Amptek Application Note XRF-1:
XRF Spectra and Spectra Analysis Software
By R.Redus, Chief Scientist, Amptek Inc, ©2008.

X-Ray Fluorescence (XRF) is a very simple analytical technique: X-rays excite atoms in a sample, which emit X-rays at energies characteristic of each element. A detector measures the energies and intensities of the emitted X-rays. From this, one deduces which elements are present and their concentrations.

Setting up a system for first order results is straightforward. One can easily identify the major constituents in simple samples and obtain approximate concentrations. This is sufficient for some applications. However, second order effects are quite important for many applications. To obtain high accuracy, identify constituents with low concentrations, or operate on a wide variety of samples, second order effects must be addressed. To remove these second order effects and obtain high accuracy, properly configuring both the hardware and software is critical.

Understanding the physical mechanisms responsible for these effects is helpful when one is configuring a system to obtain accurate results. The first section of this application note provides an overview of the physical mechanisms causing these effects. The second section illustrates how Amptek's XRS-FP software processes the spectrum to obtain the final result, the composition of the sample. As a specific example, we will use data obtained from a stainless steel alloy (310), using equipment available from Amptek, Inc.
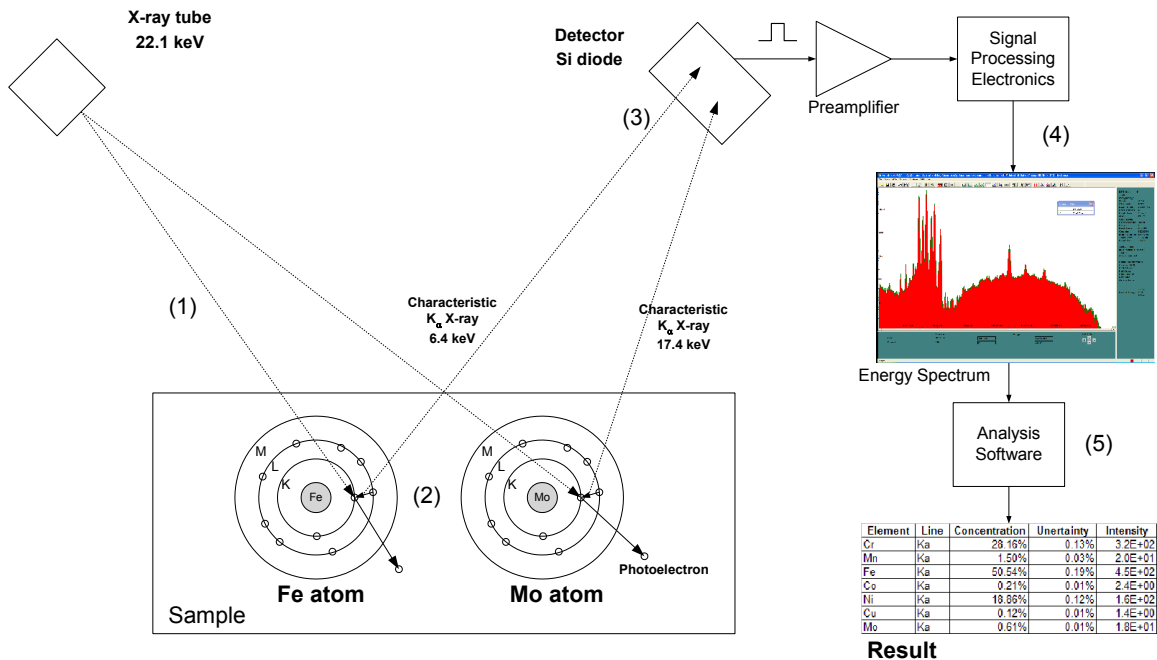
HOW DOES AN XRF ANALYZER WORK?

Figure 1. Sketch showing simple picture of ED-XRF system.

Figure 1 is a sketch of the "simple picture" of an ED-XRF system. The basic operation is:

1. Photons from the source irradiate the sample. Shown here is an X-ray tube with Ag anode.
2. Atoms in the sample are excited, then de-excite by emission of X-rays at energies characteristic of the elements in the sample. The intensities are related to the concentration of each element in the sample and to the excitation probabilities.
3. The characteristic X-rays interact in the detector. Each X-ray releases a pulse of charge, proportional to the energy deposited. The signal processing electronics measures the total charge (i.e. total energy) of each pulse, that is of each individual X-ray.
4. The signal processing electronics sum the results of many interactions into a histogram, showing the number of events in each of many energy bins. This is the energy spectrum.
5. The peaks in the spectrum correspond to the elements in the sample. The number of X-rays in each peak is proportional to the number of atoms. Analysis software detects which peaks are present, finds the intensity of each peak, and computes the concentration of each element in the sample.

Amptek Inc.

The spectra in Figure 2 illustrate this. The plot on the left shows spectra which were measured from samples of 99.9% pure copper, zinc, and lead. Each element emits a unique spectrum of X-rays characteristic of that element. Cu has a strong peak at 8.05 keV, while Zn has a strong peak at 8.64 keV. Pb has two strong peaks at 10.55 and 12.61 keV.[A] The plot on the right shows these pure element spectra along with the spectrum measured from a brass sample. Peaks are visible for all three elements. The Cu and Zn peaks are strongest, with nearly equal intensity, so this brass alloy is nearly 50:50 Cu and Zn. The Pb peaks are very weak, indicating that there is Pb in this alloy, but only a small amount.
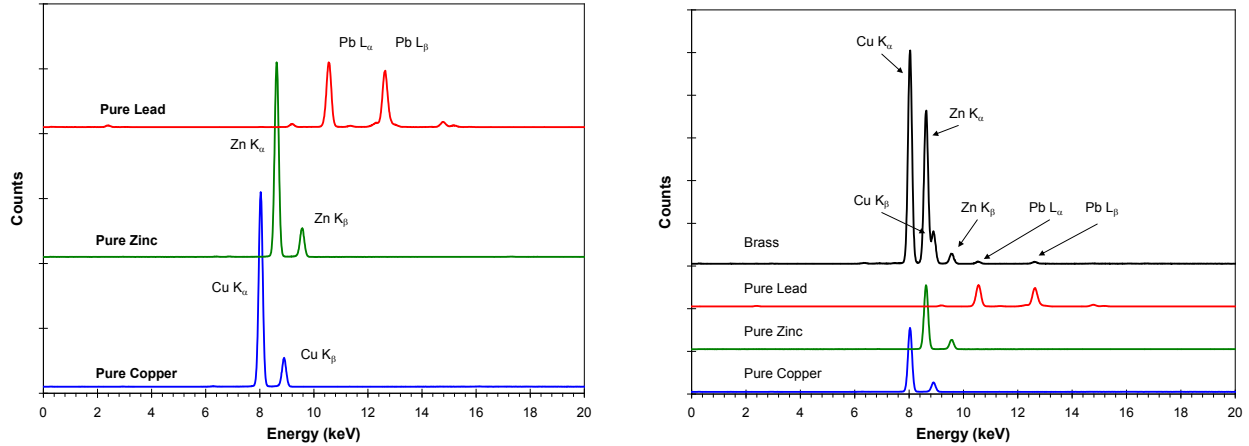

Figure 2. Left: X-ray spectra measured from samples of pure copper (blue), pure zinc (green), and pure lead (red). Right: Pure element spectra are shown on the bottom, along with a spectrum obtained from brass, in black. The brass spectrum clearly shows the characteristic X-rays of the individual elements.

