

Computational Print Control

Ján Morović^(✉), Peter Morovič, Jordi Arnabat, Xavier Fariña, Hector Gomez, Joan Enric Garcia, and Pere Gasparin

HP Inc., Sant Cugat del Vallés, Spain
jan.morovic@hp.com

Abstract. Printing may seem like a dinosaur among today's imaging technologies, since its roots stretch back to Becquerel's work that led to the first color photographs and the first mechanical color reproduction at the end of the nineteenth century. Ten years ago, we then made the fundamental discovery of a new print control domain, where instead of choices about colorant amounts that are akin to the effect of color filters used since the beginning of color printing, print can be specified by the probabilities of colorant combinations, the Neugebauer Primaries. This has led to the ability to print patterns that were previously inaccessible and consequently, by using large-scale computational optimization, to delivering more color gamut, greater ink use efficiency and greater sharpness and detail in print, while using the same materials and printing system as before. This keynote will present the basic principles of the HANS print control paradigm, review the highlights of results obtained using it to date and indicate its potential future developments.

Keywords: Printing · Imaging · Pipeline · Color · Optimization · HANS

1 Introduction

Printing is among the oldest imaging technologies used today. Its roots lie in the use of stamps to transfer patterns onto clay, as used since 3000 BC, which then led to the use of woodblocks for printing images and text onto textiles as early as 200 AD and then via Guttenberg's invention of moveable type in the 15th century to the digital printing of the late 20th century, via the invention of photomechanical color reproduction in the 19th century. Until these last two developments, printing was wholly in the realm of the crafts and relied on the individual craft person's skills and experience for success.

It was then the invention of automated color printing at the end of the 19th century that enabled an engineering solution for printing and the advent of digital printing technologies like inkjet and laser printing that opened the doors to the introduction of computation into the process. However, for decades, little has changed in terms of how printing was performed, in spite of the new potential that digital processing enabled. Even today, the printing systems used from the home, via the office and into commercial and industrial production apply the same mindset as was first introduced in the late 19th century. The key question that printing pipelines ask is: how much of each of the available colorants to use for matching each printable color. Once colorant quantities are determined, whether they be inks, dyes, toners or waxes, a pattern-building process (i.e.,

halftoning, dithering) is performed largely on a per-colorant basis and the end result is sandwiched together. This, in principle, is a computational equivalent of the photomechanical process, where colorant amounts are determined by filters and pattern formation is the result of glass “screens” where dots are formed on a photosensitive material when light is reflected from an original image and passes through filters and then screens. Even though that process today is the result of computation performed by a digital image processing pipeline, it follows the photomechanical sequence and inherits many of its fundamental constraints.

Here, the foremost limitation is that for each combination of colorant amounts, there is precisely one printed pattern that the pipeline can generate. While this may seem like no constraint at all, in this paper it will be shown that it is, in fact, a dramatic limit on an otherwise much vaster control domain.

To arrive at that insight, it is worth considering a question that, at first sight, might appear unrelated and extraneous: “Is the set of all possible spectra convex?” In other words, if we were to carry a telespectroradiometer around, measuring spectra all around the world, what would their gamut look like. Would it have concavities or not? It might seem like this is an unanswerable question, until Maxwell’s spinning top experiments are remembered [1]. There, Maxwell colored segments attached to a spinning top with different colors and found that the color that resulted from spinning it was a convex combination of the segments’ colors, with their areas as convex weights. Going back to the question about the set of all possible spectra, it can be seen that, if the spectra that occur individually in nature can be applied to segments on a spinning top, then the resulting set can be made convex. Even if the spinning top mechanism has restrictions in terms of being most directly suited for surface colors, it turns out that there is another, more universal mechanism that has the same convex properties. When viewing a scene from a certain distance, variation below a certain angular subtense cannot be resolved by the human visual system, and its content is convexly combined by optical additive mixture. Also from the perspective of measurement, the variation within the angular subtense delimited by an instrument’s aperture is equally optically added.

