

Programmable matched filter and Hadamard transform hyperspectral imagers based on micro-mirror arrays

Steven P. Love*

Space and Remote Sensing Sciences Group (ISR-2), Los Alamos National Laboratory
Los Alamos, NM, USA 87545

ABSTRACT

Hyperspectral imaging (HSI), in which each pixel contains a high-resolution spectrum, is a powerful technique that can remotely detect, identify, and quantify a multitude of materials and chemicals. The advent of addressable micro-mirror arrays (MMAs) makes possible a new class of programmable hyperspectral imagers that can perform key spectral processing functions directly in the optical hardware, thus alleviating some of HSI's high computational overhead, as well as offering improved signal-to-noise in certain important regimes (e.g. when using uncooled infrared detectors). We have built and demonstrated a prototype UV-Visible micro-mirror hyperspectral imager that is capable not only of matched-filter imaging, but also of full hyperspectral imagery via the Hadamard transform technique. With this instrument, one can upload a chemical-specific spectral matched filter directly to the MMA, producing an image showing the location of that chemical without further processing. Target chemicals are changeable nearly instantaneously simply by uploading new matched-filter patterns to the MMA. Alternatively, the MMA can implement Hadamard mask functions, yielding a full-spectrum hyperspectral image upon inverting the transform. In either case, the instrument can produce the 2D spatial image either by an internal scan, using the MMA itself, or with a traditional external push-broom scan. The various modes of operation are selectable simply by varying the software driving the MMA. Here the design and performance of the prototype is discussed, along with experimental results confirming the signal-to-noise improvement produced by the Hadamard technique in the noisy-detector regime.

Keywords: Hadamard transform, micro-mirror array, DLP, hyperspectral, MEMS, programmable filter

1. INTRODUCTION

Hyperspectral imaging typically involves the collection of large spectral/spatial datasets, followed by intensive post-processing to extract features of interest. The emergence of micro-mirror array (MMA) technology, such as the commercially-available Texas Instruments "Digital Light Processor" (DLP®) technology, presents the possibility of a whole new paradigm, in which some or all of the required processing is performed optically, within the hyperspectral instrument itself. This new approach offers a variety of advantages, depending on which particular algorithm is implemented. By using the micro-mirror array to implement, directly in hardware, spectral matched filters to detect pre-defined materials of interest, this approach offers greatly reduced data volumes and processing overhead, and significantly simplified operation. Alternatively, one might wish to maintain the versatility of a full hyperspectral dataset, but use the MMA to implement a signal-to-noise enhancing multiplexing algorithm such as the Hadamard transform. One of the advantages of a micro-mirror-based instrument is that the same hardware can be used to perform in a wide variety of modes, simply by changing the software driving the MMA.

This paper describes a working prototype of a programmable MMA-based hyperspectral imager built at Los Alamos National Laboratory. This instrument can work in either a multiplexing Hadamard transform mode, to obtain optimal signal-to-noise when the noise is dominated by detector noise (rather than photon shot noise), or in a chemical-specific matched-filter imaging mode.

2. CONCEPT OF OPERATION

The basic optical layout for a micro-mirror-based programmable spectral imager is illustrated schematically in Fig. 1. In this conceptually simplest case, light from the scene is imaged onto the first image plane, at which a slit could be placed. With a slit, the instrument images a single line across the scene. Light from that line is collimated, then dispersed, by either a grating or prism, and focused onto the micro-mirror array. In this configuration, each row of MMA pixels corresponds to a different wavelength. The role of the MMA is as a wavelength selection device. Each individually

* splove@lanl.gov; phone 505-667-0067

addressable micro-mirror can be tilted to either of two positions; in the case of the TI DLP array, these two positions are mirror tilts of $\pm 12^\circ$ from the array plane. One position is the “on” position, which routes the light onward down the optical train. The other is the “off” position, which routes the light to a beam dump. Thus any wavelength or set of wavelengths may be turned on or off as desired, simply by flipping the appropriate micro-mirrors. Light from the “on” pixels is routed to a second dispersive element, which re-combines the selected wavelengths, and finally imaged onto the detector array.




Figure 1. Basic concept for a micro-mirror array based programmable multi-band spectral filter imager. In this conceptually simple version, a slit restricts the spatial field of view to a single line across the scene. Light is spectrally dispersed and focused onto the spectrally selecting micro-mirror array, then is recombined and focused to a final image on the focal plane array (FPA).

The concept illustrated in Fig. 1 could be used as a hyperspectral imager as is, simply by implementing push-broom scanning either using an external scan mirror, or by using platform motion in the usual manner. However, it is possible to produce a full 2D spatial, 1D spectral hyperspectral imager, where the second spatial dimension is scanned internally on the MMA itself, simply by removing the slit. This is illustrated, in a bit more realistic and less schematic manner, in Fig. 2.




Figure 2. Concept of operation for the Los Alamos micro-mirror array multiplexing hyperspectral imager. Light from the scene is dispersed by a diffraction grating and imaged onto the micro-mirror array, which selects certain wavelengths and rejects others to implement an encoding scheme. The selected wavelengths are recombined by a second grating and imaged onto the detector array. Either chemically-specific matched filters or full-spectrum Hadamard transform functions can be uploaded to the micro-mirrors using the same hardware; only the software driving the micro-mirrors changes between the two options. In either case full 2D spatial imaging can be performed simply by shifting the uploaded patterns on the micro-mirrors, rather than by using an external scan mirror (but pushbroom scanning remains an option, if that is more convenient for the particular application).

