High-dynamic range projection using a steerable MEMS mirror array

by

Reynald Hoskinson

B.A., McGill University, 1996
M.Sc., The University of British Columbia, 2001

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

The Faculty of Graduate Studies

(Electrical and Computer Engineering)

The University Of British Columbia

(Vancouver)

December 2009

Abstract

This thesis describes a novel way to improve the contrast and peak brightness of conventional projectors by directing the light from the lamp away from the dark parts of the image towards the light parts before it reaches the projector's primary image modulator. A Microelectromechanical Systems (MEMS) micromirror array is inserted into the optical path between the lamp and the image forming element. Each element of the array can be tip/tilted, diverting light to the areas that need it most, at the expense of the darker parts of the image. In effect, this method will produce a low resolution approximation of the image on the image-forming element. The micromirror array will allow the projector to adapt its light source to the image being projected in order to maximize light efficiency and throughput. By directing the light away from the dark parts earlier in the display chain, the amount of light that needs to be blocked will be reduced, thus decreasing the black level of the final image. Moreover, the ability to dynamically allocate more light to the bright parts of the image will allow for peak brightnesses higher than the average maximum brightness of display. Although the text primarily refers to DLP-type (digital light projection) projectors, this technology will benefit all currently available projector types.

Employing such an mechanism within a projector's display chain requires contributions to a number of different fields related to displays. We studied the typical light path within a projector to determine how best to add micromirrors into the display chain. The tradeoffs between the distance on the screen that a light spot from a mirror (mobile light, or ML) could be moved, and its spatial extent were established. For a given micromirror tilt angle, the range of an ML can be increased at the expense of a larger blur kernel. Micromirrors suitable for this application were designed, simulated and fabricated. A novel way of optimizing the tradeoffs between tilt angle, mirror size, and mirror resonance frequency by splitting the mirrors into smaller functional subsections was employed. We developed several algorithms that determine favourable placement of the mobile lights from each of the micromirrors in the array, in order to best improve the image. From simulations, the projector average brightness could be increased by a factor of 1.2 if micromirrors were available that could be tilted to $\pm 3.5^\circ$ with the addition of this technology, without changing the projector lamp. If the requirement for perfect image reconstruction is relaxed, the improvement factor increases to 2.25. A prototype was system was developed that allows for fast control of mirror elements, and the positive effect of employing adaptive light distribution in this manner was demonstrated.

Contents

List of Tables

List of Figures

Nomenclature

$f/\#$ Ratio of focal length to diameter (f-stop)

I_m Original Image

I_p Image of light distribution at DMD

AMA Analog Micromirror array

CE Collection efficiency

CMP Chemical-mechanical polishing

CRT Cathode ray tube

DLP Digital Light Projection

DMD Digital Micromirror Device

DOF Depth of field

DRIE Deep reactive ion etching

FEM Finite element modelling

fL foot-lambert

FWHM Full width at half maximum

GP Gaussian pyramid

HDR High Dynamic Range

HID High Intensity Discharge

LCD Liquid Crystal Display

LED Light Emitting Diode

MEMS Microelectromechanical Systems

ML Mobile light source

PolyMUMPs Polysilicon Multi-User MEMS Process

PSF Point-spread function

PWM Pulse-width modulation

SCSi Single Crystal Silicon

SOI Silicon-on-Insulator

TMA Thin-film Micromirror Array

UHP Ultra-high pressure lamp

VDP Visual difference predictor

Acknowledgements

I would like to thank the many people who have guided me through my university career, as well as the friends and family who made it such an enriching experience.

My supervisor Dr. Boris Stoeber has been an extraordinary guide into the wonderful world of MEMS over the past four years, and an always available and patient mentor. Dr. Sid Fels has provided invaluable perspective throughout the sometimes windy road of my time at UBC. Helge Seetzen first alerted me to the possibility of overturning the conventional display paradigms and his confidence and ambition was infectious.

My parents, Michael and Gisele, and brothers and sister Emile, Paul, and Angela for all your constant encouragement.

Thanks to all those who have helped out on this project including Iman, Jeremy, Gerwin, Daniel, and Stefan. I'd like to acknowledge all my labmates over the years including Ian, Ramin, Dave, Mike, Allison, and Paul.

And finally, thank you, Caitlin Akai for your quiet and not-so-quiet support during my studies.

Chapter 1

Introduction

Electronic displays convert signal information representing graphical elements into a viewable image. They can be categorized based on how the image is viewed. The three main categories are direct view, virtual, and projection. In a virtual display, the image is in focus only on the retina, while direct view displays include the surface upon which the image is displayed as part of the device. A projection display forms the displayed image on an auxiliary surface. The images are typically meant to be viewed by groups of people at the same time.

Projectors are carefully engineered to channel as much of the light emanating from the lamp onto the image-forming element, and to relay the formed image through the projection optics to the screen. How efficient the projector is at doing this has major repercussions. Brightness is the primary characteristic determining projector price and quality, and projector efficiency is one key in determining the final brightness of the projected image. Simply increasing the brightness of the lamp to make the image brighter is not always an option. The lamp is typically the most expensive piece of the projector, even more so than the image forming element. Since the light that is not directed to the screen ends up as heat, a brighter lamp carries with it the need for bulkier and noisier fans and lamp electronics. No matter how carefully the lamp reflector and relay optics are engineered, a fundamental limitation on efficiency is encountered: the image is formed by letting light through to the screen for the bright parts of the image, and blocking the light for the dark regions. The image-forming element is evaluated by how well it passes light, and also how well it blocks it, because it is its ability to block light that finally creates the image seen on the screen.

1.1 Background

The two dominant types of projection displays available today are based on either Digital Light Projection (DLP), or Liquid Crystal Displays (LCD). To construct images, both employ different types of light valves, elements that selectively block light. A lamp provides uniform illumination to the light valve, and the liquid crystals in the LCD, for example, selectively reduce illumination of a pixel on the screen in order to form the dark parts of the image. The Digital Micromirror Device (DMD), inside a DLP projector functions in a similar manner. A DMD is an array of micromirrors, one for each pixel, each one of which can be tilted in one direction so that incident light reflects towards the

projection lens and then out onto the screen, or another direction so the light is reflected to a heat sink and that spot on the screen remains dark.

The light source in a conventional projector supplies a uniform brightness distribution on the light valve, limiting the maximum brightness for a displayed image. For most images, however, only a fraction of the total area is illuminated at peak brightness. A conventional projector simply blocks the light which is not necessary for a scene, thereby wasting this fraction of light, while it could be used to further illuminate the bright parts of the image.

Furthermore, the currently available light valves are "leaky" and cannot block all the light for black image areas (Dewald et al. 2004). The dynamic range of a projector is the ratio between the brightest and darkest levels it can display. The dynamic range, or contrast of a display device can be defined as the ratio $c = I_b/I_d$ between the brightest I_b and darkest I_d pixel generated by a projector. Merely increasing the illumination through the projector lamp in a conventional projector does not necessarily increase the dynamic range, as both the darkest and brightest level rise by the same relative amount. Current projector technology could benefit from an improvement in both dynamic range and peak brightness of the projected image.

Overall projector contrast is a combination of light valve contrast and the optical system. Pinning down current projector contrast norms is a difficult task because manufacturers often claim contrast ratios substantially higher than what is typically measurable in real world usage. Claims of ratios of up to 2000:1 are common for commodity projectors, while flat panel displays claim ratios much higher, above 40 000:1. Since no standardized, specific test has been accepted by the industry, these claims are often widely inflated. For the light valve itself, (Dewald et al. 2004) charts the evolution of the DMD contrast. The first DLP introduced into production in 1996 had 220:1 contrast, while the current 'Dark Metal 3' iteration has a contrast of 1000:1. Since 2004, contrast improvement effort has focused on the optical system (Dewald et al. 2004; Texas Instruments 2005b; Brennesholtz 1996).

The contrast and dynamic range of projectors is not sufficient to display many of the real-world scenes, which can have up to eight orders of magnitude of luminance range (Reinhard et al. 2005). Scenes such as sunsets, fireworks, and daylight must have their luminances tone-mapped to the limited range of current projectors. The peak brightness of a projector also limits the type of environment in which it can be used to its full potential; the brighter the room illumination, the more lower-end detail will be lost, lending a 'washed out' image.

The emerging popularity of 3D movies even further strain a projector's ability to display an image. Typically the efficiency of a theatre projector drops to 14% of what it is for 2D content (Brennesholtz 2009). Current digital cinema projectors "strain" to achieve 4.5 foot-lamberts (fl) of screen luminance, especially on larger screens, rather than the 14 fl recommended by the Society of Motion Picture Television Engineers in the standard ANSI/SMPTE 196M (SMPTE 2003). Even at 4.5 fl, lamps must be run at maximum power, leading to high electric bills and short lamp life. These limitations are felt by the

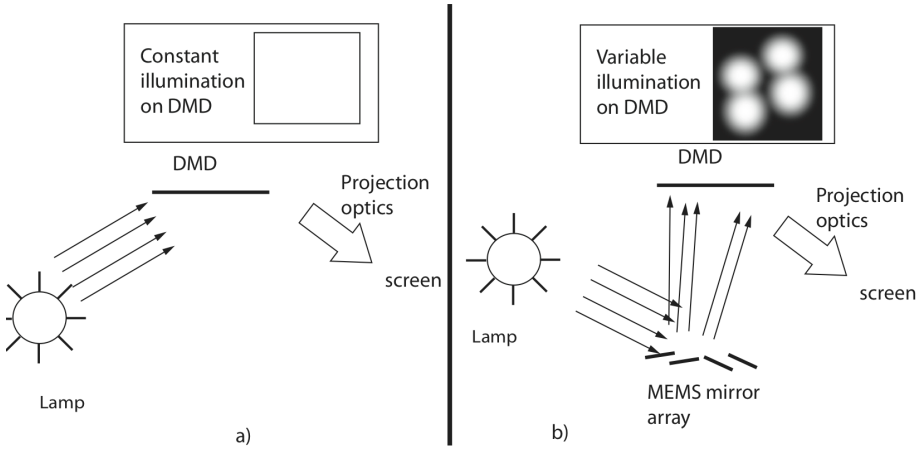

Figure 1.1: Fig. 1(a) Schematic of a conventional DLP projector, and (b) schematic of an enhanced DLP projector with second MEMS mirror array (AMA)

studios as well. Because the displayed luminances are so different, they affect the perception of colours (Brennesholtz 2009), and so studios must perform two separate expensive colour-correction processes, one for 2D and one for 3D.

Portable projectors are also appearing on the marketplace, making battery life, which is mostly dependent on the energy the lamp uses, an issue. These applications specifically, and projectors in general, would benefit from a method to use the light from the projector lamp more efficiently.

1.2 A concept for an improved projection display

This thesis details a new concept that addresses the contrast shortcomings of currently available projectors by adding a low-resolution intermediate mirror array to provide a non-homogeneous light source. This intermediate mirror device is capable of directing the uniform light from the projector lamp incident on its surface to different areas on the light valve, in effect projecting a low-resolution version of the original image onto the light valve as shown in Figure 1.1. Adding this intermediate mirror device will improve the dynamic range in two ways: by directing the light to the bright parts of the image, the achievable peak brightness will be increased. Simultaneously, the amount of light that needs to be blocked in the dark regions of the image will be reduced, thus decreasing the brightness of the black level.

This intermediate device can be realized with a low-resolution analog micromirror array (AMA), made using microelectromechanical system (MEMS) technology. The tip and tilt angle (two degrees of freedom) of the micromirrors

in the array can be set continuously in order to direct light to an arbitrary location on the light valve such as the DMD.

1.3 Other adaptive display mechanisms for projectors

In this section we describe other methods for dynamically changing illumination incident on the light valve. While these other techniques can successfully reduce the dark level of the projector under certain conditions, none increase brightness.

1.3.1 Dynamic aperture

The addition of a dynamic iris which is adjusted per frame of video is a much more limited way of adapting the projector's light source to the content (Iisaka et al. 2003) (Toyooka et al. 2005). The dynamic iris is a physical aperture near the lamp that can change size with each video frame, limiting the total illumination that reaches the screen. For dark scenes, the iris closes, limiting the light reaching the light valve, and thus increasing the contrast between images by decreasing the dark level of dark scenes while increasing illumination for brighter scenes. The image data is adjusted in real time to reflect the iris position. The system has to suppress unexpected changes in brightness, since these changes are interpreted by the human visual system as flicker.

While this method can decrease the black level of certain images, it is only a global adjustment, not allowing improving the contrast in a scene with several localized bright or dark areas. Also, while this approach decreases the black level for certain select images, it can't increase the maximum brightness of a scene like the AMA can.

1.3.2 High dynamic range projectors and displays

The concept of using two light modulators in series put forward in (Seetzen et al. 2004) has primarily been applied to flat-panel displays, but has also been applied to several projector designs to improve contrast. With the dynamic range $c_1 : 1$ of the first display, and the dynamic range $c_2 : 1$ of the second display, the theoretical contrast of the combined system is $(c_1 c_2) : 1$. Using two light modulators also results in a multiplicative increase in the number of different brightness levels. Capitalizing on the limits inherent in the human visual system's contrast sensitivity over different spatial scales, one of the modulators can be of much lower resolution than the final image.

Seetzen et al. have developed an algorithm that generates commands for each of two sequential light modulators from a given image (Seetzen et al. 2004; Trentacoste et al. 2007). The image can either be a of conventional 24-bit format, or high-dynamic range (HDR), which can represent a much higher range of luminances. The authors demonstrated the reconstruction of HDR images using two flat-panel display prototypes. In the first, a projector served as a

backlight and low-resolution modulator, illuminating an LCD panel that served as the high-resolution modulator. In their later DR-37 HDR display, an array of light emitting diodes (LEDs) are used as heterogeneous low-resolution background illumination of a conventional LCD screen, where the LCD provided the high resolution image correction (Seetzen et al. 2004). In the algorithm, the original HDR image is first downsampled to the resolution of the LED array, then the actual brightness distribution provided by the LEDs is determined by taking into account the point-spread function (physical brightness distribution) of an LED through the display optics. The high-resolution LCD correction is then generated to correct for any perceivable differences between the LED illumination and the target image.

Directly adapting these methods to projectors presents several problems. Conventional projectors use high-intensity discharge (HID) lamps as described in Section 2.3.1, which are not suitable to be placed in a tightly-packed array. Also, projector lamps typically operate efficiently only at one output intensity, so are not themselves suitable as modulators. Instead, the previous approaches described below place two subtractive modulators in series after a single projector lamp whose output does not change. While this approach can successfully reduce the dark level of the projector, it also inevitably reduces the overall brightness of the projector, because neither modulator can transmit light without loss.

(Pavlovych and Stuerzlinger 2005) added a second light modulator by focusing the light from a regular DLP projector onto an LCD panel, and then projecting the image of the LCD panel onto a display screen. Overall contrast of this projection system is improved due to the reduction in dark level by the second LCD modulator, but at the cost of a significant reduction in overall brightness, as LCD panels typically transmit only 10% or less of unpolarized incident light (Seetzen et al. 2004). Also, the extra dynamic range available in the dark regions will only be perceivable in a very dark room; the dark level increases regardless as ambient room illumination increases. Even in a moderately-lit room, the effect of the lower dark levels will thus be lost, resulting in a net decrease in overall perceivable contrast compared to an unmodified projector, because of the darker peak brightness.

(Damberg et al. 2007) similarly employ two LCD panels in series. Their prototype includes the second LCD inside the projector, and uses six panels in total – two for each colour channel. The authors report an order of magnitude increase in contrast relative to the original projector, but as in (Pavlovych and Stuerzlinger 2005), this comes at the expense of the overall brightness of the image due to the light losses from the second LCD panel.

(Kusakabe et al. 2009) have developed a high-resolution dual-modulation projector with three parallel modulators for chrominance, and one high-resolution modulator for luminance, a concept also suggested by Damberg et al. They use a LCoS device with 8192×4320 pixels for the luminance modulator, and three 4096×2160 LCoS panels for chrominance modulation. They claim a dynamic range of 1.1 million to 1, and 10-bit greyscale. Chrominance can be displayed in a much lower resolution than luminance because humans are much less sensitive

to fine detail in chrominance. The ANSI contrast ratio using a 4×4 checkerboard pattern was measured at the screen to be 250:1, which they attribute to the flare of the relay and projection lenses. As is the case with the other high-contrast projectors in the literature, the increase in contrast was solely due to the reduction in black levels, an improvement that is only relevant in extremely low brightness environments.

An AMA-projector would be particularly suited to displaying HDR images in an efficient manner. (Trentacoste et al. 2007) cite excessive power consumption as one of the reasons that Seetzen et al. moved from using a projector as a light source to an array of LEDs. The lamp of the projector had to emit enough light to illuminate the entire LCD panel at its highest possible brightness, whether it was needed or not. For high-dynamic range images that include regions with very high brightnesses, this would require a very bright lamp, especially after the efficiency of the LCD panel (16%) was taken into account. Trentecoste notes, however, that a random selection of 100 HDR image had less than 10% of the image content in the high luminance range, and the average luminance was much lower than the highest. For the images tested, the HDR projector display output a factor of between 12.5 and 4.75 too much light, depending on the image. Clearly, a projector that could adapt its spatial brightness distribution to the content of the image would be very advantageous for these types of images.

1.4 Greyscale levels

Besides peak brightness and contrast, a key attribute of any display is how many discrete steps of luminance in between the brightest and darkest settings are achievable. Conventional displays usually offer 8-bit control (256 steps) over luminance for each of the three color channels red, green and blue. While this is an adequate number of steps for conventional image formats, the much larger range of luminance of HDR displays also requires more brightness levels. Visual psychologists such as (Barten 1992) have charted the number of steps necessary to cover a given range of luminance, given the limits of the human visual system. From Barten, 962 distinct steps are sufficient for a display that can reach from $0.05 cd/m^2$ to $2700 cd/m^2$. With two modulators as proposed in (Seetzen et al. 2003), the number of steps possible in each modulator is multiplied to get the total number of luminance steps for the combined display. Two linear 8-bit devices in series theoretically provide 2562 distinct steps, which easily exceeds the requirement calculated by Barten.

An AMA-enhanced DLP projector could also provide additional grey levels from two different ways. First, multiple mobile light sources from different AMA mirrors can be directly overlapped. If only one micromirror area on the DMD is illuminated, all n AMA mirrors could divert their light to that spot, resulting in n possible grey levels from the AMA. The DMD could further modulate this light, so the maximum number of grey levels is $256n$. Secondly, the light from each AMA mirror will have an approximately Gaussian distribution on the DMD. The resulting continuously-varying intensities from the AMA mirrors will

also provide more brightness levels than would be possible using just the DMD.

There are several reasons, however, why providing extra greyscale control using the AMA is not practical. First, the intermediate brightness levels mentioned above are not fully controllable because of the limited resolution of the AMA. With an AMA for each pixel in the primarily modulator, more greyscale would be assured. For AMAs with few mirrors, however, there is not enough degrees of freedom to ensure accurate display of more greyscale levels across the entire image. The AMA we are proposing would have just 100 or fewer mirrors, due to the cost and complexity of controlling them, while the secondary modulators used in Seetzen et al. contain hundreds of individually-controllable LEDs or areas from a projector.

Another reason that added levels of greyscale are not needed is that the addition of an AMA element will not make the display radically brighter as is the case with Seetzen et al. With a baseline of 1 as the brightness a non-AMA projector would be able to show a given image, let c be the improvement factor, or how much brighter we can show the image using the same projector with an AMA added. c will depend on the content of the image, including the sum total and distribution of pixel intensities. If all intensities are within a small area in the image, the improvement factor could be large. An image where all the pixel intensities are near maximum and evenly distributed throughout, on the other hand, would not be able to to be brightened much using the AMA. For an average image, we anticipate an improvement factor of 2. This would not add enough brightness to require extra greyscale steps besides the 256 offered by the primary modulator. Because of these reasons, we will not attempt to produce more levels of greyscale than are available in the primary modulator.

1.5 Contributions

This thesis represents several novel contributions to the field of Electrical and Computer Engineering.

- A method to dynamically reallocate the light from a projector lamp from dark regions to bright regions on an image-dependent basis in order to increase projector peak brightness, contrast, and efficiency. This is also described in our paper (Hoskinson and Stoeber 2008).
- A theoretical framework for examining the tradeoffs between optical parameters that affect AMA system performance.
- The design and fabrication of an analog micromirror array suitable for an AMA projector, with composite mirrors that optimize the tradeoffs between mirror tilt angle, size, and dynamic behaviour (Hoskinson et al. 2007a; Hoskinson et al. 2007b).
- Algorithms and software implementations that allocate the mobile lights depending on the image, taking into account the physical limitations of the mirrors and projector.

- A prototype implementation demonstrating this method, showing that regions can become brighter as well as darker.

1.6 Structure of the thesis

In Chapter 2, we discuss the components that make up a projector, and the optics that need to be considered when inserting the AMA into the projector. In Chapter 3, MEMS micromirrors are discussed, and the mirrors designed and fabricated for this application are detailed. We propose algorithms that allocate the light from the AMA mirrors depending on image features in Chapter 4. In Chapter 5, the prototype is described, and results showing the proof of principle of the AMA projector are discussed. Chapter 6 concludes, and discusses potential future directions.

Chapter 2

Optical system

This chapter includes a detailed investigation of how light is propagated through a projector. Very generally, projection displays use an optical imaging system to magnify a small picture created by modulating the light from an illumination system with a two-dimensional light valve. The light valve can modulate incident light independently for each pixel in the image.

2.1 Conventional projection display systems

Looking at the light path in its entirety, all projectors can be broken down into a number of functional subsections. Figure 2.1 show the subsections of single chip DLP projectors. For purposes of clarity only the light reaching one pixel of the DMD is shown; in reality there is a ray bundle that reaches each pixel of the DMD. First in the light path, a reflector collects the light from a small arc lamp, (usually metal halide or halogen), and directs it into the illumination optics. In a single-chip DLP projector, the lamp reflector minimizes the spot size of the light at the colour wheel. After the colour wheel is the integrator, which spatially redistributes the image of the arc from a highly-peaked to a more uniform distribution with an aspect ratio that matches that of the light valve. This affects the final distribution of the light on the screen. For DLP projectors, the integrator is usually a rod, made of hollow mirrored tunnels. From the integrator rod, the light travels through relay/folding optics, which form an image of the integrator rod face on the DMD. The image of the DMD is then transmitted to the screen using a projection lens system.

2.2 Étendue

Before a detailed look at the main components of a DMD projector, we first define étendue, as it informs much of the discussion of the performance limits in projector design. Étendue has been described as "the optical engineer's version of the second law of thermodynamics"(Brennesholtz and Stupp 2008). It is the geometric capability of an optical system to transmit light. It can also refer to the optical beam itself, as a product of its divergence and cross-sectional area perpendicular to the propagation direction (Brennesholtz 1996).

Étendue can be defined as

$$E = \int\int \cos\theta dA d\Omega, \tag{2.1}$$

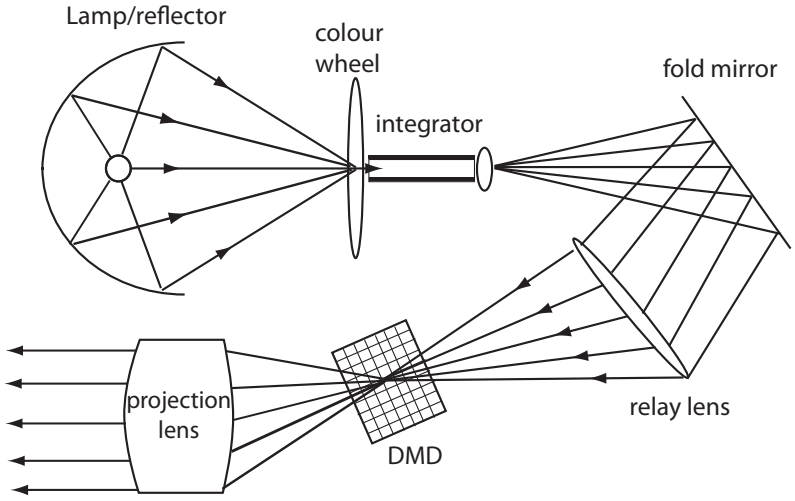

Figure 2.1: One ray bundle traversing a projection system

where E is integrated over the area of interest. The angle θ is between the centroid of the solid angle element $d\Omega$ and the normal to the surface element dA, as shown in Figure 2.2. Note that there is no term in this equation relating to optical intensity: étendue is solely a geometric property. The units for étendue are mm^2 steradian.

When a beam is modified by a well-corrected optical element, étendue is preserved. For example, when a lens focuses a beam to a spot, the area of the beam is reduced but the convergence angle of the beam increases, so étendue is preserved. The étendue of a ray bundle of light can never decrease; in an area that involves scattering, it will increase.

2.3 Projector optics

This section explains the major components of a projector, with special attention paid to the properties that constrain addition of the AMA system. We also concentrate on how optical design choices limit projector brightness and efficiency.

2.3.1 Lamp

The first element in the projector's optical path is the lamp, which today is typically an arc lamp, also referred to as an HID lamp (high-intensity discharge). A particular type of HID lamp, the UHP (ultra-high pressure) lamp by Phillips has come to dominate all but the very high-end projector market (Brennesholtz and Stupp 2008). It is usually sold as an integrated lamp and reflector with

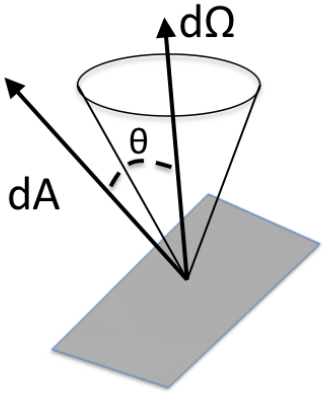

Figure 2.2: The angle θ between the area of interest A,and the centroid of the solid angle element $d\Omega$ in étendue.

a power supply designed to maximize the lamp life. Lamps typically last from 400-2000 hours. The arc length is typically 1.0 to 1.3 mm, which, along with the reflector, determines the étendue of the resulting optical beam. The smaller the arc length, the smaller the étendue, as long as the reflector can efficiently collect the luminance.

While HID lamps remain the most popular projection lighting source, LEDs and lasers have also been used. LEDs are used in small, portable projectors that can achieve up to 50 lumens, a projection market area that has seen much recent development (Conner 2006; Van Giel et al. 2007; Pan et al. 2008; Kanayama et al. 2006). LEDs have a much lower lumen output per unit of étendue than HID lamps, and so have not been used to date in many larger projectors.

Lasers have also been used as a light source for projectors. See Brennesholtz (Brennesholtz 2007) for a concise description of laser projection technologies. Because of the small divergence angles of laser beams, they have very small étendue values. A laser source also offers the possibility for much higher (10^5) brightnesses than HID lamps (Brennesholtz 2007). Lasers can also have a much narrower emission spectrum, allowing for more saturated colours and higher colour efficiency than HID lamps.

Lasers have been incorporated as projector light sources in three main ways:

1. As a raster scanner, with a laser for each primary colour, combined to one spot showing one pixel at a time. The spot moves over the image quickly enough so that each image is integrated by a human observer. The laser intensity is modulated at video rate.

2. Linear arrays such as the grating light valve (Bloom 1997), which produce the image through diffractive elements, one line at a time. A scanning mirror moves the single line through the image over time.

3. As a conventional light source, with a DMD or other light valve technology. An example is the projection television the Mitsubishi Laservue™ (Mitsubishi).

Despite being proposed for displays since 1960 (Brennesholtz 2007), lasers have predominantly been used only in very small 'pico' projectors such as from the company Microvision (Yalcinkaya et al. 2006), and large, institutional systems such as that from Evans & Sutherland (Sutherland 2009). The primary reason they have so far been kept out of commodity displays is cost per lumen: laser sources bright enough for non-pico commodity projectors continue to be commercially cost-prohibitive. While there have been attempts at producing inexpensive laser sources for displays such as from the company Novalux (Niven and Mooradian 2006), that company has since run into financial problems and has been sold. Another barrier is safety: there needs to be a solution for high-brightness front projection systems to protect from eye damage. The horizon for common laser projection continues to appear as far (or as close) as in 1960.

Multiple lamps

Using multiple lamps to create a higher-brightness image is an option both for traditional projectors and those with an AMA. The light from the lamps are typically combined via an arrangement of lenses, prisms and reflectors. Another arrangement in LED projectors such as the Mitsubishi PK10 is to have multiple LEDs of different colours combined in series using dichroic mirrors.

Multiple HID lamps are currently used in some high-end projectors where brightness requirements are paramount over cost. Using multiple light sources increases the étendue of resulting beam compared to a single source, because the effective surface area of the combined lamp and the resulting beam's divergence angle is much higher. For this reason, multiplying the number of lamps does not multiply the resulting luminance of the projector by the same factor; one published value for two lamps gave an improvement factor 1.5 (Mang et al. 2008).

Regardless of the method used to combine the light from separate lamps, the AMA could be used with the resulting beam in the same way as with the single lamp situation described in this thesis.

2.3.2 Light collection

As much of the light from the lamp is collected and channeled forwards towards the rest of the display in as small and collimated a beam as possible using a reflector. Because of aberrations, the étendue of the beam after the collection optics is much larger than the original étendue of the lamp. The main strategies for light collection fall into two categories: imaging and non-imaging optics.

In **imaging optics**, which is used in almost all non-LED projectors, an image of the arc is produced somewhere on the optical path. Two main subcategories of imaging optics for projection are critical illumination, in which the

image of the lamp is at the light valve, and Köhler illumination, in which the image of the lamp is at the entrance pupil of the projection lens. Köhler illumination has the advantage of being inherently uniform at the light valve, but is less efficient, even taking into account the separate homogenizing step needed in critical illumination. Critical illumination is thus the standard strategy for light collection in data projectors.

In **non-imaging optics**, no image of the arc is produced. The emphasis is instead on the optical beam and its aspect ratio, area, and divergence. An overview of techniques used in this branch of optics can be found in (Winston et al. 2005). There has been research into collection systems for HID-lamp projectors with non-imaging optics such as (Jacobson et al. 1998). Utilizing a non-imaging collector allowed the authors to concentrate on optimizing collection efficiency tradeoffs, in particular the tradeoff between the collection efficiency of the reflector and the resulting étendue. Efficiency in this case refers to how much of the light from the lamp is collected and sent forward to the rest of the projector. In general, the greater efficiency, the larger the étendue. All collection systems in use in HID lamp projectors today are imaging, however, because optimizing this tradeoff with non-imaging optics as they propose requires a different lamp design and reflector optics, practical implementation considerations that have kept their ideas out of consumer projectors thus far.

Non-imaging collection optics, however, are widely used with LED projectors, so perhaps there is room for some of this work to re-enter the arc-lamp world. As we shall see below, the most conceptually difficult part of incorporating the AMA into a projector is that we are adding a non-imaging optic device into a imaging-optic system.

2.3.3 Integrator

All modern projectors include some sort of integrator, such as an integrating rod or a lenslet integrator in order to homogenize the light output from the lamp/reflector, and change the circular cross-section of the lamp output beam to the rectangular aspect ratio of the light valve.

A light tunnel, which is used for DMD projectors, is a rectangular element with reflective interior surfaces. It can either be air filled and dependent on mirrors for reflection, or solid, utilizing total-internal reflection. The tunnel has the same aspect ratio as the light valve it is meant to illuminate. Light from the lamp/reflector module is focused on the entrance to the light tunnel, and each ray of light is reflected a number of times until it reaches the exit. With each reflection, each ray becomes less and less correlated with its input position, so that with an adequately-sized tunnel, the light at the output plane will be spatially uniform. The magnitude of the divergence angle (the angle between the optical axis and the ray) of each ray does not change. An image of exit pupil of the integrator is placed with a lens on the DMD at the appropriate magnification using a relay lens.

Employing an integrator incurs some light losses (Brennesholtz and Stupp 2008). An integrator rod introduces from 4-6 extra optical surfaces, each re-

sulting in a loss of beam intensity. At 0.75% loss per surface, this represents a loss of 4.5%. Another source of loss is from overfill: the light body produced by the integrator is larger than the light valve in order for the edges of the pictures to be illuminated evenly. Typically, the overfill is 5-10% per linear dimension. Overfill is needed because the intensity of the light leaving the integrator does not have an entirely flat distribution. It falls off on the edges. To obtain homogeneous illumination of the light valve, it is placed so that it occupies only the central, relatively homogeneous cross-section of the beam. Another reason for overfill is in case there is any system misalignment between the illumination path and the light valve; without overfill some pixels around the edge of the light valve would be permanently under-illuminated.

2.3.4 Light valve

The light valve in a DLP projector is a DMD, developed by Texas Instruments Inc., and manufactured by them worldwide (Hornbeck 1983) (Hornbeck 1996). DLP projectors usually have either one or 3 DMD chips as light valves. In one-DMD systems, colour is achieved by means of a rotating colour wheel. Three images are shown per frame, one for each colour. Each mirror in the DMD is approximately 16.3μm square, with a 17μm pitch, and can be positioned in two angular orientations, $\pm 12^\circ$ (10° in older versions, some newest models have $\pm 14^\circ$). Zero degrees is not addressable. Because the mirror size is fixed, the size of the array is larger for higher resolutions. For instance, a 1280×1024 pixel array is 21.8 mm $\times$ 17.4 mm.

Figure 2.3 shows a simplified version of how the individual DMD mirrors reflect light either to a heat sink or out to the projection lens. When the mirror in this diagram is set to the on state of -12 degrees, light is reflected to the projection lens (which is centred at zero degrees in the diagram's reference frame), while a +12 degree mirror tilt sends light to the 'light dump' (centred at 48 degrees), where it is emitted as heat. Figure 2.3 also shows the main source of stray light that effects contrast: flat-state light. Although the DMD is only bistable, flat-state light is still present from the integrated energy falling in the area between on and off states. This light comes partly from the transition between the on and off states of the mirrors, but since the mirrors can switch at a rate of more than 1kHz, this does not form a significant amount compared to other sources of stray light. A larger proportion comes from the reflections from the mirror vias, which attach the mirror surface to the hinge mechanism underneath it, and the gaps between the mirrors (Pinho 2003). The flat state light also includes fixed contributions from the DMD package, such as window reflectance and border metal (Texas Instruments 2005b).

The DMD itself is just an array of mirrors, which reflect light coming from any direction, without an inherent étendue. To fulfill its function as a light modulator, the incoming/outgoing light has to be limited to an angle less than that of its full tilt angle of 24 degrees (-12 to +12). That way, the light going to the heat sink can be fully separated from that going to the projection lens pupil. The étendue of the DMD is therefore defined as a product of the total

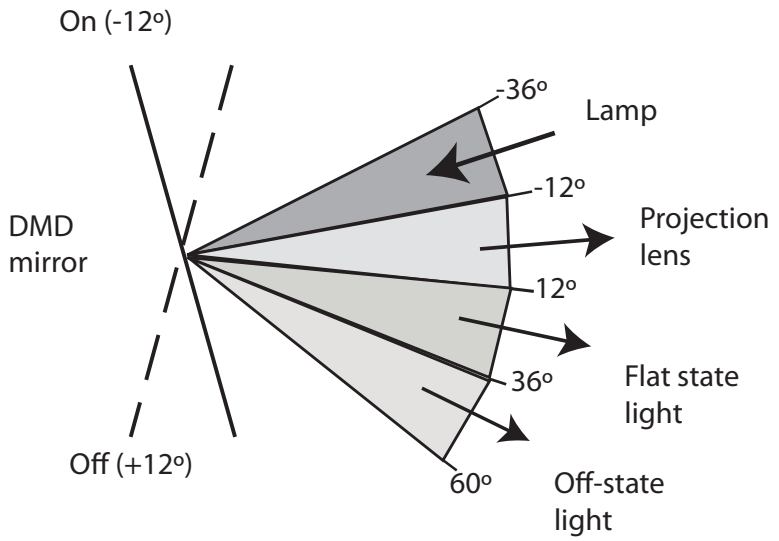

Figure 2.3: Simplified optical function Of DMD device, illustrating the constraints on input angle from the lamp to the DMD

area of the DMD and the angle of light it can accept and still fulfill its function.

