

Influence of simulation parameters on the speed and accuracy of Monte Carlo calculations using PENEPMA

X Llovet¹ and F Salvat²

¹ Scientific and Technological Centres. Universitat de Barcelona. Lluís Solé i Sabarís, 1-3. ES-08028 Barcelona, Spain

² Facultat de Física (FQA). Universitat de Barcelona. Diagonal 645, ES-08028 Barcelona, Spain

E-mail: xavier@ccit.ub.edu

Abstract. The accuracy of Monte Carlo simulations of EPMA measurements is primarily determined by that of the adopted interaction models and atomic relaxation data. The code PENEPMA implements the most reliable general models available, and it is known to provide a realistic description of electron transport and X-ray emission. Nonetheless, efficiency (i.e., the simulation speed) of the code is determined by a number of simulation parameters that define the details of the electron tracking algorithm, which may also have an effect on the accuracy of the results. In addition, to reduce the computer time needed to obtain X-ray spectra with a given statistical accuracy, PENEPMA allows the use of several variance-reduction techniques, defined by a set of specific parameters. In this communication we analyse and discuss the effect of using different values of the simulation and variance-reduction parameters on the speed and accuracy of EPMA simulations. We also discuss the effectiveness of using multi-core computers along with a simple practical strategy implemented in PENEPMA.

1. Introduction

The interaction of an electron beam with a sample is the basis of different spectroscopic techniques used for materials analysis, such as electron probe microanalysis (EPMA) [1, 2]. After impinging on the sample, the beam electrons are scattered and slowed down by the atoms of the material until they are stopped or escape from the sample. The interactions of the beam also produce secondary electrons and photons. In EPMA one is interested in the characteristic X-rays emerging from the sample. The elemental composition and stoichiometry of the probed volume of the material can be derived from the measured characteristic-line intensities. To do this, a quantitative understanding of the different processes involved in the electron transport and X-ray generation in the sample is required.

The Monte Carlo (MC) method is now widely recognised as the most accurate method to model the transport of radiation in matter. This is so because it can incorporate realistic physics interaction models and also because of its flexibility to handle complex geometries, making it well suited for the simulation of EPMA measurements, specially for heterogeneous samples such as small particles, inclusions, interfaces or multilayer films. Unfortunately, due to the random nature of the method, the results from MC simulation are affected by statistical



uncertainties which, in principle, can only be reduced at the expense of increasing the simulation time. However, significant progress has been made in recent years in the application of the so-called variance-reduction techniques which, combined with the increasing availability of fast computers, has enhanced the attractiveness of MC simulation.

Our group at the University of Barcelona develops and maintains the general-purpose MC subroutine package PENELOPE, which performs simulations of electrons, photons (and positrons) in arbitrary materials, for energies ranging from 50 eV to 1 GeV [3, 4]. PENELOPE has been applied extensively by researchers worldwide over the past two decades to a wide variety of problems in electron microscopy, radiation metrology, radiotherapy, and other fields. PENELOPE uses a combination of numerical and analytical differential cross-sections (DCS) to describe the different interaction mechanisms of electrons and photons with matter. A detailed description of the cross-section models and the electron tracking algorithm implemented in PENELOPE can be found in the code manual [3].

The programme PENEPMA was developed to facilitate the use of PENELOPE in EPMA [5–7]. The distribution package includes a manual that describes the structure and operation of the programme [8]. To set up a simulation, the user has to specify the parameters that characterise the experiment (e.g., electron beam energy, sample composition, detector take-off angle, etc.), a number of simulation parameters, which define the algorithm used for tracking the electrons within the sample, and also a set of parameters that specify the variance reduction techniques that are to be applied. These parameters largely determine the speed of the simulation; tuning these parameters adequately can greatly improve the efficiency of the simulation or, equivalently, reduce the computer time needed to get results with a prescribed accuracy.

In this study we analyse and discuss the effect of using different simulation parameters on the simulation speed for typical cases of interest in EPMA. Our aim is to establish appropriate parameter values with which simulations of X-ray intensities emitted from bulk samples, thin films on substrates and of secondary X-ray fluorescence emitted near phase boundaries may be efficiently modelled with PENEPMA.

2. Simulation parameters

2.1. Simulation algorithm

Electron trajectories and photon histories are simulated from the initial energy down to the electron and photon absorption energies E_{el} and E_{ph} , respectively, to be selected by the user, at which the particles are considered to be effectively absorbed in the medium. The default absorption energies are set to 50 eV. Note that, in the case of samples consisting of several materials, the parameters E_{el} , E_{ph} can be specified for each material. PENELOPE not only transports the primary electrons, but also all secondary electrons and photons produced in interactions of the transported particles (i.e., the complete particle shower originated by each primary electron). Secondary particles (electrons and photons) emitted with initial energy larger than the corresponding absorption energy are simulated (down to their corresponding absorption energies) after completion of each primary electron track.

The absorption energies E_{el} and E_{ph} are generally determined by the characteristics of the experiment. When we are interested only in determining the intensity of a particular X-ray line, to speed up the simulation process we can assign to E_{el} a value slightly lower than the ionisation energy of the active subshell. At the same time E_{ph} can be given a value slightly lower than that of the X-ray energy, in order to avoid spending time transporting photons that are not able to contribute to the line of interest. The X-ray spectrum will then be truncated at an energy equal to E_{ph} .

