N·H

ELSEVIER

NIM **B**

Beam Interactions
with Materials & Atoms

The Guelph PIXE software package II

J.A. Maxwell, W.J. Teesdale, J.L. Campbell *

Guelph–Waterloo Program for Graduate Work in Physics, University of Guelph, Guelph, Ontario N1G 2W1, Canada

Abstract

Modifications and improvements made to the software package GUPIX in the period since it was first described in 1989 are reported. The code now handles thin, intermediate, thick and multi-layer specimens, and permits the inclusion of "invisible" elements and complexes that may be independent or stoichiometrically related to elements whose X-rays are visible in the spectrum. M X-rays have been added to the data base, and are included in secondary fluorescence computation. The nonlinear least squares spectrum fitting procedure has been altered to conclude with a linear least-squares section which provides better reproducibility. The conventional statistical weighting has been augmented by a systematic weighting scheme that deals with uncertainties in low-energy tailing features on the spectral peaks. A new, exact treatment of peak area errors gives good agreement with the earlier approximate approach. The influence of the digital filter used to remove spectrum background upon errors and detection limits is examined, and recommendations provided on the appropriate filter dimensions. Input and output options have been expanded, and various minor changes made.

1. Introduction

The software package GUPIX [1] was introduced in 1989 specifically to treat PIXE spectra from thick specimens, i.e. specimens of sufficient thickness to reduce the incident proton beam to zero energy. It provided nonlinear least-squares fitting of the spectrum, together with subsequent conversion of the fitted X-ray peak intensities to elemental concentrations via a defined standardization technique involving fundamental parameters and a user-determined instrumental constant. Full account was taken of matrix effects and secondary fluorescence contributions in both the spectrum fitting portion and the calculation of concentrations. The GUPIX package has continued to evolve, reflecting the needs both of our own research group and of a number of groups which have adopted GUPIX for their own use. The object of this paper is to summarize a number of recent changes and improvements, using examples to illustrate their impact.

Although an early version of GUPIX did exist for DEC minicomputers, this was not updated, and instead development was focused on the PC version, taking advantage of the steady growth in processing power. One outcome of the various changes and additions is that the code has become much larger. For a period of time this was dealt with by breaking the overall code into linked sections, which were moved in and out of memory as required. However, we have now rejected this approach, preferring to use extended memory; the small additional hardware cost is more than balanced by the convenience afforded to us in continued program development.

The minimum configuration for the present GUPIX94 executable code is thus a PC with 486 (or 386/7) processor with 4 Mb of memory, 7 Mb free on the disk, and EGA graphics. The code has been adapted for the Lahey F77L3 extended memory compiler. It will handle spectra of length up to 2048 channels (although 1024 is the norm) and the maximum number of elements that can be fitted in a spectrum is 60. Whereas the initial code required the spectra to be in a defined ASCII file, GUPIX94 recognizes spectrum files not only in that original format but also in Nucleus PCA binary format, the ND/Canberra 6700 binary format, and a binary format similar to that used by the ND66 pulse height analysis system.

2. Brief overview of GUPIX

GUPIX determines the intensities of characteristic X-ray peaks in a PIXE spectrum by fitting a model spectrum M_j to the measured spectrum Y_j using the nonlinear least squares technique. The model spectrum is constructed

* Corresponding author, tel. +1 519 824 4120, fax +1 519 823 2808, E-mail: jlc@physics.uoguelph.ca.

using a data base of K, L and (as of now) M X-ray energies, fluorescence and Coster–Kronig probabilities, and relative line intensities, and the latter are modified to reflect the effects of detector efficiency, absorber effects and matrix effects. The matrix effects in turn are computed using a data base of proton ionization cross sections and stopping powers, and X-ray mass attenuation coefficients. The continuous background, however, is not modelled, and so the model spectrum contains only characteristic X-ray contributions.

The continuous background is dealt with by applying a simple digital filter operator to both the measured and the model spectra prior to least-squares fitting. This operator, like a frequency band-pass filter, is designed to attenuate the low-frequency spectral components (continuum) while passing the higher frequencies (peak structure). It was introduced by Schamber [2] and has been used in X-ray emission analysis with a variety of X-ray excitation modes. To our knowledge, GUPIX is the only PIXE analysis code which employs this method. The symmetric, zero weight filter, overall width UW + 2LW, transforms Y_j to

$$Y_j^1 = \sum_{s=-t}^{t} f_s Y_{j+s},\tag{1}$$

where j represents channel number, t is UW/2 + LW, $f_s = 1/UW$ in the central lobe, and $f_s = -1/2$ LW in the outer lobes. It is shown, along with its effect on a simple spectrum, in Fig. 1. The filter operation alters the shape of the Gaussian peak and reduces the linear background to zero. By omitting a background component from the model spectrum, we are effectively assuming that the continuum varies sufficiently slowly that it is close to linear in any local region and would therefore be eliminated by the filtering operation.

In the subsequent fit, the chi-squared is obtained by comparing the filtered spectra, i.e.
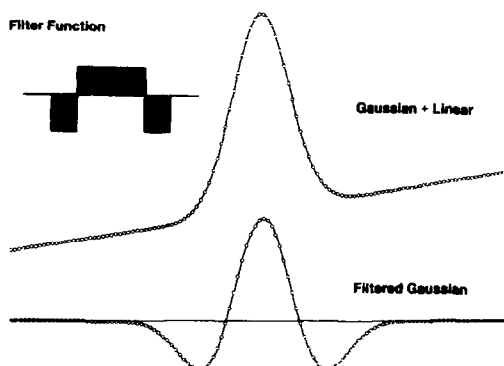
$$\chi^2 = \sum_j \left(Y_j^1 - M_j^1 \right)^2 w_j,\tag{2}$$

where the weight at each channel is w_j, and $(\sigma_j^1)^2$ is the variance of (Y_j^1):

$$\left(\sigma_j^1\right)^2 = \sum_m \left(\frac{\partial Y_j^1}{\partial Y_m}\right)^2 \sigma_m^2.\tag{3}$$

Attributing observed errors to counting statistics only $(\sigma_j^2 = Y_j)$ then

$$\left(\sigma_j^1\right)^2 = \sum_{s=-t}^{t} f_s^2 Y_{j+s}.\tag{4}$$

The fit proceeds by minimizing the value of chi-squared with respect to each of the fitting variables. These comprise the height of the principal peak of each element whose X-rays are present in the spectrum, together with the parameters (see below) that govern the relationships between X-ray energy and peak centroid and width; by allowing the latter nonlinear parameters to vary, one is able to accommodate small changes and drifts in the calibration of the pulse processing system. In addition, one or more additional variables can be introduced to represent the height of background steps due to differential absorption at absorption edges of matrix elements. While the measured spectrum need only be filtered once, the model has to be filtered prior to each loop of the fit.

From the fit results, the peak intensities and their associated errors are calculated. These are converted to elemental concentrations (either bulk concentrations or areal densities depending on whether the specimen is thick or thin) and corresponding error estimates by a procedure that is essentially a fundamental parameter approach normalized by a user-determined instrumental constant. The constant may be determined with single or multi-element standards or standard reference materials as the user prefers. In addition to concentrations and their uncertainties, the program output also presents limits of detection (LOD) for each element in the specific case at hand; this is preferable to a more generalized definition of LODs, because for a particular element the LOD can be influenced strongly by the concentrations of its neighbours in the spectrum.

Finally, the package allows the user to interpolate in the data base, and to generate detector efficiencies, absorber transmission, and other experimental parameters. Characteristic X-ray yields from a defined specimen can be computed, with the secondary fluorescence contributions identified.

3. Spectrum model and detector lineshape

The individual X-ray lines are represented, as before, by Gaussians, with a selection of four features available to describe the low-energy tailing contributions [3] of the particular Si(Li) detector; these comprise two exponential tails, a long-term shelf extending to zero energy and a short-term shelf with user-defined length. Each of these

Fig. 1. Top-hat filter and its effect on a spectrum comprising a Gaussian peak and a linear background.

features is automatically convoluted with a unit-area Gaussian resolution function to round off non-physical discontinuities.