The étendue of a DMD can be calculated as follows. We begin with the general case of a flat surface normal to the optical axis, and a uniform divergence angle $\theta_{1/2}$. Equation 2.1 can then be integrated in closed form as

$$E = n^2 A\Omega = n^2 A\pi sin^2\theta_{1/2}, \tag{2.2}$$

where n is the refractive index of the medium (air in this case), A the emission area, $\theta_{1/2}$ the half angle of the emission cone, and Ω the projected solid angle (Brennesholtz 1996).

The numerical aperture,

$$NA = n \sin\theta_{1/2} \tag{2.3}$$

is related to the $f/\#$ (pronounced f-stop, or f-number) of the projection lens, the ratio of its focal length to diameter D

$$f/\# = \frac{f}{D}. \tag{2.4}$$

When the NA is small,

$$NA = \frac{1}{2f/\#}. \tag{2.5}$$

In this case, the half angle is limited by the aperture of the projection lens, so the étendue can also be written as

$$E = \frac{\pi A}{4(f/\#)^2}. \tag{2.6}$$

Because we know that the acceptance angle of the projection lens pupil for a DMD projector is 24 degrees, we can work out the $f/\#$ of the projection lens as 2.4. In projectors, the NA of the projection pupil has to be small enough to prevent overlapping flat and on-state light. The relay optics use appropriate magnification to match the NA at the integrator with the NA of the projection optics. From the relay optics, the light hits the DMD, often first being reflected through a total-internal-reflection (TIR) prism. The light then travels either to a heat sink or the optical projection system, which magnifies the image to the screen while maintaining throughput and uniformity.

In a projector, smaller étendue is better for the light source, while bigger is better for all of the other optical components such as relay lenses, fold mirrors, and light valve. If the étendue of an optical element is larger than the étendue of the light beam, it is capable of using the entire beam. Usually, the light valve is the element with the smallest étendue, due to its higher cost relative to the other optical components. An optical system that crops some of the light is referred "étendue limited"(Brennesholtz 1996).

2.4 Design of the optical system of the HDR projector

The AMA must be inserted between the lamp and the DMD in order to change the illumination distribution on the DMD. Here we call the distribution of light from one mirror of the AMA at the DMD a mobile light source (ML). We would like to achieve the optimal spot size of each ML, while maintaining adequate light coverage of the DMD overall. The image of the DMD with this variable illumination will then be projected to the screen by the projection lens. As well as minimizing the AMA spot size, we would like to maximize the spot displacement for a given mirror tilt angle. However, as we shall see, these two goals are in opposition.

A conventional projector has one or more lenses between the output of the integrator and the DMD. The lenses serve to relay the image of the integrator onto the DMD. We have the option of either removing these lenses entirely when adding the AMA, or replacing them with ones that fulfill the considerations listed above. Without a lens at all, the light will blur too quickly, given that adequate physical separation is needed to separate the incoming and outgoing beams from the AMA. We explain the rate of blur below. Figure 2.4 shows the traversal of one bundle of rays as it traverses through a projector with an AMA.

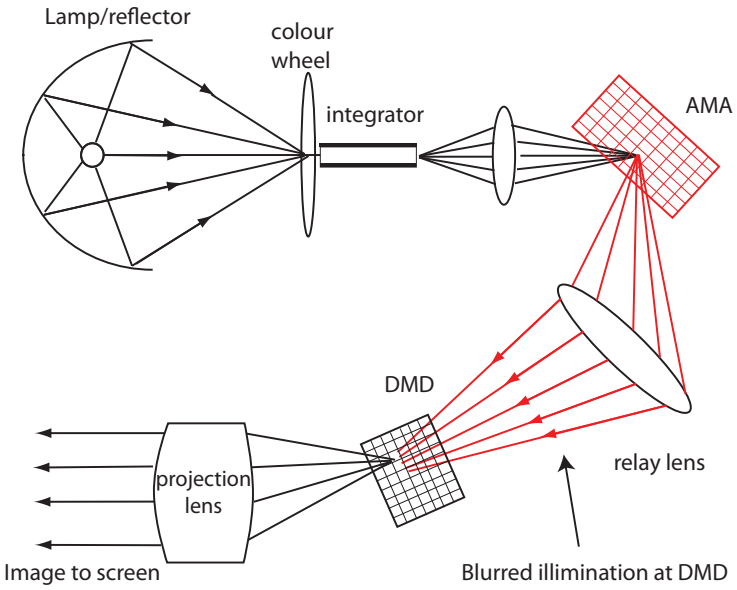

Figure 2.4: Lightpath of projector with an AMA, showing one ray bundle.

2.4.1 Circle of confusion

In photography, the range of scene depths that appear in focus in an image is called the depth of field (DOF). In this section, we adapt some of the basic calculations of depth of field to estimate what the shape of the beam from one mirror of the AMA will be when it reaches the DMD. Readers who would like a more in-depth overview of depth-of-field could consult (Jacobson et al. 2000; Hecht 2002; Nagahara et al. 2008). This is a first-order, paraxial analysis of the AMA projector system intended only to illustrate the basic relationships between the parameters of AMA size, ML size, and ML range.

A lens between the AMA and the DMD should optimally relay the light at the proper magnification onto the DMD. To determine the optical lens focal length and placement, we start with the simple lens formula (Hecht 2002)

$$\frac{1}{f} = \frac{1}{u} + \frac{1}{v}, \tag{2.7}$$

which gives a relationship between distances of object, image, and lens focal length, where f is the focal length of the lens, u the distance between the object and the lens, and v the distance between the lens and the image, as shown in Figure 2.5a. If we put the AMA at the object plane u and the DMD at the image plane v, we will get a perfectly in-focus image of the AMA on the screen, minimizing the spot size of the AMA mirror. However, since every point on the

object plane is mapped to a corresponding point on the image plane, tilting the mirrors will not move the light from one region to another when the AMA is in focus on the DMD. If anything, the light from a tilted mirror of the AMA will be blocked by the aperture of this relaying lens and not make it to the DMD at all.

To achieve the desired effect of redirecting light from one region to another, the AMA is placed at a distance d' from the object plane of the lens, as shown in Figure 2.5b. By the time the light from a point on the AMA reaches a distance u from the lens, it describes a circle c_1, which in turn is imaged onto the DMD plane, forming the circle c_2. This has the effect of blurring the image of the AMA at the DMD plane.

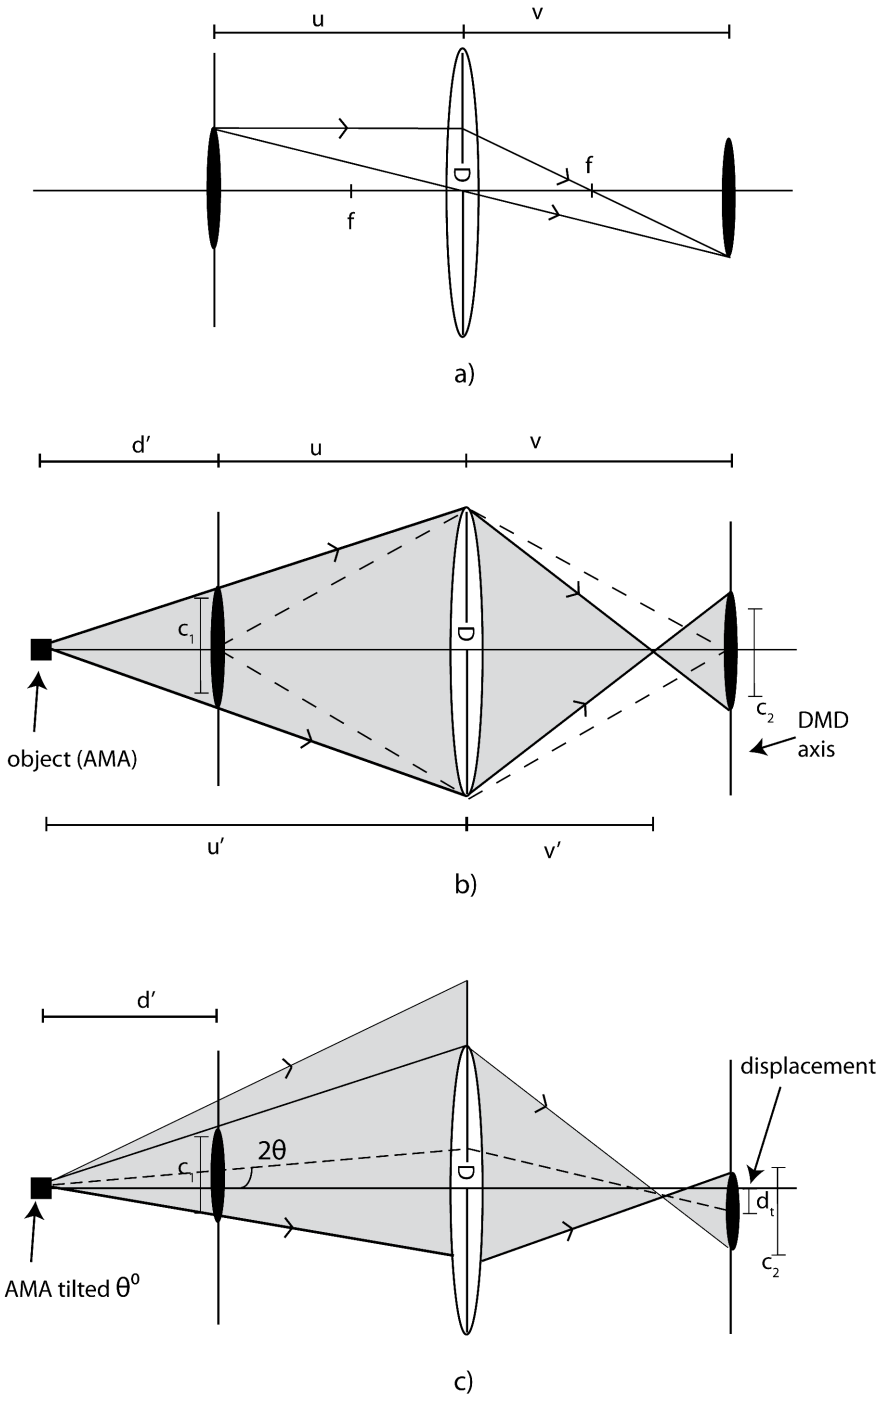

Figure 2.5: Illustration of circle of confusion. A point on the AMA spreads to a region c_1, that in turn is imaged onto the DMD as the circle c_2. In c), the mirror is tilted by θ, causing the light cone to be shifted by 2θ. The dotted line in c) represents the shifted principle ray of the cone.

Using simple plane geometry, we can show that c_1 is proportional to the lens (aperture) diameter and the separation between AMA and focal plane, and is independent of the lens focal length. It can be calculated as

$$c_1 = \frac{D}{u'}|(u - u')| = \frac{D}{u'}d'. \tag{2.8}$$

The diameter c_2 of the circle of confusion on the image side of the lens is then simply

$$c_2 = |m|c_1, \tag{2.9}$$

where m is the magnification of the lens system, obtained by the formula

$$m = -\frac{v}{u}. \tag{2.10}$$

The sign convention is that for real inverted images, m is negative, and for virtual upright images, m is positive.

The final calculation for the diameter of the circle of confusion becomes

$$c_2 = |Dm\frac{d'}{u'}|. \tag{2.11}$$

D is also the lens aperture, limiting the angle of incoming light from the AMA. To lower the rate of increase of blur diameter as the disparity increases, we could reduce the aperture of the lens. However, reducing the aperture too much will negatively affect the system efficiency of the projector. D should be of a size that makes its numerical aperture nearly equal to that of the DMD. The tradeoffs between illumination aperture and system efficiency are described in more detail in Section 2.5.

When an AMA mirror is tilted as shown in Figure 2.5c, the incident light is redirected for a distance d' before it reaches the object plane. Points on the plane at u will be imaged onto corresponding points on the plane at v, so the tilt angle of the AMA only displaces light up until d'. The displacement d_t at the DMD plane can thus be calculated as

$$d_t = md'tan(2\theta). \tag{2.12}$$

The angle of incoming light is limited by the projector's étendue, and as we described in Section 2.2, we can expect a beam divergence half angle of 12 degrees. From Equations 2.8 and 2.12 it is evident that as we increase the separation d' to increase the displacement of the ML on the DMD, the circle of confusion also grows, so we are blurring the light from the AMA mirror.

The magnification affects both the displacement and the blur. Increasing the magnification m of the AMA on the DMD would increase d', but would also increase the blur c_2 by the same factor. Likewise, if the AMA was much bigger than the DMD, $m < 1$ would decrease the blur, but also decrease d', and thus reduce the range of the ML on the DMD. Magnification is therefore not a useful way of manipulating blur in this application. Figure 2.6 summarizes how the blurred image of the AMA is imaged onto the DMD.

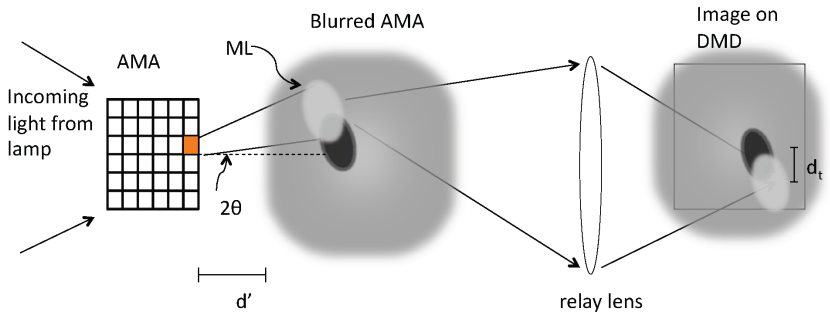

Figure 2.6: Illustration of how blurred light from an AMA mirror reaches the DMD. The AMA tilts θ, and the resulting light is redirected by 2θ. Over d' distance, this moves the ML a distance d_t.

2.5 Luminance clipping due to AMA tilt

When its mirrors are tilted, the AMA increases the étendue of the beam as a whole because the divergence angle of the beam increases. The implication is that not all light diverted from one section of the DMD to another will reach the screen. Some of it will be blocked by the DMD entrance pupil because although it would be directed to the right location spatially, some angular components will be outside of the acceptance cone of the DMD, and therefore will be clipped, not contributing to the added brightness of the projector. This is illustrated in Figure 2.7.

To determine the extent of any clipping, we must determine what effect a tilted AMA mirror will have on a ray as it travels to the DMD. As a first approximation, we trace an affected ray's path through a single thin-lens system using the matrix method described in (Halbach 1964). By representing a ray by a vector of its height r from the optical axis and the angle θ in radians the rays makes with the optical axis, we can trace the ray through an optical system by multiplying the vector with two-by-two matrices that represent discrete events such as translation, or propagation through a lens surface. A simplified situation is represented in Figure 2.8, where a ray originally travelling in the direction θ is deflected by a tilted AMA mirror by α, then travels through a lens with focal length f

$$\begin{bmatrix} r' \\ \theta' \end{bmatrix} = \begin{bmatrix} 1 & d_2 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ -1/f & 1 \end{bmatrix} \begin{bmatrix} 1 & d_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} r \\ \theta \end{bmatrix}. \tag{2.13}$$

Equation 2.13 can be simplified to

$$\begin{bmatrix} r' \\ \theta' \end{bmatrix} = \begin{bmatrix} r - rd_2/f + \theta(d_1 + d_2 - d_1 d_2/f) \\ -r/f + \theta - \theta d_1/f \end{bmatrix}. \tag{2.14}$$

With this formulation, the angular difference α' resulting from diverting the initial ray by α using the AMA is $\alpha - (\alpha d_1)/f$. We can use the Newtonian

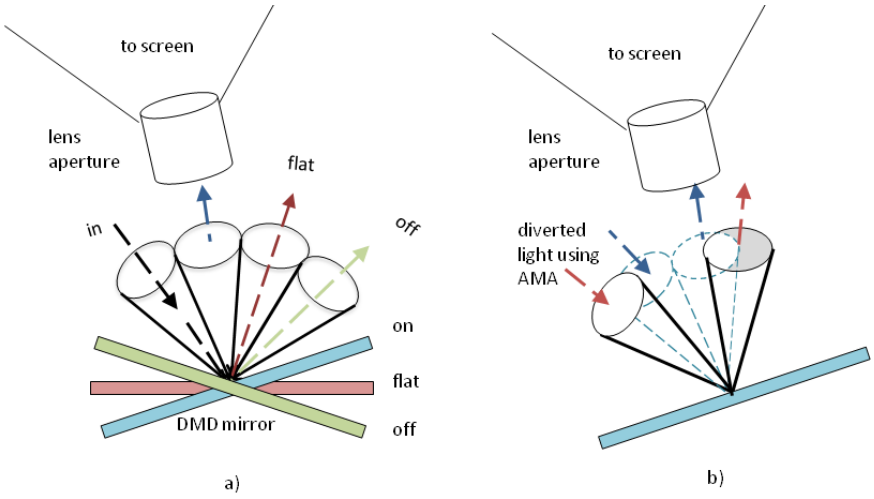

Figure 2.7: a) The regular behaviour of a DMD mirror showing the incoming light cone, the cone in the on position reaching the projection lens, and the unwanted flat and off-state light. In b) some light has been diverted using the AMA, but because the cone is now not entirely directed towards the projection pupil, not all the light makes it to the projection lens.

expression for magnification (Hecht 2002)

$$m = -f/x_o, \tag{2.15}$$

where

$$x_o = d_1 - f \tag{2.16}$$

to arrive at the expression for the change in angle of a diverted ray as

$$\alpha' = \frac{\alpha}{m}. \tag{2.17}$$

Equation 2.17 shows that if we begin with a large AMA, so that the magnification onto the DMD is less than 1, the angular change at the DMD is increased, leading to more clipping at the DMD aperture. A large magnification, on the other hand, adversely affects the circle of confusion as calculated in Equation 2.11. The desired tilt angle of the AMA therefore becomes a compromise between the flexibility of large displacements and the necessity to limit clipping.

To some degree, clipping is mitigated by the non-uniform angular distribution of the light source in a projector. Moving away from the central ray of the illumination cone, the light intensity decreases. Cutting off the edges of this cone thus does not have as big an effect as if the same area was cut off in the middle of the cone. Projector lamp research papers such as (Derra et al. 2005)

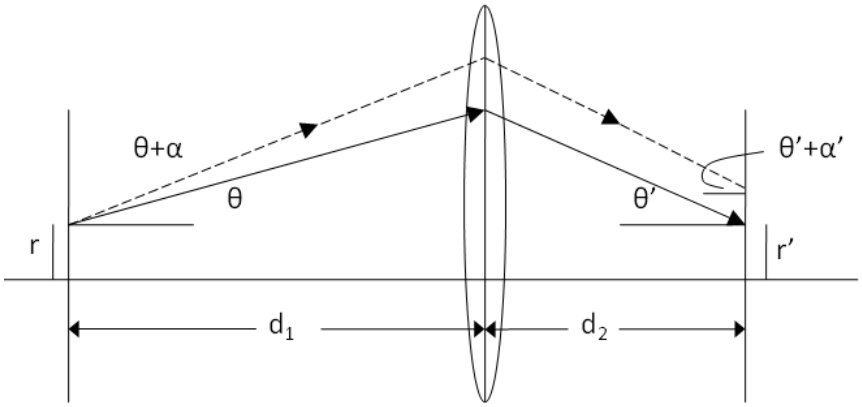

Figure 2.8: A ray of height r (solid) and angle θ is diverted by the AMA by α

include an étendue vs. collected lumens curve, showing how efficient the projection system will be at transferring the light from lamp to screen given other limiting system parameters. Heuristic descriptions of the lamp architecture can allow one to calculate the collection efficiency CE for a given étendue E and arc gap d. For the popular UHP (ultra-high pressure) lamp,

$$CE = arctan(\frac{E}{3.8d^2 + 0.9d + 0.8}). \tag{2.18}$$

The curve resulting from this equation for the typical 1mm UHP arc length is shown in Figure 2.9. The étendue ($20.17mm^2str$) of a 1024×768 pixel DMD, with dimensions of $14.08mm \times 10.5mm$ and a $f/\#$ of 2.4 (half-angle of 12°) is also included for reference as a dashed line. The estimated system collection efficiency from a lamp with this étendue is 75%. Limiting that DMD's $f/\#$ to 3.0, which is a half-angle of 9.5°, would reduce the étendue to 12.65, but the collection efficiency only 8%. This would allow the AMA mirrors, assuming a magnification of 1, to tilt the light in any direction up to 2.5° with little light loss. It would also substantially increase the normal contrast ratio of the projector.

2.5.1 Clipping quantified

To quantify the effect of a non-uniform angular light source on AMA clipping, we start by expressing étendue as defined in Equation 2.2 in terms of the cone aperture 2θ by holding the DMD area A constant, and the index of refraction $n = 1$:

$$E(\theta) = n^2 A\pi sin^2\theta. \tag{2.19}$$

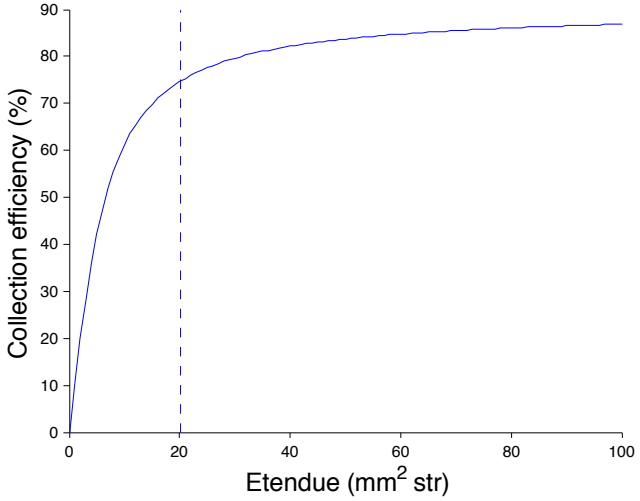

Figure 2.9: Collection efficiency of a 1mm arc length UHP lamp as a function of system étendue. The dotted line represents the étendue of a 1024×768 pixel DMD.

The lamp arc length d is also fixed, so Equation 2.18 can also now be expressed in terms of θ.

$$CE(\theta) = arctan(\frac{E(\theta)}{3.8d^2 + 0.9d + 0.8}). \tag{2.20}$$

Here the collection efficiency $CE(\theta)$ represents the percentage light collected within the cone. The light collected e_n in a concentric ring n is

$$e_n = CE(\theta_n) - CE(\theta_{n-1}) \tag{2.21}$$

For a lambertian light source, no rings are needed; the entire aperture will have uniform intensity. All others will have a different weight e associated with each ring.

At any point on the DMD, the light making it through the projection aperture from a tilted AMA mirror has an angular distribution that intersects that of the light originally meant for that space on the DMD, as shown in Figure 2.10. To find the amount of light that makes it through the aperture, we sum the areas of intersection of each weighted ring e_n and the aperture. The area A of intersection between two circles of radius r and R separated by a distance d is

$$\begin{aligned} A(r,R,d) = r^2 \cos^- 1\frac{d^2 + r^2 - R^2}{2dr} + R^2 cos^- 1\frac{d^2 + R^2 - r^2}{2dR} \\ -1/2\sqrt{(-d+r+R)(d+r-R)(d-r+R)(d+r+R)} \end{aligned} \tag{2.22}$$

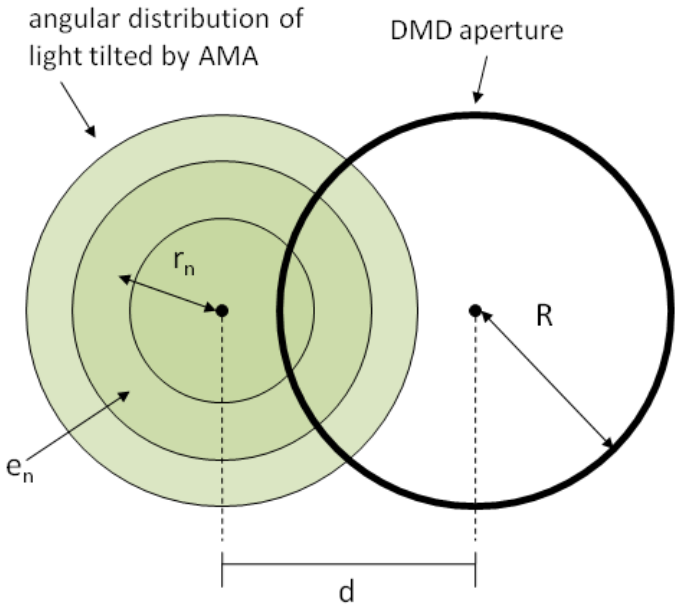

Figure 2.10: As the tilt angle d of the AMA increases, the angular distribution of the light at the DMD also changes. Only the light within the DMD aperture reaches the screen.

(Weisstein 2007). The area of intersection A_{rn} in a particular ring n is $A_n - A_{n-1}$. After accounting for overlapping areas, and weighting each area by e_r, we can determine the expected amount of clipping for a given tilt angle and light distribution. The estimated loss as a percentage of total light incident on the tilted AMA mirror is given in Figure 2.11.

For convenience, Figure 2.11 shows losses as a function of the mechanical tilt angle of the AMA mirror. If we look at the actual change in angle of the light α, it is twice that of the mechanical tilt, but the result is the same. However, α is modified by the magnification of the system as shown in Equation 2.17 before it reaches the DMD, so in a real system the x-axis will be scaled by system magnification.

As expected, the loss reaches 100% as the AMA tilt reaches the magnitude of the DMD tilt angle. At the lower degrees of tilt, the magnitude of loss is relatively low, with a 10% loss at 2.5°, a 20% loss at 3.5°, and a 34% loss at 4.5°. Care will have to be taken to allocate mirrors with these losses in mind in order to maximize the possible overall improvement using an AMA projector. They are taken into account when estimating the brightness improvement of the AMA in Chapter 4. On the other hand, in the cases where blocking light is desirable, changing the angular distribution of the light incident on the DMD would be a most effective way of blocking it.

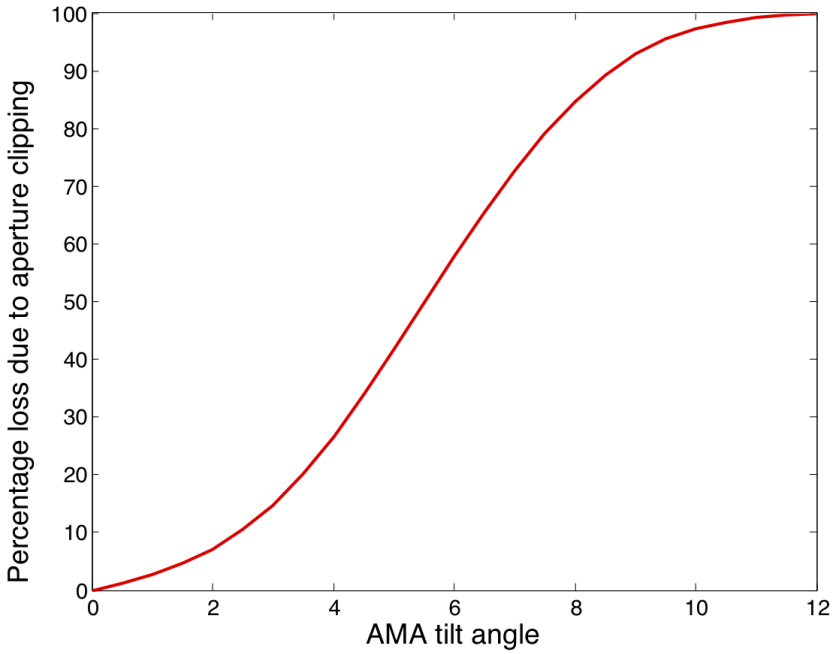

Figure 2.11: Estimated losses for AMA tilts due to the DMD aperture.

One major caveat is that these losses only need to occur if the light is diverted by the AMA towards the "off-state" light of the DMD. There is no reason for an aperture in any other direction (Janssen and Shimizu 1995). Dewald (Dewald et al. 2004) obtains improved results for both contrast and brightness when the aperture is shaped to best fit the on-state light from the DMD and block the off-state and flat-state light. A 'cat eye' aperture shape elongated in the direction orthogonal to the DMD tilt angle was found to give the best compromise between light loss and contrast increase for their system. A projector aperture specially fit to accommodate the increased aperture of the AMA would mitigate most of these clipping losses.

2.5.2 Étendue versus contrast

So far we have described only the situation where the incoming light cone fully covers the available étendue of the DMD, as in Figure 2.7a. In a real system, however, this is a worst-case scenario. Completely filling the available aperture leads to significant contrast reduction as scattering from the DMD extends the flat-state light into the on-state area, and thus becomes unwanted light that makes it through the projection pupil to the screen (Dewald et al. 2004).

To avoid this, the on-state light can be shifted in relation to the projection pupil by increasing the illumination angle, as shown in Figure 2.12. This directs the scattered light farther from the pupil, thus increasing contrast. However, the a fraction of the on-state light is also cut off, asymmetrically reducing étendue, and thus reducing the total brightness of the display.

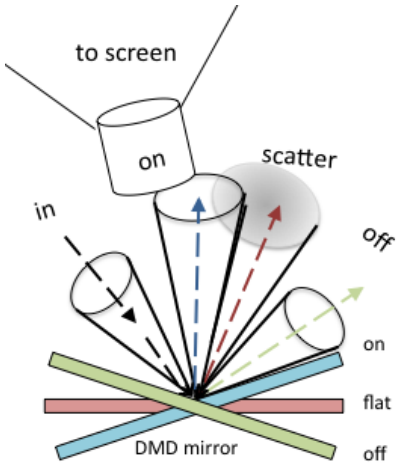

Figure 2.12: By increasing the angle of the on-state light, the unwanted scattered light can be further directed away from the projection pupil, thus increasing contrast. Some of the on-state light also is cut off, leading to a lower overall projector brightness.

The contrast ratio of the projector is significantly improved at the expense of some overall brightness loss. In an application report (Texas Instruments 2005b), Texas Instruments recommends this approach, showing that increasing the illumination angle by 2° for a 10° DMD, so that the numerical aperture is reduced by 10%, increases the contrast of a projector by 14%, while available lumens are decreased by only 4%. The losses are directly analogous to the clipping that occurs due to AMA tilt, discussed in Section 2.5.1. Our estimate loss from the clipping of 2° given in Figure 2.11 is a more pessimistic at 7%, showing that there is room for improvement in our model.

Because they assume full aperture coverage, the clipping losses quantified in Figure 2.11 would be a worst-case scenario. In an actual system, the tradeoff

between contrast and brightness can be optimized in interesting ways due to the benefits to both contrast and peak brightness of an AMA projector. Since the AMA enhances contrast by re-directing unwanted light, there is no need to artificially constrain the illumination aperture beyond the native étendue of the DMD, allowing for higher native brightness. The asymmetric case detailed in Figure 2.12 could be used to enhance native contrast, allowing the AMA to selectively redirect this extra blocked light into the projection pupil.

In summary, the key tradeoffs in pairing the AMA with light source and light valve is between the blur kernel on one side, and the system efficiency and ML range on the other. A smaller blur kernel corresponds to a combination of smaller range and/or reduced system efficiency. The tilt angle of the AMA should also be minimized to reduce extra system losses. In Chapter 4, we examine how these factors affect the possible improvement achievable by the AMA in order to determine how to balance these factors.

2.5.3 Distribution of light from AMA mirror

The distribution of light energy with the circle of confusion is called the optical point-spread function (PSF). The characteristics of the PSF are of interest to a number of different fields. Here we adapt work in the paper (Nagahara et al. 2008), which details a method to extend the depth of field of a camera by adding a microactuator that translates the detector of the camera along the optical axis, then uses deconvolution to obtain an image with an extended depth of field. The deconvolution step they do requires a estimate of the PSF size and shape, which we adapt here to our circumstances.

An idealized model for characterizing the PSF is the pillbox function is given in (Nagahara et al. 2008) as

$$p(r,c) = \frac{4}{\pi c^2}\Pi(\frac{r}{c}), \tag{2.23}$$

where c is the diameter of the circle of confusion, r is the distance of an image point from the center of the blur circle, and $\Pi(x)$ is the rectangle function, with the value 1 if $|x| < 1/2$ and 0 otherwise. (Nagahara et al. 2008) approximate the effect of aberrations by instead using a Gaussian function

$$p(r,b) = \frac{2}{\pi(gc)^2}e^{\frac{-2r^2}{(gc)^2}}, \tag{2.24}$$

where g is a constant. The parameter g between 0 and 1 has a large effect over the shape of the Gaussian, with smaller values giving a more narrow peak. This parameter should be determined through an optical simulation for a given light source, or empirically through optical measurements. It will be a function of the lamp and reflector characteristics, which determine the distribution of intensity as a function of deviation from the chief ray. We define the chief ray to be the central ray from the bundle that emanates from each point on the object. For the light sources we are working with, the chief ray will also be the ray with

the highest intensity from its bundle. Collection systems that point more of the light straight will have higher peaked Gaussian, while sources that have more of a lambertian distribution, with a gradual dropoff of intensity as the deviation angle increases, will have a broader Gaussian profile of the PSF.

Once we have come up with an estimate of the PSF of one spot on the AMA, we can estimate the distribution of light from one mirror in the array on the DMD, referred to here as the mobile light (ML) by convolving the area of the mirror with the PSF. We can then use the ML to get an estimate of the light distribution from the entire AMA as it appears on the DMD.

Let $g(x, y)$ be the ML, the luminance distribution on the DMD from one micromirror, and $f(x, y)$ be the distribution if the ML was in focus (although still with the displacement due to mirror tilt), and $p(x, y)$ be the point spread function estimated according to Equation 2.24. Then

$$g = p * f, \tag{2.25}$$

where $*$ is the convolution operator. Since convolution in the spatial domain is equivalent to multiplication in the Fourier domain, we can write Equation 2.25 as

$$G = PF, \tag{2.26}$$

where $G(\omega, \nu), F(\omega, \nu)$, and $P(\omega, \nu)$ are the Fourier transforms of g, f, and p respectively (Subbarao 1987). So to obtain the final light distribution from the AMA, we calculate the displacement and magnification of each AMA mirror, place each mirror's displaced and magnified spot into an image, and then convolve (or multiply, in the Fourier domain) the entire image by the PSF. This is the approach taken in Chapter 4.

If the DMD and AMA are parallel, and disregarding any local lens aberrations or apodization of the light incident to the AMA, the MLs will all be equal. If they are not parallel, the distribution of each ML will need to be calculated separately, because each will have a different PSF, due to the differing distances between points on the AMA and points on the DMD. For the purposes of this thesis, we assume the AMA and DMD to be parallel.

Chapter 3

Micro-electromechanical mirrors

Microelectromechanical systems (MEMS) are small devices that are fabricated with many of the same processes as integrated circuits. Often starting with a silicon wafer, materials are deposited, patterned and etched in a sequence of steps, producing a complex three-dimensional structure (Senturia 2001). Unlike an integrated circuit, MEMS devices usually incorporate mechanical structures, some of which may be free to move. Because they use IC batch-processing techniques, MEMS are usually relatively inexpensive to produce in bulk because many are fabricated in parallel.

Micromirrors have had an important role throughout the history of MEMS; they were one of the first applications of MEMS technology in general (Peterson 1982). The primary application domains for MEMS in optics, ordered by approximate precedence, have been projection displays (Hornbeck 1997), components for optical fiber communications (sources, switches, cross-connects, routers, etc.) (Bishop et al. 2002), and optical sensing and imaging (Kim et al. 2004). The optical efficiency of micromirrors can be very high, limited only by the non-reflective gaps between the individual mirrors in an array. The ratio of optical area to overall mirror area is referred to as fill-factor. A micromirror array with a fill-factor of less than one will reflect less light in total than a plain mirror with equal surface quality. Losses from the fill-factor of the AMA will take away from the brightness gains made by optical redirection. For this reason, structures the mirror needs to function must be built underneath or into the mirror itself, so that adjacent mirrors can be placed as close as possible. Mircromirror arrays have been demonstrated with a fill factor of 99% (Jung et al. 2006).

Because of their small size, accuracy of control, and speed, micromirrors are ideal for positioning the projector light. Below is a description of the principles behind actuating micromirrors, design issues which drove the development process of a custom MEMS micromirror array for implementation in a projector, and a description of the fabricated micromirrors.

