Info

C: 75%   C: 75%
M: 68%   M: 68%
Y: 67%   Y: 67%
K: 90%   K: 90%

X: 3.970   W:
Y: 3.129   H:

Doc: 17.M/11.6M

Click image to choose new foreground color. Use Shift, Opt and Cmd for additional options.

The Question of K:
A Primer on Gray Component Replacement

For a theoretically perfect inkjet printer, combining cyan (C), magenta (M) and yellow (Y) inks would produce a perfect black. The combined spectral reflectance of each ink would be sufficient to absorb virtually all visible light such that very little would be reflected back for your eye to detect. Reality, however, is not perfect. The substrates we print on are rarely perfect reflectors, and they typically reflect certain wavelengths more so than others. Likewise, our inks are not perfect spectral absorbers, so the pigments used tend to absorb some wavelengths more than others. When printed in even ratios, C+M+Y rarely produces a neutral gray — more common is a process gray with a sickly green or deep magenta cast.

That alone, however, is not an adequate reason to add black to the printing process. The C, M and Y ratios can be adjusted to produce a neutral process gray. Performing G7 calibration, for example, prior to profiling (characterization) will achieve that and produce a common visual appearance on virtually any four-color process device. So, why not just ditch the K and print with only the three primaries?

The question is a logical one. Three-color chromolithographic prints using photographic separations were introduced in Europe and the US in the latter part of the nineteenth century, based on earlier lithographic printing technology invented in Germany over a half-century before. Color lithographs were not uncommon throughout most of that century, but used well over a dozen or more colors that typically overprinted a black base image. Colors were often printed as solids, but as techniques, ink chemistry, equipment and registration methods improved, color tints were overprinted to produce intermediate colors and reduce the number of ink colors actually required on press.

Chris Brown, Engineering Manager, Mutoh

Technological Advancements

Mechanical separations using photographic plates became a common method of color printing by the beginning of the twentieth century. The earlier methods of color printing, which printed solid colors on a base image, employed opaque inks. Color printing using photographic separations required semi-transparent inks that could be printed one atop the other. Rather than a black base image, a white substrate was used and black was produced by stacking the three ink solids. A wide range of colors, including photographic reproductions, could be produced from these three primaries.

These were exponential advances in color printing, and their development in the late 1800s was no accident. The Scottish physicist James Clerk Maxwell had published his electromagnetic field theory in 1865, and made a number of advancements in the color photography field. Physicists at this point were coming to understand both light and color perception in much more intimate terms, which led to ever more innovative technologies to feed the growing demand for color printing (which continues today).

What came to be understood was that the human eye contains three types of cones, each with a unique peak responsivity that allows it to perceive color as three bands of spectral energy. These are referred to as long, medium and short, i.e. red, green and blue. When white light passes through a photographic plate, certain wavelengths are absorbed. For example, when red light is absorbed, the remaining green and blue light excites the medium and short cones to cause us to see cyan. This can be accomplished by otherwise blocking the red portion of the spectrum as it passes through a glass plate, or by absorbing the red light with a pigment as it reflects off a white substrate.

Cyan, magenta and yellow are the inverse of red, green and blue. Cyan ink printed on white paper absorbs red light, magenta absorbs green, and yellow absorbs blue. If you wish to convert white light to yellow, you remove the blue component. This is why incandescent lighting appears more yellow than compact fluorescent — the lower the color temperature (2800 vs. 5000 Kelvin), the less blue energy. It is also why the RGB LEDs found in novelty color changing lamps typically cycle red, yellow, green, cyan, blue, magenta, white and back

to red. The red, green and blue colors are produced by the LED's individual chips, whereas the intermediate colors are produced by combinations: Red and green yields yellow; green and blue, cyan; and blue and red, magenta (violet). Turn on all three chips and you see white. If measured with a spectrometer, you would see that these lamps do not produce yellow, cyan, magenta or white light, only red, green and blue. They are the perfect inverse analogy of three-color process printing.

An Imperfect Process

One problem, however, that becomes very apparent with the three-color process is that CMY-process black rarely achieves the density of an actual black ink. In colorimetric terms, the L* of the CMY black will not typically achieve the L* of the K black, which restricts the full dynamic range of the printer (whitest white to blackest black). This is the same for color monitors (referred to as contrast ratio), and film and digital photography (referred to as luminance range).

Achieving the maximum dynamic range of the printer helps to ensure the maximum color gamut is reproduced. In most cases, the media itself dictates minimum density, or Dmin. Whatever the media L*, it will define Dmin and any color printed will have a higher density than Dmin (which is why media density is generally calibrated to zero). By the same token, the black solid will define the maximum density, or Dmax, and any color printed will have a lower density than Dmax. Every printed color should have an L* value somewhere between that of Dmin and Dmax.