In the original version, the user was required to parameterize the various height and slope coefficients of the incomplete charge collection tailing components (exponential and shelf features) as functions of X-ray energy, a requirement that was somewhat inconvenient and not absolutely necessary. Now these tailing parameters are simply entered once into the detector description file as lists of the actual values corresponding to a vector of defined X-ray energies. (This file also contains quantities such as window and electrode thickness required for calculating detector efficiency). The code uses log–log interpolation to calculate the tailing parameters for the particular X-ray energies that arise in any given PIXE spectrum. No further user intervention regarding tailing is needed; the code economizes on time by neglecting tailing for any peak whose height is less than 5 counts. Economy is not however of great importance, since the code in its present configuration processes a 1024-channel spectrum with an element list of 25 (of which 10 were actually detected) in 15 s.

Various means of experimentally determining the tailing parameters have been described in the literature [3–5]; unfortunately all of them are rather laborious. A new weighting scheme, described in Section 6.4 below, allows the user to employ approximate tailing descriptions, thereby reducing this labour.

The linear energy calibration which relates the Gaussian centroids j_0 to X-ray energies E_x has been changed to a quadratic relationship to cater for very slight nonlinearities which have been reported on occasion [6]. Thus the channel-energy relationship is

$$j_0 = P_a + P_b E_x + P_c E_x^2. \tag{5}$$

The relation between peak width and energy is unchanged viz.

$$s^2 = P_d + P_e E_x. \tag{6}$$

On conclusion of each fit, this expression is used to evaluate the detector resolution, defined as the full-width at half-maximum ($2.35s$) at 5.9 keV X-ray energy. This value is compared to the expected value stored in the particular file that holds the parameters describing the detector in use; if it exceeds that expected value, prominent warnings are presented both on the computer screen and on the printed results. We have found this useful on the rare occasion when a problem such as a ground loop causes a temporary worsening of detector resolution.

Many of the characteristic X-ray lines of a given element occur in closely-spaced pairs, $K\alpha_1$ and $K\alpha_2$ being the prime example. In the earlier code we used criteria laid out by Van Espen et al. [7] to determine when such a doublet should be represented by two distinct peaks and when it would be acceptable to approximate by using a single peak. Having observed slight improvements in

quality of fit by separating previously merged doublets, we have altered these criteria. A K X-ray doublet splitting of 2 eV is now sufficient to warrant two distinct peaks. Thus the $K\alpha$ pair is treated as two distinct peaks for atomic numbers above $Z = 17$ (previously 28) and the $K\beta_{1,3}$ pair is treated as two distinct peaks above $Z = 29$ (previously 45). In the case of L X-rays, members of a doublet are merged if they are separated by less than 5 eV. These changes result in a slightly larger X-ray energy and intensity database.

The relative intensities of the X-ray lines of each element are taken from the database and adjusted to reflect matrix effects (including secondary fluorescence), transmission through absorbers, and the intrinsic efficiency of the detector. "Funny filters", i.e. absorbers with a central circular portion removed, have been added to the repertoire of absorbers. Each filter is specified by quoting its atomic number (compounds such as air and Mylar are assigned "atomic numbers" above $Z = 100$), thickness and the hole fraction. Up to five filters, one of which can be "funny", may be specified.

In addition to Si(Li) detectors, the code now caters also for Ge detectors. For this purpose a parameterization of measured Ge K escape peak intensity (R) relative to that of the parent line was developed by Prozesky [8]; this is

$$R = c_1 \left\{ 1 - c_2 \exp\left[c_3 \log\left(1 + \frac{1}{c_2 E^{c_3}} \right) \right] \right\}, \tag{7}$$

with $c_1 = 24.8962$, $c_2 = 1.1903$, $c_3 = 1.6815$ for $K\alpha$, and only c_1 differing for $K\beta$ ($c_1 = 4.0083$). This parameterization is also incorporated into the expression for the intrinsic efficiency of the Ge detector.

Pile-up peaks were represented in the original code by the "pile-up element" model of Johansson [9] and the calculation was limited to double pile-up. Because many of the geochemical and materials science specimens that we encounter show triple pile-up, the calculation has been extended to describe this. The code also determines the pile-up resolving time T by comparing intensities of the strongest pile-up peaks with those of their parents. This value is necessary to calculate a pile-up corrected peak area which is then used to compute the element concentration. If the parent peaks are not contained in the region of fit, or if the operator fails to request inclusion of pile-up peaks, recourse is had to a default T value, which is stored in the detector description file. The pile-up correction factor was changed from the approximation of $1/(1 - 2rT)$ to $\exp(2rT)$ for a paralyzable system; this is better than the previous approximation when $rT > 0.1$. Finally, the fitted peak areas provided by GUPIX94 do not contain pile-up contributions; rather these contributions are combined with the parent peak intensities just prior to generation of the concentrations.

Escape peaks are still represented by Gaussians. Our experience is that in some Si(Li) detectors the escape

peaks have a more pronounced tailing than do photopeaks at the same energy. There is a need to investigate this tailing and to incorporate it into the escape peak description in GUPIX.

4. Data base

The data base has not changed from that used in 1989 as regards proton stopping powers and ionization cross sections. The K and L X-ray energies and relative intensities have only changed in the sense that doublets such as K α are now regarded as two peaks rather than one over a much broader range of atomic number than was the case before.

M X-rays can now be dealt with. ECPSSR ionization cross sections based on DHS wavefunctions for the five M subshells were taken from the tables of Chen and Crasemann [10] for 10 elements in the range $70 < Z < 92$. The proton energy (E) range covered in these tables is 0.06–2 MeV. Dr. Chen [11] very kindly provided us with further cross section values covering the proton energy range 2.25–3.25 MeV. At each of the 18 proton energies in the range 0.2 to 3.25 MeV, the five subshell cross sections were fitted to a $\ln(\sigma)$ versus Z relationship using a fourth order polynomial. This provided cross sections for all 23 elements. In a second step, the $\sigma(E)$ values for each element (and each subshell i) were fitted to the so-called universal expression

$$\ln\left(\sigma_i U_i^2\right) = \sum_{j=0}^{5} b_{ij}\left[\ln\left(\frac{E}{\lambda U_i}\right)\right]^j,\tag{8}$$

where u_i is the binding energy for the subshell in question [12] and $\lambda = 1836.1514$. The resulting coefficients b_i were then tabulated and incorporated in the GUPIX94 database. (This approach to parameterization of ECPSSR cross sections was taken by us before [1,13] for the K shell and the L subshells.) In the present case, there were discrepancies of several percent between the fitted M cross section values and those of the initial tables. These arise because Eq. (8) is not completely adequate over the entire proton energy range. The solution was to repeat the process using two different proton energy ranges, viz. 1–3 MeV and 0.1–1 MeV; this brought the maximum discrepancy down to 2% and the typical discrepancy down to under 0.5%. For conventional PIXE analysis with 2–3 MeV incident protons, a combination of the two sets is used. For low-energy PIXE with incident proton energies below 1 MeV, the second set is used. Chen et al. [14,15] provide theoretical DHS fluorescence yields and Coster–Kronig probabilities for the five M subshells for ten elements in the range $67 < Z < 95$. The fluorescence yields were fitted to the expression

$$\left(\frac{\omega}{1 - \omega}\right)^{1/4} = \sum_{i=0}^{4} a_i Z^i\tag{9}$$

to provide values across the entire range $70 < Z < 92$. The required values of the Coster–Kronig probabilities f_{ij} were interpolated manually. Chen and Crasemann [16] have tabulated M subshell X-ray transition rates in the Dirac–Fock approximation for ten elements in the range $48 < Z < 92$. Third or fourth order polynomial fits were used to determine relative X-ray line intensities at all Z values. Finally, photo-electric cross sections σ_{PE} are required for calculating M-shell secondary fluorescence. Dirac–Slater values were obtained for all Z values and a limited set of X-ray energies from the tables of Scofield [17], and fitted to the expression

$$\log \sigma_{PE} = \sum_{i=0}^{3} a_i (\log E)^i\tag{10}$$

to generate a table of coefficients a_i. Overall, this procedure for the M shell is essentially the one that we used earlier [1] for the L-shell case.

