Computer Color Reproduction

Ricardo J. Motta
Hewlett-Packard Laboratories
Palo Alto, California USA

Abstract

Computer based color imaging is rapidly moving from the realm of the specialist to the general public. To enable users the greatest benefit from advances in color image reproduction technology, the computer systems have to modified to incorporate this new technology. Two possible approaches are described—by the adoption of color management system software, and by the adoption of device independent color variables. Although both systems are likely to perform equally well when properly configured, the later is likely to be the long term choice as it will yield the most robust, lower cost and transparent system.

Keywords: Color reproduction, color management, device independent color, color matching, color imaging, computer systems, computer imaging, scanning, printing.

Introduction

Most problems in color reproduction are rather ancient. Painting has made us familiar with the difficulty of properly reproducing the shapes and colors of things and we readily recognize the great skill needed to convey verisimilitude. Technology, such as printing and photography, has delivered on its promises to automate and simplify the production of color images. The aesthetic value of images has remained beyond reach of technology, but there is now an infrastructure in place that allows anyone to capture and produce printed images at will. The next step, computational color reproduction, promises even higher levels of quality and accessibility. To become pervasive, most of the color science and technology learned in the last 100 years will have to find a home in different parts and components of computer systems. What is the ideal system design, and the architecture that will yield the greatest benefits to the general public is still open for debate.

The clear direction in color image reproduction is towards democratization of the resources. The equipment needed to capture and produce a printed image can now be owned for a few thousands of dollars. Within the past five years, millions of computer users have purchased low cost color printers capable of producing fairly satisfactory graphics. Within the next ten years, millions more will be buying digital cameras and printers capable of producing photographs and printed pages rivaling in quality that of current commercial equipment. This same process has been repeated a number of times within other fields of technology. A well known example is the emergence of desktop publishing. Typography and page layout capability have become widely pervasive, to the point of not being noticeable. In fact, desktop publishing offers us a perfect case to reflect upon when we consider the future of computer based color reproduction.

Desktop publishing was enabled by two key pieces of technology, the digitization of type and bit mapped graphics. What has become remarkable is not the technology itself, but how it is now integral to the design and use of computers. Bit mapped fonts are pervasive to all applications on computers with graphical user interfaces, and document creation, not publishing, has become the focus of the page layout technology. For the last ten years, computer architectures, hardware and software have been modified to take into account the demands of beat mapped graphics and font technology. The current PC has features that 10 years ago only the publishing industry would have thought essential, but paradoxically only a infinitesimal number of these PCs today are in any way connected to the business of the publishing industry.

This evolution shows us how the introduction of new capabilities into computers has little to do with technology, and a great deal with the business of product development. Computer designers often borrow heavily from the best practices of certain industries when introducing new technologies to computers. While at introduction it often looks like these industries control the design and purpose of these systems, such marriages of convenience are quickly dissolved. Sound, imaging, color, fonts, and graphics are few of the technologies that have undergone, or are going through, this migration.

This process has become natural to the evolution of computer systems. To reach the mass market, any new technology has to be made transparent to the mass consumer. This last is interested in acquiring new capabilities (e.g. creating documents with color pictures) and not on learning a new skill. To this end, the technology being introduced has to be recast in ways that will make it automatic, fast, low cost and robust. Of course, this is hardly possible in a single generation of hardware or software, and the process that ensues is usually tortuous. Early users and adopters of the technology tend to have interests and concerns much different than the mass market. As adoption of the technology progresses, better market focus is obtained and engineering compromises are made that often alienate these early adopters.

The application of color imaging and color reproduction to computer systems is still in its infancy. Satisfactory results, most often than not, are only achieved through experience and experimentation. As computational color reproduction technology runs its course towards pervasiveness, designers are still struggling to identify how to make the technology automatic, fast, low cost and robust. To

