

Calibration procedures for charge-coupled device x-ray detectors

S. L. Barna^{a)}

Department of Physics, Princeton University, Princeton, New Jersey 08540

M. W. Tate and S. M. Gruner^{b)}

Department of Physics, Cornell University, Ithaca, New York 14853-2501

E. F. Eikenberry

Department of Pathology, Robert Wood Johnson Medical School, Piscataway, New Jersey 08854

(Received 2 February 1999; accepted for publication 29 March 1999)

Calibration procedures are described for use with electronic x-ray detectors, with an emphasis on detectors based on fiber-optically coupled charge-coupled devices. Methods are detailed for removing spurious events, pixel pedestals, dark-current, spatial distortion, and intensity response variations for both small-angle and wide-angle applications. The accuracy of the calibration procedures is discussed. © 1999 American Institute of Physics. [S0034-6748(99)02207-8]

I. INTRODUCTION

Charge-coupled device (CCD) x-ray area detectors have been shown to have quantum limited performance.¹ In these detectors, a fluorescent screen converts the incident x-ray pattern to light which is then optically coupled to the CCD image sensor. But high sensitivity alone is insufficient to make a useable quantitative measuring device because non-uniformities and distortions in the response of the detector cause the recorded signal to be not simply proportional to the incident dose. For example, the response of a CCD detector varies with x-ray energy, angle of incidence, and position on the face of the detector. Imaging optics and intensifiers introduce distortion into the image that needs to be corrected. Detector pixels each have an associated pedestal level and dark current that must be removed through background subtraction. If these nonuniformities and distortions are stable and reproducible, they can be calibrated out of the final digital image, in which case the ultimate accuracy of the detector will be limited by the accuracy of the calibration procedures. For many diffraction images, it is also important to remove spurious signals, colloquially called “zingers,” that result from radioactive decays within the materials of the detector and from cosmic rays. This article details calibration procedures, describing both the measurement of the calibrations and the application of the corrections to data images.

Calibration procedures described in the literature differ with respect to the ways in which the calibration data are acquired, the methods used to generate and apply the calibrations and the accuracy which results.²⁻⁷ The appropriateness of a given calibration procedure varies with the x-ray application, the detector, and the desired accuracy. Thus, a high accuracy calibration procedure appropriate for small-area detectors used for monochromatic small-angle x-ray scattering may be relatively simple because, in this case,

there is little need to calibrate the dependence of either the x-ray energy or the angle of incidence. In our experience, calibration to an accuracy of a few percent in dose is easy, but calibration to a few tenths of a percent is difficult; very few, if any, of the procedures in use are accurate to small fractions of a per cent. The goal in this article is to describe procedures that have accuracy approaching 0.5%, and to do so for detectors that may be used for wide angle diffraction at various x-ray wavelengths. Although the procedures were designed to calibrate phosphor-coupled CCD detectors for varied applications, most of the procedures are applicable to other types of detectors.

II. ZINGER REMOVAL

Zingers are unwanted, localized, random events in an image that are caused by cosmic rays, decay of radioactive isotopes present in the material of the detector itself, or, at synchrotron sources, stray hard radiation from sources other than the direct beam. The intensity of zingers can vary from very bright and quite obvious to barely statistically discernible. Since these phenomena occur randomly in time, long integrations have proportionately more zingers than short ones. Since zingers occur randomly in position as well as time, they can be identified and removed by crosscomparing two or more nominally identical exposures to remove differences that exceed expected statistical thresholds. In experiments with copious redundant information or ones in which most of the image area does not contain data, such as many crystallographic measurements, the burden of taking two exposures at each position outweighs the benefits of dezingering. However, dezingering can be extremely important in small angle or diffuse x-ray scattering work.

The simplest procedure for dezingering is to compare the two images pixel by pixel against an arbitrary fixed cutoff for the maximum allowed difference. If an unacceptable comparison is found, the higher value is replaced by the lower value. This is easy and fast to apply, but has several drawbacks. Foremost is the fact that this procedure does not take into account the expected increase in the variance of the

^{a)}Current address: Photobit Corp., 135 N. Los Robles Ave., 7th floor, Pasadena, CA 91101.

^{b)}Author to whom correspondence should be addressed; electronic mail: smg26@cornell.edu.

pixel values as the intensity rises; pixels with large signal can have differences greater than the cutoff by chance alone. Such a procedure could be tuned to work well in dark parts of the image or in bright parts, but a single cutoff will not work in both. Furthermore, this procedure introduces a statistical bias by always using the lower value to replace the higher value; however, if the number of zingers is small, this effect is negligible. Finally, the procedure is not fully automatic. It will usually be necessary to adjust the cutoff value manually to produce the best results for a given image. Nonetheless, this procedure can work well in some situations.

In order to develop a more flexible method of zinger removal, a knowledge of the statistics of the detector's response to x-ray exposure is needed. Apart from zingers, the variance of a given pixel, σ_{pix}^2 , will have two components: an effective background variance, σ_0^2 , resulting from readout noise, dark current, and other detector effects; and the Poisson variance, proportional to the number of incident x rays for the given pixel, n_{pix} , resulting from counting statistics,

$$\sigma_{\text{pix}}^2 = \sigma_0^2 + g^2 n_{\text{pix}},$$

where, assuming a detector with no point spread, g is the gain of the detector to convert x-ray quanta into detector counts. [The more realistic case of a detector with a nonzero point spread function (PSF) is discussed in the Appendix. For such a detector, the PSF smoothes the variance and the apparent gain of the detector, g , is less than the actual gain. This does not change the functional form of the statistical distribution, however.] Measurements have confirmed that σ_{pix}^2 is a linear function of the recorded dose for detectors in which the recorded signal varies linearly with the incident dose.

