



ELSEVIER

Contents lists available at ScienceDirect

Ultramicroscopy

journal homepage: www.elsevier.com/locate/ultramicGentle STEM: ADF imaging and EELS at low primary energies[☆]Ondrej L. Krivanek^{a,*}, Niklas Dellby^a, Matthew F. Murfitt^a, Matthew F. Chisholm^b, Timothy J. Pennycook^b, Kazutomo Suenaga^c, Valeria Nicolosi^d^a Nion Co., 1102 8th St., Kirkland, WA 98033, USA^b Oak Ridge National Laboratory, Materials Science and Technology Division, Oak Ridge, TN 37831-6069, USA^c National Institute of Advanced Industrial Science and Technology (AIST), Tsukuba 305-8565, Japan^d Department of Materials, University of Oxford, Parks Road, Oxford OX1 3PH, UK

ARTICLE INFO

Keywords:

STEM
ADF
EELS
Aberration correction
Nanotube
Graphene

ABSTRACT

Aberration correction of the scanning transmission electron microscope (STEM) has made it possible to reach probe sizes close to 1 Å at 60 keV, an operating energy that avoids direct knock-on damage in materials consisting of light atoms such as B, C, N and O. Although greatly reduced, some radiation damage is still present at this energy, and this limits the maximum usable electron dose. Elemental analysis by electron energy loss spectroscopy (EELS) is then usefully supplemented by annular dark field (ADF) imaging, for which the signal is larger. Because of its strong *Z* dependence, ADF allows the chemical identification of individual atoms, both heavy and light, and it can also record the atomic motion of individual heavy atoms in considerable detail. We illustrate these points by ADF images and EELS of nanotubes containing nanopods filled with single atoms of Er, and by ADF images of graphene with impurity atoms.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

The first field emission STEM designed and built in Crewe's laboratory operated at 30 keV [1,2], most likely because a 30 keV instrument was easier to build from scratch than one operating at a higher energy. The final version of this instrument had an objective lens with a spherical aberration coefficient C_s of 0.3 mm and also a low chromatic aberration coefficient C_c , and it attained about 3 Å resolution. It was the first electron microscope able to resolve single heavy atoms [3], and to produce high-quality electron energy-loss spectra from atomic-dimension sample areas [4]. Aiming for better resolution, Crewe's lab embarked on building a series of aberration correctors, a 1 MeV STEM and an aberration-corrected 200 kV STEM that was supposed to give 0.5 Å resolution, but these projects were not completed [5,6]. However, higher primary energy STEMs were built by others [7–9], and the standard STEM operating energy soon became 100 keV and later on 200 and 300 keV.

100 keV and higher energy operation brought advantages such as better resolution and an ability to look at thicker samples. It allowed researchers to examine materials and their interfaces at atomic or near-atomic resolution [10–12]. However, in low *Z* materials such as carbon and boron nitride, the higher operating energy produced significant knock-on damage (e.g. [13]), which limited the usable electron doses and hence the counting statistics of the experimental data. Theoretical estimates for knock-on threshold in carbon and boron nitride are close to 80 kV [14], and operating below this threshold should therefore either eliminate or considerably reduce the radiation damage.

Prior to aberration correction, the loss of spatial resolution that would have resulted from lowering the operating energy was typically too large to permit atomic-resolution imaging of closely packed lattices. This situation has now changed: probe sizes of about 0.8 Å have been available for some time at 120 keV [15,16], and close to 1 Å probe size has become possible at 60 kV [17,18]. There are no near-neighbor atomic distances not involving hydrogen that are shorter than 1.2 Å, and 60 kV operation is therefore sufficient for most structural investigations. This is opening up a new/old sub-field of scanning transmission electron microscopy that we like to call “gentle STEM”. In this paper we illustrate what has now become possible in this field, and examine future possibilities.

We concentrate on model samples that have become available in the last few years: single wall nanotubes filled with fullerene

[☆] *Dedication:* This paper is dedicated to the memory of Professor Albert Crewe, who was the pioneer of both STEM and gentle STEM, and who recently passed away.

* Corresponding author.

E-mail addresses: krivanek@nion.com, krivanek.ondrej@gmail.com (O.L. Krivanek).

nanopods that contained one Er atom each, and single layer graphene with adatoms. These are particularly well suited to exploring detection limits and for comparing different operating modes. The conclusions drawn are expected to be applicable to a wide range of practical materials.

2. Resolution at low accelerating voltages

Evaluating the theoretical resolution has become more complicated in the aberration-corrected era. Spherical aberration no longer dominates, and the STEM resolution limit now typically comes from one or more of the following:

- (a) chromatic aberration
- (b) higher-order geometric aberrations
- (c) finite source size
- (d) finite size of the atoms
- (e) statistical noise in the images
- (f) instrumental instabilities.

Formulas for evaluating contributions (a)–(c) and (e) are given in [17]; here we just summarize the key factors.

In C_c -corrected STEMs with minimized or corrected fifth-order aberrations operating at 200 kV and below, the minimum probe size is mostly determined by the chromatic aberration, especially if the source has an energy spread >0.5 eV, as is the case for Schottky emitters. The probe size (full width at half-maximum, FWHM) due to chromatic aberration, in the limit of negligibly small source size (i.e., zero beam current), is given by [17]

$$d_{\text{chrom}} = 0.5(\lambda C_c \delta E / E_0)^{0.5}, \quad (1)$$

where λ is the electron wavelength, C_c the coefficient of chromatic aberration, δE the energy spread and E_0 the primary energy.

To image with a non-zero probe current, the size of the source projected onto the sample needs to be finite, and this is another important limit on the attainable resolution. The probe size for a finite beam current is given by

$$d_{\text{probe}} = (d_{\text{chrom}}^2 + d_{\text{source}}^2)^{0.5}, \quad (2)$$

where d_{source} is the projected (Gaussian) size of the source that gives the required probe current. Expressing the probe size in this way leads to a practical expression [17] that predicts the attainable probe size in a C_c -limited STEM of finite source brightness with about 10% accuracy

$$d_{\text{probe}}(\text{nm}) = 550 \left((1 + 7.3 \times 10^{17} I_p / B_r) C_c \delta E \right)^{0.5} / E^{*0.75}, \quad (3)$$

where I_p is the probe current (in amperes), B_r the normalized (reduced) beam brightness (in $\text{A}/(\text{m}^2 \text{sr V})$), C_c specified in mm and E^* the relativistically scaled primary energy (in eV). Using Eq. (3) with values appropriate for the Nion UltraSTEM [19] being used at 60 keV ($I_p = 50$ pA, $B_r = 10^8 \text{ A}/(\text{m}^2 \text{sr V})$, $C_c = 1.3$ mm, $\delta E = 0.35$ eV and $E^* = 63.7$ keV), gives $d_{\text{probe}} = 1.1$ Å. The illumination half-angle needed for this performance is 30 mr. Had the energy spread been only 0.3 eV, the gun brightness twice as high, or the beam current only 23 pA, the same formula would have given $d_{\text{probe}} = 1.0$ Å.