The simulation of photon tracks follows the detailed procedure, i.e., all the interaction events in a photon history are simulated in chronological order and, hence, the transport scheme is fixed and independent of user-defined parameters. In contrast, electrons can be simulated

either detailedly, interaction by interaction, or by using a so-called mixed (class II) simulation algorithm. In the latter schemes, cut-off values of the polar scattering angle (θ_c) and the energy loss (W_c) in single interactions are introduced; interactions with scattering angle $\theta \geq \theta_c$ or energy loss $W \geq W_c$ (the so called hard events), are simulated individually according to the adopted (restricted) single-scattering DCSs. These interactions are responsible for large lateral displacements or large energy-loss fluctuations, which are then described faithfully. The remaining interactions, with $\theta < \theta_c$ and $W < W_c$, are classified as soft interactions; they have a mild effect on the electron trajectories and are simulated by using simple, but still fairly accurate, multiple scattering approximations. With a judicious selection of the cut-off parameters θ_c and W_c , mixed simulation algorithms yield reliable results in moderate computation times, even for electrons with very high energies.

In PENELOPE, the user is spared from having to define the cutoff angle, which is defined internally through a recipe that involves two energy-independent used-defined, parameters, namely the average angular deflection C_1 , $C_1 \simeq 1 - \langle \cos \theta \rangle$, produced by multiple elastic scattering along a path length equal to the mean free path between consecutive hard elastic events; and the maximum average fractional energy loss C_2 in a step between hard events. The user also defines the cut-off energy loss for hard inelastic collisions W_{cc} , and the cut-off energy loss for hard bremsstrahlung emission W_{cr} , which are both independent of the energy of the electron.

The parameters C_1 and C_2 determine the simulation of elastic scattering and the tracking algorithm for electrons. Their allowed values are limited to the interval $[0, 0.2]$. If $C_1 = C_2 = 0$, the simulation of electron elastic scattering is performed in detailed mode, i.e., interaction by interaction. If the values of these parameters are not zero, the mixed simulation algorithm is switched on, and the simulation usually becomes faster. The parameters W_{cc} and W_{cr} have a very weak influence on the accuracy of the results as long as they are both smaller than the width of the bins used to tally energy distributions. When electron energy distributions are not of interest, our recommendation is to set these cut-off energies equal to 1 keV. Note that, for the sake of consistency, W_{cc} should not be larger than E_{el} , otherwise, we would exclude secondary electrons with energies larger than E_{el} . Likewise, W_{cr} should be less than E_{ph} . When $W_{cc} = 0$, the simulation of inelastic collisions is performed in detailed mode. PENELOPE also allows performing detailed simulation of the emission of bremsstrahlung photons with energies higher than 10 eV, disregarding the emission of quanta with lesser energy.

Finally, the user may define a maximum step length s_{max} of electrons in particular material bodies. This parameter has an effect only for very thin materials and when electrons are tracked using mixed simulation; it should be assigned a value of the order of one tenth of the “thickness” of the body where electrons move. This ensures that, on average, there will be more than ~ 10 soft events (hinges) along a typical electron track through that body, which is enough to “wash out” the details of the artificial distributions used to sample these events. For thick bodies, it is not necessary to limit the length of the steps. Note that s_{max} may be given different values for different “bodies” of the same material making up the sample.

2.2. Variance reduction techniques

The direct simulation of EPMA spectra may be very inefficient, even when using mixed simulation algorithms for electron transport. The reason is that radiative processes are much less probable than the dominant interactions of elastic scattering and outer-shell ionisation (or excitation). Consequently, the simulation may require very large computer times, because most electrons will not produce any photons. To cope with this difficulty, PENEPMA takes advantage of several variance-reduction techniques. The ones implemented in PENELOPE are interaction forcing, particle splitting, and Russian roulette. The most effective techniques to improve the efficiency of simulations of EPMA spectra are interaction forcing and particle splitting.

The method of interaction forcing consists of increasing the probability of a given interaction mechanism and, at the same time, assigning appropriate statistical weights to the generated secondary particles so that simulation results remain unbiased. Interaction forcing can be applied to any interaction mechanism. The practical implementation of interaction forcing consists of replacing the mean free path λ_A of the real process by a shorter one, $\lambda_{A,f} = \lambda_A/\mathcal{F}$ where $\mathcal{F} > 1$ is the forcing factor. This implies that between each pair of “real” A interactions we will have, on average, $\mathcal{F} - 1$ “forced” A interactions. We assume that the probability distribution function for the energy loss and the angular deflections (and the directions of emitted secondary particles, if any) in the forced interactions are the same as those for real interactions. When a negative forcing factor is entered, PENELOPE assumes that its absolute value represents the desired average number of forced events in a particle’s history, and determines internally the effective forcing factor that yields this number (approximately).

When a particle, with weight w_0 , is subjected to splitting, it is replaced by a number $\mathcal{S} > 1$ of identical particles with weights $w = w_0/\mathcal{S}$. Evidently, splitting leaves the simulation unbiased. It should be noted that splitting can be applied in conjunction with interaction forcing and other variance reduction techniques. In PENELOPE particle splitting is employed to increase the number of photons produced in radiative interactions of electrons. All the photons in a radiative event are sampled from the same distribution, and to compensate for the bias in the number of secondary particles, each particle is assigned a specific weight. For instance, if in a bremsstrahlung interaction the electron produces one photon, we can, instead, produce five, each with a weight of $1/5$. The benefit is that there are more photons produced per incident electron. This saves CPU time, because the simulation of the electron trajectories is generally slow.

For the practical implementation, we assume that primary particles start moving with unit weight and each secondary particle produced by a primary one is assigned an initial weight equal to that of the primary or determined by the variance-reduction trick at work. To avoid producing particles with very small weights (which would consume CPU time without contributing appreciably to the scores), interaction forcing and splitting are applied only to particles with weight in a window defined by the user. Generally, we set the weight windows so that only primary electrons and photons of the first generations undergo forced interactions.