Executable code has been provided to allow the user to change the attenuation coefficient data base. The latter is essentially an array of all mass attenuation coefficients for the characteristic X-ray lines of all elements in all absorbing elements. These values, taken from the NIST code XCOM [18], appear to be more reliable [19,20] than those of various semi-empirical schemes; a pre-1989 version of GUPIX had employed the semi-empirical scheme of Leroux and Thinh [21] but a study [20] of such schemes revealed inaccuracies in this and various other approaches. A shortcoming of XCOM (in common with the semi-empirical schemes) is that it makes no attempt to deal with the oscillations in μ/ρ that occur within 2 keV of an absorption edge. For example, in the case of the nickel Kα_1 line absorbed in iron, XCOM gives $\mu/\rho = 360.9$ cm^{2}g^{-1}, while extremely careful synchrotron radiation measurements give 381.8 cm^{2}g^{-1} at this particular energy. We reported elsewhere [22] that this particular discrepancy was just sufficient to explain a systematic error in the concentrations of the major constituents in the NIST FeNi alloy reference material SRM 1159 measured by PIXE. In this case, where the Ni K X-rays are highly absorbed and where there is significant secondary fluorescence of iron, it is particularly important to have accurate μ/ρ values. Doubtless there are analogous cases. With the new utility GUKLUP, the user may replace XCOM attenuation coefficients with preferred experimental values, and the program keeps a record of such alterations.

Finally, it may be useful to comment upon the L subshell Coster–Kronig probabilities, which play a major role in shifting L vacancies from the subshell in which ionization takes place to a higher subshell, of which the emitted L X-ray is then characteristic. We continue to use the DHS theoretical values for these quantities. However, in the case of the L_2–L_3 Coster–Kronig probability f_{23}, experimental values have converged and are systematically 5–10% below theory [23]. Papp et al. [24] have presented a detailed study of the experimental technique used to mea-

sure f_{23}. Just as in PIXE, these Si(Li) spectroscopy experiments ignore spectator vacancy satellites and natural line widths, and Papp et al. show that this neglect is probably responsible for much of the discrepancy. There are thus two choices as regards improving GUPIX. One is to use the set of experimental f_{23} data. The other is to use the theoretical values, and to improve the spectrum description by incorporating satellites and natural linewidths in the peak descriptions. We are investigating these issues, but for the moment we have retained the theoretical values in the data base.

5. Specimen type: standardization

5.1. General

The original GUPIX code [1] dealt with thick targets only, because at that point in time other codes existed for thin target PIXE analysis. Our own work on PIXE analysis of fine air particulate deposits has necessitated inclusion of a thin target option, where the proton exits from the specimen with energy loss so small that matrix effects may be neglected. Our work on thin chalcogenide films and metal coatings, which cannot be classified as either thick or thin, has necessitated an "intermediate target" option, including the possibility of determining the specimen thickness as well as its composition; this option has recently been extended to deal with multiple layers. GUPIX94 recognizes the thin and thick target options by the initial setting of exit proton energy ($E_f = E_0$ or $E_f = 0$), the intermediate thickness option by a request to the user to specify the thickness (or exit proton energy) or to request thickness iteration, and the multi-layer option by input of the number of layers.

In each case, standardization is effected through the instrumental constant approach discussed in Ref. [1] and elsewhere [22,25]. For a given target (specimen, standard or reference material) the measured intensity of proton-induced X-rays in the principal X-ray line of element Z is

$$I(Z) = HQc_z \epsilon_z^i t_z I_1(Z, M),$$ (11)

where $I_1(Z, M)$ is the computed theoretical intensity per unit concentration per steradian per unit beam charge for the particular matrix, C_Z is the concentration of Z, Q is the measured beam charge (or equivalent), ϵ_z^i is the detector's intrinsic efficiency and t_z is the X-ray transmission through any absorbers present. The instrumental constant H subsumes both the detector solid angle and any calibration factor intrinsic to whatever device is used to integrate beam current.

5.2. Thick specimens

In the case of a specimen thick enough to stop the proton beam, I_1 is given by

$$I_1(Z) = \frac{N_{av} \omega_z b_z}{A_z} \int_{E_0}^0 \frac{\sigma_z(E) T_z(E) \, dE}{SP(E)},$$ (12)

where ω_z is the relevant fluorescence yield, b_z is the fraction of the series X-ray intensity in the line of interest, N_{av} is Avogadro's number, SP is the proton stopping power, and

$$T_z(E) = \exp\left\{ \frac{-\left(\frac{\mu}{\rho}\right) \frac{\cos \alpha}{\sin \theta_{TO}} \int_{E_0}^0 \frac{dE}{SP(E)}}{SP(E)} \right\}.$$ (13)

In this expression, which describes attenuation of X-rays, μ/ρ is the mass attenuation coefficient, α is the angle of the proton beam with respect to the specimen normal, and θ is the take-off angle of X-rays relative to the specimen surface. There are further contributions to the characteristic X-ray intensities from secondary fluorescence [26]. These are calculated without approximation by GUPIX, and subtracted from the fitted peak intensities before conversion of the latter to concentrations via Eqs. (11)–(13). Tertiary fluorescence contributions, although both calculated and investigated [26], have not yet been incorporated.

We have described elsewhere [22,27] the use of the NIST molybdenum steel reference material SRM 1155 to determine simultaneously both H and the thickness of an aluminum X-ray absorber in the context of micro-PIXE analysis of mineral grains. The Mo K X-ray intensity is insensitive to the precise absorber thickness, and provides the H value; with this H value, the Fe K X-ray intensity then provides the absorber thickness. A modest iteration is needed to provide complete consistency between the two measurements. This approach is a variant of a method first described by Ryan et al. [28] and based upon a similar alloy. Any suitable standard or reference material will provide the value of H.

GUPIX94 provides separate options for: (a) determining trace element concentrations in a known matrix; and (b) determining the major elements that comprise the matrix.

(a) Trace element determination: For convenience, the program will accept the a priori matrix concentrations in various formats. These include wt. %, atom number fraction, and oxide weight percent; the last option is useful in mineralogical work, where the major elements are often pre-determined by electron probe micro-analysis as oxide percentages. GUPIX stores the default valence for each member of the element table ($1 < Z < 92$). In some cases, however, more than one oxide is possible; the most common example is iron, with the two oxide forms FeO and Fe_2O_3. To cope with this, a second element table $101 < Z < 192$ is available; the user can define a different valence for an element from this table, and may then draw two oxides, one from each table, into the matrix element list. Finally, a separate list of the elements (most often, but not exclusively, trace elements), whose X-rays are to be in-

cluded in the model PIXE spectrum is provided by the user. Eq. (11) is then solved directly.

(b) Matrix element determination: When it is the major element concentrations that are to be determined, no a priori matrix data are called for, and the user simply provides a list of the elements observed in the spectrum; these may include major, minor and trace elements (see example below). The program first performs a thin-target solution (see Section 5.3), ignoring matrix corrections, to generate initial estimates of the percentage concentrations. The $I_i(Z, M)$ values are computed from these estimates and Eq. (11) is then solved to provide a new concentration vector. This process is iterated until the concentrations are consistent; the concentrations should of course sum to 100%.

The element list may include, at the discretion of the user, elements whose X-rays are at too low an energy to be present in the spectrum. Such "invisible" elements are categorized as "dependent" or "independent" with respect to the "visible" elements. Oxygen, for example, can be introduced as an independent invisible element; its concentration will be determined in each iteration loop as

the difference between 100% and the sum of the visible element concentrations, and this concentration value is used in computing the matrix effects in the subsequent loop [25]. As an example, Fig. 2 shows a spectrum and analysis for the high-T_c superconductor $SrLaGaO_4$; the resulting atom fractions are in good agreement with the stoichiometrically expected values. In this approach to invisible elements, only one independent invisible element may be introduced. Alternatively, the amount of oxygen may be made dependent upon one or more (or all) of the visible elements by attaching it stoichiometrically to these; this follows the practice in electron probe microanalysis of minerals as referred to earlier. In this case the restriction that concentrations must sum to 100% is not present.

