

Name: _____
Partner(s): _____
Date(s): _____

Introduction to Digital Optics

Objectives: (**to be read before class**)

- To introduce basic *models* of Digital Optics devices.
- Your discussion will make use of the “[Polarization Ellipse](#)” and [Poincaré Sphere](#).
- You will apply a Jones Matrix description to understand the combined action of a (Polarizer + LCoS SLM + Analyzer).

In-class Background:

One of the Digital Optics devices you will be using is a **Digital Micro-Mirror Device (DMD)**, an optical micro-electrical-mechanical system (MEMS) that contains a large array of aluminum micro-mirrors, which are highly reflective across a very wide range of wavelengths with negligible dispersion or absorption (whereas an LCoS SLM would be irreversibly damaged by exposure to UV, which would crosslink molecules in the display) and which can *rapidly* respond to input voltages (\sim two orders of magnitude faster than an LCoS SLM) by tilting $\pm 12^\circ$ from their unpowered, flat state. These mirrors can be individually controlled to reflect light in either of these two different directions, creating binary (two-state) **amplitude modulation** along these directions. While we say that an LCoS SLM in phase-modulation mode is “re-shaping” the output beam, we say that a DMD “sculpts” the output, via selective removal (a subtractive process akin to applying a chisel to marble).

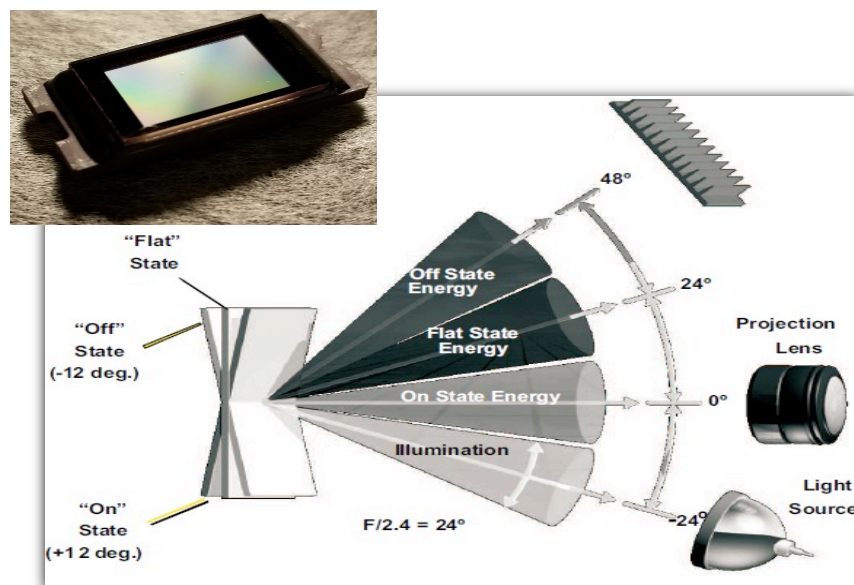


Fig. 4. Digital Micro-Mirror Device (DMD chip) [Texas Instruments Application Report DLPA022 \(July 2010\), 2-3.](#)

Example DMD project:

Anatolii Kashchuk, working in the lab of Prof. Halina Rubinsztein-Dunlop at the University of Queensland, has recently (in 2019) published a method for making “[High-speed transverse and axial optical force measurements using amplitude filter masks](#).” You may, if you are interested, pursue (and possibly extend) this approach, to take the spatially distributed light coming out of an optical trap and, by switching between the three gradient patterns imposed onto the DMD, sort the light into one of two photodiodes to be measured as amplitude signals.

