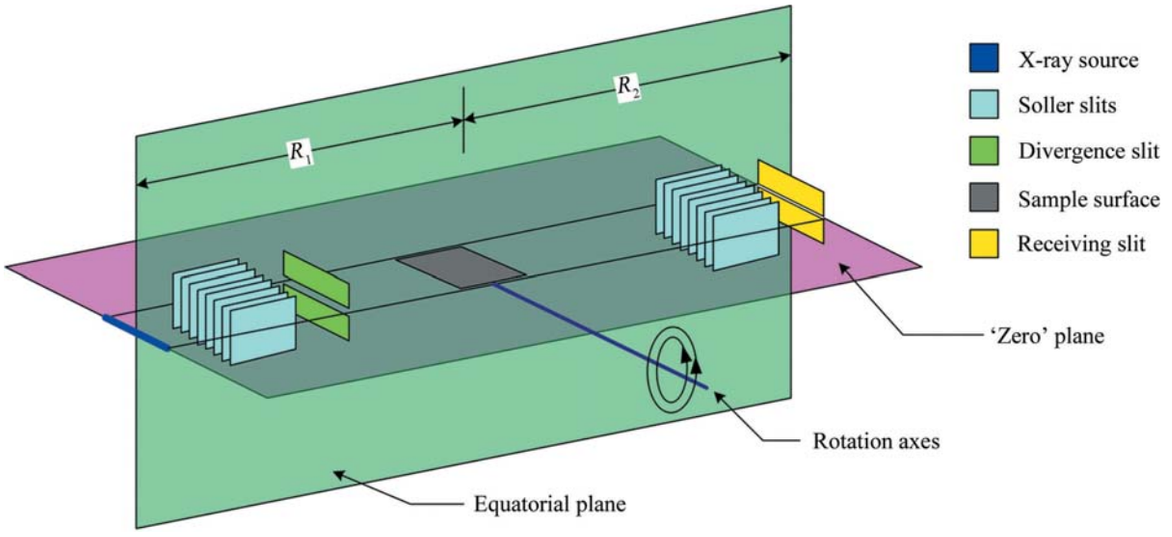

Figure 3.1.14
Diagrammatic explanation of the conditions necessary to realize a properly aligned X-ray powder diffractometer.

mathematical description of an incident spectrum. At best, a 'perfect' focusing crystal will impose an uncharacterized, though somewhat Gaussian, energy filter on the beam it diffracts. However, in certain optics the required bend radius of Johansson geometry is realized by clamping the crystal onto a curved form. The clamping restraint exists only at the edges of the optic, not in the central, active area where it is illuminated by the X-ray beam. The crystal itself however, can minimize internal stress by remaining flat; in this case an anticlastic curvature of the optic results. A 'saddle' distortion across the surface of the diffracting region of the crystal results in a complex asymmetric $K\alpha_1$ spectrum that defies accurate mathematical description. Johansson optics, however, can be bent by cementing the crystals into a preform, yielding an optic of superior perfection in curvature. Fig. 3.1.13 shows data collected from such an optic using an Si single crystal, 333 reflection, as an analyser. Parallel-beam conditions were approximated in this experiment with the use of very fine 0.05 mm incident and receiving slits. The observed symmetric emission profile of Fig. 3.1.13(*a*) can be modelled with a combination of several Gaussians. However, a Johansson optic will scatter 1–2% of high-energy radiation to a higher 2θ angle than the $K\alpha_1$ focal line of the optic. This unwanted scatter is dominated by, but not exclusive to, the $K\alpha_2$ spectrum. Louër (1992) indicated that it can be largely blocked with a knife edge aligned to just 'contact' the high-angle side of the optic's focal line. Alternatively, the NIST method is to use a slit aligned to straddle the focal line. Proper alignment of this anti-scatter slit is critical to achieving a good level of performance with the absence of '$K\alpha_2$' scatter, as illustrated in Fig. 3.1.13(*b*). As will be demonstrated, with use of any Johansson optic the elimination of the $K\alpha_2$ line is of substantial benefit in fitting the observed peaks with analytical profile-shape functions.

3.1.3. Instrument alignment

Modern instruments embody the drive towards interchangeable pre-aligned or self-aligning optics, which, in turn, has led to several approaches to obtaining proper alignment with minimum effort on the part of the user. We will not review these approaches, but instead we decribe here the methods used at NIST, which could be used to check the alignment of newer equipment. With the use of calibration methods that simply characterize the performance (which includes the errors) of the machine in an empirical manner and apply corrections, the quality of the instrument alignment may be surprisingly uncritical for a number of basic applications such as lattice-parameter refinement. However, with the use of the more advanced methods for characterization of the IPF that are based on the use of model functions, the proper alignment of the machine is critical. The models invariably include refineable parameter(s) that characterize the extent to which the given aberration affects the data; the correction is applied, and the results are therefore 'correct'. However, if the instrument is not aligned properly, the analysis attempts to model the errors due to misalignment as if they were an expected aberration. The corrections applied are therefore incorrect in degree and nature and an erroneous result is obtained.