The difference between two images, a and b , of the same diffraction pattern will have twice the variance of a single image:

$$\sigma_{a-b}^2 = \sigma_a^2 + \sigma_b^2 = 2(\sigma_0^2 + g^2 n_{\text{pix}}).$$

More practically, one has the dose in counts, $I_{\text{pix}} = gn_{\text{pix}}$, so this formula becomes

$$\sigma_{a-b}^2 = \sigma_a^2 + \sigma_b^2 = 2(\sigma_0^2 + gI_{\text{pix}}).$$

This quantity gives the expected variance of a pixel as a function of dose and can be used to construct a lookup table of allowed cutoff values at each level of x-ray exposure.

First, σ_0^2 is estimated by histogramming pixel dose versus number of occurrences of that dose within an image. The first maximum in this distribution is taken to be the background level and the width of the distribution gives σ_0^2 . In practice, the distribution is smoothed prior to calculations and the width is found by using only points below the mode, since above the mode, this distribution is skewed away from Gaussian by the signal in the image.

It is necessary to scale the two images for any difference in incident x-ray dose. One can perform a linear regression of the recorded dose in one image versus the recorded dose of the other to obtain the scale factor. Note that the application of a scale factor to one image will change the apparent gain of that image and hence the expected variance. In our

laboratory, we have chosen to take pairs of images with nominally identical exposures, therefore giving a scale factor near unity, though variations between such paired images of up to 10% in apparent exposure have been observed. The small change in variance that results in this case can be neglected in the dezingering process.

At this point, a second linear regression is performed, comparing $(a-b)^2$ to the mean of the dose, $(a+b)/2 = I_{\text{pix}}$. This regression has a slope of g and an intercept of σ_0^2 , giving a functional form for σ_{a-b}^2 versus I_{pix} . From this, a lookup table can be constructed for the allowed difference as a function of dose. We have chosen to set the cutoff value to be $5\sigma_{a-b}$, from which we expect 1 pixel in a 10^6 pixel image to exceed the cutoff by chance alone. By comparison, 66 pixels would exceed a four standard deviation cutoff by chance.

Once the lookup table has been constructed, each pair of pixels in the two images is compared. If the difference between them exceeds the cutoff corresponding to the dose of the lower one, the higher pixel is replaced with the value of the lower one plus one standard deviation, σ_{a-b} , corresponding to the expected deviation at the dose of the lower pixel value. The addition of the one standard deviation reduces the inherent statistical bias resulting from always choosing the lower value. The images then are usually added together to produce a final dezingered image.

III. BACKGROUND SUBTRACTION

Many types of x-ray detectors accumulate signal even when not exposed to radiation. This is often in the form of a dark current in the x-ray conversion device (such as leakage in a silicon detector). For CCD detectors, thermally generated charge in the CCD chip is the primary form of background signal. By cooling the CCD, the dark current can be reduced by a factor of 2 for roughly every 7°C drop in chip temperature. For normal clocking of chips, one must cool below -50°C to reduce the dark current below the $1e^-/\text{pix}/\text{s}$ level. This rate is well above the dark current for bulk silicon; most of the dark current in this mode is due to charge generated in defects near the surface of the device. Many of the current generation of CCDs can operate in the multiphase pinned (MPP) mode, which, through implants in the silicon and proper biasing, moves the charge collection region away from the surface. This results in an order of magnitude lower dark current. This biasing scheme often reduces the full well capacity of the pixel considerably, however, and is used primarily where long integrations or warmer CCD operations are necessary.

The rate of dark current accumulation, in general, varies from pixel to pixel, and hence should be measured and subtracted from an image in order to obtain proper intensity information. The signal, I , recorded at each pixel with no incident radiation usually has the form:

$$I(t, x, y) = I_0(x, y) + a(x, y)t,$$

where t is the exposure time and the constants I_0 and a both depend upon the specific pixel (as indexed by positions x and y) in question. I_0 usually includes an intentionally added

electronic offset, or “pedestal” level, to avoid negative input to the analog-to-digital converter. In practice, once a series of images with a given exposure time has been taken, a properly dezingered image of the same exposure time taken with the shutter closed can be subtracted (pixel by pixel) from each of the data images to remove the dark current.

Images with small variations in exposure time could conceivably use the same background image. One could also interpolate between background images of different times. In practice, most experiments take multiple images with only a few different integration times, so background images are taken for each exposure time used. Since both temperature change and electronic drift can affect the dark current image, it is recommended to periodically obtain new dark images. Dark current depends strongly on temperature, so the temperature must be well regulated (± 0.1 °C) if the dark current is to be subtracted accurately, especially for long exposures.

Since there is a noise associated with reading an image, as well as a shot noise in the generation of the dark current, subtraction of a dark image will add noise to the data image, an important consideration for images with very low signal. For a very few CCDs, the rate of dark accumulation is uniform from pixel to pixel to within the required accuracy, so that a single constant could be subtracted from the data image. For other CCDs, one can average many dezingered dark images together to reduce the noise introduced by subtraction of the image. For many cases, the shot noise in the x ray signal is sufficiently high that a single dezingered background image is sufficient.