Figure 3 shows a simple XRF instrumentation system. This was used to acquire the data shown here. Key components include (1) Mini-X X-ray tube (with control hardware and software), (2) the XR100-SDD X-ray detector and preamplifier, (3) the PX4 signal processor, (4) a personal computer running software for data acquisition and control and also running the XRS-FP analysis software, and (5) mechanical components, not shown in the photograph, which hold the sample in place and which shield the scattered X-rays to prevent radiation exposure. All components can be purchased from Amptek, Inc. at www.amptek.com.
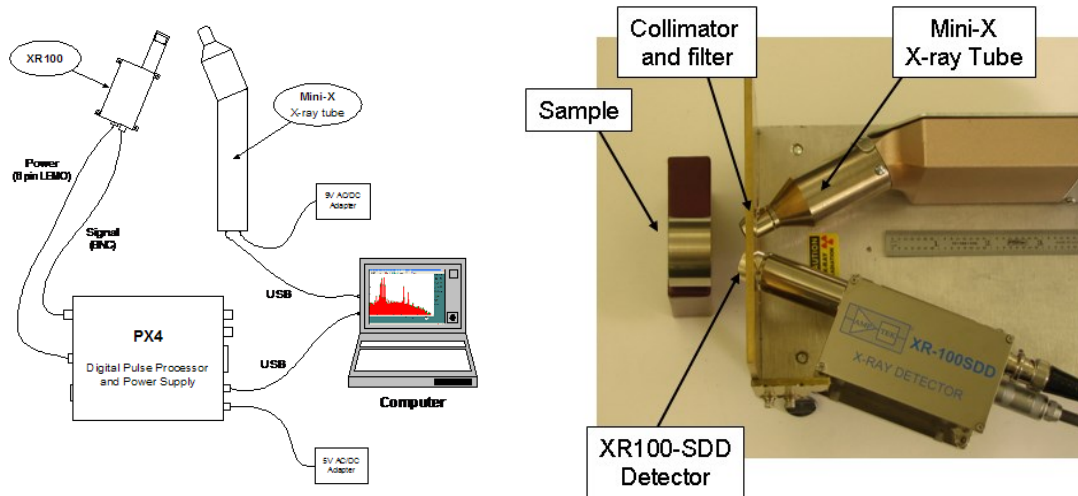

Figure 3. Block diagram (left) and photograph (right) of the setup used to take the data shown in this application note.

[A] The X-ray energies arise from the energy levels in the atoms. More specifically, the energies of a characteristic X-rays equals the difference in binding energies between atomic shells. When the primary X-ray strikes a Cu atom, it ejects an electron from the (inner) K-shell. This vacancy will be filled by an electron falling from an outer shell, and the atom emits an X-ray with energy equal to the difference in binding energies. Each element has a unique set of energy levels, hence a unique spectrum of characteristic X-rays. Because Zn has a higher atomic number than Cu, its inner shell has a higher binding energy, so it emits higher energy characteristic X-rays.

ANALYSIS OF STAINLESS STEEL 310

We now consider in more detail a specific example, the measurement of stainless steel 310. In Figure 2, we considered a qualitative analysis: which elements are present? We now want to consider a quantitative analysis: what is the concentration of each element? Table 1 shows the nominal properties of the sample used. This includes the concentration, measured by "wet chemistry", and the characteristic X-rays produced by these elements.

Element	Conc. (%)	K_α (keV)	K_α (keV)
Mo	0.4	17.48	19.61
Cu	0.1	8.05	8.90
Ni	19.3	7.48	8.26
Fe	53.2	6.40	7.06
Mn	1.7	5.90	6.49
Cr	25.3	5.41	5.95

Table 1. Nominal properties of sample.

Ideal Spectrum for Stainless Steel 310

Figure 4 shows an "idealized" spectrum for this sample, corresponding to the simple picture of Figure 1. Each element is assumed to release X-rays at precisely the characteristic energies, with intensity proportional to the nominal concentration. The spectrum consists of a series of infinitesimally thin lines.
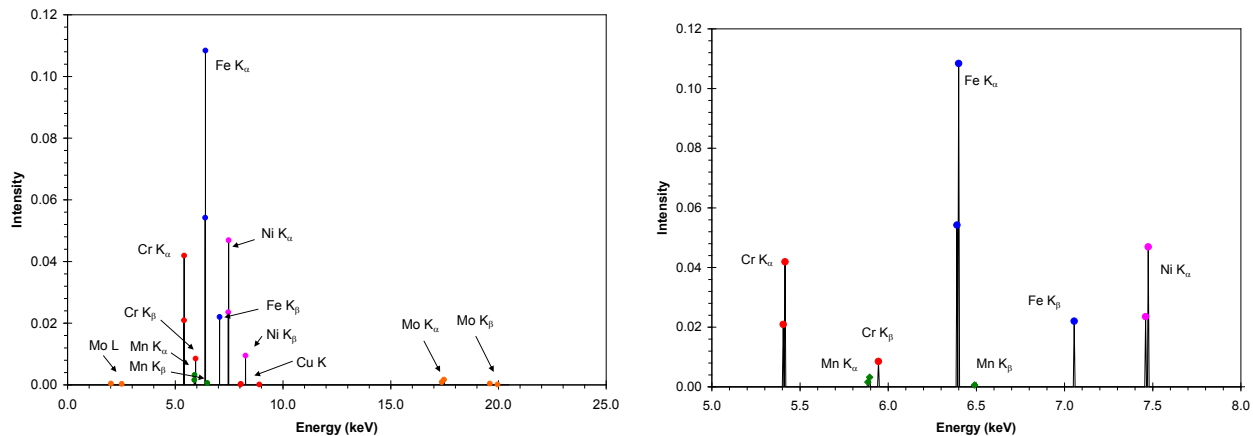

Figure 4. Idealized energy spectrum expected from stainless steel 310.

How close is this ideal to reality?

Figure 5 shows the measured energy spectrum. To a first glance, it is quite close to the idealized version. One can easily identify the major peaks (Fe, Cr, Ni, Mo) and which elements represent the most intense emission lines.
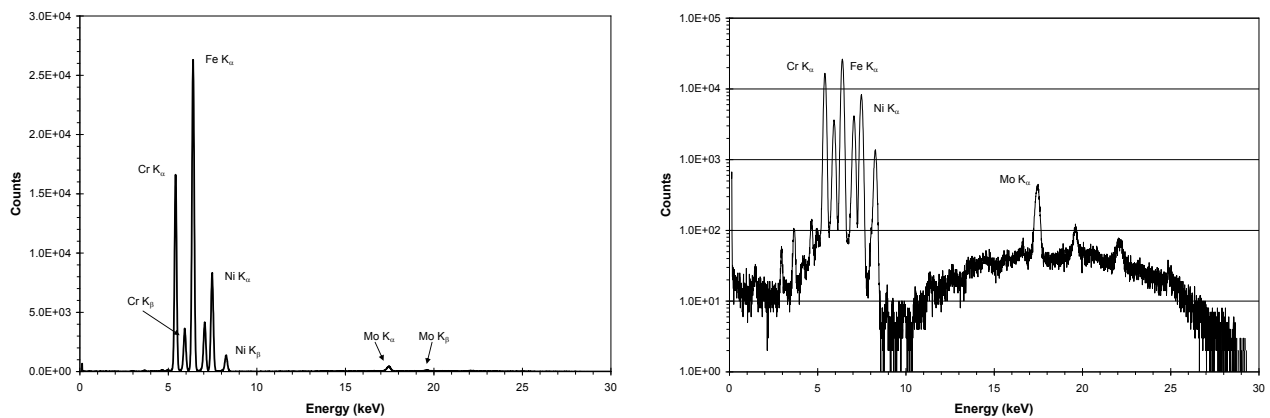

Figure 5. Measured energy spectrum from a sample of stainless steel 310. The plots show the same data, but with a linear (log) scale for the counts on the left (right).

Looking more carefully, there are several important differences between Figure 4 and Figure 5. First, the peaks have a finite width, the energy resolution. When they are close enough, they overlap so cannot be resolved. The Mn in Figure 4 (right) is a small bump adjacent to Cr and Fe and will be hard to detect.

Second, the logarithmic plot shows that there are more peaks than one would expect from the constituents of steel. This includes several peaks below 5 keV and two peaks over 20 keV. There is also a broad continuum of X-rays on which the main peaks are superimposed. Finding the correct area for peaks which overlap or sit on background can be a challenge. Third, although not obvious, the ratio of the areas of the peaks does not match their concentrations. In stainless steel 310, there is twice as much Fe as Cr, but the Fe peak is only 30% more intense.

What complicates the simple picture?

