

7.2. CCD DETECTORS

Fibre optics and lenses. Light transmission in image reduction is limited by $\text{n.a.} \times M^2$, where n.a. is the numerical aperture of the optical system and M is the linear magnification factor ($M < 1$ corresponds to reduction). Lens systems have low values of n.a. when used in reduction, so there is only 2% light transmission for a 3:1 reduction with an f/1.0 lens. With a much higher n.a., fibre optics can typically transmit 13% with the same 3:1 reduction (Coleman, 1985). Such reducing tapers are produced by locally heating, then pulling, a fused bundle of optical fibres. Each fibre within the bundle becomes tapered in this process, thereby reducing the image scale from front to back of the bundle. The bundle structure introduces a characteristic 'chicken-wire' pattern into the image which must be removed *via* intensity calibration (Section 7.2.3). Tapers up to 165 mm in diameter are available.

To obtain good resolution, extramural absorbing fibres (EMAs) must be placed in the fibre optic to absorb light that propagates between fibres. These EMAs are often a more effective absorber in the blue part of the spectrum, with the result that a red-emitting phosphor yields somewhat poorer resolution. Another concern with fibre optics is radioactivity in the glass used to produce the optic, which manifests itself as random flashes ('zingers') in both the phosphor and the CCD (Section 7.2.3).

Although fibre-optic coupling is preferred in many situations, lenses are appropriate for use with image intensification, or in cases of image magnification, such as for microtomography, using high-resolution screens.

Image intensifiers allow large demagnification ratios without undue DQE loss (Moy, 1994). In these vacuum-tube devices, visible light produces photoelectrons in the photocathode of the intensifier, which in turn are accelerated under high potential onto a secondary phosphor screen. Each photoelectron excites many photons in the secondary phosphor, giving light amplification. Image resolution is retained either by magnetic, electrostatic or proximity focusing of the electrons. In the case of microchannel plate intensifiers, photoelectrons are restricted to cascade down hollow fibres. Intensifiers often have problems with stability and linearity, and they degrade resolution. Phosphor afterglow in the secondary phosphors is also a consideration. Intensifiers typically limit the input format to less than 80 mm in diameter, although one design uses a large (230 mm) radiographic intensifier (Moy, 1994; Hammersley *et al.*, 1997). The greatest drawbacks of image intensifiers are cost, availability, susceptibility to magnetic fields and lack of robustness.

CCDs. The low noise and high sensitivity of scientific CCDs have made CCDs the preferred imaging device in X-ray applications. Visible photons are converted to charge carriers in the silicon of the CCD, with a 30–40% quantum efficiency (q.e.) in the red, dropping to 5% in the blue owing to increased absorption by the controlling gate structure (the q.e. can be considerably higher for back-illuminated devices which, however, are expensive, fragile and difficult to obtain). The 'read noise' (noise with zero charge in the pixel) is routinely less than 10 electrons for scientific grade chips with pixel read rates of less than 1 MHz. When coupled to a good phosphor screen with a modest (*circa* 3:1) fibre-optic taper, signal levels in the CCD are typically 10–30 recorded electrons per 10 keV X-ray. This is above the CCD read-noise level, thereby allowing low-dose quantum-limited X-ray imaging.

A physically large CCD chip is usually desired to maximize the detector area, given the limited image reduction ratio. Devices from 2–4 cm on an edge with 1000×1000 to 2000×2000 pixels are available. These sizes match well to the largest of the fibre-optic tapers and are therefore most often used in detectors. Larger-format CCDs are available, but are harder to obtain and are expensive.

CCDs are usually cooled well below room temperature to minimize thermally generated dark current, an unwanted source of noise. The dark current drops by a factor of 2 for every 5–7 K

drop in temperature. At 233 K, there is roughly one electron per pixel per second dark current for normal clocking operation. Since the dark current changes with temperature, the temperature must be well regulated (to within ± 0.1 K) to subtract the dark charge from an image reproducibly.

Most of the dark current comes from surface defects, not from the bulk silicon. Multiphase pinned (MPP) CCDs use charge implants within the pixel structure to move the charge collection region away from the surface, thereby reducing the dark current by several orders of magnitude. This is usually accompanied by a significant reduction in the maximum charge capacity ('full well') of a pixel. Even so, MPP CCDs are becoming the norm in most X-ray detectors.