GCR: Gray Component Replacement

As mentioned before, CMY is the inverse of RGB — take any RGB value as a percentage between zero and one, subtract it from one and the result is CMY. For example, a 20 percent red (0.20) equates to an 80 percent cyan (0.80), a 40 percent green to a 60 percent magenta, or a 30 percent blue to a 70 percent yellow. This assumes, of course, that your CMY (ink plus media) is the colorimetric inverse equivalent of your RGB, but in reality that will likely not be the case. This is why a modern color workflow converts all colors to common colorimetric equivalents, a connection space such as CIELAB. For illustrative purposes and for the sake of simplicity, we will leave that function out for now.

Returning to our CMY value, we now need to calculate a K percentage that, when added to the CMY value, will produce a colorimetrically equivalent CMYK. For those colors to be equivalent, the K must replace some portion of the CMY that is colorimetrically equivalent to the K we are adding. Any color printed using all three primaries — C, M and Y — will have a gray (achromatic) component and a color (chromatic, or tonal) component. Likewise, any value of K should have a CMY gray equivalent, at least to the maximum density (Dmax) of the CMY grayscale. This should be true of any four-color process printer.

The solution to GCR, therefore, is to identify the CMY equivalent of K, and then substitute K for the CMY gray component of the given color. One way to accomplish this is to perform a grayscale calibration of the printer prior to profiling, and the most straightforward way to do that is to calibrate the printer to G7 Aims (ANSI/CGATS TR015-2013). G7 defines neutral print density for both K and CMY process grayscale, and process gray balance (the a*, b* component of the CMY grayscale).

In the case of print density, G7 employs two relatively simple formulas to determine the print density targets between Dmin and Dmax for both the K and CMY process grayscales. This means that the CMY gray density can be matched to an equivalent K gray density, preserving L* in the original color. G7 also establishes a neutral gray balance, defining by way of formula what constitutes a neutral gray for a given media white point; it compensates for the color cast present in a given ink and media combination. If your magenta were too red, for example (too much yellow), G7 calibration would reduce the yellow channel to pull the gray balance back toward blue. The opposite would be true if your cyan was too blue, pulling it back toward green (or vice versa). The point is that G7 calibration establishes a neutral balance by adjusting the C, M and Y curves to compensate for ink and media differences.

Once a neutral gray balance has been established through G7 calibration, the next step is to identify the difference between it and the chromatic bias of your K ink. Optimally, your K ink will have no bias, the a*, b* value will both be very close to zero. If so, it should also be very close to the gray balance of your G7 calibrated process grayscale. In such case, GCR is mostly a matter of substituting CMY gray with an equivalent density of K.

If not, you (or more likely your color management software) will first need to correct for the chromatic bias of the K ink by creating a set of C, M and Y curves that generate a process gray with the same chromatic bias as your K. This operation, or at least the internal mathematics, is identical to matching process gray to the G7 gray balance aims — you simply match to your K ink (a, b) instead. The difference between these three curves (the K bias curves) and your G7 calibration curves will tell you the amount of C, M or Y ink that needs to be added back to the K channel to achieve a neutral grayscale.

The amount of C, M or Y that gets added back, if any, will be small, unless the gray balance is deliberately adjusted for a particular bias. For example, in the analog print world, press operators sometimes add C, M and Y ink to the K channel to produce rich blacks that appear blacker-than-black, or blacks biased toward blue/cyan (cool black) or red/yellow (warm black). However, these rich black blends do not always translate well to digital inkjet and can be problematic to maintain against color drift on long production runs. This is often referred to as under-color addition, or UCA. The application of UCA (and its closely related cousin, under-color removal, or UCR) is not normally applied in conjunction with GCR, and is beyond the scope of this article.

So now, by way of G7 calibration, we have established a substitution equivalence between CMY and K, such that we can remove the gray component of any CMY value and replace it with a K that will change neither the hue angle (a*, b*) nor the density (L*) of the original color. We can use the G7 curves to determine the amount of C, M and Y to remove, and we can use the same curves to determine the amount of K to add and the amount of C, M or Y to add back. This also means that many CIELAB values coming out of profile connection space will have two CMYK matches: One with no K (CMY only), and one with K. The one with K, the result of GCR, will be the one we will use.

The Why of GCR

The main advantage to GCR is the reduction of ink usage. The amount of CMY we subtract should always be more than the amount of K we add, producing an overall net reduction in ink used. This

can be particularly beneficial for media that cannot hold 300 percent ink (the equivalent of a 100 percent process-gray solid). Not applying GCR leaves us with less-than-optimal choices: Restrict the primaries with overly aggressive limits that can degrade one- and two-channel colors (and the achievable color gamut), or apply secondary ink limits that make recalibration without re-profiling (the aim of G7-based process control) exceptionally difficult.

