

# Development of hyperspectral image projectors

J. P. Rice\*, S.W. Brown, and J. E. Neira

National Institute of Standards and Technology, Gaithersburg, MD, USA 20899

## ABSTRACT

We present design concepts for calibrated hyperspectral image projectors (HIP) and related sources intended for system-level testing of instruments ranging from complex hyperspectral or multispectral imagers to simple filter radiometers. HIP, based on the same digital mirror arrays used in commercial digital light processing (DLP<sup>1</sup>) displays, is capable of projecting any combination of many different arbitrarily programmable basis spectra into each pixel of the unit under test (UUT) at video frame rates. The resulting spectral and spatial content of the image entering the UUT can simulate, at typical video frame rates and integration times, realistic scenes to which the UUT will be exposed during use. Also, its spectral radiance can be measured with a calibrated spectroradiometer, such that the hyperspectral photon field entering the UUT is well known. Use of such generated scenes in a controlled laboratory setting would alleviate expensive field testing, allow better separation of environmental effects from instrument effects, and enable system-level performance testing and validation. Example potential applications include system-level testing of complex hyperspectral imaging instruments as implemented with data reduction algorithms when viewing realistic scenes, testing the performance of simple fire-fighter infrared cameras under simulated adverse conditions, and hardware-in-the-loop testing of multispectral and hyperspectral systems.

Keywords: Hyperspectral, imaging, projector, testing, video, spectroscopy, radiometry, hardware-in-the-loop, DMD

## 1. INTRODUCTION

Performance testing of instruments with regard to imaging requirements, spectral requirements, and radiometric requirements has traditionally been performed separately, in the limits of simple spatial scenes and spectral profiles. As radiometrically-calibrated hyperspectral imagers are increasingly being developed, we have been motivated to develop testing, calibration, and validation methodologies that keep pace. For proper performance validation, the capability of the instrument, including any calibration and data reduction algorithms, in the measurement of complex scenes should be tested. These complex scenes should represent realistic scenes that the instrument will measure during its lifetime.

To address this issue, we are developing spectral sources and complex scene projectors that can project known complex spectra and hyperspectral scenes, respectively, into instruments for performance testing. We have developed visible and infrared prototypes of programmable spectral sources, as discussed previously [1,2]. In this paper several design concepts are presented using a comprehensive approach. First, we describe the concept of a spectral engine, which is used as a source having a programmable spectral profile. The design of visible and infrared spectral engines and their calibration is discussed. Then we describe the hyperspectral image projector (HIP), a complex scene projector that uses the spectral engine coupled to a spatial engine. The HIP projects an arbitrary spectrum into each spatial pixel, and as such could be used in the most general instrument testing scenarios. Finally, we consider several special cases where only a subset of the HIP, such as a spectral engine or a spatial engine, may suffice.

<sup>1</sup>Note: References are made to certain commercially available products in this paper to adequately specify the experimental procedures involved. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that these products are the best for the purpose specified. DLP and DMD are trademarks of Texas Instruments, Inc.

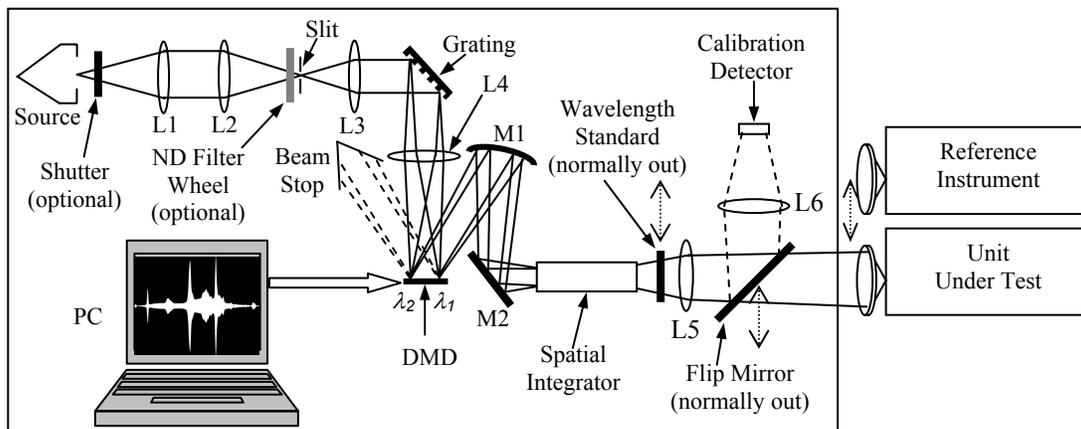
Property of the United States Government. Not subject to copyright.

\*Correspondence: joe.rice@nist.gov

## 2. SPECTRAL ENGINE

### 2.1. Basic concept

The schematic of a spectral engine is shown in Fig. 1. Light from a high intensity broadband source is collected by lens L1, focused by lens L2 through the entrance slit, and collimated by lens L3 onto a diffraction grating, as in a typical spectrograph. The spectrally dispersed output of this spectrograph is imaged by lens L4 onto a  $1024 \times 768$  element digital micromirror device (DMD), such that each of the 1024 DMD columns corresponds to a different wavelength over the spectral range,  $\lambda_1 = 380$  nm to  $\lambda_2 = 780$  nm in this example. At any instant in time, the binary spectral image sent to the DMD from the personal computer (PC) determines which DMD micromirrors are “on” and which are “off.” “On” micromirrors are tilted clockwise ( $+12^\circ$  about the micromirror diagonal) from the normal and thereby reflect the spectrograph output towards M1. “Off” micromirrors are tilted counterclockwise ( $-12^\circ$  about the micromirror diagonal) from the normal and thereby reflect the spectrograph light into the beam stop. Thus the fraction of the 768 micromirrors in each DMD column that are on determines the relative spectral intensity at the wavelength corresponding to that column. Mirror M1 collects the light from the on-state DMD micromirrors and focuses it into the spatial integrator. Mirror M2 is simply acting as a fold mirror. The spatial integrator is an integrating rod, sphere, or other optic that serves to spatially mix the spectral components, thereby presenting a spatially uniform, spectrally profiled beam to the Unit Under Test (UUT) at the output of the spectral engine. The wavelength standard, flip mirror, calibration detector, and the reference instrument are used for calibration as described later, and they are normally not in the beam.



**Figure 1.** Schematic of a spectral engine. The DMD is shown edge-on, and the user-controllable binary spectral image such as that shown on the PC, when displayed on the DMD, determines the instantaneous output spectrum presented to the UUT. The purpose of the optional shutter and the neutral density (ND) filter wheel will be explained later.

Previously we have discussed results from our visible and infrared prototypes of spectral engines similar to Fig. 1. [1,2,3]. Work since that time has focused on automating the operations and optimizing the design. Here we discuss several design issues that we have been considering.