Directly exposed CCDs. CCDs can directly image X-rays, although typical CCDs are not very efficient, since the charge collection region in the silicon is very thin ($< 10 \mu\text{m}$). Thicker depletion regions can be fabricated in high-resistivity silicon wafers, improving X-ray collection efficiency to greater than 30% at 8 keV. Direct conversion of X-rays in silicon produces a large signal with excellent energy resolution. It is therefore possible to use the chip as an X-ray counting detector with the ability to discriminate in energy. To retain X-ray-energy measurement capability, however, there must be less than one X-ray per pixel per frame to avoid signal overlap. This requires a fast readout as well as large amounts of disk storage to handle the large number of files. The X-rays damage the CCD's electronic structure, resulting in a higher dark current within the exposed pixels. Care must be taken to shield non-imaging parts of the CCD, such as the output amplifier, which will adversely affect chip performance at even lower radiation dose.

7.2.3. Calibration and correction

Inhomogeneities within the detector components introduce non-uniformities in the output image of several per cent or more, both as geometric distortion and as nonuniformity of response. The response of the system varies not only with position, but also with the angle of incidence and X-ray energy. Optimal calibration of the detector should take into account the parameters of the X-ray experiment, seeking to mimic the experimental conditions as closely as possible: a uniform source of X-rays of the proper energy positioned in place of the diffracting crystal would be ideal. Realizing such a source is somewhat problematic, so the calibration procedure is often broken down into several independent steps. Calibration procedures are detailed in Barna *et al.* (1999) and are summarized below.

7.2.3.1. Dark-current subtraction

It is important to remove both the electronic offset and the accumulated dark charge from an X-ray image. Since the integrated dark current varies from pixel to pixel and with time, a set of images needs to be taken (with no X-rays), matched in integration time to the X-ray exposures. With a properly temperature-stabilized detector, the background images may be acquired in advance and used throughout an experiment. Because the background image has noise, it is common to average a number of separate backgrounds to minimize the noise.

7.2.3.2. Removal of radioactive decay events

Cosmic rays and radioactive decay of actinides in the fibre-optic glass produce large-amplitude isolated signals ('zingers') within an X-ray image. These accumulate randomly in position and in time. For the short exposures typical at synchrotron sources and for data sets with highly redundant information, the few diffraction spots

7. X-RAY DETECTORS

affected by these events can be discarded with statistical analysis. For longer exposures, or for data with less redundancy, two (or more) nominally identical exposures can be taken and then compared pixel by pixel to remove spurious events (see Barna *et al.*, 1999).

7.2.3.3. Geometric distortion

Geometric distortions arise in the optical coupling of the system. If they are stable, the distortions can be mapped and corrected. Long-term stability is possible for a phosphor fibre-optically coupled directly to a CCD, since all distortions are mechanically fixed. Intensifier-based systems are subject to changes in magnetic and electric fields, hence stability is more of a problem.

Geometric distortions may be either continuous or discontinuous. Fibre optics often have shear between adjacent bundles of fibres. In this defect, one group of fibres will not run parallel to a neighbouring group, causing a discontinuity in the image. Rather than dealing with such discontinuities, tapers with low shear (less than one pixel maximum) are usually specified. Even with low shear, there is a continuous distortion (several per cent), which varies slowly over the face of the detector. Such distortion is inevitable, as the temperature profile cannot be precisely controlled in the large block of glass comprising the fibre optic as it is processed. To map this distortion, an image is taken of a regular array of spots. Such an image can be made by illuminating a shadow mask of equally spaced holes with a flood field of X-rays. Holes 75 μm in diameter spaced on a 1×1 mm square grid are adequate for mapping the distortions present in most fibre-optic tapers. Such masks have been lithographically fabricated in an X-ray opaque material, such as 50 μm tungsten foil (Barna *et al.*, 1999).

Given an image produced with this X-ray mask, the displacement map for every pixel in the original image can be computed as follows: find the centroid of each mask spot and its displacement relative to an ideal lattice. The array of spot positions and associated displacements can then be interpolated to find the displacement for each pixel in the original distorted image.

The displacement of a pixel from original to corrected image will not in general be a whole number. Rather, the intensity of a pixel will be distributed in a local neighbourhood of pixels in the corrected image centred about the position given in the displacement map. The size of the neighbourhood depends on the local dilation or contraction of the image; typically the intensity will be distributed in one to nine pixels. This distribution procedure yields a smooth intensity mapping. Applying corrections to mask images that have been arbitrarily displaced shows that the distortion correction algorithm is good to better than 0.25 pixels (Barna *et al.*, 1999).

The geometrical distortion is tied intimately to the correction of the response of the system (see below). Since distortions produce local regions of dilation or contraction of the image, pixels will, in general, correspond to varying sizes. This variation must be included in the flat-field calibration.

