SIGGRAPH 2001 Course 35

THE TECHNOLOGY OF DIGITAL CINEMA (D-CINEMA)

The SIGGRAPH community has been at the forefront of digital creation and processing of motion picture images for well over a decade. However, live-action principal photography has been acquired on 35 mm motion picture film, and the end result has been displayed on 35 mm motion picture film. HDTV equipment using the 1920×1080 image format now achieves spatial resolution comparable to film; efforts are nearly complete to adapt HDTV cameras and recorders to the 24 frames per second of film. In addition, electronic projection has now reached the brightness and spatial resolution of the cinema. This course describes the technology of D-cinema acquisition and display, and discusses the interface between D-cinema and digital image creation and editing.

Charles Poynton
(Organizer)

www.inforamp.net/~poynton
poynton@poynton.com
tel +1 416 413 1377
fax +1 416 413 1378

Fred Meyers
Lucasfilm Ltd.

Chuck Harrison
Far Field Assoc.

Bill Werner
Texas Instruments

Arjun Ramamurthy
Warner Bros.

Isaac Kerlow
Walt Disney Company

Copyright © 2001-05-11 Charles Poynton

SIGGRAPH 2001 Course 35

THE TECHNOLOGY OF DIGITAL CINEMA (D-CINEMA)

TABLE OF CONTENTS

1 Lightness terminology 1
(*Poynton*)

2 DLP Cinema™ technology 3
(*Werner*)

3 Application of DLP technology to digital electronic cinema – a progress report 17
(*Werner*)

4 Color representation for D-cinema 25
(*Harrison*)

5 A review of photographic color science, and its application to digital cinema exhibition 33
(*Harrison*)

6 Device-independent color in SMPTE standardization efforts 49
(*Harrison*)

7 Color management for D-cinema 55
(*Ramamurthy*)

8 Mastering movies for a digital release 75
(*Kerlow*)

APPENDIX

A Color image coding for digital projection and D-cinema 77

Charles Poynton works to integrate video technology, particularly high definition television and accurate color reproduction, into computer workstations. While at Sun Microsystems, from 1988 to 1995, he initiated Sun's HDTV research project, and introduced color management technology to Sun. Prior to joining Sun, Mr. Poynton designed and built the digital video equipment used at NASA's Johnson Space Center to convert video from the Space Shuttle into NTSC for recording and distribution. He is a Fellow of the Society of Motion Picture and Television Engineers (SMPTE), and an Honorary Member of the BKSTS. In 1994 he was awarded SMPTE's David Sarnoff Gold Medal for his work integrating video technology with computing and communications. He has organized and presented many popular courses and seminars, including *HDTV Technology* at SIGGRAPH 91, *Concepts of Color, Video and Compression* at ACM Multimedia 93, and courses on color technology at SIGGRAPHs in 1994, 1996, 1997, 1998, and 1999. His book, *A Technical Introduction to Digital Video*, was published in 1996 by John Wiley & Sons, and is now in its fourth printing.

Fred Meyers is Principal Engineer at ILM in San Rafael, California. He has just returned from Sydney, Australia, where digital acquisition of principal photography for Star Wars, Episode 2 was just completed. Fred was in charge of technical aspects of the 1080*p*24 digital acquisition system.

Bill Werner is systems engineering and electronics design team leader for DLP Cinema development at Texas Instruments. Werner joined TI in 1982, after receiving a BSEE degree from the University of Michigan. In 1987 he received a MSEE degree from Southern Methodist University; he joined the Digital Imaging Group of TI in 1993. Werner led development of the first DLP projector based on the 1280×1024 Digital Micromirror Device (DMD). He has had various systems engineering assignments within the large venue products group at TI. Werner is a member of SMPTE, BKSTS, and IEEE.

Chuck Harrison is a partner in Far Field Associates in Seattle, where he is working on standards for color representation, and security and encryption, for D-cinema. He has 30 years of experience in electronic and optical engineering. He has designed equipment for motion picture original photography, post-production, and projection.

Arjun Ramamurthy has developed state-of-the-art software and systems solutions for computer imaging for the past 15 years. He received his B.E. in Electronics and Instrumentation from Annamalai University, India in 1987, and an M.S. degree in Electrical and Computer Engineering from Clarkson University, Potsdam, N.Y. in 1989. In 1989, he joined EOIS Corp. developing machine vision systems for paint inspection, and 3-D gauging Systems. In 1996, he joined PhotoResearch Inc., where he developed color imaging systems and hand-held instruments for photometry and colorimetry. Ramamurthy is currently the Manager of Software Engineering at Warner Bros. Feature Animation, with interests in graphics, computer-assisted animation, imaging and color management systems. He is a co-author of ANSI Standard IT7.328-1995, *Image Evaluation for Projection*

Systems. He has published several papers on machine vision and image processing.

Isaac Victor Kerlow is Director of Digital Production and Talent in New Technology and New Media Group at The Walt Disney Company. He plays a key role in efforts across the company related to digital art production and digital artists working in a variety of wide-ranging new media. Prior to this post Isaac led the group of digital artists and animators at Disney Interactive responsible for developing and producing CD-ROM, on-line and platform games. Before joining Disney in 1995 Isaac spent a decade at Pratt Institute where he was the founding chairman of the Department of Computer Graphics and Interactive Media and a tenured professor. Isaac has shown his digital artwork internationally and is also the author of several books, including the second edition of "The Art of 3D Computer Animation and Imaging" published in 2000 by John Wiley and Sons.

This is an unpublished excerpt from the forthcoming book *Digital Video and HDTV: Pixels, Pictures, and Perception,* in press, to be published in August, 2001. This material is Copyright © 2001-05-10 Charles Poynton.

In a grayscale image, each pixel value represents what is loosely called *brightness*. However, brightness is defined formally as *the attribute of a visual sensation according to which an area appears to emit more or less light*. This definition is obviously subjective, so *brightness* is an inappropriate metric for image data.

Intensity refers to radiant power in a particular direction; *radiance* is intensity per unit projected area. These terms disregard wavelength composition. However, if color is involved, wavelength matters! Neither of these quantities is a suitable metric for color image data.

The term *luminance* is often used carelessly and incorrectly to refer to *luma;* see below.

In image reproduction, we are usually concerned not with (absolute) luminance, but with *relative luminance*.

Luminance is radiance weighted by the spectral sensitivity associated with the brightness sensation of vision. Luminance is proportional to intensity. Imaging systems rarely use pixel values proportional to luminance; usually, we use values nonlinearly related to luminance.

Lightness – formally, CIE L^* – is the standard approximation to the perceptual response to luminance. It is computed by subjecting luminance to a nonlinear transfer function that mimics vision. A few grayscale imaging systems code pixel values in proportion to L^*.

Regrettably, many practitioners of computer graphics, and of digital image processing, have a cavalier attitude toward these terms. In the *HSB, HSI, HSL,* and *HSV* systems, *B* allegedly stands for brightness, *I* for intensity, *L* for lightness, and *V* for value. None of these systems computes brightness, intensity, luminance, or value according to any definition that is recognized in color science!

Value refers to measures of lightness apart from CIE L^*. Imaging systems rarely, if ever, use *Value* in any sense consistent with accurate color.

Color images are sensed and reproduced based upon *tristimulus values*, whose amplitude is proportional to intensity, but whose spectral composition is carefully chosen according to the principles of color science. As their name implies, tristimulus values come in sets of 3.

Accurate color imaging starts with values, proportional to radiance, that approximate *RGB* tristimulus values. (I call these values *linear-light*.) However, in most imaging systems, *RGB* tristimulus values are subjected to a nonlinear transfer function – *gamma correction* – that mimics the perceptual response. Most imaging systems use *RGB* values that are *not* proportional to intensity. The notation *R′G′B′* denotes the nonlinearity.

See *YUV and luminance considered harmful*, www.inforamp.net/~poynton/

Luma (*Y′*) is formed as a suitably-weighted sum of *R′G′B′*; it is the basis of luma/color difference coding. Luma is comparable to lightness; it is often carelessly and incorrectly called *luminance* by video engineers.

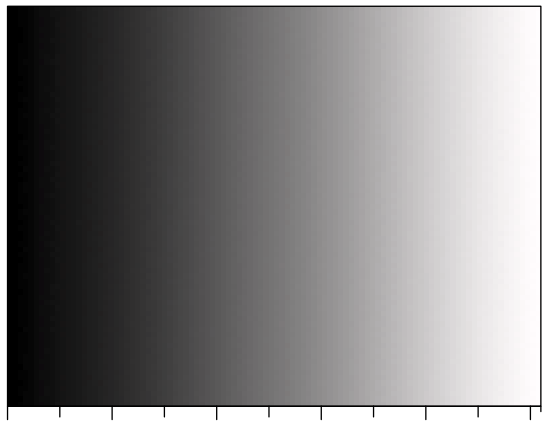

Grayscale ramp on a CRT display is generated by writing successive integer values 0 through 255 into the columns of a framebuffer. When processed by a digital-to-analog converter (DAC), and presented to a CRT display, a perceptually uniform sweep of lightness values results. A naive experimenter might conclude – mistakenly! – that code values are proportional to intensity.

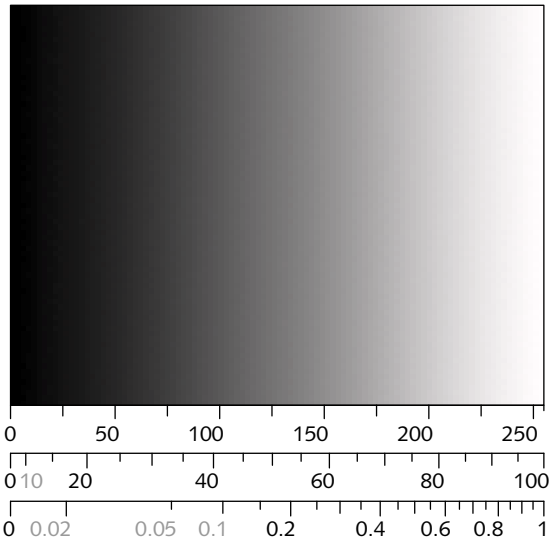

Grayscale ramp augmented with CIE lightness (L^*, on the middle scale), and CIE relative luminance (Y, proportional to intensity, on the bottom scale). The point midway across the screen has lightness value midway between black and white. There is a near-linear relationship between code value and lightness. However, luminance at the midway point is only about 20 percent of white! Luminance produced by a CRT is approximately proportional to the 2.5-power of code value. Lightness is roughly proportional to the 0.4-power of luminance. Amazingly, these relationships are near inverses. Their near-perfect cancellation has led many workers in computer graphics to misinterpret the term *intensity*, and to underestimate the importance of nonlinear transfer functions.

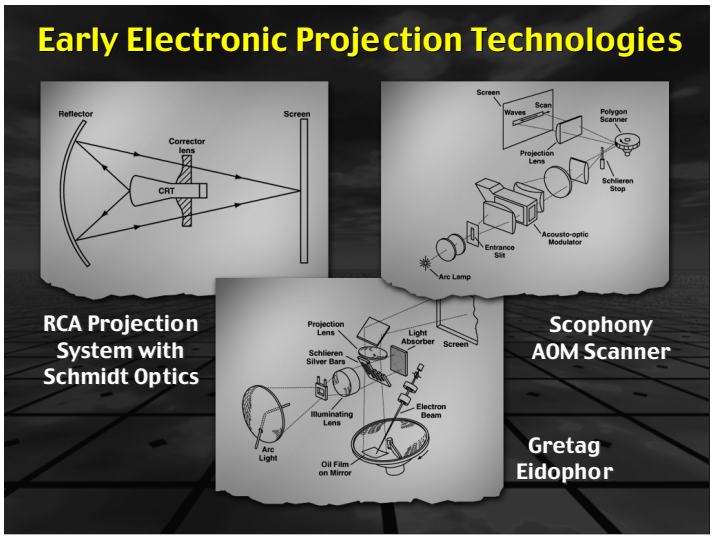

Early electronic projection technology has roots going back into the 1930s. The Scophony projector was first demonstrated in 1936 and used an acousto-optic modulator and a polygon scanner. The RCA CRT projector followed in 1940 and the Eidophor in 1943.

Electronic Projector Considerations

- Image Stability
- Optical efficiency
- Luminous flux
- Uniformity
- Support large screens
- Support multiple formats (1.85:1, 2.39:1)
- Flicker
- Bit depth
- Image lag
- Support multiple frame rates
- Contrast ratio
- Resolution (addressable vs. perceived)
- Pixelation (aliasing)
- Color gamut
- Weight
- Power consumption
- Booth compatibility
- Ease of maintenance
- Reliability
- Cost

Copyright © 2001-05-10 Texas Instruments, Inc.

Projection Technology

- **DLP Cinema™ (Digital Light Processing)**
  - **Texas Instruments**
- **ILA → D–ILA (Direct Drive Image Light Amplifier)**
  - **JVC Visual Systems**
- **GLV (Grating Light Valve)**
  - **Silicon Light Machines**
- **Laser**
  - **Principia Lightworks**
  - **COLOR Corp**

DLP Cinema Technology was developed by Texas Instruments, based on the Digital Micromirror Device.

DLP Cinema Technology

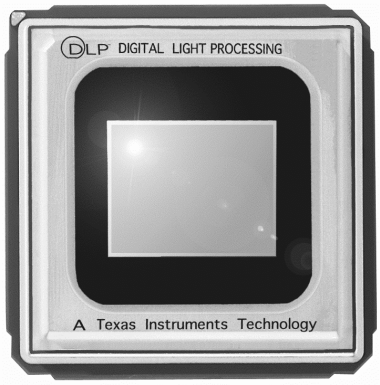

Details of DMD™ Superstructure Process

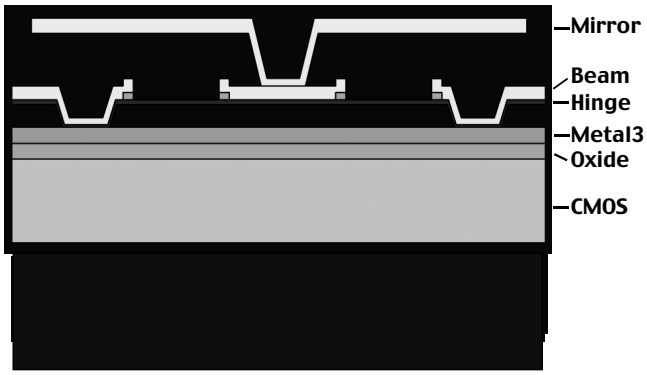

DMD Technology. The Digital Micromirror Device (DMD) is classified as a microelectromechanical system or MEMS. The micromirror superstructure is monolithically fabricated on a standard CMOS wafer. Each mirror is physically located over a SRAM based memory cell.

Close-up of Mirror Substructure

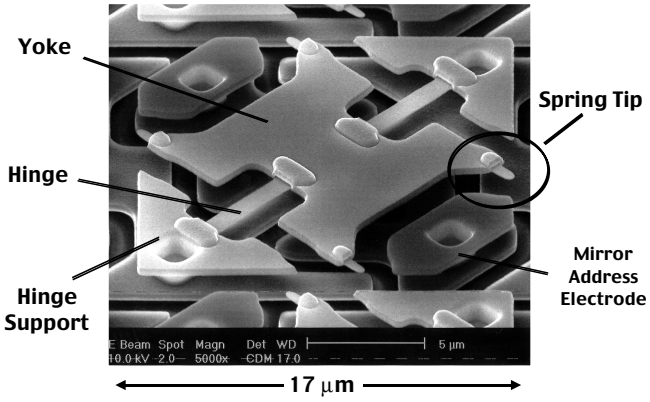

Mirror Substructure. Each mirror is attached to an underlying yoke which is suspended by two torsion hinges. Landing tips on each yoke make contact with the substrate as the mirror is rotated. Each mirror is independently operated in a bistable mode.

Early Pixel Architecture (1993)

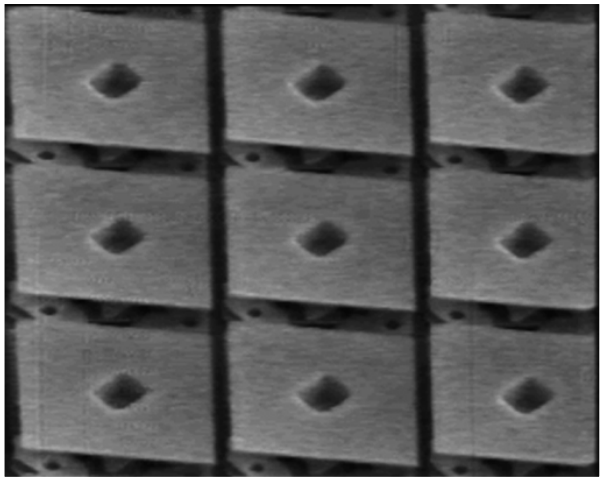

DMD Mechanical Switching. Each mirror is 16 µm square and rotates ±10° based on the contents of the underlying memory cell. Electrostatic forces established by the electrodes rotate the mirror about the hinges. Mechanical switching speeds of ~15 µs and optical switching speeds of ~5 µs are achieved.

DMD Mechanical Switching

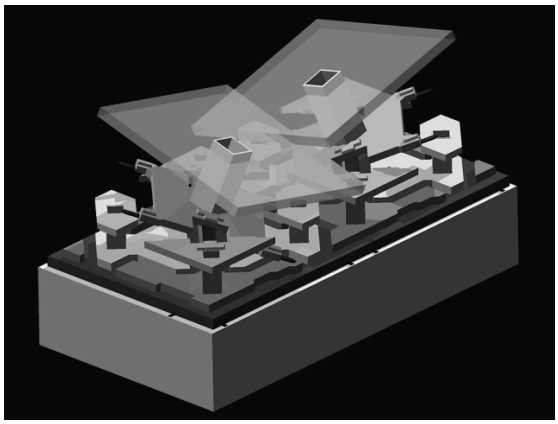

- 16 micron aluminum mirrors
- ± 10 degree rotation
- Hard rotation stops
- Switching time 15 µS

DMD mechanical switching

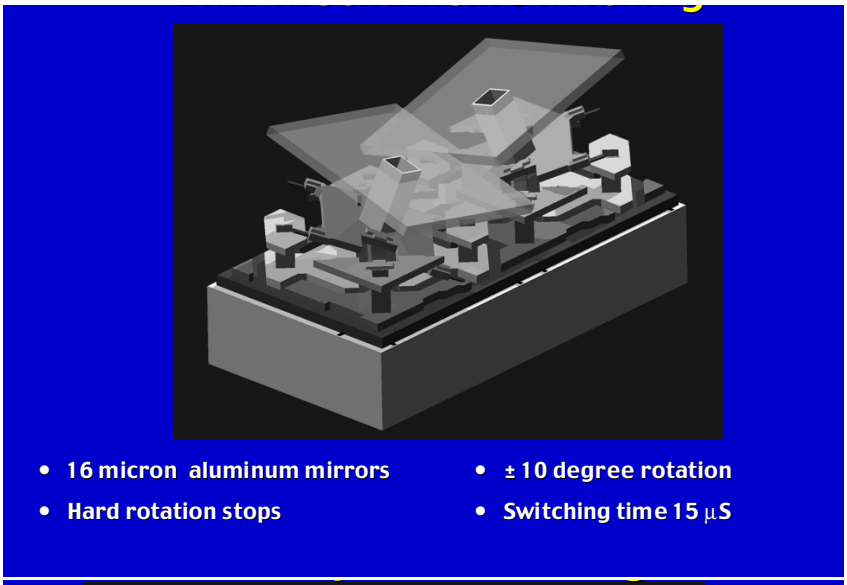

DMD optical switching

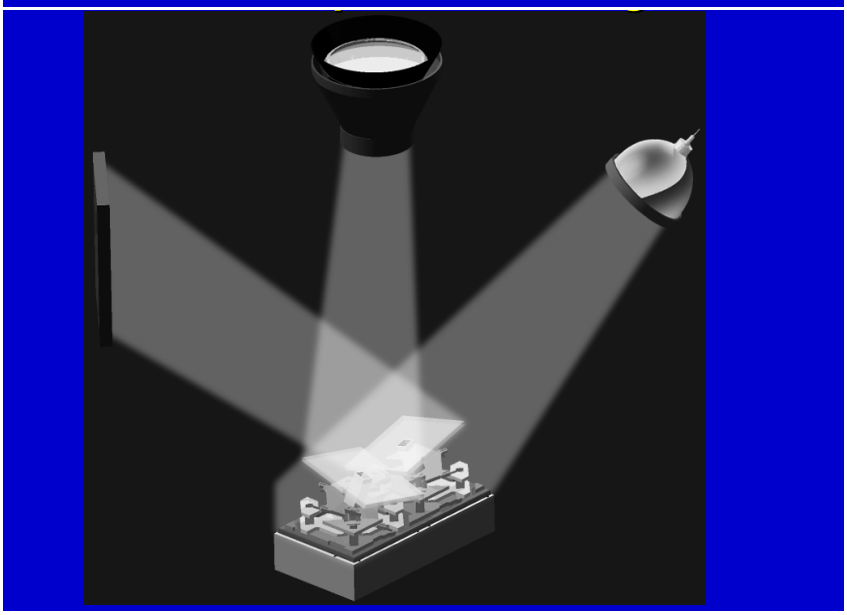

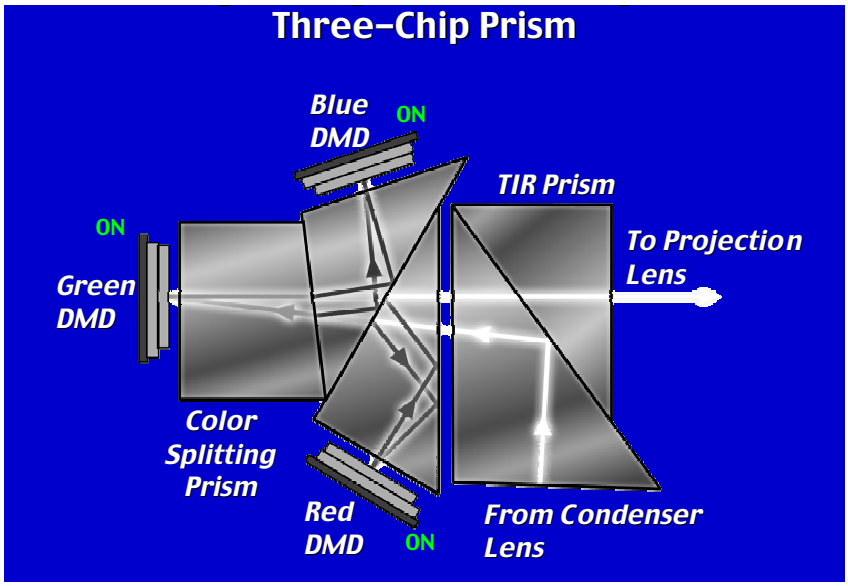

3-Chip DLP Projectors. High-end applications such as digital cinema utilize a three DMD chip configuration. Light from a Xenon lamp is passed through an integrating rod for improved uniformity. The white light is then directed into a color splitting prism assembly where it is modulated by red, green, and blue DMDs. The modulated light is recombined and projected on to the viewing screen.

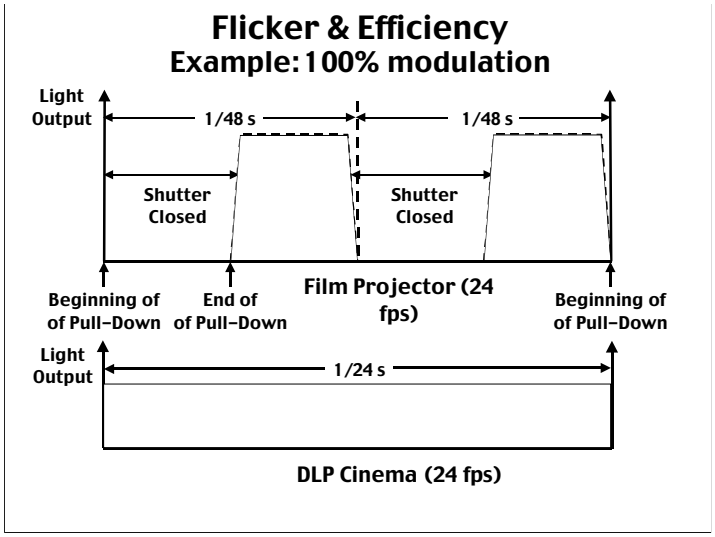

Flicker & Efficiency. Comparing film projection to DLP, approximately 50% of the available light is absorbed by the film projector shutter. Typically, a double shutter is used to flash the image on the screen twice per frame. This raises the flicker frequency to 48 Hz, a more tolerable but still visible level. DLP technology can utilize available light over the entire frame time since there is no requirement to douse the light to mechanically pull the next film frame into the projection aperture.

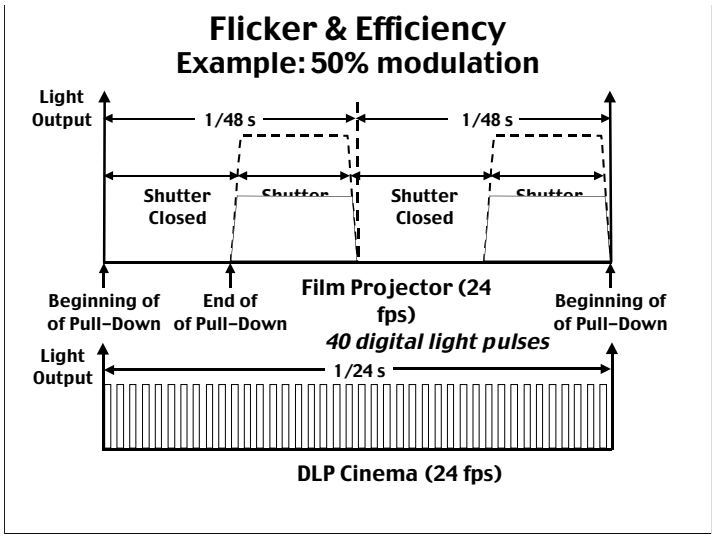

Resolution Delivered to Audience

"Because of the statistical nature of frame unsteadiness and its attendant resolution loss, the vast majority of motion-picture viewers have been enjoying their entertainment experience with a good deal less than 700 TVL (*lines per screen height)*."

*** Source: A. Kaiser, et. al.,"Resolution Requirements for HDTV Based Upon the Performance of 35mm Motion-Picture Films for Theatrical Viewing," in *SMPTE Journal*, pp. 654-659, June 1985.**

Film Resolution Degradation
Contact Printing

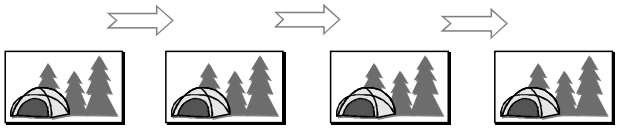

Original Negative → Interpositive → Internegative (Dupe Negative) → Release Print

- MTF generational loss
- Image unsteadiness
  - Frame-to-frame registration
  - Perforation tolerances
  - Film stock shrinkage
  - Increases with number of generations

Film Resolution Degradation
Film Projection

- Projection lens MTF
- Image Unsteadiness
  - Focus flutter
  - Jump
  - Weave

Image Unsteadiness

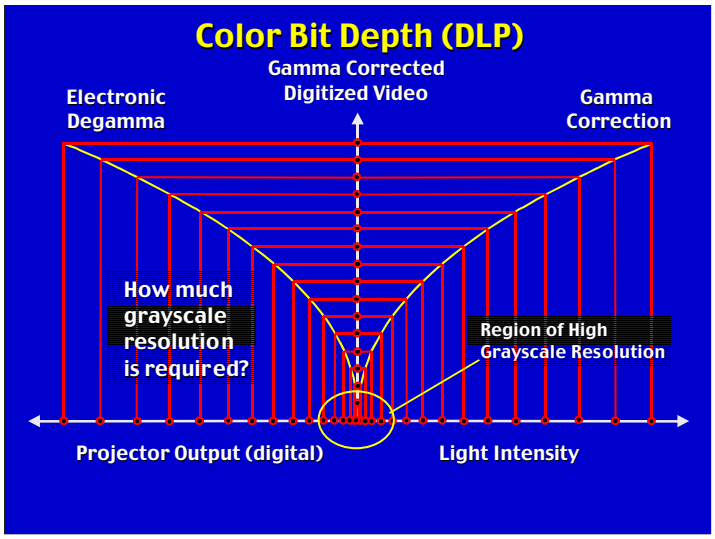

Color Bit Depth. For compatibility and coding efficiency reasons, gamma corrected data is received by the DLP system. Since the DMD has a linear response characteristic, the data must be mapped from gamma corrected into linear form. For perceptual reasons, sufficient quantization levels must be available, particularly in the low end, to avoid contouring or banding artifacts.

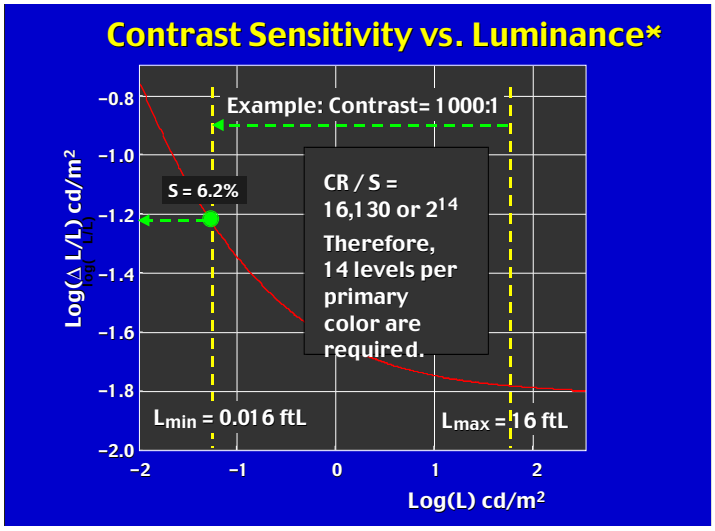

Contrast Sensitivity vs. Luminance. From the graph, given a system contrast and maximum luminance level, the minimum number of quantization levels can be determined to avoid objectionable artifacts.

