1915, Nyquist did his undergraduate work at University of North Dakota, where he also received his Maser's degree. He then went to Yale, where he received his Ph.D. in 1917. Nyquist lived and worked in the era of telegraphy, and was employed after completing his doctoral work by AT&T, the American Telephone

Test E	Test A	Test D	Test B	Test C
1:1	1.25:1	1.75:1	1.5:1	2.0:1
USM: 50%	USM: 70%	USM: 75%	USM: 80%	USM: 100%



A conventional halftone sampled at 120 pixels per inch



A simulated stochastic halftone sampled at 75 pixels per inch



A conventional halftone sampled at 90 pixels per inch



A conventional halftone sampled at 75 pixels per inch

and *Telegraph* Company, at their research labs in New Jersey.

Prior to Nyquist's work, there was no mathematical standard for signal strength calculations. Radio engineers would transmit large amounts of energy to get a signal to a receiver successfully. Nyquist's research showed that twice the noise value needed to be broadcast to get a signal to a receiver successfully. Any more than twice the noise value was wasted.

It was a great relief for broadcast engineers to rely on Nyquist's Theorem, saving much uncertainty and energy in getting their signals from Point-A to Point-B through "the ether."

More is not always better

Many people, believing that if 2:1 yields a good halftone (it does) believe that 4:1 will yield *even better* halftones (it doesn't). Nyquist tells us that twice the sampling value is needed, and my testing shows that *any more than 2:1 is wasted*; halftone quality does not improve with higher sampling rates. Ratios beyond 2:1 simply waste time and money (Adobe engineers, fully aware of Nyquist, put the 2.5:1 resolution limit in PostScript for this exact reason).

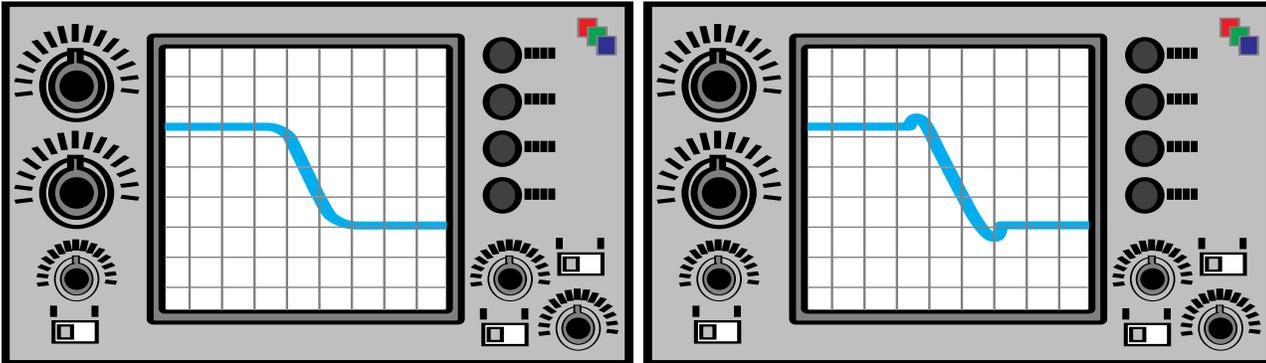
The *concept* of Nyquist's Theorem is valid: we need to sample *four pixels* per halftone dot (a 2:1 linear ratio) *or slightly fewer* in order to achieve success in halftone reproduction. But how many fewer?

Experimentation will demonstrate that ratios significantly lower than 2:1 work quite well. I have found that this is often determined by image content; *people-pictures* reproduce nicely with lower ratios, while highly detailed images of machinery, furniture or fabrics usually require higher ratios.

Oversampling ratio tests

One way to determine the best ratio of pixels to halftone dots is to test the process with a series of scans, and then print the results. Shown at the bottom of page 2 is a chart of recommended scans. They range from 1:1 to 2:1. Notice that each scan is named with a letter that does not relate to the

original (the same image is split by a prism and filtered through color separation filters). The fourth PMT sensor sees a monochromatic image, filtered only by a neutral-density filter, but with a larger aperture than the other three. The larger aperture gives a microscopic “advance look” at tonal changes in the image being scanned. These



An oscilloscope monitoring the neutral gray photomultiplier tube on a drum scanner sees tonality as a function of varying voltages. When the voltage changes quickly, an edge is sensed. By adding a momentary spike to the high and low points on the signal, the scanner changes the density at that edge by drawing larger and smaller halftone dots. The technique is called Unsharp Masking. The same technique takes place in programs like Adobe Photoshop by analysis of pixels and their neighbors. When tonal differences are found, the edges are enhanced by a computer algorithm that darkens and lightens the pixels along the edges of tonal change.

image size. This is handy when evaluating the images after printing. Prepare the test by scanning a good image at various resolutions.