Positively answering the question about whether the gamut of all spectra is convex then leads to asking why the gamuts of all printing systems are not. A defining feature of print color gamuts are their concavities, typically as the darkest printable color is approached, resulting in a characteristic funnel shape. Why is it that prints, whose colors are the result of colorant combinations being printed in patterns and then viewed such that those patterns are combined optically, end up having gamut concavities? Looking at print color formation, as understood already in the 1930s by Hans Neugebauer [2], it can be seen that it too is fundamentally convex and that it depends on the color of colorant combinations – the Neugebauer Primaries – and on the relative, local areas that they each occupy in print, where these relative local area coverages add up to one (since they are relative) and therefore constitute convex weights (Fig. 1). From first principles, print gamuts ought to be convex, but today they aren’t.

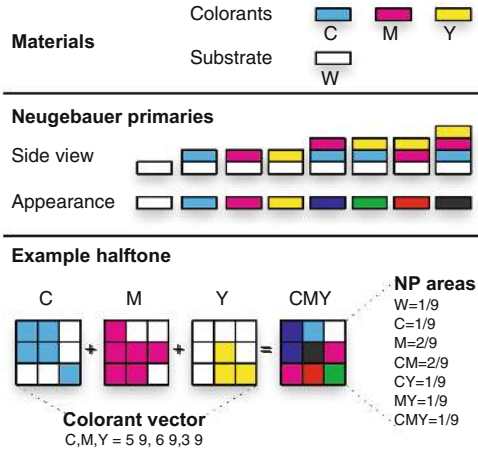


Fig. 1. Colorants and substrate (top), Neugebauer Primaries (center) and an example of how colorant amounts and Neugebauer Primary area coverages relate in a half-tone (bottom).

This is the thought process that led to the discovery of a new print control paradigm – HANS – in November 2007, whose details, properties, benefits and applications will be set out in the remainder of this paper and which has led to a change in print control that also allows for taking full advantage of computational optimization.

2 HANS Basics

The discovery of missing convexity in print has led to a new print control paradigm, where it is no longer colorant amounts but the relative area coverages of Neugebauer Primaries that are an image processing pipeline’s control domain. This new paradigm is called Halftone Area Neugebauer Separation – HANS [3] and its basic principles will be set out next.

2.1 Colorant Versus Neugebauer Primary Space

In a two-colorant printing system, e.g., using cyan (C) and magenta (M) colorants, a colorant-based approach results in some part of the two-dimensional CM colorant space being accessible. The constraints here are the maximum total colorant amount (mT) that avoids colorant–substrate interaction artifacts (e.g., bleeding, coalescence, etc. for ink; peeling, etc. for toner) and the maximum amounts (mC and mM) of the two colorants by themselves that result in a colorant amount – color relationship that is usable (e.g., monotonic) when forming colorant vectors (Fig. 2).




Fig. 2. Print control spaces for a two-colorant (C and M), bi-level printing system: (a) the colorant space and (b) the Neugebauer Primary space.

Each of the colorant vectors from the accessible polygon can be specified at a pixel that is an input to the subsequent halftoning. The combined result of the available colorant space, the halftoning, the two colorants and the substrate is a curved surface in a color space (e.g., CIE XYZ, CIE LAB, etc.), to which input colors need to be gamut mapped [4].

However, these two inks result in four NPs: W (the blank substrate), C, M and the CM overprint. If more than two levels of ink can be specified at each halftone pixel, then the space is of higher dimensionality still. Without constraints from colorant-media interactions, the accessible part of the NP space is a k^n -dimensional simplex – i.e., in the bi-level, CM case it has the 100% area coverages of the W, C, M and CM NPs as its vertices (Fig. 2b). The effect of colorant-substrate imposed constraints is the intersection of this simplex with half-spaces defined by m_C , m_M and m_T . The use of the NP space gives access to any pattern made up of a system's NPs and printed output becomes analogous to mosaics assembled from individual NP tiles.