In this case, each column of pixels in the first image plane is re-imaged to a corresponding column on the focal plane, as one would hope. But a minor complication arises in the spectrally-dispersed image plane at which the MMA is located. The problem is that each line of pixels in the scene has a different spectral calibration. A full 2D image with the proper matched filter calibration throughout can still be produced, however, by simply reading out one FPA line at a time, while shifting the pattern on the MMA array into the appropriate position for each line of pixels in turn. The result is an instrument that effectively pushbroom-scans the image, but does so internally, without a moving platform or any macroscopic moving parts. This is particularly desirable in fixed, ground-based applications where small instrument size and simplicity are important.

2.1. Chemical- or Material-Specific Matched Filter Imaging

A full-fledged hyperspectral imager generates massive data cubes requiring enormous data handling and processing resources. But for many applications, a full hyperspectral cube is not required, for instance when one is interested in detecting only one, or a few, specific chemicals or materials. In such cases, the micro-mirror instrument can be programmed to implement a chemical matched filter directly in the optical hardware, simply by telling the micro-mirror array which wavelengths to pass and reject for the chemical or material of interest. The result is much-reduced data set, consisting of an image showing the location of the chemical. At its simplest, the micro-mirror matched filter would simply be programmed to pass the wavelengths corresponding to the target material's absorption features. This in itself may be sufficient for many relatively low-sensitivity applications. Further sensitivity is gained by performing a background subtraction, with the micro-mirrors set to reject the chemical's absorption features and instead pass the background between the chemical bands. It is possible in principle to do both simultaneously, if two detector arrays are used, since the micro-mirrors don't actually throw away any light, but merely direct it into one of two possible directions. Figure 3 illustrates the basic idea in a long-wave infrared example, using ammonia vapor as the chemical to be detected.




Figure 3. Using the micro-mirror array as a programmable chemical matched filter. The micro-mirror array is simply programmed to pass wavelengths corresponding to the absorption lines of the chemical of interest (in this case, ammonia), as in pattern B above. This by itself will provide reasonably good sensitivity to the chemical. Further sensitivity is gained by subtracting from this on-line image a complementary image that specifically avoids the chemical's absorption lines and passes instead the background radiance between the lines (pattern A). Since the light rejected by the micro-mirror array is not actually lost, but merely redirected, it is possible in principle to collect any pattern and its complement simultaneously by using two detector arrays instead of one.

2.2. Hadamard Transform Hyperspectral Imaging

The chemical-specific matched filter imaging discussed above is just one possible class of algorithms that can be implemented on the micro-mirror array. While the matched filter scheme offers simplicity and reduced data volume, it does throw away spectral information. If full-spectrum hyperspectral imaging is desired instead, this instrument concept can accomplish this as well, at considerably higher signal-to-noise ratios than a standard single-slit spectrometer, by programming the micro-mirrors to implement the Hadamard functions in the Hadamard transform technique. The Hadamard functions are a complete orthonormal set of binary functions, closely related to Walsh functions, that span the spectral space and thus can form the basis for a transform-and-inverse spectral encoding scheme somewhat analogous to the Fourier transform. The advantage lies in the fact that the Hadamard technique presents the detector with light from many spectral channels at once, greatly increasing the signal level; for an n -spectral-band system, each Hadamard function presents the detector with approximately $n/2$ times more light than does a simple slit spectrometer. This “multiplex” advantage can lead to considerable gains in signal-to-noise in the regime where detector (and electronics) noise is dominant over photon shot noise. This is the regime, for instance, appropriate for infrared imaging using uncooled microbolometer detector arrays. Indeed, it was the possibility of making uncooled microbolometer arrays practical for long-wave infrared (LWIR) hyperspectral imaging that first motivated us to experiment with this multiplexing scheme.

To understand how multiplexing works, and appreciate how it can improve the signal-to-noise ratio, let us consider a generalized multiplexing spectrometer. In general, a multiplexing instrument employs some sort of optical device — an interferometer, dispersing element, mask, etc. — to modulate the intensities of the various spectral components of the incoming signal in some well-understood and reproducible manner, and then, following this modulation, combines all these wavelengths and presents them to the detector simultaneously. (Alternatively, one can construct a spatially-multiplexing instrument, in which the optics apply a modulation to light arriving from various spatial locations, then combines and presents light from all these sources simultaneously to the detector. We consider just the spectrally-multiplexed case here.) For the m^{th} configuration of the optical modulator, the intensity of the light arriving at the detector element is

$$I_m = \sum_{i=1}^n a_{m,i} S(\lambda_i) \quad (1)$$

where n is the number of spectral bands, and $S(\lambda_i)$ is the spectrum of incoming light — i.e., the brightness S at the i^{th} wavelength λ_i . The coefficients $a_{m,i}$ embody the details of the modulation scheme. They might be ones and zeroes in the case of a mask-based modulation scheme, or they could be cosines or more complex functions in the case of interferometer-based modulations. In any case, the optical modulation scheme and its associated coefficients are chosen so that, after measuring the complete set of multiplexed intensities, I_1, I_2, \dots, I_M , there is sufficient information to invert the transformation and solve for the spectral intensities $S(\lambda_i)$. Since the M equations represented by Eq. (1) must be solved for the n unknowns, $S(\lambda_i)$, it follows that M must be greater than or equal to n . The inverse transform will take the general form

$$S(\lambda_i) = \sum_{m=1}^M b_{i,m} I_m \quad (2)$$

and it is to be understood that the inversion coefficients $b_{i,m}$ are known *a priori* through the mathematics of the relevant transformation.

