Guido Schneider

Colour Grading

Opportunities and Challenges in the Transition from
Film Lab Based Colour Timing to Data-Centric Grading

HOCHSCHULE MITTWEIDA (FH)

UNIVERSITY OF APPLIED SCIENCES

Bibliographic Description:

Schneider, Guido:
Colour Grading – Opportunities and Challenges in the Transition from Film Lab Based Colour Timing to Data-Centric Grading. – 2007. –135 pages
Mittweida, Hochschule Mittweida (FH), Fachbereich Medien, Diplomarbeit, 2007

Abstract:

This paper attempts to develop a workflow for data-centric colour grading. It illustrates and develops requirements for all the steps involved and explains their effect on the colours. It also contrasts the new workflow with the traditional process of colour timing in the film lab in order to ascertain advantages and drawbacks of data-centric grading. Finally, the challenge of keeping consistent colours throughout different media and colour spaces is introduced and possible solutions expounded. The target audience are individuals working in the postproduction industry, who have a working knowledge of postproduction as a whole. However, a detailed knowledge of all stages involved in the colour grading process is not required. The secondary target audience are interested individuals who are using colour grading as a service and want to know more about the mechanisms and challenges involved.

Contents

1. Introduction

Colour grading began as the process of colour correction, where technical faults resulting in colour and brightness deviations were corrected. The film lab inherently offers possibilities to at least alter the densities of film by varying certain film processing parameters, so colour correction possibilities were available early on. As the first films were black and white, only density correction was necessary. With the invention of colour film the additional need to correct colours became apparent. The film printing machine that is used to strike a positive from a negative is used for this process. By altering the colours of the light source in the printing machine the overall colour of the positive image can be altered. Although several improvements have been made, the basic process of colour grading in the film lab has stayed the same until the present day.

With the invention of television the need for film to video transfer arose. The telecines used for this process evolved rapidly and soon offered many colour correction possibilities that were unavailable in the film lab. By utilising the full potential of the telecine colour correction evolved to a process of colour grading, where the achievement of a “look” or a visual style of the film takes precedence over a mere correction of colour faults. Until recently these possibilities were only available if a film was transferred to video tape. Films intended for cinematic release were still only able to use the film lab with its limited capabilities. However in the last decade, technologies have been developed that allow films to be completely processed in a computer, which offers even more options for colour manipulations than the telecine. This process brought with it a change in the involved challenges. In the film lab it requires great experience to assess the alterations that have to be made in order to achieve a certain result. However, the resulting colour graded film is the final end product and can therefore be evaluated with great confidence. In a computer based grading system, on the other hand, alterations to the image can be made comparatively easily and several versions of a grade can be tried. However, the image is assessed on a monitor or projector which is not a true representation of the end product. The evaluation of the result is therefore more complicated and a sizable amount of colour science has to be employed to simulate the look of film on the viewing device used. This process requires a thorough understanding, because it is critical for a good result.

To understand the challenges involved in colour grading and its present possibilities, the following questions, which also build the foundations of the next chapters, have to be answered.

- How does human vision work?
- What is the workflow in a traditional film lab and how does it influence digital colour grading today?
- What are the characteristics of a digital colour correction workflow?
- What processes do the film images undergo in a digital workflow and how does each stage influence the colours?
- What are the fundamentals of colour science and what are their implications to the practitioner in colour grading?
- How does colour management achieve a consistent look between the image viewed on a monitor and the final image on film?

1.1. Aims and Objectives

Due to the flexibility of a digital workflow, colour grading can be performed at several different stages. Hence, there are many routes one can choose to grade a film, but not all of them will offer the same quality and flexibility. This diploma thesis aims to develop a workflow in which digital colour grading can be performed with as much flexibility and confidence in accurate colour representation as possible, while achieving a high quality result.

All colour correction processes are based on human vision. Chapter 2, therefore, introduces how the human eye works and how the brain interprets its signals. It also lays the foundation for the understanding of issues relating to colour management and the representation of colours.

Chapter 3 introduces the classic film lab workflow. This helps to put the new digital colour grading workflow into perspective. Furthermore, whenever a film recording is required at the end of a digital grade, the film lab is still used to process the negative and make distribution copies. It is therefore vital for anyone involved in digital colour grading to have a thorough understanding of the workings of a film lab.

The next chapters contrast the classic film lab with a non-linear digital workflow. Chapter 4 familiarizes the reader with the characteristics of this way of working and chapter 5 gives a detailed description of the steps film images undergo and their impact on colours. Chapter 6 then introduces basic concepts of colour science that are necessary to understand one of the most important processes of digital colour grading – colour management, which is explained in chapter 7. Finally, chapter 8 gives an outlook into the future of digital colour grading.

In its core this diploma thesis is an evaluation of the non-linear digital colour grading workflow contrasted with the colour correction workflow of the film lab. However, the main focus lies on digital colour grading. Because of its flexibility there are many different ways of working and there are still debates amongst practitioners as to what is the most efficient and qualitatively best one. By means of analysing expert literature, this diploma thesis develops a workflow for digital colour grading that is both workable and yields excellent results.

The target audience for this diploma thesis are practitioners working in the field of colour grading. As colour grading becomes ever more prominent in postproduction, individuals from different backgrounds and with different levels of knowledge become involved. This paper offers a comprehensive overview of the digital as well as the classic film lab colour grading process. It requires a general understanding of postproduction as a whole. However, it recognizes the different backgrounds of readers and therefore provides a thorough explanation of all processes that are directly related to colour grading.

2. The Eye

Baylor (1995: 103) points out that "the rich colours we see are inventions of the nervous system rather than properties of light itself. Colour, like beauty, is in the eye and brain of the beholder". Hubel (1988: 33) argues that "no human inventions, including computer-assisted cameras, can begin to rival the eye". In summary, the eye is the all important final judge for every colour correction process and thus a thorough understanding of its function is vital.

2.1. Anatomy of the Eye

Although colour is created in the eye and brain, the cause for colour is the light that hits the eye. Light that falls into the eye is focussed on the retina by the cornea and the lens. The cornea provides two thirds of the bending necessary to focus the light. The lens provides the other one third, but its main task is to make the necessary adjustments to enable us to see objects in varying distances. Focussing is not done by adjusting the distance between lens and retina, but by altering the shape of the jellylike lens by means of flexing or relaxing the tendons the hold it. The amount of light that enters the eye is adjusted by opening or closing the pupil (Hubel, 1988).

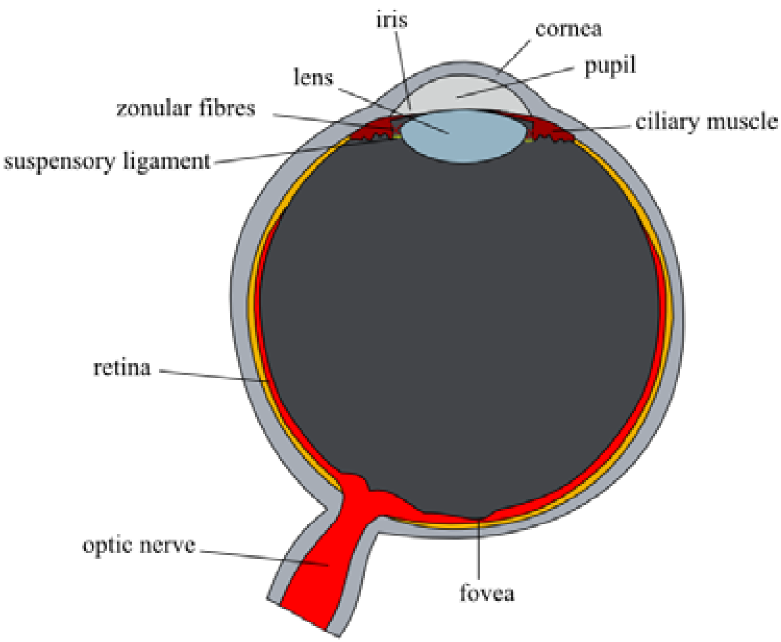

Figure 2-1 Schematic Diagram of the Human Eye

The eye is capable of seeing light with wavelengths between 400 and 700nm. The perceived hue of light changes with the wavelength, starting from violet at approx. 400nm to blue then green, yellow and then red at 700nm. Small changes in the wavelength produce noticeable colour changes. For instance, a 5% change from 550nm to 580nm changes the colour from green to yellow. It is therefore essential for colour vision to sense wavelength separately from the intensity of light (Baylor, 1995: 104).