The effectiveness of the techniques of interaction forcing and particle splitting relies on the adopted values of the corresponding forcing and splitting parameters. In an EPMA spectrum, characteristic X-ray lines arise from radiative transitions of atoms that have been ionised in an inner shell either by electron impact, or by photoelectric absorption or incoherent scattering of characteristic X-rays or bremsstrahlung photons. Therefore we will consider the forcing factors for electron-impact ionisation \mathcal{F}_{si} , bremsstrahlung emission \mathcal{F}_b , photoelectric absorption \mathcal{F}_{ph} and incoherent scattering \mathcal{F}_{co} , as well as the splitting factors for electron-impact ionisation (and the subsequent characteristic X-ray emission) \mathcal{F}_{ch} and for bremsstrahlung emission \mathcal{S}_b . Note that in the case of samples consisting of different bodies (of the same or different material), the variance-reduction parameters can be different for each body.

2.3. X-ray detector definition

In addition to the simulation and variance-reduction parameters described above, in sections 2.1 and 2.2, the speed and accuracy of the simulation of an EPMA measurement also depends on the “size” of the defined X-ray detectors. In EPMA measurements, X-ray detectors cover a very small solid angle about a specific direction (θ_d, ϕ_d) , a situation that makes the simulation of a real detector very inefficient. In PENELOPE, the user can define up to 25 ideal photon detectors, each one covering a solid angle corresponding to a “rectangle”, $(\theta_1, \theta_2) \times (\phi_1, \phi_2)$, on the surface of the unit sphere centred at the origin of the reference frame. Note that $\theta_d = \pi/2 - \chi$, where χ is the conventional take-off angle relative to the sample surface. Of course, the efficiency of

the simulation can be improved by considering a detector with a wider angular window. Notice that the different detectors considered in PENEPMA may overlap, a feature that can be used to analyse the effect of replacing the actual detector by angular windows of various sizes in a single simulation.

3. Application to EPMA simulations

In this Section, we analyse the effect of using different simulation and variance-reduction parameters on the speed and accuracy of MC simulations of typical EPMA measurements. These include the simulation of the characteristic X-ray intensity emitted from bulk samples and thin film samples, and of secondary X-ray fluorescence emitted across the boundary of two bulk materials. The simulations were performed on an 8-core personal computer (Intel Core i7 processor having a clock frequency of 2.93 GHz) with 8 GB RAM, running under Windows 7. Unless stated otherwise, all simulations were performed in serial mode.

The different simulation parameters that can be defined in PENEPMA are listed in table 1. To assess the influence of the size of the detector window on the simulation efficiency, the corresponding angular apertures $\Delta\theta$ and $\Delta\varphi$, defined as $\Delta\theta = \theta_2 - \theta_1$ and $\Delta\varphi = \varphi_2 - \varphi_1$, respectively, will be considered as parameters.

Table 1. Summary of simulation parameters.

| Parameter | Description |
|---------------------------|--|
| E_{el} | Electron absorption energy (in eV) |
| E_{ph} | Photon absorption energy (in eV) |
| C_1 | Average angular deflection in a step |
| C_2 | Maximum average fractional energy loss in a step |
| W_{cc} | Cut-off energy loss for hard inelastic collisions (in eV) |
| W_{cr} | Cut-off energy loss for hard bremsstrahlung collisions (in eV) |
| s_{max} | Maximum step length of electrons in material body (in cm) |
| \mathcal{F}_{si} | Forcing factor for inner-shell ionisation |
| \mathcal{F}_{b} | Forcing factor for bremsstrahlung emission |
| \mathcal{F}_{ph} | Forcing factor for photoelectric absorption |
| \mathcal{F}_{co} | Forcing factor for incoherent scattering |
| \mathcal{S}_{ch} | Splitting factor for characteristic X-rays |
| \mathcal{S}_{b} | Splitting factor for bremsstrahlung emission |
| $\Delta\theta$ | Polar aperture (in deg) |
| $\Delta\varphi$ | Azimuthal aperture (in deg) |

When variance reduction techniques are employed, the CPU time T required for the simulation of a given number of histories may increase. A practical measure of the effectiveness of a MC simulation run is provided by the efficiency ϵ defined by

$$\epsilon = \left(\frac{I}{\sigma}\right)^2 \frac{1}{T}, \quad (1)$$

where I is the calculated intensity of the considered X-ray line, σ is the standard deviation of I (considering only statistical uncertainties) and T is the CPU time spent in that run, excluding the PENEPMA initialisation time, which is of the order of 5 seconds. Note that PENEPMA reports statistical uncertainties at a confidence level of 99.7 % (3σ level), i.e., the result of the simulation

is expressed as $I \pm 3\sigma$. It is also worth pointing out that in EPMA the statistical uncertainty of evaluated concentrations is often given at 1σ level, e.g., at the confidence level of 68 %.

To assess the accuracy of simulated intensities, we will also give the relative uncertainty ΔI , defined as

$$\Delta I = \frac{I - I_0}{I_0} \quad (2)$$

where I_0 is the true intensity or a given reference value (see below).

3.1. Characteristic X-ray intensity from bulk, homogeneous samples

Table 2 shows the CPU time required to achieve a 1 % statistical uncertainty (at both 1σ and 3σ levels) in the intensity of the Fe K-L₂ (K α ₂) line emitted from a bulk Fe sample bombarded with 20 keV electrons at normal incidence, for different combinations of simulation parameters. The simulation efficiencies and relative errors are also listed. In all simulations, we took advantage of the axial symmetry about the beam axis and used an annular detector ($\Delta\varphi = 360^\circ$) that covers an angular aperture $\Delta\theta = 10^\circ$ around a direction of 40° relative to the sample surface. The forcing factors \mathcal{F}_{ph} and \mathcal{F}_{co} , not included in table 2, were both set to 5.