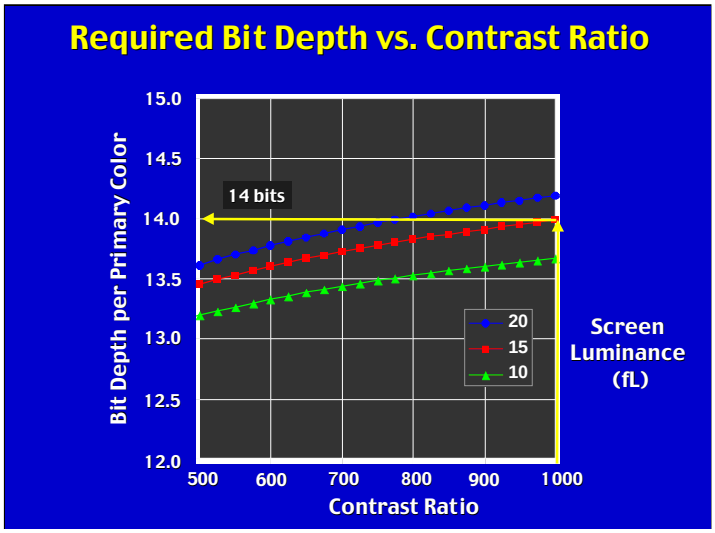

Required bit depth vs. contrast ratio. Using data from Wyszecki & Stiles, a series of curves can be generated for various maximum screen luminance levels to yield minimum bit depth requirements for a given contrast ratio.

Color Gamut. The expanded color gamut of the DLP Cinema projector compares closely with the color gamut of film (Equivalent Neutral Density, END = 2.5).

DLP Cinema™ Prototype Projector
Color Gamut

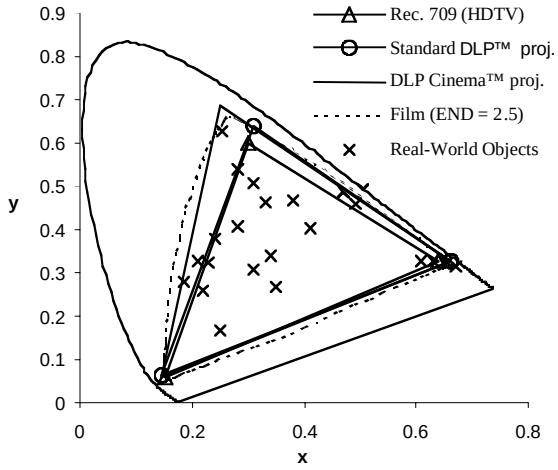

DLP Cinema™ Prototype Projector
Key Elements

- **Leverage TI investments in DLP & DMD technology.**
- **Standard film projector lamp housing.**
- **Same optical axis as film head.**
- **Anamorphic projection for flat (1.85:1) and *scope* (2.39:1) using 1280 × 1024 DMD.**
- **Display at 24 fps for "cinematic look."**
- **Increase power handling capacity to "50 feet".**
- **Meet or exceed contrast of typical release print.**
- **Improve color processing capabilities.**
- **Improve effective bit depth to 42 bits.**

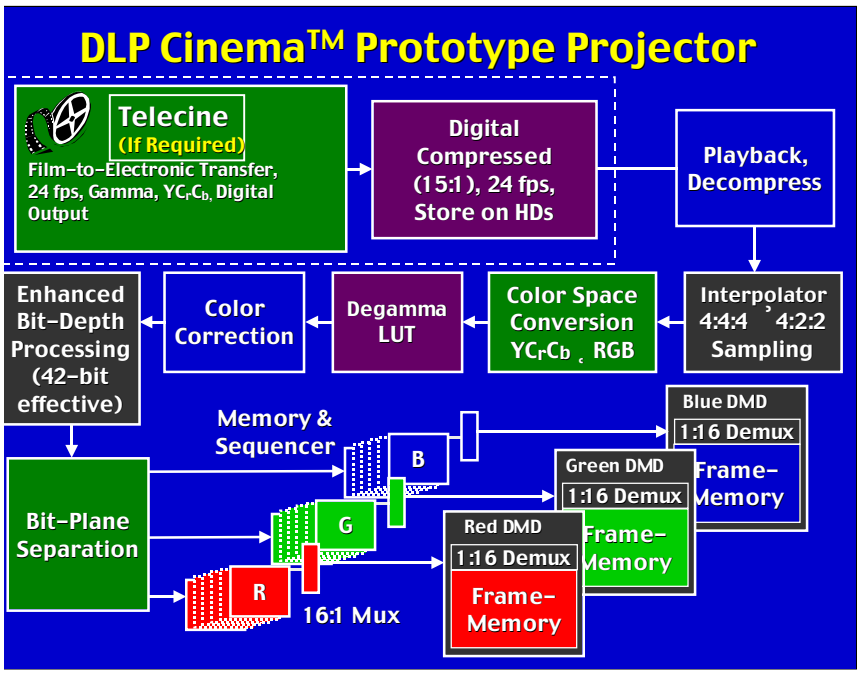

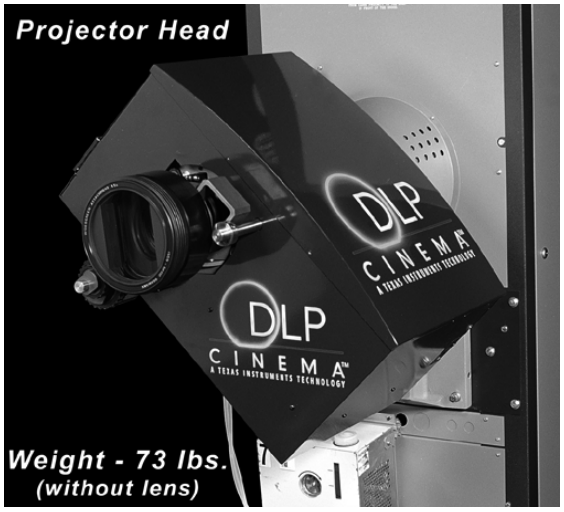

DLP Cinema Mark V Prototype Projector mounted on a conventional film projector lamp house. The standard f/2.0 lamp reflector is replaced with a f/1.5 reflector. The projection head has been tested and approved by UL for power levels up to 7 kW.

Inline Light Path. To maintain compatibility with existing theatre projection ports, the lightpath into the projector is inline with the projected light output.

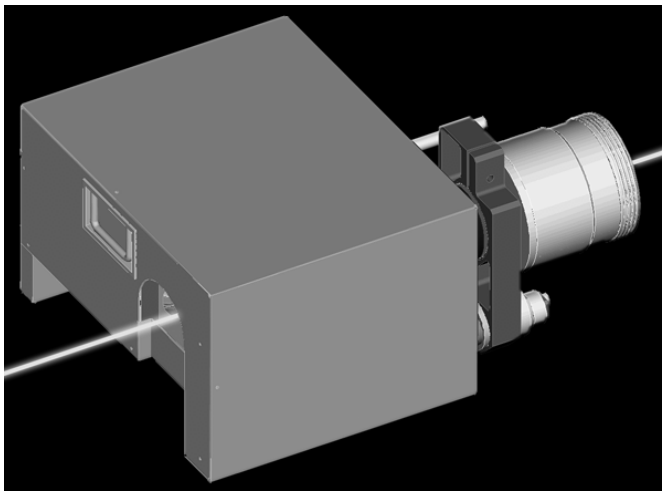

Compact internal light path. The light path is folded within the projection head for a compact package design.

DLP Cinema™ Projection Head

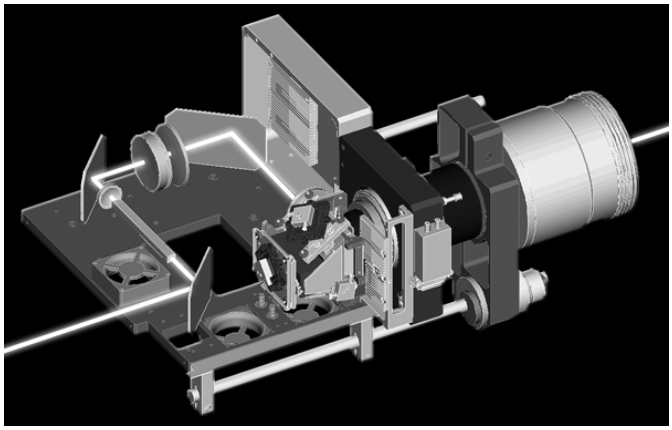

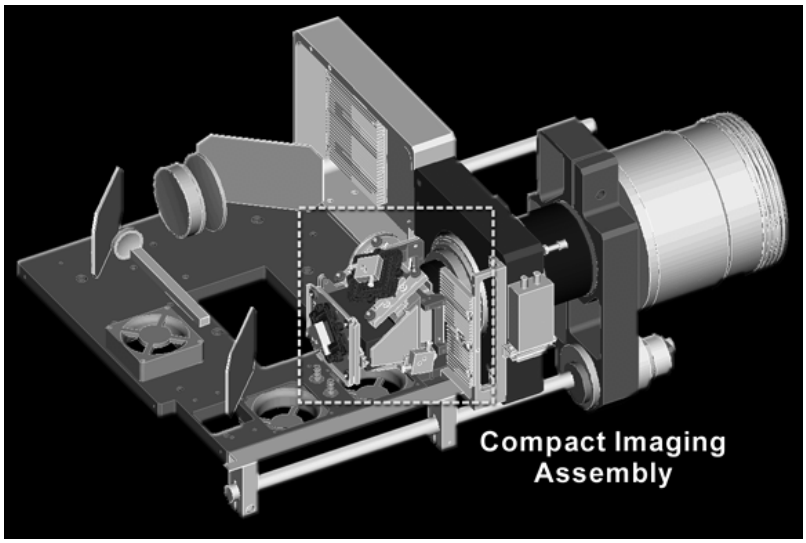

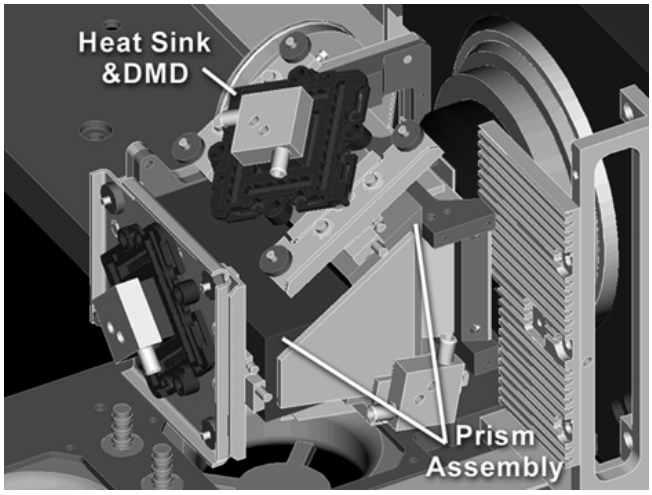

Prism Assembly

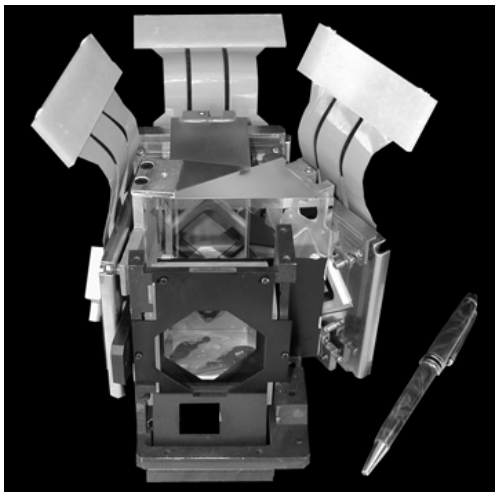

Prism Assembly. Liquid cooled heat sinks are attached to the back side of each DMD to provide adequate cooling for the high intensity illumination.

Anamorphic Projection. To achieve the two common cinema formats (flat and scope), anamorphic projection lenses are used. For flat presentations, a 1.5x anamorphic lens is used to achieve the 1.85:1 aspect ratio. For scope, a 1.9x anamorphic lens is used to achieve a 2.39:1 aspect ratio. Conventional film projection uses a 2x anamorphic lens for scope and standard spherical optics for flat presentations.

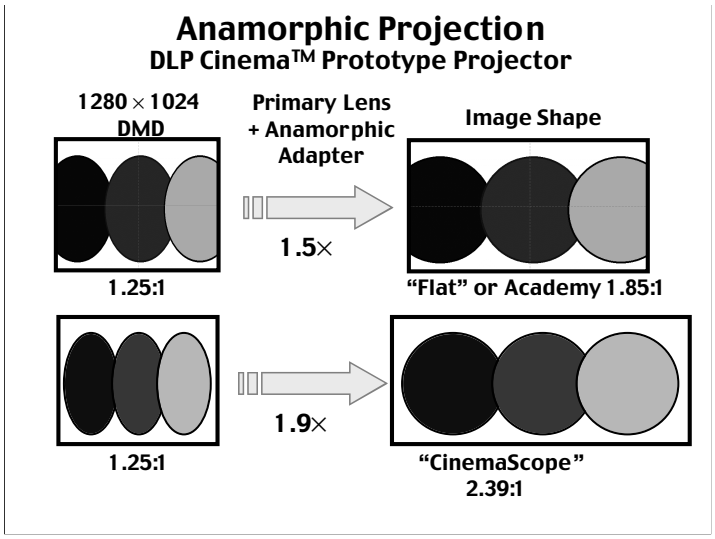

Native Image Size. In comparing the size of the image contained on a standard scope release print to that of the image generated on the DMD, the two are within 1 mm in each dimension.

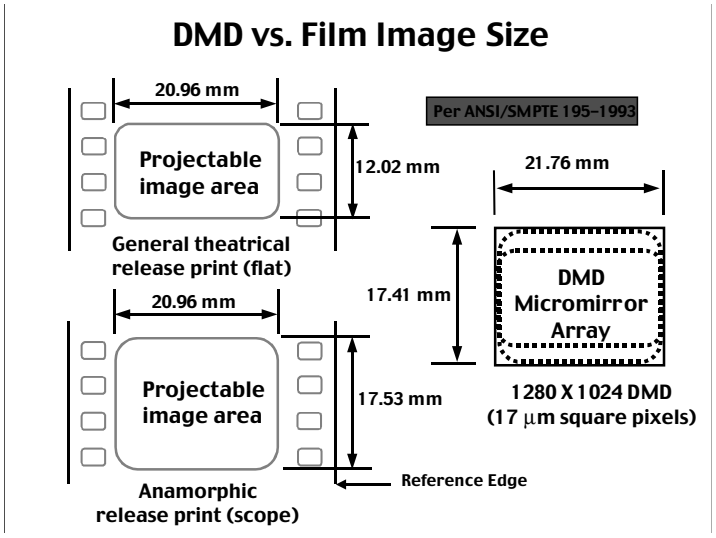

DLP Cinema™ Prototype Projector Performance

- Resolution 1280 × 1024 discrete pixels
- Display frame rate 24 fps
- Brightness 10,000+ lumens
  - (w/5,000 watt lamp)
- Contrast >1000:1 full-on/full-off
- Luminance 12.5 fL in scope
  - (50 ft screen, S/G = 1.3) (SMPTE 196M: 16 ± 2 fL)
- Luminance Uniformity 15% (SMPTE 196M: <25%)
- Color Temperature 5400 to 6500 K
  SMPTE 196M: 5400 ± 400 K)
- Effective Bit Depth 14 bits/component – linear
- Weight (projector head) 73 lbs. without lens

DLP Cinema™ Technology Advantages

- Digital fidelity
  - Color and grayscale determined by the precise division of time.
  - No lag or motion smearing as in analog light valves.
- Digital stability
  - No photodegradation effects.
  - No thermal or aging effects.
  - Quick setup and low maintenance.

Application of DLP™ Technology to Digital Electronic Cinema—A Progress Report

William B. Werner and D. Scott Dewald
Texas Instruments Incorporated
Plano, Texas

Abstract

In the entertainment industry's quest for the "suspension of disbelief," two primary sensory stimuli have been engaged: sight and sound. Since the advent of digital audio, a number of new multitrack digital sound formats have appeared in the motion picture industry, all in an attempt to achieve the ultimate audio experience. However, visual presentation has not advanced at the same pace as the aural experience. Many problems associated with conventional motion picture presentation and distribution can be overcome through a digital form of delivery. The missing link, however, has been an electronic display system that delivers cinema-quality images. Certainly the ultimate achievement in electronic projected image quality would be to match or exceed that of motion picture film.

Texas Instruments is studying the possible application of DLP™ technology to electronic cinema. This effort has involved a close working relationship with the motion picture industry to assess DLP performance from technical and artistic perspectives. Knowledge gained from this process has led to improvements in DLP technology that may someday lead to a DLP-based cinema application. This paper addresses some of the advances in image quality made in the quest to achieve the ultimate cinematic experience.

Introduction

The "Holy Grail" of the high-end electronic projector manufacturer has been to become as good or better than 35 mm film projection—to replace the motion picture film projector. The motion picture industry has embraced the digital age in the production of movies, yet motion picture distribution and presentation use technology that remains virtually unchanged from that of the early 1900s. As far as the image is concerned, the inherently analog process of depositing dyes on a film base remains the technology of choice. Recently, there has been a dramatic evolutionary change in motion picture sound reproduction from the early monaural optical tracks to the more recent multitrack digital formats, such as Dolby™ Digital, DTS™ and SDDS™. But why has the visual experience not experienced this same degree of evolutionary change?

Clearly, enabling technology exists today to distribute motion picture material digitally, whether via satellite, telephone company, or a physical medium such as optical disc. The missing link has been an electronic projector capable of reproducing the high-quality images demanded by movie-goers. Electronic projection's rather slow start can be attributed to economic, as well as technological factors. Past attempts at high-brightness projection used expensive and often difficult-to-maintain solutions. These include oil-film light valves, optically addressed liquid-crystal light valves, and high-power laser projection. Image quality aside, the maintenance and implementation costs for these technologies will probably prohibit wholesale adoption for general cinema exhibition anytime soon.

In 2Q97, Texas Instruments Digital Light Processing™ (DLP™) business began a technology investigation into electronic cinema projection using DLP technology. Based on the digital nature of our technology, the study involved an all-digital implementation for electronic cinema or in this case, *DLP Cinema.* The purpose of this effort was twofold: first, identify the customers and their key care-abouts; second, determine relative positioning of DLP technology compared to film projection picture quality and other electronic projection technologies.

Background

In 1995, work began on a high-resolution demonstration projector based on a 1280 × 1024 mirror array—the Digital Micromirror Device (DMD™), also known as DMD1210. This projector, called the 1210 Technology Demonstration Projector, was intended to demonstrate the high-resolution capabilities of the DMD1210 by displaying 24-bit/pixel, SXGA graphics through a unique PC graphics board designed with a custom digital interface. The system was based on a three-chip DLP configuration (Figure 1) and was illuminated by a 1 kW xenon lamp. A 3.5:1 fixed focal length projection lens was designed, along with the electronics necessary to process and display a 2.6 Gbps data stream from the SXGA graphics driver (Figure 2). This system gave superior performance in what was perhaps the first real-time end-to-end *all-digital* display, where the human eye serves as the final digital-to-analog converter.

Reprinted with permission from SMPTE, *Proceedings of the 140th Technical Conference*, Pasadena, Calif., Oct. 28–31, 1999 (White Plains, N.Y.: SMPTE, 1999), 406–412.

Copyright © 2001-05-11 SMPTE

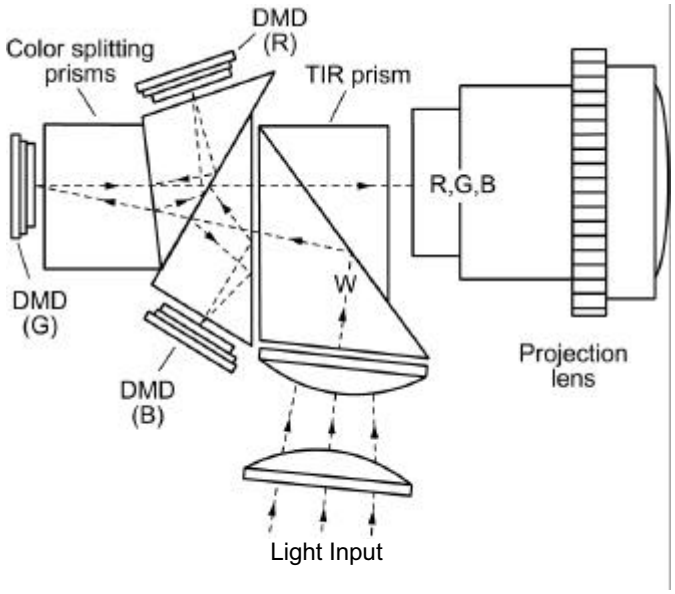

Figure 1. Three-chip DLP projector configuration (simplified for clarity)

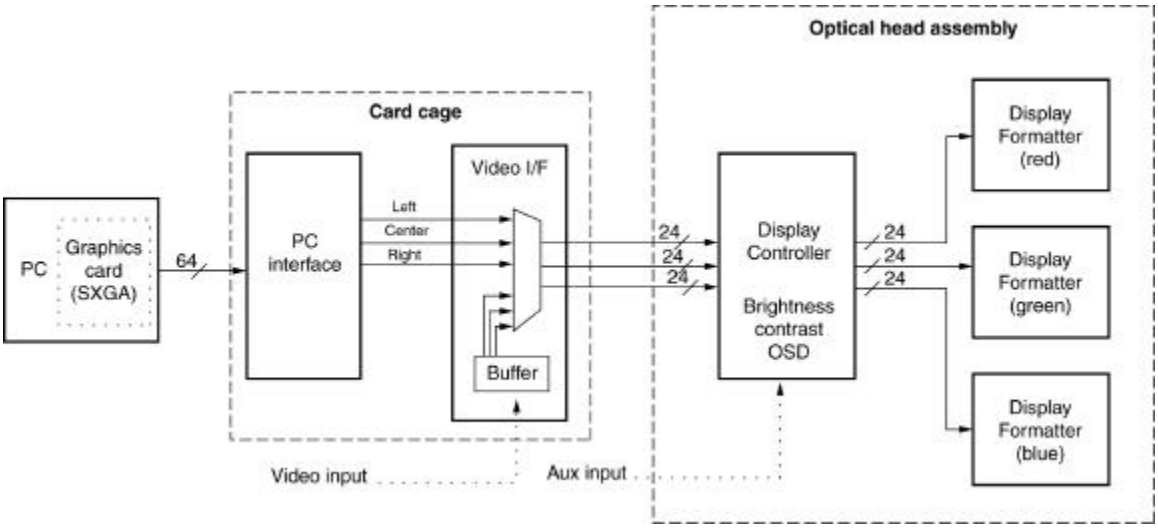

Figure 2. 1210 Technology Demonstration Projector electronics

The DLP Cinema Technology Demonstration Projector

The success of the 1210 Technology Demonstration Projector helped fuel the desire to push the performance envelope and explore new venues. To help understand and define DLP technology's potential, work began on a DLP Cinema Technology Demonstration Projector. The goal was rapid upgrade of our existing 1210 Technology Demonstration Projector to produce high-quality, high-brightness digital images for evaluation in cinema applications. The design and evaluation of this projector would foster a close working relationship with technical and creative members of the motion picture industry. The result was an improved understanding of overall system performance requirements.

Work began in May 1997 on the 1210 system upgrades, with the goal of a September 1997 demonstration to key industry representatives. The upgrades included:

- Complete redesign of the optical illumination system (Figure 3)
  - Increase lamp power from 1 to 2.5 kW using off-the-shelf xenon lamps currently used in the motion picture theater industry
  - Custom design of a dichroic lamp reflector with infrared absorption properties
  - Addition of an optical integrator to improve field uniformity.
- Design of an anamorphic attachment to the projection lens to permit the American film industry standard 1.85:1 aspect ratio images to be displayed from a 1.25:1 device (Figure 4)

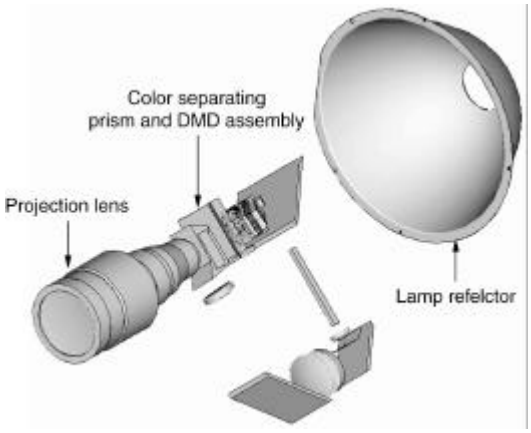

Figure 3. DLP Cinema Technology Demonstration Projector illumination optics design

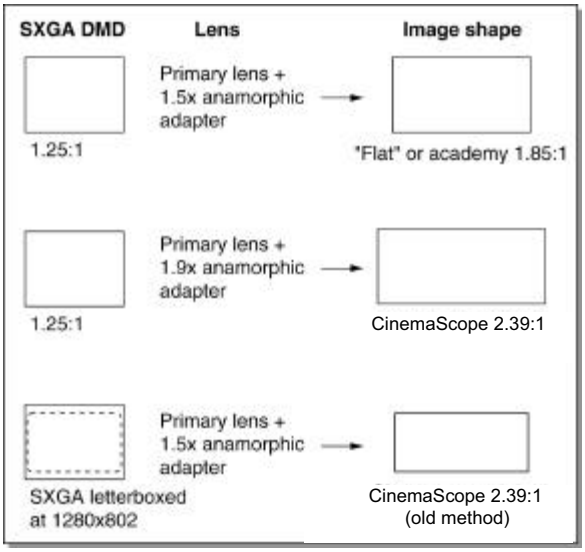

Figure 4. Use of anamorphic lenses

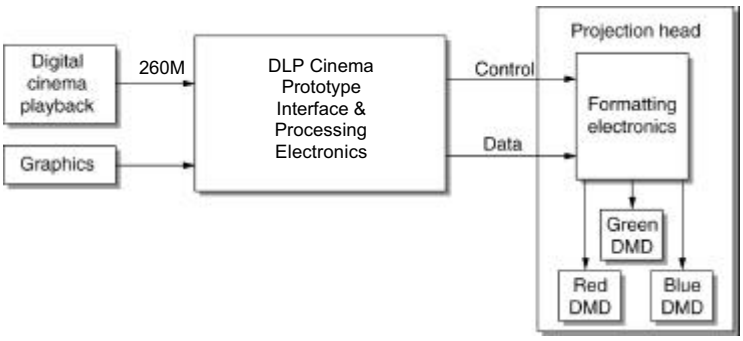

Figure 5. DLP Cinema Technology Demonstration Projector electronics

- Addition of a high-definition digital video interface based on the SMPTE 260M bit-parallel digital interface (Figure 5)
  - Implementation of a reverse 3:2 pull-down process to produce the original 24 fps motion picture images from 60 field/second input data
  - Development of techniques for displaying 24 fps images by the DMD.

Because of long lead times for optical fabrication, the team used the existing prism assembly and projection lens for the new demonstration projector. There was time to redesign the illumination optics and to obtain lamps and reflectors to increase the projector brightness and an anamorphic projection lens attachment. Anamorphic projection techniques are not new to the motion picture industry, as many films are released in CinemaScope, a film format

using a 2:1 anamorphic projection lens. The illumination system was designed to increase contrast and image uniformity. Cinema-type xenon lamps were also integrated into the optical system to achieve the desired brightness.

To more closely emulate the *film-look* in the projected image, a unique capability of the DMD was used to display the images at the standard motion picture industry rate of 24 fps. Since DLP displays are a continuous display technology, flicker-free display rates down to 24 fps are possible without the need for temporal up-sampling. Furthermore, the sub-20 µs switching speed of the mirrors makes possible displays that are free from image lag. Another advantage of DLP technology is that the image is not shuttered to blank the display as the next frame is pulled down into the film gate, as in a conventional film projector. This advantage can yield greater projector efficiency and reduce artifacts at high brightness levels.

Commonly available high-definition recording and playback equipment operates at a 60 field/second rate. A process known as 3:2 pull-down is employed to increase the temporal resolution from the 24 fps film rate to the 60 field/second video rate during the film-to-video conversion (Figure 6). To obtain and display the original 24 fps data, a technique was developed to reverse the 3:2 pull-down processed data. This produced a more film-like experience, eliminating artifacts associated with 3:2 pull-down processed displays.

The overall performance parameters of the DLP Cinema Technology Demonstration Projector were:

Resolution:	1280 × 1024 discrete pixels
Brightness:	3100 lumens with 2500 W lamp
Contrast:	240:1 full-on/full-off
Luminance uniformity:	15%
Color temperature:	5400 to 5800 K
Optical distortion:	5% pincushion with anamorphic adapter
Bit depth:	8 bits/component with gamma, 10 bits/component linear
Display rate:	24 fps

This Technology Demonstration Projector was used from September 1997 through 2Q98, conducting performance evaluations with key members of the motion picture industry. We worked closely with some of the major movie studios, exhibitors, cinematographers, directors, and technology suppliers, all of whom shared a great deal of insight and a diverse perspective on our technology. From this, we added to our knowledge of the technical aspects of projection system requirements, in addition to the artistic aspects of the display system.