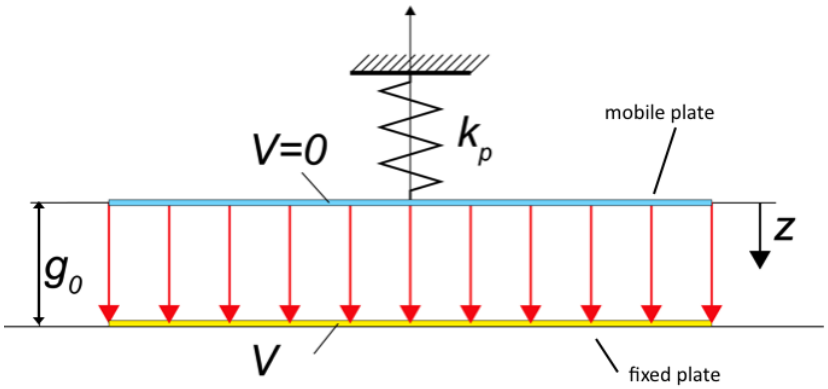

Figure 3.1: Parallel-plate actuation

3.1 Mirror positioning using electrostatic actuation

The primary mode of actuation of micromirrors has traditionally been electrostatic, due to its scalability and low power consumption. Other actuation methods include magnetic (Judy and Muller 1997) and thermal (Tuantranont et al. 2000) actuation, which are more difficult to confine, or require more power respectively compared to electrostatic actuation. When a voltage is applied between two separated surfaces, opposite charges build up in the two surfaces. These opposite charges attract, creating an attractive electrostatic force (Senturia 2001). Typically the mobile surface is attached to a spring system that provides a restoring force when the mobile surface (the mirror) approaches the fixed surface.

We can categorize the different types of micromirrors by the arrangement of springs used to separate the mirror from the substrate, and the topology of the electrostatic surfaces themselves. In this thesis we focus on three cases of electrostatic actuation: parallel-plates, torsional motion, and comb-drives, all of which have been used in micromirrors.

Parallel-plate actuation

In the simplest type, parallel-plate actuation, one fixed surface provides one electrode, and the other, mobile surface is the mirror itself, as well as the second electrode, as seen in Figure 3.1. The force pulls the mirror straight down toward the fixed electrode. This 'piston' movement can be used to change the phase of incident light (Cowan et al. 1998) for adaptive optics applications. The electrostatic force can be expressed as

$$F_p = \frac{\epsilon_0 A V^2}{2g^2} \tag{3.1}$$

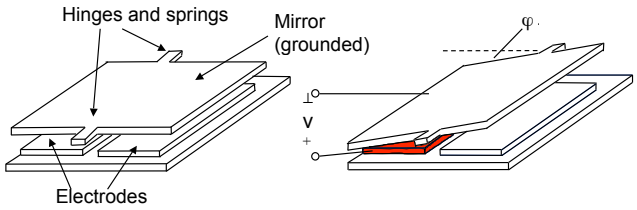

Figure 3.2: Micromirror actuation through parallel-plate electrostatics.

where ϵ_0 is the permittivity of free space ($\epsilon_0 = 8.85pF/m$ for air), A is the area of the parallel plate, V is the voltage, and g is the gap between the mirror and the electrode (Senturia 2001). The mirror is attached to one or more springs which resist the electrostatic attraction. If the gap is defined as

$$g = g_0 - z$$

where g_0 is the initial gap, and z the displacement, then the gap size is

$$g = g_0 - \frac{\epsilon_0 A V^2}{2k_p g^2} \tag{3.2}$$

where k_p is the parallel-plate spring constant. As the gap gets smaller, the electrostatic force increases according to Equation 3.1. The critical voltage at which the electrostatic force becomes larger than the restoring mechanical force is commonly referred to as the *pull-in* voltage. For voltage-controlled parallel-plate actuators, this corresponds to exactly $g = 2/3g_0$, or a displacement of one-third of the way towards the electrode. For voltages below the pull-in voltage, the system can be operated at a stable point, allowing for continuous displacement control. For voltages above the pull-in, the mirror quickly snaps to the attracting electrode, making the last two thirds of the distance to the electrode unusable for variable displacement.

Torsional motion by electrostatic actuation

Figure 3.2 shows how torsional springs suspending a one-degree of freedom mirror provide an axis of rotation, around which the mirror can tilt to one side or the other, depending on which electrode is activated. In parallel-plate actuation as described above the electrostatic field is constant over the entire mirror surface. For a tilting micromirror, the gap and hence the magnitude of the electrostatic torque changes over the mirror surface, as shown in Figure 3.3.

The electrostatic torque of a tilted micromirror can be thought of as composed of an infinite number of infinitesimally-small parallel-plate capacitors, each with a force calculated from Equation 3.1. The torque can thus be written as

$$\tau_e = x \int \frac{\epsilon_0 V^2}{2(g_0 - z(x))^2} L dx, \tag{3.3}$$

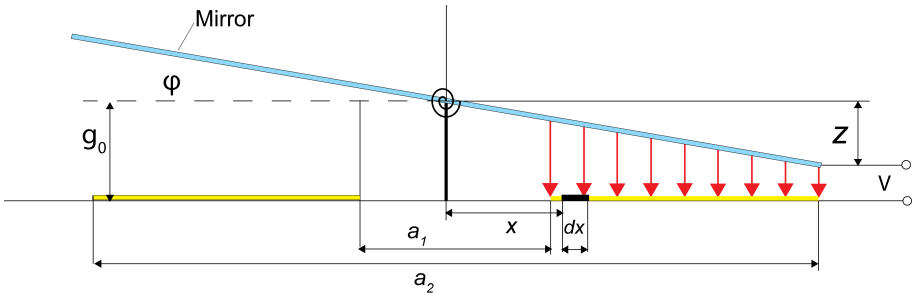

Figure 3.3: Electrostatic field applied to a mirror with torsional springs

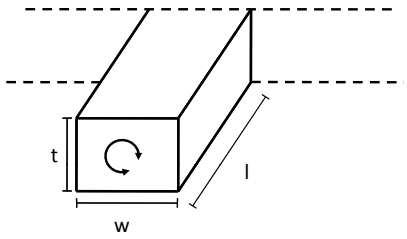

Figure 3.4: The dimensions of a micromirror spring

where x and dx represent the position and width of the infinitesimal capacitor used for integration, g_0 is the original gap between mirror and electrode, and $z(x)$, the position-dependent displacement of the tilted mirror from its original location, equal to $xtan(\varphi)$, according to Figure 3.3. L is the effective electrode length (the length of the electrode under the mirror), into the plane of Figure 3.3. See (Zhang et al. 2001) for an approximation to the numerical solution of τ_e.

The electrostatic torque τ_e causes the micromirror to rotate, which in turn causes an opposing mechanical spring torque in the beams suspending the mirror above the substrate. The torque M_t can be expressed in terms of the torsional spring constant k_t;

$$M_t = k_t \varphi, \tag{3.4}$$

φ being the mirror angle. In the case of MEMS mirrors, the shape of the springs can often be approximated by a beam with a rectangular cross-section, with one dimension, for instance the thickness t much larger than the other dimension (width w), $t >> w$. In which case k_t can be calculated analytically as

$$k_t = \frac{2Gwt^3}{3l}\left(1 - \frac{192t}{\pi^5 w} tanh(\frac{\pi w}{2t})\right), t >> w \tag{3.5}$$

where G is the shear modulus of the beam material, and l the length of the

torsion beam (Young and Budynas 2001), as shown in Figure 3.4.

For the case of a square cross-section where $t \approx w$, this can be simplified (Hibbeler 2008) to

$$k_t = \frac{t^4 G}{7.1l}. \tag{3.6}$$

The shear modulus

$$G = \frac{E}{2(1+v)} \tag{3.7}$$

is related to the Young's modulus of elasticity, E, by the Poisson's ratio, v. Equilibrium is reached when the electrostatic torque equals the opposing mechanical restoring moment

$$\tau_e = M_t. \tag{3.8}$$

Equilibrium allows the mirror to stay tilted at a certain angle, as long as the voltage is applied.

Similar to the piston-type configuration, the range of angles that the mirror can reach at equilibrium is not necessarily the full range of motion available to the mirror before it comes in physical contact with the substrate layer. While the mechanical torsion spring force varies linearly with deflection, the electrostatic force is a non-linear function of deflection, and it varies with the square of the voltage

$$\tau_e \sim v^2. \tag{3.9}$$

This means that as the applied voltage is increased, there is a point where the electrostatic force overcomes the mechanical force of the torsion springs. Theoretically, the pull-in angle or equivalently, snap-down point can be determined by looking at the tilt angle as a function of applied voltage, solving Equation 3.8 for V. The full derivation is not given here; see (Zhang et al. 2001), for example, for details.

Unlike in the parallel-plate case, where the pull-in is always $2/3g_0$, the pull-in point of this torsional system depends on the placement of the electrodes. If a is total width of the micromirror, then a_1, shown in Figure 3.3 can be defined as αa, and a_2 as βa. If the mirror angle when the mirror is tilted all the way to the electrode is φ_{max}, then the normalized angle can be defined as $\varphi_0 = \varphi/\varphi_{max}$. The solution can be calculated analytically (Zhang et al. 2001) when $\alpha = 0$ as

$$\beta\varphi_0 = 0.4404 \ (\alpha = 0, 0 \leq \beta, \theta \leq 1). \tag{3.10}$$

This equation shows that the pull-in angle depends greatly on where the electrodes are placed. α has negligible impact on this calculation because it corresponds to the electrode area at the very centre of the mirror, the area where the electrostatic torque has the smallest effect. A small value of β will mean a large tilt angle. This corresponds geometrically to the case where the electrodes are offset inwards from the edge of the mirror. If $\beta \leq 0.4404$, $\varphi_0 = 1$, and the mirror will be able to be tilted over the entire course of its range until it contacts the substrate.

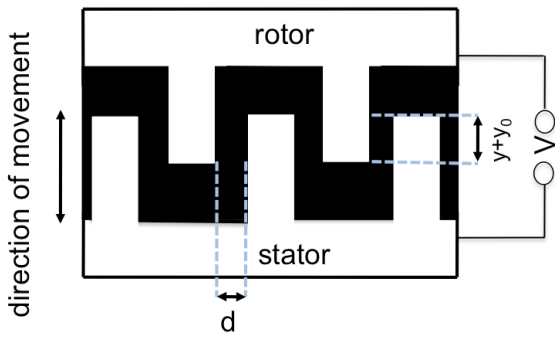

Figure 3.5: Schematic of in-plane comb drive

However, this offset removes the electrode area that creates the largest torque, increasing the voltage needed to drive the mirror. A value of $\beta = 1$ will extend the electrodes to the edge of the mirror, and make the pull-in 0.4404 of the maximum angle. As will be discussed below, we are limited to mirror sizes small enough that we cannot easily sacrifice electrode area for tilt angle, because of limits to the maximum voltage we can apply to the mirror. We can make the springs softer by lengthening them to reduce the required voltage, but that adversely affects the fill-factor if the springs are fabricated in the reflective layer. We therefore limit ourselves to the case where $\beta = 1$.

Comb drives

Another type of electrostatic actuation for micromirrors uses comb drives, which are interdigitated capacitors. Deflection of comb-like structures parallel to the plane of the substrate were originally introduced in 1989 in (Tang et al. 1989). Comb drives are typically arranged with one movable set of fingers, the rotor, and one stationary set, the stator. An electrostatic force between both sets of comb fingers engages the mobile fingers further between the stationary fingers, while the lateral distance between the fingers remains constant. Neglecting fringe fields, the electrostatic force in the direction of movement can be expressed as

$$F_c = \frac{n\epsilon_0 h}{d} V^2, \tag{3.11}$$

where V is the applied voltage between stator and rotor, ϵ_0 is the dielectric constant, n the number of fingers, h the height of the comb fingers, and d the gap between fingers (Legtenberg et al. 1996), as shown in Figure 3.5.

To apply comb drives to micromirrors, out-of-plane motion is required to produce the tilt. Several strategies have been used to deflect the micromirrors in an out-of-plane axis. Lateral actuation can be used for tilting micromirrors by using a leverage mechanism as in (Milanovic et al. 2001; Kiang et al. 1998). Another strategy is to introduce a vertical offset between the moving fingers (rotor) and the fixed fingers (stator) for out-of-plane rotation, referred to as

a vertical comb-drive actuator (Hah et al. 2004b). Because of the increased surface area compared to only two flat comb fingers of the same length, comb drives create more force for the same amount of voltage. If X is the maximum displacement for a given drive voltage, the ratio between the force of a comb drive, and the force of an equivalently-sized parallel-plate actuator F_p is

$$\frac{F_c}{F_p} = 2\frac{g_0^2}{y_0^2} = \frac{9X^2}{2w^2}, \tag{3.12}$$

where w is the minimum line width determined by the fabrication technology limit (Motamedi 2005). Especially when w is small or a large X is desired, vertical comb drive actuators can produce more force for a given voltage than parallel-plate actuation.

Theoretically, if aligned perfectly comb drives do not suffer from pull-in. Practically, comb drives always have some misalignment, which causes lateral instability, or side pull-in, which constrains motion in the same way as vertical pull-in (Borovic et al. 2006). A critical aspect of comb-drive design is the spacing between adjacent comb teeth (Krishnamoorthy et al. 2003) because the generated force is inversely proportional to the gap. The smaller the gap, however, the more dramatic the effect of misalignment, which leads to instabilities. Soft suspensions and large forces, designed to achieve large traveling range, can exacerbate this problem.

Self-alignment can mitigate this effect for vertical comb-drives: up to 98% of the theoretical maximum travel before pull-in has been reported (Krishnamoorthy et al. 2003). Comb drives tend to increase the complexity of a micromirror design because they typically require many layers in order to hide the comb drives under the mirror surface to maintain adequate fill factor. Examples of micromirrors that use comb drives include (Jung et al. 2006) which has a high reported fill-factor of 99%, but very complex fabrication with multiple Deep Reactive Ion Etching (DRIE) steps, and flip-chip bonding. Another example is (Tsai et al. 2008), which does not employ gimbals, but instead uses a cross-bar spring structure. The relevant characteristics for these mirrors are given in Table 3.1.

3.2 Micromirror design considerations

The design of the individual micromirrors which constitute the array must trade-off multiple competing considerations, such as mirror surface quality, or how optically flat and reflective the mirror is, the maximum tilt angle achievable, the actuation speed, and cost. The weighting of these parameters for a particular micromirror design depends on the intended application. For example, applications that require just one mirror, such as laser scanning, offer high tilt angles for single mirrors (Tsang and Parameswaran 2005). To achieve this, extensive use is made of the chip area around the mirror. This makes such designs unsuitable for applications where multiple closely-packed mirrors are required, such as adaptive optics.

When an array of micromirrors is considered, other design factors not present in single mirror designs become important. For arrays of mirrors, design trade-offs focus on tilt angle and fill-factor, as defined in the introduction to Chapter 3.

Another practical constraint is the nature of the fabrication steps used to create the mirror design. Custom-designed processes offer the highest flexibility, but require expensive equipment in-house, and substantial development time. Using standardized multi-user MEMS processes allows for quick and inexpensive prototyping compared to custom microfabrication processes. However, each process has its own limitations, which impose additional design constraints. Micromirror arrays that have been fabricated using standard processes include mirrors that tilt like a seesaw (Hornbeck 1983), and linear arrangements of mirrors that can tip or tilt, meaning tilt in two orthogonal directions (Tsai et al. 2004b).

To characterize the behaviour of the micromirror in motion, the micromirror can be modeled as a second-order mechanical system (Rao 2003; Conant 2002). The main aspects that pertain to micromirrors are briefly related here. For a rotational system,

$$\tau = k_\varphi \varphi + b\dot{\varphi} + I\ddot{\varphi}, \tag{3.13}$$

where φ is the mechanical mirror angle, τ the applied torque, k_φ the torsional stiffness, b the damping constant, and I_φ the polar moment of inertia around the axis of rotation. In the ideal case, this system is linear, so in response to a non-harmonic excitation, we can calculate and superimpose the responses to each frequency component of the excitation waveform. At steady-state, the dynamic response to a sinusoidal torque of magnitude τ_0 at frequency ω_d is

$$\varphi = \varphi_0 \sin(\omega_d t + \phi) \tag{3.14}$$

The resonant frequency ω is the frequency at which the amplitude of response is greatest, and is given by (Conant 2002)

$$\omega = \sqrt{\frac{k_\varphi}{I_\varphi} - \frac{b^2}{2I_\varphi^2}}. \tag{3.15}$$

Single scanning-type micromirrors that continuously tilt over a range are usually driven at resonance. For applications such as projectors, we want static control of micromirrors, where a mirror tilt angle is specified, addressed, and held until the next tilt angle is desired. The dynamic properties of the mirror are still important, however, because they govern the time it takes between when the signal is given and the mirror settles at its desired location. The higher the resonance frequency of the mirror, the faster it can be positioned.

We can approximate the shape of the mirror as a solid block in order to simplify the calculation of its moment of inertia I_φ. A block of length l_1, width w_1 and thickness t_1, rotating around the length axis can be approximated as

$$I_\varphi = \frac{1}{12} m(w_1^2 + t_1^2). \tag{3.16}$$

3.3 Previous micromirror designs

The most successful optical MEMS component to date is the Digital Micromirror Device (DMD), used in Digital Light Processing (DLP) projectors (Kessel et al. 1998), invented in 1987 and used for projectors since 1996. In most configurations, there is a one-to-one correspondence between pixels on the display and micromirrors on the DMD. The mirrors on a DMD are made of aluminum, and are approximately 16 μm square. Electrostatic actuation is used to rotate the mirror from -12° to +12°. These extremes are the only two addressable angles, so greyscale is achieved through pulse-width-modulation (PWM) – shifting between the on and off states at high rates, the ratio of which determines the brightness level.

The DMD chip cannot be used as an AMA in this application because it has no intermediate positions between its two discrete states. However, both before and after the DMD was invented there has been a substantial amount of research into variable-angle micromirrors for use in optical switching and other applications (Bishop et al. 2002; Tsai et al. 2004b; Dutta et al. 2000; Dokmeci et al. 2004; Taylor et al. 2004).

Variable-angle mirrors have many other applications besides projectors. For instance, in the telecommunication industry, transparent switching systems are a candidate to replace expensive high speed signal regeneration opto-electronics. All-optical switching requires micromirror arrays to steer optical beams from one input port to any output port with little propagation loss. The passive nature of these systems permits routing of optical signals independent of their wavelength, modulation, and polarization. One of the earliest serious ventures using scanning mirrors in the networking industry is the now-discontinued Lucent LambdaRouter (Bishop et al. 2002), also described in (Aksyuk et al. 2003). Such an array is not practical as an AMA because of the small fill-factor of the design (estimated at 53%), based on the need to only reflect multiple narrow fibre-optic beams.

Limiting previous work to mirror designs with a high fill-factor significantly reduces the field, underscoring the difficulty in designing a mirror with interesting mechanical properties that can still be tightly packed. Many current novel mirror designs feature exterior actuators and springs that dwarf the mirror itself, such as (Wang et al. 2003). While these types of mirrors illustrate interesting actuation and control ideas, they are impractical for use in an array of micromirrors.

Some mirror arrays have a high fill-factor in one direction only, such as (Taylor et al. 2004; Tsai et al. 2004a; Hah et al. 2004a). The mirrors in these configurations can be stacked tightly in one dimension, but extended components to the sides of the mirrors prevent them being stacked tightly in two dimensions.

The Thin-film Micromirror Array (TMA) developed by Daewoo, intended to be a competitor to the DMD, is one design that achieves a fill-factor of over 90% (Hwang et al. 1998). Thin-film piezoelectric actuators are used to obtain linear control over each mirror in only one tilt direction. The angle of tilt determines

the grey level, making PWM as is done with the DMD unnecessary. The TMA design seems to have been shelved, and is not commercially available.

(Tuantranont et al. 1999) attempt to increase the fill-factor of a micromirror array by using a lenslet array to collect the light from the inactive portions of the micromirrors. The authors report that these mirrors significantly improve the far-field diffraction pattern. However, while the lenslet array can increase the fill-factor, it does have other adverse effects on the overall system performance, such as introducing a limited depth of focus, chromatic aberration and distortion of the image. It also affects the complexity of the system; lenslets must be aligned to the mirrors, which increases the cost of assembling and adjusting the system.

Many mirror designs in the literature use common MEMS fabrication processes available from foundries. The Polysilicon Multi-User MEMS Process (PolyMUMPs) is probably the most-used and cited foundry process in the micromirror literature. It is a surface micromachining process that uses two structural polysilicon layers and two sacrificial phosphosilicate glass layers, as detailed in Section 3.5.1. One important property of the MUMPs process that is particularly important for mirror design is that it is a conformal process; features of underlying layers can be clearly seen though the topmost layers. Those particularly concerned with optical properties must take this effect, also known as "print-through" into account when considering where to place the electrodes that will drive them.

(Comtois et al. 1995) designed a hexagonal micromirror for the MUMPs process, with mirrors 50 μm across, and arranged with 75 μm centre-to-centre spacing between nearest micromirrors. The mirrors had a piston mode for adaptive optics applications such as wavefront correction, with one big electrode underneath that pulls the entire mirror down. A later paper (Comtois et al. 1999) lists many of the considerations a MEMS designer wishing to maximize fill-factor has to deal with, some of them specific to MUMPs but most generalizable to all surface-micromachining processes.

Sandia Lab's SUMMiT-V is another surface micromachining multi-user process. Besides having two more polysilicon layers than MUMPs (5 instead of 3), the SUMMiT-V also employs chemical-mechanical polishing (CMP), which allows for, among other things, flatter surface features because the polishing reduces the effect of "print through." (Cowan et al. 1998) provide a comparison of fill-factor achieved on various mirror designs using the MUMPs and Sandia SUMMiT-V process. The paper found that the topography of the mirrors is at least as important to the overall reflectance properties of a micromirror array as the fill-factor. A 2D micromirror array with a relatively high fill-factor built using the SUMMiT-V process is detailed in (Tsai et al. 2004b). The authors claim relatively large continuous scan angles ($\pm 4^\circ$ and $\pm 3.4^\circ$ degrees for the two orthogonal degrees of freedom), with a fill-factor of 96%. To achieve this, the torsional springs are placed under the mirrors. The electrodes are terraced upwards towards the middle of the mirror, ostensibly to reduce the actuation voltage. The terraced electrodes are made from the bottom four polysilicon layers; the top layer is reserved for the mirror. Maximum scan angle is achieved at over 90V, which is relatively high compared to many other designs.

A variation on (Tsai et al. 2004b) by the same authors uses a ring of comb drive sections (Tsai et al. 2008), along with a cross-bar arrangement of torsional springs to achieve 2DOF actuation without gimbals. The springs and comb drives are all hidden below the mirror to maintain fill-factor. They achieve tilt angles of $\pm 5.4^\circ$ at 42 V, and $\pm 2.3^\circ$ at 61 V for rotations about the x and y axis, respectively.

(Dagel et al. 2006) describes a hexagonal tip/tilt/piston mirror array with an array fill-factor of 95%. The micromirrors tilt using a novel leverage mechanism underneath the mirror. The base of the lever is attracted to a electrode on the substrate, pulling the portion of the lever on the other side of a torsional spring upwards. Bumper features on the substrate keep the leverage mechanism operating as the torsional springs bend at high voltages by acting as a pivot.

Many tip/tilt mirror systems use gimbals to suspend the mirrors, such as (Bishop et al. 2002; Lin et al. 2001; Wen et al. 2004). Usually, a frame surrounds the mirror and is attached to it by two torsional springs, forming an axis of rotation. The frame itself is then attached to the surrounding material by two springs in orthogonal directions, allowing the mirror to tip and tilt on torsion springs.

Although the basic structure of the 2-DOF mirror using gimbals is simple, there has been a fair amount of literature on the optimal design of these mirrors, such as the placement of the electrodes underneath the mirror/gimbal structure. Most gimbal designs in the literature use four electrodes, some of which have pairs of electrodes separated, with two placed under the mirror and another two under the frame. Placing the electrodes under the frame means that the frame has to be relatively thick, which negatively affects fill-factor.

Table 3.1: Published characteristics for 2-DOF micromirrors

Design	Tilt X	Tilt Y	process	fill-factor	mirror size
(Tsai et al. 2008)	$\pm 5.4^\circ$	$\pm 2.3^\circ$	SUMMiT V	96% (est)	$96\mu m$
(Jung et al. 2006)	$\pm 1.8^\circ$	0.2°	custom	99%	$360\mu m$
(Dagel et al. 2006)	$\pm 2.5^\circ$	$\pm 2.5^\circ$	SUMMiT V	98% (est)	$500\mu m$
(Aksyuk et al. 2003)	± 4.4	± 4.4	custom	53% (est)	$600\mu m$

Table 3.1 shows selected micromirrors reported in the literature. Only mirror configurations that could be put in a 2D array are listed.

3.4 Mirror designs for large-angle deflection

Due to the complexity of MEMS micromirror fabrication and the limited facilities available at UBC, we limited ourselves to designs that could be fabricated using publicly-available established multi-user MEMS processes. We chose two processes available through the CMC Microsystems (CMC 2009), a Canadian organization that provides microsystems resources to researchers at many Canadian academic institutions. As outlined below, the two fabrication processes used in this project were the surface micromachining process PolyMUMPs

(Koester et al. 2003), offered by the company MEMSCAP, and Micragem (CMC 2005), a Silicon-on-Insulator based process offered by the company Micralyne through CMC Microsystems until 2008.

Both of these fabrication processes have a limited maximum gap between the mirror layer and the substrate. Along with the mirror size, this gap limits the maximum tilt angle of the mirror. Using multiple smaller mirrors in place of larger mirrors allows for greater angles, as shown in Figure 3.6.

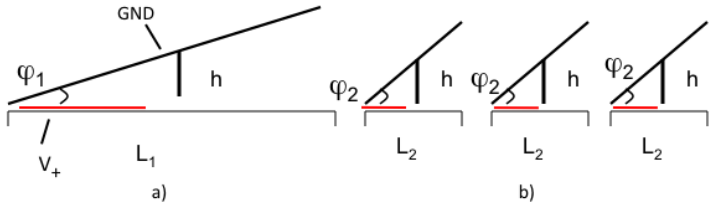

Figure 3.6: Concept of increased deflection angle through composite mirror design

A contribution in this thesis is therefore to subdivide each mirror in the array into several smaller mirrors, and address all of the smaller mirrors with the same leads. The subdivided mirrors become composite mirrors, each treated as one mirror in the larger array. This increases the scan angle compared to a single mirror with the same surface area, while maintaining the gap size which is given by the process. It also increases the resonance frequency when compared to the single mirror case; see Equation 3.15.

3.5 PolyMUMPs micromirrors

3.5.1 PolyMUMPs micromirror design

Fabrication of the first prototype set of composite micromirror arrays was done using the Polysilicon Multi-user MEMS process PolyMUMPs (Koester et al. 2003). PolyMUMPs is a surface micromachining process where successive thin layers of structural and sacrificial material are deposited, lithographically patterned, and etched to form a 3D structure. Polysilicon is used as a structural material, and thin films of phosphosilicate glass as sacrificial layers. There are 3 polysilicon layers in total, the first of which (Poly0), is not releasable from the silicon wafer on which it is deposited. A thin, non-structural metal layer is deposited last for surface reflectance, and to provide high conductivity.

For this micromirror design, which is also described in (Hoskinson et al. 2007a) all of the electrodes and leads were formed from the Poly0 layer, which allows for $3\mu m$ features. The releasable Poly1 layer between Poly0 and Poly2 was not used in order to increase the gap between the mirror (Poly2) and the electrodes (Poly0). The Poly2 ($1.5\mu m$ thick) and Metal layers ($0.5\mu m$ thick)

are used for the mirror surface. $2.75\mu m$, the combined height of the two oxide layers between Poly0 and Poly2 defines the gap between the mirror and the 3 polysilicon electrodes. An illustration of the cross-section of one mirror is shown in Figure 3.7, and the electrodes are depicted in Figure 3.8. Three electrodes were used as that is the minimum number that still provides two degrees of freedom of tilt angle, and also minimizes the number of electrical connections needed for each mirror.

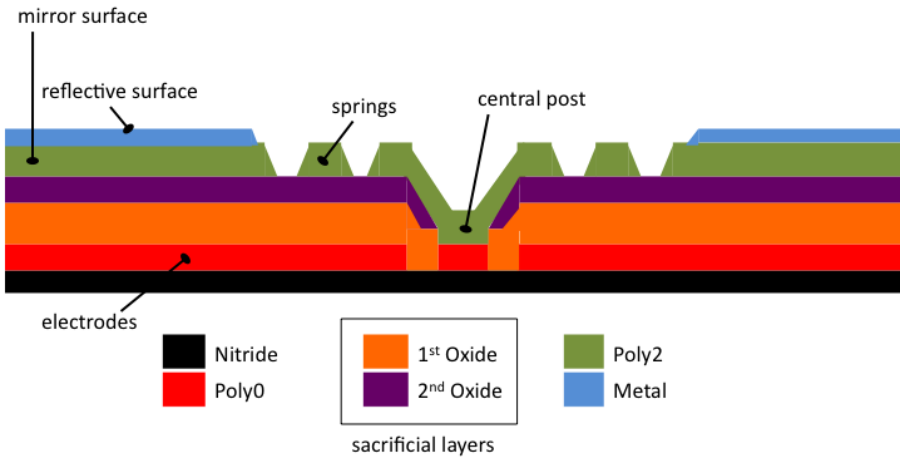

Figure 3.7: PolyMUMPs micromirror, cross-section. Both oxide layers are removed in the release step, allowing the micromirror to tilt.

To ensure the best possibility that designs made with the MUMPs process are successfully fabricated, there are a number of design rules, given in (Koester et al. 2003), which set minimum features sizes and minimum overlap between features on different layers. These rules have profound consequences on mirror design choices. For instance, the design rules imply that vertical structures have to be very stiff, so flexible spring elements have to be fabricated in horizontal layers. In this design, a central post is used to support the mirror. In order to decrease spring stiffness, the springs are cut into the mirror surface, emanating from the centre anchor. A model of the top of the mirror is shown in Figure 3.9.

The mirror surface is electrically grounded through a connection to the suspending posts, facilitating electrostatic actuation. In the micromirror array design, 62 hexagonal mirrors, each $100\mu m$ across their maximum width, constitute one composite mirror.

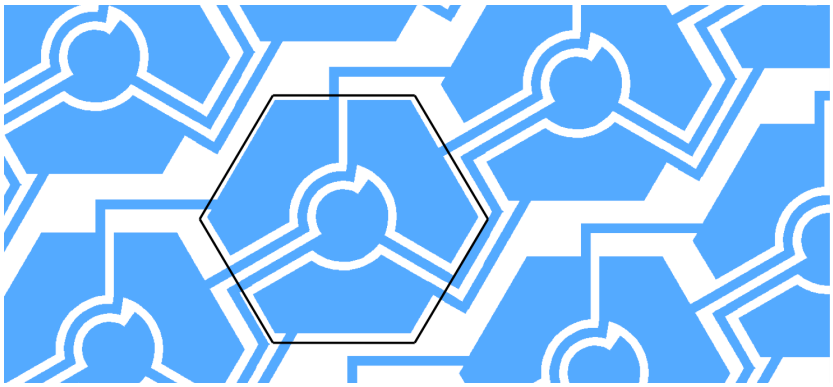

Figure 3.8: An illustration of the electrodes (blue) under a part of one section of a composite mirror. One mirror's outline is shown in black. All electrodes that cause the mirrors to tilt in the same direction are connected together. The central portion under the mirror connects the mirror surface to ground through the central post.

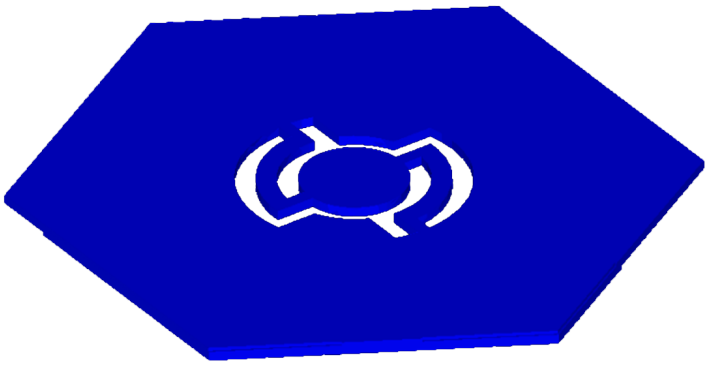

Figure 3.9: PolyMUMPs micromirror, top view of model

3.5.2 PolyMUMPs mirror characterization

Figure 3.10 shows a photograph through a microscope of several mirrors fabricated with the PolyMUMPs process. Because the polysilicon is conformal (each successive layer follows the vertical contours of those that have been applied before it), the electrodes and leads providing voltage are visible. Etch holes in the mirror are also visible. Etch holes are required when using the PolyMUMPs process for proper release of large polysilicon structures such as these mirrors. Figure 3.11 shows a one composite mirror and the surrounding bond pads.

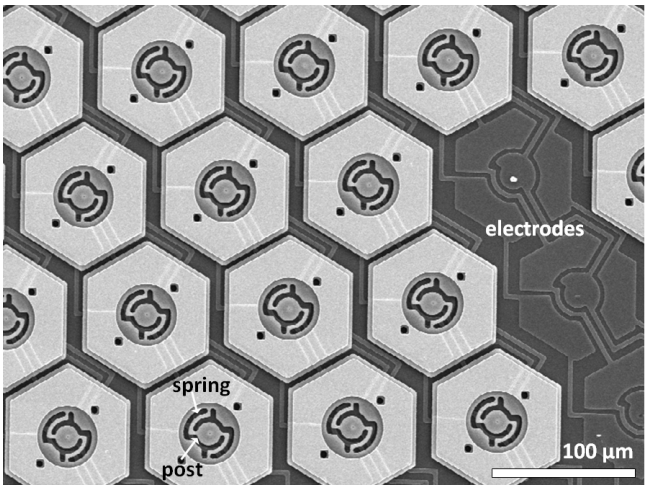

Figure 3.10: Scanning-electronic microscope image of PolyMUMPs mirrors. Some mirrors have been removed to show underlying electrode structure.

Figure 3.12 shows the static deflection angle as a function of voltage from measurements using a white-light interferometer. The maximum angle observed before pull-in was 0.75°. Pull-in was observed at 87.3 volts. This deflection angle compares favourably to other tip/tilt micromirrors fabricated via the PolyMUMPs process such as the 12 mrad (.00021°) achieved in (Lin et al. 2001) for a 130μm square mirror.

Resonance frequencies for the mirrors were measured using a Laser Doppler Vibrometer in two ways. The first involved a frequency sweep in which the velocity amplitude of the mirror's response to a sinusoidal 20V peak-to-peak voltage input, in addition to DC bias voltage of 40V, was measured at discrete frequency points. Dividing the velocity amplitude measured at each point by its input frequency gives the displacement amplitude. The second method was to excite the mirror with a step input, and again measure the resulting velocity. The Fast-Fourier Transform (FFT) of this response gives the displacement spectrum. Figure 3.13 shows that the results from both measurements match well, and both show the resonance frequency at approximately 95 kHz. Similar results have been achieved for different bias voltages.

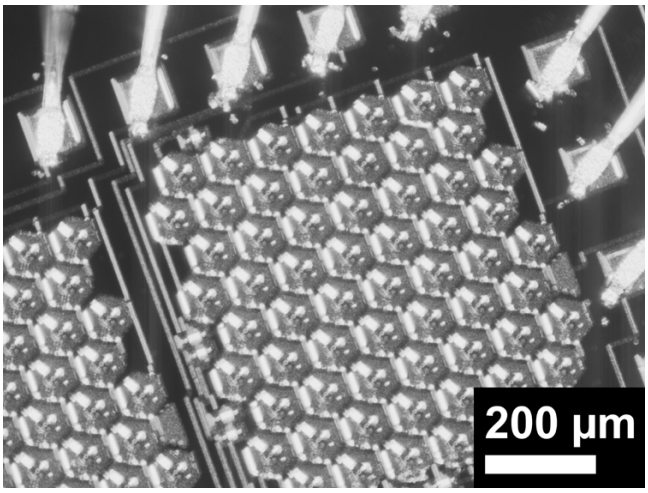

Figure 3.11: Photograph of composite mirror using the PolyMUMPs process. The bond pads used to electrically connect the wire bonds to the mirrors are visible in the top and right sides of the picture

.