2.2. The Retina

Light detection happens at the retina at the back of the eyeball. The retina can be regarded as an outgrowth from the brain, because it contains typical brain cells that also take over a part of the data processing (Gregory, 1998: 53). It is laid out in a rather strange way. The light sensitive cells are situated at the back of the retina facing away from the lens (Hubel, 1988: 37). As a result, the light has to pass through three layers of other nerve cells first to reach them. To put it simply “optically, the retina is inside out, like a camera film put in the wrong way round” (Gregory, 1998: 53).

There are two types of photoreceptor cells: the rods, which are long and slender and three types of cones, which are short and tapered. Rods are very sensitive. They can detect a single photon and are therefore used to see in dim light. Due to their high sensitivity they become fully saturated in normal lighting conditions and thus further increases in lighting levels do not have further effects in the rods (Hubel, 1988: 49). As a result, rods exclusively work at dim light levels. They are, however, “colour-blind“; only the cones used in normal light conditions can sense colour.

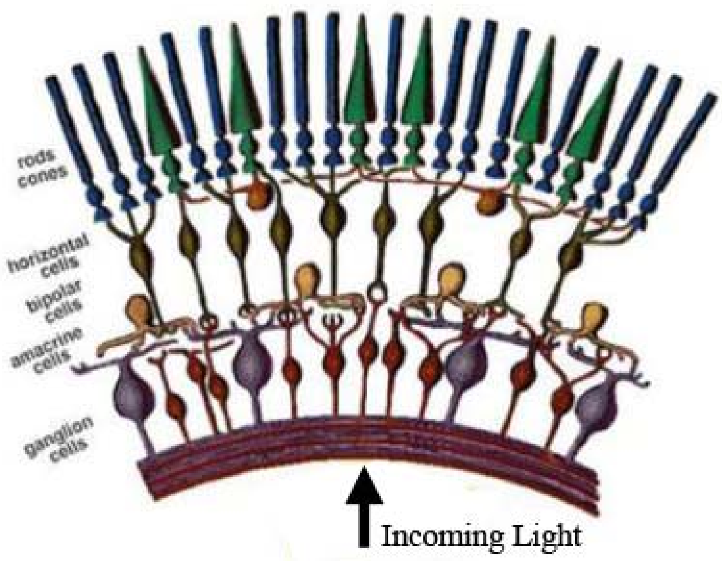

Figure 2-2 Anatomy of the Retina (Funke and Hennig, 1997)

Human vision works over a large luminance range. The range of vision covers about 8 orders of magnitude (Poynton, 2001: 196). The rods are active for the first 4 decades and vision at these low light levels is called scotopic. For the other luminance ranges the cones are active and vision is called photopic. The eye can adapt throughout the luminance range by using the pupil and also photochemical and neural processes in the eye. When the eye has reached a particular state of adaptation it can distinguish luminances across a 1:1000 range (Poynton 2001: 197).

Rods and cones are not evenly distributed over the retina. The cones are concentrated on the fovea, a spot opposite the lens with a diameter of approximately 2mm. This spot contains 100,000 cones but no rods at all. Over the fovea the other three layers of nerve cells usually covering the rods and cones are displaced to the side in order to provide maximum sharpness. This spot is covered with a brownish-yellow pigment, which protects the cells by absorbing ultra-violet and blue light (Rossotti, 1983: 112). The remainder of the retina contains 7 million cones and 120 million rods. The cones are concentrated around the fovea, but the concentration of the rods grows with increasing distance from it. There is also one spot that does not contain any light-sensitive cells at all – the blind spot, which lies at the junction with the optic nerve.

The light reaching the retina is absorbed in a visual pigment molecule, which is different for rod cells and the three cone cells. Light causes a change in the structure of this molecule and this is the only direct effect of light in the process of seeing (Baylor, 1995: 107). This transformation of the molecule triggers a chain of chemical events, which in the end causes a voltage change in the cell. The more light the pigment molecules absorb, the more negative the cell becomes. These voltage changes are detected in bipolar and horizontal cells and passed on to the brain via the optic nerve. Although the probability of absorption in rod cells depends on the wavelength of the photons, the voltage change is absolutely identical for photons of different wavelengths. Therefore, the rod cell's response only informs the brain of the quantity of photons absorbed. However, the different likelihoods of absorption mean that rods are more sensitive to certain wavelengths and less to others that is they see certain wavelengths darker. Their peak sensitivity lies in the blue-green region of the spectrum. Consequently, when a holly tree is viewed in weak moonlight the dark green leaves appear to be a lighter shade of grey than the bright-red berries, because the rods are less sensitive in the red region of the spectrum. In some cases the rods also contribute to vision in bright conditions. When an object is viewed at the edge of our visual field, it is perceived to be monochrome because the light falls on the periphery of the retina, which does not contain any cones. If the object is brought a bit further into our visual field, its light will fall onto an area on the retina which contains a few cones and rather many rods. Therefore, the perceived colour will be different from the colour perceived when the object is viewed in the centre of our visual field. At dusk rods and cones also work together to produce colours. The total cone response is strongest for light of about 550nm, but rods are most sensitive to light of about 500nm. Thus, we perceive colours turning gradually bluer before they turn grey at night.

Colour vision is achieved by the collaboration of the three cones. Each type of cone contains a slightly different light absorbing pigment that preferentially absorbs light in one region of the spectrum – namely the long, middle or short wavelengths. Therefore, these types of cones react differently to the same light stimulus and the differing cell responses are representative of the colour. Any light with ‘normal’ intensity activates the three cones in some ratio A:B:C. This ratio produces the perceived colour. For instance, if the cones are triggered in the ratio 17:44:39, the perceived colour is a shade of blue, while the ratio of 73:27:0 will give the impression of orange (Rossotti, 1983: 118). This mechanism also explains why colours can be matched by overlaying three lights sources with red, green and blue light. As long as the light of an arbitrary colour and the mixture of certain amounts of red, green and blue light excite the cones in the same ratio, the eye will see both lights as the same colour.

Our system of colour vision has evolved in different stages and as a result it consists of the overlay of an ancient and a more recent system of colour vision (Mollon, 1995: 128). Hence, today’s colour vision, instead of being a completely new and carefully developed system, is just an enhancement of a dated one. The first system discriminated between warm and cold colours and thus only had two different types of cones. The first type of cone had its peak sensitivity in the middle of the spectrum and the second one in the short-wave section. The short-wave cones were very scarce, because of the chromatic aberration which is strong for blue light. In other words, the blue component of the retinal image is permanently out of focus. Consequently, the short-wave cones were only used for colour discrimination but not for the discrimination of form or movement. This shortage of blue cones is still evident in the human eye. As a result, humans still have a poor spatial resolution for short-wavelength light. Evolution later caused a differentiation in the abundant cones that were sensitive in the middle of the spectrum. They evolved into two types of cones, one sensitive in the green region and the other one sensitive in the yellow-green region of the spectrum. Consequently we have a better spatial resolution for these two types of cones. It also has to be noted that there is no cone for the colour red as such (Mollon, 1995: 133). Red is merely characterised by a high ratio of long-wave to middle-wave cone responses.