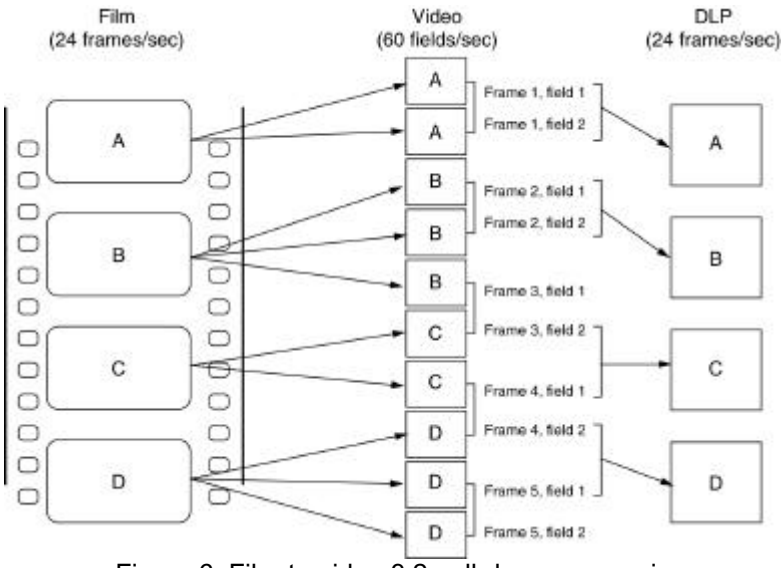

Figure 6. Film-to-video 3:2 pull-down processing

Customer feedback yielded the following key areas for improvement and near-term technology development:

- Greater system contrast
- Increased light output to support 50-foot screen widths
- Improved film-to-digital conversion process.

These results led to the next phase of development.

The DLP Cinema Prototype System

As a result of the experience and knowledge gained on the DLP Cinema Technology Demonstration Projector, work began in 4Q97 on a DLP Cinema Prototype System, a system-level approach to an all-digital "cinema of tomorrow." The prototype system will include the following key aspects:

- Post-production authoring of the digital media
  - High-quality telecine process for conversion of a film master into digital form
  - Integration of audio data
  - Data compression
  - Data encryption.
- Medium for distributing data
  - Satellite transmission
  - Fiber-optic link
  - Optical media.
- Presentation of motion picture material with emphasis on preservation of artistic content
  - Playback
  - Decryption
  - Decompression
  - Audio reproduction
  - Image projection.

The above process will emphasize preservation of artistic as well as technical content. This process allows the audience to be immersed in the experience without the distractions caused by disturbing or unnatural artifacts in the image or sound. At the same time, it enhances the creative artist's freedom of expression. These ingredients are essential in maintaining the *suspension of disbelief* and, ultimately, the success of any electronic cinema solution.

Our approach to the *all-digital* system based on DLP technology is to use as much of the existing infrastructure and technology as practical, further leveraging the existing investments made by Texas Instruments. Development work on the DLP Cinema Technology Demonstration Projector indicates that existing film projector lamp housings could be used with a DLP-based projector head, provided minor modifications are made to the lamp reflector. In this case, an exhibitor could do the same when upgrading to a DLP-based cinema projector. The DLP Cinema projection head could attach to a lamp housing in the same manner that a film projector does, as illustrated in Figure 7. To further reduce financial and remodeling impact on theater installations, the DLP projector will not deviate from the optical axis currently used by film projectors, thus allowing the continued use of existing and often quite small projection windows. This minimizes or avoids modifications to existing theater projection booths.

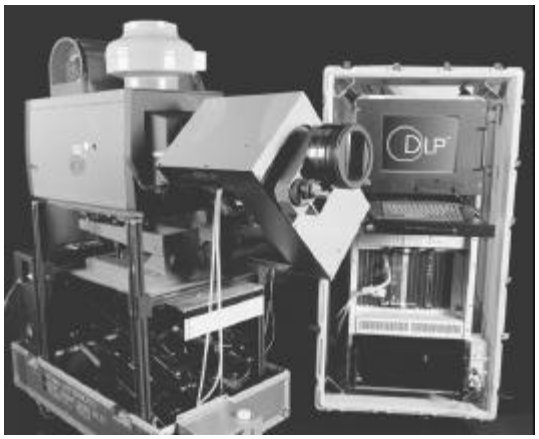

Figure 7. DLP Cinema Prototype System projection head, interface electronics and lamp housing

To date, we have completed design and construction of the DLP Cinema Prototype System projection head, interface, and display electronics. The projection head design employs many optical technology breakthroughs not available in the previous projector. These improvements increase the power-handling capability, contrast, and brightness of the projector several-fold. The system contrast ratio has been increased 300%, while the overall efficiency of the DLP Cinema projector has exceeded that of 35 mm film, as shown in Figure 8. Current DLP operation using a 5 kW lamp projects 14 ftL onto a typical 50-foot screen in flat format. Assuming a film transmissivity of 87%, the SMPTE 196M specified 16 ftL equates to approximately 14 ftL from a film projector running with clear (D-min) film.

New projection lenses for the two common North American film formats were designed. Currently, motion picture equipment projects the two common formats (flat and CinemaScope, or scope) using different projection lenses. The flat format has a 1.85:1 aspect ratio and is projected with a conventional spherical lens, while the CinemaScope format has a 2.39:1 aspect ratio and uses an anamorphic or cylindrical projection lens with a 2× horizontal expansion. For efficient reproduction of these two formats with a projector based on the 1280 × 1024 DMD and an aspect ratio of 1.25:1, two anamorphic expansion ratios are needed: 1.5× for the flat format, and 1.9× for the scope format. It should be noted that the actual size of the 1280 × 1024 DMD is within 1 mm of the projectable image size of a scope film print (Figure 9).

The projector's color processing capabilities were improved in several ways. We increased the nonlinear bit resolution from 8 to 10 bits/component and the linear space bit resolution from 10 to 14 bits/component. A multitap FIR filter was employed to up-convert the subsampled chrominance data from 4:2:2 to 4:4:4. The 3 × 3 matrix multiplier used for color-space conversion from component YCrCb to RGB was increased from 8 to 12 bits/component, and a second 3 × 3 matrix multiplier with 14-bit/component capability was added to provide linear space color correction control. We also added more processing electronics to allow for a future image optimization capability. A fiber-optic data link was used to transport high-bandwidth data between the playback and interface electronics and the projection head. This facilitates remote location of the projector as well as optical isolation of the projection head from playback electronics.

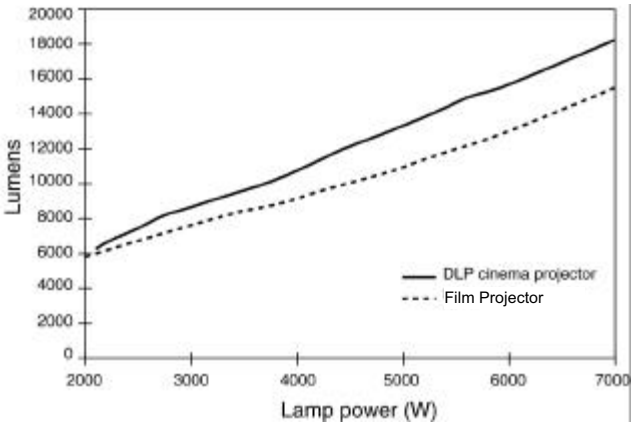

Figure 8. Film versus DLP efficiency

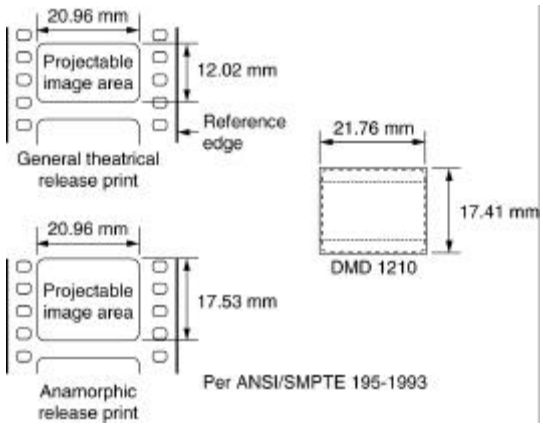

Figure 9. DMD versus film image size

The overall performance parameters of the current DLP Cinema Prototype System Projector are:

Resolution:	1280 × 1024 discrete pixels
Brightness:	13,000+ lumens with 5000 W lamp
Contrast:	>800:1 full-on/full-off
Luminance Uniformity:	15%
Color Temperature:	5800 to 6500 K
Optical Distortion:	5% pincushion with 1.5× anamorphic adapter
Bit Depth:	10 bits/component with gamma, 14 bits/component linear
Display Rate:	24 fps

Summary

In the past 18 months, we have acquired much knowledge regarding possible future applications of DLP technology. The performance of the DLP Cinema Technology Demonstration Projector exceeded the expectations of many. The follow-on DLP Cinema Prototype System has pushed the performance to an even higher level. System contrast increased to better than 800:1 and light output to date has been boosted to over 13,000 lumens. Anamorphic projection at both 1.5:1 and 1.9:1 has proven successful as well. The use of standard, high-power, off-the-shelf xenon lamps has opened the door to cost-effective projector retrofitting. Clearly, these breakthroughs are essential for DLP technology to succeed in a cinema application. However, one cannot overlook the possibility that leading-edge developments resulting from this work will probably find applications in existing and yet-to-be-developed DLP product lines.

The future of motion picture film distribution is unclear. However, it is clear that the advantages in terms of cost savings to the studios for duplication and distribution of material, not to mention the security aspects through data encryption, strongly favor a digital rather than film-based scenario. Arguably, advantages in terms of image quality also favor a digital form of delivery, overcoming the multitude of problems associated with typical release prints:

- Accumulation of dirt and scratches on the film
- Jump, weave, and travel ghosting as the film passes through the projector
- Color fading and focus flutter caused by high-power lamp illumination.

With a DLP-based display system, a highly stable, highly repeatable presentation can be delivered with no image lag or degradation over time, temperature, or brightness level. Every theater can produce the exact same answer print quality presentation as the artist intended, without the variations often found in film release prints. Furthermore, the image can be displayed at the original 24 fps rate without the need for double shuttering or temporal data up-sampling to avoid display flicker. A DLP-based digital cinema display system can open the door to new opportunities for the exhibitor, as well. Live performances and events can be presented, as well as new forms of between-show advertising.

The new knowledge we gain from the DLP Cinema Prototype System will lend further insight into the market and requirements for a product solution. The response to DLP technology in this application has been very favorable. However, much work still lies ahead in the areas of data compression, encryption, storage, and media delivery. Further blind test studies are also needed. Until now, no electronic projection technology could offer the visual clarity necessary to compete with film while also providing stability and repeatability.

Trademarks

DLP is a registered trademark of Texas Instruments Incorporated
Dolby is a registered trademark of Dolby Laboratories
DTS is a registered trademark of Digital Theater Systems
SDDS is a registered trademark of Sony Cinema Products Corporation
DMD is a registered trademark of Texas Instruments Incorporated

References

1. G. Hewlett and W. Werner, "Analysis of Electronic Cinema Projection with the Texas Instruments Digital Micromirror Device™ Display System," *Proceedings of the 137th Society of Motion Picture and Television Engineers Technical Conference*, September 1995.

2. L.J. Hornbeck, "Current Status and Future Applications for DMD™-Based Projection Displays," to be published in the *Proceedings of the IDW '98 – the 5TH International Display Workshops*, Kobe, Japan, December 1998.

3. Mark Davis, "The AC-3 Multichannel Coder," *Proc. of the 95th Audio Engineering Society Convention*, October 1993.

4. S.M.F Smyth et al., "DTS Coherent Acoustics. Delivering High-Quality Multichannel Sound to the Consumer," *Proceedings of the 100th Audio Engineering Society Convention*, May 1996.

5. Uzi Efron, "Spatial Light Modulator Technology," Marcel Dekker, Inc., 1995.

6. M.A. Karim, "Electro-Optical Displays," Marcel Dekker, Inc., 1992.

7. L. Bernard Happe, *Your Film and the Lab,* Second Edition, Focal Press, 1983.

8. Steven Bernstein, *Film Production,* Second Edition, Focal Press, 1994.

9. Digital Light Processing™ web site: http://www.ti.com/dlp

The technology of D-cinema:

Color Representation

Chuck Harrison
Far Field Associates

What is film color ? . . .

- Gamut — range of colors
  - Chromaticity gamut (2-dimensional)
  - Color gamut
- Tonal resolution — smoothness of grays
- Spatiotemporal texture

and how do we represent it?

- There are several color spaces
  - Original scene
  - Production intermediate
  - Theatrical projection
  - Home viewing
- Device independence
- Transfer functions

Chromaticity gamut

$$x = \frac{X}{X+Y+Z}$$

$$y = \frac{Y}{X+Y+Z}$$

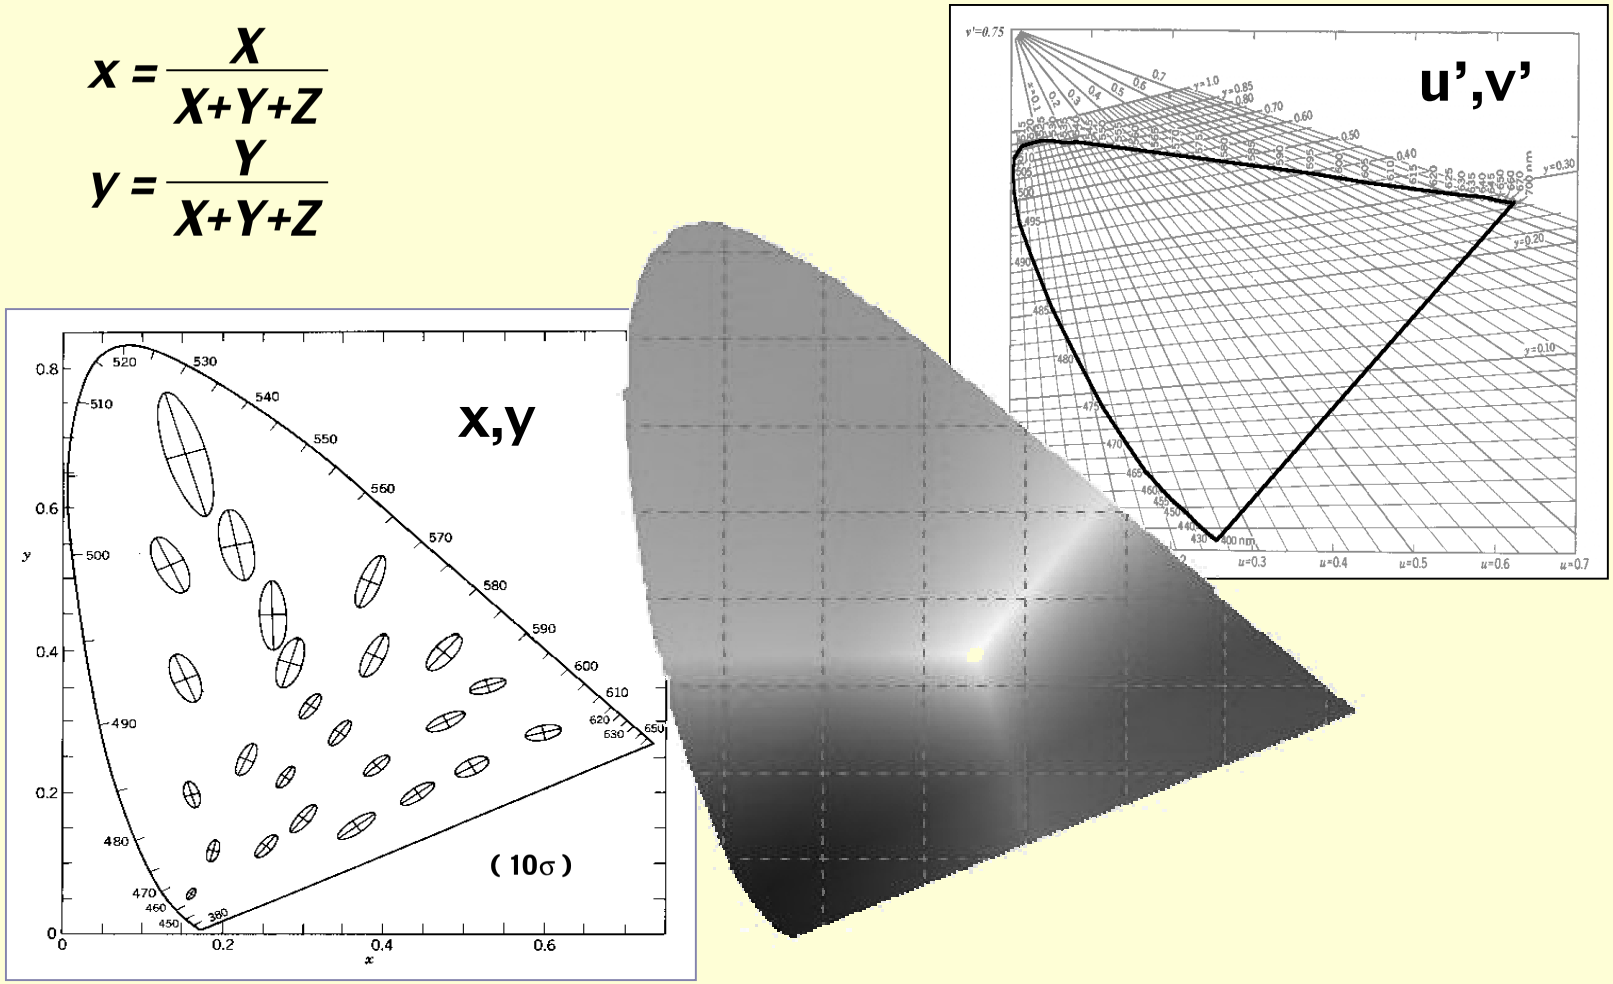

2-dimensional gamuts

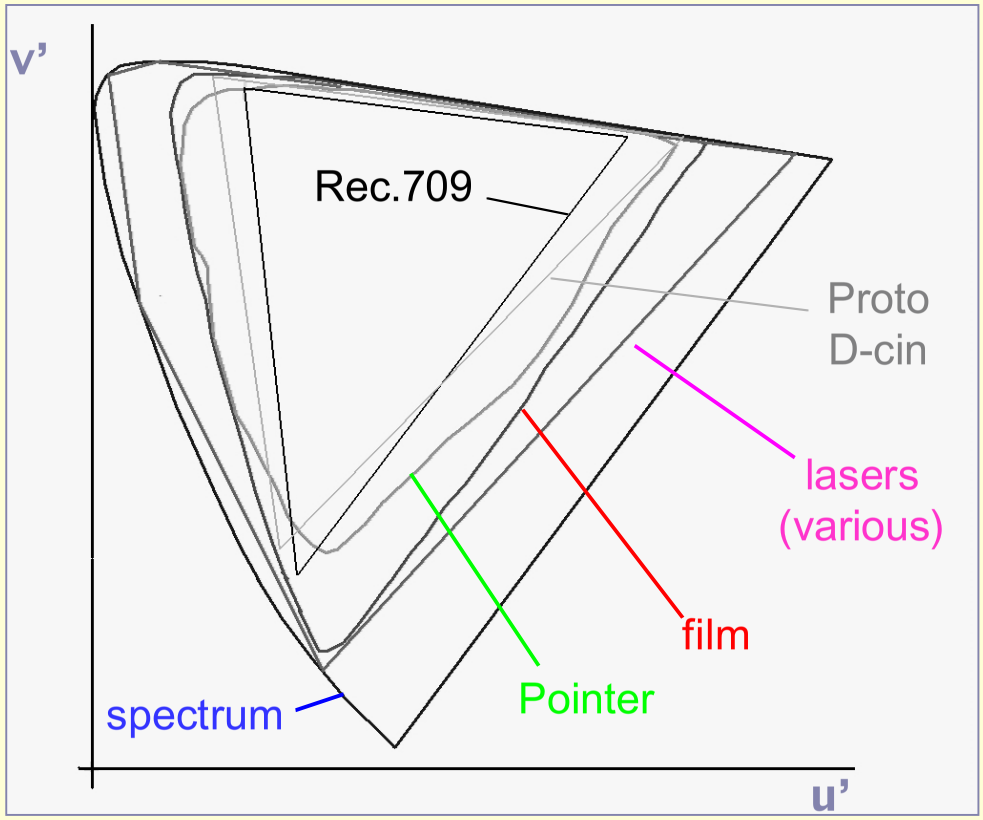

but pure (saturated) colors are dim

3-dimensional gamut

u -v -Y, for lights

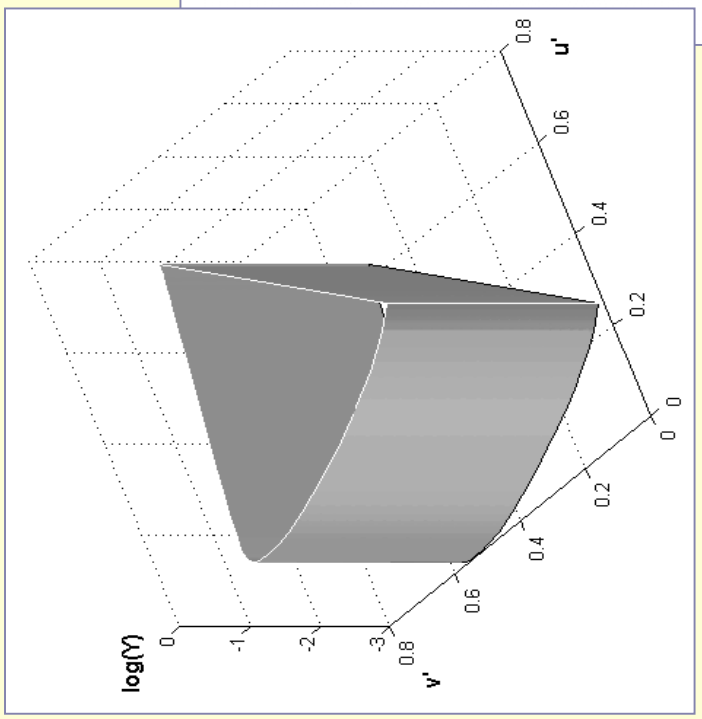

Munsell, for surfaces

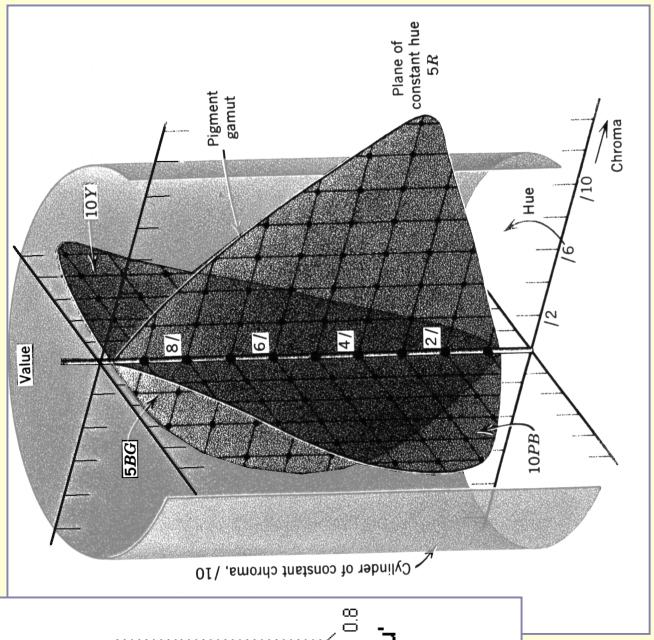

after Judd & Wyszecki

3-D gamuts in pictures . . .

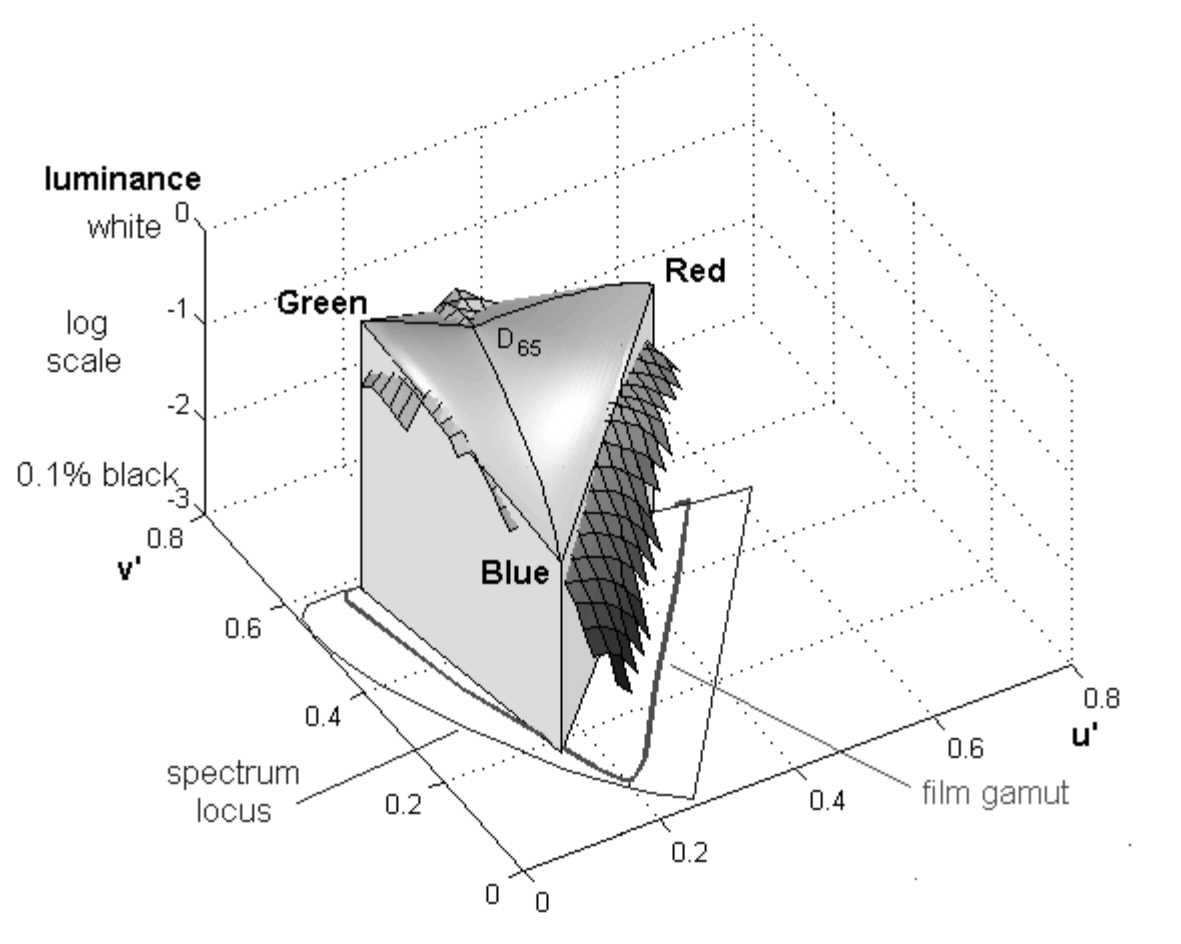

. . . and in words

- Most of the gamut is shared
- Film better:
  - dense colors
  - CMY
  - deep shadows
- D-cinema better:
  - light colors
  - RGB

Tonal resolution

- The issue is *contouring*
- Scale-sensitive

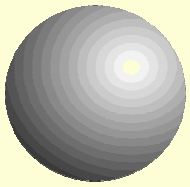

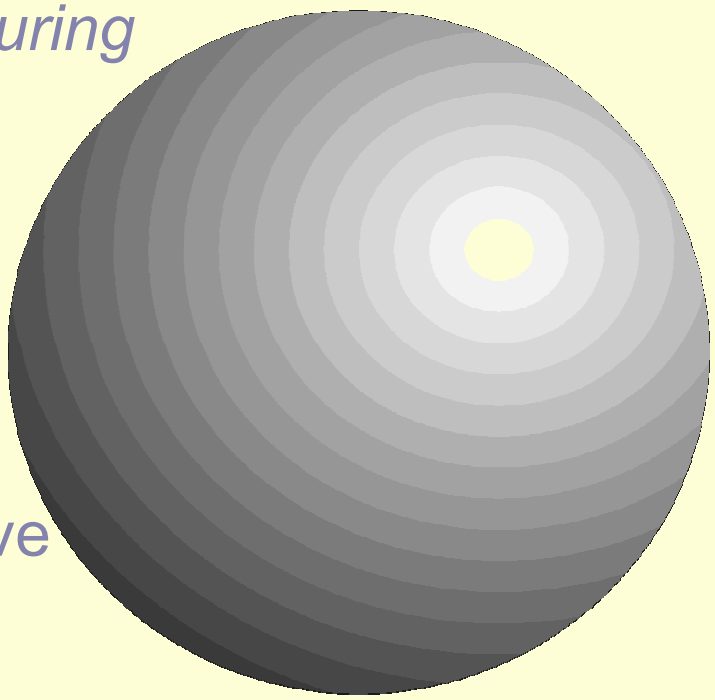

- Brightness-sensitive
- (chromaticity, too!)

Contrast sensitivity

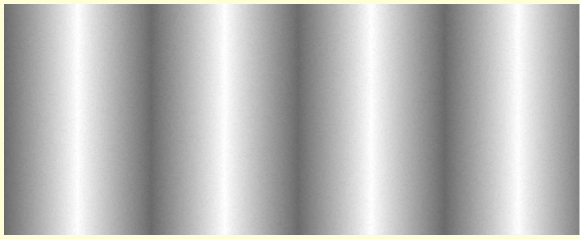

$S(u) = 2\, L_{avg} / \Delta L_{pk\text{-}pk}$ *(u = sinusoidal spatial frequency)*

$$S(u) = \frac{M_{opt}(u)}{k \cdot \sqrt{\frac{2}{T} \cdot \left\{ \frac{1}{X_o^2} + \frac{1}{X_{max}^2} + \frac{u^2}{N_{max}^2} \right\} \cdot \left[\frac{1}{\eta \cdot p \cdot E} + \frac{\Phi_o}{1 - exp\left[-\left(\frac{u}{u_o} \right)^2 \right]} \right]}}$$

Barten 1999
ISBN 0-8194-3496-5

- Modelled on shot noise & gain noise
- Important factors:

  Brightness, grating pitch, size of pattern, color

A Review of Photographic Color Science, and its Application to Digital Cinema Exhibition

Chuck Harrison, Far Field Associates

20 Feb 2000

1. Introduction

"Digital Cinema" plans to use digital projection techniques to bring to an audience an experience similar to conventional movie projection. Clearly, the image a digital projector puts up on the screen must closely match a conventional film projector in color and brightness. To effectively achieve this goal we must understand the color produced by photographic film.