In order to attain a probe size not much bigger than 1 Å at 60 keV, a large number of parameters besides the gun brightness needs to be optimized. These include the stability of the microscope high voltage and current power supplies, accuracy of tuning, stabilities of the microscope's environment and of its sample stage, the efficiency of its detection system, and the quality of its vacuum. Not optimizing the parameters results in the resolution becoming worse, the data becoming noisier,

the atomic images becoming “squiggly”, or the sample being destroyed prematurely. The experimental results shown below demonstrate that this multi-faceted task can now be accomplished.

It is also important to keep in mind that the resolution obtained in an experimental image results from the convolution of the probe with “ideal” images of atoms that would be obtained with an infinitesimally small probe. Electrons that provide the annular dark field signal come from Rutherford scattering from the deep potential well surrounding the atomic nucleus. Aberration-free (and also diffraction limit-free and noise limit-free) ADF images of single atoms would show the potential well rather than the electron orbitals, and would therefore be very small—typically less than 0.3 Å in diameter [16]. Their size would be partly due to the finite dimension of the potential well and partly due to the thermal vibrations of the atomic nuclei. With a probe size greater than 1 Å, this contribution can usually be neglected and the image resolution taken as equal to the probe size. With probe sizes smaller than 1 Å, however, such an approximation may produce significant errors, and the expected ADF resolution then needs to be worked out either by summing the squares of the probe size and the atom size, or by a full calculation of the expected images. (The above representation is most useful for ultra-thin samples, in which electron channeling can be neglected. For thicker samples, the effects of probe channeling need to be taken into account too.)

The resolution in EELS maps is determined by probe size and statistical noise, plus the spatial spread of inner shell loss scattering. This scattering involves the interaction of the incident fast electrons with individual electrons orbiting an atom [20], and is more spread out than Rutherford scattering. The spreading is often called “delocalization”. It is well described by a formula for the diameter d_{50} of the area that contains 50% of the scattering events, derived by Egerton [21]:

$$d_{50} = 0.5\lambda / (\Delta E / E_0)^{3/4}, \quad (4)$$

where ΔE is the energy loss. The formula assumes that the EELS collection angle is large enough for the resolution not to be limited by the diffraction limit arising from a small range of scattering angles forming the inelastic image, and this can be a problem in energy-filtered TEM imaging. In the STEM, this condition is met when the illumination angle or the EELS collection angle is large. In the present work, both the angles were large enough not to be limiting at the 1 Å resolution level.

The expected resolution in EELS maps is then

$$d_{\text{EELS}} = (d_{\text{probe}}^2 + d_{50}^2)^{0.5}. \quad (5)$$

With a 1 Å probe at 50–100 keV primary energy, the delocalization is the major factor determining the attainable EELS mapping resolution for energy losses up to about 1 keV energy loss. For a 200 eV energy loss, d_{50} becomes about 3 Å, which makes energy losses <200 eV unsuitable for resolving closely spaced atoms. Eq. (4) also predicts a weak dependence of the delocalization on the primary energy: $d_{50} = 1.0$ Å for a 1 keV energy loss when operating at 100 kV, whereas $d_{50} = 0.9$ and 0.7 Å for the same loss when operating at 60 and 20 kV. This has not yet been verified experimentally, but it may be an additional advantage of operating at lowered primary energies.

3. EELS and ADF imaging

An advantage resulting from the poorer localization of the inelastic signal in real space is that it is better localized in reciprocal space. This makes the inner shell loss scattering largely forward-peaked, and the EELS signal can therefore be collected

efficiently by a spectrometer accepting half-angles slightly larger than the illumination half-angle, e.g. 40 mr for a 30 mr illumination half-angle. Rutherford scattering of the electrons used for ADF imaging makes their scattering angles much higher, and the ADF signal can typically be collected by an annular detector accepting half-angles from about 60 mr to about 200 mr. The 60 mr lower cut-off is lower than the 80–90 mr lower half-angle that we normally employ for high-angle annular dark field (HAADF) with heavier and thicker samples than the ones examined here. In order to distinguish between these two collection geometries, we call the detection mode mostly used in the present work “medium angle” ADF (=MAADF).

The integrated cross-sections of inner shell losses used for microanalysis vary widely, depending on the energy of the edge and also on the number of the electrons in the inner shell giving rise to the edge, from about 10^{-8} to 10^{-3} \AA^2 at 60 keV [20]. The inner shell loss signal also has a strong background arising from the extended tails of lower-energy edges and other inelastic scattering mechanisms. Moreover, as discussed above, it typically comes from an area that is larger than the atom. These factors mean that elemental EELS maps, obtained by processing spectra recorded at every pixel in a spectrum-image [22], need larger electron dose if they are to approach single-atom sensitivity. Doses of 10^5 to 10^7 electrons/ \AA^2 electrons should be sufficient for favorable edges such as the Er $N_{4,5}$ edge and the C K-edge used here. For weaker edges at higher energies such as Na to Al ($Z=11$ – 13) K-edges, Zn to Sr ($Z=30$ – 38) L-edges, and Lu to Bi ($Z=71$ – 83) M-edges, which are the only suitable edges for these elements, but occur at energy losses > 1 keV, the needed doses are of the order of 10^9 electrons/ \AA^2 and more. Such a dose is equivalent to resting an electron probe of 1 \AA^2 area and 100 pA probe current on the location of interest for 1.6 s. This is now achievable in modern STEMs, but not many samples are able to withstand such a dose without undergoing major changes, and not many experimentalists have the patience to wait for a 256×256 pixel elemental map acquired at 1.6 s per pixel (1.2 days). But for the more favorable elements, single atom EELS mapping is readily attainable, and has already been demonstrated (e.g. [23]).

The cross-sections for elastic scattering are typically significantly higher: about 10^{-3} \AA^2 for one carbon atom and about 10^{-1} \AA^2 for a uranium atom at 60 keV [20]. A 100 pA, 1 \AA^2 electron probe positioned over a single atom therefore produces some 10^6 detected scattering events per second from a carbon atom, and 10^8 events from a uranium one. Resting the probe over the atom for 1 ms can therefore produce a signal with 3% statistical noise for the C atom and 0.3% noise for the uranium one.

The integrated intensity of an individual atom in an ADF image is proportional to $Z^{1.5}$ to $Z^{1.8}$ [24]. For low- Z atoms such as C and N, the difference in image intensity for atomic species differing by $\Delta Z=1$ is therefore about 30%, whereas for heavy atoms such as Au it is about 2%. In the presence of 3% statistical noise for the light atoms and 0.3% for the heavy ones, it should therefore be possible to identify the chemical type of single non-overlapping atoms of any Z with good reliability, and with the superior resolution of ADF imaging. Such an ability has recently been demonstrated experimentally for the case of light atom impurities in monolayer BN [18].