To illustrate the increased range of patterns accessible via HANS, a particular ink vector: $[c_C, c_M] = [0.5, 0.5]$ can be considered. From the point of view of the NP space, each ink vector corresponds to a continuum of NP area coverage vectors, which satisfy the following constraints:

$$a_C + a_{CM} = c_C \quad (1)$$

$$a_M + a_{CM} = c_M \quad (2)$$

where a_X represents the area coverage of NP X and c_Y represents the area coverage of ink Y.

Figure 3 shows halftone patterns that match the example ink vector but differ in their NPac vectors and colors when printed on a HP Designjet Z3100 printer on Hahnemühle Smooth Fine Art paper. The extremes of the NPac range differ by 26 ΔE_{2000} color difference units [5].




Fig. 3. Continuum of halftone patterns corresponding to a single ink vector: $[cC,cM] = [0.5, 0.5]$. (Color figure online)

For an ink space approach, halftoning results in one of the patterns shown in Fig. 3, while HANS can access the entire continuum. Since even a two colorant system already has 4 dimensions in the HANS domain, this direct NP control gives access to metamers (i.e., halftone patterns that differ from one another but result in the same color under given viewing conditions) because colorimetry is 3 dimensional (leaving one degree of freedom), while ink-based color separation only attains this possibility with four or more inks.

2.2 The Neugebauer Primary Area Coverage Vector

An NP is a discrete combination of colorant units (e.g., ink drops) and an NPac expresses a convex combination of multiple NPs, with weights having a $[0, 1]$ range and expressing relative area coverages. Such k^n -dimensional vector notation is convenient since multiple NPacs can be further convexly combined thanks to the domain’s associativity [6].

In its most basic instantiation, halftone print color formation consists of the convex combination of the colors of the Neugebauer Primaries (NPs). E.g., a pattern C that is formed by combining some of a printing system’s NPs, can be characterized by its NP area coverage (NPac) vector – $NPac_C$, which is of the following form:

$$T(NPac_C) = \sum_{i=1}^{k^n} (w_{Ci} * T(NP_i)) \tag{3}$$

where k is the number of colorant levels per colorant per pixel, n is the number of colorants, $\sum_{i=1}^{k^n} w_{Ci} = 1$ (i.e., the weights are convex), NP_i is the i -th NP, and $T()$ is color (in this case its Yule-Nielsen corrected colorimetry).

The key insight here is that this convex combination of relative area coverage weighted NP colors can also be seen as the convex combination of constituent sub-patterns – for example CA and CB (e.g., a pair of patterns that were individually determined and that can give rise to a continuum of patterns between them in the way laid out below) – such that:

$$w_{Ci} = w_{CBi} + w_{CAi} \tag{4}$$

Furthermore, each of the constituent sub-patterns of C (i.e., CA and CB) also has an equivalent, full pattern (A and B), whose weights relate to the sub pattern’s weights as follows for A (and equivalently for B):

$$w_{Ai} = \frac{w_{CAi}}{\sum_{j=1}^{k^n} w_{CAi}} \tag{5}$$

Having arrived at the weights of A and B (corresponding to the CA and CB constituents of C), the pattern C can now be expressed as a convex combination not of NPs but of the NPacs A and B:

$$T(NPac_C) = x * T(NPac_A) + y * T(NPac_B) \tag{6}$$

where

$$x = \sum_{j=1}^{k^n} w_{CAi}, x = \sum_{j=1}^{k^n} w_{CBi} \text{ and } x + y = 1 \tag{7}$$

As a consequence of associativity, it is possible to perform convex combinations not only of at-pixel states (i.e., the Neugebauer Primaries), but also of a pattern’s sub-patterns (Fig. 4), which is going to be important both for transitioning and optimization.




Fig. 4. Associativity of pattern convex combination.

