

Maximizing Neugebauer Primary Control: Beyond Halftoning

Peter Morovic, Hector Gomez, Utpal Sarkar, Sergio Etchebehere, Ján Morovic, HP Inc.

Abstract

A pattern that is to be formed by a colorant on a substrate is defined in terms of the halftone image that instructs a printing system. This image contains information at each pixel as to which colorants to use at which drop states, i.e., Neugebauer Primaries (NPs) such as CC (two drops of cyan ink) or CM (a drop of cyan and a drop of magenta ink). However, the transformation from this digital input to a final printed pattern involves other steps as well as physical, colorant and substrate interaction effects. An important contributor to the final output is the partitioning of a halftone image into its complementary sub-images that are printed in each pass of a multi-pass or multi-bar printing system (common in inkjet). Typically, great attention is given to the goodness of a halftone image – good pattern distribution, good NP and ink choices – however, partitioning can also have a significant impact on the output both in terms of avoiding unwanted physical interactions (e.g., coalescence, alignment artifacts) and exercising control over the pattern. This can lead to closer correspondence to the halftone and improved print properties (such as grain and robustness) and can be achieved by co-optimizing halftoning (responsible for the spatial distribution of the halftone image) and partitioning (responsible for distributing it among the N passes or bars of a given printing system). The results are not only good overall halftone images as has been the case already, but also good partial halftone sub-images whose gradual printing in different passes results in maximized NP control and thereby its effect on the printed color.

Introduction

Color and other print-quality attributes are the result of both careful choices in the digital domain and the conditions under which it is printed. For example, using light inks, dilutions or lower drop weights for lighter colors is critical for good print quality in those colors, in particular for grain, as is using dispersed-dot or blue-noise patterns, especially on inkjet systems. Likewise, printing at slower speeds, with more passes on a more stable set-up will result in better quality even for the same input halftone image as compared to printing at high speeds with few passes and large variability. Whether a digital pattern results in the best possible output depends both on how successfully the halftone is optimized and on how it is then printed.

On the digital input side, with the introduction of Halftone Area Neugebauer Separation (HANS), the solution space has been vastly expanded and new options for optimization have been demonstrated. For example, it has been shown that careful choices of the digital content can have a significant impact on the output image quality [5] as well as other attributes such as ink-use [6]. However, these have been shown without particular regard to the printing conditions and on printers ranging from high-end photo-quality ones (with great repeatability and low levels of system

noise) to industrial signage printers (with much larger degrees of variability and sources of noise).

Physical pattern formation

A key challenge in printing is the complexity of printed pattern formation resulting from colorants being deposited on a substrate. Instead of such physical patterns matching the digital pattern driving a printing system's imaging engine, it is the result of a complex interplay of factors. The colorant's rheology (if it is a liquid) [1], electrical properties (for toners), particle shapes, sizes and distributions [2] and chemical composition, the relationship of these among the colorants and their relationship with the macro- and micro-structure and chemistry of the substrate [3] all mean that even when a fixed colorant pattern is deposited on different substrates, the resulting color can vary dramatically.

A second level of complexity is introduced in how the colorant pattern is obtained in response to the digital patterns generated to drive a print imaging system. Here the speed at which the imaging engine moves relative to the substrate, the time sequence of the printed patterns being built up and the specific local geometry of the pattern all play decisive roles. An important element here is the way in which the total printed pattern is divided into parts that are printed sequentially. The structure of these parts, coupled with the impact of drying and absorption taking place between their printing, has a major effect on whether and how much the resulting physical structures coalesce as the total pattern is built up [4]. This, in turn, impacts the homogeneity and perceived grain of the resulting print. The total effect of such complexity and inter-relatedness is that a given halftone pattern can result in a variety of colors, uniformities and levels of grain depending on the materials used to print it and on the temporal sequence and spatial structure of the total and partial patterns involved.

HANS paradigm

One of the principles of the HANS paradigm is that ink-interactions defined through Neugebauer Primary area coverage (NPac) vectors and spatial distribution are decoupled. Once a mapping from input color to NPac is defined, the role of halftoning is then only that of spatially distributing the NPac according to the NP probabilities it dictates, without the ability to introduce new NPs or alter those defined at the input NPac. While several halftoning approaches are possible, such as device state error diffusion [6] based processes, it has been shown previously that a careful choice of pattern encoded in a tileable selector-value array used in the Parallel Random Area Weighted Coverage Selection (PARAWACS) approach results in the highest levels of IQ as well as deterministic pattern control and additional benefits [7]. There, the focus was on the goodness of pattern design to generate the highest levels of IQ, the property of pattern determinism has not been exploited so far and turns out to be a key enabler for extending print control beyond the halftone image.