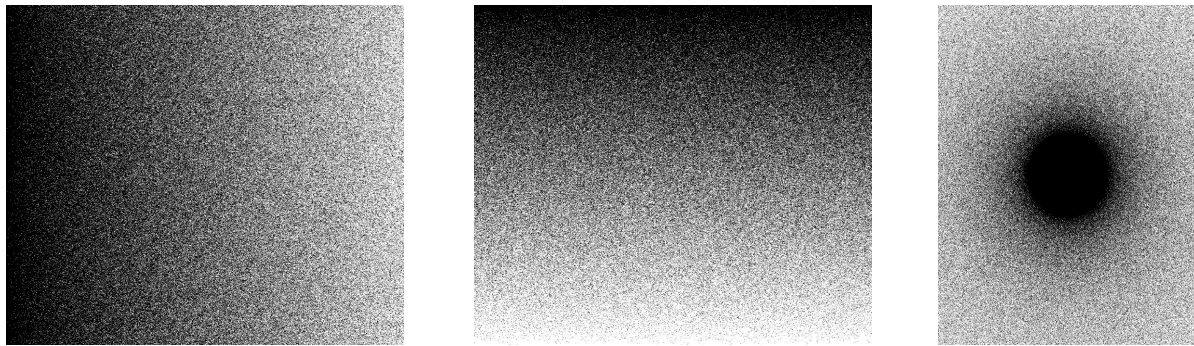


Fig. 5. Variable **transmission function** programming done by students at IWU: shown here are three patterns sent, cyclically, to a DMD: (A) linear gradient along the x -direction, (B) similar gradient along the y -direction, and (C) a (nominally) radial gradient.

Fig. 5 shows bitmaps describing the control signals sent to a DMD. While each pixel is either 100% “on” or “off,” for any input beam that is large compared to the size of the micro-pixels, the net effect is a smooth, linear gradient from black to white (from left to right, from top to bottom, and from middle outward radially, respectively). Since, in the application described above, the distribution of the light energy incident on the DMD can be treated as reflective of the centroid position of an optically trapped particle, it follows that as the image of an optically trapped micro-particle moves across the DMD, the imposed gradient results in a signal at a (downstream) PhotoDiode that is proportional to the centroid position of the optically trapped particle. If the optical trap behaves like a parabolic potential well, then from Hooke’s Law we can convert from particle displacements to (ultrafast) optical force measurements along a particular direction (limited in speed only by the bandwidth of the downstream PhotoDiode, a key point, which may allow you to access *new science*). By cycling between the three gradient patterns that are shown in Fig. 5, one can measure force components acting along the x -, y -, and z -directions, respectively (at a rate limited by the response time of the DMD array). Our LightCrafter 4500 DMD is just a fraction of the cost of the research-grade ViALUX systems, but can be updated at **4400 Hz**. (The ViALUX systems update even faster and provide more on-board memory but, as described above, this project merely uses the DMD to toggle between three distinct patterns.)

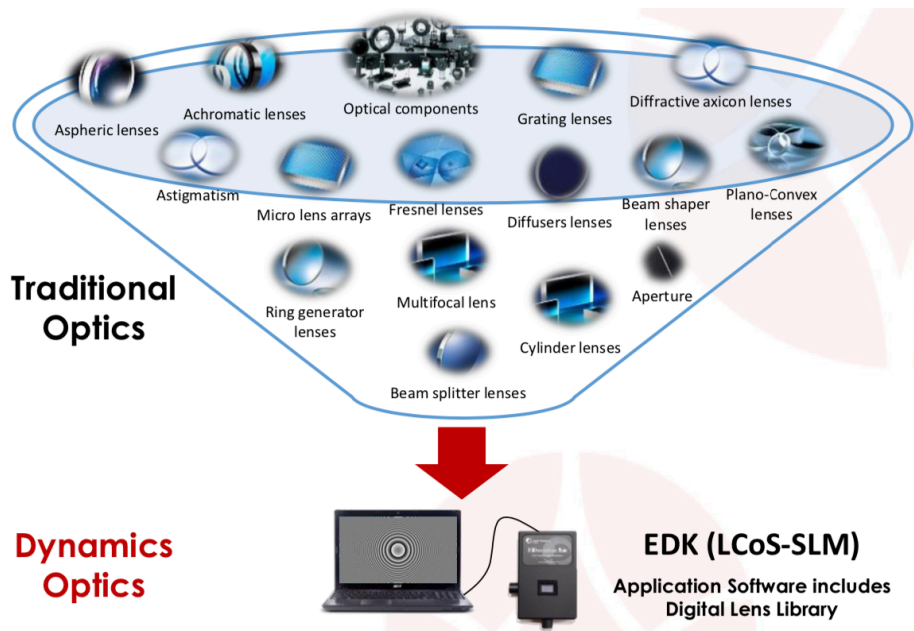


Fig. 1. Replacing *fixed* optical elements with a *dynamically* programmable device ([Jasper Display](#))