There are several variations in the design details that offer a wide variety of tradeoffs to be considered for any particular application. The grating could be replaced by a prism [4]. The combination of L3, the plane grating, and L4 could be replaced by a concave grating. A slit fiber could be used instead of the slit. While we have used cylindrical lenses for L2, L3, and L4, mirrors could be used instead, or an anamorphic concentrator could be used for higher efficiency coupling from an extended source such as a lamp filament into the slit [4]. Fiber-based supercontinuum laser sources are now available commercially with broadband output in the visible to short-wave infrared, and they represent a promising, though more expensive, alternative to the lamp as a source. In these sources the light is produced in single mode fiber, which due to its small diameter enables high efficiency coupling through very narrow slits. Rather than relying on an integrating rod for the spatial integrator of Fig. 1, it is also possible to spatially integrate the output by

passing it through an integrating sphere or a liquid light guide [2]. For DLP projectors, total-internal-reflection prisms are sometimes used to illuminate the DMD and project at the required angle, and these could be used here instead of the angled collection mirror approach shown in Fig. 1. The input and output can be fiber-coupled or coupled through a liquid light guide. Also, the spatial light modulator itself could be based on a liquid crystal display (LCD) instead of a DMD.

Rather than the forward spectrograph mode shown in Fig. 1, a reverse spectrograph mode is possible whereby the order of the elements from the slit to the DMD, inclusive, is reversed [1]. Then the DMD is illuminated directly by the full broadband light source before the grating, and the light emerges from the slit spatially integrated. While the geometric throughput is the same for both modes, the forward mode has the advantage of less power hitting the DMD for a given output power, and would also be expected to have much less stray light. The reverse mode has the advantage that the light is spatially integrated, in principle without requiring an integrating rod or sphere. A design has also been proposed whereby the DMD is positioned such that after traversing the spectrograph in the forward direction, the light traverses the same spectrograph in the reverse direction [5]. Like a double monochromator, this arrangement would be expected to have both the forward spectrograph advantages of lower power loading on the DMD and lower stray light, and the reverse spectrograph advantage of an inherent spatially-integrated output. However, its efficiency would go as the square of the grating efficiency, as for any double monochromator, and hence output power levels might be more of an issue. From this point of view, a double-pass prism-based spectral engine is attractive due to its higher efficiency.

While Fig. 1 shows the layout used for our visible-band prototype, infrared or ultraviolet spectral engines are conceivable as well with minor changes, as detailed in Section 2.4 for the infrared.

## **2.2. Calibration**

There are two general modes for calibrating the spectral engine shown in Fig. 1: external and internal. For each mode there are two types of calibration: wavelength and intensity. All calibration processes described below can be automated, and need not take very long because of the speed of the DMD.

### **2.2.1. External calibration**

External calibration involves measuring the output of the spectral engine using a well-calibrated spectroradiometer as the reference instrument in Fig. 1. This instrument is moved into position in place of the UUT to view the spectral engine output with a geometry that replicates that of the UUT.

#### **2.2.1.1. External wavelength calibration**

For external wavelength calibration, all micromirrors in a contiguous small subset of one or more DMD columns are turned on, with all others off. This produces a single narrow line spectrum on the spectroradiometer, from which the wavelength corresponding to this column (or set of columns) can be determined by analysis. This process is repeated with a different set of contiguous DMD columns turned on (and the previous set off), until the entire DMD has been sampled and the wavelength corresponding to all column sets has been deduced. For our grating-based prototypes, the wavelength vs. column relationship generally fits a straight line very well. The process of data acquisition and analysis, implemented manually in our previous work, is now completely automated for the prototype of Fig. 1.

#### **2.2.1.2. External intensity calibration**

The purpose of intensity calibration is to determine, for each DMD column (or wavelength), the spectral engine output intensity (more properly called spectral radiance) as a function of which micromirror elements are turned on in a given column. This function will generally be slightly nonlinear, because the beam from the spectrograph illuminating the DMD is not perfectly spatially uniform. Rather, the DMD illumination is generally greatest in the center and falls off towards the edges. Thus, spectral images are usually designed to take advantage of the symmetry of the illumination profile by folding the spectrum about the horizontal line through the center of the DMD, as shown in the example

spectral image on the PC display in Fig. 1. External intensity calibration is similar to external wavelength calibration, except that instead of individual small sets of contiguous DMD *columns* being turned on individually, it is DMD *rows* that are turned on individually, starting from the center of the DMD. Again using the calibrated spectroradiometer, a spectrum is acquired for each set, and the cycle is repeated until all DMD rows have been sampled. The intensity associated with each row can then be determined by plotting intensity vs. row position for each wavelength. However, external intensity calibration can proceed in two possible ways: discrete or integrated. The difference is that for discrete, each set of rows is *turned off* after each cycle of measuring the associated spectrum. For integrated, the rows (starting from the center) are *left on* after each cycle, so the intensity measured is cumulative and the last measurement for integrated intensity calibration is with all DMD mirrors on. Thus the degree to which the system is additive can be checked: ideally, for a system that is perfectly additive, the analytic sum of the measured discrete intensity spectra will equal the measured integrated intensity spectrum. Negative effects such as stray light, diffraction, detector noise, and, in the thermal-infrared, uncorrected self-background all act to force the system away from the ideal additive condition, so comparison of discrete versus integrated intensity calibration can be used to quantify and separate relative influences of these negative effects for a particular spectral engine design. By properly incorporating the information available from external intensity calibration measurements into the spectral engine calibration algorithm, improved matches to arbitrary spectra will be obtained.

### **2.2.2. Internal Calibration**

Internal calibration is for rapid and routine (perhaps daily) spectral engine re-calibration for the usual situation where the reference spectroradiometer is not available or it is otherwise just not practical or cost effective to carry around with the spectral engine unit. Also, it potentially allows calibration without ever requiring a reference spectroradiometer.

#### **2.2.2.1. Internal Wavelength Calibration**

For internal wavelength calibration, a transmissive wavelength standard (Fig.1) is placed into the beam. Also, a flip mirror is placed into the beam to divert the output to the calibration detector. The wavelength standard transmits spectral line structure having known wavelengths. The calibration detector is a simply a broadband detector, such as a silicon photodiode for the visible/near-infrared range. The DMD is operated in monochromator mode, where a contiguous set of one or more columns is turned on, calibration detector signal is measured, and the process is repeated with a different set of columns turned on (and the previous set turned off) until the entire DMD (and hence spectral range) is sampled. Then the wavelength standard is removed and the DMD monochromator mode scan is repeated. By computing the ratio of these “standard in” to the “standard out” measurements, the spectral transmittance of the wavelength standard is obtained. This is compared to the known spectral transmittance of the standard, and the measured spectral lines are correlated with the DMD columns, forming the wavelength calibration.

#### **2.2.2.2. Internal Intensity Calibration**

Internal intensity calibration relies on having a stable calibration detector system, which includes the calibration detector itself plus any optics in front of it, such as the flip mirror and lens L6 in Fig. 1. It requires either pre-calibration of the calibration detector system before installation into the spectral engine unit, or transfer from an external intensity calibration. Pre-calibration of the calibration detector system prior to installation in the spectral engine unit can be performed in principle by the NIST Spectral Irradiance and Radiance Responsivity Calibrations with Uniform Sources (SIRCUS) facility [6], the NIST Spectral Comparator Facility, or equivalent.