By exploiting the determinism of this halftoning approach – i.e., that the output halftone image is uniquely dependent on the structure of a selector-value array (which encodes the type of pattern to be produced in the halftone) – it turns out to be possible to extend the pixel-level control that HANS introduced also further down the imaging pipeline. Here, the next step is that of also designing optimal patterns in each pass of a multi-pass printing system. While the final halftone remains the same as before, the sequence in which it is laid down can be very different and turns out to have a significant impact on print attributes such as color, grain and more fundamental substrate and ink interactions such as coalescence. This is achieved not only by means of more robust and pleasing patterns in each pass, but extends also to spatio-temporal control over the reproduction of NPs.

Partitioning

The problem of partitioning [8] a halftone into complementary sub-halftones to be printed in each pass (or with each bar of a multi-pass or multi-bar system) is that of 1. maintaining (or improving) the properties already present in the halftone also during the printing process and 2. mitigating any adverse effects of the printing process such as alignments (between passes/bars, media advance etc.), ink-substrate interactions, firing-frequency effects, etc.. For example, whether the sub-halftone is well distributed (for a single ink) will impact coalescence: higher clustering of dots in a single pass when colorants are fresh or wet will make them coalesce and therefore behave differently compared to them being well spread out without unintended ink-ink interactions. Likewise, whether NPs are printed in the same pass (e.g., whether CM is printed in one pass or in two separate passes) will also impact how likely these are to result in a printed version of cyan and magenta overprinting.

Figure 1 shows an example of a naïve partitioning that is agnostic of the halftone contents. Starting with a halftone (left), it is first considered as per-colorant layers for the two colorants used, cyan and magenta and then partitioned into two passes as shown.




Figure 1. From halftone (left) to colorant layers (cyan and magenta) to per-pass partitioned halftones (for a 2-pass example).

The task of partitioning in this example is that of distributing each pass as well as possible (given some constraints) and considering colorant interactions. As can be seen, this partitioning scheme does not take into account whether inks are meant to interact – i.e., partitioning is done in a colorant independent way, and furthermore results in apparent clustering (e.g., see pass 0 for both magenta and cyan). Figure 2 shows an alternative partitioning that takes into account both the clustering of neighboring pixels as well

as uses the same strategy for both cyan and magenta colorants, thus printing their combination in the same pass.




Figure 2. From colorant layers (cyan and magenta, same as in Figure 1) to per-pass partitioned halftones (for a 2-pass example) where both ink-interaction (i.e. printing CM in the same pass) and clustering (i.e., avoiding printing neighboring pixels in the same pass) is taken into account.

The challenge of such partitioning schemes lies in the fact that traditional imaging pipelines either do not possess the aforementioned determinism of halftoning (e.g., error-diffusion halftoning is content dependent and often incorporates deliberate randomization) or are known to have limitations affecting image quality (e.g., traditional matrix based halftoning). Given the lack of determinism or predictability, designing a static partitioning pattern that when applied to such variable halftones results in the desired outcome is, in general, an unsolvable problem.

Without loss of generality and assuming the objective to be blue noise partitions (i.e., avoiding clustering and avoiding voids while exhibiting a pleasing randomness) the partitioning problem can be stated as follows:

1. For an arbitrary (unknown) blue noise pattern (the halftone),
2. design a sequence of static blue-noise patterns (the partitioning), as many as passes or print-bars,
3. such that an AND operation between the halftone and each partitioning also results in blue-noise characteristics (the partitioned per-pass halftone patterns).

Two blue-noise patterns combined are not guaranteed to result in blue-noise and in fact may result in something closer to a white-noise pattern, in that it will typically contain plenty of low-frequency components. It may both contain voids, regions that are sparser, and clusters, regions that are denser than would be expected for blue noise of the same density. Fig. 3 shows an example of a halftone image (left), a blue-noise partitioning designed for an 8 pass print-mode (thus covering up to 1/8 of the density) and the resulting combined pattern for one pass. What can clearly be seen is that the goodness of the original halftone pattern is lost after partitioning is applied and while the final, blue noise halftone will be fully reconstructed after the 8 passes, it will be done so with unintended ink and substrate interactions and white noise patterns in each pass. Combinations of patterns like those shown on the right will also likely create further undesired artifacts when exposed to misalignment and other physical sources of noise during the printing process.