Traditional optical elements can be, for a number of experiments, quite difficult to “align” and, for a *systematic* study, may need to be sequentially replaced and *re*-aligned — a workflow that can become drudgery rather quickly, unless you are very “zen.” With Digital Optics we *replace* this series of fixed optical elements with one dynamically programmable device; so, after one initial alignment, no further alignment will be required and, at least once you’ve completed this course, you’ll see that, with Digital Optics, even that initial alignment can be *automatically* optimized. Clearly, a Digital Optics device can be programmed to replace any one of a set of traditional optical elements (see Fig. 1), but what’s more, it can even be programmed to behave like a *superposition* of a number of different optical elements, all at once, including a wide variety of very expensive *custom* components that you would be unlikely to go out and purchase.

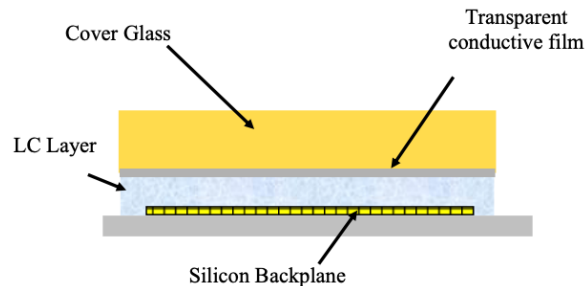


Fig. 2. Digital optics devices, such as the reflective Liquid Crystal-on-Silicon (LCoS) above ([Meadowlark Optics](#)) are based upon an array of independently addressable “pixels.” Unfortunately, *even when the device is turned off*, this pixelation means that output light will be structured, in a manner that depends upon diffraction, due to the size of each pixel, and interference, depending upon the periodic spacing between pixels. The resulting effect on the output beam is sometimes referred to as “**diffraction noise**” (the topic of *next* week’s lab).

In addition to the DMD, our exploration of Digital Optics devices includes liquid-crystal-on-silicon (LCoS) “Spatial Light Modulators” (SLMs), shown in Fig. 2, which are highly reflected grid-like pixelated structures with a liquid crystal solution atop the array of pixel electrodes. LCoS-based SLMs in our lab are driven by an 8-bit CMOS circuit, allowing you to select from $2^8 = 256$ discrete voltage levels that can be applied to any individual pixel.

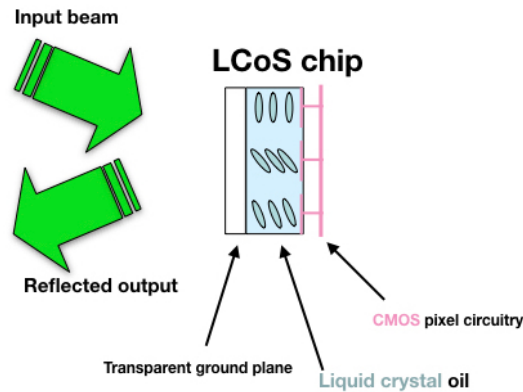


Fig. 3. A reflective Liquid Crystal-on-Silicon (LCoS) chip includes a very large array of independently addressable “pixels,” of which only a few are shown: the mirror-like electrodes on the right side of this schematic form, in conjunction with the transparent ground plane at left, a capacitor. Between the plates of each capacitor are long, rigid-rod-like molecules with polar groups that tend to align. The surfaces of the device have been made either hydrophobic or hydrophilic, to ensure that in the absence of an applied voltage there is a default orientation. Applying a small voltage creates a small tilt of nearby molecules. Applying a somewhat larger voltage yields a somewhat larger tilt.