Similarly, invisible complexes (e.g. SO_4, SiO_2) may be defined. Again, there may be one independent complex, with the proviso that it sums with the visible elements to 100% concentration. Alternatively, one or more dependent invisible complexes may be defined, using the appropriate stoichiometry to link these to particular elements in the fit list. If valence other than the default valence stored in the database is desired, then the element concerned and its
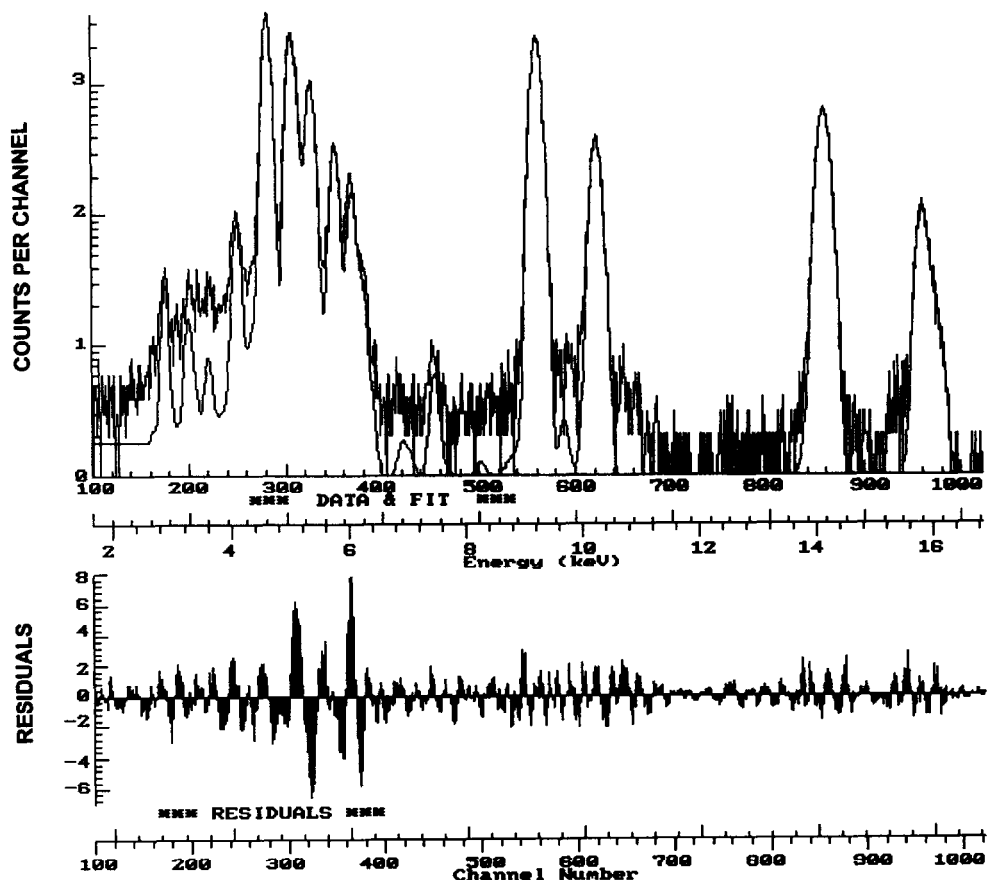

Fig. 2. PIXE spectrum and analysis of superconductor $SrLaGaO_4$. The mean concentrations given by GUPIX for five spots (stoichiometric values in brackets) are: Sr 24.61% (24.32%); La 38.71% (38.56%); Ga 19.32% (19.35%); O 17.36% (17.77%).
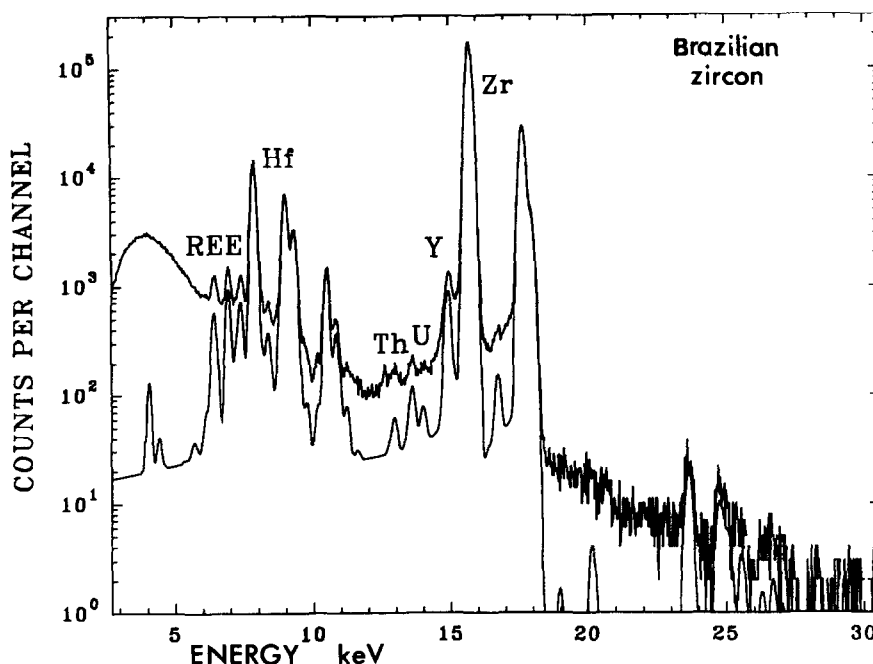
Fig. 3. PIXE spectrum of a zircon; concentrations from GUPIX are Zr 51.1%; Hf 0.76%; Y 0.22%; Fe 0.20%; SiO_4 47.72%; Th 362 ppm; U 582 ppm; sum of rare earths 986 ppm.

valence are specified by the user in the auxiliary element table $101 < Z < 192$ that was mentioned earlier. As an example, the PIXE spectrum of zircon, $ZrSiO_4$, in Fig. 3 contains Zr as the major element, Hf as a minor element, and Sc, Y, Th, U and various light rare earths as trace elements; the stoichiometry tells us that (SiO_4) is also present, although not visible in the spectrum. The spectrum was fitted using an element list that comprised all these elements together with the SiO_4 complex, summing to 100% concentration.

Finally, the option of designating certain elements observed in a thick-target spectrum as residing on the surface (as opposed to being homogeneously distributed in the

bulk) has been retained. This can be useful in situations where a thick target has an exterior coating of some kind.

5.3. Thin specimens

Here one works in terms of areal density and the unit of I_1 is X-ray intensity per steradian per unit collected charge per mg/cm^2 of element. MicroMatter standards [30] are commonly used to provide so-called sensitivity curves relating X-ray yield to areal density [13], and these relationships were employed in the earlier GUPIX. The current version uses the H value method instead, in the interests of consistency. I_1 is given by

$$I_1(Z) = \frac{\sigma_z(E_0)\,\omega_z b_z N_{av}}{A_z}.$$ (14)

Thin specimens, by definition, exhibit no matrix effects, and the question of secondary fluorescence does not arise. H values are provided by MicroMatter standards or by appropriate thin film reference materials such as the NIST SRM 1832 and SRM 1833. A typical set of measurements of H from our laboratory is in Fig. 4; this shows good consistency between the MicroMatter data and the results from NIST thin glass standards.

5.4. Specimens of intermediate thickness

If the specimen thickness is known, e.g. via proton scattering measurements [13], or if the corresponding exit
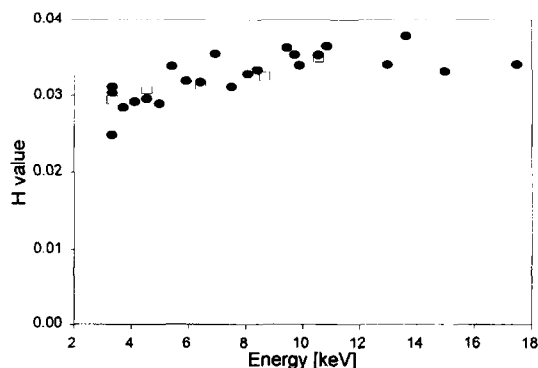

Fig. 4. Instrumental constant H measured with MicroMatter standards (filled circles) and NIST 1833 thin glass film standard.

proton energy E_f is measured, then the integration in Eq. (12) can be carried out between the limits E_0 and E_f, and the problem solved essentially as for a thick specimen. All the options described in Section 5.2 are available in precisely the same manner.