Table 2. CPU time (min) required to achieve a 1% statistical uncertainty (at 1σ and 3σ levels) in the simulation of the Fe K-L2 X-ray intensity from a bulk Fe, for different values of the simulation parameters.

| Case | E_{el} (keV) | E_{ph} (keV) | C_1 | C_2 | W_{cc} (keV) | W_{cr} (keV) | \mathcal{F}_{si} | \mathcal{F}_{br} | \mathcal{S}_{ch} | \mathcal{S}_{br} | ΔI (%) | $\epsilon(3\sigma)$ | $T(3\sigma)$ (min) | $\epsilon(1\sigma)$ | $T(1\sigma)$ (min) |
|------|--------------------------|--------------------------|-------|-------|--------------------------|--------------------------|---------------------------|---------------------------|---------------------------|---------------------------|-------------------|---------------------|-----------------------|---------------------|-----------------------|
| 1 | 0.05 | 0.05 | 0 | 0 | 0 | 0.01 | -10 | -5 | - | - | +0.00 | 0.01 | 16699 | 0.1 | 1855 |
| 2 | 1 | 1 | 0.1 | 0.1 | 1 | 1 | -10 | -5 | - | - | -0.69 | 3.8 | 43.8 | 34.2 | 4.87 |
| 3 | 6 | 6 | 0.1 | 0.1 | 1 | 1 | -10 | -5 | - | - | +0.16 | 41.8 | 3.98 | 376.3 | 0.44 |
| 4 | 6 | 6 | 0.2 | 0.2 | 1 | 1 | -10 | -5 | - | - | -0.25 | 60.1 | 2.77 | 540.6 | 0.30 |
| 5 | 6 | 6 | 0.2 | 0.2 | 6 | 6 | -10 | -5 | - | - | -0.52 | 56.5 | 2.94 | 508.5 | 0.32 |
| 6 | 6 | 6 | 0.2 | 0.2 | 1 | 1 | -10 | -5 | 4 | 4 | -0.50 | 134.2 | 1.24 | 1207.9 | 0.13 |
| 7 | 6 | 6 | 0.2 | 0.2 | 1 | 1 | -10 | -5 | 16 | 16 | -0.88 | 81.7 | 2.04 | 735.1 | 0.22 |
| 8 | 6 | 6 | 0.2 | 0.2 | 1 | 1 | -100 | -5 | 16 | 16 | +0.06 | 133.9 | 1.24 | 1204.9 | 0.13 |

We take as a reference the case of a detailed simulation with default absorption energies of 50 eV (case 1). Notice that for all the considered sets of parameters, the accuracy of the results is practically the same as that of the detailed simulation; in fact, $\Delta I < \pm 1\%$ in all the studied cases. We can see that a precision of 1 % (at 3σ level) can be achieved in about one minute CPU time by setting C_1 and C_2 to their maximum allowed values 0.2, adjusting E_{el} and E_{ph} to be slightly lower than the ionisation energy of the Fe K-shell, setting W_{cc} and W_{cr} to about 1 keV, as well as by using moderately large forcing and splitting factors (case 8). As compared with the case of detailed simulation, the use of these ‘‘optimised’’ parameters represents a gain in speed by a factor of $\sim 13,400$.

It is worth noting that the efficiency of case 6 ($\mathcal{F}_{\text{si}} = -10$, $\mathcal{S}_{\text{in}} = 4$, $\mathcal{F}_{\text{br}} = 4$) is essentially the same as that of case 8 ($\mathcal{F}_{\text{si}} = -100$, $\mathcal{S}_{\text{si}} = 16$, $\mathcal{F}_{\text{br}} = 16$). This shows that an increase of the forcing and splitting factors does not necessarily lead to an increase in efficiency, as more calculations might be required to generate each shower. Also, the increase in W_{cc} and W_{cr} up to 6 keV (case 5), although not altering significantly the simulation results, does not provide a relevant gain in efficiency, as compared with using the recommended values $W_{\text{cc}} = W_{\text{cr}} = 1$ keV (case 4).

If the uncertainty is evaluated at 1σ level, a 1 % statistical accuracy can be achieved in only 0.13 min (case 8). This shows that PENEPMA is suitable for being used as the basis of an iterative quantification method in practical EPMA work.

The effect of both forcing and splitting factors on the simulation efficiency is shown in more detail in fig. 1, which displays the efficiency and relative error as functions of \mathcal{F}_{si} , for different values of the splitting factor \mathcal{S}_{ch} . In these simulations, the following parameters were used: $C_1 = C_2 = 0.2$, $E_{\text{el}} = E_{\text{ph}} = 6$ keV, and $W_{\text{cc}} = W_{\text{cr}} = 1$ keV. No other forcing factors or splitting factors were applied. We can see that the efficiency reaches a maximum at around $\mathcal{F}_{\text{si}} \sim -75$ and then decreases. It is also evident that there is an important gain in efficiency by applying a splitting factor \mathcal{S}_{ch} of 2. However, increasing the splitting factor from $\mathcal{S}_{\text{ch}} = 2$ to $\mathcal{S}_{\text{ch}} = 5$ shows no gain. Therefore, before increasing the value of any of the simulation parameters it is advisable to perform short test simulations to verify that with the augmented parameter value not only the results remain essentially unaltered but also that there is indeed a gain in the efficiency of the simulation.

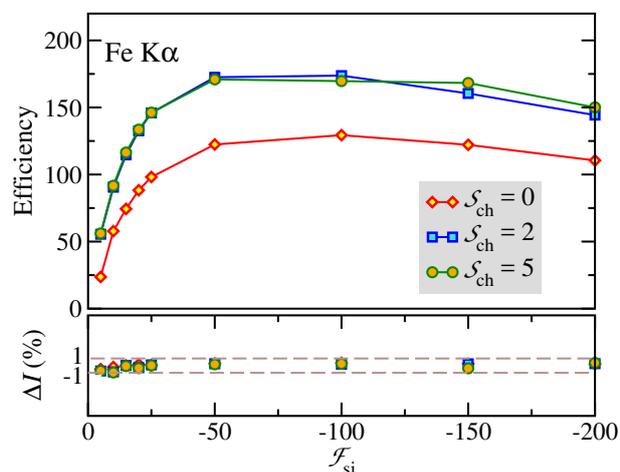


Figure 1. Efficiency and relative error as a function of forcing factor \mathcal{F}_{si} for different splitting factors \mathcal{S}_{ch} in the simulation of the Fe K-L₂ X-ray intensity from a bulk Fe sample.