The conditions for proper alignment of a Bragg–Brentano diffractometer (see Fig. 3.1.14) are:

(1) the goniometer radius, defined by the source-to-rotation-axes distance, R_1, equals that defined by the rotation-axes-to-receiving-slit distance, R_2 (to ± 0.25 mm);
(2) the X-ray line source, sample and receiving slit are centred in the equatorial plane of diffraction (to ± 0.25 mm);
(3) the goniometer rotation axes are co-axial and parallel (to ± 5 μm and <10 arc seconds);
(4) the X-ray line source, specimen surface, detector slit and goniometer rotation axes are co-planar, in the 'zero' plane, at the zero angle of θ and 2θ (to ± 5 μm and $\pm 0.001^\circ$); and
(5) the incident beam is centred on both the equatorial and 'zero' planes (to $\pm 0.05^\circ$).

The first three conditions are established with the X-rays off, while conditions (4) and (5) are achieved with the beam present, as it is actively used in the alignment procedure. Neither incident- nor diffracted-beam monochromators are considered; they are simply added on to the Bragg–Brentano arrangement and have no effect on the issues outlined here. Also, in order to execute this procedure, a sample stage that can be rotated by 180° in θ is required. However, this does not need to be the sample stage used for data collection. Before any concerted effort to achieve proper alignment, it is advisable to check the mechanical integrity of the equipment. Firmly but gently grasp a given component of the diffractometer, such as the tube shield, receiving-slit assembly or sample stage, and try to move it in a manner inconsistent with

its proper mounting and function. The number of defects, loose bolts *etc.*, that can be found this way, even with quite familiar equipment, can be surprising.

Let us briefly review the development of diffraction equipment and the subsequent impact on alignment procedures. The goniometer assemblies used for powder diffractometers utilize a worm/ring gear to achieve rotation of the θ and 2θ axes while allowing for the $\sim$0.002° resolution with the use of a stepper or servo motor actuating the worm gear. 'Home' switches, with a coarse one on the ring gear and a fine one on the worm shaft, allow the software to locate the reference angle(s) of the goniometer assembly to a repeatability of the stepper motor resolution. With the first generation of these automated goniometers, the zero angles were fixed relative to the home positions. With such a design the invariant reference was the receiving slit, and the operator adjusted the height of the tube shield and the angle of the θ stage to realize alignment condition (4). Second-generation machines offered the ability to set the zero angles relative to the home positions (or those of optical encoders) *via* software, in which case the exact angular position of either the X-ray tube focal line or of the receiving slit in θ–2θ space is arbitrary. The operator simply determines the positions where the θ and 2θ angles are zero, and then sets them there. There is no technical reason why the older designs cannot be aligned to the accuracy of newer ones. In practice, however, with older equipment the patience of the operator tends to become exhausted, and a less accurate alignment is accepted. An important consideration in evaluating modern equipment is that it is often the incident optic, not the X-ray source (focal line), that is used as the reference. Which situation is the case can be readily discerned with an inspection of the hardware: if the incident optic is anchored to the instrument chassis, then it is the reference. If it is attached to the tube shield, however, then the source establishes the reference. The NIST equipment has the latter design.

Condition (1) is that the goniometer radius, defined by the source-to-rotation-axis distance, R_1, equals that defined by the rotation-axis-to-receiving-slit distance, R_2. This condition is required for proper focusing and is generally realized with the use of rulers to achieve a maximum permissible error of $R \pm 0.25$ mm for a nominal $R = 200$ mm diffractometer. Condition (2) concerns the centring of the components in the plane of diffraction or equatorial plane. This condition is assured with the use of straightedges and rulers and, again for a line focus with an 8 to 12 mm source length, the maximum permissible error for deviations along the equatorial plane is $\pm$0.25 mm. One can also consider the takeoff angle at this time; this is the angle between the surface of the X-ray tube anode and the equatorial centre line of the diffractometer incident-beam path. As this angle decreases the resolution is improved at the expense of signal intensity, and *vice versa*, as a consequence of the variation in the size of the source that the specimen 'sees'. However, with modern fine-focus tubes, this is not a major effect. Qualitative experiments at NIST indicate that the exact angle is not critical; a 6° takeoff angle is reasonable.

The third issue concerns the concentricity of the θ and 2θ rotation axes of the goniometer assembly; this is a matter of underappreciated concern. It is not, however, one over which the end user has a great deal of control. Measurement of axes centricity requires the construction of some fairly complex and stiff structures capable of measuring displacements of the order of 1 to 2 μm and rotations of seconds of arc. The objective is to measure both the offset between the two axes and the angle between them. Concentricity errors affect XRPD data in a manner analogous to that of sample displacement; hence a 5 μm concentricity error is of concern. Worse yet is the possibility that some degree of precession occurs between the two axes with the operation of the goniometer. In this case, the performance of the machine will challenge description using established models.