Figure 3. Example blue-noise halftone (left) and independent blue-noise partitioning for 1/8 of the density (e.g., for an 8-pass print-mode) and the resulting pattern in one of the passes.

The problem changes when halftoning is deterministic, since now the pattern with which a partitioning scheme needs to interact well is a single, static one, instead of being content dependent. Hence, the problem statement can be reformulated as follows:

1. Given a known blue-noise pattern (defined by the selector-value array used by PARAWACS halftoning),
2. design a sequence of static blue-noise patterns (the partitioning), as many as passes or print-bars,
3. such that the partitioned halftones for each pass preserve blue-noise properties and enable control over ink-combinations.

PARAOMASKING

The approach taken here is that of building partitions (also referred to as masks or shingles) that are optimized based on the known PARAWACS halftone pattern, as defined by the selector-value array. The assumption here is that this array is an input to the process and there is a rich body of literature on how to generate such arrays for a variety of applications [9].

What is specific to PARAWACS is that the array needs to contain the same number of values for each tone, so for example, a 12-bit PARAWACS selector-value array will have tones from [0 to 4095] and each of these tones needs to be represented by the same number of pixels in the array, resulting in a uniform probability distribution. Furthermore, each tone individually as well as all possible sub-ranges of tones should be as close to the target blue-noise pattern as possible. While this is an over-constrained problem and as such only approximately solvable, one way to build such patterns for PARAWACS is to prioritize the goodness of sub-ranges. For example, the pattern may be built such that lowest and the highest tones have the best possible patterns – and will be mapped to use the lightest and darkest NPs (determined by ordering NPs by lightness) – while sacrificing some of the goodness of NPs surrounded by both lighter and darker ones.

Then, let's assume a PARAWACS selector-value array has been provided, and that it has been generically optimized for the quality of the final image. Analogously to PARAWACS being independent of the specific input, the partitioning too needs to be input independent, but at the same time optimize the content printed in each pass.

The final number of drops on paper at every position will not depend on the partitioning, only their distribution per pass does (and therefore their order as well). There is a strong correlation between print quality and mask properties however. The following assumptions are taken, which hold generally for inkjet printers:

- To avoid coalescence, it is beneficial to maximize distances between nearby drops. This also helps to reduce the highest firing frequencies.
- Regular distributions visually amplify variability of the printing system, since small deviations get replicated in a regular pattern.

The combination of homogeneity without regularity essentially defines the aforementioned blue noise distributions: these are characterized by the absence of low frequency components, and exhibit noisiness and isotropy in high frequencies. Such patterns are visually experienced as smooth and pleasing, so it is often an objective of halftoning (see Fig. 3 earlier for examples of blue-noise patterns). Because of the assumptions listed above, it will also be taken as the objective for partitioning here.

To produce or evaluate blue noise patterns, the potential of the pattern will be used. This is a common approach, [9, 10], and has a clear intuitive explanation: imagine drops to be repelling, then a low potential corresponds to a homogeneous distribution. Introducing or maintaining an element of randomness will take care of the isotropy and noisiness.

The experiments here have been conducted using a Gaussian point potential, i.e.:

$$V(r) = e^{-r^2/2s^2}$$

but variations are possible. Here s is a scale factor, that depends on the density of pixels, and will be varied throughout the execution of the algorithm, and r is the distance, so this potential is inherently isotropic, as will the final result be, assuming sufficient randomness is maintained.

The optimization algorithm is then as follows:

Assume NPs are ordered so that the lowest PARAWACS values correspond to the darkest NPs (the aforementioned lightness based NP ordering). The algorithm to optimize the masks for each halftone level is as follows:

- 1) Generate a list of all entries added in this level and put it in random order.
- 2) Produce a binary pattern corresponding to all preceding levels.
- 3) Partition this to obtain a pattern for each of the passes and compute the potential for each of these by convolving it with the potential kernel.
- 4) For each position in the entry list:
 - a. Select the pass to assign the drop to, based on the potential for that pass, for good results the pass for which the potential is lowest is picked at that position in the mask at the corresponding position.
 - b. Update the potential for that pass.