Figure 3.14 shows the surface characteristics of one fabricated mirror from measurements taken with a white-light interferometer. This data shows the "print through" effect from the underlying lines and electrodes. The surface curvature due to residual stress from the metallization is also evident. The radius of curvature of this mirror was measured at $2.8mm$ by plotting a cross-sectional slice of the interferometer data.

While the print-through is unavoidable with the PolyMUMPs process, the surface curvature could have been largely avoided. The final metallization step of the MUMPs process has been shown to be the primary cause of surface curvature. We chose to use the optional metal layer on these mirrors despite that because (Cowan et al. 1998) have shown that it nearly doubles the optical efficiency of the mirror compared to the case without any metallization.

With metallization, the curvature can be reduced by combining the Poly1 and Poly2 layers to provide a $3.5\mu m$ mirror layer rather than the $1.5\mu m$ Poly2 layer alone. The increased thickness of the two combined polysilicon layers would have substantially stiffened the mirror against curvature. However, this also would have reduced the electrostatic actuation gap size to $2\mu m$ from $2.75\mu m$, which would in turn reduce the maximum tilt angles of the mirrors. We therefore decided to only use the Poly2 layer to gain the added tilt angle, at the cost of increased mirror curvature.

In retrospect, the design for this mirror array was perhaps too ambitious for a first iteration. To provide the three separate voltages to the mirror electrodes for the 62 mirrors within one composite mirror, conductor bridges had to be designed across other conductors. With 25 composite mirrors, there were 1550

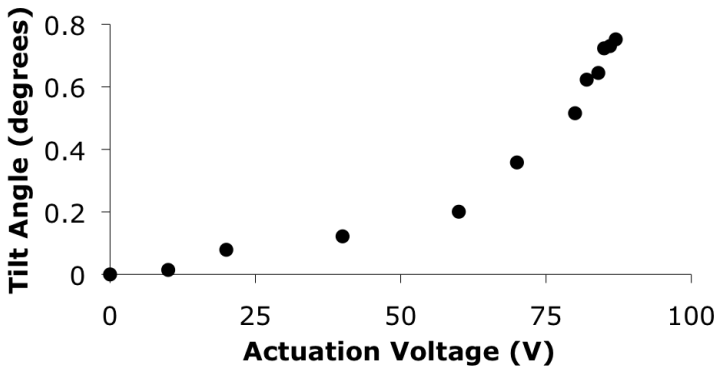

Figure 3.12: Static mirror deflection of PolyMUMPs mirror for a constant voltage supply to one electrode.

mirrors in total on the array, and a lot of potential for errors by the designer, leading to false connections. As a consequence, not all composite mirrors could be controlled independently, and some sections of the composite mirrors did not move with the other sections.

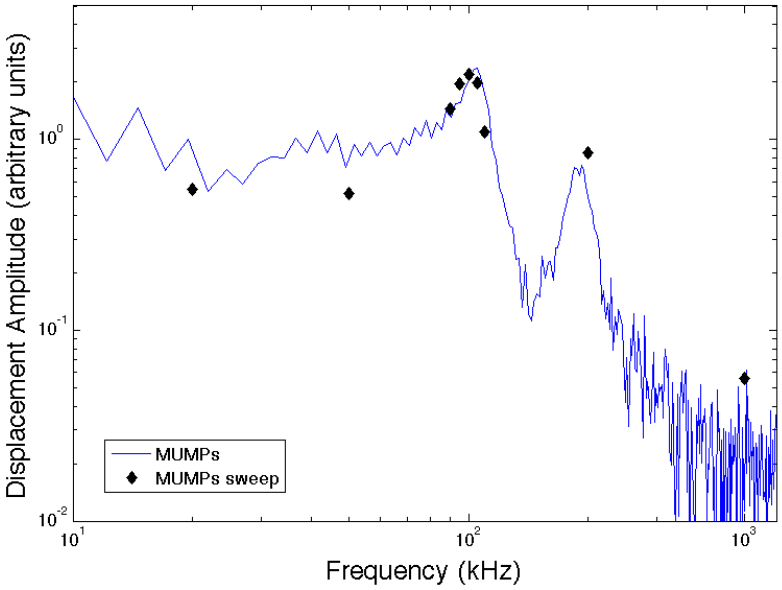

Figure 3.13: Two measures of displacement amplitude vs. frequency for PolyMUMPs mirror (Hoskinson et al. 2007a). Using a 40V bias voltage signal, a 20V peak-to-peak AC voltage amplitude was used for the sweep, and 20V was also used for the step excitation.

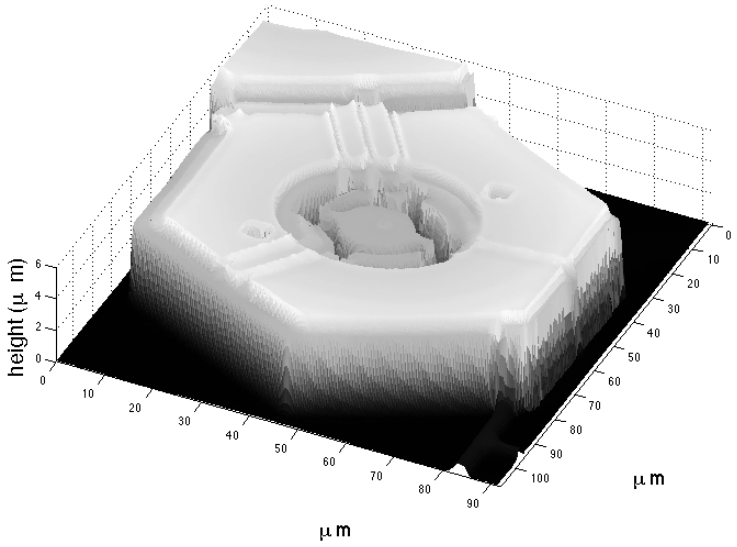

Figure 3.14: One MUMPs mirror as measured by a white-light interferometer. The radius of curvature was measured to be 2.8mm.

3.6 Micragem mirror arrays

3.6.1 Micragem mirror designs

In the Micragem process, 10 or 12μm deep cavities in glass carry structured metallization used for electrodes, leads and bond pads. The metal can be placed in the cavities or on the glass surface. A Silicon-on-Insulator (SOI) wafer is then anodically bonded to the glass, and etched so that only a 10μm thick Single Crystal Silicon (SCSi) layer is left. Another metal layer is then deposited and patterned on the SCSi surface, after which the SCSi is patterned through deep reactive ion etching (DRIE).

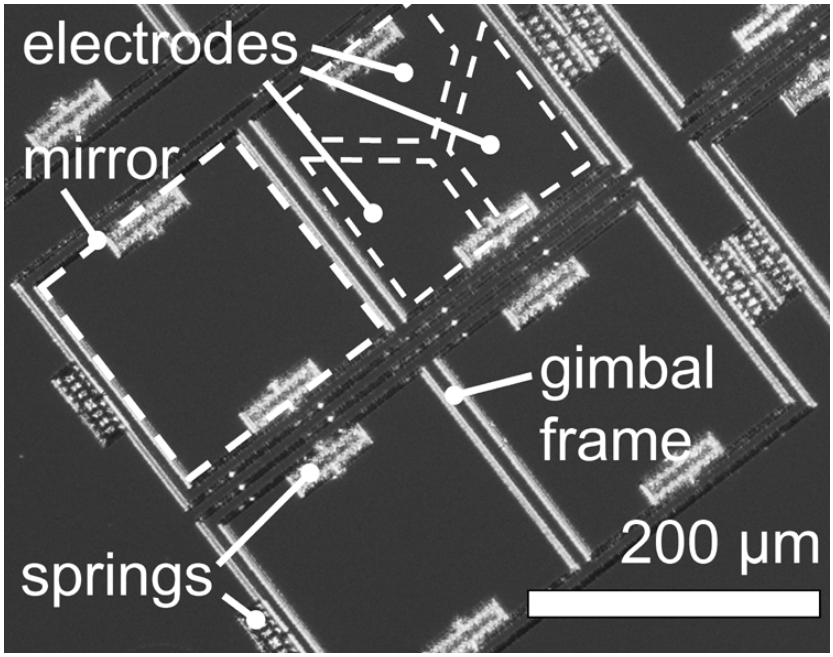

Figure 3.15: Picture of 4 micromirrors, forming one pixel of an early design micromirror array made in using the Micragem process

The constraints of the Micragem process design rules dictated many of the choices made in our design process. There is only one mechanical layer, so all of the hinges and frames supporting the mirror have to be incorporated into the same level as the mirror surface. All of these necessary support structures directly subtract from the reflective mirror surface, and therefore directly impact the overall fill-factor in the array.

The micromirrors we have designed using the Micragem process employ a thin gimbal system, which suspends square mirrors over the 12μm cavity. As with the MUMPs design, multiple mirrors are controlled simultaneously to form one AMA composite mirror by linking their electrodes. We chose to take ad-

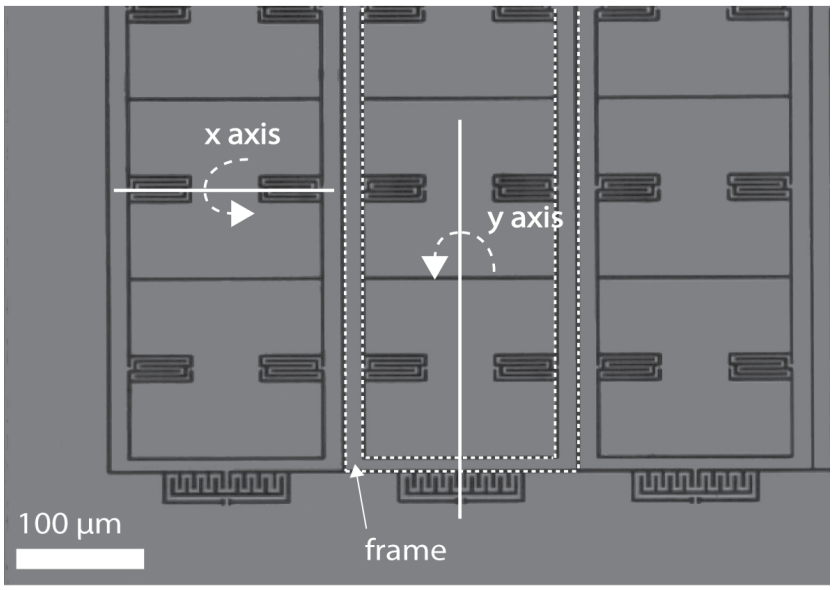

Figure 3.16: Several micromirrors with a portion of the thin gimbal frame system. Each frame suspends six square mirrors over a $12\mu m$ cavity. Each mirror can be tilted around the two indicated orthogonal axes x and y. Six rows of six mirrors (36 total) are controlled together to form a composite mirror.

vantage of this during the design of the micromirrors themselves. In our design, each row of mirrors within the same composite mirror shares the same gimbal frame, which improves fill-factor and promotes homogeneity between submirror deflections. An early design, shown in Figure 3.15, has two mirrors per row, while later designs had 5 or 6, a portion of which is shown in Figure 3.16. Multiple mirrors within one frame minimizes the ratio between frame area and mirror area. Instead of the loss of reflective area from four frame sides and two frame hinges per mirror, the loss from the frame area and outer hinges are amortized over a larger number of mirrors. The two gimbal springs and two of the four sides are only needed once per group rather than once per mirror. Figure 3.17 shows a photograph of several of the composite mirrors in the array.

The other main consideration affected by the design rules was the difference in minimum feature sizes between the SCSi layer used for the mirror surface and the metal layer used for the electrodes. In the SCSi layer, the minimum feature size (both length and width) and dark size (cut) is $2\mu m$. For the metal layer used for the electrodes, both the minimum feature and dark size is $10\mu m$. This had a substantial impact on our array design, because wires $10\mu m$ thick and $10\mu m$ apart had to connect each electrode and its corresponding bond pad on the outside of the array. The maximum vertical gap distance of the process ($12\mu m$) meant that the mirrors had to be relatively small to achieve enough tilt

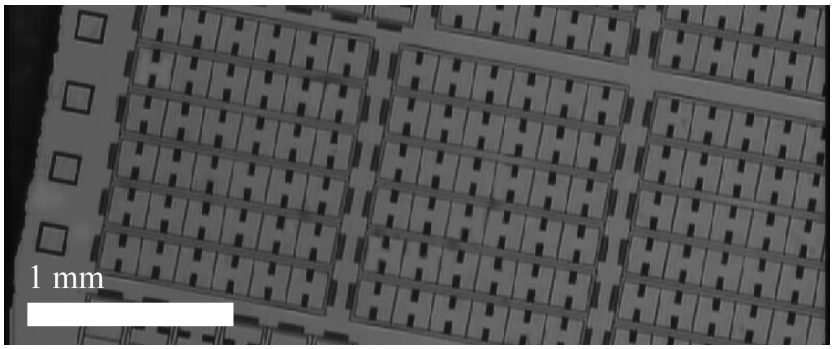

Figure 3.17: Array of composite mirrors in Micragem row design

angle to move the light by a sufficient amount. There was also a practical limit as to the number (84) of bond pads available in a convenient package for the prototype array.

Placing a row of mirrors within a gimbal frame opened up design possibilities for the electrode arrangement. Electrodes could be routed in daisy-chain form within the rows. The minimum number of electrodes needed to achieve two-degree of freedom actuation is 3. If the electrodes are structured in a way that allows them to be connected in rows, the electrode paths can be minimized. With one wire between electrodes, it is also possible to connect rows together in a compact manner, as shown in Figure 3.18. This results in a square or rectangular array of mirrors in a composite mirror. Each row of mirrors is in one frame. Each frame also constitutes a row in the electrode layer.

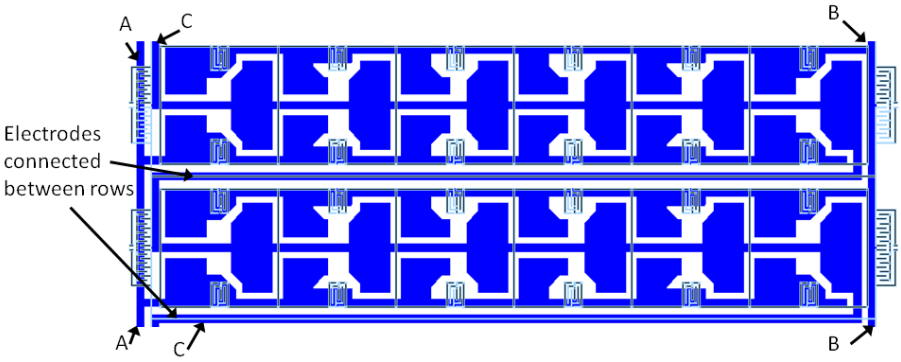

Figure 3.18: Two rows of electrodes, showing row-wise connection, and how two rows are connected to each other. The outlines of the mirror features are also shown.

3.6.2 Simulations of Micragem mirrors

Coupled electrostatic-mechanical simulations were undertaken to predict and analyze the behaviour of the micromirrors using the commercial finite-element modelling (FEM) software ANSYS^{TM}(ANSYS 2009). The EMTGEN macro in ANSYS was used to generate TRANS126 elements between the nodes on the bottom surface of the micromirror and the electrodes. TRANS126 elements are useful for simulating electrostatic-mechanical coupling between a MEMS device and a plane, and are valid if the gap between the mirror and the electrode is small compared to the dimensions of the mirror, as is the case with this design.

Figure 3.19 shows the result of a simulation of one row of micromirrors that we have designed, under a potential of 120V applied to the middle electrode C underneath each mirror. The simulation predicted a maximum negative and positive deflection of $-3.1\mu m$ and $1.65\mu m$, respectively.

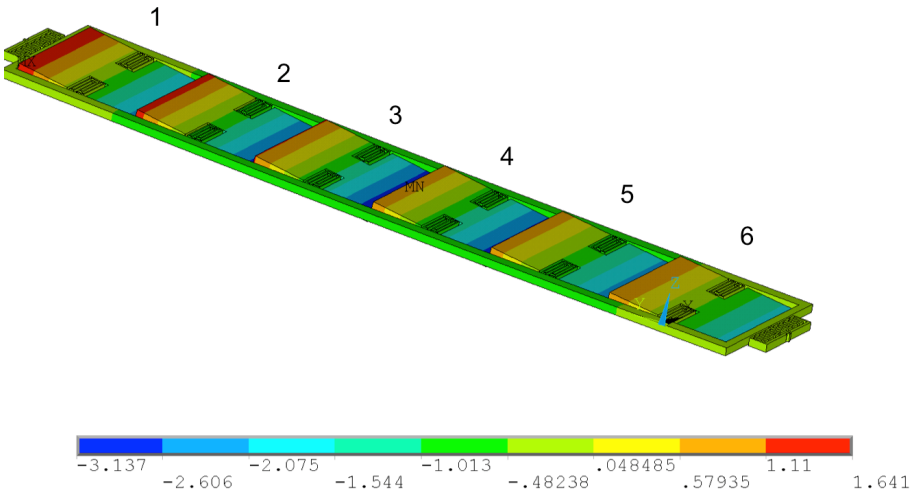

Figure 3.19: Model of micromirror produced in ANSYS.

Figure 3.19 shows that the mirror frame itself bends towards the electrodes as electrostatic force is applied. The longer the frames are for a given width, the more they are able to bend. When the frame bends, the mirrors closest to the middle of the frame become closer to the electrodes underneath. They therefore are subject to higher electrostatic forces than those towards the edges. The mirrors in the middle of the frame therefore tilt more for a given tilt voltage than those on the edges.

This effect is not symmetric, because half of the mirrors in the row are working with the bend and half working against it. The effect is to tilt the mirrors slightly towards the centre of the frame. When a voltage is applied, the flat mirrors in Figure 3.20 a) bend at different rates depending on their distance from the electrode and what side of the centre of the row they are on, as shown in Figure 3.20 b).

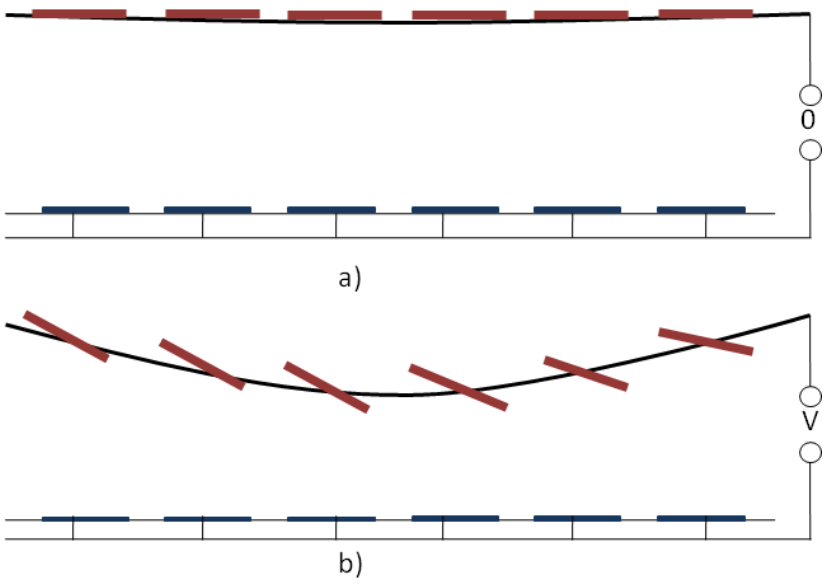

Figure 3.20: (a) The gimbal frame holding a row of mirrors can bend. (b) This causes an applied electrostatic field to be unequally converted into force, depending on the initial orientation of the mirror due to frame bending

The bending of the frame could be mitigated by making the frames wider, or by putting fewer mirrors within one frame, and thus making the frames shorter. A more compliant spring on the frames also be of help to rotate the row around the y-axis. Frame bending does introduce some interesting side effects. When the mirror tilts about the x axis, the effect on the reflected light beam is an asymmetrical compression. This effect is controllable in the sense that an additional bias can be applied to all 3 electrodes to exaggerate the compression. The asymmetrical nature of the compression is not controllable, because the total electrode size is fixed for each individual mirror.

The ANSYS simulation results for x-axis rotation as well as experimental results are shown in Figure 3.21. The experimental results were obtained using a white-light interferometer.

When we compare Figure 3.21 *a* and *b*, we see that the simulated mirrors actuate at smaller voltages than the measured values. This could be due to an overestimation in the simulation of the amount of over-etch. The Micragem design rules (CMC 2005) specify a minimum feature size for the SOI layer as $2\mu m$, and a undercut of $\pm 1\mu m$, which can significantly affect the spring constant of the micromirror's torsional springs. Depending on the rate of variation (which is not specified in the design rules), this can have a significant impact on the performance of the composite mirrors, which depend on having every mirror within the group having very similar performance. Figure 3.22 shows two springs imaged with a scanning electron microscope, and shows that even within one

torsional spring, the widths of the individual serpentine coils can change. It is evident the that the springs were assumed to have smaller widths than they did in actuality, and were thus more stiff than those in the simulation. Thus in the simulation, less voltage was needed to produce a given deflection than was measured with the fabricated mirrors.

Figure 3.23 shows a fabricated mirror under an applied voltage as reconstructed by the data from a WYKO white-light interferometer. It is evident that the individual mirrors of this array show a slightly different tilt, particularly Mirror 5.

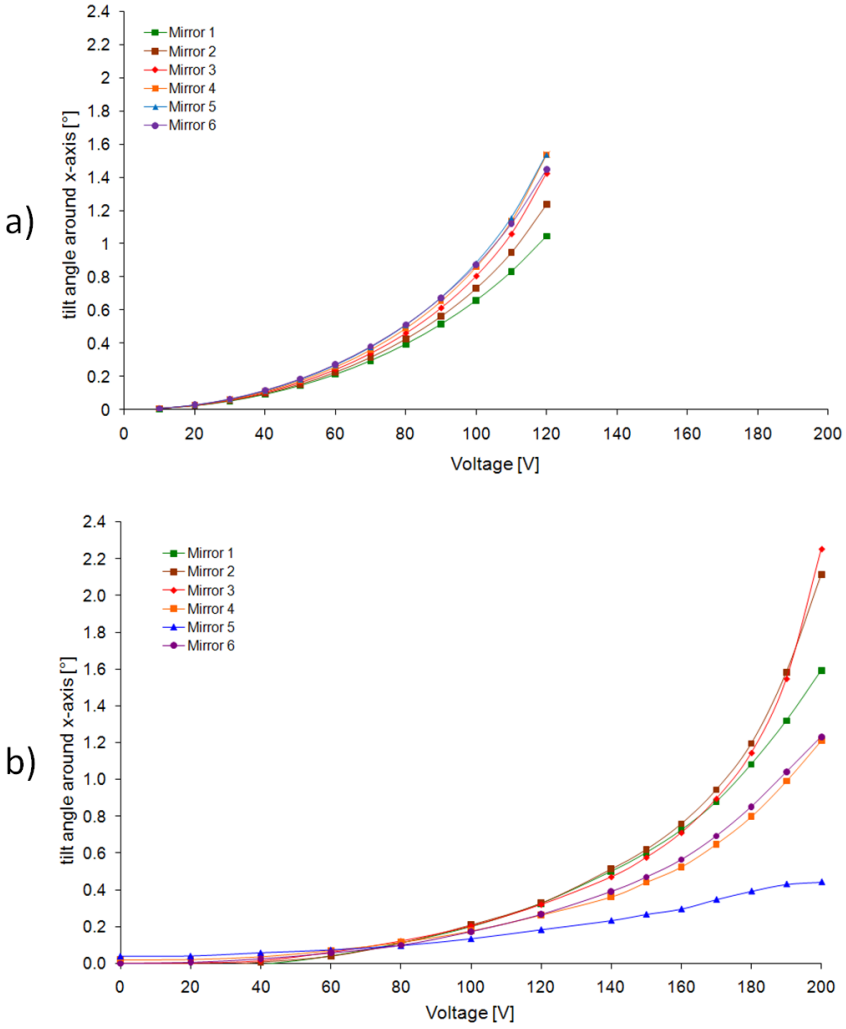

Figure 3.21: Simulated (a) and measured (b) individual mirror tilt around X axis. The mirror numbers correspond to the labels in Figure 3.19.

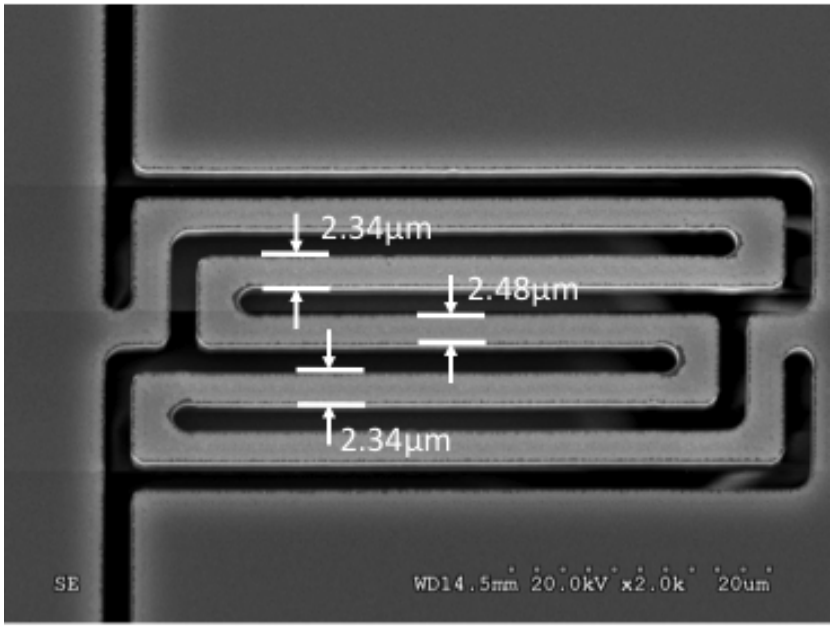

Figure 3.22: A scanning-electron microscope picture of micromirror springs, showing the differences in spring width.

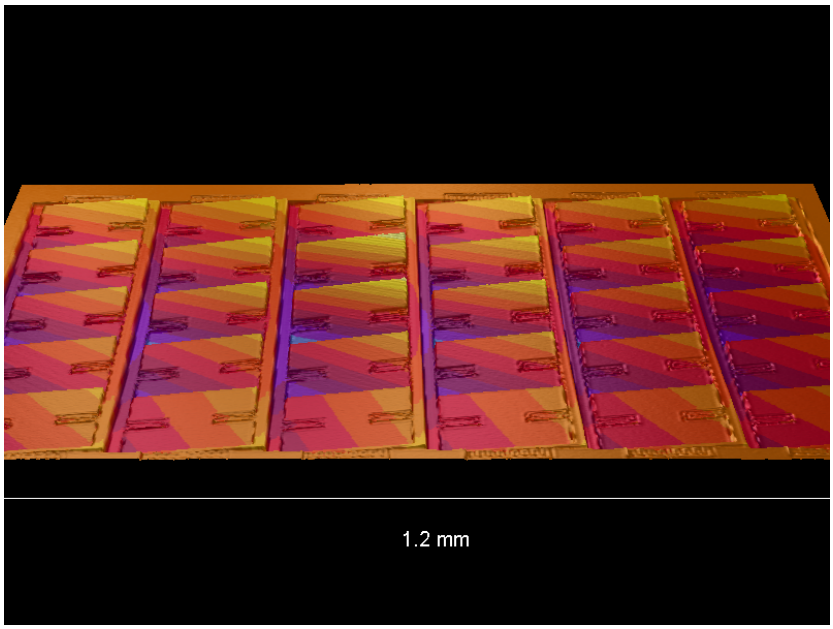

Figure 3.23: Image from WYKO white light interferometer of one composite mirror tilting.

3.7 Proposed new Micragem design

A future mirror design better suited for this application could use the centre spring concept of the PolyMUMPs mirrors with the flat surfaces achievable with the Micragem process. To illustrate this, we simulated a new mirror design using the Micragem process. Again we chose to use three electrodes to provide three degrees of freedom. Minimizing the number of electrodes is especially important for the Micragemg process because of the relatively large minimum feature size of 10 μm in the electrode layer, compared to $2\mu m$ in the SCSi layer. Equally important, a hexagon can be regularly tiled, fit together without any spaces. Similarly to Figure 3.18, an unlimited number of each of the three types of electrodes can be connected together to form composite mirrors without the need for bridges. The middle electrodes for each of the composite mirrors are connected in a row, and rows are connected via loop between rows. The other two electrodes are also connected into rows, and these rows are connected together at the ends.

Figure 3.25 shows the mirror shape and the shape of the three equal-sized electrodes underneath. The mirror is attached by springs to a central $10 \times 10\mu m$ post that suspends the mirror above the etched glass layer and electrodes. The springs are designed to provide nearly equal spring force in every direction. Each exterior side of the hexagon is 125 μm. From any corner to its opposite corner, the distance is 250 μm. The minimum distance from centre to edge is $125\sqrt{3}$, or $216.5\mu m$.

This design has an estimated fill factor of 92%. The calculated tilt angles for the four figures are as follows: Figure 3.25: 2.09° for electrode A at 180V, Figure 3.26: 1.91° for electrode B at 180V, Figure 3.27: 2.08° for electrodes A and B at 200V, and Figure 3.28: 1.64° for electrode A at 170V, and electrode B at 190V. These are highest tilt angles achieved before the mirrors reached pull-in during the simulation.

Unlike the previous Micragem design, the mirrors do not need mechanical support from the rest of the SCSi surface, but they still do need a common electrical connection to provide a voltage signal that opposes that of the electrodes underneath. Because of the smaller feature size of the SCSi surface, we chose to include this connection as fixed rods in the mirror layer. One such a rod is shown in Figure 3.26. It is connected to the central post, so it remains stationary and is under no force as the mirror is actuated. As shown in the simulations, the small cut in the mirror surface this introduces does not cause the mirror to bend in an uneven manner. In a larger array, these rods could be connected to small connecting spaces between mirrors, as illustrated in Figure 3.24.

Unfortunately, the Micragem process used to fabricate these mirrors is no longer being offered as a public service by neither Micralyne nor CMC, so the process must be replicated elsewhere before these mirrors can be fabricated.

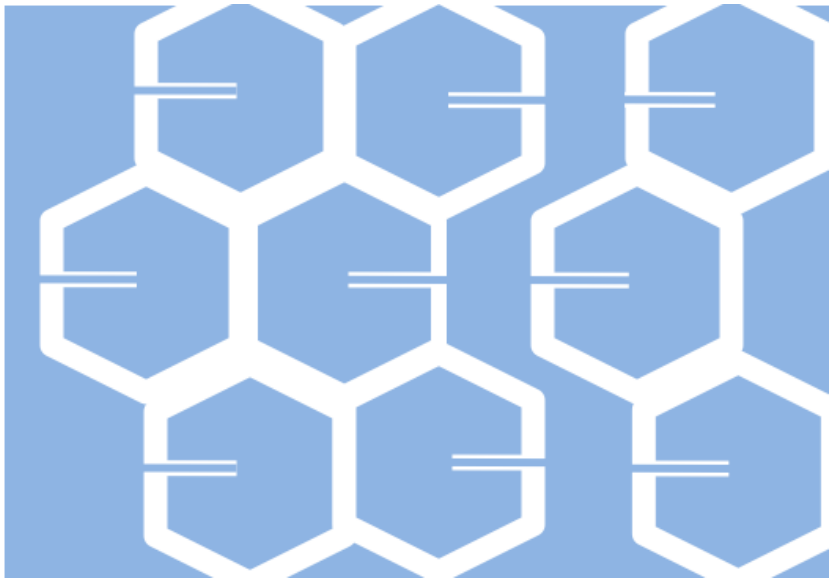

Figure 3.24: Illustration of mirrors electrically connected together by fixed rods anchored to the central posts of the mirrors.

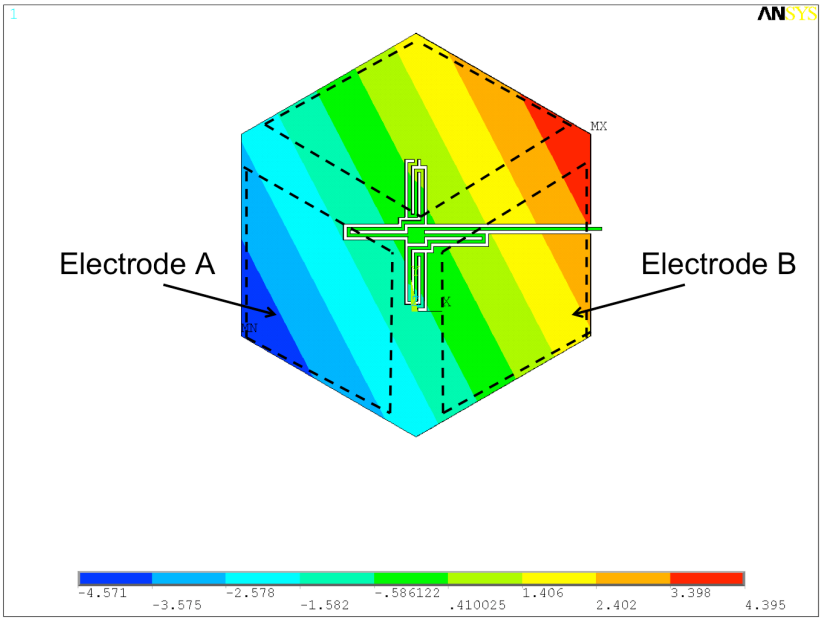

Figure 3.25: Electromechanical simulation of micromirror with electrode A actuated with 180V. The three electrodes underneath are also indicated.

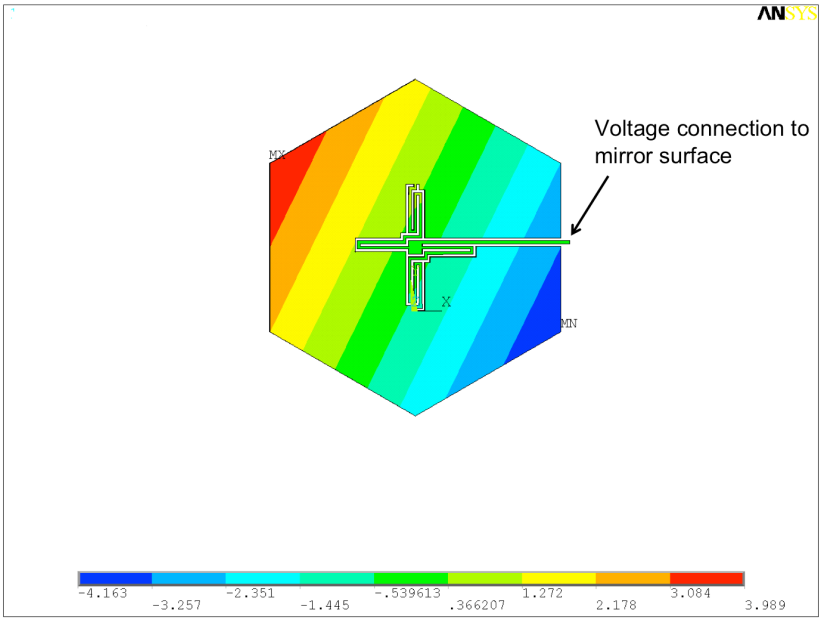

Figure 3.26: Simulation of micromirror with electrode B at 180V.