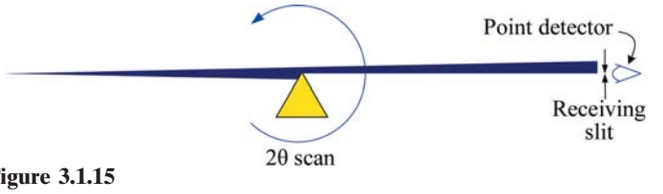

Figure 3.1.15
Diagrammatic view illustrating the use of a knife edge to determine the 2θ zero angle.

Subsequent experiments are performed with the X-rays present in order to achieve conditions (4) and (5). The criteria for proper alignment are universal, but there is a range of experimental approaches by which they can be realized. The specific approach may well be based on the age and make of the equipment as well as the inclinations of the operator. The essence of the experimental design remains constant, however: the operator uses optics mounted in the sample position that will either pass or block the X-ray beam in such a way as to tell the operator if and when the desired alignment condition has been realized. One approach is to use a knife edge mounted as shown in Fig. 3.1.15; a 2θ scan is performed using a point detector with a narrow receiving slit. When the intensity reaches 50% of the maximum, the X-ray source (focal line), the rotation axes of the goniometer and the 2θ (zero) angle are coplanar. However, the problematic presumption here is that the sample stage is aligned so exactly that the rotation axes of the goniometer assembly bisect the specimen surface, and therefore the knife edge, to within a few micrometres. This is equivalent to the z height being zero. The verification of this level of accuracy in stage alignment would be exceedingly difficult *via* direct measurements on the sample stage itself. While many would be inclined to trust the instrument manufacturer to have correctly aligned the stage, at NIST we use an alternative approach.

A straightforward means of addressing this problem is to use a stage that can be inverted, and perform the 2θ zero angle experiment in both orientations. 2θ scans of a knife edge in the normal and inverted positions can be compared to determine the true 2θ zero angle, independent of any z-height issue associated with the stage. It is often useful to draw a diagram of the results in order to avoid confusion; half the difference between the two measured zero angles yields the true one. With this information, the final alignment involves adjusting the specimen z height in the desired stage, which need not be invertible, until what is known to be the true 2θ zero angle is realized. The knife edge can also be used to centre the beam on the rotation axes, as per condition (5). Determination of the θ stage zero angle can be performed using a precision ground flat. An alternative optic to the knife edge is a rectangular 'tunnel' through which the X-ray beam passes. The entrance window of said tunnel may measure 20 to 40 μm in height and 10 mm in width, while the tunnel itself is 5 cm long. It is mounted in the beam path as illustrated in Fig. 3.1.16, with the 20 to 40 μm dimension defining the width of the beam and the 10 mm dimension describing the beam's length. Optics like this can be made of metal but are often made of glass. This optic will pass an X-ray beam only if it is parallel to the direction of the tunnel and can be used to determine both θ and 2θ zero angles. These are the optics used at NIST, *via* an experimental approach that will be discussed below.

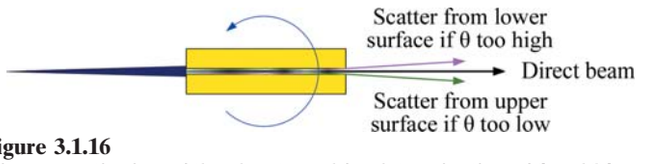

Figure 3.1.16
Diagrammatic view of the glass tunnel for determination of θ and 2θ zero angles.

If a diffractometer is being commissioned for the first time, or if major components have been replaced, it is appropriate to use fluorescent screens to achieve a rough alignment and to ensure that the incident beam does indeed cross the goniometer rotation axes and enter the detector; otherwise one may waste time looking for the beam. It is critical that these experiments are performed with the tube at operating power and that the equipment is at thermal equilibrium. Thermal effects will cause the anode to expand and contract, which will typically cause the position of the source to change. This is particularly critical when using optics to prepare the incident beam, as the performance of the optics can change markedly with movement of the source.

The objective of the first experiment using X-rays is to achieve parallelism between the line source of the tube anode, or focal line of the Johansson optic, and the receiving slit. A 5 µm platinum pinhole, which was originally manufactured as an aperture for transmission electron microscopy, is mounted in the sample position and used to image the focal line of the source onto the receiving slit (Fig. 3.1.17). This experiment is the one exception to the operating-power rule, as otherwise *Bremsstrahlung* will penetrate the platinum foil of the pinhole and produce confounding results. Success can be realized with settings of 20 kV and 10 mA; these reduced power settings are not thought to affect the angle between the tube anode and receiving slit (which is the issue addressed in this experiment). The incident slit is opened to the point at which the line source itself is imaged, not the incident slit. The Soller slits, and the post-monochromator if there is one, must also be removed to allow for the axial divergence that is needed for the success of this experiment. The pinhole images the line source onto the receiving slit; as the angle between the two decreases, progressively larger lengths of the receiving slit are illuminated during a 2θ scan. The tilt of the X-ray tube shield is varied and sequential 2θ scans are collected. As parallelism is approached, the profiles will exhibit a progressive increase in the maximum intensity value, with corresponding decreases in breadth. Conclusive results are shown in Fig. 3.1.18. It should be noted that this is a very difficult experiment to perform because the beam is essentially open and scatter is abundant. Shielding must be installed such that the detector can see only the signal that passes through the pinhole. The pinhole itself should also be shielded to minimize the area of (relatively transparent) platinum exposed to the direct beam.

