Powder3D IP 0.1
Tutorial

Bernd Hinrichsen, Robert E. Dinnebier and Martin Jansen
Max Planck Institute for Solid State Research

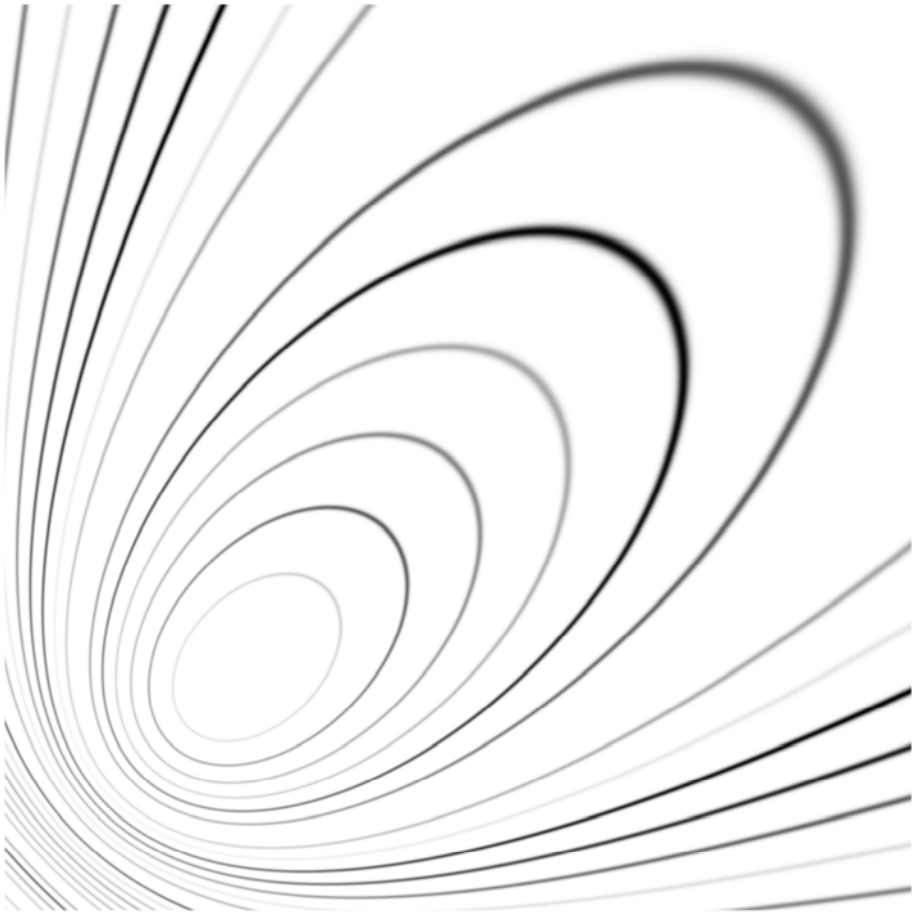

Introduction 3
Image loading 3
Opening and saving 4
Calibration 4
Filters 10
  Fractile filter 10
  Mask growth 12
  Beam stop mask 12
Intensity corrections 13
  Lorentz correction 13
  Polarization correction 14
  Incidence angle correction 14
  Background correction 15
  Display properties 15
Intensity analyses 16
  Area selection 16
  Intensity displays 17
Integration 19
References 23

Introduction

Powder3D IP is a program designed to integrate powder diffraction images from two-dimensional detectors to standard powder diffractograms. In the following a short introduction is given to the functionalities of the software. This software is similar in functionality to Fit2D(Hammersley *et al.*, 1996) and uses the same projection functions.

Installation

Powder3D IP has up to now no installation routine, but it does have one general prerequisite: the IDL virtual machine. This is similar to the Java virtual machine and can be downloaded from the site of "ITT Visual Information Solutions" for free. Once that software has been installed starting Powder3D IP is only a matter of unzipping and saving the program files to a convenient directory and double-clicking on the Powder3DIP.sav file

Image loading

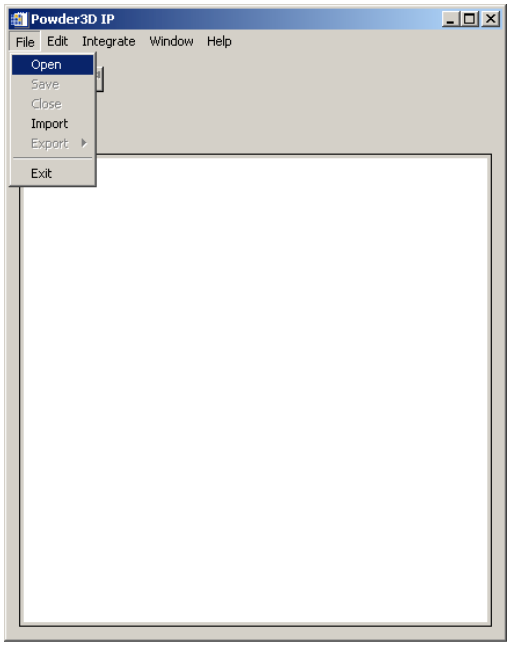

Figure 1: The initial view of Powder3D IP showing the file menu.

Loading an image can either be done using the "Open" or the "Import" function found in the "File" menu. The "Open" function can only open files written by Powder3D IP, whereas the "Import" function can import various different image types. These include binary, mar345, Stoe IPDS, tiff and bmp image formats.

Once a file has been selected and successfully loaded the screen should resemble figure 2.

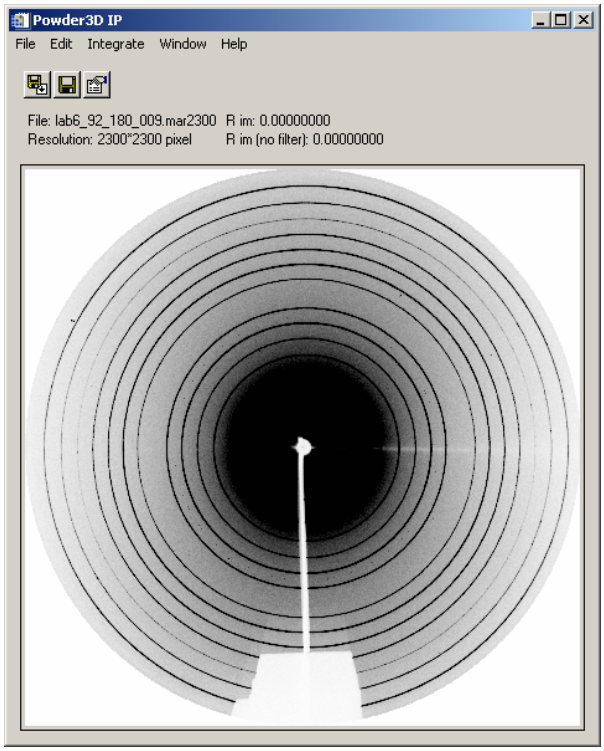

Figure 2: A calibration image has been loaded.

Opening and saving

The current data and settings can be saved using the save option "File">"Save". These files can become extremely large! They can be loaded using the open option "File">"Open".

Exporting

Images and data can be exported using the "File">"Export" functionality. Data is saved as a diffractogram, the image is saved directly from the main image display into a graphical file format.

Calibration

The first step when analyzing two-dimensional data is to calibrate the detector parameters. To do that a carefully taken image of a well prepared sample is needed. The calibration parameters are then determined as exactly as possible. These are applied to all subsequent images to extract the standard powder diffractograms. As this step affects all the subsequent data, much care should be taken during this step.

Open the calibration dialog (figure 3) by selecting "Calibrate" from the "Edit" menu. You are prompted for information on experimental details. The d-spacings of the sample are required for the calibration. These should be well known, their precision is important for a successful calibration. The wavelength and detector to sample distance (measured along the primary beam, not normal to the detector) should also be entered in the first window.

