

High-Frame-Rate Reconstruction of a Dynamic 2-D Scene From Continuous Orthogonal Projections

Robert Newsome, *Student Member, IEEE*, and Umran Inan, *Fellow, IEEE*

Abstract—We present a method for high-frame-rate reconstruction of a dynamic 2-D scene from two effectively continuous orthogonal projections. A natural application of this technique is in the high-speed imaging of transient luminous events—e.g., sprites, elves, gigantic jets, blue jets—in the Earth’s atmosphere, where 1-D multianode photometers are used to improve temporal resolution and optical sensitivity over that offered by traditional video cameras. To characterize the performance of the method, we use sprite-image sequences to produce multicolumn and multirow projection data and then use the projection data to produce 2-D “reconstructed” sprite images at the projection data frame rate to compare with the original sprite-image sequence. We also test the performance limits and the spatial resolution of the algorithm using artificially constructed dynamic images.

Index Terms—Image reconstruction, inversion, sprites, total variation, transient luminous events (TLEs).

I. INTRODUCTION

WE PRESENT a technique with which continuous measurements from two orthogonal projections of a dynamic 2-D scene can be used to augment video images to produce a high-frame-rate representation of the scene. The reconstruction is accomplished by interleaving the video frames with continuous orthogonal scene projections, modeling the process of data capture as a linear observation process, and then performing an inversion by numerically solving a total variation minimization problem subject to data fidelity constraints. The benefit of this technique is that it allows one to effectively increase the frame rate of a 2-D video camera by a significant factor (e.g., factors of 14, 16, and 25 are achieved in the examples considered in this paper).

This technique is particularly beneficial in situations where continuous (or effectively continuous) orthogonal projections of the scene are readily available. One application is in high-speed imaging of transient luminous events (TLEs) in the Earth’s atmosphere, where pairs of 1-D multianode photometers are used to improve the frame rate and optical sensitivity (albeit at the expense of spatial resolution) of conventional

Manuscript received October 25, 2007; revised February 15, 2008. First published April 28, 2009; current version published July 23, 2009. This work was supported in part by a Texas Instruments Stanford Graduate Fellowship, by the Office of Naval Research (ONR) under Grant N00014-03-1-0333 to Stanford University, and by the Defense Advanced Research Projects Agency and the High Frequency Active Auroral Research Program under ONR Grant N00014-05-1-0854 to Stanford University.

The authors are with the Space, Telecommunications, and Radioscience Laboratory, Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA (e-mail: rnewsome@stanford.edu; inan@novo.stanford.edu).

Color version of one or more figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TGRS.2009.2015289

video images. An example of a TLE imaging instrument with these capabilities is the PIPER instrument built at Stanford University [1].

Inversion problems are common in geophysics and are generally solved to find the distribution of some geophysical parameter from a set of sensor measurements [2], [3]. In this paper, we use the total variation minimization technique to determine the spatio-temporal distribution of light in a region, with particular attention paid to the application of imaging TLEs. The rest of this paper is divided into four sections. Section II introduces the TLE imaging application that motivated the development of this technique. Section III describes our observation model and explains how to set up the total variation minimization problem and reconstruct the higher frame-rate scene representation. Section IV presents results using sprite video data captured at Langmuir Lab near Socorro, NM, during an observation campaign in the summer of 2004. The final section summarizes conclusions and provides a discussion of the limitations of the technique.

II. MOTIVATION

In the last few decades, several types of TLEs occurring in the Earth’s atmosphere, including sprites, elves, gigantic jets, and blue jets, have been catalogued. Because of their quick and sometimes faint nature, imaging such TLEs is often difficult. It is desirable to use an imaging device that is capable of capturing very low light levels (much lower than the minimum light levels detectable by the human eye) in a 2-D field of view¹ at very high frame rates and spatial resolutions. In practice, however, the goals of optical sensitivity, high temporal resolution, and high spatial resolution are competitive: achieving superior performance in one category generally comes at the expense of performance in at least one (if not all) of the other categories.

To date, the two most popular TLE imaging techniques have been those based on charge-coupled devices (CCDs) and those based on photometric devices. The first category comprises video-rate cameras and high-speed cameras and has been the most often used technique in sprite imaging. Recent high-speed camera measurements of sprites at 5000, 7200, and 10 000 frames/s have been reported [4], [5]. These cameras provide high temporal resolution and high spatial resolution but do not offer the level of optical sensitivity that photometric devices offer.

¹Or even a 3-D field of view in the case of TLEs captured by multiple imaging devices from multiple viewing angles.

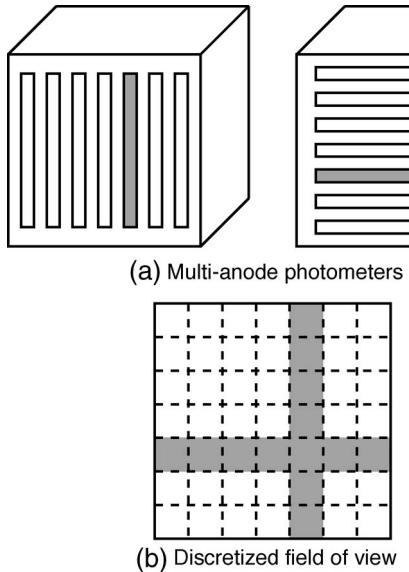


Fig. 1. (a) Two multianode photometers focused on the same field of view. (b) Each anode of each multianode photometer captures separate parts of the field of view. For any two anodes (one from each photometer), there will be one region in the field of view contributing to the response of both anodes, i.e., the intersection of the gray areas.

Photometric devices, on the other hand, offer better optical sensitivity and even higher temporal resolution at the expense of spatial resolution. A single photometer, for instance, captures its field of view in what is effectively one pixel but at arbitrarily high “frame” rates. The temporal resolution is limited by the response time of the photometer, which is often measured in nanoseconds, so single-pixel “frame” rates of millions of frames per second are conceivable. Some spatial resolution can be recovered by using multianode photometers, which typically consist of 4, 8, or 16 photomultiplier tubes arranged in a line, operating in parallel. Thus, 1-D spatial information about the field of view can be recorded by a very sensitive multianode photometer at frame rates that meet or exceed the frame rates of the highest speed CCD cameras (as presented in [6]).

The Stanford University PIPER instrument is an example of a photometric TLE imaging device. To recover some of the spatial resolution that is naturally lost by choosing to use photometers over CCDs, it uses pairs of orthogonally oriented 1-D multianode photometers, reaping the benefit of the high temporal resolutions and optical sensitivities of a photometer without sacrificing 2-D spatial information. Measurements recorded at 25 000 frames/s by both photometers are interpreted alongside each other to discern the 2-D distribution of intensity in a scene. Fig. 1 shows this idea. A pair of vertically and horizontally oriented multianode photometers is shown, focused on a common field of view. The anodes of each photometer are long and thin and are arranged in a line perpendicular to their longer dimension. If a point in the field of view is illuminated, it is detected by both photometers. While each photometer individually only determines the location of the illuminated point along one dimension (horizontally or vertically), the use of both photometer measurements simultaneously provides information about the 2-D location of the illuminated point within the field of view. Thus, the outputs of the two photometers

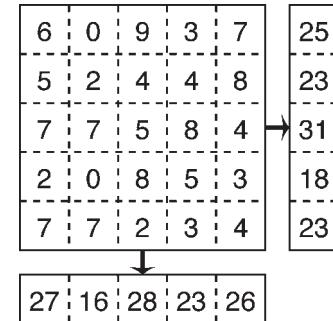


Fig. 2. Measurements from the horizontal and vertical projections represent the field-of-view sums of columns and sums of rows, respectively. The sums of the rows are recorded by the vertically oriented projection (represented on the right), while the sums of the columns are recorded by the horizontally oriented projection (represented on the bottom). In the context of TLE imaging, the numbers in the discretized field of view can be interpreted as photon counts or intensity.

in each pair (considering the high speeds at which they are acquired) act as nearly continual orthogonal projections of the field of view.

The technique described in this paper can be applied to the problem of TLE imaging by using PIPER photometer measurements to augment traditional CCD video camera data. In this scenario, the entire imaging system records the 2-D scene at the 25 000-frame/s frame rate of the PIPER photometers. Most of these high-speed “frames” consist only of 1-D photometer measurements, which are orthogonal projections of the frame (e.g., recordings of the sums of the frame’s columns and the sums of the frame’s rows as in Fig. 2). Each actual video frame is interpreted as an integration of all the missing high-speed frames occurring between it and the previous video frame. The reconstruction technique described herein is then used to fill in the missing high-speed frames using their projections in space (the PIPER photometer data) and time (the lower speed video frames) to obtain a higher frame-rate (relative to the original video camera data) 2-D representation of the field of view.