Colour vision varies considerably between different individuals. Some of those deviations are individual while others are racial or related to age. The main reason for these deviations is that the genes for the long-wave and middle-wave cones are susceptible to defects (Rossotti, 1983: 122). For example 2% of men are dichromatic. They either lack the medium-wave pigment or the long-wave pigment. More common is a misalignment of one of the two pigments which results in anomalous trichromacy. However, the exact sensitivity peak of the cones varies even in people with normal colour vision. Mollon (1995: 137) explains: “You and I may live out our lives in slightly different perceptual worlds. Although we may both enjoy colour vision that is officially normal, coloured objects that look alike to you may look distinctly different to me, and those that look different to you may look identical to me”.

2.3. Receptive Fields and their Effect on Perception

Rods and cones are cross-connected. Their signals are interpreted by ganglion cells, which are connected to a number of rods or cones in a certain region of the retinal surface. These connections form receptive fields (Gregory, 1998: 55). Receptive fields consist of rods and cones in a circular inner region and a ring surrounding it (Hubel, 1988: 40). These two regions have the opposite effect. For example shining light on the inner region but not the outer one will increase the firing of impulses of the nerve cells they are connected to and conversely shining light on the outer region but not the inner one will decrease the rate of impulse firing. There are also receptive fields that react exactly the other way round, with an increase of impulse firing when light shines on the outer region. For both kinds of receptive fields shining light on both regions at the same time will cause the nerve cells to fire in the same pattern as when no light is falling on them at all and they are idle. This shows that our vision is much more concerned with contrast and borders than with areas of uniform brightness.

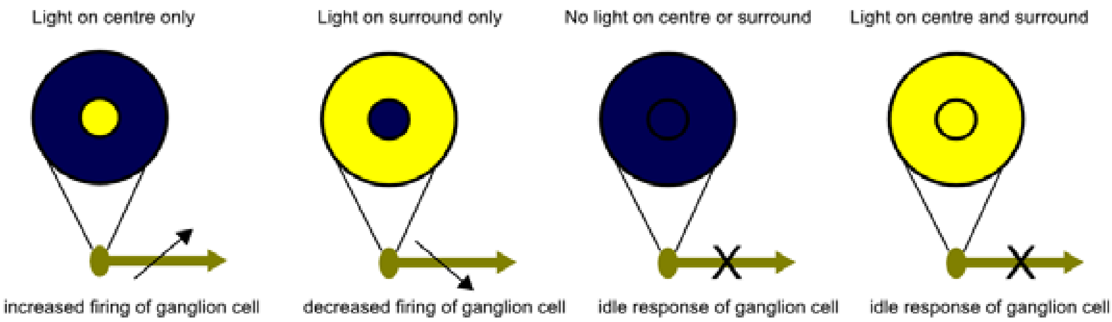

Figure 2-3 Response of On-Centre Receptive Field to Light

Neighbouring receptive fields overlap each other to a great extent and are formed of only slightly different arrays of receptors (Hubel, 1988: 43). Receptive fields are also smaller in the fovea and get progressively bigger towards the borders of the retina. As a result, acuity is greatest in the fovea. Hubel (1988: 55) argues that receptive fields explain our tendency to see brightness and colours in relation to their surround. For example, the green of a tree is likely to change if it is first seen against the blue of the sky and then against the brown of the earth. Mollon (1995: 150) concludes that our brain is not designed to analyse the true spectrum of light reaching our eye, but rather extract from it the properties of objects and their tendency to reflect one wavelength more or less than another.

One example of this behaviour is the phenomenon of colour constancy, which describes the fact that the eye can compensate for different coloured illuminations (Hubel, 1988: 178). Light with the same wavelength composition can therefore be perceived as different colours depending on the surround. This is the combination of the previously discussed physiological effect of receptive fields and a psychological one. The physiological effect is only partial, but the observer subconsciously discards any remaining colour cast. For example, a white sheet of paper will always look white, whether it is viewed in the yellowish light of a tungsten lamp or bluish daylight. However, although the eye is capable of compensating for some of those changes resulting from a different illumination, there are some colours that vary considerably when the illumination changes. For example certain mauve colours look much redder when they are illuminated with tungsten-filament lighting (Hunt, 1975: 36). Likewise, this adaptation does not work when viewing a picture of a scene. The picture is usually viewed against some background and therefore only fills a small portion of the field of view. The eye takes clues from the background to determine the correct white and not from the image. Thus, the eye does not adapt to a wrong white in a picture. One exception to this is the cinema, because its interior is black apart from the projected picture. The cinema does therefore not provide any surround to which the eye can adapt and the image also fills a much bigger area of the field of view. As a result, the eye adapts to the white point of projected films in a cinema. However, sometimes the degree of deviation from the correct white point is dependent on the tonal value or even the geometrical position in the picture. The eye only applies its correction globally, so these inaccuracies cannot be corrected by the eye. Consequently, using daylight and tungsten light in the same scene will lead to problems, even though an observer of the original scene might not notice any colour casts. The camera can only white balance on one of the two light sources, so that parts of the scene lit by the other light will appear to have a colour cast.

A second example of our tendency to rely on surrounding conditions is the eye's ability to let us see even small differences in light levels but not the absolute intensity of light falling on the retina (Hubel, 1988: 55). We can distinguish two objects if they vary in brightness by only 1%. This means that the perception of luminance is not linear but roughly logarithmic. On the other hand, judging the light levels in order to set the f-stop on a camera is a very difficult task and usually a light metre is used to help us. In other words, we perceive whether an object is black or white by judging the luminance from the scene that is surrounding it. For example, a screen of a switched-off television will look greyish. However, when it is turned on we can suddenly see intense blacks despite the fact that the smallest amount of light the television set can send is the grey we saw when it was turned off. In summary, the surround conditions influence the perceived lightness of objects. A dark background lets the object in front of it appear lighter and a light background has the opposite effect and therefore makes them appear darker (Hunt, 1975: 39). This effect is especially important for the cinema industry. Films in cinemas are viewed with a very dark surround, whereas ordinary rooms have a much lighter background. As a result, images that have been colour corrected in normal office lighting conditions will look different when viewed in a cin-

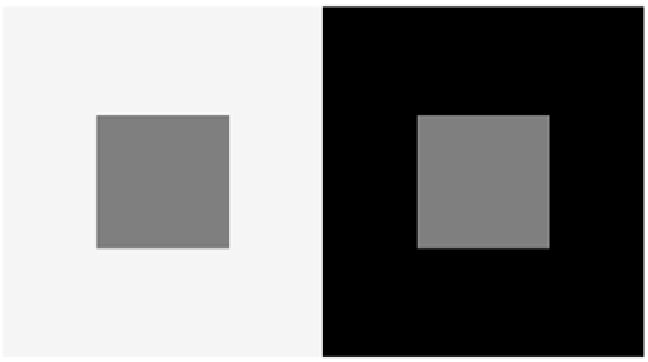

Figure 2-4 Surround Effect
The grey square on the right looks lighter than the grey square on the left, although it is exactly the same colour.

ema. Moreover, the saturation of colours varies even more than lightness. In other words, colour reproductions become more vivid when they are more intensely illuminated.