If the specimen thickness is not known, GUPIX offers an iterative thickness determination; this is subject to the proviso that there is no independent invisible element, since otherwise there would be too many variables. In the first step of this solution, the code temporarily assumes the specimen to be a thin target. Matrix corrections and secondary fluorescence are ignored, the spectrum is fitted, and the resulting peak intensities are converted to areal densities via Eq. (14). These areal densities are summed to provide an initial estimate of the specimen thickness. Next, a double iteration procedure begins, with an outer loop involving concentrations and an inner loop involving spectrum fitting. Using the first estimate of concentrations to provide matrix and secondary fluorescence effects, the spectrum fitting loop is then run to provide the best fit between measured and model spectra. The resulting new concentrations and thickness are compared with the previous values, and if convergence has not occurred, the fitting procedure is run again. This continues until the thickness and the concentrations are consistent.

5.5. Accuracy of standardization method

In the case of thin targets, it is obvious that the accuracy of the instrumental constant approach to standardization rests on the accuracy of the relationship between $\sigma(E_0)$ and Z. In the case of thick and intermediate targets, the bulk of X-ray production occurs as the proton energy falls from E_0 to about $0.9E_0$, and so the accuracy depends again directly upon the accuracy of the σ versus Z relationship at energies very close to the incident energy, which is typically 2.5–3 MeV. It does not depend strongly upon the proton energy-dependence of the cross section. We have demonstrated accuracies of a few percent both in micro-PIXE trace element analysis of homogeneous well-characterized geological reference materials [27], and in PIXE analysis of NIST standard reference alloys [22,26].

In both the thin and the thick specimen cases, measurements of H across a broad range of atomic number using a set of single-element standards results in an H value that is essentially independent of Z as is expected [22]. However, we have observed slight differences between the H values determined using K and L X-rays respectively; presumably this reflects residual inadequacies in the L X-ray data base. And in addition, at X-ray energies below 5 keV and above 25 keV, we have noted slight departures (see Fig. 4) from constancy, which we must attribute in the main to imperfect characterization of the Si(Li) detector; the cause at low X-ray energy is probably inaccurate knowledge of window thickness, Si dead layer thickness and ice layer thickness; the cause at high energy is proba-

bly inaccuracy in crystal thickness. GUPIX94 deals with the first of these two problems by allowing the user to define three separate instrumental constants H for K, L and M X-rays. It deals with the second by providing the option of defining a vector of H values with a corresponding vector of K, L and M X-ray energies; this manoeuvre effectively absorbs the inadequacies in the data base or the detector characterization into an $H(E_x)$ that has a slight dependence upon X-ray energy E_x.

5.6. Multiple-layer specimens

Multi-layered specimens are dealt with as an extension of the intermediate thickness case. The individual layer thicknesses may be specified or they may be unknown, in which case an extension of the double iteration procedure described in Section 5.4 is invoked in order to determine both the thickness and the major element composition of each layer. Several constraints are necessary if a multiple layer problem is to be solved. The elements comprising any layer must be unique to that layer. In any given layer of known thickness, there may be either dependent or independent invisible elements or complexes. But in a layer whose thickness is to be determined by iteration, there can be no independent invisible elements or complexes; this reflects the fact that the 100% concentration sum is used to constrain the thickness determination. There is one modest exception to this case: if the layer in question is the last one and if it is infinitely thick, then an independent invisible element can be included in it. These restrictions do not apply to dependent invisible elements or complexes, which are associated via stoichiometry with particular visible elements. Finally, there is an option to define the last layer in any given case as being infinitely thick in order to represent substrates.

The multi-layer calculations take full account of secondary fluorescence within each layer and between each pair of layers. In the event that the total fitted intensity of peaks of visible elements in a layer does not exceed 1 count, execution terminates and the user is asked to delete that layer from the model before proceeding.

In this case, the accuracy is a much more complex issue, as we can illustrate using the simple example of a film on a substrate, with $E_0 = 3$ MeV and the proton energy at the interface being 1 MeV. If H has been determined using 3 MeV protons, then we now require in addition of the data base that it describes accurately the energy-dependence of ionization cross sections from 3 MeV down to low proton energies. Some effort needs to be expended along the lines of the early thick-target PIXE studies [31,32] to provide quantitative assessments of the accuracy of both thickness and concentrations measured by the multi-layer PIXE method. This would be done in terms of the accuracy of the various components of the data base such as cross sections, stopping powers and attenuation coefficients. We have performed initial tests of the accu-

racy of thickness measurement using the NIST thin film standard SRM 1132A (Ni film on Fe substrate). The measured value of 9.36 ± 0.1 μm agreed well with the recommended value of 9.4 ± 0.4 μm. However, more extensive testing is desirable and we have this in progress.

6. The fitting procedure

6.1. Nonlinear parameters

As indicated earlier, there are now five nonlinear parameters A_1–A_5 describing the energy and peak-width calibrations, and these are determined by the nonlinear fitting process along with the height of the principal line of each element in the model spectrum.

6.2. Termination of the fit and elimination of weak peaks

It was noticed that when the earlier GUPIX was fitting spectra that contained one or two very strong peaks, the intensities of peaks that were weak relative to the local background could exhibit a dependence upon initial parameters of the fit. In addition, very weak peaks, despite being visible in the spectrum, were occasionally being assigned zero intensity by the fitting procedure. Previously, we had constrained the rate at which the fitting code could reduce the intensity of a principal line, in order to avoid elimination of weak peaks at an early stage where the overall fit quality is poor; peaks were not permitted to fall ultimately below 0.1 counts in height. When a minimum chi-squared was reached, peaks of height 0.2 to 1.0 were eliminated, and the nonlinear process was re-started. This process was repeated until no more elements were dropped. This last step has now been replaced by a constrained linear least-squares fit of the height parameters; at the outset of this particular computation, the height parameters are set at the values determined by the nonlinear procedure, and the nonlinear calibration parameters are held fixed at the values given by the nonlinear fit. Each height parameter is varied individually (one at a time). If, as a result of this linear step, any heights are determined to be negative, they are set to zero for that loop. The heights are iterated until no height parameter is changing by more than the relative stopping criterion parameter STEP (user-defined). This change has resulted in slight improvements in the overall quality of fits to spectra.

6.3. Background removal by top-hat filter

The continuum background is removed by convoluting the spectrum with a digital top-hat filter, as shown in Fig. 1. The optimum filter dimensions UW and LW have been discussed by Schamber [2], McCarthy and Schamber [33] and Statham [34]. A compromise has to be found among issues of concern such as statistical accuracy (and hence

limit of detection) for weak peaks, sensitivity to curvature of the continuum and sensitivity to errors in the peak model. The best compromise among these factors appears to be integers which are close to satisfying UW = fwhm and LW = 0.5fwhm. In GUPIX, the value of fwhm is taken as that at the centre of the spectrum. In the initial GUPIX, UW and LW were calculated from the initial estimates of the parameters A_i. A change has been made so that as the calibration parameters A_1–A_5 vary during the fitting iterations, the values of UW and LW are adjusted if necessary and the fit restarted. For reasons to be given in Section 6.6, the user is also provided with the option of choosing different values of UW and LW than the recommended ones.

Although this filter concept is a simple one, the overall fitting problem becomes rather more complex insofar as the chi-squared comparison is now effected between the filtered spectra (measured and model) rather than the directly measured and modelled spectra. This has ramifications for various issues, which are dealt with in the following sections.

6.4. Modified weighting schemes

6.4.1. Smoothing of statistical weights

Various authors [35–37] have remarked on the tendency of fitting programs to under-estimate the intensities of peaks that are only a few counts high and are superimposed upon weak backgrounds. This general effect arises because of the large fluctuations from one channel to the next in the weights, which are usually taken as the inverse of the channel counts. (It is to be distinguished from the particular effect described in Section 6.2: that effect was specific to our particular fitting procedure, and has been dealt with.) One solution [36] is to adopt the fit function rather than the measured data to provide the weights, but this method is not available to us, because GUPIX does not employ a full fit function that comprises both peaks and background. Another solution [35] is to derive the weights from a running three-point average, which provides a smoothing effect. The use in GUPIX of the top-hat filter should provide such a smoothing effect, but this issue was not investigated in our earlier paper on GUPIX [1].