The chief advantage of this technique over the traditional use of a stand-alone CCD camera is that it can significantly increase the frame rate of a CCD camera while retaining (an albeit limited amount of) 2-D spatial resolution. In Section IV, we present examples where this technique increases the frame rates of camera data by factors of 14, 16, and 25.

III. HIGHER FRAME-RATE DYNAMIC 2-D SCENE RECONSTRUCTION

A. Imaging Configuration

We use two 1-D projections of the scene, oriented perpendicular to each other, and a video camera recording of the scene, all focused on the same field of view. Each projection contains N_p discrete measurements arranged in a line that divide the field of view into N_p equal segments.

For the purpose of this paper, we assume that one of the projections is oriented horizontally and that the other is oriented vertically. Thus, the horizontal projection represents an integration across the “columns” of the field of view, while the vertical projection represents an integration across the “rows”

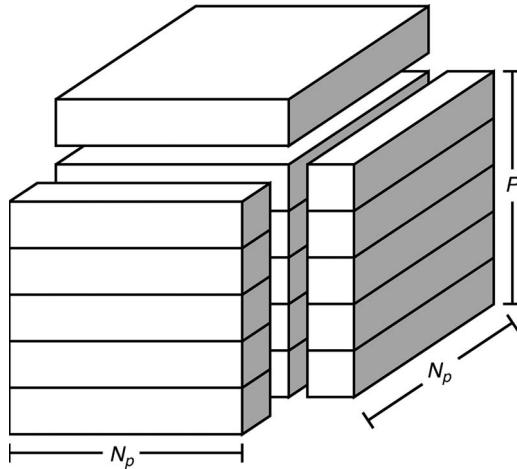


Fig. 3. Schematic representation of the “data prism.” The prism is formed by stacking the P high-speed frames vertically. The horizontally and vertically oriented projection measurements are projections of the data prism through the front and side faces, while the integrated video frame represents a projection through the top face.

of the field of view. Fig. 2 shows the relationship between the projections and their fields of view.

The video camera uses an $M_v \times N_v$ grid of imaging elements to divide the field of view into $M_v N_v$ equal segments (i.e., pixels). The overlap between two adjacent camera imaging elements’ fields of view is assumed to be negligible, and a photon coming out of any region of the field of view is registered by one and only one imaging element.

For the purposes of our discussions herein, we refer to the instruments measuring the scene’s projections as the projection measuring devices. We refer to the pair of orthogonal projections recorded at a particular instant in time as a “projection frame,” and we refer to the frames that are ultimately to be created by the reconstruction technique as “high-speed” frames (regardless of their actual frame rate—the phrase “high-speed” is a reminder that these frames have a higher frame rate than the video camera frames). Therefore, the reconstruction technique we present takes projection frames (along with video camera frames) as inputs and produces high-speed frames as outputs.

The projection measuring devices acquire samples at a rate P higher than the video camera. Thus, many projection frames (the measured sums of rows and columns of the field of view) are recorded between each video frame. Each video frame is an integration of all the high-speed frames occurring between it and the previous video frame. For each high-speed frame to be reconstructed, we have in hand only the two projections (the sums of the rows and the sums of the columns of the high-speed frame) and use these along with the next video camera frame to reconstruct the high-speed frame itself. From Fig. 3, we see that the P high-speed frames, if laid on top of each other, form a rectangular prism of data. The two projections are orthogonal projections through the rectangular prism to each of two perpendicular side walls. The video camera frame represents a third vertical projection through the prism. The task of reconstructing the high-speed frames can thus be thought of as that of reconstructing the interior of the data prism given its projection onto the three walls.

B. Reconstruction of the High-Speed Frames

Rather than reconstruct each high-speed frame individually, we reconstruct the high-speed frames in groups of P using the observations of each projection of each high-speed frame and the video camera frame for the entire set of P high-speed frames. The actual video camera frames could potentially be of a much higher spatial resolution ($M_v \times N_v$ pixels) than the high-speed frames ($N_p \times N_p$ pixels), so the video camera frames may first be cropped and subsampled down to the spatial resolution of the high-speed frames.

We can express our imaging process as

$$y = Ax + \epsilon \quad (1)$$

where $y \in \mathbb{R}^{2N_p P + N_p^2}$ is the vector of observations of the P high-speed frames to be reconstructed, $x \in \mathbb{R}^{N_p^2 P}$ is the vector of actual high-speed frames to be reconstructed, $A \in \mathbb{R}^{(2N_p P + N_p^2) \times (N_p^2 P)}$ is the matrix of observation, and $\epsilon \in \mathbb{R}^{2N_p P + N_p^2}$ is a noise vector that should be small in magnitude.

The vectors x and y are formed by use of the matrix $\text{vec}()$ operator. Specifically, x is formed as

$$x = \text{vec}([F_{p1} \quad \dots \quad F_{pP}]) \quad (2)$$

where $F_{pi} \in \mathbb{R}^{N_p \times N_p}$ is the i th high-speed frame. The vector y is formed similarly from the projections of each high-speed frame and the integrated video camera frame

$$y = \text{vec}([v_{p1} \quad h_{p1} \quad \dots \quad v_{pP} \quad h_{pP} \quad F_c]) \quad (3)$$

where $F_c \in \mathbb{R}^{N_p \times N_p}$ is the (cropped and subsampled) integrated video camera frame, and $v_{pi} \in \mathbb{R}^{N_p}$ and $h_{pi} \in \mathbb{R}^{N_p}$ are the vertical and horizontal projections of the i th high-speed frame, respectively.

For typical values of N_p and P , (1) represents an under-determined system of equations with many exact solutions in the case of no noise (i.e., assuming $\epsilon = 0$). To settle upon a reasonable reconstruction x , we therefore minimize a total variation objective that is subject to a constraint of data fidelity [7]. Specifically, we seek to solve the following optimization problem:

$$\begin{aligned} & \text{minimize} && \|G_x x\|_1 + \|G_y x\|_1 \\ & \text{subject to} && \|Ax - y\|_2 \leq \alpha \|y\|_2 \end{aligned} \quad (4)$$

where the objective function to be minimized is the total variation objective and the constraint represents a data fidelity condition. Typical values for α may range from 10^{-2} to 10^{-3} .

The total variation objective function of (4) is composed of two terms: the l_1 -norms of the x - and y -components of the gradients of each high-speed frame in the sequence of high-speed frames. The l_1 -norm is chosen over the l_2 -norm because it does not penalize large variations in x (i.e., large values in either $G_x x$ or $G_y x$) as severely as the l_2 -norm. This choice thus

allows a reconstruction of x to be found with smooth regions of nearly uniform intensity marked off by edges of sharp variation. While the use of the l_2 -norm would reduce the problem to a straightforward quadratic program with inequality constraints, the presence of sharp edges would be discouraged too greatly, resulting in reconstructed high-speed frames that appear fuzzy.

G_x and G_y are constructed in such a way as to produce sequences of images that are the x - and y -components of the gradients of the high-speed frames represented by x . For both G_x and G_y , we approximate spatial derivatives locally with forward differences (the difference between the next neighbor of a point and itself) rather than central differences (the difference between the neighbors of a point). We do not attempt to estimate the gradient at the boundaries of a high-speed frame since the values of the gradient at these locations do not meaningfully contribute to the total variation in the frame.