2.3 NPac Metamer Sets

The concept of metamers – i.e., stimuli that under one set of circumstances result in the same visual response, but under another differ from each other – is key to HANS, since it allows for the grouping of NPacs that result in the same color into metamer sets [7] as follows (Fig. 5):

1. Given a set of p NPs and their colorimetries (measured or predicted), form all polyhedra with up to n vertices ($n \in [4, p]$) in the colorimetric space.

- For a given color S , whose metamer set is to be found, compute its barycentric coordinates with respect to each of the polyhedra from step 1. In the case of tetrahedra with V_i ($i \in [1, 4]$) vertices it is done as follows:

$$\begin{pmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \end{pmatrix}^T = \begin{pmatrix} S_X \\ S_Y \\ S_Z \\ 1 \end{pmatrix}^T \times \left(\begin{pmatrix} V_{1X} & V_{2X} & V_{3X} & V_{4X} \\ V_{1Y} & V_{2Y} & V_{3Y} & V_{4Y} \\ V_{1Z} & V_{2Z} & V_{3Z} & V_{4Z} \\ 1 & 1 & 1 & 1 \end{pmatrix}^T \right)^{-1} \quad (8)$$

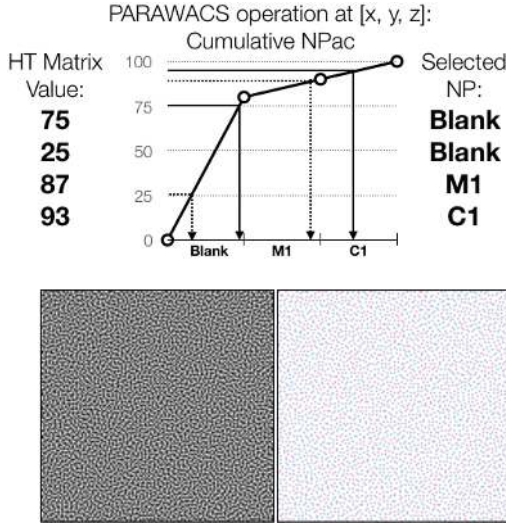


Fig. 5. (Top) PARAWACS halftoning by comparing selector matrix values with cumulative NPac intervals and (bottom) example of using a blue-noise matrix (left) with the same NPac to give a halftone (right) that honors NPac area coverages and has the spatial structure of the chosen matrix.

Equation 8 is the inverse of the system formed by three equations expressing how S is the convex combination of V_1 to V_4 and the equation of b_1 to b_4 summing to one. If all members of b are from $[0,1]$, S is inside the tetrahedron. Furthermore, since b represents normalized volumes of sub-tetrahedra formed by sets of three original vertices and S (for the original vertex that S substitutes), they are also the relative area coverages with which tetrahedron vertex NPs need to be combined. For practical purposes, up to 16 vertex polyhedra are typically considered, but this number could be increased to form more complex patterns. The result of step 2 is the set of all metamers formed by different sets of four NPs.

- The set of NPacs obtained from step 2 delimits the metamer set of color S in NPac space. Note that due to the convexity of the NPac space, the metamer set can be further sampled by generating convex combinations of NPacs obtained from step 2, thanks to the associativity of the NPac domain.

2.4 Optimization of Print Attributes

Given metamer sets for a set of colors sampling the available convex hull, color separation can be optimized by deciding which of the available metamers to choose for each color. Here any print attribute that can be determined or predicted from NPacs, can be explicitly taken into account and multiple attributes can be directly weighted depending on the needs of particular applications.

A simple attribute to consider is how much ink is used. Knowing the amount of ink involved in the NPs and by them using the metamer NPacs' area coverages, the minimum within each metamer set can be found. Using the NPac involving least ink for all available colors then leads to printing with the lowest possible material costs.

