QuadTech

Our passion is *your success.*™

Calculating a Target Image for Image-Based Color Control

By

Patrick Noffke

QuadTech, Inc.

TAGA TECHNICAL ASSOCIATION OF THE GRAPHIC ARTS

This paper was originally published by the Technical Association of the Graphic Arts (TAGA) in March of 2013.

White Paper 012013

Contents

Executive Summary

Image-based closed-loop color control systems offer unique advantages over color bar-based systems. With those advantages come several new challenges for both the printer and the color control system designer.

Calculation of the target image is impacted by all parts of a printer's workflow. Each step must be carefully analyzed, and any modifications to pre-press images must be input to the image-based system.

Selection of the target color profile is not always straightforward. For some printing conditions, there is no color standard that accurately represents the printed product, and compromises must be made.

Various types of target image calculation errors and their impact to print quality are discussed. Methods to account for TVI compensation and plate setter linearization curves are detailed. A non-linearity arising from the conversion from CMYK to CIELAB is presented. Signal-processing methods, which facilitate control of fine detail, such as text, are also discussed.

Image-Based Color Control

Closed-loop color control systems for newspaper and commercial web offset printing presses have been available for close to twenty years. Traditionally, these systems required printing special color measurement targets, such as a color bar.

More recently, image-based closed-loop color control systems have become available. These systems offer unique advantages over color bar-based systems, while posing new challenges for both the printer and the closed-loop system designers.

An obvious advantage of an image-based system is that special color measurement targets are not necessary. For commercial offset printing, color bars are usually placed in a trim area, and they are not visible in the final product. However, with newspaper and retail insert printing, there is no trim, and color measurement targets are visible and can be distracting. On newspapers, gray bars or gray dots are often used in favor of full color bars, in order to minimize the distraction from the primary content. Stanton & Seymour (2011) show that color bars, and in particular gray bars, do not accurately predict the rest of the work. Retail inserts often contain a small number of color swatches, which do not span the full extent of the page. Without measurement targets for each color in each ink key zone, it is not possible for a closed-loop color control system to control all ink keys on a press.

Seymour et al. (2010) group image-based systems into three primary categories: “color bar in the work,” ink-based, and CIELAB optimization, and demonstrate that CIELAB optimization systems strive to achieve the best color while overcoming limitations of the former two.

Calculation of the target image is an aspect of CIELAB optimization systems that affects both the printer and the control system designer. Intuitively, the concept is simple. The image-based system calculates a target CIELAB image that represents the desired color of the final printed product. The system then scans the printed image during production, performs a pixel-by-pixel comparison against the target image, and decides how to adjust the ink levels to make the two images match. Many challenges lie in the first part – calculating a target CIELAB image that represents the desired color of the printed product, with sufficient accuracy for the control system.

CIELAB Optimization

In order to fully appreciate the accuracy requirements of the target image, it helps to have a basic understanding of how CIELAB optimization systems function. The goal of the system is to adjust the ink keys within each ink key zone until the color (CIELAB value) of every pixel within an ink key zone equals that of the target. Clearly, then, the system needs to know the impact moving an ink key will have on the CIELAB values of the pixels.

First, the system needs to know the composition of inks at each pixel. If a given ink is not present at a pixel, then moving the key for the corresponding key zone will have no impact on the color of that pixel.

Second, the system needs to know, for the inks that are present, how much the color will change versus ink key opening. For a given change in ink key opening, a pixel with a tone value near 100% will have a larger color shift than a pixel with 10% tone value. Furthermore, a key zone with higher overall coverage will require more ink to achieve the target color than a key zone with lower coverage.

The ink composition at each pixel can be determined from pre-press images (page files). These images are used as input to a model of color shift versus key opening, which allows the system to decide which colors to adjust (and in what direction), and how much of a change to make for each color. This model must be matched with the response of the press when tuning the control system for optimal performance.

The QuadTech® Color Control and Web Inspection System with AccuCam™ is a CIELAB optimization system that converts the required change in CIELAB values within an ink key zone to a corresponding change in solid ink density for each of the inks present in that key zone. The change in solid ink density is then converted to a change in ink key opening by the control system component. This separation of density change and key move allows the control system to function identically for both color bar-based and image-based systems. It will also serve to provide familiar and consistent units when discussing the impact of the target image calculation errors.

The CIELAB optimization system strives to minimize the color error on all pixels within an ink key zone. This is accomplished through a least-squares solution, which gives all pixels equal weight when calculating the required change in density. For further details on the least-squares solution, refer to Appendix A.

The above discussion is primarily related to how quickly the system will converge to the best result achievable on-press, hereafter referred to as the optimized image. Regardless of the speed of convergence, as long as the system eventually does converge, the optimized image will always have the same color (within the process variation).

Target Accuracy

Ideally the optimized image is a perfect match with the target image. However, there are many variables that can make the ideal unattainable. For certain types of deviations from the ideal printing condition, we are interested to know how the optimized image will differ from the target image. Similarly, for certain types of errors in the target image calculation (and with a nominal printing condition), we would like to know the expected color of the optimized image.

Obviously, it is preferable to remove any errors in the target image calculation. Indeed, the primary motivation for writing this paper is to help readers identify the types of errors that can exist, with the expectation the printer and the system designer can work together to remove those errors. There are some cases, however, where it is impractical to remove all such errors. In those cases, workarounds may be used to ensure the error in the optimized image is predictable.

The types of target image calculation errors that can exist are listed below. Each error type will be discussed in the sections that follow.

1. GCR.
2. TAC limiting.
3. Curve errors.
4. Profile mismatch.
5. Image processing artifacts.

GCR

If the printer is applying GCR (gray-component replacement), sometimes referred to as ink optimization, the images sent to the image-based system must have the same GCR applied. Clearly, if the system is told (through the images it receives) that a gray region is comprised of cyan, magenta, and yellow inks, but on press only the black ink is used for that region, the system will incorrectly try to adjust cyan, magenta, or yellow when it should instead adjust black.

This is a fairly simple error to resolve, but if not handled correctly, the result will be quite bad. GCR is typically performed as part of the RIP, which usually produces high-resolution 1-bit TIFFs that are sent to the plate setter. These 1-bit TIFFs will have the GCR applied, and may be suitable as input to the image-based system.

GCR may be applied further upstream through an ICC profile. If an image is created as RGB, it must be converted to CMYK prior to printing. If PDF/X-1a or PDF/X-3 files are created as part of the printer's workflow, the conversion to CMYK is built-in to these files. The PDF/X-1a format requires all images be converted to CMYK, which is done through the use of an ICC profile. PDF/X-3 files may contain RGB images, but these images must have an associated ICC profile that defines the conversion to CMYK. In either case, some level of GCR may be built-in to the associated ICC profile.

If GCR is applied further upstream and possibly again during the RIP, use of the 1-bit TIFFs output from the RIP is still suitable for the image-based system, since those TIFFs will have all GCR steps applied. If it is desirable to use images further upstream as input to the image-based system, it is important to identify any steps downstream that may apply GCR and ensure the color separations used to calculate the target CIELAB image have the same level of GCR applied as the plates.

TAC Limiting

The main benefit of TAC (total area coverage) limiting is to stay within the limits of the press. The printer may also wish to apply TAC limiting in order to achieve ink savings and to reduce scumming artifacts.

Color standards such as ISO 12647, SWOP, and SNAP recommend maximum TAC levels suitable for various printing conditions. These levels typically represent the darkest color a substrate can accept. Adding ink beyond the TAC limit will tend to reduce the trapping instead of the color becoming more saturated.

For an image-based system, the printer *should* apply TAC limiting. If a region has a higher TAC than the maximum for that printing condition, the system may try to push the press beyond its limits. For example, the darkest color achievable for a given printing condition may be $L^* = 30$, $a^* = 0$, $b^* = 0$. If a region has 400 TAC, the target color for that region may be $L^* = 26$, $a^* = 0$, $b^* = 0$. Regardless of how much ink is added, the L^* value as measured by the system never gets below 30. But the image-based system will see a persistent error, so it will continue to open the keys for all four colors. This ink doesn't get taken up by the substrate, but it has to go somewhere. It may cause the surrounding colors to become more saturated, or it may start to enter the non-inked areas and create scumming.

