

ENHANCING THE SPATIAL RESOLUTION OF SOLAR CORONAGRAPH OBSERVATIONS USING DYNAMIC IMAGING

T. S. ZACCHEO AND M. KAROVSKA

Harvard-Smithsonian Center for Astrophysics, Cambridge, MA 02138

J. W. COOK, R. A. HOWARD, G. E. BRUECKNER, AND C. M. KORENDYKE

O. O. Hulburt Center for Space Research, Naval Research Laboratory, Washington, DC 20375-5352

AND

R. SCHWENN

Max-Planck-Institut für Aeronomie, 37189, Katlenburg-Lindau, Germany

Received 1996 March 11; accepted 1996 June 10

ABSTRACT

The Large Angle Spectrometric Coronagraph (LASCO) C1 coronagraph on board the *Solar and Heliospheric Observatory* (SOHO) is designed to image the corona from 1.1 to 3.0 R_{\odot} . The resolution of C1 is defined by the size of its CCD pixels, which correspond to 5"6, and not by the diffraction limit of the optical system, which may be as small as 3". The resolution of C1 can be improved using the technique of "dynamic imaging"—the process of acquiring successive images of the same scene using sub-pixel displacements of the steerable primary mirror. We developed a technique we call the fractional pixel restoration (FPR) algorithm that utilizes these observations to construct an image with improved resolution. Simulations were used to test this algorithm and to explore its limitations. We also applied the direct co-addition and FPR algorithms to laboratory preflight images of a wire mesh grid. These results show that the resolution of the C1 coronagraph can be significantly enhanced, even in the presence of noise and modest differences between successive images. In some cases, the results can even reach the diffraction limit of the telescope.

Subject headings: instrumentation: miscellaneous — space vehicles — Sun: corona —
 techniques: image processing

1. INTRODUCTION

The *Solar and Heliospheric Observatory* (SOHO), launched on 1995 December 2, is a scientific spacecraft collaboratively developed by ESA and NASA for exploring solar phenomena. Its primary objectives are to perform extensive helioseismology studies and to examine the processes that drive the heating of the corona and acceleration of the solar wind. These studies will be accomplished by collecting and analyzing data obtained from a variety of instruments on board the satellite. One such instrument is the Large Angle Spectrometric Coronagraph (LASCO), designed to observe the solar corona from 1.1 to 32 R_{\odot} (Brueckner et al. 1995). LASCO contains three individual coronagraphs, C1, C2, and C3, which produce overlapping views of the corona. The highest spatial resolution will be achieved with the C1 coronagraph, which will produce images of the inner corona from 1.1 to 3.0 R_{\odot} . These data will provide a close link to the disk observations of other instruments in the SOHO complement.

C1, which is described in detail in Brueckner et al. (1995), is a mirror coronagraph unlike the conventional externally occulted coronagraphs flown earlier on *OSO-7*, *Skylab*, P78-1 (SOLWIND coronagraph), or *SMM*. This design was chosen in order to view (unvignetted by an external occulter) as close to the solar limb as possible. Figure 1 shows a conceptual diagram of the C1 coronagraph. Light entering the aperture A0 is imaged by the objective mirror M1 on the annular field mirror M2. The central circular hole in M2 acts as an occulter for the solar disk. The remaining optical path is comprised of a collimating mirror M3; a Lyot stop A1; a Fabry-Perot filter for obtaining line

observations in Fe x, Fe xiv, Ca xv, and H α , as well as white-light images; and a telephoto lens.

The images are recorded using a Tektronix CCD detector containing 1024 \times 1024 pixels, each of which is 21 μm or 5"6 on a side. Even though the diffraction limit of the C1 telescope is $\approx 3''$, the nominal resolution of the C1 coronagraph is limited by the angular size corresponding to 2 CCD pixels or $\approx 11''$. This results in images that are under-sampled by a factor of approximately 4, since according to the Nyquist criterion, a pixel size of 1"5 or smaller is needed to obtain diffraction-limited observations. This limitation is the result of a compromise between the desired field of view and the size of available CCD detectors. Improved resolution of C1 would be helpful in studying the spatial characteristics of the fine structures in the lower corona and their evolution.

The resolution of C1 may be improved using a technique called "dynamic imaging"—a process of acquiring and co-adding multiple images of an identical scene separated by fractions of a pixel. Dynamic imaging with the C1 coronagraph will be carried out by taking a series of images separated by a fraction of a CCD pixel. These fractional pixel steps will be introduced by moving the steerable primary mirror M1 (see Fig. 1). The mirror's orthogonal tip, tilt, and fore-and-aft motions are driven by low-voltage piezo-electric transducers. Typically, subpixel displacements of one-half pixel will be used to produce four independent images, and occasionally one-fourth pixel steps may be used to acquire 16 observations. This is based on the assumption that a given scene will remain relatively stationary over the entire exposure and readout period for 4