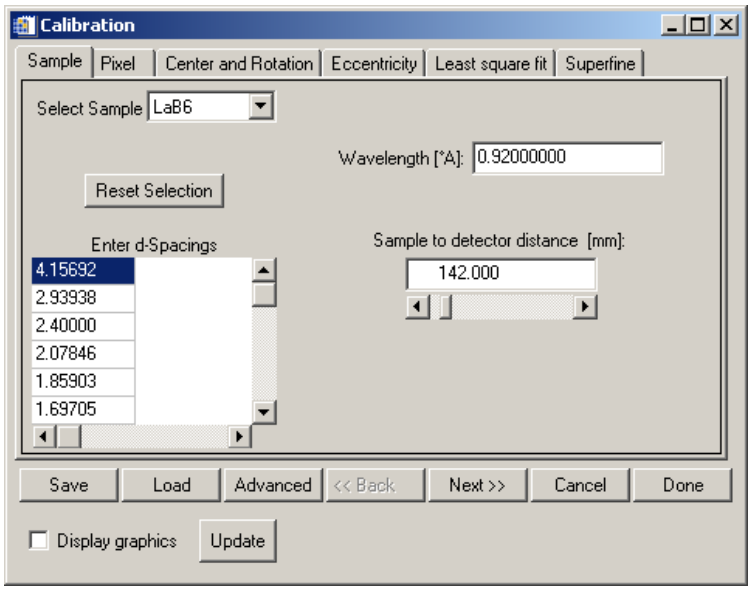

Figure 3: The calibration wizard. Setting d-spacings, wavelength and detector distance.

Press "Next" to enter the effective pixel size. This value is sometimes contained in the file header, and if it was read out it will be given here.

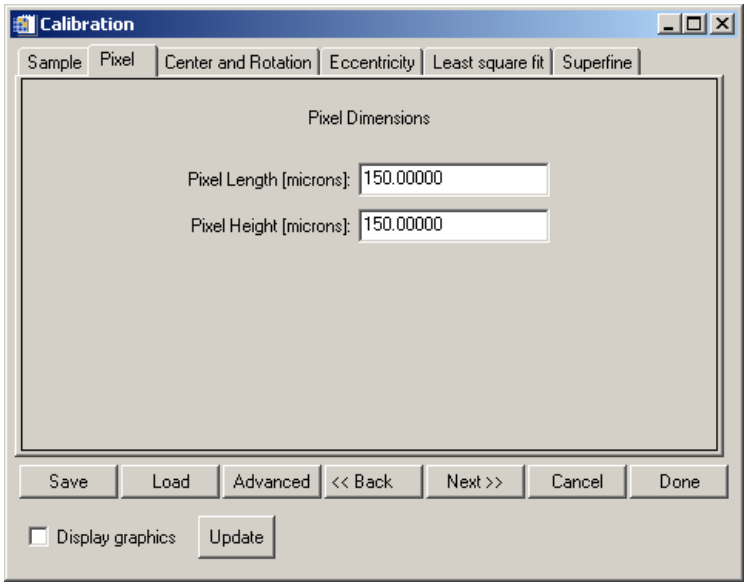

Figure 4: The calibration wizard. Setting the effective pixel size.

In the next dialog you can enter the intersection point of the primary beam with the detector. An automated function can help locate the centre precisely should it be on the image. The rotation will then be set to a sensible starting value.

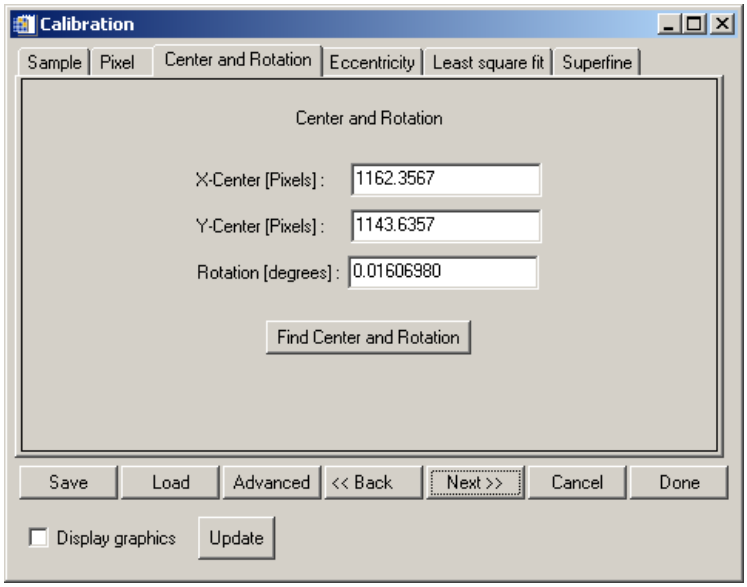

Figure 5: The calibration wizard. Setting the centre and rotation. These can be found automatically in most images.

The value of the detector tilt is then given in the following dialog. This is the tilt of the detector out of the ideal orthogonal setting. Normally this value is rather small.

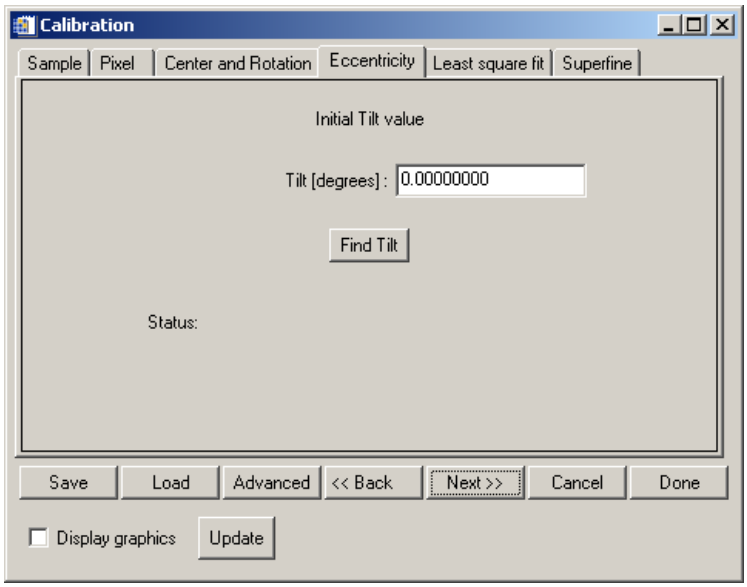

Figure 6: The calibration wizard. Setting the detector tilt.

The last step is to refine the starting values using the intersection of radial lines with the diffraction ellipses. These are refined automatically using the values determined so far. Once all the intersections have been computed they are used to refine the calibration parameters.

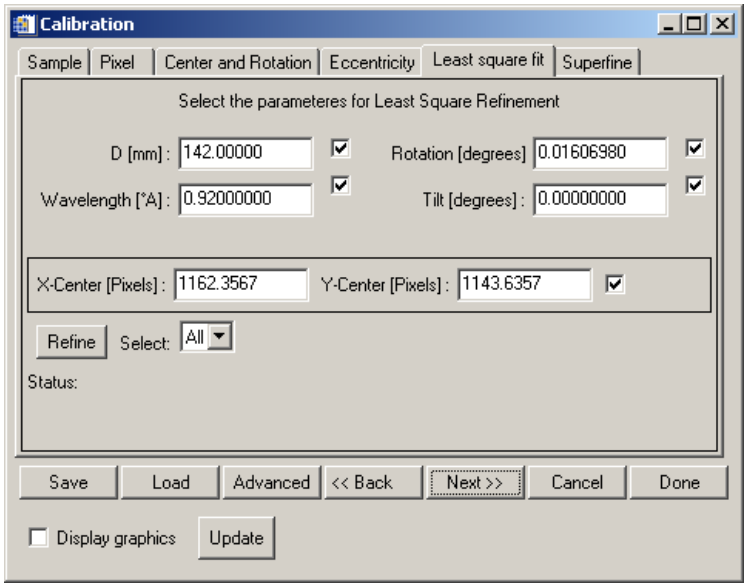

Figure 7: The calibration wizard. Refinement settings for the calibration image.

