Forward modeling method for microstructure reconstruction using x-ray diffraction microscopy: Single-crystal verification

R. M. Suter
Department of Physics and Department of Materials Science and Engineering, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

D. Hennessy and C. Xiao
Department of Physics, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

U. Lienert
Advanced Photon Source, Argonne National Laboratory, Argonne, Illinois 60439

We describe and illustrate a forward modeling method for quantitatively reconstructing the geometry and orientation of microstructural features inside of bulk samples from high-energy x-ray diffraction microscopy data. Data sets comprise charge-coupled device images of Bragg diffracted beams originating from individual grains in a thin planar section of sample. Our analysis approach first reduces the raw images to a binary data set in which peaks have been thresholded at a fraction of their height after noise reduction processing. We then use a computer simulation of the measurement and the sample microstructure to generate calculated diffraction patterns over the same range of sample orientations used in the experiment. The crystallographic orientation at each of an array of area elements in the sample space is adjusted to optimize overlap between experimental and simulated scattering. In the present verification exercise, data are collected at the Advanced Photon Source beamline 1-ID using microfocused 50 keV x rays. Our sample is a thin silicon wafer. By choosing the appropriate threshold fraction and convergence criteria, we are able to reconstruct to ≤ 10 μm errors the subregion of the silicon wafer that remains in the incident beam throughout the rotation range of the measurement. The standard deviation of area element orientations is $\approx 0.3°$. Our forward modeling approach offers a degree of noise immunity, is applicable to polycrystals and composite materials, and can be extended to include scattering rules appropriate for defected materials. © *2006 American Institute of Physics.* [DOI: 10.1063/1.2400017]

I. INTRODUCTION

X-ray diffraction microscopy[1,2] constitutes a set of techniques that can be used to probe the microstructure deep inside of bulk materials. Because the measurements are nondestructive, they give access to the dynamic response of ensembles of grains to stimuli such as heat, stress, or chemistry. Using focused beams of high-energy photons with $E \geq 40$ keV that penetrate millimeters to centimeters of materials, these techniques can spatially resolve local crystallographic orientation, grain shapes, and strain states. While high-energy x rays are available at a number synchrotron beam lines around the world, the undulator radiation from third generation sources has the spectral range, brilliance, and source size most desirable for the experimental requirements. At present, there is an apparatus at ID-11 at the European Synchrotron Radiation Facility (ESRF) that is dedicated to such measurements[1] and, as illustrated here, a similar capability is being developed at the Advanced Photon Source. Studies to date have probed, for example, *in situ* response to stress,[3,4] internal grain growth in real time,[5,6] and observation of phase transformations.[7]

In this paper, we are concerned with the problem of reconstructing the location and shape of ensembles of grains in three dimensions. This is done by isolating and imaging diffracted beams originating from individual grains and tracking their path through space in order to locate their position of origin.[8,9] The position and shape resolution are expected to reach ~ 1 μm while orientations can be measured to well below a degree. However, the reconstruction problem is formidable even for ideal samples with sharp Bragg scattering and it is more so in the case of samples with a significant defect content.[1,9] We report here on a new approach involving forward modeling or computer simulation and we demonstrate its application using a simple, single-crystal sample of silicon. It should be possible to generalize this approach to treat not only simple polycrystals but defected and composite materials as well.

The development of a three-dimensional polycrystal mapping ability is timely. For example, recent work using electron probes[10,11] has provided new statistical measures of grain boundary character distributions in a variety of polycrystals and has revealed that systematic anisotropic distributions emerge from simple thermal grain growth treatments. Volumetric measurements are now needed to understand, on a grain-by-grain basis, how these distributions arise, how they are affected by different starting conditions, and how they can be influenced by defect content, impurities, inclusions, and other factors. Recent three-dimensional grain
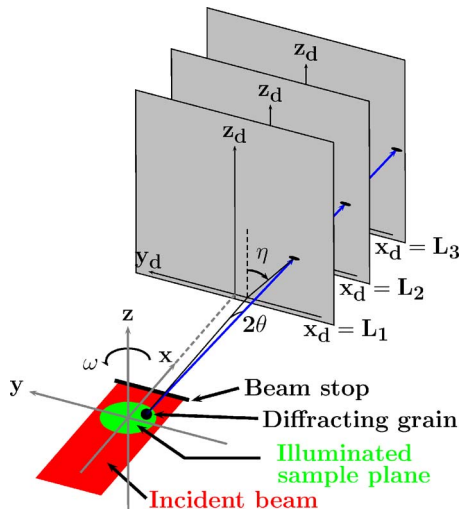
FIG. 1. Schematic of the x-ray beam, sample, and detector. We use a single detector and move it sequentially to the positions shown. Diffraction from one particular grain is shown together with notation (2θ and η) for specifying the direction of the diffracted wave vector, $\mathbf{k}_f = k\hat{\mathbf{v}}$, with $k = 2\pi/\lambda$. Note that a circular grain in the illuminated plane generates an elliptical spot on the detector due to the projection along $\hat{\mathbf{v}}$. The coordinate system origin is at the intersection of the rotation axis and the beam plane and the beam is incident in the $\hat{\mathbf{x}}$ direction.

growth simulations[12,13] using model driving forces appear to yield grain boundary character distributions that are statistically similar to experimental measurements. Measurements of the evolution of real materials can be used to test and constrain the growth rules used in such simulations.

Figure 1 shows a schematic of the experimental geometry. This is configuration B as described by Poulsen, and we use the same coordinate system conventions.[1] A line focused, monochromatic beam of x rays illuminates a thin section ($\sim 1 \ \mu$m thick by ~ 1 mm wide) of the sample that is mounted on a rotation stage, ω, with the axis normal to the illuminated plane. A beam stop prevents the transmitted beam from striking the downstream area detector. The high spatial resolution detector (for example, 4 μm pixels) images diffracted beams emanating from crystallites that happen to satisfy a Bragg condition at orientation ω. The detector is positioned close to the sample ($\sim 5-10$ mm) so that the position of a diffracted beam spot is sensitive to both the diffracted beam direction (specified by 2θ and η) and the position of origin within the sample. This information is encoded in the data by measuring the diffraction pattern at a number, $N_L \sim 3$, of rotation axis-to-detector distances, L. In order to observe as many diffracted beams as possible, continuous coverage of the rotation ω is needed. Detector images are therefore collected while the sample is rotated through an interval, $\delta\omega \sim 1°$, and N_ω adjacent intervals are measured so as to cover a range $\Delta\omega = N_\omega \delta\omega$ large enough to ensure observation of several Bragg peaks from grains of arbitrary orientation.

Another technique for nondestructive probing of internal microstructure is the polychromatic x-ray microbeam method developed by Ice and Larson.[14] This method, using a spot focused beam, has better (submicrometer) spatial resolution but has less penetration depth, making it appropriate for thin and small grained structures. Recent work, for ex-

ample, has characterized deformation microstructures of individual grains.[15] Volumetric data collection requires roughly an image per voxel whereas the current method collects data in transmission from planar sections through many grains in parallel.[8]

The next section gives a detailed account of our data analysis and map construction approach along with comparisons to other approaches. This is followed by a description of the apparatus at the Advanced Photon Source and a verification exercise using the simplest possible microstructure: a single crystal of silicon. Evaluation of the results, implications for applications to polycrystal samples, and prospects for further development are discussed in Sec. IV.

II. ANALYSIS METHOD

Reconstruction of a three-dimensional microstructure requires the processing of thousands of detector images. Each two dimensional slice requires 100–300 images and many slices are required to extend into three dimensions. In a polycrystal, each image may include 10–100 diffraction spots originating from different regions of the illuminated sample space. Since a primary purpose of the nondestructive measurement is to allow observation of the response of the microstructure to external stimuli, many three-dimensional maps may be needed. It is clear that analysis of such data sets must be extensively automated.

We describe here an analysis method in which we simulate or forward model the entire measurement process including the x-ray beam and detector, experimental geometry, and sample microstructure. The sample microstructure is described on a discretized grid representing the illuminated plane of the sample with each area element having an associated phase, orientation, and, possibly, defect structure. Software can adjust parameters describing both the sample and experiment to make the simulated diffraction match the observed data as well as possible. As presently implemented, this is an *ab initio* approach that requires no initial cross correlation of the diffraction images and no prior knowledge of orientations present in the sample. What is required is a reasonably well-defined set of experimental parameters and a known nominal crystal structure. This initial information must be sufficient that at least some regions in the sample can be made to generate scattering that overlaps a minimal set of experimentally observed diffracted beams. With this initial success, experimental parameters can be refined, as can the spatial resolution of the simulated or fitted microstructure. One advantage of this approach is that the simulation can use and adjust realistic scattering physics appropriate for the sample under consideration. Lattice bending that gives rise to a mosaic spread in the observed data can be modeled directly in the crystal frame of reference. Also, detector point spread functions can be modeled so as to avoid the complications of attempting to deconvolve observed data sets.