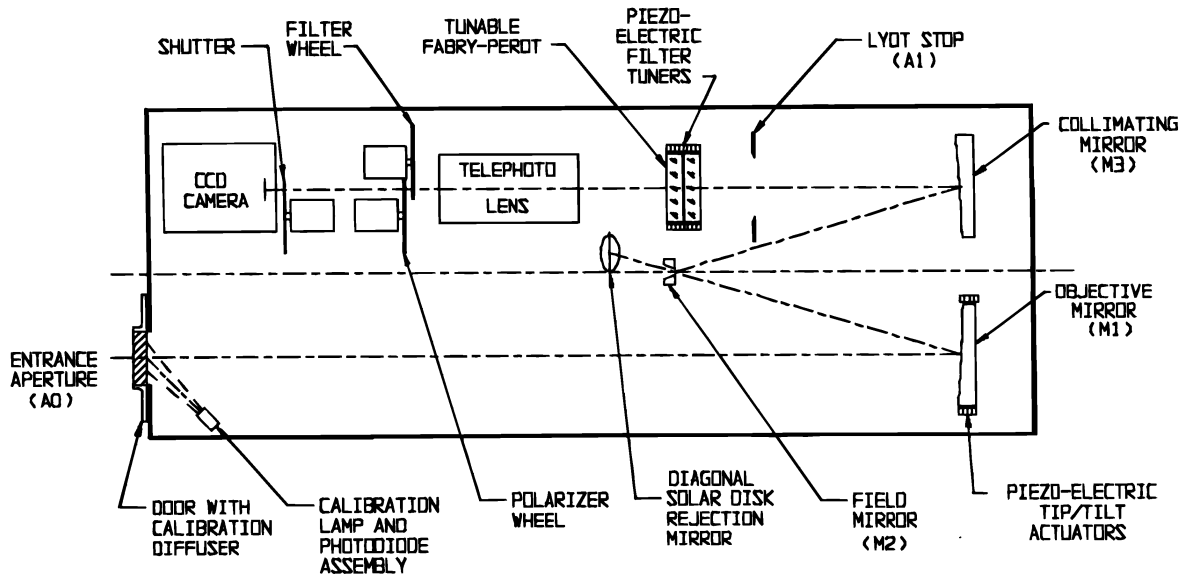


FIG. 1.—Conceptual diagram of the LASCO C1 coronagraph. The M1 objective mirror can be tilted at controlled subpixel angles in the two orthogonal axes of the CCD to acquire a dynamic imaging sequence of observations.

images (approximately 10 minutes). In order to obtain 16 images, both the individual exposure and readout intervals must be minimized to overcome the expected changes in the scene from the evolution of coronal structures. Dynamic imaging will be implemented in a limited number of the C1 observations and is not the default mode of operation.

In this paper, we present various methods for enhancing the spatial resolution of C1 observations using dynamic imaging. We provide an overview of co-addition techniques and describe a method that we have developed for enhancing the resolution of LASCO C1 by combining substepped images. Finally, we discuss the results of applying various co-addition techniques to sets of simulated and preflight laboratory observations.

2. METHODS FOR CONSTRUCTING SUBPIXEL ESTIMATES

Various algorithms have been developed to enhance the resolution of undersampled systems by constructing composite images from observations taken fractions of a pixel apart (subpixel images). The simplest method is direct co-addition. In this case, each of the undersampled images are resampled on a finer grid, aligned with respect to a common reference, and then added to estimate or reconstruct a subpixel image of the desired scene. The resulting resolution of these estimates is directly related to the number of co-added images and the fractional displacement between each. In most instances, each subpixel estimate contains the same number of pixels as there are in the entire data set. An obvious exception is when the input images are acquired without any fractional displacements. In this case, the final resolution can only equal that of the observed data.

Even though estimates constructed using direct co-addition contain more spatial frequency information than the subpixel images, they are not equivalent to a properly sampled observation. They represent the desired image degraded by the subsampling process. Various algorithms have been developed to deal with this problem as well as other aspects of undersampled data sets. In general, these algorithms can be subdivided into two classes depending on whether the input data are acquired at evenly spaced intervals. In the case of evenly spaced observations, the data are

normally co-added, and the distortions due to under-sampling are removed. Restorations based on data acquired with arbitrary substeps are produced using various methods for estimating the value at each subpixel position in the newly formed space. Rigorous treatments of this general problem and various solutions are presented in Tsai & Huang (1984), Kim, Bose, & Valenzuela (1990), and Ur & Gross (1992). A number of other techniques have been developed and applied to *Hubble Space Telescope (HST)* imagery (Dickinson & Fosbury 1995; Hook & Adorf 1995). The algorithms that are currently being used in *HST* research are the Richardson-Lucy co-addition algorithm (Hook & Lucy 1992), the method of projections on convex sets (POCS) (Adorf 1995a), and “subpixel dithering” (Adorf 1995b). In this paper, we describe a method that we developed for constructing enhanced estimates from multiple substepped images acquired with the C1 coronagraph. Our technique, “fractional pixel restorations” (FPR), is similar to those described by Ur & Gross (1992) and Adorf (1995b).

2.1. Fractional Pixel Restorations

The FPR algorithm is a two-step process. First, the images are directly co-added and then the known distortion due to the undersampling, and possibly the distortions associated with the optics, are removed using a standard deconvolution technique. The rationale for selecting this algorithm over others is that the steerable mirror in the C1 telescope was designed so that a sequence of images separated by multiples of one-fourth pixel ($1/4$) intervals could be acquired of the same scene. This makes it simple to obtain either four substepped images taken with one-half pixel displacements or 16 images with one-fourth pixel displacements. This procedure, unlike others, does not require preprocessing of the data to align or coregister the images, since the subpixel displacements are known and the spacecraft jitter is nominal.