Reproductions of a colour are usually judged by comparing them to the memory of previously experienced colour sensations when looking at objects similar to the one depicted. We see objects in the natural world under a lot of lighting conditions and therefore there is a great variation in hue, lightness and saturation. Consequently, there is no exact memory of the appearance of a particular object and this allows for quite a big tolerance when depicting objects (Hunt, 1975: 43). There are, however, objects that have a considerably lower margin of error. The eye is especially sensitive to the colour of skin and most food products. Another exception of the otherwise big margin for reproducing colours is the reproduction of greys.

2.4. Psychological Implications

A certain choice of colours can influence our feelings. We often link colours to archetypical associations. They can evoke certain emotions, impressions and connotations that depend on experiences of the people of a certain culture, although usually we are not consciously aware of this process (Lewandowski and Zeischegg, 2002: 140). Green, for example, represents fertility, peace and well-being and is symbolic for quietness. Blue, on the other hand, represents the sky, thoughts and meditation and is symbolic for space and eternity. Colours can also evoke hopes and anxieties. Moreover, by combining several colours with complimenting connotations, colour associations can be considerably increased (Lewandowski and Zeischegg, 2002: 146).

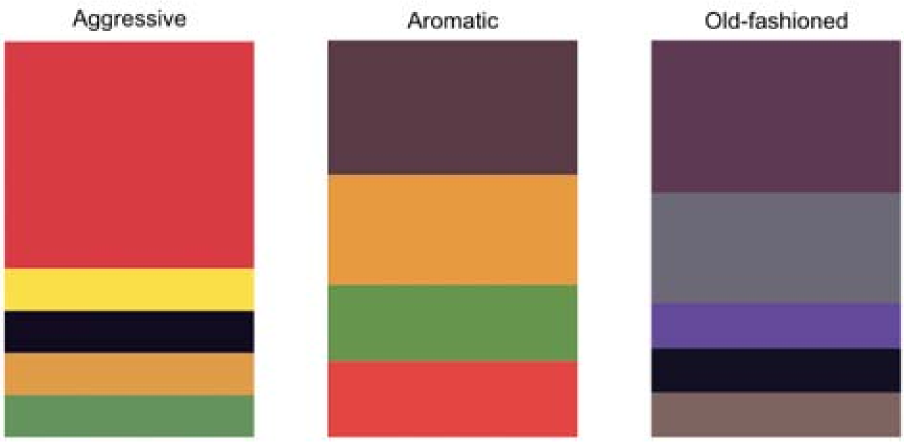

Figure 2-5 Colour Associations According to Heller (Lewandowski and Zeischegg, 2002: 147)

3. Colour Correction in the Film Lab

Before the advent of digital colour grading, film finishing was exclusively done in the film lab. Here the grader is only concerned with one medium – film – and therefore colour management is conveniently self-contained. While digital colour grading offers a lot more possibilities, the final output is still often film, which needs to be processed in the lab. Hence, the digital colour grader needs to be aware of the intricacies of this medium and its development process. As this involves various chemical reactions and film printing operations there are numerous junctures where colours are (often unintentionally) altered.

3.1. Position of the Film Lab in the Filmmaking Process

Traditionally, the film lab is positioned in the centre of the filmmaking process, providing services complementing the work of the production in the studio or on set and the editing process. It provides two main areas of operation: Front end operation and release printing. Front end operation is characterised by a close collaboration with the production and postproduction team. It can be further subdivided into rush printing and answer printing.

Rush printing encompasses all steps from the receipt of the original camera negative to the delivery of the daily rush prints to the editing room for cutting and synchronisation with sound (Happé, 1974: 4). The camera negatives are first developed and then divided into takes to be printed and the NG (no good) takes, which are not printed but set aside in case they are needed later. After development the good takes are joined together and printed by a grading operator, who determines their look for the first time. He is very dependent on the notes he receives from the DoP or the director. If he does not receive any notes it depends entirely on his judgement as to how to print the negatives. For instance, colour casts or effects that the DoP might have employed deliberately to convey the mood of the film might be cancelled out by him. As editing today takes place mostly on computer based editing systems, the negative is usually transferred to video using a telecine. In order to keep a reference to the original negative the edge numbers that are printed on it are electronically displayed in the video footage. These tapes are then digitized and the edit takes place entirely in the computer.

Answer printing takes place after the editing has finished. The editor delivers his work print to the laboratory. The edge numbers that have been printed through from the original negative are read off and a list of cuts is made up. In case of a computer based edit, the editor directly supplies a list of time-codes, which is converted to a cut list in the film lab. All the used scenes are retrieved from the negative storage and the negative cutter assembles them to produce the final film. Negative cutting has to be done by a highly experienced operator and thus, some laboratories outsource this task to a specialist company (Wheeler, 2005: 86). The information on how each of the scenes in the final film has to be printed is either taken from the way its rush prints were graded or a new grading assessment is made. A first married print is assembled and the director of photography and editor assess its quality. Any comments and changes are applied to the second answer print, which is again reviewed with the cameraman and the editor. This process might have to be repeated several times until consent on the final grading is achieved. This process is very tedious and can cost a lot of money. Wheeler (2005: 90) points out that “if too many answer prints are ordered, then beware, the producer may start to question the quality of your master negative or whether you are being just a little too fussy and consequently spending too much of their post-production budget”.

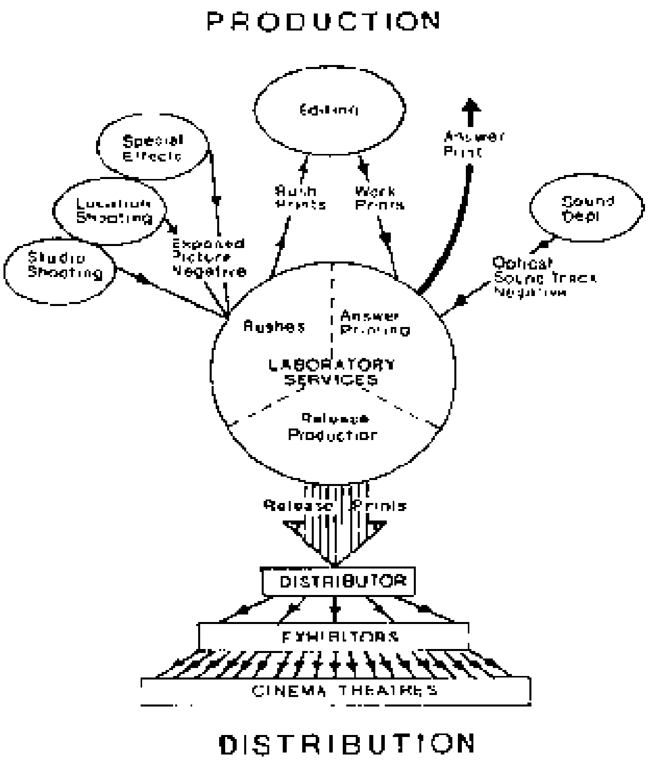

Figure 3-1 Position of the Film Lab in the Filmmaking Process (Happé, 1974: 5)

Once these creative decisions have been made, the film has to be duplicated for showings in the cinema. These release prints are often made from duplicate negatives to avoid damage to the original negative, which at this stage represents the entire budget and outcome of the production. Duplication is also useful when a change of film gauge is required (i.e. the blow-up of 16mm film to 35mm).