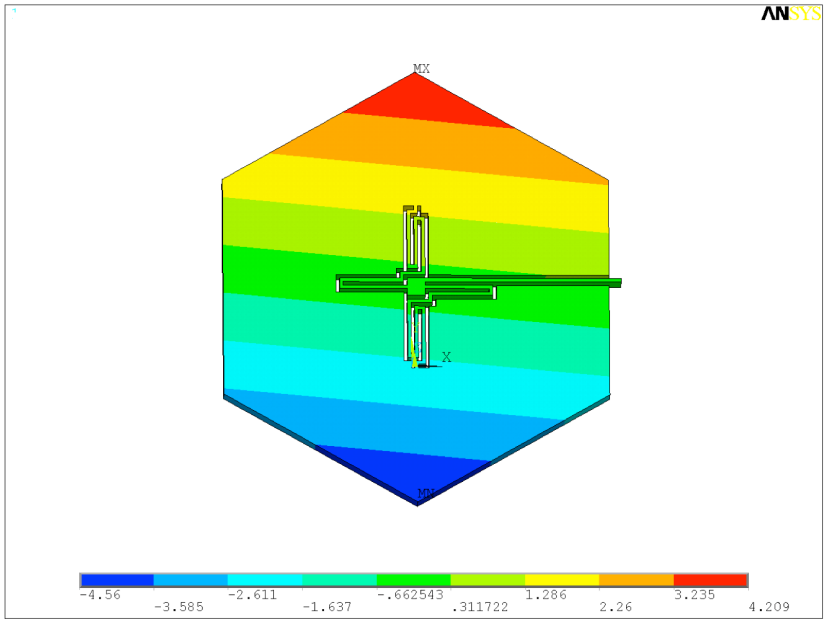

Figure 3.27: Simulation of micromirror with electrodes A and B at 200V.

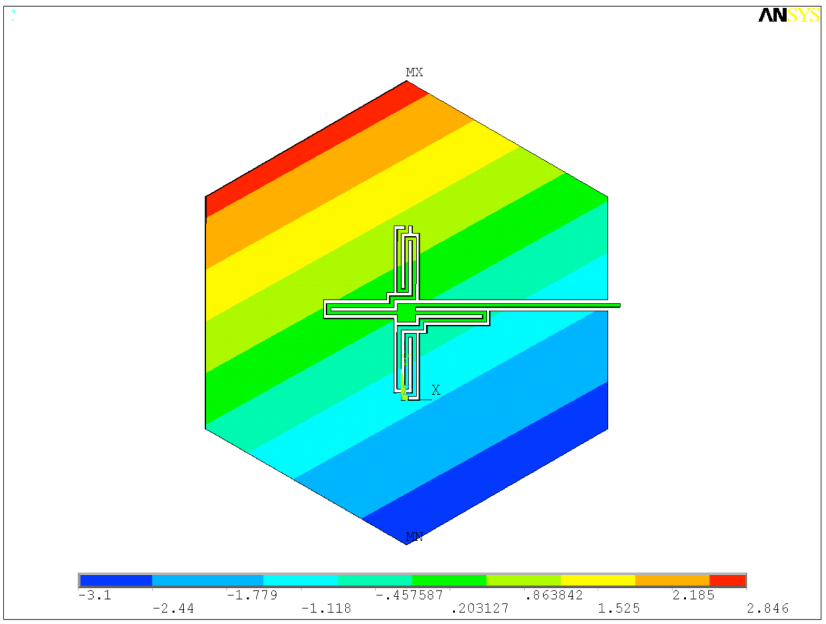

Figure 3.28: Simulation of micromirror: electrode A at 170V, and electrode B at 190V

3.8 Future mirror designs

The requirement that the mirrors must be able to be fabricated using an existing multi-user process severely limited the design space. In the PolyMUMPs designs, the mirror surfaces had a large radius of curvature and had their surfaces marred by the "print through" effect from the underlying springs. The limited depth of the PolyMUMPs process of $2.75\mu m$ maximum gap height width also limited the achievable tilt angle of the mirrors. For a tilt angle of $\pm 4^{\circ}$, the mirror would have to be less than $30\mu m$ in length, and would thus require very high voltages to actuate.

The mirrors fabricated with the Micragem process had optically flat surfaces, but the constraints on the minimum feature size of the electrode layer had large repercussions on the fill-factor of the array. While the row design mitigated this issue somewhat, it introduced a new problem of frame bending, creating non-uniform tilt angles. This makes the relationship between angle and voltage highly non-linear, and very complex to account for in the allocation algorithm. The shape of the ML also depends on the tilt angle as the frame bends, adding further complexity. While it may be possible to use the non-uniform tilt advantageously to shape the light to the characteristics of the image, the shape of the light would depend on the voltage given to each electrode under the mirror.

A future mirror design better suited for this application could use the centre spring concept of the PolyMUMPs mirrors with the flat surfaces achievable with the Micragem process. An attempt at this was made, but due to issues with fabrication the micromirror array was inoperative. Unfortunately, the Micragem process used to fabricate these mirrors is no longer being offered as a public service by neither Micralyne nor CMC. A new process, possibly Sandia's SUMMiTV (Sandia National Laboratories 2007), will have to be used for the next iteration of mirrors.

If the Sandia process is used, a design similar to that used in (Tsai et al. 2008) would provide more than adequate tilt range for the AMA device. The fill-factor of these mirrors is also high at 96%, which is crucial for this application. It could also be tiled into composite mirrors depending on the desired size of the array. The design would have to be adapted for this application, as despite its theoretically high fill-factor, it has not been demonstrated in a large array, and would also need to be adapted for use in composite mirrors.

Chapter 4

Light allocation

This chapter describes how the mirrors are allocated for a given image, how the final image is formed, and provides an evaluation of how much extra brightness can be coaxed from a projector through the addition of an AMA.

In Chapter 2, we examined the implications of adding and AMA device into a projector's optical path in terms of optics: the shape and distribution of the light reflecting from each AMA onto the DMD. In this chapter, we estimate the potential improvement to an image displayed on the projector with the addition of the AMA. To do that, we need to convert the optical specification of the ML and its range into one expressed in terms of image pixels. Section 2.5.3 gave a description of the convolution kernel that can be used to estimate the shape of the ML for a given simple physical arrangement of AMA and DMD with lenses. This is used in Section 4.3 to calculate the ML shape in pixel coordinates.

In this chapter, low-dynamic range (LDR) images with formats that are device-dependent are primarily used as examples. These image formats specify relative luminance changes on the device used to reproduce the image, rather than the absolute, scene-referred luminances of the original scene. They are also gamma-encoded so that the relative luminance changes represented by the pixel values are not linear.

4.1 Mobile light sources

A key aspect of improving dynamic range by using two modulators as described in (Seetzen et al. 2004) is that one of them can be of much lower resolution than the final displayed image. The light distribution of the low-resolution modulator is then corrected by the high-resolution modulator to achieve the desired image. The light sources in the low-resolution modulator (LRM) are in fixed locations, irrespective of where the bright parts of the image are. If the locations of the LRM light sources and the bright parts of the image coincide, the full capability of this concept can be exploited. If they do not, the closest neighboring light sources must be used. Because the intensities of the light sources fall off over distance, they will have to be set to higher intensities than in the matched case. Also, more light has to be blocked, because the neighbours will now be putting out more light than the needs of their locations. As an example, take the LED modulator detailed in (Seetzen et al. 2004). A bright spot in the image may occur in between two adjacent LEDs. In that case, the neighboring LEDs need to provide higher light intensities than the local maximum of the image.

The excess light from these LEDs would then have to be blocked by the high-resolution LCD. This would result in a lower contrast and higher power usage than would be possible if the LEDs could be positioned exactly where they are needed for each image.

In an AMA approach, the intensity of light of each mobile light source (ML) can not be changed, but the location on the high-resolution light valve is variable. This means that every ML can be targeted to exactly where it is needed. The combination of all of the virtual light sources must be sufficient to correctly display the image. The allocation scheme should also take advantage of the flexibility afforded by these MLs to the peak brightness of the image.

Similar to the approach by (Seetzen et al. 2004), the signals sent to the low-resolution modulator are calculated first, then the light distribution from the first modulator on the second modulator is simulated, and finally, the corrected image sent to the high-resolution second light modulator is calculated. Light from an ML can be described by convolution of a rectangular area of pixels with a point spread function on the DMD, as detailed in Chapter 2.

4.2 Gamma

Some discussion on how typical images are specified is needed so that we can match the representation of the MLs to the representation of luminance values in an LDR image. Pixels in most standard image formats such as HD video (Poynton 2003) are encoded with a non-linear gamma function. Gamma is the rate of change of the relative luminance of image elements of a reproduction as a function of the relative luminance of the same image elements of the original image (Fairchild 1995).

Video transmissions were first encoded with a gamma to compensate for the non-linear characteristics of cathode-ray tube (CRT) displays. CRT electron guns typically have a current/drive voltage relationship described by a power law

$$I = kV^{\gamma_d}, \tag{4.1}$$

where I is the current from the electron gun, proportional to the light intensity, k is a constant, V is the drive voltage, and γ_d is the gamma of the electron gun (Poynton 2003).

To achieve a linear lumen output, transmissions such as those specified by the National Television System Committee (NTSC) and later the sRGB format for digital images (Stokes et al. 1996) were encoded with the inverse of γ_d. $\gamma_d = 2.2$ has been used as a benchmark for the behaviour of a typical CRT, so the signal gamma, γ_s, was specified as $1/2.2$. However, actual direct-view CRTs have a γ_d in the range from 2.31 to as high as 3.1, so $\gamma_s\gamma_d = 1.27$, instead of the 'ideal' value of 1 for linear imaging. Tests show consumers prefer a nonlinear gamma (Roufs et al. 1994; Barten 1996). This may be due to the influence the surrounding luminance has on a viewer's perceived contrast of an image (Fairchild 1995).

The exponent γ_d is different for different monitors. Displays with other relationships between input signals and luminance such as LCDs typically include circuitry to do inverse gamma correction as a pre-processing step. (Kessel et al. 1998) describes how the DMD is a linear luminance modulator, and so its driver must include inverse gamma pre-processing.

The method described in Chapter 2 to estimate the distribution of light on the DMD from the AMA is in normalized, relative, linear units. An inverse gamma step is therefore performed on the incoming image so that it is in an approximately linear, relative luminance scale. Because the second modulator, such as a DMD, usually assumes a standard gamma, the corrected high-resolution signal sent to the DMD should be gamma re-encoded in anticipation of another inverse gamma step happening automatically in the DMD driver hardware. The AMA manipulates light in a linear manner, and so has a gamma of 1. If a non-linear gamma in the *output* such as 1.27 is preferred, this would be taken into account in the original image before it is decomposed into the first and second modulator components.

4.3 Mirror allocation

Finding the optimal light distribution for the AMA corresponds to choosing the locations for each of the n MLs from the array of n analog micromirrors. The procedure for exploring the solution space is as follows. The approximation starts with a representation of the light incident on the DMD, coming from the entire AMA, as a block of pixels in DMD-space. The size of the AMA light representation is the pixel resolution of the DMD with an additional overfill of approximately 20% per side. This amount of overfill is typical in projection displays (Brennesholtz and Stupp 2008), in order to maintain an approximately uniform brightness distribution over the DMD, and as a guard against any misalignment in the optics.

For an 800×600 pixel DMD, we initially represent the light incident on the DMD from the AMA as a white image of size 1000×800 pixels, centered on the DMD. The image is subdivided into n rectangles, each of which represents 1 ML. The position of the centre of each rectangle in image (DMD) coordinates is recorded. Each rectangle is then blurred using the convolution kernel specified in Equation 2.24, to get the shape and relative intensity distribution of each ML. The convolution kernel used to blur the light is arrived at using the physical parameters of the optical system, as detailed in Chapter 2.

Initially, the mirrors of the AMA are in a non-actuated, flat state so all MLs are equally distributed on the DMD. The summation of the MLs in this initial distribution should give a relatively uniform distribution on the DMD so that any image can be displayed in a conventional manner. As detailed below, different levels of blur were tested to see their effect on the overall performance of an AMA projector. The relatively-large overfill of 200 pixels per side was arrived at because it was the minimum size that gave an approximately uniform distribution over the 800×600 resolution of the projector, as shown in

Figure 4.1 for each of the levels of blur tested. With any smaller an overfill, the intensities at the edges of the area incident on the DMD were much less than those in the middle. The amount of overfill chosen was a balancing act for this simulation because the amount of blur itself affects the uniformity of the overall distribution. With a 200 pixel per-side overfill, the uniformity on the DMD ranged from 1% deviation at small levels of blur, to 15% at higher levels of blur.

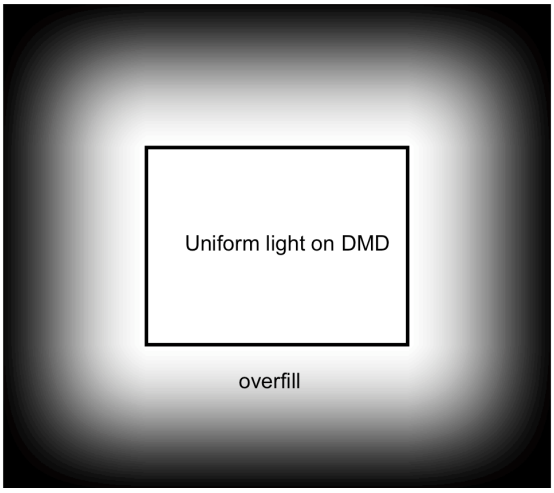

Figure 4.1: The light distribution from an non-tilted AMA, illustrating overfill. The rectangle in the middle represents I_p, the light incident on the DMD.

Once the blur of each of the MLs is established, the algorithms detailed below search for the position of each of the n MLs that provide the best improvement to the displayed image. It is assumed that there is adequate granularity in the mirror control to position the MLs at any pixel in the image. If *numpixels* is the number of pixels in the DMD, there are $numpixels^n$ total possibilities of mirror configurations, which is an astronomical number of positions to evaluate for even an 800×600 DMD and a 3×3 array of AMA mirrors using a brute force method where every solution is evaluated. Additionally, the analog nature of the AMA mirrors potentially allows for the placement of MLs between pixels, which would mean even more possibilities.

To guide the allocation algorithm, it needs to be established what constitutes an improvement. In HDR pictures where the scene is defined in terms of actual luminances, the improvement is the ability to show images correctly that have higher peak brightnesses than the original projector was able to generate. With LDR images defined in terms of relative luminance, improvement can be defined as the ability to boost the entire range of pixel brightness values; in effect to be a projector with a higher original lumen output.

We propose two different approaches to mirror allocation: Gaussian pyramids with iterative redistribution, and median cut. Both have been implemented

in Matlab and tested; their implementation and performance are detailed below.

Given that neither of these solutions is guaranteed to give the best possible improvement, an advantage of testing two algorithms is to be able to verify the effect of input parameters (range, blur level, number of mirrors) on the resulting improvement. If patterns evident from one simulation's results are echoed in the other, we can be more confident that they reflect a near-optimal solution at each parameter.

4.4 Gaussian pyramid approach

To locate luminance regions in the image to be matched with a ML, the original image and an image of the ML can be represented as Gaussian pyramids, first described in (Burt and Adelson 1983). We adapted a Gaussian pyramid implementation for Matlab described in (Simoncelli 2007) to form a pyramid in which the original image is iteratively low-pass-filtered with a binomial filter to reduce its high-frequency content, while the image is simultaneously downsampled to reduce each of its dimensions by a factor of two. The image of the ML is also reduced in the same manner, as shown in Figure 4.2. At lower resolutions, each pixel is a weighted sum of those around it at higher resolutions, so exploring the solution space at a coarse image representation allows easier recognition of the global luminance distribution in the image. Once the best ML placement is found at one level, it is refined at successively higher-resolution levels of the pyramid.

The algorithm starts with a target image to be projected, I_m, and a simulated distribution of light on the DMD, designated as I_p. I_p can be calculated from the light distribution $A(r_i, c_i)$ of the individual MLs as

$$I_p(r, c) = \sum_{i=1}^{n} A(a_i - r, b_i - c), \tag{4.2}$$

where each pixel (in row r and column c) is estimated from the sum of the contributions from all MLs of the AMA micromirrors, where (a_i, b_i) is the location of the center of the ML from mirror i. The maximum of I_p when no mirrors are tilted is 1, representing the original lumen output of the projector. The pixels in the target image I_m are also from $0 - 1$. An improvement k

$$k = min(I_p/I_m) \tag{4.3}$$

can then be defined as the minimum each pixel in the target image can be scaled and remain larger than the corresponding pixel in I_p, for every pixel in I_m. The images I_m (original image) and I_p (image from sum of MLs) are constructed for 5 levels of the Gaussian pyramid.

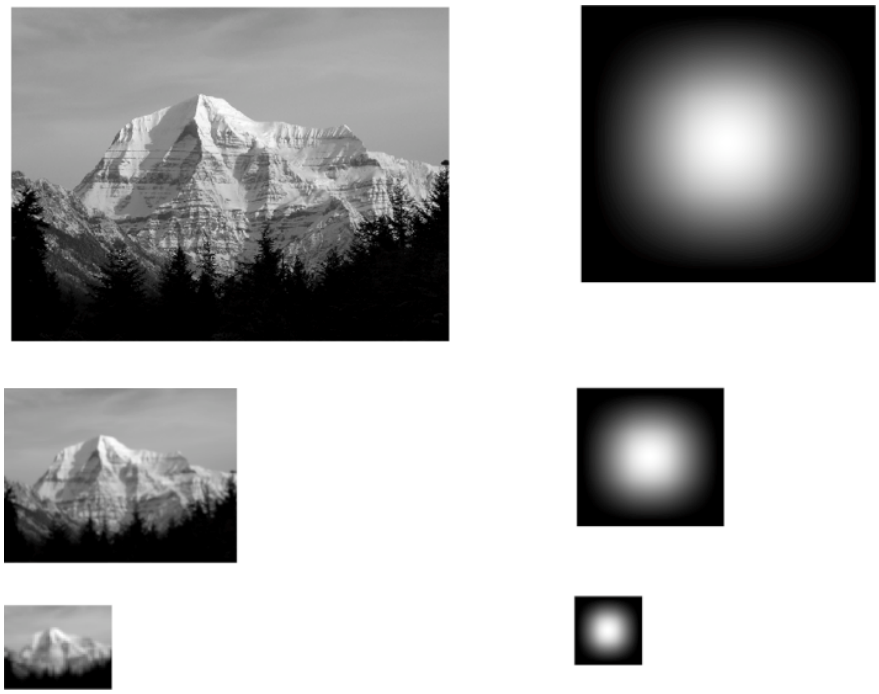

Figure 4.2: Gaussian pyramids are constructed for the image (left) and the ML. Both the ML and the image are downsampled so that each level has half the resolution of the previous.

4.5 Iterative adjustment

After the initial placement of mirrors, a "greedy" algorithm we will call iterative adjustment is used to optimize the placement of the PSFs by iterative redistribution. This algorithm uses the locally optimum choice at each iteration stage, assuming that in most cases this will lead close to the global optimum. The algorithm described below has not been optimized for real-time application, although using pyramids is a highly efficient way of performing large Gaussian blurs, and are used in other real-time applications such as (Rempel et al. 2007). This algorithm is also not guaranteed to find the optimal placement of mirrors, because it does not search the solution space exhaustively for the optimum. However, it is sufficient for a proof of concept to show the benefits of the AMA projector.

The main steps in the algorithm are outlined in Figure 4.3. It is evaluated for each level, starting at the coarsest. First, the relative dimmest pixel, which is the pixel p in I_p that has the lowest intensity relative to the target image is found, $p = min(I_p/I_m)$. To maximize k, we must reposition MLs to provide more light to p.

The closest ML (A_i) in distance to p is found, and then the point q closest

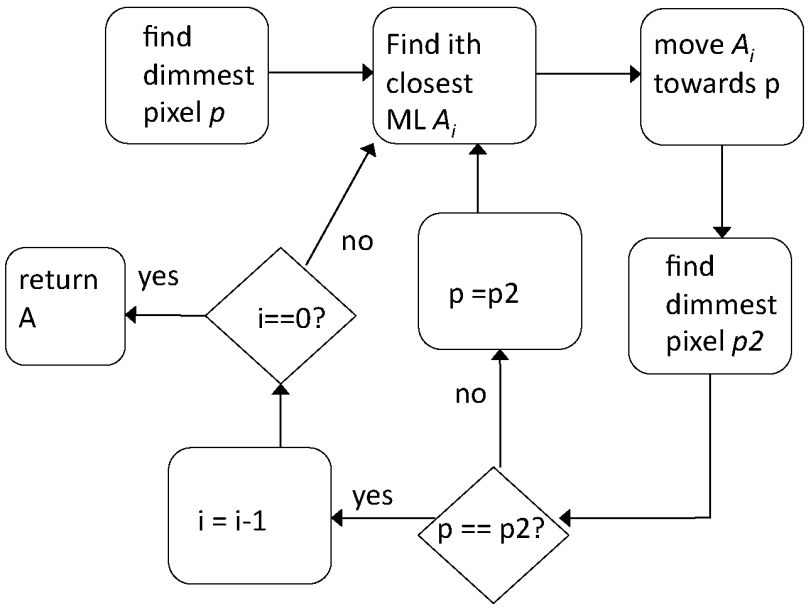

Figure 4.3: Basic steps in iterative adjustment algorithm

to p within the range of A_i is calculated. This ensures that the only solutions within the range of the ML are evaluated. The midpoint between the centre of A_i and q is also computed.

A_i is then removed from the distribution I_p, and I_p/I_m is evaluated with A_i inserted at three points successively: its original location, the midpoint, and point q. If the image obtained when A_i added is not better at q or the midpoint, A_i is left in its original location, and the next-closest ML, A_{i+1} is evaluated in the same manner. Better in this context means $min(I_p/I_m)$ becomes larger. If all of the MLs are iterated through without an improvement, the algorithm has reached a local minima, and the adjustment at this stage stops.

Finally, the placement of the MLs will provide a heterogeneous illumination I_p for the second modulator, in this case the DMD. To produce a projected resulting image close to the original LDR or HDR image, but with boost in contrast and brightness compared to a representation using a uniform illumination of the DMD, the high-resolution correction I_{dmd} for the DMD must be found. The final result for the projected image will be a per-pixel optical multiplication of the two modulators: $kI_m(i,j) = I_{dmd}(i,j)I_p(i,j)$, where k represents the improvement in brightness. I_{dmd} is thus simply equal to $kI_m(i,j)/I_p(i,j)$, quantized to the 256 possible greyscale states of the DMD.

The maximum possible improvement factor k will depend on the image's luminance distribution. If every single displayed pixel is required to have sufficient luminance, $k = min(I_p/I_m)$. At a higher k, some pixels cannot be reproduced at the luminance specified by kI_m. For a completely white image, k would be 1,

the minimum. This would still give the same brightness as a non-AMA projector, minus any extra losses introduced by the AMA. The improvement factor k may also be fixed beforehand. For instance, k could be specified to ensure that a stream of video has a constant relative brightness over time.

There may also be situations where the maximum brightness available is not needed – for instance, when rendering very dark images. In this case, the AMA could direct some light away from the DMD, discarding it before it reaches the DMD and affects the dark level.

As currently implemented, iterative adjustment is fairly computationally intensive, because the ML is represented as an image that at high levels of blur, is almost as large as the target image itself. At full resolution, moving and checking one ML requires as many additions, subtractions and comparisons as the number of pixels in the image, in this case 800×600, proportionally more for higher-resolution images. Over the course of the iterative adjustment process as many locations for the MLs are tested, this becomes the dominant source of computational cost.

4.6 Median cut approach

While performing successive iterative adjustment steps at each level of the Gaussian pyramid is an allocation method that shows the potential for improvement to projector peak brightness, it would be difficult to implement in a real-time setting, which would be necessary for on-board projector conversion given conventional video signals.

An alternative to allocating the MLs via the process of analyzing a Gaussian pyramid is to divide the original image into equal energy zones, and allocate one ML for each zone. The median cut algorithm described in (Debevec 2008), a variation from the original described in (Heckbert 1982), is an efficient method of subdividing an image into zones of approximately equal energy.

To divide the image into 2^n regions, the algorithm progresses as follows: first, the image is added to the region list as a single region. Each region in this list is subdivided along its longest dimension so that the light energy in the image becomes a constant-time operation of adding and subtracting rectangles after the construction of the table.

Although neither Debevec nor Heckbert describe it, it is possible to divide the image up into an arbitrary number of regions, not necessarily a power of two. The cut that subdivides a region into two is placed so that the two portions are unequal as the situation warrants.

The algorithm is thus slightly changed to the following:

1. Add the image into the region list as one region, along with the number of desired divisions d.

2. If $d > 2$, subdivide into *floor*$(d/2)$ (i.e. round $d/2$ down) and *ceil*$(d/2)$ (round up) regions, and add these regions to the list, keeping track of the new d for each. In order to maintain the constraint that all regions at

the final level be of approximately equal energy, choose the cutting line so that it divides the region into portions of energy *floor*$(d/2)$ and *ceil*$(d/2)$ as closely as possible, along the largest dimension of the parent region. Repeat until $d < 3$.

3. If $d = 2$, subdivide region into 2 along the longest dimension.

A useful property of the original algorithm is that the cut is always made on the long side of the rectangle, which minimizes the aspect ratio of the two divisions. To approximately maintain this property when dealing with an arbitrary desired final number of regions, the region is divided into only 2 or 3 section per iteration. The cut is always done at a place that reflects this ratio, in order to keep the energy of all regions approximately equal.

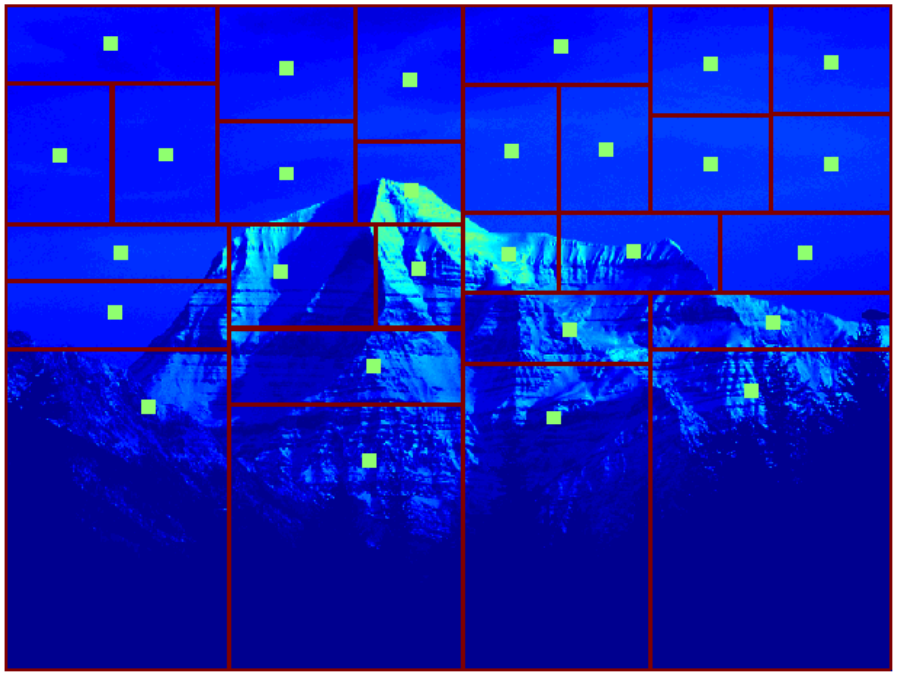

Figure 4.4: An image divided to 28 regions using the median cut algorithm, with the centroids in each region represented as dots.

Figure 4.4 shows an image cut into 28 regions. The centroids of each region are marked with squares. An ML is placed with its centre at the centroid of each region. The main advantage of dividing the image in this manner is that the image can be divided quickly into as many regions as there are MLs. With the Gaussian pyramid scheme, there are still many more pixels than there are MLs at the coarsest level of the pyramid, so the initial allocation is much less representative of the actual intensity distribution of the pixels.

One potential drawback of using this scheme for ML allocation is that the size of each region is a function only of its summed light energy, without regard to the size of the ML. Very small regions might only fit a small fraction of the total light energy of an ML, and very large regions might be larger than a single ML. Also, an equal aspect ratio of a region is not guaranteed, so some regions might be of much different shape than an ML.

4.6.1 Adjusting for limited ML range

Another drawback is that the MLs could have a limited range of movement from their original positions due to a limited tilt angle of the micromirrors, which is not taken into account in the above algorithm. In the Euclidean bipartite minimum matching problem (Agarwal and Varadarajan 2004), we are given an equal number of points from two sets, and would like to match a point from one set with a distinct point from the other set, so that the sum of distances between the paired points is minimized. We can use an algorithm that solves this problem to match each centroid from the median cut step to a location of a ML in its rest (non-tilted) state, and thus minimize the sum total distance between the pairs in the two groups. This will minimize the sum total angle that the mirrors must tilt to achieve the points specified in the median cut solution set. The Euclidean bipartite matching problem can be done in polynomial time using Kuhn's Hungarian approach (Kuhn 1955).

The bipartite minimum matching problem does not limit any one pair to be less than a given amount, only that the sum total of all pairs is minimized. This means that there might be some pairs that exceed the range of motion of an ML for a given maximum mirror tilt angle. If any of the distances between the non-tilted ML and the centroid are larger than the range, they are placed at the furthest point along the line that connects the two points that they can reach. Figure 4.5 shows the results of using the median cut algorithm on the image in Figure 4.4. Figure 4.6 shows the estimated light distribution given by the locations in Figure 4.5.

Because the MLs are only added after their locations are determined, the computational complexity of this algorithm is much less than that of the Gaussian pyramid with iterative adjustment. Many of the operations for building the median cut data structure could be performed in parallel. It is anticipated that a variant of this algorithm could be implemented on board the projector to interpret image data and direct the AMA at video rates.

4.6.2 Median cut with iterative adjustment

If speed is not an overriding concern, the solution obtained by using the median cut algorithm can be optimized by performing an iterative adjustment step as described above in Section 4.5. After the initial median cut solution is determined, the pixel with the minimum improvement is found, the closest ML moved towards it, and the image is evaluated at points along the path. Figure 4.7 shows the same example as in Figure 4.5 after such an adjustment. Some

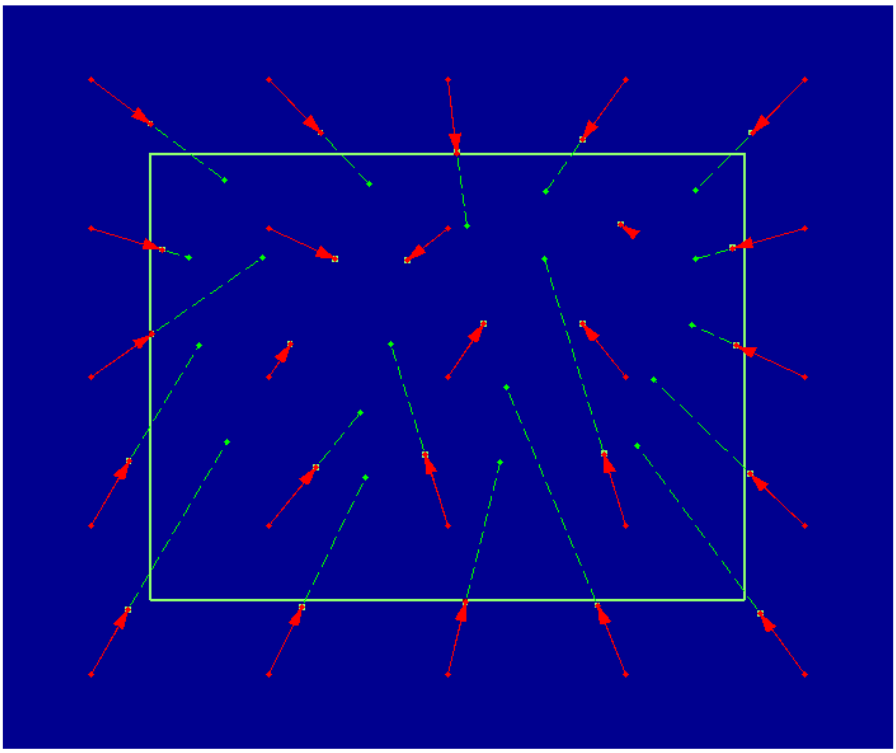

Figure 4.5: Diagram of placements of MLs for the image in Figure 4.4 using the median cut algorithm, with a mirror range of 100 pixels. The starting (small red dots) and ending (arrowheads) locations of each ML are shown, as well as where they would have gone had they had unlimited range (green dotted line). The border of the DMD is shown as a rectangle. Some of the MLs start from beyond the border because of the required overfill.

of the locations are no longer at a point directly on the path from their original location to where the median cut algorithm would have placed them. In this particular example, the improvement factor k was 1.52 for the original median cut placement, and a range 100 pixels. After optimization, k rose to 1.85 with the same range.

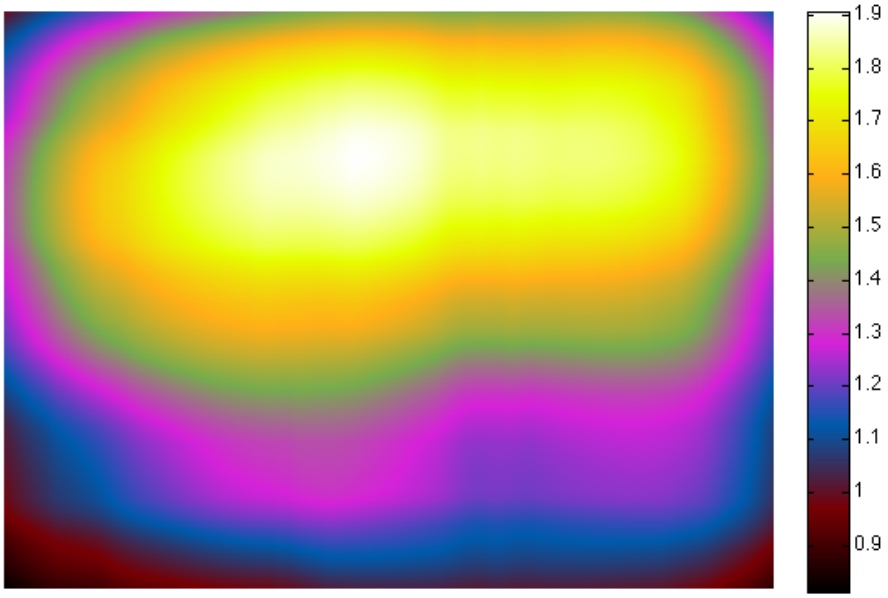

Figure 4.6: False-colour estimation of light distribution given by the ML locations in Figure 4.5. The scale bar indicates improvement factor from a baseline of 1.

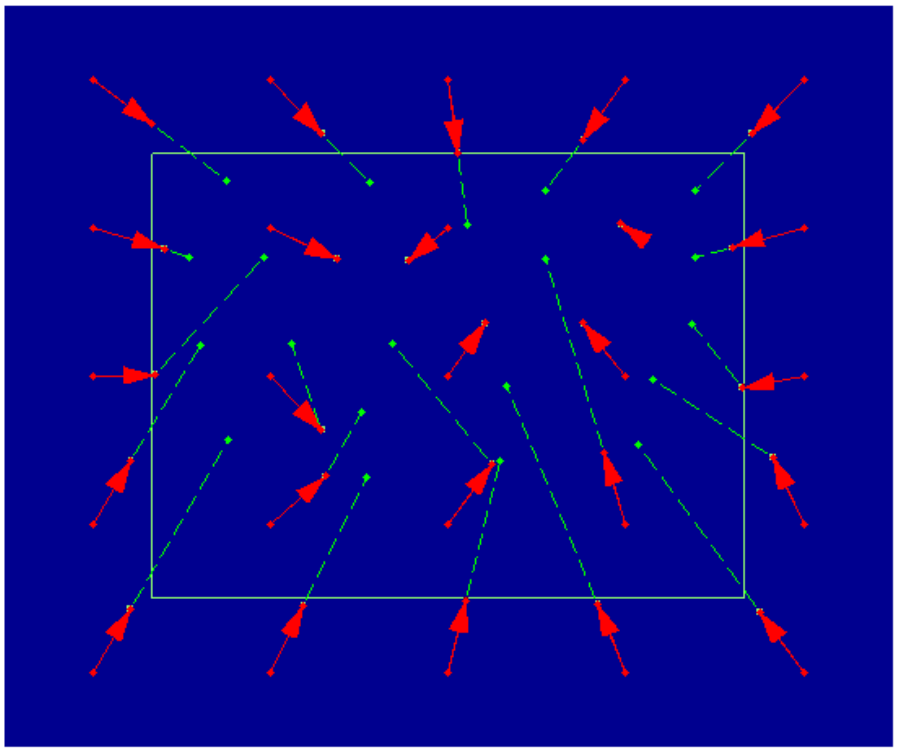

Figure 4.7: Adjusted ML locations from 4.5 after an additional optimization step. The orientation of the green dotted lines show that some of the MLs have been diverted. They are no longer all on the same path from starting position (red) to ending position (arrow terminus) to ideal position according to the median cut algorithm (end of dotted line).