To investigate this matter, 100 sequential PIXE spectra were recorded containing a weak Pd Kα peak of area 15–20 counts; this follows an approach first used in PIXE fitting studies by Ryan et al. [38]. The sum spectrum was fitted using a conventional nonlinear least-squares approach employing a spectrum model that comprised Gaussian peaks and a polynomial background; the result, 1752 ± 42 counts, provides us with a best estimate of the mean intensity in the 100 spectra viz. 17.52 ± 0.42 counts. The suite of 100 spectra was then fitted in various ways, and the results are summarized in Table 1 in terms of the mean peak area, its standard deviation, and the standard error of the mean.

Table 1
Mean fitted intensity and standard deviation for Pd Kα in a set of 100 spectra

Fit details		Pd Kα	
Background	Weights	Mean	SD
Polynomial	Normal	15.89	3.02
	Equal = 1	17.64	3.61
	Offset 10	17.22	3.48
Filter 5–10–5		17.19	6.93
10–10–10		17.15	5.31
20–10–20		17.89	4.13
40–10–40		17.56	3.57

In the conventional least-squares fit where the spectrum model included the peak together with a polynomial background, the mean fell low at 15.89, even though a three-point smoothing [33] was employed in the weight estimation. When all the weights were simply set equal to unity, an acceptable value of 17.64 resulted. Another simple solution is to add an intensity offset of 10 counts to every channel in the spectrum; this gave a result of 17.22. These two observations tend to confirm that the shortfall in the conventional fit is indeed due to the weight problem. The issue to be addressed here is the performance of GUPIX, and specifically the effect of the tophat filter treatment on the weight problem. With the filter dimensions matched to the width of the Pd peak, an acceptable result of 17.19 was generated.

GUPIX therefore continues to rely on statistical weights, with, of course, the modification made necessary by the integrating effect of the top-hat filter. In addition, new options have been introduced to allow the user to incorporate systematic errors in the overall weighting scheme. These are described in the next two sections.

6.4.2. Systematic errors in the spectrum model

Benjamin et al. [39] pointed out that the rather thick filters often necessary to suppress major element peaks in thick-target PIXE analysis are subject to errors in both their thickness (x) and mass attenuation coefficient (μ/ρ). Even small errors in these quantities can cause large errors in the resulting computed transmission exp $(-\mu x)$, and hence local misfits caused by mismatch of computed and observed Kβ/Kα intensity ratios for major elements whose X-rays are heavily absorbed. The earlier GUPIX therefore contained an optional provision to define a fractional uncertainty in the quantity μx. The resulting error in the transmission was added to the statistical errors prior to calculating the weights w_j. Because in geological and environmental specimens this effect was significant only for light major element peaks that had the additional problem of large tailing contributions, it had the added benefit of somewhat suppressing the influence of these

peaks on the chi-squared, and thereby improving the quality of fit overall.

In GUPIX94 we have rewritten the code to provide a more comprehensive treatment of possible errors in the intensity ratios of minor lines (e.g. Kβ) to the principal line (e.g. Kα_1). The user can specify an uncertainty in these database line intensity ratios. This is added to the uncertainty in the ratios arising from the specified uncertainty in μx(absorber), thus providing an overall fractional error F in, for example, the Kβ/Kα ratio. It is assumed that the model function for Kα is correct; F then provides the fractional error in the model Kβ count at any channel. This error is combined with the statistical counting error to provide the weight.

A third new systematic error has been introduced to deal with the specific issue of inadequately known low-energy tailing on intense peaks, and, if used, this error is added at each channel to the two described above, prior to calculation of the modified weights. This particular scheme is described in the next section.

If these various systematic errors are employed to modify the weights, then of course the chi-squared value is altered. In order to assist the user in assessing the systematic effects, a second chi-squared value based on the statistical weights alone is calculated (without redoing the fit) and reported in the final output. In addition, an output file comparing the statistical weights and the modified weights is available.

6.4.3. De-emphasis of intense low energy tails and artifacts

A third weighting issue, which we have recently addressed in detail elsewhere [40], needs to be mentioned here. Many specimen types, especially in the geochemical context, give rise to spectra with intense major element peaks (e.g. Fe) in the low-energy region of the spectrum and weak trace element peaks (e.g. Pd, Ag, Ba) in the high energy region. The intensities of the low-energy tails of the Kα and Kβ matrix peaks can be of the same order as the intensities of the entire trace element peaks, and are therefore as important as the latter in determining the fit between the measured and the model spectra. These tails comprise charge collection tailing and radiative Auger contributions; neither process is perfectly understood and the description of these two contributions is usually an approximate one; in addition the first component can vary with time and with counting rate. The resulting slight misfits are negligible for the major elements themselves, but can cause a small error in the parameter A_2. This results in a misalignment of the weak high-energy peaks in the model spectrum and therefore an under-estimate of their intensities. Because of this, more work is needed on low-energy tailing phenomena. In the interim, we believe that the potential for mis-fits can be minimized by use of an augmented weighting scheme which has been introduced as an option in GUPIX94.

In this approach [40], the user retains the previous possibility to define (as in Section 3 above) the low energy tailing component arising from incomplete charge collection. And the data base continues to provide a very approximate description of KMM radiative Auger satellites as single Gaussian peaks. The new option allows the user to define "tail error terms" for either or both of the ICC and the KMM features. These error terms are exponential features, defined in precisely the same way as exponential ICC tails and incorporated in the detector description file; the user specifies the tail height as a fraction of the corresponding Gaussian height and the tail slope as a fraction of the Gaussian width s. In the ICC case, the error term is constructed to provide the maximum sensible estimate of the ICC tail intensity across the relevant set of X-ray energies. In the KMM case, the error term is constructed to provide the maximum possible estimate of the entire KMM feature. The tail error contributions are now added to the statistical error in the counts at each channel, providing an augmented error value for the intensity at each channel. The weights are then calculated in the normal way, as the inverse squares of the combined error values.

In practice, this approach causes significant change in the weighting only in the tailing regions just below the most intense matrix element peaks. The effect of these regions upon the fit is de-emphasized. We have demonstrated that use of this weighting scheme eliminates the error in A_2 referred to above, and thus rectifies the under-estimate of the intensities of high-Z trace element lines. The tail error term parameters for a specified detector are input by the user into the detector description file using a text editor. This is done in precisely the same manner as for the actual tail model parameters, and so the extra complexity is not great.

6.5. Error estimates

As we pointed out before, one disadvantage of the top-hat filter treatment of continuum background is that the issue of errors in the fit parameters becomes more complex. Both Schamber [2] and Statham [34] derived an expression which could provide these errors, but at that juncture its evaluation was prohibitive in terms of computer time and memory storage. Schamber therefore presented the following simple approximation for a multiplicative correction c to the variance conventionally derived from the diagonal term of the error matrix (cross-terms were ignored):

$$c = \frac{2UW \times LW}{UW + 2LW}. \tag{15}$$

We used this in the earlier GUPIX. In tests conducted by Statham on a simple two-peak spectrum with a range of peak separations, the approximation gave an estimate of the variance within the range of $+67\%$ to -40% of the

Table 2
Estimates of peak intensity error in fits to various spectra

Specimen	Element	Peak area	Error estimate [%]			
			A [a]	B [b]	C [c]	D [d]
Basalt glass	26	8.2×10^4	0.53	0.56	0.57	0.3
	28	770	2.15	4.8	5.0	5.0
	29	1850	1.05	2.4	2.8	2.6
	30	2040	1.03	2.3	2.7	2.5
	31	470	3.7	8.2	9.0	8.7
	37	160	9.9	22.2	24.1	17.4
	38	9050	0.5	1.1	1.2	0.9
	39	485	3.3	7.4	7.7	6.0
	40	2940	1.1	2.6	2.6	2.7
	41	280	4.8	10.7	10.9	8.0
Bornite	29	3.3×10^5	0.18	0.20	0.21	0.1
	47	3950	1.0	2.1	1.9	1.6
	83	1450	1.1	2.1	2.5	4.5
SRM 1155	26	5.7×10^4	0.36	0.42	0.46	0.35
(Steel)	28	4.4×10^4	0.15	0.3	0.5	0.4
	29	1450	1.7	3.7	4.1	5.5
	42	3.7×10^4	0.3	0.6	0.5	0.4

[a] A: From diagonal element of error matrix.
[b] B: A corrected by approx. of Eq. (15).
[c] C: Full computation as per Ref. [34].
[d] D: Simple statistical estimate.

rigorously computed value. Developments in computer speed and memory are such that it is now feasible to incorporate the full calculation in GUPIX to generate the important peak height errors. These errors are then combined with the (much smaller) peak width errors, as given by the nonlinear least-squares fit including Schamber's simple correction factor, and thus estimates of the peak area errors are generated. In practice the peak height error dominates over the peak width error. GUPIX94 quotes these peak area errors as the "fit errors", and in addition it continues to provide the empirical "statistical errors", which we described previously [1] and which were based upon an idea of Clayton [41].