IV. DISTORTION CORRECTION

In the following discussion, the x-ray sensitive surface of the detector is considered to be divided into an x - y raster of equal-sized input pixels. Ideally, the recorded digital image would consist of another x - y raster of equal-sized pixels which would map directly with a one to one correspondence to the regular raster at the detector surface. This ideal is almost never achieved due to distortions in the image. Distortion correction refers to the process of numerically processing the recorded image to yield an ideal mapping.

Geometric image distortions can come from a variety of sources. Electrostatically focused image intensifiers typically introduce pincushion distortion into an image, whereas magnetically focused intensifiers have a characteristic S-shaped distortion. Smoothly varying distortion in fiber optic tapers comes primarily from the taper pulling processes as it is difficult to achieve the optimal temperature profile in a large block of glass so that all fibers in the bundle reduce uniformly. As long as distortions are stable, they can be measured and corrected *via* software. Distortions built into the fiber optics are absolutely stable over time. Distortions arising from image intensifiers vary with electric and magnetic field conditions and are much more susceptible to electronic drift and extraneous external fields. Such detectors require very stable electronics and careful attention to shielding.

The distortions discussed above are smoothly varying across the face of the detector. A much more troublesome distortion is due to discontinuous shears which may be

present in fiber optics.^{8,9} Whereas smoothly varying distortions change the shape of a recorded feature, shears can split it into two disjoint parts. Since fiber optics can be obtained that have maximum shear distortions comparable to the pixel size of many detectors (e.g., about 50 μm), the best approach is to use low-shear fiber optics. In circumstances where the resolution is too fine for available fiber optics, other calibration methods must be used. The procedures described below assume smoothly varying distortions.

To measure distortion, an x-ray shadow mask with a regular grid of holes is placed in front of the detector. Our detectors typically have distortions that are smooth on millimeter length scales; accordingly, we commissioned the lithographic fabrication of a mask with 75 μm holes arrayed on a 1 mm square grid in 50 μm thick tungsten foil (Towne Technologies, Somerville, NJ). The mask is uniformly illuminated with x rays by placing an uncollimated x-ray source 1 m away. This produces an image of a regular array of spots in the detector. To map detectors with potentially finer scale distortions would require an appropriately finer pitch mask or a procedure by which a coarser mask is translated by small amounts between exposures. It is important that the mask be as close as possible to the detector face so as to reduce parallax effects.

Using the x-ray image of the mask, the user identifies a set of four fiducial spots defining a small square near the center of the image. The centroid of each spot is computed using a user defined box size around each fiducial position. An automatic spot search is begun by iteratively extrapolating toward the edges of the image from the set of known spots and computing the centroid in a box about the predicted spot position. The position of this new spot and its row and column index are added to the set of known spots. Extrapolation along a given spot row or column is ended when the image boundary is reached or if the spot dose falls below a defined fraction of the fiducial spot's dose.

Once all of the spots in the image are located, the array is extrapolated by five spots in each direction to provide a continuous mapping at the image boundary. The array of spot positions is then smoothed using a piecewise fit to cubic polynomials. A smoothed spot position is calculated using a cubic polynomial fit versus spot index to the four previous and four following spots along the same row, repeating for each spot in the image. The smoothing is then repeated by fitting cubic polynomials along columns. Although no spot position changes by more than 0.25 pixel as a result of the smoothing, this step is necessary to produce a smooth distortion correction map. The distortion map is used to normalize pixel area in the calculation of the intensity correction (see below) and without smoothing, unphysical variations in the signal on the order of 2% were evident. These variations occurred on the length scale of the spot spacing, not on any spacing associated with the detector such as sub-bundle structure in the fiber optics. The inaccuracies inherent in calculating spot centroids may contribute to the need for smoothing. However, one should not expect the distortion resulting from a fiber optic taper to be strictly smooth if there are shear dislocations present in fiber optics. The correction of an image with shears is much more difficult and the cur-

rent distortion mask is much too coarse to map shear boundaries. Since the shears in our fiber optics have generally been small relative to the pixel size, we have chosen to ignore these distortions.

Given the array of spot positions, the distortion from an ideal lattice must be calculated. While any arbitrary ideal square lattice could be chosen, we choose to fit an ideal lattice using a least-squares fit of a square array to the spot positions. The difference between a spot's position in the recorded, uncorrected image, $(x, y)_{c,r}$, and that of the ideal lattice point is that spot's distortion vector, $(\Delta x, \Delta y)_{c,r}$. Here c and r are the spot column and row indices, respectively. Thus, the set of points $(x + \Delta x, y + \Delta y)_{c,r}$ is an ideal lattice.

Once the distortion as a function of spot index has been found, the distortion vector for each pixel in the original image is interpolated using a series of one-dimensional cubic splines. For each row of spots, spline coefficients are calculated for the y position of the spot and the x and y distortions as a function of the x position along that row. The splines for each row of spots, r , are evaluated for every pixel column, i , yielding a new set of distortion arrays $(y)_{i,r}$ and $(\Delta x, \Delta y)_{i,r}$.

A set of spline coefficients are now determined for the x and y distortions as a function of the y position for each pixel column, i , and are evaluated at each pixel row, j , giving the desired pixel-by-pixel distortion map, $(\Delta x, \Delta y)_{i,j}$. In other words, the center of each pixel (i, j) in the recorded, uncorrected image maps to $(i + \Delta x, j + \Delta y)_{i,j}$ in the distortion corrected image.