Cinematographers, color timers, and film-laboratory personnel are able to manipulate color with exquisite finesse; however they have little need to understand color science, and few do. In video practice, knowledge of colorimetry is somewhat more broadly distributed because designers of cameras, receivers, video processors, mixers, etc. must make their way through the NTSC matrix definitions in order to do their jobs. In the film world, there is little need for anyone outside the emulsion design department of a film manufacturer to pay attention to the underlying science. As we face the concatenation of additive digital projection onto a photographic production process, it is suddenly necessary to bring photographic color science into the mainstream.

In this paper we assume the reader is familiar with the CIE tristimulus colorimetry system. We concisely review color densitometry and explain the "unstable primaries" and color gamut volume of photographic film. Generally we ignore reflective materials (e.g. paper prints) and concentrate on transmissive photo materials. We then discuss how the properties of film impact on the design of projection systems, and digital image representations which are, at least theoretically, capable of reproducing the film images that cinematographers and moviegoers know so well.

2. Color Densitometry

"Color Densitometry" refers to the art and science of quantitatively describing the color of a photographic image. The field is often subdivided as follows:
(1) Spectrodensitometry
(2) Integral densitometry
(3) Analytical densitometry.

First off, we note that photographic technologists treat "density", rather than transmission or luminance, as the principal feature of interest. Density is defined as $\log_{10}(1/T)$, where T is the optical transmission. Thus a 10% gray patch has a density of 1.0D. Why this attachment to a logarithmic representation? First, because photography produces *dye* images, and dye images are (approximately) *linear* in density: for example, if we use twice as much dye as we had in the 10% gray patch, we get a 1% gray patch with a density of 2.0D -- twice as high. Second, many photographic processes are approximately *power-law* in transmission response, which makes them *linear* in density response. Third, the human eye's brightness response is somewhat logarithmic: a gray scale made up of equal density steps appears much more even than one made with equal transmission steps, for example. Thus, density notation is simply *convenient* in many ways; it doesn't make any statement about the underlying physics, which are often presented in linear-transmission notation anyway.

2.1 Spectrodensitometry

Spectrodensitometry measures the density of a patch of film at a large number of specific wavelengths. For example, Figure 1 shows spectral densitometric data for a neutral gray patch and two red patches of cine print film. Such curves are easily translated into transmission plots simply by changing the vertical scale, as in Figure 2.

We noted above that "dye images are (approximately) *linear* in density". This approximation is very good on a wavelength-by-wavelength basis in spectrodensitometry. In Figure 1 note that the “medium red” curve is precisely 1/3 the level of the “deep red” curve at every wavelength. In analytical chemistry this phenomenon is known as "Beer's Law".

A spectral plot of density versus wavelength holds a lot of data, typically 30 to 300 wavelength points; usually this is more than we wish to work with. It is desirable to distill this information down to three numbers: red, green, and blue "primary" densities.

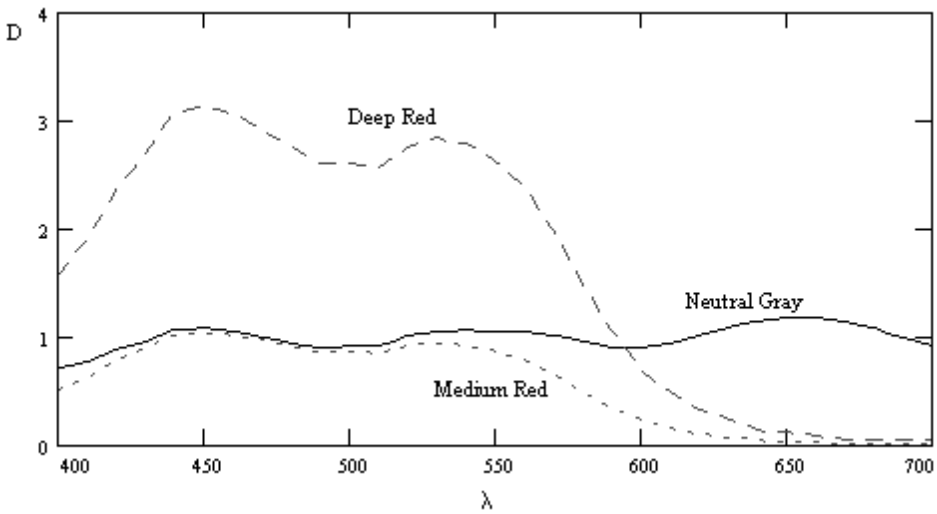

Figure 1

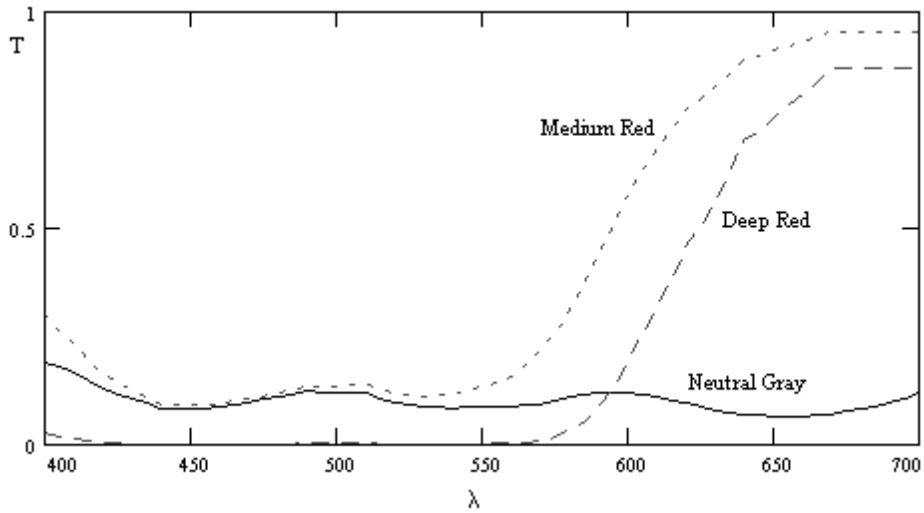

Figure 2

2.2 Integral Densitometry

Integral Densitometry is a technique for measuring three "primary" densities in a *functional* way. By "functional", we mean in a way that answers "how does the film perform in use?". Thus we might be interested in how a projection slide looks to a human observer when illuminated by tungsten light. Or we

might be interested in how a dupe negative image affects the three layers of a print stock when normally exposed. To measure the "red" density of a patch, we might simply use a red-filtered lamp, and measure the fraction of light blocked when the film is inserted into the light path. In fact, this is how most laboratory color densitometers work. The phrase "integral densitometry" comes from the fact that the instrument response is defined by an integral, for example the red density D_{red} is given by:

$$D_{red} = \log_{10}\left[\frac{\int I(\lambda)\cdot F_{red}(\lambda)\cdot S(\lambda)\, d\lambda}{\int I(\lambda)\cdot F_{red}(\lambda)\cdot S(\lambda)\cdot T(\lambda)\, d\lambda}\right]$$

where
$I(\lambda)$ is the spectral intensity of the illumination source,
$F_{red}(\lambda)$ is the spectral transmission of the red filter,
$S(\lambda)$ is the spectral responsivity of the instrument's detector, and
$T(\lambda)$ is the spectral transmission of the film sample.

The behavior of the color densitometer is defined by $I(\lambda)*F(\lambda)*S(\lambda)$; this so-called *spectral product* is the subject of various standards. For example, ISO 5-3 defines "Status A" and "Status M" spectral products primarily for measuring direct-view and negative films respectively. SMPTE RP-180 defines spectral products intended to accurately reflect the response of motion picture print films and printers in current use.

One might expect that color densitometry for direct-view materials (e.g. motion picture prints) would be precisely based on CIE colorimetry and an assumed illumination source. However, this is not the case with "Status A" measurements which are currently used for this purpose. This is because of two additional factors which have outweighed colorimetric accuracy in practice:
(1) The spectral products must be implementable: it must be possible to manufacture devices using available stable light sources, filter materials, and detectors which will retain calibration for extended periods;
(2) The measured densities must be useful for process control (e.g. adjusting printing exposures and bath chemistry) in a direct way.
Integral densitometry has evolved as a reliable tool for process control. It is not directly applicable to the problem of matching film and digital projection.

It is important to note that integral densities do *not* obey Beer's law in the way that narrow-band (spectral) densities do.

2.3 Analytical Densitometry

Spectrodensitometry and Integral Densitometry, described above, are methods for describing the empirical behavior of a piece of film. In contrast, *Analytical Densitometry* is based on a *model* of the film as three dye layers, and it attempts to quantify the amount of each dye which is present. As a model-based method, analytical densitometry is a powerful tool for predicting film color behavior beyond the specific conditions which are measured. Likewise, its weakness is that its accuracy is limited by the validity of the underlying model (a set of three dyes following Beer's law).

Since analytical densities describe dye concentrations, it is appropriate to label them by the dye colors (cyan, magenta, yellow) rather than the corresponding absorption peaks (red, green, blue); however both conventions are used. The numerical units used for analytical density require some explanation. A physical

chemist might describe the cyan dye concentration in mg/sq cm; however this is not helpful to a color scientist. An optical physicist might describe the concentration by the optical density (i.e. $\log_{10}(1/T)$) at the wavelength of the dye's peak absorption. By Beer's law, this density is proportional to the chemical concentration; however this measurement still does not depend enough on the *shape* of the dye's absorption curve to help much with evaluating color reproduction.

The unit which is actually used in analytical densitometry is called *Equivalent Neutral Density* (*END*). By definition, 1.0 END of cyan dye is the amount of cyan which, when combined with suitable amounts of magenta and yellow, produces a visually neutral transmission of 10% (i.e. it makes a perfect visual match with 10% neutral-density filter). The visual match is evaluated using the CIE "standard observer" color-matching functions $\bar{x}, \bar{y}, \bar{z}$, and a specific illuminant. END values are very appropriate for describing neutral tones of an image: for example a gray patch of density 1.7 will have cyan END 1.7, magenta END 1.7, and yellow END 1.7. Some non-intuitive aspects of END are:
(1) It is *not* linear in dye concentration: doubling the dye concentration usually *less than* doubles the END;
(2) A specific dye's END value depends not only on its absorbance spectrum and concentration, but also on the absorbance spectra of the other two dyes in the dye set;
(3) The END values depend on the assumed illumination (e.g. 5600K xenon) used for evaluating "visual neutral".

An analytical densitometry model consists of
(1) The absorbance spectra (density vs wavelength at some nominal concentration) for each of the three dyes: the *assumed dyes*;
(2) The absorption spectrum of the clear areas of film: the *stain*;
(3) The spectrum of the illumination source.
An approximate analytical densitometry model can often be built from the information on a film manufacturer's data sheet. Precise quantification of the assumed dyes and of the stain require careful laboratory measurements; photofilm manufacturers perform such measurements for their internal use. Within the model, it is possible to compute the complete spectrum of a projected film patch based on just the three analytical densities. It follows that such a model allows us to compute the *gamut,* or range of colors, which can be produced by a film.

2.4 Densitometer Geometry

Color photography forms a projected image primarily by *absorbing* light on its way from the lamphouse to the lens. However, film also scatters light, and this, too, affects the color and brightness of the projected image. To accurately measure the functional color of the film image, a densitometer must mimic the optical geometry of the projector itself. In current projectors, the optical system usually consists of a short-arc xenon lamp, deep ellipsoidal reflector, and projection lens. Both the lamphouse and projection lens have a "speed" of somewhere near f/2.

This geometry (f/2 input and f/2 output) is not common for laboratory densitometers, and does not conform with international standards. As noted above, laboratory densitometry is primarily a process-control tool; in this context, consistency is the overarching goal and the mismatch of measurement geometry to projection geometry is not important.

The analytical densitometry model does not include scattering, which follows a different optical law from absorption. This leads to some minor inaccuracies with conventional (low-scatter) color processes. Special processes (e.g. "bleach bypass", which leaves a substantial amount of granular silver in the finished film) may have much higher scattering levels. In this case, conventional analytical densitometry may not apply.

3. Unstable Primaries

Subtractive color systems such as motion picture print film, can be described as having "unstable primaries". This terminology has nothing whatever to do with long-term stability, fading, or reproducibility.

Rather, it is a way of describing what happens when you try to apply simple color-mixing computations, which *do* work precisely in additive systems, to photographic film. These color-mixing computations are the rules which describe how we get from the *color components* (R,G,B drive levels or C,M,Y dye concentrations) to the *visual color* (X,Y,Z tristimulus values).

3.1 Color Mixing in Film

Additive color devices, such as CRTs, exhibit the everyday property of "fixed primaries": for example, the hue of the red CRT phosphor remains the same at low and high drive levels. This property makes it relatively easy, using simple algebraic operations, to predict the colorimetry of a patch on a CRT faceplate given the RGB drive levels; one must know (1) the calibration curve ("gamma characteristic") relating drive level to luminous output of each phosphor, and (2) the colorimetric (x,y) coordinates of the each phosphor's emission color.

With photographic film, predicting final color given the concentrations of the three dyes is not so simple, and 3x3 matrix operations are not adequate. The necessary computation is clear enough, if tedious. It involves evaluating the transmitted light intensity at a large number (e.g. 50) different wavelengths throughout the visible spectrum:
(1) compute the optical density of each dye layer at each wavelength by multiplying the absorbance spectrum of the assumed dye by its concentration (i.e. apply Beer's law);
(2) compute the total density at each wavelength by adding the three dye densities and the stain density;
(3) convert the density values to transmission values (antilogarithm operation);
(4) compute the transmitted light intensity spectrum by multiplying the intensity of the light source at each wavelength by the transmission value of the film;
(5) convert the intensity spectrum to CIE (X,Y,Z) tristimulus values using the standard observer color-matching functions.

A crude approximation to this computation can be made by assuming virtual R,G,B primaries correlated with the absorption peaks of the cyan, magenta, and yellow dyes. As the concentration of, say, the cyan dye increases from zero to its maximum, the virtual red primary contribution drops from 100% to 0%. This is sometimes called the "1-R, 1-G, 1-B" approximation for a cyan, magenta, yellow dye system. This approximation definitely captures the general idea of a subtractive system: adding cyan dye removes red light, adding magenta dye removes green light, etc. Unfortunately, it does not capture the details correctly. The precise color of "virtual green" which the magenta dye removes depends not only on the concentration of dye which is already present, but also on the amount of other dyes which are in the film. Since the chromaticity of the virtual green primary varies with conditions, rather than being fixed, it can be called an "unstable" primary. This phrasing is just a restatement of the fact that subtractive dye systems do not follow additive color-mixing laws. There is no "shortcut" to avoid the point-by-point spectral computation described above.

3.2 Chromaticity Trajectories

It is helpful to visualize the behavior of a color reproduction system, either additive or subtractive, by plotting primary trajectories on a two-dimensional (x,y) chromaticity diagram. A "primary trajectory" is the set of chromaticities obtained as *just one* of the color components is varied over a range. Typically one starts with a neutral patch (say 10%T, 1.0D), whose chromaticity is at white point, near the center of the chromaticity chart. Increasing one phosphor's drive level (additive CRT) or reducing one dye's concentration (subtractive film) moves the chromaticity towards one of the primaries -- red, green, or blue. Reducing the drive level, or increasing dye concentration, moves the chromaticity away from the primary. In an additive system these trajectories are simple straight lines passing through the white point and heading towards the phosphor primaries.

In a film system, the primary trajectories are *not* straight lines (Figure 3). Their exact location and curvature depends on several factors: (1) the chromaticity of the illuminant (white point), (2) the spectrum of the illuminant (especially if it is spiky, like an HMI or fluorescent lamp), and (3) the density of the neutral patch

used as a starting point. An exagerated example of trajectory curvature is given in Figure 4, for film patches containing *only* a hypothetical cyan dye that has incidental absorbtion in the green.

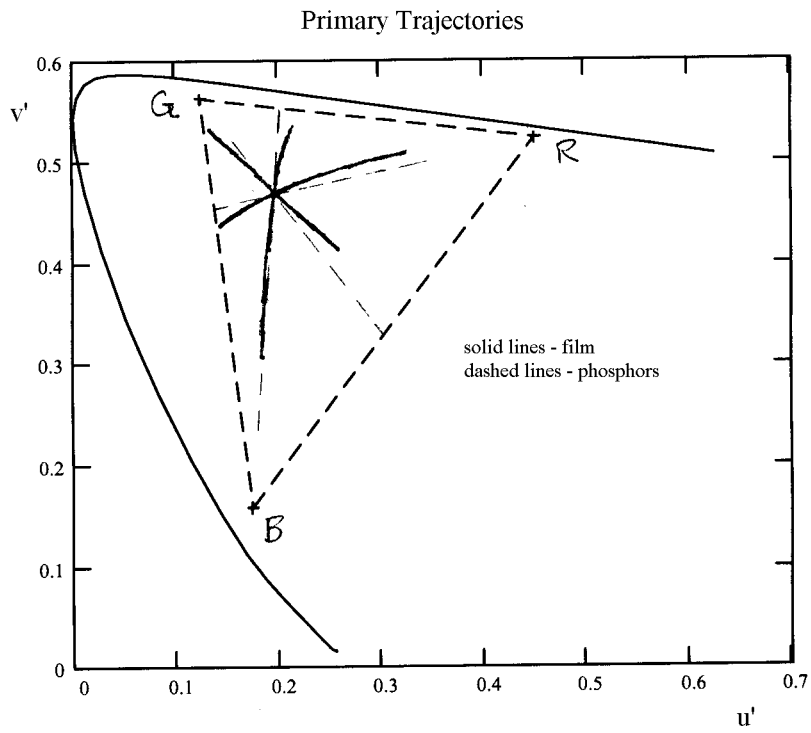

Figure3

When interpreting primary trajectories, a naive observer is likely to assume that the straight lines of an additive system correspond to constant perceived hue: e.g. that increasing the blue drive will shift the apparent color from neutral, through a pastel blue, then to more saturated blue colors, all having the same hue. Then one might assume that the curved trajectories of a film system must represent a perceptible hue shift as dye density decreases. This is a fallacy. For a real human observer, loci of constant hue are *not* straight lines on a chromaticity diagram. These loci can have marked curvature, and their shape depends quite strongly on the viewer's state of adaptation. (Fig. 5.) In some conditions the curved trajectories of photographic primaries may hold more closely to constant hue than the straight lines of an additive projector.

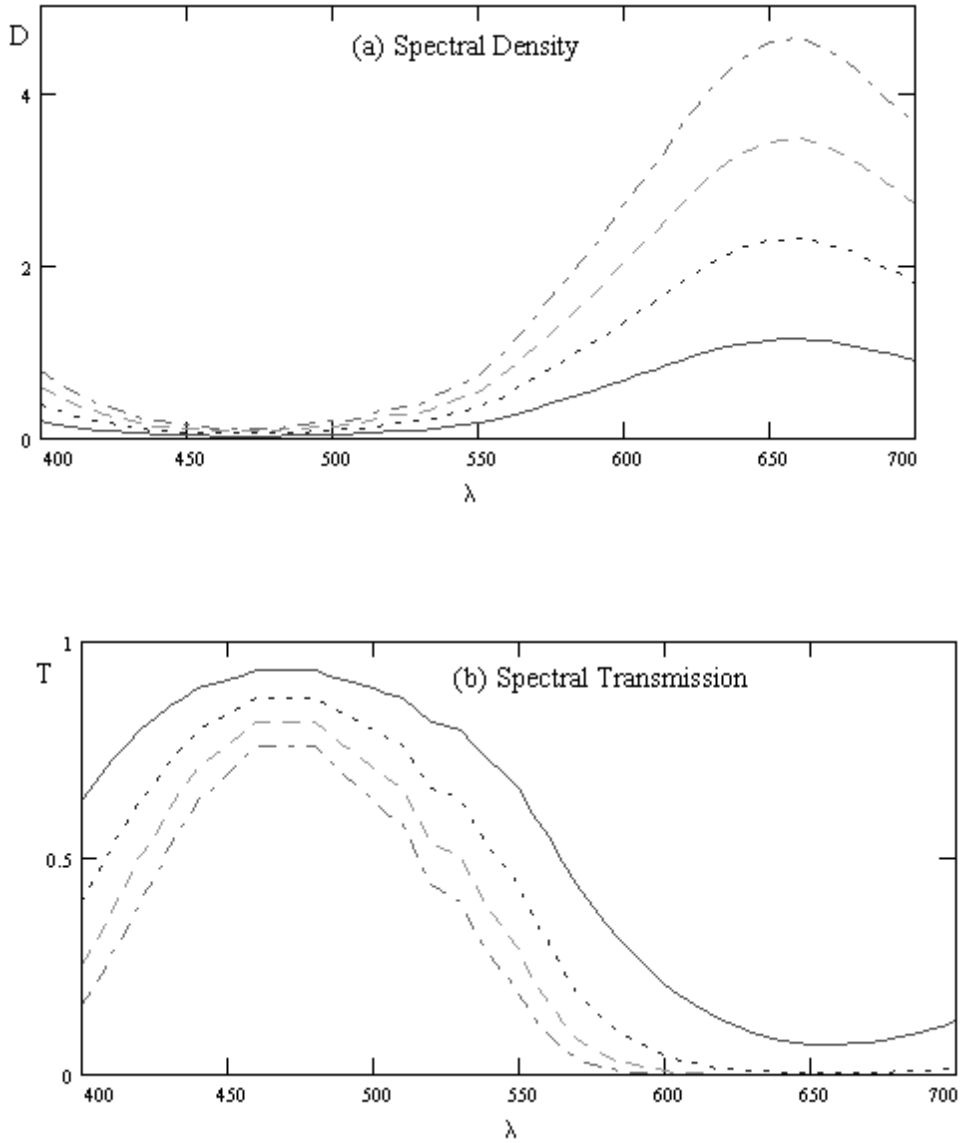

Figure 4. Four densities of cyan dye.

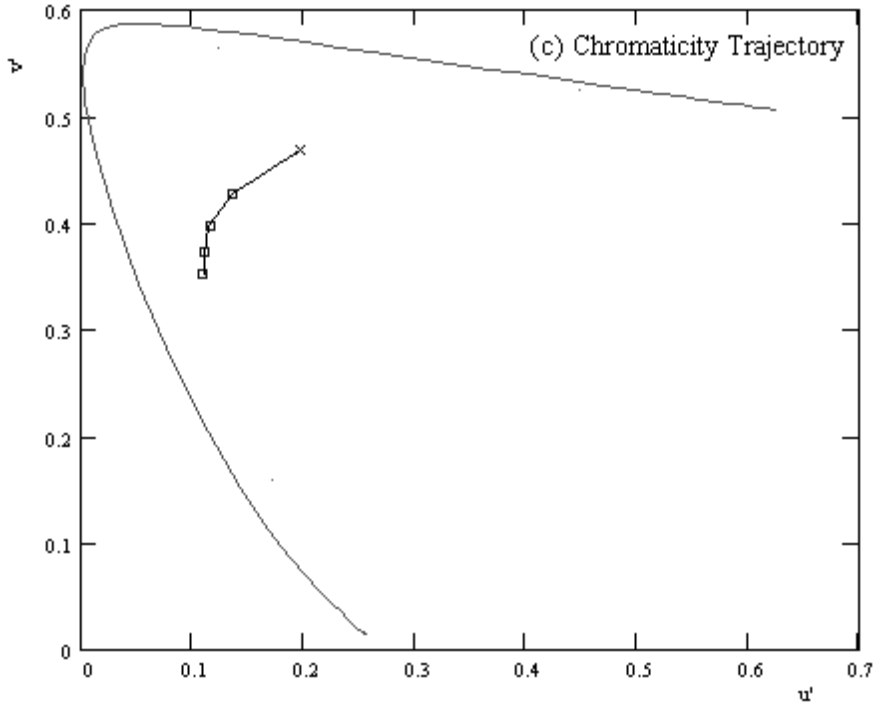

Figure 4 (*cont.*). Four densities of cyan dye.

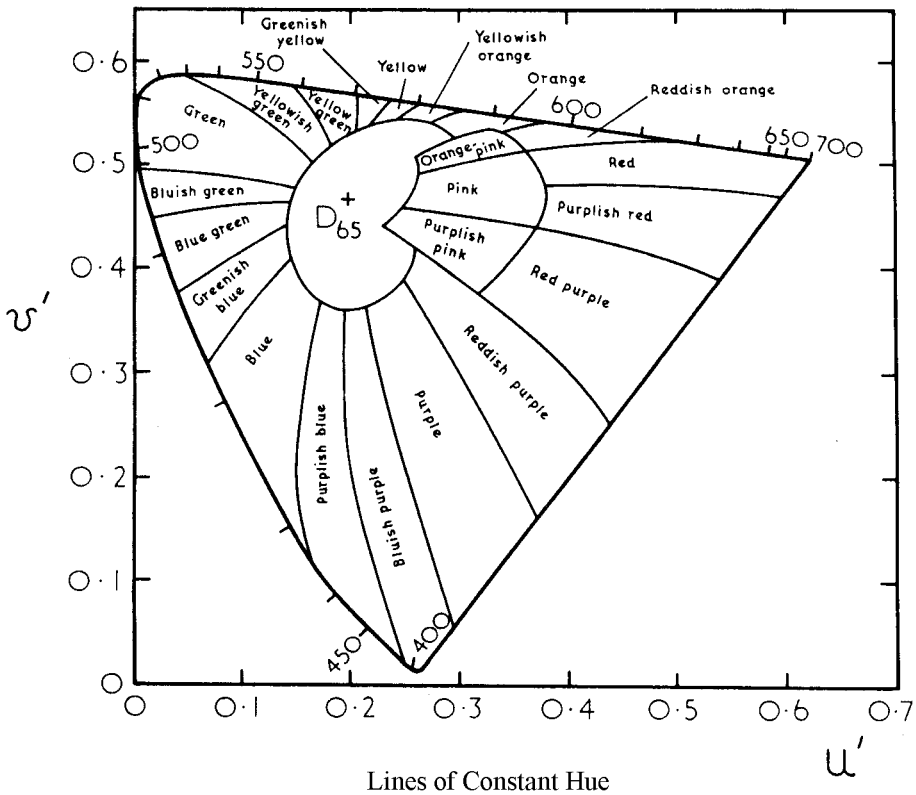

Lines of Constant Hue

colour designations for lights (Kelly, 1943) refers to observation of self-luminous areas seen against a dark background

Reproduced from RWG Hunt, The Reproduction of Colour

Figure 5

4. Color Gamut in Two and Three Dimensions

Every color reproduction system has inherent limits to the range of colors it can reproduce: its *gamut.* No present-day device can faithfully duplicate our most colorful visual experiences (e.g. the "rainbow" reflected off the surface of a CD); this would require the ability to produce narrowband light at any point in the visual spectrum.

4.1 Two-Dimensional Chromaticity Gamut

A device's color gamut can be plotted on two-dimensional CIE chromaticity coordinates (x,y), showing the range of accessible hue and saturation. It is usually preferable to place such plots in "uniform chromaticity space" (UCS)coordinates (u',v'), rather than (x,y), as distances on this chart correspond more closely with apparent visual color differences. [However, one should not take the "uniform" label too seriously: (1) there are known to be scale factor differences of 3:1 over the area of the (u',v') gamut; (2) human sensitivity to color differences is highly dependent on state of adaptation; and (3) visual estimates of larger color differences do not correlate perfectly with sums of "just-noticeable" differences. UCS is an imperfect tool.] 3-color additive color systems have a simple triangular gamut plot (Figure 6). In UCS coordinates, an additive mix of any two colors falls along the line between them; thus the triangular gamut clearly encompasses all possible mixtures of the three primaries. Film's color range can also be plotted on UCS coordinates; its gamut is roughly triangular but has a somewhat irregular outline. Some authors, seeking a

"figure of merit" to compare two color systems, compute the area covered by the gamut in two-dimensional UCS coordinates. This is of dubious value (especially when cross-comparing additive and subtractive color systems) because it ignores the third dimension - brightness.

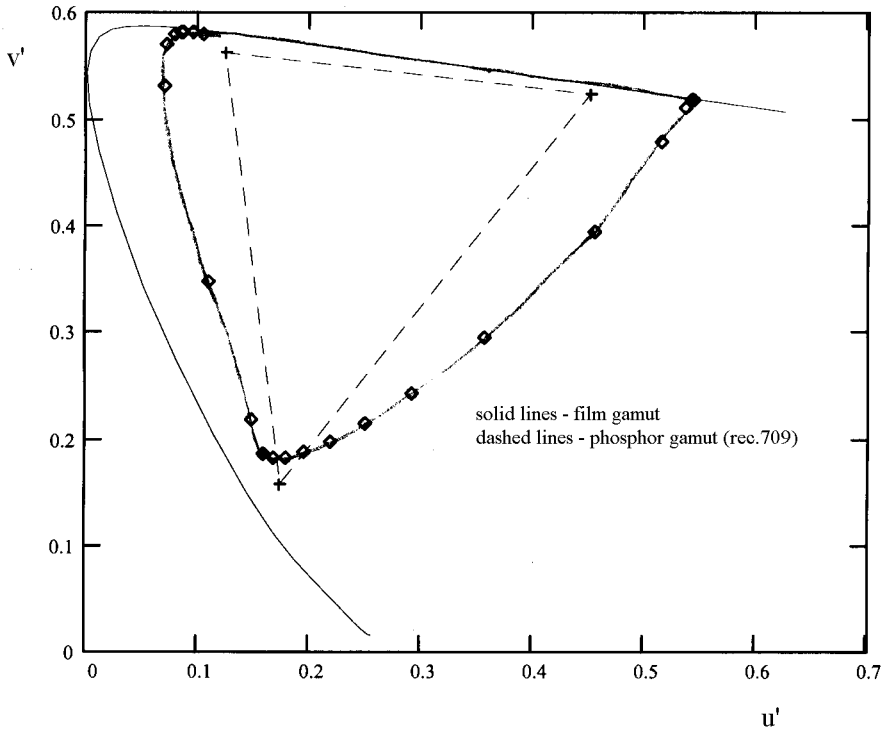

Figure 6.