As with GCR, the images sent to the image-based system must have the same TAC as the plates. Once again, the 1-bit TIFFs output from the RIP will be suitable for the image-based system. Use of images further upstream may be possible, as long as the TAC levels seen by the system agree with the levels used on press.

Curve Errors

Curve errors encompass both TVI (tone value increase) errors and plate errors.

TVI errors exist when the total process TVI (TVI compensation curves plus press TVI) does not match the expected TVI. This may occur if the press is running outside its norm, or the press has never been calibrated so the press TVI is unknown. Additionally, a TVI mismatch may exist but workflow constraints prevent implementing the correct TVI compensation curve. For example, customers may supply images separated to CMYK using a profile unknown by the printer. Or the customer may have assumed a particular TVI (through the choice of ICC profile used to make the separations) that does not match the press TVI, yet the printer may apply TVI compensation curves to all images, resulting in a process TVI different from what the customer expected.

Plate errors have an effect similar to TVI errors, but are caused by conditions such as changes to plate setter gain (i.e. a plate setter is out of calibration), lack of plate setter linearization curves, or linearization curves are applied to the pre-press images and not accounted for by the image-based system.

Curve errors can be significant. Experienced printers know that a small (2-3%) TVI error at 20% tone value is very hard to correct with a solid density move. Nonetheless, the image-based control system is going to do its best to make all pixels match the target image. With CIELAB optimization, all pixels within an ink key zone have equal weight **in CIELAB color space**. Equal weight in CIELAB color space does not translate to equal change in solid ink density. If most pixels in an ink key zone have low tone values, they have the highest contribution toward any change in ink levels. This may result in larger errors on pixels with higher tone values, since there are fewer of them.

It is possible to modify this effect through non-uniform weightings. However, determining weightings automatically is non-trivial. Use of TAC levels is inappropriate because there are always going to be cases where lower TAC levels should have higher weight than higher TAC levels, and vice versa. It is best to implement weighting through a “manual override,” where the printer, press operator, or print buyer decide which regions should have the highest weight. Even if a weighting facility is available, there will be cases where curve errors impact the optimized image. For example, the weighting facility may not be used, or a full range of tone values may be weighted equally. For simplicity, an equal weighting is assumed for the remainder of this discussion.

The following examples will demonstrate the importance of minimizing curve errors, and how, in the presence of curve errors, the optimized image depends on the ink coverage.

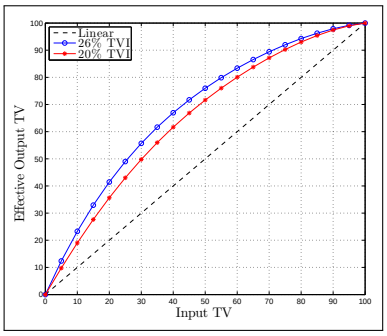

Figure 1: TVI curves for examples. The expected TVI is 26%, and the simulated process TVI is 20%.

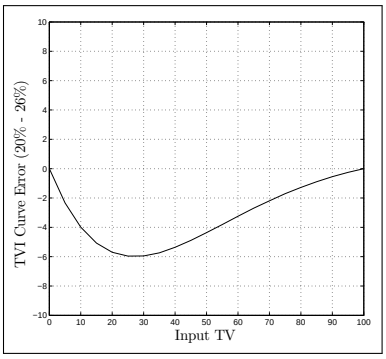

Figure 2: TVI curve error for examples. The maximum TVI difference of 6% occurs around 30% input tone value.

For these examples, a TVI error is simulated, where the expected TVI differs from the process TVI by up to 6%. The target image CIELAB values were calculated using the ICC profile ISOnewspaper26v4.icc, which is available from WAN-IFRA. The target TVI is 26% at 40% input tone value. The target CIELAB images used in the examples therefore correspond with the printing condition specified in ISO 12647-3.

A 20% process TVI was simulated by applying a -6% tone reduction to all four colors within the ICC profile, and the optimized image CIELAB values were calculated using the modified profile. For each of the examples, the optimized image defines the color that the system would converge to if the process TVI matched the 20% TVI in the modified profile, and the target matched the 26% TVI embedded in the original profile.

Figure 1 shows the expected TVI of 26% and the simulated process TVI of 20%. The difference between these two curves is shown in Figure 2. Note that the maximum TVI difference of 6% occurs around 30% input tone value.

Because of the reduced process TVI, if the solids were at the target CIELAB value, the mid-tones would then be lighter than the target. The image-based system will make a compromise that pushes the solids darker than the target in order to better match the mid-tones to the target. This type of compromise is very common with TVI (and other curve) errors. As the examples will illustrate, the extent the solid is pushed past the target depends upon the distribution of colors inline with that solid.

The examples are summarized as follows:

1. Black tone strip – A strip of black-only rectangles, with tone values 0, 10, 20, 30, 40, 50, 60, 70, 80, 90, and 100.
2. Solids only – C100, M100, Y100, and K100 solids with no other tone areas inline. A CMY100 gray region is added to the right of the four solid regions.
3. Solids and 50% – C100, M100, Y100, and K100 solids inline with 50% tone areas. The CMY100 gray region is inline with a (C = 40, M = 30, Y = 30) region.

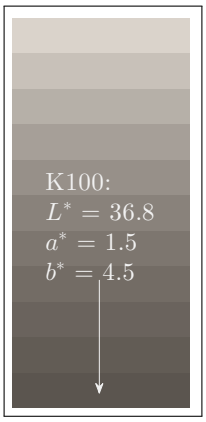

Figure 3: Example 1 - black tone strip, target image.

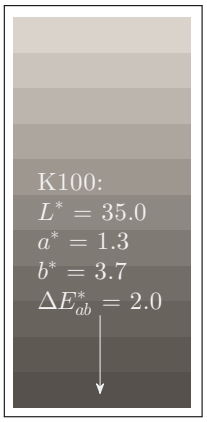

Figure 4: Example 1 - black tone strip, optimized image.

4. Solids and 20% – C100, M100, Y100, and K100 solids inline with 20% tone areas. The CMY100 gray region is inline with a (C = 25, M = 18, Y = 18) region.

5. Small solids and 20% – Smaller C100, M100, Y100, and K100 solids inline with 20% tone areas. The CMY100 gray region is inline with a (C = 25, M = 18, Y = 18) region.

Example 1 - Black Tone Strip

For this example, a series of black tone patches are inline with each other, with the tone values starting at 0 (paper) and increasing in 10% increments up to a solid. Figures 3 and 4 show the target and optimized images, respectively.

Because of the TVI error, if the solid was at the target CIELAB value, the mid-tones would be too light. The image-based system adds ink, pushing the solid darker than the target, until the overall error is balanced for all tone values present in the ink key zone.

Figure 5 best demonstrates the overall compromise that is made with CIELAB optimization. The non-optimized curve in the figure represents the color error versus tone value when the solid is at the target CIELAB value. This curve essentially demonstrates the net effect of the TVI error. The optimized curve shows the error after CIELAB optimization and color convergence. The errors in the mid-tones are reduced from the non-optimized curve, but at the expense of the shadows and the solid.

The ΔE^*_{ab} color difference between the target and optimized solid regions is 2.0, which corresponds roughly to a visual density difference of 0.04D.

Black ink primarily affects the L^* CIELAB component. Table 1 lists the target and optimized L^* values and their differences versus tone value. The greatest color difference between the target and optimized images actually occurs in the 40% tone region, where $\Delta L^* = -2.5$ (the minus sign indicates the optimized value is still lighter than the target).

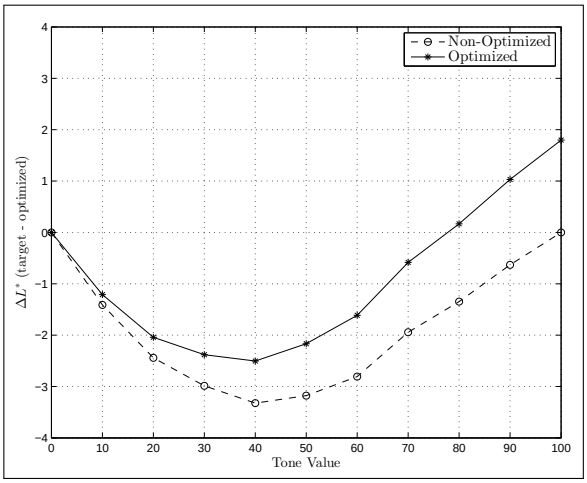

Figure 5: Example 1 - black tone strip. The non-optimized curve shows the net effect of the TVI error, if the solids were at the target CIELAB value. The optimized curve shows the color error after CIELAB optimization and color convergence. The errors in the mid-tones are reduced, but at the expense of the shadows and the solid.