The basic principles of the FPR algorithm can be modeled by representing each observation as a function of the desired result. In order to illustrate this process, we will consider the one-dimensional case in which two overlapping images are acquired one-half pixel apart. In this case,

the diffraction-limited observation, $f(x)$, represents the true object convolved with the point-spread function (PSF) of the optical system. It is assumed that the spatial frequency content of $f(x)$ is limited by the PSF and therefore can be fully described by a discrete waveform sampled at the Nyquist rate or twice the highest data frequency. The function $f(k)$ is the discrete representation of $f(x)$ sampled at this rate. The change of indices indicate that $f(k)$ is valid only for $k = 0, \dots, N$. The function $f(k)$ also represents the desired output of FPR in the case in which the PSF is unknown. However, if the PSF is known, we would prefer to estimate the true object itself.

In the one-half pixel case, two independent sets of samples are acquired at half the Nyquist rate. Each sample represents the integral of $f(x)$ or the sum of $f(k)$ over the area defined by a given pixel. The only difference between the two sequences is that their sampling periods have a one-half pixel phase difference. The discrete elements in each of the resulting waveforms are given by

$$i_1(k) = f(2k) + f(2k + 1)$$

and

$$i_2(k) = f(2k + 1) + f(2k + 2), \quad (1)$$

where $f(k)$ is the desired result. There are several schemes for combining these two functions. The method that we have chosen, based on a minimum distortion criteria, is as follows: First, both $i_1(k)$ and $i_2(k)$ are upsampled by a factor of 2 without interpolation. This produces two new images that are represented by the original data elements in $i_1(k)$ and $i_2(k)$ separated from their neighbors by single pixels with values of 0. These upsampled results are then aligned with respect to a common reference and are added to form an enhanced representation. The result of this co-addition procedure is

$$g(k) = f(k) + f(k + 1), \quad (2)$$

where $g(k)$ is an intermediate result containing $2M$ samples. In this case, M is the number of samples in both $i_1(k)$ and $i_2(k)$. Equation (2) represents the convolution in the spatial domain of $f(k)$ with a square two-sample PSF. In the spatial frequency domain, it represents the multiplication of the desired spectrum with a shifted $\sin(\omega)/\omega$ function. This PSF produces an estimate with a degraded frequency response and eliminates any information at the highest frequencies because of the fact that $\sin(\omega)/\omega = 0$ at $-\pi$ and π .

Similar models to that given by equation (2) can be derived for both one-dimensional data containing four images taken one-fourth pixel apart and two-dimensional observations. The one-fourth pixel results are represented by $f(k)$ convolved with a square four-sample PSF. This intermediate sequence is given by

$$g(k) = f(k) + f(k + 1) + f(k + 2) + f(k + 3) \quad (3)$$

in the spatial domain and is represented in the spatial frequency domain by the multiplication of the discrete Fourier transform of $f(k)$ with $\sin(2\omega)/\omega$. Multiplying by $\sin(2\omega)/\omega$ eliminates any frequency information at $\omega = \pm\pi/2$ and $\pm\pi$.

2.2. Two-dimensional Reconstructions

Direct co-addition of evenly spaced two-dimensional data is performed by applying the one-dimensional models described above to a two-dimensional array. In the two-

dimensional case, each column is upsampled followed by each newly formed row. The upsampled images are aligned with respect to both x and y and are then added to form a fractional pixel result. These results, as in the one-dimensional cases, are also represented by the convolution of the desired image, $f(x, y)$, with a blurring function, $h(x, y)$, whose coefficients are given by the outer product between two one-dimensional PSFs. In our simulations we have assumed that an equal number of samples are acquired in both the x - and y -directions. For example, in the one-half pixel cases, four images are obtained by offsetting the input by $(0, 0)$, $(0, \frac{1}{2})$, $(\frac{1}{2}, 0)$ and $(\frac{1}{2}, \frac{1}{2})$ pixels in x and y , respectively. The corresponding two-dimensional PSF for this case is then given by the outer product of the PSF described by equation (2) with itself. Even though we have chosen to use a symmetric sampling procedure, this does not preclude the possibility of constructing an estimate with different x and y spacings. As an example, one could acquire eight images that are separated by one-fourth pixel in the x -direction and one-half pixel in the y -direction. The resulting PSF would then be represented by the outer product of a one-dimensional one-fourth pixel PSF with a one-half pixel blurring function.

In the above analysis, it has been assumed that the observed scene is stationary during multiple exposures, that the fractional pixel shifts are exact, and that the data collection environment is noiseless. Any deviations from these conditions will result in less than optimal restorations. Errors will be introduced by detector noise, any mispositioning of the mirror, satellite jitter, changes in the position and intensities of the scene, and nonuniformities in the pixels. The stationary effects associated with these phenomena, such as the average variation in quantum efficiency across a pixel, can be incorporated into the linear models described above, whereas the nonstationary processes must be analyzed on a pixel-by-pixel basis or treated like signal-dependent noise sources.

Each of the estimates described above is degraded by the subsampling process. These degraded observations can in general be represented by

$$g(x, y) = h(x, y) * f(x, y), \quad (4)$$

where $h(x, y)$ is once again the subsampling system function and $*$ is the convolution operator. Standard deconvolution techniques can be used to minimize these effects. This approach is outlined in the block diagram shown in Figure 2. The deconvolution technique that we have adopted is the Richardson-Lucy (RL) algorithm (Richardson 1972; Lucy 1974). The RL algorithm is a nonlinear technique that produces positively constrained estimates, which have maximum likelihood given that the noise statistics are Poisson. Each new RL estimate is described by

$$\hat{f}(k)^{n+1} = \hat{f}(k)^n \sum_{i=0}^N h(i-k) \frac{d(i)}{\hat{g}(i)^n}, \quad (5)$$

where $d(i)$ is the data, $\hat{g}(i)$ is an estimate of $g(i)$, $n + 1$ and n denote the present and previous estimates, and the sum represents a convolution. In essence, each new estimate is simply a weighted version of the last set of values. We have chosen to adopt this approach based on the superior performance of the RL algorithm and the assumption that the noise associated with the LASCO C1 data is dominated by photon statistics.

