Creating Print Standards

Some ideas on how to create optimized, interchangeable reference standards for all printing methods

Don Hutcheson, Hutcheson Consulting

Version 032

NOTE: This document is a work in progress and will be re-issued periodically as needed. Version number is listed at the bottom of each page. Please make sure you have the latest version. Comments and suggestions will be gratefully received at (don@hutchcolor.com).

CONTENTS

Introduction

This document proposes a coordinated philosophy for print standards development, the two main principles of which are a “goals-based” approach to device calibration, and “compatibility” between multiple standards.

A primary goal of this philosophy is to reduce as much as possible the basic appearance differences between different standardized processes, and thus simplify the preparation and exchange of printable files between them.

An equally important goal is to create a standardized approach to the development process leading to print standards, so that future standards can be developed more quickly and effectively, and share maximum compatibility with existing standards.

Although we may not achieve them overnight, these goals constitute an interesting challenge with long term benefits for all, and few, if any, serious technical obstacles. Perhaps the most compelling reason to accept this challenge is so that future generations will look back at our standards efforts and realize we cared enough to try and “get it right”.

“Goals-based” standards

Most print standards efforts to date have been “observational”, rather than “goals-based”.

By “observational” I mean we throw a target on a press, set the press to some traditional parameters, and then measure whatever emerges. What we get is a characterization of that particular press run, with all its strengths and weaknesses, but not necessarily a good representation of what an “ideal” press might have done.

By “goals-based” I mean that defining a print appearance standard presents an opportunity to set some idealized performance goals for the press run(s) on which the characterization data are based.

The main question with a goals based approach is what to aim for, or more specifically, what characteristics would an “ideal” press exhibit? For most print processes (newsprint, publication, commercial, flexo, etc.) the choice of ink, substrate, and solid densities are governed by tradition or price, and not open to change just to suit the standards process. But some equally important characteristics like gray balance and print characteristic curve, which might be assumed to be tied to media choice, can be adjusted to virtually any arbitrary “standard” condition through relatively simple mechanisms such as CTP curves.

“Compatible” standards

A major benefit of a goals-based approach to standardization is that multiple different printing systems or media can be engineered to look relatively similar to each other (within gamut limits) simply by basing some, if not all, of their calibration characteristics on common goals. For example, if pictorial gray balance, contrast and density were similar from medium to medium, at least monochrome images could appear almost identical to each other. Further, if the basic CMY colorants (inks) were also similar in hue, colored images would also tend to match reasonably closely, even though a perfect "color match" might require some form of n-dimensional color transformation.

Until recently the goal of compatibility between standards was out of reach, due to limitations in pre-press, platemaking and proofing technologies. But today, thanks to widespread adoption of CTP and ICC color management, compatible standards are now quite achievable.

Benefits of compatible standards

The fundamental advantage of compatible standards is that a file prepared for one process, but reproduced on another (by accident or otherwise), will still look "reasonable", even if not exactly as intended. The more similar the overall contrast, lightness and gray balance of all printing processes, the less chance there will be of disappointment in the case of "accidental re-purposing".

Secondly, by sharing a common general appearance, the cost of deliberate conversions between processes can be reduced, or even eliminated in many cases.

Yet another advantage is that by decoupling the relatively simple "calibration" of neutrals from the complex n-dimensional "mapping" of colors, transformations between one process and another will require less complex tools. For example, if neutral tonality and gray balance are already similar, residual hue and saturation differences can often be resolved with a simple matrix or polynomial, rather than a full 4-D device-link.

Simple vs. complex functions

For the purposes of this document, the visual characteristics of any imaging process can be sorted into two groups, "simple" (one-dimensional) and "complex" (n-dimensional). The number of dimensions dictates how each characteristic can be characterized or calibrated.

The most obvious "simple" characteristics are gray balance and print density curve (PDC - a.k.a. "dot gain" curve), both of which can be defined or compensated with simple 1-dimensional graphical curves.

Complex characteristics include non-linear effects like ink hue, saturation, and trapping, whose combined effects can only be controlled with fairly complex functions, typically stored in a 3-D (or higher) LUT.