Fig. 3 schematically illustrates that the voltages applied to each pixel adjusts the alignment angle of nearby liquid crystal molecules, sandwiched between a transparent ground plane and reflective electrodes. Our resulting control over the *tilt* of the rod-like liquid crystal molecules changes the local (pixel-area) **refractive index** for any incident light along the tilt axis (which we call the x -axis). After all, the degree to which electrons can *slosh* in response to an incident field depends upon how the polarization of that input field relates to the orientation of the rod-like molecules. The spatial extent of the liquid crystal molecules is small compared to the wavelength of light utilized, so we can model them as narrow (1D) “wires,” usually represented as short straight line segments in schematic drawings, where the electronic polarizability along the long axis of the wire is larger than the polarizability along either orthogonal axis. Any atomic-scale geometry is *unresolvably* small, given the wavelengths of light we are working with, and so we are really only sensitive to the “effective medium” properties that describe *averages* over thousands of individual molecules or more. For our purposes, these effective medium properties are largely described by the index of refraction. The liquid crystal solution utilized is a transparent oil, whose viscosity limits the response time, typically to **tens of Hz** even for a *reflective* device (which, in comparison to a *transmissive* SLM requires only *half* the thickness of liquid crystal solution to achieve the same phase throw, as discussed below).

Previously, in introducing birefringent materials, we considered a generic medium characterized by one index of refraction, n_e , along the axis corresponding to the long direction of an “average” wire-like molecule in a highly oriented liquid crystal solution, ...and a different index, n_o , along the perpendicular axes. Such a material is said to be “uniaxial” in that only one direction is “different.” (By convention, the subscript on the index of refraction characterizing the other two directions indicates that they are “ordinary,” while the subscript on the other characterizes it as “extraordinary.”)

Thus, if the orientation of the liquid crystal were adjusted such that a plane-wave input beam traveled along what we’ve called the *extraordinary*-axis, then the polarization of that input beam would necessarily be *perpendicular* to that axis, and so the index of refraction would be n_o , regardless of which direction the polarization of that input beam might be rotated. So, for this orientation, the liquid crystal would appear to be an “isotropic” material. On the other hand, the SLM is constructed such that the input beam travels along one of the *ordinary* directions, so that the response will depend upon the polarization of the input beam: for a polarization that is transverse to the long direction of the molecule, the index of refraction remains n_o , but a more general polarization direction would be analyzed by breaking it into components, where the polarization along the tilt direction is characterized by an index of refraction that we can tune all the way from n_o to n_e , as the liquid crystal molecular orientation is tilted from zero to 90° . You should attempt some **sketches in your lab notebook**, to ensure that you have this concept clearly in hand.

Laser fields are normally *quite* small in comparison to the molecular fields describing the electronic interaction with the ionic cores that are in close proximity, and so the *sloshing* we induce typically corresponds to only a small perturbation away from the equilibrium distribution. Such responses are *linear*, where the displacement is *proportional* to the laser’s electric field strength, and the electrons slosh at the frequency of the applied drive (*i.e.*, the laser). Yet the index of refraction is defined to be the ratio of the speed of light in vacuum to the effective speed in the material, $n \equiv c/v$. In other words, a larger index of refraction corresponds to a slower wave speed. You also know that a simple wave travels one wavelength in one period, or $v = \lambda f$. The fact that the frequency of the applied field is equal to the frequency of the induced polarization means that there are no new frequencies to consider as we move into a slower medium (in this linear regime). From $v = \lambda f$, we conclude that as light moves into a medium with a slower wave speed (*i.e.*, characterized by a larger index of refraction) it must have a shorter wavelength, $\lambda = \lambda_0/n$, where λ_0 is the wavelength this beam would have in vacuum. This change in wavelength, in turn, changes the *effective optical path length* for any *component* of the input beam that is polarized along the x -axis. In other words, the SLM behaves like a wave plate of tunable “effective” thickness, which you may sometimes see called a variable retarder: as we tilt the liquid crystal molecules, we are **tuning the relative phase of the ordinary and extraordinary components of the output beam**. As you learned from your introduction to birefringent waveplates, such phase differences can then cause changes in the polarization state of the output beam.