The acquisition process for both cases is the same. The flip mirror is placed in the beam to direct the spectral engine output towards the calibration detector (Fig. 1). The wavelength standard is not used except for the wavelength calibration already discussed above. The DMD is then run in monochromator mode, as was done above for the “standard out” phase of the wavelength calibration, and the signal is measured by the calibration detector at each wavelength, producing a calibration spectrum. As with the external intensity calibration, the number of row elements turned on in a given column can be varied in repeated scans to build up the spectra from which the intensity vs. row element function can be fit.

In principle, the calibration detector signal can be directly converted to spectral radiance by virtue of a pre-calibration. This may potentially avoid the need for an external calibration, but it assumes a very good and relevant pre-calibration, and that the spectral engine obeys the additive condition sufficiently. This condition might generally need to be checked by characterization using an external intensity calibration anyway, at least the first time for a new spectral engine design.

For the case of transferring the internal intensity calibration from an external intensity calibration, no pre-calibration of the calibration detector system is required before it is installed. Instead, just after the external intensity calibration measurements with the spectroradiometer, before anything drifts away from its calibrated state, the flip mirror is placed in the beam, and the data acquisition for internal intensity calibration is performed using the process described above. An *effective* calibration detector spectral responsivity is then determined as the ratio of the signal measured during the internal calibration to the reference spectroradiometer measurement. Since this effective spectral responsivity includes any effects associated with the flip mirror or other parts of the calibration detector system, such as lens L6 in Fig. 1, the component reflectances or transmittances do not have to be individually known.

### **2.2.3. Hadamard Transform Technique**

The Hadamard transform technique can be used during intensity calibration for cases where it improves the signal to noise [7,8]. This technique, as applied to the DMD binary images, involves collecting a sequence of several spectra having several not-necessarily-contiguous sets of rows on at the same time, and thus represents a slight departure from the row-sequential process envisioned above. The optimal choice of which rows to turn on or off for each element of the sequence is encoded by an S-matrix [7]. Transformation of the data sequence using the Hadamard transform then recovers the original set of spectra for each row. There is a multiplex advantage afforded by using this technique, which will improve the signal to noise for intensity calibrations.

### **2.3. Spectrum Matching**

To match the spectral engine output to that of a target spectrum, an initial guess of the corresponding spectral image is first made using the wavelength calibration and intensity calibration. A spectroradiometer measurement of this initial spectrum is made and compared to the target spectrum. In general the initial spectrum may not match the target spectrum exactly due to stray light, diffraction, noise, optical aberrations, or, in the thermal-infrared, uncorrected self-background. The difference between the measured and target spectra is then fed into an algorithm that predicts the next guess for the spectral image. This algorithm uses negative feedback so as to drive the next guess closer to the target spectrum. Another spectroradiometer measurement is made, and the comparison between measured and target spectra is again made and fed back into the algorithm. This process is continued until subsequent iterations yield no further improvement to the match. Then the final matched spectral image is saved on the PC disk until needed later. We have developed a visible-band spectral engine prototype based on Fig. 1. For this, the entire matching process is automated, such that any sequence of arbitrary spectra (such as those from spectral libraries) that can be loaded on the PC in the form of text files can be matched automatically on a given spectral engine.

### **2.4. Infrared Spectral Engine**

The spectral engine concepts described above can be readily extended to the infrared (or ultraviolet, UV) with appropriate material changes to the components as necessary. The micromirrors of the commercially-available DMDs are coated with aluminum, which generally has acceptable infrared reflectance. However, there is a hermetically-sealed glass window on the DMDs, as manufactured, that serves to protect the mirrors. Anti-reflection (AR) coatings can be specified on these commercial DMD windows that are optimized for either the UV, visible/near infrared (VNIR, 380 nm to 1000 nm), or short-wavelength infrared (SWIR, 1  $\mu\text{m}$  to 2.5  $\mu\text{m}$ ), beyond which the transmission of glass is zero. For the thermal infrared (TIR, roughly 3  $\mu\text{m}$  to 12  $\mu\text{m}$ ) spectral range and beyond, the glass window on the DMD has been replaced with, for example, an AR-coated ZnSe window [9]. These have been used to make blackbody-illuminated spatial projectors [9], and laser-based spatial projectors [10, 11]. Diffraction becomes an increasingly significant issue as the wavelength increases in the TIR, since the commercially-available DMD micromirrors have

dimensions comparable to these wavelengths. For example, the  $800 \times 600$  arrays have a pitch of  $17 \mu\text{m}$ , and the  $1024 \times 768$  arrays have a pitch of  $13.68 \mu\text{m}$ .

With a few key changes, the conceptual design of an infrared spectral engine is the same as that shown in Fig. 1. In this case the source would be a high-intensity infrared type, such as a  $1000 \text{ }^\circ\text{C}$  blackbody or an argon-arc [12]. An anamorphic concentrator could be used to optimally match the geometric throughput of the source to that of the slit, thereby maximizing the spectral power that enters the spectrograph [4,13]. The shutter would be used to subtract TIR background as discussed in Section 2.6, and the neutral density (ND) filter wheel could be used for enhanced intensity resolution as described in Section 2.7 below. The beam stop would be replaced by a blackbody, perhaps even cooled below room temperature so as to reduce the TIR self-background discussed in Section 2.6. As before, an integrating rod, here simply an open rectangular tube with aluminum or gold-coated flat-plate mirrors as the internal walls, could be used after the DMD to spatially integrate the dispersed wavelengths, but a second pass through or an additional grating or prism in subtractive mode would probably give better spatial uniformity at the output. Rather than the simple silicon photodiode used for the VNIR, the calibration detector would be an appropriate detector system for the infrared spectral range. To implement internal intensity calibration using a pre-calibrated detector, the calibration detector system can be calibrated against the NIST detector-based scale at the IR-SIRCUS [14].

The reference instrument used for external calibration would be, for example, a Fourier-Transform Infrared (FTIR) spectroradiometer with an on-board radiometric calibration unit. Such units are commercially available and consist of a gold-coated scene selection mirror and two high-quality cavity-type blackbodies, one operated at a higher temperature than the other. These are used to implement a two-point FTIR spectral radiance calibration algorithm [15]. The scene selection mirror can rapidly (within a few seconds) switch between either of these two calibration blackbodies or the desired scene (the output of the spectral engine in our case).

The spectral profile of the output of the spectral engine at any instant in time is determined by the instantaneous binary spectral image on the DMD. The DMD is capable of playing back such sequences of binary images at a rate greater than  $5000 \text{ Hz}$ , and the playback of the images can be synchronized with the UUT by triggering from a signal from the UUT. In the limit that the spectral engine itself does not add any artifacts to the spectral radiance, the realism of the apparent spectra is thus determined by how well the spectral images provide a match to realistic spectra, and these are completely programmable by the user. The scenes can include, for example, arbitrary targets with arbitrary backgrounds, with arbitrary interferents, and they can be time dependent. Introduction of artifacts into the spectral radiance can be avoided by careful electronic interface with the UUT. There are two aspects of this that will be discussed in the next two sections. One concerns time sync; the other concerns TIR self-background. While the discussion assumes that the UUT is an FTIR spectrometer, it applies to other types of potential UUT systems such as dispersive spectrometers and filter radiometers.

