

Automated electron microscope tomography using robust prediction of specimen movements

David N. Mastronarde *

Boulder Laboratory for Three-Dimensional Electron Microscopy of Cells, Department of Molecular, Cellular, and Developmental Biology, University of Colorado, Boulder, CO 80309, USA

Received 5 April 2005; received in revised form 14 July 2005; accepted 20 July 2005
Available online 24 August 2005

Abstract

A new method was developed to acquire images automatically at a series of specimen tilts, as required for tomographic reconstruction. The method uses changes in specimen position at previous tilt angles to predict the position at the current tilt angle. Actual measurement of the position or focus is skipped if the statistical error of the prediction is low enough. This method allows a tilt series to be acquired rapidly when conditions are good but falls back toward the traditional approach of taking focusing and tracking images when necessary. The method has been implemented in a program, SerialEM, that provides an efficient environment for data acquisition. This program includes control of an energy filter as well as a low-dose imaging mode, in which tracking and focusing occur away from the area of interest. The program can automatically acquire a montage of overlapping frames, allowing tomography of areas larger than the field of the CCD camera. It also includes tools for navigating between specimen positions and finding regions of interest.

© 2005 Elsevier Inc. All rights reserved.

Keywords: Tomography; Automated data collection; Transmission electron microscopy; Cryo-microscopy

1. Introduction

Electron tomography has become increasingly important as a technique for examining cellular structures and macromolecules in three dimensions (Baumeister, 2002; Leapman, 2004; McIntosh et al., 2005; Steven and Aebi, 2003). The method involves taking a series of electron micrographs while tilting the specimen over a range of angles, typically at 1–2° intervals up to ±60–70°. These projections of the specimen are aligned and a three-dimensional reconstruction is computed by methods such as weighted backprojection (Frank, 1992). The recent growth of tomography has been fostered by continual improvements in the instrumentation, computers, and software needed for these tasks. In particular, much

effort has been expended to automate the acquisition of the series of tilted projections, a task that has been facilitated by improved interfaces for microscope control. Acquisition requires a sequence of operations because of imperfections in the goniometers and specimen holders now available: after tilting to a new angle, the features of interest must be recentered in the field of view, and any change in vertical (*Z*) height must be compensated by refocusing, before the final image is acquired. When done manually, these operations are slow and tedious at best, and because the sample is exposed to the beam during these steps, they are particularly difficult for beam-sensitive samples. To automate this process, an initial tracking image is taken on a CCD camera after tilting to a new angle; the image is cross-correlated with a comparable image from the previous tilt, and electronic image shift is used to recenter the specimen. Defocus is determined by taking two images with the beam tilted in opposite directions; the

* Fax: +1 303 492 7744.

E-mail address: mast@colorado.edu.

displacement between two such images is proportional to the defocus, so the objective lens can then be adjusted to bring the specimen back to the desired focus (Koster and de Ruijter, 1992; Koster et al., 1987).

The first generation of automated tomography software took one or two tracking images at each tilt, and also measured defocus once or twice on each tilt. Some software followed a strict low-dose approach by taking tracking and focusing pictures on areas displaced from the region of interest, which is essential for working with frozen-hydrated material (Dierksen et al., 1992, 1993; Rath et al., 1997). Other software was more suited to work with plastic sections, but still provided a huge savings in dose to the specimen and an increase in throughput, quality, and convenience (Fung et al., 1996; Koster et al., 1992, 1993).

With the availability of better goniometers, new approaches have been devised with the goal of dispensing with some of the tracking and focusing steps by taking advantage of the predictable properties of the stage. The method of precalibration involves measuring stage movements in X , Y , and Z through a coarse tilt series (e.g., at 5° intervals), possibly at a lower magnification, and then applying these movements to acquire one or more full tilt series at a finer interval (Ziese et al., 2002). Precalibration is the basis for the tomography package supplied by the FEI Company (Eindhoven, Netherlands). Most recently, a prediction method has been developed (Zheng et al., 2004) which assumes that each point on the specimen moves in a circle around the tilt axis as the stage is tilted and predicts changes in Z -height from recent changes in the lateral position of the specimen. This method, which will be referred to here as Z -prediction, requires essentially no tracking or focusing images during the tilt series and represents an elegant solution to tilt series acquisition in situations where the goniometer and specimen holder satisfy its assumptions.

This paper describes a different approach to skipping steps and speeding up data acquisition. Specimen position is predicted in X , Y , and Z based only on the movements during recent tilts. The prediction is relied upon, i.e., tracking or focusing is skipped, only when the statistical error of the prediction is low enough. Rather than a prediction based upon a previous trajectory, as in precalibration, or a prediction of one coordinate based upon geometrical considerations, this method uses a conservative and adaptive prediction from movements in the current tilt series. The program described here, SerialEM, was developed after experimenting with the precalibration approach and becoming concerned about its vulnerability to less than ideal circumstances: when there is specimen drift or nonreproducible holder movements, when the eucentric height is not set equivalently between a precalibration run and the actual tilt series, or when it is necessary to readjust specimen position during

the series. It will be shown here that the program succeeds in achieving the speed of precalibration where conditions are favorable, while providing the reliability of the traditional approach in less ideal cases.

This paper is organized as follows. Various components and features of the program needed for successful automation will be described first. The prediction algorithm and its implementation by a module called the “tilt series controller” are then explained. Examples are given of the behavior of the program in typical and extreme circumstances. The performance of this method relative to the precalibration and Z -prediction methods is evaluated, and data are presented indicating that the Z -prediction method is particularly susceptible to nonideal stage performance. Finally, additional features of the program are described, such as its ability to work with an energy filter and to take montaged images from overlapping CCD images. One motivation for describing some aspects of the program in detail is to highlight the various nonideal properties of current microscopes and cameras that keep automated tomography from being straightforward, and to present ways in which these problems can be overcome.

This work was reported previously in abstract form (Mastronarde, 2003).

2. Equipment and other preliminaries

2.1. Equipment

SerialEM was developed on and has been used most extensively on two Tecnai microscopes in the Boulder Laboratory for 3D Electron Microscopy of Cells (BL3DEMC). Our Tecnai F30 (FEI Co.) operates at 300 kV and has a field emission gun (FEG), a Twin objective lens, a Megascan 795 cooled CCD camera from Gatan (Pleasanton, CA) below the camera chamber, and Gatan Imaging Filter (GIF 2002), also equipped with a Megascan 795 camera. Both cameras have $2K \times 2K$ 30- μm pixels that are read out through a single port. Our Tecnai F20 operates at 200 kV and has a Twin lens, a FEG and a Gatan Ultrascan 895 CCD camera, which is a 4-port readout camera with $4K \times 4K$ 15- μm pixels. On each Tecnai, SerialEM runs on the microscope control computer, which has two 1-GHz Pentium III processors. Specimen holders used for tomography include a Gatan 670 high-tilt holder, a modified Gatan 630 high-tilt holder with in-plane rotation, a Model 2020 Advanced Tomography Holder (E.A. Fischione Instruments, Export, PA), Gatan 650 high-tilt cryo-holders, and a Gatan CT3500TR cryo-holder with in-plane rotation.

SerialEM was adapted to run on a JEOL 3100-FEF (JEOL USA, Peabody, MA) at the Lawrence Berkeley Laboratories, University of California. This 300 kV

microscope has a liquid-helium cooled stage, an in-column energy filter, and a Gatan Megascan 795 camera. The program can also control cameras from Tietz Video and Image Processing Systems (TVIPS, Gauting, Germany); this capability was tested and used for data acquisition on a Tecnai F20 at the Scripps Research Institute (La Jolla, CA).

2.2. Program architecture

SerialEM is written in C++ using Microsoft Foundation Classes (MFC). Its basic organization and several core modules were adapted from an earlier Macintosh-based version of the program used to automate tilt series acquisition on our JEOL 1000 kV electron microscope (Mastronarde et al., 2001). It controls the microscope via the COM interface provided by the manufacturer

(the Tecnai scripting adapter or the JEOL TEMCON interface). It obtains images from a Gatan camera through a custom plugin to Gatan's DigitalMicrograph software that acts as a COM server. It obtains images from a TVIPS camera through a COM and shared memory interface provided by TVIPS.