Perhaps the most familiar type of wavelength-multiplexing spectrometer is the Fourier transform spectrometer (FTS), in which a scanning-mirror interferometer, usually a Michelson interferometer, is used to modulate each wavelength component of the incoming light in a sinusoidal fashion, the frequency of the modulation as a function of mirror position being proportional to the optical frequency of the light (or inversely proportional to the wavelength). In this case, for an instrument recording single-sided symmetric interferograms, both the forward transform coefficients $a_{m,i}$ and inversion coefficients $b_{i,m}$ in Eq. (1) and Eq. (2) are cosine functions in the m^{th} mirror position x_m divided by the i^{th} wavelength. I.e.,

$$a_{m,i} = 1 + \cos \frac{2\pi x_m}{\lambda_i} = 1 + \cos 2\pi x_m \tilde{\nu}_i \quad (3)$$

$$b_{i,m} = 4 \cos \frac{2\pi x_m}{\lambda_i} = 4 \cos 2\pi x_m \tilde{\nu}_i \quad (4)$$

which, when inserted into Equations (1) and (2), produce the desired discrete Fourier transform and its inverse.

While the FTS may be the most familiar, another class of multiplexing spectrometers, with comparable multiplex advantages, can be constructed using coded masks in the spectral plane of a dispersive spectrometer.

Let us now consider noise propagation in a simple mask-based multiplexing instrument, where the spectral masks can be configured to either pass or reject each of seven spectral bands. The scene brightnesses for these seven bands are $S(\lambda_1)$, $S(\lambda_2)$, . . . $S(\lambda_7)$. Suppose we choose for the first configuration a mask that passes only the bands λ_1 , λ_2 , λ_3 , and λ_5 and rejects the rest. Combined onto a single detector element, the total intensity at the detector for this first configuration is

$$I_1 = S(\lambda_1) + S(\lambda_2) + S(\lambda_3) + S(\lambda_5). \quad (5)$$

For the next configuration we might choose to pass just λ_1 , λ_2 , λ_4 , and λ_7 , rejecting the other bands, so that

$$I_2 = S(\lambda_1) + S(\lambda_2) + S(\lambda_4) + S(\lambda_7). \quad (6)$$

Note that, assuming the intensities in the seven bands are not too different, the intensity at the detector for each of these cases is approximately four times what it would have been had this been a conventional, non-multiplexing spectrometer.

Proceeding in the same manner, we can implement a full set of seven configurations, soluble for the spectral intensities $S(\lambda_1)$, $S(\lambda_2)$, . . . $S(\lambda_7)$. The complete set can be written in matrix form as follows:

$$\begin{pmatrix} I_1 \\ I_2 \\ I_3 \\ I_4 \\ I_5 \\ I_6 \\ I_7 \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 1 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} S(\lambda_1) \\ S(\lambda_2) \\ S(\lambda_3) \\ S(\lambda_4) \\ S(\lambda_5) \\ S(\lambda_6) \\ S(\lambda_7) \end{pmatrix} \quad (7)$$

This system of seven equations can easily be solved for the spectral intensities $S(\lambda_1)$, $S(\lambda_2)$, . . . $S(\lambda_7)$, with the solution

$$\begin{pmatrix} S(\lambda_1) \\ S(\lambda_2) \\ S(\lambda_3) \\ S(\lambda_4) \\ S(\lambda_5) \\ S(\lambda_6) \\ S(\lambda_7) \end{pmatrix} = \frac{1}{4} \begin{pmatrix} 1 & 1 & 1 & -1 & 1 & -1 & -1 \\ 1 & 1 & -1 & 1 & -1 & -1 & 1 \\ 1 & -1 & 1 & -1 & -1 & 1 & 1 \\ -1 & 1 & -1 & -1 & 1 & 1 & 1 \\ 1 & -1 & -1 & 1 & 1 & 1 & -1 \\ -1 & -1 & 1 & 1 & 1 & -1 & 1 \\ -1 & 1 & 1 & 1 & -1 & 1 & -1 \end{pmatrix} \begin{pmatrix} I_1 \\ I_2 \\ I_3 \\ I_4 \\ I_5 \\ I_6 \\ I_7 \end{pmatrix} \quad (8)$$

In other words, the coefficients $b_{i,m}$ in Eq. (2) are all either $+1/4$ or $-1/4$ for this case. By measuring the seven multiplexed intensities I_m we can determine the seven spectral intensities $S(\lambda_i)$.