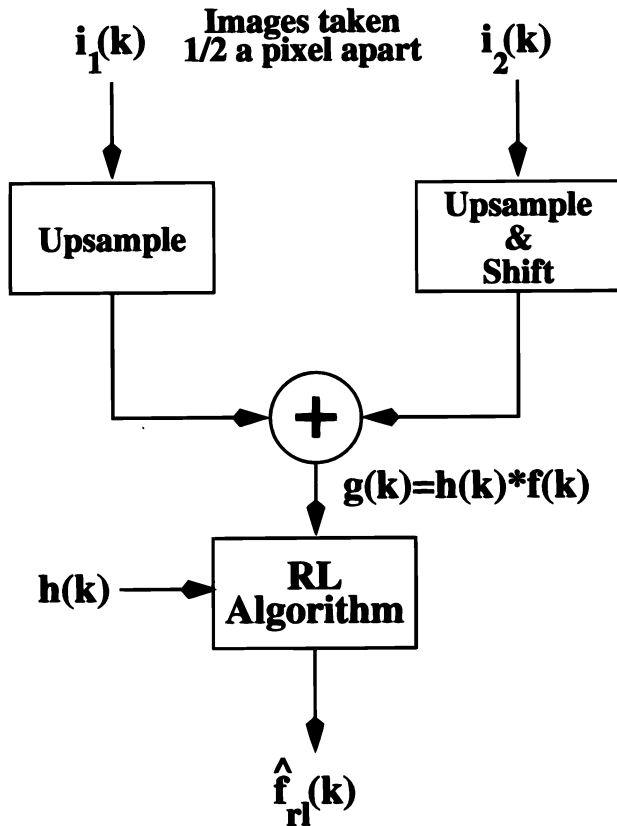


FIG. 2.—Block diagram of the fractional pixel restoration algorithm. The two samples, $i_1(k)$ and $i_2(k)$, are upsampled, aligned, and added to form the intermediate result, $g(k)$. Any distortions introduced by the sampling process or PSF are removed using the RL algorithm. The same approach can be applied to one-dimensional samples taken one-fourth pixel apart or two-dimensional images.

Deconvolution can be used to minimize not only the effects of subsampling but also those associated with a known PSF. Both sets of distortion terms can be dealt with simultaneously by replacing $h(x, y)$ in equation (4) with a composite function given by the convolution of $h(x, y)$ with the optical PSF. The extent to which this technique will improve the resolution depends on how well the optical PSF can be described in the newly formed frequency space. Little or no benefit will come from a PSF that is inadequately sampled. This may be the case for one-half pixel C1 data. However, using a PSF will most likely improve the data taken at one-fourth pixel steps.

3. SIMULATIONS

We performed extensive simulations to assess the performance of the FPR algorithm on two-dimensional data. Each of these simulations was implemented using geometric patterns to emulate the fine structure of the corona and a Gaussian model to mimic the PSF of the telescope. Initially, a simple geometric test image was used to validate the algorithm. The same image was then used to explore the effects of noise, non-equally spaced substeps, scene motion, and changes in scene intensity. Finally, we applied the FPR algorithm to a test object that was developed based on a simplified model of the C1 CCD and an average model of the corona.

The first set of simulations demonstrates the ability of the FPR algorithm to construct subpixel estimates from data

sampled at a rate similar to that set by the LASCO C1 CCD. In this case, it was assumed that the CCD under-samples the diffraction limit of C1 by a factor of 4. Simulated observations were constructed based on the test pattern shown in Figure 3a. This image contains 256×256 pixels and represents a scene sampled at 4 times the resolution of the CCD. Initially, a diffraction-limited view was formed by convolving this test image with a two-dimensional Gaussian PSF (full width half-maximum of 2 pixels) to mimic the effects of the C1 optics. Each observation was then constructed by shifting the diffraction-limited result by an integer number of pixels in both the x - and y -directions and subsampling it by a factor of 4. Shifting the input image by 1 pixel produced a shift of one-fourth pixel in the observations. Four observations separated by one-half pixel were generated by moving the diffraction-limited image by either (0, 0), (0, 2), (2, 0), and (2, 2) pixels in the x - and y -directions. A similar set of 16 observations separated by one-fourth pixel intervals was formed by shifting the input to each of the 16 possible locations and down-sampling the results. Figure 3b shows the diffraction-limited input, and the image in Figure 3c represents a typical under-sampled observation.

The FPR algorithm was applied to these undersampled observations to form subpixel estimates of the original scene. The one-half pixel observations were directly co-added by upsampling the images by a factor of 2, aligning each of them with respect to one another and adding them to form a composite image. In the one-fourth pixel case, the 16 images were upsampled by a factor of 4, shifted back to their original positions, and added. The result of co-adding four images is shown in Figure 3d. The final subpixel estimates were then formed by applying 25 iterations of the RL algorithm to the co-added images. The PSFs used by the RL algorithm were generated based on the known distortion due to sampling and, in some cases, an estimate of the optical PSF. Three sample restorations are illustrated in Figure 4. Figure 4a shows the diffraction-limited observation and is provided as a reference. Figure 4b is a typical one-half pixel restoration constructed using four images, and Figure 4c is a one-fourth pixel restoration using 16 images. The example in Figure 4d is also a one-fourth pixel estimate constructed using both the distortion due to sampling and an estimate of the optical PSF. In this case, the optical PSF was represented by the original two-dimensional Gaussian plus 10% noise (Gaussian). The expanded view in the center of Figure 4 illustrates the extent to which each approach was able to reconstruct the diffraction-limited observation. It demonstrates that the one-half pixel restoration technique improves the resolution, but it is unable to represent adequately the true object. This is not unexpected since the resulting sampling rate is still one-half that of the desired rate. The two one-fourth pixel restorations show that one can not only approach the optical diffraction limit but with some prior knowledge of the PSF can, in a noiseless environment, exceed it.