The set of ideal locations, $(i + \Delta x, j + \Delta y)_{i,j}$, is saved as a pair of floating point image files, one for the ideal x coordinates and the other for the y coordinates. Display of these files reveals interesting information, since the magnitude linearly encodes the local distortion. If the ideal location falls outside of the image boundary (a common occurrence for pixels near the edges), a bad pixel flag value is saved instead, simplifying the error checking in the distortion correction procedure and thereby speeding up the correction algorithm.

Note that the ideal coordinates of a pixel, $(i + \Delta x, j + \Delta y)_{i,j} \equiv (x', y')_{i,j}$, are generally not integers. To properly undistort an image, the dose at pixel (i, j) in the distorted image will be placed in several pixels in the neighborhood of the location $(x', y')_{i,j}$ in the undistorted image. Since image distortions will, in general, also result in dilation or contraction of the local unit area per pixel, the number of pixels over which the dose incident onto a single pixel is spread will vary across the face of the detector. For simplicity, we map each pixel, (i, j) , in the source (recorded, uncorrected) image, to a rectangular area in the corrected image of width, w , and height, h , given by

$$w_{i,j} = x'_{i,j} - x'_{i-1,j},$$

$$h_{i,j} = y'_{i,j} - y'_{i-1,j}.$$

In practice, for images without extreme distortion, the intensity of a pixel is distributed among up to nine pixels in a 3×3 square of pixels about the ideal location $(x', y')_{i,j}$ (Fig. 1). The overlap of the rectangle defined by $w_{i,j}$ and $h_{i,j}$ with each of these 9 pixels is used to compute the fraction of




FIG. 1. Example of the mapping of a source pixel into undistorted (corrected) pixel space. The distortion mapping places the location of the source pixel at the square dot, shown overlain on the 3 pixel by 3 pixel local neighborhood in corrected image space. The width and height of the source pixel are calculated from the distances to the adjacent source pixels in this undistorted space. The height and width determine how the dose within the source pixel is distributed among the pixels with which it overlaps. Each pixel receives a fraction of the dose in proportion to the overlap area with the source pixel. For example, the upper leftmost pixel receives $(\text{row}(1)\text{col}(1))/(\text{width} \times \text{height})$ of the total dose of the source pixel.

intensity distributed to each. The fractional distribution is calculated separately for columns and rows using,

$$\text{col}(1)_{i,j} = \min \left\{ \frac{1}{w_{i,j}} \left[\frac{w_{i,j}}{2} - 0.5 - (x'_{i,j} - i) \right], 0, 0 \right\},$$

$$\text{col}(3)_{i,j} = \min \left\{ \frac{1}{w_{i,j}} \left[\frac{w_{i,j}}{2} - 0.5 + (x'_{i,j} - i) \right], 0, 0 \right\},$$

$$\text{col}(2)_{i,j} = 1.0 - \text{col}(1)_{i,j} - \text{col}(3)_{i,j},$$

$$\text{row}(1)_{i,j} = \min \left\{ \frac{1}{h_{i,j}} \left[\frac{h_{i,j}}{2} - 0.5 - (y'_{i,j} - j) \right], 0, 0 \right\},$$

$$\text{row}(3)_{i,j} = \min \left\{ \frac{1}{h_{i,j}} \left[\frac{h_{i,j}}{2} - 0.5 + (y'_{i,j} - j) \right], 0, 0 \right\},$$

$$\text{row}(2)_{i,j} = 1.0 - \text{row}(1)_{i,j} - \text{row}(3)_{i,j}.$$

The fractional area in each of the nine pixels is determined by multiplying the $\text{row}(m)_{i,j}$ and $\text{col}(n)_{i,j}$ as appropriate. This fractional area for each destination pixel is multiplied by the source pixel's intensity and summed with any other contributions from neighboring pixels in the source image which may also contribute to a given destination pixel's intensity. Often either $\text{col}(1)_{i,j}$ or $\text{col}(3)_{i,j}$ and either $\text{row}(1)_{i,j}$ or $\text{row}(3)_{i,j}$ are zero, in which case the fraction is equal to zero for five of the nine pixels. If the source pixel is saturated, a value indicating saturation is stored in all pixels which have a nonzero overlap with the source pixel.

The accuracy of the correction has been measured by taking another image of the square mask which was arbitrarily displaced and rotated from its original position. Once corrected with data from the first image, no spot centroid was found to be greater than 0.25 pixels from the ideal lattice. It should also be noted that while the distortion correction preserves the overall intensity of the features in the image, it

acts as a smoothing filter in regions of the detector where the ideal location of a pixel nears pixel boundaries in the resulting image.

Some applications, such as the measurement of the total intensity of small, well-isolated crystallographic spots, do not require a distortion correction. As long as the PSF is very narrow, so that the intensity of the spot is well confined to within the measurement area of a given spot, it may simply be adequate to know the distortion map in order to predict the positions of the spots in the distorted image. Once a spot is located, a local integration of the spot and local background subtraction may be adequate.