Now suppose that the total noise is completely dominated by detector noise. Detector noise is a simple additive uncertainty that doesn't depend on the signal amplitude. Therefore, any measurement of an intensity I will have an additive error whose root mean square value is σ_{Detector} .

Following standard error propagation theory for random Gaussian noise, the uncertainty σ_S in the individual spectral intensities $S(\lambda_i)$ is given by

$$\sigma_S^2 = \sum_{m=1}^M \left(\frac{\partial S(\lambda_i)}{\partial I_m} \right)^2 \sigma_{I_m}^2 = \sum_{m=1}^M \left(\frac{\partial S(\lambda_i)}{\partial I_m} \right)^2 \sigma_{\text{Detector}}^2 \quad (9)$$

Examining Eq. (8), we find that for every I_m and every $S(\lambda_i)$ the squared partial derivative is always

$$\left(\frac{\partial S(\lambda_i)}{\partial I_m} \right)^2 = \frac{1}{4^2} \quad (10)$$

so that, summing over all seven configurations, the total uncertainty in the spectral intensity (i.e. the noise in the computed spectrum) is

$$\sigma_S = \sigma_{\text{Detector}} \frac{\sqrt{7}}{4} \approx (0.661) \sigma_{\text{Detector}} \quad (11)$$

Had we instead measured the $S(\lambda_i)$ values individually in a non-multiplexing instrument, their uncertainty would have been simply σ_{Detector} . Thus by performing this seven-band multiplexing scheme, we have reduced the noise in the quantities of interest by a factor of roughly 0.66 or increased the signal-to-noise ratio by roughly 51 percent.

The foregoing was actually a seven-band example of the Hadamard transform technique, a complete treatment of which is given in Harwit and Sloane's classic text.¹ Here we summarize some of the important points. In general, a Hadamard matrix of order n is an $n \times n$ square matrix of $+1$'s and -1 's, for which the scalar product of any two distinct rows is zero (i.e., all rows are orthogonal). The rows of a Hadamard matrix thus form an orthogonal set of basis functions that spans the spectral space of interest. A Hadamard transform is performed by multiplying the input spectral vector (i.e. the spectrum of the scene) by the appropriate order Hadamard matrix. Physically implementing the $+1$'s and -1 's in

hardware is possible in principle with a mirror-based masking scheme, but generally would require two detector arrays, one to collect the +1 contributions, the other for the -1's. However, associated with each Hadamard matrix is the related $(n-1) \times (n-1)$ S-matrix, which is formed by omitting the first row and column, then replacing all the +1's with zeroes, and all the -1's with +1's. The rows of the S-matrix, while not orthogonal, still form a basis set that spans the space of interest, so that an S-matrix transform can be performed. The S-matrix is easily implemented in hardware using a single detector array, and corresponds to the case of a simple mask with a given position either open or closed. The matrix in Eq. (7) is an example of a Hadamard S-matrix. Roughly speaking, about half the elements of any given row of a Hadamard matrix will be -1's and half will be +1's. Similarly, roughly half the elements of an S-matrix will be zeroes.

To invert the S-matrix transform is a simple matter. The inverse of an $n \times n$ S-matrix is formed taking the transpose of the S-matrix, replacing the zeroes by -1's, and multiplying the result by $2/(n+1)$. I.e.,

$$S_n^{-1} = \frac{2}{n+1} (2S_n^T - J_n) \quad (12)$$

where J_n is an $n \times n$ matrix consisting entirely of +1's. Eq. (8) illustrated this inversion for $n = 7$.

It can be shown, via the same arguments we used in Eq. (9) through Eq. (11), that the signal-to-noise improvement of an n -band S-matrix-based Hadamard transform spectrometer over an otherwise equivalent single-slit spectrometer, in the detector noise dominated regime, is

$$\frac{SNR_{S-matrix}}{SNR_{slit}} = \frac{n+1}{2\sqrt{n}} \approx \frac{\sqrt{n}}{2} \text{ for large } n. \quad (13)$$

Similarly, the improvement for a dual-detector instrument implementing the +1's and -1's of the full orthogonal Hadamard transform can be shown to be

$$\frac{SNR_{Hadamard}}{SNR_{slit}} = \sqrt{n} \quad (14)$$

The difference between the two arises because the full Hadamard uses all the available light, whereas the S-matrix implementation throws away half the light with its "closed" mask positions.

A few other aspects of Hadamard and S-matrices are worth noting. For a given order n , the Hadamard and S-matrices are in general not unique. Many realizations satisfying the definitions of Hadamard and S-matrices are possible. On the other hand, Hadamard and S-matrices do not exist for all n , only for certain values of n . Fortunately the set of orders n for which Hadamard and S-matrices can be constructed is quite dense, and a matrix close to the desired size can usually be found. Of the possible realizations of S-matrices, of particular practical interest are those, such as the matrix in Eq. (7), whose rows are cyclic permutations of each other. Cyclic S-matrices are particularly simple and convenient to implement in hardware, either as an actual physical mask, or, in our case, implemented on a micro-mirror array.

