Steven P. Love and David L. Graff, "Full-frame programmable spectral filters based on micromirror arrays," *J. Micro/Nanolith. MEMS MOEMS* **13(1)**, 011108 (2014).

http://dx.doi.org/10.1117/1.JMM.13.1.011108

Journal of
Micro/Nanolithography, MEMS, and MOEMS

SPIEDigitalLibrary.org/jm3

Full-frame programmable spectral filters based on micromirror arrays

Steven P. Love
David L. Graff

SPIE

Full-frame programmable spectral filters based on micromirror arrays

Steven P. Love
David L. Graff
Space and Remote Sensing Group (ISR-2)
Los Alamos National Laboratory
Mail Stop B244
Los Alamos, New Mexico 87544
E-mail: splove@lanl.gov

Abstract. Rapidly programmable micromirror arrays, such as the Texas Instruments Digital Light Processor (DLP®) digital micromirror device (DMD), have opened an exciting new arena in spectral imaging: rapidly reprogrammable, high spectral resolution, multiband spectral filters that perform spectral processing directly in the optical hardware. Such a device is created by placing a DMD at the spectral plane of an imaging spectrometer and by using it as a spectral selector that passes some wavelengths down the optical train to the final image and rejects others. Although simple in concept, realizing a truly practical DMD-based spectral filter has proved challenging. Versions described to date have been limited by the intertwining of image position and spectral propagation direction common to most imaging spectrometers, reducing these instruments to line-by-line scanning imagers rather than true spectral cameras that collect entire two-dimensional (2-D) images at once. Here, we report several optical innovations that overcome this limitation and allow us to construct full-frame programmable filters that spectrally manipulate every pixel, simultaneously and without spectral shifts, across a full 2-D image. So far, our prototype, which can be programmed either as a matched-filter imager for specific target materials or as a fully hyperspectral multiplexing Hadamard transform imager, has demonstrated over 100 programmable spectral bands while maintaining good spatial image quality. We discuss how diffraction-mediated trades between spatial and spectral resolution determine achievable performance. Finally, we describe methods for dealing with the DLP's 2-D diffractive effects and suggest a simple modification to the DLPs that would eliminate their impact for this application. [DOI: 10.1117/1.JMM.13.1.011108]

Subject terms: spectral imager; micromirror array; digital micromirror device; Hadamard; hyperspectral; matched filter; programmable spectral imager.

Paper 13127SSP received Jul. 15, 2013; revised manuscript received Oct. 18, 2013; accepted for publication Nov. 5, 2013; published online Jan. 2, 2014.

1 Introduction

Spectral imaging has proven indispensible for remote detection, identification, and quantification of chemical species, in applications ranging from reconnaissance to environmental monitoring, planetary science, and medical imaging. Among spectral imaging strategies, hyperspectral imaging (HSI), in which each pixel contains a high-resolution spectrum consisting of dozens to hundreds of contiguous spectral bands, offers great versatility, particularly for applications in which specific target species are not known *a priori* and one is simply searching for out-of-the-ordinary spectral signatures which can be identified later. As currently implemented, however, HSI is inherently slow, typically requiring several seconds to obtain a single spatial/spectral image cube. Monitoring phenomena occurring on timescales faster than this has necessitated sacrificing either spectral specificity (e.g., reducing the spectral component to a simple bandpass filter in front of a fast camera) or sacrificing spatial information and simply aiming a spectrometer at a point (or, at most, a line).

Conventional HSI instruments in use today typically take one of two forms. One type incorporates a dispersive (grating or prism) spectrometer, in which case only a single line of the spatial scene is imaged at any given time; one dimension of the two-dimensional (2-D) detector array is devoted to spectral information and the other to spatial, with the other spatial dimensions acquired by scanning the instrument across the scene (i.e., a "push-broom imager"). The other common type uses a Fourier-transform spectrometer, wherein each frame contains a full 2-D spatial scene, modulated in intensity by an interferometer; spectral information, in the form of an interferogram, is collected over time as a moving mirror varies the interferometer optical path difference. Either way, several hundred frames must be acquired to build up the full HSI data cube, a process typically requiring several seconds per cube, putting even moderately fast phenomena—evolving gas plumes, real-time *in vivo* biomedical signatures, combustion, explosions, etc.—out of reach of standard HSI. Likewise, imaging a large area, as is required for example in a remote-sensing broad-area search for a rare type of target, becomes prohibitively time consuming, particularly when time over the target area is limited. Moreover, full-spectrum HSI creates large volumes of data, making telemetric down-linking of data problematic. Finally, the HSI data must undergo computationally intensive processing in order to extract meaningful results, further extending the turnaround time between data acquisition and target detection.

0091-3286/2014/$25.00 © 2014 SPIE

In many applications, however, specific target species are known in advance, making the contiguous full-spectrum capability of traditional HSI unnecessary, and the potential speed can increase greatly. As a simple example, if one were fortunate enough to have a target characterized by a single strong spectral feature distinct from the background, a simple band-pass filter/camera combination could conduct a search at a rate limited only by the camera frame rate, hundreds of times faster than full hyperspectral coverage.

Most real-world applications fall between these extremes: distinguishing a predefined target of interest from potential imposters or background-clutter false alarms usually requires considerably more spectral detail than a single band but still considerably less information than a full hyperspectral image. The aim of this work is an optimal compromise: collecting enough spectral information to reliably distinguish the target of interest against a spectrally cluttered background while maintaining the speed and simplicity advantages of a simple filter. Digital micromirror devices (DMDs), as well as other addressable spatial light modulators such as microshutter arrays and LCD arrays, offer a means for implementing highly specific, spectrally complex filters having both high spectral resolution and rapid reconfigurability. Such programmable filters open a new paradigm for spectral imaging, in which much of the spectral processing necessary to produce the final data product takes place essentially instantaneously, directly in the optical hardware, rather than in postprocessing, thereby reducing both acquisition time and data volume.

Over the past several years, a number of fundamentally similar schemes for producing a DMD-based programmable spectral filter have been proposed and demonstrated.[1–10] But all of these designs share a similar shortcoming: if capable of imaging at all, they remain push-broom imagers, still requiring a time-consuming scan either of a one-dimensional single-line field of view, or, in some cases, of the spectral pattern programmed onto the DMD, in order to build up a 2-D spatial image. This article describes the design and performance of a new type of multiband programmable filter that truly behaves like a filter, allowing full-frame 2-D spatial images to be obtained in a single shot without scanning, with the same spectral band selection applied across the entire frame.

2 Basic Principles

2.1 *Previous Approaches to DMD-Based Programmable Spectral Filters*

Prior to this work, the basic design of most DMD-based programmable spectral filters proceeded as follows: light from the scene is collimated and presented to the dispersive element, usually a diffraction grating. Often, but not always, a slit is used to restrict the field-of-view to a thin line of the scene; in this case, the scene is first imaged onto the slit plane, then the light emerging from the slit is collimated and directed to the grating. The spectrally dispersed light leaving the grating is collected and focused by an imaging optic to form a spectrally dispersed image of the scene (as cropped by the slit if one is included) at the plane of the DMD. The DMD is then used as a multiband spectral selector, with columns of micromirrors programmed to the "on" orientation to direct desired wavelengths to continue down the optical train toward the detector, and other columns of micromirrors set to their "off" position to send undesired wavelengths in a different direction to a beam dump. The spectrally manipulated light is then recollimated and sent to a second grating, which reverses the dispersion of the first, enabling a final imaging optic to form a nondispersed, spectrally filtered quasi-white-light image at the detector array.