Figure 4. Example actual blue-noise halftone (left – as as that shown in Figure 3 earlier), PARAOMASKING blue-noise partitioning for 1/8 of the density (e.g., for an 8-pass print-mode) and the resulting pattern in one of the passes (left).

The outcome of the process above is illustrated in Fig. 4 showing the same halftone as before, coupled with the proposed approach to partitioning (middle) and the combined result for a given pass. As can be seen the per-pass pattern is significantly more pleasing and satisfies the aforementioned constraints of maximizing distances (avoiding clusters and therefore also avoiding voids) and not being regular.

Fig. 5 furthermore shows a comparison of a 4-pass print-mode illustrating all passes cumulatively and demonstrating the pattern goodness both by itself (as can be seen in the 1st pass) as well as together with the remaining passes. The final halftone is identical in both cases.




Figure 5. Example of 4-pass printing shown cumulatively from pass 1 to 4 – with the same resulting halftone in both cases (top part) – showing the optimized approach (left) vs a standard approach (right).

Importantly, this approach translates very directly to printed output properties. Fig. 6 shows a single pass pattern printed (left) with a typical partitioning approach (that does not co-optimize halftoning and partitioning) vs the optimized PARAOMASKING on the right. This figure also clearly shows coalescence which can be seen in the left-hand image vs being practically absent the right-hand side one.




Figure 6. Comparison (scans of printed samples) of regular partitioning that has independent blue noise properties (left) and PARAOMASKING that co-optimizes the partitioning pattern with the halftone to obtain blue noise in each pass (right). The figure shows a single pass of 12.5% coverage of a single NP. Note that the halftone input is the same between the two samples.

These diagnostic results show how partitioning can impact image quality, however it also impacts the printed color and therefore the effective NPac range – i.e. how different choices of NPacs for the same ink-vectors result in different printed output. Since printing has inherent noise in the process, this directly relates to the precision with which dots land on a substrate. For example, different colorant print-heads need to be aligned in a carriage in order for each pen to fire to the corresponding spatial locations. If such alignment were perfect, discounting aerodynamic effects, two pens firing in a single pass would result in drops landing in the same place. Another example is that between passes (or bars) additional

mechanical effects take place, such as the substrate or the printheads/carriage advancing ahead of the next pass. Since this is also not without error, the accumulated effect of drop-placement in a single pass and drop-placement between passes may also result in significant differences in the printing process. Additional sources of variation exist in the printing process, such as drop volume and shape variations.

With PARAOMASKING already showing improved pattern control (and thereby improved grain and robustness), an additional level of control becomes available, that of NP control. While the pipeline is already deliberate about which NPs to use and which to avoid, given the variation, some of this control – while available in the digital domain – gets compressed as a result of the physical processes of print.

However, if the same partitioning is applied to all colorants, then the same patterns are used in each printing pass and as a result – thanks to the determinism of halftoning and the ink-interaction control already handled by NPacs – NPs can be printed in a single pass, thereby avoiding some of the physical effects mentioned before, in particular any pass-to-pass variations and noise. Fig. 7 shows a schematic of this.

With the pipeline resources deliberate choices can be made as to whether to overprint/combine colorants or not, at the extremes of this choice lies the “fully overprinted (blank maximizing)” and “fully side-by-side (blank minimizing)” strategies as shown in the figure, for a constant amount of ink. These patterns, when printed on any printing system, will exhibit some level of color and image quality difference, however, such differences can be either reduced due to the printing process or deliberately maximized with the aforementioned pattern choices (NPacs) coupled with the optimized partitioning approach.




Figure 7. Example halftone inputs (left) maximizing overprinting (top) and minimizing overprinting (bottom) vs simulation of dots on paper (right) for each of the two cases, showing impact of drop placement and drop size variation.

Fig. 8 shows two examples of pattern strategy – fully overprinting (CM) vs fully side-by-side patterns (C and M) at two different densities or amounts of ink (25% C and M in

terms of amount of ink, in the two modalities), printed with optimized masking. These scans already show a clear difference both in terms of grain (with the overprinted, blank maximizing strategy being clearly grainier) and color (with the side-by-side, blank minimizing strategy showing more clearly the individual Cyan and Magenta drops, while the top-row CM overprinting strategy shows more 'blue' dots in comparison).