The signal-to-noise advantages of the Hadamard technique have been recognized for decades, but early implementations were limited by the difficulties in producing and precisely positioning the Hadamard coded masks. Small errors in either the manufacture or positioning of the masks can lead to unacceptably large spectral artifacts. The advent of micro-mirror arrays has made Hadamard transform spectroscopy much more practical, by providing a means for realizing highly reproducible mask patterns, and eliminating the need to physically translate (with high precision) physically cut-out masks. The potential of micro-mirror arrays for Hadamard spectroscopy has been recognized and explored, both theoretically and experimentally, by a number of groups over the last several years.²⁻⁵ In each case, the experimental goals drive the designs, leading to considerable variety. Implementations described range from an all-reflective dual Offner design⁴ that implements the basic scheme illustrated in Fig. 1, to designs in which the Hadamard transform is applied to the spatial, rather than the spectral, portion of the data, with the spectroscopy performed downstream by a conventional monochromator.⁵ The goal of our work here has been to produce a very versatile

instrument, capable of switching between Hadamard and matched- filter spectral imaging, and capable of performing all necessary spatial scanning within the micro-mirror array itself without resorting to external scan mirrors or platform push-broom motion.

3. THE WORKING PROTOTYPE INSTRUMENT

In order to demonstrate the viability this concept, we have constructed a prototype multiplexing micro-mirror-based hyperspectral imager using the commercial off-the-shelf DLP micro-mirror array produced by Texas Instruments. This array, consisting of 1024 x 768 individually addressable micro-mirrors with a 13.68 μm square pixel pitch, operates the micro-mirrors in a bi-stable mode: each micro-mirror can be tilted either $+12^\circ$ or -12° out of the device plane, with the tilt occurring about the micro-mirror diagonal. In a functioning instrument, individual micro-mirrors can be turned ‘on’ by tilting them in the direction that routes the incoming light toward the detector, and turned ‘off’ by tilting them the opposite direction so that light is routed away from the detector.

With micro-mirror dimensions this small, it is readily apparent that diffraction effects will play an important role in the interaction of the light with the micro-mirror array. An intuitive understanding of the role of diffraction can be gained by noting that the micro-mirror array is, in essence, a two-dimensional diffraction grating. The angular directions of the various diffraction orders is determined solely by the 13.68 μm spacing of the micro-mirror grid, just as the diffraction angles of a grating are determined by the groove spacing. Which of these diffraction orders will receive appreciable amounts of energy is determined by the 12° tilt of the micro-mirrors, in direct analogy with the blaze angle of a conventional diffraction grating. With all the micro-mirrors tilted the same direction, to the ‘on’ position for instance, one can think of the array as consisting of diagonal “grooves,” each with a “blaze” of 12 degrees. The effective groove spacing is one-half the diagonal dimension of the micro-mirrors, i.e.,

$$d_{\text{eff}} = \frac{13.68\mu\text{m}}{\sqrt{2}} = 9.67\mu\text{m} \quad (15)$$

Inserting this groove spacing in the standard grating equation for normal incidence

$$m\lambda = d \sin \theta \quad (16)$$

where m is the diffraction order, λ is the wavelength, and d the groove spacing, one finds, for example, that the preferred 24° diffraction angle θ , as dictated by the 12° “blaze” of the micro-mirror tilt, occurs close to 6th order for the HeNe laser wavelength of 632.8 nm. This is confirmed experimentally, as shown in Fig. 4.

One also finds that the longest wavelength for which there is a diffraction order at 24 degrees is 3.93 μm . Beyond this wavelength, the reflected/diffracted efficiency drops off rapidly as the available diffraction orders fall farther and farther from the ideal 24-degrees along-the-diagonal direction. Beyond 9.67 μm , Eq. (16) implies that no diffraction order exists along the diagonal direction, since this would require $\sin \theta$ to be greater than unity. There are diffraction spots for 9.67 μm at roughly 45 degrees, not along the diagonal, but along the rectangular axes (i.e., with $d_{\text{eff}} = 13.68 \mu\text{m}$; see the diffraction pattern in Fig. 4), but these are far from the maximum efficiency micro-mirror reflection direction.

These long-wavelength limitations can be mitigated to some extent by programming lower spatial frequency patterns onto the DLP. For example every other diagonal row might be turned on, to produce a “grating” with $d_{\text{eff}} = 19.24 \mu\text{m}$, or every third row, for $d_{\text{eff}} = 29 \mu\text{m}$. This provides diffraction orders where needed, but at the expense of throwing away half or two-thirds, respectively, of the available light.

Thus the commercial DLP array, simply by virtue of its small micro-mirror size, is less than ideally suited for use in the infrared at wavelengths much beyond 3.93 μm , particularly if one is primarily interested in using these devices to maximize signal-to-noise. Furthermore, these arrays are supplied commercially with glass windows that pass only near-IR and shorter wavelengths, so these must be replaced with an appropriate IR material.




Figure 4. Photograph showing the diffraction pattern generated by the Texas Instruments DLP micro-mirror array, illuminated by a HeNe laser (wavelength 632.8 nm), with all the micro-mirrors turned to the same orientation (e.g., the “on” state). A weak zero-order reflection is generated by the dead space between micro-mirrors, providing a useful reference. The back side of the circuit board containing the micro-mirror array can be seen at left; the micro-mirror array is facing away from the viewer in the photo. The HeNe laser beam travels between the two cardboard screens, just to the left of the zero-order reflection, and hits the array at approximately normal incidence. Most of the reflected energy (approximately 87%) goes to the bright spot at the sixth diffraction order in both the horizontal and vertical directions, corresponding to the 24-degree micro-mirror reflection angle.