4.7 Simulation results

The algorithms described above were implemented in Matlab. With them we calculate an estimated improvement factor, k, the amount we are able to increase the peak brightness of the scene, while maintaining adequate illumination for the entire image. We neglect any additional losses introduced by the AMA, such as from fill-factor.

The final achievable improvement k can be used to evaluate the choices for ML size, range and number given as initial conditions to the simulation. To determine the best combination of these parameters to guide further development, the variables examined are:

- **Range:** The distance, in pixels, along which a mobile light source can move. This depends on the optical system and the maximum tilt angle achievable by the AMA micromirrors.
- **Blur:** The PSF obtained in Chapter 2 determines the spatial localization of a mobile light source. As discussed previously, the more the disparity, the larger the blur, and also the larger the range of the mobile light source.
- **Number of mirrors:** The number of mirrors within the AMA. With more mirrors, more details in the image can be covered with separate emphasis. It is therefore estimated that a larger number of mirrors will provide more improvement. However, more mirrors also imply a greater physical cost and complexity of the mirror array and its drivers.

The range of each mirror and the amount of blur are related, and can be approximated using the methods detailed in Chapter 2. With measurements or assumptions of the blur and achievable tilt angle of an AMA mirror, the allocation system can be used to simulate the projector's effectiveness. The simulation can then be done completely on the DMD plane, so that we can evaluate the effectiveness of the allocation algorithm, and examine how the physical parameters impact the degree to which the AMA can improve projector performance.

The first step to use the allocation algorithm as an evaluative tool begins with an AMA configuration with no mirrors tilted. Five levels of blur, characterized by the size of the convolution kernel, were tested. These levels were determined by the disparity distances in Table 4.1, as explained in Chapter 2. Figure 4.8 shows each levels of blur for one ML from a 5×5 array with an original size of 200×160 pixels. As the blur increases, the intensity becomes progressively more diffuse and less localized. The peak intensity also decreases. For each of the five levels of blur, five AMA configurations were tested. The smallest number of mirrors was a 3×3 array, followed by 4×7, 5×5, 8×8, and 10×10. We included the 4×7 array because that was the dimensions of the AMA mirrors used in the prototype. In all configurations, the die size of the AMA was held constant, so the number of mirrors determines the relative size of the ML, with arrays having more mirrors also having proportionately smaller ML sizes.

Table 4.1: Disparities tested, their corresponding blur kernel sizes, and the full-width at half maximum (FWHM) of the blur kernel.

Disparity (mm)	Kernel size (pixels)	Kernel FWHM (pixels)
0	1	1
10	171	120
20	279	197
40	410	290
60	486	343
90	554	391

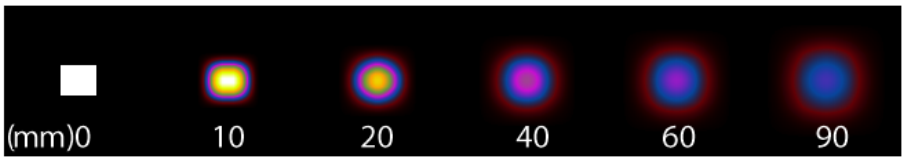

Figure 4.8: Estimated blur for one ML of a 5×5 array, with progressively increasing disparity in millimeters between focal plane and DMD.

Finally, with each combination of blur and number of AMA mirrors, different ranges of movement (in pixels) of the MLs were tested: 50, 100, 300, 500, and 1000. 1000 pixels was included to allow an ML in one corner of a 800x600 image to be moved to the opposite corner. Since the centre of some of the MLs may originate outside of the DMD entirely, this is not quite enough to allow for unfettered movement, but for the images tested this never represented a constraint.

In the physical system, the range of movement and the blur are related: the smaller the blur, the smaller the range of movement. For these simulations, we chose a larger range of values than what is physically possible given the limitations of the current prototype, in order to fully map the tradeoffs between different parameters.

Test images were chosen to give different views into the behaviour of the system. The first, "Mt. Robson," shown in Figure 4.9, has dark section at the bottom and section of high intensity in the sky and on the peak of the mountain. It is anticipated that light will be moved away from the bottom to those regions near the top of the image with the highest intensity. The second test image, "Rock Beach" shown in Figure 4.10 also has a bright sky region, but most of the darker regions are interspersed with brighter regions, which will make it more difficult to move light from one region to another without affecting fine detail. The third, "ANSI checkerboard" in Figure 4.11 is an artificial test pattern that is often used to evaluate the contrast of a projector. It has eight regions of full-intensity white interspersed with eight regions of full-intensity black. It is

Table 4.2: Image characteristics

Pixel attributes	Mt. Robson	Rocky beach	ANSI
mean	44.45	45.35	128
st. dev	43.04	67.20	128
general	dark section at bottom bright section top	intensity interspersed	full on/off sections

anticipated that an AMA projector will not be able to improve this image to a large degree due to the sharp demarcation of bright and dark areas.

Figure 4.9: Original sample image 1: Mt. Robson.

Before applying the AMA algorithm, it is useful to list some first-order image statistics to get an estimate as to potentially how much each image could be improved. After an inverse gamma has been applied to the Mt. Robson image to obtain linear luminance changes, the mean pixel value becomes 44.45 out of 255 (17%), with a standard deviation of 43.04.

For the Rocky beach image (Figure 4.10), the mean is 46.34 (18%), with a standard deviation of 67.2, a minimum pixel value of 0, and maximum of 255. For the ANSI checkerboard image in Figure 4.11, there is only two pixel values: 0 and 255. They are evenly distributed, giving a mean pixel value of 128 out of 255 (50%), and a standard deviation also of 128. These values are summarized in Table 4.2.

Figure 4.10: Original sample image 2: Rocky beach.

The goal of these simulations is to determine whether an AMA projector will still be tenable given the optical constraints given in Chapter 2 and the mechanical constraints given in Chapter 3. In terms of ease of engineering, a larger blur is easier to use than a smaller blur; a smaller range is easier to obtain than a larger range, and less mirrors would be preferable to more mirrors in the array.

All of these parameters can be modified to a certain extent. Both range and blur are constrained by optical limitations. The workable range can be extended and blur reduced by limiting the étendue of the projector lamp, but only within limits, since this will also reduce the total lumens and therefore the efficiency of the system. The number of mirrors is primarily constrained by the cost of the drivers, but also has an effect on the fill-factor of the AMA. More AMA mirrors also require more computation time, the degree of which depends on the algorithm chosen.

The following simulations are not meant to provide an exhaustive survey of all the possible parameter values that might effect an AMA projector's performance, but rather probes parameter values that seem technically feasible. They show that the improvement is possible even with limited ML range and relatively large blur kernels. They also provide a framework to use for the design of a more polished prototype. They will provide a guide to the ramifications of micromirror design choices and projector optics. At that point, more targeted simulations using smaller ranges of blur, range and mirror numbers will

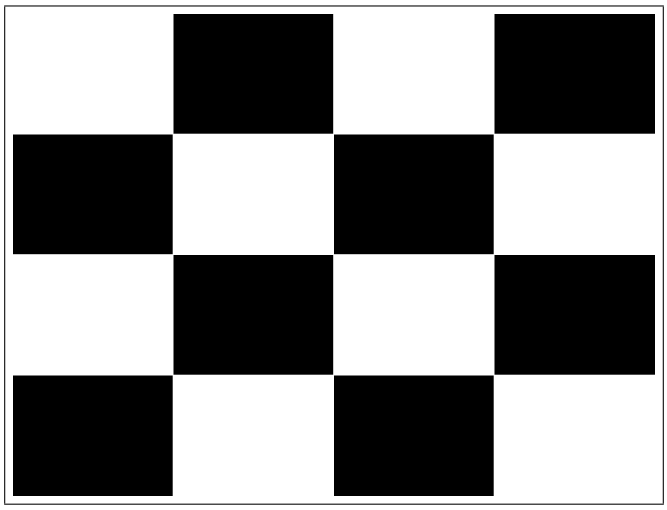

Figure 4.11: Sample image 3: ANSI checkerboard.

Table 4.3: Mechanical tilt angle (degrees) needed for select levels of blur and disparity, assuming a magnification of 1.

	blur disparity (mm)				
range (pixels)	10	20	40	60	90
50	2.0	1.0	0.5	0.3	0.2
100	4.0	2.0	1.0	0.7	0.5
300	11.4	5.9	3.0	2.0	1.3
500	17.5	9.7	5.0	3.3	2.2
1000	27.2	17.5	9.7	6.8	4.4

be needed to further optimize the system.

Table 4.3 shows the implications of tested choices of blur and range on the required tilt angle of the AMA mirrors. The numbers were calculated using Equation 2.12, assuming a magnification of 1. Combinations that include low levels of blur with high range require the most angle. Limiting the range has several implications. While 300 pixels is less than half the distance across an 800x600 image, the tip/tilt mirrors near the centre of the image could move the ML 300 pixels in any direction, for a maximum of 600 pixel range.

While a large range of movement is expected to produce the best results in the simulations, the larger tilt angles needed to produce this movement have an adverse effect on the étendue of the resulting beam. Diverting light with the AMA from one region to another changes the angle of the incoming beam. The projector lens pupil accepts a beam with an angle of 24°. Therefore, some of the more extreme possibilities such as 1000 pixel range and 10 mm blur disparity, which requires 27° tilt, resulting in a change in beam angle of 54°,

are not physically practical, but are included here for the sake of completeness in identifying trends.

4.7.1 Gaussian pyramid results

Figures 4.12, 4.13, and 4.14 show the improvement obtained for each of the test parameter combinations using the Gaussian pyramid with iterative adjustment approach.

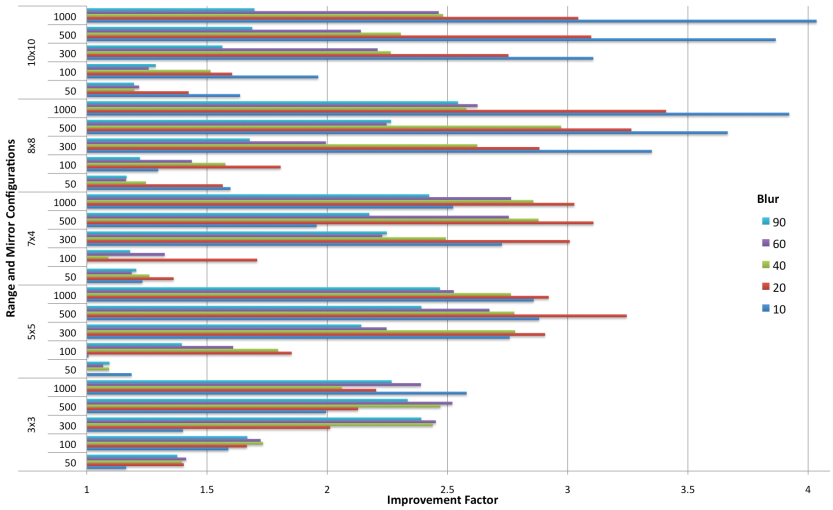

Figure 4.12: Gaussian pyramid results for the Mt. Robson image. The improvement factor k possible with various mirror number combinations, range of the mirrors in pixels, and blur disparities are shown as entries on the bar graph.

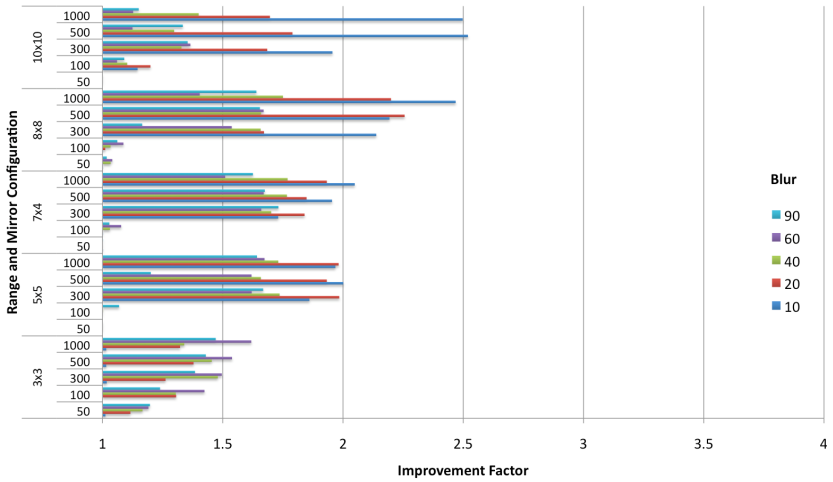

Figure 4.13: Gaussian pyramid results for the Rocky Beach image.

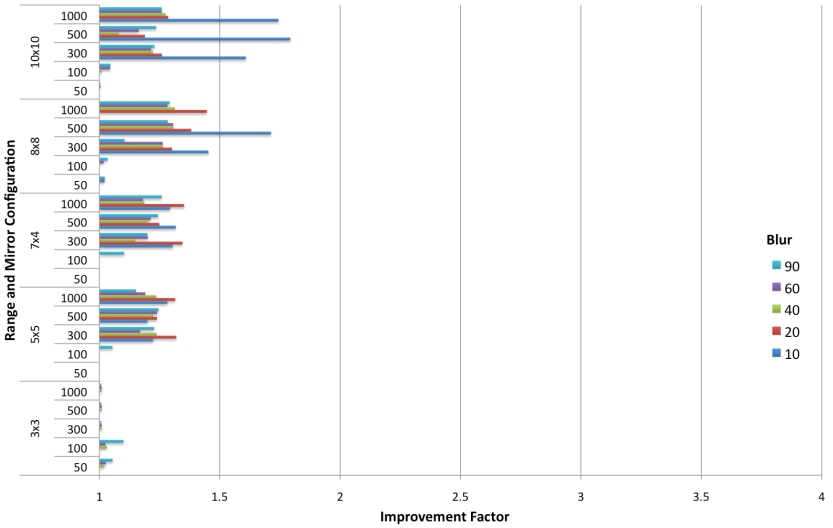

Figure 4.14: Gaussian pyramid results for the ANSI checkerboard image.

4.7.2 Median cut results

The median cut algorithm did not produce results as effective as those produced with the Gaussian pyramid approach. Because the results for the median cut algorithm alone were so small, we include only the results for the Mt. Robson image in Figure 4.15. The median cut with iterative adjustment is included in Figures 4.16, 4.17, and 4.18. Notice that in some of the conditions, notably blurs from 10 and 20mm disparity, there was no improvement, corresponding to a k of 1.

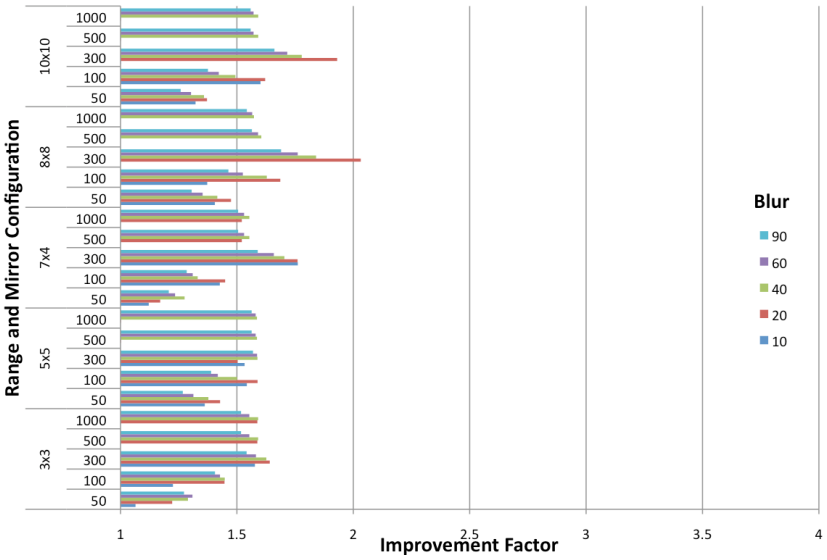

Figure 4.15: Results from allocating the MLs using the median cut algorithm for the Mt. Robson image.

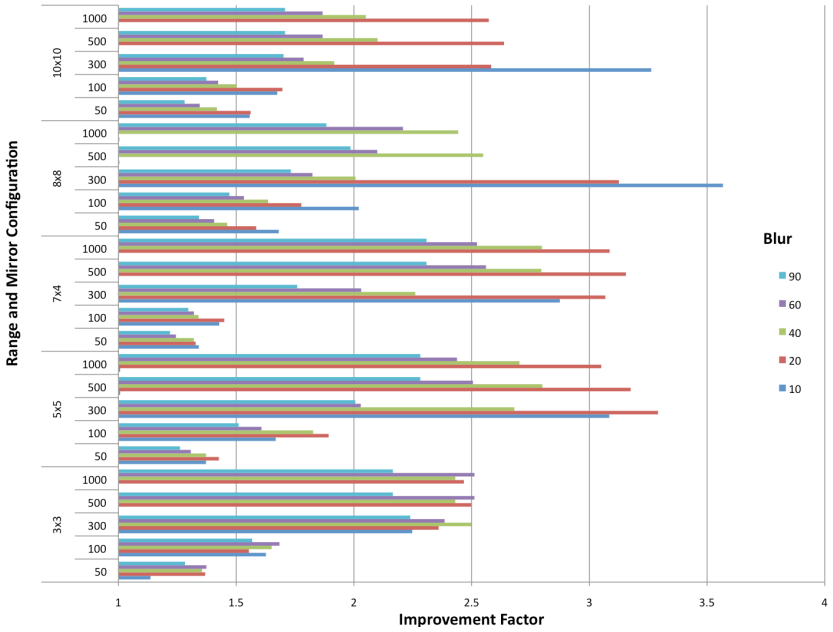

Figure 4.16: Results from allocating the MLs using the modified median cut algorithm for the Mt. Robson image.

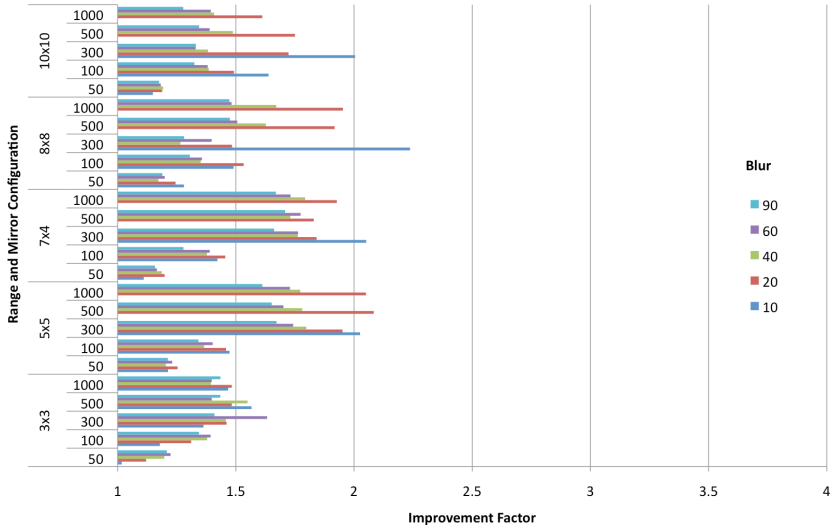

Figure 4.17: Results from allocating the MLs using the modified median cut algorithm for the Rocky Beach image.

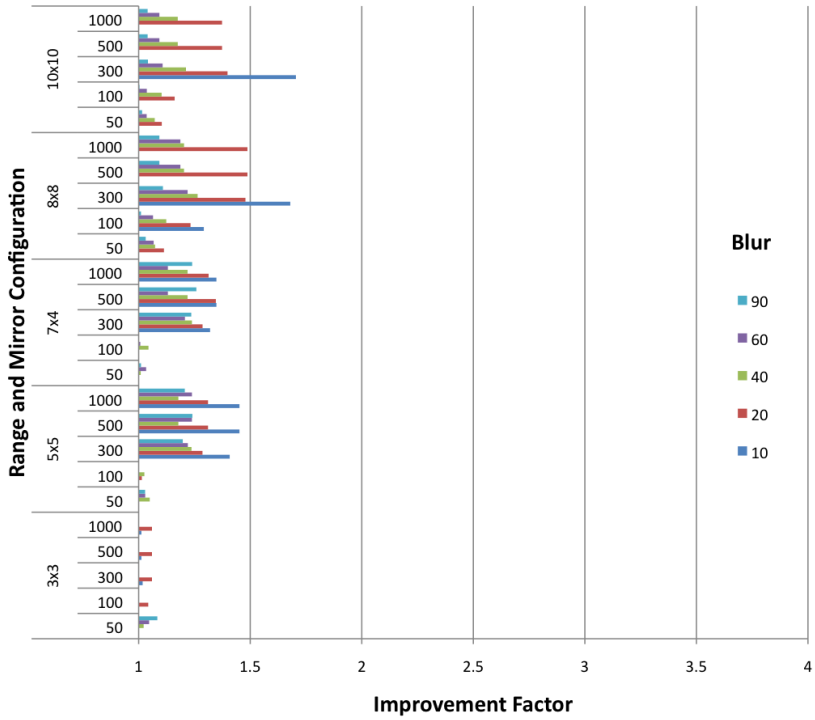

Figure 4.18: Results from allocating the MLs using the modified median cut algorithm for the ANSI checkerboard image.

4.7.3 Observations

The results above show that the Gaussian pyramid algorithm consistently returns the best solution for all three images tested. The average k over all conditions for the pyramid algorithm was 1.56, while the median cut alone was 1.29. The median cut with an additional iterative adjustment step was 1.49, although that adds the most computationally expensive part of the Gaussian pyramid approach, and thus does not provide much of the speed advantage that the median cut approach alone does. The relative differences between algorithms are consistent over the three images tested.

Because the score is set by the pixel with the least light relative to its need, outliers can have a large effect on the final outcome. In the results for the modified median cut algorithm shown in Figure 4.16, the scores for an 8x8 array for 10 blur, ranges 500 and 1000 show no improvement, even though a range of 300 has an improvement of 3.5. Because the algorithms do not do an exhaustive search, the solution that was found for 300 range is not found for 500 or 1000 in this case.

Range

Ignoring the physical implications for the moment, in terms of possible improvement factor, the more range, the better, because solutions obtained for lesser ranges will always be achievable by higher ranges (although because of the non-exhaustive search, it isn't guaranteed that the proposed algorithms will find them). While the iterative algorithm described in Section 4.4 incorporates with the maximum range at every iteration, the median cut algorithm applies the limits of the maximum range after a solution is found, so it is assumed that the smaller ranges will affect Median Cut results the most.

For all algorithms tested, limiting the mirror movement to 50 and 100 pixels severely impacts the possible improvement, irrespective of blur level or mirror numbers. Since range depends on disparity and mirror tilt angle, smaller ranges would be easier to engineer. They would also cause the least amount of an ancillary light loss due to the change of incident angle on the DMD. For this reason, the range of 300 pixels is chosen as the optimal of those we have tested. This range should be re-evaluated with every new physical prototype, when the assumptions detailed above can be validated and adjusted according to the prototype's measured parameters.

Blur

It was anticipated that the lower the blur, the more opportunity there would be for improvement. Less blur means a smaller ML size, so the intensity could be better targeted to where it is needed. The results show that with a larger number of mirrors, the best results are achieved with high range and low blur. As the number of mirrors is decreased, the differences between k scores for differing levels of blur become less pronounced.

From the results of the simulation, there was a clear difference between the pyramid algorithm and the median cut depending on blur level. For the pyramid algorithm, less blur gave better results averaged over all other conditions. For the median cut with iterative adjustment, the trend was less pronounced. Individual inspection of the results revealed that a relatively small number of under-illuminated pixels brought down the results, rendering the change in blur less of an important factor in determining k. No blur at all was also tested, but no improvement was ever achieved, so it was omitted from the graphs.

Number of mirrors

From first principles, it seems evident that the more separately-controllable mirrors are in the array, the better the opportunity for improvement, because they offer more degrees of freedom to tailor the light distribution on the DMD to the desired image. More mirrors can separately cover more features in the image. There are more opportunities to stack MLs. Although the size also depends on blur, more mirrors also means a smaller ML. Qualitatively, the bar graphs show clear upwards trends as the number of mirrors increase.

Quantitatively, the number of mirrors revealed a difference in the two main algorithms, Gaussian pyramid and median cut. Averaged over all other conditions, the pyramid algorithm did show a clear upward trend as the number of mirrors increased. The median cut algorithm, however, did not show much improvement as the number of mirrors increased.

Image type

The results show that there is a large variation between improvement factors depending on the image. The overall average improvement factors for the Mt. Robson, Rocky beach, and ANSI images are 1.81, 1.45, and 1.09 respectively. Table 4.2 shows some possible reasons for this. ANSI has very high frequency components that are not fit well by the blurred MLs, as exemplified by its high standard deviation of pixel intensities. Mt. Robson had a lower standard deviation, but still large differences in pixel intensity over the image, with a dark section at the bottom and light at top, typical of many nature scenes. Rocky beach, although with a similar mean pixel intensity, does not have the same global distribution of light and dark areas to work with.

4.7.4 Implications for prototype

Now that we have described the implications for blur and range, we can narrow down the choices to those that are physically realizable. While higher angles have been shown in the literature, $\pm 3.5^\circ$ tilt around both the X and Y axis is a reasonable limit for this discussion. As Table 4.3 shows, this is a realistic compromise that is close to that achieved in the literature with multi-user processes, and below what can be achieved with a customized fabrication process.

From the results of the simulations, the combinations of 60 blur and 500 range, or 40 blur and 300 range are closest to this maximum target. The simulations show that the improvement factor found for these conditions are relatively similar. Over all image/algorithm combinations, a blur of 60mm and range of 500 corresponds results in an average of 1.58 improvement factor, while 40/300 gives 1.61. When further broken down by algorithm, 40/300 consistently scores higher than 60/500. Using a worst-case analysis to quantify the clipping due to some of the light tilted by the AMA missing the DMD aperture, 3.5° results in a 20% clipping loss, described in Section 2.5, reducing the improvement factor to 1.26.

These results show that significant improvement is possible for a range of physically-realistic parameters. While these results show that the Gaussian pyramid and median cut approach both have promise, more tests are needed when the exact physical parameters of the projector and beam incident on the DMD are known. Some of the assumptions required to run the tests could have had some impact on the results. In particular, the shape of the Gaussian used in the blur kernel given in Equation 2.24 is influenced by the parameter g, which will have to be re-evaluated as the projector optics change. We chose a value of 0.6 for g in these simulations, as it provides an average, non-peaked distribution. When a prototype where the physical parameters can be properly obtained, these tests should be performed again to verify the performance of the algorithms.

In general, the results show that lower blur, more mirrors, and higher range, all contribute to better possible improvement. In part this depends on the algorithm chosen; the more efficient median cut algorithm works best when the blur is a bit higher, and with a medium range of pixels. For the subset of possibilities that are currently physically realizable, lower blur is more important that higher range.

4.7.5 Sub-frame positioning

The above results assume that the ML is stationary for the duration the image is displayed. A frame of video is normally shown at 30-60Hz, so the micromirrors must be able to be positioned much more quickly than that so they can be stationary for the majority of the frame, then repositioned between frames. As shown in Chapter 3, the MEMS micromirrors we propose can be positioned at 10 kHz, an order of magnitude faster than what is required for this purpose.

We could potentially exploit this additional speed by repositioning the micromirrors within a frame of video. This would allow us, for instance, to cover an area that only requires half of the intensity that an ML offers in half the time, then use the other half of the ML's intensity somewhere else.

Simulations were performed to determine the potential advantage of this sub-frame positioning. To position the MLs four times within one frame, the intensity of the ML was divided by four, and the number of MLs multiplied by four. Four sub-MLs were assigned to every location in the image that the original MLs were assigned to, and the same algorithms as above were performed. The

four sub-MLs could move independently, but were still limited by the same range constraints as before. It was found that for the conditions tested above, there is no advantage to having sub-frame positioning.

4.7.6 Allowing under-illuminated pixels

The above results show that the choice of algorithm had the most significant impact on the improvement factor. A potential reason for this is that all pixels must be illuminated properly. This means that the improvement factor is set by the pixel with the minimum ratio of light given to light required (I_p/I_m). All it takes is one outlying pixel to bring down the improvement factor. Out of the three versions of the algorithms tested, the Gaussian pyramid approach gives the most opportunity to minimize the effect of the outliers during its extensive iterative adjustment steps, perhaps explaining some of the difference in results achieved relative to the median cut approach.

If a particular pixel is an outlier, it may be that the rest of the image could have been shown at a much higher brightness if that pixel were ignored. It may be that several of the pixels will not be noticed by the viewer if they are under-illuminated. We tested how much brighter the image could get if we allowed a certain percentage of pixels to be dimmer (relatively) than the rest of the image.

Figure 4.19 shows the performance of the three algorithms tested for the Mt. Robson image, with a blur from 60 mm disparity, an AMA with 7x4 mirrors, and a range of 300 pixels. Although the Gaussian pyramid (GP) algorithm performs best for the case where all pixels are illuminated correctly, it is eventually outperformed by the two median cut (MC) versions when some pixels are allowed to be under-illuminated. For the case in which all pixels receive enough light, the pyramid representation scores highest with a 2.44 improvement factor, while median cut with iterative adjustment (MCIA) scores 2.32, and just using MC placement scores 1.9.

But at just 200 under-illuminated pixels (which, for an 800x600 image, is 0.004%), GP is overtaken by both median cut versions. GP scores 2.71, while MC alone scores 2.8, and MCIA scores 3.07. That is a 26% increase in brightness over the highest score obtained for perfect reconstruction. The increase raises to 36% if 500 pixels are allowed to be under-illuminated. The other two images tested are consistent with this trend.

Figure 4.19 shows that the median cut algorithms are sensitive to outliers, and much more overall brightness increases can be achieved if the requirements are loosened. Also, visual evidence of the under-illuminated pixels might be difficult to detect due to the inevitable presence of other sorts of display errors. The model of the ML used in the algorithm for a real projector will most likely be an approximation that differs slightly from what is physically present due to manufacturing variability. We estimate how under-illuminating pixels effects perceivable image quality in Section 4.9.

Even if allowing under-illuminated pixels is not desirable under normal conditions, these tests show that there is a lot of flexibility to tailor the performance to different conditions. For instance, when displaying a video sequence on an

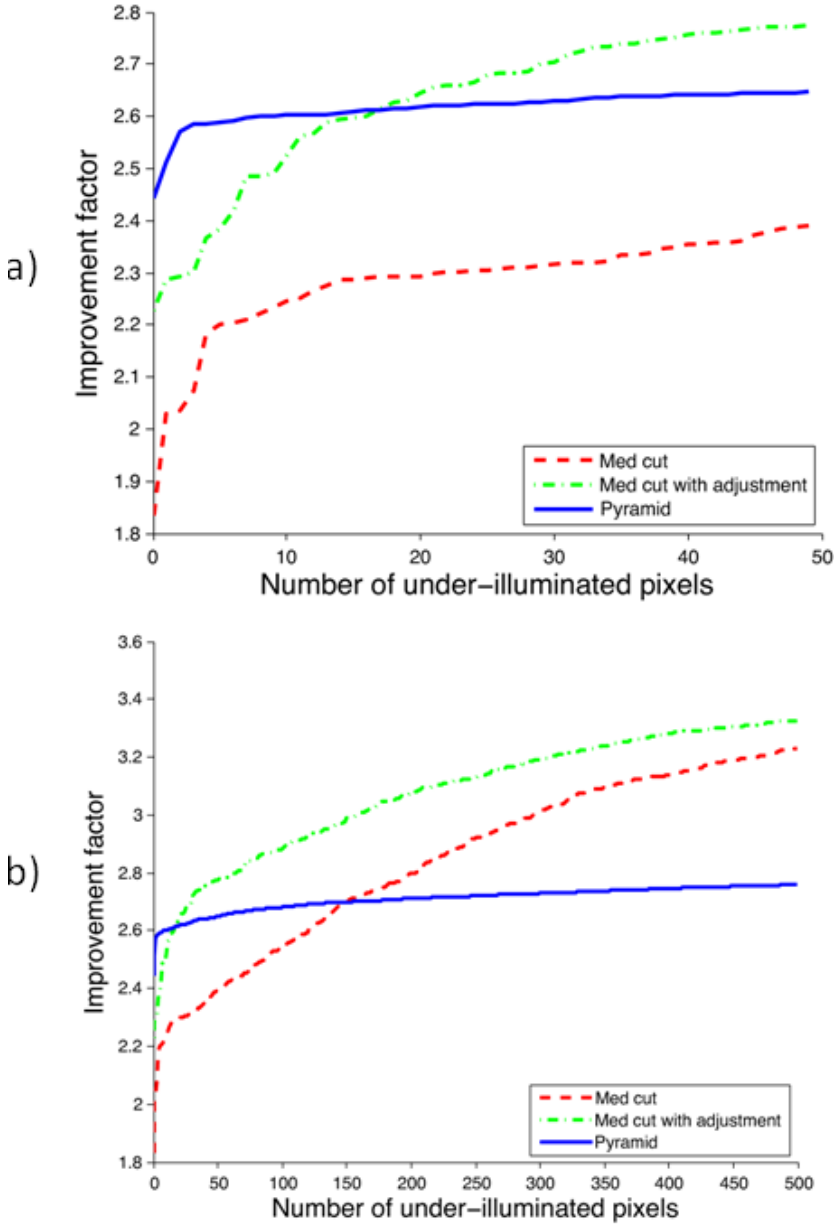

Figure 4.19: Effect of allowing some pixels to be under-illuminated on improvement factor for the Mt. Robson image, for up to a) 50 underilluminated pixels, and b) 500 pixels.

AMA projector, the potential improvement factor will vary over time as the image changes, depending on the sum total intensity in the image and its distribution. If it changes too quickly or the magnitude of the change is too abrupt, the viewer will notice the artificial change in relative brightnesses of the different features in the image. Figure 4.19 shows that fixing the improvement factor for this video sequence doesn't necessarily mean fixing it at the minimum value out of all images in the sequence. By allowing some pixels to be underilluminated, the improvement factor can be set at a higher level that matches what is achievable for an average image in the sequence, not the minimum.

4.8 Estimated system performance

4.8.1 Brightness improvement factor

Adding the AMA will incur some light losses that partially offset the gains described above. We can estimate the efficiency of a production version of the AMA at reflecting incident light by examining the statistics from the DMD. After taking into account the transmittance of the packaging window of 97%, the active area fill-factor f_D of the DMD of 88%, its mirror specular reflectivity r_D of 89.4%, and diffraction efficiency, the DMD reflects $e_D = 68\%$ of incident light (Texas Instruments 2005a).

Although the DMD is a highly-engineered device, the AMA could achieve higher efficiencies because of its different structure and use. The fill-factor of the AMA could be increased, because the AMA requires many less mirrors than the DMD, and thus could have less total space between them. A fill-factor f_A of 95% could be achievable given a different fabrication process, as it is slightly lower than the 96% achieved in (Tsai et al. 2008) for micromirrors that offer performance that would be adequate for this application.

Another factor that would mitigate the light efficiency losses due to the addition of the AMA is that the AMA will replace an existing fold mirror in the projector. We can expect the fold mirror to have a similar specular reflectivity coefficient as the AMA, so r_d can be removed from the calculation of the AMA efficiency.

As described in Section 2.5, another source of loss is due to a portion of the diverted light missing the aperture. For a $\pm 3.5°$ tilt AMA, the estimated losses due to clipping l_c is approximately 20% worst-case, if all the AMA mirrors were tilted to their maximum, and no other mitigation plans such as aperture shaping were included. For this calculation, we assume the mirrors are tilted to an average angle of 2.5°, giving a expected loss from clipping as 10% from Figure 2.11. The expected efficiency of the AMA can thus be calculated *relative* to the DMD reflectivity as

$$e_A = e_D \times \frac{(1 - l_c) f_A}{f_D r_D} \tag{4.4}$$

which results in an estimated efficiency of the AMA as $e_A = 74\%$. Although this is a significant loss of light, we still anticipate that the AMA-enabled projector

would have a net gain of peak brightness given the results from Sections 4.7.1 and 4.7.2. With 40mm blur disparity and 300 pixel range, the average improvement factor was 1.61, giving a net improvement factor of 1.20 after taking into account these losses.