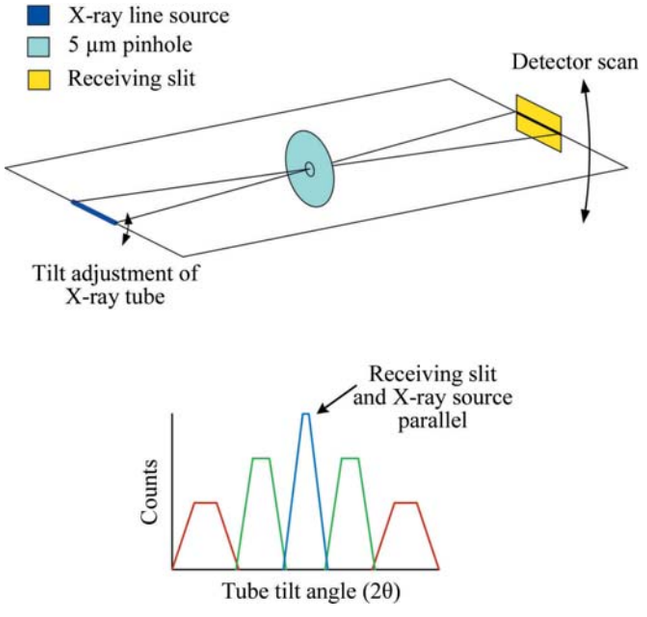

Figure 3.1.17
Design of experiments using a pinhole optic to align the X-ray source with the receiving slit.

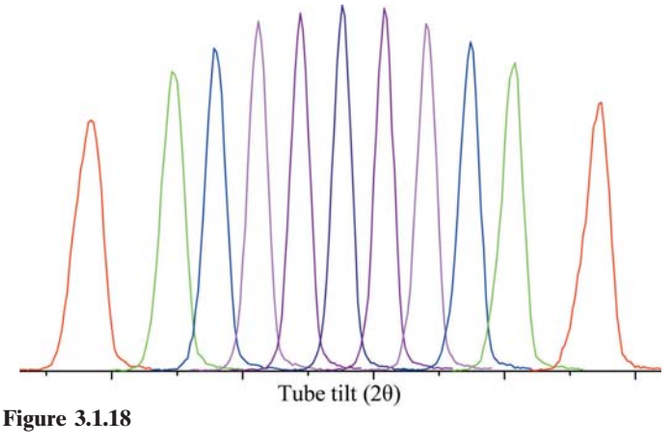

Figure 3.1.18
Successful results from the pinhole experiment showing variation in profile shape with successive adjustment of tube tilt; the central peak of highest intensity indicates the state of parallelism between the source and the receiving slit.

We now proceed to determine the θ and 2θ zero angles using the glass-tunnel optic. Initial experiments should be performed without a post-monochromator, as its presence tends to complicate finding the beam. However, it should be installed as experiments progress, as it will lead to an increase in resolution; it may alter the wavelength distribution slightly and its mass will change the torque moment on the 2θ axis. The latter two factors may alter the apparent 2θ zero by several hundredths of a degree. It is best to use a minimum slit size for the incident beam that will fully illuminate the entrance to the tunnel optic to avoid undue levels of scatter. The receiving slit should be the smallest size available, 0.05 mm in our case. The first experiment will determine a first approximation of the zero angle of θ. The tunnel optic is used, with a θ scan being performed with an open detector. Once an approximate zero angle of θ is determined, the receiving slit is installed and a 2θ scan is performed with θ at its zero point. Thus, we now have a qualitative idea of both zero angles. Then an experiment is performed as shown in Fig. 3.1.19; sequential 2θ scans are performed as θ is stepped through its zero point by very small steps (0.004° in the case of our experiment). The tunnel scatters radiation from its upper and lower surfaces when it is not parallel to the central portion of the beam, resulting in a lobe on each side of the direct beam in Fig. 3.1.19. When θ is at the desired zero angle, the direct beam is transmitted with minimum intensity in the lobes.