Despite its unsuitability for LWIR work, the DLP array is the most readily available and user-friendly micro-mirror array. We therefore chose to use the DLP array and build our initial prototype to work in the Visible/Near-UV region. Historically, our primary interest has been in remote sensing of gases, and our choice of spectral range here, 300 nm to 450 nm, provides access to several highly structured gas absorption bands, including SO₂, NO_x, and ozone.

Our criteria for the prototype design were as follows: (1) The instrument should be capable of spectrally resolving the most closely-spaced gas lines of interest, namely those of SO₂, whose absorption lines between 300 nm and 315 nm are spaced by roughly 2.1 nm. (2) The optics for the prototype should be relatively inexpensive and quick to produce, preferably off the shelf. (3) Although the micro-mirror array instrument concept naturally lends itself to extremely compact designs, compactness would not be the primary consideration for the prototype design, thus eliminating such complications as non-essential fold mirrors.

The prototype optical design, arrived at with those criteria in mind, closely mimics the schematic concept outlined in Fig. 2. It uses inexpensive off-the-shelf transmission gratings (manufactured by ThorLabs) both to initially disperse the incoming light and to recombine it after interaction with the micro-mirror array. The remainder of the optics consists of custom-specified (but still relatively inexpensive) spherical lenses of calcium fluoride and UV-grade fused silica, assembled into various achromatic combinations. The detector is a UV-sensitive CCD array manufactured by Sony, with 1280 x 960 pixels with a 4.65µm pixel pitch. The assembled prototype is illustrated in Figures 5 and 6. What is not shown in the cartoon version of Fig. 2, is the addition of a couple of optical elements that are both substantially tilted and off-center with respect to the system’s optical axis. The need for these less-than-elegant design features arises from problem that, in forming the final image, the MMA becomes a severely tilted object plane, tilted 24 degrees off-normal with respect to the light bundle leaving the array after reflecting from the micro-mirrors. These tilted elements, as well as the tilt of the final image plane on the CCD camera, can be seen in Fig. 5. Though somewhat unattractive, this strategy results in good quality (if not perfect) images, and allows us to maintain our low-cost design philosophy.




Figure 5. Photograph of the prototype programmable micro-mirror-array spectral imager. As in Fig. 2, light from the scene is dispersed by Grating 1 and imaged onto the micro-mirror array, which selects certain wavelengths and rejects others to implement an encoding scheme. The selected wavelengths are recombined by Grating 2, and imaged onto the detector array, in this case a Sony UV-sensitive CCD.




Figure 6. Another view of the prototype micro-mirror-array spectral imager. In this photograph, the micro-mirror array is visible in the upper right section of the optical assembly. The pattern displayed on the micro-mirror array mirrors that on the computer display in the background. The 45-degree tilt of the MMA is a consequence of the along-diagonal tilt orientation of the micro-mirrors.

One of the challenges in designing an instrument around the TI DLP array is the fact that the micro-mirrors in the DLP array tilt along the mirror diagonals. This results in the optical deflections being in a direction 45 degrees with respect to the rectangular axes of the array, as can be seen in Fig. 4. Therefore, in order for the instrument to have its optical components lie roughly in a horizontal plane, as is most convenient for a bread-board prototype, the DLP array must be rotated 45 degrees. It follows that the gratings must also be similarly rotated, so that the dispersion direction lies along the DLP axis. The resulting configuration, somewhat awkward but still reasonably workable, can be seen in Fig. 6.

4. RESULTS

The prototype instrument was tested in the laboratory using various light sources and materials to provide a variety of spectral signatures in the scene. A broad-band image of a typical test scene, obtained with the prototype instrument, is shown in Fig. 7. This test scene consists of a large wall calendar, illuminated by two light sources. The first of these is a mercury-vapor lamp visible near the center of the image. The Hg-vapor lamp is shielded by a makeshift aluminum foil shade that blocks the direct view of the lamp as seen from the instrument, with the exception of a small pinhole visible as a bright dot near the center of the image.




Figure 7. Broad-band image of a typical test scene obtained with the prototype micro-mirror instrument. The scene consists of a wall calendar illuminated by both a mercury-vapor lamp, visible near center, and an incandescent lamp located off the left side of the image. The Hg lamp provides a sharp line spectrum consisting of four lines in the instrument's spectral range. The shortest wavelength of these lines is strongly absorbed by the calendar's glossy coating. A sheet of uncoated white paper, which reflects all four Hg lines, is also included in the scene. The 45-degree rotation of the scene follows from the orientation of the micro-mirror array, which in turn derives from the along-diagonal tilt direction of the micro-mirrors.