The program is intended to provide not just a tool for efficient tilt series acquisition but also a complete interface for image capture, display, and storage and for control of some aspects of microscope function. It is organized around a main window that displays one image out of a stack of up to 16 buffered images (Fig. 1). This design allows one to toggle between various images in the buffer stack, making it easy to assess image alignment and compare focus or other features of images. When an image is acquired, it is placed in the first buffer and a selected number of images are rolled into higher

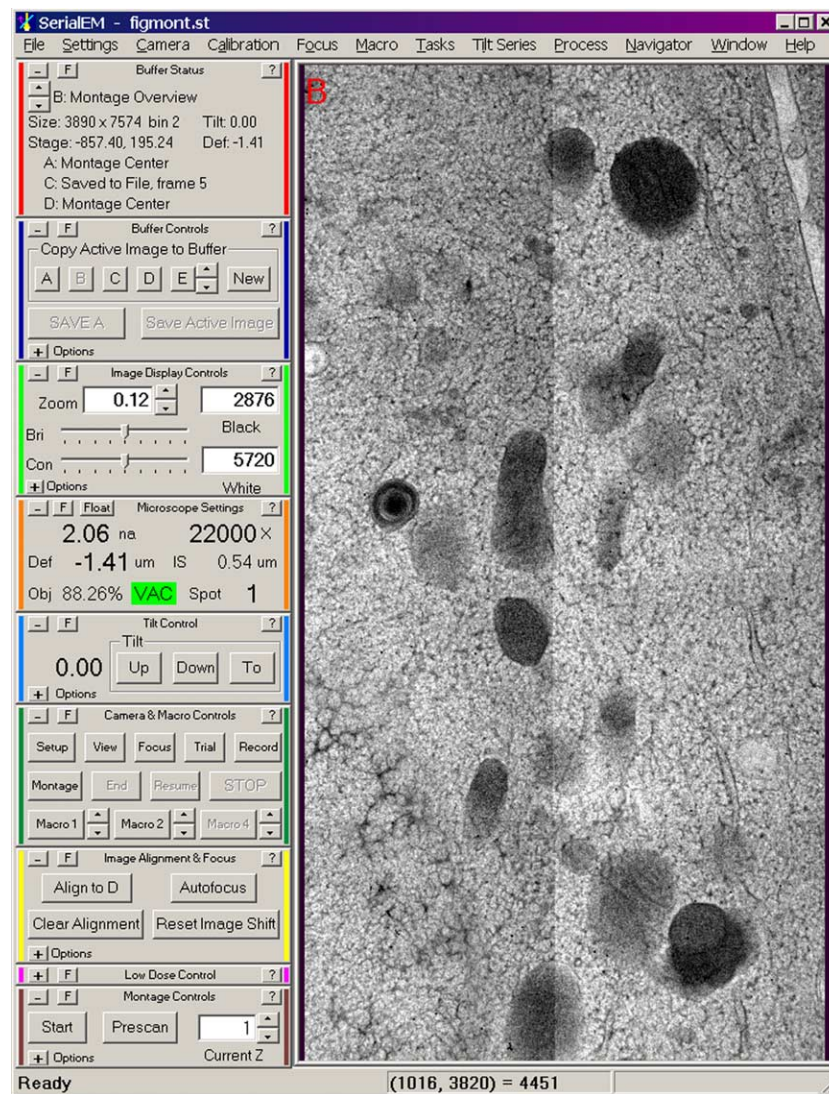


Fig. 1. Screen shot of the SerialEM program. The control panels on the left can be closed, opened further to display additional controls, or floated as independent windows and placed elsewhere on the desktop. The image window displays a stack of images that can be scrolled through. The image here is the overview shown after acquiring a 2 by 4 montage of overlapping 2K × 2K frames.

buffers. Typically, three buffers roll in this fashion, and higher buffers are used for images to be retained, such as reference images.

Beside the main window are a collection of control panels that provide the most commonly used controls and basic information about images in the buffers or about microscope status. Many of the panels can be opened up to expose lesser-used but related options. In addition, individual panels can be closed entirely or floated as free-standing windows, which is useful on two-monitor systems. The menus provide a wide variety of other functions. Dialog boxes are opened to set parameters in some situations, the most important being the image acquisition and the tilt series setup.

2.3. Image shift limitations due to the objective aperture on the 300 kV Tecnai

SerialEM was designed to cope with the limited range of image shift on the 300 kV Tecnai, a problem that becomes quite significant with the small objective aperture sizes needed when working with stained thick sections. When image-beam shift is used, the objective aperture shifts as well, apparently because it is not in the back focal plane of the objective lens for this particular microscope. With a relatively condensed beam and small objective aperture, the aperture intrudes on the field of the camera within 1–3 μm of image shift. For example, under typical conditions for working with plastic sections (magnification to give 1 nm pixel on the CCD camera, a 100 μm condenser aperture, and spot size 2) and with the beam spread to cover most of the viewing screen, the objective aperture intrudes when image shift reaches 2.1 and 3.3 μm for 20 and 30 μm apertures, respectively. With the beam condensed to about half as large, the respective limits are 0.7 and 1.8 μm . The practical limits on image shift are even smaller for several reasons: (1) the aperture shifts further under the beam tilt used for autofocus; (2) under low-dose conditions, some of the image shift range is needed to offset tracking and focusing areas; and (3) it is hard to center the aperture perfectly. Because of this constraint, the tilt series controller must be able to reset image shift when it becomes too large, move the stage by a corresponding amount, and still be able to track the specimen reliably after this operation.

2.4. Coordinate systems and measurements

In keeping with microscope conventions, the X axis corresponds to the tilt axis, the Z axis is along the beam, and the Y axis is orthogonal to the two. The X and Y positions reported below are based on the image shift settings at each step in the tilt series. These image shift values are converted to microscope coordinates by first transforming them to camera coordinates, using a calibration of the relation between image shift and camera

position at the given magnification. Camera coordinates are converted to microscope coordinates using the pixel size (calibrated from a cross-line grating) and the angle of the tilt axis on the camera. This angle is obtained from the fiducial marker alignment (Lawrence, 1992) of a typical tilt series using the IMOD software (Kremer et al., 1996; Mastronarde, 1997). The angle of image movement when the stage is moved along the X axis was found to be an inadequate estimate of the tilt axis angle, being off by about 5°. Z positions reported below are based on the focus setting needed to achieve the desired image defocus. These focus changes have been verified to be equivalent to Z height changes of the same magnitude.

3. Essential components for automated tomography

3.1. Alignment by cross-correlation

The method of aligning images by cross-correlation incorporates several features to optimize speed, reliability, and flexibility. First, images to be correlated need not be the same size or be taken at the same binning. Given two images, the alignment routine first determines what further binning is needed to bring both images to a common binning and a maximum size of at most 512 pixels. Each image is then binned if necessary, then the image at higher tilt is stretched by the ratio of the cosines of their tilt angles (Dierksen et al., 1992; Gluckenberger, 1982). The mean around the edge is measured, and intensities are linearly tapered down to the mean over 5% of the image. Both images are padded to the same size, only slightly larger than the larger image, and cross-correlated by the standard method of multiplying the Fourier transform of one with the complex conjugate of the Fourier transform of the other. The cross-correlation is band-pass filtered in frequency space by

$$F(R) = (1 - e^{-R^2/2\sigma_1^2})e^{(R-R_c)^2/2\sigma_2^2}, \quad (1)$$

where R is the radius in frequency space, σ_1 is the sigma for an inverted Gaussian to filter out low frequencies (0.03 cycle per pixel), and σ_2 is the sigma for a Gaussian starting at the cutoff frequency, R_c , to filter out high frequencies (0.05 and 0.25 cycle per pixel, respectively). Both the tapering at the edge and the low-frequency filter are important for avoiding spurious correlation; the high-frequency filter removes noise so that an accurate correlation peak position can be determined by interpolation.

With essentially unpadded images, most correlation peak positions are ambiguous; for example, a peak at (150, -100) in a correlation between two 500 \times 500-pixel images could arise from a relative shift of (150, -100), (-350, -100), (150, 400), or (-350, 400). To resolve the ambiguity, the routine correlates small samples of the

two images from corresponding positions based on each of the four possible shifts and picks the shift that gives the highest correlation. In this way, the routine can reliably detect large shifts (even larger than half of the field width) without the computational expense of correlating images padded to twice their original size.

In circumstances where the beam or aperture edge might intrude into the image, such as when correlating images taken at a lower magnification, the routine can analyze for dark regions at the borders of the image and take a centered subarea that excludes those regions. This procedure prevents the correlation from being thrown off by the strong edge.

3.2. Autofocus from beam-tilted images

Autofocusing is done with the established technique of measuring the displacement of the image upon tilting the beam (Dierksen et al., 1993; Koster et al., 1987, 1992). The basic module for focus detection takes two images, first with negative then positive beam tilt. As an option to protect against drift, it can also take a third image with negative beam tilt and use this image to estimate and subtract off the drift between images, thus isolating tilt-induced from drift-induced shift. Autofocus is calibrated by measuring the shift at 6- μm intervals over 72 μm of defocus. The image shift per defocus interval changes by up to 20% over this range, so autofocusing does not assume a linear relationship. Instead, given a measured shift in X and Y , the autofocus routine finds the nearest point on the calibration curve of Y shifts versus X shifts and interpolates to determine defocus at this point.

The autofocus routine includes a correction for beam tilt-induced image stretching similar to that recommended by Ziese et al. (2003). The correction is based upon the shift in image position per unit of Z -height or focus change and per unit of beam tilt, measured from the calibration curves and referred to as k_x and k_y for X and Y , respectively. If the tilt axis is rotated by θ from the X axis in an image, the specimen is tilted by α , and the beam tilt is τ , then a point at (x, y) in an image is located a distance $d = -x \sin \theta + y \cos \theta$ from the tilt axis, its relative Z height is $h = d \sin \alpha$ and it is displaced by $(h\tau k_x, h\tau k_y)$. The image is thus transformed by:

$$\begin{aligned} x' &= x + h\tau k_x = (1 - \tau k_x \sin \theta \sin \alpha)x + (\tau k_x \cos \theta \sin \alpha)y, \\ y' &= y + h\tau k_y = -(\tau k_y \sin \theta \sin \alpha)x + (1 + \tau k_y \cos \theta \sin \alpha)y. \end{aligned} \quad (2)$$

These equations are the basis for a linear transformation applied to each beam-tilted image before correlating.