Our forward modeling approach is an alternative to several current analysis schemes[1,7,8,14–21] including GRAINDEX, GRAINSWEEPER, and the algebraic reconstruction technique (ART). To put our work in context, we

give a very brief comparison to other methods but refer the reader to the references for detailed descriptions. GRAIN-DEX is an indexing scheme that correlates observed diffraction spots into groups that are crystallographically consistent so as to yield a list of grain orientations with center-of-mass positions in the illuminated sample plane. ART combines the set of intensity patterns corresponding to a given orientation and uses tomographic algorithms to reconstruct the grain shape. GRAINDEX and ART work primarily in the data space to deduce properties of diffracting entities in the sample; in this sense, they attempt to solve the inverse problem. These approaches are quite sensitive to the overlap of spots as diffracted beams from different sample regions cross at the detector position[18] and to broadening effects due to defect scattering.

GRAINSWEEPER and our simulation approach analyze the possible scattering from localized sample regions and search for optimum orientations over a gridded sample space; these can be characterized as forward modeling or data fitting procedures. GRAINSWEEPER identifies detector regions where a primary Bragg peak type could appear and searches for intensity therein. The detector region for a second peak, crystallographically consistent with the first, is then searched. This constrained search is computationally efficient but requires the specification of preferred Bragg peak types that must be in the observed data set for all grains. Our approach is less computationally efficient but treats a large number of peaks equally. We use Monte Carlo optimization of orientations, which uses all the scattering generated by the simulation and attempts to minimize a global cost function. An occasional overlap of spots does not limit convergence. Further, Monte Carlo variation of scattering parameters and experimental geometry is trivial to implement (at least in its simplest form). As pointed out previously,[1,9] combining an efficient initial computation with a Monte Carlo based final step is quite desirable. We have implemented this to the extent that we first perform a search over a discrete orientation set before beginning the Monte Carlo optimization. Finally, other experimental configurations, such as use of a large or "box" beam,[1] can be treated with a generalization of our current model. The flexibility of the simulation approach combined with the availability of ever more computational power should allow simulations to play a significant role in the analysis of x-ray diffraction microscope data.

A. Sample simulation

Approximating the incident x-ray beam as illuminating a thin section of the experimental sample allows an analysis to be done on a layer-by-layer basis. The planar section of the microstructure is simulated as a mesh of area elements, each with independent parameters describing the local crystallographic orientation, a switch, ϕ, to indicate the type of scattering generated by the element, and other parameters describing the state of convergence in the fitting process (Sec. II D). A list of Bragg peaks that are candidates for being generated by each area element is specified. The area covered by the mesh may be smaller or larger than the actual sample being measured. ϕ can be a binary indicator of the presence or absence of scattering from the element or, more
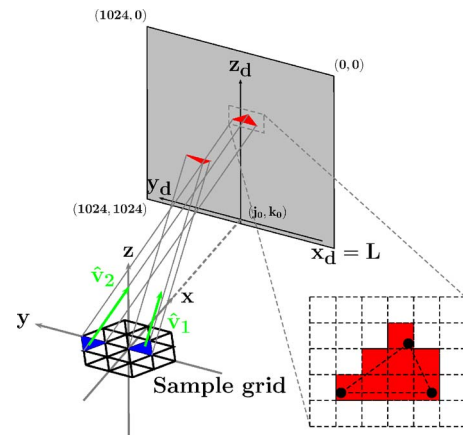

FIG. 2. Schematic diagram of the simulation geometry. The sample is represented by a hexagonal grid with each element assigned its own crystallographic phase and orientation, $(\phi_j, \hat{\mathbf{g}}_j)$. For elements that generate scattering at this ω (dark triangles), the grid element vertices are projected along a unit vector, $\hat{\mathbf{v}}$, parallel to the scattered wave vector, $\mathbf{k}_f$, onto the detector (black dots in magnified view at right). Any detector pixel inside or intersected by the projected triangle is considered to be illuminated by this Bragg peak (filled squares).

generally, can take on different values to indicate the presence of a particular crystallographic phase as appropriate for measurements of composite materials.

In our implementation, the sample mesh is a triangular grid of area elements (Fig. 2). The initial mesh can be adaptively refined by dividing selected parent triangles into four smaller ones with the offspring having the same properties as the parents. The fitting procedures can then determine optimized parameters so as to better define boundaries where, for example, orientations change rapidly. Crystallographic orientation $\hat{\mathbf{g}}$ (specified in the sample reference frame) is relative to unit cells being aligned with the laboratory axes of Fig. 1 when $\omega = 0$. The list of reciprocal lattice vectors, $\mathbf{G}_{hkl}$, is generated from elementary specification of unit cell dimensions, basis atom positions, and atomic form factors (approximated in the small scattering vector limit). A maximum G_{hkl} (i.e., 2θ) can be specified and peaks with intensities greater than a specified fraction of the maximum intensity are put into the list of candidates. This approach makes it straightforward to model any type of crystal, from elemental to complex molecular.[22]

B. Scattering calculations

For a given set of parameters, $(\hat{\mathbf{g}}, \phi)$, an area element will generate Bragg scattering in several different ω intervals. The simulation needs to determine which candidate peaks are generated in experimentally measured intervals and, of these, which diffracted beams actually hit the detectors at each position, L. The geometry of the area element is then projected onto the relevant detectors to determine the detector pixels that are illuminated. In the fitting process, we then ask whether the corresponding experimental pixels are illuminated. The detector space is four dimensional and can be parametrized as $D(j, k, i_L, i_\omega)$, where j and k are pixel indices, i_L picks one of the detector distances (Fig. 1), and i_ω specifies the ω interval in which the Bragg peak falls.

In the simplest case, the Bragg condition is evaluated and scattering is assumed to be sharp, i.e., a δ function at

$$\mathbf{Q} \equiv \mathbf{k}_f - \mathbf{k}_i = \mathbf{G}_{hkl}, \tag{1}$$

with $\mathbf{k}_f$ and $\mathbf{k}_i$ being the final and incident wave vectors. However, many generalizations are possible and desirable: the incident beam may not be perfectly monochromatic or parallel, and the crystal structure may be distorted by homogeneous strain or local defect content. If the consequent motion and/or broadening of Bragg scattering can be parametrized, then these effects can be included here and adjusted by fitting procedures.

For a particular $\mathbf{G}_{hkl}$ (specified in conventional crystal notation) the Bragg condition (1) in the laboratory frame of reference determines ω and η (the scattering angle 2θ being already defined by $G_{hkl} = 2k \sin \theta$):

$$(\mathbf{G}_{hkl}^{lab})_x = -\frac{|\mathbf{G}_{hkl}|^2}{2k}, \tag{2}$$

where the incident wave vector is $\mathbf{k}_i = k\hat{x}$ and $\mathbf{G}_{hkl}^{lab} = \Omega^{-1}(\omega)\mathbf{U}^{-1}(\hat{\mathbf{g}})\mathbf{G}_{hkl}$. Here, $\mathbf{U}^{-1}(\hat{\mathbf{g}})$ is the matrix giving the crystal lattice orientation relative to the sample coordinates and $\Omega(\omega)$ is the matrix corresponding to rotation by ω about the $\hat{z}$ axis. Equation (2) yields either zero or two Bragg values, ω_B. If $\mathbf{U}^{-1}(\hat{\mathbf{g}})\mathbf{G}_{hkl}$ lies too close to the rotation axis, the component in the xy plane will be too small to ever satisfy (2); the condition for a solution is that $\chi \gtrsim \theta_B$, where χ is the angle between the z axis and $\mathbf{U}^{-1}(\hat{\mathbf{g}})\mathbf{G}_{hkl}$ and θ_B is the Bragg angle. At high energies, many Bragg angles are small, so a large region of reciprocal space is covered by the single ω rotation in combination with an area detector. For $\chi > \theta_B$, there are two solutions to (2) having the in-plane projection of $\mathbf{U}^{-1}(\hat{\mathbf{g}})\mathbf{G}_{hkl}^{lab}$ on either side of the incident beam.

Once a relevant Bragg condition has been determined, the simulated detector image is updated. The triangle vertices are moved to laboratory coordinates with $\Omega^{-1}(\omega)$. The area element is then projected along $\mathbf{k}_f = k\hat{v}$ onto the detectors for the appropriate ω and L and the detector images are updated following the procedure outlined in Fig. 2.