Figure 5 represents a set of slices taken through the pointlike objects (impulses) shown in Figure 4. Figure 5a describes the Gaussian-like profile of the diffraction-limited data, and Figures 5b, 5c, and 5d are the corresponding slices for each of the three estimates. Each of the estimates shown in this figure was scaled so that they contained the same number of counts as the diffraction-limited object. These




FIG. 3.—Two-dimensional simulated data and the results of direct co-addition. (a) Simulated test image. (b) Diffraction-limited observation described by the convolution of (a) with a two-dimensional Gaussian PSF. (c) Typical undersampled observation. (d) Restoration formed by directly co-adding four observations taken one-half pixel apart.

plots demonstrate that the one-half pixel estimate produces a nonsymmetric result, while the two one-fourth pixel estimates are either nearly identical to the diffraction-limited observation or approach that of the true object, an impulse.

We also used the same test pattern to assess the effects of noise and various error conditions encountered during the data collection process. Figure 6 illustrates the effects of adding noise. In each case, varying amounts of Poisson or Gaussian noise were added to the undersampled observations. The images in the upper row represent one of four noisy observations used in each of the trials, and those in the lower row depict the resulting one-half pixel reconstructions. Poisson noise was introduced by setting the maximum intensity of each image and then replacing each of its pixels with a Poisson random variable whose mean was equal to that of its original intensity. Gaussian data were generated by simply adding zero-mean Gaussian random variables with variances equaling some fixed percentage of maximum pixel intensity. The percent noise figures represent the minimum amount of additive noise in each case. This figure demonstrates that either 1% Poisson or 1% Gaussian noise has little effect on the fractional pixel restorations, whereas 10% Poisson or 10% Gaussian noise degrades the results. The most severe degradations occur with 10% Gaussian noise. Even though the basic structures are still visible, the pointlike objects and most of the details are lost. Some of these simulations overestimate the

expected photon and detector noise but illustrate what might happen if there were enough random variations in intensity or motion between successive subpixel images.

Figure 7 illustrates the type of errors that occur due to imperfectly positioned subpixel images, motion in the scene, and changes in intensity in the scene over the observing period. The effects of non-evenly spaced substeps were explored by shifting the test pattern by an additional amount in both the x and y -directions and then co-adding the results as though they were sampled at exactly one-half pixel intervals. Only one-eighth of the original test image was used because of the amount of memory needed to perform the additional shift operations. Small placement errors on the order of 5% produced almost no visible differences, while 20% placement errors (Fig. 7b) produce noticeable distortions. In the restoration shown in Figure 7b, the pointlike objects are elongated and the rays not only start to collapse but also exhibit a hashed pattern.

Figure 7c demonstrates the effects due to image motion. This restoration was constructed using two subsampled observations formed from the standard test pattern and a second set based on the same image shifted by a fixed amount in the x -direction. The estimate shown in Figure 7c represents the results of moving the test pattern by one-half pixel. Small shifts on the order of one-eighth pixel seemed to have no noticeable effects on the reconstruction process, whereas larger shifts blurred the restoration in the direction of predefined motion.