3.3. Controlling pre-exposure and minimizing total specimen dose with Gatan cameras

With the high beam currents used to image thick plastic sections, CCD camera images are often blurred by an initial drift when the beam hits the specimen, especially with

sections on the plastic films used with slot grids. This problem is solved by pre-exposing the specimen until the image is stabilized. With a Gatan camera controlled by DigitalMicrograph, pre-exposure is best accomplished using the shuttering mode illustrated in Fig. 2A. In this mode, a shutter above the specimen (beam shutter) is opened at the start of the acquisition cycle while a shutter below the specimen (film shutter) is simultaneously closed. After the CCD chip has been cleared, the film shutter is reopened to expose the CCD. As a result, the CCD is illuminated only during the exposure time, while the beam is on the specimen during the whole cycle of clearing, exposure, and readout. Pre-exposure thus occurs during the clear time of the CCD chip. Other shuttering modes that would also provide pre-exposure are not suitable for high-throughput work because those modes expose

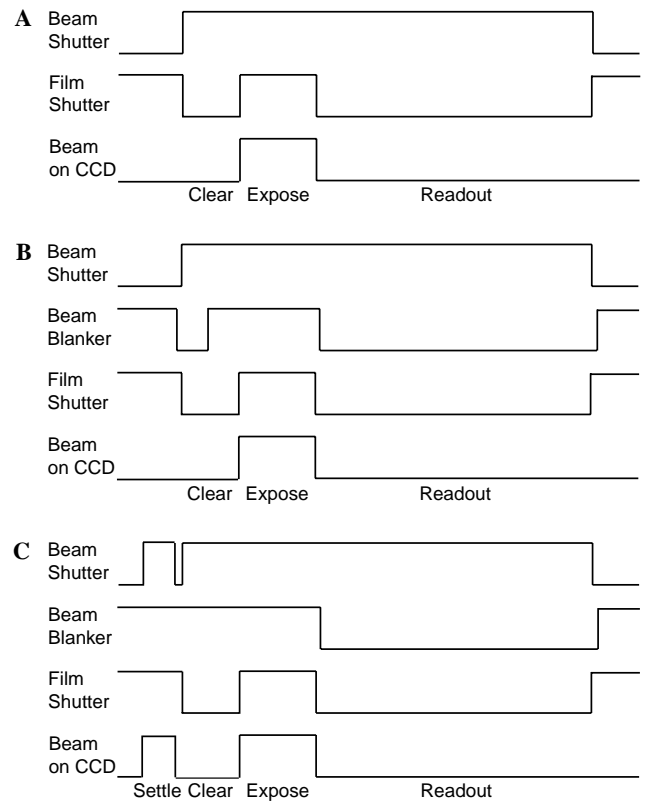


Fig. 2. Shuttering arrangements to obtain controllable pre-exposure and minimum dose to the specimen with a Gatan camera. In the traces, an upward deflection indicates an open shutter or beam present. The beam shutter and the beam blanker are above the specimen; the film shutter is below the specimen. (A) Gatan shuttering: the one shuttering mode that provides pre-exposure of the specimen without exposing the CCD while it is being cleared. The beam shutter is normally closed but is opened for the entire cycle of clearing the CCD, exposing, and reading out the image. (B) SerialEM shuttering with pre-exposure shorter than the clearing time: SerialEM blanks the beam during the first part of the clear interval; it also blanks the beam after the exposure to prevent specimen exposure during the readout. (C) SerialEM shuttering with longer pre-exposure: SerialEM opens the camera shutter for the necessary additional settling time before requesting an image.

the CCD during the clear time and thus require a matching dark reference for each image.

There are two crucial problems with this shuttering mode: the amount of pre-exposure that it provides is fixed (and indeed is too small with some 1K cameras); and the specimen is exposed during the CCD readout, which can be many seconds. (Note that the latter problem does not occur with the shuttering mode used for low-dose work, in which the beam shutter is used alone and is open only during the exposure time.) To overcome these problems, SerialEM provides a special shuttering mode. To get pre-exposure shorter than the clear time, a microscope scripting call is used to blank the beam above the specimen for the appropriate part of this interval (Fig. 2B). To get a longer pre-exposure, a camera scripting call is used to open the shutter before starting the acquisition (Fig. 2C). In either case, the beam is blanked during CCD readout to minimize the total dose. By assessing the variability in the interval from SerialEM's request for an image to the actual start of an exposure, timing parameters can be set so that this method works reliably. At least on dual processor computers, the method succeeds in eliminating all but 50–100 ms of the extra exposure during the readout.

Once the ability to blank the beam at the start of camera readout was developed, it provided an opportunity to start other time-consuming tasks at the same time, namely moving the stage, applying image shift, or changing magnification. During tilt series acquisition, this capability is most useful for tilting the stage after acquiring the final image at a tilt angle, since the time required for tilting and allowing the stage to settle can be subsumed in the image readout time.

3.4. Routines for specific tasks

The tilt series controller uses several routines, in addition to the alignment and focus routines, to perform operations needed before or during a tilt series. These routines can also be run directly by the user. Each one can reduce the magnification if necessary to obtain a camera field of view adequate for the particular task. They use binned, short-exposure images to minimize electron dose and time. Times cited below are with a 2K × 2K single-port readout Gatan 795 camera. The routines include:

1. Reset image shift and realign: image shift is set to zero and the stage is moved to compensate. Images taken before and after this operation are aligned, and a new, smaller image shift is imposed, to maintain specimen position.
2. Reverse tilt direction: the stage is tilted by 3° and back to work out the stage backlash when changing tilt direction. Images taken before and after this operation are aligned to maintain specimen position.
3. Coarse eucentricity: images are taken while tilting the stage over a small range (13°), starting with a very small increment (0.6°) and using progressively larger increments up to 5°. If specimen movements are large, the Z-height is changed during the routine to permit larger tilt increments. This routine can find the eucentric height from up to 100 μm away, but may be accurate to only 0.5–1 μm. It takes about 45 s.
4. Fine eucentricity: eight images are taken between –24° and 24°. The image shift position in Y is least-squares fit to $y = (y_0 + y_s)\cos\alpha - z_0\sin\alpha - y_0$ to determine both z_0 , the Z-height, and y_0 , the offset between tilt and optical axes, where y_s is the image shift of the specimen at zero tilt. This fit provides a value for eucentric height uncontaminated by the effects of a tilt axis offset. The result appears to be accurate to 0.1–0.25 μm, so this routine is used to set the eucentric height before starting a tilt series. It will work only for modest Z-height disparities (up to 10 μm) and may restart after adjusting Z-height if image shifts become too large. It takes about 70 s.
5. Walk up: provides for tilting to a desired angle (e.g., to start a tilt series at a high angle) while keeping the area of interest centered. Images are taken while tilting at relatively large increments ($8^\circ \cdot \cos\alpha$). An image at one intermediate tilt angle can be left in a buffer for use as a reference during a tilt series. This image is called an “anchor” and allows recovery from errors in tracking that might accumulate at the higher tilt angles. The routine takes about 70 s to tilt from 0° to 60° and an additional 30 s to reach 70°.

3.5. Beam intensity control

SerialEM incorporates control of beam intensity so that the beam can be adjusted to compensate for magnification changes in the above-described routines, and for changes in specimen thickness along the beam axis during a tilt series. A large magnification change without a change in the condenser lens presents two problems: the area illuminated by the beam can become much smaller than the camera field of view, and the intensity can increase enough to saturate the camera, since the camera exposure time might already be so short (<50 ms) that it cannot be reliably decreased to compensate for the intensity increase. The Tecnai has an “intensity zoom” feature that keeps the projected brightness approximately the same as magnification is changed, but the changes are not precise enough to rely upon, so a separate calibration of intensity was desirable. There were two reasons to use beam intensity rather than exposure time changes to adjust for the increased thickness of tilted specimens. First, when gain-normalized images are acquired from DigitalMicrograph, every change in

exposure time requires a new dark reference, which takes almost as long as the image acquisition itself. Second, the pre-exposure required to avoid drift in images from plastic sections appears to depend primarily on the total amount of beam, not on how condensed it is. The beam can thus be condensed at high tilt without requiring more pre-exposure; whereas an increased exposure time usually requires longer pre-exposure.

Beam intensity is calibrated with a routine that uses the CCD camera to measure image intensity. The procedure is started at a high magnification with a fairly condensed beam and the routine constructs a table of intensities by spreading the beam progressively with the second condenser lens, reducing intensity by $\sim 10\%$ at each step. Binning, exposure time, and magnification are changed as appropriate to keep the camera counts within a reasonable range (e.g., 100–4000 counts per pixel), with a scaling factor determined at each such change and applied thereafter, so a single continuous curve is obtained. A separate calibration is needed for each spot size, but the calibrations are fairly stable over time and do not need to be redone often. Once these calibrations are done, they can be used to reduce intensity when changing magnification and to control intensity during a tilt series. In addition, they allow the user to change intensity by controlled amounts through a menu entry, and they allow electron dose to be calculated for a given condenser lens strength and camera exposure time. The latter feature requires the user to take a single image for calibrating the beam brightness during that session.