We construct G_x and G_y by first constructing small matrices $G_{x'} \in \mathbb{R}^{N_p(N_p-1) \times N_p^2}$ and $G_{y'} \in \mathbb{R}^{N_p(N_p-1) \times N_p^2}$, which produce the x - and y -components of the gradient of a single high-speed frame. Let $x' \in \mathbb{R}^{N_p^2}$ be a vector representing an $N_p \times N_p$ high-speed frame $F_p \in \mathbb{R}^{N_p \times N_p}$ as $x = \text{vec}(F_p)$. Note that the sizes of $G_{x'}$ and $G_{y'}$ are both $N_p(N_p-1) \times N_p^2$ rather than $N_p^2 \times N_p^2$. This is because $G_{x'}$ ignores the right boundary of the high-speed frame F_p , while $G_{y'}$ ignores the bottom boundary of the high-speed frame F_p . As an example, for the case of $N_p = 3$, we have

$$G_{x'} = \begin{bmatrix} -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 1 \end{bmatrix} \quad (5)$$

$$G_{y'} = \begin{bmatrix} -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 \end{bmatrix}. \quad (6)$$

To then build the matrices G_x and G_y from $G_{x'}$ and $G_{y'}$, for example for the case of $P = 5$ high-speed frames in the sequence of high-speed frames, we have

$$G_x = \begin{bmatrix} G_{x'} & 0 & 0 & 0 & 0 \\ 0 & G_{x'} & 0 & 0 & 0 \\ 0 & 0 & G_{x'} & 0 & 0 \\ 0 & 0 & 0 & G_{x'} & 0 \\ 0 & 0 & 0 & 0 & G_{x'} \end{bmatrix} \quad (7)$$

$$G_y = \begin{bmatrix} G_{y'} & 0 & 0 & 0 & 0 \\ 0 & G_{y'} & 0 & 0 & 0 \\ 0 & 0 & G_{y'} & 0 & 0 \\ 0 & 0 & 0 & G_{y'} & 0 \\ 0 & 0 & 0 & 0 & G_{y'} \end{bmatrix}. \quad (8)$$

The objective function and constraints in (4) are convex, and thus, the convex program of (4) can be solved by a numerical convex program solver. We use the disciplined convex program-

ming modeling language cvx developed by Grant *et al.* [8], [9]. cvx allows a programmer to build up a convex program very naturally [as in the form of (4)] without regard to standard forms of canonical classes of convex programs. cvx then recasts the modeled convex program into a standard form and feeds it to SeDuMi, an interior-point solver that implements the self-dual embedding technique for optimization over self-dual homogenous cones [10], [11].

IV. RECONSTRUCTION RESULTS

We now present results of our reconstruction algorithm as applied to the TLE imaging problem described in Section II and analyze its performance. First, we examine the reconstruction of two natural data sets obtained during a sprite observation campaign in the summer of 2004 at Langmuir Laboratory near Socorro, NM [12], [13]. The first data set views the time evolution of an entire sprite, which grows and fades as a single entity for which we expect the total variation minimization scheme to work well. The second data set is zoomed in on one area of a sprite in which filamentary objects move rapidly across the screen. In this case, we also see that the total variation minimization scheme works suitably well. For the first data set, we present reconstruction results for both the l_1 -norm and l_2 -norm versions of the total variation objective function to motivate the use of the l_1 -norm version.

A. Methodology

Our procedure for testing the reconstruction technique is as follows.

- 1) Obtain a sequence of P 2-D images, where $15 \leq P \leq 30$.
- 2) Spatially subsample each of the P images down to $N_p \times N_p$ pixels.
- 3) Form horizontal and vertical projections of each subsampled image and a single integrated image that is the sum of all the subsampled images in the sequence. The projections will act as projection data, and the integrated image will act as a single video camera frame whose integration time spans the duration of the projection data.
- 4) Apply the reconstruction technique to the projection data and integrated image to form P “reconstructed” 2-D high-speed frames.
- 5) Compare the sequence of P reconstructed 2-D high-speed frames to the original sequence of P subsampled images to evaluate the technique’s performance.

A few notes on the aforementioned process need to be made. First, it should be noted that we start with 2-D images merely so that we can have something to compare our reconstructed 2-D high-speed frames to. This helps us evaluate the performance of our technique: any difference between the reconstructed high-speed frames and the original subsampled images is identified as errors in the high-speed frames. Note that the original sequence of 2-D images is not an input to the reconstruction algorithm: Only the projection data and the integrated image are inputs to the algorithm. In an actual real-world use of our technique, one would start at the third step with a set of P pairs

of orthogonal projections and a single integrated image (i.e., a single 2-D video camera frame whose integration time spans the integration times of the P pairs of orthogonal projections) and would then produce a set of P reconstructed 2-D high-speed frames.

Second, it should be noted that in a real-world use of our technique, steps 3) and 4) represent the conversion of a single video camera frame into a set of P reconstructed high-speed frames. To perform the technique on a collection of video camera frames, one would repeat steps 3) and 4) for each video frame and each group of P pairs of orthogonal projections that corresponds to the video frame.

In what follows, we only perform steps 3) and 4) once for each different investigated data set. Thus, each investigation represents a scenario where a single video camera frame (not shown in the images below) is augmented by P pairs of orthogonal projections to produce P reconstructed high-speed frames.

B. Reconstruction of a Full Sprite

Results from the first data set are shown in Fig. 4. This set of images is of a falsely colored full sprite that occurred on August 9, 2004, 4:40:31 UT, captured in high resolution at a frame rate of 60 frames/s. A total of $P = 16$ frames were recorded. Thus, we simulate a scenario where $P = 16$ pairs of 60-frame/s orthogonal projection data are used to augment a 3.75-frame/s camera to achieve an effective increase of the camera's frame rate by a factor of $P = 16$. Recall that, from the point of view of testing our reconstruction technique, the absolute frame rates involved (3.75 and 60 frames/s) are arbitrary—it is the ratio of the frame rates ($P = 16$) that is the expression of our technique's value. Had these images been originally acquired at a frame rate of 16 000 frames/s, the reconstruction outcome would be no different; it would simply correspond to a scenario where 1000-frame/s 2-D video data are augmented by orthogonal projection data to produce 16 000-frame/s 2-D reconstructed high-speed frames.

Fig. 4 shows the original image, the ideally reconstructed $N_p \times N_p$ pixel resolution high-speed frame, and the high-speed frames reconstructed with both the l_2 -norm and l_1 -norm versions of the total variation objective function for each frame of the data set. In Fig. 4, it is the last two columns that should be compared to the second column for assessment of the performance of our reconstruction method. If our reconstruction technique were to work perfectly, the reconstructed high-speed frames in the third or fourth columns of Fig. 4 would be identical to the ideally reconstructed high-speed frames in the second column.

The use of the l_1 -norm in the total variation objective function easily outperforms the use of the l_2 -norm, as expected. The l_2 -norm penalizes edges too severely, causing the solution arrived at to appear blurred. The cost in using the l_1 -norm for superior reconstruction is the computational complexity: In our experience, reconstruction with the l_1 -norm requires twice as many iterations as does reconstruction with the l_2 -norm.

In Fig. 5, we compare our reconstructed high-speed frames to their ideally reconstructed counterparts by plotting error maps

defined as

$$E_{jk}^i = \frac{|R_{jk}^i - A_{jk}^i|}{|A_{jk}^i|} \quad (9)$$

where A_{jk}^i is the (j, k) element of the i th ideally reconstructed high-speed frame, R_{jk}^i is the (j, k) element of the i th reconstructed high-speed frame, and E_{jk}^i is the (j, k) element of the error map associated with the i th reconstructed high-speed frame. Note that the error maps represent the absolute error between pixels normalized by the ideal pixel value and thus can be expressed as a percentage.

In the error maps of Fig. 5, we see that most reconstructed pixels have a normalized error less than 20% of the “true” value of the pixel (obtained from the ideally reconstructed high-speed frame). The error only rises above 20% at a few locations in a few frames. For several of these locations, the error rises mainly due to the small magnitude of the true pixel value [the denominator of (9)]. Such errors caused mainly by low signal levels are harmless and can be ignored in most applications (unless the imaging is done for the specific purpose of identifying highly localized and dim features). Examples of these errors are the yellow spots occurring on the edges of frames 1–4 in Fig. 5 and possibly the interior points in frames 3 and 4.

The locations of more interesting errors can be seen in the sixth column of the second frame of Fig. 5. These errors are symptomatic of the inability of the reconstruction technique to reproduce the highly detailed region in the left “wing” of the sprite. In the ideally reconstructed high-speed frame, a ridged structure appears in the left wing of the sprite where two small vertical dark regions are sandwiched between three brighter regions. In the reconstructed high-speed frame, these two dark regions are blurred over, introducing significant error in the reconstruction of the second dark region (in column 6). We see, however, that the error does not exceed 50%. Another example of a significant error occurs on the bottom left edge of the sprite in the third frame.

Despite the presence of a few locations where significant point errors occur, we see that no isolated islands of poor reconstruction appear in the reconstructed high-speed frames. Each of the significant errors is due to a missed edge or a blurred over region in a highly detailed area of the sprite and is not due to a phantom feature introduced strictly in the course of the reconstruction. For the purposes of most applications which would involve scientific and/or engineering uses of the reconstructed high-speed frames, the lack of such phantom features is reassuring; the reconstructed high-speed frames can be taken to be rough representations of the actual scene under view.