We have incorporated a number of correction parameters to mimic possible experimental realities. Misalignment of the detector translation is included by specifying the detector origin, (j_0, k_0), separately at each detector distance, L. Rotation stage (x, y) offsets can be entered for each ω interval. Detector plane misorientation is specified by rotations about the laboratory axes. The finite incident x-ray beam height is incorporated by summing contributions from different z elevations.

Finally, here we illustrate broadening effects due to incident beam properties. These properties are governed by the monochromator and focusing optics and can significantly affect observed scattering under realistic conditions. The effects illustrated here can be difficult to handle in data inversion approaches to analysis. We model the finite energy bandwidth, ΔE, and variations in the angle of incidence relative to the nominal beam plane, $\delta\chi$, by summing a set of discrete contributions. Figure 3 compares three incident beam models: (i) $\Delta E = \delta\chi = 0$ (ideal case), (ii) $\Delta E/E = 0.02$, with $\delta\chi = 0$, and (iii) $\Delta E/E = 0.02$ with $\delta\chi \approx 0.1°$ with χ be-

ing correlated with E in a manner appropriate for a bent silicon (111) Laue monochromator crystal.[23,24] $\Delta E/E$ is chosen as 2% in order to make effects visible in the images. Model (i) yields triangular diffraction spots that are projections, at the appropriate ω, of the triangular sample area element. The higher order spots are more extended in the z direction due to the larger projection angle. The broadened spots of model (ii) are the result of different diffraction angles for the different energy contributions. Spots progressively broaden as the sample-to-detector distance increases. The more obvious effect is that some peaks cross over to different ω intervals as the energy varies: the high order peak near the center of (b) is spread over all three images (d)–(f). The split spots appear rather sharp in (d) and (f), indicating that only a narrow part of the bandwidth contributes here. Other such split peaks can also be identified by comparison with (a)–(c). For model (iii), not surprisingly, additional broadening and crossings of ω intervals occurs.

C. Raw data and image analysis

The raw experimental data is in the form of images of diffraction patterns observed with a two-dimensional x-ray detector. Detector hardware and characterization are discussed in Sec. III A. Here, we describe the process of taking raw charge-coupled device (CCD) based images and generating a binary reduced data set that yields geometrical information about the diffracting entities in the sample. In other words, from the continuous intensity distributions in observed diffraction spots, we attempt to deduce shapes on the detector that correspond to projections of the shapes of individual crystalline grains. Because the reconstructed orientation of each area element in the sample depends on matching its diffraction to many Bragg spots, some random errors in deduced shapes can be tolerated. However, in the averaged sense, accurate shapes are critical.

Due to the combination of data collection time constraints and the inefficiency of the detector system (Sec. III A), signal levels can be quite low, and this complicates the above process. Figure 4 shows a typical diffraction pattern obtained from an aluminum polycrystal. Even with better statistics on low-order Bragg peaks, we expect to have to deal with weak signals in higher-order peaks. Because of their larger scattering angle, these peaks tend to have better projection geometries than low-order peaks, so deducing their shapes accurately can yield improved spatial resolution in the reconstruction. Thus, image processing is a critical and challenging aspect of the reconstruction process.

The image reduction process is illustrated in Figs. 5 and 6, which show a subregion of the image in Fig. 4. The uniform background in Fig. 4 or Fig. 5(a) is due to dark current in the CCD. Simple background subtraction yields improved contrast but leaves many isolated "hot" pixels. These anomalous points must be removed in order to allow consistent intensity estimates within the diffraction spots. Figure 5(c) demonstrates that this can be done using a median filter. The filter passes a 3×3 pixel mask over the image and replaces the center pixel intensity with the median of values within the mask. This nonlinear process has the effect of removing isolated high (or low) pixels without strongly affecting
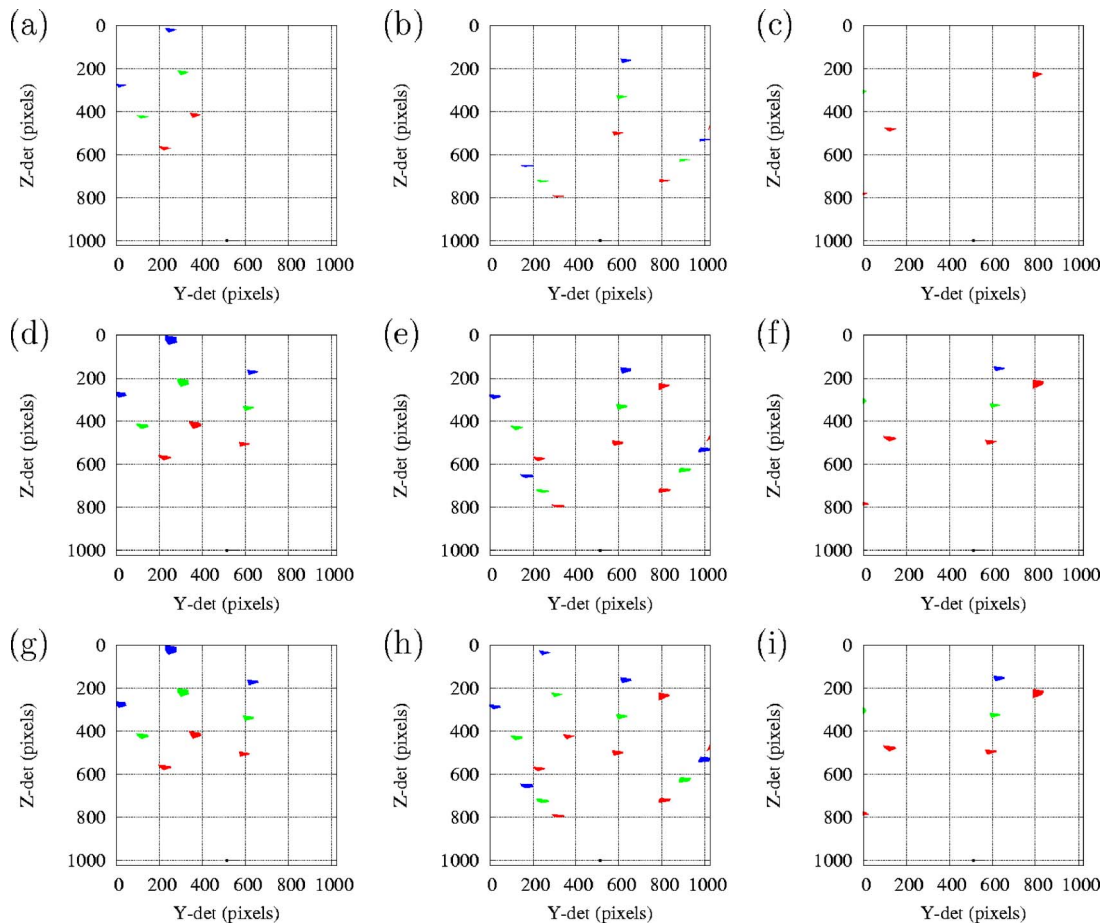
FIG. 3. Demonstration of imaging of a single triangular area element (200 μm side length) and the effect of incident beam models (see text). Each image contains scattering generated over a $\delta\omega=1°$ interval as observed at three sample-to-detector distances (6, 8, 10 mm in red, green, and blue, respectively); the detector covers a 4×4 mm^2 area. The nominal energy is 50 keV. The projection of the incident beam is shown near the bottom (centered on pixel row 1000). Each row of images shows intensity at the same three successive ω positions. (a)–(c) show model (i), (d)–(f) model (ii), and (g)–(i) model (iii).

smooth intensity variations. The next step is to identify, in images like Fig. 5(c), sets of contiguous pixels with above zero intensity. These sets are referred to as "regions of interest" or ROIs. Each ROI is searched for its maximum intensity and a first estimate of the spot shape is determined by
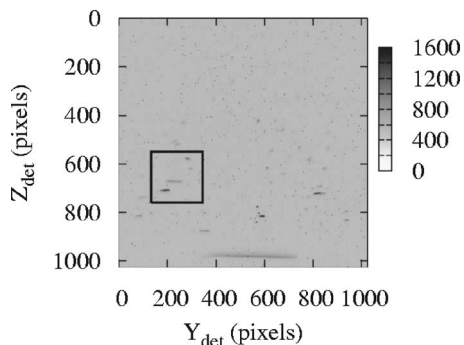

FIG. 4. A typical detector image obtained from an aluminum polycrystal. The gray scale is in CCD counts per pixel. The uniform background is slightly under 600 counts. Diffraction spots are horizontally extended (10–50 pixels) darker regions. The direct beam projects onto the image at $z_{det}=1004$, below the long horizontal streak that is stray scattering out of the direct beam. Note that the most intense points (those that saturate the gray scale) are isolated "hot" pixels rather than diffraction spots. The maximum intensity in the image is 3997. The following figures show the effects of image processing steps on the boxed region.

keeping only pixels with intensity above some fraction, f^{peak}, of the peak value. Figure 6 illustrates the effect of choosing different values for f^{peak}. This choice is clearly critical for obtaining correct grain shapes. Because the intensity distribution within a diffraction spot can be affected by strain and lattice bending effects or by overlap with other diffraction intensity, there is no universally correct value for f^{peak}. This situation motivates careful validation measurements and sample-by-sample consistency checks. Future work will also include a comparison to other cutoff criteria such as the steepest gradient condition.[9] Future incorporation of models for scattered intensity distributions (with adjustable fitting parameters) may remove the necessity for binary data sets altogether.