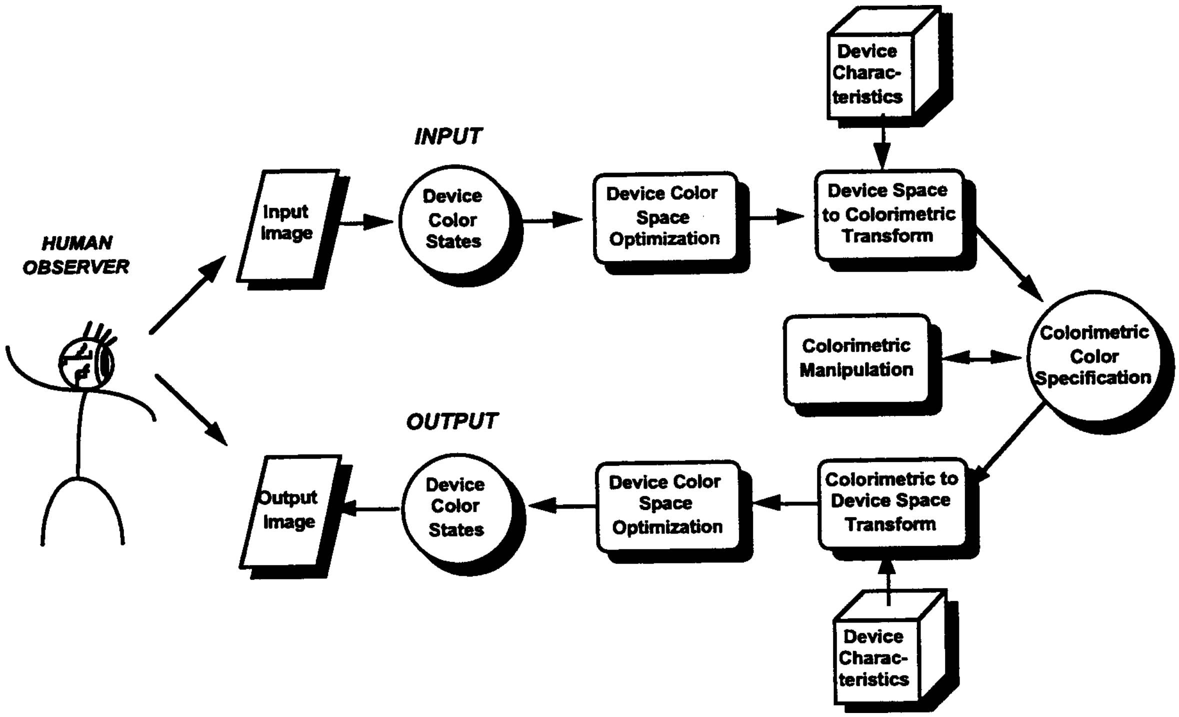

Figure 1. The device independent imaging pipeline.

achieve it, it is important to recognize what are the most relevant and stable aspects of computational color reproduction technology.

The Color Reproduction Technology

Over the past several years important incremental steps have led to a strong theoretical base for implementing color reproduction on computers. A great deal has been presented in other technical papers, including those in this volume. By far, the key development was the emergence of what is commonly called the device independent (DI) imaging pipeline as depicted in Figure 1.

The fundamental characteristic of the DI pipeline is that it assumes that the color appearance can be abstracted and represented independent of the characteristics of the original imaging device and viewing conditions. The tools that should bring this about are provided by vision and color science, but to the extent that much is still unknown about the human visual system, simplifying assumptions are often necessary, such as assuming that the original and reproduction occupy the same angular subtense, or that the viewing illuminant is the same, With this and other simplifying assumptions, it is often assumed that a colorimetric (in the broad sense that colorimetry includes some appearance transforms) match will result in an appearance match. This has resulted in many useful systems, with performance often limited only by the performance of the devices and the appropriateness of the assumptions.

Some of the key advances that enabled the implementation of the DI pipeline were:

1. CRT models and characterization methods
2. Printer models and characterization methods
3. Scanner models and characterization methods
4. Color map (device profile) generation methods
5. 3D interpolation—trilinear, tetrahedral, prism
6. Gamut compression, tone correction
7. White point transforms, correction methods
8. High speed colorimetric measurements

Despite this very promising list, some key areas remain in need of advancement. For instance: a) there is little understanding on what visual factors govern color constancy and what factors to incorporate into white point transforms; b) the effect of flare and dynamic range on imaging is little understood; c) fluorescence is not well understood and on its account color measurements frequently fail to capture the visual stimulus; d) effective gamut compression algorithms remain mostly proprietary and implementation dependent; e) metamerism in the input devices is understood, but we still lack effective ways of correcting for it.

The Color Reproduction Systems

The DI color pipeline provides a logical framework for producing color images on computers across different imaging peripherals and physical locations. It provides an order for the operations and provides a discipline as to where in the sequence should image processing and enhancements take place. It does not in itself suggest an implementation, that is—how and where the operations should be executed. This is essential to this discussion, since as we will see, it is by determining where the processing is best carried out that we will achieve our goal of a system that is automatic, fast, low cost and robust.