C. Reconstruction of Sprite Filaments

Results from the second data set are shown in Fig. 6. This set of images is taken from a 1000-frame/s video sequence of a highly magnified sprite. A total of 14 images are present. These images exhibit the small-scale filamentary structure found in the lower altitude regions of sprites [14]. Unlike the first

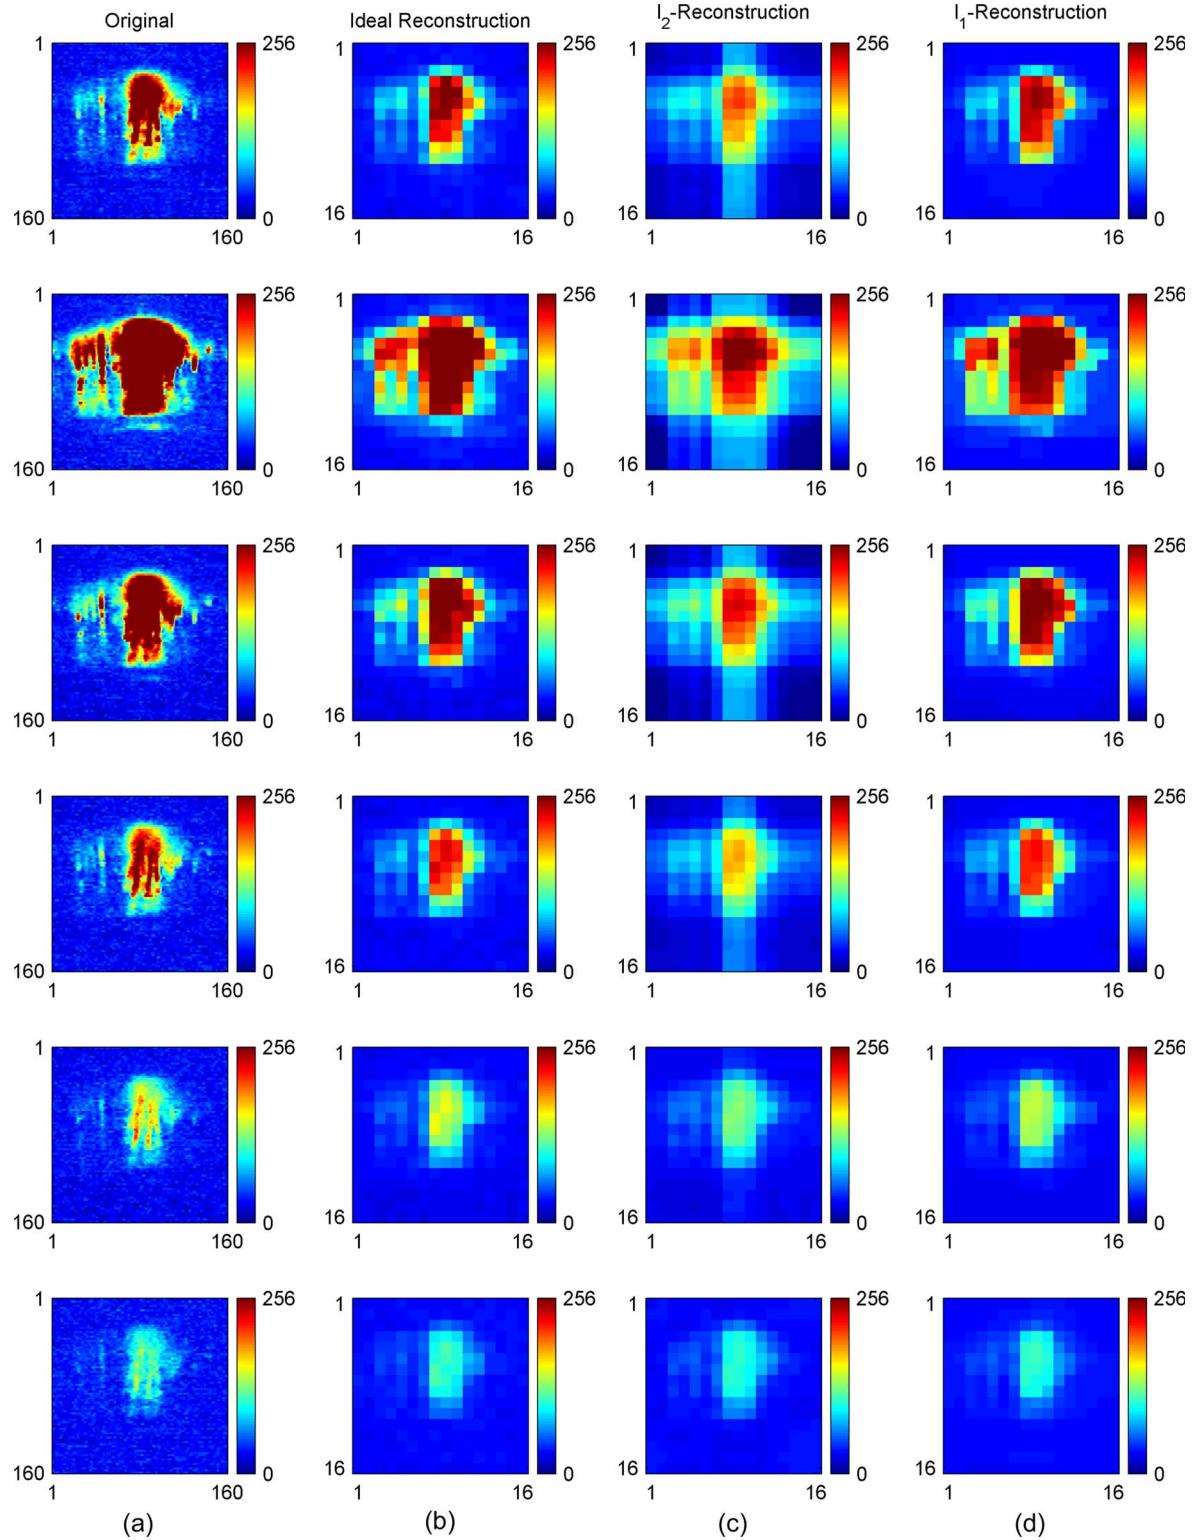


Fig. 4. Selected frames from a natural data set featuring a full sprite. The data set shown is taken from a sprite captured on August 9, 2004, 04:40:31 UT at Langmuir Laboratory, NM. The top row represents the first high-speed frame in the sequence, and subsequent rows represent subsequent high-speed frames. (a) Original (falsely colored) camera images from which the ideally reconstructed high-speed frames were generated. (b) Ideally reconstructed high-speed frames. (c) High-speed frames reconstructed using the l_2 -norm for the total variation objective. (d) High-speed frames reconstructed using the l_1 -norm for the total variation objective.

data set, these images do not exhibit a high degree of spatial smoothness but, instead, contain many small-scale structures (beads) clumped closely together, moving rapidly through the field of view.

Again, to test our reconstruction technique with this data set, we produce horizontal and vertical projections of each of the 14 images (which will act as the projection data), and we produce one integrated image that is the sum of the 14 images