D. Fitting procedure

This section describes algorithms that attempt to determine the microstructure corresponding to an observed data set. Using the gridding of sample space (Sec. II A) and the diffraction simulation (Sec. II B) described above, we seek orientations for each area element that generate simulated intensity that optimally overlaps experimental intensity (Sec. II C). The sample space grid is then adaptively refined in
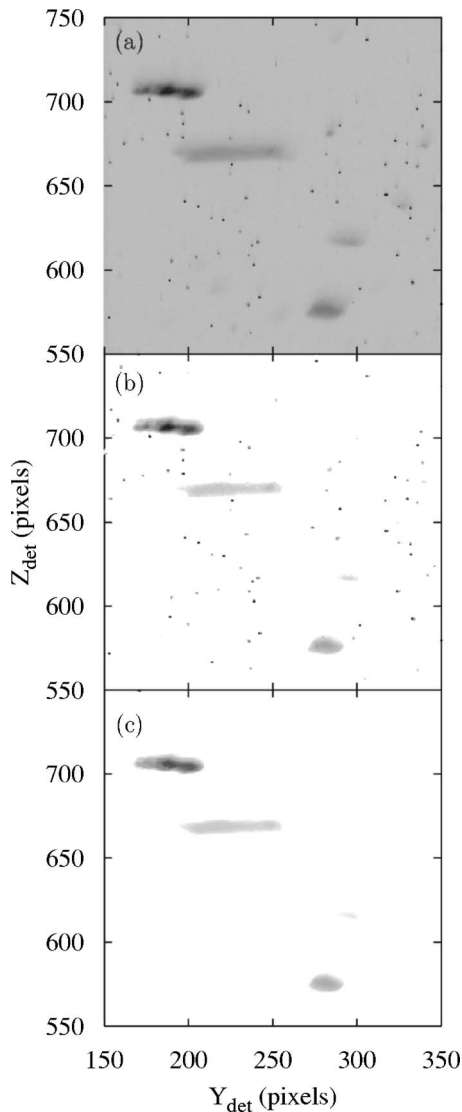
FIG. 5. Expanded view of the subregion indicated in Fig. 4 in three stages of image analysis. (a) Raw image (gray scale 0–1600; this scale saturates some pixels since the maximum in the figure is 2700), (b) background (600 counts) subtracted image (gray scale 0–1000), and (c) median filtered and background subtracted image (gray scale 0–1000).
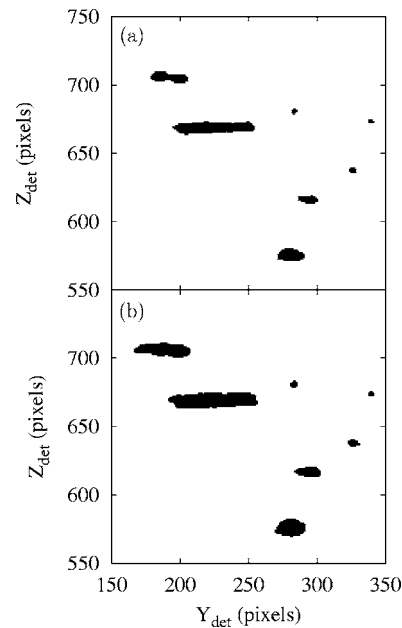

FIG. 6. Binary images of data from Fig. 5. Each region of interest has been thresholded at a fraction of the maximum intensity in the region. Thresholds are (a) $f^{\text{peak}}=0.5$ and (b) 0.25 of the peak height.

order to better resolve regions where properties change rapidly.

Data, both experimental and simulated, are stored in an array, $D(j,k,i_L,i_\omega)$, as defined in Sec. II B. The low-order bit holds the reduced experimental data while higher-order bits store simulated intensity. Of order 300 mega-bytes is required for a single layer. Typically, less than 1% of detector pixels are included in the reduced experimental data set and a comparable number of pixels are struck by the simulation.

As the orientation associated with a particular sample area element is varied, a given Bragg peak, with scattering angle, 2θ, sweeps through ω and η (Fig. 1). For some orientations, it disappears out of the experimentally accessed subspace of these variables. As it moves, this simulated peak may accidentally overlap experimental intensity coming from another part of the sample but it is unlikely to do so at all measured detector positions, L, at the relevant ω. Only for a very restricted range of orientations (for the given area

element) will multiple simulated peaks overlap experimental spots at multiple L's. This sort of overlap is the goal of the orientation search. Below, an acceptance criterion, F_{min}, is developed that measures the level of this success.

1. F_{min} acceptance criterion

A hierarchy of procedures is necessary to determine if an orientation qualifies as a valid fit for an area element. At each orientation, $\hat{\mathbf{g}}$, we calculate the fraction, $F(\hat{\mathbf{g}})$, of qualifying simulated Bragg peaks (see below) that overlap experimentally observed intensity. The basic question is, does a sufficient fraction, F_{min}, of qualifying simulated Bragg peaks overlap experimentally observed intensity? Determination of $F(\hat{\mathbf{g}})$ involves the following considerations:

1. A "qualifying simulated Bragg peak" is one that strikes the detector at least at the two smallest measured L positions (unless only a single L was measured).
2. For a simulated Bragg peak to be said to overlap experimentally observed intensity in a given detector image, it must satisfy $f^{\text{hit}} \equiv n_{\text{exp}}/N_{\text{sim}} \geq f^{\text{hit}}_{\text{min}}$. As illustrated in Fig. 7, N_{sim} is the number of detector pixels covered by the projection of the sample area element onto the detector (Fig. 2) while n_{exp} is the number of those pixels that are also illuminated in the experiment. $f^{\text{hit}}_{\text{min}}$ is an acceptance criterion typically set to 0.5.
3. For a qualifying simulated peak to be said to overlap experimental intensity, it must satisfy the criterion in item 2 at every detector position (at the relevant ω) with $L \leq L^{\text{sim}}_{\text{max}}$, where $L^{\text{sim}}_{\text{max}}$ is the maximum measured distance at which the simulated peak is on the detector. For example, if the scattering shown in Fig. 1 is generated by the simulation and it hits the detector at all three positions as shown, then it must overlap experimental intensity at all three detector positions to qualify as overlap-
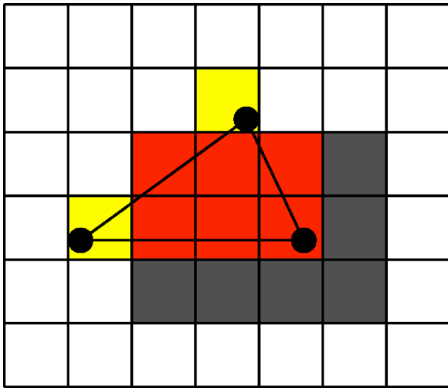
FIG. 7. An illustration of the counting scheme used in determining overlap between simulated scattering and the experiment. The square grid shows pixels in a small region of the detector. The black dots represent the projection of vertices of a triangular sample area element. Experimentally illuminated pixels not hit by the simulation are black; pixels struck in the simulation but not the experiment are yellow; overlapping intensities are red. In this example, $N_{sim}=8$ (yellow plus red), $n_{exp}=6$ (red), so $f^{hit}=0.75$. The contribution to the Monte Carlo cost function [Eq. (3)] from this region is $\delta\chi=12+8-(2\times6)=8$, which is the number of "mistakes" (black and yellow pixels). Ideally, a finer sample space grid will remove simulated intensity from the yellow pixels and neighboring area elements will cover the remaining black pixels.

ping the data. If it had hit the detector at only the closest two positions, then it would have to overlap the experiment at both of those positions. If it was on the detector only at L_1, then it would not be a "qualifying simulated Bragg peak."