Remember also that *apparent sharpness* is a factor in your resolution experiments. In the chart at the bottom of page 2 are recommendations for resolutions as well as unsharp masking amounts. Lower sample rates usually require sharpening to match the appearance of higher-ratio sampled images. Let's take a quick tangent to discuss unsharp masking....

Unsharp masking concept

Unsharp masking gets its name from a *masking film* that was used when color separations were made on process cameras and precision enlargers. The masking film was used to adjust the density of individual colors to compensate for the variations in density of the three color separation filters – red, green and blue – used to make separations. When masking films were used in contact exposures, the edges of tones were enhanced by the slightly unsharp emulsion of the masking film. Lithographers learned to use these masking films to modify the *appearance* of sharpness in the color separations they made; the process became known as *unsharp masking*.

Most drum scanners have an electronic variation on the theme of unsharp masking. High-end drum scanners feature four photo-multiplier tubes that analyze the tonality of the original image. Three of the PMT sensors see already-separated red, green and blue images of the

tonal changes are monitored by an unsharp masking logic circuit; when a tonal change of appropriate value occurs, the circuit modifies the other three signals to enhance the edge where the tonal change takes place (see illustration, above).

Thus unsharp masking is a synthetic enhancement of tonal edges to make those edges darker – or lighter – than they would have been if they had been left alone.

Almost every image intended for halftone reproduction is made better by some unsharp masking. The resulting images are visibly superior to those not sharpened.

Unsharp Masking in practice

Unsharp masking requires understanding the three variables in the Unsharp Mask control panel in Adobe Photoshop. The first element in the control is *Percentage*. The range is 1 to 500 “percent” (Photoshop's settings are odd compared the rest of the industry which uses a scale of 0-100), with reasonable values between 1 and about 200 percent.

In general, larger files need more unsharp masking. A five megabyte file will be sharpened nicely in the 75 percent range, while a 25 megabyte file will require about 125 percent for about the same visual impact.

The second variable in Photoshop unsharp masking is the *Diameter*, designated in pixels. A value of 1 is usually adequate here as larger numbers tend to draw outlines around things! As files get larger, you can increase the Diameter

setting to get an effective unsharp masking edge.

The third variable is the *Threshold* value. This setting determines how much tonal change must take place for an unsharp masking event to take place. With photos of delicate subjects, diffuse facial detail for example, a number in the 3 to 4 range is adequate. A zero value will sharpen every tonal difference, which may be too heavy-handed. You must be the judge of the correct setting for Threshold depending on the image.

Now, back to the scanning test...

The method is to scan at a variety of oversample ratios, then make halftone film from the images; proof the results and evaluate them in a blind test (where the resolution values are not known to the judges). All should be made from a good photograph, scanned on a capable scanner (you can substitute a Kodak Photo CD, and job-out the scanning in the process).

Choose a halftone frequency appropriate to your printing process. For commercial offset printing use 150 lpi; for newsprint, use 85 or 100 lpi. Quality web-printed publications (*Time* and *Newsweek* are good examples) typically use 133 lpi. If you are unsure of the correct halftone frequency, ask your printer to give you the best value for their process.

After completion, have a group of impartial judges view the results, voting for the image that seems to be the “best.” The results will likely surprise you.

If you really want to push this issue, scan a photograph of a person for one set of tests, and a photograph of wicker furniture (or something with lots of detail) for the other. You’ll then see the difference between the effects of resolution on images with different amounts of detail.

Time equals money

Time really is money in this business, as a minute saved in a typical prepress cost center is usually worth \$1.00 or more (average is \$1.25 per minute).

One advertising agency that ran this test found that the savings was really worth the effort. Prior to running the test, the firm was routinely scanning images at 3:1, resulting in full-page, full-bleed, four-color files of 79.7 megabytes!

After doing the test, and satisfying themselves that lower resolutions yielded equal results, they settled on a ratio of 1.5:1. These same ad files were reduced to about 19.9 megabytes, a savings of 60 percent *in every area of production* – disk space, network transfer time, and – most importantly, imagesetter RIP time.