Tone Value	Target L^*	Optimized L^*	ΔL^*
0	85.2	85.2	0.0
10	78.6	79.8	-1.2
20	72.2	74.2	-2.0
30	66.2	68.6	-2.4
40	60.4	62.9	-2.5
50	55.1	57.2	-2.2
60	50.1	51.7	-1.6
70	46.2	46.8	-0.6
80	42.6	42.4	0.2
90	39.6	38.6	1.0
100	36.8	35.0	1.8

Table 1: Example 1 - black tone strip, ΔL^* between optimized and target images for each tone value in the strip.

Example 2 - Solids Only

In this example, only solids are present in the image. Because there is nothing inline with the solids, the optimized image is equal to the target image, and the color differences between the solid regions are all zero. Figures 6 and 7 show the target and optimized images, respectively.

There is not much more to say about this example, except that it will be the basis for the following examples.

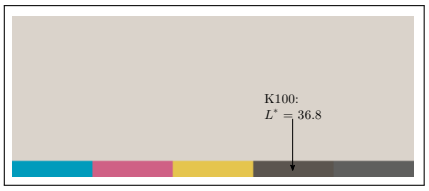

Figure 6: Example 2 - solids only. C100, M100, Y100, K100, and CMY100 regions, target image. The area above the solid regions is paper.

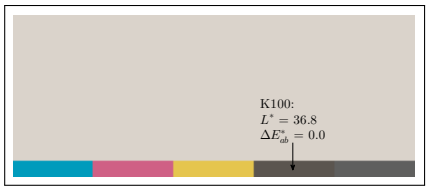

Figure 7: Example 2 - solids only. C100, M100, Y100, K100, and CMY100 regions, optimized image. Since there are no half-tones inline with the solids, the optimized image equals the target image.

Example 3 - Solids and 50%

The target and optimized images for this example are shown in Figures 8 and 9, respectively. Each solid is inline with a large 50% region. As with the black tone strip example, the system makes a compromise to balance the errors between the solids and mid-tones.

Table 2 lists the ΔE^*_{ab} of the 50% regions before and after CIELAB optimization, and the solid ΔE^*_{ab} after optimization (the solid error before optimization is zero). The errors on the 50% patches are reduced, at the expense of the solids. Because the size of the half-tone region is much larger than that of the solid, more weight is given to the half-tone region, causing the solids to be pushed well beyond the target.

Table 3 shows the solid ink density shift necessary to produce the ΔE^*_{ab} values listed in Table 2. The density shifts range from 0.05D for yellow to 0.08D for cyan and black. This is significant for newsprint.

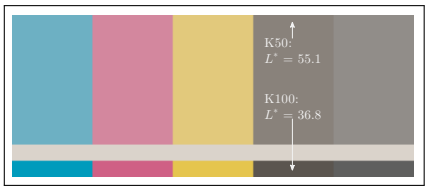

Figure 8: Example 3 - solids and 50%, target image. The solids are inline with 50% halftones.

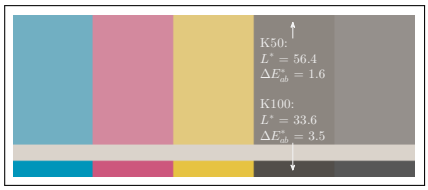

Figure 9: Example 3 - solids and 50%, optimized image. The system makes a compromise to balance the errors between the solids and the mid-tones.

	Cyan	Magenta	Yellow	Black	Gray
Non-Optimized 50%	3.0	3.6	3.7	3.2	3.2
Optimized 50%	1.3	1.8	1.1	1.6	1.3
Optimized Solid	3.8	4.0	4.7	3.5	3.1

Table 2: Example 3 - solids and 50%, ΔE^*_{ab} before and after CIELAB optimization. The mid-tone errors are reduced, at the expense of the solids.

	Cyan	Magenta	Yellow	Black
ΔD	0.08	0.07	0.05	0.08

Table 3: Example 3 - solids and 50%. Density shift (from nominal) after CIELAB optimization.

Example 4 - Solids and 20%

The target and optimized images for this example are shown in Figures 10 and 11, respectively. Now, instead of a 50% region, each solid is inline with a large 20% region.

Table 4 lists the ΔE^*_{ab} of the 20% regions before and after CIELAB optimization, and the solid ΔE^*_{ab} after optimization. The errors on the 20% patches are again reduced at the expense of the solids.

Notice how the solid errors are similar in magnitude to those from example 3. The error in the optimized K100 region is even smaller with a 20% (1.9 ΔE^*_{ab}) than with a 50% (3.5 ΔE^*_{ab}) inline with the solid. This is due in part to the fact that the non-optimized errors (i.e. the starting point for CIELAB optimization) are smaller for the 20% case. An in-depth analysis of the sensitivities used for CIELAB optimization is required to fully explain this effect, and is beyond the scope of this paper.

Table 5 shows the solid ink density shift necessary to produce the ΔE^*_{ab} values listed in Table 4. The black density shift decreased compared to example 3, corresponding to the reduced ΔE^*_{ab} just discussed. The yellow density shift increased from 0.05D to 0.07D, and its ΔE^*_{ab} is slightly higher than in example 3.

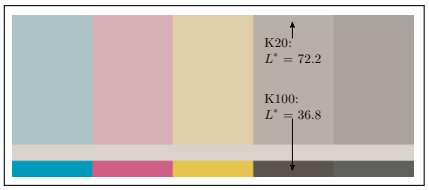

Figure 10: Example 4 - solids and 20%, target image. The solids are inline with 20% halftones.

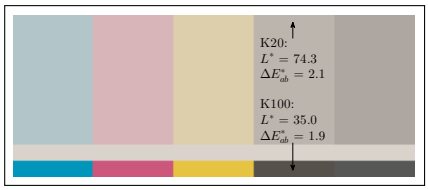

Figure 11: Example 4 - solids and 20%, optimized image. The system makes a compromise to balance the errors between the solids and the mid-tones.

	Cyan	Magenta	Yellow	Black	Gray
Non-Optimized 20%	2.6	3.2	3.0	2.4	3.0
Optimized 20%	1.9	2.4	1.8	2.0	1.7
Optimized Solid	4.0	4.1	5.5	1.9	3.3

Table 4: Example 4 - solids and 20%, ΔE^*_{ab} before and after CIELAB optimization. The mid-tone errors are reduced, at the expense of the solids.

	Cyan	Magenta	Yellow	Black
ΔD	0.08	0.07	0.07	0.04

Table 5: Example 4 - solids and 20%. Density shift (from nominal) after CIELAB optimization.

Example 5 - Small Solids and 20%

The target and optimized images for this example are shown in Figures 12 and 13, respectively. The solid region is much smaller than in the previous example, and inline with a large 20% region.

Table 6 lists the ΔE^*_{ab} of the 20% regions before and after CIELAB optimization and the solid ΔE^*_{ab} after optimization, and Table 7 shows the solid ink density shift necessary to produce the solid ΔE^*_{ab} values. Now, because of the small percentage of solid pixels inline with the 20%, very little weight is given to the solid, and the solid errors in the optimized image are correspondingly very large. The ΔE^*_{ab} range from 5.3 for black to 9.7 for yellow, and the density shifts are all at least 0.11D. This is very high for newsprint.

In the previous examples, both the ΔE^*_{ab} of the optimized image and required density shifts were already quite large, and should be sufficient to demonstrate the importance of minimizing curve errors. This example has errors that might be considered extreme, while the conditions for the example (6% TVI error and small percentage of solids present) are not all that uncommon.