The Hg lamp has four bright, sharp lines in the instrument's spectral range, at 313 nm, 365 nm, 404.7 nm, and 435.8 nm. The second light source is an ordinary incandescent lamp located just off the left side of the image. Illumination from the incandescent lamp is spectrally continuous, but still contains some broad spectral features. Two types of paper form the background. The calendar itself has a glossy coating that strongly absorbs the 313 nm line of the Hg lamp. Taped to this is a sheet of uncoated white paper, which strongly reflects all four lines. Thus, taken together, this test scene produces a fairly wide range of spectral complexity, comparable to many real-world outdoor scenes. The main thing to note in Fig. 7 is that, despite the complexities of spectrally dispersing and recombining the light in a system with off-normal and off-axis elements, and our self-imposed constraint of using only spherical lenses, the instrument provides reasonably sharp imaging across the field of view.

Figure 8 illustrates an implementation of a spectral matched filter for Hg-vapor emission. Here the micro-mirror array is programmed first, in the left frame, to reject all four of the Hg lamp emission lines (313.365, 404.7, and 435.8 nm), then, on the right, to pass just those four lines while rejecting all other wavelengths. Thus we see in left frame a scene illuminated by the broad-band incandescent lamp, while the pinhole in the Hg lamp shade just barely visible (due to some continuum emission between the sharp lines). In the right-hand frame, the result is completely dominated by the bright Hg emission emerging from the pinhole. The MMA is acting as Hg-vapor emission detector.




Figure 8. Demonstration of spectral matched-filter imaging using the programmable micro-mirror-array spectrometer prototype. Left image shows the Hg lamp scene with the micro-mirror array programmed to reject the Hg lamp emission lines; only a small amount of continuum emission permits the pinhole in the Hg lamp’s shade to be faintly visible (just to the right of “August 4” on the background calendar). The image on the right shows the same scene, under the same conditions, with the micro-mirror array programmed to pass the four Hg-vapor emission lines (313.365, 404.7, and 435.8 nm) while rejecting all other wavelengths. The result is an image dominated by the extremely bright Hg lamp emission emerging from the pinhole in the Hg lamp shade.

Next, the prototype instrument was programmed to perform full-spectrum hyperspectral measurements using the Hadamard transform technique. To verify the predicted signal-to-noise advantage of the Hadamard method, the instrument was set up to operate in two different modes, with all other variables kept constant. First, the instrument was programmed to simulate a traditional single-slit spectrometer. This is accomplished simply by turning on a column of micro-mirrors (or, to simulate a wider slit, multiple adjacent columns of micro-mirrors can be set to the ‘on’ position). To scan across the scene, pushbroom-style, the “slit” thus formed is moved by sequentially flipping the appropriate micro-mirrors, moving the simulated slit across the MMA. Next, the Hadamard technique is implemented on the MMA. An 83-band implementation was chosen for this initial test because this number of spectral bands corresponds almost exactly to a width of eight micro-mirrors for each “slit” in the Hadamard mask pattern, i.e. eight micro-mirrors per “Hadamard zone.” (The “slit” width for the single-slit simulation was therefore set to be eight micro-mirrors as well.) Because the Hadamard S-matrix implemented on the MMA is of the cyclic variety, collecting a complete set of Hadamard functions for all spatial locations is accomplished simply by translating the pattern, one Hadamard zone width per step, 83 times. The spectra for each location are then obtained by inverting the Hadamard transform (taking into account which cyclic permutation corresponds to the given spatial location).

Sample frames from the two methods are shown in Fig. 9. What is immediately apparent in this comparison is the greatly increased signal level seen in the Hadamard case. Here the Hg lamp illumination, with its four bright spectral lines, makes the Hadamard pattern readily apparent. Depending on the location in the scene, the illumination consists of various combinations of one, two, three, or all four Hg lines.




Figure 9. Comparison of signal levels obtained for a single slit spectral imager, as simulated using the micro-mirror array (left), and for an 83-band Hadamard transform spectral imager implemented on the micro-mirror array (right). Light levels, integration times, spectral bin widths, and all other parameters are identical for two cases.

The spectral results for the two methods, for an representative spatial pixel, are compared in Fig. 10. The detector gain and integration time in this experiment were deliberately set so as to provide a prominent detector noise contribution, to better show the differences between the Hadamard and single-slit methods; much better SNR is possible with more optimal detector parameter settings.




Figure 10. Weak-signal performance comparison of a conventional single-slit spectrometer and an 83-band micro-mirror Hadamard transform, both implemented on the prototype programmable micro-mirror spectral imager (the “slit” simulated here by turning on a slit-shaped region of the micro-mirror array). Scene is illuminated with a Hg-vapor lamp having four bright spectral lines in this region. All four are easily seen in the Hadamard spectrum, whereas the non-multiplexed slit spectrum barely shows the strongest line. Spectral resolution, detector gain, integration time, and illumination levels are identical for the two data sets.

In this detector-noise-dominated case, it can be seen in Fig. 10 that the single-slit method barely provides enough SNR to detect the Hg vapor lines at all; only the strongest line, at 435.8 nm, is visible, with a signal-to-noise of roughly unity.

With the Hadamard technique, however, all four lines are clearly visible above the noise. The predicted SNR improvement, from Eq. (13), is $\sqrt{83}/2$ or about 4.6, consistent with Fig. 10. Detailed comparison of many pixels verifies the theoretical prediction.