The implementation process is largely an engineering one, capable of taking many forms from the same technology base. Like many previous technologies and innovations, color reproduction has found its way into computer systems by the route that travels from applications to operating system. Since operating systems have not yet been built that take color reproduction into account, it is worthwhile to reflect on where the technology is and where should it head.

New ideas and technologies are usually introduced into computing through application programs that cater to specific audiences that can control and understand its features. For instance, color correction for producing CMYK images for printing was first available to the public in applications such as Aldus' PageMaker and Adobe's Photoshop. Another way new technology finds it way into the computing environment is through system software that certain companies will provide to enhance the operation of their peripheral devices. A good example is the color matching printer driver, which will assume that the application has no knowledge of how to prepare the images for printing, make assumptions about what the user is seeing in the CRT and modify the printed colors to match the screen. Such drivers have been available from HP and Tektronix since 1989.

Clearly the two approaches are in conflict—both the application and the device drivers might be unaware that the colors have been modified and color correction might happen two times. This has resulted in a great deal of dissatisfaction among less experienced users. For the system to operate properly one has to understand all the options of the application and device drivers, understand what is meant by the choices and make the selections that will ensure the color correction happens only once. This is by far the most common mode of failure in personal computer color printing, and the solution would, at first, appear to be a simple one—to bring color reproduction under control of the operating system, providing the discipline necessary to resolve where the color processing is executed. But deciding what is the architecture for the system, where should the processing take place, has turned out to be rather difficult. In essence, the system designer trying to address the problems of color reproduction has two choices:

1. CMS Approach—Implement the entire DI pipeline in only one place, under OS control, utilizing a Color Management System.
2. DI Approach—Distribute the color pipeline among the components, using a device independent color inter change method defined by the OS.

We will discuss these two methods.

The CMS Approach

This first approach is currently being implemented by the major OS vendors. It has the advantage of requiring minimum development effort and system redesign. It consists basically of retrofitting an existing system by building a filter that will replace the device dependent data in color commands by colors that will properly reproduce when displayed or printed. To insure that the correct operations take place a control layer is added to the system that keeps track of the correct methods for each device. This control layer includes methods of associating incoming images (e.g., acquired with scanners) with device profiles (e.g., the scanner characteristics) and methods for users and application developers to choose among color matching methods and other implementation details.

Color reproduction offers the system designers unique challenges because, unlike fonts or graphics, so much of the technology has remained proprietary and there is no objective measure for improvement. In many respects, the CMS approach is a logical step for operating system designers to pursue. Computational color reproduction technology is in its infant stages, and much experimentation and discovery will be necessary before the knowledge is stable and widely available. Thus, lacking a clear cut solution to most problems, the tendency has been to design systems with enough flexibility to accommodate all different opinions and defer the choice among competing methods (e.g. gamut compression schemes) to the user at run-time. Unfortunately, this practice burdens the users with choices that offer no objective value, forced upon him or her by the lack of best solutions at design time. Although this practice tends to be acceptable for the trained or experienced user, it is a barrier to the neophyte and most other users, who, again, want to perform a task, not learn a new skill. For example, many features of current CMSs are only relevant to the minority of experienced or professional users. These include on-monitor print proofing, out of gamut alarms, explicit choice of rendering intent, alternative color matching methods, explicit profile selection and third party color profiles:

- Monitor proofing—To most users transforming a vibrant CRT color imaging into a washed out version that attempts to emulate the printer is confusing since the CRTs do not have the dynamic range needed to really look like paper under common viewing conditions.
- Out of gamut alarms—The concept of what is color gamut has to be learned. Most users would be happy if the best available color is printed, or if they were not offered or shown an out of gamut color to begin with.
- Rendering intent—The idea that the color rendering depends on the image content is also foreign to most users.
- Profile selection—This is barrier to the user since it requires one to synchronize selections made on many different dialogue boxes. For instance, one may have to specify the same choices more than one time, like selecting a printing halftoning and later selecting the profile for that halftoning.