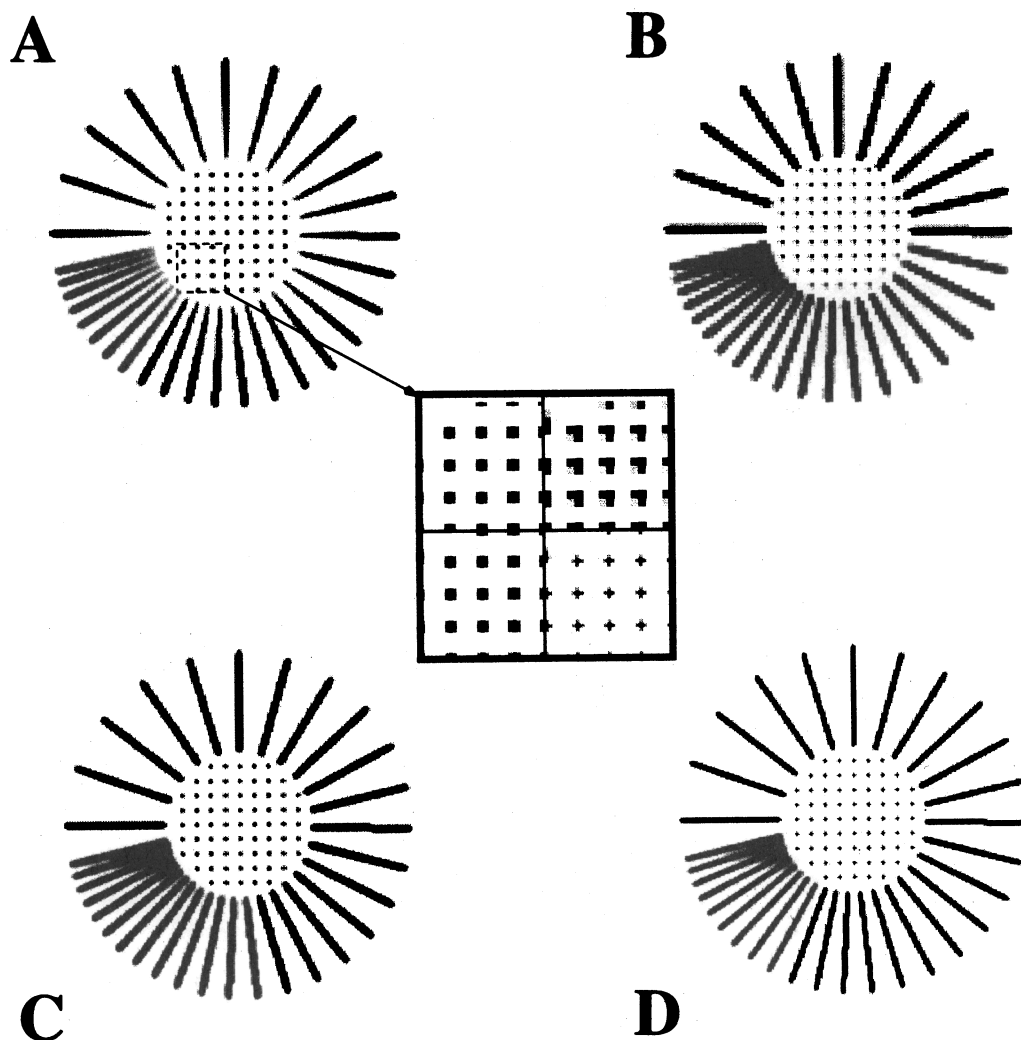


FIG. 4.—Two-dimensional fractional pixel restorations. (a) Diffraction-limited object. (b) Typical one-half pixel restoration derived from four observations and the known distortion due to sampling. (c) One-fourth pixel restoration based on 16 observations and the sampling PSF. (d) One-fourth pixel restoration formed by deconvolving the distortions due to sampling and an estimate of the optical PSF.

Two pairs of observations were also used to study the effects of changes in intensity. The first pair was generated using the standard test pattern, and the second was formed by dividing each pixel value by a fixed constant. The reconstructions shown in Figure 7*d* were generated using a difference in intensity of 50%. Changes in intensity have little effect on the pointlike objects but decrease the amount of modulation between the rays and introduce significant variations in some of the larger structures.

Finally, we developed a more realistic set of simulations in order to estimate the FPR algorithm's ability to detect fine structures in the presence of an average corona, instrumental stray light, and detector noise. This was accomplished using an analytical function to describe the average background corona and both Poisson and Gaussian random variables to represent the known noise sources. The test pattern shown in Figure 8*a* (Plate 11) was constructed using a radially symmetric function, whose intensity decreased as $r^{-3.5}$, to represent both the average background corona and the instrumental stray light intensity. Based on preflight laboratory measurements, the stray light component is roughly comparable in brightness and radial decrease to the average background corona. Therefore, for simplicity, we parameterize them similarly. Seven equally

spaced pairs of rays were used to depict the structures of interest. The simulated corona at $1.1 R_{\odot}$ was assigned a value of 100,000 counts or 60% of the estimated detector full-well capacity, and each pair of lines, which decreased in intensity radially as a function of r^{-2} , were set to 1, 2, 5, 10, 20, 50, and 100% of this maximum. Four simulated observations were then constructed by subsampling this image and adding the appropriate amount of both Poisson and Gaussian noise. Poisson random variables based on the intensity of each pixel were used to represent photon noise, and Gaussian random variables with $\sigma = 200$ counts were added to describe the noise due to the electronics, dark current, and stray light variations. Finally, these subsampled images were quantized to 16 bits, co-added, and deconvolved. A typical observation is shown in Figure 8*b*, and the resulting restoration is shown in Figures 8*c*. Figure 8*d* represents a difference image constructed by subtracting the known corona from the reconstructed image. It demonstrates that fine structures with intensities below 2% of the average corona will most likely go undetected and that in order to perform any rigorous analysis, the contrast should exceed 10%. This conclusion is valid only in the inner regions of the corona or in regions in which the number of observed photoelectric events is near the full-well capacity




FIG. 8.—Simulated C1 corona observations. (a) Original diffraction-limited image of an average corona plus seven equally spaced pairs of spokes with maximum intensities of 1, 2, 5, 10, 20, and 100% of the corona at $1.1 R_{\odot}$. (b) Typical simulated C1 observation generated by subsampling (a) and adding both Poisson and Gaussian noise. (c) Restoration based on four observations separated by one-half pixel. (d) Difference image of (c) minus the average corona. This difference image was reversed for display purposes. The dark line along the edge of the corona illustrates that the restoration algorithm is unable to fully recover all the high frequency information along the steep edge of the mask.