V. INTENSITY (FLAT-FIELD) CORRECTION

Irregularities that cannot be eliminated during the construction of a detector limit the uniformity of response over the active area and hence the accuracy to which a signal can be measured. Such irregularities exist in each part of the optical chain in a CCD detector. Phosphor screens have variations in thickness, both on a small scale due to the packing of individual phosphor grains, and on a longer scale due to variations in the deposition process. Variations in the phosphor thickness also give rise to an x-ray energy dependent response.¹⁰ Since the apparent thickness of the phosphor also changes with incident angle, there is an angle dependent response as well. Fiber optics have a distinct "chicken wire" pattern where smaller subbundles have been fused. There are also variations in transmission from edge to center in the fiber optic taper, as well as local variations in the transmission of individual fibers. As noted above, the change in area due to geometric distortions will affect the calibration procedure. The CCD will have some pixel to pixel variations in sensitivity as well.

To the extent that such irregularities are stable with time, they can be calibrated and the subsequently recorded images can be software corrected. The time scale over which the calibration is valid varies with the type of detector and circumstances of use. For example, sufficiently high x-ray doses on the detector can locally darken the glass in the fiber optics through the creation of color centers; therefore, care should be taken to shield the detector from the main beam at synchrotron sources. Other calibration drifts arise from movement of the fiber optics relative to the phosphor or the CCD or from dirt or scratches on the detector face. For the highest accuracy work, it is prudent to periodically recalibrate the detector. One reason we choose to store our calibration files in files similar to diffraction images is that calibration files can be readily compared by simply displaying the difference between two calibrations using standard display software.

The most straightforward way to calibrate system non-uniformities is to produce a uniform source of x rays that reproduces the angle and energy distribution of x rays in the experimental setup. One approach⁷ is to use an x-ray fluorescent glass sample at the same position as the diffracting sample. The resulting field of x rays is not uniform, however, and must be calibrated independently for each fluorescent sample. The response of our CCD detectors is a function of

the independent variables of local position, x-ray energy, and angle of incidence of the x rays with respect to the detector face. Since it is relatively easy to generate a uniform source of x rays with near-normal incidence, we have chosen to measure separately the pixel-to-pixel variations of the detector to normally incident x rays, the angle dependence, and the energy dependence of the response. In combination, these allow correction of images in which the sample to detector distance, x-ray energy, or angle of incidence varies without having to acquire new calibration data.

The response at normal incidence is measured by exposing the detector to a uniform field of x rays of the proper energy, accomplished by placing a low power x-ray generator at a distance of 1 m along a normal to the center of the detector face. We use a TFS-6050 Cu x-ray tube (with TCM 5000M power supply, both from TruFocus Corporation, Watsonville, CA 95076) because this tube has a 0.2 mm focal spot, which is a good approximation to a point source. The uniformity of the field was verified, via raster scans with a NaI-photomultiplier scintillation counter, to be within 0.5% over a 8 cm×8 cm area at this distance. The primary source of the remaining nonuniformity was found to arise from defects in the Be window of the x-ray tube. Uniformity was further improved by rotating the x-ray tube assembly around the generator-to-detector axis between exposures. A series of frames which accumulated more than 40 000 x-rays/pixel were taken to reduce the shot noise of the x rays to below the 0.5% level.

After background subtraction, the image is distortion corrected to compensate for any area dilation due to the distortion. A small correction is then applied to account for the angular effects due to the finite distance of the source and the absorption of the air:

$$I'(\theta) = \frac{I(\theta)}{\cos^2 \theta \times \cos \theta \times e^{-\alpha/\cos(\theta)}}$$

In this equation, θ is the angle of incidence of the ray from the source to the point in question on the detector face, the cosine squared term corrects for the $1/r^2$ falloff in intensity (i.e., the center of the detector is closer to the source than the edge), the second cosine term corrects for the apparent reduction in size of the phosphor pixel due to the angle of incidence, and the exponential term corrects for the x-ray absorption in the air. Note that for most CCD detectors at the angles involved in the setup described above, the air absorption term is negligible compared to the observed nonuniformity due to the Be window and could be ignored. The correction file is now generated by dividing a constant by the intensity found at each pixel. The constant is chosen such that the average applied intensity correction is near unity. The correction file is stored as a floating point image consisting of an array of factors to be applied to each pixel in the data image.

Ideally, a separate calibration should be made for each x-ray energy of interest. However, the dominant intensity nonuniformities are often due to the fiber optics and pixel-to-pixel sensitivity variations in the CCD, both of which are independent of x-ray energy. Most of the energy dependence to normally incident radiation comes from variations in the

phosphor and in the entrance window. To perform calibrations at different x-ray energies, one could obviously use x-ray tubes with various characteristic energies. We have found it to be far more practical to approximate the response of a particular energy by varying the accelerating voltage of the x-ray tube and hence the proportion of bremsstrahlung radiation produced. In addition, the air in the x-ray path will absorb the lower energy x rays preferentially, hardening the x-ray energy spectrum as well. To calibrate for Cu radiation, it is important to keep the accelerating potential as low as practical. For this energy, we operate the tube at no more than 10 kV. To calibrate for higher energy radiation, we have empirically found that an 18 kV setting on the above setup with copper anode closely matches the response of monochromatic 13.6 keV x-ray illumination for our 1k detector.¹¹ The validity of this procedure can be checked by correcting an image containing smoothly varying, diffuse diffraction. Correction with a source file of the incorrect energy will result in obvious features in the image due to variations in the phosphor thickness. One should also note that the wrong choice of calibration energy can make the nonuniformity worse than if no correction was applied. As we shall see below, an increase in phosphor thickness can, depending on the x-ray energy, either increase or decrease the resultant signal per x ray. Unless this dependence is properly taken into account, relatively increasing signals can be incorrectly calibrated so as to appear to be decreasing.