The parameters you wish to refine can be ticked using the tick boxes to the right of the parameter values.

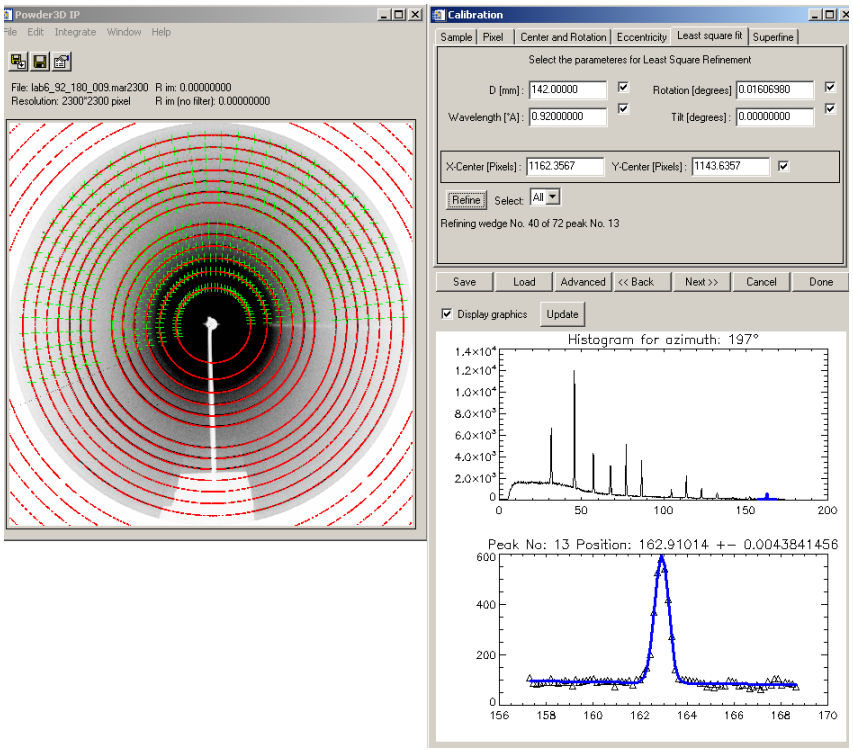

Figure 8: The calibration wizard. Refining the radial line intersections with the ellipses.

If you have chosen to have the graphics displayed while determining the intersection points you will see something like the image in figure 8.

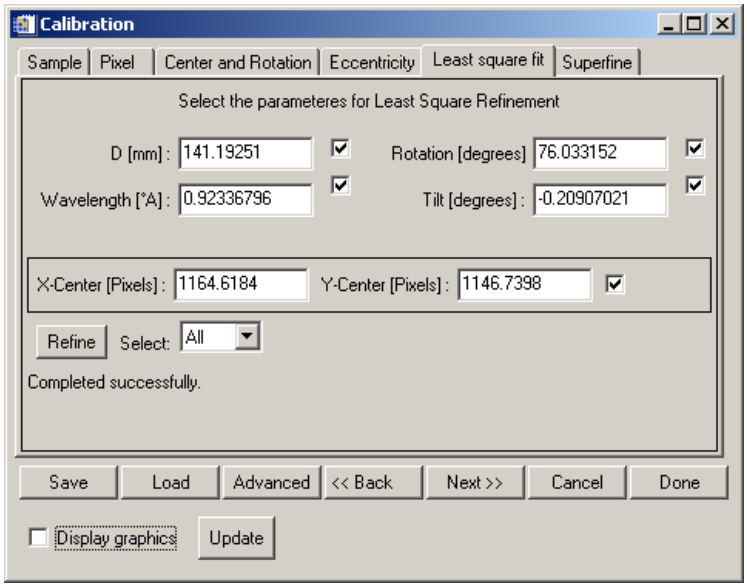

Figure 9: The calibration wizard. The refined parameters of the detector alignment.

Once the parameters have been refined the values will be updated with those from the least square refinement. If you are satisfied with the values you can now press "Done". The refined values will then be used to compute the diffraction, azimuth and incident angle for each pixel of the image. This can take some time with large images.

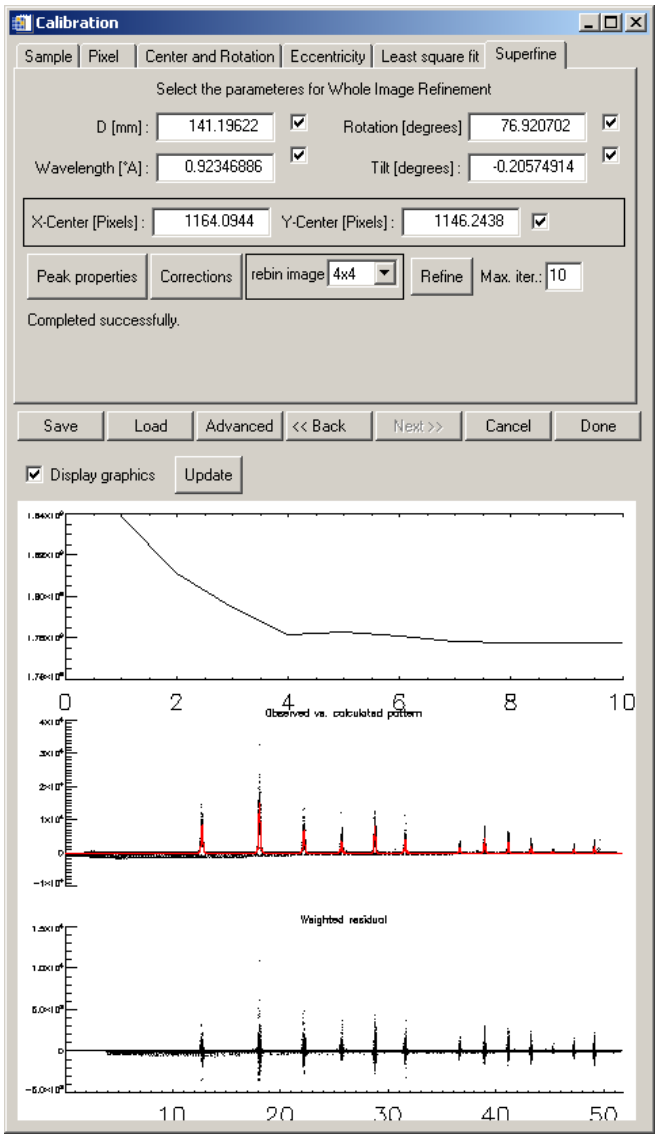

Figure 10: The calibration wizard. A converged whole image refinement.

A whole image refinement (WIR) is a rather complex refinement of an entire synthetic image against the measured image. Outliers in the measured image should already be masked, and the background should be well determined. Then the starting values from the initial refinement can be used to improve the calibration as well as refine polarization values and the detectors point spread function. Of course each peak's intensity and width has to be refined as well. This can be set by pressing the "Peak properties" button (figure 11).

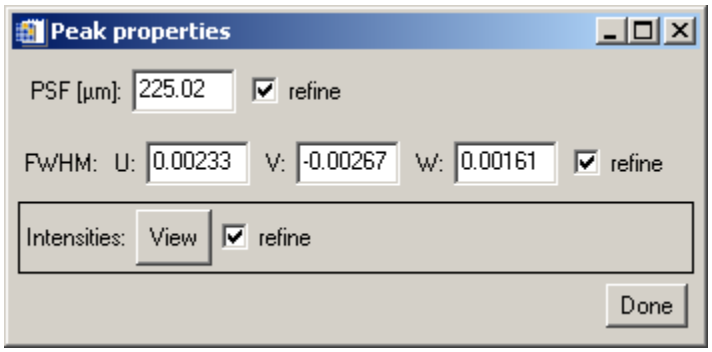

Figure 11: The calibration wizard. Setting and refining peak and detector properties.