4. Single-wall carbon nanotubes with Er-containing nanopods

Figs. 1–5 illustrate the above points with 60 keV images and spectra recorded from a single-wall carbon nanotube (SWCNT) filled with C_{82} fullerene nanopods that typically contained one erbium atom ($Z=68$) each [23]. The imaging was carried out in a Nion UltraSTEM [19] equipped with a Gatan Enfina energy loss spectrometer, at Oak Ridge National Laboratory. Fig. 1 shows a series of HAADF and MAADF STEM images, plus one bright field STEM image, which document the time evolution of a single nanotube over 53 min of observation at 60 keV. The nanotube was anchored to larger masses at both ends and did not move during the entire observation period. The vacuum around the sample was in the 10^{-9} Torr range, and there was no observable sample contamination nor beam-assisted etching due to residual gases such as water vapor. The microscope proved to be sufficiently stable for extended time-sequence observations and also for capturing short-term, rapid motion of individual atoms. The short-term stability requirement meant that the all-too-common image jumps which are caused by factors such as cooling water turbulence, acoustically transmitted vibrations, and power supply noise, and which typically give rise to “squiggly” images of atomic columns, had to be reduced to less than about 0.1 \AA peak-to-peak.

The per-pixel dwell time was 10 μs for all the images, but the pixel size varied: 0.28 \AA for images (a) and (b), 0.24 \AA for images (c) to (e) and 0.12 \AA per pixel for images (f) to (h). This meant that the electron dose was about 4×10^4 , 5×10^4 and 2×10^5 electrons per \AA^2 for the three types of images, respectively. No exact accounting was kept of the total dose administered to the area up to and including the image (h), which included recording all the images shown and also re-centering the observed area, re-focusing, etc., and also aberration tuning on adjacent amorphous carbon. It is likely that it was about 2×10^6 electrons per \AA^2 , with more than half of the dose coming during the time that the




Fig. 1. Time evolution over 53 min of observation of a single wall nanotube filled with Er-containing nanopods. (a)–(c) high angle annular dark field (HAADF) images (inner half-angle ~ 80 mr, outer half-angle ~ 200 mr); (d)–(g) medium angle annular dark field images (MAADF) (inner half-angle ~ 60 mr, outer half-angle ~ 200 mr); (h) BF image (collection half-angle ~ 5 mr).




Fig. 2. Higher-magnification prints of small areas of Fig. 1. (a) atoms in nanopod 4 in frame (f); (b) line profile through the atom pair in (a) taken between the arrows in the figure and integrated over 1.5 Å perpendicular to the profile; (c) line profile of the bottom atom in (a); (d) atoms at the boundary of nanopods 2 and 3 in frame (f); (e) atom in nanopod 10 in frame (g). The arrow in (e) points to the pixel in which the Er atom jumped back into its original position. Scale bars in (a) and (e) are 2 Å long, the magnification of (d) is the same as of (a).




Fig. 3. An ADF image (a) and a spectrum image acquired simultaneously. (b)–(d) energy slices through spectrum image; (e) erbium map; (f) carbon map.

small pixel (i.e., high-magnification) images (f) to (h) were recorded.

Image (a) was recorded about 5 min after the first observation of the area. It shows 10 intact nanopods, which have been numbered in the image. Each nanopod contained one erbium atom, with one exception: the Er atom that should have been in nanopod 1 was located outside its fullerene cage, between nanopods 1 and 2.

The time sequence documents many interesting points about atomic motion inside the nanotube. Remarkably, there was no visible damage to the nanotube itself in the 53 min of observation. The structure within the nanotube was not as stable: nanopods 8 and 9 fused together between frames (b) and (c) and then with nanopod 7 between frames (c) and (d). The Er atoms were even less stable. The atom between nanopods 1 and 2 stayed between these nanopods throughout the sequence, but kept on changing position between the left side of the nanotube and the right one. Other atoms were even more mobile: the Er atom in nanopod 8 moved up to the top of its nanopod before frame (c) and then to nanopod 5 in (d), while the other Er atoms in nanopods 4–6 stayed roughly where they were; and the top atom in nanopod 5 in frame (e) moved up to nanopod 4 before frame (f), in which it oscillated between 2 positions (see below).

The above observations are in good agreement with previous bright field EM studies of single atom-containing nanopods inside single wall nanotubes, which documented the motion of Er atoms within individual C_{92} nanopods [25], and also with studies documenting the migration of single Tb atoms out of C_{92} nanopods, as well as the merging of C_{92} nanopods and the migration of Er atoms between the merged nanopods [26]. The previous work showed that the motion of heavy atoms within a nanopod depended on the nanopod temperature more than on the electron dose, and that the merging of the nanopods did depend on the electron dose. It was observed at doses of the order of 10^5 electrons per \AA^2 with 120 kV electrons, with no major damage to the nanotube itself occurring at this dose.

The spatial and time resolutions were significantly improved in the present study. The improved time resolution has enabled us to distinguish between different types of heavy atom motions within the nanopods, as discussed below.

Many images of the Er atoms had a simple round shape, especially in the lower-dose frames (a)–(e). In the higher dose frames (f)–(h), however, their shapes became more irregular, due to their motion while the beam was over them. The beam scanned from left to right (line direction) and top to bottom (frame direction). Three types of Er atom motions were observed and are




Fig. 4. MAADF images recorded (a) just before and (b) just after the recording the spectrum of Fig. 5b, and then (c) just before and (d) just after the spectrum-image of Fig. 3. The arrow in (c) points to an atom-sized “pore” created in the side of the nanotube, a defect that subsequently healed during the acquisition of the spectrum image. The images were processed by the filtering procedure illustrated in Fig. 6.

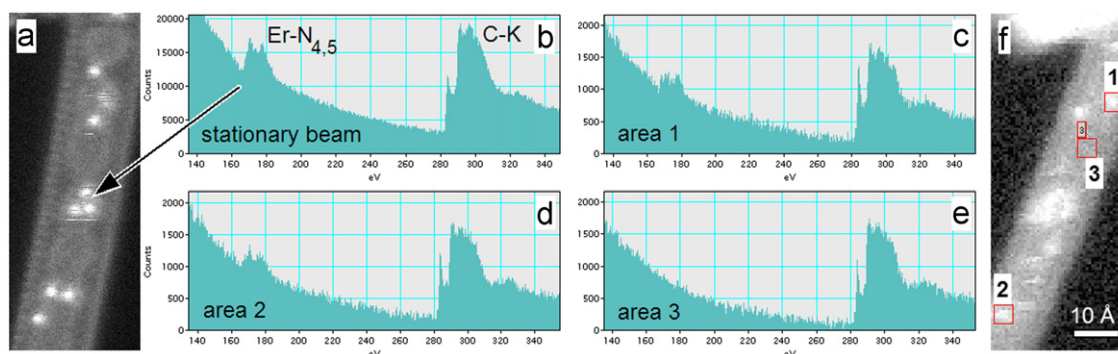


Fig. 5. Electron energy loss spectra obtained with a stationary probe and by summing over 7×7 pixel areas in the spectrum image of Fig. 3. (a) ADF image showing the location of the stationary probe; (b) spectrum obtained in a 1 s acquisition with the stationary probe; (c) to (e) spectra obtained by summing pixels in the spectrum image; (f) post-Er $N_{4,5}$ energy slice showing the location of the pixels corresponding to the SI spectra.