4.   For orientation $\hat{\mathbf{g}}$, the number of qualifying Bragg peaks (out of the full list of those being simulated—see Sec. II B) is calculated and compared to the number that overlap experimental intensity to compute $F(\hat{\mathbf{g}})$.

2. Nature of the search space and overview of search logic

The space of orientations of an asymmetric body is spanned by 4π sr specifying the orientation of an axis times 2π rad rotation about that axis, so the total space is $8\pi^2$ rad^3. Symmetries reduce the relevant volume. For example, for cubic lattices, with 24-point symmetry operations, the space can be divided into 24 zones, each of which contains all physically distinguishable orientations. A search over a single "fundamental zone" includes all possible scattering patterns from a specified crystal structure. We choose the zone with the smallest rotation angles (in an axis-angle representation) from the reference orientation specified by Euler angles (0, 0, 0). Thus, a cubic crystal which, at $\omega=0$, has its axes aligned with those of Fig. 1 is considered "unrotated."

To find an orientation with $F(\hat{\mathbf{g}})>F_{\min}$ requires searching in the three-dimensional space of orientations. Over the vast majority of this space, very little overlap of simulation with experiment occurs. As mentioned in Sec. I, an angular resolution on the order of a degree implies that a fractional subvolume of the space of orientations on the order of 10^{-6} must be located. Seeking to minimize a simple cost function (such as the one introduced in Sec. II D 4) without any starting knowledge is quite inefficient. To overcome this prob-
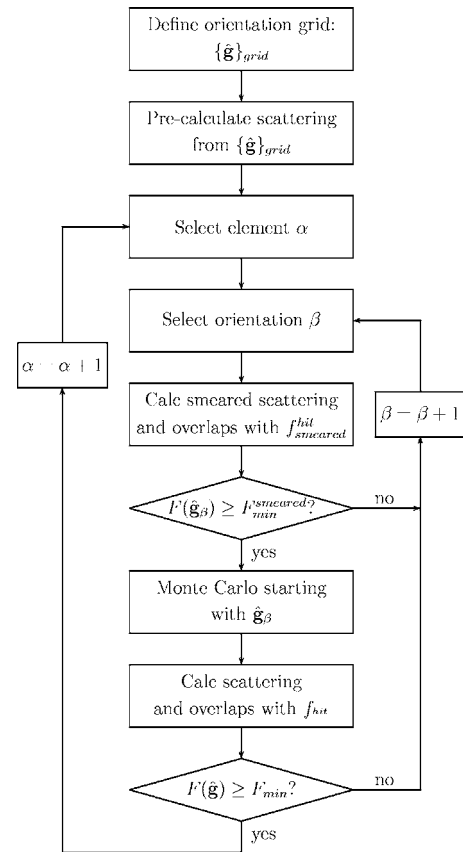

FIG. 8. Skeletal overview of fitting procedures. If β exceeds the number of members of $\{\hat{\mathbf{g}}\}_{grid}$, then element α is tagged as generating no scattering and the next element, $\alpha+1$, is processed.

lem, we first perform a search over a discrete grid of points while smearing the simulated scattering to cover the region between points in the grid. This search, with a much higher probability of accidental overlaps than described above, finds candidate orientations that are then checked by a Monte Carlo simulation that uses the scattering rules described above to home in on an optimal $\hat{\mathbf{g}}$. The basic flow of logic is shown in Fig. 8; the following sections describe the two search algorithms.

3. Orientation grid search

A uniform grid of orientations covering the fundamental zone for the relevant crystal structure is defined. Cubic symmetry and a spacing of $\Delta\psi=5°$ yields a set, $\{\hat{\mathbf{g}}\}_{grid}$, of about $(8\pi^2/24)\times(180/5\pi)^3=4950$ points. To save computing time, scattering for each orientation can be precalculated and limited to peaks occurring in the experimental range of parameters.[18] Further, it is practical to limit the Bragg peaks used in this search to those with $Q<Q_{\max}^{grid}$. These low angle peaks are more likely to appear in the data due to their relatively large intensities and they are more likely to hit detectors at several distances, L. In practice, we store for each relevant peak (i) the index (or indices) of the relevant ω interval(s), (ii) the unit vector, $\hat{v}$, giving the direction of the scattered beam in the laboratory frame, and (iii) the quantities, $\delta\tilde{y}_d=(\Delta\psi/2)\cos\eta\tan2\theta$ and $\delta\tilde{z}_d=(\Delta\psi/2)\sin\eta\tan2\theta$.

For each trial orientation for a given triangular area element, the vertices are projected along the relevant set of $\hat{v}$'s

at corresponding ω's to the measured detector distances. A rectangular bounding box is determined around the projected triangle and is extended by $\pm \delta y_d = \pm \delta \tilde{y}_d (L - x_s)$ in the y direction and similarly for the z direction. This generates a crude "smearing" to fill in the range of possible scattering corresponding to one dimension (rotation about the incident beam) of the volume associated with a point in the orientation grid. To fill the other relevant direction, this same smeared bounding box is searched in neighboring ω's spanning the range $\Delta \psi$ (the third orientation dimension corresponds to rotation about the scattering vector and is irrelevant for a single Bragg peak).

On a detector at a given ω and L, f^{hit}, as already defined, is calculated and compared to an acceptance parameter, $f^{\text{hit}}_{\text{smeared}}$, where the latter must be small compared to $f^{\text{hit}}_{\text{min}}$ due to the large search area. Correspondingly, a separate value, $F^{\text{smeared}}_{\text{min}}$ is used for the required fraction of simulated Bragg peaks hitting experimental intensity. As shown in Fig. 8, if an orientation is found for which $F(\hat{\mathbf{g}}) \geq F^{\text{smeared}}_{\text{min}}$, the grid search is suspended and Monte Carlo refinement of the orientation is performed as described in the next section. If this refinement ends without finding a qualifying orientation, then the grid search is resumed at the next orientation in the list. If the combined grid and Monte Carlo searches fail to find any qualifying orientation, then the crystallographic phase, ϕ, associated with this area element is set to zero signifying that no scattering is generated by this element.

4. Monte Carlo refinement

Once the orientation grid search has identified a candidate orientation, a Monte Carlo routine attempts to optimize the orientation based on scattering as defined in Sec. II B. Convergence criteria $f^{\text{hit}}_{\text{min}}$ and F_{min} are used to determine acceptable Bragg peaks and orientations, respectively.

F_{min} must be set appropriately for a given data set and fitting conditions. Generally, even for a good orientation, $F(\hat{\mathbf{g}}) < 1$ because the simulation generates scattering that is not observed in the data set. This can be due to blocking of low-order or $\eta \approx \pm \pi/2$ scattering by the beam block or high-order scattering for which insufficient counting statistics were obtained. Further, any errors in the experimental parameters used by the simulation or mechanical imperfections that are not mimicked by the simulation will generate less than optimal matching.

Our current Monte Carlo routine uses a simple, zero temperature minimization procedure based on a cost function, χ. Starting with a nominal orientation, a random increment is given to each of the orientation parameters. The decision to accept or reject the trial orientation is based on the function (see Fig. 7)

$$\chi = E + S - 2O, \tag{3}$$

where E is the total number of pixels in the entire experimental reduced data set, S is the total number of pixels hit by the simulated structure, and O is the number of pixels hit by both. A perfect fit to noiseless data would have $\chi = 0$. While χ is global (i.e., it depends on the entire simulated microstructure), a local variant associated with an area element, $s - 2o$, can be defined to facilitate parallel processing of area ele-

ments. Here s is the number of pixels struck by all the scattering from the element and o is the number of those pixels also in the experimental data set.

For each trial orientation of an element, the change, $\Delta \chi$, is computed. $\Delta \chi$ is always computed for two different orientations rather than involving the state with no scattering from the element. Putting in scattering from an incorrect orientation increases χ relative to the state with no scattering since it increases S without (presumably) increasing O; the minimization procedure then searches for orientations that bring χ back down by generating overlap with the experiment. It is not required that the final cost be less than that with no scattering. Thus, it is possible for an element to satisfy the convergence criterion on $F(\hat{\mathbf{g}})$ while also raising χ relative to the $\phi = 0$ state. This can occur, for instance, if the simulation generates scattering that is not observed in the experiment. Note also, however, that the cost function includes contributions from *all* scattering that strikes any detector. Peaks that strike only a single detector are used to determine optimal orientations even though such peaks are not used in counting Bragg peak overlaps. This means that high order peaks (typically those with the best projection geometry) and peaks with large $|\eta|$ contribute to the optimization.