If we allow some pixels to remain under-illuminated as explained in Section 4.7.6, the case becomes even stronger. With 200 under-illuminated pixels, the average score for the median cut algorithm with iterative adjustment was 3.07. After taking into account losses, this still leaves a 2.25 improvement.

4.8.2 Contrast improvement

An AMA projector will also improve projector contrast, both on the low-end with a reduced black level, and on the high-end with increased brightness. The AMA is primarily a redistributive system, so the contrast improvement is the square of the brightness improvement, since the black level would be reduced the same amount. If we use the brightness improvement factor of 1.2 for an image where no under-illuminated pixels allowed, the overall contrast improvement factor would be 1.4. For a system that allows under-illuminated pixels, while the brightness improvement is 2.25, the contrast improvement would be 5, a very substantial change.

Ultimately, the magnitude of improvement will depend on critical design decisions during the projector design phase, decisions that substantially affect any estimate of increased contrast due to the AMA. The major initial decision which will depend on the projector's intended use is the choice of illumination numerical aperture. Typically the DMD illumination angle is increased slightly to increase projector contrast, at the expense of some brightness. This need not be the case given the AMA's effect on contrast, which would increase the overall improvement factor but decrease the contrast improvement factor.

4.9 AMA projector image fidelity

After the introduction of the increase in performance that can be achieved when some pixels are not adequately illuminated, it is natural to wonder about the effect on visual fidelity the addition of the AMA would have on the projected image. In this section we quantify any undesired visual artifacts that are introduced when an AMA is added to the system.

No display available today can perfectly reproduce an arbitrary image. Conventional displays have limited contrast, so there will be pixels that have more or less light than they should. Other problems are display-technology specific. In projectors, there is distortion from the projection lens, and vignetting, the spatial fall-off of brightness from center to fringe. These imperfections can be aggravated by high-gain screens causing their own brightness variations. Evidence of how frequent projector imperfections arise can be found in the research done in the area of tiling multiple projectors, where imperfections must be corrected before an adequate combined image is formed. Tiling the images from multiple

projectors allows for the display of very high-resolution images over a large area, but inter and intra-projector differences can cause distracting artifacts. A good survey of projector imperfections in the context of tiling is (Majumder et al. 2008).

Conversely, external factors influence the appearance of the projected image even when it is perfectly displayed. The level of ambient illumination in the room determines the minimum dark level of the image. In an environment with significant ambient illumination, the steps between the first few darkest pixel values could be imperceptible, because they correspond to changes in luminance that are too small to perceive. Since the perception of luminance is roughly logarithmic, these same imperceptible intensity levels could be perceptible in a darker room.

In this hypothetical case, an image displayed with an AMA projector with a larger improvement factor (and thus a higher overall brightness) with some pixels under-illuminated might still have more perceivable detail than the same image displayed at a lower brightness with all pixels correctly illuminated, because the missed pixels are not perceivable at that level of ambient illumination.

Colour, which will not be directly addressed in this work, is also a major issue in display accuracy. For instance, (Heckaman 2006) show that adding a white channel to the colour wheel may increase brightness, but negatively affects the colour saturation of the projected image.

The particular errors quantified in this work are those directly attributable to the addition of the AMA. Specifically, we will examine the additional source of error that comes from having a non-homogeneous backlight of lower resolution than the DLP itself. The errors can be categorized into five main areas:

1. **Quantization**: There are only 256 greyscale values in a DLP. There will thus be errors due to quantization as the DMD cannot compensate accurately enough for the continuous, non-homogeneous illumination from the AMA.

2. **Over-illumination**. The DMD doesn't have complete control over contrast, so it will leak light, causing some pixels to be over-illuminated. Other components in the optical path also contribute to leakage from one region to its neighbours, such as scattering from the projector lens. These factors contribute to an amount of per-pixel error. Normally this is expressed as a constant loss of contrast over the whole image, because of the constant illumination source. With a projector that includes an AMA, the light that must be blocked by the DMD varies spatially, causing the amount of light leakage to also vary spatially.

   To estimate over-illumination, it is necessary to have an estimate of the amount of leakage in the projector for a given amount of luminance. This can be estimated from the contrast ratio of the device, and the peak brightness. If the peak brightness of the projector is x, and contrast ratio is $1 : y$, then the estimated brightness of a pixel set to its darkest level (0) in the DMD is x/y. If the pixel is set to a brightness of z, the estimated

result is $z+x/y$. While this is not a perfect estimate due to the variability of lens scattering and possible non-linear effects, it provides a base level to compare the difference in an AMA-projector to a conventional projector.

3. **Under-illumination**. The improvement factor, k, sets the improvement factor of the AMA projector for a given image. If k is set too high, the image cannot be reproduced perfectly because the total luminance available is constrained by the projection lamp. The result could be a loss of texture in the displayed image. Choosing k thus becomes a tradeoff between maximizing total image brightness and maintaining image accuracy.

   To maximize overall brightness, it is tempting to sacrifice some accuracy by allowing some pixels to be under-illuminated. We can ensure this does not happen by specifying to the allocation algorithm that all pixels must have sufficient light, but how many under-illuminated pixels could be present without adversely affecting the viewing experience, and under what conditions?

4. **Simulation imperfections**. In the analysis above, it is assumed that the distribution of light incident on the DMD is perfectly known. In reality, due to real-time constraints there will be approximations made that cause imperfections in the model of the shape of each individual ML, and possibly in its location due to slight micromirror differences across the array that affect their tilt angles. Although these differences can be compensated for by using more complex models and feedback mechanisms, there will always be some error due to this mismatch.

5. **Real time constraints**. In the allocations algorithms proposed above, the final result is optimized at the resolution level of the final image. In a real-time application, the bandwidth of the hardware might not be fast enough to support this. Simplified approximations and/or lower resolutions might be necessary to obtain adequate performance, but would also result in other sources of errors.

4.10 Visual difference prediction

To quantify the effect of these inaccuracies, the simulated optical combination of the AMA image on the DMD (I_P) and the DMD image are compared against the original image. To avoid extensive psychophysical user testing, we use an automated visual difference predictor (VDP) (Ramanarayanan et al. 2007), which has been extended for application to high dynamic range images as well as conventional images (Mantiuk et al. 2005). The VDP algorithm provides a metric to distinguish the subset of differences between two images that a standard human observer would be able to detect. The algorithm filters the images through various stages that mimic the light scattering in the cornea, lens, and retina, and our non-linear response to luminance and contrast. The limitations of the human visual system mean many physical differences between images are not

detectable by a human observer. See (Mantiuk et al. 2005; Trentacoste 2006) for thorough introductions to the methodology used in visual-difference estimates.

Visual-difference prediction has been used before to evaluate novel displays, notably in (Trentacoste 2006) for Brightside/Dolby's high-dynamic range display, the DR-37. Like Trentacoste, we use the implementation HDR-VDP described in (Mantiuk et al. 2005). Our objective is to establish visual equivalence between images made with an AMA-equipped projector, compared to a regular projector that has the same peak brightness as the improved image. For instance, if the peak brightness of the AMA projector when none of the mirrors are tilted is l, and the AMA projector improves the projected image to $2l$ by tilting the AMA mirrors, then this simulated image is compared via VDP to an image perfectly displayed at $2l$. We compare images of equal luminance because luminance can affect the perception of features within the image, and we wanted to keep the differences to those directly attributable to introduction of the AMA.

The VDP algorithm computes for each pixel in the image the probability of the viewer detecting distortion. However, this does not mean that the results returned are accurate on a per-pixel basis. A major step in the VDP pipeline is to process the images with a series of filters sensitive to spatial frequency and orientation that approximate the spatial processing done in the visual cortex. The filtering operations are performed in the frequency domain, so the location of a difference is dependent on its spatial frequency. For computational efficiency, images are not filtered at every spatial frequency or orientation, resulting in banded areas of detection. In reality these bands are smoother, and if more frequencies were tested, the features would be more representative of the shape of the difference (Trentacoste 2006). As long as we keep in mind to interpret the results in terms of a probabilistic area of difference, the VDP is a useful tool for evaluating whether two images would look the same to a typical viewer.

4.10.1 VDP tests on simulation results

Because a real-time performance implementation is outside the scope of this thesis, and the prototype is not in a state where simulation imperfections can be properly evaluated, we presently omit analysis of these sources of error. We first test via simulation whether the artifacts for quantization and the over-illumination issues inherent with an AMA-projector are normally detectable. We tested using a 7x4 array of mirrors, using a range of 300 pixels, and a disparity of 60 mm. With DMD pixels being approximately $14\mu m$ square, we can use simple trigonometry to calculate that the micromirror mechanical tilt angle needed to traverse 300 pixels with this disparity, assuming no magnification, is $\pm 2^\circ$, which is within the range of published micromirrors.

Quantization is achieved as follows. In our simulation, both the image of an ML and I_p, the distribution of light on the DMD from all MLs are stored as double-precision floating-point arrays. Because the DMD values are calculated

as $I_{dmd} = kI_m/I_p$, they are initially floating-point values as well, but are then quantized to 256 values, to I'_{dmd}. The resulting image is then calculated as

$$I_{mq} = I'_{dmd} * I_p. \tag{4.5}$$

In this equation and those that immediately follow, it is implicit that the calculation happens on a per-pixel basis, e.g. with each pixel in the DMD image being multiplied by the corresponding pixel in I_p.

The limited contrast of the projector is simulated by adding a portion of I_p to the resulting image.

$$I_{mc} = I_{mq} + I_p/c, \tag{4.6}$$

where c refers to the contrast ratio of the DMD. For instance, if we take a conservative value of 1 : 400 contrast ratio for the DMD, c is 400. I_{mc} is compared against kI_m, where k is the improvement factor.

The physical parameters of both the target (simulated image) and mask image (perfect image) need to be specified for the VDP algorithm. For these tests, the image was specified to be displayed at a size of 0.81 m by 0.45 m at a distance of 1.2 m.

The peak brightness of the image also must be specified. As described in Section 1.1, in the standard ANSI/SMPTE 196M the Society of Motion Picture Television Engineers recommends that the screen luminance be set at 14 foot-lamberts (SMPTE 2003). A foot-lambert (fL) is approximately $3.426cd/m^2$, so 14 fL is $48cd/m^2$.

Such relatively low peak luminances are sufficient for conventional film images shown in a theatre setting where the ambient light is minimized. For the home or business with much higher ambient light levels, higher brightness capabilities are needed. The recommended luminance for a television according to the SMPTE is 50 fL, or $170cd/m^2$. Modern LCD and plasma televisions are approximately 300 cd/m^2, whereas the DR-37 from Seetzen et al. can reach up to 8500 cd/m^2.

To evaluate an AMA projector using the VDP algorithm, we chose an intermediate level of $200cd/m^2$ to use for peak brightness. The median cut algorithm with iterative adjustment on the Mt. Robson image obtained a 2.22 improvement for 'perfect' reconstruction. Simulation of the quantization and contrast were added to the simulation results and tested with VDP. At this peak brightness, the VDP algorithm predicted that there was a probability of detection $P > 95\%$ for 0.0002% of the pixels, as shown in Figure 4.20. For an 800×600 image, this works out to 96 pixels. The affected pixels are hardly visible at the bottom right-hand corner of the image.

4.10.2 Quantifying the effect of some under-illuminated pixels

Given that quantization and contrast issues did not significantly affect the image at this peak brightness, there is potential to significantly increase the overall brightness of the image if the constraints for perfect reproduction are relaxed.

Figure 4.20: VDP results from testing the effect of quantization, limited contrast and allowing no under-illuminated pixels.

Table 4.4 shows the results of increasing the improvement factor to various levels, depend on the percentage of pixels that become under-illuminated. For 0.01% pixel error, which would allow for an improvement factor of 2.77, VDP calculated that only 0.0002% of pixels might be noticed. That is the same as the VDP result when all pixels are sufficiently illuminated, and only contrast and quantization artifacts are taken into account.

It is worth noting that we are not modifying I_p with these tests, just the image sent to the DMD. Also, it is not necessarily the brightest pixels that are under-illuminated; they could be anywhere in the image. In this particular image they start out in the bottom corner from where light has been redirected to other brighter parts of the image.

The results in Table 4.4 show that there is a significant potential for further increasing the peak brightness if the constraints for perfect results are relaxed. More than $5x$ improvement can be obtained with 5% of pixels under-illuminated. The VDP statistics suggest it is possible to obtain a lot more improvement without sacrificing visual quality. When the image is inflated too much, however, the peaks of the image no longer are as bright as they were originally, and become easily-recognizable difference features. Subjectively, there was a big jump between 1% and 2% of pixels underexposed; the extra 1% seems to come mostly from the central peak highlight, as shown in Figure 4.21.

Visual-difference prediction allows us to make educated estimates as to the

Percentage pixels under-illuminated	improvement	fraction of pixels $P > 95\%$
0.0	2.22	0.0002%
0.01	2.77	0.0002%
0.1	3.32	0.0065%
1.0	3.89	0.0923%
2.0	4.29	0.37%
5.0	5.07	1.13%

Table 4.4: Results of allowing a percentage of pixels to be under-illuminated.

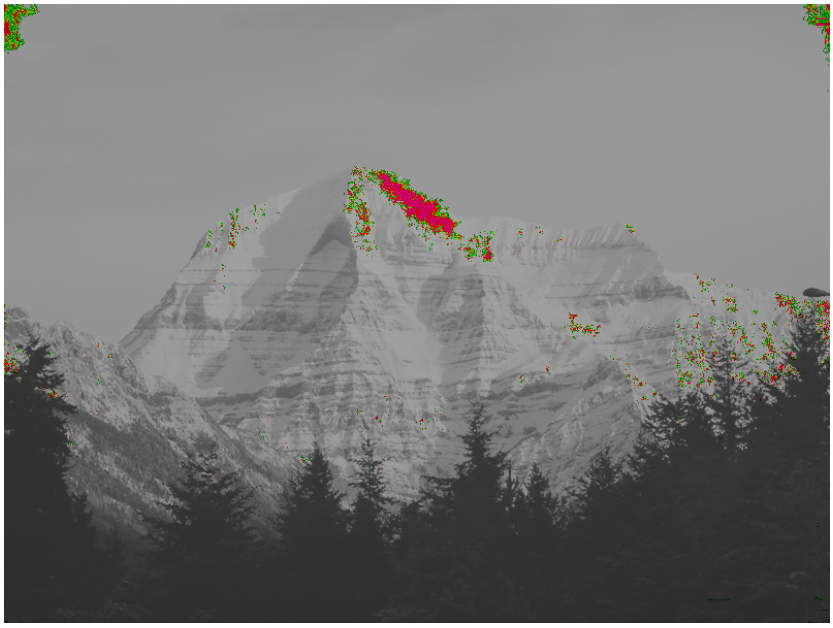

Figure 4.21: VDP results from testing the effect of quantization, limited contrast and allowing 2% under-illuminated pixels.

effect of adding an AMA to the illumination path of the projector. We have shown that significant brightness increases can be achieved with no loss in detectable picture quality. There is also room for further increases if some losses in pixel quality are acceptable. In the future, with detailed knowledge of the ambient light level and projector characteristics such as its gamma and native peak brightness, there may be opportunities to distinguish between details that are and are not perceptible, and intelligently allocate the resources from the AMA accordingly. Image details that would not be perceptible can be identified before allocation, allowing for higher, more targeted brightness increases. Perception-

driven resource allocation using an AMA could help projectors become more efficient and responsive to different environments, rather than putting the onus solely on the viewer to manually make changes to the room and the projector settings to obtain an acceptable picture.

Chapter 5

Physical implementation

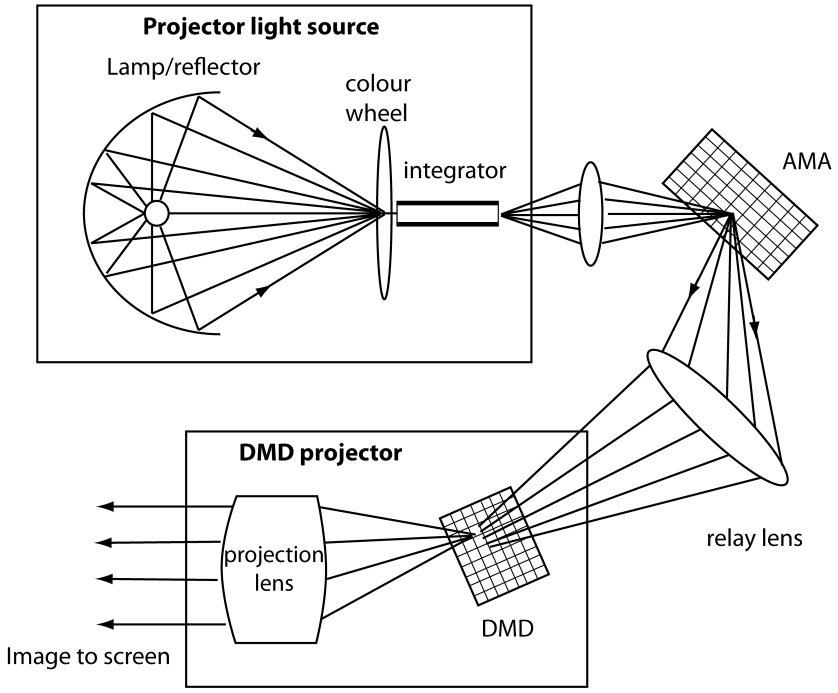

Figure 5.1: Schematic of prototype, showing the major components.

5.1 Prototype overview

We have built a prototype to demonstrate dual light modulation using an AMA and DMD. The components of prototype are mounted on an optical table to allow for precise adjustments. Figure 5.1 shows a block diagram with the prototype components.

5.1.1 Optical system

Initially, we planned on adding the AMA to an existing projector, a Knoll Systems HD102 DMD projector (identical to the Infocus model 4805), without modifying any of the other optical components. This proved impractical due to the structure of the cavity that encloses the end of the integrating rod, folding mirror and the DMD, where the AMA would have had to be placed. The cavity was too small to place the AMA and still have adequate room for a relay lens between the AMA and DMD.

Instead, we used this projector as a light source, cutting a hole to allow the light from the lamp to exit the projector before it reaches its DMD. The light is collected with a 60mm lens and directed to the AMA. A 45mm lens relays the light from the AMA to a second projector, a Mitsubishi PK20 projector that is relatively easy to open up to allow access to the DMD. The light incident to the DMD is reflected normally through a prism to the PK20 projection lens onto a screen. Figure 5.2 shows a photograph of the prototype.

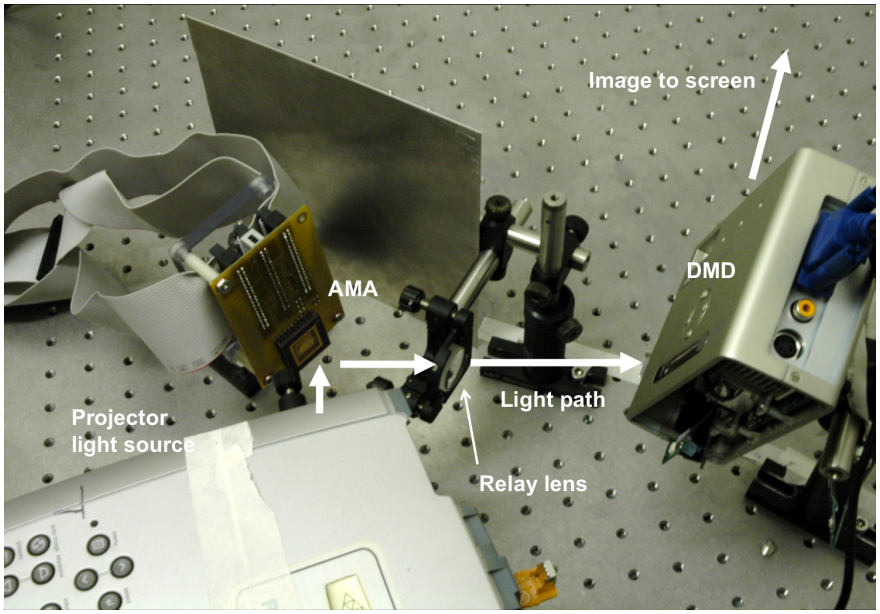

Figure 5.2: Photograph of prototype, including projector light source, AMA, relay lens, and DMD.

Both of the projectors use sequential colour. In the Knoll projector, a 6-section colour wheel of repeating red, green and blue filters is used, while the PK20 uses red, green and blue LEDs. The light from the LEDs is not used for this prototype. We use a video splitter to send the same VGA signal to both projectors, so there is frame-level synchronization. Since the VGA signal does not dictate the orders of colours in a colour-sequential display, the colours

within the frame are not synchronized, and so are not correct. For this reason we use only greyscale images with the prototype.

In the future, the colour wheel from the Knoll projector can be taken out entirely to obtain an approximately $3x$ increase in brightness for a greyscale image. Alternatively, a different light source that can be synchronized with the PK20 LEDs should be investigated to allow for colour images.

5.1.2 AMA driver

CMC Microsystems offers packaging options for the Micragem AMA, the most suitable for this application being an 84-pin pin-grid array (PGA) package. The fabricated AMA chip is mounted in the package, and connected to the pins of the package with up to 84 wire bonds. In our design, each composite mirror requires three independent analog voltages for the three electrodes that control the mirrors. With 84 pins, the maximum number of mirrors that can be controlled is thus 28. The prototype chip has 7×4 mirrors. Since 1 pin must be reserved to electrically ground to the mirror surface, one of the mirrors only has two electrodes activated.

As shown in Figure 5.2, the PGA package is electrically connected to a small printed circuit board though a socket that allows for easy chip insertion and removal. The printed circuit board (PCB) was designed to take 3 standard 40-pin ribbon cables and route their signals to the 84 pins of the socket and then the PGA. The three ribbon cables are in turn connected to three Analog Devices AD5535 digital-to-analog converter (DAC) evaluation boards which provide the high-voltage signals, as shown in Figure 5.3. The converters have 14-bit resolution between 0 and 200V. Each DAC chip provides 32 channels of analog output, so there are 96 possible controls in total, meaning that every electrode in the AMA can be controlled using this apparatus. To control the three DACs, We selected the Altera DE2 floating-point gate FPGA system as a convenient testbed for the prototype. Software was written for the DE2's Cyclone II FPGA to supply the Serial Peripheral Interface (SPI) required as input by the DAC. As implemented, the FPGA can provide a clock signal to the DAC at the AD5535's highest rate possible, 30 MHz. At this clock signal rate, the channel update rate is 1.2 MHz. All 32 channels can thus be updated at a rate of 37 kHz, which is more than sufficient given the resonance frequency of the micromirrors is approximately 10 kHz.

5.1.3 Synchronization

The Altera DE2 development board includes a VGA output that can be controlled by the FPGA. We used this to demonstrate the capability of synchronizing the AMA with a DMD. A VGA signal to the DMD is synchronized with 20 nanosecond accuracy to the AMA mirror control. This level of synchronization provides the opportunity for sub-video-frame synchronization between the AMA and the DMD. This would allow for moving the AMA mirrors to multiple positions within one frame of video, as detailed in Section 4.7.5.

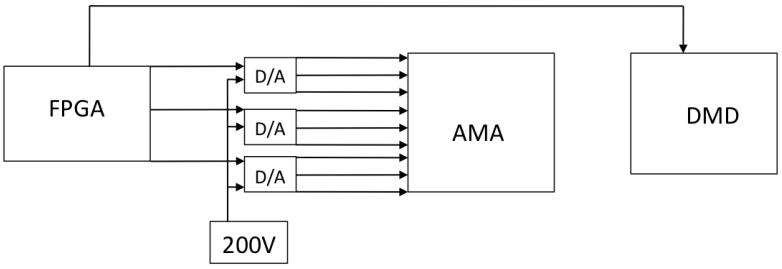

Figure 5.3: Schematic of the control signal flow in the prototype.

To examine how the frame of video could be divided into sub-frames suitable for the AMA, we measured the light output from the PK20 projector before its disassembly and inclusion in the AMA projector prototype. A photodiode connected to an oscilloscope measured the change in light intensity over time. Figure 5.4 shows two signals obtained from the projector showing its behaviour to two different input signals. We found that the each of the colours was displayed 4 times per frame. The frame was 1/60s long. Each colour had a different period within the frame: 1.78ms for green, 1.06ms for blue, and 1.38ms for red. The difference in periods for colours are possibly due to the relative strength of the red, green and blue LEDs, and that the human visual system is most sensitive to green in terms of brightness, and least to blue.

Figure 5.4: Light intensity from a Mitsubishi PK20 projector over time. In the upper signal, the horizontal axis is time, and the vertical axis is intensity. The lower signal shows the VGA sync signal. In a), the projector was given an image with full red and blue, while b) shows a signal with all three colours in varying proportions.

The FGPA writes to each DAC at every subfield change so that a different state of the AMA mirrors could be specified for each colour field. Due to data bandwidth issues, only static images are implemented for sub-frame changes at this time. Because the colour is not synchronized between the two projectors in the prototype, the image display functionality of the FPGA code is not used in

the current version of the prototype.

5.2 Prototype results

Due to the problems with the micromirror fabrication, no fully-functional AMA chips were available for integration into the prototype. The goals for this version of the prototype were therefore modest. We wanted to show that we could make regions of the image brighter or darker, while still keeping the DMD image undistorted. This would validate the main claims of the work. Net gains in brightness compared to an unmodified projector were not pursued because we were using a different light source, and we significantly change the light path in order to give ourselves enough freedom to experiment, which resulted in some light losses. At this point our goal was not to compete against professionally-optimized projectors, but simply to show the potential of this method.

An image was displayed with the AMA projector by directing a VGA signal from a computer to both to the light source (Knoll projector) and the DMD (PK20 projector). The AMA was controlled by a set of predetermined D/A voltages which can be cycled through using buttons on the Altera DE2. These voltages actuate the AMA mirrors, which in turn change the final intensity distribution of the projected image. The objective was to measure the relative luminance changes in the image as the AMA mirrors were actuated, in order to validate the concept that the AMA can affect projector contrast and peak brightness.

A Nikon D70s SLR camera was used to capture images of the projected light as the settings for disparity and AMA voltages were manipulated. The light reflecting from the AMA to the DMD and finally out of the projection lens was the sole source of illumination for the images. The Nikon camera has its own gamma response, which make the pixel values reported by the camera non-linear measurements of luminance, even when a RAW format is used. Without knowing the actual camera response, no accurate measure of relative intensity changes can be made. To obtain approximately linear measurements of luminance, we constructed HDR images from multiple images taken with different exposures using the software Photosphere by Greg Ward (Ward 2009). Photosphere uses an algorithm similar to that described in Mitsunga and Nayar (Mitsunaga and Nayar 1999) to estimate the radiometric response function of a camera from images taken at different exposures. Once an estimate of the response function is made, the estimated (linear) radiometric distribution in the scene can be calculated. Inanici and Glavin (Inanici and Galvin 2004) have shown that this technique is a useful tool for capturing luminance values within 10% of actual over a wide range of luminances, which is adequate for this initial study. The Nikon camera has an exposure bracketing function that takes multiple pictures of the same scene at differing exposure settings. The resulting images were input into Photosphere to construct the HDR images.

Figure 5.5 illustrates the effect of changing the disparity on the level of blur. As the blur increases, details of the AMA disappear, and the resulting ML

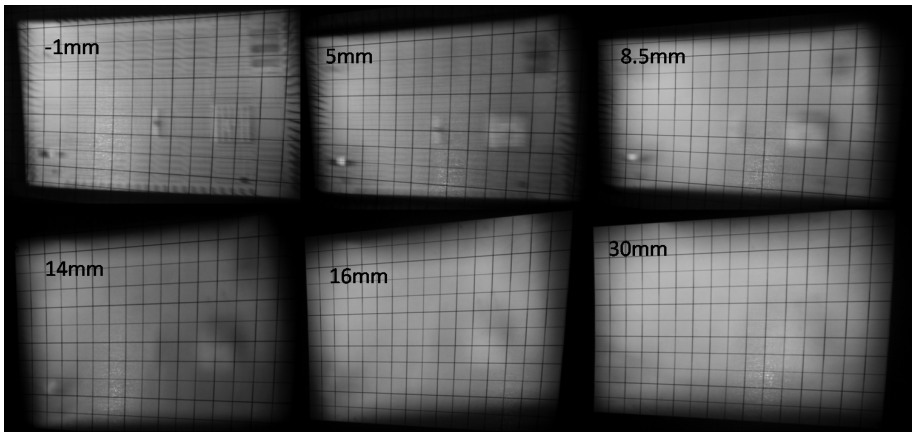

Figure 5.5: Images showing the output of the AMA prototype with different separation settings. The DMD is showing a black grid of 50 pixels squares on a white background. One AMA mirror in the centre right has been actuated.

becomes more and more disperse even as the displacement of the ML increases. The ML can not be tracked between different disparities by the grid lines alone, because although care was taken to move the DMD only along the optical axis, the alignment of the DMD in the projector made this difficult to do in practice over a wide range of disparities. For this reason, we can only make qualitative measurements between disparities.

The grid lines are from the DMD, and are in 50 pixel increments. We can thus measure the size of the image AMA in units of DMD pixels by counting the grid lines. Physically, the AMA array is 9 mm ×5 mm, while the DMD is 11.17 mm ×8.38 mm. We can see from the image in Figure 5.5 that the AMA at -1mm disparity is approximately 650×450 pixels, which gives a magnification level of 1 in the horizontal axis and 1.25 in the vertical axis. The image of the AMA gets larger as the disparity increases because of the increasing size of the blur kernel.

In Figures 5.6, 5.7, 5.8, and 5.9, we examine the region affected by the one AMA mirror tilted in Figure 5.5. To more clearly show the areas of relative change that we are concerned about, the intensity of these images is in units of percentage difference, calculated as

$$P_d = 100\frac{I_t - I_n}{I_n}, \tag{5.1}$$

where I_t is the image with actuated AMA mirrors, and I_n is the image of non-actuated mirrors, and P_d the resulting percentage difference. The calculation is done on a per-pixel basis.

The mirror shown was actuated with 200V. From Figure 3.21, we see the tilt angle is between 0.4° and 2.2° at this voltage, depending on which mirror

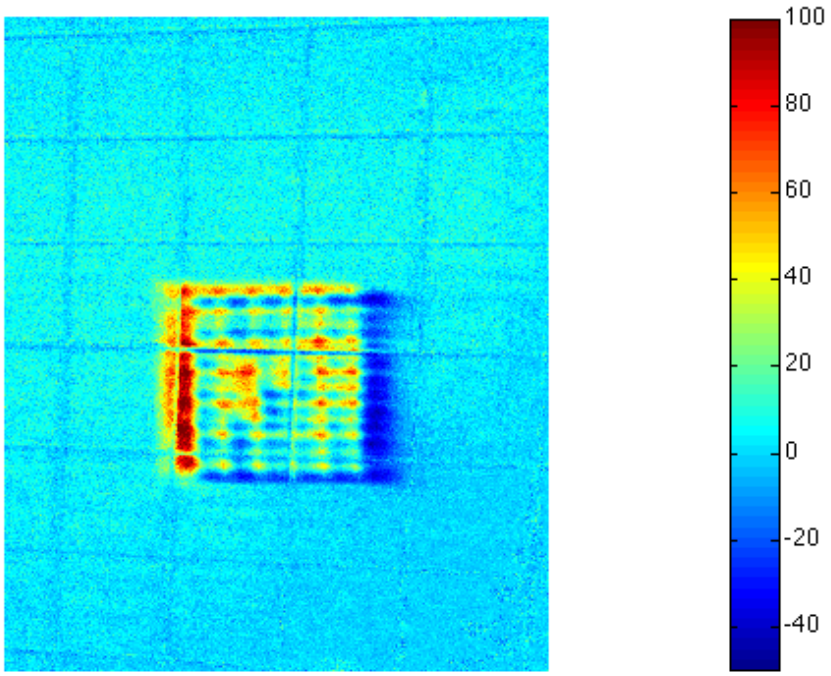

Figure 5.6: One ML at approximately -1mm disparity (the DMD is is on the relay-lens side of the focal plane). Units of percent difference.

in the composite arrangement is measured, due to their differing orientations when actuated as described in Section 3.6.2. For this reason we anticipate the shape of the ML will be spatially shortened in one dimension compared to the shape of the dark region that is left when the ML moves.

Figure 5.6 shows the ML nearly in focus, with a slight bias in light to the top left. At such low disparities, the ML hardly moves at all. The mirror is tilting the light towards the top left, showing that the separation is actually slightly negative. If the DMD had been exactly on the focal plane, the ML would not have been displaced at all, except from the bending artifacts that occur when these mirrors are actuated. The size of the ML was measured by measuring the grid in terms of image pixels, measuring the size of the ML in image pixels, and then solving for the size of the ML in terms of DMD pixels, with the knowledge that the grid is in 50 DMD pixel increments. With this process, we measured the size of the ML in this case as approximately 75 pixels per side.

Figure 5.7 shows how the ML has become blurred as the disparity increases to 8.5mm. The size of the ML has grown to 100 pixels per side, and has started to change shape due to mirror bending. The ML has been displaced 60 pixels in

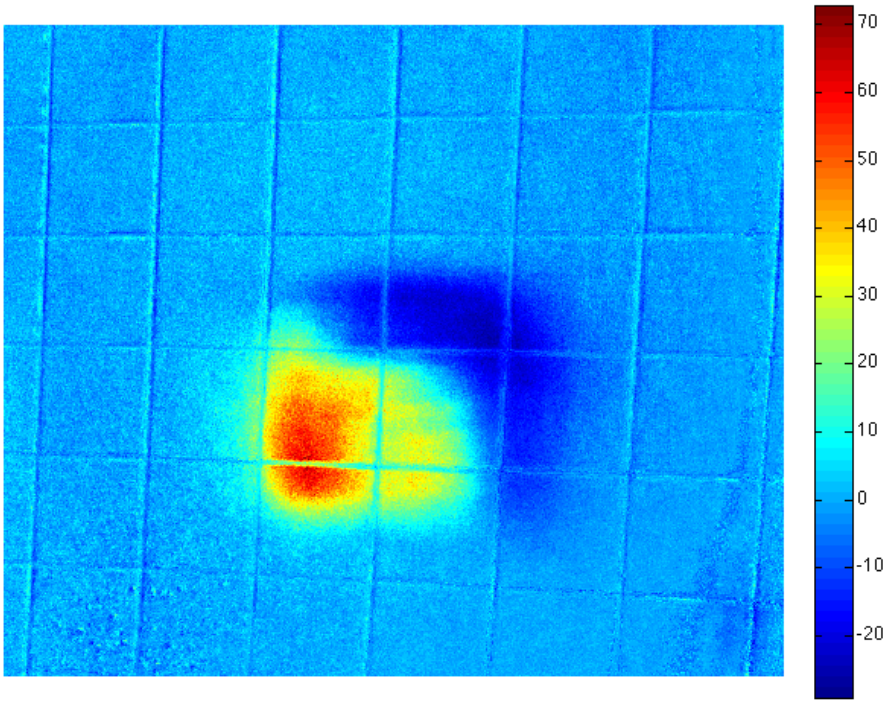

Figure 5.7: The same region with the mirror at the same tilt angle, at 8.5mm disparity. Units of percent difference.

a diagonal direction towards the bottom left. The actual ML is less wide than the region it is leaving because of the bending artifacts in the mirror. The angle of tilt is not uniform over the composite mirror, causing the ML to distort as it is tilted.

At 14mm disparity, the size of the ML is now 120 pixels on its longest side, as shown in Figure 5.8. The ML has been displaced approximately 90 pixels. Figure 5.9 shows the ML at 30mm of disparity. Although hardly visible in Figure 5.5, showing a closeup of the region in false colour with units of relative difference from the non-tilted case shows that the ML has shifted 160 pixels. The ML is further elongated, and now is spread over an area of 160 pixels per side.