If you could reduce your prepress storage, network communications and imagesetting time by 60 percent without a quality loss, wouldn’t you do it?

If you translate this firm’s savings into into productivity, this means that there is more disk space, more time available on the network, and more work completed in a shift. Translate the savings into dollars, and you discover that for every \$100 spent on the old plan there is now \$60 (*more than half* their money!) available for increased salaries, new equipment, and expansion – all better investments than digits needlessly written to a hard disk!

Running the tests takes about one business day. This may seem like a waste of time, but the result is well worth the effort.

Resolution math

Where do we start when it comes to calculating resolution sizes and scan ratios?

The math needed to calculate the resolution of any file is really quite simple. It is a three-element multiplication equation that yields a value for resolution. I’ll present three variations here for different situations:

1. You want to scan an image and need to know the resolution
2. You have an image already scanned (or a digital photograph) and need to know how large it can be in print for a given halftone frequency
3. You have a space to fill, and need to know how many pixels of data to scan or capture to fill the space.

First, the calculation for scanning resolution:

$$Q \times lpi \times \% = R$$

...where *Q* is the oversample ratio you have chosen from your test, used as a multiplier; *lpi* is the halftone frequency expressed in lines-per-inch, and *%* is the percentage of enlargement or reduction expressed as a decimal number. *R* is the resulting resolution. Following is an example of this formula at work. Use a pocket calculator to make the calculations:

A 5 x 7 reflective photo needs to be enlarged to 9 inches vertical (7 inches to 9 inches). We begin by calculating the enlargement by dividing *what we want* by *what we have*: 9/7 or 128% – which when expressed as a decimal number is 1.28.

The oversample value – *Q* – we will choose for this example is 1.5, that being a reasonable value that neither pushes the limits of quality, nor clogs the system with digital cholesterol. Our halftone frequency is 150 lpi.

Armed with these three variables, we can now complete our resolution calculation:

$$1.5 \times 150 \times 1.28 = 288$$

The Associative Law applies here, so you can

multiply these numbers *in any order* and still get the same result: 288. Use this value to set the resolution of the scanner (some scanner control panels allow you to set the percentage with a software control; for this formula to work, you must leave this setting at 100%).

Fixed-resolution calculations

With fixed-resolution images, the equation has the same elements, but the math is twisted around to accommodate the fact that we have no control over the scanner – the file has already been scanned.

Fixed-resolution images are those from CD-ROMs, files supplied by clients and service bureaus, stock photos, Internet-source images, and photographs from digital cameras.

We have an image; we need to know how large that image can be in print. The formula for fixed-resolution images is:

$$\frac{D}{(Q \times lpi)}$$

where *D* (in pixels) is the dimension you are calculating (I usually check the longest dimension first), *Q* is the oversample value, and *lpi* is the screen frequency in lines-per-inch. Following is an example:

We plan to use a digital photograph whose dimensions are 1960 x 3008. We are planning to reproduce this image in a full-page broadsheet newspaper ad (13.5 x 22 inches) in the local daily paper. The newspaper uses a screen frequency of 100 lpi. We will probably need to crop the image after we open it, so we'll see if this image is large enough. We'll use a *Q* value of 1.25 here (as low as I *ever* go with conventional halftones). Here's the math:

$$\frac{3008}{(1.25 \times 100)}$$

Mathematics rules tell us that we must execute the part in parentheses first, so we multiply the *Q* by the *lpi* (1.25 x 100) first, then perform the division. so we end up with:

$$\frac{3008}{125}$$

or 24.064 inches (vertical). This is the maximum dimension to which I can enlarge this image under these circumstances. When I test the other dimension (1960 pixels) I get a reproduction size of 15.68 inches, just enough to fill that broadsheet page (horizontally) with a small amount of room for cropping.

BUT YOU CAN'T DO THAT!!!

A person attending one of my seminars exclaimed one day that digital photos just can't be enlarged that much – but it's simply not true. Digital images are elastic; their ultimate size depends on the

variables described here: screen frequency, *Q* value, and number of pixels. If you have these facts, you can calculate the usefulness of any fixed-resolution image for any task – be it a postcard or a billboard.