III. VERIFICATION TESTS

Our experiments are carried out at APS undulator beamline XOR-1-ID. White radiation enters the B-hutch where the monochromator, sample stage, and detector system are mounted on a single optical table. We use the monochromator and focusing arrangement described in Ref. 23 with a vertically bent silicon monochromator in transmission geometry. Heavy slits are placed between the monochromator and sample to further reduce the background and to allow limitation of the horizontal extent of the beam. The detector system uses a liquid-nitrogen-cooled, lens-coupled, 16-bit, $1k \times 1k$ CCD camera focused on a Ce-doped yttrium aluminum garnet (YAG) crystal. The nominal pixel size is 4 μm, so the field of view is roughly 4×4 mm. The x-ray sensitive Ce-doped layer is thin (1 μm) so as to avoid smearing of the images of diffracted beams as they pass through the sensitive layer at different scattering angles.

The sample, detector, and coordinate system are shown schematically in Fig. 1. The sample is mounted on a conventional Huber 410 rotation stage. xy translations mounted on top of a vertical (z) translation provide longitudinal and transverse positioning relative to the rotation axis and x-ray beam. The single-crystal monochromator[23] requires that the sample and detector stages be mounted at a tilt angle of $2\theta^{\text{50 keV}}_{\text{Si(111)}} = 4.53°$. Prior to measurements, the rotation axis is aligned perpendicular to the beam plane by measuring absorption images of a steel cylinder as it is rotated in ω and translated along the axis. This alignment is required in order to keep a constant sample plane in the beam during rotation. Data collection is run through the control program SPEC using a hierarchy of macros. The macro starts data acquisition and ω motion, opens a "fast" shutter (located down-
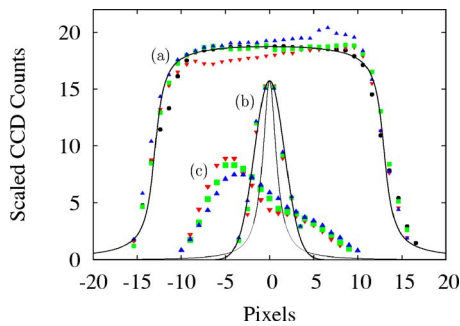
FIG. 9. Horizontal beam profile and horizontal and vertical diffraction spot profiles from a silicon wafer. (a) Circles show the horizontal incident beam profile after summing across the thin vertical width. The solid line is a guide to the eye. The full width at half-maximum is 98 μm. The downward, square, and upward pointing symbols are horizontal profiles of a (311) Bragg spot measured at L_1, L_2, and L_3, respectively. All profiles have been shifted to the same center and the incident beam has been normalized to the Bragg spots which have been scaled by 10^3. Since the single-crystal sample is wider than the beam, the horizontal extent of these spots is essentially that of the beam. (b) The vertical profile across a (311) Bragg spot [symbols as in (a)]. The light line indicates the detector response to the direct incident beam; the heavy line is a Gaussian guide to the eye showing that the vertical width of the peak is resolved. (c) Vertical profile of a higher-order, (553) or (731) Bragg peak showing the increased width due to larger deflection in the vertical direction.

stream of the monochromator) when the edge of the $\delta\omega$ interval is reached, and closes it again at the end of this interval before the motion is stopped.

The detector is positioned so that the field of view contains the direct beam near the bottom edge (as in Fig. 1). This means that we can observe the incident beam position at each L via short exposures with the beam stop removed. Also, relatively high-order diffraction will strike the detector; these peaks provide improved projection geometry.

A. Preliminary characterizations

Incident beam. Characterization of the incident beam yields a measure of the CCD detector spatial resolution as well as the width of the beam in the y direction. We first use a knife-edge scan to determine the incident beam width in z. A 0.5 μm gold film with a sharp edge is scanned across the beam while Au L-shell fluorescence is monitored with an energy dispersive detector. The film is rotated about the beam to make the film edge parallel to the line focused beam by minimizing the width of the fluorescence transition. The resultant beam profile is well characterized by a sharp central Gaussian plus Lorentzian tails and the height is found to be ≤ 2.5 μm FWHM. Figure 9(a) shows the horizontal intensity profile of the narrowed incident beam (upstream slits were set to a nominal 100 μm opening) as imaged on the high resolution detector. The intensity is uniform within 3% over the central 72 μm and the full width at half maximum is 98 μm.

Sample and measurement parameters. The sample was a small piece of a (111)-oriented silicon wafer, 5×3 mm^2 in cross section and was thinned to <0.2 mm. The sample was mounted with the (111) axis roughly parallel to the incident beam at $\omega = 0$. Data were collected using $\delta\omega = 1°$ integrations. We measured 40 such intervals over $\pm 20°$ in ω relative to normal incidence. For each $\delta\omega$, images were measured at L

≈ 5, 7, and 9 mm, so there are 120 2MB images. Each image and the associated $\delta\omega$ motion took 4 s. Data collection speed was further limited by the CCD readout time and was roughly 60 min.

For our range of incidence angles, there is a roughly diamond-shaped region of sample space that is always illuminated by the 100 μm wide incident beam as well as adjacent regions that move into and out of the beam. This presents an analysis challenge since different Bragg peaks do not originate from identical sample regions. We show below that we are nevertheless able to isolate the always illuminated region with good resolution. Measurements on polycrystals using a beam wider than the entire sample avoid this complication.

B. Silicon data set and fits

Diffracted beam images. Figure 9 shows images of diffracted beams. As expected, the y-direction profile of the diffraction spot is very close to that of the incident beam. The z-direction profiles are foreshortened projections of the wafer thickness in the x direction. The roughly tenfold compression of the x-direction geometry implies a similar loss in shape resolution in this direction. Low-order Bragg spots have a full width at half-maximum in the z direction of ~ 4 pixels or ~ 16 μm, which at the angle of observation corresponds to a sample thickness of roughly 160 μm. Higher-order spots with small $|\eta|$ have broader, fully resolved profiles that more precisely determine the x-direction dimensions of the diffracting entity.

Results. Fits were begun with no *a priori* information about sample orientation. The initially "guessed" orientation corresponds to crystalline cube axes aligned with the laboratory coordinate system at $\omega = 0$. 1594 Bragg peaks out to $|Q| \approx 14.4$ Å^{-1} or $2\theta = 33.2°$ were candidates for observation while only the lowest-order 300 of these were used in the orientation grid search. The simulation initially covered sample space with six equilateral triangles with 150 μm sides forming a hexagon centered on the rotation axis. Note that this hexagon is substantially larger than the region illuminated by the 100 μm wide beam; however, the *simulated* beam width was 200 μm so that scattering from anywhere in the hexagon can be generated by the simulation. One goal of the fitting was for the program to determine the spatial location of the actual sample, much as it must do in treating polycrystal data. Before fitting, area elements were regridded three times, yielding 384 triangular elements with 18.75 μm sides. After completion of three iterations, the smallest elements were 4.69 μm on a side. After each iteration, experimental parameters including the detector distance, L_1, and rotation axis projections onto the detector, j_{01}, j_{02}, j_{03} were adjusted. The intervals $L_2 - L_1$ and $L_3 - L_2$ were fixed to the measured values of 2 mm. The z-direction origins, k_{01}, k_{02}, k_{03}, were fixed to the incident beam positions as imaged on the detector before and after data collection. Prior fits had determined that the detector pixel pitch was 3.78 μm and the beam energy was 49.8 keV. The bandwidth was $<1\%$ and was not included in these simulations.

A variety of image analysis and fitting criteria were used. Table I shows values of f^{peak} and F_{min} used and numerical

TABLE I. Parameters used and results of three fits to silicon single-crystal data. Corresponding maps are shown in Fig. 10.

Fit	f^{peak}	F_{min}	Fitted elements	Fitted area (mm^2)	$\overline{\Psi}$ (deg)[a]	$\Delta\Psi$ (deg)[b]
a	0.25	0.5	1484	0.044	54.31	0.5
b	0.25	0.75	907	0.028	54.32	0.3
c	0.5	0.75	270	0.009	54.39	0.2

[a]Average rotation from reference orientation.
[b]Total rotation range of fitted elements.

results of fits. For all fits, $f^{hit}_{smeared}=0.1$ and $f^{hit}_{min}=0.5$. The simulation generates between 75 and 80 Bragg peaks (depending on orientation) that hit at least one detector and 38–42 "qualifying Bragg peaks" that hit detectors at multiple L's (see Sec. II D). Thus, $F_{min}=0.5$ or 0.75 mean that a grid element must generate scattering that hits ≥ 20 or 30 experimental peaks, respectively. For all criteria the software determined that a set of area elements near the center of the simulated hexagon contribute to the observed scattering. Fitted orientations are essentially independent of the criteria of Table I. The nominal Euler angles are $(\omega, \chi, \phi) = (350.4°, 34.2°, 52.5°)$, which correspond to $(\hat{\mathbf{n}}, \Psi) = (-0.55, 0.33, -0.77, 54.3°)$ in an axis-angle representation. Had the [111] wafer normal been mounted perfectly parallel with the incident beam direction, the rotation angle would have been $\cos^{-1}(1/\sqrt{3})=54.7°$. Orientation variation among fitted elements is small as indicated by the range, $\Delta\Psi$.