2.5 Summary of HANS Color Separation

Bringing together the above leads to the following method of computing a color separation to NPac space, assuming no constraints from the ink–substrate interaction and using the original Neugebauer model:

1. Print and measure all NPs.
2. Compute NP convex hull in Yule-Nielsen corrected CIE XYZ and sample it.
3. For each sample from step 2 compute the corresponding metamer set.
4. From each metamer set select the NPac that is optimal in terms of the chosen (combination of) print attributes.

The result of the above process is a set of CIE XYZ coordinates that span the full color gamut and where each has an NPac assigned to it. To compute the NPac color separation for an arbitrary, in-gamut CIE XYZ input, the CIE XYZs from step 2 need to be tessellated (e.g., using Delaunay tessellation [8, 9]), the tessellating tetrahedron enclosing the input color needs to be found and barycentric coordinates for the vertex NPacs can be computed. These barycentric coordinates then form an NPac with the vertex NPacs as its members, which can finally be re-expressed as an NPac with NP members, thanks to associativity. Applying this algorithm to an entire, in-gamut CIE XYZ image will result in a corresponding NPac image, which is then the input to the next, halftoning stage. Note also that this transformation can be performed by first computing a look-up table indexed in XYZ, which for each node contains NPacs.

2.6 From NPac to Halftone

The HANS halftoning maps NPacs to at-half-tone-pixel NPs (e.g., drops per ink) and two categories of halftoning approaches apply: device-state error diffusion [3] or PARA-WACS matrix-based halftoning [10]. In both cases, the role of halftoning is that for a sufficiently large area of a constant NPac, it needs to result in a placement of individual NPs such that when counting their frequency over the area, the original area coverage is obtained and that the resulting pattern meets expectations in terms of pleasingness, robustness, detail preservation, etc. At half-tone-pixel level, halftoning becomes the selection of a single NP from the NPac specified for that pixel.

To simplify this selection in the matrix halftoning case, called PARAWACS (Parallel Random Weighted Area Coverage Selection), an NPac can be expressed cumulatively. I.e., the NPac: [Blank (white) 80%, 1 drop M 10%, 1 drop C 10%] becomes [W = 80%, M = 90%, C = 100%] in cumulative terms, which in turn defines intervals for each of the NPs such that [0 to 80] corresponds to the Blank state, [80 to 90] to one drop of M and [90 to 100] to one drop of C. Using this representation, what is needed next is a uniform distribution of values between 0% and 100%, for example in the form of a pre-computed matrix of a desired pattern, where the matrix values then act as NP selectors. Here, matrix values are categorized according to cumulative intervals so that if the value at [x, y] is in the range [0 to 80] the halftone pixel is left blank, if it is in the range [80 to 90] a drop of the M NP is placed and if the random number is in the interval [90 to 100], a C NP is used. The diagram in Fig. 5 shows this process for four random values (corresponding to four [x, y] locations) and using a blue-noise matrix [11].

3 Applications and Results

While the above has focused on setting out the basic principles and mechanisms of the HANS computational print control paradigm, this section will provide an overview of some of the key applications and results obtained to date.

3.1 CMY Metamers

Since the HANS domain is of significantly higher dimension than the colorant domain, it changes even the basic constraints of print color reproduction. A significant result here is the ability to form metamers even in three-colorant systems, which previously were considered to lack the ability of forming metamers. When controlling such printing systems in colorant terms, it is the case that for each within-gamut color there is precisely one colorant combination that results in it.

Instead, with HANS, even in a binary, three-colorant printing system there are $2^3 = 8$ NPs, which in turn can be combined in multiple ways to match in-gamut colors. As a result, even with three colorants there is redundancy and it is possible to choose from a variety of alternatives [12]. Applying HANS to a printer with only CMY inks still resulted in 115 distinct metamers for a mid-gray color (Fig. 6) and gave access to a range of ink use where the maximum was 13% greater than the minimum for a system where colorant domain control yields no metamers. The optimization that resulted in metamers for a sampling of the CMY system's color gamut used 244 "base" NPacs with known colorimetry that involved the system's NPs and a sampling of the gamut surface. Over this set, a search was performed, resulting in the evaluation of 14,4 million tetrahedra.