3.2. How Does Film Work

3.2.1. Physical Structure of Film

Film consists of light sensitive layer(s) on a transparent base. Black and white film has only one light sensitive emulsion whereas colour film has at least three. The transparent base usually consists of cellulose acetate or polyester, and a subbing layer applied on top of it helps bonding between the base and the emulsion (Corbett, 1968: 51). The light sensitive emulsion is suspended in gelatine, which has several practical properties (Wheeler, 2005: 64). It is transparent and therefore ensures that light can hit all the light-sensitive crystals no matter how deep they are embedded in the base. Also, water easily permeates the gelatine which ensures that the chemicals used during processing can reach the light sensitive crystals. Finally, it is relatively easy to produce. In addition, a thin super-coat of gelatine is regularly applied on top of the emulsion to protect it from surface friction and abrasion. Many films also have an anti-halation layer at the back of the base. This layer prevents light that has passed through the film from being reflected in the camera and re-entering the film from the back.

3.2.2. Light Sensitive Emulsions

3.2.2.1. Black and White Emulsions and Image Formation

Of all the many layers of film the light sensitive emulsion is the most important one. It is here where the photographic image is formed. Two principal ingredients form this emulsion, silver halide crystals and gelatine. The gelatine holds apart the silver halide crystals, which form the light sensitive material. The unexposed emulsion contains silver ions Ag^+ and bromide ions Br^-. These ions are atoms with an electric charge, which is caused by loosing or gaining one of their electrons. The lattice structure of the crystal is retained by the strong attraction between the two types of ion. Photons collide with the emulsion during exposure and cause them to loose their charge, which results in the formation of free bromine and sub-microscopic specks of metallic silver (Case 1985:32). Once hit by 4 or more photons a silver bromide crystal becomes capable of full conversion to metallic silver during development. These so-called development-centres are too small to be seen, but they form a latent image of the photographed scene.

The likelihood of crystals being hit by photons is increased when bigger crystals are used, as they have a bigger surface. Furthermore, the development of larger crystals will result in more silver, because the whole crystal is converted during development. As a result, coarse grained emulsions are faster, which means that they need less exposure to form a suitable image. If an emulsion only contains grains of one size, its response to light will be limited: Above a certain exposure all grains will develop and below a certain exposure no grains will develop at all. Consequently a film with this kind of emulsion will produce an image with very high contrast. Conversely, if an emulsion contains mixed grains of different sizes its response will be much more gradual: At high light levels there will still be some small grains that do not develop and even in low light situations some large grains will develop. Thus, films containing such an emulsion will have an extended range of sensitivities. Incidentally, the graininess of a film image is not due to the visibility of single grains as they are too small to be seen (they only measure 1/100 mm across), but due to the fact that grains form larger visible clumps (Case, 1985:34).

Untreated silver bromide is only affected by light below 500nm, which is blue and UV light. This problem can be overcome by introducing colour dyes that absorb light of higher wavelengths and pass its energy on to the silver bromide. However, black and white film stocks that are used for printing and duplication are not treated with dyes and are therefore only sensitive to blue light. This makes them easier to handle and is not a problem for printing. Due to the sensitivity of silver bromide all films are sensitive to X-rays and UV light. An UV filter in front of the camera and the printing machine helps to prevent it from reaching the film. However, X-ray scans at airports are still a problem for film. If they are done repeatedly the image might get foggy.

3.2.2.2. Colour Emulsions and Image Formation

Colour dyes make black and white camera film sensitive to a wider range of colours, but in colour emulsions they are used to produce emulsions that actually reproduce the different colours. In order to achieve this, three colour sensitive layers are needed. The blue-sensitive layer is located at the top of the emulsion. Underneath it is a yellow filter which absorbs all blue light. This is followed by a green-sensitive and finally a red sensitive layer. These two layers are sensitive to blue light as well, but due to the yellow filter the blue light does not reach them. All of these layers contain silver bromide crystals. They also contain a sensitizing dye which is responsible for each layer’s colour selectivity and colour couplers that will form a coloured dye image during development. During exposure, light hits the film. Blue light exposes the topmost blue-sensitive layer and is then absorbed by the yellow filter. The remainder of the light exposes the red and green layer and is then absorbed by the anti-halation layer to prevent it from bouncing back into the emulsion. Modern colour films are made up of as many as 9 layers (Wheeler, 2005: 70). There are three super-thin emulsion layers for each colour, a highly sensitive solution, a medium sensitivity solution and a slow solution. This enables modern colour films to have a latitude of 10 stops and it also yields a better image quality.

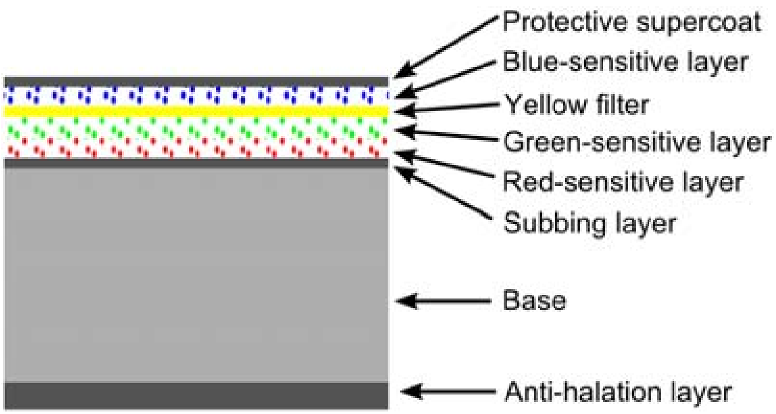

Figure 3-2 Film Structure (Case, 1985: 41)

3.3. Film Development

3.3.1. The Black and White Film Development Process

So far the image on the exposed film is not visible, but is only apparent in a slight chemical transformation of some silver bromide crystals. In the development stage this is going to change. During this process the latent image on the film is made visible. The developing agent donates electrons to the silver ions, which results in their transformation to metallic silver. The other two products of this reaction, bromine ions and oxidized developing agent, pass into the developing solution, the silver, however, is insoluble and thus stays in the emulsion. Although this reaction occurs on all crystals in the light sensitive emulsion (exposed and un-exposed), the development-centres act as a catalyst for this reaction, causing them to develop up to 200 times faster than unexposed crystals. This causes the exposed areas to be visibly darkened. The film is then passed into an acid stop bath which stops the process of development. At this stage there are still undeveloped silver halide grains in the film, which are washed away with a fixing solution. It is also important to note, that the film is thoroughly watered between the different baths to avoid contamination of the solutions. This development process creates a negative image, since all the parts of the film image that have been exposed to light become dark and the parts that have not received any light stay transparent. In order to keep this image for a long time, the film has to be washed in order for all the chemicals to be completely removed from the film.