4. Tilt series acquisition

4.1. The prediction method

The essential features of the prediction method are that the next value of a coordinate is predicted by extrapolation from previous values, and that a prediction is relied upon only when two measures of reliability are good enough. The two measures are the standard error of the extrapolated value from a linear or quadratic least-squares fit, and the difference between the last prediction and the actual current location. Fig. 3 illustrates this process for one coordinate. No predictions can be made for the first two tilts. For the third tilt, the position can be extrapolated from the previous two tilts. The specimen will be set to this predicted position, but the reliability of this prediction cannot yet be assessed, so the position is still tracked and the prediction is not relied upon. On the fourth tilt, both the actual error in the last prediction and the standard error of the new extrapolated value can be measured. In Fig. 3, both errors are relatively large, as indicated by the distance between open and filled circles, and by the error bars, respective-

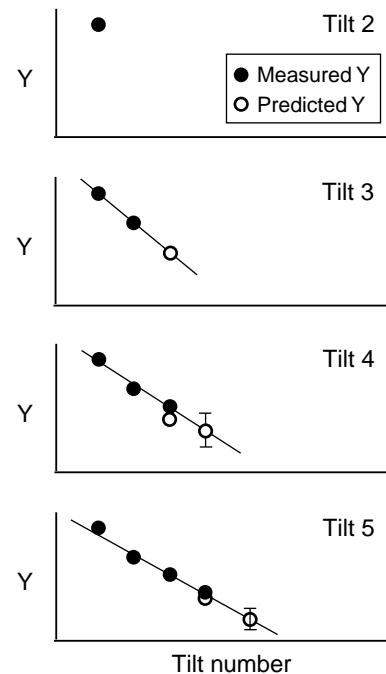


Fig. 3. Application of the prediction method to one coordinate. At each tilt angle, a line is least-squares fit to positions measured on previous tilts (filled circles) and extrapolated to predict the position at the current tilt. Open circles show positions predicted at the current or previous tilts. The error bars represent the standard error of the prediction. The program relies on the prediction and skips tracking only when the standard error of the current prediction and the actual error of the previous prediction are both small enough (e.g., tilt 5).

ly, so the prediction is again used but not relied upon. On the fifth tilt, both errors are small enough, and the program will rely on the prediction by acquiring a final image without first tracking the position.

This simple approach is embellished by a host of details. The prediction of X/Y position is somewhat more complex than indicated by Fig. 3. Predictions and standard errors are computed separately for X and Y . A tracking image is required if the actual error from the last prediction is greater than a criterion in either X or Y , or if the square root of sum of squares of the two standard errors is greater than the same criterion (about 2% of the image size). Regardless of whether a tracking image is taken, the X/Y position is measured on every tilt by aligning the final image to the reference image from the previous tilt. Thus, once predictions in X/Y become reliable, there is no need to take tracking images.

The fitting is adaptive, in that the program will use only the most recent data if it gives a substantially better fit. Specifically, it starts with data from a restricted range of tilts (e.g., 8–12°) and computes a series of predictions by dropping out progressively more of the older data, but retaining at least four points. If the standard error of the prediction using all points is more than 1.3 times the smallest standard error, it drops points until the standard error is less than 1.1 times the smallest error. This

method allows the predictions to accommodate quickly to a transition from irregular to well-behaved changes.

The aligned final image becomes the reference for the next tilt; since it is first shifted into alignment with the previous reference, an error in specimen centering in this particular image will not throw off the tracking. Nevertheless, such errors represent a potential loss of data, and users can choose to have the final image acquired again if the amount of data lost exceeds a criterion such as 2.5–5% of the image size. In computing data loss, the program ignores errors in Y that are less than the amount that the zero-degree field of view is foreshortened perpendicular to the tilt axis by tilting to the current angle.

In contrast to the X/Y position, which is measured from the final image regardless of whether tracking is done, new measurements of Z position are available only when autofocusing is done. Thus, focusing must be done periodically, typically every 5–6°, even when predictions have low errors. Focusing is required if the actual error from the last prediction or twice the standard error of the current prediction exceeds a criterion such as 0.2 μm . In addition, users can specify that autofocusing should occur every time above a certain tilt angle, which is helpful with some samples where autofocusing is less reliable at high tilts. When a prediction is available, autofocusing may be repeated if the actual value differs from the prediction by more than a criterion amount.

A variety of events can disturb the system and cause the program to ignore positions before the disturbance and start accumulating fresh data for predictions. These are primarily events that move the stage, such as resetting image shift. If the user stops the tilt series acquisition, actions such as backing up to a previous tilt angle or changing the target defocus will also count as disturbances.

4.2. Typical tilt series for dose-insensitive specimens

Fig. 4 describes a typical tilt series, showing specimen movements and the occasions on which alignment images were acquired. The user followed our typical protocols for plastic section tomography. Figs. 4A–C show the positions in X , Y , and Z , respectively, for the final acquired image at each tilt angle. Fig. 4D shows whether tracking and focusing images were taken in addition to the final image on each tilt. The program was able to skip tracking after three tilts and focusing after four tilts, with some extra tracking and focusing images required at -50° . At -38° , the program substituted the anchor image taken while tilting up for the reference image, thus disturbing the prediction of positions. Tracking and focusing images were taken until predictions became reliable again, after three tilts. Thereafter, tracking was not needed

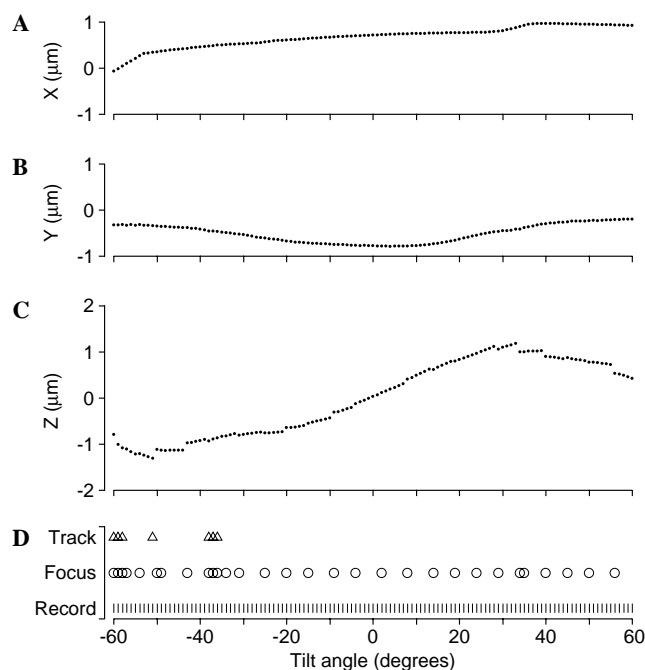


Fig. 4. Summary of specimen trajectory and images taken during a typical tilt series. (A and B) The X and Y position, based on the electronic image shift applied at each tilt. (C) The Z position, based on the focus change needed to achieve the desired defocus. (D) Log of tracking images (triangles), autofocus (circles), and final Record images (vertical lines) taken during the series. The 1° tilt series was taken by Eileen O'Toole on the Tecnai F30 microscope at a film magnification of 23 000 \times ; Record images were 2048 \times 2048 pixels with a 1.0 nm pixel size on the Gatan 795 camera. Total acquisition time was 33 min.

and focusing was done nearly exclusively at the required intervals.

Other typical tilt series often show a few angles where tracking is done because of small irregularities in the stage movement. Series at different tilt increments (e.g., 0.75–2°) require about the same number of images for focusing and tracking as in the example of Fig. 4. Thus, the prediction method allows very small tilt intervals to be used with little increase in the time spent focusing or tracking.

4.3. Tilt series at high magnification

For acquisitions at high magnification, the tilt series controller has an option to go to a lower magnification for tracking. A low-magnification reference image is acquired before the first tilt, and another image obtained after tilting is used for alignment; it becomes the new reference. Once the stage movements become regular enough so that shifts are expected to be acceptably small at the working magnification (e.g., <0.3 of the image size), the low-magnification tracking is no longer needed. This option is recommended when the camera field of view is less than $\sim 1 \mu\text{m}$.

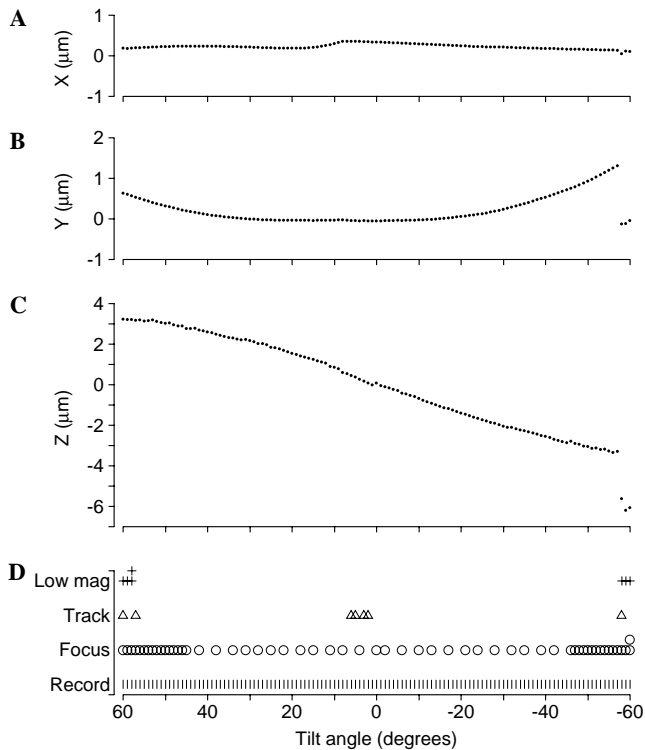


Fig. 5. Tilt series taken at relatively high magnification. (A–C) Specimen trajectory in X , Y , and Z . (D) Log of acquired images. Conventions are as in Fig. 4; in addition, crosses in (D) indicate when low-magnification tracking pictures were taken. At -58° , the image shift was reset automatically, causing large changes in Y and Z . The 1° tilt series was taken by Mary Morphew and Pamela Bjorkman on the Tecnai F30 at a film magnification of $59000\times$. Record images were 2048×2048 pixels with a 0.39 nm pixel size on the Gatan 795 camera. Total acquisition time was 42 min.

