APPLICATION OF PARALLEL COMPUTING TO THE MONTE CARLO
SIMULATION OF ELECTRON SCATTERING IN SOLIDS: A RAPID METHOD
FOR PROFILE DECONVOLUTION

SAND--90-2549C

DE91 001227

A. D. Romig, Jr.¹, S. J. Plimpton¹, J. R. Michael¹, R. L. Myklebust² and D. E. Newbury²¹Sandia National Laboratories, Albuquerque, NM 87185²National Institute of Standards and Technology, Gaithersburg, MD 20899

X-ray microanalysis by analytical electron microscopy (AEM) has proven to be a powerful tool for characterizing the spatial distribution of solute elements in materials. True compositional variations over spatial scales smaller than the actual resolution for microanalysis can be determined if the measured composition profile is deconvoluted. Explicit deconvolutions of such data, via conventional techniques such as Fourier transforms, are not possible due to statistical noise in AEM microanalytical data. Hence, the method of choice is to accomplish the *deconvolution* via *iterative convolutions*.¹ In this method, a function describing the assumed true composition profile, calculated by physically permissible thermodynamic and kinetic modeling, is convoluted with the x-ray generation function and the result compared to the measured composition profile. If the measured and calculated profiles agree within experimental error, it is assumed that the true compositional profile has been determined. If the measured and calculated composition profiles are in disagreement, the assumptions in the physical model are adjusted and the convolution process repeated. To employ this procedure it is necessary to calculate the x-ray generation function explicitly. While a variety of procedures are available for calculating this function,²⁻⁴ the most accurate procedure is to use Monte Carlo modeling of electron scattering.^{5,6}

Research Objectives: Monte Carlo simulations of electron scattering, and the calculation of the x-ray generation volume, in thin film and bulk specimens have been performed for many years.

There are many uses of Monte Carlo simulations in electron microscopy and x-ray microanalysis,

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED
MASTER

and one very important application is to use it as a tool to investigate x-ray spatial resolution. The ideal application of the Monte Carlo simulation (once it is demonstrated that a particular Monte Carlo code actually models all aspects electron scattering) is to explore a parameter space larger than one can conveniently explore experimentally. For example, if a Monte Carlo simulation accurately describes the spatial resolution for x-ray microanalysis in a thin film of some thickness, it is easier and more efficient to examine spatial resolution in specimens of other thicknesses by Monte Carlo simulation than by additional experiments. However, the widespread application of Monte Carlo modeling has never been realized in the microanalysis community for such problems because of the computation intensive nature of the Monte Carlo algorithm (e.g. many hours of computer time were required to run one simulation which might improve ones ability to interpret data; the time required was so large that it was not possible to use the calculations to improve experimental measurements in a time effective way). However, the development of parallel computation architectures has vastly improved the efficiency of Monte Carlo simulations. It was the objective, therefore, of this work to investigate the utility of the parallel computation methodology for Monte Carlo simulations of electron scattering in solids as it applies to x-ray microanalysis in bulk and thin film specimens. To attain the objective, the NIST (National Institutes of Standards and Technology) Monte Carlo algorithm^{5,6} was modified to utilize a parallel architecture. The new parallel code was tested on a variety of microanalytical problems, in bulk and thin film specimens, to demonstrate its validity. An important result of the parallel Monte Carlo simulations was the significant reduction in computation time compared to Monte Carlo simulations run on a variety of traditional computers (from workstation to supercomputer). It was clearly shown that the time advantage in solving such problems with parallel computation methods is so significant that Monte Carlo modeling becomes a useful real time tool for microanalysis. Specifically, it was demonstrated that parallel Monte Carlo calculations are a very rapid method for the deconvolution of composition profiles measured by AEM.