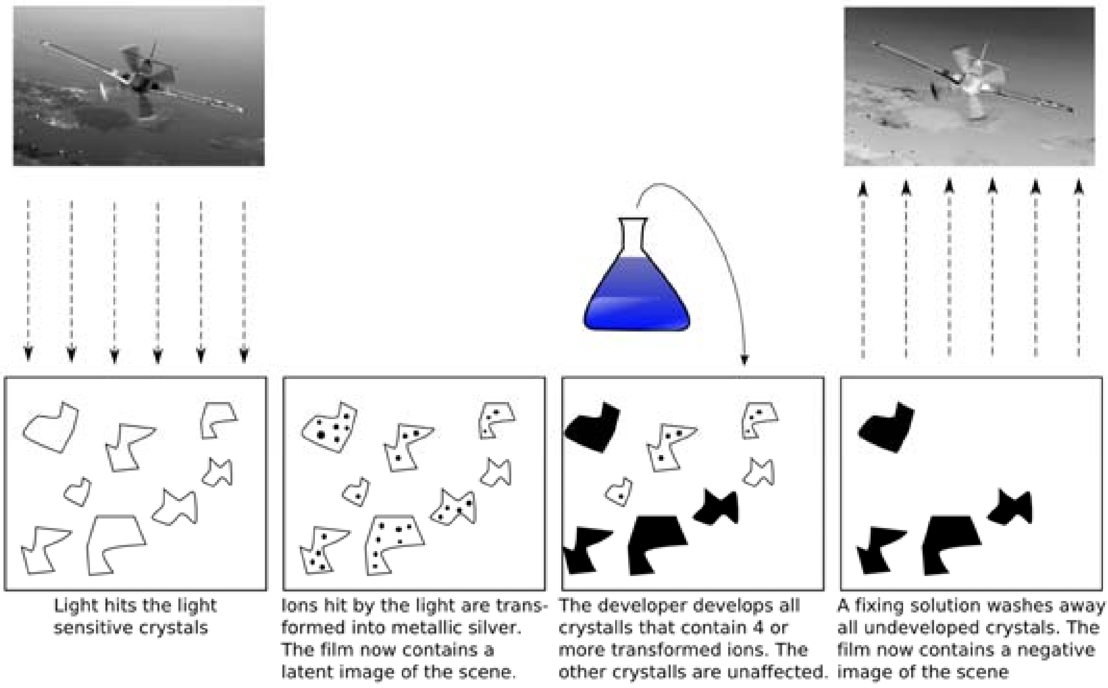

Figure 3-3 Image Formation and Development for Black and White Film

3.3.2. The Colour Film Development Process

In contrast to black and white film emulsions, each of the 3 emulsion layers in colour film contains colour couplers. In the first step of development the exposed silver grains are transformed into metallic silver similarly to black and white film. There is, however, an important difference: The oxidized developing agent formed during this reaction causes the colour couplers to produce a coloured dye. The blue sensitive layer forms a yellow image, the green sensitive layer forms a magenta image and the red sensitive layer forms a cyan image. At this point the film contains both silver and a coloured dye in the exposed areas. The film is then passed into a stop bath which stops the development. Following that, the film is put into a bleaching solution, which reverses the development process and converts the metallic silver of the exposed regions back to silver bromide. The coloured dyes are unaffected by this process. The fixing solution then dissolves away the silver bromide leaving only the coloured dyes. Similarly to the black and white development this process yields a negative image. Bright parts in the scene will be dark on the film and vice versa. Moreover, all the colours will be negative as well. For instance, blue parts of the scene will be yellow on the film and magenta parts will be green.

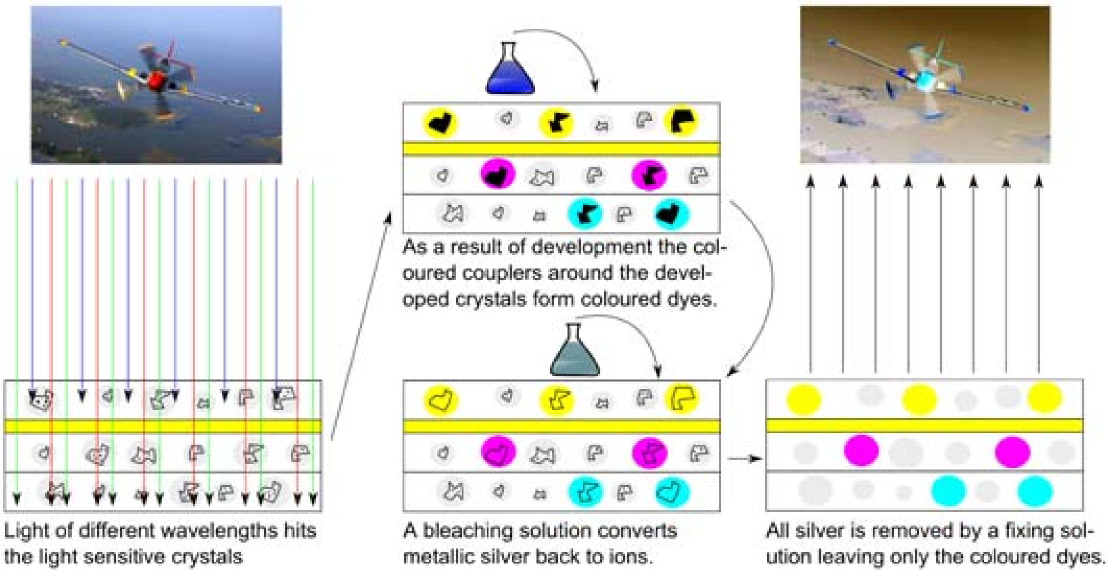

Figure 3-4 Image Formation and Development for Colour Film

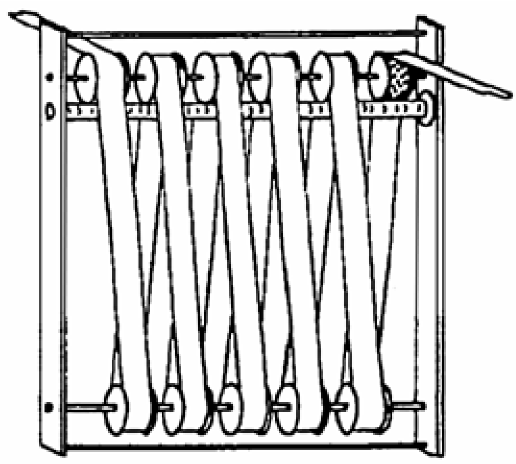

Figure 3-5 Film Processing Rack
These racks are submerged in the developing solution (Case, 1985: 53)

Film development takes place in processing machines. The film is transported over several pulleys mounted on spindles that pass the film as one continuous strip from one solution to the next. The temperature of most of those solutions has to be tightly controlled up to an accuracy of 0.1°C for the developing solution (Case, 1985:54). In addition, the time the film spends in each solution has to be precisely timed. By extending the time in the developing solution or increasing the temperature, the amount of silver produced can be increased. As this effect is greater in more exposed areas, the difference between highlights and shadows is increased and therefore the image acquires more contrast. However, the layers of colour film react slightly differently to changes in development time, which can cause a mismatch of image characteristics between the colour emulsion layers. The strong dependability of the development process on very tightly controlled conditions is also one of the factors that make this process somewhat unpredictable. As a result, the look of the developed film will depend on the exact temperature and chemical condition of the various solutions the film is passed through and the amount of time the film is kept in each one.

3.4. Characteristics of Film

From the previous paragraphs it can be seen that the formation of images on film is a very delicate process, which depends on a large number of factors. Due to the unavoidable tolerances in the development process and the medium film itself, it is paramount for film to undergo a tight quality control. To this end, several objective and measurable characteristics have been established that aid in the quality control process.

3.4.1. Density and Densitometers

Density is used to characterise exposed and developed film. It is derived from the film's transmittance, which is the amount of light that passes through a piece of film in relation to the amount of light that was incident. The reciprocal of transmittance is called opacity and if it is converted to a logarithmic measure it is denoted as density. A density of 0.3 equals a light transmittance of 50%, whereas a density of 3 equals a transmittance of 0.1%, which is effectively black. Using the logarithm is advantageous, because it coincides with the perception of luminance in the human eye. As a result, equal steps of measured density appear to the eye as equal steps of darkness.