The intensities can be viewed but not altered by pressing the "View" button.

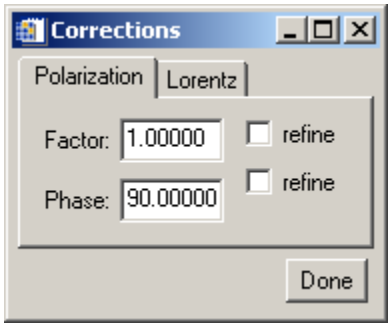

Figure 12: The calibration wizard. Setting and refining Polarization factor and phase.

The polarization factor and the phase angle can be refined, to determine these to a higher precision using WIR. The Lorentz refinement is not yet functional.

Filters

One of the most interesting features of the software is the provision of simple but powerful filters for diffraction images. In the following pages the controls will be introduced.

Fractile filter

"Edit">"Mask". A fractile filter removes a fraction of the data in an attempt to eliminate outlier data. The concept is similar to the use of the median value instead of the mean. The median is well known to be a more robust estimator than the mean, generally unaffected by strong outlier signals. In a similar vein fractile filtering ensures, by removing highest and lowest fractions of the data, that the mean value becomes a sensible estimator. The great advantage of fractile filtering is that the variance (standard deviation)^2 become more meaningful to. Fractile statistics is also referred to as quantile or percentile statistics.

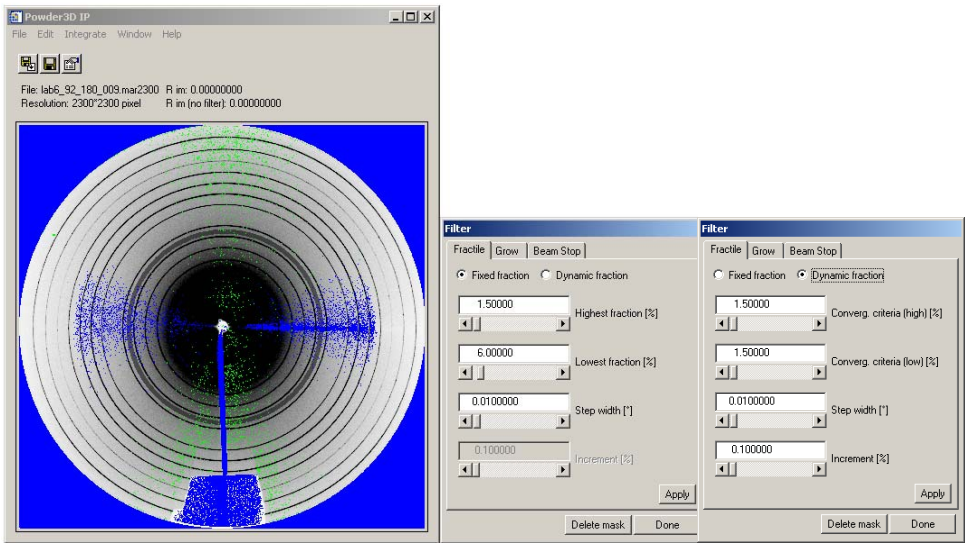

Figure 13: The mask dialog, and the effect of fractile masking on the image.

In Powder3D IP the filtering is done on the entire image with user set fractions and step width. The step width should be similar to the step width of the integration.

Another alternative offered to the setting of a fixed fraction is the dynamic fraction filter. This especially interesting when only very few striking artefacts need to be filtered. This is the case in some high temperature diffraction images in which only a few sapphire peaks (stemming from an enveloping capillary) need to be removed. Here the upper two text fields are now the convergence criteria.
This is a more computationally expensive filter which works in the following fashion. For each step width (bin) the filter fraction is enlarged incrementally by the value entered in the last field. The variance of the filtered population is compared to the previous value, and if the change is less than the convergence criteria the filtering of that bin is complete. The advantage of this filtering method is that, if the values have been chosen well, it is less aggressive than fixed fractile filtering, which removes a lot of the good signal when filtering only a few strong outliers.

Mask growth

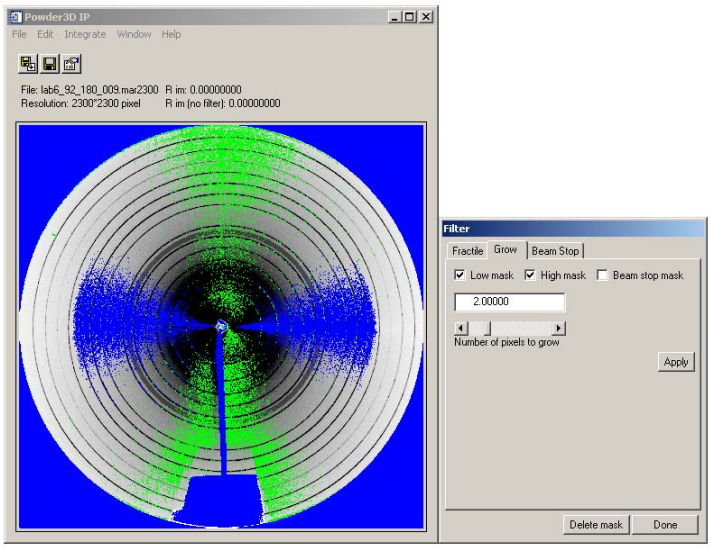

Figure 14: The mask dialog, and the effect of mask dilatation on the image.

"Edit">"Mask". Mask growth is a useful method to ensure a good coverage of outlier peaks or instrument shadow (figure 14). Generally the peak spread of the detector ensures that outlier peaks cover many neighbouring pixels. The mask might however not cover the full extent (tails) of the outlier peak. Growing the mask ensures that this happens.

Beam stop mask

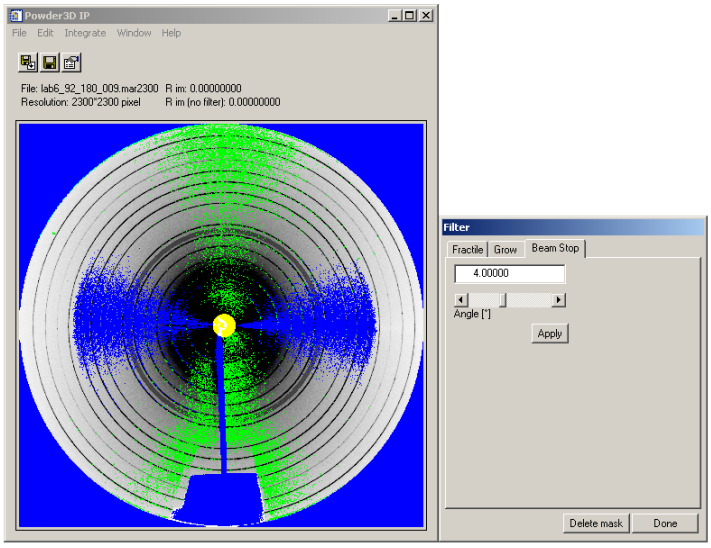

Figure 15: The mask dialog, and the effect of the primary beam stop mask on the image.

Large intensity aberrations exist close to the primary beam stop. Masking this region is done by simply selecting the first angles of the diffraction image to be disregarded (figure 15).

Intensity corrections

The intensity corrections of two-dimensional diffractograms are slightly more complex than the equatorial correction functions. It is sensible to apply these corrections before integrating the image, as they reduce the data variance as well as ensuring more accurate integrated data. These corrections then need not be applied by the Rietveld software.