Fig. 5 shows a tilt series taken at a magnification ($59000\times$) where the camera field of field was $0.8 \mu\text{m}$. Low-magnification tracking was invoked for the first three tilts, and also near the end of the series when the image shift was automatically reset to zero. A few other tracking images were taken at the working magnification of the tilt series, but through most of the series no tracking was needed. The user chose to have the program focus on every tilt above 45° . This example illustrates how the low-magnification tracking option can bolster the reliability of a tilt series at higher magnification without requiring significantly more time.

Fig. 6 shows results from a test tilt series taken at a very high magnification ($230000\times$) where the field of view of the CCD camera was only $0.2 \mu\text{m}$. To avoid the difficulties of starting at high tilt, the series was taken in two segments, first from 0° to -60° , then from 1° to 70° . Here, low-magnification tracking was invoked at the start of each segment and on only five other tilts. Regular tracking images were needed, however, through much of the tilt series. When predictions become unreliable, the program degrades

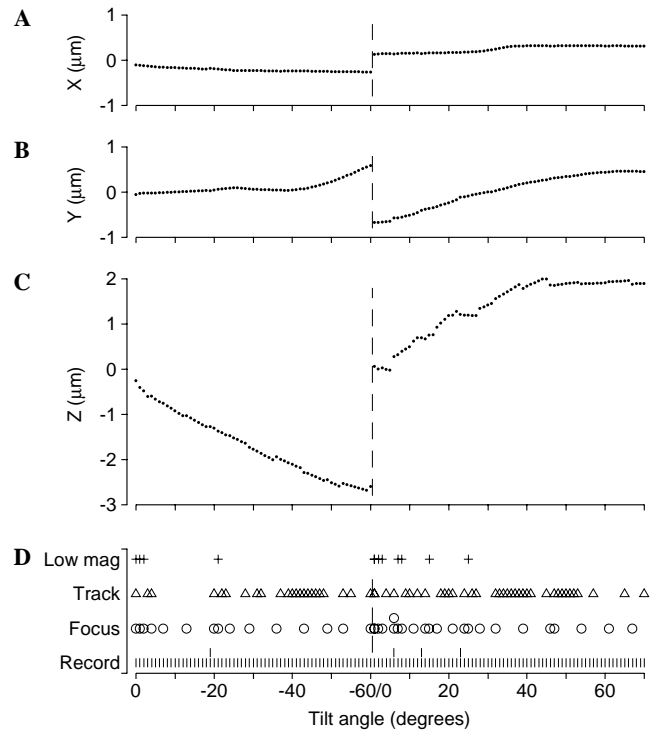


Fig. 6. Test tilt series taken at extremely high magnification. (A–C) Specimen trajectory in X , Y , and Z . (D) Log of acquired images. Conventions are as in Figs. 4 and 5. The specimen was a 200-nm plastic section of a T-cell lymphocyte. The 1° series was obtained in two segments on the Tecnai F30, first from 0° to -60° , then from 1° to 70° , at a film magnification of $230000\times$. Record images were 1024×1024 pixels with a 0.2 nm pixel size on the Gatan 795 camera. Acquisition time was 26 min (exclusive of manual steps for restarting the series at 1° for the second segment).

gracefully into the traditional method of tracking on every tilt.

4.4. Low-dose mode

For tomography of beam-sensitive specimens such as frozen-hydrated material, the standard approach has been to take focus and tracking images at a location displaced from the area of interest along the tilt axis (Dierksen et al., 1993; Nickell et al., 2005; Rath et al., 1997). To incorporate this approach into SerialEM, a low-dose mode was implemented, independent of any specialized low-dose systems on the microscope. This mode provides up to four different “areas,” and each of the five sets of camera parameters (called View, Focus, Trial, Record, and Preview) is linked to one of these areas. Each area can have an independent magnification, beam intensity, spot size, and energy filter settings. To activate an area, one can either acquire an image with the corresponding camera parameter set or select the area to be shown when the screen is down. Once an area is active, one can adjust the parameters for that area. One area, the “View” area, is intended to provide a lower magnification, low-exposure view centered on the

area of interest. Automated operations that ordinarily would be done at lower magnification, such as finding eucentricity, are generally done with View images in low-dose mode. The focusing and tracking areas are meant to be offset along the tilt axis. These areas can be positioned conveniently by clicking on a View image; the program selects the closest point along the tilt axis. Two camera parameter sets are linked to the “Record” area so that both binned, low-exposure images (Preview) and final, high-quality images (Record) can be taken from this area.

The main complication for tilt series acquisition in low-dose mode is that separate alignment references are needed for the tracking and Record areas. The former is needed when tracking is required before acquisition, whereas the latter is needed when aligning the final acquired image with the final image of the previous tilt. The references are kept in two separate buffers; the program automatically aligns an image from the tracking or Record area to the proper reference. During a tilt series, even when tracking is not needed, the program will acquire a tracking reference periodically so that one is available from a nearby tilt angle if needed (see upper row of triangles in Fig. 7D). A stack of these references is also kept so that if the user needs to back up

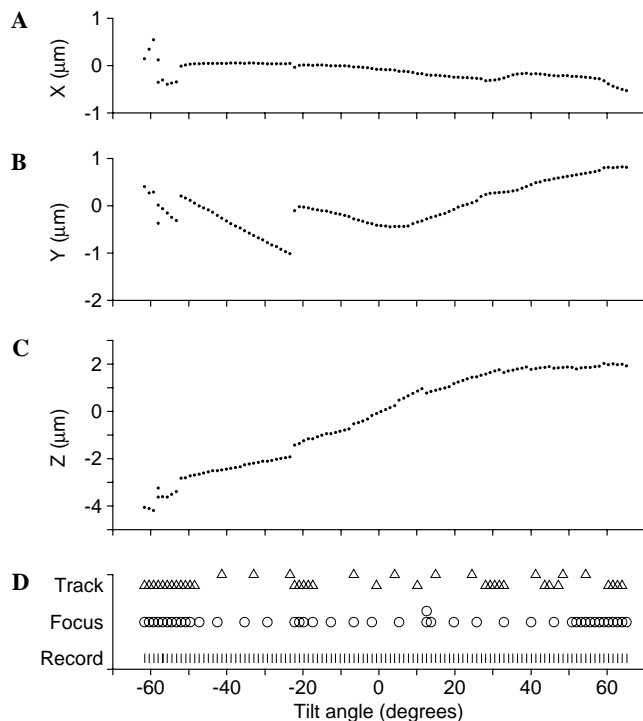


Fig. 7. Cryo-tilt series taken in low-dose mode with energy filtering. (A–C) Specimen trajectory in X, Y, and Z. (D) Log of acquired images. Conventions are as in Fig. 4. The upper row of triangles indicates images taken just to obtain a new reference image of the tracking area. The 1.2° tilt series was taken by Daniela Nicastro. Record images were 2048×2048 pixels with a 0.95 nm pixel size on the GIF camera. Total acquisition time was 56 min and total dose to the Record area was $128 \text{ electrons}/\text{\AA}^2$.

to a previous tilt angle, both a tracking reference from this stack and a Record reference from the data file are available. Maintaining two references does become problematic when they disagree, as happens when there is substantial drift between a tracking image and the final image. The disagreements are generally resolved by simply shifting the tracking reference by the same amount as the final image is shifted to align to its reference, although sometimes a new tracking reference is automatically taken to resolve the conflict.

The examples in Figs. 7 and 8 illustrate the challenges of doing tilt series with an externally inserted cryo-rod, which is susceptible to drift, noneucentric behavior, and failure of the Tecnai Compustage to hold the rod tightly at high tilt. Fig. 7 shows a relatively favorable example. Despite difficulties tracking at the start, the predictions were reliable enough through most of the tilt series to keep tracking and focusing to a minimum.

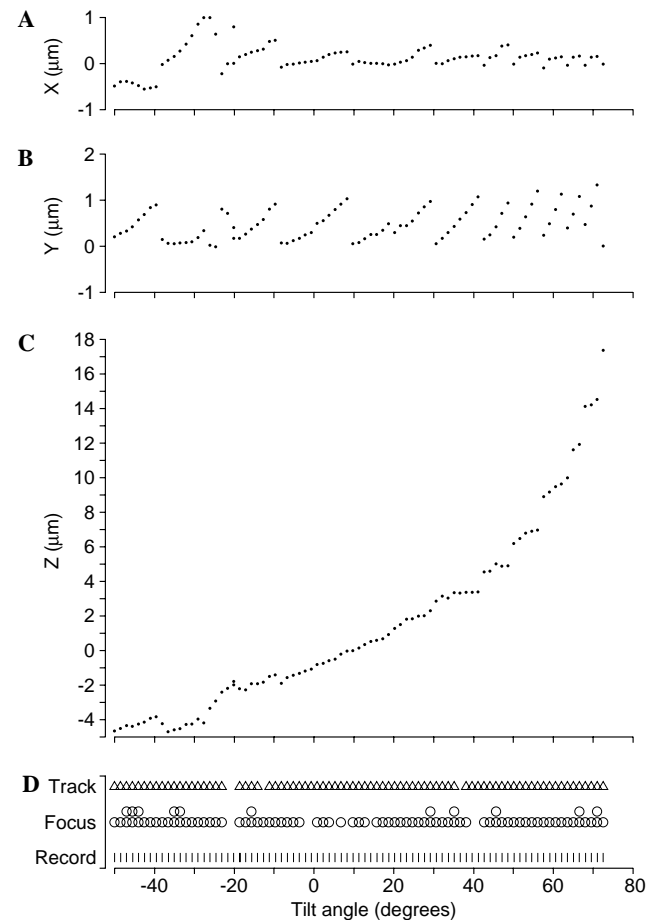


Fig. 8. Cryo-tilt series taken under unfavorable circumstances. (A–C) Specimen trajectory in X, Y, and Z. (D) Log of acquired images. Conventions are as in Fig. 7. The extreme amount of lateral shift between images required the image shift to be reset many times. The 1.5° tilt series was taken by Jason Pierson. Record images were 2048×2048 pixels with a 0.95 nm pixel size on the GIF camera. Total acquisition time was 65 min.