The dependence of the response to the angle of incidence of the x rays to the phosphor surface results from two competing processes: First, the absorption efficiency is greater for x rays incident to the phosphor at oblique angles since the maximum path length through the phosphor increases. Second, the visible light produced by the x rays is attenuated as it travels through the phosphor, so light from x rays converted near the incident surface will be attenuated more. At higher angles of incidence, the x rays will tend to convert nearer to the front surface of the phosphor, resulting in a smaller signal. For lower energy x rays, most of the x rays are converted even at normal incidence, so there will be only a marginal increase in the number converted at higher angles of incidence. The x rays will convert nearer to the surface and the net effect will be a reduction in signal/x ray. For higher energy x rays, the increased absorption at higher angles is significant and outweighs the reduction due to light attenuation. There will be an intermediate energy where the effects cancel and there is no angle dependence. The angle dependence of the response is not negligible: it can easily be a 10% effect for x rays entering at 45°, depending on the energy.

The response of the phosphor as a function of incident angle is measured by exposing the phosphor to a monochromatic beam of x rays of the appropriate energy and at various angles of incidence. The intensity of the spot is measured at each angle. The measured response can be fit well to the form:

$$I(\theta) = I(0) + a\theta^2,$$

where $I(0)$ is the intensity at normal incidence, θ is the angle of incidence in degrees, and the coefficient a describes the

angular correction. Note that a can be either positive or negative and will in general depend on phosphor type and thickness as well as x-ray energy. The coefficient a is measured at several x-ray energies (8.0, 8.9, 11.0, 13.5, and 18.0 keV) and then fit to a quartic polynomial to provide a smooth interpolation between 8 and 18 keV.¹¹

We find it convenient to develop a polynomial fit to the correction as a function of both the x-ray energy and angle of incidence. This two-dimensional surface is smooth, varies slowly with energy and angle, and is therefore readily computed from a small number of coefficients. Then, given the beam energy, specimen to phosphor distance, beam location on the face of the detector, the tilt angle between the detector and the main beam (usually 0°), and the pixel size, a multiplicative angular response correction factor is determined for each pixel of the detector. This array of multiplicative factors is then applied to each image.

The above procedure needs to be modified if the x-ray image is polychromatic, as is the case for Laue diffraction. The modifications are straight forward if the x-ray wavelength incident on each pixel are known *a priori* — one simply applies the x-ray energy correction separately. If polychromatic radiation of an unknown spectrum is incident on a given pixel, alternative correction procedures may be needed.

The accuracy of the intensity correction depends on the nature of the data to which it is applied. The point spread function of the detector limits the accuracy to which the correction factor can be determined for a given pixel. The point spread acts as a smoothing filter to the flood illumination data, effectively smoothing the corresponding correction factors which are to be applied. Taken from this flood data, the correction factor for a given pixel comes from a weighted average of the surrounding pixels. For small diffraction features, however, the surrounding pixels are not illuminated, reducing the precision of the correction. For broad features, the data resemble the conditions under which the correction was generated, leading to a more accurate measure. For the Princeton 1k detector, 75 μm spots could be measured to an absolute accuracy of 1%, while 300 μm spots could be measured to 0.3%.¹¹

Since the distortion correction acts as a smoothing filter, effectively increasing the point spread, a more precise intensity correction can be made by applying the intensity correction before the distortion correction as opposed to the method described above which performs the distortion correction first. The intensity correction to be applied before distortion correction is calculated by dividing a constant by the pixel value in the original distorted calibration image. The distortion correction files are used to normalize the area/pixel at the phosphor surface. Since it is dependent upon accurate spatial information and because it is a smooth function of position, the obliquity correction can still, for simplicity, be performed after the distortion correction. One would expect the greatest difference between the two methods of intensity correction to occur for sharp features. For the Princeton 1k detector, however, images of 75 μm spots which were intensity corrected either before or after the distortion correction step gave integrated spot intensities which agreed

to better than 0.2%. Agreement for larger features was better than 0.1%.

VI. DISCUSSION

Application of the full package of corrections described here to a 1024×1024 image requires ~8 s on a 150 MHz desktop computer. Intensities of diffraction features can be recorded to 1% or better and spatial position can be corrected to 0.25 pixels or better. This package of corrections is routinely used for both crystallographic and diffuse x-ray signals on two of the CCD detectors at the Cornell High Energy Synchrotron Source (CHESS), with excellent results.^{11–14}

To further improve the accuracy of the intensity calibration for extremely fine features, one could try to acquire the correction data in a manner that more closely mimics the data to be corrected. One could envision a procedure, applicable for instance to crystallography, where small spots are scanned across the face of the detector. This would require a stable source, accurate timing in the shutter, and many exposures to cover the entire detector. A major concern is the energy dependent response of the system. Care must be taken to calibrate with the appropriate x-ray energy spectrum so that the phosphor response is matched to the data. The angle dependence in the phosphor response has often been neglected in image corrections. It is possible that not all measurements would be adversely affected by such an omission. Thermal disorder within crystals produces a systematic reduction in the expected intensity of higher order diffraction spots as compared to an ideal crystal. Typical crystallographic data analysis attempts to fit for this disorder through the introduction of B factors. Applying the B-factor correction to diffraction intensities amounts to applying an angle-dependent scale factor. It is possible that the angle dependence in the uncorrected detector response could be taken up into the B factors as well, thereby modifying the apparent disorder within the crystal. This hypothesis has not been tested.