Lorentz correction

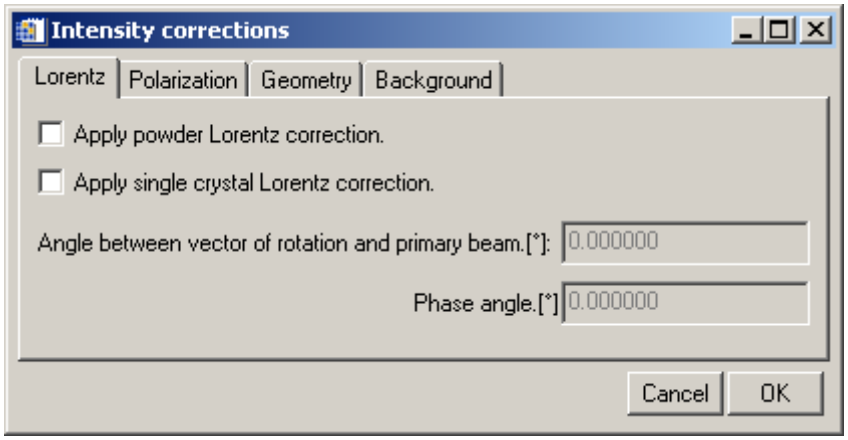

Figure 16: The separate Lorentz corrections can be made. The statistical powder correction and the angular speed (single crystal) correction can be computed separately.

"Edit">"Corrections". Traditionally Lorentz factors(Buerger, 1970, Zevin, 1990), which originally only were only angular speed corrections, have been conveniently combined in powder diffraction formalisms with diffraction probability factor $(\sin\theta)^{-1}$ and the polarization correction to the enigmatic *LP* correction. As in two dimensions no current "simple" LP corrections exist or are even appropriate, they are separated here.
The angular speed correction is called the "single crystal correction", here the angular speed is dependent the angle of the rotational vector to the primary beam (generally 90°), as well as the "phase angle". This angle is the projection of the rotational axis on the detector (more simply: the capillary shadow) relative to the detector coordinates. 0° in detector coordinates is at 3 o'clock.
It should be mentioned that the "single crystal Lorentz correction" does become undefined in the region close to the "capillary shadow". This effect is well known in single crystal area detector reduction. This region cannot be used for integration and should be excluded.
The "powder Lorentz correction" is the simple probability factor correction of $(\sin\theta)^{-1}$.

Polarization correction

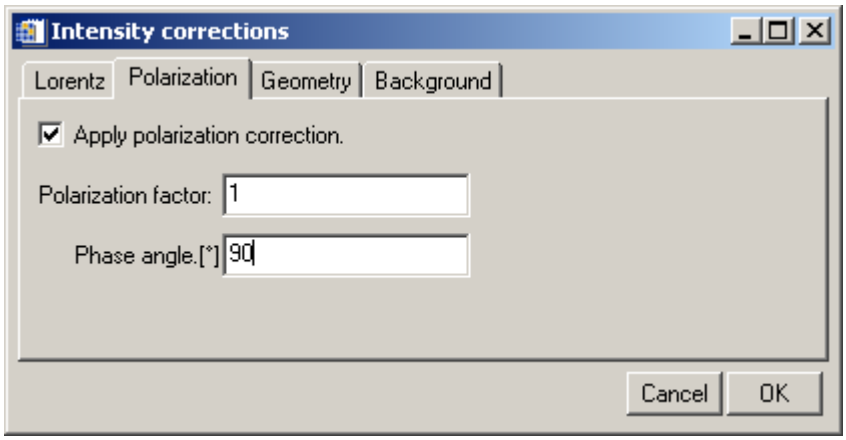

Figure 17: The polarization factor and its phase angle, which is dependent on the detector setup relative to the polarization plane.

"Edit">"Corrections". The polarization factor (Azaroff, 1955, 1956, Kahn *et al.*, 1982, Lipson & Langford, 1999, Whittaker, 1953) is determined by the intensity relations of the horizontally and vertically polarized radiation as well as the monochromator angle. For polarization factors close to 0 the polarization correction has little or no azimuthal variation, therefore the phase angle becomes unimportant. For values close to 1 (as is typical for synchrotron radiation) there is a very strong azimuthal dependence. The correct polarization phase angle is then very important. Normally it is at right angles to the detector coordinates (0°, 90°)

Incidence angle correction

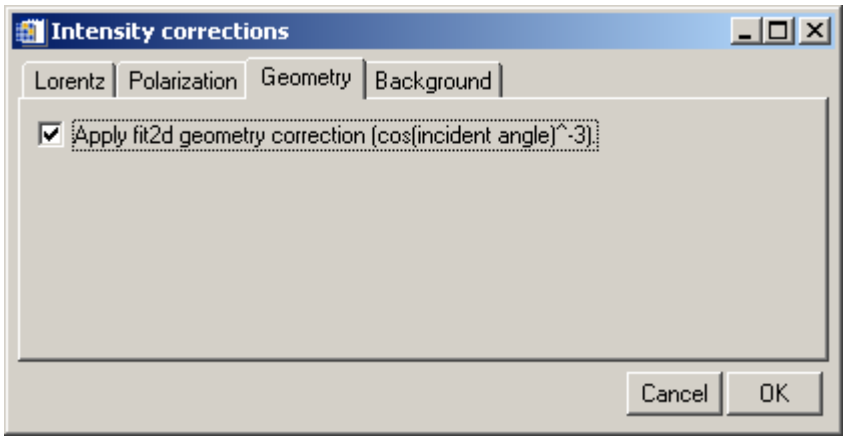

Figure 18: The incident angle correction can be applied using this tick box. This is similar to the correction in the Fit2D software.

"Edit">"Corrections". Various incidence angle correction functions are in use in single crystal diffraction (Tanaka *et al.*, 2005, Wu *et al.*, 2002, Zaleski *et al.*, 1998). The only implemented one currently is the rather simple $\cos^3$(incident angle) correction that is known as "geometry correction" in Fit2D (Hammersley *et al.*, 1996).

Background correction

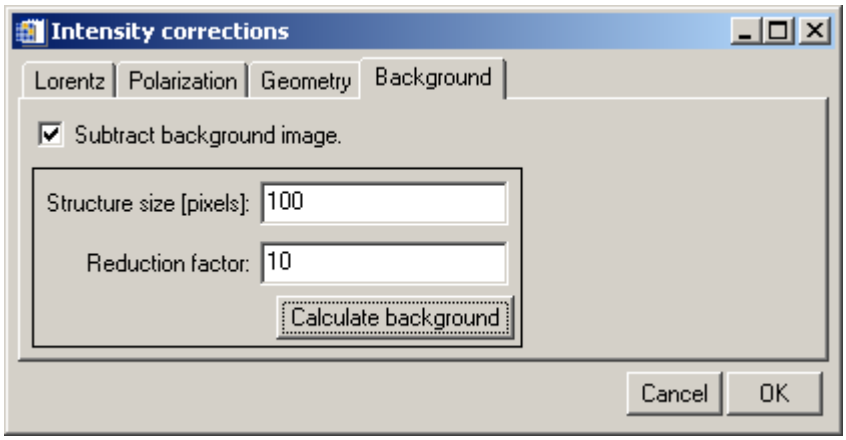

Figure 19: The background can be computation can be configured and started using this dialog.

"Edit">"Corrections". A median algorithm is used for the background determination (figure 19). The structure size is the diameter of the circle which is sampled to calculate the median. The background pixel in the centre of this circle then takes the value of the median. It is clear that the larger this value is the higher the computational cost of the background calculation is. The reduction factor is a simple method to keep this cost low. The image is initially reduced by this factor, then the background is determined, the background is then scaled up to the original image size again.