Several variations on this basic theme have been described, generally reflecting the many possible designs for imaging spectrometers: refractive versus reflective optics, flat transmission gratings versus curved reflection gratings, etc. At least one design manages to use the same grating for both the initial dispersion and final de-dispersion steps. Our own early design, which we described in a 2009 paper,[1] is illustrated in Fig. 1. In this design, because there is no slit, light from the scene, already collimated if the scene is at infinity, impinges directly onto the grating. Without a slit, and with DMD situated at a spectrally dispersed image plane, light of a given wavelength from one spatial location in the scene can overlap with light of a different wavelength from another spatial location. But because the angles of incidence onto the DMD differ for the various overlapping components, they can still be disentangled and handled independently. Essentially, each spatial location in the spectral dispersion direction has a different wavelength calibration

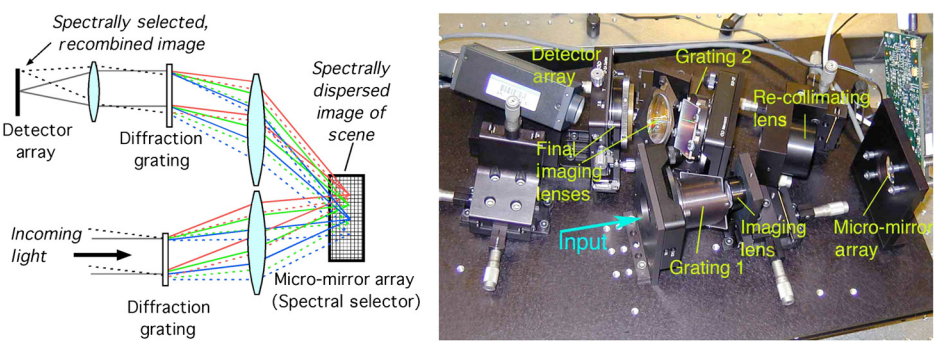

Fig. 1 The original Los Alamos micromirror array multiplexing hyperspectral imager (Ref. 1). The concept of operation is illustrated on the left: light from the scene is dispersed by a diffraction grating and imaged onto the micromirror array, which selects certain wavelengths and rejects others to implement a chemical-specific filter or Hadamard encoding scheme. The selected wavelengths are recombined by a second grating and imaged onto the detector array. Note how the spatial origin of the light (solid versus dashed rays) affects the location of spectral components at the micromirror array. The working prototype is shown at right. Either matched filters or full-spectrum Hadamard transform functions can be uploaded to the micromirrors. Full two-dimensional (2-D) spatial imaging is performed by shifting the uploaded patterns on the micromirrors, rather than by using an external scan mirror.

on the DMD. If the detector array is read out one line at a time and the spectral pattern uploaded to the DMD is shifted successively for each readout to compensate for the wavelength calibration shift, a 2-D image with the same spectral manipulation throughout can be obtained. This version of the instrument then becomes essentially a push-broom imager, but with the scanning taking place internally, on the DMD, rather than externally via instrument pointing. Indeed, all instruments based on the basic scheme outlined above require a push-broom scan, internal or external, to produce a 2-D spatial image.

The design shown in Fig. 1 clearly illustrates why designs of this sort, with the DMD situated at a spectrally dispersed image plane, cannot produce a 2-D spatial image without some sort of scan. Spatial location within the scene, in the spectral dispersion dimension (i.e., the across-slit direction in instruments employing a slit), is fundamentally intertwined with wavelength: a different spatial location implies a different angle of incidence onto the grating, which in turn implies a different direction of propagation for the dispersed light for a given wavelength.

2.2 *Programmable Spectral Filtering with Scan-Free 2-D Spatial Imaging: a Fundamentally Different Approach*

Placing the grating at a collimated pupil plane and the spectral selector—the DMD, microshutter, or LCD array—at a spectrally dispersed image plane seems a logical choice and is in accord with standard spectrometer design. This is really just a very simple modification of a traditional push-broom hyperspectral imager: the spectral selector simply sits at the location normally assigned to the focal plane detector array in an HSI and a second de-dispersing and re-imaging stage are added. But is this the only choice? The goal of producing a full-frame spectral imager, with programmable spectral processing applied simultaneously and without spectral shifts to the entire 2-D image, demands that we consider a nonstandard approach.

In order for the DMD or other spectral selectors to spectrally manipulate an entire 2-D image at once, without any wavelength shifts with image position, it must be placed at a location where light of any given wavelength converges to a well-defined location (at least to a line, but perhaps even to a point) that is unique for each wavelength, regardless of where in the 2-D scene it originated. As illustrated in Fig. 2, such a location can be produced if the grating is placed at an image plane in which light arrives at the same angle of incidence, at least in the spectral dispersion dimension, for all image points. Since in this situation, light strikes the grating at the same angle of incidence for all image points, light of a given wavelength will emerge from the grating traveling in the same unique direction (in the spectral dimension) for all image points. An imaging optic can, therefore, focus these collimated wavelength bundles to well-defined wavelength bins at the programmable selector, and a final imaging optic can subsequently recombine the wavelengths passed by the selector to a nondispersed, spectrally filtered image at the detector.

3 Realization Strategies and Early Prototypes

Given the concept outlined above, the problem of how to produce a full-frame programmable filter becomes one of how to produce the required image plane having an angle of incidence independent of location within the image. While obviously not a property of a simple lens, this can actually be accomplished relatively simply in several different ways, each with its own advantages and disadvantages.[11]

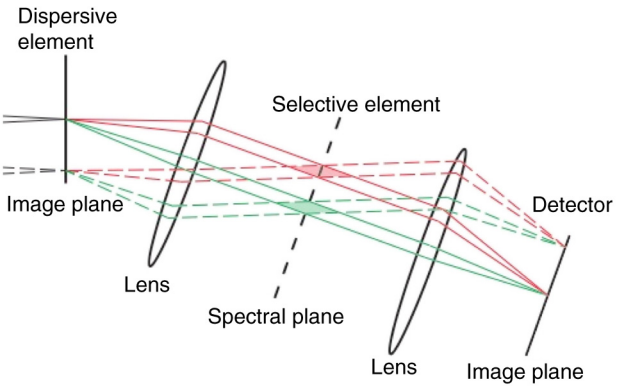

Fig. 2 Conceptual representation of a full-frame programmable spectral imager capable of simultaneous spectral manipulation of an entire 2-D image, without spectral shifts across the image. The spectral dispersive element (grating or prism) is placed at an image plane whose light has the same angle of incidence in the spectral dispersion dimension (the plane of the page in this diagram) throughout the image. Assuming this sort of image plane can be produced, light of a given wavelength emerges from the disperser propagating in the same direction in the spectral dimension, regardless of where in the image it originated. Light from all points in the 2-D image can, therefore, be focused to unique wavelength bins at a spectral plane, where a programmable selective element (DMD, LCD, microshutter array, etc.) can turn selected wavelengths "on" or "off" simultaneously and in the same manner for all image points. A final imaging optic then collects the wavelengths passed by the selector and focuses the light to a spectrally manipulated quasi-white-light image at the detector array.