Matrix effects

The most fundamental complication is that the characteristic X-rays emitted by the elements in the sample interact with other atoms in the sample. This is sketched in Figure 6 (left): Fe emits 6.4 keV K_α X-rays, just above the Cr K edge. The Fe K_α X-rays are absorbed very efficiently by Cr, which then emits additional characteristic X-rays. The Fe line is *absorbed* while the Cr line is *enhanced*. Further, the primary X-rays from the excitation source, the X-ray tube in this case, are absorbed by the sample. The composition and density of the sample determine how far the primary X-rays penetrate into the sample, hence the number of atoms excited. Absorption of the primary radiation and absorption and enhancement of characteristic lines are termed *matrix effects*. The main consequence is that, even if a "perfect" X-ray source, detector, and electronics were used, analysis software would still be needed to correct the measured intensities for these matrix effects to obtain the true concentrations.
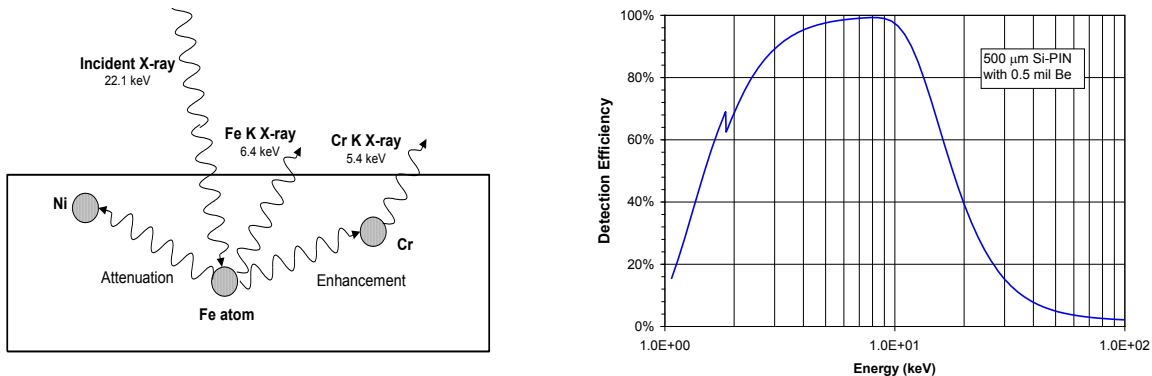

Figure 6. Left: Sketch illustrating matrix effects, where X-rays emitted by one element are absorbed by other atoms in the sample, which in turn have enhanced emissions. Right: Plot showing detection efficiency versus energy.

Attenuation of the X-rays

Some of the X-rays emitted by the sample towards the detector will not be detected. A fraction of the higher energy X-rays pass through the detector without interacting. A fraction of lower energy X-rays interact between the sample and the detector, in the detector's Be window or in air along the path. The primary consequence is that the sensitivity of the system is reduced at lower and higher energies. Figure 6 (right) shows the efficiency versus energy for a 500 μm thick Si-PIN detector with a ½ mil Be window. To obtain accurate relative intensities, software must correct for the losses.

Finite energy resolution

The detector has a finite energy resolution, turning the "delta function" response in Figure 4 into a broader peak (approximately Gaussian). The broadening largely arises from two physical effects: statistical fluctuations in charge production (the *Fano factor*) and *electronic noise* at the preamplifier's input. The Fano factor is a physical limit of the detector material. Electronic noise is greatly affected by the details of the detector but also by the settings used in the signal processing electronics. There are additional, higher order terms which affect the photopeak shape including the intrinsic line width of the characteristic X-rays, charge trapping near the surface of the detector, and charge trapping in the bulk of the detector. Many of these higher order contributions are non-Gaussian.

The primary consequence of finite energy resolution is *spectral interference*, i.e., overlapping peaks. In Figure 7 (left), the Fe K_β peak and the Ni K_α peak interfere somewhat. Both can be detected, but to determine the net area of each, this overlap must be addressed. The software cannot simply obtain net

counts in some region but must *deconvolve* the spectrum, fitting it to a sum of separate photopeaks. The Mn peaks exhibit a severe overlap: The Mn K_α is almost completely overlapped by the Cr K_β (which has twice the area), while the Mn K_β is lost under the Fe K_α.
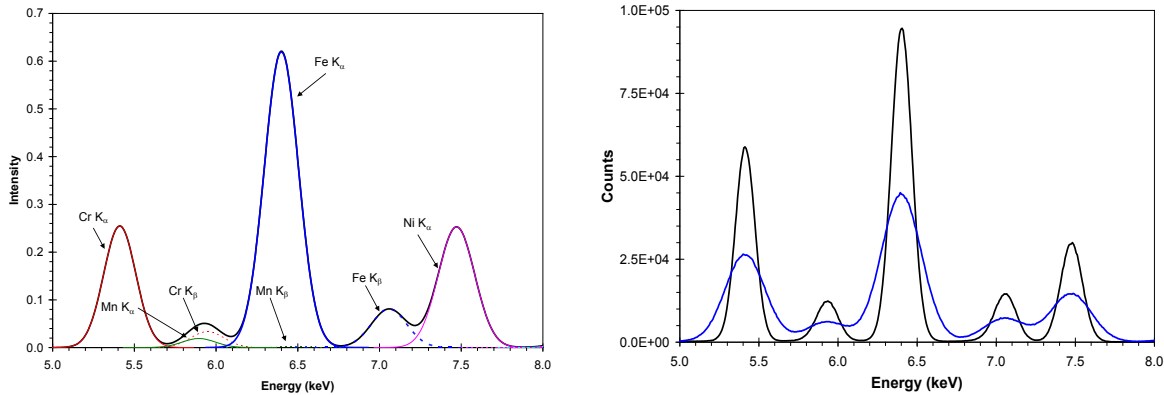

Figure 7. Effect of finite energy resolution. The pot on the left shows a computed spectrum, assuming Gaussian photopeaks due to Fano broadening and noise. The spectral interference is clear. The plot on the right shows spectra measured from stainless 310, using the same detector but configured for high (black) and low resolution (blue)

Additional X-ray Interactions in the Detector

   The main photopeaks are due to X-rays which undergo photoelectric absorption in the sample, emit a characteristic X-ray which reaches the detector, and undergo complete photoelectric absorption. When the X-ray undergoes photoelectric absorption by a Si atom, it emits a photoelectron and a Si characteristic X-ray. For the main photopeaks, both secondary particles stop in the detector so the full energy of the incident X-ray is detected. However, some of the X-rays will scatter, only depositing a portion of the incident energy. There are several different scattering processes, leading to various features in the spectrum. Figure 8 sketches the physical mechanism while Figure 9 illustrates the corresponding effects on the spectrum.

   Escape peaks: When the X-rays interact in the detector, they produce Si characteristics X-rays (1.75 keV). Some of these escape the detector. The deposited energy is reduced, e.g. a 6.4 keV X-ray (Fe K_α) deposits only 4.65 keV. Every feature in the spectrum will have an associated escape feature at 1.75 keV lower energy.

   Photoelectron escape continuum: The Si characteristic X-ray may stop in the detector, but the photoelectron may escape the active volume. The electron loses energy continuously along its path, so the energy lost depends on the path length. Escaping electrons form a continuous distribution with the minimum energy deposited at the escape peak.
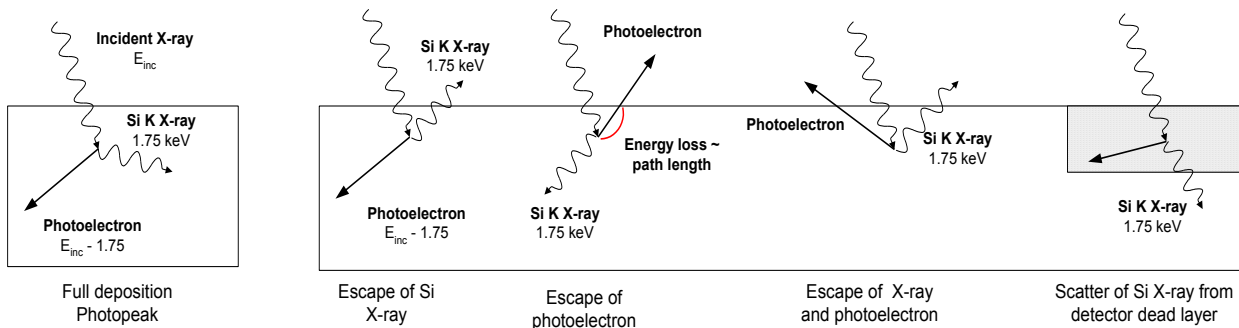

Figure 8. Sketches illustrating various X-ray interactions responsible for spectral features. The left-most diagram shows full energy deposition, responsible for the photopeak. The other sketches illustrate the escape of a Si characteristic X-ray (removing 1.75 keV), escape of the photoelectron (energy lost depends on the path length), both photoelectron and X-ray escaping, and a photon interacting in the dead layer of the detector, where the Si X-ray reaches the active volume.