## 2.5. Synchronization to the UUT

It will generally be necessary to synchronize the spectral engine to the UUT and to the reference instrument. If the DMD binary spectral image were to change during an FTIR interferometer scan, for instance, spurious temporal effects (e.g., aliasing) would result, providing apparent spectra that are highly distorted. However, there is plenty of time for the DMD binary spectral image to be changed *between* interferometer scans, since it takes less than  $24 \mu\text{s}$  for the DMD mirrors to switch from on to off or visa versa [16]. Generally, FTIR interferometer scan mirrors take much longer than this to turn around between interferometer scans, as they are much more massive than the tiny DMD micromirrors.

Synchronization and DMD timing control in our prototype is implemented by using a commercially-available DMD Discovery 1100 drive electronics board augmented with an Accessory Light-modulator Package (ALP) electronics board [17]. The ALP can either generate an internal timing reference signal or can be synchronized to an external timing reference signal. It ensures that the spectral engine DMD, UUT, and reference instrument are all synchronized to this reference signal such that, for example, the DMD binary spectral image changes only between interferometer scans.

## 2.6. Thermal-Infrared Self-Background Subtraction

The purpose of the shutter in Fig. 1 is to subtract the thermal-infrared (TIR) self-background. Note that this is an entirely separate background component from any TIR scene backgrounds, which may be desired as part of the projected scene. The undesired self-background arises from the infrared radiance emitted by the components within the projector itself, as they are at finite temperature (mostly at room temperature). This includes, for example, the room temperature infrared emission from the spectrograph, DMD, optics, etc. Even when the DMD mirrors are in their off position, and even if the cold BB is cold enough to provide effectively zero radiance, the non-zero emissivity of the DMD mirrors themselves will provide a small, though non-zero, TIR self-background. This self-background can be measured periodically by closing the shutter (located near the source in Fig 1, though far enough away that the source does not heat the shutter significantly) and acquiring one or more interferograms with the shutter closed. These interferograms can be co-added as necessary, then subtracted from any interferograms acquired when the shutter is open, producing a delta-interferogram. This process can be repeated whenever the self-background changes significantly throughout the course of the testing. The delta-interferogram will then be fed back into the UUT and used as input to the UUT interferogram-to-spectrum processing algorithm, as this represents the interferogram that would be acquired by the UUT for real scenes. Using this process, the undesired TIR self-background is always subtracted, and any desired scene background can be added by explicitly including it in the spectral sequence that is sent to spectral engine to display.

## 2.7. Intensity Resolution

Variable intensity at each wavelength in the basic spectral engine is achieved by the fraction of the 768 micromirrors that are on in a given column. For dc operation, this sets a theoretical limit of 1/768 (or 1/1024 if we spectrally disperse across rows rather than columns) to the intensity resolution for an XGA DMD. For increased intensity resolution, the time domain can be used, either by pulse width modulation or a rotating filter wheel as described next.

### 2.7.1. Pulse Width Modulation

Pulse width modulation (PWM) is commonly used to achieve variable intensity (gray scale) in DLP projectors [16]. In the version of PWM implemented by the ALP firmware, the available “picture time”  $\tau_p$ , typically several ms, is subdivided into small time intervals  $\delta t$ , given by

$$\delta t = \frac{\tau_p - \tau_0}{(2^n - 1)} \quad (1)$$

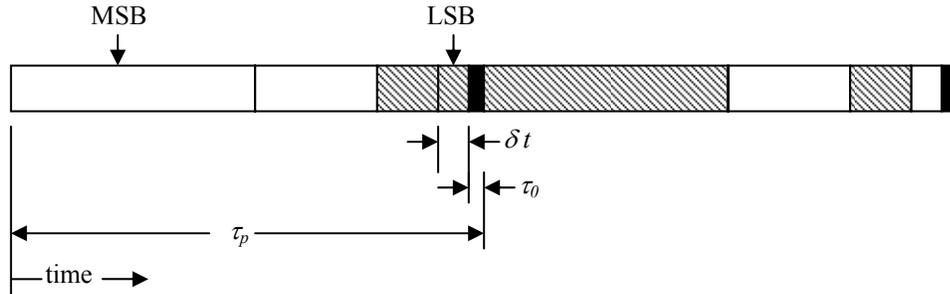
where  $\tau_0$  is fixed at about 124  $\mu$ s. As shown in Fig. 2, the most significant bit (MSB) occupies a time  $2^{(n-1)}\delta t$ , the next-to-most significant bit occupies time  $2^{(n-2)}\delta t$ , etc., down to the least significant bit (LSB), which occupies time  $\delta t$ . Integrated over the time interval  $\tau_p$ , this provides grey scale with  $n$  bits of intensity resolution.

### 2.7.2. Filter Wheel

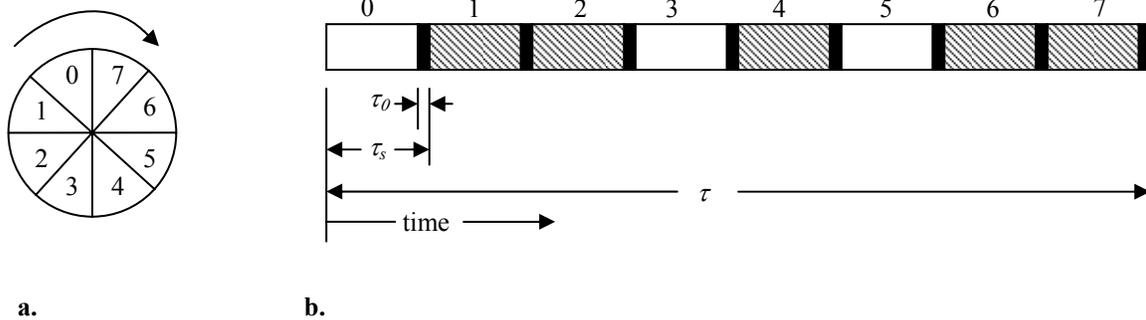
A neutral density (ND) filter wheel can be used as shown in Fig. 1 to improve the overall intensity resolution while maintaining a constant micromirror state during each UUT cycle. The light level illuminating the DMD during any cycle can be varied in  $M$  steps by using an ND filter wheel that has  $M$  equally-partitioned sectors, each consisting of a neutral density filter having a transmittance  $T_{ND}$  given approximately by

$$T_{ND} = \frac{1}{2^j} \quad \text{for } j = 0, 1, \dots, M-1 \quad (2)$$

where  $j$  enumerates the sectors as shown for  $M = 8$  in Fig. 3. The ND filter wheel spins in sync with the DMD and the interferometer scanning such that it makes one complete revolution for every  $M$  interferometer scans and the transition between sectors occurs during interferometer scan mirror turnarounds. Then the interferograms are co-added in groups of  $M$ . When a DMD mirror element is on during a given sector period, the relative intensity for that mirror element is weighted by a factor proportional to the transmittance of that sector by Eq. (2). Thus, for example in the case of  $M = 8$ , the intensity resolution is improved, compared with the basic dc-mode spectral engine, by a factor of  $2^8 = 256$ , providing an overall theoretical intensity resolution of  $1/(768 \times 256) = 1/196608$ .