illustrated in Fig. 2. Fig. 2(a) shows the atom group in nanopod 4 in frame (f), printed with a high magnification. The image contains three distinct maxima, with two of the maxima about 1.5 \AA apart and just resolved, as shown in a profile in Fig. 2b. Following the atoms from frame to frame, as was done in the discussion of Fig. 1, shows that only 2 atoms should have been present in nanopod 4 in frame (f). Moreover, the images of the two closely separated atoms give only about 50% intensity (above the local background intensity of 120) relative to the atom immediately below them, whose profile is shown in Fig. 2c, and also about 50% of the intensity of the other atoms in the entire sequence. A very probable explanation is that the close images actually depict the same Er atom, oscillating between two positions and spending about the same time in each. The oscillation had to be fast enough so that many cycles were averaged over each pixel time, i.e. it had to be faster than about 1 MHz. This kind of motion is entirely possible for an atom with two close-by sites of similar binding energy, and it would also occur if a single heavy atom was attached to a backbone of lighter atoms which was oscillating.

Fig. 2(d) shows the pair of Er atoms that were imaged at the bottom of nanopod 2 and the top of nanopod 3 in frame (f), also

printed with a high magnification. The bottom atom looks much like the bottom atom in Fig. 2(a), but the top one is very different. Its image has bright and dark streaks, and is about 4 \AA wide. The individual streaks are typically about 1.5 \AA wide and just one pixel high. The streakiness stops for about the bottom third of the image of the atom, which looks similar to other single-atom images in the sequence. We conclude that this atom was jumping between several sites in a small area, with a jump frequency of 100–1000 Hz, i.e. stationary on the time scale of a few pixels, but mobile on the time scale of each scan line, 1024 pixels long and thus taking 10 ms each. We also conclude that it stopped jumping for the bottom third of its image.

Fig. 2e shows a highly magnified print of the Er atom in nanopod 10 imaged in frame (g). This atom was mostly stationary in the preceding frames, but jumped out of the imaged area before the next frame (which is shown in this paper as Fig. 4(a)). In frame (f) it is on the left side of its nanopod, but in frame (g) it is mostly on the right side. However, about half-way down its image, it jumped away, probably just as the probe was approaching it, since there is no extra intensity indicative of an Er atom for this whole scan line going across the nanopod. In the next scan line, it was displaced 3.5 \AA to the left, stayed there while the beam was

on it, and then jumped back into its original position while the beam was scanning over the original location, just in time for increasing the brightness of the arrowed pixel. In other words, it jumped twice while the beam was in its general vicinity *but not on it*, and the second time it jumped right under the beam.

Given the high frequency of atomic motions under the electron beam documented in Figs. 1 and 2, it is clear that increasing the acquisition time per atom significantly is not likely to lead to data in which most Er atoms stay still. Unfortunately, acquiring EELS spectra with high signal-to-noise ratio (SNR) calls for the electron dose to be increased typically $10\text{--}1000\times$. The acquisition can be done simply with a stationary probe, or by spectrum-imaging (SI) [22]. We have done both on the imaged area. The nanopod and Er atom structure was heavily damaged before the start of the EELS observations and was known from the images in the sequence of Fig. 1. The intent of the EELS experiments was to check the resolution in the EELS maps and to explore different strategies of EELS spectrum acquisition in the presence of radiation damage, and a damaged structure served these purposes just as well as a pristine one.

Fig. 3 shows information extracted from a 50 by 120 pixel subarea of an 83 by 202 pixel spectrum-image whose energy range spanned 80 to 700 eV at 0.5 eV/channel, and whose pixel dwell time was 9 ms. The spectrum-image was recorded just after the image of Fig. 4(c). Because of the smaller area of the spectrum image and its long per-pixel time, the electron exposure was much higher than for the ADF images: about 10^7 electrons/Å².

Images of Fig. 3(b)–(d) show 20 eV sums of energy slices taken from just before the Er $N_{4,5}$ edge at 168 eV (b), on top of the edge (c) and just after the C K-edge at 285 eV (d). Fig. 3(e) shows an erbium elemental map obtained by simply scaling and subtracting the pre-edge Er image from the post-edge one, so that no trace of the nanotube/nanopod structure remained in the map. Such a procedure is similar to a jump ratio image, but it avoids division by zero in the vacuum next to the sample. In a thin binary material, as was the case here, a suitably scaled pre-edge background can model the contribution element B makes to the intensity collected with an energy window over an edge due to element A, and vice versa. The method did not involve fitting and extrapolating the pre-edge background, and therefore gave a higher signal-to-noise ratio in the Er map than any other processing methods that we tried. The carbon map (Fig. 3(f)) was similarly obtained by scaling and subtracting a pre-C energy image from the post-C one, so that no trace of Er remained.

The displayed part of the spectrum-image starts at nanopod “0”—one nanopod higher than the top of the images shown in Figs. 1 and 4. ADF images recorded under similar conditions to Fig. 1(g) just before and just after the spectrum image are shown in Fig. 4. They make it clear just how much the Er distribution within the nanotube was changed by the spectrum imaging, and also that the SI recording resulted in much damage to the nanopods. The 10° difference in the nanotube orientation between Figs. 3 and 4 was due to the sample drifting about 2 nm to the left during the 3-min acquisition of the spectrum image, and no drift correction being applied. (The data was recorded before the microscope room was optimized, and the room temperature was cycling with a magnitude of about 2 °C and a periodicity of about 1 h.)

Determining where the atoms were during the spectrum image acquisition cannot be done with certainty: the atoms clearly moved during the acquisition itself. An interesting group of 8 Er atoms can be seen just below the center of the ADF image of Fig. 3a, about where the remnants of nanopod 4 were in the pre-image. No such group appears in the “before” or “after” ADF images of Fig. 4, and it is likely that it was an artifact that arose because some of the atoms jumped (possibly due to the nanopods

moving suddenly) and were imaged again. The group provided a convenient test object: in the ADF image, real atomic separations (i.e., separations not due to atomic jumps, which means separations that showed no abrupt intensity changes along the line connecting the two atoms) of 2.5 and 3.7 Å were resolved in the line scan direction, between atoms on the left side of the nanotube and in the center, and between the center atoms and atoms on the right side, respectively. In the simultaneously recorded spectrum-image and the resultant Er elemental map, only the larger separation was resolved. The resolution predicted by Eqs. (4) and (5) for a 1.4 Å probe and a 175 eV energy loss at 60 keV primary energy is 3.6 Å, in good agreement with this observation. The predicted resolution for the carbon K-edge with its maximum at 300 eV is 2.6 Å. The C map shows the inner nanotube that resulted from the fusion of the nanopods resolved from the outer one, even though their separation is only about 3.4 Å in several places.