3.1 *Microlouvers and Capillary Arrays*

An ordinary imaging lens (or mirror), with its entrance pupil more or less colocated with the optic itself, produces an image in which the angle of incidence onto the image plane varies strongly with position, since in this situation the chief ray for each image point travels in a straight line from the original object point, through the center of the colocated pupil and lens, to the image point. But if the system is optically fast enough, i.e., if the f-number is small enough, each image point will have a ray traveling normal to the image plane, and a fan of rays whose projection onto the spectral dispersion dimension is normal to the image plane. One strategy for achieving an image plane with a spatially invariant angle of incidence is simply to restrict the rays reaching the image plane to those traveling in the right direction, and reject the others.

One way to do this would be to use a microlouver assembly consisting of many thin, parallel, blackened plates separated by spaces, as shown in Fig. 3, that pass rays traveling in the desired directions and absorb the rest. The tolerance with which the directional selection is performed, and the resulting range of propagation directions, will determine the achievable spectral resolution when a grating is placed at this image plane. In essence, this scheme is an array of parallel mini-spectrometers, each with its own "slit," one for each column of pixels.

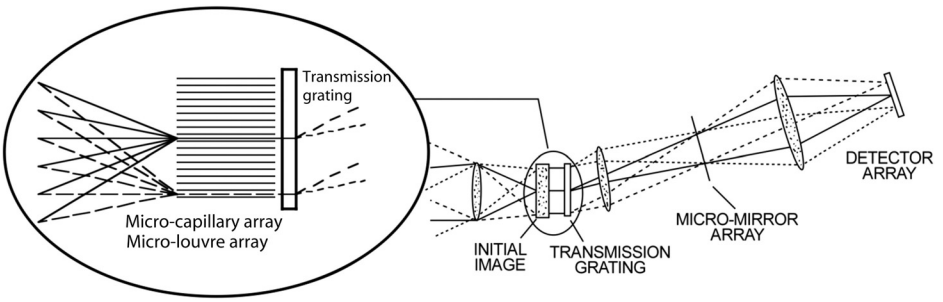

Fig. 3 Conceptual diagram of a full-frame programmable spectral imager using a conventional imaging lens to form the initial image, a transmission grating as the dispersive element, and a microlouver array or capillary array to restrict the angles of incidence onto the grating (Ref. 11). With a microlouver array, the effective collimation is in one dimension: in the spectral dimension, the plane of the page in this figure, the projection of the admitted rays' propagation vectors is normal to the grating, while in the orthogonal dimension, out of the page, a wide fan of rays is admitted to the grating. A bundle-of-tubes capillary array restricts the rays in both dimensions so that only rays normal to the grating are admitted. The degree of collimation in the relevant dimension, and hence the ultimate spectral resolution, depends on the aspect ratio of the louver gaps or capillary tubes, as well on how effective the applied blackening is a suppressing internal reflections.

A prototype based on this concept provided our first proof-of-principle demonstration of a full-frame programmable filter.[11,12] Instead of a microlouver assembly, this prototype was constructed using a bundle-of-tubes capillary array manufactured by the Schott Corporation, (Southbridge, Massachusetts). This capillary array consists of 100-μm inner diameter glass tubes, bundled into a hexagonal close-packed array $\sim$2 cm in diameter, and 1.27 cm long. The glass is blackened by Schott using a hydrogen reduction process, to absorb off-axis rays and suppress waveguide-like behavior. The spectral selector in this first prototype was an LCD spatial light modulator array (Holoeye Corp., San Diego, California) between crossed polarizers. The original design incorporated a 300-grooves/mm transmission grating, allowing this prototype filter to provide programmable bands over the entire visible spectrum. Sample images obtained with this early prototype are shown in Fig. 4.

A number of shortcomings of this early prototype immediately became apparent. First, the use of a capillary array rather than a microlouver structure (a choice driven by the ready availability of the capillary array) rejects off-axis rays in both the spectral and orthogonal dimensions, unnecessarily discarding a large fraction of the available light and making this a rather low-throughput system. Secondly, rather low spectral resolution was achieved, no better than 10 nm in the visible region, with the 300-grooves/mm grating used. This is largely because of simple geometry: the angular restriction achieved by the capillary tubes depends on their diameter/length ratio a/L, with the maximum off-axis deviation angle of a ray able to make it through a tube without hitting a wall being simply (in the small-angle approximation) $\delta\theta \approx \pm a/L$. This limits the spectral resolution, since, from the grating equation, $\delta\lambda \approx d(\delta\theta)$, where d is the groove spacing and $\delta\lambda$ is the wavelength shift caused by the angle shift $\delta\theta$. In practice, the angular spread is even larger, because the blackening of the walls is not perfect, allowing some internal reflections to further spread the exiting light. The obvious solution to the geometric limitation is to increase the length of the capillary array, but the hydrogen reduction blackening process can only penetrate a limited distance into these small tubes, limiting the practical lengths to not much more than 1 cm. Abutting two or three identical capillary arrays end-to-end does improve the spectral resolution in the expected manner, as we confirmed experimentally, but registration of the segments is extremely problematical. In any case, this approach may be regarded as primarily a proof-of-principle exercise, since far better approaches exist, as we shall discuss next.

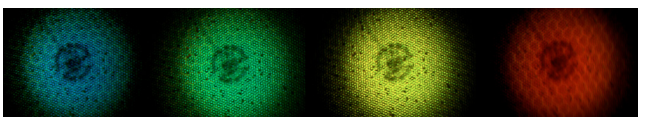

Fig. 4 Actual color images of a light bulb (note the "100 W" and manufacturer's label visible near the center of each image) obtained with the first prototype capillary-array- and LCD-based full-frame programmable spectral filter. In each image, the LCD array is programmed for a single moderate-width spectral band ($\sim$50 nm), corresponding to, from left to right, the visible colors blue, green, yellow, and red. (A color version of this figure can be seen in the electronic version of this article.) The hexagonal structure of the capillary array can be seen superimposed on the images. The key aspect to note is that the color is the same throughout each image, indicating negligible wavelength shift with position.

3.2 *Telecentric Imaging onto the Grating*

The conditions set out in Sec. 2.2—that the diffraction grating should be situated at an image plane and that the image plane should be one in which light arrives at the same angle of incidence, at least in the spectral dispersion dimension, for all image points—can be stated in a more straightforward and optically rigorous manner: The imaging system forming the image on the grating should be telecentric in image space and should have a large f ratio (hence a small range of ray angles) in the spectral dimension and an arbitrary (but preferably small, to admit the most light) f ratio in the orthogonal dimension.[11] The simplest way to construct an image-space telecentric imager is to place the aperture forming the entrance pupil at one focal length in front of an ordinary lens. To achieve the angular restriction requirements, greatly restricted in the spectral dimension and essentially unrestricted in the orthogonal dimension, the entrance pupil becomes a slit. A simple system based on this principle is illustrated in Fig. 5.