Figure 2 shows the simulation efficiency as a function of the azimuthal aperture for simulations of the Fe K-L₂ line emitted from bulk Fe and of the Al K-L₃ line emitted from bulk Al, in both cases the samples being bombarded with electron beams of 20 keV at normal incidence. The polar aperture $\Delta\theta$ was set to 10° for all cases. It is seen that the efficiency increases linearly without altering the accuracy, reaching its maximum for $\Delta\phi = 360^\circ$ (annular detector). Note that this kind of configuration can in principle only be used in cases where the flux of emerging photons has axial symmetry about the z -axis (e.g., for plane-surface samples that are homogeneous within the interaction volume, and at normal incidence).

The variation of the efficiency with polar aperture $\Delta\theta$ is also shown in fig. 2 for the same simulation cases. Here, the use of large polar apertures $\Delta\theta$, while increasing the efficiency, do alter significantly the accuracy of results, as shown by the relative error, which rapidly (at $\sim \Delta\theta = 20^\circ$) departs from the 1 % uncertainty. In general, it is recommended to keep $\Delta\theta$ below 20° , especially for materials for which absorption within the target may be important (such as Al in our case).

3.2. X-ray intensity from thin films on substrates

The simulation of the X-ray intensity emitted from a relatively thin film is generally more difficult than for bulk samples because electrons might traverse the film without producing (or producing

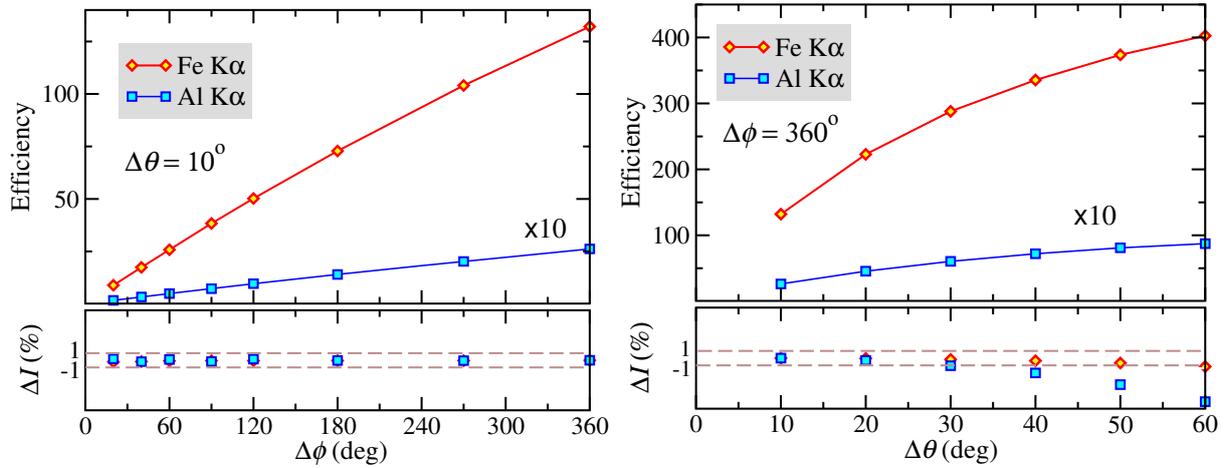


Figure 2. Efficiency and relative error as a function of azimuthal (left) and polar (right) aperture for simulations of the Fe K-L₂ line emitted from bulk Fe and of the Al K-L₃ line emitted from bulk Al.

very few) radiative interactions (see fig. 3a). Thus, the use of variance-reduction techniques is strongly recommended.

In order to determine the effect of simulation parameters on the intensity emitted from thin films, simulations were performed for a 100-nm-thick Al film deposited on Fe bombarded with 20 keV electrons at normal incidence, for different values of the simulation parameters used for both the film and the substrate. The results are tabulated in table 3, which compares the CPU time required to achieve a 1 % statistical uncertainty (at both 1σ and 3σ levels) in the Al K-L₃ ($K\alpha_1$) X-ray intensity emitted from the film. In all cases, s_{\max} was set to 10 nm for Al (1/10th of the film thickness), and the values E_{el} and E_{ph} of the electron and photon absorption energies were both set equal to 1 keV (not shown in the Table).

Table 3. CPU time (min) required to achieve an uncertainty of 1 % (at 1σ and 3σ levels) in the Al K-L₃ X-ray intensity emitted from a 100-nm-thick Al/Fe film bombarded with an electron beam of 20 keV, for different values of the simulation parameters.