Previously, the fit errors were on occasion non-physical (which was the reason for providing in addition the alternative empirical formulation). Our experience now suggests that the fit errors are reasonable, except when correlated errors in the two peak width parameters lead to overly large width error estimates. Such correlated errors can arise when the region of fit spans only a small range of X-ray energies; in such situations (which are rare), the width error is probably more reliable if the Fano factor parameter A_5 is fixed and only A_4 is allowed to vary.

Table 2 summarizes the peak area errors generated in fits to various representative spectra; for simplicity, statistical weighting was used. The differences between Schamber's early approximation and the detailed treatment of Statham are generally small. In addition, the Statham values agree quite well with the empirical statistical errors
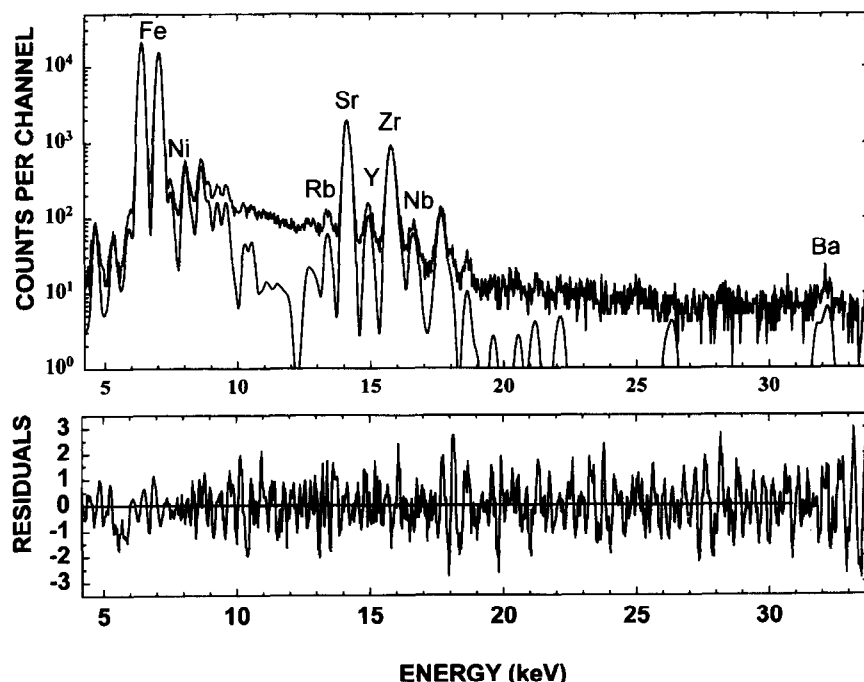
Fig. 5. PIXE spectrum of BHVO-1 basalt, showing isolated weak Ba K X-ray peak discussed in text.

calculated from peak and background intensities. There is only one divergence from these general observations. For very intense peaks, where the error is very small, the detailed error calculation gives an uncertainty that ranges up to twice as large as the empirical statistical error estimate and is clearly erroneous; hence for very intense peaks of major elements (concentration several percent by mass) the statistical error is to be preferred.

6.6. Scatter of replicate data

Whether Statham's exact error calculation or Schamber's approximate expression is used, the intensity errors

and the ensuing concentration errors are larger when background is treated by a top-hat filter than when it is represented by a mathematical model. Detection limits are therefore correspondingly larger, and there is more scatter in replicate data.

If we refer back to the example of the 100 replicate weak Pd peaks, Table 1 shows that the standard deviation in the case of the optimum filter (6.93) is twice that generated when a polynomial background is used (3.5–3.6) with a smoothed weighting scheme. Here we have a disadvantage of the top-hat filter: it worsens the precision relative to the true capability of the PIXE technique. The effect can be alleviated by increasing the extent of the

Table 3
Concentration means [a] and standard deviations for 5 elements in BHVO-1 as a function of filter dimensions

Element	Mean ± SD				Rec. [b]	Prev. [c]
	4–8–4	8–16–8	8–8–8	16–8–16		
Ba	113 ± 50	119 ± 17	116 ± 34	118 ± 17	139	115
Sr	403 ± 3	405 ± 3	404 ± 3	404 ± 3	403	404
Nb	18.7 ± 1.8	17.8 ± 0.8	18.2 ± 1.0	17.7 ± 0.8	19	19
Y	24.5 ± 1.8	25.4 ± 1.4	24.8 ± 1.5	25.1 ± 1.3	27.6	25
Zr	174 ± 3	175 ± 2	174 ± 3	175 ± 2	179	173

[a] Means and standard deviations are from 10 measurements at 10 spots on the specimen.
[b] Rec. = recommended values.
[c] Prev. = values measured in earlier PIXE work [27].

wings of the filter. In our example replacement of the (5–10–5) filter by a (20–10–20) filter returns the standard deviation to within striking distance of the optimum value.

This is further investigated by a re-analysis of some of the data on barium in the USGS standard basalt BHVO-1 which we presented elsewhere [27]. In these spectra (see Fig. 5), the Ba K lines are weak and isolated, and our earlier Ba concentrations showed a large standard deviation. The filter dimensions used in fitting were 4–8–4, corresponding to the FWHM at the centre of the spectrum. However, the most appropriate filter for the Ba lines, which are at the upper energy limit of the spectrum, would be about twice as wide. The set of ten replicate spectra was therefore fitted using a filter of twice the normally pre-scribed width. As Table 3 shows, the scatter of the Ba intensities was reduced by a factor of 3. Following the Pd example above, the same reduction could be achieved by quadrupling the outer lobes of the filter. While a large departure from the prescribed filter dimensions is not difficult to rationalize for a weak, isolated peak such as Ba, it would not be advocated for a cluster of peaks displaying a wide range of intensity such as the group in the centre of the spectrum comprising relatively intense Sr and Zr and relatively weak Nb and Y. Nonetheless it is interesting to examine their intensities and scatter also. In this case, the standard deviations again decrease although the result is less dramatic, and the mean concentrations are only modestly affected. For each filter used, the correspon-dence with "reported values" for this standard is rather

similar to our earlier results [27] which were computed prior to introduction of the linear least-squares termination described in Section 6.2 above.

Thus, in regions of isolated peaks or low background curvature, increasing the filter's overall dimensions or just its outer lobes can provide significant improvement in precision and limit of detection. It would be difficult to build an automatic facility to deal with every situation. At this time it is simply recommended that in cases such as the Ba example discussed, users perform a second fit with an expanded filter.

6.7. Weak peaks in vicinity of intense neighbours

A converse situation to that of the isolated Ba peak discussed in the preceding section is that of a weak peak situated very close to an intense neighbour. The question again arises as to the accuracy provided by the filter background technique. A topical example is that of trace nickel in garnets, where the weak Ni line close to the intense Kβ line of iron, which typically constitutes several wt.%. The Ni concentration can be used via the geother-mometer of Griffin et al. [42] to estimate the formation temperature of the garnet, and along with other mineralogi-cal data, this is a potentially useful tool in diamond exploration.