Monte Carlo Modeling: The basic principle of a Monte Carlo electron trajectory simulation is to calculate the path of each electron in a step-wise manner, accounting for both elastic and inelastic scattering.⁶ The distance between scattering events, the scattering angles and the rate of energy loss with distance are calculated from physical models. Random numbers are used to select the scattering angles so that, for a large number of electron trajectories, the resulting probability histogram closely resembles the actual angular distribution for single scattering events. To obtain physically meaningful results, a large number of electron trajectory simulations are required (usually on the order of 10^5). For electrons scattering in a solid, elastic scattering generally results in trajectory changes of less than 5° , but scattering through angles as large as 180° is possible. Elastic scattering events do not alter the energy of the electron. Inelastic scattering decreases the energy of the electron, but the scattering angle is very small, typically less than 0.1° . In this Monte Carlo model, the energy loss is described by the Bethe continuous energy loss function and any slight change in electron trajectory during an inelastic scattering event is ignored. For more detail, please refer to the review paper in this proceedings.⁷

Parallel Computation of Monte Carlo Scattering: Monte Carlo simulations are obviously very computation intensive. Even in a thin film, a very large number of calculations are required to simulate the electron beam/specimen interaction for a single analysis point (i.e., 10^5 trajectories incident at a single point). For profile deconvolution, as many as 100 incident points may be required to calculate the composition profiles and up to 10 iterations using various physical models may be required. Hence, up to 10^8 trajectories may be required to determine a single composition profile. The computation intensive nature of Monte Carlo simulations and the fact that the scattering of each primary electron is independent of all other primary electrons suggests that such problems are ideally suited to computer simulations using machines with a parallel architecture.

The NCUBE 2 parallel computer used in this work has 1024 processors; hence 1024 trajectories (each with a unique starting random number) can be simulated simultaneously. The NCUBE 2 has

a computing architecture well-suited for this approach; it is a MIMD (multiple instruction-multiple data) machine where each processor (or node) executes a copy of the simulation program independently on its own data set. Messages can be passed between nodes at the end of the simulation to sum the results. Each of the 1024 nodes on the NCUBE 2 at Sandia has a ~1.0 Mflop CPU with 1 Mbyte of memory. The nodes are interconnected in a hypercube topology.

Only two basic changes were necessary in adapting the original serial version of the Monte Carlo simulation to a parallel code. The first was a restructuring of the simulation so that the basic algorithm became the following: (1) input the simulation parameters for a run of N total trajectories and send a copy of the parameters to all P nodes, (2) simulate N/P trajectories on each node, (3) combine the binned simulation data by sending messages between nodes, and (4) output the results. Step 2 requires all the computation time; step 3 can be performed very quickly due to the high connectivity of the hypercube architecture. It should be noted that this restructuring requires no modification of the basic physics routines which calculate the electron trajectories.

The second change was to alter the random number generation scheme. In the typical serial Monte Carlo program, an initial random number seed is supplied by the user. A sequence of random numbers is generated from the seed and used for successive trajectories, i.e., the first electron uses the first few random numbers, the next trajectory uses a few more, etc. The drawback is that one does not know *a priori* what random number the *n*th trajectory will start with. Hence, a parallel simulation cannot duplicate the serial results. This is important for both debugging/developing the simulation and because, in practice, programs can be run on the NCUBE 2 on any available power-of-two subset of nodes (up to 1024 maximum).

To overcome this, a pseudo-random number tree scheme was implemented as discussed in Fox⁸ and detailed by Frederickson⁹. It insures identical results independent of the number of processors the simulation runs on. Each node is assigned one or more sets of N/1024 trajectories to simulate

(e.g., for a run of 10,000 trajectories on 64 nodes, each node simulates 16 sets of 9 or 10 trajectories each). When each set is begun, a new random number seed is used to restart the primary random number generator. The new seeds come from a list of 1024 seeds each node calculates at the beginning of the simulation run using a second random number generator. The second generator is started with the initial seed specified by the user; hence, one random number seed is still sufficient to create all the trajectories. Since each node uses the same list of 1024 seeds and uses the appropriate seeds from the list for its own set(s) of $N/1024$ trajectories, the same trajectories are simulated regardless of which processors they run on. Note that the accuracy of this method depends on the second generator producing seeds that are far apart in the sequence of random numbers created by the primary generator. As pointed out by Percus¹⁰ this is not guaranteed for all choices of generators, but in practice satisfactory results were achieved; no statistical biases were seen in the Monte Carlo results. Evidence that the random number generation scheme or, in fact, the entire processes of parallelizing the code did not affect the results of the Monte Carlo simulation is shown in Figure 1. Here, backscatter electron coefficients, in bulk targets at 20 KeV, versus atomic number are shown. The open circles represent the parallel computation results and the theoretical curve is a parametric fit to experimental data.¹¹ The agreement is excellent.