Display properties

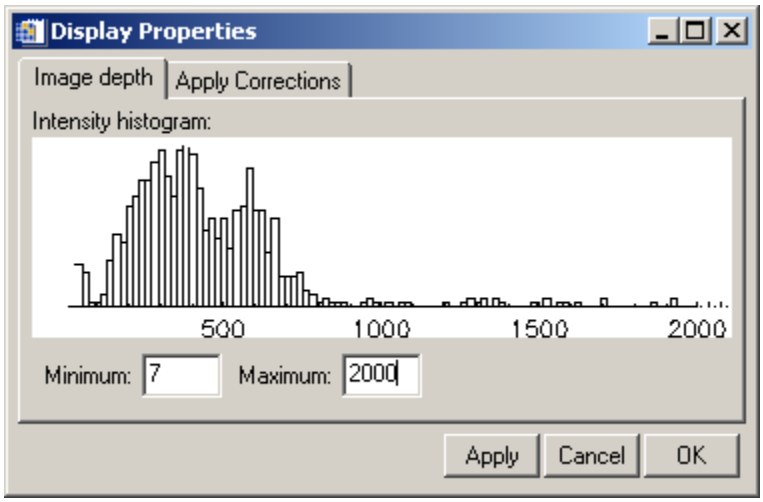

Figure 20: The display dialog for setting the optimal display range.

"Edit">"Display". To aid in selecting an appropriate intensity range for the display of powder diffraction images, an interactive histogram is made available showing the currently selected intensity distribution. By pressing the "Apply" button, the user can see the effect on the image.

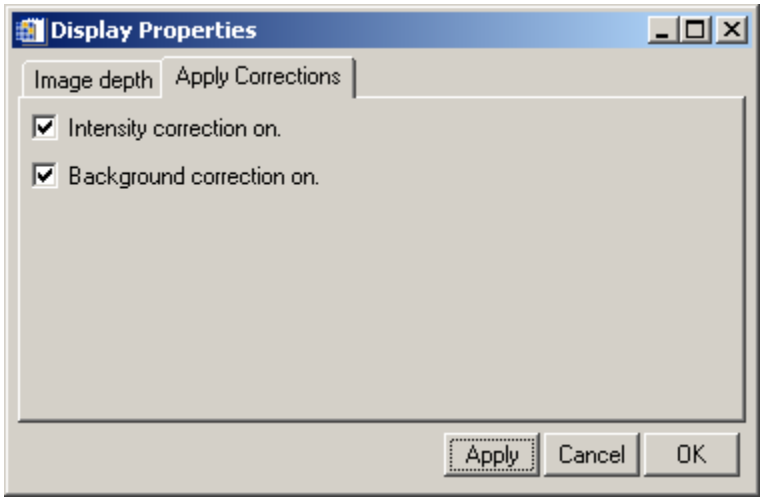

Figure 21: A direct appreciation of the effect of the intensity corrections applied to the image can be previewed by selecting the appropriate tick box.

"Edit">"Display". To estimate the correctness of the background determination as well as the polarization correction these can be applied to the displayed image. The default setting is not to display intensity corrections.

Intensity analyses

A few tools are provided to allow a closer look at the intensity distributions over the image.

Area selection

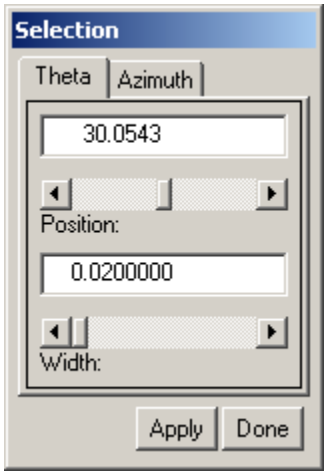

Figure 22: The area selection interface. Selecting the theta region.

The area selection tool is the most fundamental setting for intensity analysis. The central theta value is set in the top field, while the width of the theta region is set in the lower field. The selected area then covers a region from theta-width/2 to theta+width/2.

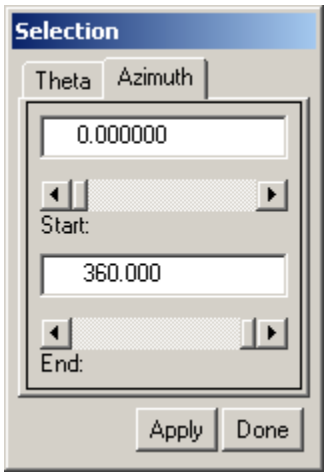

Figure 23: The area selection interface, selecting the azimuth region.

Intensity displays

Once the region of interest has been selected it is possible to view data within this region as a histogram or as a function of the azimuth.

Intensity histogram

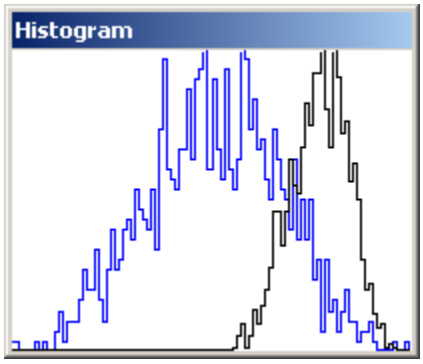

Figure 24: A histogram of two intensity distributions. The black distribution is the raw data. The blue distribution is the masked data. The two distributions do not share the same x-axis.

"Window">"Histogram". The histogram provides insight the intensity distribution of the selected region as well as the effect that filtering has on the intensities in that region. Generally the intensities describe a normal "Gaussian" distribution. Sometimes different distribution models are better. A more careful analysis of this data is possible by simply clicking the right mouse button within the histogram window and selecting "Analyze". The data is then transferred to iTools® a software provided by ITTVIS free of charge.

One the most useful tools (Hinrichsen *et al.*, 2007) within the histogram display is accessed via the context menu as well. This is the normal Pareto distribution fitter which estimates the high intensity fractile setting to achieve a more or less normal intensity

distribution. To start it is generally a good idea to filter a few percent off the top figure 25.

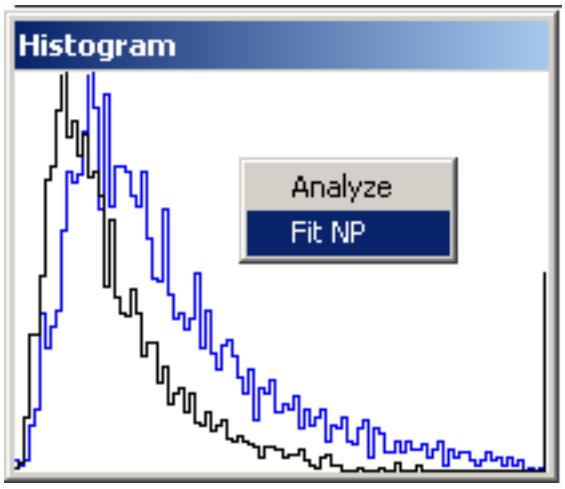

Figure 25: A histogram showing an exemplary normal Pareto distribution. Using the context menu item "Fit NP" an analysis of the distribution is made.