ACKNOWLEDGMENT

This work was supported by DOE Grant No. DE-FG02-97ER62443.

APPENDIX THE EFFECT OF THE PSF ON MEASURED GAIN

The finite point spread function of most real detectors smoothes the statistical distribution of the pixel intensities. Neighboring pixels in this case are responding to the same incident x ray. Hence, they are no longer statistically independent, and their variance will be reduced. This effect is purely geometrical, and only depends on the functional form of the point spread.

First, if there is no point spread (each pixel of the detector counts the exact number of x rays which are absorbed in the corresponding area of the phosphor), we expect the number of x rays recorded in pixel *i* to be distributed with mean, μ_i , and variance, σ_i^2 , according to

$$(\mu_i, \sigma_i^2) = (\mu_i, \mu_i),$$




FIG. 2. Geometry of the point spread function calculation. The signal increase recorded in pixel *u* for x-rays incident at point *r* is the integral of the point spread function, centered at *r* over the area of the pixel.

where we have invoked Poisson statistics equating the variance and the mean. If the detector has a finite gain per pixel, g_i , this becomes

$$(\mu_{\text{meas},i}, \sigma_{\text{meas},i}^2) = (g_i \mu_i, g_i^2 \sigma_i^2) = (g_i \mu_i, g_i^2 \mu_i).$$

If the signal in a given pixel is the sum of independent signals from different regions of the phosphor (each of which obeys Poisson statistics), we get

$$(\mu_{\text{meas},i}, \sigma_{\text{meas},i}^2) = \left(\sum_j g_j \mu_j, \sum_j g_j^2 \mu_j \right),$$

where the sum is over each contributing region of the phosphor, and g_j represents the gain for each region (which will be a function of both the detector gain and the distance to the region *j* from the pixel *i*).

Using this, we can now calculate the effect of the point spread function. If we consider an infinitesimal element of the phosphor area, $d^2\mathbf{r}$, its gain in the pixel *u* will be the overlap of its (normalized) point spread function with the pixel multiplied by the inherent gain of the detector, g_d (see Fig. 2):

$$g_j = g_d \int_{\text{Pixel } u} \text{psf}(\mathbf{r} - \mathbf{u}) d^2\mathbf{u}.$$

Also $\mu_j = \rho(\mathbf{r}) d^2\mathbf{r}$, where $\rho(\mathbf{r})$ is the x-ray flux at *r*. Thus

$$g_j \mu_j = g_d \rho(\mathbf{r}) d^2\mathbf{r} \int_{\text{Pixel } u} \text{psf}(\mathbf{r} - \mathbf{u}) d^2\mathbf{u},$$

$$g_j^2 \mu_j = g_d^2 \rho(\mathbf{r}) d^2\mathbf{r} \left[\int_{\text{Pixel } u} \text{psf}(\mathbf{r} - \mathbf{u}) d^2\mathbf{u} \right]^2.$$

Because the point spread function is negligible at large distances, the summation of the previous paragraph can be replaced with a two-dimensional integral over all space:

$$\mu_{\text{meas}} = \sum_j g_j \mu_j = \int_{\text{allspace}} g_d \rho(\mathbf{r}) d^2 \mathbf{r} \int_{\text{Pixu}} \text{psf}(\mathbf{r}-\mathbf{u}) d^2 \mathbf{u}$$

$$\sigma_{\text{meas}}^2 = \sum_j g_j^2 \mu_j = \int_{\text{allspace}} g_d^2 \rho(\mathbf{r}) d^2 \mathbf{r} \left[\int_{\text{Pixu}} \text{psf}(\mathbf{r}-\mathbf{u}) d^2 \mathbf{u} \right]^2.$$

Since an ideal detector will have a variance equal to the measured intensity, the variance for a real detector will be reduced from the expected value by a factor, β :

$$\beta = \frac{\sigma_{\text{real}}^2}{\sigma_{\text{ideal}}^2} = \frac{\sigma_{\text{meas}}^2}{\mu_{\text{meas}}}$$

$$\beta = \frac{\int_{\text{allspace}} g_d^2 \rho(\mathbf{r}) d^2 \mathbf{r} \left[\int_{\text{Pixu}} \text{psf}(\mathbf{r}-\mathbf{u}) d^2 \mathbf{u} \right]^2}{\int_{\text{allspace}} g_d \rho(\mathbf{r}) d^2 \mathbf{r} \int_{\text{Pixu}} \text{psf}(\mathbf{r}-\mathbf{u}) d^2 \mathbf{u}}.$$

This equation can be simplified for the most interesting case, uniform illumination of a detector with unity gain and a Gaussian point spread function,

$$\text{psf}(\mathbf{r}) = \frac{1}{2\pi\sigma^2} e^{-r^2/2\sigma^2}.$$

Substituting this into the previous equation, and replacing the x-ray flux with a constant:

$$\beta = \frac{\frac{\rho}{2\pi\sigma^2} \int_{\text{allspace}} d^2 \mathbf{r} \left[\int_{\text{Pixu}} d^2 \mathbf{u} (e^{-(\mathbf{r}-\mathbf{u})^2/2\sigma^2}) \right]^2}{\frac{\rho}{2\pi\sigma^2} \int_{\text{allspace}} d^2 \mathbf{r} \int_{\text{Pixu}} d^2 \mathbf{u} (e^{-(\mathbf{r}-\mathbf{u})^2/2\sigma^2})}.$$