Once the zero positions of the θ and 2θ angles are determined, the stage is inverted and this set of experiments is repeated. It is desirable to drive the stage by 180°; however, remounting the stage in an inverted position is acceptable if the mounting structure centres the stage to within a few micrometres. Again, it is often useful to draw a diagram of the results from these two zero-angle determinations to ensure that the data are interpreted

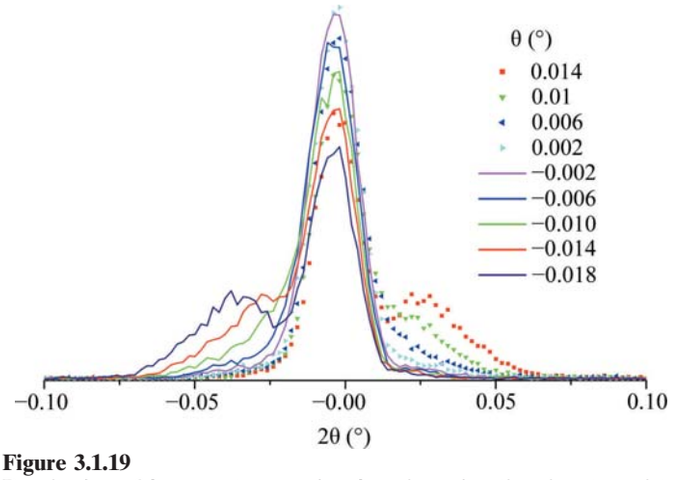

Figure 3.1.19
Results from 2θ scans at successive θ angles using the glass tunnel to determine the θ and 2θ zero angles.

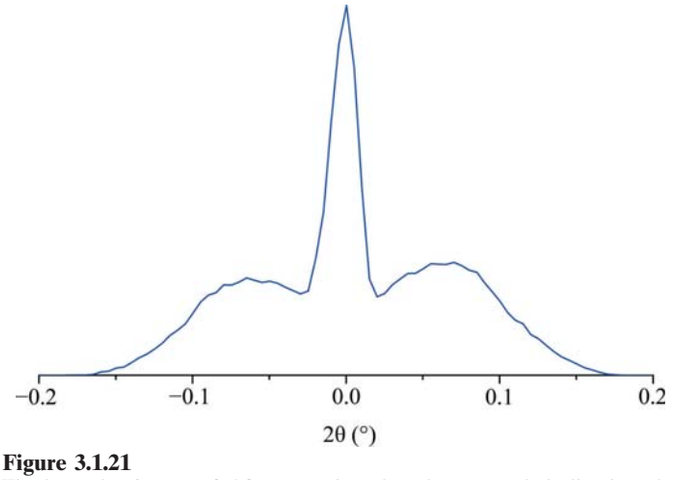

Figure 3.1.21
Final results from a θ–2θ scan using the glass tunnel, indicating the correct determination of θ and 2θ zero angles.

correctly, as shown in Fig. 3.1.20. In this example, the sample height is displaced in the positive z direction, otherwise the positions of orientation 1 and 180° from orientation 1 would be reversed. The operator should verify that fully self-consistent results are obtained with respect to the four zero angles shown in Fig. 3.1.20. Because the beam is divergent, the difference between the two θ zero angles will not be precisely 180°, as shown in Fig. 3.1.20. Again, half the difference between the two measured 2θ zero angles yields the true one, with respect only to the locations of the X-ray source and the goniometer rotation axes. Using the data of Fig. 3.1.20 and the goniometer radius, the z-height error on the stage in question could be computed and an adjustment made; this should be followed by repeating the two zero-angle measurements and checking for self-consistency to provide additional confidence in the alignment.

The final task is to mount the stage to be used in subsequent data collection and adjust its sample height until the known true 2θ zero angle is obtained. The final experiment is a θ–2θ scan of the tunnel optic to yield data of the kind shown in Fig. 3.1.21. The symmetry of the lobes on each side of the peak from the direct beam is indicative of the correct θ zero angle setting. This final high-resolution experiment is an excellent indicator of the state of the alignment of the instrument. These experiments, when used in conjunction with profile fitting, can yield measurements of the zero angles with an uncertainty for θ and 2θ of $\pm 0.001°$. Given the high certainty with which the zero angles are determined, they would then not be refined in subsequent data analyses. The alignment of the incident-beam slit, issue (5), is accomplished with a scan of the direct beam. If the machine is equipped with a variable-divergence incident-beam slit, it is important to evaluate it at several settings because changes in the centre line of the beam may occur as the divergence angle is altered. Use of an excessively narrow receiving slit should be avoided for scans of the direct beam, since the thickness of the metal blades used for the slit itself may be larger than the width of the slit, leading to a directional selectivity as the scan is performed.