The slit-shaped entrance pupil allows for one-dimensional (1-D) collimation, rather than the unnecessarily restrictive 2-D collimation of the capillary array, greatly improving the throughput. Furthermore, in this type of instrument, spatial and spectral resolutions are determined by the same

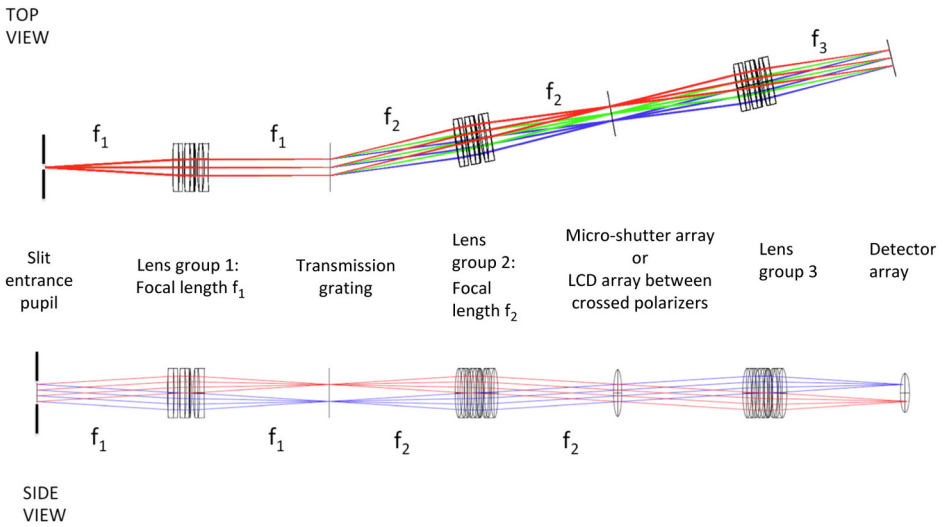

Fig. 5 Optical layout for a full-frame programmable spectral imager using telecentric imaging onto a transmission grating (Ref. 11). In general, telecentric imaging can be achieved by placing the entrance pupil aperture at one focal length in front of the imaging lens. Here, the entrance pupil is a slit, narrow in the spectral dimension, and tall in the orthogonal dimension. In the spectral dimension (top view), light strikes the grating as very narrow ray bundles, with the bundles for all image points parallel to each other. In the other dimension (side view), the bundles are still parallel to each other but span a much wider range of angles because the slit is made tall to admit as much light as possible. Note that the slit remains fixed and serves no scanning function. The dispersed light is focused to a spectral plane, where the spectral selector is located. Here, for simplicity, we show the selector as a microshutter or LCD array so that selected light passes straight through. A final lens group images the spectrally filtered scene onto the detector, forming a full 2-D image of the scene with no need for a push-broom scan.

considerations as any other spectrometer and are limited only by fundamental optics. The most important of these fundamental limitations is the diffraction-mediated trade-off between spatial and spectral resolution: narrowing the slit entrance pupil improves the spectral resolution but leads to increased diffraction that degrades the spatial resolution at the grating and the detector image planes. This trade can be quantified as follows:

Basic diffraction theory[13,14] shows that, for a rectangular aperture of full width d, the spatial resolution in the dimension corresponding to the width d is given by

$$\delta x = \frac{f\lambda}{d}, \tag{1}$$

where f is the focal length of the imaging lens, λ is the wavelength, and δx is the smallest resolvable line spacing in the image plane. (In the more familiar case of a circular aperture, a factor of 1.22 appears in the analogous expression; this reduces to unity for the rectangular case.) If we are considering the initial image on the grating in Fig. 5, then f in Eq. (1) becomes the initial focal length f_1. The final image on the detector is magnified or reduced by a factor of (f_3/f_1), so for that case, f_3 replaces f in Eq. (1).

Spectral resolution, expressed as an absolute wavelength spread, is more complicated and more instrument specific since it depends on the grating groove spacing in addition to the slit width and focal length. Let us consider instead the number of spectral bands that can be manipulated by a spectral selector (DMD or LCD). If we let f_1 equal f_2 in Fig. 5, a spectrally dispersed 1 : 1 image of the entrance slit is produced at the spectral selection plane; if f_1 and f_2 are not equal, the slit image is magnified by a factor of (f_2/f_1). The number of nonoverlapping spectral bands, N, that can be displayed across a spectral selector of width W is simply the selector width divided by the width of the slit image

$$N = \left(\frac{f_1}{f_2}\right)\frac{W}{d}. \tag{2}$$

The number of nonoverlapping spectral bands is also equal to the spectral range, $\Delta\lambda = |(\lambda_{\text{start}} - \lambda_{\text{end}})|$, that can be displayed across a selector of width W (which depends on the spectral dispersion as determined by the grating groove density and focal length), divided by the spectral resolution $\delta\lambda$, or more correctly, the spectral bin width (which is determined by the spectral dispersion and the slit width)

$$N = \frac{\Delta\lambda}{\delta\lambda}. \tag{3}$$

Combining Eqs. (1)–(3), together with the appropriate magnification factors, we arrive at a simple relationship between the spatial resolution δx in the final image at the detector and the spectral resolution $\delta\lambda$

$$(\delta\lambda)(\delta x) \approx \left(\frac{f_2 f_3}{f_1}\right)\frac{\lambda}{W}\Delta\lambda, \tag{4}$$

where $\Delta\lambda$ is spectral range that fills the width, W, of the spectral selector, λ is the average wavelength, and f_1, f_2, and f_3 are the focal lengths of the first, second, and third lens groups in Fig. 5. Alternatively, expressing this in terms of the number of spectral bands N, this becomes

$$\frac{(\delta x)}{N} \approx \left(\frac{f_2 f_3}{f_1}\right)\frac{\lambda}{W}. \tag{5}$$

Results illustrating this spatial/spectral resolution trade, for a system very similar to the one in Fig. 5, are shown in Fig. 6. This particular system used focal lengths of ~10 cm for all three lens groups, spanned the entire visible spectrum on a selector ~1 cm across, and imaged onto

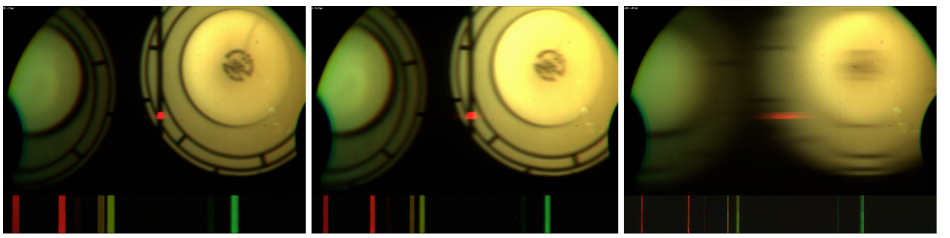

Fig. 6 Images and spectra obtained with a telecentric full-frame spectral filter system similar to that shown in Fig. 5. The scene in each image consists of two desk lamps (one partially obscured on the left) and a laser spot in the center. The laser spot, used as a spectral reference, consists of several coaligned lasers of various wavelengths. Across the bottom of each frame is the spectrum of this laser reference, obtained by moving the detector array from the image plane to the spectral plane, illustrating the spectral resolution for the settings used to obtain the image. From left to right, the spectral resolutions are 2.7 nm (150 spectral bands across the spectral range), 1.8 nm (225 bands), and 0.4 nm (1000 bands). With the exception of the green laser, a diode laser that reveals some multimode structure upon close inspection of the highest resolution figure, the laser linewidths are all much smaller than the instrument resolution. While some slight blurring is apparent in the 225-band example, only in the last case does diffraction significantly degrade the spatial image quality.