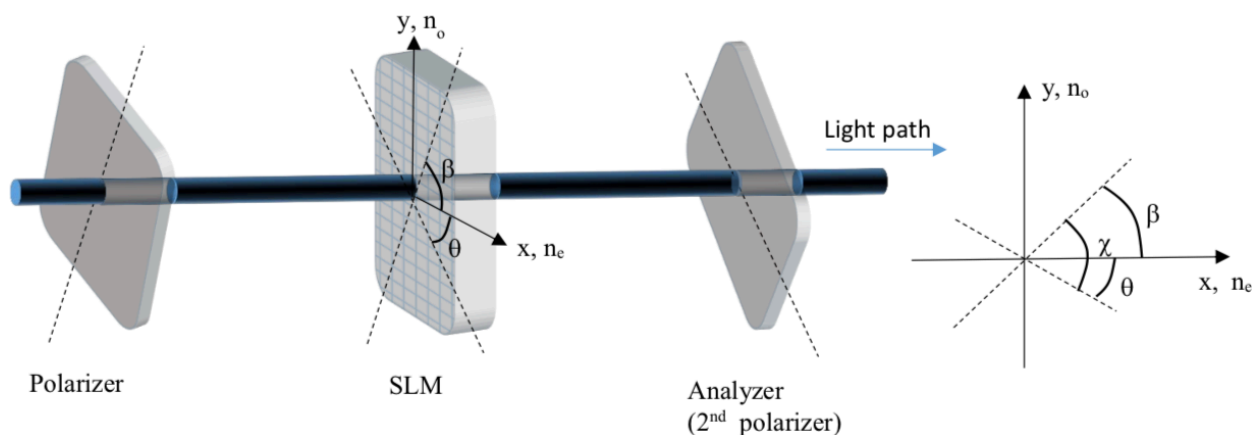


Fig. 4. Your real experimental layout will differ somewhat from this schematic (e.g., you will primarily use reflective SLMs rather than transmissive SLMs), but the heart of the setup is still described as: **Polarizer + SLM + Analyzer**. *Nota Bene*: the x - and y -axes in this schematic are fixed by the SLM orientation, and normally differ *significantly* from the horizontal and vertical. [Fig. credit: Lowell McCann, UW-River Falls]

By setting a polarizer on the input beam and then, following the SLM, an analyzer on the output (where, again, we refer to any polarizer placed just before a detector as an “analyzer”), we select the *kinds* of adaptive control that the SLM will have upon the beam (e.g., local control over the *amplitude* of the output beam, or local control over the *phase* of the output beam, or some combination of the two).

Q1) Describe how to use such an SLM for *phase-only* modulation. That is, specify β , the orientation of the transmission axis of the input polarizer, with respect to the extraordinary axis of the SLM, then specify θ , the orientation of the transmission axis of the analyzer with respect to the extraordinary axis of the SLM (or, equivalently, you could specify χ , the orientation of the transmission axis of the analyzer with respect to that of the input polarizer) [Note, however, that the extraordinary axis of the SLM differs quite a bit from the horizontal direction in the lab!]:

Q2-Part i) Describe how to use such an SLM for *amplitude-only* modulation:

Q2-Part ii) When your set-up is configured for *amplitude-only* modulation, describe changes in the output beam, as you vary a voltage applied to the SLM:

- first in terms of the evolution of the [Polarization Ellipse](#), and
- then in terms of “trajectories” on the [Poincaré Sphere](#).

[An interactive visualization is available via the [Wolfram Demonstrations Project](#); more extensive discussion of these topics may be found [here](#);

Review: (based on Sect. 1.2 - 1.3 of the Jasper EDK model SLM user manual)

(I) Jones Matrix & Uniaxial Crystals

The Jones Matrix formalism is a tool in handling problems dealing with the polarization of light. It takes advantage of the fact that a simple 2×1 matrix can be used represent the polarization state of a plane wave.

$$\vec{E} = \begin{bmatrix} E_x \\ E_y \end{bmatrix}$$

where E_x and E_y are complex numbers, encoding both amplitude and phase. Furthermore, the *action upon* that polarization state, by various optical components, can be simply represented by operating upon the input polarization state with appropriate 2×2 matrices. The resulting 2×1 matrix represents the final status (polarization direction, amplitude, phase change) of the beam after passing through the optical component.