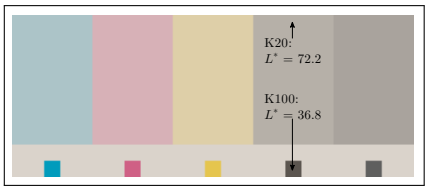

Figure 12: Example 5 - small solids and 20%, target image. The solids are smaller and inline with 20% half-tones.

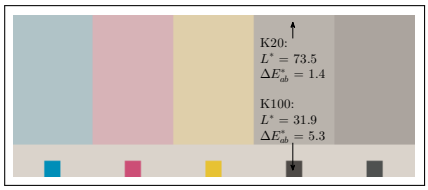

Figure 13: Example 5 - small solids and 20%, optimized image. The system makes a compromise to balance the errors between the solids and the mid-tones.

	Cyan	Magenta	Yellow	Black	Gray
Non-Optimized 20%	2.7	3.2	3.0	2.4	3.0
Optimized 20%	1.3	1.9	0.9	1.4	0.6
Optimized Solid	7.8	7.6	9.7	5.3	6.1

Table 6: Example 5 - small solids and 20%, ΔE^*_{ab} before and after CIELAB optimization. The mid-tone errors are reduced, but now the solids are pushed very far past the target.

	Cyan	Magenta	Yellow	Black
ΔD	0.15	0.13	0.11	0.12

Table 7: Example 5 - small solids and 20%. Density shift (from nominal) after CIELAB optimization.

Curve Error Compensation

Fortunately, in many cases, curve errors can be reduced or completely avoided. TVI errors are mitigated by measuring the TVI of the press and applying TVI compensation curves such that the total process TVI equals an expected value, such as set forth in ISO 12647. Plate errors are mitigated by regularly calibrating the plate setter. If a plate setter is not linear, a calibration curve may be applied during the RIP to linearize the output.

Of course, TVI is difficult to measure in practice, and some process variation will occur. This is true regardless of whether a printer is using an image-based system, a color bar-based system, or manual control. The difference among the control methodologies will lie in what kind of compromise is made to best accommodate any TVI error. The CIELAB optimization system will, *on average*, produce a result that provides the best visual match to the target. But on occasion – and depending on the tone value distribution – some colors may be driven far from the target, in order to achieve a good color match elsewhere in the image.

TVI Compensation Curves

ISO 12647 specifies the required process TVI, but it does not provide methods to correct for a TVI error. Other efforts that are profile-centric, such as ISO 15339 (albeit still in draft status), do not directly specify a TVI, but the expectation is that the printer must match a set of characterization data. However, the characterization data has a particular TVI, and if the printer's TVI doesn't match that of the data, TVI

compensation may be required. Again, the method to implement TVI compensation is not discussed.

To know whether TVI compensation curves are required, a press test must first be performed, and TVI must be measured from patches with a variety of tone values. Ideally, TVI is measured on full tone ramps, from paper to solid, in at most 10% tone increments. If the TVI does not match the values specified in the standard (or embedded in the characterization data), a TVI compensation curve is necessary.

It may be tempting to simply subtract the measured TVI from the expected TVI and add the TVI difference to each of the tone values within the tone ramp. In other words, for a particular color, if a 20% patch measures 15% TVI, and the specification is 19% TVI, the simple approach would be to add 4% to all 20% tone values for that color. That would get you close, but it is not quite correct.

Consider the graph in Figure 14 which is a plot of TVI vs. tone value as required by ISO 12647-3. According to this plot, the *expected* TVI of a 20% tone value is about 19%. If we measure 15% TVI, and decide to add 4% to all 20% tones, we will then have a 24% tone value on the plate. However, from the graph, we see a 24% tone value has about 22% TVI. So we overshot our target TVI of 19% by about 3%.

The correct way to implement a TVI compensation curve is through interpolation. The goal is to make the total process TVI (compensation curves "plus" press TVI) equal to the curve in the standard. "Plus" is in quotes because, as we just saw, addition doesn't work quite right. We achieve the target TVI (that of the standard) by "looking up" (or interpolating) the compensation curves using the measured TVI of the press. This matches what is actually happening in a printer's workflow – the plates contain curved images, and the press TVI is "applied" to these images.

To describe how to apply interpolation for TVI compensation, we'll first define a "tone curve" as the apparent or output tone value vs. input tone value. In other words, the output of a tone curve is the input tone value plus any TVI. An example tone curve is illustrated in Figure 15, which is a plot of apparent tone value vs. input tone value for ISO 12647-3.

Let TV represent a tone value ranging from 0 to 100. Let $C_1(TV)$ represent the TVI compensation tone curve, $C_2(TV)$ represent the press tone curve, and $C_S(TV)$ represent the tone curve of the standard. The images are first "run through" $C_1(TV)$, the compensation curve, followed by "running through" $C_2(TV)$, which encompasses the press TVI. Therefore, for an input tone value, TV, with a TVI compensation curve, the net (or process) tone curve on press is equal to

$$\text{Process Tone Curve} = C_2(C_1(TV)) \tag{1}$$

We want this equal to the tone curve of the standard, or

$$C_2(C_1(TV)) = C_S(TV) \tag{2}$$

We need to solve for C_1 in order to find the compensation curve that makes the total process TVI equal to that of the standard. We do this by applying the inverse of C_2, the press tone curve, to the tone curve of the standard, or

$$C_1(TV) = C_2^{-1}(C_S(TV)) \tag{3}$$

The inverse of a function $y = f(x)$ is obtained by swapping the x and y variables and solving for y. For example, if $y = x^2$, then we would write $x = y^2$, and solving for

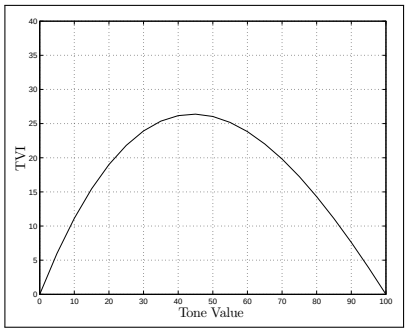

Figure 14: 26% TVI curve, difference from linear.

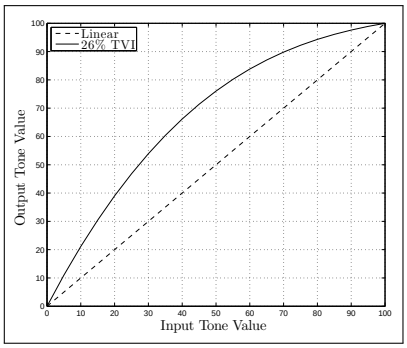

Figure 15: 26% TVI curve, absolute.

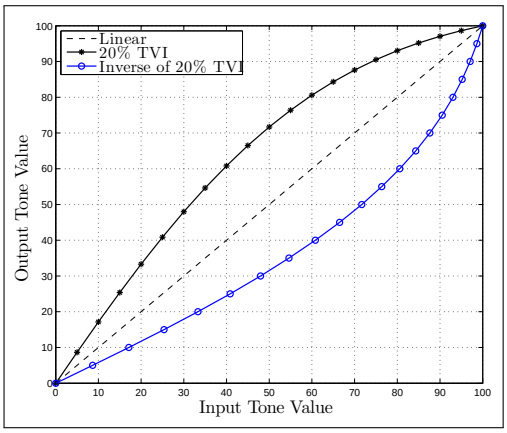

Figure 16: 20% TVI and its inverse.

y, we get the inverse function of $y = \sqrt{x}$. Note that the inverse of a function is the original function flipped about the line $y = x$.

Let's return to our example. A 20% patch was measured to have 15% TVI, and we want to find the compensation required to make the result equal to 19% TVI. Start by looking up $C_S(20)$, which we know from Figure 15 is 39. Then look up $C_2^{-1}(39)$ using the blue curve on Figure 16, giving $C_2^{-1}(39) = 23$. A 20% tone value therefore needs a 3% gain in pre-press.

Applying this back to the press TVI (the black curve in Figure 16), we see that a 23% tone value comes off the press looking like a 39% tone value. Since we started with a 20% tone (before compensation curves are applied), the net process TVI is 19%, which is what the standard requires.