| Case | Material | C_1 | C_2 (keV) | W_{cc} (keV) | W_{cr} | \mathcal{F}_{si} | \mathcal{F}_{br} | \mathcal{F}_{ph} | \mathcal{F}_{co} | \mathcal{S}_{ch} | ΔI (%) | $\epsilon(3\sigma)$ | $T(3\sigma)$ (min) | $\epsilon(1\sigma)$ | $T(1\sigma)$ (min) |
|------|----------|-------|----------------|--------------------------|-----------------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|-------------------|---------------------|-----------------------|---------------------|-----------------------|
| 1 | Al | 0 | 0 | 0 | 0.01 | -10 | -5 | - | - | - | | | | | |
| | Fe | 0.2 | 0.2 | 1 | 1 | - | - | - | - | - | +0.00 | 0.5 | 311.2 | 4.8 | 34.5 |
| 2 | Al | 0 | 0 | 0 | 0.01 | -10 | -5 | 10 | - | 2 | | | | | |
| | Fe | 0.2 | 0.2 | 1 | 1 | - | - | - | - | - | +0.04 | 1.0 | 170.6 | 8.8 | 18.9 |
| 3 | Al | 0 | 0 | 0 | 0.01 | -10 | -5 | 10 | - | 2 | | | | | |
| | Fe | 0.2 | 0.2 | 1 | 1 | -5 | -5 | - | - | - | -0.63 | 0.7 | 222.1 | 6.7 | 24.6 |
| 4 | Al | 0 | 0 | 0 | 0.01 | -20 | -5 | 10 | - | 5 | | | | | |
| | Fe | 0.2 | 0.2 | 1 | 1 | -5 | -5 | - | - | - | -0.61 | 2.9 | 58.1 | 38.3 | 4.3 |
| 5 | Al | 0 | 0 | 0 | 0.01 | -50 | -5 | 10 | 10 | 5 | | | | | |
| | Fe | 0.2 | 0.2 | 1 | 1 | -5 | -5 | 10 | 10 | - | -0.89 | 5.2 | 32.1 | 46.6 | 3.5 |
| 6 | Al | 0.2 | 0.2 | 1 | 1 | -50 | -5 | 10 | 10 | 10 | | | | | |
| | Fe | 0.2 | 0.2 | 1 | 1 | -5 | -5 | 10 | 10 | - | -0.62 | 8.0 | 20.8 | 71.9 | 2.3 |

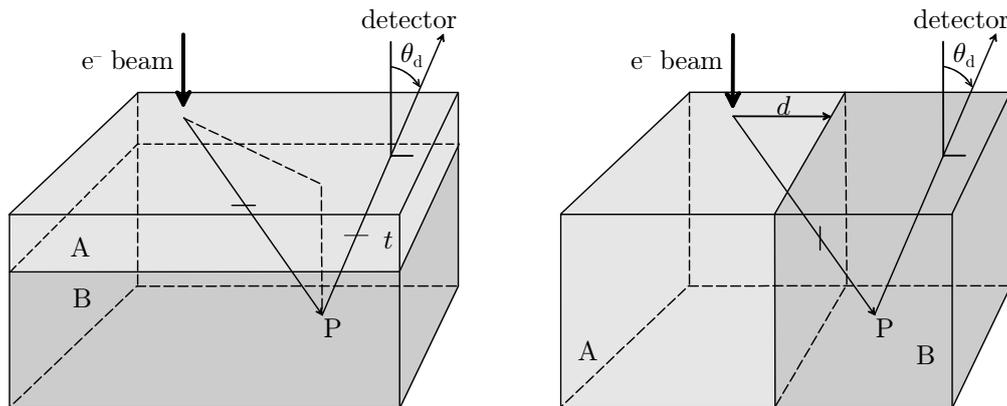


Figure 3. Geometries of the thin film on a substrate (left) and of a material couple (right) used in simulations of SF. Characteristic X-rays and bremsstrahlung photons generated in material A can undergo photoelectric absorption or incoherent scattering and produce X-ray fluorescence emission in material B.

It is apparent that when the different forcing and splitting factors are gradually increased, efficiency improves, with case 6 giving a ~ 15 times speed up as compared to case 1. We see that a considerable reduction in the CPU time can be achieved by using moderately high values of the forcing and splitting factors for both film and substrate materials. The increase in efficiency is significant when going from $\mathcal{F}_{\text{si}} = -10$ (case 3) to $\mathcal{F}_{\text{si}} = -20$ (case 4), indicating that the forcing factor \mathcal{F}_{si} for the film appears to be the most important parameter to pay attention to. It is also worth mentioning that the use of the maximum allowed values for C_1 and C_2 does not alter practically the accuracy of results, and therefore these values can be safely used even for such relatively thin films.

Case 6 represents the most efficient selection of simulation parameters, which allows the obtention of the X-ray intensity in ~ 2 minutes time (at 1σ level), showing that PENEPEMA could be used in practice as the basis of a quantification algorithm.

3.3. Secondary fluorescence from phase boundaries

A more demanding case is the simulation of secondary X-ray fluorescence emitted near the interface of two materials, for which specific simulations have been published [9]. As mentioned earlier, the characteristic lines in an EPMA spectrum arise from radiative transitions of atoms that have been ionised in an inner shell either by electron impact, or by photoelectric absorption or incoherent scattering of characteristic X-rays or bremsstrahlung photons. Characteristic and bremsstrahlung X-rays emitted from interactions of primary electrons can reach distances much larger than the electron range, and can produce X-ray fluorescence through photoabsorption and Compton scattering in a sample region (possibly having a different chemical composition) far away from the electron point of impact. This contribution is usually referred to as secondary fluorescence (SF). Petaccia *et al.* [10] have described the application of particle splitting to simulate X-ray fluorescence in EPMA using PENELOPE.

To study the relevance of simulation parameters for the calculation of SF, we will consider an ideal specimen consisting of two materials A and B separated by a plane interface perpendicular to the surface of the specimen. The electron beam impacts on the left-hand side of the specimen (material A) at a distance d from the interface (fig. 3b). Primary characteristic and bremsstrahlung photons from material A can induce emission of fluorescent photons from materials A and B.

Table 4 shows the CPU time spent to achieve an uncertainty of 1 % (at both 1σ and 3σ levels) in the simulation of the Fe K-L2 X-ray intensity emitted from a Cu/Fe couple bombarded with an electron beam of 20 keV at normal incidence, when the beam impacts on Cu (material A) at 10 μm away from the boundary, for different values of the simulation parameters. Because of the distance of the electron point of impact to the boundary, the Fe K-L2 intensity has a negligible contribution from impact ionisation by primary electrons (the range of electrons of 20-keV energy in Cu is of $\sim 1.6 \mu\text{m}$). Instead, the Fe K-L2 intensity arises from SF of Cu characteristic X-rays and bremsstrahlung generated at the Cu side of the couple, which are absorbed (and to a lesser extent incoherently scattered) at the Fe side (material B).