7.2.3.4. Flat-field corrections

Variation in response arises from nonuniformity in the phosphor, defects in the fibre optics and pixel-to-pixel variation in the sensitivity of the CCD itself. Of these, variations in the fibre optics are usually most pronounced, often with decreased transmission at the bundle interfaces, resulting in the characteristic 'chicken-wire' pattern. The response of the system is generally stable with time, although exposure to the direct beam at synchrotron sources will cause colour centres to form in the glass of the fibre optic.

Recalibration of the system will correct for the reduced transmission in the exposed spot.

Light output from the phosphor depends on a number of factors: phosphor thickness, X-ray energy, angle of incidence and depth of conversion within the screen. X-ray photons are absorbed within phosphor grains with an exponentially decaying distribution as they travel deeper into the phosphor layer. Making the phosphor layer thicker increases the probability that the X-ray photons are absorbed, yielding an increase in the signal produced. But thicker phosphor layers also scatter the visible photons which one desires to collect in the CCD. At some point, the loss of light will be greater than the increase in X-ray stopping power and the net response will decrease. For one particular phosphor preparation, this appears to happen at *circa* 85% stopping power. The phosphor-screen resolution also falls with increasing screen thickness, although surprisingly slowly owing to the diffusive nature of the light scatter. These effects are discussed in Gruner *et al.* (1993).

Consider a given phosphor layer with thickness nonuniformity. For X-ray energies where the stopping power of the phosphor is low, regions of the phosphor with a thickness larger than the mean will be brighter owing to the increased X-ray stopping power. For energies at which X-rays are strongly absorbed, the increased opacity of the thicker regions of the phosphor will cause the response to go down. This illustrates the importance of calibrating the response of the system to the particular X-ray energy of interest.

In like manner, X-rays impinging on the phosphor at an angle away from the normal are presented with a longer path length and hence an increased stopping power. Also, because of the oblique angle, the distribution of visible-light production will be shifted toward the phosphor surface. Again, for strongly absorbed X-ray energies, the increase in the optical path for the light will cause the recorded signal to fall, whereas for X-rays not as strongly absorbed, the signal will increase.

To map the nonuniformity in response, one would ideally use a uniform source of X-rays of the proper energy placed at the position of the sample. This would calibrate the detector with the proper energy and angle of incidence for the diffraction data to be corrected. Correction factors are computed from a series of images taken of this uniform source. Sufficient numbers of X-rays per pixel must be collected to reduce the shot noise in the X-ray measurement to the required level (*e.g.* 40 000 X-rays per pixel must be acquired to correct to 0.5%).

Providing a truly uniform source with an arbitrary X-ray energy and angular distribution is difficult at best. Other sources can be used, however, with good results. Amorphous samples containing a variety of elements can be fabricated which produce X-ray fluorescence at various wavelengths when excited by a synchrotron beam (Moy *et al.*, 1996). These can be placed at the position of the sample, thereby mimicking the angular distribution of X-rays from the experiment. The fluorescence is not uniform in space, however, so that the actual distribution must be mapped by some means. Once mapped, however, these samples provide a stable calibration source.

Another alternative is to separate the calibration procedure into several parts, mapping the dependence at normal incidence and treating the angular dependence as a higher-order correction. By moving an X-ray source sufficiently far away, the detector can be illuminated at near-normal incidence with excellent uniformity. For example, an X-ray tube at 1 m distance can produce a field with uniformity better than 0.5% over a 10×10 cm area. A sum of images with sufficient X-ray statistics, taken with the area dilation of each pixel computed through the geometric distortion calibration, can be used to compute a pixel-by-pixel normalization factor. Again, this should be determined for the X-ray energy of interest. In practice, the appropriate energy may be approximated by a linear combination of several energies. However, the proper coefficients

will need to be determined empirically. This can be judged in diffraction having broad diffuse features, because imperfections in the phosphor stand out clearly when the data are corrected using factors derived from the wrong X-ray energy.

7.2.3.5. Obliquity correction

It has been found empirically that the light output for a given X-ray energy has a quadratic variation with the angle of incidence. The quadratic coefficient varies with X-ray energy and may be either positive or negative. The angular dependence may be measured by illuminating the detector with a small stable spot of X-rays and recording the integrated dose at various angles of incidence. Given the placement of the detector in relation to the beam and sample, the angle of incidence at a particular pixel can be computed and can be used to find the correction factor needed. With this method, a change in the experimental setup does not require a new calibration, just the computation of a new set of coefficients. The combination of energy and obliquity sensitivity varies slowly and may be approximated by a quadratic or cubic fit to a surface as a function of X-ray energy and angle. The few coefficients defining this surface allow quick computation of the combined energy and obliquity factor with which to multiply the local flat-field correction for X-rays of known incident energy and angle.