The timing results from the simulation on all 1024 nodes of the NCUBE 2 were compared to three conventional computers, a VAX 785, a DEC 3100 workstation, and one processor of a CRAY-XMP 4/16. The algorithm is written in standard Fortran 77; all the codes were compiled using the best optimization level available on that machine. For a run of 100,000 electrons (20 keV) into bulk Cu the timings (in seconds) were as follows: NCUBE = 32.1, VAX = 64650, DEC = 23900, CRAY = 4699. Thus, the NCUBE is 146 times faster than the CRAY for bulk simulations. The NCUBE was found to be 85% efficient for bulk simulations, meaning 1024 nodes are 870 times faster ($.85 \times 1024$) than a single node. The extra 15% is lost in the waiting that

occurs as some nodes finish their trajectories a little sooner than others and in the part of the simulation that is not performed in parallel (e.g. the input/output).

For thin film runs of 100,000 electrons into a Cr/Fe/Ni film (discussed in this paper), the timings (in seconds) were as follows: NCUBE = 4.81, VAX = 1350, DEC = 425, CRAY = 56.73. Here the NCUBE is 11.8 times faster than the CRAY. The difference between this and the bulk case is that in thin film simulations there are a few high-angle scattering events that leave a high-energy secondary electron traversing the film parallel to the film surface. This event takes a much longer time to simulate than a typical trajectory where the primary electron passes quickly through the film. These anomalously long trajectories skew the computational balance so that 95% of the processors end up having to wait for the other 5% to finish. This cut the parallel computing efficiency to 13% in this instance. However, some of this speed loss is recovered in practice since (as in this paper), a thin film simulation is usually run on a series of spaced beam impact points (e.g., across an interface). The simulation assigns a subset of processors to the trajectories for each point and the whole set of impact points can then be run in parallel. This has the effect of increasing the efficiency (and thus the overall speed for the simulation) since each processor simulates more trajectories and the run times of different processors became more uniformly balanced. For example, the same thin film simulation on 16 incident points (100,000 electrons incident at each point) runs in 20.0 seconds on the NCUBE, which is 45 times faster than it runs on the CRAY.

It is worth noting that Monte Carlo trajectory simulations such as this one can be vectorized for the CRAY. However, this requires significant changes to the basic algorithm and data structures so that vectors of particles are followed simultaneously. Work on codes similar to this one have shown speed-ups on the order of 8x can be achieved by such efforts¹².

Simulation of Monolayer Impurity at a Grain Boundary: This simulation considered the calculation of a composition profile in a thin film target. The target was a 25 nm thick film of Al, with a monolayer thick slab of Cu was oriented parallel to the incident electron beam. The results of two simulations are shown in Figure 2. The first simulation is for a 1 nm probe at 100 KeV (typical of a modern cold field-emission analytical electron microscope). The apparent composition at the boundary is approximately 62 wt.% Cu. Apparent compositions of less than approximately 0.1 wt.% would be below the background and not measurable. Hence, at distances exceeding 3 nm from the boundary the Cu signal has essentially decreased to zero. The second simulation is for 20 nm probe at 200 KeV, typical of a modern thermionic source (LaB₆) analytical electron microscope. The apparent composition at the boundary is approximately 4.9 wt.% Cu. Apparent compositions of less than approximately 0.1 wt.% would be below the background and not measurable. Hence, at distances exceeding 30 nm from the boundary the Cu signal has essentially decreased to zero. An interesting conclusion to be drawn from these simulations is the order in magnitude improvement in spatial resolution for microanalysis in field-emission electron microscopes, even at 100 KeV.