   Higher order effects: There are many additional interactions possibilities. For example, if both the X- and the photoelectron escape, even less energy will be deposited. This forms a continuum below the escape peak. As another example, there is always some "dead" region in a detector, e.g. a dead layer

at the surface and dead regions near the outer boundary. X-rays interacting here produce secondaries which can reach the active volume. A Si characteristic X-ray can scatter into the active volume, producing a 1.75 keV artifact or a photoelectron can scatter from the dead layer, producing a continuum from zero up to 1.75 keV.
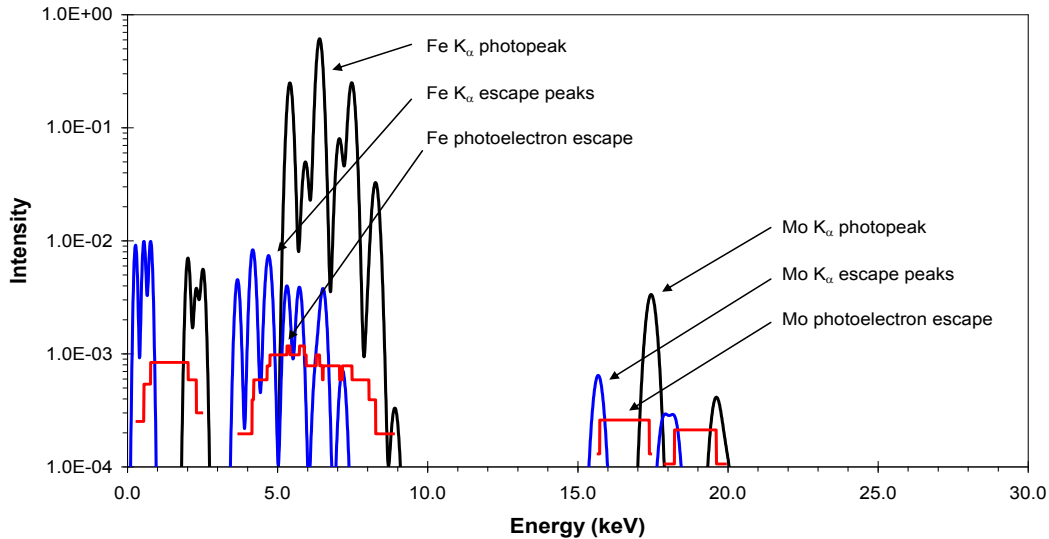

Figure 9. Sketch of the photopeaks, escape peaks, and the continuum due to escaping photoelectrons. These arise from the processes sketched in Figure 8.

Bremsstrahlung, Scattering, and Environmental Interference

Figure 1 showed the tube emitting only Ag characteristic X-rays from the anode, showed these only interacting in the sample by producing characteristic X-rays, and showed only the sample's characteristic X-rays reaching the detector. In reality, an X-ray tube will also emit Bremsstrahlung X-rays over a wide energy range. Some X-rays from the tube will undergo Compton (inelastic) scatter and some will undergo Rayleigh (elastic) scatter from the sample into the detector. This will form a continuum with a peak from the scattered Ag. A filter is usually placed in front of the tube to reduce backscatter at the energies of primary interest.
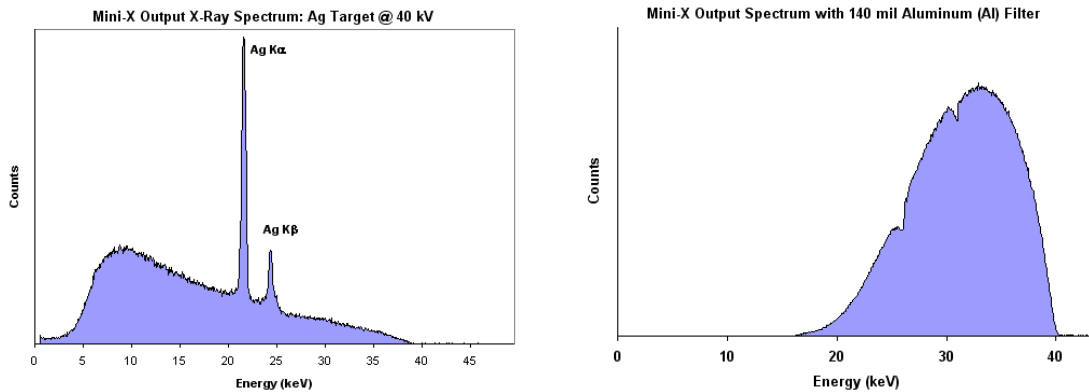

Figure 10. Spectra measured from X-ray tubes. Left: Spectrum measured using an Ag anode, without a filter, at 40 kVp. Right: Spectrum measured using a W anode, with a 140 mil Al filter, also at 40 kVp.

There are always various materials in the vicinity of the detector. X-rays will interact with these, producing *environmental interferences*, characteristic X-ray peaks and a scattering continuum. For example, Ar in the air produces a peak at 2.96 keV and Al in the detector package produces a peak at 1.48 keV. Collimators and shielding must be placed around the tube and detector, producing characteristic X-rays along with additional scattering. Figure 11 sketches these effects. The black trace illustrates the main photopeaks, with a Gaussian shape, along with detector escape features. The blue trace illustrates the backscattered tube spectrum, with a filtered Bremsstrahlung continuum and a scattered Ag peak. Also shown, in red, is a continuum formed from Compton backscattering out of the detector. In green are shown typical environmental interferences. Mechanical design of the system is critical to minimizing these effects.
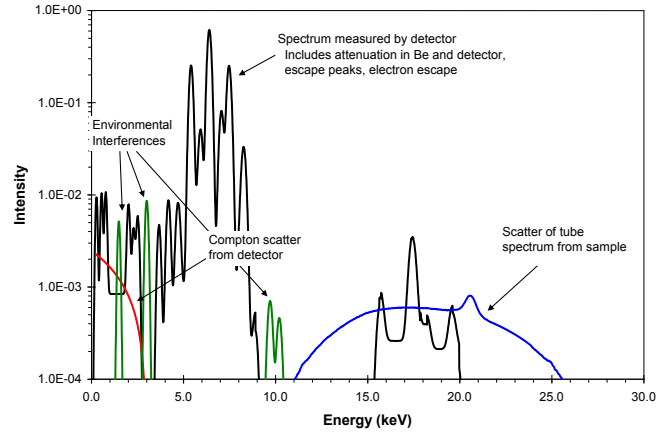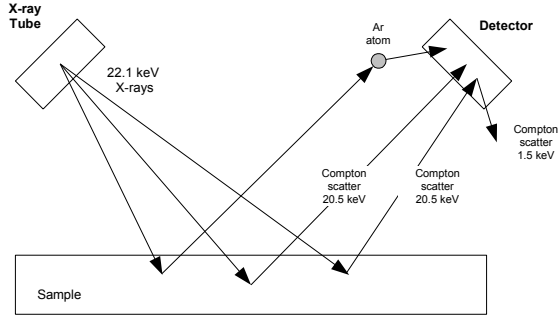
Figure 11. Sketch showing additional interactions and their effect on the spectrum.