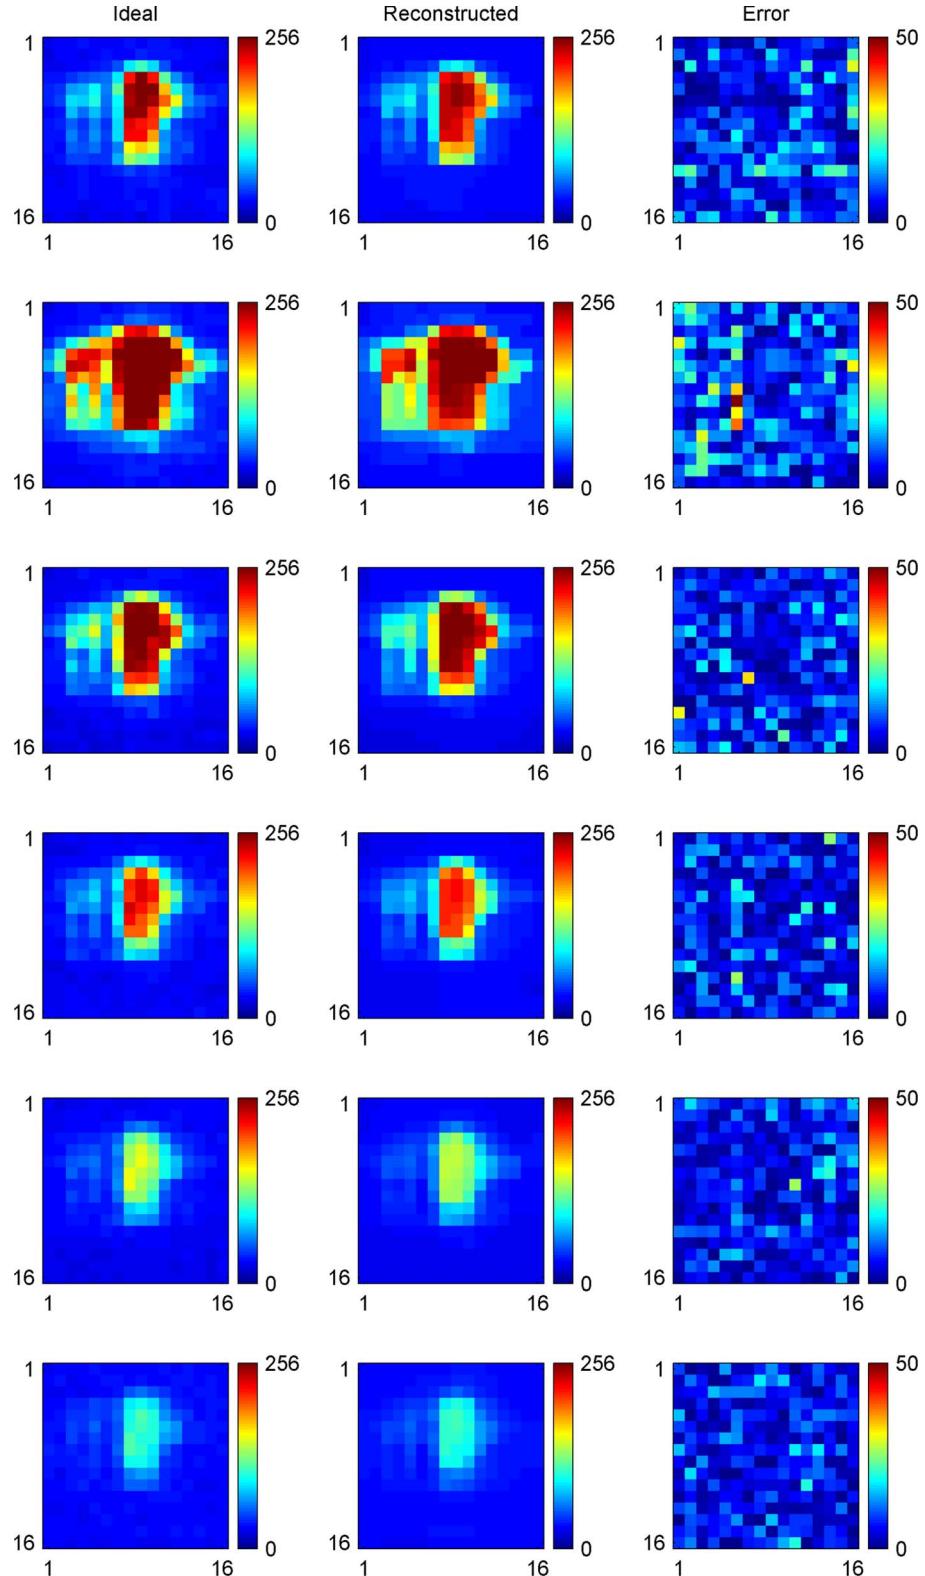


Fig. 5. Summary of reconstruction error for the data set of a full sprite captured on August 9, 2004, 04:40:31 UT. The top row represents the first high-speed frame in the sequence, and subsequent rows represent subsequent high-speed frames. (a) Ideally reconstructed high-speed frames. (b) Reconstructed high-speed frames. (c) Normalized absolute error maps (expressed as percentages).

(that will act as the single video camera frame whose integration time overlaps the occurrence of the 14 projection frames). For reference, we again create ideally reconstructed high-speed frames to which we can directly compare our re-

construction results. This time, our setup represents a scenario where 1000-frame/s projection data are used to augment a 71.4-frame/s camera to achieve an effective increase of the camera's frame rate by a factor of 14.