Comparison of Monte Carlo Simulation with Experimental Results: In order to verify the accuracy of the Monte Carlo electron trajectory simulation running on the parallel processor, the calculated results were compared with experimental measurements. Previously reported experimental data^{12,14} are ideal for this comparison since the specimen thicknesses were known, the probe size had been accurately measured¹⁵ and the specimen was well characterized. A specimen of 40 wt.% Ni - 30 wt.% Cr-Fe, which had been heat-treated to produce α precipitates (80 - 90 wt.% Cr) in a γ matrix (25 - 29 wt.% Cr) was used to generate the experimental data¹³. The α/γ interface was ideal for measuring concentration profiles since these interfaces are very straight and atomically discrete.

The 1.0 nm (full-width at half-maximum) electron probe was stepped across an α/γ interface while recording the x-ray spectrum obtained at each point. The specimen thicknesses were also measured for each profile. Four different specimen thicknesses were analyzed. As has been previously shown, the width of the concentration profiles is a measure of the amount of beam broadening which occurs as the electron probe traverses the specimen¹⁶. If the composition change across the interface is normalized to 100%, the distance between the 12% and the 88% points on the profile represents the diameter which contains 50% of the incident electrons. The diameter which contains 90% of the incident electrons may be obtained by multiplying the 50% diameter by 1.82. The 90% diameter is the most commonly used definition of x-ray spatial resolution.

Figure 3 shows an example of the actual experimental data.¹³ In this case the specimen was 112 nm thick. The analytical electron microscope was operated at 100 KV and the probe size was 1 nm. The calculated curve was generated by Monte Carlo calculation. The excellent fit between the experimental data and Monte Carlo simulation is apparent.

The Monte Carlo simulation was used to model the same specimen thicknesses (55 nm, 95 nm, 112 nm and 150 nm) as in the experimental results using a 1.0 nm FWHM electron probe diameter. The 90% spatial resolutions were calculated in the same manner as those for the experimental profiles. The Monte Carlo simulations for the four specimen thicknesses are shown in Figure 4. Table 1 compares the spatial resolutions determined by experimental measurements¹³ and the Monte Carlo simulation for the same four specimens thicknesses. All four curves calculated by the Monte Carlo simulation agree well with the measured data. A comparison of the measured profile width relative to the width determined from the Monte Carlo simulations is good. The 90% widths as determined by the Monte Carlo simulations are 2.7 nm, 4.9 nm, 6.7 nm and 10.3 nm for the four thicknesses. The measured values are 5.2, 6.6, 9.0 and 12.0 nm.

Table 1

Comparison of x-ray spatial resolution at α/γ interface in an Fe-Ni-Cr alloy. The 90% x-ray spatial resolutions determined experimentally and by Monte Carlo calculation are compared for thin films of four different thicknesses.

<u>Specimen Thickness (nm)</u>	<u>MC (nm)</u>	<u>Experimental (nm)¹³</u>
55.	3.6	5.2
95.	6.7	6.6
120.	7.2	9.0
150.	10.8	12.0

In all cases the Monte Carlo results underestimated the experimental measurements. This result is expected since the Monte Carlo calculations are not influenced by experimental difficulties which include the exact orientation of the interface parallel to the electron beam, specimen drift and errors in probe placement. These results show that the Monte Carlo electron trajectory simulation can accurately model the electron trajectories in specimens of medium atomic number.