Finally, Figure 11 summarizes a complete hyperspectral dataset obtained using the Hadamard technique. This is a true hyperspectral image cube, with full spectra comprising all 83 spectral bands recorded for all spatial location. In the center of the Fig. 11 is shown the broad-band image, obtained by summing over all spectral bands. Fig. 11 (a-d) shows spectra obtained for various spatial locations, highlighting some of the spectral diversity in the scene. Note, for example, the difference between the glossy and plain paper in Fig. 11 (a) and (c), with the 313 nm Hg line almost entirely missing from the glossy calendar's spectrum, but strongly present in that of the plain paper. Considerable broad spectral structure is seen in the glint from the incandescent lamp in Fig. 11(b), together with a substantial contribution from the 435.8 nm Hg line. This broad spectral structure from the incandescent lamp can also be seen as a weak background in Fig. 11(c).

The choice of 83 spectral bands in this Hadamard example was merely one of convenience. Given the pixel dimensions of the micro-mirror array, and the optical quality of the imaging, Hadamard imaging with several hundred spectral bands should be possible in principle.




Figure 11. Hadamard hyperspectral image from the micro-mirror spectral imager. Image in center is a broad-band image obtained with the prototype instrument. The scene is a large wall calendar printed on glossy coated paper, a page of ordinary white paper, all illuminated by both a mercury-vapor lamp (seen with its shade in silhouette near center) and an incandescent lamp. Example spectra (a-d) from this hyperspectral data cube show a variety of features. Note that the 313 nm line of the Hg lamp is absorbed by the glossy coating in (a) but not by the plain paper in (c). Spectra (b) and (d) are dominated by the incandescent lamp.

5. SUMMARY

We have built and demonstrated a working prototype of a micro-mirror-based programmable hyperspectral imager, capable of both chemical-specific matched filter imaging, and full-spectrum hyperspectral Hadamard transform imaging with multiplex-enhanced SNR. The signal-to-noise improvement predicted for the Hadamard transform technique in the noisy-detector regime was experimentally verified with this prototype. With MMA-based matched-filter imaging,

detection of a predetermined set of chemicals or materials can be accomplished in real time, without laborious post-processing. Unlike traditional interference-filter-based approaches to multi-band detection of pre-defined chemicals, the micro-mirror array instrument offers near-instantaneous reconfigurability in the field, with the target chemical changeable simply by uploading a new matched filter pattern to the MMA. This unprecedented flexibility makes this new instrument concept attractive for wide range of applications. In either the Hadamard or matched-filter mode, the MMA-based instrument described here offers the ability to perform the spatial scan entirely within the MMA itself, without resorting to external moving mirrors or platform motion, a considerable advantage for compact, field-portable instrumentation. Though this particular instrument was designed to work in the blue to near-UV spectral range, the idea can easily be extended to the mid-infrared using current commercial Texas Instruments DLP arrays (if the DLP window material is changed to an appropriate IR-transparent material). With the 13.68 μm micro-mirror pixel pitch of this DLP array, however, diffraction places a more fundamental long-wavelength limit at 3.93 μm , beyond which the device's excellent reflected/diffracted efficiency of roughly 90 percent falls off rapidly because diffraction orders no longer fall near the preferred "blaze" angle defined by the micro-mirror tilt. While the use of lower spatial frequency patterns on the DLP can produce diffraction orders where needed at longer wavelengths, this method throws away much-needed light. Extending these techniques to longer wavelengths (particularly to the important 8-12 μm LWIR atmospheric window), while maintaining a high enough throughput for the overall sensitivity to be competitive with other techniques, will therefore require MMAs with proportionally larger pixel pitches. This extension would be of particular interest, however, because the Hadamard SNR enhancement would be of special value in producing hyperspectral imagers using uncooled microbolometer arrays (which fall firmly into the noisy-detector regime). This combination of MMA-based Hadamard transform spectroscopy and uncooled microbolometer detectors could lead to a new generation of compact, low-power hyperspectral imager for the thermal infrared.

ACKNOWLEDGMENTS

The author gratefully acknowledges Mark Lyon of LANL for his work on developing the software interface to the TI DLP array.

REFERENCES

- [1] Harwit, M. and Sloane, N. J. A., *Hadamard Transform Optics*, (Academic, New York, 1979).
- [2] Wuttig, A. and Riesenberger, R., "Sensitive Hadamard transform imaging spectrometer with a simple MEMS," Proc. SPIE 4881, 167-178 (2003).
- [3] Hanf, M., Schaporin, A., Hahn, R., Dotzel, W., and Gessner, T., "Novel Hadamard transform spectrometer realized using a dynamically driven micro-mirror array as light modulator," Proc. SPIE 5717, 117-126 (2005).
- [4] Wehlburg, C. M., Wehlburg, J. C., Gentry, S. M., and Smith, J., "Optimization and characterization of an imaging Hadamard spectrometer," Proc. SPIE 4381, 506-515 (2001).
- [5] DeVerse, R. A., Hammader, R. M., and Fateley, W. G., "Hadamard transform Raman imagery with a digital micro-mirror array," *Vibrational Spectroscopy* 19, 177-186 (1999).