The alignment presented here was carried out using a scintillation detector; however, much of it could be performed using a PSD in 'picture-taking' mode. In any case, the count rates have to be monitored to ensure that they are within the linear range of the detector (5000 to 10 000 counts per second), otherwise anomalous results are obtained. Attenuating foils that are flat and in good condition can be used to reduce the intensity. It should also be stressed that if the observations made during the experiments do not meet expectations, something is wrong and the desired outcome, *i.e.* the correct alignment, will not be realized. Drawing a diagram of the X-ray beam path can be very useful for discovering the cause of apparently unexplainable observations. Also, throughout these experiments it is appropriate for the operator to try various additional settings to ensure that the machine is operating as expected. Anomalous observations can almost always be explained in a quantitative manner with appropriate investigation. Patience is required.

In the past, achieving acceptable performance with a Johansson optic was considered so problematic that they were under-used, despite the improvements in the data quality they provided. Modern instrumentation can provide their advantages with dramatically reduced effort. The NIST Johansson IBM, however, was derived from an older design that was originally supplied with a Siemens D500, *circa* 1987. It uses a Huber 611 monochromator housing that provides 5 degrees of freedom in the positioning of the optic: the *a* distance, the takeoff angle, crystal 2θ, tilt and azimuth. For aforementioned reasons, we installed a modern Johansson optic manufactured by Crismatec (now part of Saint Gobain). There are two stages to the procedure for aligning the machine equipped with the IBM: first, the crystal optic itself is aligned

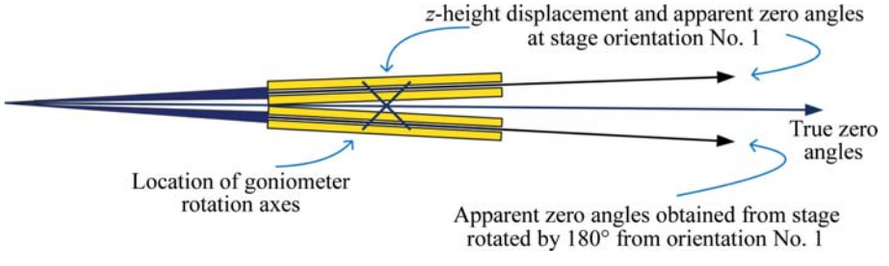

Figure 3.1.20
Diagram of hypothetical results from two zero-angle measurements (Fig. 3.1.19) with the sample stage in the normal and inverted positions to determine the true 2θ zero angle of the goniometer assembly in the absence of a z-height error from sample-stage misalignment.

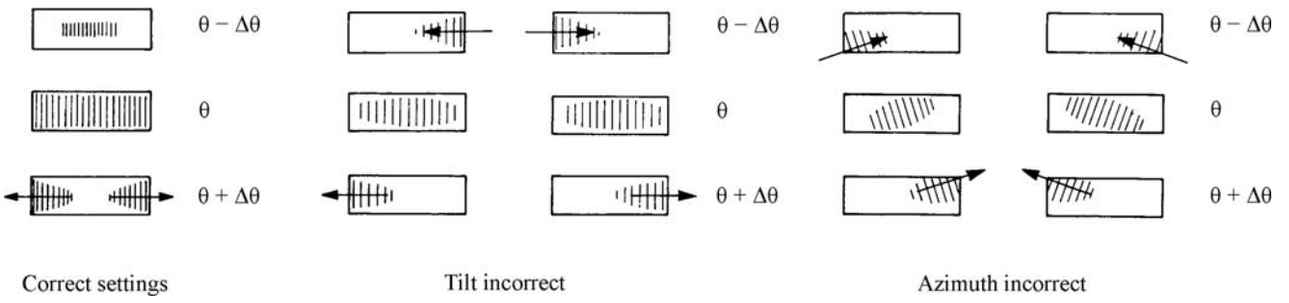

Figure 3.1.22
Figures found within the instructions for a Siemens D500 incident-beam monochromator in a Huber 611 monochromator housing, illustrating image formation and movement for correct and incorrect settings of tilt and azimuth angles (reproduced with verbal permission from Huber).

with the line source of the tube anode, and then the tube shield/IBM assembly is aligned with the goniometer. The second stage is analogous to the instrument alignment described above, so here we will discuss only the first stage (although not exhaustively).

The alignment of the Johansson optic to the X-ray source is done largely with the X-rays present. The crystal tilt and azimuth are set by using a fluorescent screen or camera to observe the diffraction images from the optic as it is rotated through its diffraction angle. Fig. 3.1.22, which is reproduced from the instructions supplied by Siemens, shows how the images form and move, informing the operator of necessary adjustments. Initially, a set of hex-drive extensions was used to drive the optic remotely through its 2θ angle. The source was operated at full power while the movement of the image was observed through a lead-impregnated window. Later, a motor drive was installed onto the 2θ actuator of the 611 housing. In the end, the incident-beam intensity realized from the optic is dependent upon the operator's ability to discern the subtleties in the image movement (Fig. 3.1.22). Blocking the axially divergent signals from the optic with a high-resolution 0.05° Soller slit dramatically improves the sensitivity of this observation to the setting of the tilt and azimuth angles. The inclusion of the Soller slit, however, will reduce the intensity markedly. A complete darkening of the room, including blocking of the shutter lights, as well as allowing time for pupil dilation, can be helpful. However, the use of an X-ray imager or a PSD in picture-taking mode improves the quality of the alignment by allowing for a more accurate interpretation of the observations.