By reversing the order of integration in the denominator, it is simple to show that it evaluates to ρA , where A is the area of a pixel. Defining the units of area so that A is equal to one, we find

$$\beta = \frac{1}{2\pi\sigma^2} \int_{\text{allspace}} d^2 \mathbf{r} \left[\int_{\text{Pixu}} d^2 \mathbf{u} (e^{-(\mathbf{r}-\mathbf{u})^2/2\sigma^2}) \right]^2.$$

This can be numerically integrated for selected values of σ . The results are graphed in Fig. 3. Even for relatively small point spreads, the measured variance in the signal of a uniform detector illumination will very strongly underestimate the actual variance. For example, an excellent detector with a PSF of only 2 pixels (full width at half maximum, FWHM) gives images with variance reduced more than an order of magnitude by this effect. Monte Carlo simulations of detector performance verify this calculation.

Although this calculation does not affect the dezingering process, it becomes significant when measuring the detective quantum efficiency (DQE) of the detector. A common procedure for doing this is to use a calibrated source to measure the gain of the detector. The detector is then illuminated with a uniform flat-field x-ray source, and the pixel-to-pixel variation is measured as a function of the mean intensity. This

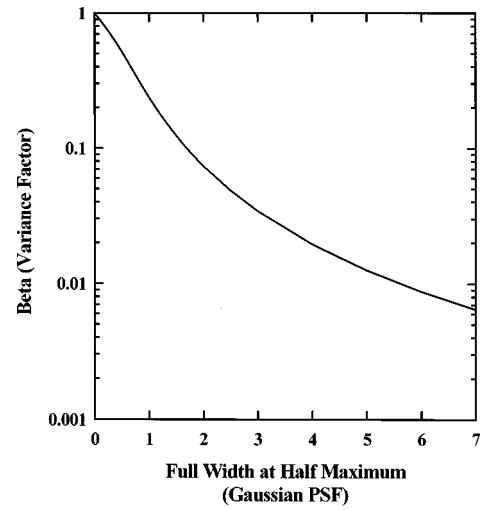


FIG. 3. Effect of the point spread function (PSF) on the variance of an image with uniform illumination. The measured pixel variance (for a detector with a given Gaussian PSF) is reduced by a factor $1/\beta$ relative to a detector with no PSF smoothing. Other functional forms for the PSF can cause the variance to fall even more rapidly.

value is then compared with Poisson statistics. Clearly, the smoothing effect will reduce this variation, and can lead to false measurements of extremely high values for the quantum efficiency (even giving unphysical values greater than 1!). It should also be noted that realistic point spread functions are non-Gaussian and can have significant tails at large distances from their centers. This will tend to increase the effect described here. Note that if there is at most one recorded quantum/x ray, the signal in adjacent pixels is statistically independent and there is no PSF reduction of the variance.

- ¹M. W. Tate, S. M. Gruner, and E. F. Eikenberry, *Rev. Sci. Instrum.* **68**, 47 (1997).
- ²S. M. Gruner and J. R. Milch, *Trans. Am. Crystallogr. Assoc.* **18**, 149 (1982).
- ³D. J. Thomas, *Proc. R. Soc. London, Ser. A* **425**, 129 (1989).
- ⁴M. Stanton, W. C. Phillips, Y. Li, and K. Kalata, *J. Appl. Crystallogr.* **25**, 549 (1992).
- ⁵A. P. Hammersley, S. O. Svensson, and A. Thompson, *Nucl. Instrum. Methods Phys. Res. A* **346**, 312 (1994).
- ⁶A. P. Hammersley, S. O. Svensson, A. Thompson, H. Graafsma, A. Kwick, and J. P. Moy, *Rev. Sci. Instrum.* **66**, 1 (1995).
- ⁷J. P. Moy, A. P. Hammersley, S. O. Svensson, A. Thompson, K. Brown, L. Claustre, A. Gonzalez, and S. McSweeney, *J. Synchrotron Radiat.* **3**, 1 (1996).
- ⁸C. I. Coleman, *Adv. Electron. Electron Phys.* **64B**, 649 (1985).
- ⁹H. W. Deckman and S. M. Gruner, *Nucl. Instrum. Methods Phys. Res. A* **246**, 527 (1986).
- ¹⁰S. M. Gruner, S. L. Barna, M. E. Wall, M. W. Tate, and E. F. Eikenberry, *Proc. SPIE* **2009**, 98 (1993).
- ¹¹M. W. Tate, E. F. Eikenberry, S. L. Barna, M. E. Wall, J. L. Lowrance, and S. M. Gruner, *J. Appl. Crystallogr.* **28**, 196 (1995).
- ¹²R. L. Walter, D. J. Thiel, S. L. Barna, M. W. Tate, M. E. Wall, E. F. Eikenberry, S. M. Gruner, and S. E. Ealick, *Structure* **3**, 835 (1995).
- ¹³D. J. Thiel, S. E. Ealick, M. W. Tate, S. M. Gruner, and E. F. Eikenberry, *Rev. Sci. Instrum.* **67**, 1 (1996).
- ¹⁴M. E. Wall, S. F. Ealick, and S. M. Gruner, *Proc. Natl. Acad. Sci. USA* **94**, 6180 (1997).