Fig. 6. A visualization of 115 metamers: each column represents an NP area coverage vector where constituent color block lengths heights are proportional to the corresponding base NPac color's relative area coverage.

3.2 Gamut and Ink Use Optimization

Among the first results of using a HANS printing pipeline was its application to a *HP Designjet L65500* printer using CMYKcm latex inks on a self-adhesive vinyl substrate [3]. Computing a mapping to NPacs, following the process outlined in Sect. 2 of this paper, resulted in 20% less ink use than the minimum achievable in a colorant domain pipeline (where that was the result of using maximum gray component replacement) with less grain than previous minimum. At the same time as delivering a reduction in ink use, the HANS pipeline also gave access to 10% more color gamut (Fig. 7).




Fig. 7. (Left) Ink use of colorant-domain and HANS pipelines, (center) corresponding color gamuts and (right) scanned prints.

3.3 Hand-Building of Color Separations

Instead of being confined only to complex, computational solutions, the HANS NPac domain also lends itself to dramatic simplification and the hand-building of printing system resources. An extreme example here is the possibility of defining a color separation by hand-selecting only eight NPacs – one for each vertex of a device RGB cube and then having NPacs computed for the entire color space by interpolation in an

analytical tessellation of the eight vertices, using tetrahedra that all share the black to white long diagonal of the RGB cube. The result is a printing pipeline that, for each printed color combines up to four of a set of eight NPacs [6]. This results both in a high degree of simplicity when generating printing resources and in a system that can be characterized and calibrated on the basis of a very small number of measurements (Fig. 8).

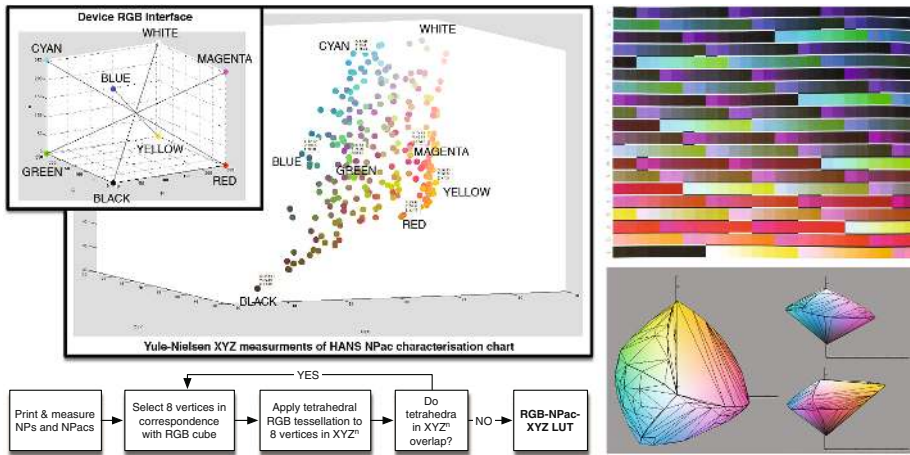


Fig. 8. Hand-building node selection putting RGB vertices and NPacs in direct correspondence. (Color figure online)



Fig. 9. 3D object printed using HANS3D pipeline to control voxel-by-voxel volumetric material composition (Model: Klein Bottle by krasul, published on July 22, 2013, www.thingiverse.com/thing:121871 under Creative Commons – Attribution license).