Accounting for Curves

If an image-based system uses high-resolution 1-bit TIFFs identical to those used to make plates, the system must account for any curves that are applied. Suppose, for example, a printer is attempting to match ISO 12647-3, and the press TVI is 20%, as in the earlier discussions. In order to meet the requirements of the standard, a TVI

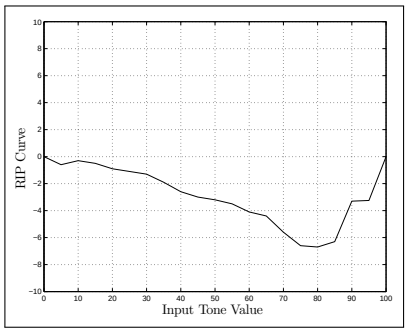

Figure 17: Plate setter linearization curve, difference from linear. This curve is applied to 1-bit TIFFs to produce linear output on a plate setter.

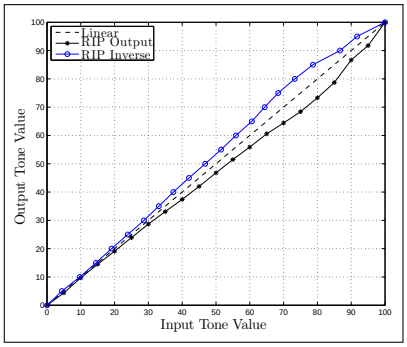

Figure 18: Plate setter linearization curve and its inverse. Since the 1-bit TIFFs have the RIP curve applied, the inverse curve must be applied prior to converting to CIELAB.

compensation curve must be applied such that tone values are increased throughout the tone scale.

It may be desirable for the image-based system to use the ISOnewspaper26v4.icc profile for calculating the target CIELAB values. However, if the images input to the image-based system have the TVI compensation applied, the target will have an extra TVI (the gain of the compensation curve) that was not intended. The image-based system must then "reverse out" the TVI compensation curve. This is done very similarly to implementing a TVI compensation curve. The system must apply the inverse of the TVI compensation curves to the images it receives, and the net result of the compensation curve "plus" the inverse compensation curve is the image prior to TVI compensation. The system may then calculate the CIELAB value using the ISOnewspaper26v4.icc profile, and the net TVI of the target will equal the 26% TVI of the standard.

Similarly, a plate setter may have a dot gain, which is a type of TVI, causing the printer to implement a plate setter linearization curve. Figure 17 shows one such linearization curve from a particular printer. The linearization curve was implemented as part of the RIP, which meant the 1-bit TIFFs used by the image-based system had these curves applied. The inverse of the linearization curve must therefore be applied to the 1-bit TIFFs prior to conversion to CIELAB. The RIP tone curve and its inverse are shown in Figure 18.

Note that 1-bit TIFFs cannot be used directly to interpolate into a tone curve. The system must first convert the images to continuous tone (e.g. 8-bit), with filtering applied to de-screen the image (i.e. convert the binary half-tone dots back to the continuous tone they originated from). The need for filtering and other image processing aspects are discussed in a later section.

In some cases, a printer's customer may perform their own RIP, and it may be impractical to apply the necessary linearization curves. This results in a process TVI that does not match the target TVI. If no curve is applied and the plate setter has a positive dot gain, the mid-tones would tend to be too dark on press. The image-based system will push the solids light to a variable degree that depends on coverage. If this variability is undesirable, it is possible to "let the error ride," such that solids in the optimized image will always match the target. This is accomplished by applying

the plate setter gain to the 1-bit TIFFs. The system then *expects* the mid-tones to be darker than the target, and variable compromises are avoided. The optimized image will not match the target, due to the lack of a linearization curve, but the color errors are well-controlled and predictable. Clearly the best solution is for the customer to implement the linearization curves, but if that is not possible, this workaround may be an acceptable trade-off.

Identifying and accounting for any curves is a crucial part of implementing and commissioning an image-based system. As demonstrated earlier, even a small TVI error of 6% can result in significant errors in the optimized images. With the printer and image-based system designer working together, any such curve errors can be minimized or completely avoided.

Profile Mismatch

The term "profile" is used loosely here to encompass any means of converting from CMYK separations to CIELAB, including press characterization data, ICC profiles, and a color model.

Press characterization data is usually obtained from a profiling test, which provides a set of CIELAB data at known points in CMYK space. The points in CMYK space are discrete, and are usually spaced with at least 10% tone increments. For overprints containing black, the tone increments tend to get even higher. The number of colors that fit on a profiling target is limited by the page size and the aperture size of the instrument used to measure the colors. Profiling targets may be spread across multiple pages, but this adds complexity to the profiling process.

Once press characterization data is obtained, it may be used to generate an ICC profile. ICC profiles for printing devices contain much useful information, such as the levels of GCR or TAC limiting for a particular printing condition, the substrate color, and "perceptual" and "saturation" rendering intents. The different rendering intents are useful for converting images, logos, etc., when the gamut of the image color space does not match the printable gamut, and some gamut compression or other mapping is desirable. However, for the purposes of calculating a target image, ICC profiles typically contain no more information than the press characterization data. For this purpose, the profiles simply contain a CMYK to CIELAB look-up table, and this is almost always interpolated onto a regular spacing in CMYK space directly from the characterization data.

Another means to convert from CMYK to CIELAB is through a color model, which factors in various parameters related to a printing condition and calculates the reflectance according to an underlying physical model of how light interacts with the ink and substrate. CIELAB is then calculated from the reflectance. The accuracy of the CIELAB values is only as good as the model, and presumably the parameters of the model were obtained by measuring physical properties of a printed target, such as color patches in a profiling chart.

Obtaining an accurate profile is extremely difficult. Printing of profiling targets is challenging because the solid ink density needs to be consistent across the extent of the chart. Uneven spatial distribution of ink coverage can result in more ink consumption in one key than its neighbor, and the solid ink density may vary through the key boundaries. White space or fiducials may be required at the edges of charts for measurement by a particular instrument. This creates a discontinuity in coverage and may result in uneven solid ink density. If no color control system is used during a profiling test, hitting the solid aim points across the entire chart can take a very long time, resulting in much waste.

Several "standard" profiles have been created, and using one of these is a good starting point for image-based color control. However, there are many more printing conditions than profiles, and finding an existing profile that is "good enough" for calculating a target CIELAB image is not always possible.

Figure 19: Profile mismatch, two-color target image, red and gray.

Figure 20: Profile mismatch, one-color target image, red only.

	Red (2-color)	Gray (2-color)	Red (1-color)
Non-Optimized	5.5	3.0	5.5
Optimized	5.1	0.6	3.9

Table 8: Profile mismatch, ΔE^*_{ab} before and after CIELAB optimization.

	Cyan	Magenta	Yellow	Black
Two color	0.04	0.00	0.07	0.00
One color	0.00	-0.03	0.06	0.00

Table 9: Profile error, density shift after CIELAB optimization.

By way of example, press characterization data was obtained for one particular printing condition, which was heatset on uncoated paper similar to newsprint. It was determined that the closest pre-existing profile was the characterization data in ISO 15339, reference printing condition #2 (or RPC-2). Much effort was expended to match RPC-2 during a profiling test. After the test was complete, some parts of the printable gamut still differed from RPC-2 by as much as 9 ΔE^*_{ab}. It was further determined that changing ink levels would not suffice to match the gamuts – the color of the particular paper and ink combination was just that different from RPC-2 (still, RPC-2 was the closest pre-existing profile).

This type of profile mismatch will impact the optimized image. Figures 19 and 20 show two different target images, one with two colors, red (M100, Y100) and gray (C100, M100, Y100), and the other with only red. For this example, the non-optimized colors were the values obtained during the profiling test. The target CIELAB image was calculated using a profile built from RPC-2. Table 8 shows the red started at 5.5 ΔE^*_{ab} from the target, and the gray non-optimized error was 3.0 ΔE^*_{ab}. For the two-color case, the red error was reduced to 5.1 ΔE^*_{ab}, and the gray error to 0.6 ΔE^*_{ab}. But for the one-color case, the red error could be further reduced to 3.9 ΔE^*_{ab}.