Of course, the CMS approach should deliver good color reproduction. As we said before, it is not the technology that will determine the usefulness of the system, but where and how the technology is implemented. Some important problems left unsolved by CMSs are:

- Maintenance—Because there will be so many components to the system, between CMM and profiles, it is very hard to determine if the system is properly con figured, what is the proper configuration, and to how perform diagnostics when things go wrong.
- Cost—Much of the color matching technology needed

to make these systems work is proprietary and has been identified as a potentially lucrative business by a few companies. Consumers may indirectly pay for the technology multiple times as they buy the OS, the applications, the device drivers and the third party enhancements, all of which could offer improved versions of the basic profiles and color matching technology.
- Performance—By remaining proprietary and diverse, the color matching modules will stay hard to accelerate in hardware or software.
- Consistency—Because performance will be an issue, different applications might elect to go different routes with images, some of them avoiding color matching because of performance, resulting in a inconsistent image presentation.
- Portability—Because the input device profile has to be attached to the file, it becomes harder to communicate across systems unless they all use the same color matching technology, which is unlikely. This is also a barrier for permanence, i.e. systems separated in time, because the maintenance of the profile standards and the image standards are with different organizations.

The DI Interchange Approach

The DI Interchange approach promises to be more robust by having the operating system describe colors and images using well defined device independent variables.* It frees the system from tracking and maintaining any control over how colors are reproduced since the burden of transforming and producing the colors accurately is transferred to the color imaging device (e.g., displays, scanners and printers) manufacturers. It requires even less modification to the systems than the CMS approach, but because it requires the collaboration of the color imaging device manufacturers, no OS vendor has embraced it yet.

Interesting enough, a similar method, dubbed "CRT is king", is the current de-facto approach. Since the current operating systems lack a precise color specification, peripheral device vendors and application developers assume that the color interchange space for the system is that of the local CRT. When properly configured by the user, this system works fine since input devices can provide colors accurate for the display, printers can reproduced screen colors accurately, and by association all components (input, display, printer) of a local system reproduce the same colors. On systems such as the Apple Macintosh that have a well defined CRT transfer function (e.g. gamma 1.8), users tend to be much happier with current color imaging equipment. An important added benefit from a consistently defined and implemented CRT transfer function is that users of the same brand computer can freely exchange images.

The CRT transfer function ranks the highest among all the variables that affect the display image appearance. The two other often cited variables, primaries and white point, have either little impact (in the case of the primaries) or are little understood (in the case of the white point) and no standard, effective methods exist to fully account for it. On Microsoft Windows and DOS based systems, as well as in Unix based systems, there is little definition of what should display manufacturers build as the transfer function on their CRTs and display cards. This problem has been alleviated with the introduction of CRT characterization devices and visual characterization software. Still, either by virtue of being expensive or requiring extra user intervention, these CRT characterization methods are not a reliable solution to the problem. Given that it is relatively simple and inexpensive to specify and manufacture displays to a prescribed transfer function, it is remarkable that is has not happened yet for all platforms.

This simple measure would address the needs of the vast majority of computer users by providing consistent image display across platforms, geographical location and time. It would also provide the input and output imaging device manufacturers with a stable target for optimizing the color reproduction of their devices hardware and software. Those manufacturers without the know-how to develop the required color reproduction transforms for their printers and scanners could utilize the work that has already been performed in their behalf by the CMS vendors.

Adopting a standard definition of the display transfer function would be an easy first step towards a truly DI system. Other possible enhancements that could gradually move all systems towards the goals of greater interoperability and ease of use, would be the adoption of standard chromaticities and white points. Many image encoding standards and proposals already cover this subject, and existing standards such as CCIR 709 from the television industry address most of the needs of an unambiguous color description that can be used for fast color displays. A complete DI color standard would probably borrow heavily from the work of the color and vision science communities towards establishing white point transforms and measurement methods that would further clarify the meaning of the chosen color space.

Conclusion

The largest barrier to the adoption of the DI approach will be alignment between the OS vendors and peripheral device manufacturers. The emergence of DI color image encoding standards should facilitate this as it will provide a clear target for the optimization of system components. It is in the best interest of the OS vendors to optimize performance and reliability, and not to assume responsibility for all the processing burden, maintenance and color reproduction failures. These should be assumed by the peripheral vendors as they are competitive advantages that can be engineered into the devices hardware and software. Consumers will benefit by the increased performance and ease of use. Most of the technology is available and is the same for the DI approach and the CMS approach. It is now up to the system designers to decide which approach that will translate into the greatest benefit to their customers.

* It should be noted that the device independency follows from the color variables being well defined, and not necessarily from the use of new variables, such as CIE based color spaces.

✰ This paper was previously published in *SPIE,* Vol. **2414**, p. 2 (1996).