The goal is the formation of an image in the centre of the beam path that splits symmetrically out to the edges with increasing crystal 2θ angle (Fig. 3.1.22). The directions supplied by Siemens and Huber allude to the fine adjustment [see Huber (2014) for movies] of the tilt and azimuth by examining the structure of the diffracted beam at the optic's focal point. A fluorescent screen located at the focal point and set at a 6° angle to the beam path is used to image the beam structure. With the use of the Soller slit for coarse alignment of tilt and azimuth, the desired final image for the fine-adjustment mode was, indeed, obtained. But it was not possible, even with a deliberate mis-setting of tilt and azimuth angles, to use the defective images at the focal point as a source of feedback for correcting the settings because they were too diffuse.

The Johansson optic is supplied with a and b distances that correspond to the angle of asymmetry in the diffraction planes and the bend radius. The instructions indicate that an incorrect setting in a will cause the optic's diffraction image to move up or down in the plane of diffraction with variation of the crystal 2θ angle. Again, a lack of sensitivity prevents the use of this effect as a feedback loop to set a. Alternative experiments for the optimization of the distance a of the optic were time consuming and not conclusive, so we decided to accept the supplied value for a. As before, we set the takeoff angle at 6°. A critical and quite difficult problem is the alignment of the slit located between the X-ray tube and the crystal optic (not shown in Fig. 3.1.3). This slit centres the beam onto the active area of the optic; misalignment leads to unwanted scatter from the optic's edges. It is aligned with the X-ray beam present, yielding an image of the shadow cast by the optic itself on one side, and one edge of the slit on the other. The optic is rotated in 2θ so that its surface is parallel to the X-ray beam, *i.e.* shadowing is minimized. The shadow from the second edge of the slit is obscured by the optic. Geometric considerations are used in conjunction with knowledge of the radius of curvature of the optic to obtain the correct location for the slit. A drawing is highly useful in this instance. After the installation of this slit, it is appropriate to re-check the tilt and azimuth settings, as the alignment of the optic is nearly complete.

The setting of the crystal 2θ is performed by evaluation of the direct beam, either with scans using a scintillation detector or by taking pictures with a PSD. With increasing crystal 2θ, the beam diffracted by the optic will build in the centre forming a broad profile; then the intensity on either side of the initial profile will rise, leading to the desired box form; and then intensity at the centre of the box will fall, followed lastly by the intensity at either side of the centre. This is consistent with Fig. 3.1.22. The process will repeat at half the $K\alpha_1$ intensity for the $K\alpha_2$ line. (Avoid tuning to the wrong line.) The crystal 2θ setting should be checked at regular intervals with a scan of the direct beam; this is the only setting on the IBM that has been observed to drift with time.

The final step in alignment of the IBM is the installation of the anti-scatter slit located at the focal line of the optic (Fig. 3.1.3). This is performed after the IBM assembly is aligned to the goniometer. Optimal performance of the anti-scatter slit can be expected only if it is located precisely at the focal line, which itself constitutes the smallest region within which a maximum of X-ray flux is transmitted. Therefore, the NIST alignment procedure includes an experiment using a narrow slit positioned by an x–y translator to evaluate the relative flux of the beam in the vicinity of the focal line. The y direction is parallel to the b direction (Fig. 3.1.3). A 0.05 mm slit is translated across the beam in the x direction, while intensity readings are recorded from an open detector. This process is repeated for a sequence of y distances. A plot of the recorded intensity *versus* x at a sequence of y settings will yield a set of profiles which broaden on either side of the true value of b; the narrowest, highest-intensity profile will indicate the location of the focal line. Thus, the experiment determines both the true b distance and the location in the x direction of the focal line. Once b is known, translational adjustment of the IBM

assembly may be required to locate the focal line precisely on the goniometer radius. The experiment also effectively measures the size of the focal line, in our case this was 0.15 mm. A slit of this dimension was fabricated, and the *x*–*y* translator was replaced with a standard slit retainer positioned at the desired location. The results are shown in Fig. 3.1.13.

3.1.4. SRMs, instrumentation and data-collection procedures

NIST maintains a suite of SRMs suitable for calibration of powder-diffraction equipment and measurements (NIST, 2015*a*,*b*,*c*,*d*). These SRMs can be divided into various categories based on the characteristic they are best at calibrating for: line position, line shape, instrument response or quantitative analysis, although some degree of overlap exists. The powder SRMs are certified in batches, typically consisting of several kilograms of feedstock, that are homogenized, riffled and bottled prior to the certification. A representative sample of the bottle population, typically consisting of ten bottles, undergoes certification measurements. The specific size of each lot is based on expected sales rates, mass of material per unit and an anticipated re-certification interval of 5 to 7 years. When the stocks of a given certification are exhausted, a new batch of the SRM is certified and a letter is appended to the code to indicate the new certification. Hence SRM 640e (2015) is the sixth certification of SRM 640, originally certified in 1973. The microstructural character of the SRM artifact and/or the certification procedure itself are expected to change (improve) with each renewal.