Fig. 8 shows a worst-case situation, in which progressive shifts in Y required the image shift to be reset many times. The program succeeded in acquiring this tilt series by reverting almost completely to the traditional approach of focusing and tracking on each tilt.

4.5. Compensating for deviation between tilt and optical axes

The specimen movements during a tilt series can be minimized by using image shift to make the effective optical axis coincide with the tilt axis (Zheng et al., 2004; Ziese et al., 2002). SerialEM includes an option to move the center of the image shift coordinates to the tilt axis. The appropriate shift is estimated by running the fine eucentricity procedure. However, this shift is not imposed at the start of a tilt series, because it would misalign the objective aperture on the 300 kV Tecnai. Instead, the shift is imposed when the user selects the option or, thereafter, when the program is started, thus allowing the user to align the aperture with the beam shifted. This option does significantly reduce the lateral movements during a tilt series when the tilt axis offset is more than $\sim 1 \mu\text{m}$.

5. Comparison with other methods

5.1. Comparison with pre-calibration method

Pre-calibration involves measuring specimen displacements through a tilt series at a large tilt angle ($\sim 5^\circ$), then using those displacements to acquire the full tilt series quickly without further tracking steps (Ziese et al., 2002). The typical tilt series done with the prediction method described here avoids the overhead of tilting through the angular range twice, requires many fewer tracking images, and requires either a comparable or a somewhat greater amount of focusing, depending on whether the user has chosen to focus every time at high tilt. On balance, the prediction method should thus be closely competitive in speed to the pre-calibration method, while providing protection against problems that can make pre-calibration fail, such as drift, variable holder behavior, and difficulty setting the eucentric height. This protection allows for more reliable acquisition even from plastic sections and is crucial for doing cryotomography, where such problems become typical.

Acquisition times from SerialEM indicate that it enables high-throughput tomography, with relatively little time spent beyond what is needed just for collecting the final images. On our F30 with a single-port readout camera, standard tilt series ($\pm 60^\circ$, 1° intervals, $2\text{K} \times 2\text{K}$ images) take 35–45 min from the first to the last picture saved. Simply acquiring and saving these final images takes 32 min. On our F20 with a 4-port read-




Fig. 9. Trajectories from three tilt series taken on serial plastic sections during one session, showing that a single pre-calibration curve would not have provided accurate focus. (A–C) The X , Y , and Z coordinates, respectively. The 1.5° tilt series were taken by Byung-Ho Kang on the Tecnai F30 at a film magnification of 15500; (1.5 nm pixel on the Gatan 795 camera).

out camera, such tilt series take 15–22 min, of which 13 min are required for acquisition and saving. Thus, in the best cases, very little time is spent on focusing and tracking images, and in general only a modest amount of extra time is needed to assure a reliable result.

If a calibration curve can be taken once and used in multiple tilt series, then pre-calibration would be more efficient than the prediction method. However, for this approach to work, stage motions need to be reproducible at different positions on the grid, which seems not to be the case, at least for our microscopes. Specimen movements were compared between tilt series taken in the same session for six user sessions. Sometimes the specimen followed similar trajectories, but sometimes it did not. An example of the latter situation appears in Fig. 9, which shows the specimen coordinates during successive tilt series from three serial sections. All three started at the same image shift value in Y , from which they should have taken the same path in Y and Z . The variation in Z positions between the three series is over $1 \mu\text{m}$ (Fig. 9C), indicating that one pre-calibration curve would not have given well-focused pictures for all of these tilt series.

5.2. Comparison with Z-prediction method

Tilt series records from SerialEM can also be used to evaluate how well the Z -prediction method of Zheng

et al. (2004) would have performed. This method predicts changes in Z-height from Y-displacements, assuming that the point of interest on the specimen moves in a circle around the tilt axis. The lateral offset of the tilt axis from the optic axis is measured by a calibration routine and assumed to be constant until the stage is disturbed by servicing. Both the vertical offset from the eucentric point and the change in Z-height with tilting are determined by least-squares fitting to equations that relate these values to the displacements in Y on recent tilts and to the tilt axis offset. Two approaches were taken to evaluating how well this method would work on our microscopes. One approach was to plot Z versus Y to assess the validity of assuming that the specimen moves in a circle. Many such plots, particularly from our F30, do not have much circular character, such as

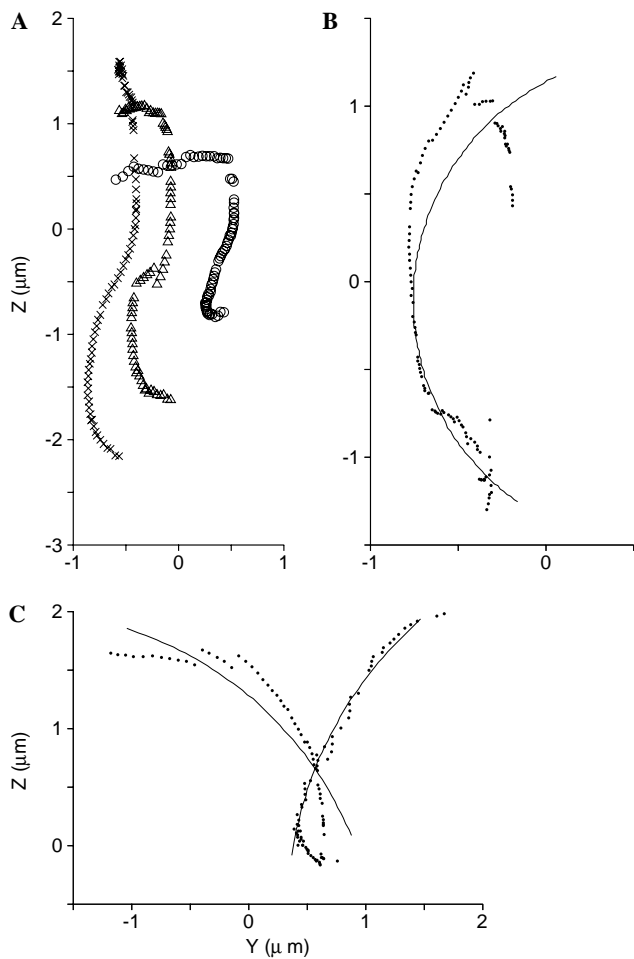


Fig. 10. Plots of Z versus Y position, to assess how well the specimen moves in a circle around a tilt axis. (A) Positions from the three tilt series in Fig. 9, which are dominated by vertical movements. (B) Positions from the tilt series of Fig. 4, which has a partially circular trajectory. The points are the measured data and the curve is the best-fitting circle for the given tilt angles. (C) Positions from a tilt series taken by Jaap Brink on a JEOL 3100-FEF, showing mostly circular movements at high-tilt angles and a large shift of the tilt axis between -12° and 20° . The curves on the left and right were fit separately to subsets of points from -60° to -12° and from 20° to 60° , respectively.

the three curves shown in Fig. 10A (based on the tilt series in Fig. 9). One set of tilt series from our F20 yielded nicely circular plots, but more often plots that are circular in parts show serious deviations from circularity in other parts, such as the example in Fig. 10B (based on the tilt series in Fig. 4).

The Z-prediction method was evaluated more directly by applying it to the Y-displacements measured in various recent tilt series from plastic sections. The method was applied in the positive and negative direction from zero tilt, with up to four values of the Y position relative to the Y value at zero tilt inserted into Eq. (6) of Zheng et al. (2004) to estimate the eucentric height and Z-displacement relative to zero tilt. A value for the lateral offset of the tilt axis was also required to apply these equations. Rather than use an actual measurement of the tilt axis offset for each microscope, a best-case estimate was obtained by measuring the root-mean-square deviation between predicted and measured Z positions for a range of offsets. For each set of