a detector array ~1.8 cm wide. The results are consistent with Eq. (5), as can be verified by close examination of the diffraction pattern of the laser spot, although the overall visual appearance is considerably better than Eq. (5) would suggest [because the resolution criterion used in Eq. (5), based on the distance to the first zero of the $\mathrm{sinc}^2(x)$ diffraction pattern, is somewhat pessimistic]. In terms of practical image usefulness, we found that, with the slit width set to produce 225 spectral bands across the selector plane (spectral resolution 1.8 nm), the spatial resolution in the final image was still quite acceptable with only minor blurring, but attempting 1000 spectral bands (spectral resolution, 0.4 nm) resulted in unacceptably degraded spatial resolution. Many variations of this basic design are possible, however, and these results only apply to one particular system. Equation (5) provides a means for optimizing a system given a set of requirements on the number of spectral bands, spatial resolution, and overall instrument size. In the absence of a size constraint, Eq. (5) shows that the way to simultaneously maximize both spatial and spectral resolutions is to make the first stage of the instrument—imaging onto the grating—large (i.e., a large grating and a large value of f_1) and then demagnify the grating image using small values for the subsequent focal lengths f_2 and f_3.

4 DMD-Based Full-Frame Programmable Filters

Compared to liquid crystal devices (LCDs), DMDs offer significant advantages as spectral selectors. While LCDs are limited to the visible spectral region and a small portion of the near infrared, commercially available DMDs can operate well from the ultraviolet into the thermal infrared. Moreover, because LCDs are polarization-based devices and need to be placed between crossed polarizers, they begin by throwing away at least half of the available light. DMDs offer a well-defined "off" state with essentially no light leakage and high throughput in the "on" state, whereas considerable off-state rejection leakage can occur in LCD setups, depending on the quality of the polarizers and field-of-view. For these reasons, the aim of this work from the outset has been to produce DMD-based full-frame programmable spectral filters.

Implementing the concepts described in the preceding sections using DMDs also presents some unique challenges, however, the most important of which is the diffraction produced by the micromirror array itself.

4.1 *Diffraction in the DLP® DMD and Its Consequences*

The Texas Instruments DLP® micromirror array used here consists of 1024×768 square micromirrors with a 13.68-μm pixel pitch. Each micromirror can be set to tilt along its diagonal at either $+12$ or -12 deg with respect to the device plane, the two orientations representing the "on" and "off" states in this application. The flat, plane-of-device orientation is not a settable position for the micromirrors, although when the device is powered down the micromirrors float at approximately this orientation.

With micromirror dimensions this small, it is readily apparent that diffraction effects will play an important role in the interaction of the light with the micromirror array. An intuitive understanding of the role of diffraction can be gained by noting that the micromirror array is, in essence, a 2-D diffraction grating. The angular directions of the various diffraction orders are determined solely by the 13.68-μm spacing of the micromirror 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 deg tilt of the micromirrors, in direct analogy with the blaze angle of a conventional diffraction grating. With all the micromirrors tilted in 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 deg. Viewed in this manner, as approximately a 1-D grating, the effective groove spacing is one-half the diagonal dimension of the micromirrors, or roughly 9.67 μm. Inserting this groove spacing in the grating equation for normal incidence, $m\lambda = d \sin\theta$, where m is the diffraction order, λ is the wavelength, and d is the groove spacing, one finds, for example, that the preferred 24 deg diffraction angle θ, as dictated by the 12 deg "blaze" of the micromirror tilt, occurs close to the sixth order for the He-Ne laser wavelength of 632.8 nm, as can be seen in Fig. 7.

In previous non full-frame, scanning push-broom types of programmable spectral filters based on the DLP, in which the DLP was situated at a spectrally dispersed image plane, diffraction from the micromirrors was not a major issue, apart from the fact that it affected the throughput of the device and whether or not it would work well at a given wavelength. Because of the image plane location of the DLP in those designs, subsequent optics could simply collect as many of the significant diffraction orders as practical, and the light

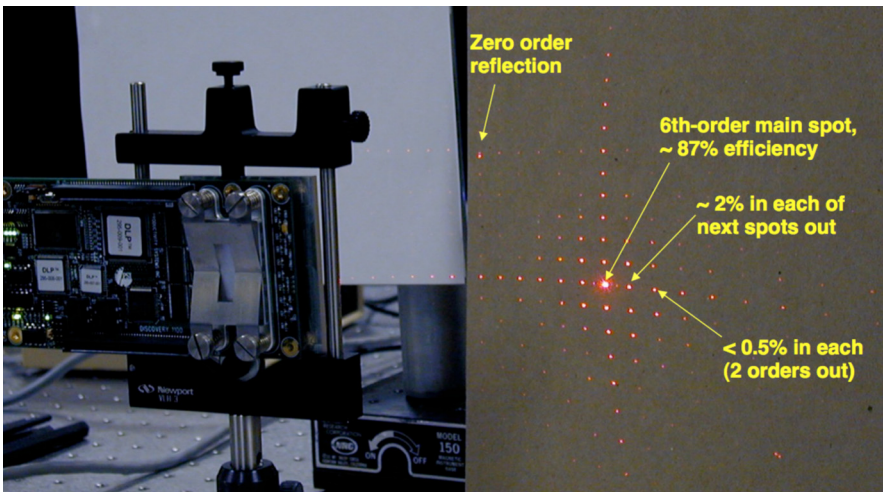

Fig. 7 Photograph showing the diffraction pattern generated by the Texas Instruments DLP® micromirror array, illuminated by a He-Ne laser (wavelength, 632.8 nm), with all the micromirrors turned to the same orientation (e.g., the "on" state). A weak zero-order reflection is generated by the dead space between micromirrors, providing a useful reference. The back side of the circuit board containing the micromirror array can be seen at left; the micromirror array is facing away from the viewer in the photo. The He-Ne 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 (~87%) goes to the bright spot at the sixth diffraction order in both the horizontal and vertical directions, corresponding to the 24-deg micromirror reflection angle.

would automatically be reassembled correctly in the final image, much as it would be if it were simply diffusely scattered by a screen. But in the present concept, the spectral plane, at which the DLP is situated, is essentially a pupil plane, or more precisely, a spectrally dispersed image of the entrance pupil. At a pupil-like location, changes in propagation direction, as occur with micromirror diffraction, translate to changes in position at the final image. Left uncompensated in our full-frame programmable filter, diffraction from the DLP would produce a set of spatially separated final images, one for each of the 2-D diffraction orders, and worse yet, their spatial positions would be wavelength dependent, making each of the multiple images a spectrally dispersed smear.

4.2 *DLP®-Based Full-Frame Spectral Imagers with Diffraction Compensation*

Fortunately, it is possible to reverse the effects of diffraction from the DLP, remove the spectral dispersion of the final image, and bring the multiple ghost images produced by the various DLP diffraction orders into registration. This can be accomplished using another 2-D diffraction grating of the same pitch and blaze angle, operating in the opposite-signed diffraction orders as the original DLP. This 2-D grating might be an actual 2-D structure manufactured for this purpose, a pair of orthogonally oriented ordinary diffraction gratings, or another DLP with all its mirrors tilted in the same direction.[11] Figure 8 illustrates the two-grating compensation scheme.