ZACCHEO et al. (see 471, 1063)

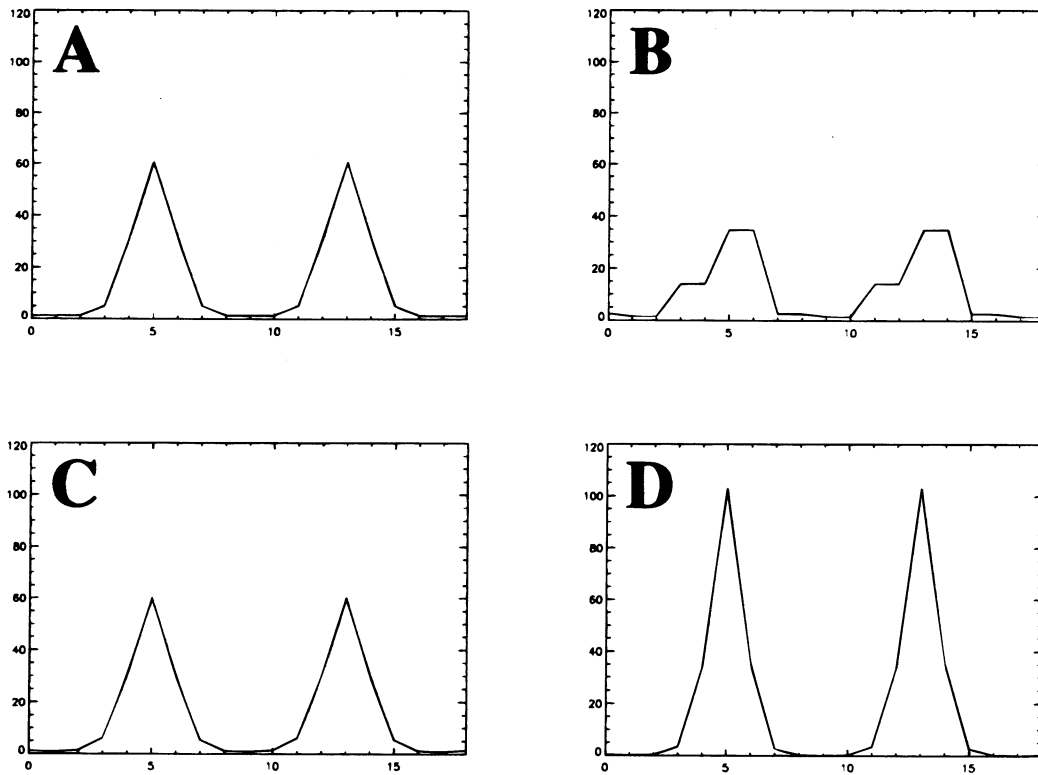


FIG. 5.—Slices through the two-dimensional fractional pixel restorations. (a) Diffraction-limited object. (b) Typical one-half pixel restoration. (c) One-fourth pixel restoration. (d) One-fourth pixel restorations obtain using the known distortions due to sampling and an estimate of the optical PSF.

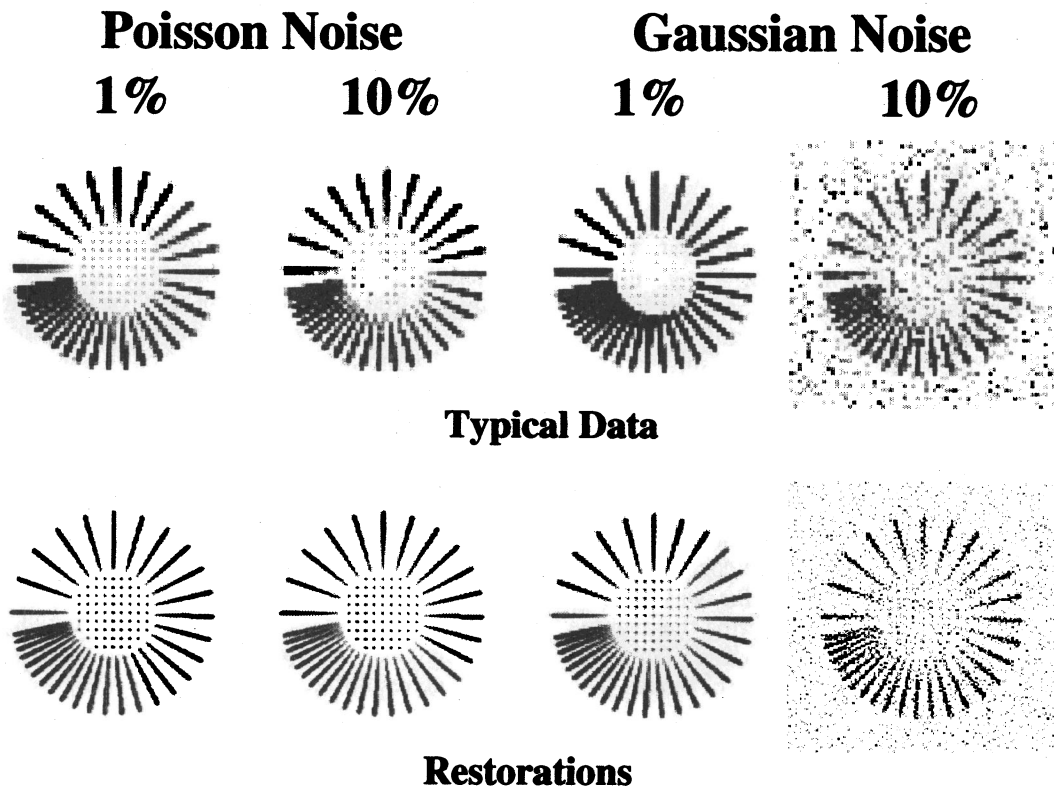


FIG. 6.—Effect of noise on the fractional pixel reconstruction algorithm. The upper row represents typical one-half pixel observations corrupted by varying amounts of either Poisson or Gaussian noise. The lower set of images describes the corresponding restorations. The Poisson data were generated based on images with maximum intensities of either 10,000 or 100 photons, respectively. The Gaussian data were formed by adding Gaussian random variables with a fixed standard deviation equaling either 1% or 10% of the maximum pixel values.