**Figure 2.** Pulse width modulation as implemented by the ALP. Shaded blocks represent intervals during which the micromirror is off (binary state 0); unshaded blocks represent on (binary state 1). The example shown is for two cycles (ALP “pictures”) at  $n = 4$ , where the first cycle encodes binary level 1100 (decimal 12) and the second cycle encodes binary level 0101 (decimal 5). The micromirror is always off during the dark time  $\tau_0$ , which occurs once per cycle and is shaded black. Any transition between the two micromirror states takes about 24  $\mu\text{s}$  to complete (not shown).



**Figure 3. a.** Neutral density filter wheel for increasing intensity resolution while maintaining a constant micromirror state during each sector interval  $\tau_s$ . **b.** Timing diagram for the example of  $M = 8$ . As in Fig. 2, shaded blocks represent the micromirror off state while unshaded blocks represent the on state, except here the ALP is run in binary mode with no PWM. If the filter transmittance varies according to Eq. (2), the integrated intensity over the filter wheel rotation period  $\tau$  can be varied with  $M$  bits of resolution. The example shown encodes binary 10010100 (decimal 148, corresponding to a relative intensity of  $148/255 = 58\%$ ). The ALP dark time could be adjusted to correspond to the transitions between the different sectors.

### 3. SPATIAL SCENE PROJECTORS

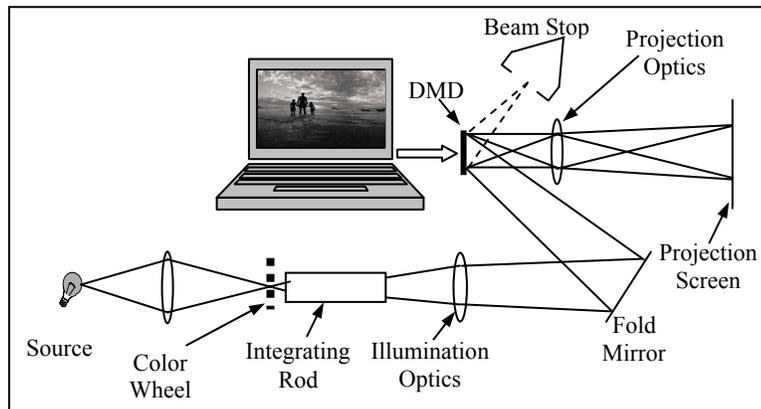
The ultimate goal of scene projection is to project realistic, controllable, and known static or dynamic scenes into the UUT, so that expensive field testing of the UUT can be replaced by less expensive lab testing. Ideally, the scenes would be simultaneously controllable, realistic, and known in several domains that are usually considered separately: spatial, spectral, radiometric, temporal, and polarization. Realistic test and calibration systems must sacrifice realism in several of these domains in order to achieve acceptable controllability and knowledge in one of them. It will be seen

below that the HIP attempts spatial, spectral, and radiometric realism, control, and knowledge by only sacrificing temporal realism within the integration time of the UUT. Polarization control, though possible in principle, is beyond the scope of this paper.

The HIP can be understood as an extension to the common DLP color projector, where the color wheel is replaced by a spectral engine. Thus, a brief review of DLP color projectors is provided first. Then the HIP itself is described, applicable to the general case where the UUT is a full hyperspectral 2-D imager. Finally several special cases are discussed, applicable to simpler cases for the UUT.

### 3.1. DLP Color Projector

The basic principles of a single-DMD DLP color projector are illustrated in Fig. 4. Light from the source, normally a high-intensity lamp, after passing through a color wheel and an integrating rod, uniformly illuminates the DMD. The instantaneous binary spatial image, comprised of the “on” pixels of the DMD, is then projected onto a diffuse screen for viewing by human observers. The light from the “off” pixels is absorbed by a beam stop and is not projected. The parts of Fig. 4 after the color wheel are collectively referred to below as the spatial engine.



**Figure 4.** Conventional single-DMD DLP projector, consisting of a source, a color wheel, and a spatial engine.

The color wheel is typically composed of three colored sectors, each having a red (R), green (G), or blue (B) filter. A fourth, smaller, clear sector to provide additional white balance is sometimes used but can be ignored for the purposes of this discussion. As the color wheel spins in sync with the video refresh rate of the screen, the red, green and blue color planes of the image are each displayed sequentially by the DMD. The video rate is fast enough that a human observer integrates the time-sequential red, green, and blue images into a single color image. Pulse width modulation of the DMD is used to achieve 8 bits of intensity resolution per color plane, providing  $2^{24}$  (more than 16 million) colors. It is important to appreciate, however, that though the projected light can be well color-matched by a human observer, the underlying basis spectra are still simply combinations of the R, G, and B filtered light, and as such the projected, time-integrated spectra cannot be spectrally matched to arbitrary spectra, even those within the same visible band. This is exactly the limitation that is overcome by the HIP, described next.

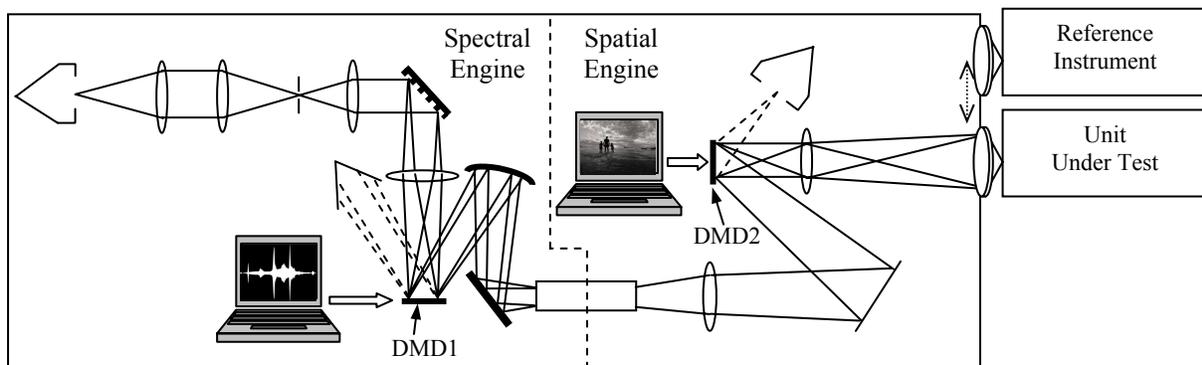
### 3.2. Hyperspectral Image Projector

#### 3.2.1. Basic Concept

The most general case considered here is to project an arbitrary spectrum into every pixel of a 2D staring array UUT. Imagine replacing the color wheel of a DLP color projector with the spectral engine described in Section 2. The result

is a HIP, consisting of a spectral engine optically in series with a spatial engine, as shown in Fig. 5. The images displayed on DMD1 are synchronized to those displayed on DMD2. This is implemented, for example, by using an ALP on each DMD, synchronized to one another and to the UUT and the reference instrument. For each frame, and hence during the UUT integration time, the spectral engine scans through a sequence of basis spectra, and the spatial engine synchronously scans through a corresponding set of spectral planes. Each spectral plane is the spatial image associated with a particular basis spectrum, analogous to the way that each color plane of an RGB color image is the spatial image associated with a red, green, or blue primary color. A wide variety of spectra can be obtained by adding combinations of the basis spectra, just as a wide variety of colors arises by adding combinations of primary colors. As the spectral planes are all displayed within the integration time of the UUT, this summation happens temporally in the UUT. The sum of the sequence of spectral planes displayed on DMD2 generates the hyperspectral spatial image projected for a given frame, just as the sum of the color planes generates the full-color spatial image during each frame in a DLP color projector. For the purposes of this paper, we consider only cases where DMD1 does not use PWM, but where DMD2 does.