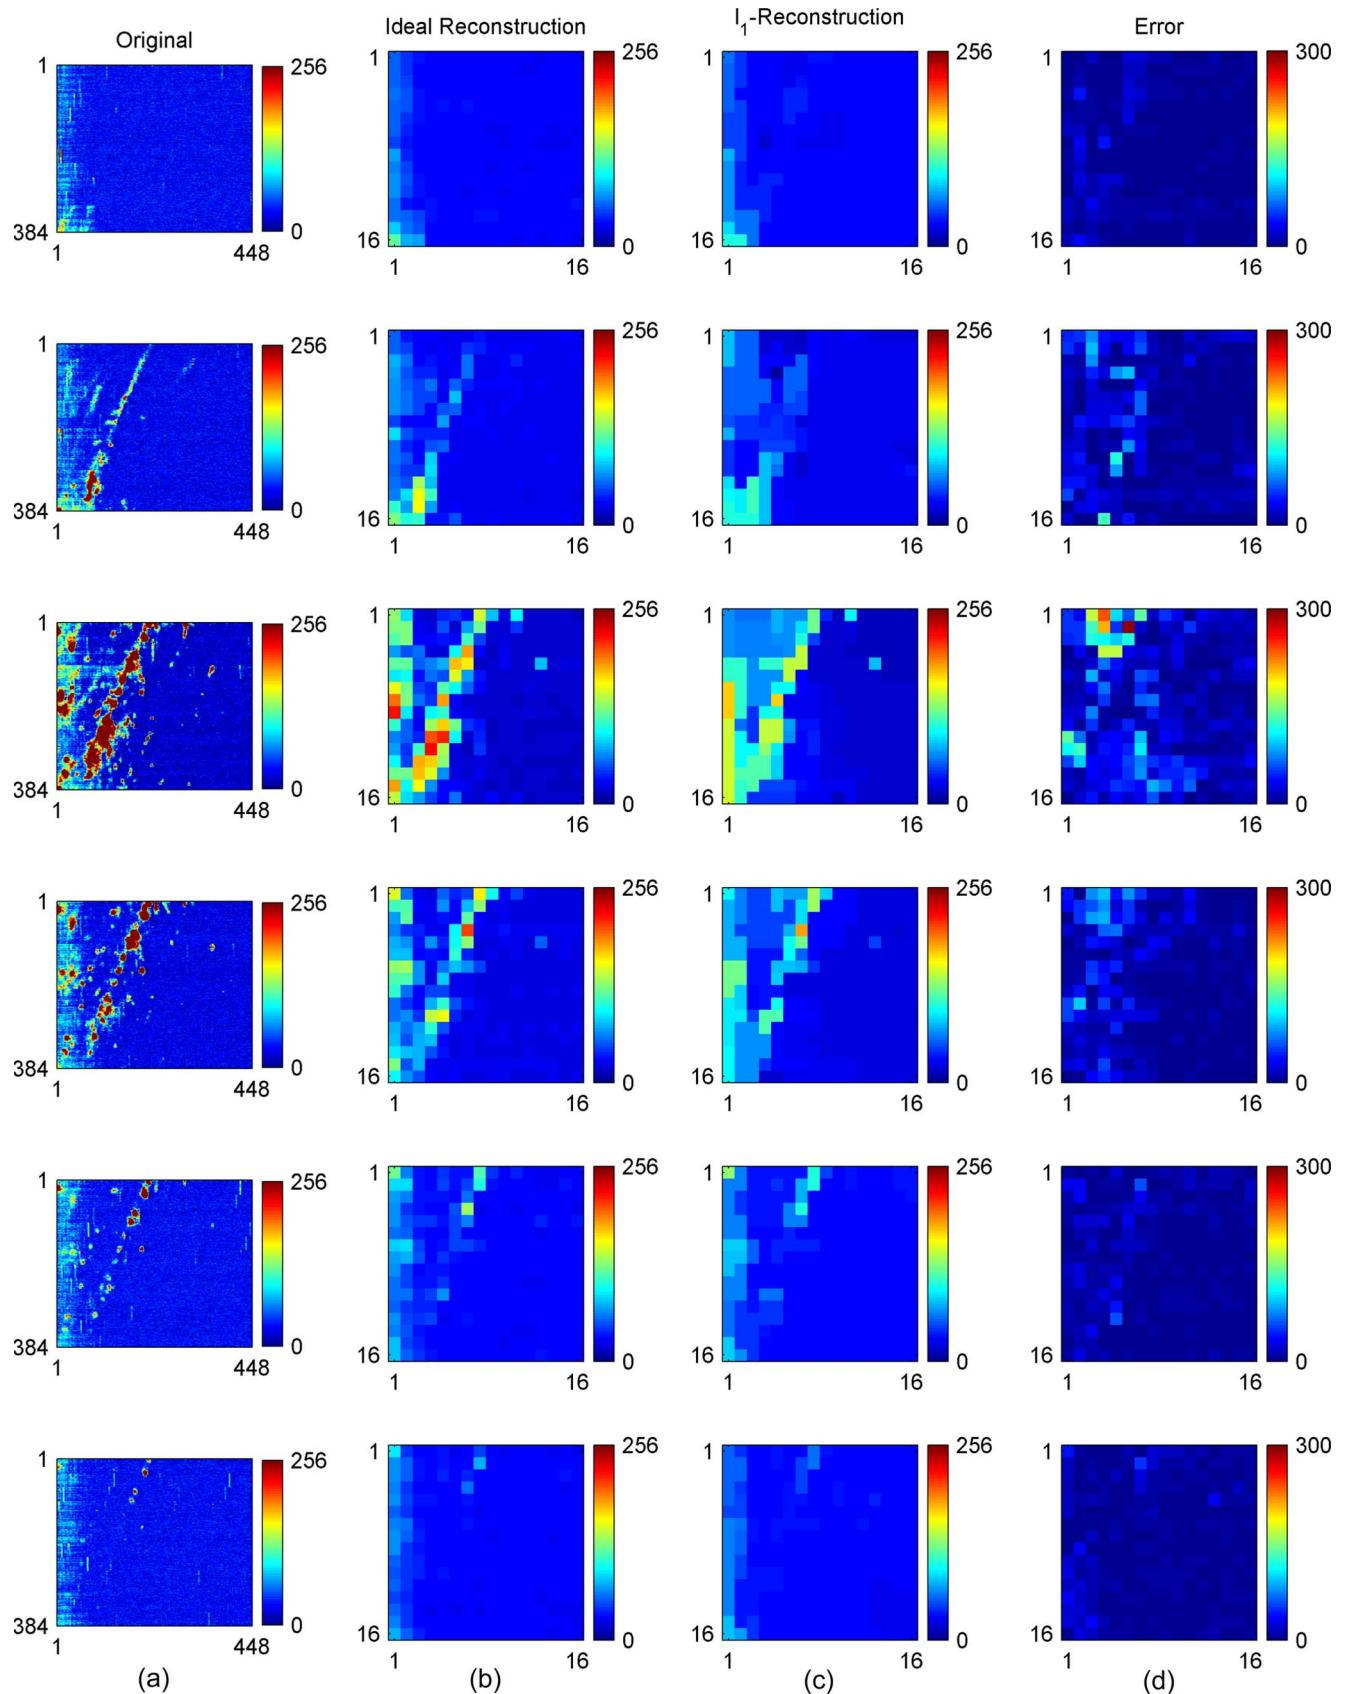


Fig. 6. Selected frames from a natural data set featuring the filamentary structure of a highly magnified sprite image. The data set shown is taken from a sprite captured on August 9, 2004, 04:29:17 UT at Langmuir Laboratory, New Mexico. The top row represents the first high-speed frame in the sequence, and subsequent rows represent subsequent high-speed frames. (a) Original (falsely colored) images from which the ideally reconstructed high-speed frames were generated. (b) Ideally reconstructed high-speed frames. (c) High-speed frames reconstructed using the l_1 -norm for the total variation objective. (d) Normalized absolute error maps.

At first thought, one would expect this type of image to challenge the minimum total variation reconstruction approach, which tends to blur out highly varying regions. As it turns out, however, we are helped by the fact that the low spatial resolution of the projection measurement devices does not resolve the individual beads and filaments very well. In other words, the ideally reconstructed $N_p \times N_p$ pixel high-speed frames are already blurred by low resolution, and the reconstruction technique (fundamentally unable to produce reconstructed high-speed frames that are superior to the ideally reconstructed high-speed frames) is thus relieved of the burden of capturing all the small-scale features in great detail.

In surveying Fig. 6, we see that the reconstruction performance is once again qualitatively satisfying. The basic structure captured in the ideally reconstructed high-speed frames is roughly preserved (with some blurring) in the reconstructed high-speed frames. Note that the color bars associated with the error maps in both Figs. 5 and 6 have greatly different scales (0%–50% in the former and 0%–300% in the latter). Comparison of the errors here with those for the previous data set suggests that the reconstruction error was, in fact, greater in the filamentary data set. However, as the error maps in Fig. 6 are mostly dark blue, we see that most reconstructed pixels in the filamentary data set have normalized absolute errors less than 50% (somewhat worse but yet still comparable to the first data set, where most pixels had normalized absolute errors less than 20%).

The largest meaningful pixel errors in Fig. 6 occur in the third high-speed frame. Comparing the third ideally reconstructed and reconstructed high-speed frames, we see that the reconstructed frames have shaded in the dark region between the bright left wall and the bright filament propagating upward and to the right (compare the top-left regions in both frames). The associated region in the third error map shows normalized absolute pixel errors in the range of 200%–300% over a space of about 7 pixels. These large errors are due to the tendency of the total variation minimization scheme to smooth over regions of high variance. Similar blurring appears in the same region in the preceding and following high-speed frames.

Aside from this one area of sharp intensity variation, most other pixel errors can be assigned to the same causes as discussed earlier, i.e., regions of low signal value (where the normalized errors are not meaningful) and poorly rendered edges (which tend to not upset the visual similarity, and hence the practical use in most applications, between ideally reconstructed and reconstructed high-speed frames).

D. Reconstruction of Hypothetical Beads

The set of images in the last section contained small-scale propagating “beads” that were deemed too small to be individually registered by the projection measurement device elements. From the point of view of reconstruction, we were satisfied in merely reconstructing the large-scale configuration of the beads and not the beads themselves. It is natural to ask how the total variation reconstruction scheme would perform if the beads were large enough to be individually registered by the projection measurement device elements. To investigate this

question, we constructed hypothetical data sets consisting of various configurations of bead arrangements, velocities, propagation directions, and brightnesses. In each data set, all the beads were roughly the size of a single projection measurement device element. Fig. 7 shows the reconstruction results for a particular data set in which 25 2-D high-speed frames are reconstructed from 25 projection frames and one integrated camera frame, representing a scenario where projection data are used to effectively increase the camera frame rate by a factor of 25.

The data set we are presenting was designed to showcase the following general results from all the hypothetical bead data sets we tried reconstructing.

- 1) Reconstruction is better for beads not propagating directly along a single projection orientation axis.
- 2) Well-separated beads are reconstructed better than nearly collocated beads.
- 3) Dimmer beads can sometimes be “masked” by brighter beads on the same row or column.
- 4) There is a tendency for the total variation scheme to cause beads to cast “shadows” across their rows and columns, sometimes creating “ghost” beads at the intersections of the rows and columns of other beads.

It should be noted that the hypothetical data set shown in Fig. 7 represents a particularly difficult data set for the total variation scheme to reconstruct. The data set consists of several discrete small-scale features moving in dissimilar directions, while the total variation scheme tends to smooth over discrete small-scale features and better captures large groups of small-scale features traveling in similar directions (like the beads in Fig. 6). Thus, for data sets similar to the hypothetical data set we are presenting, reconstruction results must be interpreted carefully. In practice, if video camera frames suggest a scenario similar to this hypothetical data set, one would want to consider minimizing some other objective function as the total variation objective function is not particularly appropriate in this situation.

The hypothetical bead data set used in Fig. 7 consists of four beads. Two beads propagate along only one projection orientation axis: the bead propagating right to left across the bottom of the field of view and the bead propagating vertically up the right side of the field of view. The other two beads start in the upper and lower left corners of the field of view and propagate diagonally across the field of view to the opposite corner in such a way that their paths pass very near each other in the middle of the field of view but do not intersect. The two diagonally propagating beads are of different brightness levels. For each original image in the sequence of images, a small amount of additive white Gaussian noise (of a variance one-twentieth of the brightness of the brightest bead) was added to help the total variation minimization scheme converge. Without the added noise, the total variation minimization scheme did not always converge for the hypothetical data sets we tried. The addition of a small amount of noise makes the data sets more realistic without affecting our ability to see how the reconstruction technique performs in different scenarios.