FIG. 7.—Effects of substepping errors, object motion and changes in intensity on one-half pixel imaging. (a) Reference image constructed using four observations separated by exactly one-half pixel. (b) Resulting reconstruction formed from four observations with 20% placement error. (c) Restoration based on observations with one-half pixel of motion in the x -direction. (d) Results of a 50% change in intensity over the observing period.

of the CCD. Similar simulations predict that the contrast ratio between the fine structures and the average background must be significantly higher in order to detect these features in the outer corona (e.g., close to $3.0 R_{\odot}$), where the pixel intensities are expected to be on the order of several thousand counts.

4. APPLICATION TO LABORATORY OBSERVATIONS

A set of four dynamic C1 images was obtained during laboratory tests conducted in early 1993. These images were acquired by moving the M1 mirror of the flight instrument (see Fig. 1). A CCD detector with $25 \mu\text{m}$ pixels was used instead of the flight CCD, whose pixels are $21 \mu\text{m}$ on a side. The target was a fine wire mesh grid placed in front of a piece of graph paper (which is noticeably out of focus in the images). We used the FPR algorithm to improve the spatial resolution of these laboratory images. These results are shown in Figure 9 (Plate 12), which contains three 256×256 sections of the complete 1024×1024 images. The upper right-hand corner of each image shows a portion of the circular hole in the M2 mirror. The wire grid is well focused, while the out-of-focus graph paper can be seen in squares of around half the size of the image dimension. Figure 9a shows one of the four original dynamic imaging observations. This image was provided for comparison with

the direct co-addition result (Fig. 9b) and the FPR estimate (Fig. 9c). The final panel shows a magnified view of identical subsections of these three images.

While there is an improvement seen between the original image and the directly co-added image, it is clear that the FPR image contains more spatial information. Both the apparent wire width and the nodes where wires cross are better resolved in the FPR image than in the co-added image.

5. SUMMARY

In this work, we explored various techniques for combining multiple substepped images and developed a suitable algorithm (FPR) for use with LASCO C1 data. We describe the results of numerous tests conducted using various simulated data sets. These simulations not only demonstrate the feasibility of the FPR algorithm but also illustrate its tolerance to known sources of errors and noise. They show that the resolution of the C1 coronagraph can be enhanced, even in the presence of significant noise and modest errors, and provide an estimate of the amount of contrast required to detect fine solar structures. The spatial resolution of the C1 coronagraph can be increased by almost a factor of 2 using four one-half pixel substepped images and under special circumstances can approach the resolution limits set by the




FIG. 9.—Laboratory LASCO C1 dynamic imaging observations of a fine wire mesh grid. (a) One of the four original images. (b) Result of directly co-adding the four original images. (c) The FPR estimate. The final image in the lower right contains magnified views of identical subsections from the three previous images.

ZACCHEO et al. (see 471, 1065)

C1 optics. Finally, we tested the FPR algorithm using pre-flight laboratory images of a fine wire mesh grid. These results show that the FPR algorithm, applied to actual data, produces images with enhanced resolution that are superior to the original data and to the directly co-added results.

The *SOHO* satellite is presently in its operational orbit around the L1 Lagrangian point between the Sun and Earth. In the coming months, the LASCO C1 coronagraph will perform an observational program that should include dynamic imaging sequences to observe the fine structures of the inner corona. These images will yield new information for determining, for example, the characteristic spatial scales of the small structure in coronal holes and streamers,

the filling factors, and the morphology of various structures as a function of their height above the limb.

This project is supported by a grant from the Naval Research Laboratory to the Harvard-Smithsonian Center for Astrophysics. M. K. is a member of the AXAF Science Center and is partially supported by NASA, contract NAS 8-39073. J. W. C., R. A. H., G. E. B., and C. M. K. are supported by NASA and by the Office of Naval Research. M. K. and T. S. Z. would like to thank Hans-Martin Adorf and Robert Fosbury of the Space Telescope-European Coordinating Facility for their lists of references on sub-pixel restoration techniques.

REFERENCES

- Adorf, H.-M. 1995a, ST-ECF Newsletter, 22, 17
 ———. 1995b, ST-ECF Newsletter, 23, 19
 Brueckner, G. E., et al. 1995, Sol. Phys., 162, 357
 Dickinson, M., & Fosbury, B. 1995, ST-ECF Newsletter, 22, 14
 Hook, R., & Adorf, H.-M. 1995, in *Calibrating Hubble Space Telescope: Post Servicing Mission*, ed. A. Koratkar & C. Leitherer (Baltimore: Space Telescope Science Institute), 341
 Hook, R., & Lucy, L. B. 1992, ST-ECF Newsletter, 17, 10
 Kim, S. P., Bose, N. K., & Valenzuela, H. M. 1990, IEEE Trans. on Acoustics, Speech, and Signal Processing, 38(4), 1013
 Lucy, L. B. 1974, AJ, 79, 745
 Richardson, W. H. 1972, J. Opt. Soc. Am., 62, 55
 Tsai, R. Y., & Huang, T. S. 1984, in *Advances in Computer Vision and Image Processing*, Vol. 1, ed. T. S. Huang (Greenwich: JAI Press), 317
 Ur, H., & Gross, D. 1992, CVGIP, 54(2), 181