Table 4. Comparison of CPU time T required to achieve an uncertainty of 1 % (at 1σ and 3σ levels) in the Fe K-L2 X-ray intensity emitted from a Fe/Cu couple bombarded with an electron beam of 20 keV at 10 μm away from the boundary on the Cu side, for different values of the simulation parameters.

| Case | Material | \mathcal{F}_{si} | \mathcal{F}_{br} | \mathcal{F}_{ph} | \mathcal{F}_{co} | \mathcal{S}_{ch} | \mathcal{S}_{br} | ΔI (%) | $\epsilon(3\sigma)$ | $T(3\sigma)$ (min) | $\epsilon(1\sigma)$ | $T(1\sigma)$ (min) | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|------|----------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|-------------------|---------------------|-----------------------|---------------------|-----------------------|---|----|-----|----|----|----|----|----|-------|-----|-------|------|------|----|---|---|----|----|---|---|---|----|-----|----|----|----|----|----|-------|-----|------|------|------|----|---|---|----|----|---|---|---|----|-----|----|----|----|----|----|-------|-----|------|------|------|----|---|---|----|----|---|---|---|----|-----|----|----|----|----|----|-------|-----|------|------|------|----|---|---|----|----|---|---|---|----|-----|----|----|----|----|----|-------|-----|------|------|------|----|
| 1 | Cu | -10 | -5 | - | - | - | - | +0.00 | 0.7 | 238.2 | 6.3 | 26.4 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | Fe | - | - | - | - | - | - | | | | | | 2 | Cu | -10 | -5 | 10 | 10 | - | - | +0.01 | 0.7 | 237.5 | 6.3 | 26.3 | Fe | - | - | 10 | 10 | - | - | 3 | Cu | -10 | -5 | 10 | 10 | 5 | 5 | +0.25 | 1.8 | 90.1 | 16.6 | 10.0 | Fe | - | - | 10 | 10 | - | - | 4 | Cu | -20 | -5 | 10 | 10 | 10 | 10 | +0.45 | 1.9 | 87.6 | 17.1 | 9.7 | Fe | - | - | 10 | 10 | - | - | 5 | Cu | -20 | -5 | 10 | 10 | 20 | 20 | -0.51 | 2.1 | 79.4 | 18.8 | 8.8 | Fe | - | - | 10 | 10 | - | - | 6 | Cu | -30 | -5 | 10 | 10 | 20 | 20 | -0.68 | 1.7 | 99.5 | 15.1 | 11.0 | Fe |
| 2 | Cu | -10 | -5 | 10 | 10 | - | - | +0.01 | 0.7 | 237.5 | 6.3 | 26.3 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | Fe | - | - | 10 | 10 | - | - | | | | | | 3 | Cu | -10 | -5 | 10 | 10 | 5 | 5 | +0.25 | 1.8 | 90.1 | 16.6 | 10.0 | Fe | - | - | 10 | 10 | - | - | 4 | Cu | -20 | -5 | 10 | 10 | 10 | 10 | +0.45 | 1.9 | 87.6 | 17.1 | 9.7 | Fe | - | - | 10 | 10 | - | - | 5 | Cu | -20 | -5 | 10 | 10 | 20 | 20 | -0.51 | 2.1 | 79.4 | 18.8 | 8.8 | Fe | - | - | 10 | 10 | - | - | 6 | Cu | -30 | -5 | 10 | 10 | 20 | 20 | -0.68 | 1.7 | 99.5 | 15.1 | 11.0 | Fe | - | - | 10 | 10 | - | - | | | | | | | | | | | | | | |
| 3 | Cu | -10 | -5 | 10 | 10 | 5 | 5 | +0.25 | 1.8 | 90.1 | 16.6 | 10.0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | Fe | - | - | 10 | 10 | - | - | | | | | | 4 | Cu | -20 | -5 | 10 | 10 | 10 | 10 | +0.45 | 1.9 | 87.6 | 17.1 | 9.7 | Fe | - | - | 10 | 10 | - | - | 5 | Cu | -20 | -5 | 10 | 10 | 20 | 20 | -0.51 | 2.1 | 79.4 | 18.8 | 8.8 | Fe | - | - | 10 | 10 | - | - | 6 | Cu | -30 | -5 | 10 | 10 | 20 | 20 | -0.68 | 1.7 | 99.5 | 15.1 | 11.0 | Fe | - | - | 10 | 10 | - | - | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 4 | Cu | -20 | -5 | 10 | 10 | 10 | 10 | +0.45 | 1.9 | 87.6 | 17.1 | 9.7 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | Fe | - | - | 10 | 10 | - | - | | | | | | 5 | Cu | -20 | -5 | 10 | 10 | 20 | 20 | -0.51 | 2.1 | 79.4 | 18.8 | 8.8 | Fe | - | - | 10 | 10 | - | - | 6 | Cu | -30 | -5 | 10 | 10 | 20 | 20 | -0.68 | 1.7 | 99.5 | 15.1 | 11.0 | Fe | - | - | 10 | 10 | - | - | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 5 | Cu | -20 | -5 | 10 | 10 | 20 | 20 | -0.51 | 2.1 | 79.4 | 18.8 | 8.8 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | Fe | - | - | 10 | 10 | - | - | | | | | | 6 | Cu | -30 | -5 | 10 | 10 | 20 | 20 | -0.68 | 1.7 | 99.5 | 15.1 | 11.0 | Fe | - | - | 10 | 10 | - | - | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 6 | Cu | -30 | -5 | 10 | 10 | 20 | 20 | -0.68 | 1.7 | 99.5 | 15.1 | 11.0 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | Fe | - | - | 10 | 10 | - | - | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

The forcing factor \mathcal{F}_{si} and the splitting factor \mathcal{S}_{ch} of the material where the electron beam impacts (in this case Cu) appears to play an important role in the optimisation of the simulation, with optimum values of $\mathcal{F}_{\text{si}} = -20$ and $\mathcal{S}_{\text{ch}} = 20$. Increasing from $\mathcal{F}_{\text{si}} = -20$ to $\mathcal{F}_{\text{si}} = -30$ results in a slight decrease of the efficiency and CPU time. It is also evident that the simulation efficiency is less sensitive to the forcing factors pertaining to photoelectric absorption and Compton scattering, for both materials B and A.