Figure 10 shows sample-space maps corresponding to the fits listed in Table I. Shading covers triangular grid elements for which qualifying orientations were determined. In the color version, the weighting of red, green, and blue (rgb) is determined by the fitted axis-angle pair, $(\hat{\mathbf{n}}, \Psi)$ for each area element. The color scale has been adjusted for each image: the full rgb range is scaled to include only the range of orientation spanned by the corresponding fit. In all cases, the orientation is uniform over almost all of the fitted region. Comparing Fig. 10(a) and 10(b), we see that the fitted region becomes more compact and regular as the F_{min} acceptance criterion increases. This trend indicates improved signal averaging as the number of matched peaks increases; equally important, it indicates that the image analysis and thresholding at a uniform fraction of ROI peak height generates a consistent set of projection contours. In Fig. 10(b), the fitted
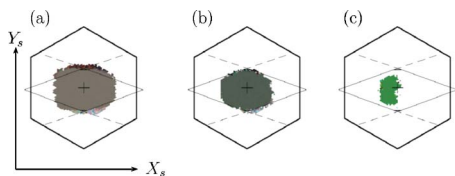

FIG. 10. Sample-space maps of the illuminated "slot" through the silicon wafer obtained under data reduction and convergence criteria of Table I. The maps are aligned so that the beam is incident from the left and the wafer is in the yz plane at $\omega=0$. The hexagons, with 150 μm sides, indicate the simulated region of sample space. The + sign indicates the simulated rotation axis. The dashed lines show the region of sample space illuminated by the *experimental* incident beam as the sample is rotated $\pm 20°$. The solid diamond outlines the region that remains in the beam throughout. White space inside the hexagons has been determined to not satisfy the fitting criteria for *any* crystallographic orientation and have been assigned $\phi=0$.
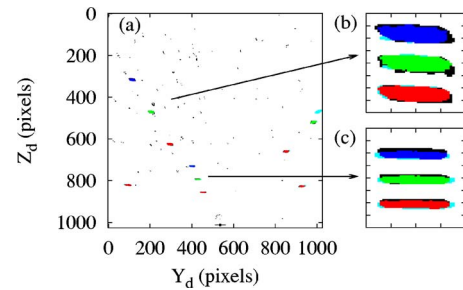

FIG. 11. Diffracted beam images at nominal orientation $\omega=-6.5°$ with fitted intensity corresponding to Fig. 10(a). In (a), three images, at L_1, L_2, L_3, have been superimposed. Experimental data are shown as black points. Colors (or gray scale) show simulated intensity: red, green and blue when overlapping experimental data, aqua when not. The line at the bottom indicates the projected position of the incident beam at L_1; the associated dot indicates the detector origin at $(j_{01}, k_{01})=(536, 1013)$. Bragg spots are seen to radiate from the incident beam position at different scattering angles, 2θ, and orientations, η. (b) and (c) show expanded views of individual Bragg spots; tic marks correspond to 20 pixel intervals; spot positions have been displaced for ease of comparison. The reduced height in (c) relative to (b) is due to the smaller vertical displacement causing a more anisotropic projection in (c). The at-most subtle broadening of experimental spots with L is consistent with an energy bandwidth $<1\%$.

region closely matches the shape of the always illuminated subarea of the sample. The left and right truncations are consistent with what is known of the thickness of the wafer—about 130 μm. The maximum deviation of edges from the always illuminated diamond region is 8 μm; the sharpness of the upstream (or left) face of the wafer is unresolved while the downstream face has deviations of $\sim$7 μm. This asymmetry is consistent with the line-shape asymmetry shown in Fig. 9: optimal position resolution is associated with sharp Bragg spot edges. Even in Fig. 10(b), deviations of orientation are seen at the edges of the fitted region—this demonstrates a boundary effect in which the orientation at the edge of a grain or sample can be adjusted slightly to obtain overlap with a sufficient number of observed peaks. In polycrystals, this effect may limit the precision of location of grain boundaries, but near boundary elements can choose which grain orientation provides the best fit. Figure 10(c) shows a fit to a data set generated with $f^{peak}=0.5$. Here, the downstream side of the sample is missing because the corresponding portion of the intensity has been eliminated (Fig. 9).

Figures 11 and 12 show the diffraction pattern observed in a single $\delta\omega$ interval, along with simulation results associated with the maps of Figs. 10(a) and 10(b), respectively.
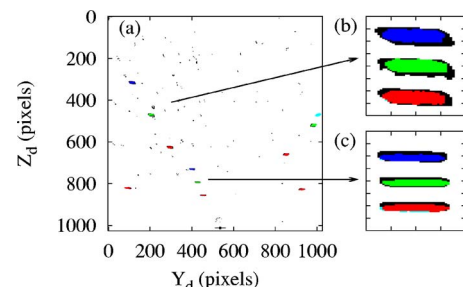

FIG. 12. Diffracted beam images at nominal orientation $\omega=-6.5°$ with fitted intensity corresponding to Fig. 10(b). All conventions are as in Fig. 11. Here, due to use of more stringent convergence criteria, the simulated intensity lies almost entirely within the experimental Bragg spots.

TABLE II. Experimental parameters obtained in the fit shown in Fig. 12.

Parameter	Units	Fitted value	Change
γ^{a}	(μm/pixel)	3.78	−0.22
$(j_0,k_0)_1$	pixels	(531.8, 1013.5)	(−3.2, 0)
$(j_0,k_0)_2$	pixels	(534.7, 1013.5)	(−0.3, 0)
$(j_0,k_0)_3$	pixels	(537.0, 1012.5)	(2.0, 0)
L_1	mm	5.04	0.04
E	keV	49.8	−0.2
χ		39605^{b}	

[a]Detector pixel pitch calibration.
[b]Reduced data hit 48685 pixels.

The chosen ω interval is atypical in that most intervals contain only one or no diffraction peaks from the single crystal sample. Here, five diffracted beams are present at L_1, three at L_2, and two at L_3. One high-order simulated peak (at L_1) that is not in the experimental data set occurs at the right edge; this peak has $2\theta = 29.3°$ with Miller indices (775) or (11,1,1) and was probably weak enough to be missed in the data reduction. Simulated diffraction spots are expected and seen to be smaller than the experimental ones: experimental spots are projections from the entire illuminated volume of the sample at each ω, whereas the simulation includes scattering only from those grid elements that remain in the beam and diffract at many other ω positions. At large $|\omega|$ more of the sample is in the beam and there is more "missed" intensity than near $\omega = 0$. As f_{hit} increases, the simulated spots become smaller as can be seen in the figures. The entire data set and these fits can be seen at http://www.andrew.cmu.edu/user/suter/3dxdm/3dxdm.html.

IV. DISCUSSION

We have described and verified a method that should allow nondestructive internal microstructure mapping in a wide variety of sample types. Our current analysis software closely mimics the measurement and can account for a variety of experimental and sample complexities. Here, we have demonstrated the ability to resolve, to ~ 10 μm, the shape of a region of sample with "soft" edges. Three dimensional maps of more interesting samples are a straightforward extension of this work; in fact, we have recently succeeded in reconstructing a significant volume of microstructure in an aluminum polycrystal sample.[25] Many applications suggest themselves, including fundamental studies of grain growth coupled to three-dimensional computer simulations,[12,13] response of microstructure to stress of a variety of types, and the structure of polyphase materials and solid-state phase transformations. In each case, it should be possible to study the response of three-dimensional, interconnected ensembles of grains.