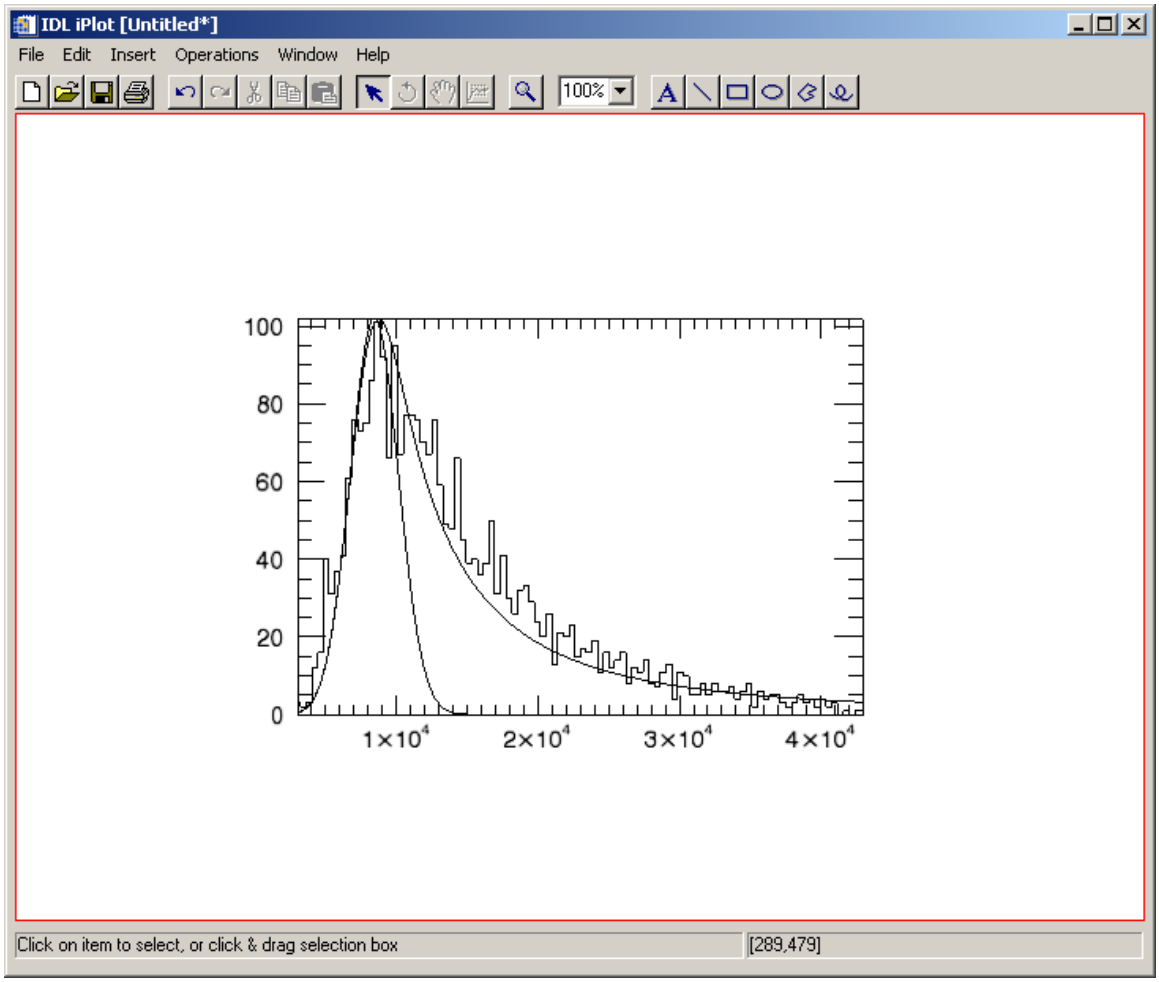

Figure 26: The distribution curves computed during the analysis are displayed for closer perusal.

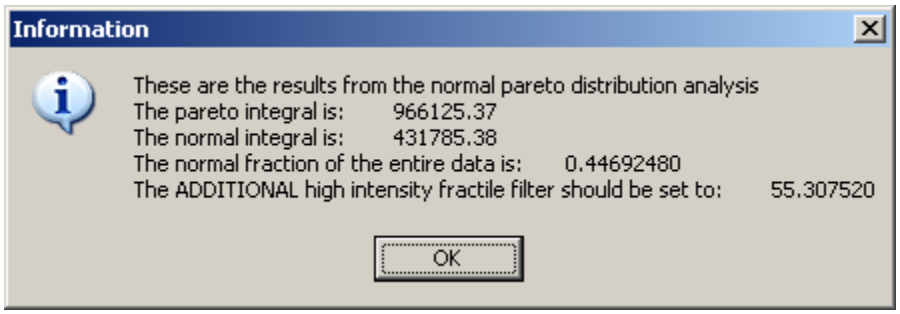

Figure 27: The really nitty gritty information is displayed by an information message box. In this case it recommends us to cut off an ADDITIONAL 55% off the top intensity to achieve a normal distribution. This leads to a whopping 59% = 55% + 4% (initial filter) top fractile.

After clicking on the "Fit NP" context menu item the window shown in figure 26 is displayed. The reason for this is to give you an idea of the reliability of the filtering suggestion. If the normal Pareto curve (the skew one) describes your data well, then the chances are good that the suggestion in the information box figure 27 is sensible. Of course the Gaussian curve the automatic routine tries to fit into the normal Pareto curve should fit snugly, as in figure 26.

Azimuthal intensity display

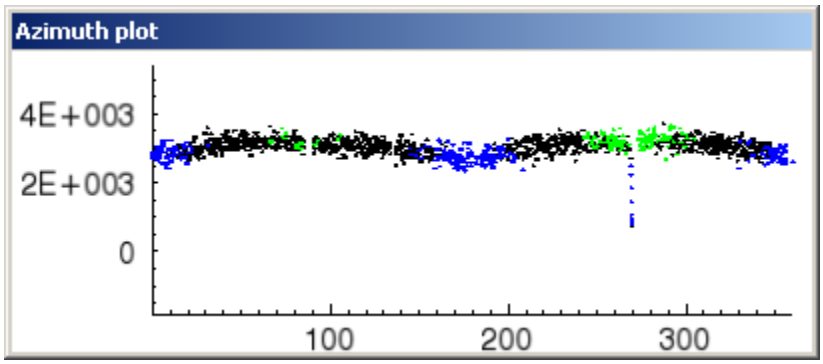

Figure 28: An azimuth plot of the intensities within the selected area

"Window">"Azimuth". In order to see which and how azimuthal corrections should be made to the data, the azimuthal plot is most useful. The coloured intensities are those masked by the filters. Here, similarly to the histogram window, the data can be analyzed with iTools by a right click.

Integration

The integration is the final step of the two-dimensional powder data reduction. In general a large number of images need to be reduced to powder diffractograms. These can be named as well as the output directory and output format. Filters and intensity corrections can be applied to all images in the process. An arbitrary number of integration bins can

be selected. One image can be reduced to a number of diffractograms using wedge integration. Lastly the method of intensity extraction can be switched between mean and median.

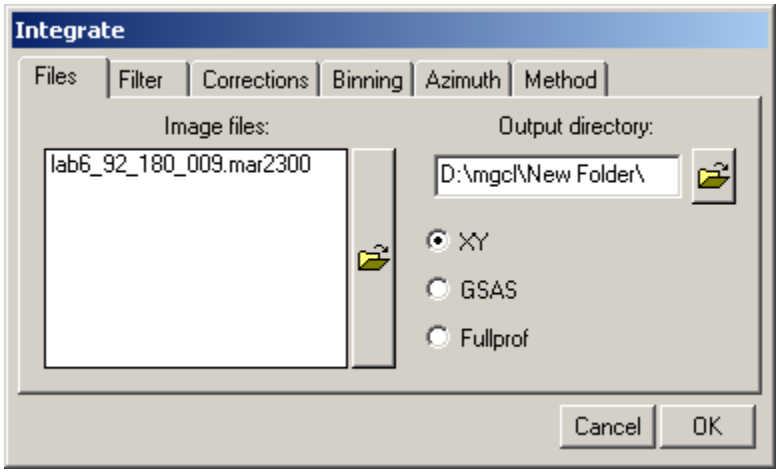

Figure 29: The integration interface. On the left the input files can be entered. On the right the output directory and the output format are set.