Also evident from Figures 5.6 5.7, 5.8, and 5.9 is that the peak intensity change attenuates with the blur. At very low levels of blur such as 8.5mm, there is a 75% improvement in peak brightness, and a 25% decrease in peak darkness. At 30mm of disparity, the peak has been reduced to a 40% increase, and 25% decrease. The non-symmetrical nature of these increases are likely

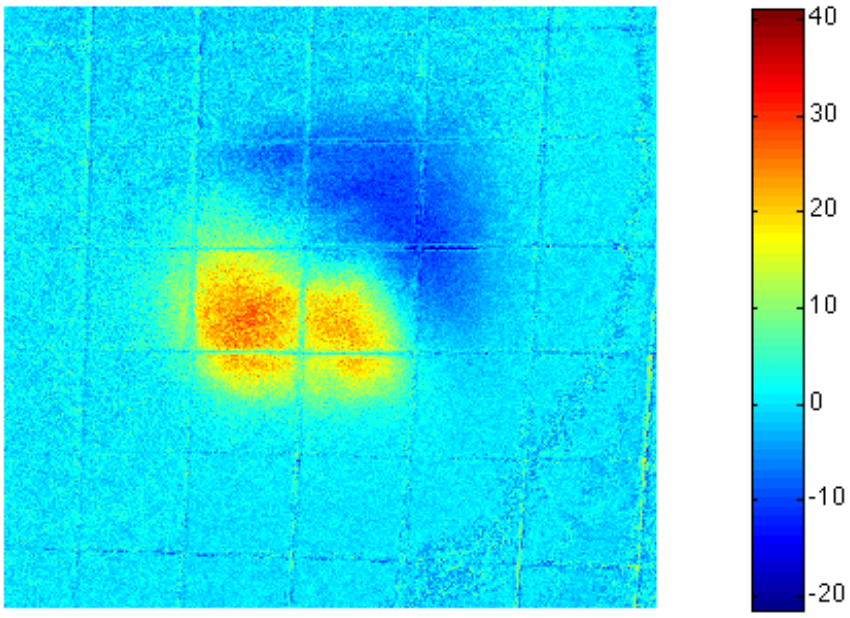

Figure 5.8: 14mm disparity, relative change. Units of percent difference.

caused by the lens-effect of the bending AMA mirrors, as the ML becomes clearly smaller than the space it leaves as the disparity increases.

Using Equation 2.12, we can use the measurements above to compare the estimated tilt angle of the micromirrors to the measurements shown in Figure 3.21. A DMD pixel is $14\mu m$ square, and we have shown above that the magnification is approximately 1. Using the pixel measurements above, the 8.5mm disparity works out to 2.7°, 14mm disparity to 2.6°, and the 30mm disparity to 2.1°. These discrepancies can be attributed to measure error of the ML range, the slight magnification of the vertical dimension of the AMA, the bending of the AMA mirrors themselves, and the differences in spring values between AMA mirrors in the array. However, all are within 0.5° of the value measured with the white-light interferometer for 200V actuation.

Figure 5.10 shows multiple mirrors actuated, creating a non-homgenous light field across most of the image. As the ML range increases with the blur disparity, some of the MLs blend into each other to form larger regions both brighter and darker than the original luminance.

Overall, these measurements show that the approach of using an AMA is successful at redistributing light from one region of the projected image to another. We also demonstrated that the AMA does not geometrically distort the image from the DMD in any way. The effect that disparity has on both range

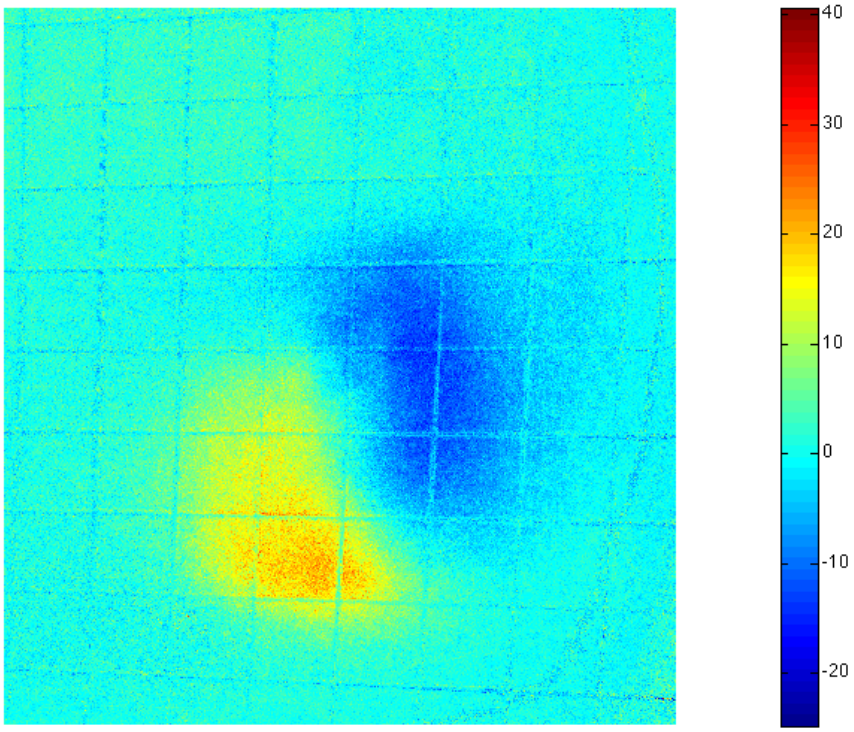

Figure 5.9: 30mm disparity, relative change. Units of percent difference.

and blur was shown, and validates the approach to simulations taken in Chapter 4. While difficulties with micromirror fabrication severely impact the attainable tilt angle, it was adequate to show that the AMA can make areas of the image both brighter and darker as the mirrors are actuated.

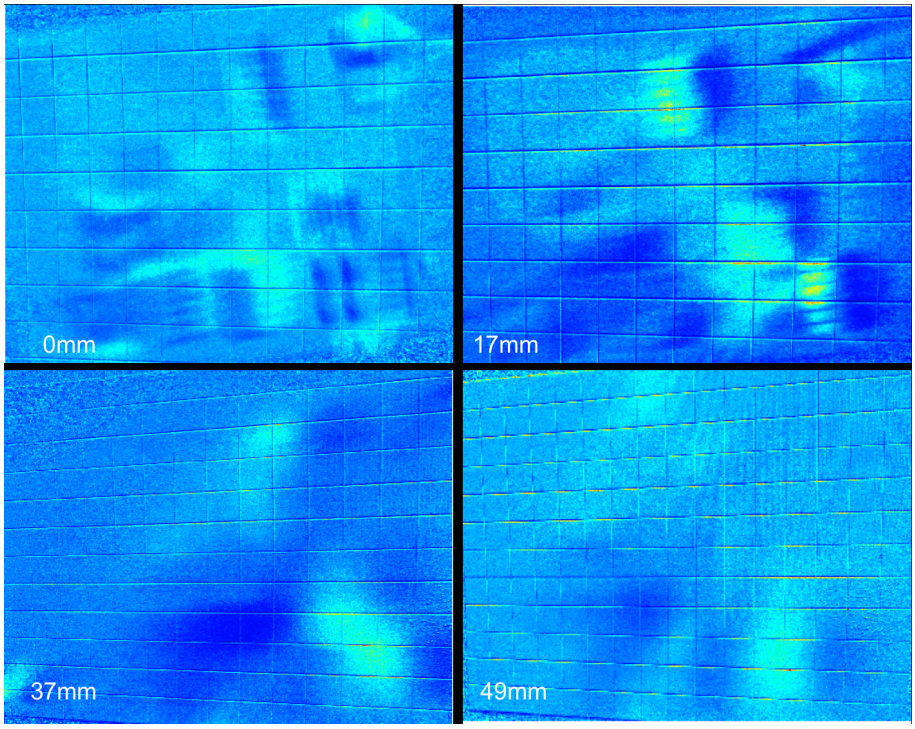

Figure 5.10: Relative change for four different disparity settings, showing multiple mirrors actuated.

Chapter 6

Conclusions and future work

6.1 Summary and conclusions

This thesis has presented a method to significantly improve the brightness and contrast capability of a projector through the addition of an analog micromirror array. By channeling light to where it is needed and away from where it is not, an AMA projector makes better use of its light source, which is one of the most expensive components of today's projectors. After the addition of an AMA, the light reaching the primary image modulator, such as a DMD, is not considered to be uniform. Instead, the distribution from the AMA mirrors is simulated, and the compensation for the non-homogeneity is applied to the original image before it is sent to the DMD. The result is an image of higher contrast and peak brightness than would otherwise be possible with the same projection lamp.

Dispensing with the assumption of homogeneity creates other opportunities for increases in efficiency. To obtain a uniform distribution, the light from the lamp extends past the borders of the DMD so that apodization is minimized. This overfill light is normally wasted. In an AMA projector, however, the overfill can be used to further illuminate the image because homogeneity is no longer assumed.

The optical issues that arise when adapting a projector with an AMA were analyzed. Because of the geometrical nature of the light from a non-laser projector lamp, light from each mirror of the AMA is blurred before reaching the DMD. A methodology for simulating the effects of blur, and its relation to mirror range, and number of mirrors was implemented. It was found that a larger range of movement for the mobile light from an AMA mirror (ML) implies a larger blur. This can be mitigated up to a point by increasing the maximum mechanical tilt angle of the mirrors.

Two-degree of freedom analog micromirror arrays were designed, fabricated and analyzed for this purpose using multi-user MEMS fabrication processes. A novel way of optimizing the tradeoffs between tilt angle, mirror size, and mirror resonance frequency by splitting the mirrors into smaller functional subsections that move synchronously was employed.

We developed several algorithms that determine favourable placement of the mobile lights from each of the micromirrors in the array, in order to best improve the image. Results from these implementations show that with a mirror array

of 28 mirrors, the average brightness could be increased by a factor of 1.2 and the contrast by 1.4 if micromirrors were available that could be tilted to $\pm 3.5^\circ$ with the addition of this technology, without changing the projector lamp. The improvement factor rises to 2.25 and contrast factor by 5 if a minimal number of pixels are allowed to be under-illuminated. Additional micromirrors in the array would also increase this improvement.

Finally, we constructed a physical prototype adapted from a commodity projector. Mirror control electronics were designed to provide a test-bed for the AMA projector prototype. We verified with the prototype that the AMA can increase the peak brightness and contrast of the projector.

6.2 Summary of Contributions

In summary, this thesis has examined a novel method to dynamically reallocate the light from a projector lamp from dark regions to bright regions on an image-dependent basis in order to increase projector peak brightness, contrast, and efficiency. This is also described in our paper (Hoskinson and Stoeber 2008). This approach required novel work in a number of areas:

- A theoretical framework for examining the tradeoffs between optical parameters that affect AMA system performance.
- The design and fabrication of an analog micromirror array suitable for an AMA projector, with composite mirrors that optimize the tradeoffs between mirror tilt angle, size, and dynamic behaviour (Hoskinson et al. 2007a; Hoskinson et al. 2007b).
- Algorithms and software implementations that allocate the mobile lights depending on the image, taking into account the physical limitations of the mirrors and projector.
- A prototype implementation demonstrating this method, showing that regions can become brighter as well as darker.

6.3 Future work

The addition of the AMA into a projector involves several added costs. Further research and development will have to be done to minimize them. For instance, the AMA could be made with a process that facilitates the incorporation of high-voltage CMOS that can be connected via through-holes to the mirror electrodes. This would remove the need for separate D/A converters and wire-bonding, and allow for more efficient use of electrode space, since routing leads would no longer be needed. This would also allow for many more individually-controlled mirrors. This would involve a substantial research effort to design a new fabrication process that best combine the high-voltages CMOS with the MEMS materials such as single-crystal silicon.

Shaping the DMD aperture to let in as much on-state light given the ability of the AMA to divert unwanted light could mitigate losses that occur in conventional non-AMA projectors from contrast control. Normally, some of the usable illumination is clipped before reaching the DMD to prevent stray unwanted light from affecting contrast, but with the added ability to control where light reaches the DMD, and at what angle, more initial stray light can be tolerated without affecting overall contrast. Shaping the DMD aperture would also help minimize AMA clipping loss.

This work analyzed only static images. The sample images tested showed that the improvement factor can vary substantially from image to image. A system must be implemented to ensure that this does not cause perceptible artifacts as the global brightness changes over time for video. The adaptive iris (Toyooka et al. 2005) included in some projectors faces the same problem, and has been mitigated to some degree.

Since the distribution of light on the DMD is already non-homogeneous, the integrating rod normally present in a projector could be omitted, further reducing the cost of the device. The AMA could be thought of as an "active integrator"; in an integrator, light that is originally near the centre of the distribution is moved to the exteriors. The AMA instead moves this light to wherever it is needed in the image, not necessarily the centre. The non-homogeneous nature of the incident light could be taken into account during the normal image-processing step of the projector. This would also increase system efficiency versus an equivalent non-AMA projector. In this case, because the light incident on the AMA would be non-homogeneous, a separate model of every ML would be needed.

The large gap in improvement factor between perfect reconstruction and that with even only 200 under-illuminated pixels (1.3 vs. 2.5) shows that there is work to do in optimizing the allocation algorithms to avoid outliers. Further research could also be done to determine whether outliers could be less perceptible than a similar error on a regular projector. It may be that the greater dynamic range of an AMA-projected image inherently masks errors to some degree.

The visual difference prediction employed in Section 4.9 was only for validation purposes, but it may be possible to employ a small subsection as part of the allocation algorithm in a predictive capacity. If the pixels could be weighted by a 'perceptibility' metric, then the AMA projector could make intelligent decisions as to the potential impact of under-illuminated regions during the allocation algorithm. Light could even be diverted away from regions of low perceptual significance towards regions of high perceptual significance for a greater overall improvement than if every image pixel had an identical weighting.

Another avenue for future work is with the non-homogeneous correction mechanism needed to produce accurate images in an AMA projector. It could be extended to cover ambient light conditions with the addition of a camera. Various methods for compensating for room illumination have been suggested in the literature, such as (Nayar et al. 2003; Wang et al. 2005; Ashdown et al. 2006; Park et al. 2008). All involve capturing the state of the surface projected upon using a camera, and compensating for any non-uniformities by changing the

image projected. Each approach comes at the cost of dynamic range, since some of the image bits normally used for greyscale values are now used for radiometric compensation. The output of a projector can saturate if the compensation exceeds the projector's dynamic range, resulting in perceptible artifacts in the compensated image.

A similar technique for radiometric compensation could be used with a camera-enabled AMA projector that could alleviate dynamic range limitations. Images from a camera-equipped AMA projector could be fed into the image processing algorithm, and non-homogeneous screen brightnesses accounted for by the image processing algorithm when determining the control sequences for the AMA and DMD. The added dynamic range introduced with the AMA would alleviate the problems of previous approaches.

This work has demonstrated that an AMA-enhanced projector can make projectors intelligent allocators of their light sources. This ability opens up several interesting research directions, as the projector could conceivably take into account the ambient lighting in the room, the psychophysical limitations of the viewers, and the content being presented to display the best possible image over a wide range of conditions.

Bibliography

Agarwal, P., and Varadarajan, K. 2004. A near-linear constant-factor approximation for euclidean bipartite matching? In *Proceedings of the twentieth annual symposium on Computational geometry*, ACM New York, NY, USA, 247–252.

Aksyuk, V., Basavanhally, N. R., Bishop, D. J., Bolle, C. A., Chang, C. C., Frahm, R., Gasparyan, A., Gates, J. V., George, R., Giles, C. R., Kim, J., Kolodner, P. R., Lee, T. M., Neilson, D. T., Nijander, C., Nuzman, C. J., Paczowski, M., Papzian, A. R., Pardo, F., Ramsey, D. A., Ryf, R., Scotti, R. E., Shea, H., and Simon, M. E. 2003. 238x238 micromechanical optical cross connect. *IEEE Photonics Technology Letters 15*, 4, 587–589.

ANSYS, 2009. ANSYS. http://www.ansys.com. Accessed 03/24/09.

Ashdown, M., Okabe, T., Sato, I., and Sato, Y. 2006. Robust Content-Dependent Photometric Projector Compensation. *Proceedings of the Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'06), Projector-Camera Systems (PROCAMS)*, 6.

Barten, P. G. J. 1992. Physical model for the contrast sensitivity of the human eye. *Proceedings of SPIE 1666*, 57–72.

Barten, P. 1996. Effect of Gamma on Subjective Image Quality. *SID Intl. Symp. Digest of Technical Papers*.

Bishop, D. J., Giles, C. R., and Austin, G. P. 2002. The Lucent LambdaRouter: MEMS technology of the future here today. *IEEE Communications Magazine 40*, 3, 75–79.

Bloom, D., 1997. The grating light valve: revolutionizing display technology.

Borovic, B., Lewis, F., Liu, A., Kolesar, E., and Popa, D. 2006. The lateral instability problem in electrostatic comb drive actuators: modeling and feedback control. *Journal of Micromechanics and Microengineering 16*, 7, 1233–1241.

Brennesholtz, M. S., and Stupp, E. H. 2008. *Projection Displays*, 2nd ed. Wiley.

Brennesholtz, M. S. 1996. Light collection efficiency for light valve projection systems. In *Projection Displays II Proceedings of SPIE*, vol. 2650, 71–79.

Brennesholtz, M. 2007. Invited Paper: The evolution of projection displays: From mechanical scanners to microdisplays. *Journal of the Society for Information Display 15*, 759.

Brennesholtz, M. S., 2009. 3D: Brighter is Better, http://displaydaily.com/2009/05/27/3d-brighter-is-better/ Last accessed 05/06/2009.

Burt, P., and Adelson, E. 1983. A multiresolution spline with application to image mosaics. *ACM Transactions on Graphics 2*, 4, 217–236.

CMC. 2005. Introduction to Micragem: A silicon-on-insulator based micromachining process. Tech. rep., Micralyne Inc. Version 4.0.

CMC, 2009. Canadian microelectronics corporation, www.cmc.ca.

Comtois, J. H., Bright, V. M., Gustafson, S. C., and Michalicek, M. A. 1995. Implementation of hexagonal micromirror arrays as phase-mostly spatial light modulators. *Proceedings of SPIE 2641*, 76–87.

Comtois, J., Michalicek, A., Cowan, W., and Butler, J. 1999. Surface-micromachined polysilicon moems for adaptive optics. *Sensors and Actuators 78*, 54–62.

Conant, R. 2002. *Micromachined Mirrors*. PhD thesis, University of California, Berkeley.

Conner, A. 2006. LEDs for Microdisplay Illumination. *SID International symposium digest of technical papers 37*, 2, 1819.

Cowan, W. D., Bright, V. M., Lee, M. K., Comtois, J. H., and Michalicek, M. A. 1998. Design and testing of polysilicon surface-micromachined piston micromirror arrays. *Proceedings of SPIE 3292*, 60–70.

Dagel, D., Cowan, W., Spahn, O., Grossetete, G., Grine, A., Shaw, M., Resnick, P., and Jokiel Jr, B. 2006. Large-stroke MEMS deformable mirrors for adaptive optics. *Journal of Microelectromechanical Systems 15*, 3, 572–583.

Damberg, G., Seetzen, H., Ward, G., Heidrich, W., and Whitehead, L. 2007. High Dynamic Range Projection Systems. *Proceedings of the 2006 Society for Information Display Annual Symposium*.

Debevec, P. 2008. A median cut algorithm for light probe sampling. In *International Conference on Computer Graphics and Interactive Techniques*, ACM New York, NY, USA.

Derra, G., Moench, H., Fischer, E., Giese, H., Hechtfischer, U., Heusler, G., Koerber, A., Niemann, U., Noertemann, F., Pekarski, P., et al. 2005. UHP lamp systems for projection applications. *JOURNAL OF PHYSICS-LONDON-D APPLIED PHYSICS 38*, 17, 2995.

Dewald, D., Segler, D., and Penn, S. 2004. Advances in contrast enhancement for DLP projection displays. *Journal of the Society for Information Display 11*, 177–181.

Dokmeci, M., Pareek, A., Bakshi, S., Waelti, M., Fung, C., Heng, K. H., and Mastrangelo, C. 2004. Two-axis single-crystal silicon micromirror arrays. *Journal of Microelectromechanical Systems 13*, 6, 1006–1017.

Dutta, S. B., Ewin, A. J., Jhabvala, M. D., Kotecki, C. A., Kuhn, J. L., and Mott, D. B. 2000. Development of individually addressable micromirror arrays for space applications. In *Proceedings of SPIE – Volume 4178*, 365–371.

Fairchild, M. 1995. Considering the surround in device-independent color imaging. *Color Res. Appl 20*, 6, 352–363.

Hah, D., Huang, S. T.-Y., Tsai, J.-C., Toshiyoshi, H., and Wu, M. C. 2004. Low-voltage, large-scan angle MEMS analog micromirror arrays with hidden vertical comb-drive actuators. *Journal of Microelectromechanical Systems 13*, 2, 279–289.

Hah, D., Patterson, P. R., Nguyen, H. D., Toshiyoshi, H., and Wu, M. C. 2004. Theory and experiments of angular vertical comb-drive actuators for scanning micromirrors. *IEEE Jorunal of Selected Topics in Quantum Electronics 10*, 3, 505–513.

Halbach, K. 1964. Matrix representation of Gaussian optics. *American Journal of Physics 32*, 90.

Hecht, E. 2002. *Optics*, 4th ed. Addison Wesley, San Francisco, CA, USA.

Heckaman, R. 2006. Effect of DLP projector white channel on perceptual gamut. *Journal of the SID*, 755.

Heckbert, P. 1982. Color image quantization for frame buffer display. *ACM SIGGRAPH Computer Graphics 16*, 3, 297–307.

Hibbeler, R. 2008. *Mechanics of Materials*, 7th ed. Pearson Prentice Hall, New Jersey.

Hornbeck, L. J. 1983. 128 x 128 deformable mirror device. *IEEE Transactions on Electron Devices ED-30*, 539–545.

Hornbeck, L. J., 1996. Active yoke hidden hinge digital micromirror device.

Hornbeck, L. 1997. Digital Light Processing TM for High-Brightness, High-Resolution Applications. *Electronic Imaging EI '97 Projection Displays III, IS&T and SPIE, San Jose CA 3013*, 27–40.

Hoskinson, R., and Stoeber, B. 2008. High-dynamic range image projection using an auxiliary mems mirror array. *Optics Express 16*, 7361–7368.

Hoskinson, R., Mansoor, I., and Stoeber, B. 2007. High mechanical bandwidth polysilicon micromirrors for large area / high deflection angle micromirror arrays. In *Canadian Conference on Electrical and Computer Engineering*.

Hoskinson, R., Mansoor, I., and Stoeber, B. 2007. Submirror pixels for large scan angle micromirror arrays. In *Canadian Workshop on MEMS and Microfluidics*.

Hwang, K.-H., Song, Y.-J., and Kim, S.-G. 1998. Thin-film micromirror array for high brightness projection displays. *Japanese Journal of Applied Physics 37*, 12B, 7074–7077.

Iisaka, H., Toyooka, T., Yoshida, S., and Nagata, M. 2003. Novel Projection System Based on an Adaptive Dynamic Range Control Concept. In *Procoeedings Int Disp Workshops*, vol. 10, 1553–1556.

Inanici, M., and Galvin, J. 2004. Evaluation of High Dynamic Range Photography as a Luminance Mapping Technique. Tech. rep., Lawrence Berkeley National Laboratory, Lighting Research Group, LBNL-Report.

Jacobson, B., Gengelbach, R., Stewart, C., and Rutan, D. 1998. Metal Halide Lighting Systems and Optics for High Efficiency Compact LCD Projectors. In *Proc. SPIE*, vol. 3296, 38–45.

Jacobson, R., Ray, S., Attridge, G. G., and Axford, N. 2000. *The Manual of photography: Photographic and Digital Imaging*, 9th ed. Focal Press.

Janssen, P., and Shimizu, J., 1995. High contrast illumination system for video projector. Patent, Aug. 15. US Patent 5,442,414.

Judy, J., and Muller, R. 1997. Magnetically actuated, addressable microstructures. *Microelectromechanical Systems, Journal of 6*, 3, 249–256.

Jung, I., Krishnamoorthy, U., and Solgaard, O. 2006. High fill-factor two-axis gimbaled tip-tilt-piston micromirror array actuated by self-Aligned vertical electrostatic combdrives. *Microelectromechanical Systems, Journal of 15*, 3, 563–571.

Kanayama, H., Maeda, M., Miwa, T., Ikeda, T., Murata, H., and Chihara, K. 2006. Ultra Small Projector with High Efficiency Illumination System. In *International Conference on Consumer Electronics (ICCE'06)*.

Kessel, P. V., Hornbeck, L., Meier, R., and Douglass, M. 1998. A MEMS-based projection display. In *Proceedings of the IEEE*, vol. 86, 1687–1704.

Kiang, M., Solgaard, O., Lau, K., and Muller, R. 1998. Electrostatic combdrive-actuated micromirrors for laser-beamscanning and positioning. *Microelectromechanical Systems, Journal of 7*, 1, 27–37.

Kim, S., Barbastathis, G., and Tuller, H. 2004. MEMS for optical functionality: Special issue on electroceramics in micro-electro-mechanical systems. *Journal of Electroceramics 12*, 1-2, 133–144.

Koester, D., Cowen, A., Mahadevan, R., Stonefield, M., and Hardy, B. 2003. PolyMUMPs design handbook. Tech. rep., MEMSCAP. Revision 10.0.

Krishnamoorthy, U., Lee, D., and Solgaard, O. 2003. Self-aligned vertical electrostatic combdrives for micromirror actuation. *Microelectromechanical Systems, Journal of 12*, 4, 458–464.

Kuhn, H. 1955. The Hungarian method for the assignment and transportation problems. *Naval Research Logistics Quarterly 2*, 83–97.

Kusakabe, Y., Kanazawa, M., Nojiri, Y., Furuya, M., and Yoshimura, M. 2009. A high-dynamic-range and high-resolution projector with dual modulation. In *Proceedings of SPIE*, vol. 7241.

Legtenberg, R., Groeneveld, A., and Elwenspoek, M. 1996. Combdrive actuators for large displacements. *J. Micromech. Microeng 6*, 3, 320–329.

Lin, J. E., Michael, F. S. J., and Kirk, A. G. 2001. Investigation of improved designs for rotational micromirrors using multi-user MEMS processes. *MEMS Design, Fabrication, Characterization, and Packaging SPIE 4407*, 202–213.

Majumder, A., Bhasker, E., and Juang, R. 2008. Advances towards high-resolution pack-and-go displays: A survey. *Journal of the Society for Information Display 16*, 481.

Mang, O., Chou, C., and Yuan, O. 2008. Optimized design and analysis of dual-lamp illumination modules compatible with existing projector systems (Journal Paper). *Optical Engineering 47*, 07.

Mantiuk, R., Daly, S., Myszkowski, K., and Seidel, H.-P. 2005. Predicting visible differences in high dynamic range images - model and its calibration. In *Human Vision and Electronic Imaging X, IS&T/SPIE's 17th Annual Symposium on Electronic Imaging (2005)*, B. E. Rogowitz, T. N. Pappas, and S. J. Daly, Eds., vol. 5666, 204–214.

Milanovic, V., Last, M., and Pister, K. 2001. Torsional micromirrors with lateral actuators. *Transducers 2001*.

Mitsubishi. Laservue.

Mitsunaga, T., and Nayar, S. 1999. Radiometric self calibration. In *Computer Vision and Pattern Recognition, 1999. IEEE Computer Society Conference on.*, vol. 1.

Motamedi, M., Ed. 2005. *MOEMS: micro-opto-electro-mechanical systems*. SPIE Press.

Nagahara, H., Kuthirummal, S., Zhou, C., and Nayar, S. 2008. Flexible Depth of Field Photography. In *European Conference on Computer Vision (ECCV)*.

Nayar, S., Peri, H., Grossberg, M., and Belhumeur, P. 2003. A Projection System with Radiometric Compensation for Screen Imperfections. *First IEEE International Workshop on Projector-Camera Systems (PROCAMS-2003)*.

Niven, G., and Mooradian, A. 2006. Trends in laser light sources for projection display. In *13th International Display Workshops (IDW '06)*.

Pan, J., Wang, C., Lan, H., Sun, W., and Chang, J. 2008. High efficiency pocket size projector with a compact projection lens and LED based light source system. *Applied optics 47*, 19, 3406–3414.

Park, H., Lee, M., Seo, B., Park, J., Jeong, M., Park, T., Lee, Y., and Kim, S. 2008. Simultaneous Geometric and Radiometric Adaptation to Dynamic Surfaces With a Mobile Projector-Camera System. *Circuits and Systems for Video Technology, IEEE Transactions on 18*, 1, 110–115.

Pavlovych, A., and Stuerzlinger, W. 2005. A High-Dynamic Range Projection System. *Progress in biomedical optics and imaging 6*, 39.

Peterson, K. E. 1982. Silicon as a mechanical material. *Proceedings of the IEEE 70*, 5, 420–457.

Pinho, G. P. 2003. Optics of digital cinema. *Projection Displays IX (Proceedings of SPIE/IS&T Volume 5002)*, 123–131.

Poynton, C. 2003. *Digital Video and HDTV: Algorithms and Interfaces*. Morgan Kaufmann.

Ramanarayanan, G., Ferwerda, J., Walter, B., and Bala, K. 2007. Visual equivalence: towards a new standard for image fidelity. *ACM Transactions on Graphics (TOG) 26*, 3.

Rao, S. S. 2003. *Mechanical Vibrations*, 4th ed. Prentice Hall.

Reinhard, E., Ward, G., Pattanaik, S., and Debevec, P. 2005. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting (The Morgan Kaufmann Series in Computer Graphics)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

Rempel, A., Trentacoste, M., Seetzen, H., Young, H., Heidrich, W., Whitehead, L., and Ward, G. 2007. Ldr2Hdr: on-the-fly reverse tone mapping of legacy video and photographs. *ACM Transactions on Graphics (TOG) 26*, 3.

Roufs, J., Koselka, V., and van Tongeren, A. 1994. Global brightness contrast and the effect on perceptual image quality. In *Proceedings of SPIE*, vol. 2179.

Sandia National Laboratories, 2007. Sandia Ultra-planar, Multi-level MEMS Technology 5 (SUMMiT V), http://www.mems.sandia.gov/tech-info/summit-v.html.

Seetzen, H., Whitehead, L., and Ward, G. 2003. A high dynamic range display using low and high resolution modulators. *Society for Information Display (SID) Digest of Technical Papers*, 1450–1453.

Seetzen, H., Heidrich, W., Stuerzlinger, W., Ward, G., Whitehead, L., Trentacoste, M., Ghosh, A., and Vorozcovs, A. 2004. High dynamic range display systems. *ACM Transactions on Graphics (Siggraph 2004) 23*, 3, 760–768.

Senturia, S. D. 2001. *Microsystem design*. Kluwer Academic Publishers, Norwell, MA, USA.

Simoncelli, E. P., 2007. matlabPyrTools, http://www.cns.nyu.edu/ lcv/software.html.

SMPTE. 2003. *196M Indoor Theater and Review Room Projection – Screen Luminance and Viewing Conditions*.

Stokes, M., Anderson, M., Chandrasekar, S., and Motta, R. 1996. A Standard Default Color Space for the Internet–sRGB. *Microsoft and Hewlett-Packard Joint Report, Version 1*.

Subbarao, M. 1987. Direct recovery of depth-map I: Differential methods. In *Proc. IEEE Computer Society Workshop on Computer Vision, Miami Beach*.

Sutherland, E. ., 2009. Untitled, http://www.es.com/ Accessed 18/04/2009.

Tang, W., Nguyen, T., and Howe, R. 1989. Laterally driven polysilicon resonant microstructures. *Micro Electro Mechanical Systems, 1989, Proceedings,'An Investigation of Micro Structures, Sensors, Actuators, Machines and Robots'. IEEE*, 53–59.

Taylor, W. P., Brazzle, J. D., Osenar, A. B., Corcoran, C. J., Jafri, I. H., Keating, D., Kirkos, G., Lockwood, M., Pareek, A., and Bernstein, J. J. 2004. A high fill factor linear mirror array for a wavelength selective switch. *Journal of Micromechanics and Microengineering 14*, 1, 147–152.

Texas Instruments. 2005. *DMD 0.7 XGA 12 DDR DMD Discovery.*

Texas Instruments. 2005. Single-panel DLP projection system optics. Tech. rep. Application Report – Discovery DLPA002.

Toyooka, T., Yoshida, S., and Iisaka, H. 2005. Illumination control system for adaptive dynamic range control. *Journal of the Society for Information Display 13*, 105.

Trentacoste, M., Heidrich, W., Whitehead, L., Seetzen, H., and Ward, G. 2007. Photometric image processing for high dynamic range displays. *Journal of Visual Communication and Image Representation 18*, 5, 439–451.

Trentacoste, M. 2006. Photometric image processing for high dynamic range displays (MSc thesis). *The Faculty of Graduate Studies (Computer Science), The University Of British Columbia, Vancouver*.

Tsai, J.-C., Fan, L., Hah, D., , and Wu, M. 2004. A high fill-factor, large scan-angle, two-axis analog micromirror array driven by leverage mechanism. In *International Conference on Optical MEMS and Their Applications*.

Tsai, J.-C., Huang, S., and Wu, M. 2004. High fill-factor two-axis analog micromirror array for $1xN^2$ wavelength-selective switch. In *IEEE International Conference on Micro Electro Mechanical Systems*, 101–104.

Tsai, J., Chiou, S., Hsieh, T., Sun, C., Hah, D., and Wu, M. 2008. Two-axis MEMS scanners with radial vertical combdrive actuators—design, theoretical analysis, and fabrication. *Journal of Optics A: Pure and Applied Optics 10*, 044006.

Tsang, S.-H., and Parameswaran, M. 2005. Self-locking vertical operation single crystal silicon micromirrors using silicon-on-insulator technology. In *Electrical and Computer Engineering, 2005. Canadian Conference on*, 429–432.

Tuantranont, A., Bright, V. M., Zhang, W., and Lee, Y. C. 1999. Packaging of lenslet array on micromirrors. *Proc. SPIE 3631*, 156–164.

Tuantranont, A., Bright, V., Liew, L., Zhang, W., and Lee, Y. 2000. Smart phase-only micromirror array fabricated by standard CMOSprocess. In *Micro Electro Mechanical Systems, 2000. MEMS 2000. The Thirteenth Annual International Conference on*, 455–460.

Van Giel, B., Meuret, Y., Bogaert, L., Murat, H., De Smet, H., and Thienpont, H. 2007. Efficient and Compact Illumination in LED Projection Displays. In *SID International symposium digets of technical papers*, vol. 38, Society for information display.

Wang, K., Böhringer, K. F., and Sinclair, M. 2003. Low voltage and pull-in state adjustable dual servo scanning mirror. In *The 5th Pacific Rim Conference on Lasers and Electro-Optics*.

Wang, D., Sato, I., Okabe, T., and Sato, Y. 2005. Radiometric Compensation in a Projector-Camera System Based on the Properties of Human Vision System. *Proceedings of IEEE International Workshop on Projector-Camera Systems 2005 2*, 8.

Ward, G., 2009. Photosphere, http://www.anyhere.com/.

Weisstein, E. W., 2007. Circle-circle intersection. From MathWorld–A Wolfram Web Resource. http://mathworld.wolfram.com/Circle-CircleIntersection.html.

Wen, J., Hoa, X. D., Kirk, A. G., and Lowther, D. A. 2004. Analysis of the performance of a MEMS micromirror. In *IEEE Transactions on Magnetics*, vol. 40, 1410–1413.

Winston, R., Minano, J., and Benitez, P. 2005. *Nonimaging Optics*. Academic Press.

Yalcinkaya, A., Urey, H., Brown, D., Montague, T., and Sprague, R. 2006. Two-axis electromagnetic microscanner for high resolution displays. *J. Microelectromech. Syst 15*, 786–94.

Young, W. C., and Budynas, R. 2001. *Roark's formula for stress and strain*. McGraw-Hill, September.

Zhang, X., Chau, F., Quan, C., Lam, Y., and Liu, A. 2001. A study of the static characteristics of a torsional micromirror. *Sensors & Actuators: A. Physical 90*, 1-2, 73–81.