Table 9 shows the solid ink density shift required to change from the non-optimized to the optimized CIELAB values. For the two-color case, the presence of gray prevented the system from moving magenta. In the one-color case, a -0.03D magenta shift moved the red closer to the RPC-2 target color. The optimized image once again depends upon the distribution of colors within an ink key zone.

Image Processing Artifacts

Signal Processing Considerations

The signal processing aspects of an imaging device are quite complex, particularly for a line-scan imaging device. To thoroughly treat the signal processing aspects requires a level of mathematics that, while perhaps interesting, is not necessary to support the main argument of this section:

In order to accurately measure and control fine detail, such as text or thin lines, an image-based system must remove any sources of aliasing. Furthermore, the signal characteristics of the target image must match that of the live image. In particular, the frequency response of the two images must be closely matched.

Aliasing occurs when two different signals can be sampled to have the same result. The two signals are said to be *aliases* of one another. Another effect of aliasing is that the same signal can be sampled to have completely different results, depending on the phase (time shift, or in the case of images, spatial offset) of the sampling. The latter case is illustrated in Figure 21, where a 5 Hz signal is sampled at 10 Hz for two different samplings. The two samplings are offset from one another by 50 ms. The first sampling correctly yields a 5 Hz signal, but the second sampling incorrectly yields a DC (constant) signal. In this case, the DC sampling would be aliased with a true DC signal.

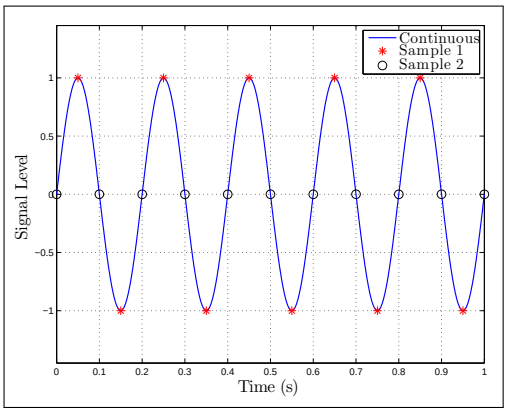

Figure 21: Aliasing can cause the same signal to be sampled with completely different results. The first sample (red) correctly yields a 5 Hz signal, but the second sample (black) incorrectly yields a DC (constant) signal.

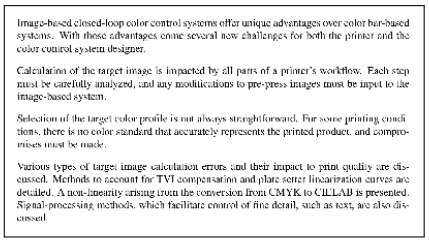
Image-based closed-loop color control systems offer unique advantages over color bar-based systems. With those advantages come several new challenges for both the printer and the color control system designer.

Calculation of the target image is impacted by all parts of a printer's workflow. Each step must be carefully analyzed, and any modifications to pre-press images must be input to the image-based system.

Selection of the target color profile is not always straightforward. For some printing conditions, there is no color standard that accurately represents the printed product, and compromises must be made.

Various types of target image calculation errors and their impact to print quality are discussed. Methods to account for TVI compensation and plate setter linearization curves are detailed. A non-linearity arising from the conversion from CMYK to CIELAB is presented. Signal-processing methods, which facilitate control of fine detail, such as text, are also discussed.

Figure 22: Downsampled image of text, no filtering, unshifted.

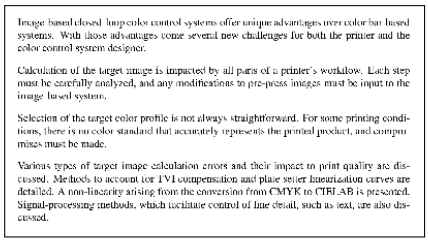
Image based closed-loop color control systems offer unique advantages over color bar based systems. With those advantages come several new challenges for both the printer and the color control system designer.

Calculation of the target image is impacted by all parts of a printer's workflow. Each step must be carefully analyzed, and any modifications to pre-press images must be input to the image based system.

Selection of the target color profile is not always straightforward. For some printing conditions, there is no color standard that accurately represents the printed product, and compromises must be made.

Various types of target image calculation errors and their impact to print quality are discussed. Methods to account for TVI compensation and plate setter linearization curves are detailed. A non-linearity arising from the conversion from CMYK to CIELAB is presented. Signal-processing methods, which facilitate control of fine detail, such as text, are also discussed.

Figure 23: Downsampled image of text, no filtering, shifted one-half pixel before downsampling. The text is visibly different from the text in the unshifted image.

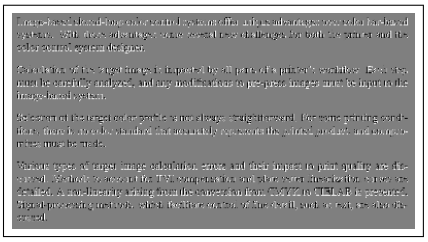

Figure 24: Difference between shifted and unshifted images, no filtering. Any pixels not gray are errors.

Figures 22 through 24 demonstrate aliasing of a rasterized image of text. The text was rasterized at 300 dpi and downsampled to 75 dpi, for a downsampling factor of 4. No filtering was applied to the image prior to downsampling. Figure 22 shows the 75 dpi image. Figure 23 shows the same text image at 75 dpi, but prior to downsampling, the 300 dpi image was shifted in both directions by 2 pixels, or one-half of a 75 dpi pixel. Figure 24 shows the difference between the shifted and unshifted images. Clearly, there are very large errors throughout the image (anything not gray is an error).

Figures 25 through 27 show the same text images at 75 dpi, but with anti-aliasing filters applied prior to downsampling. Now, the difference image is much closer to gray, indicating the errors are much smaller. Any residual error is primarily due to the fact that the images were not registered prior to taking the difference.

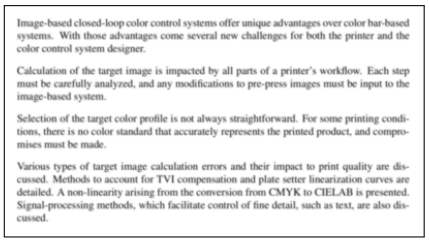

Figure 25: Downsampled image of text, with filtering, unshifted.

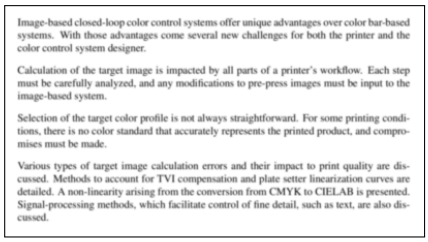

Figure 26: Downsampled image of text, with filtering, shifted one-half pixel before downsampling.

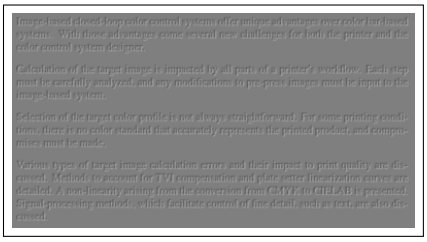

Figure 27: Difference between shifted and unshifted images, with filtering. The image is much closer to gray (no error). Any residual error is due to the fact the images were not registered.

In the audio world, it is very common to apply anti-aliasing filters prior to sampling. This is typically done with an analog low-pass filter circuit. The output of the filter is input to an analog-to-digital converter (ADC), where sampling occurs.

Apparently, then, an image-based system simply has to apply anti-aliasing filters in the signal processing chain. Indeed, the Nyquist-Shannon sampling theorem states that, for a sampling frequency of f_s, a signal is completely determined by sampling if it contains no frequencies greater than or equal to $f_s/2$. Then if an image-based system samples at, say, 75 dpi, the image must not contain frequency content greater than or equal to 37.5 lpi (lines per inch). But half-tone dots contain frequencies much higher than 37.5 lpi, as will any high-contrast gradients such as edges of text or line art, or images with any amount of detail.

In the imaging world, **the only way to remove aliasing is to optically blur the image**. If frequencies greater than or equal to the spatial sampling frequency are not removed, the simple act of sampling an image on a CCD or CMOS sensor will produce aliasing. Blurring the image is the only way to remove high frequencies from an optical signal.