Atomic streaks show the same patterns and widths in both the ADF image and in the erbium map (especially in the bottom 30% of Fig. 3(a) and (e)). These are not indicative of the resolution of the EELS map: they resulted from atoms moving around and being present for some pixels but not for others. But they demonstrate the high signal-to-noise ratio in the Er EELS map, in which it is possible to tell whether an Er atom was present or not at *each individual pixel*.

Fig. 4 shows ADF images of the nanotube recorded just before and just after the spectrum of Fig. 5(b) was obtained (Fig. 4(a) and (b)), and just before and just after the spectrum image of Fig. 3 (Fig. 4(c) and (d)). The first three images were acquired with pixel size of 0.12 Å, and the last one was acquired with pixel size of 0.4 Å, i.e. with about 10% of the electron dose of the other 3 images. Its statistical precision was therefore considerably inferior. About 5 spectra aiming to capture the EELS signature of single Er atoms with a stationary probe were recorded between the images (b) and (c), and the damage due to them produced the large changes between the two images.

The images of Fig. 4 were processed using the probe tail-subtracting filter explained in Fig. 6. They are displayed slightly non-linearly, with enhanced contrast for the nanotubes relative to the Er atoms, so that nanotube details can be seen without the Er atoms becoming very oversaturated. The structure of the nanotube and of the damaged nanopods inside is partly resolved in the high-dose images (a) to (c). Image (c) shows a very small hole, which appeared in the side of the nanotube, and which annealed away during the spectrum imaging. The nanotube wall appears to curve inward by about 0.7 Å right next to the hole. The structure looks very much like a small pore, 2.5 Å wide. The distance between the pairs of atoms on the opposite sides of each carbon hexagon is 2.46 Å, and in the $\langle 10\bar{1}0 \rangle$ projection, there is a pair of atoms in the next hexagon that projects into this gap. The pore probably arose because a carbon atom pair aligned with the $\langle 10\bar{1}0 \rangle$ direction went missing. Being able to resolve a pore created by 2 missing carbon atoms, in profile, is a nice illustration of the sensitivity of aberration-corrected ADF imaging.

Fig. 5 compares a spectrum obtained with a stationary probe resting for 1 s over nanopod 4 with its 3 Er atoms (Fig. 4(a) and (b)) to three energy loss spectra obtained by summation over 7×7 pixel areas in the spectrum-image, as shown in Fig. 5(f). Two of the summed spectra came from areas containing one Er atom each, and the third one from a nearby area containing only carbon. The stationary probe spectrum was obtained with a total dose of 3×10^8 electrons i.e. about 2×10^8 electrons per Å². Despite the very high dose given to nanopod 4, the before–after comparison shown in Fig. 4(a) and (b) makes it clear that the 3 Er atoms were present in the nanopod for the entire acquisition period.




Fig. 6. ADF images of single and multiple layer graphene. (a) as recorded; (b) FFT of the image (only the central area is shown); (c) profile of a smoothing plus de-fogging filter; (d) the image FFT multiplied by the filter; (e) smoothed and de-fogged image. The insert in (e) shows line profile A-A' that starts in the vacuum just outside the monolayer area and continues through the monolayer into a double layer area. The long white arrow in (e) points to a single graphene sheet that curled over at its edge, and the short white arrow points to a streak that may have been a carbon atom dangling off the edge that run away while the beam was scanning in its area.

Interestingly, there were more major changes in the nanotube outside the probed nanopod 4 than within it. The Er $N_{4,5}$ edge resulting from the 3 atoms is easy to see, and shows the pre-edge Fano dip [27].

The stationary probe spectrum had a better signal-to-noise ratio than the spectra extracted from the spectrum-image, which nevertheless showed the fine structure of the Er and C edges, and hinted at the Fano dip. The extracted spectra correspond to an acquisition time of $49 \times 9 \text{ ms} = 0.44 \text{ s}$. Halving the acquisition time should have increased the fractional shot noise $1.4 \times$, and reducing the number of contributing atoms $3 \times$ should have reduced the signal $3 \times$, for an overall decrease of the signal-to-noise ratios of the SI spectra by about $4 \times$ relative to the fixed-beam one. This is approximately borne out in practice. A noteworthy aspect of the summed spectra is that their absolute channel-to-channel variation does not reduce much in the least intense parts of the spectra (e.g., just in front of the C K-edge in Fig. 5(e)), even though a reduction would be expected for purely statistical noise. This probably occurred because the EELS CCD fixed pattern subtraction was not perfect, and adding 49 fixed patterns produced a dominant noise term for spectral regions of low intensity. But the problem is not overwhelming, and the spectra demonstrate that it is now possible to detect single atoms and even study the fine structure of their EELS edges, with a spatial resolution comparable to interatomic distances, provided of course that the atoms can withstand the needed electron dose.

5. Graphene edge with adatoms

Graphene can be thought of as an unfolded single wall nanotube, and imaging it is therefore closely related to imaging nanotubes. We examined samples prepared by liquid phase

exfoliation of bulk graphite powders in N-methyl-2-pyrrolidone (NMP), which gave monolayer dispersions with good yield [28]. Full details of the sample preparation are given in [18].

The exfoliation produced graphite flakes with small monolayer regions in various locations at or near the flake edges. The size of the regions varied. Smaller monolayer areas of around 10 by 10 carbon hexagons, surrounded by thicker regions, were typically more stable under the beam and were therefore more suitable for observation. Fig. 6(a) shows a part of a $1 \text{ k} \times 1 \text{ k}$ MAADF image of one such area. The pixel size was 0.12 \AA , the per-pixel dwell time $64 \mu\text{s}$, and the rest of the parameters were as for the nanotube MAADF images shown in Fig. 1.