While it is unreasonable to expect complex functions to behave the same way on printing processes with widely differing colorants, it is both reasonable and relatively easy to "force" simple characteristics, like gray balance and tone shape, to behave similarly from system to system, regardless of ink color, substrate or other variables.

The only real question is, how do we determine (and agree on) universal standards of tone curve and gray balance that can be applied to virtually any printing standard, without drastically altering traditional expectations?

Print Density Curve

Today's separate imaging paths for plates and proofs make it easy to force any proofing or printing process to imitate virtually any arbitrary print-density, or "dot gain" curve.

In seeking a suitable common PDC, the big questions are;

- How do we determine the "ideal" PDC?
- How can one PDC be applied to devices with different dynamic ranges?

To gain or not to gain?

By default, CTP eliminates film-induced 'dot gain', but some dot gain needs to be simulated to remain compatible with legacy files. There are conflicting opinions on what the 'right' legacy dot gain curve should look like. Should it simulate analog proofing or an actual press? If a press, positive or negative plates? Should printers with no legacy requirements be forced to print with legacy gains, or should legacy files all be converted up-stream so we can all print with linear plates?

The argument that zero dot gain is somehow 'better' than traditional dot gain is false, because dot gain does not limit the color gamut of a device. But by the same token, emulating traditional dot gain with a forced CTP curve is no more inherently "right".

Whether the CTP process is set to 'linear', (with far less dot gain than before), or to emulate traditional dot gain is relatively unimportant, so long as the pre-press color conversion process (e.g. an ICC profile) is calibrated to that condition.

A strong argument can be made that if some printers don't need to emulate traditional dot gain, but others are obliged to for legacy reasons, some amount of emulated gain should be standardized. The question remains, "what gain curve or PDC should a CTP press emulate?" Let's look at the main options.

Option 1: Plate-linear

A 'plate inear' CTP system makes plates on which a 50% file value creates a 50% dot on the plate (measured by dot area) and so on for all other dot percentages. In practice this is not truly a "linear" printing situation because there is still some dot gain on press, but all traditional film-induced dot gain is eliminated.

Specifying "plate linear" as part of a standard does not in fact define a specific PDC, because the overall print density curve is affected by other factors, however it is attractive from a simplicity viewpoint.

On the down side, linear plates are generally incompatible with 'legacy' CMYK files (separated for traditional dot gain) because legacy files will typically print much lighter than expected, and cannot be mixed with files separated for a linear CTP system. Even today, most files supplied from "unknown" sources are adjusted by default to allow for conventional film-based gain values.

The danger of over-linearizing

A potential problem with linear plates is that the very act of "linearizing" can over-compensate for small dot loss amounts near the highlight point, and result in unnaturally heavy printing in the lowest percent range. This in turn can introduce profile errors if the dot-density relationship is not linear between white paper and the first tone value in the profiling target (usually about 12%).

Option 2: Press-linear

It is possible to calibrate the plate/press combination to truly print with zero physical dot gain, in other words so that a 50% file value produces a measured dot percentage on press of 50%. This requires a custom CTP curve to eliminate whatever dot gain occurs on press, but produces much lighter results than traditional printing, or even “plate-linear” printing and is therefore clearly not a good choice.

Option 3: Matching legacy analog curves

If many legacy files must be printed on a CTP system, a logical option is to adjust the CTP curves to match the analog proofer on which those files were approved, or to match the curves of the press on which they were printed from analog (film-based) plates.

The advantage of this policy is that legacy files and files supplied with legacy proofs will look the same printed on both a CTP and an analog press. Most digital halftone proofing systems seem to be calibrated by default to match the PDC of their analog predecessors, so matching legacy analog proof curves is already a kind of unofficial standard. Many printers initially starting with CTP have adjusted their process to match traditional gains, and there seems to be a comfort value, as well as a simple logic, in reproducing what has been done in the past.

The big question is, does matching a legacy analog proof or press curve cause any serious problems or future incompatibility? Color gamut is entirely unaffected by choice of PDC, so long as solid densities and trapping remain the same. The only problems I can see with this option are that (a) not everyone is teaching the same idea, and (b) each legacy system has a different natural PDC, hence there is still no “standard” curve.