Density is measured with a densitometer, which consists of three main parts. A light and an optical system provide a fixed illumination of the film. A measuring device, which usually is a photocell, measures the amount of light that is transmitted through the film (Corbett, 1968:100). An indicating circuit then accepts the signal from the measuring device and displays it. As the photocell produces a linear reading (that is it measures opacity instead of density), this circuit has to transform the signal to a logarithmic output, to allow for a convenient reading of the density. Due to slight variations in the electronic devices used in densitometers, it is unlikely that two densitometers will produce the same result throughout the entire range, even if they have been calibrated against the same standard. Hence, it is paramount that all measurements for the comparison of different films have to be carried out with the same device (Case, 1985:66).

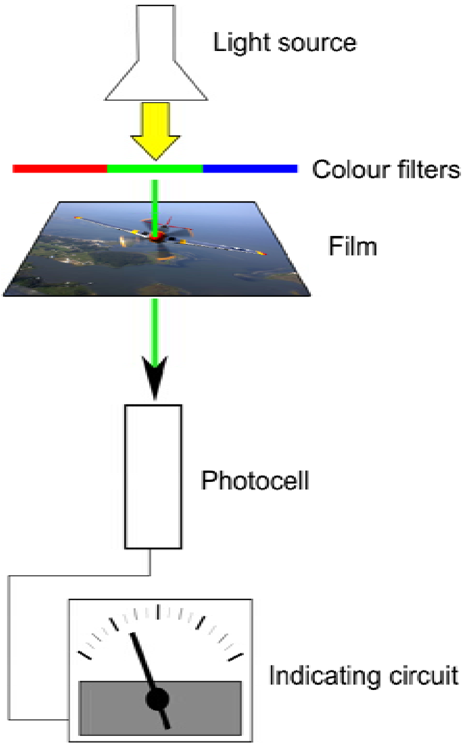

Figure 3-6 Schematic of a Colour Densitometer

For colour film each layer's density has to be read separately. Thus, three passes are required, each of which uses the primary light source of the dye layer to be measured. This prevents interference from the other two layers, because the primary light source of one layer is only absorbed by this layer and not by the other two. Nonetheless, all layers have unwanted absorptions throughout the spectrum, so all layers will add a small amount of absorption to the measured dye. Hence, the so-called integral density will be slightly higher than the analytical density that would be measured if only one dye were present. However, for routine control work this measurement is sufficiently accurate. In addition, dyes in films do not transmit all of the wavelengths in their range equally and neither do the filters that are used to colour the light used for the measurement. Consequently, densitometric results vary greatly depending on the filters used. Films intended for printing are measured with filters that have transmission characteristics similar to the sensitivities of print raw stock. These filters are called status M filters and the densities that are measured with them are called status M densities. For films intended for projection, filters that are closer to the eye in their transmission characteristics are used. These filters are called status A filters and the respective densities are called status A densities.

3.4.2. Exposure

Exposure determines how much light falls on the photographic emulsion and it is thus not actually a characteristic of the film itself. Even though, it is used to characterise how film reacts to different amounts of light. It is the product of illumination and time and therefore a change in exposure can be achieved by altering either of the two. In other words, a short exposure to bright light produces the same photographic effect as a long exposure to a light which is less intense[1]. As with other human sensations, Weber's law holds true for exposure also: A series of uniform steps is achieved by doubling or halving the exposure instead of adding or subtracting a fixed amount. Because of this behaviour, exposure is often dealt with in logarithmic values, log E. Every doubling of the exposure is represented by adding 0.3 log E.

3.4.3. Sensitometry

Different exposures will result in different densities on a given film and sensitometry is the study of this relationship (Corbett, 1968: 84). The density/log exposure diagram of a film is also called its characteristic curve, the sensitometric curve or the H & D curve (after Hurter and Driffield, who are the originators of this technique)[2]. All factors which affect this S-shaped curve are subject of the field of sensitometry. The curve shows that small exposures do not increase the density at all, but after a certain threshold exposure the density slowly rises. In the region of higher exposures the curve forms a straight line. Above a certain exposure the curve begins to flatten again. The density cannot increase anymore above this point, because all the silver has already been developed. The exposed picture should fall within the straight line section of the curve. Exposures within this region produce a faithful image, because the densities vary in exact proportion to the brightness in the original scene (Case, 1985: 74). However, many modern films have an extended toe region, where the slope of the curve although reduced is still adequate to produce good tonal variation. The result is a longer curve and image detail that penetrates deeper into the dark shadow area.

[1] However the actual relationship between time and illuminance does not exactly follow this scientific law. For extreme values of time or illuminance the emulsions become less effective. This phenomenon is called reciprocity failure. In practice it means that for a very long or a very short exposure time a greater illuminance than expected is needed for the same effect on film (Corbett, 1968:89).

[2] Log exposure is generally used instead of exposure, because in a density/exposure graph the highlights and shadows are not well represented, which makes this type of graph virtually meaningless and less useful for the cinematographer (Wheeler, 2005: 77).

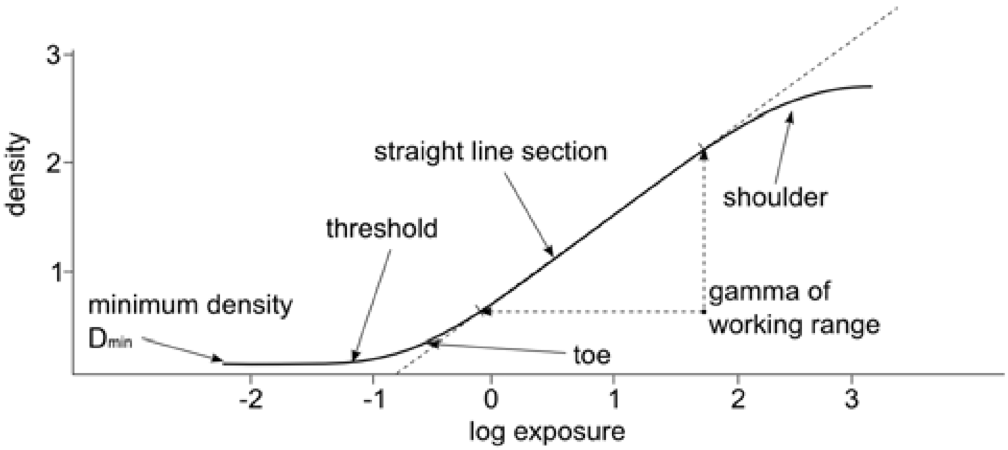

Figure 3-7 Relationship between Density and Log Exposure

In order to produce a sensitometric curve for a specific film, a series of exposures has to be recorded on that film stock. A sensitometer is used to produce a series of accurate and repeatable exposures on a strip of film. This can be achieved with a step wedge containing 21 densities with increments of 0.15 log E, which produces an exposure range of 10 stops. The resulting densities on the film are then measured with a densitometer and a graph similar to figure 3-7 that relates the imposed exposures to the measured densities can be created and analysed. As the sensitometer is used for process control it must be treated as a precision instrument, because all tests carried out in the lab assume a constant exposure from the sensitometer.

3.4.4. Gamma

In the field of sensitometry gamma denotes the gradient of the linear section of the sensitometric curve. The gamma value shows at what rate the density of a film changes when a certain change of exposure is applied. However, this is only true if the linear section of the film covers most of the working range (Corbett 1968:86). As many modern films have virtually no straight line section, an average gradient measurement is to be preferred. This is done by measuring the gradient of a line joining two points that cover the working range of the film.