Subsequent steps in the image processing chain, such as downsampling to a lower resolution image, must also incorporate anti-aliasing filters. Since this processing is completely digital (i.e. it occurs after sampling and digitization), digital anti-aliasing filters are appropriate.

The type of shift discussed in the text image example is a regular occurrence when imaging devices are scanning a moving web. The web can shift by a fraction of a pixel between image acquisitions. If this phase shift does not align exactly with the

phase of the target image, and anti-aliasing filters (particularly optical blur) are not used, large errors will exist between the sampled and target images.

The remaining challenge of an image-based system designer is to match the calculated target CIELAB image to the final scanned images used for color measurement. The target image calculation may start from high-resolution 1-bit TIFFs, such as those used to make the plates, and must incorporate any optical blurring or further downsampling done to the scanned image. Once the frequency characteristics of the two images are well-matched, accurate color measurement and control of fine detail is possible with an image-based system.

Image Conversion Non-Linearities

When calculating target CIELAB images, the image-based system must start from pre-press images (e.g. CMYK separations), and filter and resample them to match the resolution and frequency characteristics of the scanner. At some point in that process, the images are converted from CMYK to CIELAB. If the conversion to CIELAB is done at too-low of a resolution, a non-linearity can be introduced that creates significant errors around high-contrast regions.

Figures 28 through 30 demonstrate this non-linearity with the text image used in the previous section. In Figure 28, the text image is downsampled from 300 dpi to 75 dpi in CMYK color space, then converted to CIELAB at 75 dpi (the rendered image is the CIELAB image converted to sRGB). In Figure 29, the image is first converted to CIELAB at 300 dpi, then the image is downsampled to 75 dpi. When converted to CIELAB at the lower resolution of Figure 28, the text pixels appear darker than if converted to CIELAB at a higher resolution, as in Figure 29.

If the darker image were used as the target image, the image-based system would add more ink. In many cases, the scanned image will never be dark enough to match the target, and the image-based system may continue to add ink until scumming or other bad things start to happen.

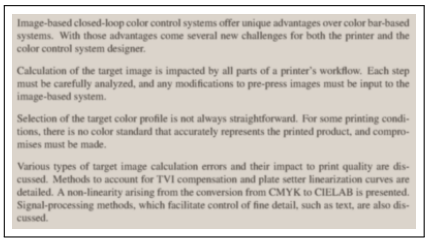

Figure 28: 75 dpi image of text, converted to CIELAB at 75 dpi.

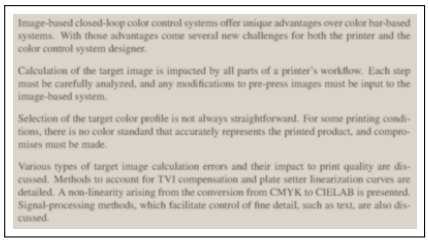

Figure 29: 75 dpi image of text, converted to CIELAB at 300 dpi.

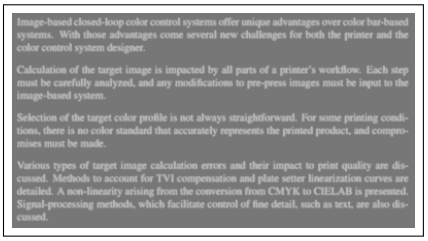

Figure 30: Difference between the two above images (right minus left). When converting to CIELAB at 75 dpi (left), the text appears darker due to the non-linearities in the conversion to CIELAB. When converting to CIELAB at 300 dpi (right), the effect of the non-linearity is reduced, and the target image better matches the scanned image.

The magnitude of the darkening effect is best illustrated through a one-dimensional example. Figure 31 shows both a filtered and unfiltered black "pulse" signal of tone value vs. pixel. The unfiltered signal is black from pixels 15 to 25, and white elsewhere. This signal could represent a solid black line, a slice through text, etc.

When converting to CIELAB prior to downsampling, the unfiltered version of the tone signal is used, and the CIELAB signal is filtered. When converting to CIELAB after downsampling, the filtered version of the tone value signal is used, and no filtering is done in CIELAB space. Figure 32 shows both of these signals converted to L^* using the ISOnewspaper26v4.icc profile, and the difference (ΔL^*) between the two L^* signals is shown in Figure 33. The signals differ by up to 6 ΔL^* at the black edges, where there is a high contrast.

Figures 34 and 35 show the same plots for a black pulse function, but the conversion to CIELAB is done using the 20% TVI ICC profile used earlier in the paper. Now, the signals differ by only 3 ΔL^* at the black edges.

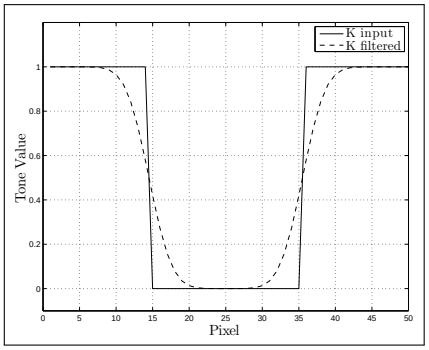

Figure 31: Filtering applied to a black "pulse" signal. The unfiltered signal is used when converting to CIELAB at a high resolution, and the filtered signal is used when converting to CIELAB at a low resolution.

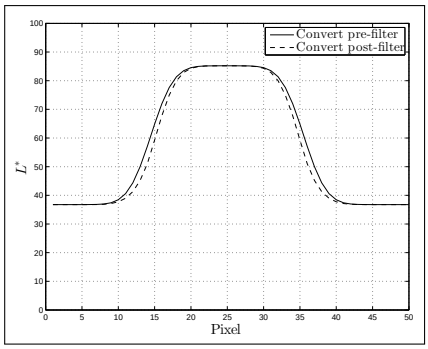

Figure 32: Converting black to L^* before and after filtering, 26% TVI.

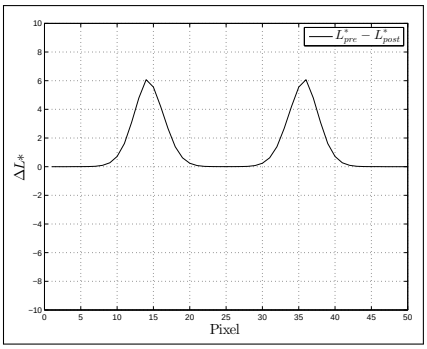

Figure 33: ΔL^* at black edges due to non-linearities, 26% TVI.

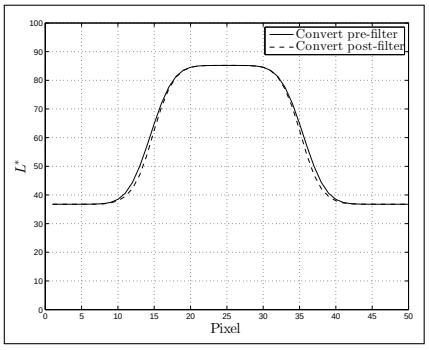

Figure 34: Converting black to L^* before and after filtering, 20% TVI.

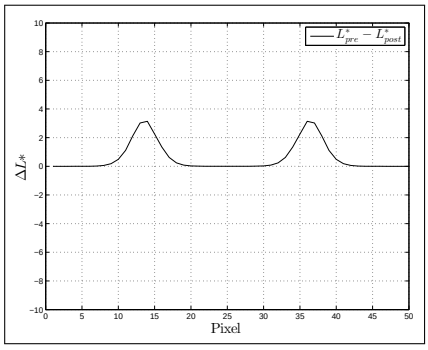

Figure 35: ΔL^* at black edges due to non-linearities, 20% TVI.

The reason for this difference in ΔL^*, and the primary source of the non-linearity, is due to the relationship between CIELAB color components and tone value. Figures 36 and 37 plot L^* vs. tone value for black, using 26% TVI and 20% TVI, respectively. The 20% TVI profile results in a curve that is closer to linear than the 26% TVI curve. If this relationship was exactly linear, there would be no difference between computing CIELAB for black before or after downsampling. In most

cases, however, the relationship between CIELAB and tone value is not linear, so the conversion to CIELAB must be done at a higher resolution.