While grating-based compensation works with off-the-shelf gratings only approximately matching the DLP pixel pitch (the mismatch can be offset with slight tilts of the gratings), as illustrated in Fig. 8, construction and alignment of such a system are difficult. On the other hand, obtaining gratings exactly matching the DLP's pixel pitch and blaze is expensive. The most accurate match to a DLP's diffraction behavior can be obtained, of course, with another identical DLP. This exact match leads to much simpler and more elegant designs, as illustrated in Figs. 9 and 10, which show our current working prototype.

It is important to recognize that these diffraction compensation schemes cannot recover and re-register all the light thrown into the DLP's many diffraction orders. The problem is that the compensating DLP or gratings produce multiple diffraction orders of their own. This is not a big problem for wavelengths for which a diffraction order falls at or very near the 24-deg micromirror reflection angle, as in Figs. 7 and 11(a), since in those cases, most of the light will be in one diffraction order, both for the original DLP and the compensator. But for wavelengths for which 24 deg falls between diffraction orders, as in Fig. 11(b), considerable light will be lost to the final image. Consider, for example, the wavelength at 525 nm as illustrated in Fig. 11(b). For this wavelength, the 24-deg reflection angle falls halfway between the seventh and eighth diffraction orders, or more precisely, denoting the 2-D diffraction orders as (x, y) pairs, it falls in the middle of the square formed by the 2-D diffraction orders (8,7), (7,8), (7,7), and (8,8), and the light energy will be more or less equally distributed among these four brightest orders (fortunately, the fraction of the energy going into orders farther from this square is fairly small). The compensator will use its $(-8, -7)$ order to place the DLP's (8,7) order in the correct location in the image plane, $(-7, -7)$ will compensate (7,7), and so on. The problem is that the light from the DLP's (8,7) order will go into the compensator's other orders as well: roughly one fourth going into each of $(-7, -7)$, $(-8, -8)$, and $(-7, -8)$, none of which direct light to the correct location in the final image plane. So at best, assuming further-removed orders are negligible, the compensated image at 525 nm will only be roughly one fourth as bright as at wavelengths, like 655.7 or 562 nm, having a diffraction order falling exactly at 24 deg. The result is that the instrument's response function will oscillate strongly as a function of wavelength, with peaks at the on-order wavelengths and troughs at the between-order wavelengths, as will be seen in the results.

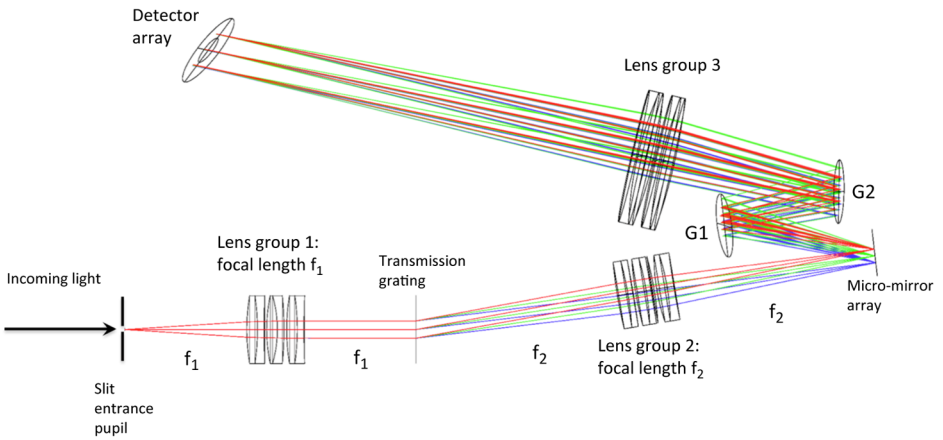

Fig. 8 Optical layout for a DLP-based full-frame programmable spectral filter, illustrating DLP diffraction compensation by means of two diffraction gratings (G1 and G2; Ref. 11). The system uses a slit entrance pupil and telecentric imaging onto a transmission grating, followed by quasi-telecentric imaging of the dispersed pupil onto the DLP micromirror array. The two gratings used in this design have a groove density of 75 grooves/mm, the closest match available in off-the-shelf gratings to the DLP's pixel pitch. The two gratings have their grooves oriented roughly 90 deg to each other in order to compensate both dimensions of the DLP's 2-D diffraction pattern.

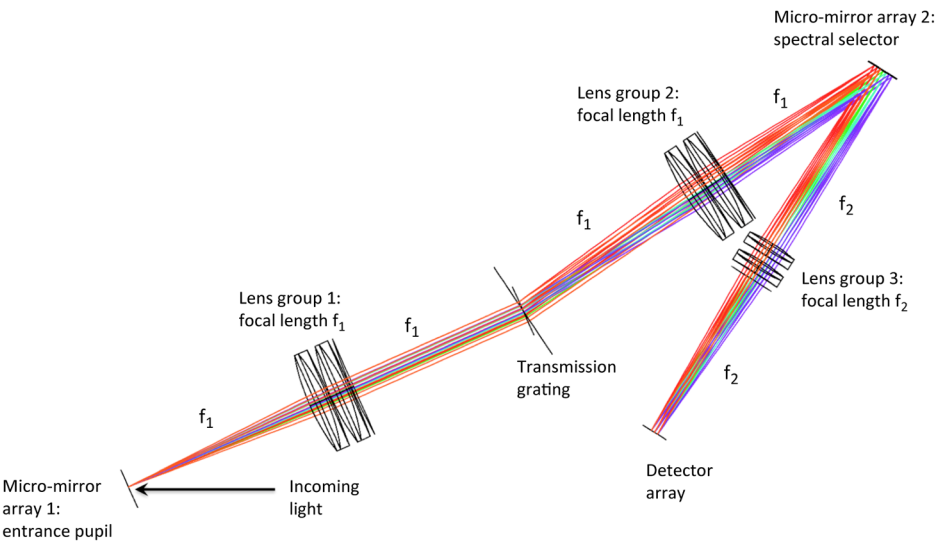

Fig. 9 Optical layout for a full-frame DLP-based programmable spectral imager, with diffraction compensation by means of a second DLP (Ref. 11). Incoming light strikes the compensating DLP first, introducing a "precompensating" diffractive dispersion. This first DLP, programmed to display a fixed slit function (no scanning is necessary), also serves as the entrance pupil for the quasi-telecentric imaging onto the grating. The spectral selection DLP reverses the effects of the precompensating DLP to produce a final nondispersed image at the detector.

It should also be noted that the complexity associated with micromirror diffraction compensation is really a consequence of the idiosyncrasies of a particular type of micromirror array (albeit the only type readily available, commercially, at this time). The fact that the DLP® array only offers two settable mirror positions, both tilted with respect to the plane of the array, means that for both the "on" and "off" states, the DLP behaves like a grating. Suppose instead that there was a settable mirror position in which the micromirror is oriented parallel to the array plane. In that case, regions of the array programmed to this "flat" position would behave as a simple mirror, not as a grating. This state could become the "on" position in applications of this sort so that the selected light would undergo simple, diffraction-free, wavelength-independent specular reflection. Regions in the "off" state would necessarily still involve tilted micromirrors and associated diffraction, but this would be irrelevant since this light is rejected anyway. Such a DMD would greatly improve the performance for spectral applications in which the DMD must sit at a pupil-like plane rather than an image plane, eliminating the need for diffraction compensation and improving throughput. DMDs of this type have been produced, particularly for astronomical applications,[15] but are not commercially available. We are currently working on fabricating such a DMD of our own design, with large micromirrors optimized for the long-wave infrared region (8 to 12 μm).