A noteworthy feature of MAADF images from these kinds of samples was that the intensity in the center of the holes in the monolayer hexagons was typically 50–70% of the intensity of the carbon atoms defining the hexagons, rather than 0%, as is would have been with a sharp and tail-less probe. This was the case even though the carbon-carbon distance of 1.42 \AA was typically well resolved, which proved that the full-width of the probe at half-maximum (FWHM) was less than 1.42 \AA , i.e. that the probe half-width was less than 0.71 \AA . The distance from the hexagon atoms to the center of the hexagon is the same as the carbon bond length: 1.42 \AA . This shows that the probe had a relatively narrow central maximum and a more extended tail, which contributed around 10% of each carbon atom's peak intensity to the centers of the carbon hexagons. We note that the effects of the probe tail were well visible because of the quantitative nature of ADF imaging, in which the intensity in the vacuum next to the sample goes to zero, and this gives a baseline to which the intensity in the center of the hexagons can be compared. In bright field (BF) phase contrast imaging, by comparison, the intensity in the vacuum is 1. A tail in a BF image of an atom results in a slight change in the image pattern and a reduction of the overall phase contrast. But there is no readily visible change in the DC level of any image

area, and the contrast is easily boosted back up, rendering the reduction invisible.

A second feature of the unprocessed images was that they showed a lot of statistical (shot) noise. This was mostly due to the fact that the image of each atom was spread over an area of 10×10 pixels or more, i.e. the image data was oversampled about $5 \times$ relative to the resolution attained. The signal per pixel was therefore quite small, and the shot noise a large percentage of the signal. Another way to view this is to realize that the useful image information is concentrated in the central area of the image's fast Fourier transform (FFT), which amounts to only 4% of the whole FFT area in a $5 \times$ oversampled image, whereas the shot noise is spread evenly over the FFT area. This means that $5 \times$ oversampling contributes $25 \times$ more shot noise to an image than if the sampling was such that the highest spatial frequency captured in the image was equal to the Nyquist limit.

Despite the above, oversampled images are greatly preferable, because they contain extra information about the motion of the atoms, lattice distortions, etc. The oversampling does not carry any noise penalty if the ADF detector is sufficiently noise-free, so that dividing the atomic image into many small pixels does not add extra instrumental noise due to the many additional readouts. This was the case with our experimental set-up. The added shot-noise is easy to smooth out: one simply filters out spatial frequencies higher than those containing the useful information, after first checking the raw images for information such as streaks indicative of atomic motions as for instance analyzed in Fig. 2.

The filtering can be modified so that at the same time as the image is smoothed, the probe tails are largely removed. The second step removes the extra intensity that the probe tails spread from each atom to its neighbors and also into holes and beyond the edge of the sample. It yields crisper, de-fogged images of essentially the same resolution. This was done with the image of Fig. 6(a), and the result is shown in Fig. 6(e). Both these images were displayed with a slightly non-linear contrast scale designed to show the monolayer structure more clearly while decreasing the saturation of the images of heavier atoms.

Fig. 6(b)–(d) illustrates the tail removal methodology. The shape of the probe was not known exactly, and it probably varied a little from image to image and especially from one autotuning operation to the next. We therefore chose a particularly simple filtering procedure, in which the filter consisted of two Gaussians: a positive one which was relatively wide and which gradually cut-off spatial frequencies beyond those actually transferred by the microscope's optics, and a negative one which was narrower and which essentially added a weak and spread-out negative skirt to the image of every atom. The width of the positive Gaussian was chosen to match the largest useful spatial frequencies visible in the Fourier transform, and the width and magnitude of the negative Gaussian were chosen so that intensity profiles of atoms at the edge of the sample went smoothly and rapidly to zero, without undershooting (positive total tail) or overshooting (negative tail). The procedure thus amounted to low-pass filtering to eliminate noise at spatial frequencies higher than those actually captured in the image, and then subtracting an even more smoothed version of the image from the original, so as to eliminate the "fog" that probe tails create in images. Because the filtering was rotationally symmetric and had no sharp cut-offs that might have caused "ringing" in the processed image, the probability of creating misleading artifacts out of random noise was small.

Fig. 6(b) shows the central portion of the FFT of the original $1 \text{ k} \times 1 \text{ k}$ image, a part of which is shown in Fig. 6(a). The FFT was then multiplied by the rotationally symmetric filter whose profile is shown in Fig. 6(c). This produced the FFT shown in Fig. 6(d), which was inverse Fourier-transformed to give the filtered image, shown in Fig. 6(e).

The insert in Fig. 6(e) shows the line profile between the black arrows marked A and A'. It starts in the vacuum next to the adatom at the edge of the graphene sheet, crosses two monolayer area carbon dumbbells, and then goes into a double-layer area at the upper right of the image, in which the profile intensity is simply twice the intensity of the single layer. The atoms in the double layer area are well aligned in projection, i.e. the projected stacking is A–A rather than the A–B stacking of bulk graphite, in which atoms in the second layer normally lie over the centers of the carbon hexagons in the first layer. This is probably due to the thicker and partly amorphous layers bordering the small double layer patch positioning the second layer in this way. However, it needs to be noted that because the interlayer spacing in graphite is 3.4 \AA whereas the in-layer interatomic distance is only 1.42 \AA , a mistilt by 23° of a double layer relative to the beam direction changes the projection from A–B to A–A or vice versa. In order to be able to ascertain the stacking type from the projected image, one therefore needs to know the angle between the layer normal and the beam direction to better than about 10° . This was not the case here.

Two other interesting features that became clear in the filtered image have been marked by white arrows. The long arrow points to monolayer graphene which curled up so that its edge was parallel to the beam, i.e. it simulated one quarter of a conical nanotube, whose atomic structure is readily distinguishable. The short arrow points to a streak that probably came from a single carbon atom dangling off the graphene edge, which ran away while the beam was scanning in its area. (Such a dangling carbon atom was in fact seen in the image recorded just before Fig. 6 and shown in Fig. 7(a)). Other interesting features in the image included an impurity atom at the edge of the central monolayer, impurity atoms located away from the edge, and the detailed atomic arrangements at the monolayer's edge. These are discussed below.

Fig. 7 shows the central portion of the image of Fig. 6 at higher magnification (Fig. 7(b)), and the same part of the sample imaged immediately before (Fig. 7(a)). Both the images were processed by the de-fogging filter and displayed slightly non-linearly. The impurity atom at the graphene edge stayed in its place, which was only about 50% probable, as could be seen by observing, with the same electron dose, the mobility of impurity edge atoms at other locations in the same sample. The structure of the edge itself had undergone major modifications. In the left image, a variety of atomic arrangements is seen at the edge: two five-fold rings (indicated by single arrows), a single dangling carbon atom (indicated by a double arrow), a distorted "armchair" (in which a complete carbon hexagon sits right at the sample's edge) just above the bottom five-fold ring, and some atoms that were moving and contributed only streaks to the image. In the right image, the edge terminates in 4 regular armchairs. The rearrangement required the addition of just one carbon atom below the impurity atom and the removal of one carbon atom above the impurity atom. The armchair-terminated edge is similar to graphene edges imaged by bright field phase contrast TEM [29], but the observations of a 5-fold ring at graphene's edge and of a single dangling carbon atom appear to be new.