The obliquity correction is often ignored, since the solution of structures from X-ray diffraction typically includes a temperature factor which also varies with angular position. Uncorrected angular dependence of detector response will be convoluted with the true temperature factor and often does not impede the solution of the structure.

These procedures do not allow correction to arbitrary accuracy, however. The calibration data are taken with uniform illumination, whereas diffraction spots are localized. Given a nonzero point spread in the detector, the computed correction factor arises from a weighted average of many illuminated pixels. The signal from a diffraction spot only illuminates a few pixels, so the true factor might well be different. This should be less of a problem as diffraction spots become larger, becoming more like a uniform illumination. Measurement for one detector showed 75 μm spots could be measured to 1% accuracy, whereas 300 μm spots could be measured to 0.3% accuracy (Tate *et al.*, 1995).

7.2.3.6. Modular images

The size of available fibre optics and CCDs and the inefficiencies of image reduction limit the practical imaging area of a single CCD system. Closely stacked fibre-optic taper CCD modules can be used to cover a larger area. Although the image recorded from each module could be treated as independent in the analysis of the X-ray data, merging the sub-images into one seamless image facilitates data processing. Each module will have its own distortion and intensity calibration. It is no longer possible to choose an arbitrary lattice onto which each distorted image will be mapped: the displacement and scaling must be consistent between the modules. This would be accomplished most easily by having a distortion mask large enough to calibrate all modules together, although it is possible to map the intermodule spacing with a series of mask displacements.

Flat-field correction proceeds as in the case of a single module detector after proper scaling of the gain of each unit is performed. There can potentially be a change in the relative scale factors between modules, since each is read through an independent amplifier chain. Multimodule systems emphasize the need for enhanced stability and ease of recalibration.

7.2.4. Detector system integration

Hardware interfacing. CCDs are operated, during both data acquisition and readout, by a dedicated hardware controller attached to a host computer, generally a PC. The requirements for the controller are complex and quite stringent in order to obtain low-noise operation. As with any high-speed electronics, noise increases with speed. Typically, pixel read rates of 100–500 kHz are used, although higher read rates can be used (and still preserve low-noise performance) for CCDs with multistage on-chip amplifiers. Some CCDs have multiple output amplifiers that increase pixel throughput by using parallel digitization channels. The entire CCD can also be read at reduced resolution by analogue summation of adjacent rows and/or columns on the chip (binning). This has the added benefit of increasing the signal-to-noise ratio for signals with low spatial frequency since there are fewer digitizations. Binning is highly recommended whenever the reduced resolution can be tolerated.

The time resolution of the detector can be further increased if the CCD array is used as frame storage. In this case, a portion of the imaging area is masked to X-rays, making it available for storage. The exposed area can be shifted rapidly into the masked area and a second exposure begun. Storage for five to ten subframes can easily be configured before readout is necessary. The time resolution is ultimately limited by the phosphor decay time and the time needed to shift the image. Although most CCDs are capable of being operated in very flexible ways, flexible CCD controllers are expensive. The consequence is that few commercial CCD X-ray detectors permit use of all the available options.

The detector itself is contained in a cryostat with the low-noise parts of the controller nearby, either in a separate box connected by a short cable or mounted inside the cryostat itself. A longer cable carries the time-multiplexed digitized data to the computer. High-speed serial data technologies are under investigation to simplify this connection and will become imperative for the much larger format detectors that are being developed.

Several installations have constructed a safety shield in front of the detector that opens only when data are being collected. This device helps to protect the delicate front surface of the detector and is highly recommended.

Data acquisition software. There is a wide spectrum of computer configurations surrounding CCD detectors. The major tasks to be performed are operating the detector, controlling the beamline, storing raw data, correcting images and analysing diffraction patterns. In home laboratories, where exposure times are relatively long, a single PC typically handles these tasks. In another arrangement, the detector controller is really an embedded system, mostly unseen by the operator, making the detector a remote image server. The raw data or corrected images come to the user's workstation where subsequent analysis is performed. This circumvents the problem that the detector computer may be running a different operating system from the workstation. At storage-ring sources, where the data volume is very large, the detector is almost always configured as a remote image server; the user's workstation does not even need to be nearby. Clusters of remote computers that can perform tasks in parallel become attractive for streamlined data collection, correction and analysis from large data sets. Remote analysis over the internet is being explored by several storage-ring facilities.

Control software should be easy to use, but flexible and extensible. It should be easy to set up experiments and sequence the individual steps in an experiment: exposure, readout, correction, storage and crystal movement, and wavelength change for MAD experiments. Extensible software would permit a user-written macro to be run at each step in place of the detector primitive that is provided. For instance, if it were desired to collect two