Random Timing Effects

X-rays are produced and interact in the detector at random times, which has several important consequences. First, there will always be statistical fluctuations in the number of events, governed by Poisson statistics. If N counts are in a peak, or a channel, the standard deviation is √N. In the case of the Mn and Cr overlapping peaks, the uncertainty in the Cr peak may be larger than the total Mn counts! Second, there is a finite processing time for the electronics. When one event occurs, there is some minimum *dead time* interval during which a second event cannot be accurately measured. The random timing implies that some pulses will not be recorded due to *dead time*, so the output count rate is less than the input count rate. This must be corrected when absolute concentrations are needed. Third, if two pulses occur sufficiently close in time, then the electronics record them as a single event. Two 6.4 keV X-rays interacting within a nanosecond are indistinguishable from a 12.8 keV X-ray. They form a *sum peak*, with energy equal to the sum of the energies of the incident photons.
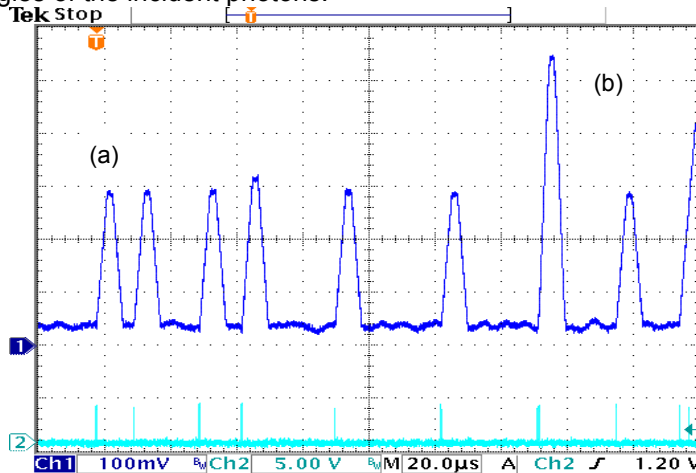

Figure 12. Oscilloscope trace of pulses measured for a 5.9 keV monoenergetic source. The light blue trace shows ICR, a logic indicator that a pulse was detected, while the dark blue trace shows the signal, the shaped pulses. The peak amplitude indicates the energy of the X-ray. Pulse (a) is typical, while pulse (b) shows two pulses which are piled up. The logic pulses are distinguishable, so the spectrometer will reject these two, so these are part of the dead time loss. They must be closer in time for the system to record an event adding to sum peak.
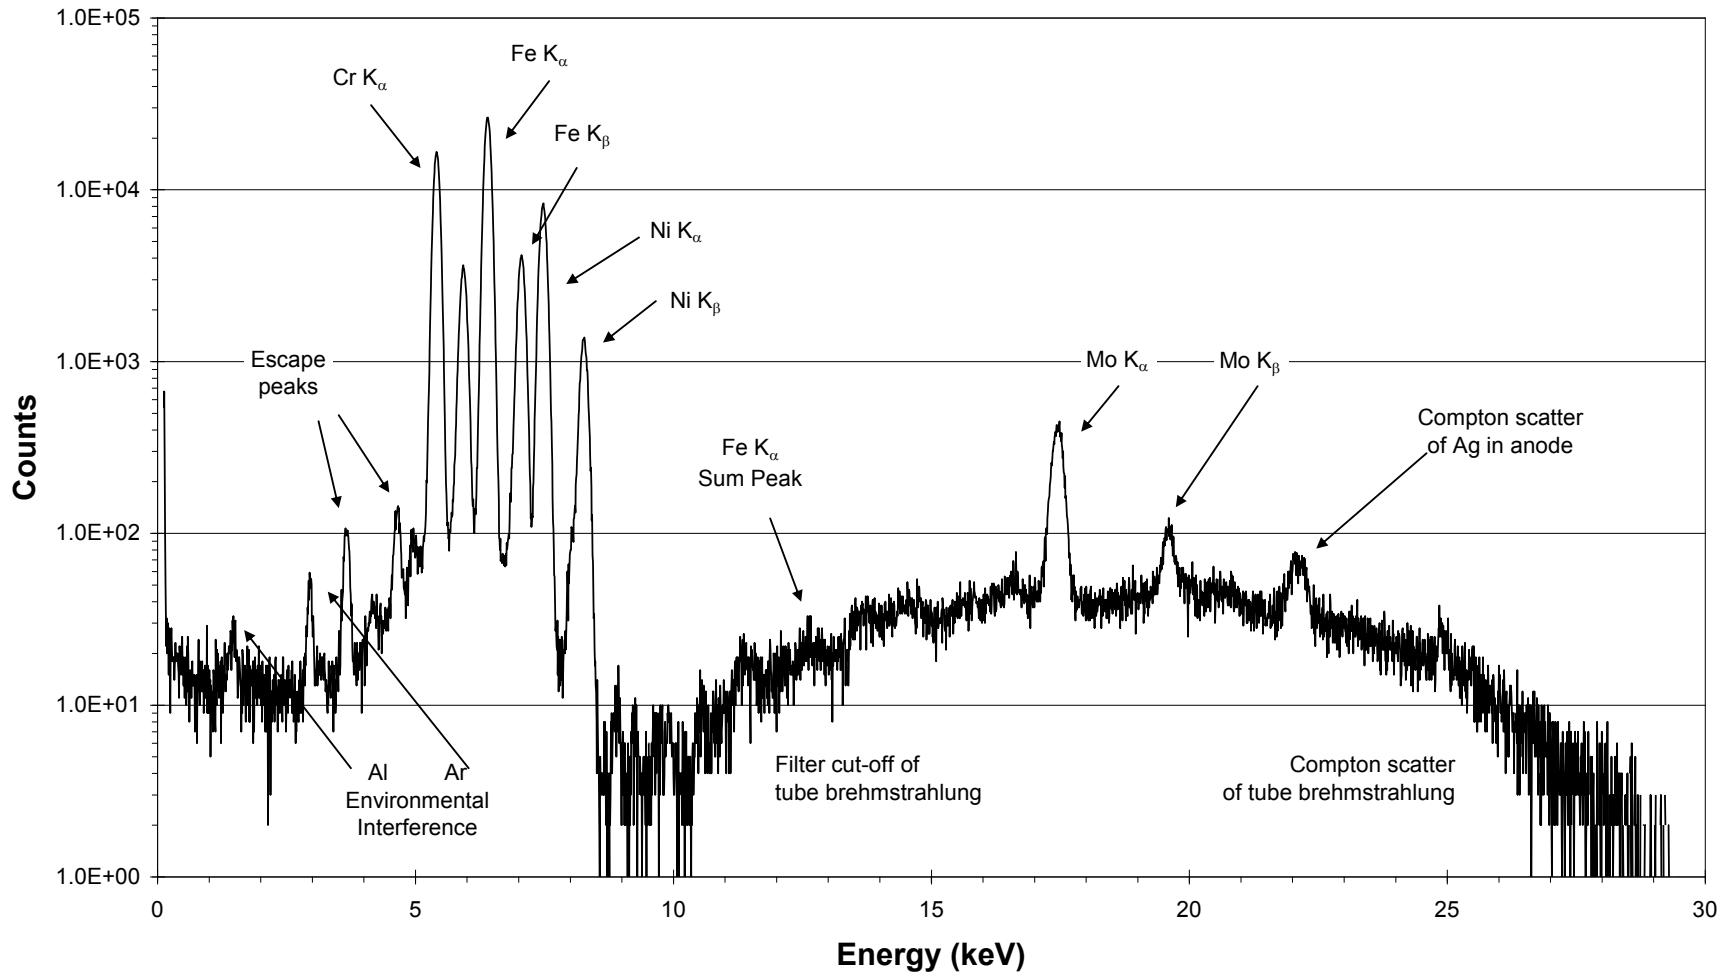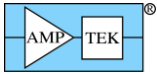
Figure 13. Spectrum illustrating the major processes that impact an XRF measurement. This spectrum was measured using an Amptek Mini-X tube (Ag anode, 30 kVp, Al filter and Al collimator), an Amptek X-123SDD detector and signal processing electronics, and a stainless steel 310 sample placed at approximately 2 cm distance with 67.5° geometry. Data were taken for 15 minutes.

HOW DOES THE ANALYSIS SOFTWARE WORK?

The XRS-FP analysis software uses the measured spectrum, such as that shown in Figure 13, to deduce the concentration of elements in the spectrum. It essentially inverts the effects described in the previous section. There are three major steps to the analysis:

1. Unravel the detector response to recover the incident photopeaks. The output of this step is a processed spectrum, ideally showing only the incident photopeaks. Figure 14 shows this.

2. Deconvolve the photopeaks to determine the intensity of the X-rays interacting in the detector. The output of this step is a table of the intensities in each photopeak to be analyzed.