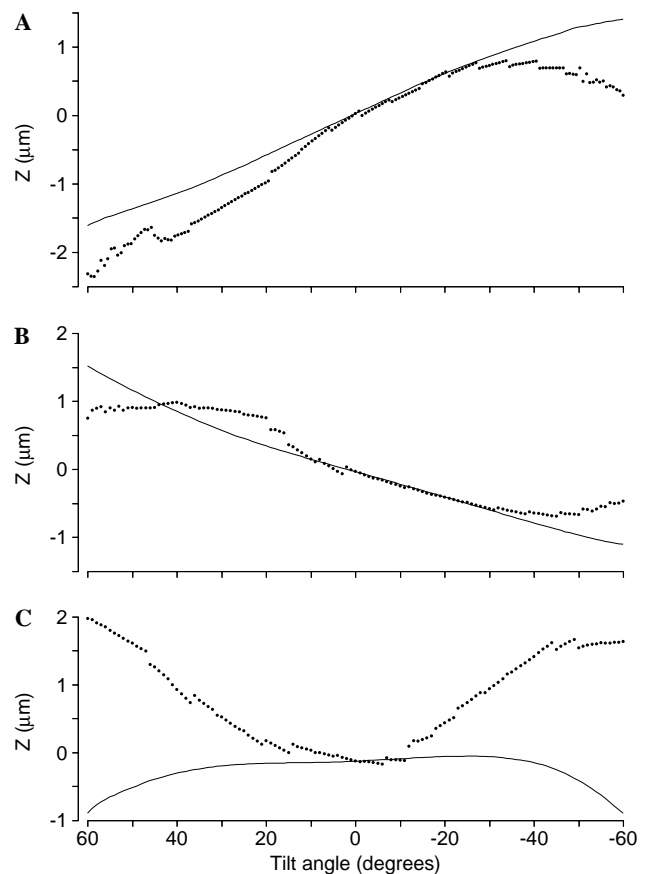


Fig. 11. Plots of Z position versus tilt angle. Points are measured Z positions; curves show predictions from Eq. (6) of Zheng et al. (2004) using the Y image shift displacements (see text for details). (A) Tilt series with extreme deviation from the Z-prediction, obtained by Charles Robertson on the Tecnai F20. (B) Tilt series with a typical amount of deviation, obtained by Zachary Gergely on the Tecnai F30. (C) Tilt series with extreme deviations from JEOL 3100-FEF (Fig. 10C).

tilt series from a session, the offset was found that minimized the mean deviation averaged over all of the tilt series. The mean deviation averaged $0.26\ \mu\text{m}$ and the maximum deviation averaged $0.75\ \mu\text{m}$ over the 45 data sets; these results would have been even larger if a fixed value for the lateral offset were assumed for each microscope. Predicted and actual Z values for one of the most deviant tilt series are plotted in Fig. 11A. Since $0.25\ \mu\text{m}$ is a reasonable expectation for focus accuracy, these results can be characterized by the fraction of tilts on which the predicted focus was off by more than $0.25\ \mu\text{m}$. In Fig. 11A, this occurred on 51% of tilts. The fraction of tilts on which the prediction was off by more than $0.25\ \mu\text{m}$ had a mean of 0.28, a standard deviation of 0.17, and a range of 0.0–0.58. Results from a tilt series that is typical based on this measure are shown in Fig. 11B. Similar results have been obtained from a Tecnai F30 at the University of Queensland (Brisbane, Australia) and more striking deviations have been seen on a JEOL 3100-FEF (Figs. 10C and 11C). These analyses indicate that the Z -prediction method would not perform up to expectation on these microscopes, presumably due to imperfect stage behavior.

6. Other program features

6.1. Montaging

Some projects require reconstructions of areas larger than can fit in a camera field at the necessary resolution, even with a 4K CCD camera. To answer this need, SerialEM can acquire images at an array of overlapping positions to form a large, montaged image. It uses image shift to move to each position in turn then acquire and save an image.

SerialEM implements two features to alleviate the difficulty of working with an area larger than the field of the camera. First, it composes a single montaged image by subsampling pixels from all of the frames, which allows one to visualize the entire area captured. The default subsampling produces an image no bigger than $1\text{K} \times 1\text{K}$, but the user can select lower sampling to see the overview at higher resolution (e.g., Fig. 1). Second, there is a “prescan” mode in which frames are not saved and are acquired as rapidly as possible using the highest binning that will yield a montage of at least 1K pixels in the longest dimension. This method relatively quickly shows the exact area being captured, which would be difficult to do by taking a lower magnification picture.

When the montage contains an even number of frames in one or both dimensions, there is no frame at the center of the image. To avoid having to take an extra image to use for alignment, the program composes a center image, the same size as one of the frames, using the appropriate portions of two or four overlapping

frames. This image is suitable for alignment because the frames are shifted into registration using the analysis of overlaps between frames described below.

Tilt series acquisition with montaging is straightforward and essentially the same as with single-frame images. Of course, it takes longer to acquire montages, but the yield of data per unit time is somewhat higher, because tilting and any needed focusing are done only once per montaged image. For example, one user on our F20 acquired numerous $\pm 60^\circ$ tilt series at 1.5° intervals; the single-frame series took ~ 18 min, 1×2 montages took ~ 32 min, and 1×3 montages took ~ 41 min. Tracking can go less smoothly when montaging because the longer time between tilts allows drift to make a larger contribution to the specimen displacements upon tilting.

To be useful for tomography, the separate frames of a montage must be shifted into precise registration where they overlap and fused into a single image. The required frame shifts are determined by using correlation to assess the relative shift between images in each individual overlap zone. SerialEM saves all of the frames individually so that these operations can be performed later with the IMOD package. This approach has two advantages: it allows occasional errors in correlation-based frame registration to be corrected manually; and it allows the final aligned images of the tilt series to be produced directly from the raw data, with only one interpolation step.

6.2. Energy filtering

Energy filtering has been found to be valuable for imaging both thick frozen-hydrated specimens (Grimm et al., 1997; Koster et al., 1997) and stained plastic sections (Bouwer et al., 2004; Han et al., 1996). Filtering out inelastically scattered electrons and forming an image from zero-energy-loss electrons generally improves contrast and signal-to-noise ratio. Filtering is accomplished by spreading electrons out based upon their energies and introducing a slit that passes electrons in a selected energy range. SerialEM provides the basic controls necessary for using an energy filter: slit insertion, slit width, and slit offset from the zero-loss peak. It also implements solutions for several problems that arise when working with energy filters.

One problem is that the zero-loss peak alignment can shift by ~ 10 eV when changing magnification. This is not only inconvenient for users when they align the peak then change magnification, but it also can cause failures in the various operations that switch to lower magnification. The latter difficulty could be solved by removing the slit at lower magnification, but then the program would have to anticipate the potentially much larger intensity of the unfiltered beam. Instead, a procedure was developed to calibrate the energy shift of the zero-loss peak across magnifications, by scanning a wide range of energies (48 eV) with a narrow slit width to

localize the peak at each magnification. With this calibration, whenever the magnification differs from that at which the zero-loss peak was aligned, an appropriate adjustment factor is automatically applied to the user's selected energy loss. (To keep track of changes, the program detects when the zero-loss peak is aligned through DigitalMicrograph.) Since the slit on the GIF opens from one side, changing its width also changes its center, so the program adds an adjustment for width changes as well.

A second problem is that the zero-loss peak can drift out of the slit over the course of a tilt series that takes more than 20–30 min. A routine was thus developed to refine the alignment of the zero-loss peak using a minimal dose. Starting at the last known position of the zero-loss peak and stepping to negative energy offsets, it takes highly binned, short-exposure images with a moderately narrow slit (20 eV) and seeks the point where the recorded image intensity drops the fastest. It stops when both the image intensity and the change in intensity are sufficiently smaller than their peak values. If these conditions are not satisfied, the search is restarted at -20 eV then 20 eV from the original position. Once the peak position is found, it is added to the adjustments applied to the user's offset when the slit is in. The procedure is accurate to ~ 0.5 eV and works surprisingly well even for very thick samples where the zero-loss peak contains only a few percent of the electrons. This routine can be run automatically at a selected interval during a tilt series; a 15 min interval has been adequate in our experience.

The GIF can be difficult to work with because it has an internal magnification of 20–40 times. The Tecnai has a special lens mode, EFTEM, that provides low magnifications with the objective lens at normal strength, but in this mode the beam and specimen are very small on the viewing screen. SerialEM provides several features to alleviate these inconveniences. It can automatically switch between regular and EFTEM mode when the screen is raised. During this switch, it can change to the magnification that gives the closest match between pixel sizes on cameras before and after the GIF.

6.3. Grid mapper and navigator

Finding specimen areas of interest can be as time-consuming as acquiring data from them. To improve the efficiency of this task, SerialEM includes a Navigator window that provides controls for recording, annotating, selecting, and returning to stage positions, transforming positions when the grid position changes in the specimen holder, and acquiring maps of a grid or of selected areas at desired magnifications. Maps are images at a known stage position that can be reloaded from disk file and redisplayed simply by clicking on an entry in the Navigator's list. The mapping feature uses

the montaging capabilities of the program, which required several enhancements to be suitable. The primary change was to allow montages to be acquired with stage movements instead of with image shift; in addition, the montage overview of a map is composed without subsampling so that the user can inspect an entire map on the screen at its full resolution.

The Navigator is meant to be used at two levels of resolution, to avoid the time and computer resources required to acquire and manage maps of large areas at high resolution. Thus, one first acquires a map of up to $10K \times 10K$ pixels at a low enough magnification (~ 300 – $600\times$) to cover the region of interest. On this map, one can draw borders around an area to be captured by a montage at an intermediate magnification (e.g., $\sim 3000\times$). One can also mark a series of additional locations at which intermediate maps of that size are desired. A series of montages can then be acquired automatically at these marked locations. These maps are then inspected by the user to identify, mark, and move the stage to features of interest.