Option 4: Matching TR001 “Highlight Contrast”

The option I have used with great success for over four years is to adjust the CTP curve so that the press sheet emulates the PDC or contrast of TR001 from paper white to mid-tone. If the density range of the new device does not match the DR of TR001, variable compression/expansion is applied from mid-tones to solid density. This option results in a virtually perfect contrast match between quite diverse devices in the visually-critical lighter tones, with a very acceptable “best fit” match in darker tones.

If the maximum ink density of the new process is less than the TR001 spec, some shadow contrast will be lost, but highlights will not lose density. If the new process has higher solid ink densities, shadows will look better, but highlights and mid-tones will not gain in density. In 36 years of pre-press experience, I’ve found this to be the best way to maintain an apparent “match” between an original and a proof, or between devices with different available contrast.

This option has the advantage that many major printers and publishers already support TR001, while many companies use TR001-based separation parameters and proofing systems in the absence of a true commercial printing standard.

As TR001 is based on a press run made with analog plates, this option is very similar to matching analog plate curves, and therefore "legacy compatible", at least in a negative plate environment.

By coincidence or luck, the paper-subtracted TR001 neutral density curve appears to match very closely the result of printing un-linearized CTP plates on commercial stock at commercial densities. It also closely matches the European traditional positive plate curve, making it an ideal basis for an international curve standard.

Option 5: Matching a modeled function

The option that is most likely to please both practical printers and theoreticians is to match the CTP curve to a mathematically calculated curve derived from either TR001 or data from multiple analog proofing systems and/or press runs. The mathematical calculations "smooth out" any unwanted bumps in real measured data, making it easier to obtain a good ICC profile.

The ideal 'standard' print density curve

It seems from the above logic that the best option is to maintain the same basic PDC established in TR001, as it provides maximum legacy compatibility and is already accepted as a US defacto standard. By adjusting it to fit a simple mathematical formula or function it can also serve as the basis for an agreeable international curve standard.

Gray balance

Today's presses and pre-press proofing systems vary quite noticeably in their reproduction of "gray balance". Creating a new printing standard is a great opportunity to define gray balance more precisely than before.

As every other color reproduction industry (except printing) has known for decades, gray is the most important "color" in any color reproduction system, largely because it is the most unambiguous "memory color", and the yardstick by which the human visual system judges other colors.

Gray balance or white balance is the first characteristic to be controlled in ...

- video camera calibration
- photographic printing
- slide film processing
- digital photography
- cinematography

... however amazingly, the definition and control of gray balance has been largely neglected by offset lithography. For example, SWOP defines the vital characteristic of gray balance only indirectly, as a by-product of individual ink hue and dot gain specifications, or as a CMY combination that matches the color of a tint of "black" ink, (which itself may not be neutral). As a result, a nominally "gray" dot percentage combination (e.g. 50 cyan, 40 magenta, 40 yellow) can produce a very different visual "color" on different so-called "SWOP presses" or "SWOP-certified proofing" systems. This does not invalidate SWOP as a general-purpose "specification", but it does show why SWOP is no longer a satisfactory substitute for a real printing "standard".

An industry newly-aware of CIE color spaces and ICC color management needs a universally-agreed, unambiguous definition of how to make gray on any printing press.

What makes gray special?

Gray has some unique qualities that justify special emphasis in any print standard;

- Since gray lacks any hue or saturation, it is a visual yardstick by which we judge the accuracy of pastel colors.
- Unlike other colors, Gray can be controlled simply by adjusting simple one-dimensional functions, such as the individual C, M and Y calibration curves present in all CTP or digital proofing systems.
- If gray balance is controlled through colorant-specific functions such as global CMY curve functions, it does not affect the total color gamut of the system.
- Gray balance can be 'pre-corrected' as an isolated variable, separate from the typically more complicated functions needed to correct hue and saturation.
- Gray balance can typically be monitored more quickly and easily than other variables. A simple 3-color gray patch can often indicate the overall health of a total imaging system, including the cumulative effects of ink densities, trapping and dot gain.