Optical Element	Jones Matrix
Horizontal Polarizer	$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$
Vertical Polarizer	$\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$
Linear polarizer, transmission axis at θ w.r.t. horizontal	$\begin{bmatrix} \cos^2 \theta & \cos \theta \sin \theta \\ \cos \theta \sin \theta & \sin^2 \theta \end{bmatrix}$
Quarter-wave plate, fast axis at $\pm 45^\circ$	$\frac{1}{\sqrt{2}} \begin{bmatrix} 1 & \mp i \\ \mp i & 1 \end{bmatrix}$
Quarter-wave plate, fast axis at θ w.r.t. horizontal	$\begin{bmatrix} \cos^2 \theta + i \sin^2 \theta & (1 - i) \sin \theta \cos \theta \\ (1 - i) \sin \theta \cos \theta & \sin^2 \theta + i \cos^2 \theta \end{bmatrix}$
Half-wave plate, fast axis at θ w.r.t. horizontal	$\begin{bmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{bmatrix}$

Fig. 5. A few commonly encountered Jones matrices

For example, let's apply this formalism to describe an SLM where the liquid crystal molecular direction (characterized by n_e) is, at the moment, aligned with the x -axis (and the y -axis is characterized by n_o). Referring back to Fig. 4, our first step might be to describe the polarization of the input beam in terms of its components along the x - and y -axes. The normalized Jones vector for modeling the polarization state of the light incident upon the SLM is set, by the first polarizer, to be:

$$\begin{bmatrix} \cos(\beta) \\ \sin(\beta) \end{bmatrix}$$

If this input beam must traverse a *total* path length of ℓ in the liquid crystal medium (which, for normal incidence upon a *reflective* SLM corresponds to twice the thickness of the liquid crystal solution), then the action of the SLM upon this beam is modeled by the following 2×2 matrix:

$$\begin{bmatrix} e^{-i2\pi\frac{\ell}{\lambda_0/n_e}} & 0 \\ 0 & e^{-i2\pi\frac{\ell}{\lambda_0/n_o}} \end{bmatrix} = \begin{bmatrix} e^{-i\left(\frac{2\pi}{\lambda_0}\ell\right)n_e} & 0 \\ 0 & e^{-i\left(\frac{2\pi}{\lambda_0}\ell\right)n_o} \end{bmatrix}$$

The polarization of the light exiting the analyzer is then predicted to be:

$$\begin{bmatrix} E_x^f \\ E_y^f \end{bmatrix} = \begin{bmatrix} \cos^2 \theta & \cos \theta \sin \theta \\ \cos \theta \sin \theta & \sin^2 \theta \end{bmatrix} \begin{bmatrix} e^{-i\left(\frac{2\pi}{\lambda_0}\ell\right)n_e} & 0 \\ 0 & e^{-i\left(\frac{2\pi}{\lambda_0}\ell\right)n_o} \end{bmatrix} \begin{bmatrix} \cos(\beta) \\ \sin(\beta) \end{bmatrix}$$

Note: the Jones vector that appears all the way on the right side in the equation above describes what we get from the first polarizer. This is then acted upon by the variable retarder (*i.e.*, the SLM), and finally is acted upon by the analyzer. In other words, the operator corresponding to the *final* optical element appears furthest to the left (so that it is the *final* operator to act).

We began with a *normalized* Jones matrix, so the square modulus of our result above is simply the *fraction* of the intensity that makes it through, which is to say the transmittance, T :

$$T = \left[\begin{matrix} (E_x^f)^* & (E_y^f)^* \end{matrix} \right] \begin{bmatrix} E_x^f \\ E_y^f \end{bmatrix}$$

SLM HW#1: *Show* this simplifies to:

$$T = \cos^2(\theta - \beta) - \sin(2\theta)\sin(2\beta)\sin^2\left[\frac{\pi\ell}{\lambda}(n_e - n_o)\right]$$

There are two special configurations that are of great importance when setting up an SLM:

$$\text{Case A: } \beta = 45^\circ \text{ and } \theta = -45^\circ$$

$$\text{Case B: } \beta = \theta = 0^\circ$$

In Case (A), the input polarizer and the output analyzer are orthogonal, so no light would be transmitted except for the retardation provided by the SLM. Thus, this configuration allows for amplitude modulation. On the other hand, in Case (B), the input is polarized along the extraordinary direction, and so tilting the liquid crystal molecules will not rotate the polarization (meaning that all of the light should pass through the analyzer), but tilting the liquid crystal molecules will still have an effect: this configuration allows for phase modulation.

SLM HW#2: *Write* the expression for the transmittance for Case (A).

SLM HW#3: Confirm for Case (B) the transmittance is one, and *write* the phase shift.