Fig. 7 shows eight selected images, ideally reconstructed high-speed frames, and reconstructed high-speed frames from

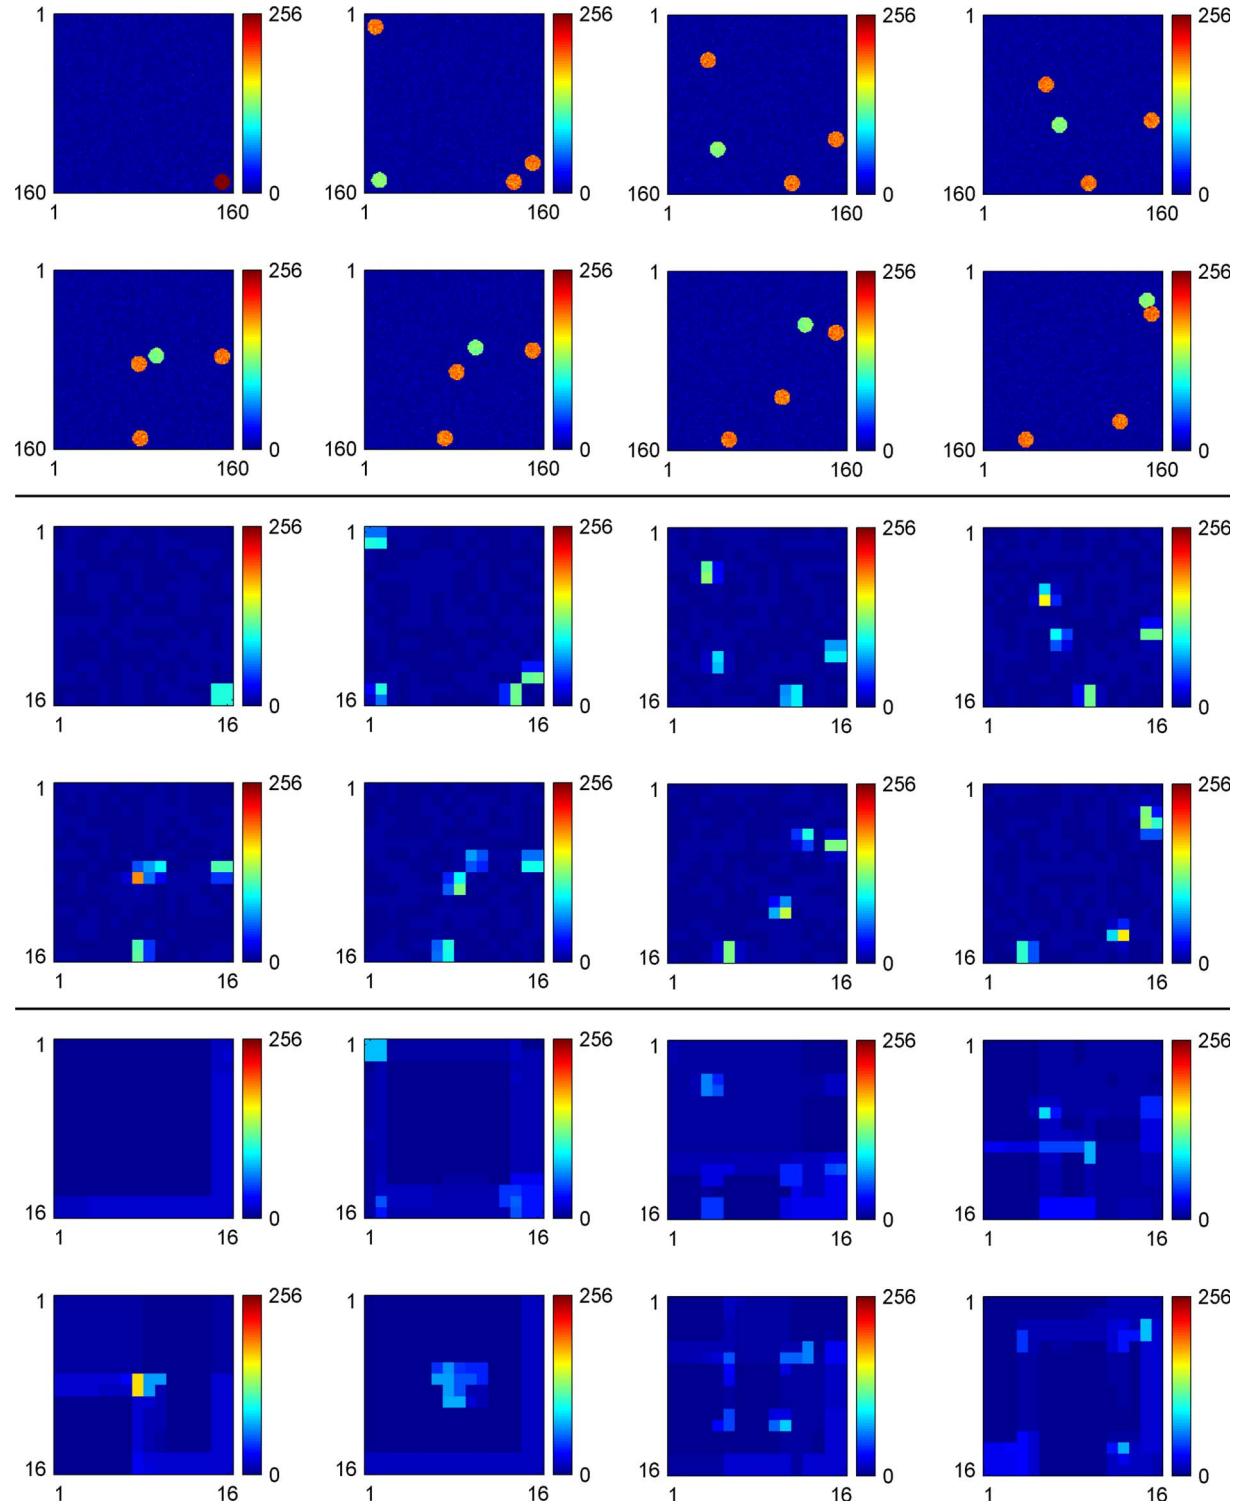


Fig. 7. Selected frames from an hypothetical data set featuring small propagating beads roughly the size of a projection measurement device element. The top panel shows the original hypothetical low-resolution images (with frames progressing from left to right, then top to bottom). The middle panel shows the ideally reconstructed high-speed frames associated with the images of the top panel. The bottom panel shows the reconstructed high-speed frames. Of the 25 high-speed frames in the data set, frames 1, 4, 8, 11, 14, 15, 18, and 21 are presented. These frames are a representative sample of the 17 absent high-speed frames.

the 25-image data set used. The eight frames are typical of the other 17 absent frames and were chosen to display the strengths and weaknesses of the reconstruction technique. Normalized error maps are not shown as they do not add insight beyond what can be gained from simple visual comparison of the ideally reconstructed and reconstructed high-speed frames.

We discuss Fig. 7 frame by frame. In the first high-speed frame, we see a good example of shadow casting. The single bead in the lower right corner (which is really the superposition of the horizontally and vertically propagating beads) is reconstructed as a dim line along both the last row and column of the reconstructed high-speed frame. The bead itself is not

even distinct from its shadow, although this is not always the case (as in the next high-speed frame, where the two beads in the lower right corner are distinct from their shadows in the reconstructed high-speed frame).

In the second high-speed frame, we see that the two beads in the lower right corner (which have now separated themselves from each other) are reconstructed as a single blur. Nearly collocated beads are commonly reconstructed as single blurs, while well-separated ones are often reconstructed more accurately. The two diagonally moving beads have also entered the frame, but the dimmer (green) bead has been masked in the reconstructed frame by either the other diagonally moving bead (in the top left) or the horizontally moving bead (in the bottom right).

In the third high-speed frame, we see three examples of ghost beads at the intersections of the shadows of real beads. These same three ghost beads appear again in the fourth high-speed frame. We note in the fourth high-speed frame that one of the ghost beads appears, in fact, brighter than the physical beads.

The fifth and sixth high-speed frames are examples of good and bad collocated bead reconstructions, respectively. In the fifth high-speed frame, the brighter bead appears to the left of the dimmer bead in the reconstructed frame, which is true to reality. In the sixth high-speed frame, the arrangement of the two nearly collocated beads is not well reconstructed. The amorphous blob in the middle of the sixth reconstructed high-speed frame is actually three beads, the upper left of which is a ghost bead.