Figure 8. Scans of printed samples with optimized partitioning, showing difference between fully overprinted (top) vs fully side-by-side (bottom) strategies.

Results and Discussion

The described approach has been tested on a HP Latex-ink based prototype printing system to validate the described properties of a co-optimized partitioning, PARAOMASKING as compared against standard partitioning.

The set-up used 6 color inks (Cyan, Magenta, Yellow, Black and dilutions of cyan and magenta) as well as two non-colored fluids (a pre-treatment fluid that helps color-ink and media interactions and an overcoat fluid that acts as a coating in order to improve handling and durability) and a vinyl substrate, with focus on an 8-pass print-mode – typical for applications that have higher image quality expectations. All prints were done consecutively and have been repeated multiple times in randomized order to ensure results being robust.

Color measurements were performed using a Minolta FD-9 spectrophotometer with a M0 measurement setting and a standard 1400 patch ECI target (used for CMYK color profiling) was used to evaluate color differences, using the CIE ΔE_{2000} color difference metric.

First qualitative results will be reported, showing the impact of the partitioning approach to grain in particular (Fig. 9).

Next, the impact on color is evaluated. To confirm the increased range of NP control as described earlier, the same color chart (the ECI target) was processed with two extreme NP strategies, mimicking the earlier example shown in Figure 8:

- Minimize overprinting (and minimize blank space)
- Maximize overprinting (and maximize blank space)

These resulted in two different halftone images, which in turn were printed with a traditional (uncontrolled) partitioning approach as well as PARAOMASKING. Then, a comparison of color differences between the two alternative halftone strategies and the two partitioning approaches was done and is summarized in the Tab. 1.




Figure 9. Comparison (scans of printed samples) of a color transition (top), a section of a London Tube Map (middle) and a skin tone image (bottom) with standard partitioning top, and optimized partitioning bottom in each case, showing grain improvements with the optimized PARAOMASKING scheme.

Table 1. Color differences between different alternatives of halftone structure (side-by-side or overprinting) and partitioning approaches (standard and optimized).

CIE ΔE_{2000}	SBS vs OP (standard)	SBS vs OP (optimized)	Factor
Min	0.013	0.016	1.2x
Median	0.290	0.485	1.7x
Mean	0.323	0.603	1.9x
95 th %tile	0.669	1.542	2.3x
99 th %tile	0.949	1.957	2.1x
Max	1.745	2.849	1.6x

The optimization approach, which aims to print NPs in the same pass and thereby reduces the contribution of various mechanical variations, significantly increases the color differences between the two halftone approaches as compared to the standard partitioning that does not control the halftoning-partitioning interaction and suffers both from system variability and lack of deliberate NP interactions in the printing process. The results show an increase in range between 1.6 and 2.3x.

As is to be expected, the differences also follow a non-uniform distribution in color space as can be seen in Figure 10, whereby larger differences can be seen both for lighter colors and for less chromatic colors. This is so since in both cases for lighter colors, the halftone and partitioning options are much greater as compared to darker or more chromatic colors. E.g.,

a full saturation Cyan (represented by a high or full coverage of the highest drop state of Cyan), has no difference between a side-by-side or overprinting strategy since there is a single choice. While partitioning does play a role here in terms of avoiding artifacts such as coalescence, the overall color will be less affected in these cases. This explains both the color distribution in Figure 10 as well as the relative factors in Table 1, which are higher at the tail of the distribution.




Figure 10. Comparison of CIE ΔE_{2000} color differences between a side-by-side (white minimizing) and overprinting (white maximizing) strategy using standard partitioning (top) and optimized partitioning (bottom). The plots correspond to results in Table 1.

Conclusions

This paper introduced an approach to co-optimize halftoning and partitioning with the objective of preserving the goodness of patterns designed for halftone content as well as mitigating some of the artifacts caused by the printing process. Key to the approach is the use of deterministic halftoning, such as PARAWACS as well as the NP-level control afforded by the HANS paradigm that enables deliberate choices of ink interactions.

Results clearly confirm the improved print control through PARAOMASKING both in terms of image quality (improved grain) and NP control (improved color range) thus expanding the effective pattern range available to the pipeline and opening up new levels of possible optimization for attributes such as grain, color as well as robustness.