4.2 Three-Dimensional Color Gamut

No color reproduction system achieves maximum brightness and maximum saturation at the same time. The purest (most saturated) colors must be dim. To describe this behavior, it is useful to plot several gamut outlines on the same chart, for different levels of brightness (luminosity). See Figure 7. For a 3-color additive system, the usual triangle applies for all luminosities below a certain point. As the desired luminosity increases, the blue corner of the triangle becomes "clipped", then the red, and finally the green. At luminosities quite close to peak white, the chromaticity gamut is a small triangle, in a rotated orentation. With film, the change of gamut with luminosity is more continuous: it starts with a quite large gamut accessible at very low light level, shrinking simultaneously on all sides to a small region surrounding peak white (Figure 8). We can visualize the complete gamut of these systems as a three-dimensional solid (the *gamut solid*) with coordinates (u',v') in the horizontal plane and with the vertical coordinate being luminosity Y. In this visualization the nested gamut curves we have drawn in Figure 8 are "topographical contour lines" or "level curves" of the color solid. This color solid is shaped somewhat like a mountain peak, with a broad base.

The shape of this (u',v',Y) gamut solid is misleading. The broad base corresponds to colorimetrically-pure but dim colors which do not appear "colorful" to an observer and which are not very distinct from one another. Another generation of "uniform color space" is required to recognize this situation, and the CIE L*u*v* and L*a*b* systems are popular. In these systems, "lightness" L* is scaled from Y (using a perceptually-based "gamma"); the adjusted chromaticity (u*,v*) represents a *deviation* from white which has been multiplied by L*. This multiplication makes the bottom (dark) end of a gamut solid come to a conical point; gamut solids in CIELUV or CIELAB spaces have pointed "light" and "dark" ends, and an

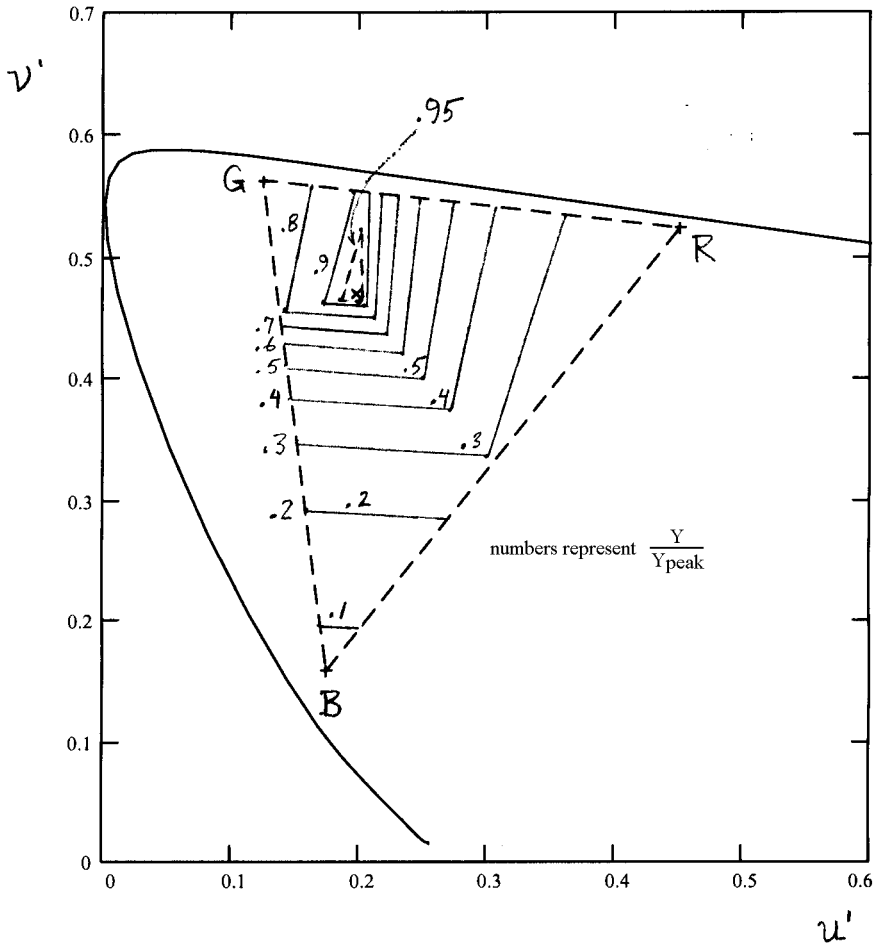

Figure 7. Variation of CRT gamut with brightness level.

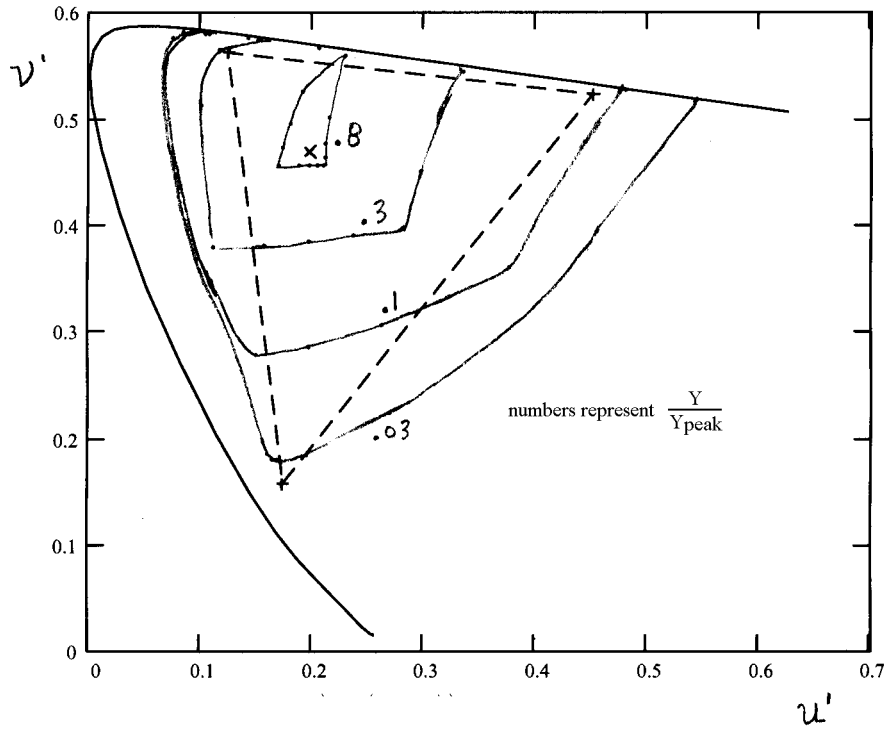

Figure 8. Variation of film gamut with brightness level.

asymmetric "belly" in the middle. The geometric volume of the gamut solid (which can be expressed in cubic-deltaE units) is a rough measure of the number of discernible colors which the process can reproduce. As such, it is sometimes useful as a figure of merit.

The CIELAB and CIELUV systems were developed to assist with important commercial problems in color-matching of well-lit surface colors (e.g. paints and textiles). Some caution is advisable when evaluating color imaging systems using these tools. First, as noted before, the "uniformity" implied in these systems is quite crude. Second, they only attempt to represent "discernible" differences, and cannot represent the "importance" of certain colors (blue sky, skin tones) or their value to artistic composition. Third, they assume a well-defined "white point" to which the observer is adapted. The adaptation white point is pertinent to reflection photographs and television, which are viewed with ambient illumination. Theatrical presentation ("dark surround") does *not* provide a stable adaptation condition. For example, with appropriate lighting and composition a cinematographer routinely makes objects appear to sparkle, glare, or glow -- obtaining "brighter than white" appearance. It is this author's opinion that, when "reference white" is taken as "clear film base", the CIELUV gamut-volume method grossly undervalues the low-brightness region of the color space. Cinema can provide good tonal rendition in an overall-dark scene and dramatic dark-to-light sequential contrast; this depends upon a darkened presentation venue, the wide dynamic range of film, and human visual adaptation. System designers relying too heavily on CIELUV-type models may throw away this powerful portion of the cinematographer's palette.

5. Electronic Projection Systems to Emulate Film

It is clear that one benchmark for Digital Cinema is exact emulation of all the on-screen colors produced by a good photographic print. Technically, this requires that the (u',v',Y) gamut solid of the digital projector must completely surround the gamut solid of film. This benchmark has not yet been reached, and is in fact rather difficult using any of the well-known 3-color additive technologies: DMD, light-valve, or LCD.

The challenges of matching film's color performance are twofold: chromaticity range and contrast range. As discussed in section 4.2, the (u',v',Y) gamut of an additive system resembles a triangular prism with a faceted point "whittled" at the top. The film gamut more resembles a volcanic mountain with a broad footprint, tapering up to a peak. It is difficult for the 3-primary additive triangle to encompass the entire footprint of the film gamut, even if most of the "mountain" is enclosed. Most practical digital projectors use a broadband xenon arc light source, and derive the additive primaries by subdividing the spectrum. To get a large additive gamut, the three primary chromaticities would need to be near the spectrum locus, i.e. narrowband; this leads to inefficient use of the light source and is usually avoided. Potential technical solutions for ultimately achieving full chromaticity range include efficient narrow-band primaries (lasers) and deploying 4 or more primaries.

Contrast range is at least as important for a "film-like" exhibition as chromaticity range. Loosely, contrast is the ratio between the luminance of the brightest white and darkest black that can be presented on the projection screen. The contrast performance that a projector delivers depends on the composition of the scene. For example:
(1) A generally bright scene must still hold reasonable detail in small shadow regions. This calls for low flare light and good *simultaneous contrast.*
(2) A cut to black should be dead black. This calls for good *sequential contrast.*
(3) A generally dark scene, to which the viewers adapt, must retain tonal detail. This requires adequate tonal resolution (bit depth). The dark scene may contain localized "glare" highlights, well above adaptation level.
(4) A slow "fade to black" must be smooth and total. This stresses tonal resolution and sequential contrast.
Motion picture film in a well-darkened theater can achieve sequential contrast of 10000:1 or better. *Simultaneous* contrast, on the other hand, is often limited by audience bounce light and projector lens flare to much more modest values.

Different models of digital projector will vary in gamut and contrast range, but no projector yet fully meets the performance of film. Inevitably, some subtleties are lost in transferring a film to digital projection. To

obtain a visually acceptable "film emulation" under these circumstances will require scene-by-scene corrections to the program.

6. Digital Image Representation

In Digital Cinema, the moving images are encoded as a digital data stream. The format of this data must be defined to accurately represent photographic images, and to interface efficiently to projectors. An industry-standard format is essential, in order to facilitate interoperability and color consistency over the duration of the Digital Cinema age. There are several desirable, if conflicting, features of the standard data format:
(1) A single digital version of the program should provide consistent color and contrast in every theater;
(2) The digital projection should match the film original as closely as possible;
(3) The digital program should be bandwidth-efficient;
(4) The format should be easily extensible to accommodate future technical improvements.

The most difficult problem is how to achieve consistency --goal (1) -- with varying models of projectors in the field. One approach is to precisely define "standardized projector performance" based on the technology of mid-2000, and insist that all projectors operate with the defined gamut and contrast ratio until further notice. This seems shortsighted, and saddles the industry with a "least common denominator" performance which may be consistent, but which isn't very good and doesn't match film very well. It seems preferable to interpret the "consistent color" goal in a more flexible way that permits improvements over the years.

The second goal -- film-like performance -- would seem to require that the digital program must encode the actual color of the film original. For this we might look to present-day "best practice" in feature film work: 10-bit RGB "Cineon" format. But if the digital version of the program represents the full film gamut, it may be necessary to encode corrections required by today's projectors as an added layer. As projector technologies evolve, they will be better able to match the film-based "director's intent".

The third goal -- bandwidth efficiency -- creates a potential conflict with both "film-like performance" (goal 2) and "extensibility" (goal 4). These two goals both suggest that the format should be designed to carry *more* information than today's projectors are able to reproduce. It is also true that many image formats which do well at representing wide gamuts (e.g. CIELUV) do *not* make efficient use of data bits.

6.1 Color Conversion Profiles - Perfect Reproduction

We have reviewed the science of photographic color reproduction in some detail. It is apparent that accurately translating a photographic original into an equivalent additive color image is not trivial. If we were starting from an answer print, it would be technically plausible to proceed as follows:
(1) scan the original print with three narrowband (R,G,B) densitometry channels
(2) use a 3x3 matrix calculation to compute the three dye concentrations at each pixel
(3) use the assumed dye spectra to compute the absorption spectrum of the film at about 50 wavelengths
(4) multiply the absorption spectrum by the spectrum of a nominal illumination source
(5) use standard CIE computations to determine (X,Y,Z) for each pixel
[Note: At this point we have a *device-independent* representation of the projected film color.]
(6) use a 3x3 matrix calculation to compute the drive levels for the three additive primaries.
Needless to say, this is an impractical amount of computation to perform at video rates. It is more reasonable to embed the computations in a precalculated lookup table. This "lookup table" needs to map an (r,g,b) scanner triple into an (r,g,b) projector triple: a 3-dimensional problem. 3-D color lookup is a well-known process (for example, see the ICC printer device profile specification). 3-D lookup tables can require a large amount of memory, and it is challenging to perform 3-D table interpolation at video rates.

The situation is actually more complicated than described above, because the transfer usually begins with a negative or interpositive rather than a print. To predict the projected film color, we must accurately model the behavior of the photographic printing process(es) which intervene between the scanned source material (OCN or IP) and the answer print. Then the film color can be computed in device-independent form and

mapped to the electronic projector as before. All this is a daunting task; however the results of modeling may still be expressed as a 3-dimensional lookup table in a standardized "ICC device link profile" format.

Thus far, we have described what is necessary to digitally represent the "pure" film image such that an additive projector can accurately duplicate it, *assuming* that it is within its capabilities to do so. Unfortunately this is not usually the case.

6.2 Color Conversion Profiles - Gamut Adaptation (or Compression)

With today's projection technology, the system designer faces the situation that the electronic projector's gamut solid does *not* enclose film's gamut solid; in fact the projector's gamut may be much smaller than film's. *Gamut compression* -- color correction which fits the original film image into the projector's capabilities -- is then required. This process is similar to the color manipulation that is normally required in a telecine suite when tranferring film to video release. However, a standard video transfer is not suitable for Digital Cinema because (1) the capabilities of the digital projector exceed standard video; and (2) the Digital Cinema release will be viewed in a theatrical (dark-surround) environment.

Most telecine color corrections - both primary and secondary - are applied uniformly over the full image frame. As such, they can be described as 3-D mappings from input (r,g,b) to output (r,g,b). Such adjustments can be encoded as a "correction layer" applied to the "pure" film image data. Different color correction parameters are applied from scene to scene, so a separate 3-D mapping is allocated for each scene. This may be a useful technique in Digital Cinema mastering. Scene-to-scene corrections are normally performed by a skilled colorist, who understands which portions of the scene's composition are most visually important. As long as electronic projectors' gamut remains much smaller than film's, it is unlikely that this process will be fully automated.

We can anticipate that, over time, commercial electronic projectors with improved gamuts will appear. In these projectors we may find that color space can be divided as follows:
(1) colors produced by both the projector and by film (most "ordinary" colors)
(2) colors produced by film but impossible for the projector
(3) colors produced by the projector but impossible for film
(4) colors which are impossible for either the projector or film (e.g. laser lines, pitch blackness).
The new projector's gamut will be *different* from film, but not necessarily smaller. It this case it is more appropriate to speak of *gamut adaptation* rather than *gamut compression,* nonetheless the goal (a film-like appearance) and the process (3-D mapping) remain unchanged. The skills of a colorist are still pertinent.

In video-transfer practice, it is occasionally necessary to do "fancy" color correction. For example, if the key level changes markedly during a scene, the color-correction parameters may require *temporal dynamics.* Sometimes two regions of the image with similar tonal levels require different treatment - comparable to the "dodging" or "burning in" that still photographers use (*region isolation*). It will be challenging to represent such operations properly during Digital Cinema mastering.

7. Conclusions

We have reviewed the underlying subtractive color science of photographic film. Any digital projection process which starts with scanned film must take into account the nonlinear nature of the medium in order to achieve accurate reproduction. The necessary computations (based on spectral characteristic of dyes and illuminants) are cumbersome, but there is no accurate shortcut. The most practical approach is to precompute a large number of points in the color space and build a 3-dimensional lookup table. This is in accordance with the ICC device profile methodology and it provides an *objective* description of film color. Using these techniques it is possible to digitally represent the true color of projected film imagery as a goal for Digital Cinema projection.

Today's digital projectors cannot accurately reproduce film color over its full gamut solid. (Particular difficulties appear in deep shadows and in saturated blue-greens.) Therefore gamut compression (or gamut adaptation) must be incorporated into Digital Cinema distribution. The ideal way to do this is not immediately apparent. In this paper we have presented some of the conflicting ideals (e.g. consistency, accuracy, extensibility) which must be considered on our way to a useful Digital Cinema standard.

Author Contact:
Chuck Harrison
Far Field Associates, LLC
18815 111th Pl SE
Snohomish, WA 98290
+1 360 863 8340 (voice) 863 8220 (fax)
chuck_harrison@iname.com

Device-Independent Color in SMPTE Standardization Efforts

v.3 9 Nov 00
C. Harrison

1. Overview

Digital representation of image data has become widespread. Digital formats facilitate "open systems", with data easily exhanged between different platforms. In practice, this has resulted in unpredictable results as color images are transferred between legacy systems which were originally developed in a closed, "end-to-end" engineering model.

These problems have stimulated action by the computer-based prepress industry and, more recently, by the online (internet) media sector. Their solutions are based on the concept of *device-independent color*. When an image is represented in a device-independent encoding, its appearance is specified without specifying particular physical input or output devices. For example, this makes it possible for two different printers to meaningfully use two different sets of inks to accurately reproduce the same digital image.

SMPTE's purview extends over two "legacy systems" which have provided reliable color reproduction for many decades: motion pictures and television. For many years, color television stood as the nearest approximation to "device-independent color imaging" and as a model to other industries. However, the practice of digital color representation has advanced dramatically in the last decade, and SMPTE's approach to color encoding is now obsolescent. This has become particularly apparent in recent attempts to prepare standards for Digital Cinema.

Reliable digital interoperation of motion picture, television, and other media can only be forseen if all participants adopt the device-independent methodology. Other industries are making great strides in this direction; motion picture and television practicioners must do likewise.

2. Color Television

The color television standards defined by the NTSC in 1953 were a masterful effort to apply the science of color to end-to-end image reproduction. In this paper we are not concerned with the subcarrier modulation techniques and bandwidth limitations which give rise to many modern complaints. We are, rather, concerned with the transfer functions and encoding techniques which permit three scalar signals (either analog or digital) to represent the perceived color of a visual field. The techniques defined by NTSC, and carried over into most subsequent color television systems, are based on sound colorimetry; they can, in principle, achieve accurate reproduction. However, it is in the definition of "accurate reproduction" that the seeds of trouble are planted.

As expressed by Pritchard [1977], "The overall system philosophy is to linearly reproduce the light input to the camera at the output of the kinescope," and "The original assumption, for colorimetric reproduction fidelity reasons, was that the camera filters and reproducer light-emitting characteristics were essentially 'matched'." In other words, the designers aimed to reproduce the colorimetry of the input scene precisely at the CRT display, and the system specification allowed this to be done. In this model, there is only one "scene", and its color is the same at the camera and receiver.

From the outset, it was clear that this model does not apply to the vast majority of television applications. The three most important deviations are:

(1) The contrast at the display is higher than at the original scene. An "overall system gamma" of about 1.3 is found to give a good visual appearance in typical dim-surround home viewing conditions.

(2) The effect of the original scene illuminant color must be removed in order to mimic the human observer's "color constancy" effect. This is called "white balance" at the camera.
(3) The camera sensitivity curves which would provide colorimetrically perfect reproduction are impractical to implement and would cause artifacts when scene colors outside the CRT gamut are presented.

The history of television since 1953 includes a continuing effort by both camera designers and receiver designers to improve color appearance in the context of a closed, "end-to-end" system. These efforts have been highly effective, but have taken television practice farther and farther from the early philosophy of colorimetrically accurate reproduction. As a result, the "input color space" and "output color space" are now, quite intentionally, significantly different. Yet a great deal of television terminology and standardization remains based on the false premise that there is only one color space of concern.

3. Motion Pictures

Motion pictures are a branch of photography, and the colorimetry of color photographs has been extensively studied. The technical constraints of photochemical processes gave rise early on to a recognition that color reproduction is approximate, and that photosystem design must be treated as a complex global optimization problem.

Even more so than television, photography embodied "closed systems", where characteristics of the negative were designed specifically to work with or compensate for characteristics of the print stock. Spectral sensitivities, dye compositions, and complex proprietary transfer functions and interlayer interactions have been designed based on sophisticated modeling of the photographic process.

A very important aspect of the photographic process is that the final print rarely reproduces the full tonal scale of the original scene. Camera negative stocks are designed with broad *latitude*, and can capture a range of scene tones from deep shadow to sparkling highlight. Later in the production process, the cinematographer can choose to make a "dark print" or a "light print", either of which retains only a fraction of the tonal range originally recorded. A substantial range of postproduction color correction is also available.

When digital image processing became practical for feature film work, it was clear that the full range of the camera negative had to be represented, in order to support important creative processes in digital postproduction. Thus it has become commonplace to digitize the red, green, and blue printing density over a range of 2 log units ("Cineon coding"). The resulting file contains a great deal of information which, after a creative decision, will prove to be "whiter than white" or "blacker than black".

SMPTE documents (prominently 268M and RP-180) precisely define a digital representation for images in the postproduction environment, which has proven immensely useful. However, this representation is device-dependent, as its relationship to visible scenes is entirely dependent on the undocumented physical behavior of particular motion picture film stocks.

It is not useful to tie the negative-density representation directly to either an acquisition color space or a display color space. In acquisition, the cinematographer chooses a combination of film stock, lighting, filtration, and push or pull processing to achieve the *negative* that he wants. There is no "standard" relationship between the colorimetry of the scene in front of the camera lens and the density of the negative. On the other hand, there is no "standard" relationship between the negative and the print. Although we could define a nominal "one-light" print characteristic, the fact that printing is a lossy compression operation which always destroys a large amount of shadow and highlight information makes this approach useless in practice.

4. Color Management Concept

Historically, color management practices arose in the prepress industry in an attempt to obtain consistent color among "soft proofs" (CRT displays), instant proofs (e.g. color xerographic prints), and production color lithography. Color management as deployed today must be considered a work in progress; the goal of consistent reproduction across widely different media is a challenge, indeed. Nonetheless color management organisations (e.g., the International Color Consortium, ICC) have made good progress in defining the key problems and proposing methods of approach.

The most important contribution of the color management movement is to elucidate the fact that every device has an inherent color space. A camera's color space is different from a printer's color space. The operation of color reproduction is a *mapping* from a source space to a destination space. The ICC has defined the *device profile* as a standardized way to describe the relationship between numerical values in a device's color space and the actual scene being acquired or displayed. In pure additive RGB systems, such as idealized television, a device profile need only define a primary set, white point, and gamma characteristic. In dye-based subtractive systems, a device profile is likely to be more complicated and nonlinear.

The ICC has also defined the idea of a "profile connection color space", a device-independent representation which describes the color of an image without reference to any input or output device. CIELAB is a widely used connection space. The concept is that, for example, an input device profile relates the CIELAB coordinates of the original scene to the RGB data of a camera; an output device profile relates the CIELAB coordinates of a CRT display to its RGB drive levels; a mathematical operation can combine the two profiles to create colorimetrically perfect reproduction.

5. Applying Color Management

The concept of "perfect colorimetric reproduction" described above immediately faces the same issues that originally destroyed this model for color television. There are refinements to the ICC profile methodology which address (albeit imperfectly) the fact that the desired output color is usually *not* the same as the input color, even when that is technically feasible. The ICC tools partially address the need to express, unambiguously, a precisely defined mapping from one color space into another. Some weaknesses exist in the definition of the profile connection space (dynamic range, perceptual uniformity, absolute colorimetry for non-reflection media) and in the concept of "rendering intent".

It is also the case that, for general scenes, a camera device profile is poorly defined. An unambiguous transformation to a profile connection space is possible only if the camera has front-end spectral sensitivity curves which are linear combinations of the CIE standard observer curves. Otherwise, two colors which appear different to a human observer can encode as the same RGB values from the camera, or vice versa. For various practical reasons, such spectral sensitivity curves are rarely implemented. Certainly the response of existing television cameras and photographic film emulsions makes them incapable of perfect colorimetric acquisition. Thus it must be recognized that an input device *cannot* be characterized by its "primary chromaticities"; only a tabulation of the three spectral responsivity functions is capable of describing its performance. This fact has been well known in photography for over a century, but modern digital image practicioners persist in ignoring it.

Color scanning devices (e.g., drum scanners, telecines, etc.) which digitize photographic images face a simpler task than cameras. While the spectral composition of a region in a real-life scene has an enormous number of degrees of freedom, the transmission of a strip of color film offers only three. Therefore, unlike a camera profile, a scanner profile can be precisely standardized. SMPTE RP-180 would form the foundation for a telecine profile, if other issues (mentioned in section 3 above) are resolved.

The idea of camera "white balance" adjustment is a simple approximation which cannot be completely justified. In television practice, it is often said that adjusting the gains of the RGB channels will make a scene recorded under 3250K studio lights appear "as it would" if illuminated by CIE illuminant D65. In terms of color science, such a statement is technically false. Once the taking-environment's *spectral products* (i.e., the wavelength-by-wavelength products of illuminant and camera sensitivity curves) have been applied, massive amounts of information about the surface reflectances in the scene have been lost. It is impossible to precisely predict the colorimetry of a scene under a different illuminant. (A scene including self-luminous objects – e.g., lamps, windows, or flames – cannot be properly "white balanced" at all.) This basic fact about colorimetry does not prevent "white balancing" from working adequately well in practice. However, it adds to the difficulty of defining an input profile for a real-world camera device. Kodak's "Photo YCC" standard relies on an abstract "reference image capture device" normalized to a particular whitepoint, but it fails to provide a precise definition of "normalizing".

6. Scope of Work for SMPTE

SMPTE standards should allow digitized television and film images to be reliably interchanged and transferred to other media (including the incipient Digital Cinema medium), without ambiguity in color or tonal scale. This demands a transition to device-independent specifications which clearly state what a particular digital RGB or YCrCb triple means in terms of measurable optical parameters at a particular viewable point in the process.

Existing standards from the ICC and others reveal useful models for this work. For television, we should consider whether a simple ICC methodology based on additive camera and receiver color spaces can codify the existing industry practices in a useful way. However, in the motion picture arena, the high standards of art require considerable extensions to the concepts underlying present-day color management systems. The key areas of concern are:

(1) The acquisition process (film photography) is variable and provides an important aspect of creative control. A device-independent characterization of this process would require a complex model with many parameters.
(2) The film intermediate (negative) used for post-production incorporates considerable latitude which is needed for creative control. The relationship between negative density and answer print color is ambiguous until final timing (grading) decisions are made.
(3) Once the scene is timed, there is ideally a unique relationship between density and answer print color. However, this relationship is nonlinear and its precise characteristics are not publicly known.

These three issues relate to preparing meaningful device-independent specifications for feature-quality digital production practices.

An additional range of issues arise when images are transferred from one medium to another. Different media have different available tonal ranges and chromaticity ranges. Equally important, human color perception is deeply affected by the conditions in which the image is viewed. As a result, an image originated on film for theatrical projection requires heavy modification to appear its best on television. Film programs transferred to Digital Cinema form face related, but lesser, difficulties. In color management terminology, this translation process is called "gamut mapping" and is the subject of much research, including a reportership under the CIE. Workable gamut mapping techniques could provide great benefits to the efficiency and quality of image transfers in our industry; however SMPTE activity in this area should be considered a long-term project due to the immature state of the art.

While automatic gamut mapping decisions are considered to be a goal in many prepress applications, this appears to be a distant (or irrelevant) goal for motion picture work. Gamut mapping in feature film production (which already exists under other names) is a creative tool, and the principal benefit of SMPTE standardization would be to improve the precision, repeatability, and interoperability of these creative tools in the digital domain.

Quantitative definitions of gamut mapping, such as might be desired for Recommended Practices in Digital Cinema, must be based on the device-independent color representations discussed earlier. Therefore an effort directed at formalizing device-independent color representations for film and Digital Cinema projectors is an important step towards obtaining improved quality and efficiency.

References

D.H. Pritchard, "U.S. Color Television Fundamentals - A Review", *IEEE Trans. Consum Electron.,* **CE-23**:467-478 (Nov. 1977). Reprinted in
T. Rzeszewski, *Color Television*, IEEE Press, 1983. ISBN 0-87942-168-1.

R.W.G. Hunt, *The Reproduction of Colour, 5e*, Fountain Press, 1995. ISBN 0-86343-381-2.

Color Management for Digital Workflows

Arjun Ramamurthy
Warner Bros.

- **Introduction**
- **Current Practices**
- **ICC Overview**
- **Architecture**
- **Implementation**
- **Results**
- **Benefits**
- **Limitations & Pitfalls**
- **Q/A**

Color Management for Film (CMS). The need for a CMS is discussed as are implementation details and workflow integration. This talk highlights the successes we have had at Warner Bros. Feature Animation, since the development and deployment of a CMS in 1996, and provides a guide on how to implement a CMS at a digital production facility. This presentation has a disclaimer. Various manufacturer's and their products are mentioned. None of these references should be regarded as an endorsement of either the manufacturer or the product.