3. Account for attenuation and matrix effects to determine the concentrations of the elements in the sample. The output of this step is a table of concentrations, as in Table 2.
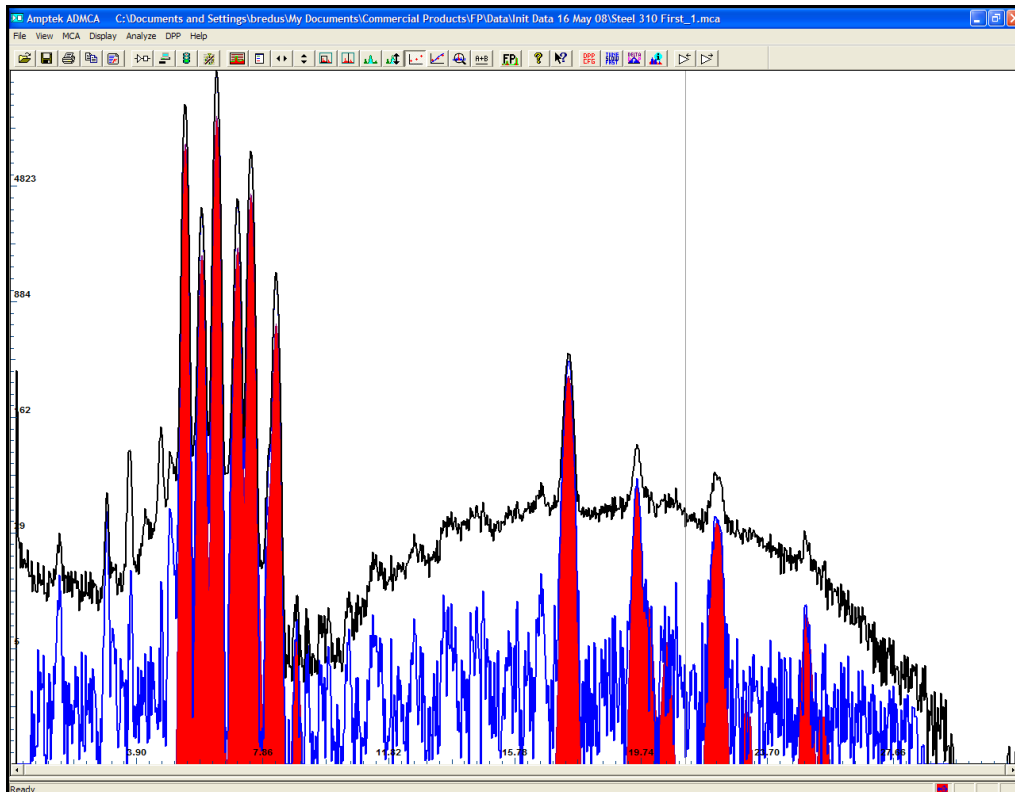

Figure 14. Plot showing the processing of the raw spectrum shown in Figure 13. This is a screen capture from Amptek's ADMCA and XRS-FP software. The black trace is the original raw spectrum. The blue trace shows the processed spectrum, after correcting for escape peaks, sum peaks, background removal, etc. The red peaks are the result of the deconvolution, showing the fitted Gaussian photopeaks. The intensities of the characteristic lines are the net area of the peaks.

Element	Intensity	Known Concentration	Measured Concentration			Chi²
Cr	316.6	25.3%	28.20%	±	0.13%	32.8
Mn	19.6	17.2%	1.51%	±	0.03%	9.9
Fe	444.4	53.2%	50.49%	±	0.20%	87.0
Co	2.5	0.0%	0.21%	±	0.01%	6.7
Ni	154.2	19.3%	18.85%	±	0.12%	24.3
Cu	1.4	0.1%	0.12%	±	0.01%	0.5
Mo	17.7	0.4%	0.61%	±	0.01%	0.2

Table 2. Results of a "standardless" analysis. The intensity column represents the net area of the deconvolved, processed peaks. The known concentration column represents the result of a chemical analysis on the standard material. The measured concentration is the result of the XRF analysis.
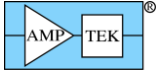
XRS-FP STEP BY STEP

We will now examine in more detail the steps by which XRS-FP goes from the raw spectrum to the final result. Before one can analyze a sample, there are several key actions required. They include (1) setting up the hardware, (2) launching and configuring the software, (3) performing an energy calibration, (4) calibrating the FP software using reference materials, and (5) acquiring the spectrum to be analyzed. This example assumes all these have been carried out, the system is configured and calibrated, and a spectrum has been acquired from the sample.

Unraveling detector response

1. Smooth fluctuations

   The first step in processing is to apply a filter which smoothes the spectrum. This reduces the effects of counting statistics and will improve the ability of the software to recognize the background and to remove the true peaks from background.
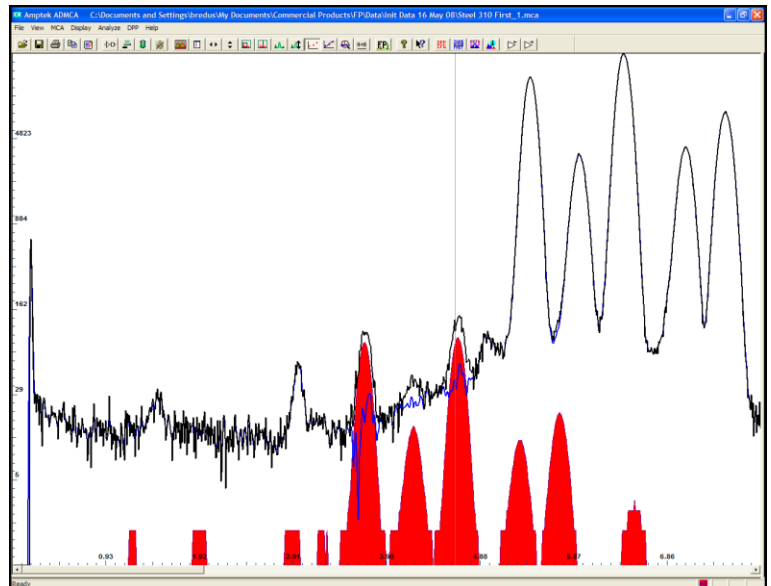
   The plot on the right illustrates (in red) the smoothed spectrum derived from the raw data (black trace).
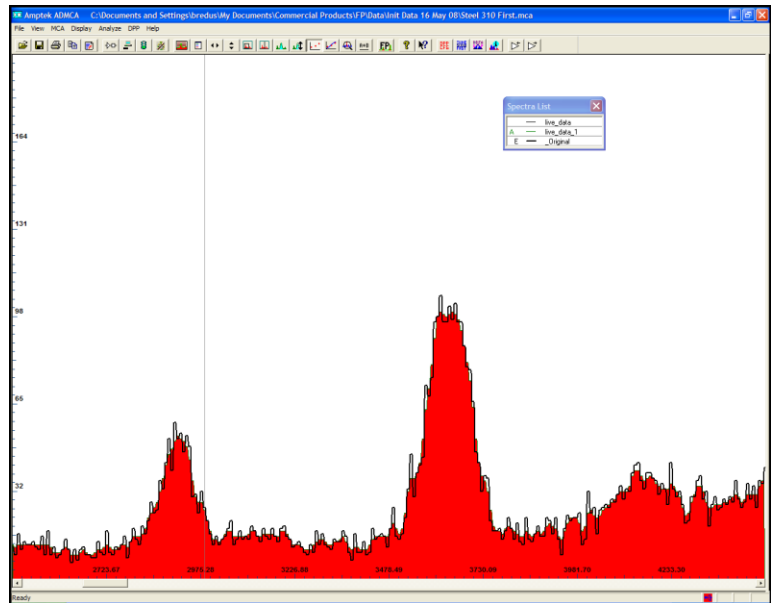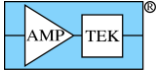

2. Correct for escape peaks

   The software applies known corrections to remove the escape peaks and put the counts into the photopeaks. This removes possible interferences and corrects the photopeak net area for escapes.

   The plot on the right illustrates (in red) the escape peaks. The black trace depicts the original spectrum, while the blue trace shows the result of the escape peak removal. These false peaks have been removed and their counts added back into the photopeaks.

3. Remove sum peak

The software applies known corrections to remove the sum peaks and put the counts into the photopeaks. This removes possible interferences and corrects the photopeak net area.