6.4. Program setup and calibration

In addition to basic calibrations such as for image shift and autofocus, the program has a relatively large number of other calibrations and parameters settings; many of them associated with the special features and solutions to problems described above. Complete program setup is thus rather daunting, but basic function for tilt series acquisition can be achieved with a more palatable subset, as described in the setup instructions. Calibrations are meant to be done by a program administrator and are saved in a file that is accessed by all users. On our Tecnaï, it has not been necessary for general users to redo calibrations, except for image shift at one magnification when montaging and absolute beam intensity when dose calibration is desired. Calibrations have been quite stable over time, with subsets of them needing to be redone only upon rare events such as a camera change or a gun or column realignment associated with major service. Program maintenance should thus take only ~ 1 – 2 days/year of operation.

7. Discussion

The method described here for predicting specimen position during a tilt series allows for rapid but robust data acquisition under a variety of conditions. By assessing whether a prediction is reliable, it is possible to skip tracking and focusing steps only when appropriate and to fall back toward the traditional approach of tracking and focusing on every step when necessary. This approach is about as fast as other methods for achieving rapid tilt series acquisition and is reliable in cases where

microscope and specimen behavior are not optimal. The method has been implemented in a mature and richly featured program which is currently producing about 1100 tilt series per year (~1 TB of raw image data) on one of our microscopes.

Some automated tomography software starts at 0° and acquires the tilt series in two directions (Fung et al., 1996; Zheng et al., 2004). SerialEM's approach of taking the series from one extreme angle to the other has the advantage that it produces a set of images in the right order, with consistent changes from one image to the next. It also allows the user to verify that imaging conditions are adequate for the highest tilt angles at the start of the tilt series. Starting a series at the highest tilt angle is more demanding because autofocus is less reliable, the specimen movements upon tilting are larger, and it can be difficult to center the feature of interest. However, these are minor problems for typical tilt series, and the Walk-up procedure for getting to high tilt alleviates the latter difficulty. Nevertheless, for challenging situations it might be advantageous to start the series from zero, as was done for the very high magnification series shown in Fig. 6. This strategy requires that the data be restacked into the proper order and that the user take care not to assume any continuity across 0° when aligning the images. In particular, if there is any stage backlash in tilting, the tilt interval will be less than the nominal amount at the transition point.

Changing beam intensity to compensate for changing specimen thickness during a tilt series may appear problematic because it changes the amount of beam convergence and requires a deviation from parallel illumination. Beam convergence places the tilted specimen in a magnification gradient (Fan et al., 1995), which means that the specimen is tilting through a changing distortion field when the beam intensity is changed during the series. On the Tecnai microscopes, parallel illumination cannot generally be achieved at a useful beam intensity, so always doing tomography under parallel illumination conditions is not a practical option. Magnification gradients and image rotation as a function of Z-height have been assessed for different condenser lens settings and found to be much smaller than the effects reported by Fan et al. (1995) on a JEOL 4000, generally less than 0.2% magnification change and ~0.02° rotation per μm of Z-height. Corrections for these distortions have been implemented in IMOD, but the effects are apparently small enough that correcting them does not give consistent or significant improvements in tilt series alignment. It thus appears that changing intensity, which is done for the practical reasons given above, is an acceptable procedure on a Tecnai. Nevertheless, it would be desirable to implement periodic changes in exposure time to minimize the effects of changing intensity. To reduce the problems with changing exposure time

when working with plastic sections that require pre-exposure, it would also be helpful to develop automatic tuning of the pre-exposure needed to get a high quality image.

SerialEM is freely available for nonprofit organizations at <http://bio3d.colorado.edu/SerialEM>, which also has further details on program features and a copy of the program's on-line help system.

Acknowledgments

I thank Ken Downing, J. Richard McIntosh, Daniela Nicastro, and Eileen O'Toole for critically reading the manuscript, and Brad Marsh and Adam Costin for providing tilt series logs. James Kremer wrote some utility classes for the HVEM version of SerialEM and Tobin Fricke modified the program to run on a JEOL. This work was supported by NIH/NCRR Grant RR00592 and NIH Program Project Grant P01GM61306 to J. Richard McIntosh.

References

- Baumeister, W., 2002. Electron tomography: towards visualizing the molecular organization of the cytoplasm. *Curr. Opin. Struct. Biol.* 12, 679–684.
- Bouwer, J.C., Mackey, M.R., Lawrence, A., Deerinck, T.J., Jones, Y.Z., Terada, M., Martone, M.E., Peltier, S., Ellisman, M.H., 2004. Automated most-probable loss tomography of thick selectively stained biological specimens with quantitative measurement of resolution improvement. *J. Struct. Biol.* 148, 297–306.
- Dierksen, K., Typke, D., Hegerl, R., Baumeister, W., 1993. Towards automatic electron tomography. II. Implementation of autofocus and low-dose procedures. *Ultramicroscopy* 49, 109–120.
- Dierksen, K., Typke, D., Hegerl, R., Koster, A.J., Baumeister, W., 1992. Towards automatic electron tomography. *Ultramicroscopy* 40, 71–87.
- Fan, G.Y., Young, S.J., Miller, P., Ellisman, M.H., 1995. Conditions for electron tomographic data acquisition. *J. Electron Microsc.* (Tokyo) 44, 15–21.
- Frank, J., 1992. *Electron Tomography*. Plenum Press, New York.
- Fung, J.C., Liu, W., de Ruijter, W.J., Chen, H., Abbey, C.K., Sedat, J.W., Agard, D.A., 1996. Toward fully automated high-resolution electron tomography. *J. Struct. Biol.* 116, 181–189.
- Gluckenberger, R., 1982. Determination of a common origin in the micrographs of tilt series in three-dimensional microscopy. *Ultramicroscopy* 9, 167–174.
- Grimm, R., Barmann, M., Hackl, W., Typke, D., Sackmann, E., Baumeister, W., 1997. Energy filtered electron tomography of ice-embedded actin and vesicles. *Biophys. J.* 72, 482–489.
- Han, K.F., Gubbens, A.J., Sedat, J.W., Agard, D.A., 1996. Optimal strategies for imaging thick biological specimens: exit wavefront reconstruction and energy-filtered imaging. *J. Microsc.* 183 (Pt. 2), 124–132.
- Koster, A.J., Braunfeld, M.B., Fung, J.C., Abbey, C.K., Han, K.F., Liu, W., Chen, H., Sedat, J.W., Agard, D.A., 1993. Towards automatic three dimensional imaging of large biological structures using intermediate voltage electron microscopy. *MSA Bull.* 23, 176–188.

- Koster, A.J., Chen, H., Sedat, J.W., Agard, D.A., 1992. Automated microscopy for electron tomography. *Ultramicroscopy* 46, 207–227.
- Koster, A.J., de Ruijter, W.J., 1992. Practical autoalignment of transmission electron microscopes. *Ultramicroscopy* 40, 89–107.
- Koster, A.J., Grimm, R., Typke, D., Hegerl, R., Stoschek, A., Walz, J., Baumeister, W., 1997. Perspectives of molecular and cellular electron tomography. *J. Struct. Biol.* 120, 276–308.
- Koster, A.J., van den Bos, A., van der Mast, K.D., 1987. An autofocus method for a TEM. *Ultramicroscopy* 21, 209–222.
- Kremer, J.R., Mastronarde, D.N., McIntosh, J.R., 1996. Computer visualization of three-dimensional image data using IMOD. *J. Struct. Biol.* 116, 71–76.
- Lawrence, M.C., 1992. Least-squares method of alignment using markers. In: Frank, J. (Ed.), *Electron Tomography*. Plenum Press, New York, pp. 197–204.
- Leapman, R.D., 2004. Novel techniques in electron microscopy. *Curr. Opin. Neurobiol.* 14, 591–598.
- Mastronarde, D.N., 1997. Dual-axis tomography: an approach with alignment methods that preserve resolution. *J. Struct. Biol.* 120, 343–352.
- Mastronarde, D.N., 2003. SerialEM: a program for automated tilt series acquisition on Tecnai microscopes using prediction of specimen position. *Microsc. Microanal.* 9 (Suppl. 2), 1182–1183.
- Mastronarde, D.N., Marsh, B.J., Otegui, M., 2001. Automated montaging HVEM tomography of large cellular volumes. *Microsc. Microanal.* 7, 90–91.
- McIntosh, R., Nicastro, D., Mastronarde, D., 2005. New views of cells in 3D: an introduction to electron tomography. *Trends Cell Biol.* 15, 43–51.
- Nickell, S., Forster, F., Linaroudis, A., Net, W.D., Beck, F., Hegerl, R., Baumeister, W., Plitzko, J.M., 2005. TOM software toolbox: acquisition and analysis for electron tomography. *J. Struct. Biol.* 149, 227–234.
- Rath, B.K., Marko, M., Radermacher, M., Frank, J., 1997. Low-dose automated electron tomography: a recent implementation. *J. Struct. Biol.* 120, 210–218.
- Steven, A.C., Aebi, U., 2003. The next ice age: cryo-electron tomography of intact cells. *Trends Cell Biol.* 13, 107–110.
- Zheng, Q.S., Braunfeld, M.B., Sedat, J.W., Agard, D.A., 2004. An improved strategy for automated electron microscopic tomography. *J. Struct. Biol.* 147, 91–101.
- Ziese, U., Geerts, W.J., Van Der Krift, T.P., Verkleij, A.J., Koster, A.J., 2003. Correction of autofocusing errors due to specimen tilt for automated electron tomography. *J. Microsc.* 211, 179–185.
- Ziese, U., Janssen, A.H., Murk, J.L., Geerts, W.J., Van der Krift, T., Verkleij, A.J., Koster, A.J., 2002. Automated high-throughput electron tomography by pre-calibration of image shifts. *J. Microsc.* 205, 187–200.