As a first example, consider the operation of the HIP in the RGB limit where it emulates a DLP color projector. In this limit, PWM of DMD2 is used to achieve 8 bits of intensity resolution per spectral plane, and at 7 ms per spectral plane provides a frame period of 21 ms, which appears flicker-free to a human observer for the particular conditions under which our prototype was operated. Meanwhile, DMD1 in the RGB limit is operated with 1 bit of intensity resolution (binary) per spectral plane, and the three binary spectral images are simple bands that correspond to selecting the red, green, or blue portions of the spectrum as dispersed across DMD1 in the spectral engine. Fig. 6 shows the timing diagram for another, simpler example.



**Figure 5.** Schematic of the Hyperspectral Image Projector (HIP). A spectral engine optically in series with a spatial engine provides a time-integrated 2D image to the UUT whereby each pixel is presented with an arbitrary spectrum.

### 3.2.2. HIP Frame Formats

More generally, let  $N'$  represent the number of basis spectra, each generated in the spectral engine as a binary image on DMD1. The sequence of these  $N'$  images is displayed within the UUT spatial frame time,  $\tau$ , which is the reciprocal of the frame rate,  $f$ . Upon display on DMD1 of each spectral image within the sequence, ALP1 sends an electrical pulse to ALP2, which triggers ALP2 to have DMD2 synchronously display the spatial image for the corresponding spectral plane. All spectral planes must be displayed within one spatial frame, so that temporal summation can occur at the UUT. Thus  $N'$  is also the number of spectral planes per spatial frame.

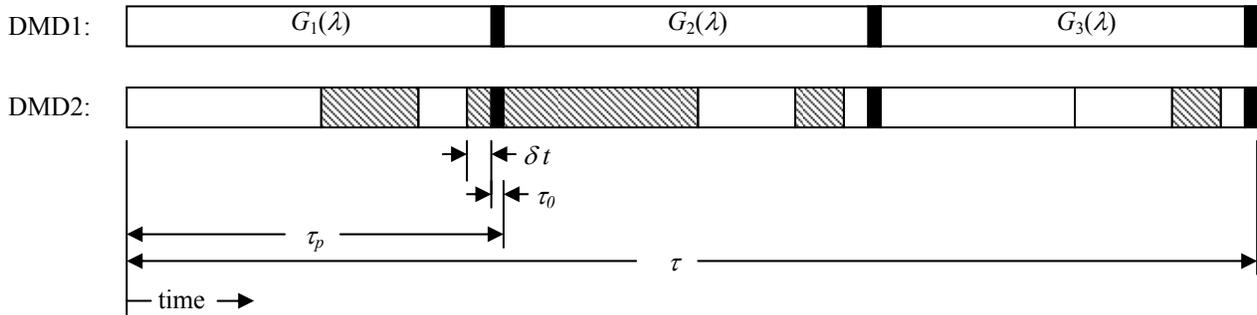
Let  $n$  represent the number of bit planes per spatial DMD picture. The term “picture” has a very specific meaning within the context of using an ALP [17]: the spatial DMD displays one picture per spectral plane. Thus one frame is comprised of  $N'$  pictures. Within each picture, PWM is used by ALP2 to provide intensity resolution  $n$  for that picture, which it does by displaying  $n$  binary images, each for a time period of that is half of the previous one. The current

generation of the ALP that we are using supports a firmware-implemented PWM algorithm where  $n$  can be set to any value from 1 to 8, and this value can be changed with a simple software command.

Some examples of possible frame formats are considered in Table 1. Both  $N'$  and  $n$  are independent variables that can be re-configured by software as necessary to meet given requirements of UUT frame rate, intensity resolution, and length of dynamic video playback. Tradeoffs between these parameters result from the DMD/ALP timing and memory space constraints. The number of ALP2 binary images per spatial frame is  $nN'$ , whereas for ALP1 it is only  $N'$ . Thus ALP2 presents the timing bottleneck. The maximum ALP2 display rate  $p$ , in pictures per second, decreases with increasing  $n$  as shown in Table 1 [17]. Thus, the maximum UUT frame rate  $f_M$  is given by

$$f_M = \frac{p}{N'} \quad (3)$$

and the minimum UUT integration time is the reciprocal of this. The examples of Table 1 all provide frame rates above 45 Hz.



**Figure 6.** Timing diagram for a HIP example where  $N' = 3$ ,  $n = 4$ . PWM on this DMD2 pixel will provide an apparent spectrum  $f(\lambda) = \{10G_1(\lambda) + 5G_2(\lambda) + 13G_3(\lambda)\}/15$  for this spatial pixel, for this example frame as encoded.

| Format Name | Number of Basis Spectra $N'$ | # of Bit Planes per Spatial ALP2 Picture $n$ | # of ALP2 Binary Images per Spatial Frame | Max. Spatial ALP Display Rate, $p$ (pictures/s) | Max. Frame Rate $f_M$ (frames/s) | Min. UUT Integration Time (ms) | Max. # of On-Board Spatial ALP2 Pictures | Max. # of On-Board Spatial ALP Frames | Dynamic Sequence Repeat Period, $\tau_r$ (s) |
|-------------|------------------------------|--|---|---|----------------------------------|--------------------------------|--|---------------------------------------|--|
| Binary      | 138                          | 1  | 138                                       | 6918  | 50.1                             | 19.9                           | 2730                                     | 20                                    | 0.395  |
| Binary2     | 100                          | 1  | 100                                       | 6918  | 69.2                             | 14.5                           | 2730                                     | 27                                    | 0.395  |
| RGB         | 3                            | 8  | 24  | 162   | 54.0                             | 18.5                           | 341                                      | 114                                   | 2.105  |
| HIP1        | 32                           | 4  | 128                                       | 1451  | 45.3                             | 22.1                           | 682                                      | 21                                    | 0.470  |
| HIP2        | 20                           | 5  | 100                                       | 936   | 46.8                             | 21.4                           | 546                                      | 27                                    | 0.583  |
| HIP3        | 10                           | 6  | 60  | 557   | 55.7                             | 18.0                           | 455                                      | 46                                    | 0.817  |

**Table 1.** Example HIP frame formats, including formats for running a HIP in the binary and RGB limits.

For dynamic scene projection, memory issues must also be considered. The current generation of the ALP has 2 GB of on-board SDRAM [17], which translates into the ability to store a maximum of 2730 binary frames at XGA spatial resolution ( $1024 \times 768$ ). A sequence of these frames can be repeatedly played back without interruption, but loading new frames from the PC over the USB 2.0 interface is limited to much less than real time and requires interrupting playback. Despite this limitation of the ALP USB interface relative to standard video interfaces such as Digital Visual Interface (DVI), short sequences of dynamic scenes can be uploaded to the ALPs and repeatedly projected, enabling dynamic testing of the UUT. The repeat period,  $\tau_r$ , of an uninterrupted dynamic sequence cycle depends on  $N'$  and  $n$ ,

since they determine the required number of binary images that must be loaded into ALP memory. The maximum number of ALP2 pictures listed in Table 1 is determined as the greatest integer of  $2730/n$ , and this is divided by the factor  $N'$  to determine the maximum number of frames in Table 1. Dividing this by the frame rate gives the sequence repeat period. This can be summarized as

$$\tau_r = \text{int}\left(\frac{2730}{nN'f_M}\right), \quad (4)$$

where “int” represents the greatest integer function. Referring to Table 1, in the RGB limit, the repeat period is greater than 2 seconds, but in the typical HIP cases it is on the order of a half of a second.