3.4 HANS3D

Finally, the HANS paradigm of controlling the use of at-pixel states by assigning relative area coverages to them can also be extended to the control of 3D printing systems. There, it is the set of all possible voxel contents that are the atomic building blocks of a print control pipeline and their use is governed by assigning relative volumetric coverages to

them. Since in 3D printing, color is only one of many properties that can be controlled, the equivalent there of NPs are material vectors (Mvecs) and instead of NPacs, it is material vector volume coverages (Mvocs) that determine local printed object composition. The resulting HANS3D pipeline [13] can take advantage of the breadth of HANS benefits and can further deliver control over volumetric material use (e.g., tuning layer-to-layer complementarity) and allow for a co-optimization of multiple object properties via the determination and intersection of per-property metamer sets (Fig. 9).

4 Summary

Controlling printing systems by determining the use of their atomic building blocks opens up a vast new space of printable patterns, whose properties go beyond what conventional, colorant-space pipelines give access to. The underlying convexity of colorant and material combinations is uniquely suited for computational optimization and has already yielded access to metamers where previously there were none, to improvements in ink use efficiency and color gamut and to a native, volumetric, voxel-by-voxel control of 3D printing systems. While the overview presented here draws on 10 years of theoretical and experimental work, the HANS print control paradigm is still only in its infancy and promises fertile ground for future exploration and evolution.

Acknowledgements. The authors would like to acknowledge the support of their colleagues at HP Inc.: Africa Real, Albert Serra, Andrew Fitzhugh, Annarosa Multari, David Gaston, Gary Dispoto, Jay Gondek, Juan Manuel Garcia Reyero, Lihua Zhao, Rafa Gimenez, Ramon Pastor, Sascha de Peña, Tsuyoshi Yamashita, Virginia Palacios and Yan Zhao.

References

1. Maxwell, J.C.: Experiments on colour, as perceived by the eye, with remarks on colour blindness. *Trans. Roy. Soc. Edinb.* **21**(2), 275–298 (1855)
2. Neugebauer, H.E.J.: Die theoretischen Grundlagen des Mehrfarbenbuchdrucks. *Zeitschrift für wissenschaftliche Photographie, Germany* **36**(4), 73–89 (1937)
3. Morovič, J., Morovič, P., Arnabat, J.: HANS – controlling inkjet print attributes via neugebauer primary area coverages. *IEEE Trans. Image Process.* **21**(2), 688–696 (2011)
4. Stollnitz E.J.: *Reproducing Color Images with Custom Inks*, Ph.D. thesis, University of Washington (1998)
5. CIE: CIE 142:2001: Improvement to industrial colour difference evaluation, CIE Central Bureau, Vienna (2001)
6. Morovič, J., Morovič, P., Rius, M., García-Reyero, J.M.: 8 vertex HANS: An ultra-simple printer color architecture, In: *IS&T/SID 21st Color and Imaging Conference*, Albuquerque, NM, pp. 210–214 (2013)
7. Finlayson, G.D., Morovic, P.: Metamer sets. *J. Opt. Soc. Am. A:* **22**(5), 810–819 (2005)
8. Delaunay, B.: Sur la sphère vide. *Izvestia Akademii Nauk SSSR, Otdelenie Matematicheskikh i Estestvennykh Nauk* **7**, 793–800 (1934)
9. Preparata, F.P., Shamos, M.I.: *Computational Geometry: An Introduction*. Springer, New York (1985)

10. Morovič, P., Morovič, J., Gondek, J., Ulichney, R.: Direct pattern control halftoning of neugebauer primaries. *IEEE Trans. Image Process.* (2017) (submitted for publication)
11. Ulichney, R.: Dithering with Blue Noise. *Proc. IEEE* **76**(1), 56–79 (1988)
12. Morovič, P., Morovič, J., García-Reyero, J.M.: HANS: Meet The CMY Metamers, In: *IS&T/SID Color and Imaging Conference 2011*, San Jose, CA, 7–11 November 2011, pp. 229–233 (2011)
13. Morovič, P., Morovič, J., Tastl, I., Gottwals, M.: HANS3D: a volumetric voxel-control print content processing pipeline, In: *ACM SIGGRAPH 2017*, July 2017, Anaheim, CA (2017) (submitted for publication)