The present approach focused on specific patterns commonly used in high quality printing, namely blue-noise, and explored its performance in real experiments. The approach, however, is not limited to blue-noise optimization and is open to other forms of halftone-partitioning co-optimization in the future.

Acknowledgments

The authors would like to thank Elizabeth Zapata and Andreu Gonzalez for their support of this work.

References

- [1] R. Mai, A. Pekarovicova, P. D. Fleming, "Correlation between Ink Rheology and Press Performance of Water-based Flexographic Inks", *Flexo*, 32. (2007).
- [2] M. Ataeefard, A. Shadman, M. R. Saeb, Y. Mohammadi, "A hybrid mathematical model for controlling particle size, particle size distribution, and color properties of toner particles", *Appl. Phys. A* **122**, 726 (2016). <https://doi.org/10.1007/s00339-016-0242-1>
- [3] A. Lavery, J. Provost. "Color-media interactions in ink jet printing," *IS&T International Conference On Digital Printing Technologies*, 437-442. 1997.
- [4] J. W. Boley, G. T.-C. Chiu, "Print Mask Design for Maximum Throughput Subject to Print Quality Constraints", *IS&T Int'l Conf. on Digital Printing Technologies and Digital Fabrication (NIP26)*, 646-651. (2010). <https://doi.org/10.2352/issn>
- [5] Morovic P., Gomez H., Morovic J., Gasparin P., Ramirez T., Fariña X., Etchebehere S. (2020) Revisiting Print-attribute Optimization: A Direct Pattern Generation Approach, 28th IS&T Color and Imaging Conference, 299-306.
- [6] Morović J., Morović P. and Arnabat J. (2011) HANS – Controlling Inkjet Print Attributes Via Neugebauer Primary Area Coverages, *IEEE Transactions on Image Processing*, 21(2):688-696.
- [7] Morović P., Morović J., Gondek J., Ulichney R. (2017) Direct Pattern Control Halftoning of Neugebauer Primaries, *IEEE Transactions on Image Processing*, 26(9):4404-4413.
- [8] J. Hanson, and J. R. Sanders, "Review of image enhancement techniques for monochrome and color printing", *Proc. SPIE 2171, Color Hard Copy and Graphic Arts III*, 9 May 1994; doi: 10.1117/12.175291
- [9] R. A. Ulichney, "The void-and-cluster method for dither array generation," in *Proceedings SPIE, Human Vision, Visual Processing, Digital Displays IV*, B. E. Rogowitz and J. P. Allebach, Eds., 1993, vol. 1913, pp. 332-343.
- [10] D. Lieberman and J. P. Allebach, On the relation between DBS and void and cluster, *roc. IS&T, International Conf. Digital Printing Technologies*, Toronto, Canada, vol. 14 (1998).

Acknowledgements

The authors would like to thank Elizabeth Zapata and Andreu Gonzalez at HP Inc. for their support.

Author Biography

Peter Morović received his Ph.D. in computer science from the University of East Anglia (UK) in 2002 and holds a B.Sc. in theoretical computer science from Comenius University (Slovakia). He has been a senior color and imaging scientist at HP Inc. since 2007, has published 65+ scientific articles and has 180+ US patents

filed (140 granted) to date. His interests include color science, image processing, color vision, computational photography, computational geometry.

Hector Gomez received his master degree in mathematics from the University of Barcelona in 1998. Since 2007 has been working in HP as a writing systems engineer dealing with image quality problems, servicing and color calibration algorithms. His interests include data science, machine learning and image processing.

Sergio Etchebehere recieved his MSc in color science from University of Eastern Finland and Universidad de Granada in 2016, where then he later worked as a researcher. Since 2017 he is a Color and Imaging Scientist at HP Inc. His research interests are color science, image processing, human vision and computer vision.

Utpal Sarkar received his M.Sc. in Mathematics from the University of Utrecht, and in Physics from the University of Barcelona. He works at the Large Format Division of HP in Barcelona on imaging and printing pipelines and algorithms.

Ján Morovič received his Ph.D. in color science from the University of Derby (UK) in 1998, where he then worked as a lecturer. Since 2003 he has been at Hewlett-Packard in Barcelona as a senior color scientist and later master technologist. He has also served as the director of CIE Division 8 on Image Technology and Wiley and Sons have published his 'Color Gamut Mapping' book. He is the author of over 100 papers and has filed 180+ US patents (142 granted).