The plot on the right illustrates this, using data from a different spectrum, a pure Mn target measured at a 50% dead time. The black trace shows the raw data. The red peaks are the computed sum peaks, which are removed by XRS-FP.
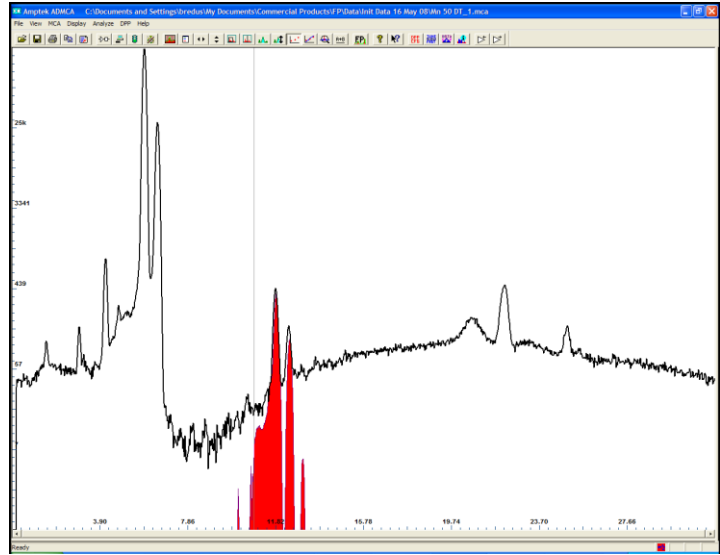

4. Remove continuum

4.1. Background removal

The first background removal step applies what is essentially a low pass filter to estimate the background, using a proprietary function to model the slowly changing background. Without attempting to separate the distinct physical mechanisms, the background is removed.

The plot on the right shows the raw spectrum as the black trace, the modeled background as the thick blue trace, and the result in red. The smoothly varying continuum has been removed.
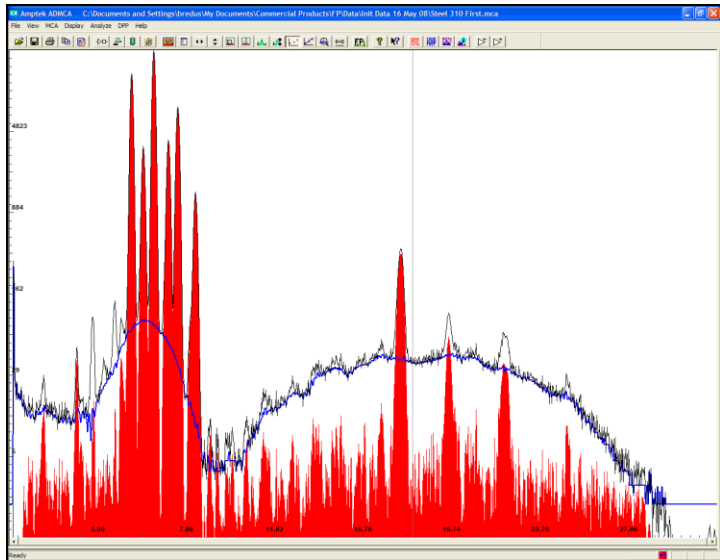
4.2. Blank subtraction

Blank subtraction removes artifacts which do not vary significantly from one spectrum to another and cannot be handled with other processing. For example, it can remove Ar peaks from air, tube scatter peaks, or some contaminant peaks from the XRF system itself.

In the plot on the right, the blue trace shows the result of blank subtraction. The Al and Ar peaks at 1.5 and 3.0 keV have been removed, along with some of the residual features above 10 keV.
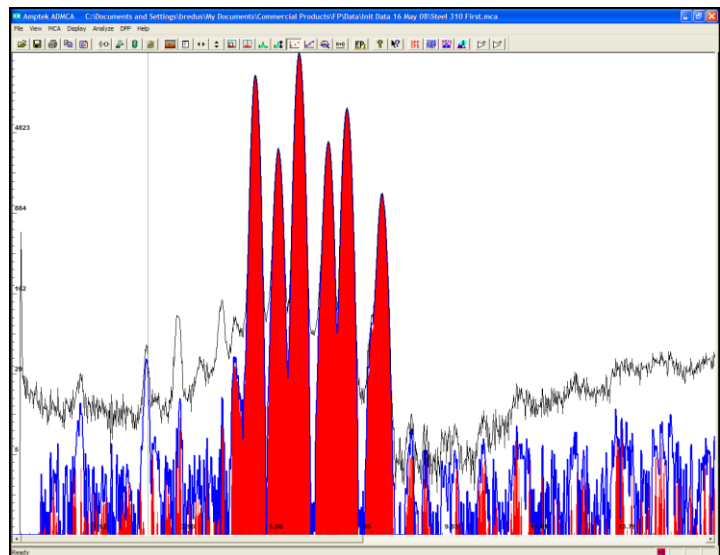
4.3. Compton peak removal

Compton peak processing estimates the intensity of the Compton and Rayleigh scatter peaks. This is used during some of the deconvolution and quantitative analysis methods, depending on the options chosen. The spectrum is not shown here since this does not appear as a change in the spectrum. This processing step produces certain parameters used in deconvolution.

Deconvolution

Deconvolution is the process of assigning the areas to the peaks of interest. The output of this step is a list of "elemental intensities", shown in Table 2 as "intensity". This is depicted in the spectrum display as the Gaussian peaks, with the appropriate areas and centroids. The results of the deconvolution for stainless steel 310 are shown in Figure 15. On the left, one can see the Gaussian fits to the major constituent elements: Fe, Cr, Ni, etc. The various other peaks and bumps are ignored. On the right, one can see the fit to Mo and also to the scatter peaks arising from the Ag anode.
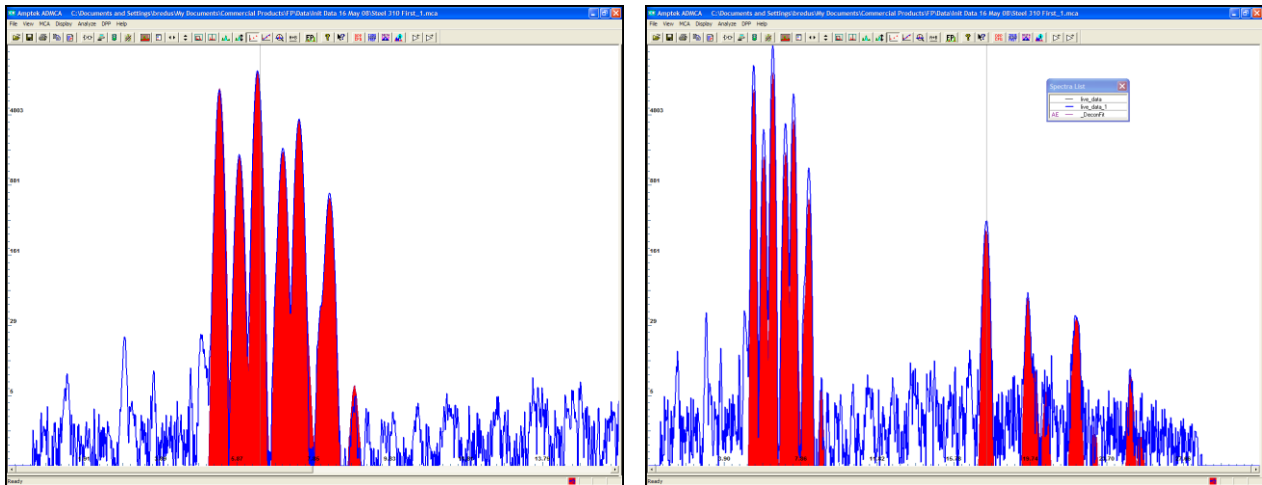

Figure 15. Plots showing the results of a nonlinear Gaussian deconvolution for steel 310.

There are several options in the software which have a major impact on this deconvolution.

First, note that XRS-FP does not automatically find the peaks in the spectrum. The user must supply a list of peaks to be analyzed (in the Specimen Component Table). The user also selects which peak is the basis for the deconvolution, i.e. K_α line. The other lines will be shown and used in analysis, but the selected line is at the core of the intensity calculation. XRS-FP can deconvolve as many as 30 peaks. Limiting the peaks to those anticipated in the sample will generally provide the best results.