Many carbon hexagons are seen to be somewhat distorted, and the distortion of the same hexagon is typically different in the two images. There were three principal causes for the distortions: (a) statistical noise, which randomly enhanced different parts of the spread-out atomic images, and thus caused the smoothed images of individual atoms to shift randomly from frame to frame, (b) sample movement, which translated into the displacement of some parts of the image but not others, and (c) real distortions present in the carbon sheet, plus apparent distortions caused by the fact that the sheet was not aligned perpendicular to the beam




Fig. 7. (a) MAADF image of the edge of monolayer graphene; (b) same area imaged about 2 min later. The single arrows in (a) point to 5-fold rings at the graphene's edge, the double arrow to a single atom of carbon dangling off the graphene edge. A-A' profile through the impurity atom at the edge is shown as an insert in image (b).

and was probably also slightly buckled. The best way to separate the random distortions from the real ones is to image the same area in a sequence of images, and the two images shown here indicate that most of the distortions in the present case were caused by statistical noise. In stable samples the statistical noise grows smaller at larger electron doses, and our practical experience [18] indicates that they can be kept as small as about 0.1 Å if the dose is increased about $4 \times$ relative to the one used here.

There were several impurity adatoms, which gave much stronger contrast than the carbon atoms. Adatoms on the right side of the images were located on top of the carbon sheet and were moving frequently, and this made their analysis difficult. The single adatom at the graphene's edge was stationary, and formed the apex of a 5-fold ring, with larger separation from its neighbors than the apex atom in the carbon-only 5-fold ring seen just above the adatom in Fig. 7(a). A profile through the adatom (insert in Fig. 7(b)) shows its intensity to be $3.6 \times$ larger than that of the C images. Using the $I=Z^{1.64}$ dependence of the atomic intensity I on the atomic number Z that we derived previously [18] from experimental images of B, C, N and O atoms imaged under essentially the same conditions as here, gave $Z_{\text{impurity}}=6 \times 3.6^{1/1.64}=13.1$, and we therefore tentatively identified the atom as aluminum. However, the extrapolation to $Z=13$ based on experimental data obtained for $Z=5-8$ is a stretch, and it is therefore possible that the impurity atom was Mg or Si, or even Na or P. EELS was tried on the atom and similar intensity impurity atoms in the vicinity, but it was not conclusive: the atoms were not strongly attached and tended to run away under the beam.

6. Discussion

The observations shown here illustrate several important points. First, the signal-to-noise ratio (SNR) in ADF images of

individual atoms as light as carbon is now high enough to make it possible to tell whether the atom was in place for each pixel in an atomic image spanning an area consisting of 100 or more pixels. This is allowing atomic motions to be studied on a time scale corresponding to the per-pixel dwell time, in the present case 10 μs (for single Er atoms) and 64 μs (for single carbon atoms).

Second, a detailed exploration of atomic motions was possible here because the microscope itself was stable enough so that images of well anchored parts of the samples did not “jitter” by more than about 0.1 Å peak-to-peak. Many different requirements had to be met for this level of short-term stability to be achieved. The results presented here show that the effort was very worthwhile.

Third, heavy atoms in nanotubes and $Z \sim 13$ atoms on graphene were seen to be mobile even when the beam was not directly over them. The observation that the atoms were more stationary when the beam dose was smaller suggests that the beam had to be in the general vicinity in order for the atoms to move, but the present study was not detailed enough to settle this point. Further studies in which individual nanopods containing heavy atoms are imaged at a rapid rate (e.g., 100 frames per second is attainable for a 32×32 pixel image with 10 μs per pixel) at different sample temperatures should prove illuminating.

Fourth, nanotubes and the inside parts of graphene sheets show little radiation damage at irradiation doses up to about 10^8 electrons/Å² when observed at 60 kV, in agreement with general considerations of knock-on displacement damage [30]. But fullerenes and the edges of graphene sheets are more fragile. The binding energy for edge atoms is expected to be lower than for atoms within the sheet, and the edge's fragility is therefore expected. The binding energies of carbon atoms in a nanotube and in a fullerene nanopod are similar, and the 60 kV beam's ability to damage one but not the other is therefore surprising at first sight. However, nanopods may damage due to inelastic scattering followed by an electronic damage mechanism [31], which can

occur at 60 kV and lower primary voltages, whereas monolayer sheets and nanotubes seem to be impervious to this type of damage. Another damage mechanism that may operate at primary voltages of 60 kV and lower is hydrogen-mediated knock-on displacement [32], in which hydrogen atoms act as an impedance-matching medium able to absorb a higher energy from the fast electrons and then transfer it to neighboring heavier atoms. This mechanism needs hydrogen atoms on or within the sample, and its cross-section is likely to be very small. It is therefore likely to be rare even at the high doses used in the present work.

Fifth, because the Er atoms in the nanopods and the impurity atoms on graphene were quite mobile, spurious findings would have resulted if spectrum imaging results were taken at face value. E.g., one of the nanopods would have been thought to contain eight atoms, when in fact no nanopod contained more than three. This artifact provides a nice illustration of the dilemma presented by radiation damage: one has a choice between using a small dose and thereby acquiring noisy and thus possibly misleading data from the correct structure, or using a high dose and acquiring relatively noise-free data from a possibly misleading structure. In the future, we intend to test operation at primary energies lower than 60 keV, which may allow atoms of these types to be analyzed without disturbing them excessively.

Sixth, the images presented here became easier to interpret when their statistical noise occurring at high spatial frequencies was filtered out, and the probe tail contributions were suppressed by a “de-fogging” procedure that consisted of subtracting a low-pass filtered version of the image. Some information was lost during the filtering, such as telltale signs showing that atoms moved during the exposure, and it was therefore best to look at both the raw data and the processed images. The de-fogging procedure largely avoided producing misleading artifacts, but as with all filtering operations, the danger of digging up false “signals” from statistical noise was certainly present. It is therefore important to use good judgment when performing the filtering, and to avoid enhancing spatial frequencies for which the true signal is weaker than the statistical noise. Principal component analysis [33] can quantify the probability of the extracted structures being real, and we intend to use it in the future. The best general precaution is nevertheless the standard one when operating in the presence of statistical noise: repeat the experiment several times, and compare the results.

Seventh, ADF imaging is very useful for indicating the chemical type of individual atoms, especially when large electron doses cannot be used, but it is not as element-specific as spectroscopic techniques. We have recently used ADF imaging to identify, with better than 99% confidence, all the atoms except one (which was identified with a 94% confidence level), in an 82-atom area of monolayer BN that contained many atomic substitutions [18]. In the present work, ADF imaging was able to suggest the atomic type of a particularly noticeable edge adatom, but it could not determine it with certainty. We nevertheless note that the information collected about the atom was sufficient to determine its type in principle, as explained below.