### 3.2.3. HIP Effective Intensity Resolution

During each UUT frame period  $\tau$ , the spectral engine runs through the display of its  $N'$  basis functions  $G_i(\lambda)$ , as shown in Fig. 6. These are synchronized to the display of spatial engine spectral planes such that the intensity level of the apparent spectrum,  $f(\lambda)$ , as projected into an arbitrary UUT spatial pixel, is given by

$$f(\lambda, y_1, \dots, y_{N'}) \equiv \sum_{i=1}^{N'} g_i(\lambda, y_i), \quad (5)$$

where the summation is inherently performed by the UUT over its finite integration time period  $\tau$ , and where  $g_i(\lambda, y_i)$  is the scaled intensity of basis function  $G_i(\lambda)$ . The scaling is controlled by using PWM at the spatial DMD with a bit depth  $n$ , to derive a gray scale with 0 as the minimum intensity and  $G_i(\lambda)$  as the maximum,

$$g_i(\lambda, y_i) = \frac{y_i}{2^n - 1} G_i(\lambda). \quad (6)$$

Here  $y_i$  is an integer between 0 and  $2^n - 1$ , and indicates the gray level to be displayed (for an arbitrary pixel of the spatial engine) during the time that the spectral engine is set to  $G_i(\lambda)$ . For each spectral plane considered individually, the intensity resolution  $\delta g_i/G_i$  is determined only by  $n$ :

$$\frac{\delta g_i}{G_i} \equiv \frac{g_i(\lambda, y_{i+1}) - g_i(\lambda, y_i)}{G_i(\lambda)} = \frac{1}{2^n - 1} \quad (7)$$

where Eq. (6) was used to arrive at this result.

However, the effective bit depth,  $n_e$ , at any wavelength will in general be greater than  $n$ , because the basis spectra may be chosen to spectrally overlap. This can be seen by considering the effective intensity resolution,  $\delta F/F$ , where  $F$  is defined by

$$F(\lambda) \equiv \sum_{i=1}^{N'} G_i(\lambda) \quad (8)$$

Computation of  $\delta F/F$  in a way analogous to Eq. (7) could be done directly by enumerating all the  $f$  levels and computing differences between adjacent  $f$  levels. However, even for spectrally flat basis functions the number of levels is very large since it scales as  $2^{nN'}$ , and the wavelength dependence does not cancel as it did in Eq. (7), making direct computation of intensity resolution non-trivial for realistic cases.

Some insight, as well as a general method for guaranteeing a specified intensity resolution, can be derived by considering the limit of just two basis functions that completely overlap such that one is just an attenuated version of the other,

$$G_2(\lambda) = \gamma G_1(\lambda). \quad (9)$$

Here  $\gamma$  is less than or equal to 1, and is called the flash pairing ratio since it literally means that we will flash a given basis spectrum twice on the spectral engine during the same frame, once for  $G_1(\lambda)$  and a second time for  $G_2(\lambda)$ . The levels can be enumerated by applying Eq. (5) and Eq. (6) to this special case,

$$f(\lambda, y_1, y_2) = \frac{G_1(\lambda)}{2^n - 1} (y_1 + \gamma y_2) \quad (10)$$

from which the intensity resolution can be computed as the minimum difference between non-degenerate  $f$  levels, considering all possible combinations of  $y_1$  and  $y_2$ . The result, for this special case, is

$$\frac{\delta f}{F} = \frac{\gamma}{(2^n - 1)(1 + \gamma)}. \quad (11)$$

As we generalize this to the normal case where additional basis functions are displayed during the frame, and where these additional basis functions do not necessarily have the same spectral shape as  $G_1(\lambda)$ , the attainable intensity resolution will only improve. This is because in spectral regions where these other basis functions overlap with  $G_1(\lambda)$ , they will provide additional intensity that can be controlled. Thus Eq. (11) represents an upper limit, worst case estimate for the intensity resolution as long as the basis functions are always paired such that for each basis function  $G_i(\lambda)$  there exists a paired basis function  $G_i'(\lambda)$ , where

$$G_i'(\lambda) = \gamma_i G_i(\lambda). \quad (12)$$

Here we include the index  $i$  on  $\gamma$  since the value for the flash pairing ratio can be separately chosen for each basis function. Thus to guarantee a given intensity resolution in the general case, we can use the method of flash pairing where during every frame in which the spectral engine flashes  $g_i(\lambda)$ , it also flashes its pair,  $g_i'(\lambda) = \gamma_i g_i(\lambda)$ . Solving Eq. (11) for the value of the flash pairing ratio  $\gamma_i$  that will guarantee a required intensity resolution, we obtain

$$\gamma_i = \frac{2^n - 1}{2^{n_e} - 2^n + 1} \quad (13)$$

where the effective bit depth  $n_e$  is defined implicitly in terms of the intensity resolution by

$$2^{n_e} \equiv \frac{F}{\delta f}. \quad (14)$$

In summary, the effective bit depth,  $n_e$ , at any wavelength will generally be greater than  $n$ , because the basis spectra may be chosen to spectrally overlap. To guarantee a required bit depth without generally constraining the shapes of the basis spectra, flash pairing can be used.

Table 2 shows examples of applying flash pairing to the three HIP frame format examples of Table 1 to achieve effective 8-bit intensity resolution. The number  $N'$  indicates the total number of basis spectra. When flash pairing is used, the number of spectrally unique basis functions,  $N$ , is reduced by a factor of 2. Also, since intensity quantization

from the 768 rows of the DMD will set in for very low values of  $\gamma_i$ , there is probably a practical lower limit for  $\gamma_i$  of about 0.1 or so, depending on the shape of the basis function.

| Format Name | Number of Basis Spectra $N'$ | Number of Spectrally Dissimilar Basis Spectra $N$ | Spectral Plane Bit Depth $n$ | Effective Bit Depth $n_e$ | Flash Pairing Ratio $\gamma_i$ |
|-------------|------------------------------|---|------------------------------|---------------------------|--------------------------------|
| HIP1        | 32                           | 16  | 4                            | 8                         | 0.062                          |
| HIP2        | 20                           | 10  | 5                            | 8                         | 0.138                          |
| HIP3        | 10                           | 5   | 6                            | 8                         | 0.326                          |

**Table 2.** Application of flash pairing to the HIP formats of Table 1 to improve the effective bit depth to 8 bits.

### 3.3. Special UUT Cases

#### 3.3.1. Scanned Single Pixel UUT

When the UUT is a single pixel mechanically scanned across the spatial scene, as in some remote sensing instruments and certain chemical warfare agent (CWA) detectors, a spectral engine alone, rather than a full HIP, may suffice. For this case, the UUT mechanical scanner would be parked so that the UUT constantly stares at output of the spectral engine. A sequence of spectra, played back on the spectral engine DMD, will simulate the different spectral radiance profiles entering the UUT as if it were spatially scanning real scenes. Such sequences could be pre-loaded from the PC to ALP local memory, and would need to be synchronized to the UUT 2D scanner signals so that proper simulation of scanned spatial imagery could be achieved.