Gray balance options

For the purposes of this document, gray balance is defined as the CMY percentage dot ratio (or absolute percentage dot values) needed to make a "neutral gray" when the cyan dot value is 50%, where "neutral gray" is defined as $a^* = b^* = 0.0$

Note that gray balance should ideally be defined for a complete range of CMY percentages, and both with and without the effects of black ink, however the "at 50% cyan" convention is usually sufficient to define the "general" gray balance of any system. The influence of paper color and black ink will be discussed later.

There are at least five logical options for measuring or defining a "general" gray balance;

Option 1: Device-specific

Allow each ink/substrate combination to have its own gray balance, based on the natural effects of substrate, ink hue, ink densities, ink sequence, etc.

Standards-specific logic: low

Chance of widespread acceptance: low

Main objection: "Gray balance" will vary from system to system, conflicting with the goal of "common standards" and making it difficult to match a reference image or proof.

Option 2: Match the substrate

Define gray balance as the ratio of CMY that matches the color (a* and b* values) of the substrate.

Standards-specific logic: low

Chance of widespread acceptance: medium

Main objection: Paper stocks are often non-neutral, hence gray balance will vary from system to system, conflicting with the goal of "common standards" and making it difficult to match a reference image or proof.

Option 3: Match the black ink

Define gray balance as the ratio of CMY that matches the color (a* and b*) of a tint of the black ink at the same lightness level.

Standards-specific logic: low

Chance of widespread acceptance: medium

Main objection: Black ink is often not neutral, hence gray balance will vary from system to system, conflicting with the goal of "common standards" and making it difficult to match a reference image or proof.

Option 4: Normalized traditional

Define gray balance as 50 cyan, 40 magenta, 40 yellow, as per the average of most offset traditions.

Standards-specific logic: medium

Chance of widespread acceptance: high

Main objection: Ink-on-paper printing (offset, letterpress, gravure, etc.) is the only imaging branch I know of that did NOT start with the assumption that neutral grays should be produced with equal quantities of each primary colorant, which is a basic rule of color photography, video, cinematography, etc.

I have demonstrated numerous times that, simply by increasing cyan density or pigmentation, a neutral gray can be printed with equal CMY percentages of otherwise normal offset inks. If a modified CMY inkset CAN achieve gray balance with equal values, why do we continue to print with what is effectively a weak cyan ink?

I have heard countless explanations of why offset gray balance requires unbalanced CMY values, but most sound more like apologies for the fact that cyan is traditionally a weaker ink, than good explanations for WHY we still run with weak cyan.

If we must keep to this embarrassing tradition, can we at least call it from now on "gray imbalance"?

Option 5: Equal CMY

Define gray balance as C = M = Y (e.g. 50 C, 50 M, 50 Y).

The main incentive to convert from traditional "gray imbalance" to an "equal CMY definition of gray balance is that anything else is illogical. One practical advantage of "Equal CMY" gray balance is that CTP curves and/or proofer curves could be altered for CMY together without affecting gray balance. Another advantage is that the calculation of gray balance, UCR and GCR would be simplified, and could be more easily moved to an application-level function or an in-RIP process, rather than being embedded in every CMYK profile.

Standards-specific logic: high

Chance of widespread acceptance: low (initially at least)

Recommendation

The “normalized traditional” approach is more likely to be accepted by end users in the offset printing world. However it perpetuates a rather illogical concept of gray balance of which, as an industry, we should be ashamed. If normalized traditional is officially adopted as a common standards goal, such adoption should at least be tagged with a note to future workers about the desirability to move to equal CMY.

My own preference is to make a cold-turkey break with tradition and define gray as equal CMY, but I doubt that will win much support amongst traditionalists. Certainly any equal CMY standard must allow for legacy files and staggered industry adoption and, amongst other things;

- contain a well-crafted explanation for why equal CMY was adopted,
- embed an automated identification tag in files converted for equal CMY,
- provide for simple and efficient conversions between standards.

A possible compromise would be to move towards equal CMY in stages, perhaps increasing M and Y by 2% a year for five years. Another suggestion is to allow two parallel standards, normalized traditional and equal CMY, with the option to convert between the two via an ICC link or profiles.

None of these options is ideal. The question is, which offers the most short term and long term benefits.