To understand the role of an SRM in the calibration of XRPD measurements and equipment, it is helpful to discuss briefly the documentation accompanying an SRM [see also Taylor & Kuyatt (1994), GUM (JCGM, 2008*a*) and VIM (JCGM, 2008*b*)]. NIST SRMs are known internationally as certified reference materials. Accompanying an SRM is a certificate of analysis (CoA), which contains both certified and non-certified quantity values and their stated uncertainties. Certified quantity values are determined by NIST to have metrological traceability to a measurement unit – often a direct linkage to the SI. Non-certified values (those lacking the word certified, as presented within a NIST CoA) are defined by NIST as best estimates of the true value provided by NIST where all known or suspected sources of bias have not been fully investigated. Both certified and non-certified quantity values are stated with an accompanying combined expanded (k = 2) uncertainty. Expanded uncertainty is defined as the combined standard uncertainty values for a given certified value multiplied by a coverage factor, k, of 2 and represents a 95% confidence interval for a given value. The combined standard uncertainties are determined by applying standard procedures for the propagation of uncertainty. The distinguishing characteristic of a NIST-certified quantity value is that all known instrumental measurement uncertainties have been considered, including the uncertainties from the metrological traceability chain. NIST defines uncertainties in two contexts: type A and type B. Type A are the random uncertainties determined by statistical methods, for example the standard deviation of a set of measurements. Type B uncertainties are systematic in nature and their extent is usually based on scientific judgment using all relevant information available on possible biases of the experiment. Assessing the technical origin and magnitude of these type-B uncertainties is a dominant part of the NIST X-ray metrology program.

XRPD SRM-certified quantity values are used primarily for calibration of XRPD measurement systems. The calibration data collected on test instruments also contain the two types of errors: random and systematic. It is the systematic measurement errors, or so-called instrument bias, that can be corrected with a calibration. Calibration is a multi-step process. First, certified quantity values are related to test instrument data. This is done by computing, from these values, what would constitute an 'ideal' data set from the 'measurement method' to be calibrated. The 'method' in this case would include the test instrument, its configuration settings and the data-analysis method to be used in subsequent measurements. Then a data set from the SRM is collected and analysed under the conditions of the method. Lastly, a calibration curve is generated by comparing the 'ideal' data set to the measured one. This would establish a correction to the instrument data and yield a calibrated measurement result. For XRPD, this correction has classically taken the form of a calibration function shifting the apparent 2θ indications. There is also the possibility that comparing the 'ideal' instrument response with the observed one indicates a mechanical, optical or electrical malfunction of the instrument. This, of course, requires further investigation and repair, rather than simply applying a calibration curve.

The generation of a calibration curve as just described can be thought of as a 'classical' calibration, and is applicable when the data-analysis procedure(s) use empirical methods to parameterize the observations. More recent, advanced methods such as the FPA use model functions that relate the form of the data directly to the characteristics of the diffraction experiment. The parameters of the model describing the experiment are refined in a least-squares context in order to minimize the difference between a computed pattern and the observed one. With the use of methods that use model functions, the calibration takes on a different form, as the collection and analysis of data can be thought of as replacing the aforementioned multi-step process. The calibration is completed by comparing the results of the refinement with certified quantity values from an appropriately chosen SRM and the known physical-parameter values that describe the optical configuration of the test instrument.

Random measurement error, describing the variation of data for a large set of measurements, can be estimated by repeating measurements over an extended period and computing the variance in the data. Furthermore, over time, one could recalibrate the system and look at the variance of the systematic bias for a given instrument, *i.e.* the rate of drift in the instrument. One would also have to investigate the sensitivity of both the random error and the variance in the systematic bias to environmental variables such as ambient temperature, power fluctuations *etc*. This systematic error variance, combined with the prior determined random error variance and the certified value and its uncertainty, provides an instrumental measurement uncertainty that can be applied to all measurements from a given instrument. Such an in-field study, however, would take years to complete. Instead, the instrumental measurement uncertainties for a given commercial XRPD measurement system are typically provided by the manufacturer, with the stated caveat that periodic calibrations should be performed *via* factory specifications. The instrumental measurement uncertainties determined through such a study are invariably much larger than those of the NIST-certified quantity values, as they contain both the instrument measurement errors (systematic and random) combined with certified quantity value uncertainties.

NIST maintains a suite of more than a dozen SRMs for powder diffraction. However, one often encounters discussions of non-institutionally-certified standards such as 'NAC' ($Na_2Ca_3Al_2F_{14}$), annealed yttrium oxide and silver behenate. Our discussions here