With GCR applied, however, CMY-process gray will be largely substituted with K, so a 100 percent process-gray solid should print with far less than 300 percent ink. In fact, it should be less than half that, depending on the amount of added C, M or Y. The same will be true of any other three-channel combinations — the CMYK equivalents will print with less ink than the CMY originals. Aside from the cost savings, that can also give us more vibrant color reproduction.

We often see a tendency for dark colors to fall toward black more quickly than we would like. Dark blues and violets, and even sometimes very dark reds and greens can turn black due to an oversaturation of ink; excess dot overlap results from high concentrations of ink, absorbing too much light and causing dark colors to lose all perception of color (or appear as a muddled gray). By reducing the amount of ink through K substitution, we maintain the same low L* without the excess dot overlap, allowing the chromatic component to remain discernable. This can result in smoother dark gradients with stronger color and more subtle shading.

Another problem with high ink concentrations is loss of detail. Excess ink in very dark areas can be drawn into lighter areas, leading to poorly defined edges with some fine-detail areas being completely flooded with unwanted ink. This and the poor reproduction of dark colors are the main visual objections to three-color process inkjet, and are generally described as a lack of contrast and poor shadow detail.

Beyond this, and possibly somewhat counter-intuitive, is the impact of K on the achievable color gamut. On most inkjet printers, the K solid density will exceed the CMY process solid density by upwards of 30 percent in some cases. This means that not only can K be used to substitute for the CMY gray component, but when mixed with higher ratios of two other channels (CM, CY or MY), can

produce colors not achievable with CMY alone, colors with lower L* values that still retain perceptible chroma. This is the result of having an expanded L* axis; the greater dynamic range causes the entire gamut to increase to as near its maximum volume as possible.

One caveat to GCR is its impact on very light colors, and the speckling that can occur if GCR is used to replace all CMY process gray in an image. Depending on the physical resolution of your printer, and the dot size used, the high density of the K ink makes its dots very obvious if printed at very low percentages. This is most obvious when low percentages of K are printed alone on white media, or in conjunction with high concentrations of yellow. For this reason, GCR may not be utilized if the process gray component is below a certain threshold — better to print a three percent gray with CMY rather than K alone, since the CMY will print in higher concentrations with less discernable dots.

The point at which GCR is applied is primarily resolution dependent, but because GCR will have minimal impact on colors with ink concentrations well below the media saturation point, the general rule is that starting a bit late is much better than starting too early. Once GCR is applied, it can be phased in through the mid-tone range so as to achieve its maximum impact at a point just short of saturation. For example, many inkjet media hit a maximum saturation between 240 and 260 percent ink, so using GCR to keep the maximum ink level around 200 percent should avoid K speckling in highlights while still achieving its maximum impact in darker midtones and shadows.

Reiterating the G7 Curves
Printers drift; over the course of a production run color can shift due to ink, media and environmental changes. Every new ink refill and every new roll of media can lead to a color shift. When this occurs, recalibrating back to G7 Aims by reiterating the active G7 calibration curves will reestablish the gray balance on printer. But, how does this relate to GCR?

By calibrating to G7 Aims as an initial step, we establish a neutral gray balance, and we establish the relationship between our K grayscale density and our CMY-process grayscale density. If a media or ink change introduces a color shift (a cast), reiterating the G7 calibration curves will remove the cast and return

the system to its original neutral balance. That means that, while the ratio of C, M and Y substituted for K may change, the colorimetric equivalence of our process gray to K gray will remain intact. In other words, the reiterated curves will adjust the values of C, M and Y such that they produce the same neutral gray with the same density as when the system was originally characterized, and that will maintain the CMY to K equivalency of our GCR function. This is also what ensures the viability our ICC profile — if our GCR remains valid, so too should our profile.

Christopher Brown is the Engineering Manager in charge of Mutoh's Advanced Engineering Group at Mutoh America's US headquarters in Phoenix, Arizona, a position he has held for the past 17 years. Part of Mutoh's global research and development network, AEG focuses on development of color management and digital image processing technologies, coordinated through Mutoh's Research & Development Center located in Shimosuwa, Nagano, Japan. Mr. Brown led the team that developed the VM-10 spectrophotometer, and designed

Mutoh's G7 Calibrator software. His primary work is with micro spectrometers and spectral sensors, sensor arrays and embedded instrumentation. He is a software developer with a background in image processing, color management and digital halftoning; an IDEAlliance Certified G7 Expert; and, holds multiple patents related to spectrophotometer methods, spectral