Analyses were done at 15 different spots on a single homogeneous garnet grain, using 2.5 μC of charge at each spot. A typical spectrum is shown in Fig. 6. The 15 spectra
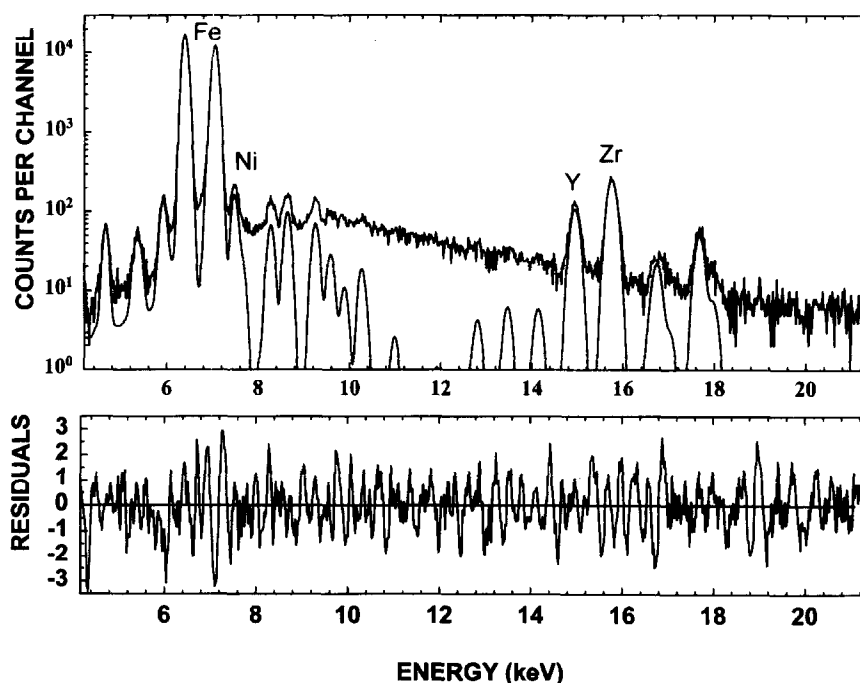

Fig. 6. PIXE spectrum of a garnet containing approximately 120 ppm nickel.

Table 4
Effect on top-hat filter dimensions on intensity of a weak Ni peak

UW–LW–UW	Ni [ppm]	
	$Q = 2.5 \, \mu C$	$Q = 37.5 \, \mu C$
5–9–5	117.7	120.7
4–8–4	119.0	120.1
3–6–3	118.7	122.0
4–7–4	119.2	120.3
3–5–3	118.6	123.2

were fitted using the 5–9–5 filter dimensions dictated by the central energy of the spectrum, dimensions that differ somewhat from those optimum for nickel, which are 4–7–4 or 4–8–4. The spectra were therefore fitted with a range of filters as indicated in Table 4. The results given for the 2.5 μC charge are means over the 15 spectra, and the statistical error for the Ni result in any one of these is 4.5%. The final column gives the result of fitting the sum spectrum, with its improved statistics. The nickel concentration provided by GUPIX turns out to be rather robust over a significant range of filter dimensions.

7. Summary

Modifications and improvements to the GUPIX package for PIXE analysis have been presented. All of these are incorporated in the current version of the program. A variety of practical examples has been used to illustrate various aspects of the program, with particular emphasis paid to the influence of our background treatment upon analytical accuracy and precision. Development is continuing; at the time of proof-reading, natural (Lorentzian) widths of X-ray lines have been incorporated in the peak description.

Acknowledgements

References

[1] J.A. Maxwell, J.L. Campbell and W.J. Teesdale. Nucl. Instr. and Meth. B 43 (1989) 218.

[2] F.H. Schamber, in: X-ray fluorescence analysis of environmental samples, ed. T.G. Dzubay (Ann Arbor Science, Michigan, 1977).

[3] J.L. Campbell, Nucl. Instr. and Meth. B 50 (1990) 189.

[4] N.P.-O. Larsson, U.A.S. Tapper and B.G. Martinsson, Nucl. Instr. and Meth. B 43 (1989) 574.

[5] J. Bohsung, Diplomarbeit, Universität Heidelberg, 1993.

[6] W. Maenhaut and J. Vandenhaute, Bull. Soc. Chim. Belg. 95 (1986) 407.

[7] P. Van Espen, H. Nullens and F. Adams, X-ray Spectrom. 16 (1980) 126.

[8] V.M. Prozesky, private communication.

[9] G.I. Johansson, X-Ray Spectrom. 11 (1982) 194.

[10] M.H. Chen and B. Crasemann, At. Data Nucl. Data Tables 41 (1989) 257.

[11] M.H. Chen, private communication.

[12] C.M. Lederer and V.S. Shirley, Table of Isotopes, 7th ed. (Wiley, New York, 1978).

[13] S.A.E. Johansson and J.L. Campbell, PIXE: a novel technique for elemental analysis (Wiley, Chichester, 1988).

[14] M.H. Chen, B. Crasemann and H. Mark, Phys. Rev. A 21 (1980) 449.

[15] M.H. Chen, B. Crasemann and H. Mark, Phys. Rev. A 27 (1983) 2989.

[16] M.H. Chen and B. Crasemann, Phys. Rev. A 30 (1984) 170.

[17] J.H. Scofield, Lawrence Livermore National Laboratory Report UCRL-51326 (1973).

[18] M.J. Berger and J.H. Hubbell, National Bureau of Standards Report NBSIR 87-3597.

[19] E.B. Saloman, J.H. Hubbell and J.H. Scofield, At. Data Nucl. Data Tables 38 (1988) 1.

[20] J.L. Campbell and D. Deforge, X-Ray Spectrom. 18 (1989) 235.

[21] J. Leroux and T.P. Thinh, Revised Tables of Mass Attenuation Coefficients (Corporation Scientifique Claisse, Quebec, 1977).

[22] J.L. Campbell, D. Higuchi, J.A. Maxwell and W.J. Teesdale, Nucl. Instr. and Meth. B 77 (1993) 95.

[23] W. Jitschin, Proc. 15th Int. Conf. on X-Ray and Inner-Shell Ionization (X90), Conf. Proc. No. 125, eds. T.A. Carlson, M.O. Krause and S.T. Manson (AIP, New York, 1990) p. 408.

[24] T. Papp, J.L. Campbell and S. Raman, Phys. Rev. A 49 (1994) 729.

[25] J.L. Campbell, W.J. Teesdale and J.-X. Wang, Nucl. Instr. and Meth. B 50 (1989) 189.

[26] J.L. Campbell, J.-X. Wang, J.A. Maxwell and W.J. Teesdale, Nucl. Instr. and Meth. B 43 (1989) 539.

[27] G.K. Czamanske, T.W. Sisson, J.L. Campbell and W.J. Teesdale, Amer. Mineral. 78 (1993) 893.

[28] C.G. Ryan, D.R. Cousens, S.H. Sie, W.L. Griffin, G.F. Suter and E. Clayton, Nucl. Instr. and Meth. B 47 (1990) 55.

[29] J.L. Campbell and W.J. Teesdale, Nucl. Instr. and Meth. B 74 (1993) 503.

[30] MicroMatter Co., Deer Harbour, Washington, USA.

[31] J.L. Campbell, J.A. Cookson and H. Paul, Nucl. Instr. and Meth. 212 (1983) 427.

[32] J.L. Campbell and J.A. Cookson, Nucl. Instr. and Meth. B 3 (1984) 185.

[33] J.J. McCarthy and F.H. Schamber, NBS Special Publication 604 (1981).

[34] P.J. Statham, Anal. Chem. 49 (1977) 2149.

[35] G.W. Phillips and K.W. Marlow, IEEE Trans. Nucl. Sci. NS-24 (1977) 154.

[36] T. Awaya, Nucl. Instr. and Meth. 165 (1979) 317.

[37] E. Clayton and C.G. Ryan, Nucl. Instr. and Meth. B 49 (1990) 161.

[38] C.G. Ryan, E. Clayton, W.L. Griffin, S.H. Sie and D.R. Cousens, Nucl. Instr. and Meth. B 49 (1990) 115.

[39] T.M. Benjamin, C.J. Duffy and P.S.Z. Rogers, Nucl. Instr. and Meth. B 30 (1988) 454.

[40] J.A. Maxwell, W.J. Teesdale and J.L. Campbell, Nucl. Instr. and Meth. B 94 (1994) 172.

[41] E. Clayton, PIXAN: The Lucas heights PIXE Analysis Computer Package, AAEC Report M113, Australian Atomic Energy Commission (1986).

[42] W.L. Griffin, D.R. Cousens, C.G. Ryan, S.H. Sie and G.F. Suter, Contr. Mineral. Petrol. 103 (1989) 199.