As a further numerical experiment, the same simulations were repeated with a 10 nm probe, as shown in Figure 5. In this case the effect of thickness on resolution is less dramatic. It clearly illustrates that probe size is the **most** important parameter for high resolution x-ray microanalysis in the analytical electron microscope

Conclusions: The simulation of electron scattering in solids using Monte Carlo techniques is well suited to parallel computation. Significant gains in computation time are realized making

explicit calculation of convoluted composition profiles possible. Computation time is sufficiently shortened so that such simulations may be used in a real-time experimental environment. Because such simulations break naturally into independent computational pieces requiring little intercommunication, they are ideal candidates for fast parallel implementation on a MIMD machine such as the NCUBE 2. Similar performance gains should be possible for other kinds of Monte Carlo transport simulations (ion or photon), as well as codes modeling other electron-solid interactions, such as Auger electron production. Even on smaller, cheaper parallel machines, run times for such codes can be competitive with traditional supercomputers like the CRAY.

Acknowledgements:

A portion of this work performed at Sandia National Laboratories and supported by the US DOE under contract no. DE-AC04-76DP00789. The computational methods used here were developed with the support of the Applied Mathematical Science Program, DOE Office of Energy Research. The National Institute of Standards and Technology is a U.S. Department of Commerce facility.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References:

1. D. B. Williams and A. D. Romig, Jr., "Studies of interfacial segregation in the analytical electron microscope: a brief review," *Ultramicroscopy* 30: 38, 1989.
2. J. I. Goldstein, J. L. Costley, G. W. Lorimer and S. J. B. Reed, "Quantitative x-ray analysis in the electron microscope," *SEM/1977 I*, 315.
3. A. D. Romig, Jr. and M. J. Cieslak, "Solute Segregation to phase interfaces and grain boundaries: studies by analytical electron microscopy and profile deconvolution," *Analytical Electron Microscopy - 1987*, ed. D. C. Joy, San Francisco Press, 25.
4. P. Doig, D. Lonsdale and P. E. J. Flewitt, "The spatial resolution of x-ray microanalysis in the scanning transmission electron microscope," *Phil. Mag.* 41A: 759, 1980.
5. D. E. Newbury, "Electron beam-specimen interactions in the analytical electron microscope," in *Principles of Analytical Electron Microscopy*, D. C. Joy, A. D. Romig, Jr. and J. I. Goldstein, eds., New York: Plenum Press, 1986, 1.
6. D. E. Newbury and R. L. Myklebust, "Monte Carlo electron trajectory simulation of beam spreading in thin foil targets," *Ultramicroscopy* 3: 391, 1979.
7. D. E. Newbury and R. L. Myklebust, "Monte Carlo electron trajectory simulation of x-ray emission from films supported on substrates," *Microbeam Analysis - 1990*, ed. D. B. Williams, P. Ingram and J. R. Michael, San Francisco Press, 1991. (this volume)
8. G. Fox, et. al., *Solving Problems on Concurrent Processors*, New Jersey: Prentice Hall, 1988, 210.
9. P. Fredickson, R. Hiromoto and J. Larson, "A parallel Monte Carlo transport algorithm using a pseudo-random tree to guarantee reproducibility," *Parallel Computing* 4: 281, 1987.
10. O. Percus and M. Kalos, "Random number generators for MIMD parallel processors," *J. Parallel and Distributed Computing* 6:477, 1989.
11. J. I. Goldstein, et. al., *Practical Scanning Electron Microscopy and Microanalysis*, New York: Plenum Press, 1981, 76.

12. K. Miura, "EGS4V: vectorization of the Monte Carlo cascade shower simulation code EGS4," *Comp. Phys. Comm.* 45:127, 1987.
13. J. R. Michael, D. B. Williams, C. F. Klien and R. Ayer, "The measurement and calculation of the x-ray spatial resolution in the analytical electron microscope," *J. Microscopy* 160: 41, 1990.
14. J. R. Michael, "The effect of specimen thickness and electron probe characteristics on the x-ray spatial resolution in the analytical electron microscope," *Microbeam Analysis - 1989*, ed. P. E. Russell, San Francisco Press, 496.
15. J. R. Michael, "Probe size measurement in a field-emission STEM," *Microbeam Analysis - 1988*, ed. P. E. Russell, San Francisco Press, 60.
16. J. R. Michael and D. B. Williams, "A consistent definition of probe size and spatial resolution in the analytical electron microscope," *J. Microscopy* 147: 289, 1987.