In the seventh and eighth high-speed frames, we see two more ghost beads, both shadows of the two diagonally propagating beads.

Through all the reconstructed high-speed frames, we see that the horizontally and vertically propagating beads are usually not as distinctly reconstructed as the diagonally propagating beads.

The earlier discussion has focused on examples of nonphysical artifacts of reconstruction. We note, however, that in all the high-speed frames, the diagonally moving beads are always present in the reconstructed high-speed frames (except for the one example in which the dim bead is masked by brighter beads on the same row or column). Thus, the task of the observer in interpreting the reconstructed high-speed frames of Fig. 7 is mainly to distinguish between ghost beads and physical beads with the understanding that dim beads in some reconstructed high-speed frames may briefly “disappear” due to masking by brighter beads. In practice, if the video frames of a data set suggest a scenario more like Fig. 7 (highly discrete small-scale features moving in dissimilar directions) than Fig. 4 (a single large-scale feature) or Fig. 6 (a large configuration of small-scale features moving in similar directions), then interpretation of the reconstructed high-speed frames must be done carefully.

V. CONCLUSION

A technique is presented for the reconstruction of a higher frame-rate representation of a dynamic 2-D scene from continuous measurements of two orthogonal projections of the scene. This reconstruction is accomplished by modeling the

observation of the scene as an underdetermined linear observation process. Reconstruction of the higher speed frames is performed by minimizing a total variation objective function subject to data fidelity constraints. A natural application of this technique to the imaging of TLEs in the Earth’s atmosphere is reviewed.

The performance of the reconstruction technique is demonstrated on three data sets. The first data set is generated from a sequence of wide-field-of-view images of a sprite that was observed from Langmuir Laboratory, New Mexico, in August 2004. Application of the method to this data set demonstrates the ability of the technique to capture the time evolution of large-scale features in a dynamic 2-D scene.

The second data set is generated from a sequence of narrow-field-of-view images of a different sprite observed during the same evening, which contained small-scale sprite structures (beads) propagating across the field of view. The reconstruction technique captures the large-scale configuration of the beads but fails to register the individual beads as they were smaller than the size of the projection measurement device elements.

The third data set is a hypothetical set of images containing a number of beads propagating across the field of view, where each bead is roughly the size of a projection measurement device element. The set of resulting reconstructed high-speed frames reveals that the reconstruction technique performs better on small-scale structures that are well separated and not propagating strictly along only one of the projection orientation axes. They also reveal that dim features may sometimes be “masked” by brighter features located on the same projection axis and that most reconstructed features cast “shadows” along each projection axis that often form “ghost” features at their interactions.

There are several practical aspects of high-speed frame reconstruction which are beyond the scope of this paper and which serve as avenues for further research. Such issues include handling pixel misregistration between the projection measurement devices and the video camera frames, correcting for differences in the optical sensitivities of the projection measurement devices and the video camera, and accounting for variance in the optical responses of the individual elements of each projection measurement device.

ACKNOWLEDGMENT

The authors would like to thank R. A. Marshall for providing the sprite-image sequences used in this paper.

REFERENCES

- [1] R. A. Marshall, R. T. Newsome, and U. S. Inan, “Photometric imaging using orthogonal linear photometer arrays,” *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 11, pp. 3885–3893, Nov. 2008.
- [2] Y. Yu and L. Carin, “Three-dimensional Bayesian inversion with application to subsurface sensing,” *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 5, pp. 1258–1270, May 2007.
- [3] D. Wen, Y. Yuan, J. Ou, K. Zhang, and K. Liu, “A hybrid reconstruction algorithm for 3-D ionospheric tomography,” *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 6, pp. 1733–1739, Jun. 2008.
- [4] S. A. Cummer, N. Jaugey, J. Li, W. A. Lyons, T. E. Nelson, and E. A. Gerken, “Submillisecond imaging of sprite development and structure,” *Geophys. Res. Lett.*, vol. 33, no. 4, p. L04104, Feb. 2006. DOI:10.1029/2005GL024969.

- [5] M. G. McHarg, H. C. Stenbaek-Nielsen, and T. Kammae, "Observations of streamer formation in sprites," *Geophys. Res. Lett.*, vol. 34, no. 6, p. L06 804, Mar. 2007. DOI: 10.1029/2006GL027854.
- [6] M. G. McHarg, R. K. Haaland, D. Moudry, and H. C. Stenbaek-Nielsen, "Altitude-time development of sprites," *J. Geophys. Res.*, vol. 107, no. A11, p. 1364, Nov. 2002. DOI: 10.1029/2001JA000283.
- [7] W. C. Karl, "Regularization in image restoration and reconstruction," in *Handbook of Image and Video Processing*, A. Bovik, Ed. San Francisco, CA: Academic, 2000, pp. 141–160.
- [8] M. Grant, S. Boyd, and Y. Ye, "Disciplined convex programming," in *Global Optimization: From Theory to Implementation (Nonconvex Optimization and its Application Series)*, L. Liberti and N. Maculan, Eds. Dordrecht, The Netherlands: Kluwer, 2005.
- [9] M. Grant, S. Boyd, and Y. Ye, *cvx User's Guide, Version 1.0*, Feb. 27, 2007. [cited 2007 May 11]. [Online]. Available: <http://www.stanford.edu/~boyd/cvx/>
- [10] J. F. Sturm, *Implementation of Interior Point Methods for Mixed Semidefinite and Second Order Cone Optimization Problems*, Aug. 2002, Center Econ. Res., Tilburg Univ., EconPapers 73.
- [11] J. F. Sturm, "Using SeDuMi 1.02, a matlab toolbox for optimization over symmetric cones," *Optim. Methods Softw.*, vol. 11/12, pp. 625–653, 1999.
- [12] R. A. Marshall and U. S. Inan, "High-speed telescopic imaging of sprites," *Geophys. Res. Lett.*, vol. 32, no. 5, p. L05 804, Mar. 2005. DOI: 10.1029/2004GL021988.
- [13] R. A. Marshall and U. S. Inan, "High-speed measurements of small-scale features in sprites: Sizes and lifetimes," *Radio Sci.*, vol. 41, no. 6, p. RS6 S43, Oct. 2006. DOI: 10.1029/2005RS003353.
- [14] E. A. Gerken and U. S. Inan, "Observations of decameter-scale morphologies in sprites," *J. Atmos. Sol.-Terr. Phys.*, vol. 65, no. 5, pp. 567–572, Mar. 2003.



Robert Newsome (S'09) was born in Knoxville, TN, on June 28, 1982. He received the B.S.E.E. degree from the University of Texas, Austin, in 2005, and the M.S.E.E. degree from Stanford University, Stanford, CA, in 2007, where he is currently working toward the Ph.D. degree in electrical engineering in the Space, Telecommunications, and Radioscience Laboratory, Department of Electrical Engineering, under Prof. U. Inan.

His research interests include the measurement and characterization of optical signatures of ionospheric processes and their application toward ionospheric process modeling. Such ionospheric processes include transient luminous events, e.g., sprites, and charged-particle precipitation.

Mr. Newsome is a member of Tau Beta Pi and Eta Kappa Nu.



Umran Inan (S'76–M'77–SM'99–F'06) was born in Turkey on December 28, 1950. He received the B.S. and M.S. degrees in electrical engineering from the Middle East Technical University, Ankara, Turkey, in 1972 and 1973, respectively, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 1977.

He is currently a Professor of electrical engineering with Stanford University, where he serves as the Director of the Space, Telecommunications, and Radioscience Laboratory, Department of Electrical Engineering. He actively conducts research in electromagnetic waves in plasmas, lightning discharges, ionospheric physics, and very low frequency remote sensing. He has served as the Principal Ph.D. Dissertation Advisor for 28 students.

Dr. Inan is a member of Tau Beta Pi, Sigma Xi, and the Electromagnetics Academy. He is a Fellow of the American Geophysical Union. He is the recipient of the 2007 Stanford University Allan V. Cox Medal for Faculty Excellence in Fostering Undergraduate Research and the 1998 recipient of the Stanford University Tau Beta Pi Award for Excellence in Undergraduate Teaching. He has served as the Chair of the U.S. National Committee, International Union of Radio Science (URSI) and the International Chair of Commission H (Waves in Plasmas) of URSI.