Second, XRS-FP has three models for the photopeak shape: Gaussian, reference, and integrate. "Gaussian" is best under most circumstance. It does a Gaussian fit to the photopeak. This is depicted in Figure 15. Even if the peak is somewhat non-Gaussian, the high quality of the fit to the Gaussian portion can provide excellent results. "Reference" is helpful when a peak is significantly non-Gaussian, such as higher energy peaks measured with CdTe detectors. "Integrate" sums the total counts within an ROI. For simple spectra, this approach is faster and can be quite precise. The user can select which of the three photopeak models will be used for any particular peak.

Third, XRS-FP permits the user to select linear vs. nonlinear deconvolution. In a "linear deconvolution", the energies, peak widths, and ratios of the intensities are all fixed. The software then does a linear regression to fit the net area (intensity). In the "nonlinear deconvolution", the software is permitted to adjust the energies, peak widths, and line ratios (within limits) to improve the match.

Fourth, If Compton peak removal is used, and C/R ratio is checked, then XRS-FP will distinguish (and separate deconvolve) peaks due to source backscatter versus the elemental peaks. In Figure 15 (right), these options were checked and so the scatter peaks are included in the deconvolution.
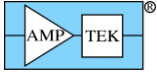
Quantitative Analysis

This is the step which relates the net intensity of each peak to the concentration in the sample. This processing includes attenuation (including the detector's sensitivity and losses in the Be window and in air) along with matrix effects, absorption and enhancement in the sample. The result of the quantitative analysis step is not shown in the spectrum, but in the output table, such as in Table 2. In XRS-FP there are three distinct correction methods that can be used and three different calibration options.

Correction methods

- FP (Fundamental Parameters): This method uses a set of parameters for each element and a matrix for inter-elemental effects. These parameters are based on the "physics" of the materials and geometry, though they may be tuned by calibration.

- SIR-FP (Scattered Intensity Ratio): This method uses information from the Rayleigh and Compton scatter peaks to provide additional information about elements (typically with low atomic numbers) that cannot be analyzed, because their x-ray lines are not detectable. Knowledge of these low Z elements is important because it permits one to better treat for matrix effects, thereby improving the accuracy of the elements which are detectable.

- SLQ (Simple Least Squares): This method uses a calibrated set of parameters and does a simple least square fit. This method requires calibration of the quantitative analysis parameters.
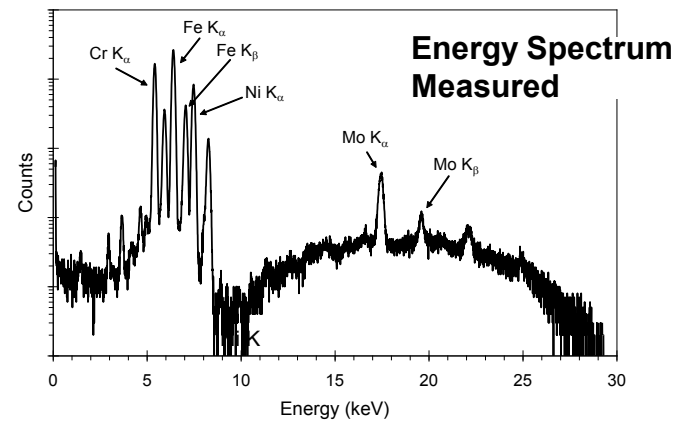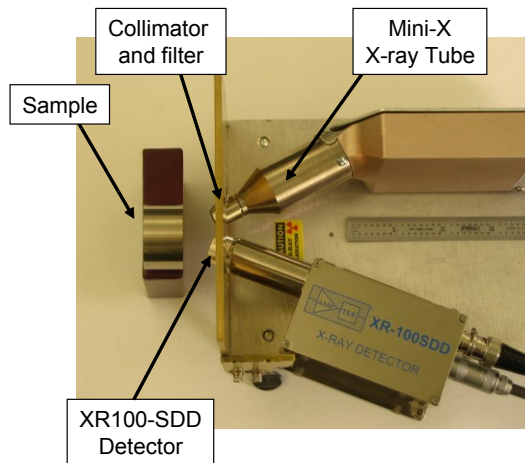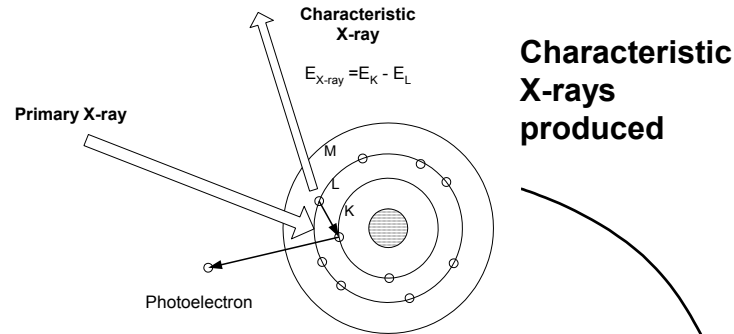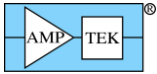
Calibration Options

Each option uses certain parameters, which account for attenuation and enhancement. There are default parameters which can be adjusted for a particular configuration. This adjustment is "calibration" of the quantitative analysis parameters.

- Standardless: Uses default values for thickness, scattering, attenuation, etc. These are derived from the physics of the materials. It is extremely critical that the set-up information be accurate!

- Single Standard: Uses measured results from a known sample to "tune" parameters for the actual setup. Actually, it will use those elements in that standard to tune parameters for only those elements; others get analyzed via standardless parameters. For example, if one is measuring steel, you can start with a known steel standard, calibrate with that, and then get pretty good results from other, similar steel alloys.

- Multiple standards (using single element standards): One starts with many samples, each made of a "pure" element. You read them in, producing an array of coefficients.

- Multiple standards (multi-element standards): In this option, you have several standards, each with several elements (and with the same elements in several standards). Each standard will give some nominal parameters; the calibration sequence combines the result for the best fit combination.

There are two distinct calibration steps needed in XRS-FP. First, one must perform an energy calibration (also called a spectrometer calibration). This assigns an energy scale to the histogram recorded by the spectrometer. To do an energy calibration, one needs a spectrum with peaks at known energies. Second, one can calibrate the parameters used in the quantitative analysis. This requires one or more samples of reference materials, with known composition.

Standardless analysis is fairly quick and easy to set up but only gives approximate results. Using standards to calibrate the analysis requires more time to get started but gives much better results, as long as the sample is similar to the standards. If one uses several steel samples to calibrate, then XRS-FP will provide excellent results on most steel alloys. However, the results will be less accurate for a brass alloy, or for jewelry, for example.

Characteristic X-rays produced

Characteristic X-ray

$E_{X-ray} = E_K - E_L$

Primary X-ray

M
L
K

Photoelectron

Collimator and filter

Mini-X X-ray Tube

Sample

XR-100SDD X-RAY DETECTOR

XR100-SDD Detector

Energy Spectrum Measured

Cr K$_\alpha$ Fe K$_\alpha$ Fe K$_\beta$ Ni K$_\alpha$ Mo K$_\alpha$ Mo K$_\beta$

Counts

0   5   10   15   20   25   30
Energy (keV)

Analysis Results

Element	Line	Concentration	Uncertainty	Intensity
Cr	Ka	28.16%	0.13%	3.2E+02
Mn	Ka	1.50%	0.03%	2.0E+01
Fe	Ka	50.54%	0.19%	4.5E+02
Co	Ka	0.21%	0.01%	2.4E+00
Ni	Ka	18.86%	0.12%	1.6E+02
Cu	Ka	0.12%	0.01%	1.4E+00
Mo	Ka	0.61%	0.01%	1.8E+01

Analysis Software Processes Spectrum

This plot illustrates the key elements in an XRF system. The key hardware includes an X-ray source, an X-ray detector, signal processing and control electronics (not shown), collimation and filters, a sample to be tested, and (not shown) radiation shielding and a mount to hold the sample in a known geometry. The incident X-rays excite the atoms in the sample, which emit X-rays at energies characteristic of the atomic transitions, which are unique to the elements in the sample. The detector measures the spectrum of X-rays emitted by the sample. Analysis software processes this spectrum to correct for various effects which distort the spectrum, ultimately determine the relative intensities of the incident characteristic X-ray lines. This is analyzed to produce the final result, the relative concentrations of the elements in the sample.