5 Results from a DLP®-Based Full-Frame Programmable Filter

With the dual-DLP design shown in Fig. 9, we constructed a full-frame programmable spectral imager covering the entire visible spectrum and capable of well over 100 programmable

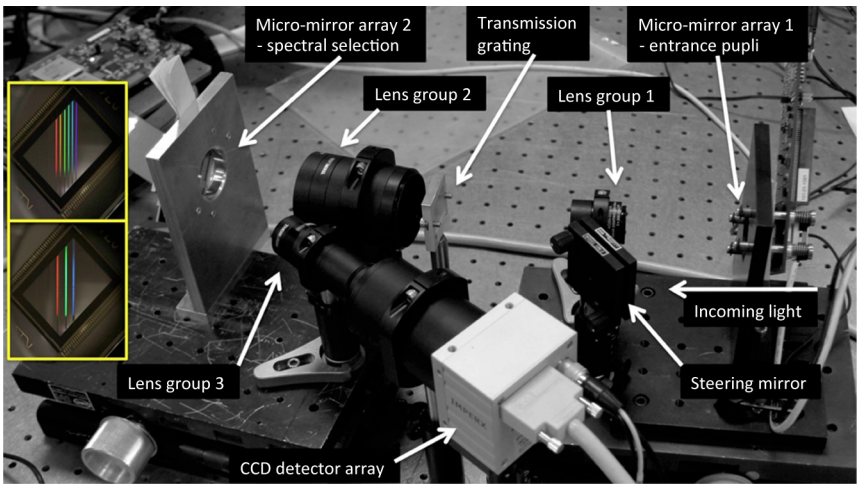

Fig. 10 Photograph of a working prototype full-frame DLP-based programmable spectral imager, with diffraction compensation by means of a second DLP (Ref. 11), based on the design shown in Fig. 9. Inset shows the spectral selection DLP array, programmed with two example spectral band selections, as viewed from the detector position. Because the DLP micromirrors tilt along their diagonals, both DLP arrays are mounted, for convenience, with a 45-deg rotation about their axes, so that incident and reflected light remains parallel to the table throughout the instrument.

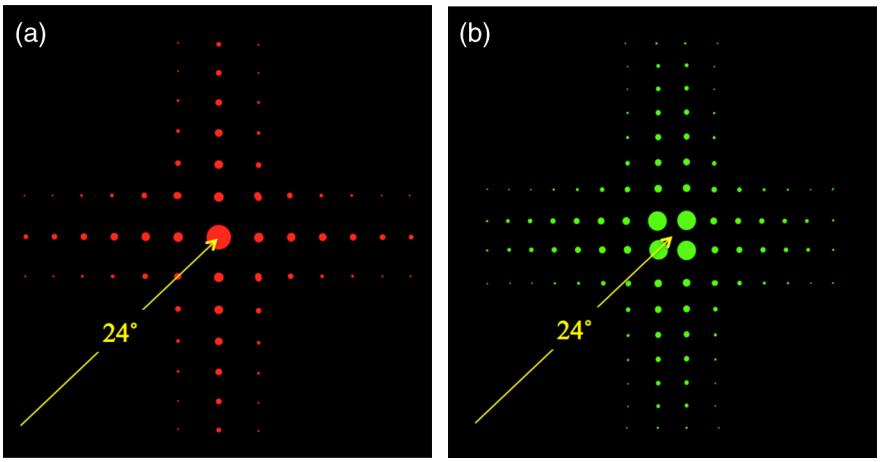

Fig. 11 Conceptual diagram illustrating the DLP diffraction patterns for two wavelengths, (a) 655 nm, which has a diffraction order falling at the DLP's preferred "blaze reflection" angle of 24 deg, and (b) 525 nm, for which the 24-deg preferred angle falls between diffraction orders. In the first case, most of the diffracted energy, nearly 90%, is contained in the preferred angle diffraction order. In the second case, diffracted energy is distributed primarily amongst the four nearest-neighbor diffraction orders surrounding the preferred reflection angle.

spectral bands. With this prototype, we have been able to demonstrate both full HSI using the Hadamard transform technique and background-suppressing matched-filter imaging of a faint target signature.

The Hadamard transform technique uses a set of N binary basis functions, with roughly half the wavelength bands "on" and half "off" in any given Hadamard function, to encode the N bands of the complete spectra. The spectra can then be recovered with a simple linear inverse transform. As with previously described DMD-based programmable filters, the Hadamard transform technique is always an option in our instrument since the Hadamard encoding functions are simply another set of spectral band selections that can be uploaded to the DMD. The Hadamard technique is discussed in some detail in our earlier paper[1] and in great detail in Harwit and Sloane's classic book on the subject.[16] Here, we use it to obtain full spectral information on a scene prior to constructing background-suppressing matched-filter band sets for particular targets. Examples of Hadamard-derived spectra and background-suppressing matched-filter imaging are shown in Fig. 12.

Figure 12 shows how the prototype spectral filter can extract a weak target signature from a strong background. In this example, the continuous-spectrum incandescent bulb can be thought of as the "background" and the reflection of the discrete-spectrum compact fluorescent lamp (CFL) can be thought of as the target. A matched filter for the CFL spectrum can be produced as the difference of two sets of bands, one selected to pass the CFL emission lines and the other selected to pass wavelengths falling between those lines. If the bandwidths for two band sets are chosen such that the integrated intensity of the incandescent bulb is the same in both sets of spectral bands, the incandescent bulb will cancel out in the difference, leaving only light from the CFL, either direct or reflected. Doing so in Fig. 12, the incandescent bulb (the "background") is so effectively suppressed that it appears to be off and visible only because of the reflected illumination from the CFL. This process is