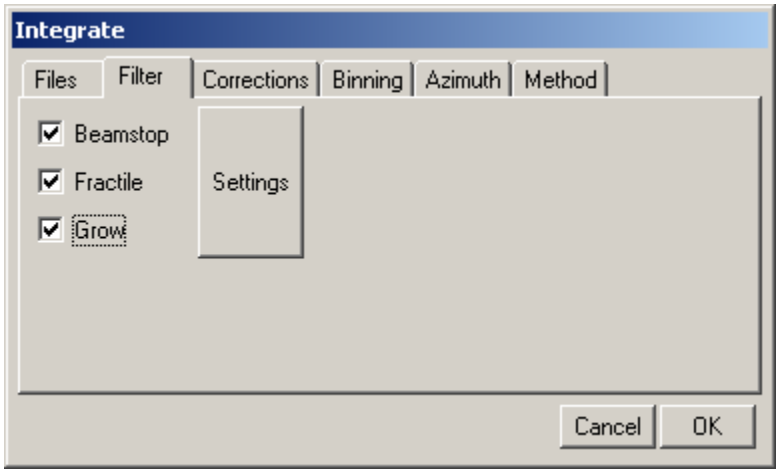

Figure 30: On the second integration tab the applied filters can be set. By pressing the "Settings" button the details of the filters can be set. These are applied individually to each image before integration.

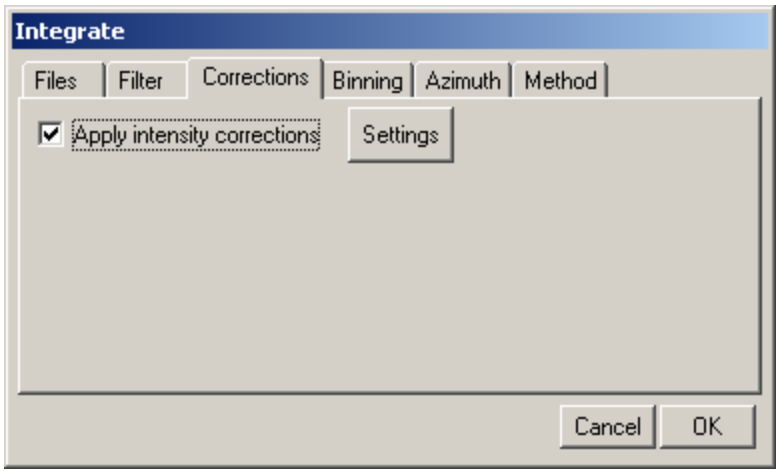

Figure 31: The intensity corrections that have been computed are applied to each image before integration. These are computed only once after the calibration and are applied quickly.

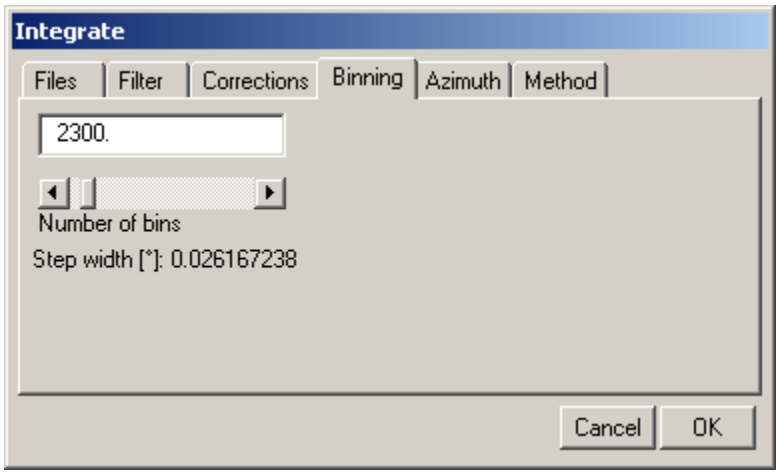

Figure 32: In this tab the number of bins is set. As an aid the effective step width is computed and displayed.

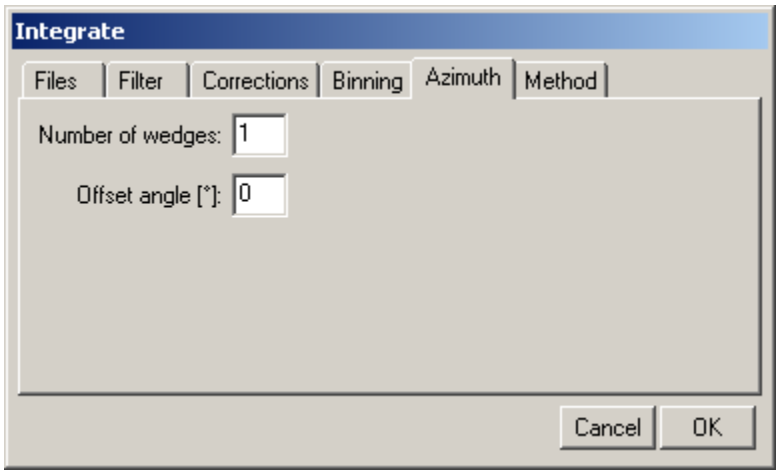

Figure 33: The integration can be divided into any number of wedges. An offset angle can be set if required. 0° is 3 o'clock.

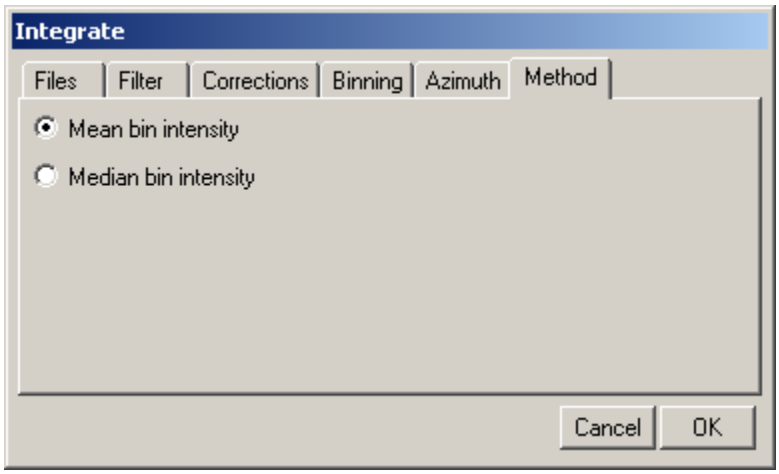

Figure 34: The intensity of each bin can be computed as mean or as median value. The median is known to be a more robust estimator than the mean. However the default setting is the mean.

References

Azaroff, L. (1955). Acta Crystallographica 8, 701-704.

Azaroff, L. (1956). Acta Crystallographica 9, 315.

Buerger (1970). Contemporary Crystallography. McGraw-Hill.

Hammersley, A. P., Svensson, S. O., Hanfland, M., Fitch, A. N. & Häusermann, D. (1996). High Pressure Research 14, 235-248.

Hinrichsen, B., Dinnebier, R. E. & Jansen, M. (2007). Journal of Applied Crystallography submitted.

Kahn, R., Fourme, R., Gadet, A., Janin, J., Dumas, C. & Andre, D. (1982). Journal of Applied Crystallography 15, 330-337.

Lipson, H. & Langford, J. I. (1999). Trigonometric intensity factors, International Tables for Crystallography, Vol. C, edited by A. J. C. Wilson & E. Prince, pp. 590-591.

Tanaka, K., Yoshimi, T. & Morita, N. (2005). Acta Crystallographica Section A 61, C146.

Whittaker, E. (1953). Acta Crystallographica 6, 222-223.

Wu, G., Rodrigues, B. L. & Coppens, P. (2002). Journal of Applied Crystallography 35, 356-359.

Zaleski, J., Wu, G. & Coppens, P. (1998). Journal of Applied Crystallography 31, 302-304.

Zevin, L. (1990). Acta Crystallographica Section A 46, 730-734.

Written by Bernd Hinrichsen, 20.02.2007
Last updated by Bernd Hinrichsen, 31.03.2008