Negative colour stock has a gamma of around 0.5. Hence, the gradient of the film is rather flat and a greater scene brightness range can be accommodated within the density range on the film. Conversely, colour positive stock has a very high gamma of about 3.0. The gradient is very steep and therefore a wide range of densities is produced for a small range of exposures. This produces an image with high contrast and saturated colours. The combined gamma of colour negative film which is printed onto colour positive film is 0.5 * 3.0 = 1.5. This gamma produces a screen image that is much preferred by viewers (see chapter 6.1).

The characteristic curve of an emulsion and hence the gamma varies with the conditions of development (Corbett, 1968:124). When the development time is increased the gamma increases as well. However, this only works within certain limits. If the development time is set too high the fog level will rise but the gamma will not increase any further. Other development conditions such as a variation in temperature or a modification in the concentration of the chemicals may also affect the shape of the characteristic curve. Increasing the development time is referred to as forced development. The gamma of colour negative film is only slightly affected by this process, but the fog levels increase and the amount of grain increases noticeably. The most useful result of this process is that under-exposed mid-tones are returned to their normal density.

3.5. Colour Grading on Film Printers

During filming negative stock is exposed and thus the brightness and all the colours are inverted. This negative is printed on another strip of film in order to be correctly displayed. As this process inverts the brightness and colours again, a positive is produced. Printing is the only process in the traditional film postproduction where brightness and colours and can be adjusted. Protection copies and optical effects are also achieved through printing, although colour grading is not necessarily involved in these steps.

3.5.1. Types of Printing Machines

Printing takes place in many phases throughout the postproduction process. Hence, there are different printing machines that cater for different requirements.

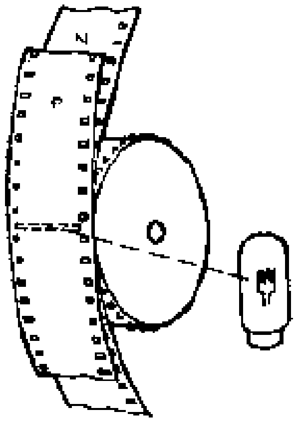

Figure 3-8 Continuous Contact Printer (Happé, 1974: 45)

In contact printers the positive and the negative are pressed together and then pass a large rotating sprocket with an illuminated slit. The light passes through the negative onto the positive film, which is exposed with the content of the negative. Due to the physical contact between negative and positive, prints made on contact printers are of a very high optical quality (Wheeler, 2005: 84). These printers can run at high speeds and are therefore suitable for bulk printing of rushes and release prints.

In optical printers the negative is lit by a lamp and the image is projected through a lens on the positive film. Usually, the film is transported one frame at a time. This machine can be used for up- or downscaling of film, for instance for the blow-up of 16mm to 35mm.

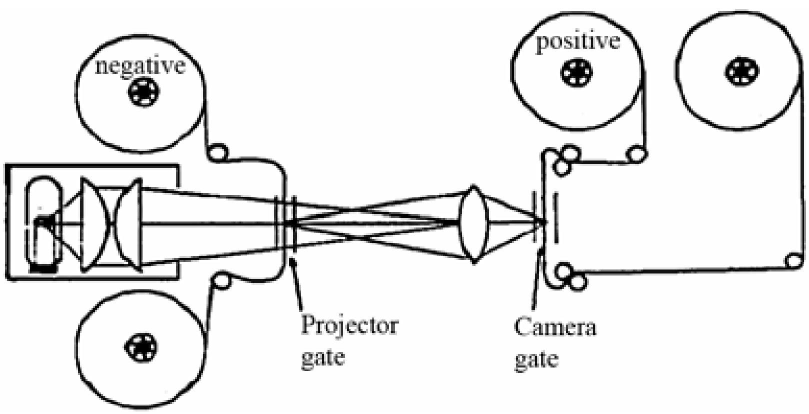

Figure 3-9 Intermittent Optical Printer (Case, 1985:99)

3.5.2. Printer Control

The light that is used to illuminate the films in the printer can be altered in colour and brightness. As a result, the brightness and colour balance of the film to be printed can be adjusted. However, only an overall change in brightness or colour is achievable. It is not possible to separately alter shadows, mid-tones and highlights or other isolated areas of the picture.

The adjustable light is achieved by using an additive light head. Here, dichroic filters divide white light into its red, green and blue components. Each of these components is then passed through a light valve before they are recombined. A light valve has 75 opening settings. The distance between them is 0.025 log E, which yields a range of just over 1.80 log E for each colour (Case, 1985:102). The opening of the valves is controlled by a combination of a manual trimmer setting and the grading data. The manually operated trimmer, which has 24 steps, is used for day-to-day printer balance adjustments. This is used to compensate for changes in stock and film processing parameters. The grading data is fed to the printer on a punched paper tape or as an ASCII file, which is prepared by the colour grading operator. He can choose between 51 grading values for each colour. A well-exposed negative should print correctly with the printer lights set close to R25, G25, B25. Additionally, colour filters can be inserted in the light beam. This is only used for major setup changes and to bring the trim controls in range.

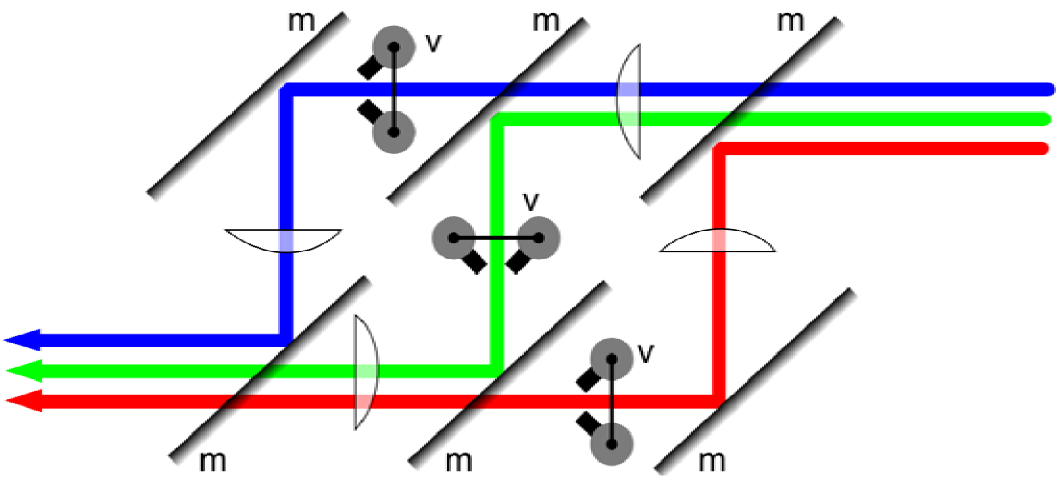

Figure 3-10 Additive Light Head with Dichroic Mirrors (m) and Light Valves (v) (Case, 1985: 103)

It is important that the printer is checked over regularly to compensate for variations in printer lamp output, transmission of lenses and sensitivity of stock. The light output of a printer can be measured with a photometer. The most common routine test, however, involves the printing of a strip of test film. This test film usually consists of an area with 18% grey and a close-up face for visual reference. Instead of the 18% grey card a LAD patch can be inserted. LAD stands for Laboratory Aim density and is a control patch that specifies densities that lie in the middle of those usually acquired for a normal camera exposure. This patch is printed at normal light onto the film that is tested. The printer is then adjusted so that a given set of densities for the printed film is achieved.

3.5.3. Colour Correction on Printing Machines

The colour grading capabilities of the printing machine are mainly used during the rush printing and the answer printing stages. During rush printing the colours of the