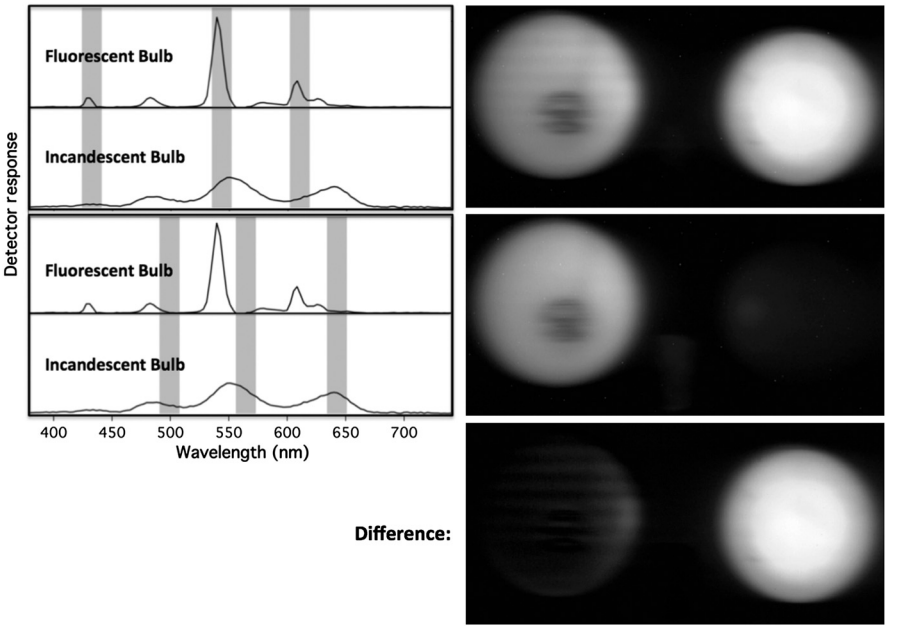

Fig. 12 An example of matched filter imaging with background suppression using the full-frame programmable spectral imager prototype shown in Figs. 9 and 10. The scene for this example consists of two light bulbs: an incandescent bulb on the left and a compact fluorescent lamp (CFL) on the right. Shown on the left side are 128-band spectra of the two bulbs, obtained with the prototype using the Hadamard technique. The CFL spectrum consists of several sharp lines, whereas the incandescent bulb has a continuous spectrum, modulated here by the instrument response function (see text). For the top image, the DMD is programmed to pass three bands (gray bars in the upper spectrum on left) corresponding to CFL emission lines; the middle image is obtained with three different programmed bands (gray bars in the lower spectrum), chosen to miss the CFL emission. The difference between these two images is shown at bottom. Note that, in the difference, the incandescent emission is so effectively suppressed that the only illumination visible is the reflection of the CFL.

analogous to the one that could be implemented in the infrared for detecting a trace gas against a spectrally complex background. The example shown in Fig. 12 is a fairly crude, hand-constructed version of a matched filter, in which we manually varied the bandwidths until the incandescent bulb canceled effectively. Much more sophisticated and mathematically rigorous background-suppressing matched filters can be implemented with this instrument, filters that make use of the duty cycle-based gray-scale capability of the DLP array to optimally handle spectrally complex backgrounds, as is discussed in companion papers[17,18] by Graff and Love.

6 Summary

Although micromirror-based programmable filters historically have been limited to 1-D spatial imaging, with a push-broom scan required to sweep out the other spatial dimension, this limitation is not fundamental. We have shown here that by placing the dispersive element of the spectrometer at a telecentric image plane and constraining the focal ratio in the dispersion dimension, it is possible to create a spectral plane at which a micromirror array or the other programmable spectral selector can be placed to manipulate the spectral content of an image simultaneously for all image points, without the spectral shifts with image position that are unavoidable when the dispersive element is placed at the more customary collimated pupil location. We have demonstrated several versions of full-frame programmable spectral filters based on this principle, with the telecentric image on the disperser formed either conventionally, with a slit-shaped entrance pupil placed a focal length in front of the imaging optic, or by means of an off-axis-rejecting capillary array. This strategy does lead to complications when used with spectral selection devices, like the DLP array, that exhibit strong diffractive behavior, but, as we have demonstrated here, the DLP diffraction can readily be compensated with a second DLP array or pair of gratings so that a high quality nondispersed final image can be formed. Using these techniques, we have demonstrated full-frame programmable spectral imagers with over a 100 programmable spectral bands, capable of implementing real-time background-suppressing matched filter imaging in two spatial dimensions at once, as well as full-spectrum Hadamard HSI.

Acknowledgments

The authors gratefully acknowledge Tom Hale for his valuable assistance in constructing the early capillary-array/LCD prototype. This work was performed as part of an Exploratory Research project (20110267ER) under the Laboratory-Directed Research and Development (LDRD) program at Los Alamos National Laboratory.

References

1. S. P. Love, "Programmable matched filter and Hadamard transform hyperspectral imagers based on micro-mirror arrays," *Proc. SPIE* **7210**, 721007 (2009).
2. E. P. Wagner, II et al., "Construction and evaluation of a visible spectrometer using digital micromirror spatial light modulation," *Appl. Spectrosc.* **49**(11), 1715–1719 (1995).
3. K. J. Kearney, M. Corio, and Z. Ninkov, "Imaging spectroscopy with digital micromirrors," *Proc. SPIE* **3965**, 11–20 (2000).
4. C. M. Wehlburg et al., "Optimization and characterization of an imaging Hadamard spectrometer," *Proc. SPIE* **4381**, 506–515 (2001).
5. A. Wuttig and R. Riesenberg, "Sensitive Hadamard transform imaging spectrometer with a simple MEMS," *Proc. SPIE* **4881**, 167–178 (2003).
6. M. Hanf et al., "Novel Hadamard transform spectrometer realized using a dynamically driven micro-mirror array as light modulator," *Proc. SPIE* **5717**, 117–126 (2005).

7. N. T. Quyen et al., "New Raman spectrometer using a digital micromirror device and a photomultiplier detector for rapid on-line industrial analysis," *Appl. Spectrosc.* **62**(3), 273–278 (2008).
8. N. Goldstein et al., "DMD-based adaptive spectral imagers for hyperspectral imagery and direct detection of spectral signatures," *Proc. SPIE* **7210**, 721008 (2009).
9. N. Goldstein et al., "Infrared adaptive spectral imagers for direct detection of spectral signatures and hyperspectral imagery," *Proc. SPIE* **8618**, 86180D (2013).
10. J. D. Newman et al., "MEMS-based spectral-polarimetric imaging and target tracking architecture for airborne broad-area search," *Proc. SPIE* **8053**, 805302 (2011).
11. S. P. Love and D. L. Graff, International Patent Application PCT/US2012/000417, "Full-frame programmable spectral imager," Filed September 28, 2012 (original provisional US patent application filed September 30, 2011).
12. S. P. Love and D. L. Graff, "Full-frame programmable spectral filters based on micromirror arrays," *Proc. SPIE* **8618**, 86180C (2013).
13. E. Hecht and A. Zajac, *Optics*, Addison-Wesley, Reading, Massachusetts (1974).
14. M. Born and E. Wolf, *Principles of Optics*, 7th ed., Cambridge University Press, Cambridge, UK (1999).
15. M. D. Canonica et al., "The two-dimensional array of 2048 tilting micromirrors for astronomical spectroscopy," *J. Micromech. Microeng.* **23**, 055009 (2013).
16. M. Harwit and N. J. A. Sloane, *Hadamard Transform Optics*, Academic, New York (1979).
17. D. L. Graff and S. P. Love, "Real-time matched-filter imaging for chemical detection, using a DMD-based programmable filter," *Proc. SPIE* **8618**, 86180F (2013).
18. D. L. Graff and S. P. Love, "Toward real-time spectral imaging for chemical detection with a digital micromirror device-based programmable spectral filter," *J. Micro/Nanolithogr. MEMS MOEMS* **13**(1), 1–8 (2013).

Steven P. Love is a physicist in the Space and Remote Sensing Group at Los Alamos National Laboratory. He received his BS in physics from Washington State University in 1983 and his doctorate in experimental solid-state physics from Cornell University in 1991. He has led the development of several new remote sensing technologies at Los Alamos, including wide-angle imaging lidar for climate research, ultra-compact long-wave infrared hyperspectral imagers, and most recently, MEMS-based programmable spectral imagers.

David L. Graff received his PhD in chemical physics from the Ohio State University in 2010, developing new instruments and methods for high resolution spectroscopy in the microwave and submillimeter spectral regions. He has most recently worked at Los Alamos National Laboratory as a postdoctoral research assistant in the Space and Remote Sensing Group developing spectral imaging sensors based on micromirror technology. His general interests include remote sensing instrument and sensor development.