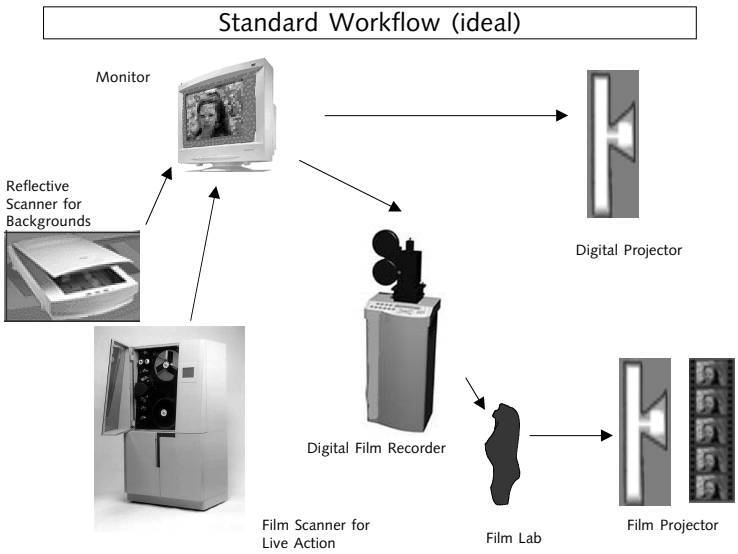

In the **Standard Workflow**, Digital Production Studios (Live-action and Animation) acquire images from a scanning device or generate images directly on the monitor, artistically style and add effects and transfer them to film using a digital film recorder. Either Artwork or film is scanned, combined with digital imagery, color styled, composited and transferred to film. After recording, the film is developed and printed at the processing lab, and subsequently projected. For Digital Cinema, digital files may be directly delivered from the compositing station.

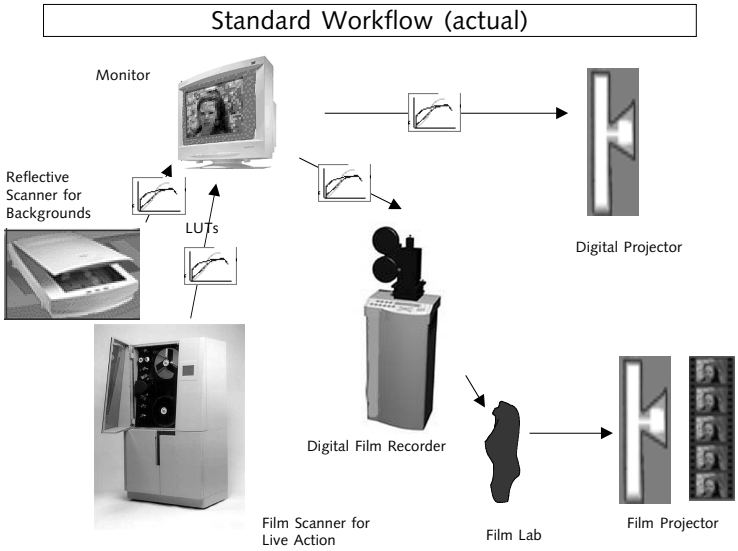

Standard Workflow (actual). Color, Contrast and Saturation of an image are not reproduced very well as images are transferred between devices, and therefore correction mechanisms, most commonly Look Up Tables (LUTs) are employed.

Copyright © 2001-05-10 Warner Bros.

Acquisition Systems for opaque objects (Backgrounds and Overlays) feature a light source with reflective optics and a filter wheel to image individual colors on the sensing elements (CCDs). Numerous other designs exist, such as no filter wheel and tricolor CCDs, area arrays etc.8 or 12-bit output, with calibration on manufacturer supplied D_{MAX} and D_{MIN}. Issues of concern are lamp stability and uniformity. For Film, light is passed directly through the negative, the filter wheel and onto the sensing element. CCDs (linear and array) or Photomultiplier tubes are generally used as the detectors in this case. 12-bit data, with white point calibration on film base (D_{MIN}). Camera negative (or internegative) is scanned and a "generic" 3x3 matrix multiply is carried out to convert to positive.

Display Systems are primarily Cathode Ray Tube (CRT) based relying on phosphors to emit light when struck by electron beam. Important issues to keep in mind are gamma, white point, and phosphor primaries. Other issues are uniformity and convergence.

Acquisition Systems

Reflective Scanner

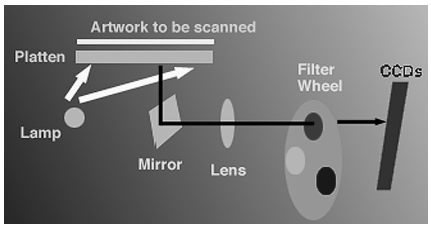

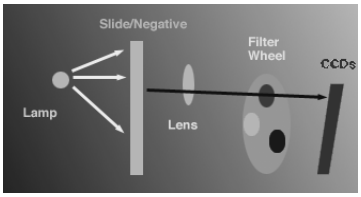

Transmissive Scanner

Display Systems

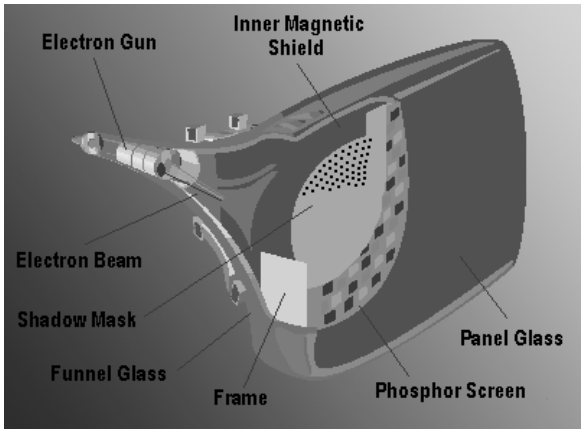

Film Recorders

Cathode Ray Tube Type

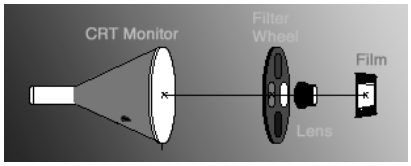

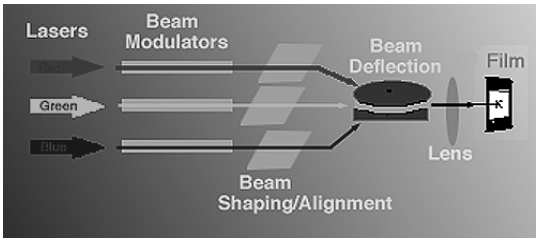

Laser Type

Film Recorders are CRT and Laser based. CRT recorders employ a high resolution monochrome CRT to trace individual channels on a flat screen. This image is passed through a filter wheel and focussed onto the film. Film transport is carried out using a standard camera shuttle made by Oxberry®. Celco® Extreme MPX recorders image a 5 µm spot onto 35 mm film, (max 4k resolution) with recording times of 20 seconds, at 12-bit/channel resolution. Mostly use camera negative (EXR 5245) or B/W (5231) stock. Laser recorders employ RGB lasers that are acousto-optical modulated to convert pixel data to film exposure. Beam shaping optics align and converge the three beams onto the Beam deflector that scans each line onto the film. For flatness across scan lines, the film gate is curved and matched to bean deflector. Arrilaser® specs callout a 6 µm spot size, with 4k resolution, recording times of <6 seconds and 12-bit resolution. We use IP/IN (5244) stock.

Film Media Chain

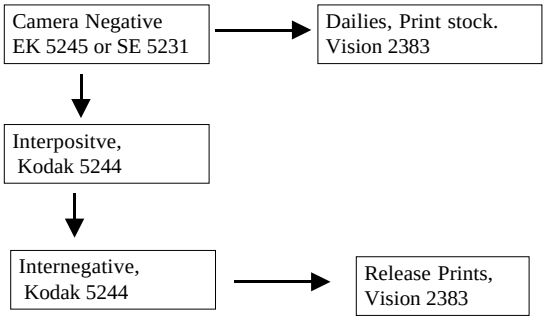

The **Film Media Chain** takes the camera negative to the final release print. The camera negative is printed to produce "dailies" or one-light prints to provide the artists with immediate feedback. Preferred stock is Kodak® Vision 2383. On approval, the negative is "cut" and used to produce a small number of prints on interpositive (IP). (Kodak 5244). Color adjustments are "timed" into this step. These few IP copies are copied again to produce a larger number of internegatives (IN) (Kodak 5244). These are used as the production masters and final release prints are made from this set. Additional color adjustments may be required to ensure proper continuity between sequences at this stage. Spectral response, Dye Sensitivity and Exposure response for each film stock is different, adding to additional image divergence from the original. Response curves may be obtained from http://www.kodak.com. Negative development is carried out by "washing" the film in a series of chemicals that are closely monitored for pH balance, temperature etc. However, daily variations at mid grey (18%) can be in the order of ±0.05. Print variations may be as large as ±0.15 points (or two stops).

So **How how do we deliver** images to each stage of the production and still keep the artists happy? One common practice, for achieving a match between devices is to apply custom Look Up Tables (LUTs) to each channel of the image. Another approach is to modify the hardware to make one device emulate another. Film Printing modeling approaches that represent what the color timer may be asked to do have been attempted. The brute force approach is to manually color correct each frame and if that is too tedious, just let the Color Timer worry about it!

Current Practices

How do we deliver ?

- Manually color correct each sequence.
- Make Monitor Emulate Film.
- Model the film printing process.
- Iteratively apply LUTs until it looks right.
- Let the Color Timer worry about it.

Drawbacks. Asking users to manually color correct each scan is clearly time-consuming, not very repeatable and labor intensive. The approaches that use custom LUTs compromise the overall color range (color gamut) and the tonal range available from either device. In addition to having been created in an iterative fashion, requiring a fair amount of labor, the current schemes suffer from the inherent problems of LUTs, namely, precision loss and data artifacts. They also require operator interaction and are hence non-repeatability and do not handle color gamut compression issues or the differences in the RGB primaries (color space differences). Modifying device behavior reduces the color gamut and can dramatically curtail device life. In some studios, monitors are replaced every 800 hours (6-8 months). Lastly, such processes require operators develop knowledge of the media's behavior making integration of new technologies very slow.

Current Practices (contd.)

Drawbacks

- Do not handle color primary differences.
- Reduces device (monitor) life.
- Labor intensive. Iterative.
- Requires knowledge of media behavior.

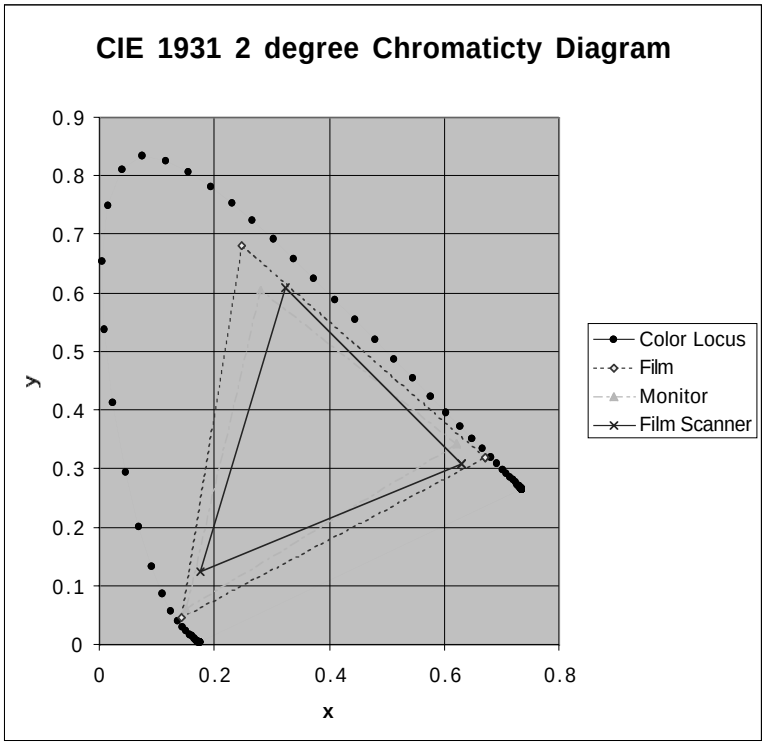

Gamut differences between Film Recorder's output on Film (Kodak® EK-5245) and a Barco® 120T Reference Calibrator Monitor and a Kodak® Genesis Plus scanner (for the same film) are shown on the CIE 1931 2° Observer Diagram. All data is from devices currently operated at WB or vendors. On the CIE diagram, the gamut of a device may be ascertained by connecting the primaries with straight lines. Notice film reproduced colors outside the monitor's gamut and also that the scanner has a "further" red and green primary.

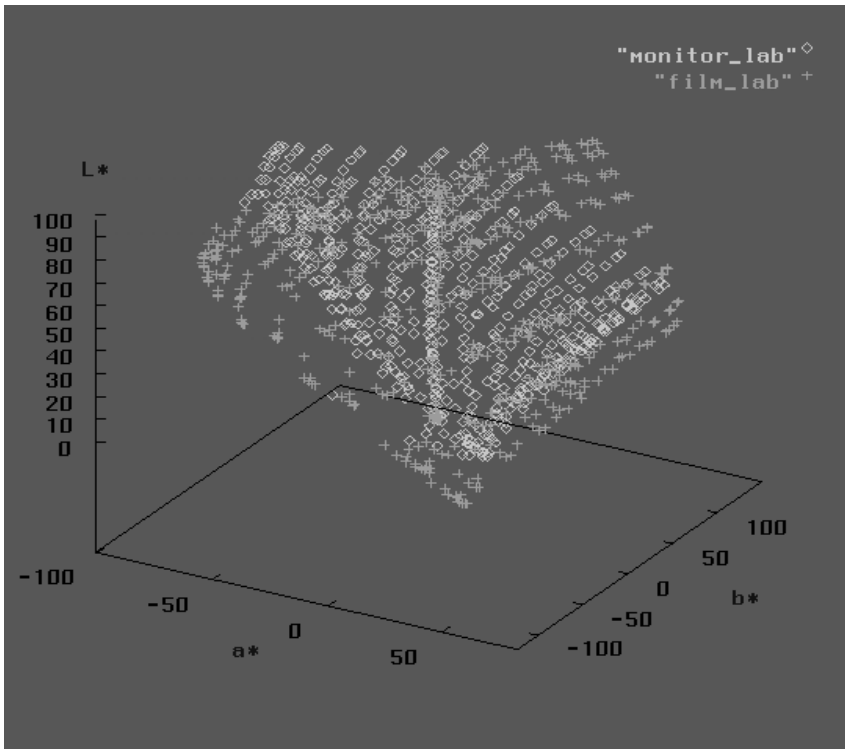

Gamut differences in 3D between Film Recorder's output on Film (Kodak® EK-5245) and a Barco® 120T Reference Calibrator Monitor illustrate why simple 3x3 matrix matrices cannot be employed to rotate and scale one set of primaries onto another. Gamut shapes are different, because though the primaries may be shown as points, linked by straight lines, on the CIE 1931 color space, the true gamut is a 3-dimensional solid and is not necessarily a uniform "extruded" triangle.

A **Color Management Systems** (CMS) should conform to the following **specifications**:

• Color transforms should be carried out in a repeatable and visually convincing fashion.
• Color transforms should be computed in a Device Independent color space (CIE color space) that fully encompasses the entire range of colors perceivable by humans.
• It should rely on instrumentation that is capable of measuring in CIE color space.
• It should not require devices to be altered or adjusted to characterize them, rather be operated as per manufacturer's recommended practices.
• The CMS must operates in a manner that is transparent to the users, with the least amount of intrusion and minimal or no change in workflow.
• The CMS must be capable of handling gamut compression and gamut mapping faithfully.
• The underlying technology should be such that it can handle future technological advances. Color transformations and computations should be carried out with floating point precision with support for 16 bits per color plane.
• The quality of color match must be measurable and quantifiable.

Color Management Systems

Requirements/Specifications

- Color transforms in CIE Space.
- Operate on devices in "ideal" state.
- Minimal Change to workflow.
- Computations should not introduce artifacts.
- Handle Gamut mapping.
- Allow easy integration of new devices.
- Quality of color match must be measurable.

Quantifiable Color match

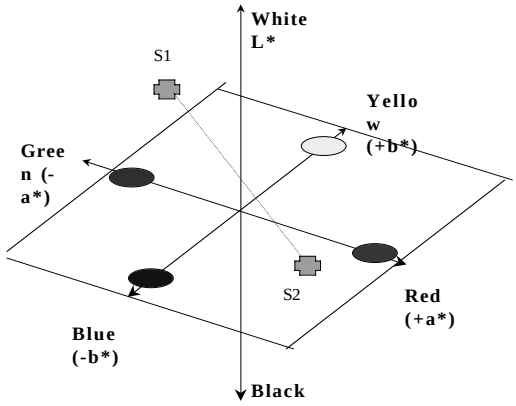

$$\Delta E = \sqrt{(L^*_1 - L^*_2)^2 + (a^*_1 - a^*_2)^2 + (b^*_1 - b^*_2)^2}$$

A **Quantifiable Color Match** can be carried out in a color space that is perceptually uniform (within reason). CIE developed the notion of CIE L*a*b* (or CIELAB) to approximate the *lightness* response of human vision. In the CIE 1976 color space the following defines the transformation from XYZ to LAB:

$L^* = 116 * f^*(Y/Y_n) - 16$
$a^* = 500 * (f(X/X_n) - f(Y/Y_n))$
$b^* = 200 * (f(Y/Y_n) - f(Z/Z_n))$
where: $f(X/X_n) = (X/X_n)^{1/3}$ for $(X/X_n) >= 0.008856$ and
$f(X/X_n) = 7.787 * (X/X_n) + 16/116$ for $(X/X_n) < 0.008856$
The same being true for $f(Y/Y_n)$ and $f(Z/Z_n)$

The values X_n, Y_n, Z_n are the CIE Tristimulus values for the perfect reflecting or transmitting diffuser. At D_{50}, These are: $X_n = 96.422$, $Y_n = 100.00$, $Z_n = 82.521$. Delta E, ΔE, the color difference metric provides a measure of separation between the sample and the target color. One ΔE represents the smallest color difference the average human eye can perceive. However, current research in color science is beginning to indicate that ΔE values of up to three are perceived as being acceptable by most users.

The International Color Consortium

Established to create, promote and encourage a vendor-neutral cross platform CMS.

- Founded in 1993, under FOGRA.
- Current Membership exceeds 60 companies.
- Developed the ICC profile to characterize devices.
- Initial Profile Spec based on Apple Colorsync™ 2.0.
- Current version of Spec is 3.4

The CMS relies on the framework established by the **International Color Consortium** (ICC, http://www.color.org). Formed in 1993, under the auspices of the German Graphic Arts technical association, FOGRA, the ICC's purpose is to create and promote a standarized, open, vendor-neutral, cross-platform color management system architecture. The ICC's proposal called for devices to be individually characterized("profiled"), and then allow a color engine to perform the conversion of image data from one color space to another.

The **components of a color management system** (CMS) are
• the CMS Framework.
• Profiles
• Transforms and
• Color Matching Methods (or Color Management Modules), the color "engines."

Components of a Color Management System

- CMS Framework
  Establishes a correspondence between other parts.
- Profile
  Describes a device's color or Abstract color space.
- Transforms
  Carries out conversions from one color space to another.
- Color Matching Engine

The standard intermediate color space is termed the *Profile Connection Space* or *PCS* in ICC parlance. The color conversion route from the source device color space, through the PCS, to the destination device color space is termed a *Color Transform*. Color data is seldom actually converted into the PCS and back out. Rather, the source and destination profiles are composed by the color management system into a single transform object that performs the source to destination color conversion.

ICC Profiles and Transforms

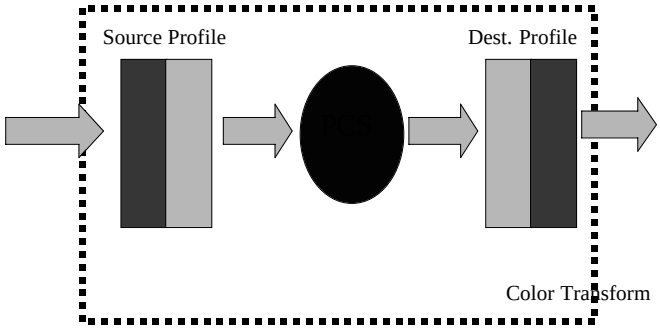

An ICC Profile is a mathematical description of how to convert device coordinates into and out of the profile connection space. An ICC profile's structure is similar to that of a TIFF file: It contains a 128-byte fixed-length header and a tag-based structure consisting of public (required and optional fields) and private tags. For a detailed description of the ICC Profile Format, see the specification document available from the ICC at www.color.org.

ICC Profile File Format

- 128 Byte Fixed-Length Header
- Tag-Based like TIFF
- Public Required Tags
- Public Optional Tags
- Private Tags

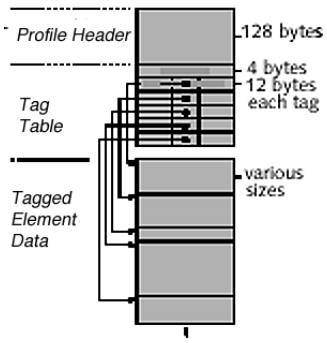

ICC Profile Types

- Input Devices. (Scanners, Digital Cameras etc).
- Display Devices. (Monitors, LCD etc.)
- Output Devices. (Printers, Film recorders)
- Device Link. (Dedicated, Device to Device)
- Color Space Conversions. (CIE XYZ, LAB etc)
- Abstract Profiles (Adjustments in a profile, etc)
- Named Color Profiles. (Pantone®, Trumatch® etc)

The ICC Specification defines seven types of profiles with different required tags. Most are "bidirectional": They contain the necessary processing tables for both forward and backward transforms between the device color space and the PCS. The exceptions are Input profiles, DeviceLink profiles, and Abstract profiles. Input profiles only contain the device-to-PCS pathway. DeviceLink profiles are one-way, device-to-device profiles. Abstract profiles are PCS-to-PCS profiles.

ICC Profile Connection Space

Two Profile Connection Spaces.

- CIEXYZ PCS. (Used for spaces like monitor RGBs)
- CIELAB PCS. (Used for more complex color gamuts)

Both are based on the CIE 1931 Standard Observer, with a D50 Illuminant and relative colorimetry.

Color Measurements should be carried out at 0/45 or 45/0 Reflectance geometry.

There are two Profile Connection spaces defined for ICC profiles. They support the two processing models used in ICC profiles. The XYZ PCS enables simple processing model for spaces like monitor RGBs that are just a lut and 3x3 matrix away from CIEXYZ. CIELAB is a better space for characterizing devices with more complex gamuts. Its (relative) perceptual uniformity allows for efficient population of the large, multidimensional interpolation tables used in printer and film recorder profiles.

Profile Processing Models

Two Profile Processing Models.

- Shaper/Matrix model.

  Used for monitor (i.e. Simple Input and Display profiles)
  Uses the CIE XYZ PCS. Small, Simple and Fast.

- Shaper/MFT model.

  Used for complex RGB and CMYK Input and Output profiles. Also Used for other Profile types.
  Large, Complex and Slower.

Profile Processing Models

Shaper/Matrix Model Processing consists of 2 phases: (A) linearize the tone curves, and (B) compute the transformation to CIE XYZ. Utilizing the individual colorant information, the tone curves are linearized as 1 D “shaper” Look-Up-Tables (gamma tables). These “shaper” tables are denoted as TRCTags in the ICC header. These tags re defined to be of type “curveType”, which contains a 32-bit count and a 1D table of 16-bit entries. A count of 0 is interpreted as a ramp with a slope of 1.0. If the count is 1, the single entry is interpreted as a gamma value, otherwise the input should be used to interpolate an output value from the entries. In the second phase, the 3 TRCs (gamma curves) linearize the raw values with respect to Luminance (Y) and the 3x3 matrix converts these linearized values into CIE XYZ. The embedded media white point tag allows source to D_{50} white point scaling.

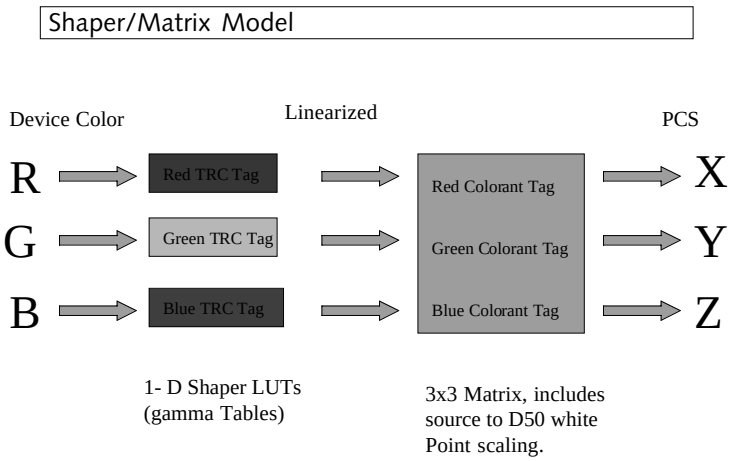

Shaper/MFT processing Model consists of building multi-dimensional Multifunction tables that can contain 8-bit or 16-bit/component data and are of ICC type lut8Type or lut16Type. If the PCS is not XYZ, the 3x3 matrix must be the identity matrix.

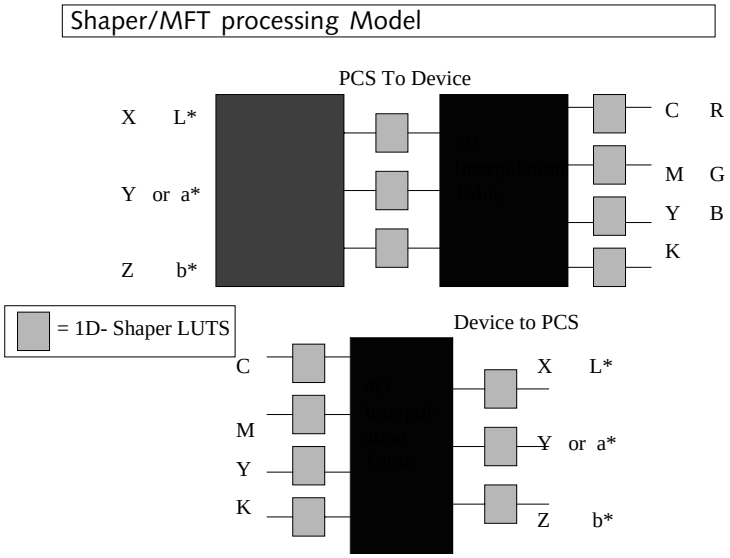

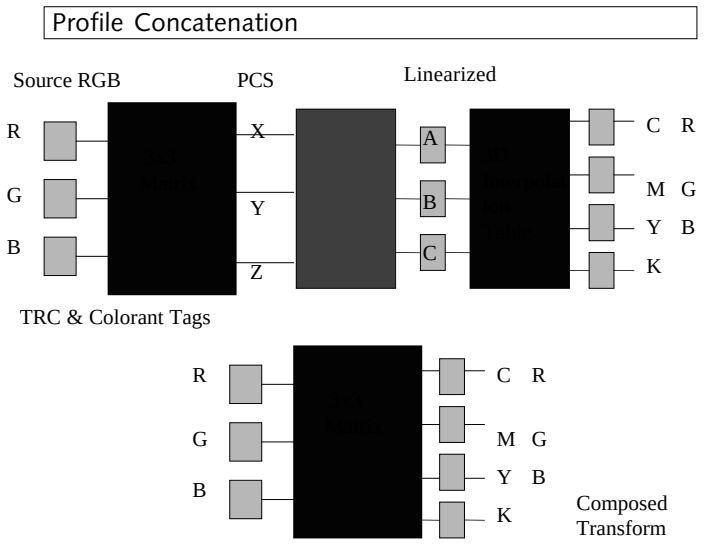

CMMs seldom explicitly convert color data in and out of the PCS during processing. Usually, the tables from the profile sequence are premultiplied and renormalized by the CMM during transform creation. This results in a single interpolation table through which the color data can be processed. This is both faster and more accurate than the series of interpolations that would otherwise be required.

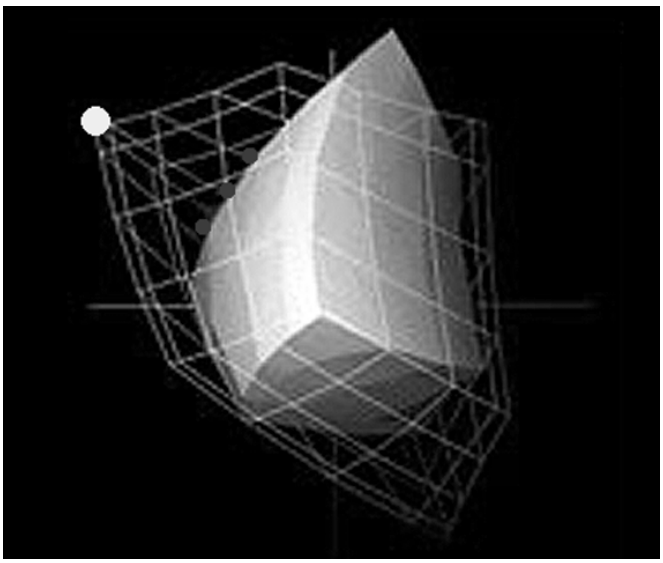

The issue of **Gamut Mapping** relates to how to map the range of colors from one device to another, given the different color gamuts (both in number of colors and range) of different devices. An ICC profile requires three gamut mappings to be provided. Perceptual, Saturation, and Relative Colorimetric. Perceptual is the most common rendering intent, especially used for the photographic reproduction of images. A Saturation intent is provided to get mostly saturated colors, even at the expense of hue and lightness accuracy. This intent is mostly used for business graphics, where the color of a pie chart does not have to be accurate, just colorful. In the Relative colorimetric intent, the intention is to obtain the highest accuracy within the device's gamut. Gamut mapping is selected as a “rendering” intent in the software and is utilized in the color engine. It should be noted that implementation of these rendering intents varies between color engines.