We expect that future theoretical and experimental work will improve the attainable precision of the ADF-based atomic identification. Intensity ratios for images of different types of atoms, recorded with identical illumination and signal collection geometries, as well as identical image processing, are almost certain to be very reproducible from experimental session to session. Determining these ratios experimentally for a particular set-up of a particular microscope should result in experimental “Z-factors” similar in their nature to the well-known K-factors that helped quantify X-ray microanalysis [34]. Had the Z-factors been previously worked out for our experimental set-up for the

candidate atoms (Mg to Si), we probably would have been able to determine the edge adatom unambiguously, without needing the larger dose necessary for an EELS-based identification.

The ADF-based atomic identification can of course only work for images of isolated atoms on a known background and preferably no background at all, as when the atom is at the edge or is incorporated into a one-atom thick sheet. However, when there are more atoms to analyze, EELS may be able to obtain the required information at a lower dose that does not make the atoms run away. ADF and EELS (or energy-dispersive X-ray spectroscopy) are therefore expected to play complimentary roles in the analysis of ultra-thin samples: the spectroscopy technique can indicate what atoms are present in general, at a spatial resolution that is not high enough to resolve atoms less than 2–3 Å apart, and ADF can resolve and identify individual atoms, by comparing their intensity to known atomic types imaged elsewhere on the sample.

7. Conclusion

Aberration correction has allowed STEM imaging and analysis to progress remarkably in the last few years. Being able to reach atomic resolution at operating energies that avoid knock-on displacement damage in light atom materials has been a particularly important development. It is now possible to image and identify single light atoms in samples that until recently either damaged heavily (if imaged at higher energies) or gave little indication of their precise atomic structure (if imaged at lower energies).

Lowering the primary energy even further may result in additional advantages, such as a reduced mobility of the adatoms. Maintaining close to 1 Å resolution at energies significantly lower than 60 kV will require additional instrumental developments, such as STEM chromatic aberration correction, or the development of more mono-energetic electron sources, possibly using monochromation. The resultant “gentle STEM” approach has a good chance of reaching the goal that many of us have had for electron microscopy for some time: being able to determine the structure of general, non-periodic matter such as complicated molecules very simply and directly: by imaging them atom-by-atom. It will certainly be an exciting area to work in and to watch.

Acknowledgements

We are grateful to our colleagues at Nion, especially George Corbin, James Woodruff and Dr. Zoltan Szilagy for their part in the design and construction of the microscope used here, and to Dr. S.J. Pennycook for the use of the microscope and for initiating the microscope project. Research at Oak Ridge National Laboratory (MFC and TJP) was sponsored by the Division of Materials Sciences and Engineering of the US Department of Energy. TJP was also supported by the US Department of Energy Grant DE-FG02-09ER46554.

References

- [1] A.V. Crewe, J. Wall, L.M. Welter, *J. Appl. Phys.* 39 (1968) 5861.
- [2] A.V. Crewe, J. Wall, *J. Mol. Biol.* 48 (1970) 375.
- [3] A.V. Crewe, J. Wall, J. Langmore, *Science* 168 (1970) 1338.
- [4] M. Isaacson, D. Johnson, *Ultramicroscopy* 1 (1975) 33.
- [5] A.V. Crewe, *Microsc. Microanal.* 10 (2004) 414.
- [6] A.V. Crewe, Cold field emission and the scanning transmission electron microscope, in: P.W. Hawkes (Ed.), *Advances in Imaging & Electron Physics*, vol. 159, Elsevier, Amsterdam, 2009, p. 1.

- [7] I.M.R. Wardell, P.E. Bovey, Cold field emission and the scanning transmission electron microscope, in: P.W. Hawkes (Ed.), *Advances in Imaging & Electron Physics*, vol. 159, Elsevier, Amsterdam, 2009, p. 221.
- [8] H.S. von Harrach, Cold field emission and the scanning transmission electron microscope, in: P.W. Hawkes (Ed.), *Advances in Imaging & Electron Physics*, vol. 159, Elsevier, Amsterdam, 2009, p. 287.
- [9] H. Inada, et al., Cold field emission and the scanning transmission electron microscope, in: P.W. Hawkes (Ed.), *Advances in Imaging & Electron Physics*, vol. 159, Elsevier, Amsterdam, 2009, p. 123.
- [10] N.D. Browning, M.F. Chisholm, S.J. Pennycook, *Nature* 366 (1993) 143.
- [11] D.A. Muller, Y. Tzou, R. Raj, J. Silcox, *Nature* 366 (1993) 725.
- [12] P.E. Batson, *Nature* 366 (1993) 727.
- [13] J.C. Meyer, A. Chuvilin, G. Algara-Siller, J. Biskupek, U. Kaiser, *Nano Lett.* 9 (2009) 2683.
- [14] A. Zobelli, A. Gloter, C.P. Ewels, G. Seifert, C. Colliex, *Phys. Rev. B* 75 (2007) 245402.
- [15] P.E. Batson, N. Dellby, O.L. Krivanek, *Nature* 418 (2002) 617.
- [16] P.E. Batson, *Ultramicroscopy* 106 (2006) 1104.
- [17] O.L. Krivanek, et al., Aberration-corrected electron microscopy, in: P.W. Hawkes (Ed.), *Advances in Imaging and Electron Physics*, vol. 153, Elsevier, Amsterdam, 2008, p. 121.
- [18] O.L. Krivanek et al., *Nature*, in press, doi:10.1038/nature08879.
- [19] O.L. Krivanek, et al., *Ultramicroscopy* 108 (2008) 179.
- [20] R.F. Egerton, in: *EELS in the Electron Microscope*, 2nd Edn., Plenum Press, New York, 1996.
- [21] R.F. Egerton, *Ultramicroscopy* 107 (2007) 575.
- [22] C. Jeanguillaume, C. Colliex, *Ultramicroscopy* 78 (1989) 252.
- [23] K. Suenaga, et al., *Nature Chemistry* 1 (2009) 415.
- [24] P. Hartel, H. Rose, C. Dignes, *Ultramicroscopy* 63 (1996) 93.
- [25] K. Suenaga, et al., *Nano Lett.* 3 (2003) 1395.
- [26] K. Urita, et al., *Nano Lett.* 4 (2004) 2451.
- [27] P. Rez, *Microsc. Microanal.* 7 (2001) 356.
- [28] Y. Hernandez, et al., *Nature Nanotechnol.* 3 (2008) 563.
- [29] Ç.O. Girit, et al., *Science* 323 (2009) 1705.
- [30] R.F. Egerton, R. McLeod, F. Wang, M. Malac, this issue (2010).
- [31] R.F. Egerton, M. Takeuchi, *Appl. Phys. Lett.* 75 (1999) 1884.
- [32] G.M. Bond, I.M. Robertson, F.M. Zeides, H.K. Birnbaum, *Phil. Mag. A* 55 (1987) 669.
- [33] I.T. Jolliffe, in: *Principal Component Analysis*, 2nd Edn., Springer, New York, 2002.
- [34] G. Cliff, G.W. Lorimer, *J. Microsc.* 103 (1975) 203.