When you have a space, but need to know the number of pixels

If you have a space in a page layout, and need to calculate how much of a digital file you need to fill the space, the formula is:

$$Q \times lpi \times inches = pixels$$

The product of this equation is the number of pixels needed to fill the space. Let's see an example:

We have designed a full-page ad, and allowed a space for a photo that is 9 inches tall by 3.25 inches wide. We are printing by web-fed offset, at 133 lpi. The *Q* value chosen is 1.41† and we will multiply them all together to get the number of pixels we need in the image.

$$1.41 \times 133 \times 9 = 1687.7$$

You can't have a partial pixel, so it's best to seek 1688 pixels (or 1700 to be more cavalier!). In this example we'll use a scanner to prepare the file: We're looking for an image with *at least* 1700 pixels in the vertical dimension to fill the 9-inch space.

Most scanners have controls for the resolution of the resulting image. We should set the scanner to produce an image with these dimensions, or slightly more, and the result will be satisfactory – or better.

■ This is one of a series of essays on the use of computers in the graphic arts. Others are available on a variety of subjects from black and white scanning to managing color images for the World Wide Web. You are welcome to download the others from my FTP site, and reproduce them for your associates and clients. You may not charge for them, nor change the content in any way.

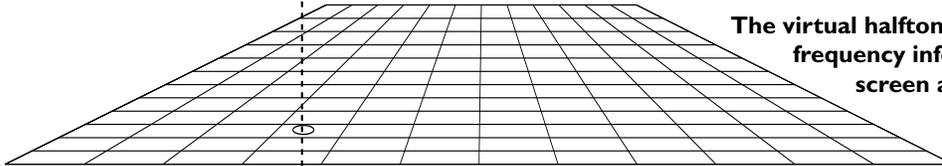
Brian P. Lawler
Graphic Arts Consultant
6045 Madbury Court
San Luis Obispo, California 93401 USA
Telephone: 805 544-8814
Fax: 805 544-8445
Email: brian@thelawlers.com

<http://www.thelawlers.com>

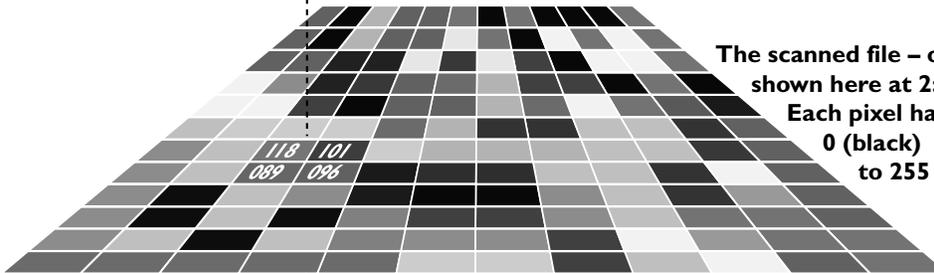
†1.41 is the hypotenuse (diagonal) of a square. Several industry experts have proffered this value as a valid *Q* number because it is the diagonal of a potential halftone cell. The value works fine, but I think that's a happy coincidence.



A “virtual” light source shines through a “virtual” halftone screen onto a grid of pixels that were scanned from the original image.



The virtual halftone screen contains frequency information and screen angle.



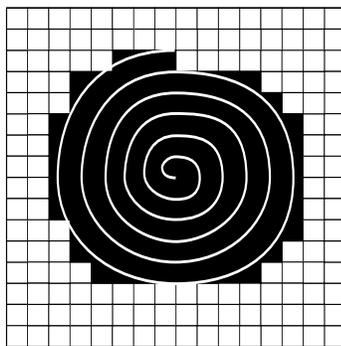
The scanned file – or bit-map – of the image shown here at 2:1 halftone frequency. Each pixel has a numeric value from 0 (black) to 255 (white).



Average of 118, 089, 101, 096 = 101
(or 101/256 or 39.45 white or 38.45% reflective black)



This average value is carried to the halftone generator...

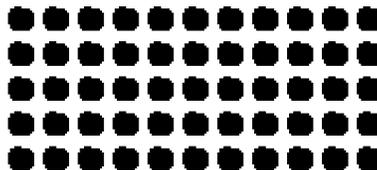


...where a spot function is applied to convert the value into a halftone dot of the required size and shape.

This example shows a spiral-shaped dot, growing from the center out, and the imagesetter spots that will be turned on to mark that dot on the film.

Different spot functions result in different shapes of halftone dots.

A pattern of these dot-formations creates a halftone pattern – the illusion of tonality – when they are small enough to be *confused* by our brain into tones.



...or at smaller sizes