At WBFA, **Perceptual Rendering Intent** is the intent of choice. At profile generation time, the software embeds a media white point tag, which is the actual measurement of either the film substrate or display white. The entire profile is calculated relative to this substrate, so that display-white or film density from an input profile would be mapped onto the media white on any output. The human eye (and brain) is extremely adaptive, and we do not perceive the color of the base media unless it is very dark or colorful. In the color engine co-developed with ITEC®. Perceptual rendering attempts to preserve hue angle, then relative lightness. Saturation is then set to whatever is achievable on the device at the selected hue-lightness combination. To compensate for hue variations in the output gray ramp and color corruption in the red and blue regions at high saturation (typical of animation and digital effects where color values can reach film gamut boundaries easily) a hue correction algorithm is employed. This algorithm simulates the way human vision processes colors from gray to very saturated and results in visually acceptable color reproductions in saturated areas.

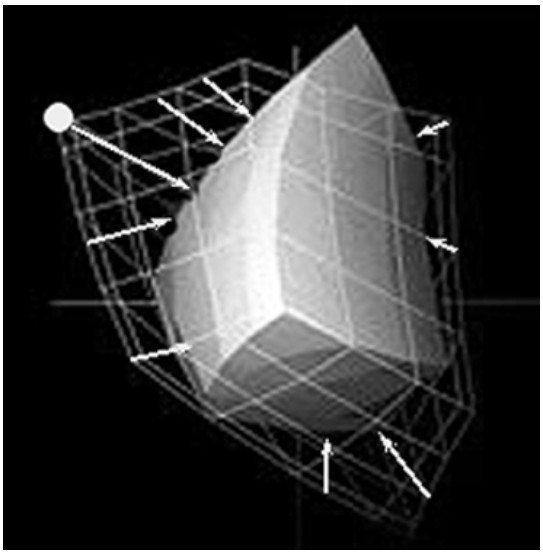

To generate a **Device Profile for a monitor,** firstly calibrate the monitor. Next, measure the colorimetry of a set of colors from the monitor. Then, invoke the profile making application giving it the set of RGB values and their corresponding CIE Values. Ideally, the profile making application should automatically drive this process. At WBFA, critical color stations are equipped with Barco® self-calibrating monitors, calibrated to a 5400K (correlated color temperature) white point, with max luminance of 80 cd/m^2. Calibration is carried out weekly.

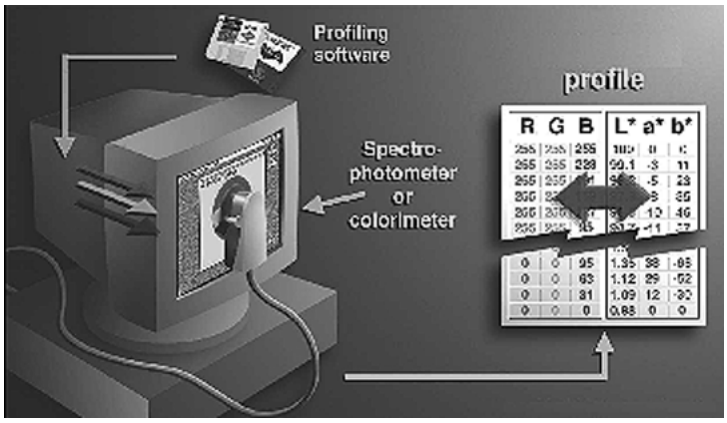

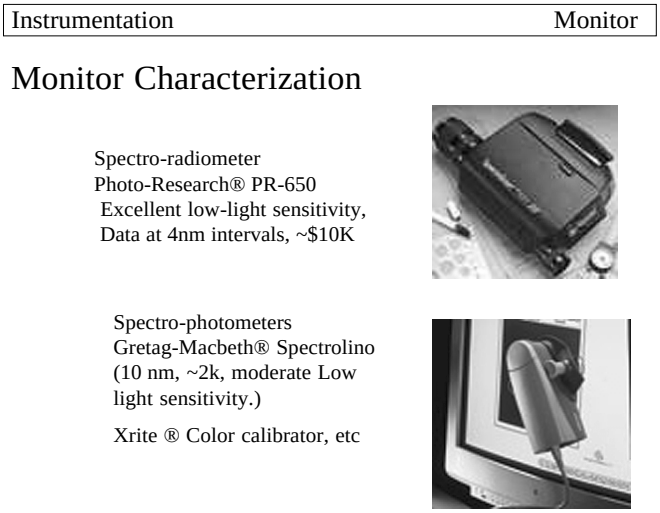

To **Characterize (profile) Monitors,** one may use either a spectroradiometer (ideal) or a spectrophotometer that can measure emissive colors. Cost and accuracy (primarily at low-light levels) are the traditional tradeoffs. Additionally, for closed loop automation, support of a device from the profile making application will also be a determining factor. At WBFA, Monitor Profiles have been generated using ColorBlind® from ITEC, with the PhotoResearch® PR-650 and the Gretag-Macbeth® Spectrolino. Monitor profiles have also been generated using Barco® software in conjunction with onboard sensor, the Optisense®. Given production constraints, profiles were generated using the Optisense® of all "color critical" monitors, and the profile that was closest to the "mean" was deemed as "*The Monitor Profile*".

To **Profile Scanners,** reflective or film, a recommended practice is to obtain a scan target and its reference file from the film's manufacturer.

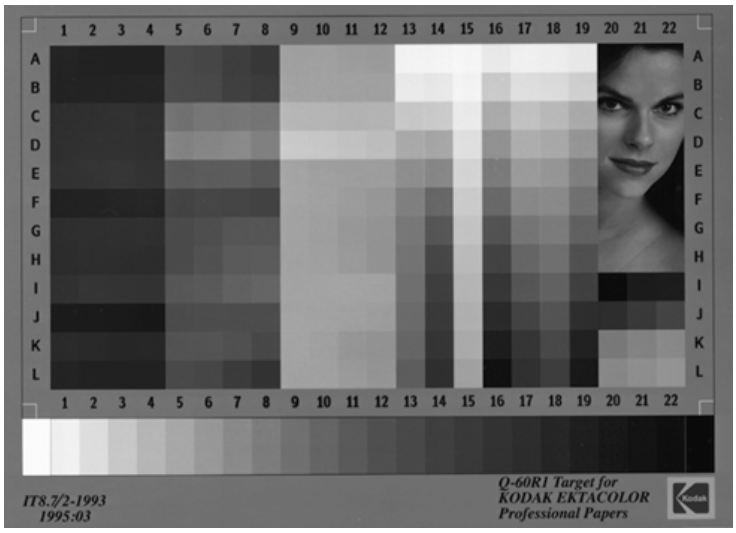

The **Kodak Q-60™ developed by the IT8. 7/2** ANSI committee is shown alongside. A similar transmissive target (4x5 and 35mm) was generated by the IT8. 7/1 ANSI committee. For 35mm versions, the target is spread over 7 slides. Both targets use 12 Hue angles (A-L) and 3 lightness values at each hue angle. At each hue angle and lightness combination there are 4 chroma (or saturation) levels. In addition, 12 flesh tones and a visual image are provided on the RHS. The bottom scale has a 24 step gray scale. Unfortunately, profiles generated using this target did not yield acceptable matches.

The **Failure of the IT8** target at WBFA can be attributed to a variety of causes. Firstly, the media. The reflective target is made on photographic paper and this media has characteristics that are different from paint pigments on Background board and Overlay Cels. Similarly the Ektachrome target is different from camera negative and it is a positive image. Secondly, based on the above, the IT8s do not provide information for the full range of production colors. Thirdly, targets are mass-produced and reference data may be quite different from actual data on the target. Lastly, metamarism caused by the scanner filter set/lamp were amplified because of the different spectral characteristics of paint pigment and photographic media.

Failure of IT8

Why IT8 Target does not work

- Photographic paper (IT8.7/2) or Ektachorme film +ve (IT8.7/1).
- Incomplete gamut characterization.
- Variability of reference data.
- Metamaric failure of pigments.

Metamerism is the phenomenon whereby two colors can match under one set of viewing conditions, but not under another. Under Tungsten lighting, a dark blue (S1) and black (S2) may appear to be similar, but the difference is apparent under lighting, say fluorescent, which has a larger blue component in its spectrum.

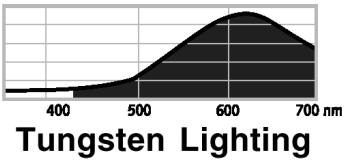

Tungsten Lighting

S1 S2

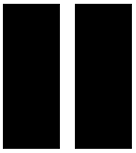

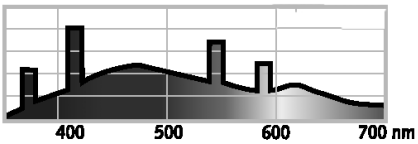

Fluorescent Lighting

S1 S2

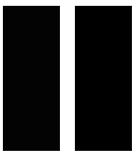

A **Reflective Scanner Target** was generated by hand at WBFA. Background artists were asked to paint 244 swatches of color, representative of the colors favored in the production, and also representative of colors on the IT8 7/2. To generate a manageable target, 1 inch square patches were cut-out and pasted in an arrangement similar to that of the IT8 7/2. The colors were tracked closely during a production for aging and fading and new targets were created for a new production.

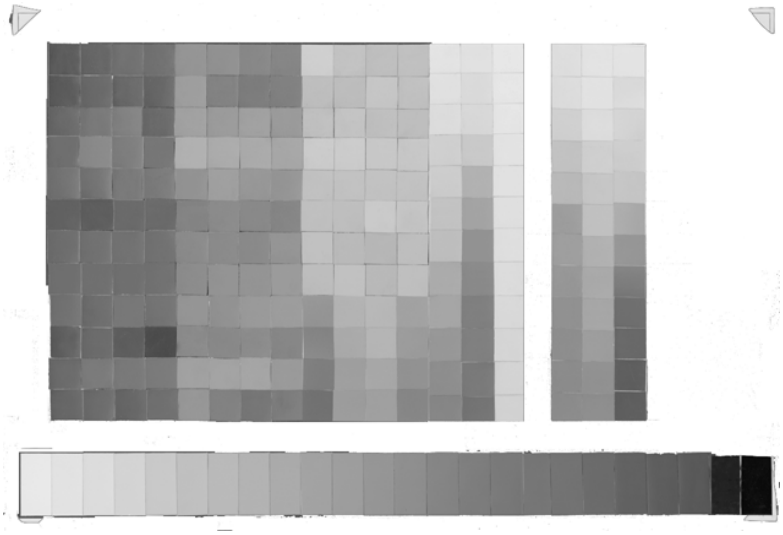

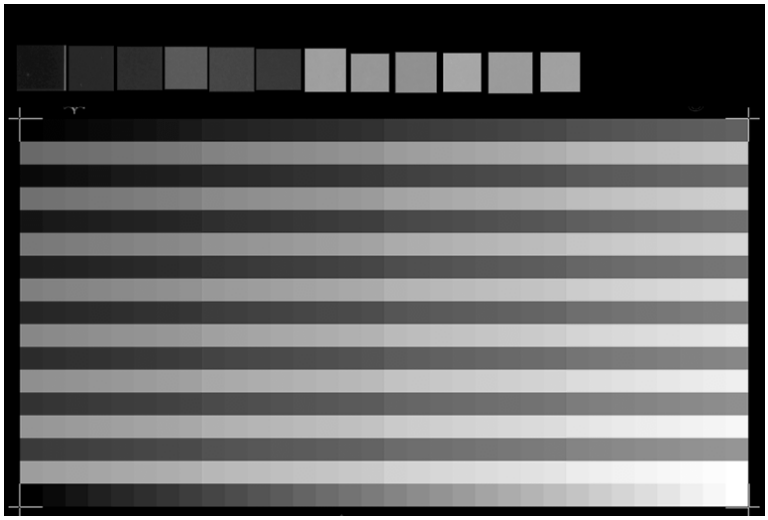

A **Film Scanner Target** was generated by film recording 556 color patches, one patch to a frame. The target is shown here on one frame for conciseness. 12 patches represented "digital" flesh tones and the remainder were color ramps of RGB and combinations.
To generate the profile, the negative was developed and printed. The print was measured to obtain CIE data, while the negative was scanned and converted to positive RGB in the film scanner. The profiling software was then invoked to generate a profile with these 2 data sets. The assumption here is that the negative to positive conversion with the scanner is valid and representative of the film printing process. Results yielded were acceptable. Other, more robust processes have also been implemented resulting in better profiles (lower ΔE).

Instrumentation Scanners

Painted Patches are measured using Gretag-Macbeth ® Spectrolino mated with XY table (Spectrascan)

Film measured using Xrite® Colortron mated with Spotlight® Light Table.

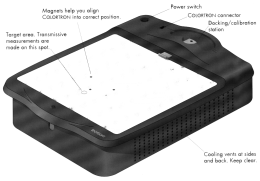

Instrumentation. For the reflective target, the Gretag-Macbeth® Spectrolino mated with the XY table (SpectraScan) was used for obtaining CIE reference data. This instrument also allowed us to monitor the integrity of the target. For the Film Scanner target, the Xrite® Colortron was mated with the Spotlight® light table to measure each film frame and obtain CIE values. The Spotlight Light table has a continuous spectrum calibrated light source suited for these measurements. A new light table attachment to the SpectraScan is now available, permitting similar measurements.

For Film Recorder Profiling a set of patches evenly distributed in the output color space are generated and printed. The target patch values are selected to derive the most information at places where precision is critical. This means small steps in highlights and shadows, larger steps in between. These patches are then measured to provide colorimetric data. From the measured data points, over one million colors are filled in between through spline interpolation. This process ensures precise prediction of where in color space additional printed patches would have resulted in device independent color space. To create an ICC printer profile the software has to create 2 comprehensive tables, one four dimensional table which describes a CIE value for every possible CMY (K)/RGB (i.e., either CMY, CMYK or RGB, depending on the device) combination. The other table consists of a three-dimensional table that maps every possible CIELAB value to a CMY (K)/RGB equivalent. To fully utilize the dimensions of each of the tables, linearization tone curves are generated internally. This optimization allows linear matrix multiplication of a 3x3 matrix to transfer RGB data to CIE data.

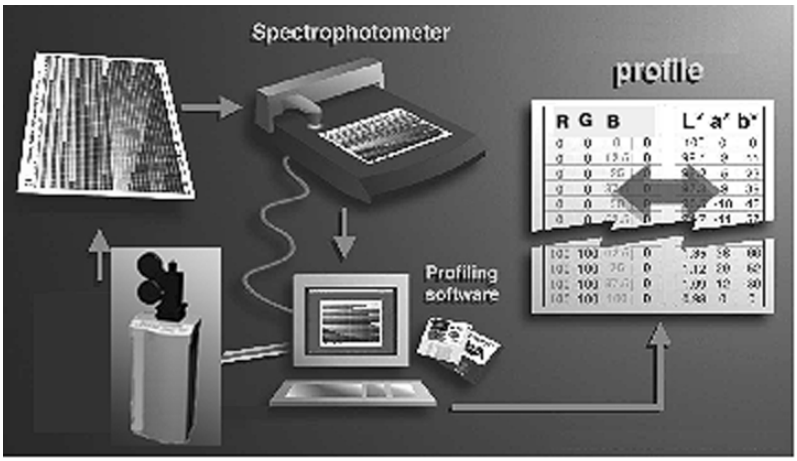

The **Film Recorder Profiling target** is shown alongside. The custom film target for film recorder profiling is composed of 540 patches, consisting of the RGB primaries and combinations of the 3 colors (8x8x8, i.e., 8 steps of R, G, and B and combinations) and a 32 step gray ramp. The target patches are imaged onto film using a digital film recorder and subsequently the film is developed and processed to obtain a positive print.

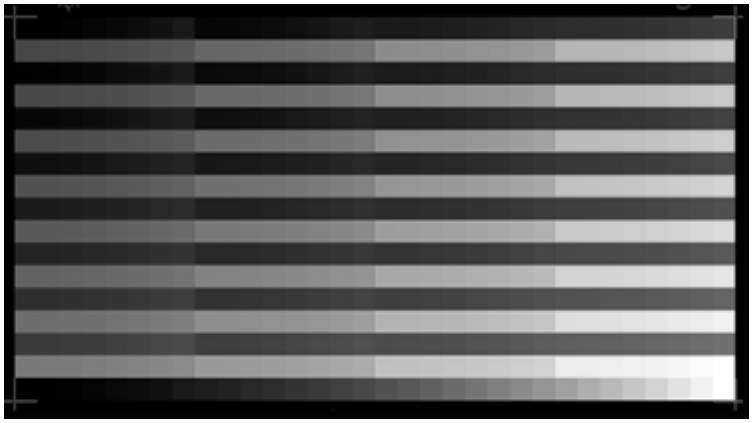

Profiling Issues for Film

- CRT Flare and Scanner crosstalk.
- Scanner filter set inaccuracies.
- Negative development variations.
- Variability in positive film print.
- Density variations within a Frame.

Profiling Issues for Film. Initially, the entire target was on put on 1 frame, with the intention of using a "profiled" scanner as a colorimeter. However, this proved futile because CRT flare caused one color to bleed into another and the Scanner filter sets did not capture the entire film gamut. After multiple approaches, the decision was made to image each color on a separate frame. This negative was then developed. Film lab variations caused variations of ±0.05 density points, requiring us to closely track the negative, and re-shoot the frames so that we may get a "mean" bath day. For the printing process, the lab variations are closer to ±0.15, requiring us to re-print the "chosen" negative over the period of 1 week and then selecting the "mean" print. To minimize the variations within a frame, the central 5 sq. mm was chosen as the measurement area.

Instrumentation Film

Film Calibration universally done using Xrite ® 310 Densitometer.

Film profiling carried out using Xrite® Colortron mated with Spotlight® Light Table.

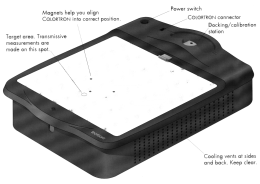

Instrumentation. For film recorder calibration, the Xrite® 310 densitometer is the standard in the industry. For profiling, the Xrite® Colortron mated with the Spotlight light table was used.

Revisiting the **Standard Workflow (ideal),** the desired intention is to transfer images between devices, seamlessly and transparently, maintaining color fidelity throughout the process.

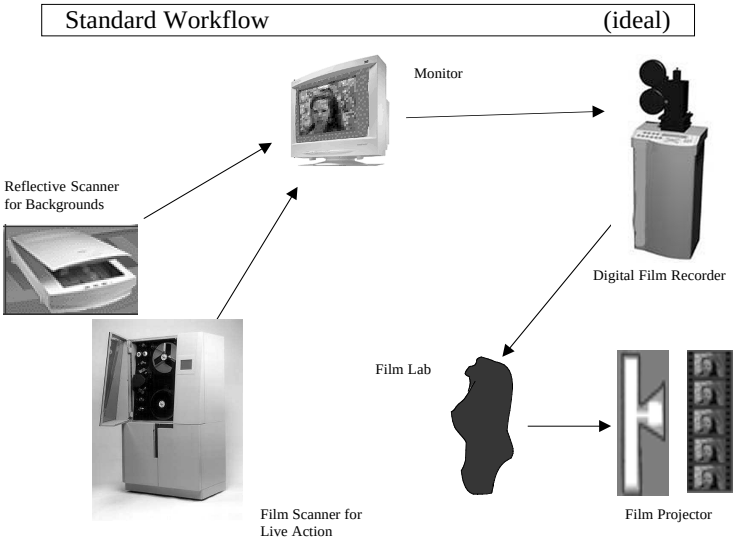

In the **Workflow Model with CMS,** Images are transferred from each device's space to the CIE space, and then out of CIE space to the output/display device, with appropriate gamut mapping applied through rendering intents. The color engine concatenates input and output ICC profiles, creating a composite transform, requiring only one pass for all linear interpolations. This reduces data artifacts and data loss. Internal computations are carried out in floating point precision, with support for 16-bit data. The computation employs tetrahedral interpolation to compute CIE XYZ values. Tetrahedral interpolation processes four points at any one time and changes no more than three points during a transition, resulting in rapid computation times and seamless transitions.

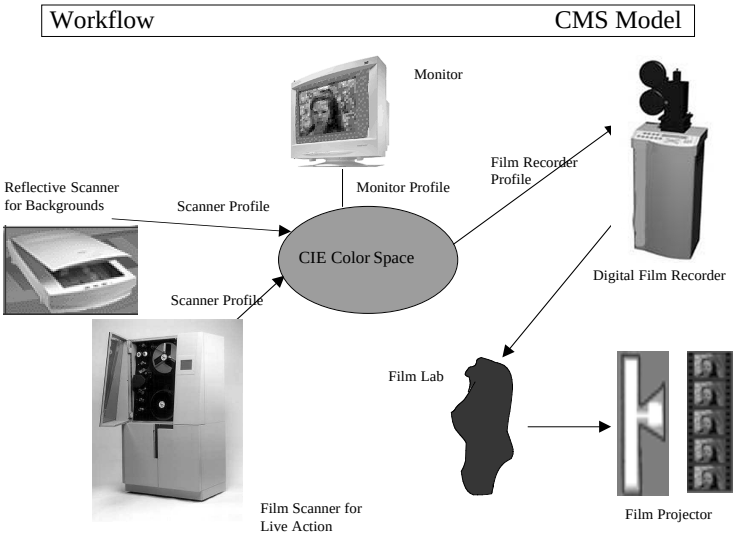

The Workflow at Warner Bros. Feature Animation, consists of images and effects being digitally created on the monitor, with approvals also carried out on the monitor. Final approval is still carried out on film. For production reasons, images are transferred to and maintained in Monitor Space. Given film archiving, restoration or just scan-record workflows, the images may be maintained in either Scanner space or Film Recorder Space and transfer between them accomplished with an ICC device link directly between Scanner and Monitor.

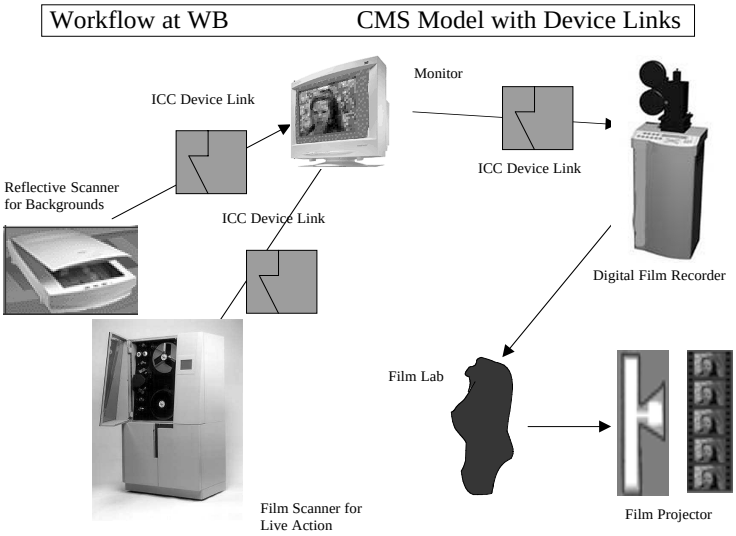

Results

- Compute Mathematical Fit, ΔE.

Device	Medium	Avg. ΔE
Monitor (Barco ® 120 T)		1.8
Film Recorder (Celco ® MPX)	EK, 5245	2.3
Film Recorder (Celco ® MPX)	SE, 5231	2.8
Reflective Scanner (Dicomed ®)	Paint on Cel and Board	1.75
Film Scanner (Kodak ® Genesis)	EK, 5245	2.3

Results. Using the ΔE as the metric of color reproduction accuracy, the quality of color match may be measured. It is possible to compute the ΔE for a particular device by examining the quality of fit between the forward and reverse profile of a device. Mathematically, we can pass "known" data through the forward profile, then using the reverse profile, come back to the original space. Using the statistics modules ColorBlind, the figure shows ΔE values for the devices used at WBFA. Visual comparison, however, of images should be regarded as the ultimate judge of the quality of color match.

Limitations of a CMS

- Be-aware of "private" tags.
- Do not confuse Calibration and Profiling.
- Tight control over drift and processes is Important.
- Not all colors will match perfectly. There is a ΔE.
- Lighting conditions affect quality of color match.
- Currently no profiling capability for negative film.
- Is CIELAB the ideal model ?
  Conical Model has problem at the edges.
  Fluorescent colors do not interpolate well..

Limitations of a CMS. Among the numerous limitations it should be noted that although the ICC's framework is well published, the inclusion of Private Tags in profiles makes some profiles perform better than others depending on the color engine chosen.Also be prepared to "edit"/modify the profile, because based on visual checks color matches may not be as good as you would like. Many times calibration and profiling are confused and device behavior is changed after a profile has been generated. This will adversely affect the quality of color match. The same applies to Gamma changes, LUT applied after profiling, etc. Bear in mind that there is a ΔE, hence error in color match and this will be more apparent in certain colors than others. Color matches will only be as good as the instrumentation, the devices, and the consistency with which they reproduce color. Processes and devices should be closely monitored and any drift immediately corrected, otherwise new profiles will have to be constantly made and edited. Lighting Conditions should be similar when quality of color match between devices is evaluated. This is perhaps the most overlooked aspect of color match evaluation. Viewing booths, Light tables, and ambient light conditions should be properly setup and checked regularly. Currently no profiling capability exists for negative film, causing us to transfer film to positive to make color measurements.

CIELAB space limitations prevent good interpolations for colors that are near gamut boundaries of a device like the film recorder. The lack of any information on the Background around the image and the Surround Field around the background shows additional deficiencies in the CIELAB appearance model.

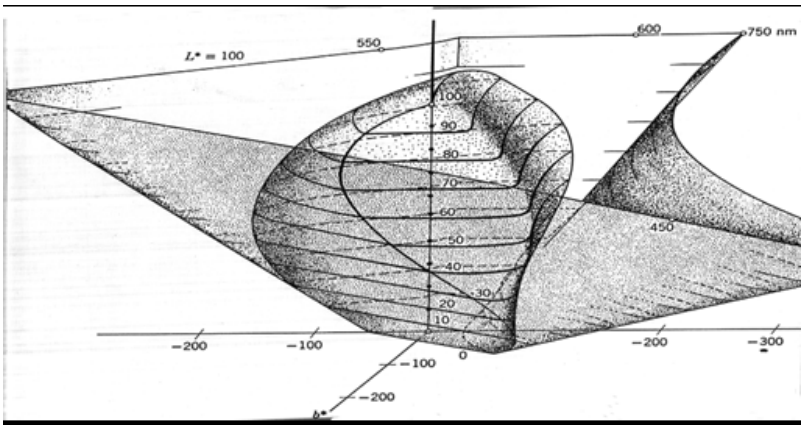

From "Color in Business, Science and Industry", D.B. Judd and G. Wyszecki

From the production trenches, the Benefits of a Color Management System are very clear and have been realized over the course of three feature productions. Additionally, other digital production studios have also begun to implement portions in their workflow. At Warner Bros., we have realized savings in time, labor and material. Prior to CMS, 2 hours were budgeted for manual color correction on each background (>1500); after CMS was implemented, this had dropped to almost zero. Camera retakes, de to incorrect color match, were also eliminated. Equally important, the CMS was able to provide artists with the creative freedom to select from a wider gamut of colors than previously available and have the confidence that the colors would reproduce faithfully.

Benefits of a CMS (from the trenches)

- "Quest for Camelot" ,"The Iron Giant" & "Osmosis Jones" (2001, Live Action + Anim).
- Savings in Time, Labor and Material.
- No Manual Color Correction for film output.
- No Manual Color Correction for scanned images.
- Camera re-takes due to color match down to 0.
- Greater Director and Artist confidence, larger gamut.
- No media memory required for Creative Crew.
- Dramatic Time Savings in POST production.

Mastering movies for a digital release 8

Isaac Kerlow
New Technology and New Media
The Walt Disney Company

We present an overview of the process of preparing animated movies and digital live action movies for a digital release. This presentation examines some of the practical details and key issues in the process. Examples of D-cinema releases will be screened and some lessons learned will be shared with the audience.

1 Overview of the D-cinema effort at Disney

2 Overview of the D-cinema pipeline

3 Differences between animation and live action

4 Short chronology of recent developments

5 Case studies

6 Overall compression issues and image quality

7 Digital color timing issues and the cinematographers' point of view

8 Lessons learned

Copyright © 2001-05-10 The Walt Disney Company

Color Image Coding for Digital Projection and D-Cinema

Dave A LeHoty, Silicon Light Machines

March 7 2000

1. Introduction

The gamut of reality, film, or a SMPTE d-cinema recommendation defines a volume which is spanned by color and luminance according to the CIE tristimulus model for human color vision. Given the ability to translate content to tristimulus values, the values can be digitized, stored, recalled, and processed as needed to support the d-cinema workflow. The color image coding can change in support of the unique characteristics of the d-cinema workflow steps.

The purpose of this paper is to stimulate discussion on color image coding within the SMPTE D-Cinema Adhoc Group on Colorimetry. To that end, this paper includes:

- several color image coding systems and their defining characteristics
- luminance contour plots for ITU-R BT.709, ITU-R BT.1361, and a film stock
- a color image coding system for efficient perceptual color difference encoding
- comments on representing d-cinema gamut with ITU-R BT.1361

We assume the reader is familiar with the principles of color science.

2. Representation of a Gamut Including Luminance

Gamut is related to the CIE tristimulus values through linear equations, transfer functions ("gamma characteristics"), or both. Some representations, such as Cineon and DPX, characterize dye concentrations measured directly from motion picture film. Representations based on these non-fixed primaries cannot be related to CIE tristimulus values through linear equations and/or transfer functions. At some point in the d-cinema workflow where the data moves from a studio-defined representation to a SMPTE standard representation, it is desirable that the data be related to tristimulus values in a straightforward way.

One solution is to apply an ICC device profile to map an RGB film scanner triplet into an additive linear-light representation such as non-truncated RGB values or tristimulus values. This may be awkward if some correction layers are much more straightforward to specify or apply in the scanner space. This approach is desirable, though, because it enables the correction layers to be expressed independent of the nonlinear details of the combination of the specific film and scanner. Also, careful examination of the linear-light representation shows the gamut possible within the representation. And, careful examination of the ICC device profile shows the gamut limits of the combination of the specific film and scanner within the linear-light representation limits. Cineon and DPX files can hold this data based on fixed primaries if the data represents additive RGB values.