Simulations of SF are CPU intensive. However, as shown here, they can benefit from the application of variance reduction techniques, to the extent that a typical 10-point simulation profile, with an uncertainty of 1 % (at 1σ level), can be obtained in around 1 h CPU time. Therefore, for such kind of simulations, it is highly recommended not to use the default simulation parameters. Likewise, performance comparisons with other SF calculations should be performed under optimised conditions [11].

4. Parallel execution

The PENEPMA distribution package contains an auxiliary programme called PENEPMA-SUM, which allows the combination of results from independent simulation runs of the same problem. Two simulation runs are effectively independent if the sequences of random numbers used in the two runs are uncorrelated. In PENEPMA this is ensured by initialising each run with seeds of the random number generator (a pair of integer values) taken from a pre-calculated list that describes states of the generator separated by 10^{14} calls in the same long sequence (see e.g., [12]).

That is, the various runs use parts of a long sequence of random numbers that are sufficiently far from each other to avoid overlaps.

This auxiliary programme can be used as a simple parallelisation device in order to improve the efficiency of a simulation problem by running several versions of the same simulation task simultaneously (in parallel) on a multi-core computer. After completing the individual simulations, the programme PENEPMASUM combines and averages the partial results producing a unique set of output files. Here we should mention that true parallelised versions of PENELOPE have been coded by other groups (see e.g., the code PENORN [13], developed for application to SEM analysis).

In order to evidence the effectiveness of the parallel execution of PENEPMASUM, 2, 4, 6, and 8 versions of the same simulation task (differing only in the seeds of the random number generator) were executed in parallel on a 8-core computer. Specifically, we simulated case 6 from Sec. 3.3 (table 4), that is, the calculations of SF (Fe K-L2 X-rays) emitted from a Fe/Cu couple when a 20-keV-electron beam impacts on the Cu side of the couple at 10 μm from the boundary. The results, shown in table 5, indicate that by running simultaneously 8 versions of the same simulation on a 8-core computer, a gain in efficiency of about a factor of 4.6 is achieved, compared to one single execution (case 1). Thus, by running PENEPMASUM in parallel on a 8-core computer, it will only take 2.4 minutes to simulate the SF intensity with a precision of 1 % (at 1σ level). Notice that, in contrast with more sophisticated parallelised versions of PENELOPE (e.g., the code PENORN [13]), no modification of the original PENEPMASUM code is required.

Table 5. Simulation results as a function of the number of parallel executions.

| Case | Number of parallel executions | $\epsilon(3\sigma)$ | $T(3\sigma)$ (min) | $\epsilon(1\sigma)$ | $T(1\sigma)$ (min) |
|------|-------------------------------|---------------------|-----------------------|---------------------|-----------------------|
| 1 | 0 | 1.7 | 99.5 | 15.1 | 11.0 |
| 2 | 2 | 3.3 | 49.7 | 30.1 | 5.6 |
| 3 | 4 | 6.1 | 27.3 | 54.9 | 3.0 |
| 4 | 6 | 7.1 | 23.3 | 64.2 | 2.6 |
| 5 | 8 | 7.8 | 21.4 | 70.0 | 2.4 |

5. Conclusion

In this study, we have shown that it is possible to perform MC simulations of typical EPMA experiments using PENEPMASUM within reasonably short CPU times by adequately choosing the simulation parameters that define the electron transport algorithm, the solid angle covered by the detector, and the adopted variance reduction techniques. The simulation time can further be reduced by taking advantage of multi-core computers by means of a very simple multiple-run strategy.

Acknowledgements

Financial support from the Spanish Ministerio de Economía y Competitividad and ERDF (projects nos. FPA2013-44549-P and FPA2016-77689-C2-2-R) and from the Generalitat de Catalunya (grant 2014 SGR 846) is gratefully acknowledged.

References

- [1] Reed S J B 1993 *Electron Microprobe Analysis* (Cambridge: Cambridge University Press)
- [2] Rinaldi R and Llovet 2015 *Microsc. Microanal.* **21** 1053

- [3] Salvat F 2015 *PENELOPE-2014. A code system for Monte Carlo simulation of electron and photon transport*. (Issy-les-Moulineaux, France: OECD/NEA Data Bank)
- [4] PENELOPE is publicly available from the OECD/NEA Data Bank, <http://www.oecd-nea.org>
- [5] Llovet X, Fernández-Varea J M, Sempau J and Salvat F 2005 *Surf. Interface Anal.* **37** 1054
- [6] Salvat F, Llovet X, Fernández-Varea J M and Sempau J 2006 *Microchim. Acta* **155** 67
- [7] Llovet X and Salvat F 2017 *Microsc. Microanal.* **23** 634
- [8] Because of its specificity, PENEPMA is not distributed with the PENELOPE package, but it is publicly available from the authors
- [9] Llovet X, Pinard P T, Donovan J J and Salvat F 2012 *J. Phys. D: Appl. Phys.* **45** 225301
- [10] Petaccia M, Segui S and Castellano G 2015 *Microsc. Microanal.* **21** 753
- [11] Ritchie N W M 2017 *Microsc. Microanal.* **23** 618
- [12] Badal A and Sempau J 2006 *Comp. Phys. Commun.* **175** 440
- [13] Bekar K B, Miller T M, Patton B W, Weber C F 2015 *Proc. Joint Int. Conf. on Mathematics and Computation, Supercomputing in Nuclear Applications and the Monte Carlo Method* (LaGrange Park, Illinois: American Nuclear Society)