Figure Captions:

Figure 1 - Comparison of backscatter electron coefficients for bulk targets at 20 KeV, illustrating the validity of the new parallel simulation. Theory line is parametric fit to experimental data.¹¹

Figure 2 - Calculation of the spatial distribution of x-rays from an Al specimen 25 nm thick. The specimen contains a one monolayer (0.3 nm thick) Cu film oriented parallel to the incident primary electron beam. Each point represents 10^5 electron trajectories. For a 1 nm probe at 100 KeV, the "measured" boundary composition is approximately 62 wt.% Cu. For a 20 nm beam at 200 KeV, the "measured" boundary composition is approximately 4.9 wt.% Cu. The distance at which the boundary film can no longer be detected is 3 nm for the 1 nm probe and 30 nm for the 20 nm probe.

Figure 3 - Experimental data for a thin film in which a precipitate phase containing 89.1 wt.% Cr is in contact with a matrix containing 26.2 wt.% Cr.¹³ The film is 112 nm thick. The 100 KeV electron probe size was 1 nm. A profile was measured across the phase interface. The experimental data is compared to a Monte Carlo simulation.

Figure 4 - Comparison of composition profiles as a function of thickness in specimens as described in Figure 3. The 100 KeV electron probe size was 1 nm. Four specimen thickness, 55, 95, 112 and 150 nm, are shown. The 1.8x column shows the measured spread of the profile (for 90% of the electrons) as determined by the Monte Carlo simulation. The Exp column are the results of the experimental measurement.¹² The Monte Carlo simulations consistently underestimate the experimental data, as discussed in the text. All units in the data table insert are in nm.

Figure 5 - Repeat of the numerical Monte Carlo simulations described in Figure 4. In this case, a 10 nm electron probe is used. A comparison of the results shown in Figures 4 and 5 illustrate the paramount importance of probe size in high spatial resolution x-ray microanalysis. All units in the data table insert are in nm.

Figure 1

Comparison of backscatter electron coefficients for bulk targets at 20 KeV, illustrating the validity of the new parallel simulation. Theory line is parametric fit to experimental data.¹¹

Figure 2

Calculation of the spatial distribution of x-rays from an Al specimen 25 nm thick. The specimen contains a one monolayer (0.3 nm thick) Cu film oriented parallel to the incident primary electron beam. Each point represents 10^5 electron trajectories. For a 1 nm probe at 100 KeV, the "measured" boundary composition is approximately 62 wt.% Cu. For a 20 nm beam at 200 KeV, the "measured" boundary composition is approximately 4.9 wt.% Cu. The distance at which the boundary film can no longer be detected is 3 nm for the 1 nm probe and 30 nm for the 20 nm probe.

Figure 3

Experimental data for a thin film in which a precipitate phase containing 89.1 wt.% Cr is in contact with a matrix containing 26.2 wt.% Cr.¹³ The film is 112 nm thick. The 100 KeV electron probe size was 1 nm. A profile was measured across the phase interface. The experimental data is compared to a Monte Carlo simulation.

Figure 4

Comparison of composition profiles as a function of thickness in specimens as described in Figure 3. The 100 KeV electron probe size was 1 nm. Four specimen thickness, 55, 95, 112 and 150 nm, are shown. The 1.8x column shows the measured spread of the profile (for 90% of the electrons) as determined by the Monte Carlo simulation. The Exp column are the results of the experimental measurement.¹² The Monte Carlo simulations consistently underestimate the experimental data, as discussed in the text. All units in the data table insert are in nm.

Figure 5

Repeat of the numerical Monte Carlo simulations described in Figure 4. In this case, a 10 nm electron probe is used. A comparison of the results shown in Figures 4 and 5 illustrate the paramount importance of probe size in high spatial resolution x-ray microanalysis. All units in the data table insert are in nm. All units in the data table insert are in nm.

-END-

DATE FILMED

11 / 20 / 90