If multiple processing steps are to occur in the linear light domain, a linear light representation would mitigate the compounding of errors due to conversion between nonlinear representations such as gamma characteristics and linear light representations during the creative process. The data could be converted to some standard nonlinear representation just before or as part of the compression process. The motivation for applying some gamma characteristic or other nonlinear transfer function is to reduce the bit depth required to store, transport, and/or project a quality image to the d-cinema audience.

Eight color image coding systems and their defining characteristics are listed below. In addition to the representations directly related to tristimulus values through linear equations and transfer functions, representations with separate luminance and chromaticity coordinates are also included.

Copyright © 2001-05-10 Silicon Light Machines

Throughout, the term luminance (Y) is used for a quantity proportional to linear-light intensity. The term luma (Y') is used for a quantity formed from the weighted sum of gamma-modified (non-linear-light) red, green, and blue; for more see [Poynton, Charles, *A Technical Introduction to Digital Video,* New York: John Wiley & Sons, 1996]. Standard CIE nomenclature is used with XYZ for tristimulus values, xy for chromaticity values, and u'v' for the 1976 system of approximately perceptually uniform chromaticity values. We also conform to the convention of using a prime to indicate gamma-modified (non-linear-light) values except for the entrenched conflicting conventions for u'v', C_BC_R, and P_BP_R.

For each representation, some salient characteristics are listed. When the interpretation of a gamut coordinate is not well-defined, its range is not listed. In general, the range of a gamut coordinate may restrict the volume of luminance and chromaticity spanned. The characteristics which can define a representation or differentiate two representations using similar coordinates are designated with the symbol *. These characteristics together with some scaling factors define the representation and its absolute reference with respect to the tristimulus model.

An arbitrary point in this volume can be represented by triplets which include:

- [X, Y, Z] or XYZ or CIE tristimulus values
  - this is a linear-light representation
  - where the X range is ~ 0 .. 1.1
  - where the Y range is the luminance range*
  - where the Z range is ~ 0 .. 1.8

- [x, y, Y] or xyY
  - this is <u>not</u> a linear-light representation
  - where the x range is ~ 0 .. 0.75
  - where the y range is ~ 0 .. 0.85
  - where the Y range is the luminance range*

- [u', v', Y] or u'v'Y
  - this is <u>not</u> a linear-light representation
  - where the u' range is ~ 0 .. 0.65
  - where the v' range is ~ 0 .. 0.60
  - where the Y range is the luminance range*

- [R, G, B] or RGB
  - this is a linear-light representation <u>when not clipped</u>
  - the red xy coordinates*
  - the green xy coordinates*
  - the blue xy coordinates*
  - the whitepoint xy coordinates*
  - the luminance [Y] of a specific non-zero [R, G, B] triplet*
  - the range of R, G, and B*
  - where the [R, G, B] values can be negative enough to span the colors and luminance
  - where the [R, G, B] values can be large enough to span the colors and luminance
  - the primaries, whitepoint, and extents may follow ITU-R BT.709 or 1361

- [R', G', B'] or R'G'B'
  - this is <u>not</u> a linear-light representation
  - the defining characteristics of the underlying [R, G, B] -space*
  - the gamma characteristic*
  - the gamma characteristic may follow ITU-R BT.709 or ITU-R BT.1361

- [Y', B' - Y', R' - Y'] or Y' |B'-Y' |R'-Y'
  - this is <u>not</u> a linear-light representation
  - the defining characteristics of the underlying [R, G, B] -space*
  - the gamma characteristic*
  - the equation for relating luma [Y'] to [R', G', B'] *
  - Y' C_BC_R and Y' P_BP_R are linearly related to Y' |B'-Y' |R'-Y' and also R'G'B'
  - ITU-R BT.709 and ITU-R BT.1361 are Y'C_BC_R representations

- [Y, B - Y, R - Y] or Y|B-Y|R-Y
  - this is a linear-light representation <u>when not clipped</u>
  - a gamma characteristic is <u>not</u> applied
  - this representation is consistent with constant luminance coding
  - the defining characteristics of the underlying [R, G, B] -space*
  - the equation for relating Y to [R, G, B]*
  - for constant luminance the equation for Y is derived from the [R, G, B] -space definition

- [(Y)', (B - Y)', (R - Y)'] or (Y)'|(B-Y)'|(R-Y)'
  - this is <u>not</u> a linear-light representation
  - a gamma characteristic is applied <u>after</u> the luminance and differences are formed
  - this representation is consistent with constant luminance coding
  - the defining characteristics of the underlying [Y, B - Y, R - Y] -space*
  - the gamma characteristic*

The three linear-light representations listed above (XYZ, RGB, Y|B-Y|R-Y) are mathematically equivalent, within a linear 3x3 matrix transformation, if all of the defining details are known and the coordinate ranges are not clipped. The other representations require the application of gamma characteristics. If the details enumerated above are known, one can translate between any of these representations and a tristimulus triplet. Since most of the representations are built on the foundation of linear-light RGB triplets, let's investigate the colors and luminances of these representations.

3. Luminance Contour Plots of RGB Representations in u'v'-space

When the RGB values are non-negative, the colors which can be represented lie within a triangle in u'v' space. The vertices of this triangle correspond to the specific colors used to define the red, green, and blue coordinates. At any color within or on the triangle, there is a maximum luminance which can be represented by the RGB values. This maximum luminance can be plotted in three-space with respect to two color coordinates. Alternatively, maximum luminance contours can be projected onto the plane spanned by the two color coordinates much like a topographical map. Even though gamut is often plotted in two dimensions, it is a three dimensional space or volume which can be spanned by luminance and two color coordinates.

The shape of the luminance contours are a function of the following:
- the u'v' coordinates of the red, green, and blue primaries
- the whitepoint u'v' coordinates
- the ranges of the RGB values
- the limits, if any, imposed by encoding

In addition, in order to faithfully or consistently project an image given RGB triplets with such a pedigree, some additional details are necessary:
- the absolute luminance or acceptable luminance range of some white, not necessarily D_{65}
- the intended rendering environment (such as a dark theatre)
- the color and luminance used when the RGB triplets are not realizable by the projector

Below are some luminance contour plots for a variety of assumptions. The primaries and whitepoint conform to ITU-R BT.709; so, R=G=B yields the color of the whitepoint, D_{65}. The source material need not be created with a whitepoint of D_{65}, since the source can be translated into this coordinate system after creation. If the translation is done well, the whites of the source material will be represented, just not with R=G=B. On the other hand, a projector optimized for maximum luminance at D_{65} will produce a lower luminance at any other color including a different white of interest such as D_{55}.

3.1. Luminance Contour Plots for ITU-R BT.709 and ITU-R BT.1361

The luminance contour plots that follow generally have a consistent legend. Green is used to show the spectral locus and the line of purples. Red is used to show the triangle spanned by ITU-R BT.709 primaries. There is a red cross at D_{65}. The contours for 3%, 20%, 40%, 60%, and 80% of max luminance are shown in magenta. The contours for 10%, 30%, 50%, 70%, and 90% of max luminance are shown in blue.

For ITU-R BT.709, the linear-light RGB values are restricted to the range zero to one, inclusive. Therefore, the luminance contours are constrained to lie within the primaries of that standard. Note that the 3% contour also traces the triangle bounded by the primaries.

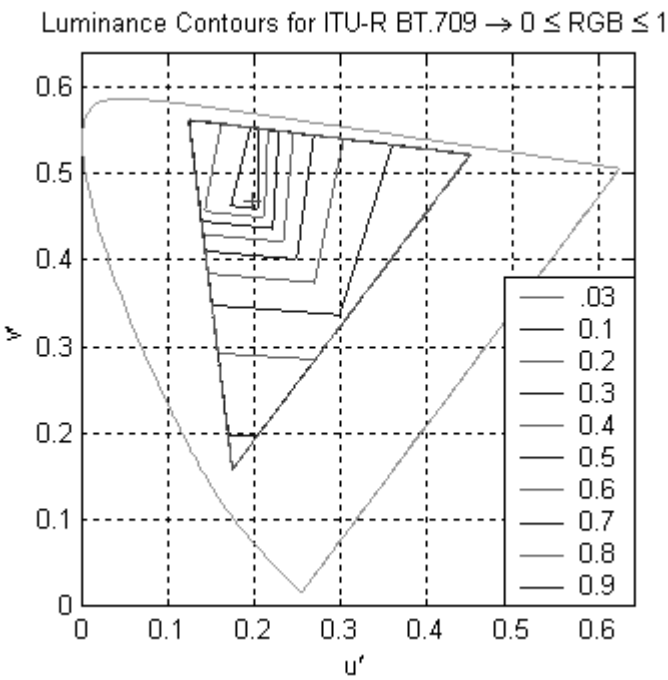

Figure 1.

For ITU-R BT.1361, the RGB values are restricted to the range -0.25 to 1.33, inclusive. Since negative values are allowed by the representation, the luminance contours extend beyond the triangle formed by the ITU-R BT.709 primaries. This enables the RGB representation to include almost all colors within the 3% luminance contour except those blues and purples requiring wavelengths below about 460 nanometers.

The luminance contour for ITU-R BT.1361 shown in Figure 2 does <u>not</u> include the constraints of limited excursions in luma (Y') and chroma (C_BC_R) implied by the standard. The limits are:

$$\frac{1-16}{219} \le Y' \le \frac{254-16}{219}; \qquad \frac{1-128}{224} \le C_B, C_R \le \frac{254-128}{224}.$$

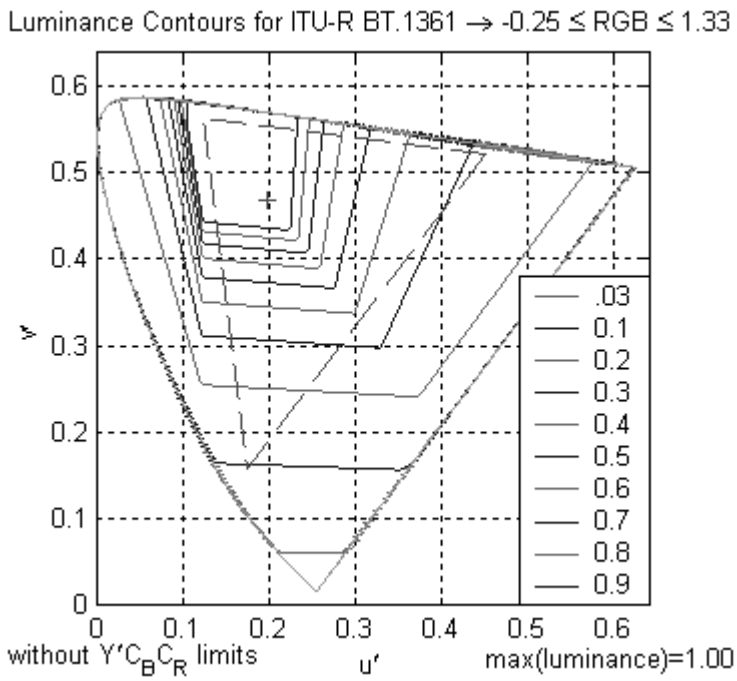

Figure 2.

Any color can be represented when the RGB values span positive and negative values; but, the luminance which can be represented for an arbitrary color may be very small or even negative. For ITU-R BT.709 primaries, luminance is related to RGB values by: $Y = 0.2126R + 0.7152G + 0.0722B$.

The 90% Luminance contour for ITU-R BT.1361 is much larger than that of ITU-R BT.709 since the RGB extents have been increased from unity to 1.33. Figure 3 shows luminance contours for ITU-R BT.709 with one modification -- the minimum RGB value has been extended from zero to -0.25. Here we can see the additional colors available just by allowing some negative RGB values.

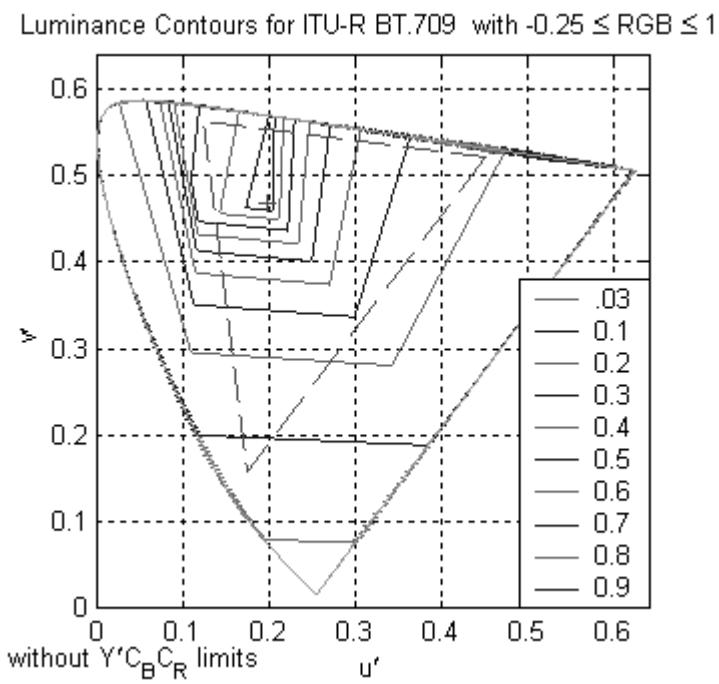

Figure 3.

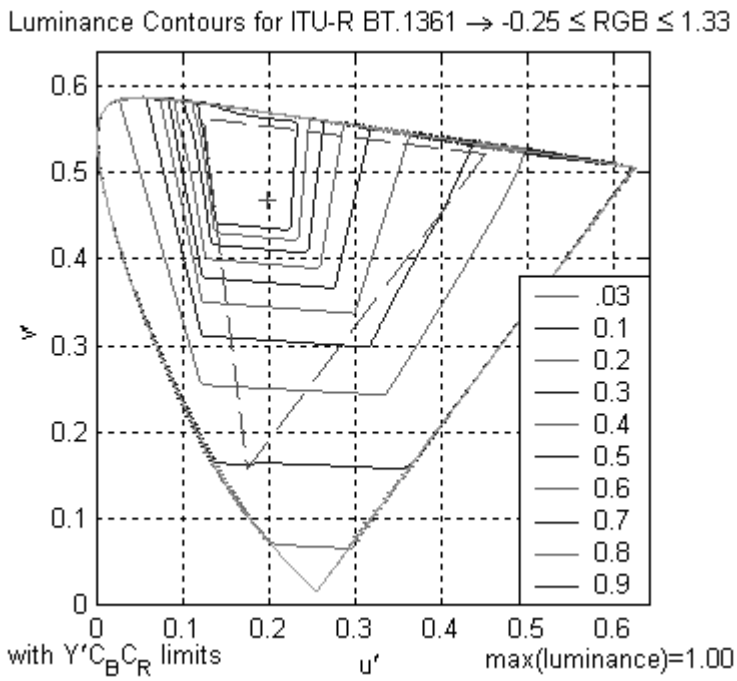

Figure 4.

Figure 4 shows the luminance contours for ITU-R BT.1361. Including the effects of the $Y'C_BC_R$ limits have not changed the contours significantly. Compare Figure 4 to Figure 2. The largest changes occur at the 20% luminance contour. As the absolute RGB values get smaller, the corresponding absolute C_BC_R values also get smaller and are less likely to extend beyond the C_BC_R limits.

Within the primary triangle, the maximum luma extent restricts the maximum luminances. For example, the luma limit restricts the maximum luminance of the whitepoint color to ~1.087 compared to the maximum RGB extent of 1.33. One interpretation of the opportunity to set RGB values above unity is to increase the 90% luminance contour as opposed to increasing the maximum luminance at the whitepoint color. Since the largest contour plotted is at 90% of RGB set to unity, the effect of limited maximum luma does not change the contour plot much.

3.3. Luminance Contour Plots Using Narrowband Primaries

Some future d-cinema projectors are likely to create images using lasers. At Silicon Light Machines our current development projector uses solid-state lasers for the red, green, and blue primaries. The corresponding laser wavelengths are 642, 532, and 457 nanometers. These were convenient wavelengths available to us in off-the-shelf solid state lasers.

The Silicon Light Machines Grating Light Valve® technology can be used in a number of applications requiring an expanded color gamut, one of which is d-cinema. The primary wavelengths under consideration span 608..650, 525..550, and 450..470 nanometers. The lasers used for a particular application are a function of total system cost for the desired brightness and gamut delivered to the viewers. Based on our tradeoffs of film gamut and available lasers, a d-cinema system should probably have wavelengths within the ranges mentioned above.

Primaries located at these wavelengths are certainly more saturated than the ITU-R BT.709 primaries. Even without negative RGB values, the laser primaries span a gamut close to film but with generally higher luminance. These primaries also encompass the primary triangle of ITU-R BT.709. For consistency with the other plots, the whitepoint is D_{65}. The RGB extents are limited to zero and one. The colors of this narrowband primary representation or equivalent projector are bounded by the narrowband primaries as shown in Figure 5.

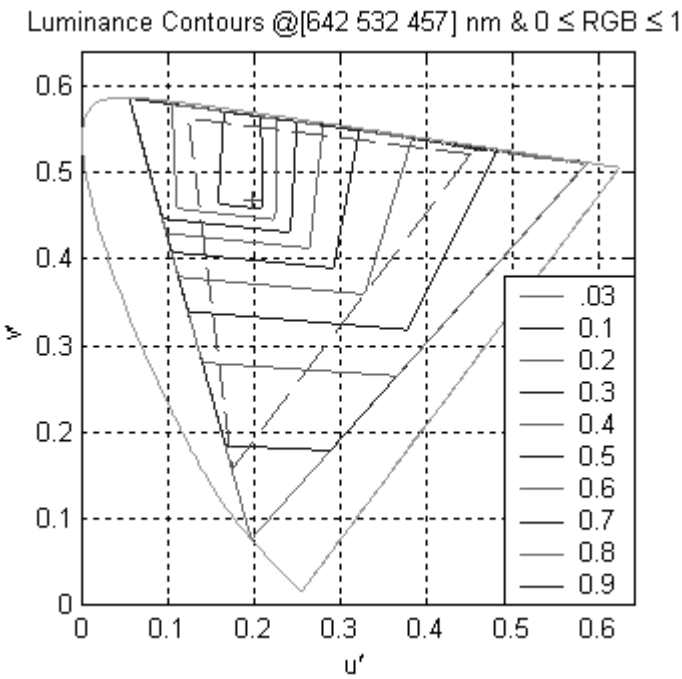

Figure 5.

4. The Gamut of D-Cinema

We can choose to limit the gamut representation for d-cinema color and luminance to that spanned by:

- the possible human experience (i.e., no limit to the gamut),
- the typical human experience,
- the typical human experience which can be recorded today,
- the possible human experience which can be recorded today,
- the movie film and projectors in widest use today,
- the movie film and projectors available to the movie industry today,
- the anticipated movie film and projectors in widest use in the future, or
- the anticipated movie film and projectors available to the movie industry in the future.

ITU-R BT.1361 is titled *Worldwide Unified Colorimetry and Related Characteristics of Future Television and Imaging Systems*. Quoting Annex I of this recommendation, "A realistic approach is to limit reproduction of the gamut of real surface colors as determined by Pointer." This may be a defendable compromise for television showing real surfaces, but what about the gamut of movie film and projectors which record the raw source and project the finished product to cinema audiences? Certainly our imagination and creativity are not limited to real surfaces.

Alternatively, we suggest that we start with the reproduction of the gamut of the film experience available today and plan for an expansion of the gamut available in the future to the creative community and d-cinema audiences.

4.1. The Current Gamut of Film

Chuck Harrison contributed to the adhoc group a luminance contour plot for film based on an approximate analytical densitometric model for an obsolete film stock [Harrison, Chuck, *A Review of Photographic Color Science, and Its Application to Digital Cinema Exhibition*, 20 February 2000; see Figure 8: Variation of Film Gamut with Brightness Level]. The data, reprinted with permission, is shown below in Figure 6.

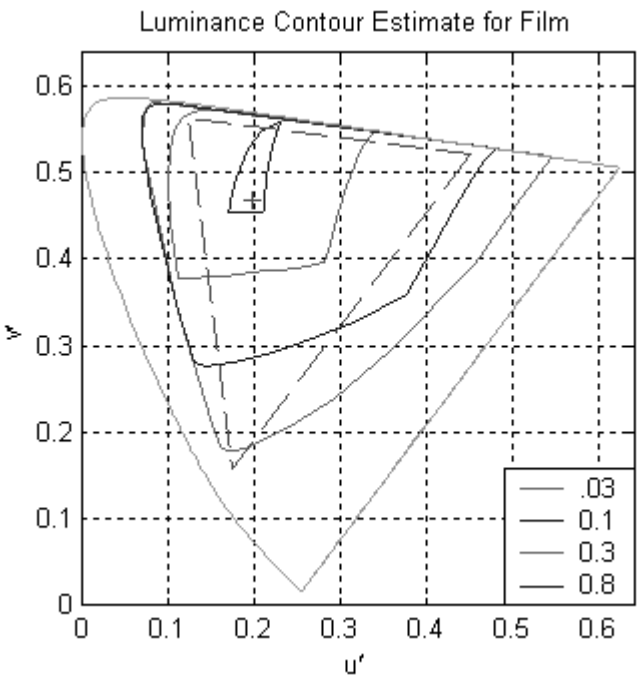

Figure 6.

Harrison believes that the film stock in use today has more pure primaries and therefore a somewhat larger gamut; so, current film should have luminance contours that span more colors than shown in the plot.

Using this plot as an estimate for the luminance contours of current film, we can make comparisons to the luminance contours for ITU-R BT.1361 and other RGB representations shown in section 3.

Luminance contour estimates and measured values for cinematic films in use today are needed to evaluate the needs and standards for d-cinema.

4.2. The Future Gamut of D-Cinema

The future gamut of d-cinema may be expanded to include the additional colors and intensities that are available to d-cinema projectors.

5. Suggestions for Alternate Gamut Representations in D-Cinema

Let's consider alternatives to RGB-based representations of gamut. Other feasible representations are XYZ, xyY, and u'v'Y. These representations can span the visible gamut. So, let's compare the compressibility, processing requirements, bit depth, and perceptual color differences.

5.1. XYZ

Given tristimulus values, multiplication by a three-by-three color space transformation matrix is needed to generate the RGB values for a particular three primary (additive) projector; this transform requires nine multiplies and six adds per pixel. Given the relationships between xyY or u'v'Y to XYZ, the operations, input/output, and lookup memory can be estimated for converting to XYZ and ultimately RGB.

For the luma/chroma system in use today in SMPTE 292M, five multiplies and four adds can be used for luma/chroma decode from $Y'C_BC_R$ to R'G'B'. Given encoded tristimulus values, these five multiplies and four adds are not needed. But, some decoding to XYZ is likely to be necessary. And, furthermore, a multiplication-by-a-constant is much easier to implement in hardware than a multiplication of two

variables. Multiplications-by-a-constant can be used for luma/chroma decode if the decode coefficients are standardized.

Unfortunately, tristimulus values may not compress efficiently since they do not separate into luminance and two color coordinates.

5.2. xyY

The relationship between xyY and XYZ is:

$$X = \frac{1}{y}(xY); \qquad Z = \frac{1}{y}(1 - x - y)Y$$

Given:

- encoded x
- encoded y
- encoded Y

Generate XYZ with the follow steps:

- use encoded x to calculate or lookup x
- use encoded y to calculate or lookup (1/y)
- use encoded Y to calculate or lookup Y
- perform two adds and four multiplies per the equations above

Note that if x and y coordinates were in perceptual space, encoding might not be necessary; so, it would not be necessary to decode the x input to find x. Remember that the u'v' coordinates are in perceptual space approximately; so, fewer gates or less memory may be needed with a u'v'Y representation.

Of course, given sufficient memory resources, one could lookup the value for say (R/Y) directly from a table using an address formed from the concatenation of encoded x and encoded y. For a given perceptual color difference goal, the lookup table may be smaller using color coordinates such as u'v' which map more closely to perception.

5.3. u'v'Y

The relationship between u'v'Y and XYZ is:

$$X = \frac{1}{4v'}(9u'Y); \quad Z = \frac{1}{4v'}\left((12 - 3u') - 5\right)Y$$

Given:

- encoded u' = u'
- encoded v' = v'
- encoded Y

Generate XYZ with the follow steps:

- use encoded v' to calculate or lookup 1/(4v')
- use encoded Y to calculate or lookup Y
- perform two adds and six multiplies per the equations above

Note that two of the six multiplications are by constants (each with only two non-zero bits, so the additional processing throughput is minimal). Careful reuse of the (3Y)/(4v') term will simplify two of the multiplies. For even greater throughput, the calculation of X and Z could be combined with the color space transform operation. Also, the only lookups which may be necessary are for encoded v' to 1/(4v') and for encoded Y to Y. If we choose an encoding scheme for Y which is efficient to decode in real-time, then a u'v'Y representation can be implemented in real-time today with current technology!

A Log(Y)u'v' standard has been implemented. This is the SGILOG compression type in Sam Leffler's popular TIFF I/O library. See [*LogLuv Encoding for Full-Gamut, High Dynamic Range Images*; Vol. 3, No. 1: 15-31 of the Journal of Graphics Tools; Greg Ward Larson]. This representation uses eight bits for u', eight bits for v', and sixteen bits for Log(Y). Quoting from the paper by Ward:
"this encoding has the following desirable properties. It

- covers the entire visible color gamut;
- covers the full range of perceivable luminances (over 38 orders of magnitude);
- uses imperceptible step sizes in a perceptually uniform space;
- may be calibrated to absolute luminance and color;
- enables optimal visual fidelity on any output device."

For d-cinema, we may prefer to use an encoding for Y which is efficient to decode in real-time without resorting to a large lookup table. We are unlikely to need to cover 38 orders of magnitude in luminance for d-cinema.

A variation of the SGILOG representation would fit in thirty bits per pixel. This variation uses twelve bits for log luminance and nine bits each for u' and v'. This representation is consistent with the Principle of Constant Luminance. The solution is shown below using lookup tables in lieu of multipliers and adders.

Converting Log(Y)u'v' to Linear-Light RGB Aided by Tables

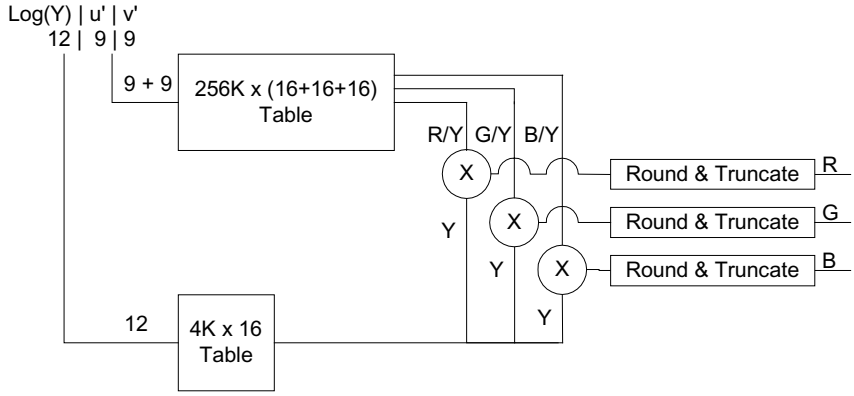

Figure 7.

The maximum step size of this 30 bit representation is about one unit of ΔE_{uv}. Since this representation uses luminance and color encoding which approximate brightness and color perception, other representations of the same bit depth may have a higher ΔE_{uv} value and higher perceptual color differences resulting from encoding quantization. This is more likely as the luminances and chromaticities represented increases.

Poynton states "If ΔE_{uv} is unity or less, the color difference is taken to be imperceptible. However, L*u*v* does not achieve perceptual uniformity, it is merely an approximation. ΔE_{uv} values between 1 and 4 may or not be perceptible, depending upon the region of color space being examined. ΔE_{uv} values greater than 4 are likely to be perceptible; whether such differences are objectionable depends upon circumstances" [Chapter 20, *The CIE System of Colorimetry*, in Poynton, Charles, *Digital Video and HDTV: Pixels, Pictures, and Perception*, in press].

If the requested luminance and chromaticity are not realizable by the d-cinema projector, what luminance and color or RGB values are projected? This question arises with any extended gamut image coding such as ITU-R BT.1361 or Log(Y)u'v'. We could simply truncate the values at zero and one; this is straightforward to implement, but the resulting look may not match the intent. SMPTE should address this question so that the projection of non-realizable luminance and chromaticity is predictable and acceptable.

6. Summary of Comments on Representing D-Cinema Gamut with ITU-R BT.1361

ITU-R BT.1361 can be related to tristimulus values through linear equations and gamma characteristics. In order to faithfully or consistently project an image represented in this standard some additional details are necessary:

- the absolute luminance or acceptable luminance range of some white, not necessarily D_{65}
- the intended rendering environment (such as a dark theatre)
- the color and luminance used when the RGB triplets are not realizable by the projector

The primaries and whitepoint of ITU-R BT.1361 are those of ITU-R BT.709; so, R=G=B yields the color of the whitepoint, D_{65}. The source material need not be created with a whitepoint of D_{65}, since the source can be translated into this coordinate system after creation. The whites of the source material will be represented, just not with R=G=B. A projector optimized for maximum luminance at D_{65} will produce a lower luminance at any other color including a different white of interest such as D_{55}.

As mentioned above, if the requested RGB triplets are not realizable by the d-cinema projector, what luminance and color or RGB values are projected? This question arises with any extended gamut image coding such as ITU-R BT.1361 or Log(Y)u'v'. We could simply truncate the values at zero and one; this is straightforward to implement, but the resulting look may not match the intent. SMPTE should address this question so that the projection of non-realizable luminance and chromaticity is predictable and acceptable.

We have provided many luminance contour plots for comparison, including contours for ITU-R BT.709, ITU-R BT.1361, and a film stock. Comparing the plots, there is some mismatch in the ability of ITU-R BT.1361 to represent some of the 30% luminance contour estimated for film. More luminance contour estimates and measured values for cinema films in use today are needed to evaluate the needs and proposals for d-cinema.

In addition to evaluating the ability of a d-cinema image coding to contain the gamut for film, the perceptual color differences from encoding quantization should also be examined.