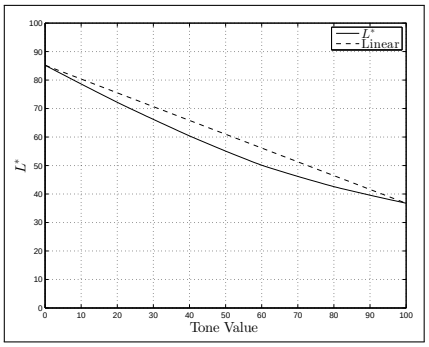

Figure 36: L^* vs. black tone value, 26% TVI.

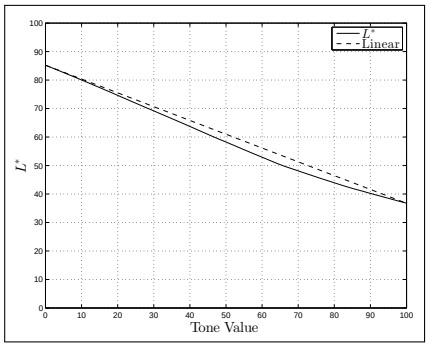

Figure 37: L^* vs. black tone value, 20% TVI. The 20% TVI conversion is closer to linear than for 26% TVI, causing the darkening effect to be less.

Both the signal processing considerations and image-conversion non-linearities mandate that the image processing be performed by software designed specifically for the image-based system, using high-resolution images. In other words, a low-resolution image output from a printer's RIP or other software will not suffice for image-based color control. Most software packages use very "sharp" low-pass digital filters for downsampling, which means some amount of aliasing may still be present, and the blurring characteristics of the imaging device are not accounted for. Furthermore, the conversion to CIELAB must be done at a high resolution to avoid the non-linearities just discussed.

Summary

CIELAB optimization systems strive to achieve the best visual match to the target image. Several issues must be addressed to ensure an accurate target image and allow such a system to perform optimally.

The system needs the ink composition of each pixel. The images sent to the system must accurately represent GCR levels and TAC limits. The 1-bit TIFFs used for making plates are appropriate as input to the image-based system.

TVI and plate errors can result in significant color errors in the optimized image. Most of these errors can be reduced or completely removed. A mismatch between the process TVI and target TVI can be mitigated through TVI compensation curves. When curve errors are unavoidable, simply accounting for them results in a predictable optimized image.

A mismatch between a profile and a printing condition may result in variability in the optimized image. If the variability is too great, a profiling test may be performed to define a new target. This solution is not without challenges.

Removing certain image processing artifacts facilitates accurate control of fine detail such as text, line art, and many images. Aliasing is prevented by optically blurring the image and using anti-aliasing filters when downsampling. A non-linearity in the conversion from CMYK to CIELAB is mitigated by doing the conversion at a resolution much higher than that of the imaging device.

Appendix A

The color-control-in-the-work (CCIW) problem definition and its solution are presented here. For further details, see Seymour (1999).

For CCIW, we wish to determine the amount to change each ink given a color error for a number of pixels. If the ink changes are expressed in terms of density, then we want to find

$$D = \begin{bmatrix} \Delta D_C \\ \Delta D_M \\ \Delta D_Y \\ \Delta D_K \end{bmatrix} \tag{4}$$

where ΔD_C, ΔD_M, ΔD_Y, ΔD_K, are the required solid ink density changes for cyan, magenta, yellow, and black, respectively.

For N pixels, the measurement error vector is expressed as

$$E = \begin{bmatrix} \Delta L_1^* \\ \Delta a_1^* \\ \Delta b_1^* \\ \Delta L_2^* \\ \Delta a_2^* \\ \Delta b_2^* \\ \vdots \\ \Delta L_N^* \\ \Delta a_N^* \\ \Delta b_N^* \end{bmatrix} \tag{5}$$

where $(\Delta L_i^*, \Delta a_i^*, \Delta b_i^*)$ is the error vector (target minus measured) for pixel i.

Using a model or press data, the relationship between change in density and change in CIELAB color may be determined. This provides a set of twelve sensitivity factors

for each pixel. For N pixels, the sensitivity matrix is written as

$$S = \begin{bmatrix} \frac{dL_1^*}{dC} & \frac{dL_1^*}{dM} & \frac{dL_1^*}{dY} & \frac{dL_1^*}{dK} \\ \frac{da_1^*}{dC} & \frac{da_1^*}{dM} & \frac{da_1^*}{dY} & \frac{da_1^*}{dK} \\ \frac{db_1^*}{dC} & \frac{db_1^*}{dM} & \frac{db_1^*}{dY} & \frac{db_1^*}{dK} \\ \frac{dL_2^*}{dC} & \frac{dL_2^*}{dM} & \frac{dL_2^*}{dY} & \frac{dL_2^*}{dK} \\ \frac{da_2^*}{dC} & \frac{da_2^*}{dM} & \frac{da_2^*}{dY} & \frac{da_2^*}{dK} \\ \frac{db_2^*}{dC} & \frac{db_2^*}{dM} & \frac{db_2^*}{dY} & \frac{db_2^*}{dK} \\ \vdots & \vdots & \vdots & \vdots \\ \frac{dL_N^*}{dC} & \frac{dL_N^*}{dM} & \frac{dL_N^*}{dY} & \frac{dL_N^*}{dK} \\ \frac{da_N^*}{dC} & \frac{da_N^*}{dM} & \frac{da_N^*}{dY} & \frac{da_N^*}{dK} \\ \frac{db_N^*}{dC} & \frac{db_N^*}{dM} & \frac{db_N^*}{dY} & \frac{db_N^*}{dK} \end{bmatrix} \tag{6}$$

where $\frac{dL_1^*}{dC}$ is the change in L^* with respect to cyan density for pixel 1, $\frac{dL_1^*}{dM}$ is the change in L^* with respect to magenta density for pixel 1, etc.

For at least two pixels, we have an over-determined set of equations,

$$E = S \cdot D$$

The solution is given by

$$D = S^+ \cdot E$$

where S^+ is the Moore-Penrose pseudo-inverse of S. The singular value decomposition of S is given by

$$S = UWV^T$$

where W is a diagonal matrix. The pseudo-inverse is given by

$$S^+ = VW^+U^T$$

where the element at the ith row and jth column of W^+ is given by

$${w^+}_{ij} = \left\{ \begin{array}{c} \frac{1}{w_{ij}}, i = j, w_{ij} \neq 0 \\ 0, otherwise \end{array} \right\}$$

In other words, the pseudo-inverse is computed by first computing the singular value decomposition of S, and then inverting the non-zero elements of W. This matrix is then multiplied by E to determine the solution for D, or

$$D = VW^+U^T \cdot E$$

Alternatively, the solution can be obtained by multiplying both sides of the first equation from the left by S^T

$$S^T E = S^T SD$$

Now if we multiply from the left by the inverse of $S^T S$, and move D to the left-hand side, we get

$$D = (S^T S)^{-1} S^T E$$

If $S^T S$ is not invertible, that means there are an infinite number of solutions to the CCIW equation. The pseudo-inverse will then choose the one that minimizes the norm of D, and the least squares solution is instead

$$D = (S^T S)^+ S^T E$$

This form is more practical for CCIW, since $S^T S$ is 4x4, and the pseudo-inverse calculation is less computationally intensive. It can be shown that this form of the solution is mathematically equivalent to the first form presented.

Literature Cited

ISO 12647-2 (2004). Graphic technology – process control for the production of half-tone colour separations, proofs and production prints – Part 2: Offset lithographic processes.

ISO 12647-3 (2005(E)). Graphic technology – process control for the production of half-tone colour separations, proofs and production prints – Part 3: Coldset offset lithography on newsprint.

ISO 15339 (2012). Graphic technology – printing from digital data across multiple technologies – Part 1: Principles and characterized reference printing conditions, draft international standard ISO/DIS 15339-1.

Seymour, J. (1999). Markless color control in a printing press. *United States Patent No. 5,967,050.*

Seymour, J., Tiltman, S., & Nelson, A. (2010). Closed-loop color control in the work. In *TAGA Proceedings*.

SNAP (2011). Specifications for newsprint advertising production, cold set printing standards.

Stanton, A., & Seymour, J. (2011). An analysis of the current status of process control for color reproduction in newspapers. In *TAGA Proceedings*.