#### 3.3.2. Scanned Linear Array UUT

The case of scanned linear arrays, common in remote sensing instruments, could also be handled, in principle, by a simplified version of HIP involving only one DMD. For example, in one proposed design of a hyperspectral scene generator [5], the spectrum is dispersed along one axis of the DMD, as in our spectral engine. However, rather than using the second DMD axis for intensity resolution, it is instead used for the spatial axis that would be aligned with the spatial axis of the UUT linear array. PWM would then be required to provide all of the intensity resolution. As in the above case, the UUT would be operated in stare mode, with its mechanical scanner parked. But for this case, each DMD image would generate the spectra along one axis and one of the spatial dimensions along the other axis. The sequence of DMD images played back would then simulate the second spatial dimension as if the UUT mechanical scanner were actually scanning.

#### 3.3.3. Single-Band 2D Focal Plane Array UUT

In the case where the UUT only has a single band but is a 2D imager, a spatial engine alone, rather than a full HIP, may suffice, depending on the nature of the application. This is the traditional hardware-in-the-loop scenario, and indeed the Micromirror Array Projection System (MAPS), a blackbody-illuminated spatial engine, provides scene projection for this application [9]. One of the applications that we are working on involves projection of imagery for testing of fire-fighter imaging cameras in a way that does not require actual fires. Fire-fighter cameras are usually single-band imagers operating in the 7  $\mu\text{m}$  to 14  $\mu\text{m}$  spectral region. It has been shown that imagery can be projected with, for example, a CO<sub>2</sub> laser operating at 10.6  $\mu\text{m}$  illuminating a MAPS DMD array [10,11]. To first order, as long as the spectrally-integrated irradiance at each pixel of the focal plane is matched by the laser light compared to when the camera views a real scene, the images recorded by the camera will simulate real fires. As a special case of our infrared

HIP development, we are currently developing a laser-illuminated spatial engine to serve this application and determine its range of utility for testing one camera against another and enable comparison to imaging metrics.

#### 4. SUMMARY

We have described the concepts of a spectral engine and of a hyperspectral image projector. We have built visible and infrared spectral engine prototypes, and a visible HIP, and are working towards developing this technology further for a variety of applications involving performance testing and validation of electro-optical sensors and imagers.

#### ACKNOWLEDGEMENTS

This work was funded in part by the NIST Office of Law Enforcement Standards (OLES) and by the NIST Advanced Technology Program. We thank Jerry Fraser, Carol Johnson, Nick Paulter, Francine Amon, and Anthony Hamins for useful discussions of particular applications that motivated this work.

#### REFERENCES

1. J. P. Rice, S. W. Brown, B. C. Johnson, and J. E. Neira, "Hyperspectral image projectors for radiometric applications," *Metrologia* **43**, S61-S65 (2006).
2. S. W. Brown, J. P. Rice, J. D. Jackson, J. E. Neira, B. C. Johnson, "Spectrally tunable sources for advanced radiometric applications," *J. Res. Natl. Inst. Stand. Technol.*, in press (2006).
3. S. W. Brown, J. P. Rice, J. E. Neira, and B. C. Johnson, "Hyperspectral image projector for advanced sensor characterization," *Proc. SPIE* in press (2006).
4. N. MacKinnon, U. Stange, P. Lane, C. MacAulay, and M. Quatrevalet, "Spectrally programmable light engine for in vitro or in vivo molecular imaging and spectroscopy," *Appl. Opt.* **44**, 2033-2040 (2005).
5. N. R. Nelson, P. Bryant, and R. Sundberg, "Development of a hyperspectral scene generator," *Proc. SPIE* **5151**, 480-484 (2003).
6. S. W. Brown, G. P. Eppeldauer, J. P. Rice, J. Zhang, K. R. Lykke, "Spectral irradiance and radiance responsivity calibrations using uniform sources (SIRCUS) facility at NIST," *Proc. SPIE* **5542**, 363-374 (2004).
7. M. Harwit and N. J. A. Sloane, *Hadamard Transform Optics* (Academic Press, New York, 1979).
8. R. A. DeVerse, R. M. Hammaker, and W. G. Fateley, "Realization of the Hadamard Multiplex Advantage Using a Programmable Optical Mask in a Dispersive flat-field near-infrared spectrometer," *Applied Spectroscopy* **54**, 1751-1758 (2000).
9. D. B. Beasley, M. Bender, J. Crosby, T. Messer, and D. A. Saylor, "Dynamic IR scene projector based upon the digital micromirror device," *Proc. SPIE* **4366**, 96-102 (2001).
10. W. R. Folks, J. M. López-Alonso, B. Monacelli, A. Weeks, G. Zummo, D. Mullally, and G. Boreman, "Characterization of a digital-micromirror device-based infrared scene projector," *Opt. Eng.* **44**, pp-pp (2005).
11. W. R. Folks, D. Mullally, G. Zummo, A. Weeks, and G. Boreman, "DMD-based infrared scene projection: a comparison of MWIR and LWIR modulation transfer function," *Proc. SPIE* **5408**, 199-202 (2004).
12. J. M. Bridges and A. L. Migdall, "Characterization of argon arc source in the infrared," *Metrologia* **32**, 625-628 (1996/96).
13. N. Davidson and L. Khaykovich, "High-resolution spectrometry for diffuse light by use of anamorphic concentration," *Optics Letters* **24**, 1835-1837 (1999).
14. G. P. Eppeldauer, J. P. Rice, J. Zhang, and K. R. Lykke, "Spectral irradiance responsivity measurements between 1  $\mu\text{m}$  and 5  $\mu\text{m}$ ," *Proc. SPIE* **5543**, 248-257 (2004).
15. H. E. Revercomb, H. Buijs, H. B. Howell, D. D. LaPorte, W. L. Smith, and L. A. Sromovsky, "Radiometric calibration of IR Fourier transform spectrometers: solution to a problem with the High-Resolution Interferometer Sounder," *Appl. Optics* **27**, 3210-3218 (1988).
16. D. Dudley, W. M. Duncan, and J. Slaughter, "Emerging digital micromirror device (DMD) applications," *Proc. SPIE* **4985**, 14-25 (2003).
17. R. Höfling, and E. Ahl, "ALP: universal DMD controller for metrology and testing," *Proc. SPIE* **5289**, 322-329 (2004).