An essential part of obtaining images such as in Fig. 10 and matching diffraction spots as shown in Figs. 11 and 12 is the ability to optimize uncertain experimental parameters during data fitting. This optimization is done by minimizing the cost function, χ, via zero temperature Monte Carlo. Several parameters are listed in Table II along with their fitted values and the change from initial estimates. Note that parameters must be known *a priori* sufficiently that at least

some sample area elements can be fitted initially; our algorithm performs Monte Carlo adjustment of these parameters after each iteration of fitting to orientations. Note that the k_0 parameters were not adjusted in these fits since measurements of the position on the detector of the incident beam define this precisely. The final value of the Monte Carlo cost function, χ, yields ~ 330 "missed" pixels per detector image (on average) which is consistent with the background noise and uncovered experimental pixels seen in Fig. 12. The simulation overlaps on the order of 1400 pixels in each diffraction spot.

The number of parameters that can be adjusted independently may be surprising. Optimization of the detector origins (j_0,k_0) at each measurement distance makes experimental Bragg spots colinear in L. It is assumed that the origins are constant across all ω so these adjustments are made against many fitted spots. L_1 adjustment, keeping the mechanically defined intervals to other L's fixed, makes the straight line paths with fixed 2θ values extrapolate consistently to the sample and peaks from individual grains extrapolate to a common position in the illuminated plane. The detector calibration constant and the x-ray photon energy are independent: γ and E both adjust scattering angles but E also changes the distance in ω between simulated peaks, as shown in Fig. 3.

Underlying all this work is the assumption that diffraction spots can be consistently analyzed to obtain geometrical information. It is encouraging that for the silicon data presented here a simple thresholding scheme yields quite precise geometry. Silicon has classes of Bragg peaks with intensities different by a factor of 2 and intensities vary with Q due to form factor and lattice vibration effects, so an independent analysis of each ROI in the raw images is essential. Complex line shapes in interesting materials most likely preclude any universally correct threshold fraction, f^{peak}. Some self-consistency tests can be imagined: in polycrystals, too high a threshold should yield maps with holes, so one may be able to adjust f^{peak} downward to just obtain completion. Without doubt, comparisons of x-ray data with orientation imaging microscopy[9,10] on sectioned layers of the same sample will be necessary. Finally, realistic intensity calculations may allow the use of information that we are now throwing away—the intensity distributions within diffraction spots.

Combining Monte Carlo simulation, physically based line-shape models, and the inclusion of peak intensity distributions (with modeling of detector point-spread functions) holds the possibility of yielding extremely detailed microstructure information. Including in the measurement a strain sensitive detector well downstream of the sample[26] (combined with a semitransparent high-resolution imaging detector) could yield intragrain strain maps with geometric and orientational information about intergranular environments. Of course, complex models imply additional computation and a multitude of degrees of freedom. The Monte Carlo allows for the making of correlated moves on sets of parameters that cross area elements such as would be generated by a continuous defect density model for a grain. Again, the forward model has the robust feature that changing one scat-

tering parameter will generate changes in all the scattering generated by the grain or area element. Finally, a dynamically optimized finite temperature Monte Carlo[27] provides a means for gaining efficiency, determining correlations among parameters, and estimating errors.

Facilities at the Advanced Photon Source, while adequate for the silicon measurements presented here and preliminary work on polycrystals, are currently being upgraded. A high-precision air bearing rotation stage with submicrometer errors is being commissioned. Refractive focusing optics have been developed that allow further reduction in the energy band pass. Purchase of a CCD camera with more rapid readout is planned as is the design of a semitransparent scintillator system that will allow combined mapping and strain measurements. In the longer term, an undulator source tailored for high energies is planned. These upgrades will speed data collection and further improve spatial and angular resolution.

ACKNOWLEDGMENTS

[1] H. F. Poulsen, *Three-Dimensional X-ray Diffraction Microscopy*, Springer Tracts in Modern Physics, Vol. 205, edited by G. Hohler (Springer, New York, 2004).

[2] H. F. Poulsen, X. Fu, E. Knudsen, E. M. Lauridsen, L. Margulies, and S. Schmidt, Mater. Sci. Forum **467–470**, 1363 (2004).

[3] B. Jakobsen, H. F. Poulsen, U. Lienert, J. Almer, S. D. Shastri, H. O. Sørensen, C. Gundlach, and W. Pantleon, Science **312**, 889 (2006).

[4] L. Margulies, G. Winther, and H. F. Poulsen, Science **291**, 2392 (2001); see also the associated "Perspective" by F. Heidelbach, *ibid.* **291**, 2330 (2001).

[5] S. Schmidt, S. F. Nielsen, C. Gundlach, L. Margulies, X. Huang, and D. Juul Jensen, Science **305**, 229 (2004).

[6] S. E. Offerman, N. H. van Dijk, J. Sietsma, S. Grigull, E. M. Lauridsen, L. Margulies, H. F. Poulsen, M. Th. Rekveldt, and S. van der Zwaag, Science **298**, 1003 (2002).

[7] T. Kubo, E. Ohtani, and K. Funakoshi, Am. Mineral. **89**, 285 (2004).

[8] E. M. Lauridsen, S. Schmidt, R. M. Suter, and H. F. Poulsen, J. Appl. Crystallogr. **34**, 744 (2001).

[9] H. F. Poulsen, S. F. Nielsen, E. M. Lauridsen, S. Schmidt, R. M. Suter, U. Lienert, L. Margulies, T. Lorentzen, and D. Juul Jensen, J. Appl. Crystallogr. **34**, 751 (2001).

[10] G. S. Rohrer, Annu. Rev. Mater. Res. **35**, 99 (2005).

[11] See also http://mimp.materials.cmu.edu/ and publications listed thereon.

[12] J. Gruber, D. C. George, A. P. Kuprat, G. S. Rohrer, and A. D. Rollett, Scr. Mater. **53**, 351 (2005).

[13] D. Kinderlehrer, J. Lee, I. Livshits, and S. Taasan, in *Continuum Scale Simulation of Engineering Materials: Fundamentals—Microstructures—Process Applications*, edited by D. Raabe, F. Roters, F. Barlat, and L.-Q. Chen (Wiley-VCH, Berlin, 2004), pp. 356–367.

[14] G. E. Ice, B. C. Larson, W. Yang, J. D. Budai, J. Z. Tischler, J. W. L. Pang, R. I. Barabash, and W. Liu, J. Synchrotron Radiat. **12**, 155 (2005).

[15] L. E. Levine, B. C. Larson, W. Yange, M. E. Kassner, J. Z. Tischler, M. A. Delos-Reyes, R. J. Fields, and W. Liu, Nat. Mater. **5**, 619 (2006).

[16] H. O. Sørensen, B. Jakobsen, E. Knudsen, E. M. Lauridsen, S. F. Nielsen, H. F. Poulsen, S. Schmidt, G. Winther, and L. Margulies, Nucl. Instrum. Methods Phys. Res. B **246**, 232 (2006).

[17] H. F. Poulsen and F. Xiaowei, J. Appl. Crystallogr. **36**, 1062 (2003).

[18] S. Schmidt, H. F. Poulsen, and G. B. M. Vaughan, J. Appl. Crystallogr. **36**, 326 (2003).

[19] T. Markussen, X. Fu, L. Margulies, E. M. Lauridsen, S. F. Nielsen, S. Schmidt, and H. F. Poulsen, J. Appl. Crystallogr. **37**, 96 (2004).

[20] X. Fu, E. Knudsen, H. F. Poulsen, G. T. Herman, T. Gabor, B. M. Carvalho, and H. Y. Liao, Proc. SPIE **5535**, 261 (2004).

[21] A. Alpers, E. Knudsen, G. T. Herman, and H. F. Poulsen, J. Appl. Crystallogr. **39**, 582 (2006).

[22] N. K. Roy and R. M. Suter (unpublished).

[23] U. Lienert, C. Schulze, V. Honkimaki, T. Tschentscher, S. Garbe, O. Hignette, A. Horsewell, M. Lingham, H. F. Poulsen, N. B. Thomsen, and E. Ziegler, J. Synchrotron Radiat. **5**, 226 (1998).

[24] We take the variation in angle of incidence, $\delta\chi$ to be the difference in scattering angles for radiation from the (111) planes of a silicon monochromator crystal: $\delta\chi = 2\theta - 2\theta_0 = 2G_{111}^{Si}(1/2k - 1/2k_0)$. Energies greater than the nominal have $\delta\chi < 0$, corresponding to having $k_i^z < 0$. This formula approximates the incoming white beam as being parallel.

[25] R.M. Suter, C. Hefferan, D. Hennessy, C. Xiao, U. Lienert, in preparation.

[26] R. V. Martins, L. Margulies, S. Schmidt, H. F. Poulsen, and T. Leffers, Mater. Sci. Eng., A **387–389**, 84 (2004).

[27] D. Bouzida, S. Kumar, and R. H. Swendsen, Phys. Rev. A **45**, 8894 (1992).